

Preface

My primary goal in writing *Understanding Analysis* was to create an elementary one-semester book that exposes students to the rich rewards inherent in taking a mathematically rigorous approach to the study of functions of a real variable. The aim of a course in real analysis should be to challenge and improve mathematical intuition rather than to verify it. There is a tendency, however, to center an introductory course too closely around the familiar theorems of the standard calculus sequence. Producing a rigorous argument that polynomials are continuous is good evidence for a well-chosen definition of continuity, but it is not the reason the subject was created and certainly not the reason it should be required study. By shifting the focus to topics where an untrained intuition is severely disadvantaged (e.g., rearrangements of infinite series, nowhere-differentiable continuous functions, Fourier series), my intent is to restore an intellectual liveliness to this course by offering the beginning student access to some truly significant achievements of the subject.

The Main Objectives

In recent years, the standard undergraduate curriculum in mathematics has been subjected to steady pressure from several different sources. As computers and technology become more ubiquitous, so do the areas where mathematical thinking can be a valuable asset. Rather than preparing themselves for graduate study in pure mathematics, the present majority of mathematics majors look forward to careers in banking, medicine, law, and numerous other fields where analytical skills are desirable. Another strong influence on college mathematics is the ongoing calculus reform effort, now well over ten years old. At the core of this movement is the justifiable goal of presenting calculus in a more intuitive way, emphasizing geometric arguments over symbolic ones. Despite these various trends—or perhaps because of them—nearly every undergraduate mathematics program continues to require at least one semester of real analysis. The result is that instructors today are faced with the task of teaching a difficult, abstract course to a more diverse audience less familiar with the nature of axiomatic arguments.

The crux of the matter is that any prevailing sentiment in favor of marketing mathematics to larger groups must at some point be reconciled with the fact

that theoretical analysis is extremely challenging and even intimidating for some. One unfortunate resolution of this dilemma has been to make the course easier by making it less interesting. The omitted material is inevitably what gives analysis its true flavor. A better solution is to find a way to make the more advanced topics accessible and worth the effort.

I see three essential goals that a semester of real analysis should try to meet:

1. Students, especially those emerging from a reform approach to calculus, need to be convinced of the need for a more rigorous study of functions. The necessity of precise definitions and an axiomatic approach must be carefully motivated.
2. Having seen mainly graphical, numerical, or intuitive arguments, students need to learn what constitutes a rigorous mathematical proof and how to write one.
3. There needs to be significant reward for the difficult work of firming up the logical structure of limits. Specifically, real analysis should not be just an elaborate reworking of standard introductory calculus. Students should be exposed to the tantalizing complexities of the real line, to the subtleties of different flavors of convergence, and to the intellectual delights hidden in the paradoxes of the infinite.

The philosophy of *Understanding Analysis* is to focus attention on questions that give analysis its inherent fascination. Does the Cantor set contain any irrational numbers? Can the set of points where a function is discontinuous be arbitrary? Are derivatives continuous? Are derivatives integrable? Is an infinitely differentiable function necessarily the limit of its Taylor series? In giving these topics center stage, the hard work of a rigorous study is justified by the fact that *they are inaccessible without it*.

The Structure of the Book

This book is an introductory text. Although some fairly sophisticated topics are brought in early to advertise and motivate the upcoming material, the main body of each chapter consists of a lean and focused treatment of the core topics that make up the center of most courses in analysis. Fundamental results about completeness, compactness, sequential and functional limits, continuity, uniform convergence, differentiation, and integration are all incorporated. What is specific here is where the emphasis is placed. In the chapter on integration, for instance, the exposition revolves around deciphering the relationship between continuity and the Riemann integral. Enough properties of the integral are obtained to justify a proof of the Fundamental Theorem of Calculus, but the theme of the chapter is the pursuit of a characterization of integrable functions in terms of continuity. Whether or not Lebesgue's measure-zero criterion is treated, framing the material in this way is still valuable because it is the questions that are important. Mathematics is not a static discipline. Students

should be aware of the historical reasons for the creation of the mathematics they are learning and by extension realize that there is no last word on the subject. In the case of integration, this point is made explicitly by including some relatively recent developments on the generalized Riemann integral in the additional topics of the last chapter.

The structure of the chapters has the following distinctive features.

Discussion Sections: Each chapter begins with the discussion of some motivating examples and open questions. The tone in these discussions is intentionally informal, and full use is made of familiar functions and results from calculus. The idea is to freely explore the terrain, providing context for the upcoming definitions and theorems. A recurring theme is the resolution of the paradoxes that arise when operations that work well in finite settings are naively extended to infinite settings (e.g., differentiating an infinite series term-by-term, reversing the order of a double summation). After these exploratory introductions, the tone of the writing changes, and the treatment becomes rigorously tight but still not overly formal. With the questions in place, the need for the ensuing development of the material is well-motivated and the payoff is in sight.

Project Sections: The penultimate section of each chapter (the final section is a short epilogue) is written with the exercises incorporated into the exposition. Proofs are outlined but not completed, and additional exercises are included to elucidate the material being discussed. The point of this is to provide some flexibility. The sections are written as self-guided tutorials, but they can also be the subject of lectures. I have used them in place of a final examination, and they work especially well as collaborative assignments that can culminate in a class presentation. The body of each chapter contains the necessary tools, so there is some satisfaction in letting the students use their newly acquired skills to ferret out for themselves answers to questions that have been driving the exposition.

Building a Course

Teaching a satisfying class inevitably involves a race against time. Although this book is designed for a 12–14 week semester, there are still a few choices to make as to what to cover.

- The introductions can be discussed, assigned as reading, omitted, or substituted with something preferable. There are no theorems proved here that show up later in the text. I do develop some important examples in these introductions (the Cantor set, Dirichlet's nowhere-continuous function) that probably need to find their way into discussions at some point.
- Chapter 3, Basic Topology of \mathbf{R} , is much longer than it needs to be. All that is required by the ensuing chapters are fundamental results about open and closed sets and a thorough understanding of sequential compactness. The characterization of compactness using open covers as well

as the section on perfect and connected sets are included for their own intrinsic interest. They are not, however, crucial to any future proofs. The one exception to this is a presentation of the Intermediate Value Theorem (IVT) as a special case of the preservation of connected sets by continuous functions. To keep connectedness truly optional, I have included two direct proofs of IVT, one using least upper bounds and the other using nested intervals. A similar comment can be made about perfect sets. Although proofs of the Baire Category Theorem are nicely motivated by the argument that perfect sets are uncountable, it is certainly possible to do one without the other.

- All the project sections (1.5, 2.8, 3.5, 4.6, 5.4, 6.6, 7.6, 8.1–8.4) are optional in the sense that no results in later chapters depend on material in these sections. The four topics covered in Chapter 8 are also written in this project-style format, where the exercises make up a significant part of the development. The only one of these sections that might require a lecture is the unit on Fourier series, which is a bit longer than the others.

The Audience

The only prerequisite for this course is a robust understanding of the results from single-variable calculus. The theorems of linear algebra are not needed, but the exposure to abstract arguments and proof writing that usually comes with this course would be a valuable asset. Complex numbers are never used in this book.

The proofs in *Understanding Analysis* are written with the introductory student firmly in mind. Brevity and other stylistic concerns are postponed in favor of including a significant level of detail. Most proofs come with a fair amount of discussion about the context of the argument. What should the proof entail? Which definitions are relevant? What is the overall strategy? Is one particular proof similar to something already done? Whenever there is a choice, efficiency is traded for an opportunity to reinforce some previously learned technique. Especially familiar or predictable arguments are usually sketched as exercises so that students can participate directly in the development of the core material.

The search for recurring ideas exists at the proof-writing level and also on the larger expository level. I have tried to give the course a narrative tone by picking up on the unifying themes of approximation and the transition from the finite to the infinite. To paraphrase a passage from the end of the book, real numbers are approximated by rational ones; values of continuous functions are approximated by values nearby; curves are approximated by straight lines; areas are approximated by sums of rectangles; continuous functions are approximated by polynomials. In each case, the approximating objects are tangible and well-understood, and the issue is when and how well these qualities survive the limiting process. By focusing on this recurring pattern, each successive topic

builds on the intuition of the previous one. The questions seem more natural, and a method to the madness emerges from what might otherwise appear as a long list of theorems and proofs.

This book always emphasizes core ideas over generality, and it makes no effort to be a complete, deductive catalog of results. It is designed to capture the intellectual imagination. Those who become interested are then exceptionally well prepared for a second course starting from complex-valued functions on more general spaces, while those content with a single semester come away with a strong sense of the essence and purpose of real analysis. Turning once more to the concluding passages of Chapter 8, “By viewing the different infinities of mathematics through pathways crafted out of finite objects, Weierstrass and the other founders of analysis created a paradigm for how to extend the scope of mathematical exploration deep into territory previously unattainable.”

This exploration has constituted the major thrill of my intellectual life. I am extremely pleased to offer this guide to what I feel are some of the most impressive highlights of the journey. Have a wonderful trip!

Acknowledgments

The genesis of this book came from an extended series of conversations with Benjamin Lotto of Vassar College. The structure of the early chapters and the book’s overall thesis are in large part the result of several years of sharing classroom notes, ideas, and experiences with Ben. I am pleased with how the manuscript has turned out, and I have no doubt that it is an immeasurably better book because of Ben’s early contributions.

A large part of the writing was done while I was enjoying a visiting position at the University of Virginia. Special thanks go to Nat Martin and Larry Thomas for being so generous with their time and wisdom, and especially to Loren Pitt, the scope of whose advice extends well beyond the covers of this book. I would also like to thank Julie Riddleberger for her help with many of the figures. Marian Robbins of Bellarmine College, Steve Kennedy of Carleton College, Paul Humke of Saint Olaf College, and Tom Kriete of the University of Virginia each taught from a preliminary draft of this text. I appreciate the many suggested improvements that this group provided, and I want to especially acknowledge Paul Humke for his contributions to the chapter on integration.

My department and the administration of Middlebury College have also been very supportive of this endeavor. David Guertin came to my technological rescue on numerous occasions, Priscilla Bremser read early chapter drafts, and Rick Chartrand’s insightful opinions greatly improved some of the later sections. The list of students who have suffered through the long evolution of this book is now too long to present, but I would like to mention Brooke Sargent, whose meticulous class notes were the basis of the first draft, and Jesse Johnson, who has worked tirelessly to improve the presentation of the many exercises in the book. The production team at Springer has been absolutely first-rate. My sincere thanks goes to all of them with a special nod to Sheldon

Axler for encouragement and advice surely exceeding anything in his usual job description.

In a recent rereading of the completed text, I was struck by how frequently I resort to historical context to motivate an idea. This was not a conscious goal I set for myself. Instead, I feel it is a reflection of a very encouraging trend in mathematical pedagogy to humanize our subject with its history. From my own experience, a good deal of the credit for this movement in analysis should go to two books: *A Radical Approach to Real Analysis*, by David Bressoud, and *Analysis by Its History*, by E. Hairer and G. Wanner. Bressoud's book was particularly influential to the presentation of Fourier series in the last chapter. Either of these would make an excellent supplementary resource for this course. While I do my best to cite their historical origins when it seems illuminating or especially important, the present form of many of the theorems presented here belongs to the common folklore of the subject, and I have not attempted careful attribution. One exception is the material on the previously mentioned generalized Riemann integral, due independently to Jaroslav Kurzweil and Ralph Henstock. Section 8.1 closely follows the treatment laid out in Robert Bartle's article "Return to the Riemann Integral." In this paper, the author makes an italicized plea for teachers of mathematics to supplant Lebesgue's ubiquitous integral with the generalized Riemann integral. I hope that Professor Bartle will see its inclusion here as an inspired response to that request.

On a personal note, I welcome comments of any nature, and I will happily share any enlightening remarks—and any corrections—via a link on my webpage. The publication of this book comes nearly four years after the idea was first hatched. The long road to this point has required the steady support of many people but most notably that of my incredible wife, Katy. Amid the flurry of difficult decisions and hard work that go into a project of this size, the opportunity to dedicate this book to her comes as a pure and easy pleasure.

Middlebury, Vermont
August 2000

Stephen Abbott

Contents

Preface	v
1 The Real Numbers	1
1.1 Discussion: The Irrationality of $\sqrt{2}$	1
1.2 Some Preliminaries	4
1.3 The Axiom of Completeness	13
1.4 Consequences of Completeness	18
1.5 Cantor's Theorem	29
1.6 Epilogue	33
2 Sequences and Series	35
2.1 Discussion: Rearrangements of Infinite Series	35
2.2 The Limit of a Sequence	38
2.3 The Algebraic and Order Limit Theorems	44
2.4 The Monotone Convergence Theorem and a First Look at Infinite Series	50
2.5 Subsequences and the Bolzano–Weierstrass Theorem	55
2.6 The Cauchy Criterion	58
2.7 Properties of Infinite Series	62
2.8 Double Summations and Products of Infinite Series	69
2.9 Epilogue	73
3 Basic Topology of \mathbf{R}	75
3.1 Discussion: The Cantor Set	75
3.2 Open and Closed Sets	78
3.3 Compact Sets	84
3.4 Perfect Sets and Connected Sets	89
3.5 Baire's Theorem	94
3.6 Epilogue	96
4 Functional Limits and Continuity	99
4.1 Discussion: Examples of Dirichlet and Thomae	99
4.2 Functional Limits	103
4.3 Combinations of Continuous Functions	109

4.4	Continuous Functions on Compact Sets	114
4.5	The Intermediate Value Theorem	120
4.6	Sets of Discontinuity	125
4.7	Epilogue	127
5	The Derivative	129
5.1	Discussion: Are Derivatives Continuous?	129
5.2	Derivatives and the Intermediate Value Property	131
5.3	The Mean Value Theorem	137
5.4	A Continuous Nowhere-Differentiable Function	144
5.5	Epilogue	148
6	Sequences and Series of Functions	151
6.1	Discussion: Branching Processes	151
6.2	Uniform Convergence of a Sequence of Functions	154
6.3	Uniform Convergence and Differentiation	164
6.4	Series of Functions	167
6.5	Power Series	169
6.6	Taylor Series	176
6.7	Epilogue	181
7	The Riemann Integral	183
7.1	Discussion: How Should Integration be Defined?	183
7.2	The Definition of the Riemann Integral	186
7.3	Integrating Functions with Discontinuities	191
7.4	Properties of the Integral	195
7.5	The Fundamental Theorem of Calculus	199
7.6	Lebesgue's Criterion for Riemann Integrability	203
7.7	Epilogue	210
8	Additional Topics	213
8.1	The Generalized Riemann Integral	213
8.2	Metric Spaces and the Baire Category Theorem	222
8.3	Fourier Series	228
8.4	A Construction of \mathbf{R} From \mathbf{Q}	243
	Bibliography	251
	Index	253

Chapter 1

The Real Numbers

1.1 Discussion: The Irrationality of $\sqrt{2}$

Toward the end of his distinguished career, the renowned British mathematician G.H. Hardy eloquently laid out a justification for a life of studying mathematics in *A Mathematician's Apology*, an essay first published in 1940. At the center of Hardy's defense is the thesis that mathematics is an aesthetic discipline. For Hardy, the applied mathematics of engineers and economists held little charm. "Real mathematics," as he referred to it, "must be justified as art if it can be justified at all."

To help make his point, Hardy includes two theorems from classical Greek mathematics, which, in his opinion, possess an elusive kind of beauty that, although difficult to define, is easy to recognize. The first of these results is Euclid's proof that there are an infinite number of prime numbers. The second result is the discovery, attributed to the school of Pythagoras from around 500 B.C., that $\sqrt{2}$ is irrational. It is this second theorem that demands our attention. (A course in number theory would focus on the first.) The argument uses only arithmetic, but its depth and importance cannot be overstated. As Hardy says, "[It] is a 'simple' theorem, simple both in idea and execution, but there is no doubt at all about [it being] of the highest class. [It] is as fresh and significant as when it was discovered—two thousand years have not written a wrinkle on [it]."

Theorem 1.1.1. *There is no rational number whose square is 2.*

Proof. A rational number is any number that can be expressed in the form p/q , where p and q are integers. Thus, what the theorem asserts is that no matter how p and q are chosen, it is never the case that $(p/q)^2 = 2$. The line of attack is indirect, using a type of argument referred to as a proof by contradiction. The idea is to assume that there *is* a rational number whose square is 2 and then proceed along logical lines until we reach a conclusion that is unacceptable. At this point, we will be forced to retrace our steps and reject the erroneous

assumption that some rational number squared is equal to 2. In short, we will prove that the theorem is true by demonstrating that it cannot be false.

And so assume, for contradiction, that there exist integers p and q satisfying

$$(1) \quad \left(\frac{p}{q}\right)^2 = 2.$$

We may also assume that p and q have no common factor, because, if they had one, we could simply cancel it out and rewrite the fraction in lowest terms. Now, equation (1) implies

$$(2) \quad p^2 = 2q^2.$$

From this, we can see that the integer p^2 is an even number (it is divisible by 2), and hence p must be even as well because the square of an odd number is odd. This allows us to write $p = 2r$, where r is also an integer. If we substitute $2r$ for p in equation (2), then a little algebra yields the relationship

$$2r^2 = q^2.$$


But now the absurdity is at hand. This last equation implies that q^2 is even, and hence q must also be even. Thus, we have shown that p and q are both even (i.e., divisible by 2) when they were originally assumed to have no common factor. From this logical impasse, we can only conclude that equation (1) *cannot* hold for any integers p and q , and thus the theorem is proved. \square

A component of Hardy's definition of beauty in a mathematical theorem is that the result have lasting and serious implications for a network of other mathematical ideas. In this case, the ideas under assault were the Greeks' understanding of the relationship between geometric *length* and arithmetic *number*. Prior to the preceding discovery, it was an assumed and commonly used fact that, given two line segments \overline{AB} and \overline{CD} , it would always be possible to find a third line segment whose length divides evenly into the first two. In modern terminology, this is equivalent to asserting that the length of \overline{CD} is a rational multiple of the length of \overline{AB} . Looking at the diagonal of a unit square (Fig. 1.1), it now followed (using the Pythagorean Theorem) that this was not always the case. Because the Pythagoreans implicitly interpreted number to mean rational number, they were forced to accept that number was a strictly weaker notion than length.

Rather than abandoning arithmetic in favor of geometry (as the Greeks seem to have done), our resolution to this limitation is to strengthen the concept of number by moving from the rational numbers to a larger number system. From a modern point of view, this should seem like a familiar and somewhat natural phenomenon. We begin with the *natural numbers*

$$\mathbf{N} = \{1, 2, 3, 4, 5, \dots\}.$$

The influential German mathematician Leopold Kronecker (1823–1891) once asserted that “The natural numbers are the work of God. All of the rest is

Figure 1.1: $\sqrt{2}$ EXISTS AS A GEOMETRIC LENGTH.

the work of mankind.” Debating the validity of this claim is an interesting conversation for another time. For the moment, it at least provides us with a place to start. If we restrict our attention to the natural numbers \mathbf{N} , then we can perform addition perfectly well, but we must extend our system to the *integers*

$$\mathbf{Z} = \{ \dots, -3, -2, -1, 0, 1, 2, 3, \dots \}$$

if we want to have an additive identity (zero) and the additive inverses necessary to define subtraction. The next issue is multiplication and division. The number 1 acts as the multiplicative identity, but in order to define division we need to have multiplicative inverses. Thus, we extend our system again to the *rational numbers*

$$\mathbf{Q} = \left\{ \text{all fractions } \frac{p}{q} \text{ where } p \text{ and } q \text{ are integers with } q \neq 0 \right\}.$$

Taken together, the properties of \mathbf{Q} discussed in the previous paragraph essentially make up the definition of what is called a *field*. More formally stated, a field is any set where addition and multiplication are well-defined operations that are commutative, associative, and obey the familiar distributive property $a(b+c) = ab+ac$. There must be an additive identity, and every element must have an additive inverse. Finally, there must be a multiplicative identity, and multiplicative inverses must exist for all nonzero elements of the field. Neither \mathbf{Z} nor \mathbf{N} is a field. The finite set $\{0, 1, 2, 3, 4\}$ is a field when addition and multiplication are computed modulo 5. This is not immediately obvious but makes an interesting exercise (Exercise 1.3.1).

The set \mathbf{Q} also has a natural *order* defined on it. Given any two rational numbers r and s , exactly one of the following is true:

$$r < s, \quad r = s, \quad \text{or} \quad r > s.$$

This ordering is transitive in the sense that if $r < s$ and $s < t$, then $r < t$, so we are conveniently led to a mental picture of the rational numbers as being laid out from left to right along a number line. Unlike \mathbf{Z} , there are no intervals of empty space. Given any two rational numbers $r < s$, the rational number




Figure 1.2: APPROXIMATING $\sqrt{2}$ WITH RATIONAL NUMBERS.

$(r+s)/2$ sits halfway in between, implying that the rational numbers are densely nestled together.

With the field properties of \mathbf{Q} allowing us to safely carry out the algebraic operations of addition, subtraction, multiplication, and division, let's remind ourselves just what it is that \mathbf{Q} is lacking. By Theorem 1.1.1, it is apparent that we cannot always take square roots. The problem, however, is actually more fundamental than this. Using only rational numbers, it is possible to *approximate* $\sqrt{2}$ quite well (Fig. 1.2). For instance, $1.414^2 = 1.999396$. By adding more decimal places to our approximation, we can get even closer to a value for $\sqrt{2}$, but, even so, we are now well aware that there is a “hole” in the rational number line where $\sqrt{2}$ ought to be. Of course, there are quite a few other holes—at $\sqrt{3}$ and $\sqrt{5}$, for example. Returning to the dilemma of the ancient Greek mathematicians, if we want every length along the number line to correspond to an actual number, then another extension to our number system is in order. Thus, to the chain $\mathbf{N} \subseteq \mathbf{Z} \subseteq \mathbf{Q}$ we append the *real numbers* \mathbf{R} .

The question of how to actually construct \mathbf{R} from \mathbf{Q} is rather complicated business. It is discussed in Section 1.3, and then again in more detail in Section 8.4. For the moment, it is not too inaccurate to say that \mathbf{R} is obtained by filling in the gaps in \mathbf{Q} . Wherever there is a hole, a new *irrational* number is defined and placed into the ordering that already exists on \mathbf{Q} . The real numbers are then the union of these irrational numbers together with the more familiar rational ones. What properties does the set of irrational numbers have? How do the sets of rational and irrational numbers fit together? Is there a kind of symmetry between the rationals and the irrationals, or is there some sense in which we can argue that one type of real number is more common than the other? The one method we have seen so far for generating examples of irrational numbers is through square roots. Not too surprisingly, other roots such as $\sqrt[3]{2}$ or $\sqrt[5]{3}$ are most often irrational. Can all irrational numbers be expressed as algebraic combinations of n th roots and rational numbers, or are there still other irrational numbers beyond those of this form?

1.2 Some Preliminaries

The vocabulary necessary for the ensuing development comes from set theory and the theory of functions. This should be familiar territory, but a brief review

of the terminology is probably a good idea, if only to establish some agreed-upon notation.

Sets

Intuitively speaking, a *set* is any collection of objects. These objects are referred to as the *elements* of the set. For our purposes, the sets in question will most often be sets of real numbers, although we will also encounter sets of functions and, on a few rare occasions, sets whose elements are other sets.

Given a set A , we write $x \in A$ if x (whatever it may be) is an element of A . If x is not an element of A , then we write $x \notin A$. Given two sets A and B , the *union* is written $A \cup B$ and is defined by asserting that

$$x \in A \cup B \text{ provided that } x \in A \text{ or } x \in B \text{ (or potentially both).}$$

The *intersection* $A \cap B$ is the set defined by the rule

$$x \in A \cap B \text{ provided } x \in A \text{ and } x \in B.$$

Example 1.2.1. (i) There are many acceptable ways to assert the contents of a set. In the previous section, the set of natural numbers was defined by listing the elements: $\mathbf{N} = \{1, 2, 3, \dots\}$.

(ii) Sets can also be described in words. For instance, we can define the set E to be the collection of even natural numbers.

(iii) Sometimes it is more efficient to provide a kind of rule or algorithm for determining the elements of a set. As an example, let

$$S = \{r \in \mathbf{Q} : r^2 < 2\}.$$

Read aloud, the definition of S says, “Let S be the set of all rational numbers whose squares are less than 2.” It follows that $1 \in S$, $4/3 \in S$, but $3/2 \notin S$ because $9/4 \geq 2$.

Using the previously defined sets to illustrate the operations of intersection and union, we observe that

$$\mathbf{N} \cup E = \mathbf{N}, \quad \mathbf{N} \cap E = E, \quad \mathbf{N} \cap S = \{1\}, \text{ and } E \cap S = \emptyset.$$

The set \emptyset is called the *empty set* and is understood to be the set that contains no elements. An equivalent statement would be to say that E and S are *disjoint*.

A word about the equality of two sets is in order (since we have just used the notion). The *inclusion* relationship $A \subseteq B$ or $B \supseteq A$ is used to indicate that every element of A is also an element of B . In this case, we say A is a *subset* of B , or B *contains* A . To assert that $A = B$ means that $A \subseteq B$ and $B \subseteq A$. Put another way, A and B have exactly the same elements.

Quite frequently in the upcoming chapters, we will want to apply the union and intersection operations to infinite collections of sets.

Example 1.2.2. Let

$$\begin{aligned} A_1 &= \mathbf{N} = \{1, 2, 3, \dots\}, \\ A_2 &= \{2, 3, 4, \dots\}, \\ A_3 &= \{3, 4, 5, \dots\}, \end{aligned}$$

and, in general, for each $n \in \mathbf{N}$, define the set

$$A_n = \{n, n + 1, n + 2, \dots\}.$$

The result is a nested chain of sets

$$A_1 \supseteq A_2 \supseteq A_3 \supseteq A_4 \supseteq \dots,$$

where each successive set is a subset of all the previous ones. Notationally,

$$\bigcup_{n=1}^{\infty} A_n, \quad \bigcup_{n \in \mathbf{N}} A_n, \quad \text{or} \quad A_1 \cup A_2 \cup A_3 \cup \dots$$

are all equivalent ways to indicate the set whose elements consist of any element that appears in at least one particular A_n . Because of the nested property of this particular collection of sets, it is not too hard to see that

$$\bigcup_{n=1}^{\infty} A_n = A_1.$$

The notion of intersection has the same kind of natural extension to infinite collections of sets. For this example, we have

$$\bigcap_{n=1}^{\infty} A_n = \emptyset.$$

Let's be sure we understand why this is the case. Suppose we had some natural number m that we thought might actually satisfy $m \in \bigcap_{n=1}^{\infty} A_n$. What this would mean is that $m \in A_n$ for *every* A_n in our collection of sets. Because m is not an element of A_{m+1} , no such m exists and the intersection is empty.

As mentioned, most of the sets we encounter will be sets of real numbers. Given $A \subseteq \mathbf{R}$, the *complement* of A , written A^c , refers to the set of all elements of \mathbf{R} not in A . Thus, for $A \subseteq \mathbf{R}$,

$$A^c = \{x \in \mathbf{R} : x \notin A\}.$$

A few times in our work to come, we will refer to De Morgan's Laws, which state that

$$(A \cap B)^c = A^c \cup B^c \quad \text{and} \quad (A \cup B)^c = A^c \cap B^c.$$

Proofs of these statements are discussed in Exercise 1.2.3.

Admittedly, there is something imprecise about the definition of set presented at the beginning of this discussion. The defining sentence begins with the phrase “Intuitively speaking,” which might seem an odd way to embark on a course of study that purportedly intends to supply a rigorous foundation for the theory of functions of a real variable. In some sense, however, this is unavoidable. Each repair of one level of the foundation reveals something below it in need of attention. The theory of sets has been subjected to intense scrutiny over the past century precisely because so much of modern mathematics rests on this foundation. But such a study is really only advisable once it is understood why our naive impression about the behavior of sets is insufficient. For the direction in which we are heading, this will not happen, although an indication of some potential pitfalls is given in Section 1.6.

Functions

Definition 1.2.3. Given two sets A and B , a *function* from A to B is a rule or mapping that takes each element $x \in A$ and associates with it a single element of B . In this case, we write $f : A \rightarrow B$. Given an element $x \in A$, the expression $f(x)$ is used to represent the element of B associated with x by f . The set A is called the *domain* of f . The *range* of f is not necessarily equal to B but refers to the subset of B given by $\{y \in B : y = f(x) \text{ for some } x \in A\}$.

This definition of function is more or less the one proposed by Peter Lejeune Dirichlet (1805–1859) in the 1830s. Dirichlet was a German mathematician who was one of the leaders in the development of the rigorous approach to functions that we are about to undertake. His main motivation was to unravel the issues surrounding the convergence of Fourier series. Dirichlet’s contributions figure prominently in Section 8.3, where an introduction to Fourier series is presented, but we will also encounter his name in several earlier chapters along the way. What is important at the moment is that we see how Dirichlet’s definition of function liberates the term from its interpretation as a type of “formula.” In the years leading up to Dirichlet’s time, the term “function” was generally understood to refer to algebraic entities such as $f(x) = x^2 + 1$ or $g(x) = \sqrt{x^4 + 4}$. Definition 1.2.3 allows for a much broader range of possibilities.

Example 1.2.4. In 1829, Dirichlet proposed the unruly function

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

The domain of g is all of \mathbf{R} , and the range is the set $\{0, 1\}$. There is no single formula for g in the usual sense, and it is quite difficult to graph this function (see Section 4.1 for a rough attempt), but it certainly qualifies as a function according to the criterion in Definition 1.2.3. As we study the theoretical nature of continuous, differentiable, or integrable functions, examples such as this one will provide us with an invaluable testing ground for the many conjectures we encounter.

Example 1.2.5 (Triangle Inequality). The *absolute value function* is so important that it merits the special notation $|x|$ in place of the usual $f(x)$ or $g(x)$. It is defined for every real number via the piecewise definition

$$|x| = \begin{cases} x & \text{if } x \geq 0 \\ -x & \text{if } x < 0. \end{cases}$$

With respect to multiplication and division, the absolute value function satisfies

$$(i) \quad |ab| = |a||b| \text{ and}$$

$$(ii) \quad |a + b| \leq |a| + |b|$$

for all choices of a and b . Verifying these properties (Exercise 1.2.4) is just a matter of examining the different cases that arise when a , b , and $a+b$ are positive and negative. Property (ii) is called the *triangle inequality*. This innocuous looking inequality turns out to be fantastically important and will be frequently employed in the following way. Given three real numbers a , b , and c , we certainly have

$$|a - b| = |(a - c) + (c - b)|.$$

By the triangle inequality,

$$|(a - c) + (c - b)| \leq |a - c| + |c - b|,$$

so we get

$$(1) \quad |a - b| \leq |a - c| + |c - b|.$$

Now, the expression $|a - b|$ is equal to $|b - a|$ and is best understood as the *distance* between the points a and b on the number line. With this interpretation, equation (1) makes the plausible statement that “the distance from a to b is less than or equal to the distance from a to c plus the distance from c to b .” Pretending for a moment that these are points in the plane (instead of on the real line), it should be evident why this is referred to as the “triangle inequality.”

Logic and Proofs

Writing rigorous mathematical proofs is a skill best learned by doing, and there is plenty of on-the-job training just ahead. As Hardy indicates, there is an artistic quality to mathematics of this type, which may or may not come easily, but that is not to say that anything especially mysterious is happening. A proof is an essay of sorts. It is a set of carefully crafted directions, which, when followed, should leave the reader absolutely convinced of the truth of the proposition in question. To achieve this, the steps in a proof must follow logically from previous steps or be justified by some other agreed-upon set of facts. In addition to being valid, these steps must also fit coherently together to form a cogent argument. Mathematics has a specialized vocabulary, to be sure,

but that does not exempt a good proof from being written in grammatically correct English.

The one proof we have seen at this point (to Theorem 1.1.1) uses an indirect strategy called *proof by contradiction*. This powerful technique will be employed a number of times in our upcoming work. Nevertheless, most proofs are direct. (It also bears mentioning that using an indirect proof when a direct proof is available is generally considered bad manners.) A direct proof begins from some valid statement, most often taken from the theorem's hypothesis, and then proceeds through rigorously logical deductions to a demonstration of the theorem's conclusion. As we saw in Theorem 1.1.1, an indirect proof always begins by negating what it is we would like to prove. This is not always as easy to do as it may sound. The argument then proceeds until (hopefully) a logical contradiction with some other accepted fact is uncovered. Many times, this accepted fact is part of the hypothesis of the theorem. When the contradiction is with the theorem's hypothesis, we technically have what is called a *contrapositive* proof.

The next proposition illustrates a number of the issues just discussed and introduces a few more.

Theorem 1.2.6. *Two real numbers a and b are equal if and only if for every real number $\epsilon > 0$ it follows that $|a - b| < \epsilon$.*

Proof. There are two key phrases in the statement of this proposition that warrant special attention. One is “for every,” which will be addressed in a moment. The other is “if and only if.” To say “if and only if” in mathematics is an economical way of stating that the proposition is true in two directions. In the forward direction, we must prove the statement:

(\Rightarrow) *If $a = b$, then for every real number $\epsilon > 0$ it follows that $|a - b| < \epsilon$.*

We must also prove the converse statement:

(\Leftarrow) *If for every real number $\epsilon > 0$ it follows that $|a - b| < \epsilon$, then we must have $a = b$.*

For the proof of the first statement, there is really not much to say. If $a = b$, then $|a - b| = 0$, and so certainly $|a - b| < \epsilon$ no matter what $\epsilon > 0$ is chosen.

For the second statement, we give a proof by contradiction. The conclusion of the proposition in this direction states that $a = b$, so we assume that $a \neq b$. Heading off in search of a contradiction brings us to a consideration of the phrase “for every $\epsilon > 0$.” Some equivalent ways to state the hypothesis would be to say that “for all possible choices of $\epsilon > 0$ ” or “no matter how $\epsilon > 0$ is selected, it is always the case that $|a - b| < \epsilon$.” But assuming $a \neq b$ (as we are doing at the moment), the choice of

$$\epsilon_0 = |a - b| > 0$$

poses a serious problem. We are assuming that $|a - b| < \epsilon$ is true for *every* $\epsilon > 0$, so this must certainly be true of the particular ϵ_0 just defined. However, the statements

$$|a - b| < \epsilon_0 \quad \text{and} \quad |a - b| = \epsilon_0$$

cannot both be true. This contradiction means that our initial assumption that $a \neq b$ is unacceptable. Therefore, $a = b$, and the indirect proof is complete. \square

One of the most fundamental skills required for reading and writing analysis proofs is the ability to confidently manipulate the quantifying phrases “for all” and “there exists.” Significantly more attention will be given to this issue in many upcoming discussions.

Induction

One final trick of the trade, which will arise with some frequency, is the use of *induction* arguments. Induction is used in conjunction with the natural numbers \mathbf{N} (or sometimes with the set $\mathbf{N} \cup \{0\}$). The fundamental principle behind induction is that if S is some subset of \mathbf{N} with the property that

- (i) S contains 1 and
- (ii) whenever S contains a natural number n , it also contains $n + 1$,

then it must be that $S = \mathbf{N}$. As the next example illustrates, this principle can be used to define sequences of objects as well as to prove facts about them.

Example 1.2.7. Let $x_1 = 1$, and for each $n \in \mathbf{N}$ define

$$x_{n+1} = (1/2)x_n + 1.$$

Using this rule, we can compute $x_2 = (1/2)(1) + 1 = 3/2$, $x_3 = 7/4$, and it is immediately apparent how this leads to a definition of x_n for all $n \in \mathbf{N}$.

The sequence just defined appears at the outset to be increasing. For the terms computed, we have $x_1 \leq x_2 \leq x_3$. Let's use induction to prove that this trend continues; that is, let's show

$$(2) \quad x_n \leq x_{n+1}$$

for all values of $n \in \mathbf{N}$.

For $n = 1$, $x_1 = 1$ and $x_2 = 3/2$, so that $x_1 \leq x_2$ is clear. Now, we want to show that

$$\text{if we have } x_n \leq x_{n+1}, \text{ then it follows that } x_{n+1} \leq x_{n+2}.$$

Think of S as the set of natural numbers for which the claim in equation (2) is true. We have shown that $1 \in S$. We are now interested in showing that if $n \in S$, then $n+1 \in S$ as well. Starting from the induction hypothesis $x_n \leq x_{n+1}$, we can multiply across the inequality by $1/2$ and add 1 to get

$$\frac{1}{2}x_n + 1 \leq \frac{1}{2}x_{n+1} + 1,$$

which is precisely the desired conclusion $x_{n+1} \leq x_{n+2}$. By induction, the claim is proved for all $n \in \mathbf{N}$.

Any discussion about why induction is a valid argumentative technique immediately opens up a box of questions about how we understand the natural numbers. Earlier, in Section 1.1, we avoided this issue by referencing Kronecker's famous comment that the natural numbers are somehow divinely given. Although we will not improve on this explanation here, it should be pointed out that a more atheistic and mathematically satisfying approach to \mathbf{N} is possible from the point of view of axiomatic set theory. This brings us back to a recurring theme of this chapter. Pedagogically speaking, the foundations of mathematics are best learned and appreciated in a kind of reverse order. A rigorous study of the natural numbers and the theory of sets is certainly recommended, but only after we have an understanding of the subtleties of the real number system. It is this latter topic that is the business of real analysis.

Exercises

Exercise 1.2.1. (a) Prove that $\sqrt{3}$ is irrational. Does a similar argument work to show $\sqrt{6}$ is irrational?

(b) Where does the proof of Theorem 1.1.1 break down if we try to use it to prove $\sqrt{4}$ is irrational?

Exercise 1.2.2. Decide which of the following represent true statements about the nature of sets. For any that are false, provide a specific example where the statement in question does not hold.

(a) If $A_1 \supseteq A_2 \supseteq A_3 \supseteq A_4 \cdots$ are all sets containing an infinite number of elements, then the intersection $\bigcap_{n=1}^{\infty} A_n$ is infinite as well.

(b) If $A_1 \supseteq A_2 \supseteq A_3 \supseteq A_4 \cdots$ are all finite, nonempty sets of real numbers, then the intersection $\bigcap_{n=1}^{\infty} A_n$ is finite and nonempty.

(c) $A \cap (B \cup C) = (A \cap B) \cup C$.

(d) $A \cap (B \cap C) = (A \cap B) \cap C$.

(e) $A \cap (B \cup C) = (A \cap B) \cup (A \cap C)$.

Exercise 1.2.3 (De Morgan's Laws). Let A and B be subsets of \mathbf{R} .

(a) If $x \in (A \cap B)^c$, explain why $x \in A^c \cup B^c$. This shows that $(A \cap B)^c \subseteq A^c \cup B^c$.

(b) Prove the reverse inclusion $(A \cap B)^c \supseteq A^c \cup B^c$, and conclude that $(A \cap B)^c = A^c \cup B^c$.

(c) Show $(A \cup B)^c = A^c \cap B^c$ by demonstrating inclusion both ways.

Exercise 1.2.4. Verify the triangle inequality in the special cases where

(a) a and b have the same sign;

(b) $a \geq 0$, $b < 0$, and $a + b \geq 0$.

Exercise 1.2.5. Use the triangle inequality to establish the inequalities

(a) $|a - b| \leq |a| + |b|$;

(b) $||a| - |b|| \leq |a - b|$.

Exercise 1.2.6. Given a function f and a subset A of its domain, let $f(A)$ represent the range of f over the set A ; that is, $f(A) = \{f(x) : x \in A\}$.

(a) Let $f(x) = x^2$. If $A = [0, 2]$ (the closed interval $\{x \in \mathbf{R} : 0 \leq x \leq 2\}$) and $B = [1, 4]$, find $f(A)$ and $f(B)$. Does $f(A \cap B) = f(A) \cap f(B)$ in this case? Does $f(A \cup B) = f(A) \cup f(B)$?

(b) Find two sets A and B for which $f(A \cap B) \neq f(A) \cap f(B)$.

(c) Show that, for an arbitrary function $g : \mathbf{R} \rightarrow \mathbf{R}$, it is always true that $g(A \cap B) \subseteq g(A) \cap g(B)$ for all sets $A, B \subseteq \mathbf{R}$.

(d) Form and prove a conjecture about the relationship between $g(A \cup B)$ and $g(A) \cup g(B)$ for an arbitrary function g .

Exercise 1.2.7. Given a function $f : D \rightarrow \mathbf{R}$ and a subset $B \subseteq \mathbf{R}$, let $f^{-1}(B)$ be the set of all points from the domain D that get mapped into B ; that is, $f^{-1}(B) = \{x \in D : f(x) \in B\}$. This set is called the *preimage* of B .

(a) Let $f(x) = x^2$. If A is the closed interval $[0, 4]$ and B is the closed interval $[-1, 1]$, find $f^{-1}(A)$ and $f^{-1}(B)$. Does $f^{-1}(A \cap B) = f^{-1}(A) \cap f^{-1}(B)$ in this case? Does $f^{-1}(A \cup B) = f^{-1}(A) \cup f^{-1}(B)$?

(b) The good behavior of preimages demonstrated in (a) is completely general. Show that for an arbitrary function $g : \mathbf{R} \rightarrow \mathbf{R}$, it is always true that $g^{-1}(A \cap B) = g^{-1}(A) \cap g^{-1}(B)$ and $g^{-1}(A \cup B) = g^{-1}(A) \cup g^{-1}(B)$ for all sets $A, B \subseteq \mathbf{R}$.

Exercise 1.2.8. Form the logical negation of each claim. One way to do this is to simply add “It is not the case that...” in front of each assertion, but for each statement, try to embed the word “not” as deeply into the resulting sentence as possible (or avoid using it altogether).

(a) For all real numbers satisfying $a < b$, there exists an $n \in \mathbf{N}$ such that $a + 1/n < b$.

(b) Between every two distinct real numbers, there is a rational number.

(c) For all natural numbers $n \in \mathbf{N}$, \sqrt{n} is either a natural number or an irrational number.

(d) Given any real number $x \in \mathbf{R}$, there exists $n \in \mathbf{N}$ satisfying $n > x$.

Exercise 1.2.9. Show that the sequence (x_1, x_2, x_3, \dots) defined in Example 1.2.7 is bounded above by 2; that is, prove that $x_n \leq 2$ for every $n \in \mathbf{N}$.

Exercise 1.2.10. Let $y_1 = 1$, and for each $n \in \mathbf{N}$ define $y_{n+1} = (3y_n + 4)/4$.

(a) Use induction to prove that the sequence satisfies $y_n < 4$ for all $n \in \mathbf{N}$.

(b) Use another induction argument to show the sequence (y_1, y_2, y_3, \dots) is increasing.

Exercise 1.2.11. If a set A contains n elements, prove that the number of different subsets of A is equal to 2^n . (Keep in mind that the empty set \emptyset is considered to be a subset of every set.)

Exercise 1.2.12. For this exercise, assume Exercise 1.2.3 has been successfully completed.

(a) Show how induction can be used to conclude that

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

for any finite $n \in \mathbf{N}$.

(b) Explain why induction *cannot* be used to conclude

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

It might be useful to consider part (a) of Exercise 1.2.2.

(c) Is the statement in part (b) valid? If so, write a proof that does not use induction.

1.3 The Axiom of Completeness

What exactly is a real number? In Section 1.1, we got as far as saying that the set \mathbf{R} of real numbers is an extension of the rational numbers \mathbf{Q} in which there are no holes or gaps. We want every length along the number line—such as $\sqrt{2}$ —to correspond to a real number and vice versa.

We are going to improve on this definition, but as we do so, it is important to keep in mind our earlier acknowledgment that whatever precise statements we formulate will necessarily rest on other unproven assumptions or undefined terms. At some point, we must draw a line and confess that this is what we have decided to accept as a reasonable place to start. Naturally, there is some debate about where this line should be drawn. One way to view the mathematics of the 19th and 20th centuries is as a stalwart attempt to move this line further and further back toward some unshakable foundation. The majority of the material covered in this book is attributable to the mathematicians working in the early and middle parts of the 1800s. Augustin Louis Cauchy (1789–1857), Bernhard Bolzano (1781–1848), Niels Henrik Abel (1802–1829), Peter Lejeune Dirichlet, Karl Weierstrass (1815–1897), and Bernhard Riemann (1826–1866) all figure prominently in the discovery of the theorems that follow. But here is the interesting point. Nearly all of this work was done using intuitive assumptions about the nature of \mathbf{R} quite similar to our own informal understanding at this point. Eventually, enough scrutiny was directed at the detailed structure of \mathbf{R} so that, in the 1870s, a handful of ways to rigorously *construct* \mathbf{R} from \mathbf{Q} were proposed.

Following this historical model, our own rigorous construction of \mathbf{R} from \mathbf{Q} is postponed until Section 8.4. By this point, the need for such a construction will be more justified and easier to appreciate. In the meantime, we have many proofs to write, so it is important to lay down, as explicitly as possible, the assumptions that we intend to make about the real numbers.

An Initial Definition for \mathbf{R}

First, \mathbf{R} is a set containing \mathbf{Q} . The operations of addition and multiplication on \mathbf{Q} extend to all of \mathbf{R} in such a way that every element of \mathbf{R} has an additive inverse and every nonzero element of \mathbf{R} has a multiplicative inverse. Echoing

the discussion in Section 1.1, we assume \mathbf{R} is a *field*, meaning that addition and multiplication of real numbers is commutative, associative, and the distributive property holds. This allows us to perform all of the standard algebraic manipulations that are second nature to us. We also assume that the familiar properties of the ordering on \mathbf{Q} extend to all of \mathbf{R} . Thus, for example, such deductions as “If $a < b$ and $c > 0$, then $ac < bc$ ” will be carried out freely without much comment. To summarize the situation in the official terminology of the subject, we assume that \mathbf{R} is an *ordered field*, which contains \mathbf{Q} as a subfield. (A rigorous definition of “ordered field” is presented in Section 8.4.)

This brings us to the final, and most distinctive, assumption about the real number system. We must find some way to clearly articulate what we mean by insisting that \mathbf{R} does not contain the gaps that permeate \mathbf{Q} . Because this is the defining difference between the rational numbers and the real numbers, we will be excessively precise about how we phrase this assumption, hereafter referred to as the *Axiom of Completeness*.

Axiom of Completeness. *Every nonempty set of real numbers that is bounded above has a least upper bound.*

Now, what exactly does this mean?

Least Upper Bounds and Greatest Lower Bounds

Let’s first state the relevant definitions, and then look at some examples.

Definition 1.3.1. A set $A \subseteq \mathbf{R}$ is *bounded above* if there exists a number $b \in \mathbf{R}$ such that $a \leq b$ for all $a \in A$. The number b is called an *upper bound* for A .

Similarly, the set A is *bounded below* if there exists a *lower bound* $l \in \mathbf{R}$ satisfying $l \leq a$ for every $a \in A$.


Definition 1.3.2. A real number s is the *least upper bound* for a set $A \subseteq \mathbf{R}$ if it meets the following two criteria:

- (i) s is an upper bound for A ;
- (ii) if b is any upper bound for A , then $s \leq b$.

The least upper bound is also frequently called the *supremum* of the set A . Although the notation $s = \text{lub } A$ is still common, we will always write $s = \sup A$ for the least upper bound.

The *greatest lower bound* or *infimum* for A is defined in a similar way (Exercise 1.3.2) and is denoted by $\inf A$ (Fig. 1.3).

Although a set can have a host of upper bounds, it can have only one *least* upper bound. If s_1 and s_2 are both least upper bounds for a set A , then by property (ii) in Definition 1.3.2 we can assert $s_1 \leq s_2$ and $s_2 \leq s_1$. The conclusion is that $s_1 = s_2$ and least upper bounds are unique.

Figure 1.3: DEFINITION OF $\sup A$ AND $\inf A$.

Example 1.3.3. Let

$$A = \left\{ \frac{1}{n} : n \in \mathbf{N} \right\} = \left\{ 1, \frac{1}{2}, \frac{1}{3}, \dots \right\}.$$

The set A is bounded above and below. Successful candidates for an upper bound include 3, 2, and $3/2$. For the least upper bound, we claim $\sup A = 1$. To argue this rigorously using Definition 1.3.2, we need to verify that properties (i) and (ii) hold. For (i), we just observe that $1 \geq 1/n$ for all choices of $n \in \mathbf{N}$. To verify (ii), we begin by assuming we are in possession of some other upper bound b . Because $1 \in A$ and b is an upper bound for A , we must have $1 \leq b$. This is precisely what property (ii) asks us to show.

Although we do not quite have the tools we need for a rigorous proof (see Theorem 1.4.2), it should be somewhat apparent that $\inf A = 0$.

An important lesson to take from Example 1.3.3 is that $\sup A$ and $\inf A$ may or may not be elements of the set A . This issue is tied to understanding the crucial difference between the maximum and the supremum (or the minimum and the infimum) of a given set.

Definition 1.3.4. A real number a_0 is a *maximum* of the set A if a_0 is an element of A and $a_0 \geq a$ for all $a \in A$. Similarly, a number a_1 is a *minimum* of A if $a_1 \in A$ and $a_1 \leq a$ for every $a \in A$.

Example 1.3.5. To belabor the point, consider the open interval

$$(0, 2) = \{x \in \mathbf{R} : 0 < x < 2\},$$

and the closed interval

$$[0, 2] = \{x \in \mathbf{R} : 0 \leq x \leq 2\}.$$

Both sets are bounded above (and below), and both have the same least upper bound, namely 2. It is *not* the case, however, that both sets have a maximum. A maximum is a specific type of upper bound that is required to be an element of the set in question, and the open interval $(0, 2)$ does not possess such an element. Thus, the supremum can exist and not be a maximum, but when a maximum exists then it is also the supremum.

Let's turn our attention back to the Axiom of Completeness. Although we can see now that not every nonempty bounded set contains a maximum, the Axiom of Completeness asserts that every such set does have a least upper bound. We are not going to prove this. An *axiom* in mathematics is an accepted assumption, to be used without proof. Preferably, an axiom should be an elementary statement about the system in question that is so fundamental that it seems to need no justification. Perhaps the Axiom of Completeness fits this description, and perhaps it does not. Before deciding, let's remind ourselves why *it is not a valid statement about \mathbf{Q}* .

Example 1.3.6. Consider again the set

$$S = \{r \in \mathbf{Q} : r^2 < 2\},$$

and pretend for the moment that our world consists only of rational numbers. The set S is certainly bounded above. Taking $b = 2$ works, as does $b = 3/2$. But notice what happens as we go in search of the *least* upper bound. (It may be useful here to know that the decimal expansion for $\sqrt{2}$ begins 1.4142... .) We might try $b = 142/100$, which is indeed an upper bound, but then we discover that $b = 1415/1000$ is an upper bound that is smaller still. Is there a smallest one?

In the rational numbers, there is not. In the real numbers, there is. Back in \mathbf{R} , the Axiom of Completeness states that we may set $\alpha = \sup S$ and be confident that such a number exists. In the next section, we will prove that $\alpha^2 = 2$. But according to Theorem 1.1.1, this implies α is not a rational number. If we are restricting our attention to only rational numbers, then α is not an allowable option for $\sup S$, and the search for a least upper bound goes on indefinitely. Whatever rational upper bound is discovered, it is always possible to find one smaller.

The tools needed to carry out the computations described in Example 1.3.6 depend on some results about how \mathbf{Q} and \mathbf{N} fit inside of \mathbf{R} . These are discussed in the next section.

We now give an equivalent and useful way of characterizing least upper bounds. Recall that Definition 1.3.2 of the supremum has two parts. Part (i) says that $\sup A$ must be an upper bound, and part (ii) states that it must be the smallest one. The following lemma offers an alternative way to restate part (ii).

Lemma 1.3.7. *Assume $s \in \mathbf{R}$ is an upper bound for a set $A \subseteq \mathbf{R}$. Then, $s = \sup A$ if and only if, for every choice of $\epsilon > 0$, there exists an element $a \in A$ satisfying $s - \epsilon < a$.*

Proof. Here is a short rephrasing of the lemma: Given that s is an upper bound, s is the least upper bound if and only if any number smaller than s is not an upper bound. Putting it this way almost qualifies as a proof, but we will expand on what exactly is being said in each direction.

(\Rightarrow) For the forward direction, we assume $s = \sup A$ and consider $s - \epsilon$, where $\epsilon > 0$ has been arbitrarily chosen. Because $s - \epsilon < s$, part (ii) of Definition 1.3.2 implies that $s - \epsilon$ is *not* an upper bound for A . If this is the case, then there must be some element $a \in A$ for which $s - \epsilon < a$ (because otherwise $s - \epsilon$ would be an upper bound). This proves the lemma in one direction.

(\Leftarrow) Conversely, assume s is an upper bound with the property that no matter how $\epsilon > 0$ is chosen, $s - \epsilon$ is no longer an upper bound for A . Notice that what this implies is that if b is any number less than s , then b is not an upper bound. (Just let $\epsilon = s - b$.) To prove that $s = \sup A$, we must verify part (ii) of Definition 1.3.2. (Read it again.) Because we have just argued that any number smaller than s cannot be an upper bound, it follows that if b is some other upper bound for A , then $b \geq s$. \square

It is certainly the case that all of our conclusions to this point about least upper bounds have analogous versions for greatest lower bounds. The Axiom of Completeness does not explicitly assert that a nonempty set bounded below has an infimum, but this is because we do not need to assume this fact as part of the axiom. Using the Axiom of Completeness, there are several ways to prove that greatest lower bounds exist for bounded sets. One such proof is explored in Exercise 1.3.3.

Exercises

Exercise 1.3.1. Let $\mathbf{Z}_5 = \{0, 1, 2, 3, 4\}$ and define addition and multiplication modulo 5. In other words, compute the integer remainder when $a + b$ and ab are divided by 5, and use this as the value for the sum and product, respectively.

(a) Show that, given any element $z \in \mathbf{Z}_5$, there exists an element y such that $z + y = 0$. The element y is called the *additive inverse* of z .

(b) Show that, given any $z \neq 0$ in \mathbf{Z}_5 , there exists an element x such that $zx = 1$. The element x is called the *multiplicative inverse* of z .

(c) The existence of additive and multiplicative inverses is part of the definition of a *field*. Investigate the set $\mathbf{Z}_4 = \{0, 1, 2, 3\}$ (where addition and multiplication are defined modulo 4) for the existence of additive and multiplicative inverses. Make a conjecture about the values of n for which additive inverses exist in \mathbf{Z}_n , and then form another conjecture about the existence of multiplicative inverses.

Exercise 1.3.2. (a) Write a formal definition in the style of Definition 1.3.2 for the *infimum* or *greatest lower bound* of a set.

(b) Now, state and prove a version of Lemma 1.3.7 for greatest lower bounds.

Exercise 1.3.3. (a) Let A be bounded below, and define $B = \{b \in \mathbf{R} : b \text{ is a lower bound for } A\}$. Show that $\sup B = \inf A$.

(b) Use (a) to explain why there is no need to assert that greatest upper bounds exist as part of the Axiom of Completeness.

(c) Propose another way to use the Axiom of Completeness to prove that sets bounded below have greatest lower bounds.

Exercise 1.3.4. Assume that A and B are nonempty, bounded above, and satisfy $B \subseteq A$. Show $\sup B \leq \sup A$.

Exercise 1.3.5. Let $A \subseteq \mathbf{R}$ be bounded above, and let $c \in \mathbf{R}$. Define the sets $c + A$ and cA by $c + A = \{c + a : a \in A\}$ and $cA = \{ca : a \in A\}$.

- Show that $\sup(c + A) = c + \sup A$.
- If $c \geq 0$, show that $\sup(cA) = c \sup A$.
- Postulate a similar type of statement for $\sup(cA)$ for the case $c < 0$.

Exercise 1.3.6. Compute, without proofs, the suprema and infima of the following sets:

- $\{n \in \mathbf{N} : n^2 < 10\}$.
- $\{n/(m+n) : m, n \in \mathbf{N}\}$.
- $\{n/(2n+1) : n \in \mathbf{N}\}$.
- $\{n/m : m, n \in \mathbf{N} \text{ with } m+n \leq 10\}$.

Exercise 1.3.7. Prove that if a is an upper bound for A , and if a is also an element of A , then it must be that $a = \sup A$.

Exercise 1.3.8. If $\sup A < \sup B$, then show that there exists an element $b \in B$ that is an upper bound for A .

Exercise 1.3.9. Without worrying about formal proofs for the moment, decide if the following statements about suprema and infima are true or false. For any that are false, supply an example where the claim in question does not appear to hold.

- A finite, nonempty set always contains its supremum.
- If $a < L$ for every element a in the set A , then $\sup A < L$.
- If A and B are sets with the property that $a < b$ for every $a \in A$ and every $b \in B$, then it follows that $\sup A < \inf B$.
- If $\sup A = s$ and $\sup B = t$, then $\sup(A + B) = s + t$. The set $A + B$ is defined as $A + B = \{a + b : a \in A \text{ and } b \in B\}$.
- If $\sup A \leq \sup B$, then there exists an element $b \in B$ that is an upper bound for A .

1.4 Consequences of Completeness

The first application of the Axiom of Completeness is a result that may look like a more natural way to mathematically express the sentiment that the real line contains no gaps.

Theorem 1.4.1 (Nested Interval Property). *For each $n \in \mathbf{N}$, assume we are given a closed interval $I_n = [a_n, b_n] = \{x \in \mathbf{R} : a_n \leq x \leq b_n\}$. Assume also that each I_n contains I_{n+1} . Then, the resulting nested sequence of closed intervals*


$$I_1 \supseteq I_2 \supseteq I_3 \supseteq I_4 \supseteq \cdots$$

has a nonempty intersection; that is, $\bigcap_{n=1}^{\infty} I_n \neq \emptyset$.

Proof. In order to show that $\bigcap_{n=1}^{\infty} I_n$ is not empty, we are going to use the Axiom of Completeness (AoC) to produce a single real number x satisfying $x \in I_n$ for every $n \in \mathbf{N}$. Now, AoC is a statement about bounded sets, and the one we want to consider is the set

$$A = \{a_n : n \in \mathbf{N}\}$$

of left-hand endpoints of the intervals.



Because the intervals are nested, we see that every b_n serves as an upper bound for A . Thus, we are justified in setting

$$x = \sup A.$$

Now, consider a particular $I_n = [a_n, b_n]$. Because x is an upper bound for A , we have $a_n \leq x$. The fact that each b_n is an upper bound for A and that x is the least upper bound implies $x \leq b_n$.

Altogether then, we have $a_n \leq x \leq b_n$, which means $x \in I_n$ for every choice of $n \in \mathbf{N}$. Hence, $x \in \bigcap_{n=1}^{\infty} I_n$, and the intersection is not empty. \square

The Density of \mathbf{Q} in \mathbf{R}

The set \mathbf{Q} is an extension of \mathbf{N} , and \mathbf{R} in turn is an extension of \mathbf{Q} . The next few results indicate how \mathbf{N} and \mathbf{Q} sit inside of \mathbf{R} .

Theorem 1.4.2 (Archimedean Property). (i) *Given any number $x \in \mathbf{R}$, there exists an $n \in \mathbf{N}$ satisfying $n > x$.*

(ii) *Given any real number $y > 0$, there exists an $n \in \mathbf{N}$ satisfying $1/n < y$.*

Proof. Part (i) of the proposition states that \mathbf{N} is not bounded above. There has never been any doubt about the truth of this, and it could be reasonably argued that we should not have to prove it at all. This is a legitimate point of view, especially in light of the fact that we have decided to assume other familiar properties of \mathbf{N} , \mathbf{Z} , and \mathbf{Q} as given.

The counterargument is that we will prove it because we can. A set can possess the Archimedean property without being complete— \mathbf{Q} is a fine example—but a demonstration of this fact requires a good deal of scrutiny into the axiomatic construction of the ordered field in question. In the case of \mathbf{R} , the Axiom of Completeness furnishes us with a very short argument. A large number of deep results ultimately depend on this relationship between \mathbf{R} and \mathbf{N} , so having a proof for it adds a little extra certainty to these upcoming arguments.

And so to the proof. Assume, for contradiction, that \mathbf{N} is bounded above. By the Axiom of Completeness (AoC), \mathbf{N} should then have a least upper bound,

and we can set $\alpha = \sup \mathbf{N}$. If we consider $\alpha - 1$, then we no longer have an upper bound (see Lemma 1.3.7), and therefore there exists an $n \in \mathbf{N}$ satisfying $\alpha - 1 < n$. But this is equivalent to $\alpha < n + 1$. Because $n + 1 \in \mathbf{N}$, we have a contradiction to the fact that α is supposed to be an upper bound for \mathbf{N} . (Notice that the contradiction here depends only on AoC and the fact that \mathbf{N} is closed under addition.)

Part (ii) follows from (i) by letting $x = 1/y$. □


This familiar property of \mathbf{N} is the key to an extremely important fact about how \mathbf{Q} fits inside of \mathbf{R} .

Theorem 1.4.3 (Density of \mathbf{Q} in \mathbf{R}). *For every two real numbers a and b with $a < b$, there exists a rational number r satisfying $a < r < b$.*

Proof. To simplify matters, let's assume $0 \leq a < b$. The case where $a < 0$ follows quickly from this one (Exercise 1.4.1). A rational number is a quotient of integers, so we must produce $m, n \in \mathbf{N}$ so that

$$(1) \quad a < \frac{m}{n} < b.$$

The first step is to choose the denominator n large enough so that consecutive increments of size $1/n$ are too close together to “step over” the interval (a, b) .



Using Theorem 1.4.2, we may pick $n \in \mathbf{N}$ large enough so that

$$(2) \quad \frac{1}{n} < b - a.$$

Multiplying inequality (1) by n gives $na < m < nb$. With n already chosen, the idea now is to choose m to be the smallest natural number greater than na . In other words, pick $m \in \mathbf{N}$ so that

$$m - 1 \stackrel{(3)}{\leq} na \stackrel{(4)}{<} m.$$

Now, inequality (4) immediately yields $a < m/n$, which is half of the battle. Keeping in mind that inequality (2) is equivalent to $a < b - 1/n$, we can use (3) to write

$$\begin{aligned} m &\leq na + 1 \\ &< n \left(b - \frac{1}{n} \right) + 1 \\ &= nb. \end{aligned}$$

Because $m < nb$ implies $m/n < b$, we have $a < m/n < b$, as desired. □

Theorem 1.4.3 is paraphrased by saying that \mathbf{Q} is *dense* in \mathbf{R} . Without working too hard, we can use this result to show that the irrational numbers are dense in \mathbf{R} as well.

Corollary 1.4.4. *Given any two real numbers $a < b$, there exists an irrational number t satisfying $a < t < b$.*

Proof. Exercise 1.4.3. □

The Existence of Square Roots

It is time to tend to some unfinished business left over from Example 1.3.6 and this chapter's opening discussion.

Theorem 1.4.5. *There exists a real number $\alpha \in \mathbf{R}$ satisfying $\alpha^2 = 2$.*

Proof. After reviewing Example 1.3.6, consider the set

$$T = \{t \in \mathbf{R} : t^2 < 2\}$$

and set $\alpha = \sup T$. We are going to prove $\alpha^2 = 2$ by ruling out the possibilities $\alpha^2 < 2$ and $\alpha^2 > 2$. Keep in mind that there are two parts to the definition of $\sup T$, and they will both be important. (This always happens when a supremum is used in an argument.) The strategy is to demonstrate that $\alpha^2 < 2$ violates the fact that α is an upper bound for T , and $\alpha^2 > 2$ violates the fact that it is the least upper bound.

Let's first see what happens if we assume $\alpha^2 < 2$. In search of an element of T that is larger than α , write

$$\begin{aligned} \left(\alpha + \frac{1}{n}\right)^2 &= \alpha^2 + \frac{2\alpha}{n} + \frac{1}{n^2} \\ &< \alpha^2 + \frac{2\alpha}{n} + \frac{1}{n} \\ &= \alpha^2 + \frac{2\alpha + 1}{n}. \end{aligned}$$

But now assuming $\alpha^2 < 2$ gives us a little space in which to fit the $(2\alpha + 1)/n$ term and keep the total less than 2. Specifically, choose $n_0 \in \mathbf{N}$ large enough so that

$$\frac{1}{n_0} < \frac{2 - \alpha^2}{2\alpha + 1}.$$

This implies $(2\alpha + 1)/n_0 < 2 - \alpha^2$, and consequently that

$$\left(\alpha + \frac{1}{n_0}\right)^2 < \alpha^2 + (2 - \alpha^2) = 2.$$

Thus, $\alpha + 1/n_0 \in T$, contradicting the fact that α is an upper bound for T . We conclude that $\alpha^2 < 2$ cannot happen.

Now, what about the case $\alpha^2 > 2$? This time, write

$$\begin{aligned} \left(\alpha - \frac{1}{n}\right)^2 &= \alpha^2 - \frac{2\alpha}{n} + \frac{1}{n^2} \\ &> \alpha^2 - \frac{2\alpha}{n}. \end{aligned}$$

The remainder of the argument is requested in Exercise 1.4.6. \square

A small modification of this proof can be used to show that \sqrt{x} exists for any $x \geq 0$. A formula for expanding $(\alpha + 1/n)^m$ called the binomial formula can be used to show that $\sqrt[m]{x}$ exists for arbitrary values of $m \in \mathbf{N}$.

Countable and Uncountable Sets

The applications of the Axiom of Completeness to this point have basically served to restore our confidence in properties we already felt we knew about the real number system. One final consequence of completeness that we are about to present is of a very different nature and, on its own, represents an astounding intellectual discovery. The traditional way that mathematics gets done is by one mathematician modifying and expanding on the work of those who came before. This model does not seem to apply to Georg Cantor (1845–1918), at least with regard to his work on the theory of infinite sets.

At the moment, we have an image of \mathbf{R} as consisting of rational and irrational numbers, continuously packed together along the real line. We have seen that both \mathbf{Q} and \mathbf{I} (the set of irrationals) are dense in \mathbf{R} , meaning that in every interval (a, b) there exist rational and irrational numbers alike. Mentally, there is a temptation to think of \mathbf{Q} and \mathbf{I} as being intricately mixed together in equal proportions, but this turns out not to be the case. In a way that Cantor made precise, the irrational numbers far outnumber the rational numbers in making up the real line.

Cardinality

The term *cardinality* is used in mathematics to refer to the size of a set. The cardinalities of finite sets can be compared simply by attaching a natural number to each set. The set of Snow White's dwarfs is smaller than the set of United States Supreme Court Justices because 7 is less than 9. But how might we draw this same conclusion without referring to any numbers? Cantor's idea was to attempt to put the sets into a 1–1 correspondence with each other. There are fewer dwarfs than Justices because, if the dwarfs were all simultaneously appointed to the bench, there would still be two empty chairs to fill. On the other hand, the cardinality of the Supreme Court is the same as the cardinality of the set of fielders on a baseball team. This is because, when the judges take the field, it is possible to arrange them so that there is exactly one judge at every position.

The advantage of this method of comparing the sizes of sets is that it works equally well on sets that are infinite.

Definition 1.4.6. A function $f : A \rightarrow B$ is one-to-one (1-1) if $a_1 \neq a_2$ in A implies that $f(a_1) \neq f(a_2)$ in B . The function f is *onto* if, given any $b \in B$, it is possible to find an element $a \in A$ for which $f(a) = b$.

A function $f : A \rightarrow B$ that is both 1-1 and onto provides us with exactly what we mean by a 1-1 correspondence between two sets. The property of being 1-1 means that no two elements of A correspond to the same element of B (no two judges are playing the same position), and the property of being onto ensures that every element of B corresponds to something in A (there is a judge at every position).

Definition 1.4.7. Two sets A and B have the same cardinality if there exists $f : A \rightarrow B$ that is 1-1 and onto. In this case, we write $A \sim B$.

Example 1.4.8. (i) If we let $E = \{2, 4, 6, \dots\}$ be the set of even natural numbers, then we can show $\mathbf{N} \sim E$. To see why, let $f : \mathbf{N} \rightarrow E$ be given by $f(n) = 2n$.

$$\begin{array}{cccccccc} \mathbf{N} : & 1 & 2 & 3 & 4 & \cdots & n & \cdots \\ & & \downarrow & \downarrow & \downarrow & \downarrow & \cdots & \downarrow \\ E : & 2 & 4 & 6 & 8 & \cdots & 2n & \cdots \end{array}$$

It is certainly true that E is a proper subset of \mathbf{N} , and for this reason it may seem logical to say that E is a “smaller” set than \mathbf{N} . This is one way to look at it, but it represents a point of view that is heavily biased from an overexposure to finite sets. The definition of cardinality is quite specific, and from this point of view E and \mathbf{N} are equivalent.

(ii) To make this point again, note that although \mathbf{N} is contained in \mathbf{Z} as a proper subset, we can show $\mathbf{N} \sim \mathbf{Z}$. This time let

$$f(n) = \begin{cases} (n-1)/2 & \text{if } n \text{ is odd} \\ -n/2 & \text{if } n \text{ is even.} \end{cases}$$

The important details to verify are that f does not map any two natural numbers to the same element of \mathbf{Z} (f is 1-1) and that every element of \mathbf{Z} gets “hit” by something in \mathbf{N} (f is onto).

$$\begin{array}{cccccccc} \mathbf{N} : & 1 & 2 & 3 & 4 & 5 & 6 & 7 & \cdots \\ & & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \\ \mathbf{Z} : & 0 & -1 & 1 & -2 & 2 & -3 & 3 & \cdots \end{array}$$




Figure 1.4: $(-1, 1) \sim \mathbf{R}$ USING $f(x) = x/(x^2 - 1)$.

Example 1.4.9. A little calculus (which we will not supply) shows that the function $f(x) = x/(x^2 - 1)$ takes the interval $(-1, 1)$ onto \mathbf{R} in a 1-1 fashion (Fig. 1.4). Thus $(-1, 1) \sim \mathbf{R}$. In fact, $(a, b) \sim \mathbf{R}$ for any interval (a, b) .

Countable Sets

Definition 1.4.10. A set A is *countable* if $\mathbf{N} \sim A$. An infinite set that is not countable is called an *uncountable* set.

From Example 1.4.8, we see that both E and \mathbf{Z} are countable sets. Putting a set into a 1-1 correspondence with \mathbf{N} , in effect, means putting all of the elements into an infinitely long list or sequence. Looking at Example 1.4.8, we can see that this was quite easy to do for E and required only a modest bit of shuffling for the set \mathbf{Z} . A natural question arises as to whether *all* infinite sets are countable. Given some infinite set such as \mathbf{Q} or \mathbf{R} , it might seem as though, with enough cleverness, we should be able to fit all the elements of our set into a single list (i.e., into a correspondence with \mathbf{N}). After all, this list is infinitely long so there should be plenty of room. But alas, as Hardy remarks, “[The mathematician’s] subject is the most curious of all—there is none in which truth plays such odd pranks.”

Theorem 1.4.11. (i) *The set \mathbf{Q} is countable.* (ii) *The set \mathbf{R} is uncountable.*

Proof. (i) For each $n \in \mathbf{N}$, let A_n be the set given by

$$A_n = \left\{ \pm \frac{p}{q} : \text{where } p, q \in \mathbf{N} \text{ are in lowest terms with } p + q = n \right\}.$$

The first few of these sets look like

$$A_1 = \left\{ \frac{0}{1} \right\}, \quad A_2 = \left\{ \frac{1}{1}, \frac{-1}{1} \right\}, \quad A_3 = \left\{ \frac{1}{2}, \frac{-1}{2}, \frac{2}{1}, \frac{-2}{1} \right\},$$

$$A_4 = \left\{ \frac{1}{3}, \frac{-1}{3}, \frac{3}{1}, \frac{-3}{1} \right\}, \quad \text{and} \quad A_5 = \left\{ \frac{1}{4}, \frac{-1}{4}, \frac{2}{3}, \frac{-2}{3}, \frac{3}{2}, \frac{-3}{2}, \frac{4}{1}, \frac{-4}{1} \right\}.$$

The crucial observation is that each A_n is *finite* and every rational number appears in exactly one of these sets. Our 1–1 correspondence with \mathbf{N} is then achieved by consecutively listing the elements in each A_n .

$$\begin{array}{ccccccccccccccc} \mathbf{N} : & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 & 11 & 12 & \cdots \\ & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \downarrow & \\ \mathbf{Q} : & \frac{0}{1} & \frac{1}{1} & \frac{-1}{1} & \frac{1}{2} & \frac{-1}{2} & \frac{2}{1} & \frac{-2}{1} & \frac{1}{3} & \frac{-1}{3} & \frac{3}{1} & \frac{-3}{1} & \frac{1}{4} & \cdots \\ & \underbrace{\hspace{1.5cm}} & \underbrace{\hspace{1.5cm}} & & \underbrace{\hspace{2.5cm}} & & & & \underbrace{\hspace{2.5cm}} & & & & & \\ & A_1 & A_2 & & A_3 & & & & A_4 & & & & & \end{array}$$

Admittedly, writing an explicit formula for this correspondence would be an awkward task, and attempting to do so is not the best use of time. What matters is that we see why every rational number appears in the correspondence exactly once. Given, say, $22/7$, we have that $22/7 \in A_{29}$. Because the set of elements in A_1, \dots, A_{28} is finite, we can be confident that $22/7$ eventually gets included in the sequence. The fact that this line of reasoning applies to any rational number p/q is our proof that the correspondence is onto. To verify that it is 1–1, we observe that the sets A_n were constructed to be disjoint so that no rational number appears twice. This completes the proof of (i).

(ii) The second statement of Theorem 1.4.11 is the truly unexpected part, and its proof is done by contradiction. Assume that there *does* exist a 1–1, onto function $f : \mathbf{N} \rightarrow \mathbf{R}$. Again, what this suggests is that it is possible to enumerate the elements of \mathbf{R} . If we let $x_1 = f(1)$, $x_2 = f(2)$, and so on, then our assumption that f is onto means that we can write


$$(1) \quad \mathbf{R} = \{x_1, x_2, x_3, x_4, \dots\}$$

and be confident that every real number appears somewhere on the list. We will now use the Nested Interval Property (Theorem 1.4.1) to produce a real number that is not there.

Let I_1 be a closed interval that *does not* contain x_1 . Next, let I_2 be a closed interval, contained in I_1 , which does not contain x_2 . The existence of such an I_2 is easy to verify. Certainly I_1 contains two smaller *disjoint* closed intervals, and x_2 can only be in one of these. In general, given an interval I_n , construct I_{n+1} to satisfy

$$(i) \quad I_{n+1} \subseteq I_n \text{ and}$$

$$(ii) \quad x_{n+1} \notin I_{n+1}.$$



We now consider the intersection $\bigcap_{n=1}^{\infty} I_n$. If x_{n_0} is some real number from the list in (1), then we have $x_{n_0} \notin I_{n_0}$, and it follows that

$$x_{n_0} \notin \bigcap_{n=1}^{\infty} I_n.$$

Now, we are assuming that the list in (1) contains every real number, and this leads to the conclusion that

$$\bigcap_{n=1}^{\infty} I_n = \emptyset.$$

However, the Nested Interval Property (NIP) asserts that $\bigcap_{n=1}^{\infty} I_n \neq \emptyset$. By NIP, there is at least one $x \in \bigcap_{n=1}^{\infty} I_n$ that, consequently, cannot be on the list in (1). This contradiction means that such an enumeration of \mathbf{R} is impossible, and we conclude that \mathbf{R} is an *uncountable* set. \square

What exactly should we make of this discovery? It is an important exercise to show that any subset of a countable set must be either countable or finite. This should not be too surprising. If a set can be arranged into a single list, then deleting some elements from this list results in another (shorter, and potentially terminating) list. This means that countable sets are the smallest type of infinite set. Anything smaller is either still countable or finite.

The force of Theorem 1.4.11 is that the cardinality of \mathbf{R} is, informally speaking, a larger type of infinity. The real numbers so outnumber the natural numbers that there is no way to map \mathbf{N} onto \mathbf{R} . No matter how we attempt this, there are always real numbers to spare. The set \mathbf{Q} , on the other hand, is countable. As far as infinite sets are concerned, this is as small as it gets. What does this imply about the set \mathbf{I} of irrational numbers? By imitating the demonstration that $\mathbf{N} \sim \mathbf{Z}$, we can prove that the union of two countable sets must be countable. Because $\mathbf{R} = \mathbf{Q} \cup \mathbf{I}$, it follows that \mathbf{I} cannot be countable because otherwise \mathbf{R} would be. The inescapable conclusion is that, despite the fact that we have encountered so few of them, the irrational numbers form a far greater subset of \mathbf{R} than \mathbf{Q} .

The properties of countable sets described in this discussion are useful for a few exercises in upcoming chapters. For easier reference, we state them as some final propositions and outline their proofs in the exercises that follow.

Theorem 1.4.12. *If $A \subseteq B$ and B is countable, then A is either countable, finite, or empty.*

Theorem 1.4.13. (i) *If A_1, A_2, \dots, A_m are each countable sets, then the union $A_1 \cup A_2 \cup \dots \cup A_m$ is countable.*

(ii) *If A_n is a countable set for each $n \in \mathbf{N}$, then $\bigcup_{n=1}^{\infty} A_n$ is countable.*

Exercises

Exercise 1.4.1. Without doing too much work, show how to prove Theorem 1.4.3 in the case where $a < 0$ by converting this case into the one already proven.

Exercise 1.4.2. Recall that \mathbf{I} stands for the set of irrational numbers.

(a) Show that if $a, b \in \mathbf{Q}$, then ab and $a + b$ are elements of \mathbf{Q} as well.

(b) Show that if $a \in \mathbf{Q}$ and $t \in \mathbf{I}$, then $a + t \in \mathbf{I}$ and $at \in \mathbf{I}$ as long as $a \neq 0$.

(c) Part (a) can be summarized by saying that \mathbf{Q} is closed under addition and multiplication. Is \mathbf{I} closed under addition and multiplication? Given two irrational numbers s and t , what can we say about $s + t$ and st ?

Exercise 1.4.3. Using Exercise 1.4.2, supply a proof for Corollary 1.4.4 by applying Theorem 1.4.3 to the real numbers $a - \sqrt{2}$ and $b - \sqrt{2}$.

Exercise 1.4.4. Use the Archimedean Property of \mathbf{R} to rigorously prove that $\inf\{1/n : n \in \mathbf{N}\} = 0$.

Exercise 1.4.5. Prove that $\bigcap_{n=1}^{\infty} (0, 1/n) = \emptyset$. Notice that this demonstrates that the intervals in the Nested Interval Property must be closed for the conclusion of the theorem to hold.

Exercise 1.4.6. (a) Finish the proof of Theorem 1.4.5 by showing that the assumption $\alpha^2 > 2$ leads to a contradiction of the fact that $\alpha = \sup T$.

(b) Modify this argument to prove the existence of \sqrt{b} for any real number $b \geq 0$.

Exercise 1.4.7. Finish the following proof for Theorem 1.4.12.

Assume B is a countable set. Thus, there exists $f : \mathbf{N} \rightarrow B$, which is 1–1 and onto. Let $A \subseteq B$ be an infinite subset of B . We must show that A is countable.

Let $n_1 = \min\{n \in \mathbf{N} : f(n) \in A\}$. As a start to a definition of $g : \mathbf{N} \rightarrow A$, set $g(1) = f(n_1)$. Show how to inductively continue this process to produce a 1–1 function g from \mathbf{N} onto A .

Exercise 1.4.8. Use the following outline to supply proofs for the statements in Theorem 1.4.13.

(a) First, prove statement (i) for two countable sets, A_1 and A_2 . Example 1.4.8 (ii) may be a useful reference. Some technicalities can be avoided by first replacing A_2 with the set $B_2 = A_2 \setminus A_1 = \{x \in A_2 : x \notin A_1\}$. The point of this is that the union $A_1 \cup B_2$ is equal to $A_1 \cup A_2$ and the sets A_1 and B_2 are disjoint. (What happens if B_2 is finite?)

Now, explain how the more general statement in (i) follows.

(b) Explain why induction *cannot* be used to prove part (ii) of Theorem 1.4.13 from part (i).

(c) Show how arranging \mathbf{N} into the two-dimensional array

$$\begin{array}{cccccc} 1 & 3 & 6 & 10 & 15 & \cdots \\ 2 & 5 & 9 & 14 & \cdots & \\ 4 & 8 & 13 & \cdots & & \\ 7 & 12 & \cdots & & & \\ 11 & \cdots & & & & \\ \vdots & & & & & \end{array}$$

leads to a proof of Theorem 1.4.13 (ii).

Exercise 1.4.9. (a) Given sets A and B , explain why $A \sim B$ is equivalent to asserting $B \sim A$.

(b) For three sets A, B , and C , show that $A \sim B$ and $B \sim C$ implies $A \sim C$. These two properties are what is meant by saying that \sim is an *equivalence relation*.

Exercise 1.4.10. Show that the set of all finite subsets of \mathbf{N} is a countable set. (It turns out that the set of *all* subsets of \mathbf{N} is not a countable set. This is the topic of Section 1.5.)

Exercise 1.4.11. Consider the open interval $(0,1)$, and let S be the set of points in the open unit square; that is, $S = \{(x, y) : 0 < x, y < 1\}$.

(a) Find a 1–1 function that maps $(0, 1)$ into, but not necessarily onto, S . (This is easy.)

(b) Use the fact that every real number has a decimal expansion to produce a 1–1 function that maps S into $(0, 1)$. Discuss whether the formulated function is onto. (Keep in mind that any terminating decimal expansion such as .235 represents the same real number as .234999... .)

The Schröder–Bernstein Theorem discussed in Exercise 1.4.13 to follow can now be applied to conclude that $(0, 1) \sim S$.

Exercise 1.4.12. A real number $x \in \mathbf{R}$ is called *algebraic* if there exist integers $a_0, a_1, a_2, \dots, a_n \in \mathbf{Z}$, not all zero, such that

$$a_n x^n + a_{n-1} x^{n-1} + \cdots + a_1 x + a_0 = 0.$$

Said another way, a real number is algebraic if it is the root of a polynomial with integer coefficients. Real numbers that are not algebraic are called *transcendental* numbers. Reread the last paragraph of Section 1.1. The final question posed here is closely related to the question of whether or not transcendental numbers exist.

(a) Show that $\sqrt{2}$, $\sqrt[3]{2}$, and $\sqrt{3} + \sqrt{2}$ are algebraic.

(b) Fix $n \in \mathbf{N}$, and let A_n be the algebraic numbers obtained as roots of polynomials with integer coefficients that have degree n . Using the fact that every polynomial has a finite number of roots, show that A_n is countable. (For each $m \in \mathbf{N}$, consider polynomials $a_n x^n + a_{n-1} x^{n-1} + \cdots + a_1 x + a_0$ that satisfy $|a_n| + |a_{n-1}| + \cdots + |a_1| + |a_0| \leq m$.)

(c) Now, argue that the set of all algebraic numbers is countable. What may we conclude about the set of transcendental numbers?

Exercise 1.4.13 (Schröder–Bernstein Theorem). Assume there exists a 1–1 function $f : X \rightarrow Y$ and another 1–1 function $g : Y \rightarrow X$. Follow the steps to show that there exists a 1–1, onto function $h : X \rightarrow Y$ and hence $X \sim Y$.

(a) The range of f is defined by $f(X) = \{y \in Y : y = f(x) \text{ for some } x \in X\}$. Let $y \in f(X)$. (Because f is not necessarily onto, the range $f(X)$ may not be all of Y .) Explain why there exists a *unique* $x \in X$ such that $f(x) = y$. Now define $f^{-1}(y) = x$, and show that f^{-1} is a 1–1 function from $f(X)$ onto X .

In a similar way, we can also define the 1–1 function $g^{-1} : g(Y) \rightarrow Y$.

(b) Let $x \in X$ be arbitrary. Let the *chain* C_x be the set consisting of all elements of the form

$$(1) \quad \dots, f^{-1}(g^{-1}(x)), g^{-1}(x), x, f(x), g(f(x)), f(g(f(x))), \dots$$

Explain why the number of elements to the left of x in the above chain may be zero, finite, or infinite.

(c) Show that any two chains are either identical or completely disjoint.

(d) Note that the terms of the chain in (1) alternate between elements of X and elements of Y . Given a chain C_x , we want to focus on $C_x \cap Y$, which is just the part of the chain that sits in Y .

Define the set A to be the union of all chains C_x satisfying $C_x \cap Y \subseteq f(X)$. Let B consist of the union of the remaining chains not in A . Show that any chain contained in B must be of the form

$$y, g(y), f(g(y)), g(f(g(y))), \dots,$$

where y is an element of Y that is *not* in $f(X)$.

(e) Let $X_1 = A \cap X$, $X_2 = B \cap X$, $Y_1 = A \cap Y$, and $Y_2 = B \cap Y$. Show that f maps X_1 onto Y_1 and that g maps Y_2 onto X_2 . Use this information to prove $X \sim Y$.

1.5 Cantor's Theorem

Cantor's work into the theory of infinite sets extends far beyond the conclusions of Theorem 1.4.11. Although initially resisted, his creative and relentless assault in this area eventually produced a revolution in set theory and a paradigm shift in the way mathematicians came to understand the infinite.

Cantor's Diagonalization Method

The proof presented for Theorem 1.4.11 (ii) is different from any of the arguments that Cantor gave for this result. It was chosen because of how directly it reveals the connection between the concepts of uncountability and completeness, and because the technique of using nested intervals will be used several more times in our work ahead.

Cantor initially published his discovery that \mathbf{R} is uncountable in 1874, but in 1891 he offered another proof of this same fact that is startling in its simplicity. It relies on decimal representations for real numbers, which we will accept and use without any formal definitions.

Theorem 1.5.1. *The open interval $(0, 1) = \{x \in \mathbf{R} : 0 < x < 1\}$ is uncountable.*

Exercise 1.5.1. Show that $(0, 1)$ is uncountable if and only if \mathbf{R} is uncountable. This shows that Theorem 1.5.1 is equivalent to Theorem 1.4.11.

Proof. As with Theorem 1.4.11, we proceed by contradiction and assume that there does exist a function $f : \mathbf{N} \rightarrow (0, 1)$ that is 1–1 and onto. For each $m \in \mathbf{N}$, $f(m)$ is a real number between 0 and 1, and we represent it using the decimal notation

$$f(m) = .a_{m1}a_{m2}a_{m3}a_{m4}a_{m5}\dots$$

What is meant here is that for each $m, n \in \mathbf{N}$, a_{mn} is the digit from the set $\{0, 1, 2, \dots, 9\}$ that represents the n th digit in the decimal expansion of $f(m)$. The 1–1 correspondence between \mathbf{N} and $(0, 1)$ can be summarized in the doubly indexed array

\mathbf{N}	$(0, 1)$										
1	$\longleftrightarrow f(1)$	$=$	$\mathbf{.a_{11}}$	a_{12}	a_{13}	a_{14}	a_{15}	a_{16}	\dots		
2	$\longleftrightarrow f(2)$	$=$	$.a_{21}$	$\mathbf{a_{22}}$	a_{23}	a_{24}	a_{25}	a_{26}	\dots		
3	$\longleftrightarrow f(3)$	$=$	$.a_{31}$	a_{32}	$\mathbf{a_{33}}$	a_{34}	a_{35}	a_{36}	\dots		
4	$\longleftrightarrow f(4)$	$=$	$.a_{41}$	a_{42}	a_{43}	$\mathbf{a_{44}}$	a_{45}	a_{46}	\dots		
5	$\longleftrightarrow f(5)$	$=$	$.a_{51}$	a_{52}	a_{53}	a_{54}	$\mathbf{a_{55}}$	a_{56}	\dots		
6	$\longleftrightarrow f(6)$	$=$	$.a_{61}$	a_{62}	a_{63}	a_{64}	a_{65}	$\mathbf{a_{66}}$	\dots		
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\ddots		

The key assumption about this correspondence is that *every* real number in $(0, 1)$ is assumed to appear somewhere on the list.

Now for the pearl of the argument. Define a real number $x \in (0, 1)$ with the decimal expansion $x = .b_1b_2b_3b_4\dots$ using the rule

$$b_n = \begin{cases} 2 & \text{if } a_{nn} \neq 2 \\ 3 & \text{if } a_{nn} = 2. \end{cases}$$

Let's be clear about this. To compute the digit b_1 , we look at the digit a_{11} in the upper left-hand corner of the array. If $a_{11} = 2$, then we choose $b_1 = 3$; otherwise, we set $b_1 = 2$.

Exercise 1.5.2. (a) Explain why the real number $x = .b_1b_2b_3b_4\dots$ cannot be $f(1)$.

(b) Now, explain why $x \neq f(2)$, and in general why $x \neq f(n)$ for any $n \in \mathbf{N}$.

(c) Point out the contradiction that arises from these observations and conclude that $(0, 1)$ is uncountable.

□

Exercise 1.5.3. Supply rebuttals to the following complaints about the proof of Theorem 1.5.1.

(a) Every rational number has a decimal expansion so we could apply this same argument to show that the set of rational numbers between 0 and 1 is uncountable. However, because we know that any subset of \mathbf{Q} must be countable, the proof of Theorem 1.5.1 must be flawed.

(b) A few numbers have *two* different decimal representations. Specifically, any decimal expansion that terminates can also be written with repeating 9's. For instance, $1/2$ can be written as $.5$ or as $.4999\dots$. Doesn't this cause some problems?

Exercise 1.5.4. Let S be the set consisting of all sequences of 0's and 1's. Observe that S is not a particular sequence, but rather a large set whose elements are sequences; namely,

$$S = \{(a_1, a_2, a_3, \dots) : a_n = 0 \text{ or } 1\}.$$

As an example, the sequence $(1, 0, 1, 0, 1, 0, 1, 0, \dots)$ is an element of S , as is the sequence $(1, 1, 1, 1, 1, 1, \dots)$.

Give a rigorous argument showing that S is uncountable.

Having distinguished between the countable infinity of \mathbf{N} and the uncountable infinity of \mathbf{R} , a new question that occupied Cantor was whether or not there existed an infinity "above" that of \mathbf{R} . This is logically treacherous territory. The same care we gave to defining the relationship "has the same cardinality as" needs to be given to defining relationships such as "has cardinality greater than" or "has cardinality less than or equal to." Nevertheless, without getting too weighed down with formal definitions, one gets a very clear sense from the next result that there is a hierarchy of infinite sets that continues well beyond the continuum of \mathbf{R} .

Power Sets and Cantor's Theorem

Given a set A , the *power set* $P(A)$ refers to the collection of all subsets of A . It is important to understand that $P(A)$ is itself considered a set whose elements are the different possible subsets of A .

Exercise 1.5.5. (a) Let $A = \{a, b, c\}$. List the eight elements of $P(A)$. (Do not forget that \emptyset is considered to be a subset of every set.)

(b) If A is finite with n elements, show that $P(A)$ has 2^n elements. (Constructing a particular subset of A can be interpreted as making a series of decisions about whether or not to include each element of A .)

Exercise 1.5.6. (a) Using the particular set $A = \{a, b, c\}$, exhibit two different 1–1 mappings from A into $P(A)$.

(b) Letting $B = \{1, 2, 3, 4\}$, produce an example of a 1–1 map $g : B \rightarrow P(B)$.

(c) Explain why, in parts (a) and (b), it is impossible to construct mappings that are *onto*.

Cantor’s Theorem states that the phenomenon in Exercise 1.5.6 holds for infinite sets as well as finite sets. Whereas mapping A into $P(A)$ is quite effortless, finding an *onto* map is impossible.

Theorem 1.5.2 (Cantor’s Theorem). *Given any set A , there does not exist a function $f : A \rightarrow P(A)$ that is onto.*

Proof. This proof, like the others of its kind, is indirect. Thus, assume, for contradiction, that $f : A \rightarrow P(A)$ is onto. Unlike the usual situation in which we have sets of numbers for the domain and range, f is a correspondence between a set and its power set. For each element $a \in A$, $f(a)$ is a particular *subset* of A . The assumption that f is onto means that every subset of A appears as $f(a)$ for some $a \in A$. To arrive at a contradiction, we will produce a subset $B \subseteq A$ that is not equal to $f(a)$ for any $a \in A$.

Construct B using the following rule. For each element $a \in A$, consider the subset $f(a)$. This subset of A may contain the element a or it may not. This depends on the function f . If $f(a)$ does not contain a , then we include a in our set B . More precisely, let

$$B = \{a \in A : a \notin f(a)\}.$$

Exercise 1.5.7. Return to the particular functions constructed in Exercise 1.5.6 and construct the subset B that results using the preceding rule. In each case, note that B is not in the range of the function used.

We now focus on the general argument. Because we have assumed that our function $f : A \rightarrow P(A)$ is onto, it must be that $B = f(a')$ for some $a' \in A$. The contradiction arises when we consider whether or not a' is an element of B .

Exercise 1.5.8. (a) First, show that the case $a' \in B$ leads to a contradiction.

(b) Now, finish the argument by showing that the case $a' \notin B$ is equally unacceptable.

□

Exercise 1.5.9. As a final exercise, answer each of the following by establishing a 1–1 correspondence with a set of known cardinality.

(a) Is the set of all functions from $\{0, 1\}$ to \mathbf{N} countable or uncountable?

(b) Is the set of all functions from \mathbf{N} to $\{0, 1\}$ countable or uncountable?

(c) Given a set B , a subset \mathcal{A} of $P(B)$ is called an *antichain* if no element of \mathcal{A} is a subset of any other element of \mathcal{A} . Does $P(\mathbf{N})$ contain an uncountable antichain?

1.6 Epilogue

The relationship of having the same cardinality is an *equivalence relation* (see Exercise 1.4.9), meaning, roughly, that all of the sets in the universe can be organized into disjoint groups according to their size. Two sets appear in the same group, or *equivalence class*, if and only if they have the same cardinality. Thus, \mathbf{N} , \mathbf{Z} , and \mathbf{Q} are grouped together in one class with all of the other countable sets, whereas \mathbf{R} is in another class that includes the interval $(0, 1)$ among other uncountable sets. One implication of Cantor's Theorem is that $P(\mathbf{R})$ —the set of all subsets of \mathbf{R} —is in a different class from \mathbf{R} , and there is no reason to stop here. The set of subsets of $P(\mathbf{R})$ —namely $P(P(\mathbf{R}))$ —is in yet another class, and this process continues indefinitely.

Having divided the universe of sets into disjoint groups, it would be convenient to attach a “number” to each collection which could be used the way natural numbers are used to refer to the sizes of finite sets. Given a set X , there exists something called the *cardinal number* of X , denoted $\text{card } X$, which behaves very much in this fashion. For instance, two sets X and Y satisfy $\text{card } X = \text{card } Y$ if and only if $X \sim Y$. (Rigorously defining $\text{card } X$ requires some significant set theory. One way this is done is to define $\text{card } X$ to be a very particular set that can always be uniquely found in the same equivalence class as X .)

Looking back at Cantor's Theorem, we get the strong sense that there is an *order* on the sizes of infinite sets that should be reflected in our new cardinal number system. Specifically, if it is possible to map a set X into Y in a 1–1 fashion, then we want $\text{card } X \leq \text{card } Y$. Writing the strict inequality $\text{card } X < \text{card } Y$ should indicate that it is possible to map X into Y but that it is impossible to show $X \sim Y$. Restated in this notation, Cantor's Theorem states that for every set A , $\text{card } A < \text{card } P(A)$.

There are some significant details to work out. A kind of metaphysical problem arises when we realize that an implication of Cantor's Theorem is that there can be no “largest” set. A declaration such as, “Let U be the set of all possible things,” is paradoxical because we immediately get that $\text{card } U < \text{card } P(U)$ and thus the set U does not contain everything it was advertised to hold. Issues such as this one are ultimately resolved by imposing some restrictions on what can qualify as a set. As set theory was formalized, the axioms had to be crafted so that objects such as U are simply not allowed. A more down-to-earth problem in need of attention is demonstrating that our definition of “ \leq ” between cardinal numbers really is an ordering. This involves showing that cardinal numbers possess a property analogous to real numbers, which states that if $\text{card } X \leq \text{card } Y$ and $\text{card } Y \leq \text{card } X$, then $\text{card } X = \text{card } Y$. In the end, this boils down to proving that if there exists $f : X \rightarrow Y$ that is 1–1, and if there exists $g : Y \rightarrow X$ that is 1–1, then it is possible to find a function $h : X \rightarrow Y$ that is both 1–1 and onto. A proof of this fact eluded Cantor but was eventually supplied independently by Ernst Schröder (in 1896) and Felix Bernstein (in 1898). An argument for the Schröder–Bernstein Theorem is outlined in Exercise 1.4.13.

There was another deep problem stemming from the budding theory of cardinal numbers that occupied Cantor and which was not resolved during his lifetime. Because of the importance of countable sets, the symbol \aleph_0 (“aleph naught”) is frequently used for $\text{card } \mathbf{N}$. The subscript “0” is appropriate when we remember that countable sets are the smallest type of infinite set. In terms of cardinal numbers, if $\text{card } X < \aleph_0$, then X is finite. Thus, \aleph_0 is the smallest infinite cardinal number. The cardinality of \mathbf{R} is also significant enough to deserve the special designation $\mathfrak{c} = \text{card } \mathbf{R} = \text{card}(0, 1)$. The content of Theorems 1.4.11 and 1.5.1 is that $\aleph_0 < \mathfrak{c}$. The question that plagued Cantor was whether there were any cardinal numbers strictly in between these two. Put another way, does there exist a set $A \subseteq \mathbf{R}$ with $\text{card } \mathbf{N} < \text{card } A < \text{card } \mathbf{R}$? Cantor was of the opinion that no such set existed. In the ordering of cardinal numbers, he conjectured, \mathfrak{c} was the immediate successor of \aleph_0 .

Cantor’s “continuum hypothesis,” as it came to be called, was one of the most famous mathematical challenges of the past century. Its unexpected resolution came in two parts. In 1940, the German logician and mathematician Kurt Gödel demonstrated that, using only the agreed-upon set of axioms of set theory, there was no way to disprove the continuum hypothesis. In 1963, Paul Cohen successfully showed that, under the same rules, it was also impossible to prove this conjecture. Taken together, what these two discoveries imply is that the continuum hypothesis is undecidable. It can be accepted or rejected as a statement about the nature of infinite sets, and in neither case will any logical contradictions arise.

The mention of Kurt Gödel brings to mind a final comment about the significance of Cantor’s work. Gödel is best known for his “Incompleteness Theorems,” which pertain to the strength of axiomatic systems in general. What Gödel showed was that any consistent axiomatic system created to study arithmetic was necessarily destined to be “incomplete” in the sense that there would always be true statements that the system of axioms would be too weak to prove. At the heart of Gödel’s very complicated proof is a type of manipulation closely related to what is happening in the proofs of Theorems 1.5.1 and 1.5.2. Variations of Cantor’s proof methods can also be found in the limitative results of computer science. The “halting problem” asks, loosely, whether some general algorithm exists that can look at every program and decide if that program eventually terminates. The proof that no such algorithm exists uses a diagonalization-type construction at the core of the argument. The main point to make is that not only are the implications of Cantor’s theorems profound but the argumentative techniques are as well. As a more immediate example of this phenomenon, the diagonalization method is used again in Chapter 6—in a constructive way—as a crucial step in the proof of the Arzela–Ascoli Theorem.

Chapter 2


Sequences and Series

2.1 Discussion: Rearrangements of Infinite Series

Consider the infinite series

$$\sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{n} = 1 - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \frac{1}{5} - \frac{1}{6} + \frac{1}{7} - \frac{1}{8} + \dots$$

If we naively begin adding from the left-hand side, we get a sequence of what are called *partial sums*. In other words, let s_n equal the sum of the first n terms of the series, so that $s_1 = 1$, $s_2 = 1/2$, $s_3 = 5/6$, $s_4 = 7/12$, and so on. One immediate observation is that the successive sums oscillate in a progressively narrower space. The odd sums decrease ($s_1 > s_3 > s_5 > \dots$) while the even sums increase ($s_2 < s_4 < s_6 < \dots$).



$$s_2 < s_4 < s_6 < \dots < S \dots < s_5 < s_3 < s_1$$

It seems reasonable—and we will soon prove—that the sequence (s_n) eventually hones in on a value, call it S , where the odd and even partial sums “meet”. At this moment, we cannot compute S precisely, but we know it falls somewhere between $7/12$ and $5/6$. Summing a few hundred terms reveals that $S \approx .69$. Whatever its value, there is now an overwhelming temptation to write

$$(1) \quad S = 1 - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \frac{1}{5} - \frac{1}{6} + \frac{1}{7} - \frac{1}{8} + \dots$$

meaning, perhaps, that if we could indeed add up *all* infinitely many of these numbers, then the sum would *equal* S . A more familiar example of an equation

of this type might be

$$2 = 1 + \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \frac{1}{16} + \frac{1}{32} + \frac{1}{64} + \cdots,$$

the only difference being that in the second equation we have a more recognizable value for the sum.

But now for the crux of the matter. The symbols $+$, $-$, and $=$ in the preceding equations are deceptively familiar notions being used in a very unfamiliar way. The crucial question is whether or not properties of addition and equality that are well understood for finite sums remain valid when applied to infinite objects such as equation (1). The answer, as we are about to witness, is somewhat ambiguous.

Treating equation (1) in a standard algebraic way, let's multiply through by $1/2$ and add it back to equation (1):

$$\begin{aligned} \frac{1}{2}S &= \quad \frac{1}{2} \quad -\frac{1}{4} \quad +\frac{1}{6} \quad -\frac{1}{8} \quad +\frac{1}{10} \quad -\frac{1}{12} \quad +\cdots \\ + S &= 1 - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \frac{1}{5} - \frac{1}{6} + \frac{1}{7} - \frac{1}{8} + \frac{1}{9} - \frac{1}{10} + \frac{1}{11} - \frac{1}{12} + \frac{1}{13} - \cdots \\ \hline (2) \quad \frac{3}{2}S &= 1 \quad +\frac{1}{3} - \frac{1}{2} + \frac{1}{5} \quad +\frac{1}{7} - \frac{1}{4} + \frac{1}{9} \quad +\frac{1}{11} - \frac{1}{6} + \frac{1}{13} \quad \cdots \end{aligned}$$

Now, look carefully at the result. The sum in equation (2) consists *precisely* of the same terms as those in the original equation (1), only in a different order. Specifically, the series in (2) is a rearrangement of (1) where we list the first two positive terms ($1 + \frac{1}{3}$) followed by the first negative term ($-\frac{1}{2}$), followed by the next two positive terms ($\frac{1}{5} + \frac{1}{7}$) and then the next negative term ($-\frac{1}{4}$). Continuing this, it is apparent that every term in (2) appears in (1) and vice versa. The rub comes when we realize that equation (2) asserts that the sum of these rearranged, but otherwise unaltered, numbers is equal to $3/2$ its original value. Indeed, adding a few hundred terms of equation (2) produces partial sums in the neighborhood of 1.03. Addition, in this infinite setting, is not commutative!

Let's look at a similar rearrangement of the series

$$\sum_{n=0}^{\infty} (-1/2)^n.$$

This series is geometric with first term 1 and common ratio $r = -1/2$. Using the formula $1/(1-r)$ for the sum of a geometric series (Example 2.7.5), we get

$$1 - \frac{1}{2} + \frac{1}{4} - \frac{1}{8} + \frac{1}{16} - \frac{1}{32} + \frac{1}{64} - \frac{1}{128} + \frac{1}{256} \cdots = \frac{1}{1 - (-\frac{1}{2})} = \frac{2}{3}.$$

This time, some computational experimentation with the "two positives, one negative" rearrangement

$$1 + \frac{1}{4} - \frac{1}{2} + \frac{1}{16} + \frac{1}{64} - \frac{1}{8} + \frac{1}{256} + \frac{1}{1024} - \frac{1}{32} \cdots$$

yields partial sums quite close to $2/3$. The sum of the first 30 terms, for instance, equals .666667. Infinite addition is commutative in some instances but not in others.

Far from being a charming theoretical oddity of infinite series, this phenomenon can be the source of great consternation in many applied situations. How, for instance, should a double summation over two index variables be defined? Let's say we are given a *grid* of real numbers $\{a_{ij} : i, j \in \mathbf{N}\}$, where $a_{ij} = 1/2^{j-i}$ if $j > i$, $a_{ij} = -1$ if $j = i$, and $a_{ij} = 0$ if $j < i$.

$$\begin{bmatrix} -1 & \frac{1}{2} & \frac{1}{4} & \frac{1}{8} & \frac{1}{16} & \cdots \\ 0 & -1 & \frac{1}{2} & \frac{1}{4} & \frac{1}{8} & \cdots \\ 0 & 0 & -1 & \frac{1}{2} & \frac{1}{4} & \cdots \\ 0 & 0 & 0 & -1 & \frac{1}{2} & \cdots \\ 0 & 0 & 0 & 0 & -1 & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix}$$

We would like to attach a mathematical meaning to the summation

$$\sum_{i,j=1}^{\infty} a_{ij}$$

whereby we intend to include every term in the preceding array in the total. One natural idea is to temporarily fix i and sum across each row. A moment's reflection (and a fact about geometric series) shows that each row sums to 0. Summing the sums of the rows, we get

$$\sum_{i,j=1}^{\infty} a_{ij} = \sum_{i=1}^{\infty} \left(\sum_{j=1}^{\infty} a_{ij} \right) = \sum_{i=1}^{\infty} (0) = 0.$$

We could just as easily have decided to fix j and sum down each column first. In this case, we have

$$\sum_{i,j=1}^{\infty} a_{ij} = \sum_{j=1}^{\infty} \left(\sum_{i=1}^{\infty} a_{ij} \right) = \sum_{j=1}^{\infty} \left(\frac{-1}{2^{j-1}} \right) = -2.$$

Changing the order of the summation changes the value of the sum. One common way that double sums arise (although not this particular one) is from the multiplication of two series. There is a natural desire to write

$$\left(\sum a_i \right) \left(\sum b_j \right) = \sum_{i,j} a_i b_j,$$

except that the expression on the right-hand side makes no sense at the moment.

It is the pathologies that give rise to the need for rigor. A satisfying resolution to the questions raised will require that we be absolutely precise about what we mean as we manipulate these infinite objects. It may seem that progress is slow at first, but that is because we do not want to fall into the trap of letting the biases of our intuition corrupt our arguments. Rigorous proofs are meant to be a check on intuition, and in the end we will see that they actually improve our mental picture of the mathematical infinite. As a final example, consider something as intuitively fundamental as the associative property of addition applied to the series $\sum_{n=1}^{\infty} (-1)^n$. Grouping the terms one way gives

$$(-1 + 1) + (-1 + 1) + (-1 + 1) + (-1 + 1) + \cdots = 0 + 0 + 0 + 0 + \cdots = 0,$$

whereas grouping in another yields

$$-1 + (1 - 1) + (1 - 1) + (1 - 1) + (1 - 1) + \cdots = -1 + 0 + 0 + 0 + 0 + \cdots = -1.$$

Manipulations that are legitimate in finite settings do not always extend to infinite settings. Deciding when they do and why they do not is one of the central themes of analysis.

2.2 The Limit of a Sequence

An understanding of infinite series depends heavily on a clear understanding of the theory of sequences. In fact, most of the concepts in analysis can be reduced to statements about the behavior of sequences. Thus, we will spend a significant amount of time investigating sequences before taking on infinite series.

Definition 2.2.1. A *sequence* is a function whose domain is \mathbf{N} .

This formal definition leads immediately to the familiar depiction of a sequence as an ordered list of real numbers. Given a function $f : \mathbf{N} \rightarrow \mathbf{R}$, $f(n)$ is just the n th term on the list. The notation for sequences reinforces this familiar understanding.

Example 2.2.2. Each of the following are common ways to describe a sequence.

- (i) $(1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \cdots)$,
- (ii) $(\frac{1+n}{n})_{n=1}^{\infty} = (\frac{2}{1}, \frac{3}{2}, \frac{4}{3}, \cdots)$,
- (iii) (a_n) , where $a_n = 2^n$ for each $n \in \mathbf{N}$,
- (iv) (x_n) , where $x_1 = 2$ and $x_{n+1} = \frac{x_n+1}{2}$.

On occasion, it will be more convenient to index a sequence beginning with $n = 0$ or $n = n_0$ for some natural number n_0 different from 1. These minor variations should cause no confusion. What is essential is that a sequence be an *infinite* list of real numbers. What happens at the beginning of such a list is of

little importance in most cases. The business of analysis is concerned with the behavior of the infinite “tail” of a given sequence.

We now present what is arguably the most important definition in the book.

Definition 2.2.3 (Convergence of a Sequence). A sequence (a_n) converges to a real number a if, for every positive number ϵ , there exists an $N \in \mathbf{N}$ such that whenever $n \geq N$ it follows that $|a_n - a| < \epsilon$.

To indicate that (a_n) converges to a , we write either $\lim a_n = a$ or $(a_n) \rightarrow a$.


In an effort to decipher this complicated definition, it helps first to consider the ending phrase “ $|a_n - a| < \epsilon$,” and think about the points that satisfy an inequality of this type.

Definition 2.2.4. Given a real number $a \in \mathbf{R}$ and a positive number $\epsilon > 0$, the set

$$V_\epsilon(a) = \{x \in \mathbf{R} : |x - a| < \epsilon\}$$

is called the ϵ -neighborhood of a .


Notice that $V_\epsilon(a)$ consists of all of those points whose distance from a is less than ϵ . Said another way, $V_\epsilon(a)$ is an interval, centered at a , with radius ϵ .



Recasting the definition of convergence in terms of ϵ -neighborhoods gives a more geometric impression of what is being described.

Definition 2.2.3B (Convergence of a Sequence: Topological Version).

A sequence (a_n) converges to a if, given any ϵ -neighborhood $V_\epsilon(a)$ of a , there exists a point in the sequence after which all of the terms are in $V_\epsilon(a)$. In other words, every ϵ -neighborhood contains all but a finite number of the terms of (a_n) .



Definition 2.2.3 and Definition 2.2.3B say precisely the same thing; the natural number N in the original version of the definition is the point where the sequence (a_n) enters $V_\epsilon(a)$, never to leave. It should be apparent that *the value of N depends on the choice of ϵ* . The smaller the ϵ -neighborhood, the larger N may have to be.

Example 2.2.5. Consider the sequence (a_n) , where $a_n = 1/\sqrt{n}$.

Our intuitive understanding of limits points confidently to the conclusion that

$$\lim \left(\frac{1}{\sqrt{n}} \right) = 0.$$

Before trying to prove this not too impressive fact, let's first explore the relationship between ϵ and N in the definition of convergence. For the moment, take ϵ to be $1/10$. This defines a sort of "target zone" for the terms in the sequence. By claiming that the limit of (a_n) is 0, we are saying that the terms in this sequence eventually get arbitrarily close to 0. How close? What do we mean by "eventually"? We have set $\epsilon = 1/10$ as our standard for closeness, which leads to the ϵ -neighborhood $(-1/10, 1/10)$ centered around the limit 0. How far out into the sequence must we look before the terms fall into this interval? The 100th term $a_{100} = 1/10$ puts us right on the boundary, and a little thought reveals that

$$\text{if } n > 100, \quad \text{then } a_n \in \left(-\frac{1}{10}, \frac{1}{10} \right).$$

Thus, for $\epsilon = 1/10$ we choose $N = 101$ (or anything larger) as our response.

Now, our choice of $\epsilon = 1/10$ was rather whimsical, and we can do this again, letting $\epsilon = 1/50$. In this case, our target neighborhood shrinks to $(-1/50, 1/50)$, and it is apparent that we must travel farther out into the sequence before a_n falls into this interval. How far? Essentially, we require that

$$\frac{1}{\sqrt{n}} < \frac{1}{50} \quad \text{which occurs as long as } n > 50^2 = 2500.$$

Thus, $N = 2501$ is a suitable response to the challenge of $\epsilon = 1/50$.

It may seem as though this duel could continue forever, with different ϵ challenges being handed to us one after another, each one requiring a suitable value of N in response. In a sense, this is correct, except that the game is effectively over the instant we recognize a *rule* for how to choose N given an *arbitrary* $\epsilon > 0$. For this problem, the desired algorithm is implicit in the algebra carried out to compute the previous response of $N = 2501$. Whatever ϵ happens to be, we want

$$\frac{1}{\sqrt{n}} < \epsilon \quad \text{which is equivalent to insisting that } n > \frac{1}{\epsilon^2}.$$

With this observation, we are ready to write the formal argument.

We claim that

$$\lim \left(\frac{1}{\sqrt{n}} \right) = 0.$$

Proof. Let $\epsilon > 0$ be an arbitrary positive number. Choose a natural number N satisfying

$$N > \frac{1}{\epsilon^2}.$$

We now verify that this choice of N has the desired property. Let $n \geq N$. Then,

$$n > \frac{1}{\epsilon^2} \quad \text{implies} \quad \frac{1}{\sqrt{n}} < \epsilon \quad \text{and hence} \quad |a_n - 0| < \epsilon.$$

□

Quantifiers

The definition of convergence given earlier is the result of hundreds of years of refining the intuitive notion of limit into a mathematically rigorous statement. The logic involved is complicated and is intimately tied to the use of the quantifiers “for all” and “there exists.” Learning to write a grammatically correct convergence proof goes hand in hand with a deep understanding of why the quantifiers appear in the order that they do.

The definition begins with the phrase,

“For all ϵ , there exists $N \in \mathbf{N}$ such that ...”

Looking back at our first example, we see that our formal proof begins with, “Let $\epsilon > 0$ be an arbitrary positive number.” This is followed by a construction of N and then a demonstration that this choice of N has the desired property. This, in fact, is a basic outline for how every convergence proof should be presented.

TEMPLATE FOR A PROOF THAT $(x_n) \rightarrow x$:

- “Let $\epsilon > 0$ be arbitrary.”
- Demonstrate a choice for $N \in \mathbf{N}$. This step usually requires the most work, almost all of which is done prior to actually writing the formal proof.
- Now, show that N actually works.
- “Assume $n \geq N$.”
- With N well chosen, it should be possible to derive the inequality $|x_n - x| < \epsilon$.

Example 2.2.6. Show

$$\lim \left(\frac{n+1}{n} \right) = 1.$$

As mentioned, before attempting a formal proof, we first need to do some preliminary scratch work. In the first example, we experimented by assigning specific values to ϵ (and it is not a bad idea to do this again), but let us skip straight to the algebraic punch line. The last line of our proof should be that for suitably large values of n ,

$$\left| \frac{n+1}{n} - 1 \right| < \epsilon.$$

Because

$$\left| \frac{n+1}{n} - 1 \right| = \frac{1}{n},$$

this is equivalent to the inequality $1/n < \epsilon$ or $n > 1/\epsilon$. Thus, choosing N to be an integer greater than $1/\epsilon$ will suffice.

With the work of the proof done, all that remains is the formal writeup.

Proof. Let $\epsilon > 0$ be arbitrary. Choose $N \in \mathbf{N}$ with $N > 1/\epsilon$. To verify that this choice of N is appropriate, let $n \in \mathbf{N}$ satisfy $n \geq N$. Then, $n \geq N$ implies $n > 1/\epsilon$, which is the same as saying $1/n < \epsilon$. Finally, this means

$$\left| \frac{n+1}{n} - 1 \right| < \epsilon,$$

as desired. □

Divergence

Significant insight into the role of the quantifiers in the definition of convergence can be gained by studying an example of a sequence that does not have a limit.

Example 2.2.7. Consider the sequence

$$\left(1, -\frac{1}{2}, \frac{1}{3}, -\frac{1}{4}, \frac{1}{5}, -\frac{1}{5}, \frac{1}{5}, -\frac{1}{5}, \frac{1}{5}, -\frac{1}{5}, \frac{1}{5}, -\frac{1}{5}, \frac{1}{5}, -\frac{1}{5}, \dots \right).$$

How can we argue that this sequence does not converge to zero? Looking at the first few terms, it seems the initial evidence actually supports such a conclusion. Given a challenge of $\epsilon = 1/2$, a little reflection reveals that after $N = 3$ all the terms fall into the neighborhood $(-1/2, 1/2)$. We could also handle $\epsilon = 1/4$. (What is the smallest possible N in this case?)

But the definition of convergence says “*For all* $\epsilon > 0$...,” and it should be apparent that there is no response to a choice of $\epsilon = 1/10$, for instance. This leads us to an important observation about the logical negation of the definition of convergence of a sequence. To prove that a particular number x is *not* the limit of a sequence (x_n) , we must produce a single value of ϵ for which no $N \in \mathbf{N}$ works. More generally speaking, the negation of a statement that begins “For all P, there exists Q...” is the statement, “For at least one P, no Q is possible...” For instance, how could we disprove the spurious claim that “At every college in the United States, there is a student who is at least seven feet tall”?

We have argued that the preceding sequence does not converge to 0. Let’s argue against the claim that it converges to $1/5$. Choosing $\epsilon = 1/10$ produces the neighborhood $(1/10, 3/10)$. Although the sequence continually revisits this neighborhood, there is no point at which it enters and never leaves as the definition requires. Thus, no N exists for $\epsilon = 1/10$, so the sequence does not converge to $1/5$.

Of course, this sequence does not converge to any other real number, and it would be more satisfying to simply say that this sequence does not converge.

Definition 2.2.8. A sequence that does not converge is said to *diverge*.

Although it is not too difficult, we will postpone arguing for divergence in general until we develop a more economical divergence criterion later in Section 2.5.

Exercises

Exercise 2.2.1. Verify, using the definition of convergence of a sequence, that the following sequences converge to the proposed limit.

(a) $\lim \frac{1}{(6n^2+1)} = 0.$

(b) $\lim \frac{3n+1}{(2n+5)} = \frac{3}{2}.$

(c) $\lim \frac{2}{\sqrt{n+3}} = 0.$

Exercise 2.2.2. What happens if we reverse the order of the quantifiers in Definition 2.2.3?

Definition: A sequence (x_n) *verconges* to x if *there exists* an $\epsilon > 0$ such that *for all* $N \in \mathbf{N}$ it is true that $n \geq N$ implies $|x_n - x| < \epsilon$.

Give an example of a vercongent sequence. Can you give an example of a vergonent sequence that is divergent? What exactly is being described in this strange definition?

Exercise 2.2.3. Describe what we would have to demonstrate in order to disprove each of the following statements.

(a) At every college in the United States, there is a student who is at least seven feet tall.

(b) For all colleges in the United States, there exists a professor who gives every student a grade of either A or B.

(c) There exists a college in the United States where every student is at least six feet tall.

Exercise 2.2.4. Argue that the sequence

$$1, 0, 1, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0, 0, 1, (5 \text{ zeros}), 1, \dots$$

does not converge to zero. For what values of $\epsilon > 0$ does there exist a response N . For which values of $\epsilon > 0$ is there no suitable response?

Exercise 2.2.5. Let $[[x]]$ be the greatest integer less than or equal to x . For example, $[[\pi]] = 3$ and $[[3]] = 3$. Find $\lim a_n$ and supply proofs for each conclusion if

(a) $a_n = [[1/n]],$

(b) $a_n = [[(10+n)/2n]].$

Reflecting on these examples, comment on the statement following Definition 2.2.3 that “the smaller the ϵ -neighborhood, the larger N may have to be.”

Exercise 2.2.6. Suppose that for a particular $\epsilon > 0$ we have found a suitable value of N that “works” for a given sequence in the sense of Definition 2.2.3.

(a) Then, any larger/smaller (*pick one*) N will also work for the same $\epsilon > 0$.

(b) Then, this same N will also work for any larger/smaller value of ϵ .

Exercise 2.2.7. Informally speaking, the sequence \sqrt{n} “converges to infinity.”

(a) Imitate the logical structure of Definition 2.2.3 to create a rigorous definition for the mathematical statement $\lim x_n = \infty$. Use this definition to prove $\lim \sqrt{n} = \infty$.

(b) What does your definition in (a) say about the particular sequence $(1, 0, 2, 0, 3, 0, 4, 0, 5, 0, \dots)$?

Exercise 2.2.8. Here are two useful definitions:

(i) A sequence (a_n) is *eventually* in a set $A \subseteq \mathbf{R}$ if there exists an $N \in \mathbf{N}$ such that $a_n \in A$ for all $n \geq N$.

(ii) A sequence (a_n) is *frequently* in a set $A \subseteq \mathbf{R}$ if, for every $N \in \mathbf{N}$, there exists an $n \geq N$ such that $a_n \in A$.

(a) Is the sequence $(-1)^n$ eventually or frequently in the set $\{1\}$?

(b) Which definition is stronger? Does frequently imply eventually or does eventually imply frequently?

(c) Give an alternate rephrasing of Definition 2.2.3B using either frequently or eventually. Which is the term we want?

(d) Suppose an infinite number of terms of a sequence (x_n) are equal to 2. Is (x_n) necessarily eventually in the interval $(1.9, 2.1)$? Is it frequently in $(1.9, 2.1)$?

2.3 The Algebraic and Order Limit Theorems

The real purpose of creating a rigorous definition for convergence of a sequence is *not* to have a tool to verify computational statements such as $\lim 2n/(n+2) = 2$. Historically, a definition of the limit like Definition 2.2.3 came 150 years after the founders of calculus began working with intuitive notions of convergence. The point of having such a logically tight description of convergence is so that we can confidently state and prove statements *about convergence sequences in general*. We are ultimately trying to resolve arguments about what is and is not true regarding the behavior of limits with respect to the mathematical manipulations we intend to inflict on them.

As a first example, let us prove that convergent sequences are bounded. The term “bounded” has a rather familiar connotation but, like everything else, we need to be explicit about what it means in this context.

Definition 2.3.1. A sequence (x_n) is *bounded* if there exists a number $M > 0$ such that $|x_n| \leq M$ for all $n \in \mathbf{N}$.


Geometrically, this means that we can find an interval $[-M, M]$ that contains every term in the sequence (x_n) .

Theorem 2.3.2. *Every convergent sequence is bounded.*

Proof. Assume (x_n) converges to a limit l . This means that given a particular value of ϵ , say $\epsilon = 1$, we know there must exist an $N \in \mathbf{N}$ such that if $n \geq N$, then x_n is in the interval $(l - 1, l + 1)$. Not knowing whether l is positive or negative, we can certainly conclude that

$$|x_n| < |l| + 1$$

for all $n \geq N$.



We still need to worry (slightly) about the terms in the sequence that come before the N th term. Because there are only a finite number of these, we let

$$M = \max\{|x_1|, |x_2|, |x_3|, \dots, |x_{N-1}|, |l| + 1\}.$$

It follows that $|x_n| \leq M$ for all $n \in \mathbf{N}$, as desired. \square

This chapter began with a demonstration of how applying familiar algebraic properties (commutativity of addition) to infinite objects (series) can lead to paradoxical results. These examples are meant to instill in us a sense of caution and justify the extreme care we are taking in drawing our conclusions. The following theorems illustrate that sequences behave extremely well with respect to the operations of addition, multiplication, division, and order.

Theorem 2.3.3 (Algebraic Limit Theorem). *Let $\lim a_n = a$, and $\lim b_n = b$. Then,*

- (i) $\lim(ca_n) = ca$, for all $c \in \mathbf{R}$;
- (ii) $\lim(a_n + b_n) = a + b$;
- (iii) $\lim(a_nb_n) = ab$;
- (iv) $\lim(a_n/b_n) = a/b$, provided $b \neq 0$.

Proof. (i) Consider the case where $c \neq 0$. We want to show that the sequence (ca_n) converges to ca , so the structure of the proof follows the template we described in Section 2.2. First, we let ϵ be some arbitrary positive number. Our goal is to find some point in the sequence (ca_n) after which we have

$$|ca_n - ca| < \epsilon.$$

Now,

$$|ca_n - ca| = |c||a_n - a|.$$

We are given that $(a_n) \rightarrow a$, so we know we can make $|a_n - a|$ as small as we like. In particular, we can choose an N such that

$$|a_n - a| < \frac{\epsilon}{|c|}$$

whenever $n \geq N$. To see that this N indeed works, observe that, for all $n \geq N$,

$$|ca_n - ca| = |c||a_n - a| < |c|\frac{\epsilon}{|c|} = \epsilon.$$

The case $c = 0$ reduces to showing that the constant sequence $(0, 0, 0, \dots)$ converges to 0. This is addressed in Exercise 2.3.1.

Before continuing with parts (ii), (iii), and (iv), we should point out that the proof of (i), while somewhat short, is extremely typical for a convergence proof. Before embarking on a formal argument, it is a good idea to take an inventory of what we *want* to make less than ϵ , and what we are *given* can be made small for suitable choices of n . For the previous proof, we wanted to make $|ca_n - ca| < \epsilon$, and we were given $|a_n - a| < \text{anything we like}$ (for large values of n). Notice that in (i), and all of the ensuing arguments, the strategy each time is to bound the quantity we want to be less than ϵ , which in each case is

$$|(\text{terms of sequence}) - (\text{proposed limit})|,$$

with some algebraic combination of quantities over which we have control.

(ii) To prove this statement, we need to argue that the quantity

$$|(a_n + b_n) - (a + b)|$$

can be made less than an arbitrary ϵ using the assumptions that $|a_n - a|$ and $|b_n - b|$ can be made as small as we like for large n . The first step is to use the triangle inequality (Example 1.2.5) to say

$$|(a_n + b_n) - (a + b)| = |(a_n - a) + (b_n - b)| \leq |a_n - a| + |b_n - b|.$$

Again, we let $\epsilon > 0$ be arbitrary. The technique this time is to divide the ϵ between the two expressions on the right-hand side in the preceding inequality. Using the hypothesis that $(a_n) \rightarrow a$, we know there exists an N_1 such that

$$|a_n - a| < \frac{\epsilon}{2} \quad \text{whenever} \quad n \geq N_1.$$

Likewise, the assumption that $(b_n) \rightarrow b$ means that we can choose an N_2 so that

$$|b_n - b| < \frac{\epsilon}{2} \quad \text{whenever} \quad n \geq N_2.$$

The question now arises as to which of N_1 or N_2 we should take to be our choice of N . By choosing $N = \max\{N_1, N_2\}$, we ensure that if $n \geq N$, then $n \geq N_1$ and $n \geq N_2$. This allows us to conclude that

$$\begin{aligned} |(a_n + b_n) - (a + b)| &\leq |a_n - a| + |b_n - b| \\ &< \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon \end{aligned}$$

for all $n \geq N$, as desired.

(iii) To show that $(a_n b_n) \rightarrow ab$, we begin by observing that

$$\begin{aligned} |a_n b_n - ab| &= |a_n b_n - ab_n + ab_n - ab| \\ &\leq |a_n b_n - ab_n| + |ab_n - ab| \\ &= |b_n| |a_n - a| + |a| |b_n - b|. \end{aligned}$$

In the initial step, we subtracted and then added ab_n , which created an opportunity to use the triangle inequality. Essentially, we have broken up the distance from $a_n b_n$ to ab with a midway point and are using the sum of the two distances to overestimate the original distance. This clever trick will become a familiar technique in arguments to come.

Letting $\epsilon > 0$ be arbitrary, we again proceed with the strategy of making each piece in the preceding inequality less than $\epsilon/2$. For the piece on the right-hand side ($|a| |b_n - b|$), if $a \neq 0$ we can choose N_1 so that

$$n \geq N_1 \quad \text{implies} \quad |b_n - b| < \frac{1}{|a|} \frac{\epsilon}{2}.$$

(The case when $a = 0$ is handled in Exercise 2.3.7.) Getting the term on the left-hand side ($|b_n| |a_n - a|$) to be less than $\epsilon/2$ is complicated by the fact that we have a variable quantity $|b_n|$ to contend with as opposed to the constant $|a|$ we encountered in the right-hand term. The idea is to replace $|b_n|$ with a worst-case estimate. Using the fact that convergent sequences are bounded (Theorem 2.3.2), we know there exists a bound $M > 0$ satisfying $|b_n| \leq M$ for all $n \in \mathbf{N}$. Now, we can choose N_2 so that

$$|a_n - a| < \frac{1}{M} \frac{\epsilon}{2} \quad \text{whenever} \quad n \geq N_2.$$

To finish the argument, pick $N = \max\{N_1, N_2\}$, and observe that if $n \geq N$, then

$$\begin{aligned} |a_n b_n - ab| &\leq |a_n b_n - ab_n| + |ab_n - ab| \\ &= |b_n| |a_n - a| + |a| |b_n - b| \\ &\leq M |a_n - a| + |a| |b_n - b| \\ &< M \left(\frac{\epsilon}{M2} \right) + |a| \left(\frac{\epsilon}{|a|2} \right) = \epsilon. \end{aligned}$$

(iv) This final statement will follow from (iii) if we can prove that

$$(b_n) \rightarrow b \quad \text{implies} \quad \left(\frac{1}{b_n} \right) \rightarrow \frac{1}{b}$$

whenever $b \neq 0$. We begin by observing that

$$\left| \frac{1}{b_n} - \frac{1}{b} \right| = \frac{|b - b_n|}{|b| |b_n|}.$$

Because $(b_n) \rightarrow b$, we can make the preceding numerator as small as we like by choosing n large. The problem comes in that we need a worst-case estimate on the size of $1/(|b||b_n|)$. Because the b_n terms are in the denominator, we are no longer interested in an upper bound on $|b_n|$ but rather in an inequality of the form $|b_n| \geq \delta > 0$. This will then lead to a bound on the size of $1/(|b||b_n|)$.

The trick is to look far enough out into the sequence (b_n) so that the terms are closer to b than they are to 0. Consider the particular value $\epsilon_0 = |b|/2$. Because $(b_n) \rightarrow b$, there exists an N_1 such that $|b_n - b| < |b|/2$ for all $n \geq N_1$. This implies $|b_n| > |b|/2$.

Next, choose N_2 so that $n \geq N_2$ implies

$$|b_n - b| < \frac{\epsilon|b|^2}{2}.$$

Finally, if we let $N = \max\{N_1, N_2\}$, then $n \geq N$ implies

$$\left| \frac{1}{b_n} - \frac{1}{b} \right| = |b - b_n| \frac{1}{|b||b_n|} < \frac{\epsilon|b|^2}{2} \frac{1}{|b| \frac{|b|}{2}} = \epsilon.$$

□


Limits and Order

Although there are a few dangers to avoid (see Exercise 2.3.8), the Algebraic Limit Theorem verifies that the relationship between algebraic combinations of sequences and the limiting process is as trouble-free as we could hope for. Limits can be computed from the individual component sequences provided that each component limit exists. The limiting process is also well-behaved with respect to the order operation.

Theorem 2.3.4 (Order Limit Theorem). *Assume $\lim a_n = a$ and $\lim b_n = b$.*

- (i) *If $a_n \geq 0$ for all $n \in \mathbf{N}$, then $a \geq 0$.*
- (ii) *If $a_n \leq b_n$ for all $n \in \mathbf{N}$, then $a \leq b$.*
- (iii) *If there exists $c \in \mathbf{R}$ for which $c \leq b_n$ for all $n \in \mathbf{N}$, then $c \leq b$. Similarly, if $a_n \leq c$ for all $n \in \mathbf{N}$, then $a \leq c$.*

Proof. (i) We will prove this by contradiction; thus, let's assume $a < 0$. The idea is to produce a term in the sequence (a_n) that is also less than zero. To do this, we consider the particular value $\epsilon_0 = |a|$. The definition of convergence guarantees that we can find an N such that $|a_n - a| < |a|$ for all $n \geq N$. In particular, this would mean that $|a_N - a| < |a|$, which implies $a_N < 0$. This contradicts our hypothesis that $a_n \geq 0$. We therefore conclude that $a \geq 0$.



(ii) The Algebraic Limit Theorem ensures that the sequence $(b_n - a_n)$ converges to $b - a$. Because $b_n - a_n \geq 0$, we can apply part (i) to get that $b - a \geq 0$.

(iii) Take $a_n = c$ (or $b_n = c$) for all $n \in \mathbf{N}$, and apply (ii). \square

A word about the idea of “tails” is in order. Loosely speaking, limits and their properties do not depend at all on what happens at the beginning of the sequence but are strictly determined by what happens when n gets large. Changing the value of the first ten—or ten thousand—terms in a particular sequence has no effect on the limit. Theorem 2.3.4, part (i), for instance, assumes that $a_n \geq 0$ for all $n \in \mathbf{N}$. However, the hypothesis could be weakened by assuming only that there exists some point N_1 where $a_n \geq 0$ for all $n \geq N_1$. The theorem remains true, and in fact the same proof is valid with the provision that when N is chosen it be at least as large as N_1 .

In the language of analysis, when a property (such as non-negativity) is not necessarily true about some finite number of initial terms but is true for all terms in the sequence after some point N , we say that the sequence *eventually* has this property. (See Exercise 2.2.8.) Theorem 2.3.4, part (i), could be restated, “Convergent sequences that are eventually nonnegative converge to nonnegative limits.” Parts (ii) and (iii) have similar modifications, as will many other upcoming results.

Exercises

Exercise 2.3.1. Show that the constant sequence (a, a, a, a, \dots) converges to a .

Exercise 2.3.2. Let $x_n \geq 0$ for all $n \in \mathbf{N}$.

- (a) If $(x_n) \rightarrow 0$, show that $(\sqrt{x_n}) \rightarrow 0$.
- (b) If $(x_n) \rightarrow x$, show that $(\sqrt{x_n}) \rightarrow \sqrt{x}$.

Exercise 2.3.3 (Squeeze Theorem). Show that if $x_n \leq y_n \leq z_n$ for all $n \in \mathbf{N}$, and if $\lim x_n = \lim z_n = l$, then $\lim y_n = l$ as well.

Exercise 2.3.4. Show that limits, if they exist, must be unique. In other words, assume $\lim a_n = l_1$ and $\lim a_n = l_2$, and prove that $l_1 = l_2$.

Exercise 2.3.5. Let (x_n) and (y_n) be given, and define (z_n) to be the “shuffled” sequence $(x_1, y_1, x_2, y_2, x_3, y_3, \dots, x_n, y_n, \dots)$. Prove that (z_n) is convergent if and only if (x_n) and (y_n) are both convergent with $\lim x_n = \lim y_n$.

Exercise 2.3.6. (a) Show that if $(b_n) \rightarrow b$, then the sequence of absolute values $|b_n|$ converges to $|b|$.

(b) Is the converse of part (a) true? If we know that $|b_n| \rightarrow |b|$, can we deduce that $(b_n) \rightarrow b$?

Exercise 2.3.7. (a) Let (a_n) be a bounded (not necessarily convergent) sequence, and assume $\lim b_n = 0$. Show that $\lim(a_n b_n) = 0$. Why are we not allowed to use the Algebraic Limit Theorem to prove this?

(b) Can we conclude anything about the convergence of $(a_n b_n)$ if we assume that (b_n) converges to some nonzero limit b ?

(c) Use (a) to prove Theorem 2.3.3, part (iii), for the case when $a = 0$.

Exercise 2.3.8. Give an example of each of the following, or state that such a request is impossible by referencing the proper theorem(s):

(a) sequences (x_n) and (y_n) , which both diverge, but whose sum $(x_n + y_n)$ converges;

(b) sequences (x_n) and (y_n) , where (x_n) converges, (y_n) diverges, and $(x_n + y_n)$ converges;

(c) a convergent sequence (b_n) with $b_n \neq 0$ for all n such that $(1/b_n)$ diverges;

(d) an unbounded sequence (a_n) and a convergent sequence (b_n) with $(a_n - b_n)$ bounded;

(e) two sequences (a_n) and (b_n) , where $(a_n b_n)$ and (a_n) converge but (b_n) does not.

Exercise 2.3.9. Does Theorem 2.3.4 remain true if all of the inequalities are assumed to be strict? If we assume, for instance, that a convergent sequence (x_n) satisfies $x_n > 0$ for all $n \in \mathbf{N}$, what may we conclude about the limit?

Exercise 2.3.10. If $(a_n) \rightarrow 0$ and $|b_n - b| \leq a_n$, then show that $(b_n) \rightarrow b$.

Exercise 2.3.11 (Cesaro Means). Show that if (x_n) is a convergent sequence, then the sequence given by the averages

$$y_n = \frac{x_1 + x_2 + \cdots + x_n}{n}$$

also converges to the same limit.

Give an example to show that it is possible for the sequence (y_n) of averages to converge even if (x_n) does not.

Exercise 2.3.12. Consider the doubly indexed array $a_{m,n} = m/(m+n)$.

(a) Intuitively speaking, what should $\lim_{m,n \rightarrow \infty} a_{m,n}$ represent? Compute the “iterated” limits

$$\lim_{n \rightarrow \infty} \lim_{m \rightarrow \infty} a_{m,n} \quad \text{and} \quad \lim_{m \rightarrow \infty} \lim_{n \rightarrow \infty} a_{m,n}.$$

(b) Formulate a rigorous definition in the style of Definition 2.2.3 for the statement

$$\lim_{m,n \rightarrow \infty} a_{m,n} = l.$$

2.4 The Monotone Convergence Theorem and a First Look at Infinite Series

We showed in Theorem 2.3.2 that convergent sequences are bounded. The converse statement is certainly not true. It is not too difficult to produce an example of a bounded sequence that does not converge. On the other hand, if a bounded sequence is *monotone*, then in fact it does converge.


Definition 2.4.1. A sequence (a_n) is *increasing* if $a_n \leq a_{n+1}$ for all $n \in \mathbf{N}$ and *decreasing* if $a_n \geq a_{n+1}$ for all $n \in \mathbf{N}$. A sequence is *monotone* if it is either increasing or decreasing.

Theorem 2.4.2 (Monotone Convergence Theorem). *If a sequence is monotone and bounded, then it converges.*

Proof. Let (a_n) be monotone and bounded. To prove (a_n) converges using the definition of convergence, we are going to need a candidate for the limit. Let's assume the sequence is increasing (the decreasing case is handled similarly), and consider the *set* of points $\{a_n : n \in \mathbf{N}\}$. By assumption, this set is bounded, so we can let

$$s = \sup\{a_n : n \in \mathbf{N}\}.$$

It seems reasonable to claim that $\lim(a_n) = s$.



To prove this, let $\epsilon > 0$. Because s is the least upper bound of $\{a_n : n \in \mathbf{N}\}$, $s - \epsilon$ is not an upper bound, so there exists a point in the sequence a_N such that $s - \epsilon < a_N$. Now, the fact that (a_n) is increasing implies that if $n \geq N$, then $a_N \leq a_n$. Hence,

$$s - \epsilon < a_N \leq a_n \leq s < s + \epsilon,$$

which implies $|a_n - s| < \epsilon$, as desired. \square

The Monotone Convergence Theorem is extremely useful for the study of infinite series, largely because it asserts the convergence of a sequence without explicit mention of the actual limit. This is a good moment to do some preliminary investigations, so it is time to formalize the relationship between sequences and series.

Definition 2.4.3. Let (b_n) be a sequence. An *infinite series* is a formal expression of the form

$$\sum_{n=1}^{\infty} b_n = b_1 + b_2 + b_3 + b_4 + b_5 + \cdots .$$

We define the corresponding *sequence of partial sums* (s_m) by

$$s_m = b_1 + b_2 + b_3 + \cdots + b_m,$$

and say that the series $\sum_{n=1}^{\infty} b_n$ *converges to* B if the sequence (s_m) converges to B . In this case, we write $\sum_{n=1}^{\infty} b_n = B$.

Example 2.4.4. Consider

$$\sum_{n=1}^{\infty} \frac{1}{n^2}.$$

Because the terms in the sum are all positive, the sequence of partial sums given by

$$s_m = 1 + \frac{1}{4} + \frac{1}{9} + \cdots + \frac{1}{m^2}$$

is increasing. The question is whether or not we can find some upper bound on (s_m) . To this end, observe

$$\begin{aligned} s_m &= 1 + \frac{1}{2 \cdot 2} + \frac{1}{3 \cdot 3} + \frac{1}{4 \cdot 4} + \cdots + \frac{1}{m^2} \\ &< 1 + \frac{1}{2 \cdot 1} + \frac{1}{3 \cdot 2} + \frac{1}{4 \cdot 3} + \cdots + \frac{1}{m(m-1)} \\ &= 1 + \left(1 - \frac{1}{2}\right) + \left(\frac{1}{2} - \frac{1}{3}\right) + \left(\frac{1}{3} - \frac{1}{4}\right) + \cdots + \left(\frac{1}{(m-1)} - \frac{1}{m}\right) \\ &= 1 + 1 - \frac{1}{m} \\ &< 2. \end{aligned}$$

Thus, 2 is an upper bound for the sequence of partial sums, so by the Monotone Convergence Theorem, $\sum_{n=1}^{\infty} 1/n^2$ converges to some (presently unknown) limit less than 2.

Example 2.4.5 (Harmonic Series). This time, consider the so-called *harmonic series*

$$\sum_{n=1}^{\infty} \frac{1}{n}.$$

Again, we have an increasing sequence of partial sums,

$$s_m = 1 + \frac{1}{2} + \frac{1}{3} + \cdots + \frac{1}{m},$$

that upon naive inspection appears as though it may be bounded. However, 2 is no longer an upper bound because

$$s_4 = 1 + \frac{1}{2} + \left(\frac{1}{3} + \frac{1}{4}\right) > 1 + \frac{1}{2} + \left(\frac{1}{4} + \frac{1}{4}\right) = 2.$$

A similar calculation shows that $s_8 > 2\frac{1}{2}$, and we can see that in general

$$\begin{aligned}
 s_{2^k} &= 1 + \frac{1}{2} + \left(\frac{1}{3} + \frac{1}{4}\right) + \left(\frac{1}{5} + \cdots + \frac{1}{8}\right) + \cdots + \left(\frac{1}{2^{k-1}+1} + \cdots + \frac{1}{2^k}\right) \\
 &> 1 + \frac{1}{2} + \left(\frac{1}{4} + \frac{1}{4}\right) + \left(\frac{1}{8} + \cdots + \frac{1}{8}\right) + \cdots + \left(\frac{1}{2^k} + \cdots + \frac{1}{2^k}\right) \\
 &= 1 + \frac{1}{2} + 2\left(\frac{1}{4}\right) + 4\left(\frac{1}{8}\right) + \cdots + 2^{k-1}\left(\frac{1}{2^k}\right) \\
 &= 1 + \frac{1}{2} + \frac{1}{2} + \frac{1}{2} + \cdots + \frac{1}{2} \\
 &= 1 + k\left(\frac{1}{2}\right),
 \end{aligned}$$

which is unbounded. Thus, despite the incredibly slow pace, the sequence of partial sums of $\sum_{n=1}^{\infty} 1/n$ eventually surpasses every number on the positive real line. Because convergent sequences are bounded, the harmonic series diverges.

The previous example is a special case of a general argument that can be used to determine the convergence or divergence of a large class of infinite series.

Theorem 2.4.6 (Cauchy Condensation Test). *Suppose (b_n) is decreasing and satisfies $b_n \geq 0$ for all $n \in \mathbf{N}$. Then, the series $\sum_{n=1}^{\infty} b_n$ converges if and only if the series*

$$\sum_{n=0}^{\infty} 2^n b_{2^n} = b_1 + 2b_2 + 4b_4 + 8b_8 + 16b_{16} + \cdots$$

converges.

Proof. First, assume that $\sum_{n=0}^{\infty} 2^n b_{2^n}$ converges. Theorem 2.3.2 guarantees that the partial sums

$$t_k = b_1 + 2b_2 + 4b_4 + \cdots + 2^k b_{2^k}$$

are bounded; that is, there exists an $M > 0$ such that $t_k \leq M$ for all $k \in \mathbf{N}$. We want to prove that $\sum_{n=1}^{\infty} b_n$ converges. Because $b_n \geq 0$, we know that the partial sums are increasing, so we only need to show that

$$s_m = b_1 + b_2 + b_3 + \cdots + b_m$$

is bounded.

Fix m and let k be large enough to ensure $m \leq 2^{k+1} - 1$. Then, $s_m \leq s_{2^{k+1}-1}$ and

$$\begin{aligned}
 s_{2^{k+1}-1} &= b_1 + (b_2 + b_3) + (b_4 + b_5 + b_6 + b_7) + \cdots + (b_{2^k} + \cdots + b_{2^{k+1}-1}) \\
 &\leq b_1 + (b_2 + b_2) + (b_4 + b_4 + b_4 + b_4) + \cdots + (b_{2^k} + \cdots + b_{2^k}) \\
 &= b_1 + 2b_2 + 4b_4 + \cdots + 2^k b_{2^k} = t_k.
 \end{aligned}$$

Thus, $s_m \leq t_k \leq M$, and the sequence (s_m) is bounded. By the Monotone Convergence Theorem, we can conclude that $\sum_{n=1}^{\infty} b_n$ converges.

The proof that $\sum_{n=0}^{\infty} 2^n b_{2^n}$ diverges implies $\sum_{n=1}^{\infty} b_n$ diverges is similar to Example 2.4.5. The details are requested in Exercise 2.4.1. \square

Corollary 2.4.7. *The series $\sum_{n=1}^{\infty} 1/n^p$ converges if and only if $p > 1$.*

A rigorous argument for this corollary requires a few basic facts about geometric series. The proof is requested in Exercise 2.7.7 at the end of Section 2.7 where geometric series are discussed.

Exercises

Exercise 2.4.1. Complete the proof of Theorem 2.4.6 by showing that if the series $\sum_{n=0}^{\infty} 2^n b_{2^n}$ diverges, then so does $\sum_{n=1}^{\infty} b_n$. Example 2.4.5 may be a useful reference.

Exercise 2.4.2. (a) Prove that the sequence defined by $x_1 = 3$ and

$$x_{n+1} = \frac{1}{4 - x_n}$$

converges.

(b) Now that we know $\lim x_n$ exists, explain why $\lim x_{n+1}$ must also exist and equal the same value.

(c) Take the limit of each side of the recursive equation in part (a) of this exercise to explicitly compute $\lim x_n$.

Exercise 2.4.3. Following the model of Exercise 2.4.2, show that the sequence defined by $y_1 = 1$ and $y_{n+1} = 4 - 1/y_n$ converges and find the limit.

Exercise 2.4.4. Show that

$$\sqrt{2}, \sqrt{2\sqrt{2}}, \sqrt{2\sqrt{2\sqrt{2}}}, \dots$$

converges and find the limit.

Exercise 2.4.5 (Calculating Square Roots). Let $x_1 = 2$, and define

$$x_{n+1} = \frac{1}{2} \left(x_n + \frac{2}{x_n} \right).$$

(a) Show that x_n^2 is always greater than 2, and then use this to prove that $x_n - x_{n+1} \geq 0$. Conclude that $\lim x_n = \sqrt{2}$.

(b) Modify the sequence (x_n) so that it converges to \sqrt{c} .

Exercise 2.4.6 (Limit Superior). Let (a_n) be a bounded sequence.

(a) Prove that the sequence defined by $y_n = \sup\{a_k : k \geq n\}$ converges.

(b) The *limit superior* of (a_n) , or $\limsup a_n$, is defined by

$$\limsup a_n = \lim y_n,$$

where y_n is the sequence from part (a) of this exercise. Provide a reasonable definition for $\liminf a_n$ and briefly explain why it always exists for any bounded sequence.

(c) Prove that $\liminf a_n \leq \limsup a_n$ for every bounded sequence, and give an example of a sequence for which the inequality is strict.

(d) Show that $\liminf a_n = \limsup a_n$ if and only if $\lim a_n$ exists. In this case, all three share the same value.

2.5 Subsequences and the Bolzano–Weierstrass Theorem

In Example 2.4.5, we showed that the sequence of partial sums (s_m) of the harmonic series does not converge by focusing our attention on a particular *subsequence* (s_{2^k}) of the original sequence. For the moment, we will put the topic of infinite series aside and more fully develop the important concept of subsequences.

Definition 2.5.1. Let (a_n) be a sequence of real numbers, and let $n_1 < n_2 < n_3 < n_4 < n_5 < \cdots$ be an increasing sequence of natural numbers. Then the sequence

$$a_{n_1}, a_{n_2}, a_{n_3}, a_{n_4}, a_{n_5}, \cdots$$

is called a *subsequence* of (a_n) and is denoted by (a_{n_j}) , where $j \in \mathbf{N}$ indexes the subsequence.

Notice that the order of the terms in a subsequence is the same as in the original sequence, and repetitions are not allowed. Thus if

$$(a_n) = \left(1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \frac{1}{5}, \frac{1}{6}, \cdots\right),$$

then

$$\left(\frac{1}{2}, \frac{1}{4}, \frac{1}{6}, \frac{1}{8}, \cdots\right) \quad \text{and} \quad \left(\frac{1}{10}, \frac{1}{100}, \frac{1}{1000}, \frac{1}{10000}, \cdots\right)$$

are examples of legitimate subsequences, whereas

$$\left(\frac{1}{10}, \frac{1}{5}, \frac{1}{100}, \frac{1}{50}, \frac{1}{1000}, \frac{1}{500}, \cdots\right) \quad \text{and} \quad \left(1, 1, \frac{1}{3}, \frac{1}{5}, \frac{1}{7}, \frac{1}{9}, \cdots\right)$$

are not.

Theorem 2.5.2. *Subsequences of a convergent sequence converge to the same limit as the original sequence.*

Proof. Exercise 2.5.1 □

This not too surprising result has several somewhat surprising applications. It is the key ingredient for understanding when infinite sums are associative (Exercise 2.5.2). We can also use it in the following clever way to compute values of some familiar limits.

Example 2.5.3. Let $0 < b < 1$. Because

$$b > b^2 > b^3 > b^4 > \cdots > 0,$$

the sequence (b^n) is decreasing and bounded below. The Monotone Convergence Theorem allows us to conclude that (b^n) converges to some l satisfying $b > l \geq 0$. To compute l , notice that (b^{2n}) is a subsequence, so $(b^{2n}) \rightarrow l$ by Theorem 2.5.2. But $(b^{2n}) = (b^n)(b^n)$, so by the Algebraic Limit Theorem, $(b^{2n}) \rightarrow l \cdot l = l^2$. Because limits are unique, $l^2 = l$, and thus $l = 0$.

Without much trouble (Exercise 2.5.5), we can generalize this example to conclude $(b^n) \rightarrow 0$ whenever $-1 < b < 1$.

Example 2.5.4 (Divergence Criterion). Theorem 2.5.2 is also useful for providing economical proofs for divergence. In Example 2.2.7, we were quite sure that

$$\left(1, -\frac{1}{2}, \frac{1}{3}, -\frac{1}{4}, \frac{1}{5}, -\frac{1}{5}, \frac{1}{5}, -\frac{1}{5}, \frac{1}{5}, -\frac{1}{5}, \frac{1}{5}, -\frac{1}{5}, \frac{1}{5}, -\frac{1}{5}, \cdots\right)$$

did not converge to any proposed limit. Notice that

$$\left(\frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \frac{1}{5}, \cdots\right)$$

is a subsequence that converges to $1/5$. Also,

$$\left(-\frac{1}{5}, -\frac{1}{5}, -\frac{1}{5}, -\frac{1}{5}, -\frac{1}{5}, \cdots\right)$$


is a different subsequence of the original sequence that converges to $-1/5$. Because we have two subsequences converging to two different limits, we can rigorously conclude that the original sequence diverges.

The Bolzano–Weierstrass Theorem

In the previous example, it was rather easy to spot a convergent subsequence (or two) hiding in the original sequence. For *bounded* sequences, it turns out that it is always possible to find at least one such convergent subsequence.

Theorem 2.5.5 (Bolzano–Weierstrass Theorem). *Every bounded sequence contains a convergent subsequence.*

Proof. Let (a_n) be a bounded sequence so that there exists $M > 0$ satisfying $|a_n| \leq M$ for all $n \in \mathbf{N}$. Bisect the closed interval $[-M, M]$ into the two closed intervals $[-M, 0]$ and $[0, M]$. (The midpoint is included in both halves.) Now, it must be that at least one of these closed intervals contains an infinite number of the points in the sequence (a_n) . Select a half for which this is the case and label that interval as I_1 . Then, let a_{n_1} be some point in the sequence (a_n) satisfying $a_{n_1} \in I_1$.



Next, we bisect I_1 into closed intervals of equal length, and let I_2 be a half that again contains an infinite number of points of the original sequence. Because there are an infinite number of points from (a_n) to choose from, we can select an a_{n_2} from the original sequence with $n_2 > n_1$ and $a_{n_2} \in I_2$. In general, we construct the closed interval I_k by taking a half of I_{k-1} containing an infinite number of points of (a_n) and then select $n_k > n_{k-1} > \cdots > n_2 > n_1$ so that $a_{n_k} \in I_k$.

We want to argue that (a_{n_k}) is a convergent subsequence, but we need a candidate for the limit. The sets

$$I_1 \supseteq I_2 \supseteq I_3 \supseteq \cdots$$

form a nested sequence of closed intervals, and by the Nested Interval Property there exists at least one point $x \in \mathbf{R}$ contained in every I_k . This provides us with the candidate we were looking for. It just remains to show that $(a_{n_k}) \rightarrow x$.

Let $\epsilon > 0$. By construction, the length of I_k is $M(1/2)^{k-1}$ which converges to zero. (This follows from Example 2.5.3 and the Algebraic Limit Theorem.) Choose N so that $k \geq N$ implies that the length of I_k is less than ϵ . Because x and a_{n_k} are both in I_k , it follows that $|a_{n_k} - x| < \epsilon$. \square

Exercises

Exercise 2.5.1. Prove Theorem 2.5.2.

Exercise 2.5.2. (a) Prove that if an infinite series converges, then the associative property holds. Assume $a_1 + a_2 + a_3 + a_4 + a_5 + \cdots$ converges to a limit L (i.e., the sequence of partial sums $(s_n) \rightarrow L$). Show that any regrouping of the terms

$$(a_1 + a_2 + \cdots + a_{n_1}) + (a_{n_1+1} + \cdots + a_{n_2}) + (a_{n_2+1} + \cdots + a_{n_3}) + \cdots$$

leads to a series that also converges to L .

(b) Compare this result to the example discussed at the end of Section 2.1 where infinite addition was shown not to be associative. Why doesn't our proof in (a) apply to this example?

Exercise 2.5.3. Give an example of each of the following, or argue that such a request is impossible.

(a) A sequence that does not contain 0 or 1 as a term but contains subsequences converging to each of these values.

(b) A monotone sequence that diverges but has a convergent subsequence.

(c) A sequence that contains subsequences converging to every point in the infinite set $\{1, 1/2, 1/3, 1/4, 1/5, \dots\}$.

(d) An unbounded sequence with a convergent subsequence.

(e) A sequence that has a subsequence that is bounded but contains no subsequence that converges.

Exercise 2.5.4. Assume (a_n) is a bounded sequence with the property that every convergent subsequence of (a_n) converges to the same limit $a \in \mathbf{R}$. Show that (a_n) must converge to a .

Exercise 2.5.5. Extend the result proved in Example 2.5.3 to the case $|b| < 1$. Show $\lim(b^n) = 0$ whenever $-1 < b < 1$.

Exercise 2.5.6. Let (a_n) be a bounded sequence, and define the set

$$S = \{x \in \mathbf{R} : x < a_n \text{ for infinitely many terms } a_n\}.$$

Show that there exists a subsequence (a_{n_k}) converging to $s = \sup S$. (This is a direct proof of the Bolzano–Weierstrass Theorem using the Axiom of Completeness.)

2.6 The Cauchy Criterion

The following definition bears a striking resemblance to the definition of convergence for a sequence.

Definition 2.6.1. A sequence (a_n) is called a *Cauchy sequence* if, for every $\epsilon > 0$, there exists an $N \in \mathbf{N}$ such that whenever $m, n \geq N$ it follows that $|a_n - a_m| < \epsilon$.

To make the comparison easier, let's restate the definition of convergence.

Definition 2.2.3 (Convergence of a Sequence). A sequence (a_n) *converges* to a real number a if, for every $\epsilon > 0$, there exists an $N \in \mathbf{N}$ such that whenever $n \geq N$ it follows that $|a_n - a| < \epsilon$.

As we have discussed, the definition of convergence asserts that, given an arbitrary positive ϵ , it is possible to find a point in the sequence after which the terms of the sequence are all closer to the limit a than the given ϵ . On

the other hand, a sequence is a Cauchy sequence if, for every ϵ , there is a point in the sequence after which the terms are all closer *to each other* than the given ϵ . To spoil the surprise, we will argue in this section that in fact these two definitions are equivalent: Convergent sequences are Cauchy sequences, and Cauchy sequences converge. The significance of the definition of a Cauchy sequence is that there is no mention of a limit. This is somewhat like the situation with the Monotone Convergence Theorem in that we will have another way of proving that sequences converge without having any explicit knowledge of what the limit might be.

Theorem 2.6.2. *Every convergent sequence is a Cauchy sequence.*

Proof. Assume (x_n) converges to x . To prove that (x_n) is Cauchy, we must find a point in the sequence after which we have $|x_n - x_m| < \epsilon$. This can be done using an application of the triangle inequality. The details are requested in Exercise 2.6.2. \square

The converse is a bit more difficult to prove, mainly because, in order to prove that a sequence converges, we must have a proposed limit for the sequence to approach. We have been in this situation before in the proofs of the Monotone Convergence Theorem and the Bolzano–Weierstrass Theorem. Our strategy here will be to use the Bolzano–Weierstrass Theorem. This is the reason for the next lemma. (Compare this with Theorem 2.3.2.)

Lemma 2.6.3. *Cauchy sequences are bounded.*

Proof. Given $\epsilon = 1$, there exists an N such that $|x_m - x_n| < 1$ for all $m, n \geq N$. Thus, we must have $|x_n| < |x_N| + 1$ for all $n \geq N$. It follows that

$$M = \max\{|x_1|, |x_2|, |x_3|, \dots, |x_{N-1}|, |x_N| + 1\}$$

is a bound for the sequence (x_n) . \square

Theorem 2.6.4 (Cauchy Criterion). *A sequence converges if and only if it is a Cauchy sequence.*

Proof. (\Rightarrow) This direction is Theorem 2.6.2.

(\Leftarrow) For this direction, we start with a Cauchy sequence (x_n) . Lemma 2.6.3 guarantees that (x_n) is bounded, so we may use the Bolzano–Weierstrass Theorem to produce a convergent subsequence (x_{n_k}) . Set

$$x = \lim x_{n_k}.$$

The idea is to show that the original sequence (x_n) converges to this same limit. Once again, we will use a triangle inequality argument. We know the terms in the subsequence are getting close to the limit x , and the assumption that (x_n) is Cauchy implies the terms in the “tail” of the sequence are close to each other. Thus, we want to make each of these distances less than half of the prescribed ϵ .

Let $\epsilon > 0$. Because (x_n) is Cauchy, there exists N such that

$$|x_n - x_m| < \frac{\epsilon}{2}$$

whenever $m, n \geq N$. Now, we also know that $(x_{n_k}) \rightarrow x$, so choose a term in this subsequence, call it x_{n_K} , with $n_K \geq N$ and

$$|x_{n_K} - x| < \frac{\epsilon}{2}.$$

To see that n_K has the desired property (for the original sequence (x_n)), observe that if $n \geq n_K$, then

$$\begin{aligned} |x_n - x| &= |x_n - x_{n_K} + x_{n_K} - x| \\ &\leq |x_n - x_{n_K}| + |x_{n_K} - x| \\ &< \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon. \end{aligned}$$

□

The Cauchy Criterion is named after the French mathematician Augustin Louis Cauchy. Cauchy is a major figure in the history of many branches of mathematics—number theory and the theory of finite groups, to name a few—but he is most widely recognized for his enormous contributions in analysis, especially complex analysis. He is deservedly credited with inventing the ϵ -based definition of limits we use today, although it is probably better to view him as a pioneer of analysis in the sense that his work did not attain the level of refinement that modern mathematicians have come to expect. The Cauchy Criterion, for instance, was devised and used by Cauchy to study infinite series, but he never actually proved it in both directions. The fact that there were gaps in Cauchy’s work should not diminish his brilliance in any way. The issues of the day were both difficult and subtle, and Cauchy was far and away the most influential in laying the groundwork for modern standards of rigor. Karl Weierstrass played a major role in sharpening Cauchy’s arguments. We will hear a good deal more from Weierstrass, most notably in Chapter 6 when we take up uniform convergence. Bernhard Bolzano was working in Prague and was writing and thinking about many of these same issues surrounding limits and continuity. For whatever reason, his historical reputation does not seem to do justice to the impressive level of his insights.

Completeness Revisited

In the first chapter, we established the Axiom of Completeness (AoC) to be the assertion that “sets bounded above have least upper bounds.” We then used this axiom as the crucial step in the proof of the Nested Interval Property (NIP). In this chapter, AoC was the central step in the Monotone Convergence Theorem (MCT), and NIP was central to proving the Bolzano–Weierstrass Theorem

(BW). Finally, we needed BW in our proof of the Cauchy Criterion (CC) for convergent sequences. The list of implications then looks like

$$\text{AoC} \Rightarrow \begin{cases} \text{NIP} \\ \text{MCT} \end{cases} \Rightarrow \text{BW} \Rightarrow \text{CC}.$$

But this list is not the whole story. Recall that in our original discussions about completeness, the fundamental problem was that the rational numbers contained “gaps.” The reason for moving from the rational numbers to the real numbers to do analysis is so that when we encounter a sequence of numbers that looks as if it is converging to some number—say $\sqrt{2}$ —then we can be assured that there is indeed a number there that we can call the limit. The assertion that “sets bounded above have least upper bounds” is simply one way to mathematically articulate our insistence that there be no “holes” in our number field, but it is not the only way. Instead, we could have taken NIP to be an axiom and used it to prove that least upper bounds exist, or we could have assumed MCT and used it to prove NIP and the rest of the results. In fact, *any* of these results could have served as our starting point. The situation is reminiscent of the adage, “Which came first, the chicken or the egg?” All of the preceding statements are equivalent in the sense that once we have assumed any one of them to be true, it is possible to derive the rest. (These implications are explored in Exercise 2.6.6.) However, because we have an example of an ordered field that is not complete—namely, the set of rational numbers—we know it is impossible to prove any of them using only the field and order properties. Just how we decide which should be the axiom and which then become theorems depends on preference and context and in the end does not matter all that much. What is important is that we understand all of these results—AoC, NIP, MCT, BW, and CC—as belonging to the same family, each asserting the completeness of \mathbf{R} in its own particular language.

Exercises

Exercise 2.6.1. Give an example of each of the following, or argue that such a request is impossible.

- A Cauchy sequence that is not monotone.
- A monotone sequence that is not Cauchy.
- A Cauchy sequence with a divergent subsequence.
- An unbounded sequence containing a subsequence that is Cauchy.

Exercise 2.6.2. Supply a proof for Theorem 2.6.2.

Exercise 2.6.3. (a) Explain how the following pseudo-Cauchy property differs from the proper definition of a Cauchy sequence: A sequence (s_n) is *pseudo-Cauchy* if, for all $\epsilon > 0$, there exists an N such that if $n \geq N$, then $|s_{n+1} - s_n| < \epsilon$.

(b) If possible, give an example of a divergent sequence (s_n) that is pseudo-Cauchy.

Exercise 2.6.4. Assume (a_n) and (b_n) are Cauchy sequences. Use a triangle inequality argument to prove $c_n = |a_n - b_n|$ is Cauchy.

Exercise 2.6.5. If (x_n) and (y_n) are Cauchy sequences, then one easy way to prove that $(x_n + y_n)$ is Cauchy is to use the Cauchy Criterion. By Theorem 2.6.4, (x_n) and (y_n) must be convergent, and the Algebraic Limit Theorem then implies $(x_n + y_n)$ is convergent and hence Cauchy.

(a) Give a direct argument that $(x_n + y_n)$ is a Cauchy sequence that does not use the Cauchy Criterion or the Algebraic Limit Theorem.

(b) Do the same for the product $(x_n y_n)$.

Exercise 2.6.6. (a) Assume the Nested Interval Property (Theorem 1.4.1) is true and use a technique similar to the one employed in the proof of the Bolzano–Weierstrass Theorem to give a proof for the Axiom of Completeness. (The reverse implication was given in Chapter 1. This shows that AoC is equivalent to NIP.)

(b) Use the Monotone Convergence Theorem to give a proof of the Nested Interval Property. (This establishes the equivalence of AoC, NIP, and MCT.)

(c) This time, start with the Bolzano–Weierstrass Theorem and use it to construct a proof of the Nested Interval Property. (Thus, BW is equivalent to NIP, and hence to AoC and MCT as well.)

(d) Finally, use the Cauchy Criterion to prove the Bolzano–Weierstrass Theorem. This is the final link in the equivalence of the five characterizations of completeness discussed at the end of Section 2.6.

2.7 Properties of Infinite Series

Given an infinite series $\sum_{k=1}^{\infty} a_k$, it is important to keep a clear distinction between

(i) the sequence of *terms*: (a_1, a_2, a_3, \dots) and

(ii) the sequence of *partial sums*: (s_1, s_2, s_3, \dots) , where $s_n = a_1 + a_2 + \dots + a_n$.

The convergence of the series $\sum_{k=1}^{\infty} a_k$ is defined in terms of the sequence (s_n) . Specifically, the statement

$$\sum_{k=1}^{\infty} a_k = A \quad \text{means that} \quad \lim s_n = A.$$

It is for this reason that we can immediately translate many of our results from the study of sequences into statements about the behavior of infinite series.

Theorem 2.7.1 (Algebraic Limit Theorem for Series). *If $\sum_{k=1}^{\infty} a_k = A$ and $\sum_{k=1}^{\infty} b_k = B$, then*

(i) $\sum_{k=1}^{\infty} ca_k = cA$ for all $c \in \mathbf{R}$ and

$$(ii) \sum_{k=1}^{\infty} (a_k + b_k) = A + B.$$

Proof. (i) In order to show that $\sum_{k=1}^{\infty} ca_k = cA$, we must argue that the sequence of partial sums

$$t_m = ca_1 + ca_2 + ca_3 + \cdots + ca_m$$

converges to cA . But we are given that $\sum_{k=1}^{\infty} a_k$ converges to A , meaning that the partial sums

$$s_m = a_1 + a_2 + a_3 + \cdots + a_m$$

converge to A . Because $t_m = cs_m$, applying the Algebraic Limit Theorem for sequences (Theorem 2.3.3) yields $(t_m) \rightarrow cA$, as desired.

(ii) Exercise 2.7.8. □

One way to summarize Theorem 2.7.1 (i) is to say that infinite addition still satisfies the distributive property. Part (ii) verifies that series can be added in the usual way. Missing from this theorem is any statement about the *product* of two infinite series. At the heart of this question is the issue of commutativity, which requires a more delicate analysis and so is postponed until Section 2.8.

Theorem 2.7.2 (Cauchy Criterion for Series). *The series $\sum_{k=1}^{\infty} a_k$ converges if and only if, given $\epsilon > 0$, there exists an $N \in \mathbf{N}$ such that whenever $n > m \geq N$ it follows that*

$$|a_{m+1} + a_{m+2} + \cdots + a_n| < \epsilon.$$

Proof. Observe that

$$|s_n - s_m| = |a_{m+1} + a_{m+2} + \cdots + a_n|$$

and apply the Cauchy Criterion for sequences. □

The Cauchy Criterion leads to economical proofs of several basic facts about series.

Theorem 2.7.3. *If the series $\sum_{k=1}^{\infty} a_k$ converges, then $(a_k) \rightarrow 0$.*

Proof. Consider the special case $n = m + 1$ in the Cauchy Criterion for Convergent Series. □

Every statement of this result should be accompanied with a reminder to look at the harmonic series (Example 2.4.5) to erase any misconception that the converse statement is true. Knowing (a_k) tends to 0 does not imply that the series converges.

Theorem 2.7.4 (Comparison Test). *Assume (a_k) and (b_k) are sequences satisfying $0 \leq a_k \leq b_k$ for all $k \in \mathbf{N}$.*

(i) *If $\sum_{k=1}^{\infty} b_k$ converges, then $\sum_{k=1}^{\infty} a_k$ converges.*

(ii) If $\sum_{k=1}^{\infty} a_k$ diverges, then $\sum_{k=1}^{\infty} b_k$ diverges.

Proof. Both statements follow immediately from the Cauchy Criterion for Series and the observation that

$$|a_{m+1} + a_{m+2} + \cdots + a_n| \leq |b_{m+1} + b_{m+2} + \cdots + b_n|.$$

Alternate proofs using the Monotone Convergence Theorem are requested in the exercises. \square

This is a good point to remind ourselves again that statements about convergence of sequences and series are immune to changes in some finite number of initial terms. In the Comparison Test, the requirement that $a_k \leq b_k$ does not really need to hold for *all* $k \in \mathbf{N}$ but just needs to be *eventually* true. A weaker, but sufficient, hypothesis would be to assume that there exists some point $M \in \mathbf{N}$ such that the inequality $a_k \leq b_k$ is true for all $k \geq M$.

The Comparison Test is used to deduce the convergence or divergence of one series based on the behavior of another. Thus, for this test to be of any great use, we need a catalog of series we can use as measuring sticks. In Section 2.4, we proved the Cauchy Condensation Test, which led to the general statement that the series $\sum_{n=1}^{\infty} 1/n^p$ converges if and only if $p > 1$.

The next example summarizes the situation for another important class of series.

Example 2.7.5 (Geometric Series). A series is called *geometric* if it is of the form

$$\sum_{k=0}^{\infty} ar^k = a + ar + ar^2 + ar^3 + \cdots .$$

If $r = 1$ and $a \neq 0$, the series evidently diverges. For $r \neq 1$, the algebraic identity

$$(1 - r)(1 + r + r^2 + r^3 + \cdots + r^{m-1}) = 1 - r^m$$

enables us to rewrite the partial sum

$$s_m = a + ar + ar^2 + ar^3 + \cdots + ar^{m-1} = \frac{a(1 - r^m)}{1 - r}.$$

Now the Algebraic Limit Theorem (for sequences) and Example 2.5.3 justify the conclusion

$$\sum_{k=0}^{\infty} ar^k = \frac{a}{1 - r}$$

if and only if $|r| < 1$.

Although the Comparison Test requires that the terms of the series be positive, it is often used in conjunction with the next theorem to handle series that contain some negative terms.

Theorem 2.7.6 (Absolute Convergence Test). *If the series $\sum_{n=1}^{\infty} |a_n|$ converges, then $\sum_{n=1}^{\infty} a_n$ converges as well.*

Proof. This proof makes use of both the necessity (the “if” direction) and the sufficiency (the “only if” direction) of the Cauchy Criterion for Series. Because $\sum_{n=1}^{\infty} |a_n|$ converges, we know that, given an $\epsilon > 0$, there exists an $N \in \mathbf{N}$ such that

$$|a_{m+1}| + |a_{m+2}| + \cdots + |a_n| < \epsilon$$

for all $n > m \geq N$. By the triangle inequality,

$$|a_{m+1} + a_{m+2} + \cdots + a_n| \leq |a_{m+1}| + |a_{m+2}| + \cdots + |a_n|,$$

so the sufficiency of the Cauchy Criterion guarantees that $\sum_{n=1}^{\infty} a_n$ also converges. \square

The converse of this theorem is false. In the opening discussion of this chapter, we considered the *alternating harmonic series*

$$1 - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \frac{1}{5} - \frac{1}{6} + \cdots .$$

Taking absolute values of the terms gives us the harmonic series $\sum_{n=1}^{\infty} 1/n$, which we have seen diverges. However, it is not too difficult to prove that with the alternating negative signs the series indeed converges. This is a special case of the Alternating Series Test, whose proof is outlined in Exercise 2.7.1.

Theorem 2.7.7 (Alternating Series Test). *Let (a_n) be a sequence satisfying,*

- (i) $a_1 \geq a_2 \geq a_3 \geq \cdots \geq a_n \geq a_{n+1} \geq \cdots$ and
- (ii) $(a_n) \rightarrow 0$.

Then, the alternating series $\sum_{n=1}^{\infty} (-1)^{n+1} a_n$ converges.

Proof. Exercise 2.7.1. \square

Definition 2.7.8. If $\sum_{n=1}^{\infty} |a_n|$ converges, then we say that the original series $\sum_{n=1}^{\infty} a_n$ *converges absolutely*. If, on the other hand, the series $\sum_{n=1}^{\infty} a_n$ converges but the series of absolute values $\sum_{n=1}^{\infty} |a_n|$ does not converge, then we say that the original series $\sum_{n=1}^{\infty} a_n$ *converges conditionally*.

In terms of this newly defined jargon, we have shown that

$$\sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{n}$$

converges conditionally, whereas

$$\sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{n^2}, \quad \sum_{n=1}^{\infty} \frac{1}{2^n} \quad \text{and} \quad \sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{2^n}$$

converge absolutely. In particular, any convergent series with (all but finitely many) positive terms must converge absolutely.

The Alternating Series Test is the most accessible test for conditional convergence, but several others are explored in the exercises. In particular, Abel's Test, outlined in Exercise 2.7.14, will prove useful in our investigations of power series in Chapter 6.

Rearrangements

Informally speaking, a rearrangement of a series is obtained by permuting the terms in the sum into some other order. It is important that all of the original terms eventually appear in the new ordering and that no term gets repeated. In an earlier discussion, we formed a rearrangement of the alternating harmonic series by taking two positive terms for each negative term:

$$1 + \frac{1}{3} - \frac{1}{2} + \frac{1}{5} + \frac{1}{7} - \frac{1}{4} + \cdots$$

There are clearly an infinite number of rearrangements of any sum; however, it is helpful to see why neither

$$1 + \frac{1}{2} - \frac{1}{3} + \frac{1}{4} + \frac{1}{5} - \frac{1}{6} + \cdots$$

nor

$$1 + \frac{1}{3} - \frac{1}{4} + \frac{1}{5} + \frac{1}{7} - \frac{1}{8} + \frac{1}{9} + \frac{1}{11} - \frac{1}{12} + \cdots$$

is considered a rearrangement of the original alternating harmonic series.

Definition 2.7.9. Let $\sum_{k=1}^{\infty} a_k$ be a series. A series $\sum_{k=1}^{\infty} b_k$ is called a *rearrangement* of $\sum_{k=1}^{\infty} a_k$ if there exists a one-to-one, onto function $f : \mathbf{N} \rightarrow \mathbf{N}$ such that $b_{f(k)} = a_k$ for all $k \in \mathbf{N}$.

We now have all the tools and notation in place to resolve an issue raised at the beginning of the chapter. In Section 2.1, we constructed a particular rearrangement of the alternating harmonic series that converges to a limit different from that of the original series. This happens because the convergence is *conditional*.

Theorem 2.7.10. *If $\sum_{k=1}^{\infty} a_k$ converges absolutely, then any rearrangement of this series converges to the same limit.*

Proof. Assume $\sum_{k=1}^{\infty} a_k$ converges absolutely to A , and let $\sum_{k=1}^{\infty} b_k$ be a rearrangement of $\sum_{k=1}^{\infty} a_k$. Let's use

$$s_n = \sum_{k=1}^n a_k = a_1 + a_2 + \cdots + a_n$$

for the partial sums of the original series and use

$$t_m = \sum_{k=1}^m b_k = b_1 + b_2 + \cdots + b_m$$

for the partial sums of the rearranged series. Thus we want to show that $(t_m) \rightarrow A$.

Let $\epsilon > 0$. By hypothesis, $(s_n) \rightarrow A$, so choose N_1 such that

$$|s_n - A| < \frac{\epsilon}{2}$$

for all $n \geq N_1$. Because the convergence is absolute, we can choose N_2 so that

$$\sum_{k=m+1}^n |a_k| < \frac{\epsilon}{2}$$

for all $n > m \geq N_2$. Now, take $N = \max\{N_1, N_2\}$. We know that the finite set of terms $\{a_1, a_2, a_3, \dots, a_N\}$ must all appear in the rearranged series, and we want to move far enough out in the series $\sum_{n=1}^{\infty} b_n$ so that we have included all of these terms. Thus, choose

$$M = \max\{f(k) : 1 \leq k \leq N\}.$$

It should now be evident that if $m \geq M$, then $(t_m - s_N)$ consists of a finite set of terms, the absolute values of which appear in the tail $\sum_{k=N+1}^{\infty} |a_k|$. Our choice of N_2 earlier then guarantees $|t_m - s_N| < \epsilon/2$, and so

$$\begin{aligned} |t_m - A| &= |t_m - s_N + s_N - A| \\ &\leq |t_m - s_N| + |s_N - A| \\ &< \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon \end{aligned}$$

whenever $m \geq M$. □

Exercises

Exercise 2.7.1. Proving the Alternating Series Test (Theorem 2.7.7) amounts to showing that the sequence of partial sums

$$s_n = a_1 - a_2 + a_3 - \cdots \pm a_n$$

converges. (The opening example in Section 2.1 includes a typical illustration of (s_n) .) Different characterizations of completeness lead to different proofs.

(a) Prove the Alternating Series Test by showing that (s_n) is a Cauchy sequence.

(b) Supply another proof for this result using the Nested Interval Property (Theorem 1.4.1).

(c) Consider the subsequences (s_{2n}) and (s_{2n+1}) , and show how the Monotone Convergence Theorem leads to a third proof for the Alternating Series Test.

Exercise 2.7.2. (a) Provide the details for the proof of the Comparison Test (Theorem 2.7.4) using the Cauchy Criterion for series.

(b) Give another proof for the Comparison Test, this time using the Monotone Convergence Theorem.

Exercise 2.7.3. Let $\sum a_n$ be given. For each $n \in \mathbf{N}$, let $p_n = a_n$ if a_n is positive and assign $p_n = 0$ if a_n is negative. In a similar manner, let $q_n = a_n$ if a_n is negative and $q_n = 0$ if a_n is positive.

(a) Argue that if $\sum a_n$ diverges, then at least one of $\sum p_n$ or $\sum q_n$ diverges.

(b) Show that if $\sum a_n$ converges conditionally, then *both* $\sum p_n$ and $\sum q_n$ diverge.

Exercise 2.7.4. Give an example to show that it is possible for both $\sum x_n$ and $\sum y_n$ to diverge but for $\sum x_n y_n$ to converge.

Exercise 2.7.5. (a) Show that if $\sum a_n$ converges absolutely, then $\sum a_n^2$ also converges absolutely. Does this proposition hold without absolute convergence?

(b) If $\sum a_n$ converges and $a_n \geq 0$, can we conclude anything about $\sum \sqrt{a_n}$?

Exercise 2.7.6. (a) Show that if $\sum x_n$ converges absolutely, and the sequence (y_n) is bounded, then the sum $\sum x_n y_n$ converges.

(b) Find a counterexample that demonstrates that part (a) does not always hold if the convergence of $\sum x_n$ is conditional.

Exercise 2.7.7. Now that we have proved the basic facts about geometric series, supply a proof for Corollary 2.4.7.

Exercise 2.7.8. Prove Theorem 2.7.1 part (ii).

Exercise 2.7.9 (Ratio Test). Given a series $\sum_{n=1}^{\infty} a_n$ with $a_n \neq 0$, the Ratio Test states that if (a_n) satisfies

$$\lim \left| \frac{a_{n+1}}{a_n} \right| = r < 1,$$

then the series converges absolutely.

(a) Let r' satisfy $r < r' < 1$. (Why must such an r' exist?) Explain why there exists an N such that $n \geq N$ implies $|a_{n+1}| \leq |a_n| r'$.

(b) Why does $|a_N| \sum (r')^n$ necessarily converge?

(c) Now, show that $\sum |a_n|$ converges.

Exercise 2.7.10. (a) Show that if $a_n > 0$ and $\lim(na_n) = l$ with $l \neq 0$, then the series $\sum a_n$ diverges.

(b) Assume $a_n > 0$ and $\lim(n^2 a_n)$ exists. Show that $\sum a_n$ converges.

Exercise 2.7.11. Find examples of two series $\sum a_n$ and $\sum b_n$ both of which diverge but for which $\sum \min\{a_n, b_n\}$ converges. To make it more challenging, produce examples where (a_n) and (b_n) are positive and decreasing.

Exercise 2.7.12 (Summation by Parts). Let (x_n) and (y_n) be sequences, and let $s_n = x_1 + x_2 + \cdots + x_n$. Use the observation that $x_j = s_j - s_{j-1}$ to verify the formula

$$\sum_{j=m+1}^n x_j y_j = s_n y_{n+1} - s_m y_{m+1} + \sum_{j=m+1}^n s_j (y_j - y_{j+1}).$$

Exercise 2.7.13 (Dirichlet's Test). Dirichlet's Test for convergence states that if the partial sums of $\sum_{n=1}^{\infty} x_n$ are bounded (but not necessarily convergent), and if (y_n) is a sequence satisfying $y_1 \geq y_2 \geq y_3 \geq \cdots \geq 0$ with $\lim y_n = 0$, then the series $\sum_{n=1}^{\infty} x_n y_n$ converges.

(a) Let $M > 0$ be an upper bound for the partial sums of $\sum_{n=1}^{\infty} x_n$. Use Exercise 2.7.12 to show that

$$\left| \sum_{j=m+1}^n x_j y_j \right| \leq 2M|y_{m+1}|.$$

(b) Prove Dirichlet's Test just stated.

(c) Show how the Alternating Series Test (Theorem 2.7.7) can be derived as a special case of Dirichlet's Test.

Exercise 2.7.14 (Abel's Test). Abel's Test for convergence states that if the series $\sum_{n=1}^{\infty} x_n$ converges, and if (y_n) is a sequence satisfying

$$y_1 \geq y_2 \geq y_3 \geq \cdots \geq 0,$$

then the series $\sum_{n=1}^{\infty} x_n y_n$ converges.

(a) Carefully point out how the hypothesis of Abel's Test differs from that of Dirichlet's Test in Exercise 2.7.13.

(b) Assume that $\sum_{n=1}^{\infty} a_n$ has partial sums that are bounded by a constant $A > 0$, and assume $b_1 \geq b_2 \geq b_3 \geq \cdots \geq 0$. Use Exercise 2.7.12 to show that

$$\left| \sum_{j=1}^n a_j b_j \right| \leq 2Ab_1.$$

(c) Prove Abel's Test via the following strategy. For a fixed $m \in \mathbf{N}$, apply part (b) to $\sum_{j=m+1}^n x_j y_j$ by setting $a_n = x_{m+n}$ and $b_n = y_{m+n}$. (Argue that an upper bound on the partial sums of $\sum_{n=1}^{\infty} a_n$ can be made arbitrarily small by taking m to be large.)

2.8 Double Summations and Products of Infinite Series

Given a doubly indexed array of real numbers $\{a_{ij} : i, j \in \mathbf{N}\}$, we discovered in Section 2.1 that there is a dangerous ambiguity in how we might define $\sum_{i,j=1}^{\infty} a_{ij}$. Performing the sum over first one of the variables and then the other is referred to as an *iterated* summation. In our specific example, summing the rows first and then taking the sum of these totals produced a different result than first computing the sum of each column and adding these sums together. In short,

$$\sum_{j=1}^{\infty} \sum_{i=1}^{\infty} a_{ij} \neq \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} a_{ij}.$$

There are still other ways to reasonably define $\sum_{i,j=1}^{\infty} a_{ij}$. One natural idea is to calculate a kind of partial sum by adding together finite numbers of terms in larger and larger “rectangles” in the array; that is, for $m, n \in \mathbf{N}$, set

$$(1) \quad s_{mn} = \sum_{i=1}^m \sum_{j=1}^n a_{ij}.$$

The order of the sum here is irrelevant because the sum is finite. Of particular interest to our discussion are the sums s_{nn} (sums over “squares”), which form a legitimate sequence indexed by n and thus can be subjected to our arsenal of theorems and definitions. If the sequence (s_{nn}) converges, for instance, we might wish to define

$$\sum_{i,j=1}^{\infty} a_{ij} = \lim_{n \rightarrow \infty} s_{nn}.$$

Exercise 2.8.1. Using the particular array (a_{ij}) from Section 2.1, compute $\lim_{n \rightarrow \infty} s_{nn}$. How does this value compare to the two iterated values for the sum already computed?

There is a deep similarity between the issue of how to define a double summation and the topic of rearrangements discussed at the end of Section 2.7. Both relate to the commutativity of addition in an infinite setting. For rearrangements, the resolution came with the added hypothesis of *absolute* convergence, and it is not surprising that the same remedy applies for double summations. Under the assumption of absolute convergence, each of the methods discussed for computing the value of a double sum yields the same result.

Exercise 2.8.2. Show that if the iterated series

$$\sum_{i=1}^{\infty} \sum_{j=1}^{\infty} |a_{ij}|$$

converges (meaning that for each fixed $i \in \mathbf{N}$ the series $\sum_{j=1}^{\infty} |a_{ij}|$ converges to some real number b_i , and the series $\sum_{i=1}^{\infty} b_i$ converges as well), then the iterated series

$$\sum_{i=1}^{\infty} \sum_{j=1}^{\infty} a_{ij}$$

converges.

Theorem 2.8.1. Let $\{a_{ij} : i, j \in \mathbf{N}\}$ be a doubly indexed array of real numbers. If

$$\sum_{i=1}^{\infty} \sum_{j=1}^{\infty} |a_{ij}|$$

converges, then both $\sum_{i=1}^{\infty} \sum_{j=1}^{\infty} a_{ij}$ and $\sum_{j=1}^{\infty} \sum_{i=1}^{\infty} a_{ij}$ converge to the same value. Moreover,

$$\lim_{n \rightarrow \infty} s_{nn} = \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} a_{ij} = \sum_{j=1}^{\infty} \sum_{i=1}^{\infty} a_{ij},$$

where $s_{nn} = \sum_{i=1}^n \sum_{j=1}^n a_{ij}$.

Proof. In the same way that we defined the “rectangular partial sums” s_{mn} above in equation (1), define

$$t_{mn} = \sum_{i=1}^m \sum_{j=1}^n |a_{ij}|.$$

Exercise 2.8.3. (a) Prove that the set $\{t_{mn} : m, n \in \mathbf{N}\}$ is bounded above, and use this fact to conclude that the sequence (t_{nn}) converges.

(b) Now, use the fact that (t_{nn}) is a Cauchy sequence to argue that (s_{nn}) is a Cauchy sequence and hence converges.

We can now set

$$S = \lim_{n \rightarrow \infty} s_{nn}.$$

In order to prove the theorem, we must show that the two iterated sums converge to this same limit. We will first show that

$$S = \sum_{i=1}^{\infty} \sum_{j=1}^{\infty} a_{ij}.$$

Because $\{t_{mn} : m, n \in \mathbf{N}\}$ is bounded above, we can let

$$B = \sup\{t_{mn} : m, n \in \mathbf{N}\}.$$

Let $\epsilon > 0$ be arbitrary. Because B is the least upper bound for this set, we know there exists a particular $t_{m_0 n_0}$ that satisfies

$$B - \frac{\epsilon}{2} < t_{m_0 n_0} \leq B.$$

Exercise 2.8.4. (a) Argue that there exists an $N_1 \in \mathbf{N}$ such that $m, n \geq N_1$ implies $B - \frac{\epsilon}{2} < t_{mn} \leq B$.

(b) Now, show that there exists an N such that

$$|s_{mn} - S| < \epsilon$$

for all $m, n \geq N$.

For the moment, consider $m \in \mathbf{N}$ to be fixed and write s_{mn} as

$$s_{mn} = \sum_{j=1}^n a_{1j} + \sum_{j=1}^n a_{2j} + \cdots + \sum_{j=1}^n a_{mj}.$$

Our hypothesis guarantees that for each fixed row i , the series $\sum_{j=1}^{\infty} a_{ij}$ converges absolutely to some real number r_i .

Exercise 2.8.5. (a) Use the Algebraic Limit Theorem (Theorem 2.3.3) and the Order Limit Theorem (Theorem 2.3.4) to show that for all $m \geq N$

$$|(r_1 + r_2 + \cdots + r_m) - S| \leq \epsilon.$$

Conclude that the iterated sum $\sum_{i=1}^{\infty} \sum_{j=1}^{\infty} a_{ij}$ converges to S .

Exercise 2.8.6. Finish the proof by showing that the other iterated sum, $\sum_{j=1}^{\infty} \sum_{i=1}^{\infty} a_{ij}$, converges to S as well. Notice that the same argument can be used once it is established that, for each fixed column j , the sum $\sum_{i=1}^{\infty} a_{ij}$ converges to some real number c_j .

□

One final common way of computing a double summation is to sum along diagonals where $i + j$ equals a constant. Given a doubly indexed array $\{a_{ij} : i, j \in N\}$, let

$$d_2 = a_{11}, \quad d_3 = a_{12} + a_{21}, \quad d_4 = a_{13} + a_{22} + a_{31},$$

and in general set

$$d_k = a_{1,k-1} + a_{2,k-2} + \cdots + a_{k-1,1}.$$

Then, $\sum_{k=2}^{\infty} d_k$ represents another reasonable way of summing over every a_{ij} in the array.

Exercise 2.8.7. (a) Assuming the hypothesis—and hence the conclusion—of Theorem 2.8.1, show that $\sum_{k=2}^{\infty} d_k$ converges absolutely.

(b) Imitate the strategy in the proof of Theorem 2.8.1 to show that $\sum_{k=2}^{\infty} d_k$ converges to $S = \lim_{n \rightarrow \infty} s_{nn}$.

Products of Series

Conspicuously missing from the Algebraic Limit Theorem for Series (Theorem 2.7.1) is any statement about the product of two convergent series. One way to formally carry out the algebra on such a product is to write

$$\begin{aligned} \left(\sum_{i=1}^{\infty} a_i \right) \left(\sum_{j=1}^{\infty} b_j \right) &= (a_1 + a_2 + a_3 + \cdots)(b_1 + b_2 + b_3 + \cdots) \\ &= a_1 b_1 + (a_1 b_2 + a_2 b_1) + (a_3 b_1 + a_2 b_2 + a_1 b_3) + \cdots \\ &= \sum_{k=2}^{\infty} d_k, \end{aligned}$$

where

$$d_k = a_1 b_{k-1} + a_2 b_{k-2} + \cdots + a_{k-1} b_1.$$

This particular form of the product, examined earlier in Exercise 2.8.7, is called the *Cauchy product* of two series. Although there is something algebraically natural about writing the product in this form, it may very well be that computing the value of the sum is more easily done via one or the other iterated summation. The question remains, then, as to how the value of the Cauchy product—if it exists—is related to these other values of the double sum. If the two series being multiplied converge absolutely, it is not too difficult to prove that the sum may be computed in whatever way is most convenient.

Exercise 2.8.8. Assume that $\sum_{i=1}^{\infty} a_i$ converges absolutely to A , and $\sum_{j=1}^{\infty} b_j$ converges absolutely to B .

(a) Show that the set

$$\left\{ \sum_{i=1}^m \sum_{j=1}^n |a_i b_j| : m, n \in \mathbf{N} \right\}$$

is bounded. Use this to show that the iterated sum $\sum_{i=1}^{\infty} \sum_{j=1}^{\infty} |a_i b_j|$ converges so that we may apply Theorem 2.8.1.

(b) Let $s_{nn} = \sum_{i=1}^n \sum_{j=1}^n a_i b_j$, and use the Algebraic Limit Theorem to show that $\lim_{n \rightarrow \infty} s_{nn} = AB$. Conclude that

$$\sum_{i=1}^{\infty} \sum_{j=1}^{\infty} a_i b_j = \sum_{j=1}^{\infty} \sum_{i=1}^{\infty} a_i b_j = \sum_{k=2}^{\infty} d_k = AB,$$

where, as before, $d_k = a_1 b_{k-1} + a_2 b_{k-2} + \cdots + a_{k-1} b_1$.

2.9 Epilogue

Theorems 2.7.10 and 2.8.1 make it clear that absolute convergence is an extremely desirable quality to have when manipulating series. On the other hand, the situation for conditionally convergent series is delightfully pathological. In the case of rearrangements, not only are they no longer guaranteed to converge to the same limit, but in fact if $\sum_{n=1}^{\infty} a_n$ converges conditionally, then for *any* $r \in \mathbf{R}$ there exists a rearrangement of $\sum_{n=1}^{\infty} a_n$ that converges to r . To see why, let's look again at the alternating harmonic series

$$\sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{n}.$$

The negative terms taken alone form the series $\sum_{n=1}^{\infty} (-1)/2n$. The partial sums of this series are precisely $-1/2$ the partial sums of the harmonic series, and so march off (at half speed) to negative infinity. A similar argument shows that the sum of positive terms $\sum_{n=1}^{\infty} 1/(2n-1)$ also diverges to infinity. It is not too difficult to argue that this situation is *always* the case for conditionally convergent series (Exercise 2.7.3). Now, let r be some proposed limit, which, for

the sake of this argument, we take to be positive. The idea is to take as many positive terms as necessary to form the first partial sum greater than r . We then add negative terms until the partial sum falls below r , at which point we switch back to positive terms. The fact that there is no bound on the sums of either the positive terms or the negative terms allows this process to continue indefinitely. The fact that the terms themselves tend to zero is enough to guarantee that the partial sums, when constructed in this manner, indeed converge to r as they oscillate around this target value.

Perhaps the best way to summarize the situation is to say that the hypothesis of absolute convergence essentially allows us to treat infinite sums as though they were finite sums. This assessment extends to double sums as well, although there are a few subtleties to address. In the case of products, we showed in Exercise 2.8.8 that the Cauchy product of two absolutely convergent infinite series converges to the product of the two factors, but in fact the same conclusion follows if we only have absolute convergence in one of the two original series. In the notation of Exercise 2.8.8, if $\sum a_n$ converges absolutely to A , and if $\sum b_n$ converges (perhaps conditionally) to B , then the Cauchy product $\sum d_k = AB$. On the other hand, if both $\sum a_n$ and $\sum b_n$ converge conditionally, then it is possible for the Cauchy product to diverge. Squaring $\sum (-1)^n/\sqrt{n}$ provides an example of this phenomenon. Of course, it is also possible to find $\sum a_n = A$ conditionally and $\sum b_n = B$ conditionally whose Cauchy product $\sum d_k$ converges. If this is the case, then the convergence is to the right value, namely $\sum d_k = AB$. A proof of this last fact will be offered in Chapter 6, where we undertake the study of *power series*. Here is the connection. A power series has the form $a_0 + a_1x + a_2x^2 + \dots$. If we multiply two power series together as though they were polynomials, then when we collect common powers of x the result is

$$\begin{aligned} & (a_0 + a_1x + a_2x^2 + \dots)(b_0 + b_1x + b_2x^2 + \dots) \\ &= a_0b_0 + (a_0b_1 + a_1b_0)x + (a_0b_2 + a_1b_1 + a_2b_0)x^2 + \dots \\ &= d_0 + d_1x + d_2x^2 + \dots, \end{aligned}$$

which is the Cauchy product of $\sum a_nx^n$ and $\sum b_nx^n$. (The index starts with $n = 0$ rather than $n = 1$.) Upcoming results about the good behavior of power series will lead to a proof that convergent Cauchy products sum to the proper value. In the other direction, Exercise 2.8.8 will be useful in establishing a theorem about the product of two power series.

Chapter 3

Basic Topology of \mathbf{R}

3.1 Discussion: The Cantor Set

What follows is a fascinating mathematical construction, due to Georg Cantor, which is extremely useful for extending the horizons of our intuition about the nature of subsets of the real line. Cantor's name has already appeared in the first chapter in our discussion of uncountable sets. Indeed, Cantor's proof that \mathbf{R} is uncountable occupies another spot on the short list of the most significant contributions toward understanding the mathematical infinite. In the words of the mathematician David Hilbert, "No one shall expel us from the paradise that Cantor has created for us."

Let C_0 be the closed interval $[0, 1]$, and define C_1 to be the set that results when the open middle third is removed; that is,

$$C_1 = C_0 \setminus \left(\frac{1}{3}, \frac{2}{3} \right) = \left[0, \frac{1}{3} \right] \cup \left[\frac{2}{3}, 1 \right].$$

Now, construct C_2 in a similar way by removing the open middle third of each of the two components of C_1 :

$$C_2 = \left(\left[0, \frac{1}{9} \right] \cup \left[\frac{2}{9}, \frac{1}{3} \right] \right) \cup \left(\left[\frac{2}{3}, \frac{7}{9} \right] \cup \left[\frac{8}{9}, 1 \right] \right).$$

If we continue this process inductively, then for each $n = 0, 1, 2, \dots$ we get a set C_n consisting of 2^n closed intervals each having length $1/3^n$. Finally, we define the *Cantor set* C (Fig. 3.1) to be the intersection

$$C = \bigcap_{n=0}^{\infty} C_n.$$

It may be useful to understand C as the remainder of the interval $[0, 1]$ after the iterative process of removing open middle thirds is taken to infinity:

$$C = [0, 1] \setminus \left[\left(\frac{1}{3}, \frac{2}{3} \right) \cup \left(\frac{1}{9}, \frac{2}{9} \right) \cup \left(\frac{7}{9}, \frac{8}{9} \right) \cup \dots \right].$$




Figure 3.1: DEFINING THE CANTOR SET; $C = \bigcap_{n=0}^{\infty} C_n$.

There is some initial doubt whether anything remains at all, but notice that because we are always removing *open* middle thirds, then for every $n \in \mathbf{N}$, $0 \in C_n$ and hence $0 \in C$. The same argument shows $1 \in C$. In fact, if y is the endpoint of some closed interval of some particular set C_n , then it is also an endpoint of one of the intervals of C_{n+1} . Because, at each stage, endpoints are never removed, it follows that $y \in C_n$ for all n . Thus, C at least contains the endpoints of all of the intervals that make up each of the sets C_n .

Is there anything else? Is C countable? Does C contain any intervals? Any irrational numbers? These are difficult questions at the moment. All of the endpoints mentioned earlier are rational numbers (they have the form $m/3^n$), which means that if it is true that C consists of only these endpoints, then C would be a subset of \mathbf{Q} and hence countable. We shall see about this. There is some strong evidence that not much is left in C if we consider the total length of the intervals removed. To form C_1 , an open interval of length $1/3$ was taken out. In the second step, we removed two intervals of length $1/9$, and to construct C_n we removed 2^{n-1} middle thirds of length $1/3^n$. There is some logic, then, to defining the “length” of C to be 1 minus the total

$$\frac{1}{3} + 2 \left(\frac{1}{9} \right) + 4 \left(\frac{1}{27} \right) + \cdots + 2^{n-1} \left(\frac{1}{3^n} \right) + \cdots = \frac{\frac{1}{3}}{1 - \frac{2}{3}} = 1.$$

The Cantor set has *zero length*.

To this point, the information we have collected suggests a mental picture of C as a relatively small, thin set. For these reasons, the set C is often referred to as Cantor “dust.” But there are some strong counterarguments that imply a very different picture. First, C is actually *uncountable*, with cardinality equal to the cardinality of \mathbf{R} . One slightly intuitive but convincing way to see this is to create a 1–1 correspondence between C and sequences of the form $(a_n)_{n=1}^{\infty}$, where $a_n = 0$ or 1. For each $c \in C$, set $a_1 = 0$ if c falls in the left-hand component of C_1 and set $a_1 = 1$ if c falls in the right-hand component. Having established where in C_1 the point c is located, there are now two possible components of C_2 that might contain c . This time, we set $a_2 = 0$ or 1 depending on whether c falls in the left or right half of these two components of C_2 . Continuing in this




Figure 3.2: MAGNIFYING SETS BY A FACTOR OF 3.

way, we come to see that every element $c \in C$ yields a sequence (a_1, a_2, a_3, \dots) of zeros and ones that acts as a set of directions for how to locate c within C . Likewise, every such sequence corresponds to a point in the Cantor set. Because the set of sequences of zeros and ones is uncountable (Exercise 1.5.4), we must conclude that C is uncountable as well.

What does this imply? In the first place, because the endpoints of the approximating sets C_n form a countable set, we are forced to accept the fact that not only are there other points in C but there are uncountably many of them. From the point of view of *cardinality*, C is quite large—as large as \mathcal{R} , in fact. This should be contrasted with the fact that from the point of view of *length*, C measures the same size as a single point. We conclude this discussion with a demonstration that from the point of view of *dimension*, C strangely falls somewhere in between.

There is a sensible agreement that a point has dimension zero, a line segment has dimension one, a square has dimension two, and a cube has dimension three. Without attempting a formal definition of dimension (of which there are several), we can nevertheless get a sense of how one might be defined by observing how the dimension affects the result of magnifying each particular set by a factor of 3 (Fig. 3.2). (The reason for the choice of 3 will become clear when we turn our attention back to the Cantor set). A single point undergoes no change at all, whereas a line segment triples in length. For the square, magnifying each length by a factor of 3 results in a larger square that contains 9 copies of the original square. Finally, the magnified cube yields a cube that contains 27 copies of the original cube within its volume. Notice that, in each case, to compute the “size” of the new set, the dimension appears as the exponent of the magnification factor.

Now, apply this transformation to the Cantor set. The set $C_0 = [0, 1]$ becomes the interval $[0, 3]$. Deleting the middle third leaves $[0, 1] \cup [2, 3]$, which is where we started in the original construction except that we now stand to produce an additional copy of C in the interval $[2, 3]$. Magnifying the Cantor set by a factor of 3 yields *two* copies of the original set. Thus, if x is the dimension of C , then x should satisfy $2 = 3^x$, or $x = \ln 2 / \ln 3 \approx .631$ (Fig. 3.3).

	dim	$\times 3$	new copies
point	0	\rightarrow	$1 = 3^0$
segment	1	\rightarrow	$3 = 3^1$
square	2	\rightarrow	$9 = 3^2$
cube	3	\rightarrow	$27 = 3^3$
C	x	\rightarrow	$2 = 3^x$

Figure 3.3: DIMENSION OF C ; $2 = 3^x \Rightarrow x = \ln 2 / \ln 3$.

The notion of a noninteger or fractional dimension is the impetus behind the term “fractal,” coined in 1975 by Benoit Mandelbrot to describe a class of sets whose intricate structures have much in common with the Cantor set. Cantor’s construction, however, is over a hundred years old and for us represents an invaluable testing ground for the upcoming theorems and conjectures about the often elusive nature of subsets of the real line.

3.2 Open and Closed Sets

Given $a \in \mathbf{R}$ and $\epsilon > 0$, recall that the ϵ -neighborhood of a is the set

$$V_\epsilon(a) = \{x \in \mathbf{R} : |x - a| < \epsilon\}.$$

In other words, $V_\epsilon(a)$ is the open interval $(a - \epsilon, a + \epsilon)$, centered at a with radius ϵ .

Definition 3.2.1. A set $O \subseteq \mathbf{R}$ is *open* if for all points $a \in O$ there exists an ϵ -neighborhood $V_\epsilon(a) \subseteq O$.

Example 3.2.2. (i) Perhaps the simplest example of an open set is \mathbf{R} itself. Given an arbitrary element $a \in \mathbf{R}$, we are free to pick any ϵ -neighborhood we like and it will always be true that $V_\epsilon(a) \subseteq \mathbf{R}$. It is also the case that the logical structure of Definition 3.2.1 requires us to classify the empty set \emptyset as an open subset of the real line.

(ii) For a more useful collection of examples, consider the open interval

$$(c, d) = \{x \in \mathbf{R} : c < x < d\}.$$

To see that (c, d) is open in the sense just defined, let $x \in (c, d)$ be arbitrary. If we take $\epsilon = \min\{x - c, d - x\}$, then it follows that $V_\epsilon(x) \subseteq (c, d)$. It is important to see where this argument breaks down if the interval includes either one of its endpoints.

The union of open intervals is another example of an open set. This observation leads to the next result.

Theorem 3.2.3. (i) *The union of an arbitrary collection of open sets is open.*
(ii) *The intersection of a finite collection of open sets is open.*

Proof. To prove (i), we let $\{O_\lambda : \lambda \in \Lambda\}$ be a collection of open sets and let $O = \bigcup_{\lambda \in \Lambda} O_\lambda$. Let a be an arbitrary element of O . In order to show that O is open, Definition 3.2.1 insists that we produce an ϵ -neighborhood of a completely contained in O . But $a \in O$ implies that a is an element of at least one particular $O_{\lambda'}$. Because we are assuming $O_{\lambda'}$ is open, we can use Definition 3.2.1 to assert that there exists $V_\epsilon(a) \subseteq O_{\lambda'}$. The fact that $O_{\lambda'} \subseteq O$ allows us to conclude that $V_\epsilon(a) \subseteq O$. This completes the proof of (i).

For (ii), let $\{O_1, O_2, \dots, O_N\}$ be a finite collection of open sets. Now, if $a \in \bigcap_{k=1}^N O_k$, then a is an element of each of the open sets. By the definition of an open set, we know that, for each $1 \leq k \leq N$, there exists $V_{\epsilon_k}(a) \subseteq O_k$. We are in search of a *single* ϵ -neighborhood of a that is contained in every O_k , so the trick is to take the smallest one. Letting $\epsilon = \min\{\epsilon_1, \epsilon_2, \dots, \epsilon_N\}$, it follows that $V_\epsilon(a) \subseteq V_{\epsilon_k}(a)$ for all k , and hence $V_\epsilon(a) \subseteq \bigcap_{k=1}^N O_k$, as desired. \square

Closed Sets

Definition 3.2.4. A point x is a *limit point* of a set A if every ϵ -neighborhood $V_\epsilon(x)$ of x intersects the set A in some point other than x .

Limit points are also often referred to as “cluster points” or “accumulation points,” but the phrase “ x is a limit point of A ” has the advantage of explicitly reminding us that x is quite literally the limit of a sequence in A .

Theorem 3.2.5. *A point x is a limit point of a set A if and only if $x = \lim a_n$ for some sequence (a_n) contained in A satisfying $a_n \neq x$ for all $n \in \mathbf{N}$.*

Proof. (\Rightarrow) Assume x is a limit point of A . In order to produce a sequence (a_n) converging to x , we are going to consider the particular ϵ -neighborhoods obtained using $\epsilon = 1/n$. By Definition 3.2.4, every neighborhood of x intersects A in some point other than x . This means that, for each $n \in \mathbf{N}$, we are justified in picking a point

$$a_n \in V_{1/n}(x) \cap A$$

with the stipulation that $a_n \neq x$. It should not be too difficult to see why $(a_n) \rightarrow x$. Given an arbitrary $\epsilon > 0$, choose N such that $1/N < \epsilon$. It follows that $|a_n - x| < \epsilon$ for all $n \geq N$.

The reverse implication is requested as Exercise 3.2.4. \square

The restriction that $a_n \neq x$ in Theorem 3.2.5 deserves a comment. Given a point $a \in A$, it is always the case that a is the limit of a sequence in A if we are allowed to consider the constant sequence (a, a, a, \dots) . There will be occasions where we will want to avoid this somewhat uninteresting situation, so it is important to have a vocabulary that can distinguish limit points of a set from *isolated points*.

Definition 3.2.6. A point $a \in A$ is an *isolated point* of A if it is not a limit point of A .

As a word of caution, we need to be a little careful about how we understand the relationship between these concepts. Whereas an isolated point is always an element of the relevant set A , it is quite possible for a limit point of A not to belong to A . As an example, consider the endpoint of an open interval. This situation is the subject of the next important definition.

Definition 3.2.7. A set $F \subseteq \mathbf{R}$ is *closed* if it contains its limit points.

The adjective “closed” appears in several other mathematical contexts and is usually employed to mean that an operation on the elements of a given set does not take us out of the set. In linear algebra, for example, a vector space is a set that is “closed” under addition and scalar multiplication. In analysis, the operation we are concerned with is the limiting operation. Topologically speaking, a closed set is one where convergent sequences within the set have limits that are also in the set.

Theorem 3.2.8. A set $F \subseteq \mathbf{R}$ is closed if and only if every Cauchy sequence contained in F has a limit that is also an element of F .

Proof. Exercise 3.2.6. □

Example 3.2.9. (i) Consider

$$A = \left\{ \frac{1}{n} : n \in \mathbf{N} \right\}.$$

Let’s show that each point of A is isolated. Given $1/n \in A$, choose $\epsilon = 1/n - 1/(n+1)$. Then,

$$V_\epsilon(1/n) \cap A = \left\{ \frac{1}{n} \right\}.$$

It follows from Definition 3.2.4 that $1/n$ is not a limit point and so is isolated. Although all of the points of A are isolated, the set does have one limit point, namely 0. This is because every neighborhood centered at zero, no matter how small, is going to contain points of A . Because $0 \notin A$, A is not closed. The set $F = A \cup \{0\}$ is an example of a closed set and is called the *closure* of A . (The closure of a set is discussed in a moment.)

(ii) Let’s prove that a closed interval

$$[c, d] = \{x \in \mathbf{R} : c \leq x \leq d\}$$

is a closed set using Definition 3.2.7. If x is a limit point of $[c, d]$, then by Theorem 3.2.5 there exists $(x_n) \subseteq [c, d]$ with $(x_n) \rightarrow x$. We need to prove that $x \in [c, d]$.

The key to this argument is contained in the Order Limit Theorem (Theorem 2.3.4), which summarizes the relationship between inequalities and the limiting

process. Because $c \leq x_n \leq d$, it follows from Theorem 2.3.4 (iii) that $c \leq x \leq d$ as well. Thus, $[c, d]$ is closed.

(iii) Consider the set $\mathbf{Q} \subseteq \mathbf{R}$ of rational numbers. An extremely important property of \mathbf{Q} is that its set of limit points is actually *all of \mathbf{R}* . To see why this is so, recall Theorem 1.4.3 from Chapter 1, which is referred to as the density property of \mathbf{Q} in \mathbf{R} .

Let $y \in \mathbf{R}$ be arbitrary, and consider any neighborhood $V_\epsilon(y) = (y - \epsilon, y + \epsilon)$. Theorem 1.4.3 allows us to conclude that there exists a rational number $r \neq y$ that falls in this neighborhood. Thus, y is a limit point of \mathbf{Q} . The density property of \mathbf{Q} can now be reformulated in the following way.

Theorem 3.2.10 (Density of \mathbf{Q} in \mathbf{R}). *Given any $y \in \mathbf{R}$, there exists a sequence of rational numbers that converges to y .*

Proof. Combine the preceding discussion with Theorem 3.2.5. □

The same argument can also be used to show that every real number is the limit of a sequence of irrational numbers. Although interesting, part of the allure of the rational numbers is that, in addition to being dense in \mathbf{R} , they are countable. As we will see, this tangible aspect of \mathbf{Q} makes it an extremely useful set, both for proving theorems and for producing interesting counterexamples.

Closure

Definition 3.2.11. Given a set $A \subseteq \mathbf{R}$, let L be the set of all limit points of A . The *closure* of A is defined to be $\overline{A} = A \cup L$.

In Example 3.2.9 (i), we saw that if $A = \{1/n : n \in \mathbf{N}\}$, then the closure of A is $\overline{A} = A \cup \{0\}$. Example 3.2.9 (iii) verifies that $\overline{\mathbf{Q}} = \mathbf{R}$. If A is an open interval (a, b) , then $\overline{A} = [a, b]$. If A is a closed interval, then $\overline{A} = A$. It is not for lack of imagination that in each of these examples \overline{A} is always a closed set.

Theorem 3.2.12. *For any $A \subseteq \mathbf{R}$, the closure \overline{A} is a closed set and is the smallest closed set containing A .*

Proof. If L is the set of limit points of A , then it is immediately clear that \overline{A} contains the limit points of A . There is still something more to prove, however, because taking the union of L with A could potentially produce some new limit points of \overline{A} . In Exercise 3.2.8, we outline the argument that this does not happen.

Now, *any* closed set containing A must contain L as well. This shows that $\overline{A} = A \cup L$ is the smallest closed set containing A . □

Complements

The mathematical notions of open and closed are not antonyms the way they are in standard English. If a set is not open, that does not imply it must be closed. Many sets such as the half-open interval $(c, d] = \{x \in \mathbf{R} : c < x \leq d\}$ are neither

open nor closed. The sets \mathbf{R} and \emptyset are both simultaneously open and closed although, thankfully, these are the only ones with this unsightly property. There is, however, an important relationship between open and closed sets. Recall that the *complement* of a set $A \subseteq \mathbf{R}$ is defined to be the set

$$A^c = \{x \in \mathbf{R} : x \notin A\}.$$

Theorem 3.2.13. *A set O is open if and only if O^c is closed. Likewise, a set F is closed if and only if F^c is open.*

Proof. Given an open set $O \subseteq \mathbf{R}$, let's first prove that O^c is a closed set. To prove O^c is closed, we need to show that it contains all of its limit points. If x is a limit point of O^c , then *every* neighborhood of x contains some point of O^c . But that is enough to conclude that x cannot be in the open set O because $x \in O$ would imply that there exists a neighborhood $V_\epsilon(x) \subseteq O$. Thus, $x \in O^c$, as desired.

For the converse statement, we assume O^c is closed and argue that O is open. Thus, given an arbitrary point $x \in O$, we must produce an ϵ -neighborhood $V_\epsilon(x) \subseteq O$. Because O^c is closed, we can be sure that x is *not* a limit point of O^c . Looking at the definition of limit point, we see that this implies that there must be some neighborhood $V_\epsilon(x)$ of x that does not intersect the set O^c . But this means $V_\epsilon(x) \subseteq O$, which is precisely what we needed to show.

The second statement in Theorem 3.2.13 follows quickly from the first using the observation that $(E^c)^c = E$ for any set $E \subseteq \mathbf{R}$. \square

The last theorem of this section should be compared to Theorem 3.2.3.

Theorem 3.2.14. (i) *The union of a finite collection of closed sets is closed.*
(ii) *The intersection of an arbitrary collection of closed sets is closed.*

Proof. De Morgan's Laws state that for any collection of sets $\{E_\lambda : \lambda \in \Lambda\}$ it is true that

$$\left(\bigcup_{\lambda \in \Lambda} E_\lambda \right)^c = \bigcap_{\lambda \in \Lambda} E_\lambda^c \quad \text{and} \quad \left(\bigcap_{\lambda \in \Lambda} E_\lambda \right)^c = \bigcup_{\lambda \in \Lambda} E_\lambda^c.$$

The result follows directly from these statements and Theorem 3.2.3. The details are requested in Exercise 3.2.10. \square

Exercises

Exercise 3.2.1. (a) Where in the proof of Theorem 3.2.3 part (ii) does the assumption that the collection of open sets be *finite* get used?

(b) Give an example of an infinite collection of nested open sets

$$O_1 \supseteq O_2 \supseteq O_3 \supseteq O_4 \supseteq \cdots$$

whose intersection $\bigcap_{n=1}^{\infty} O_n$ is closed and nonempty.

Exercise 3.2.2. Let

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

- Find the limit points of B .
- Is B a closed set?
- Is B an open set?
- Does B contain any isolated points?
- Find \overline{B} .

Exercise 3.2.3. Decide whether the following sets are open, closed, or neither. If a set is not open, find a point in the set for which there is no ϵ -neighborhood contained in the set. If a set is not closed, find a limit point that is not contained in the set.

- \mathbf{Q} .
- \mathbf{N} .
- $\{x \in \mathbf{R} : x > 0\}$.
- $(0, 1] = \{x \in \mathbf{R} : 0 < x \leq 1\}$.
- $\{1 + 1/4 + 1/9 + \dots + 1/n^2 : n \in \mathbf{N}\}$.

Exercise 3.2.4. Prove the converse of Theorem 3.2.5 by showing that if $x = \lim a_n$ for some sequence (a_n) contained in A satisfying $a_n \neq x$, then x is a limit point of A .

Exercise 3.2.5. Let $a \in A$. Prove that a is an isolated point of A if and only if there exists an ϵ -neighborhood $V_\epsilon(a)$ such that $V_\epsilon(a) \cap A = \{a\}$.

Exercise 3.2.6. Prove Theorem 3.2.8.

Exercise 3.2.7. Let $x \in O$, where O is an open set. If (x_n) is a sequence converging to x , prove that all but a finite number of the terms of (x_n) must be contained in O .

Exercise 3.2.8. Given $A \subseteq \mathbf{R}$, let L be the set of all limit points of A .

- Show that the set L is closed.
- Argue that if x is a limit point of $A \cup L$, then x is a limit point of A . Use this observation to furnish a proof for Theorem 3.2.12.

Exercise 3.2.9. (a) If y is a limit point of $A \cup B$, show that y is either a limit point of A or a limit point of B (or both).

- Prove that $\overline{A \cup B} = \overline{A} \cup \overline{B}$.
- Does the result about closures in (b) extend to infinite unions of sets?

Exercise 3.2.10 (De Morgan's Laws). A proof for De Morgan's Laws in the case of two sets is outlined in Exercise 1.2.3. The general argument is similar.

- Given a collection of sets $\{E_\lambda : \lambda \in \Lambda\}$, show that

$$\left(\bigcup_{\lambda \in \Lambda} E_\lambda \right)^c = \bigcap_{\lambda \in \Lambda} E_\lambda^c \quad \text{and} \quad \left(\bigcap_{\lambda \in \Lambda} E_\lambda \right)^c = \bigcup_{\lambda \in \Lambda} E_\lambda^c.$$

- Now, provide the details for the proof of Theorem 3.2.14.

Exercise 3.2.11. Let A be bounded above so that $s = \sup A$ exists. Show that $s \in \overline{A}$.

Exercise 3.2.12. Decide whether the following statements are true or false. Provide counterexamples for those that are false, and supply proofs for those that are true.

- (a) For any set $A \subseteq \mathbf{R}$, \overline{A}^c is open.
- (b) If a set A has an isolated point, it cannot be an open set.
- (c) A set A is closed if and only if $\overline{A} = A$.
- (d) If A is a bounded set, then $s = \sup A$ is a limit point of A .
- (e) Every finite set is closed.

(f) An open set that contains every rational number must necessarily be all of \mathbf{R} .

Exercise 3.2.13. Prove that the only sets that are both open and closed are \mathbf{R} and the empty set \emptyset .

Exercise 3.2.14. A set A is called an F_σ set if it can be written as the countable union of closed sets. A set B is called a G_δ set if it can be written as the countable intersection of open sets.

- (a) Show that a closed interval $[a, b]$ is a G_δ set.
- (b) Show that the half-open interval $(a, b]$ is both a G_δ and an F_σ set.
- (c) Show that \mathbf{Q} is an F_σ set, and the set of irrationals \mathbf{I} forms a G_δ set. (We will see in Section 3.5 that \mathbf{Q} is *not* a G_δ set, nor is \mathbf{I} an F_σ set.)

3.3 Compact Sets

Definition 3.3.1. A set $K \subseteq \mathbf{R}$ is *compact* if every sequence in K has a subsequence that converges to a limit that is also in K .

Example 3.3.2. The most basic example of a compact set is a closed interval. To see this, notice that if (a_n) is contained in an interval $[c, d]$, then the Bolzano–Weierstrass Theorem guarantees that we can find a convergent subsequence (a_{n_k}) . Because a closed interval is a closed set (Example 3.2.9, (ii)), we know that the limit of this subsequence is also in $[c, d]$.

What are the properties of closed intervals that we used in the preceding argument? The Bolzano–Weierstrass Theorem requires boundedness, and we used the fact that closed sets contain their limit points. As we are about to see, these two properties completely characterize compact sets in \mathbf{R} . The term “bounded” has thus far only been used to describe sequences (Definition 2.3.1), but an analogous statement can easily be made about sets.

Definition 3.3.3. A set $A \subseteq \mathbf{R}$ is *bounded* if there exists $M > 0$ such that $|a| \leq M$ for all $a \in A$.

Theorem 3.3.4 (Heine–Borel Theorem). *A set $K \subseteq \mathbf{R}$ is compact if and only if it is closed and bounded.*

Proof. Let K be compact. We will first prove that K must be bounded, so assume, for contradiction, that K is not a bounded set. The idea is to produce a sequence in K that marches off to infinity in such a way that it cannot have a convergent subsequence as the definition of compact requires. To do this, notice that because K is not bounded there must exist an element $x_1 \in K$ satisfying $|x_1| > 1$. Likewise, there must exist $x_2 \in K$ with $|x_2| > 2$, and in general, given any $n \in \mathbf{N}$, we can produce $x_n \in K$ such that $|x_n| > n$.

Now, because K is assumed to be compact, (x_n) should have a convergent subsequence (x_{n_k}) . But the elements of the subsequence must necessarily satisfy $|x_{n_k}| > n_k$, and consequently (x_{n_k}) is unbounded. Because convergent sequences are bounded (Theorem 2.3.2), we have a contradiction. Thus, K must at least be a bounded set.

Next, we will show that K is also closed. To see that K contains its limit points, we let $x = \lim x_n$, where (x_n) is contained in K and argue that x must be in K as well. By Definition 3.3.1, the sequence (x_n) has a convergent subsequence (x_{n_k}) , and by Theorem 2.5.2, we know (x_{n_k}) converges to the same limit x . Finally, Definition 3.3.1 requires that $x \in K$. This proves that K is closed.

The proof of the converse statement is requested in Exercise 3.3.2. □

There may be a temptation to consider closed intervals as being a kind of standard archetype for compact sets, but this is misleading. The structure of compact sets can be much more intricate and interesting. For instance, one implication of the Heine–Borel Theorem (Theorem 3.3.4) is that the Cantor set is compact (Exercise 3.3.3). It is more useful to think of compact sets as generalizations of closed intervals. Whenever a fact involving closed intervals is true, it is often the case that the same result holds when we replace “closed interval” with “compact set.” As an example, let’s experiment with the Nested Interval Property proved in the first chapter.

Theorem 3.3.5. *If $K_1 \supseteq K_2 \supseteq K_3 \supseteq K_4 \supseteq \cdots$ is a nested sequence of nonempty compact sets, then the intersection $\bigcap_{n=1}^{\infty} K_n$ is not empty.*

Proof. In order to take advantage of the compactness of each K_n , we are going to produce a sequence that is eventually in each of these sets. Thus, for each $n \in \mathbf{N}$, pick a point $x_n \in K_n$. Because the compact sets are nested, it follows that the sequence (x_n) is contained in K_1 . By Definition 3.3.1, (x_n) has a convergent subsequence (x_{n_k}) whose limit $x = \lim x_{n_k}$ is an element of K_1 .

In fact, x is an element of *every* K_n for essentially the same reason. Given a particular $n_0 \in \mathbf{N}$, the terms in the sequence (x_n) are contained in K_{n_0} as long as $n \geq n_0$. Ignoring the finite number of terms for which $n_k < n_0$, the same subsequence (x_{n_k}) is then also contained in K_{n_0} . The conclusion is that $x = \lim x_{n_k}$ is an element of K_{n_0} . Because n_0 was arbitrary, $x \in \bigcap_{n=1}^{\infty} K_n$ and the proof is complete. □

Open Covers

Defining compactness for sets in \mathbf{R} is reminiscent of the situation we encountered with completeness in that there are a number of equivalent ways to describe this phenomenon. We demonstrated the equivalence of two such characterizations in Theorem 3.3.4. What this theorem implies is that we could have decided to *define* compact sets to be sets that are closed and bounded, and then *proved* that sequences contained in compact sets have convergent subsequences with limits in the set. There are some larger issues involved in deciding what the definition should be, but what is important at this moment is that we be versatile enough to use whatever description of compactness is most appropriate for a given situation.


To add to the delight, there is a third useful characterization of compactness, equivalent to the two others, which is described in terms of *open covers* and *finite subcovers*.

Definition 3.3.6. Let $A \subseteq \mathbf{R}$. An *open cover* for A is a (possibly infinite) collection of open sets $\{O_\lambda : \lambda \in \Lambda\}$ whose union contains the set A ; that is, $A \subseteq \bigcup_{\lambda \in \Lambda} O_\lambda$. Given an open cover for A , a *finite subcover* is a finite subcollection of open sets from the original open cover whose union still manages to completely contain A .

Example 3.3.7. Consider the open interval $(0, 1)$. For each point $x \in (0, 1)$, let O_x be the open interval $(x/2, 1)$. Taken together, the infinite collection $\{O_x : x \in (0, 1)\}$ forms an open cover for the open interval $(0, 1)$. Notice, however, that it is impossible to find a finite subcover. Given any proposed finite subcollection

$$\{O_{x_1}, O_{x_2}, \dots, O_{x_n}\},$$

set $x' = \min\{x_1, x_2, \dots, x_n\}$ and observe that any real number y satisfying $0 < y \leq x'/2$ is not contained in the union $\bigcup_{i=1}^n O_{x_i}$.



Now, consider a similar cover for the closed interval $[0, 1]$. For $x \in (0, 1)$, the sets $O_x = (x/2, 1)$ do a fine job covering $(0, 1)$, but in order to have an open cover of the closed interval $[0, 1]$, we must also cover the endpoints. To remedy this, we could fix $\epsilon > 0$, and let $O_0 = (-\epsilon, \epsilon)$ and $O_1 = (1 - \epsilon, 1 + \epsilon)$. Then, the collection

$$\{O_0, O_1, O_x : x \in (0, 1)\}$$

is an open cover for $[0, 1]$. But this time, notice there is a finite subcover. Because of the addition of the set O_0 , we can choose x' so that $x'/2 < \epsilon$. It follows that $\{O_0, O_{x'}, O_1\}$ is a finite subcover for the closed interval $[0, 1]$.

Theorem 3.3.8. *Let K be a subset of \mathbf{R} . All of the following statements are equivalent in the sense that any one of them implies the two others:*

- (i) K is compact.
- (ii) K is closed and bounded.
- (iii) Any open cover for K has a finite subcover.

Proof. The equivalence of (i) and (ii) is the content of Theorem 3.3.4. What remains is to show that (iii) is equivalent to (i) and (ii). Let's first assume (iii), and prove that it implies (ii) (and thus (i) as well).

To show that K is bounded, we construct an open cover for K by defining O_x to be an open interval of radius 1 around each point $x \in K$. In the language of neighborhoods, $O_x = V_1(x)$. The open cover $\{O_x : x \in K\}$ then must have a finite subcover $\{O_{x_1}, O_{x_2}, \dots, O_{x_n}\}$. Because K is contained in a finite union of bounded sets, K must itself be bounded.

The proof that K is closed is more delicate, and we argue it by contradiction. Let (y_n) be a Cauchy sequence contained in K with $\lim y_n = y$. To show that K is closed, we must demonstrate that $y \in K$, so assume for contradiction that this is not the case. If $y \notin K$, then every $x \in K$ is some positive distance away from y . We now construct an open cover by taking O_x to be an interval of radius $|x - y|/2$ around each point x in K . Because we are assuming (iii), the resulting open cover $\{O_x : x \in K\}$ must have a finite subcover $\{O_{x_1}, O_{x_2}, \dots, O_{x_n}\}$. The contradiction arises when we realize that, in the spirit of Example 3.3.7, this finite subcover cannot contain all of the elements of the sequence (y_n) . To make this explicit, set

$$\epsilon_0 = \min \left\{ \frac{|x_i - y|}{2} : 1 \leq i \leq n \right\}.$$

Because $(y_n) \rightarrow y$, we can certainly find a term y_N satisfying $|y_N - y| < \epsilon_0$. But such a y_N must necessarily be excluded from each O_{x_i} , meaning that

$$y_N \notin \bigcup_{i=1}^n O_{x_i}.$$

Thus our supposed subcover does not actually cover all of K . This contradiction implies that $y \in K$, and hence K is closed and bounded.

The proof that (ii) implies (iii) is outlined in Exercise 3.3.8. □

Exercises

Exercise 3.3.1. Show that if K is compact, then $\sup K$ and $\inf K$ both exist and are elements of K .

Exercise 3.3.2. Prove the converse of Theorem 3.3.4 by showing that if a set $K \subseteq \mathbf{R}$ is closed and bounded, then it is compact.

Exercise 3.3.3. Show that the Cantor set defined in Section 3.1 is a compact set.

Exercise 3.3.4. Show that if K is compact and F is closed, then $K \cap F$ is compact.

Exercise 3.3.5. Decide which of the following sets are compact. For those that are not compact, show how Definition 3.3.1 breaks down. In other words, give an example of a sequence contained in the given set that does not possess a subsequence converging to a limit in the set.

- (a) \mathbf{Q} .
- (b) $\mathbf{Q} \cap [0, 1]$.
- (c) \mathbf{R} .
- (d) $\mathbf{R} \cap [0, 1]$.
- (e) $\{1, 1/2, 1/3, 1/4, 1/5, \dots\}$.
- (f) $\{1, 1/2, 2/3, 3/4, 4/5, \dots\}$.

Exercise 3.3.6. As some more evidence of the surprising nature of the Cantor set, follow these steps to show that the sum $C + C = \{x + y : x, y \in C\}$ is equal to the closed interval $[0, 2]$. (Keep in mind that C has zero length and contains no intervals.)

The observation that $\{x + y : x, y \in C\} \subseteq [0, 2]$ passes for obvious, so we only need to prove the reverse inclusion $[0, 2] \subseteq \{x + y : x, y \in C\}$. Thus, given $s \in [0, 2]$, we must find two elements $x, y \in C$ satisfying $x + y = s$.

(a) Show that there exist $x_1, y_1 \in C_1$ for which $x_1 + y_1 = s$. Show in general that, for an arbitrary $n \in \mathbf{N}$, we can always find $x_n, y_n \in C_n$ for which $x_n + y_n = s$.

(b) Keeping in mind that the sequences (x_n) and (y_n) do not necessarily converge, show how they can nevertheless be used to produce the desired x and y in C satisfying $x + y = s$.

Exercise 3.3.7. Decide whether the following propositions are true or false. If the claim is valid, supply a short proof, and if the claim is false, provide a counterexample.

- (a) An arbitrary intersection of compact sets is compact.
- (b) Let $A \subseteq \mathbf{R}$ be arbitrary, and let $K \subseteq \mathbf{R}$ be compact. Then, the intersection $A \cap K$ is compact.
- (c) If $F_1 \supseteq F_2 \supseteq F_3 \supseteq F_4 \supseteq \dots$ is a nested sequence of nonempty closed sets, then the intersection $\bigcap_{n=1}^{\infty} F_n \neq \emptyset$.
- (d) A finite set is always compact.
- (e) A countable set is always compact.

Exercise 3.3.8. Follow these steps to prove the final implication in Theorem 3.3.8.

Assume K satisfies (i) and (ii), and let $\{O_\lambda : \lambda \in \Lambda\}$ be an open cover for K . For contradiction, let's assume that no finite subcover exists. Let I_0 be a closed interval containing K , and bisect I_0 into two closed intervals A_1 and B_1 .

(a) Why must either $A_1 \cap K$ or $B_1 \cap K$ (or both) have no finite subcover consisting of sets from $\{O_\lambda : \lambda \in \Lambda\}$.

(b) Show that there exists a nested sequence of closed intervals $I_0 \supseteq I_1 \supseteq I_2 \supseteq \cdots$ with the property that, for each n , $I_n \cap K$ cannot be finitely covered and $\lim |I_n| = 0$.

(c) Show that there exists an $x \in K$ such that $x \in I_n$ for all n .

(d) Because $x \in K$, there must exist an open set O_{λ_0} from the original collection that contains x as an element. Argue that there must be an n_0 large enough to guarantee that $I_{n_0} \subseteq O_{\lambda_0}$. Explain why this furnishes us with the desired contradiction.

Exercise 3.3.9. Consider each of the sets listed in Exercise 3.3.5. For each one that is not compact, find an open cover for which there is no finite subcover.

Exercise 3.3.10. Let's call a set *clompact* if it has the property that every *closed* cover (i.e., a cover consisting of closed sets) admits a finite subcover. Describe all of the clompact subsets of \mathbf{R} .

3.4 Perfect Sets and Connected Sets

One of the underlying goals of topology is to strip away all of the extraneous information that comes with our intuitive picture of the real numbers and isolate just those properties that are responsible for the phenomenon we are studying. For example, we were quick to observe that any closed interval is a compact set. The content of Theorem 3.3.4, however, is that the compactness of a closed interval has nothing to do with the fact that the set is an *interval* but is a consequence of the set being bounded and closed. In Chapter 1, we argued that the set of real numbers between 0 and 1 is an uncountable set. This turns out to be the case for any nonempty closed set that does not contain isolated points.

Perfect Sets

Definition 3.4.1. A set $P \subseteq \mathbf{R}$ is *perfect* if it is closed and contains no isolated points.

Closed intervals (other than the singleton sets $[a, a]$) serve as the most obvious class of examples of perfect sets, but in fact it is not too difficult to prove that the Cantor set from Section 3.1 is another example.

Theorem 3.4.2. *The Cantor set is perfect.*

Proof. The Cantor set is defined as the intersection $C = \bigcap_{n=0}^{\infty} C_n$, where each C_n is a finite union of closed intervals. By Theorem 3.2.14, each C_n is closed, and by the same theorem, C is closed as well. It remains to show that no point in C is isolated.

Let $x \in C$ be arbitrary. To convince ourselves that x is not isolated, we must construct a sequence (x_n) of points in C , different from x , that converges to x . From our earlier discussion, we know that C at least contains the endpoints of the intervals that make up each C_n . In Exercise 3.4.3, we sketch the argument that these are all that is needed to construct (x_n) . \square

One argument for the uncountability of the Cantor set was presented in Section 3.1. Another, perhaps more satisfying, argument for the same conclusion can be obtained from the next theorem.

Theorem 3.4.3. *A nonempty perfect set is uncountable.*

Proof. If P is perfect and nonempty, then it must be infinite because otherwise it would consist only of isolated points. Let's assume, for contradiction, that P is countable. Thus, we can write

$$P = \{x_1, x_2, x_3, \dots\},$$

where *every* element of P appears on this list. The idea is to construct a sequence of nested compact sets K_n , all contained in P , with the property that $x_1 \notin K_2$, $x_2 \notin K_3$, $x_3 \notin K_4$, \dots . Some care must be taken to ensure that each K_n is nonempty, for then we can use Theorem 3.3.5 to produce an


$$x \in \bigcap_{n=1}^{\infty} K_n \subseteq P$$

that *cannot* be on the list $\{x_1, x_2, x_3, \dots\}$.

Let I_1 be a closed interval that contains x_1 in its interior (i.e., x_1 is not an endpoint of I_1). Now, x_1 is not isolated, so there exists some other point $y_2 \in P$ that is also in the interior of I_1 . Construct a closed interval I_2 , centered on y_2 , so that $I_2 \subseteq I_1$ but $x_1 \notin I_2$. More explicitly, if $I_1 = [a, b]$, let

$$\epsilon = \min\{y_2 - a, b - y_2, |x_1 - y_2|\}.$$

Then, the interval $I_2 = [y_2 - \epsilon/2, y_2 + \epsilon/2]$ has the desired properties.



This process can be continued. Because $y_2 \in P$ is not isolated, there must exist another point $y_3 \in P$ in the interior of I_2 , and we may insist that $y_3 \neq x_2$. Now, construct I_3 centered on y_3 and small enough so that $x_2 \notin I_3$ and $I_3 \subseteq I_2$. Observe that $I_3 \cap P \neq \emptyset$ because this intersection contains at least y_3 .

If we carry out this construction inductively, the result is a sequence of closed intervals I_n satisfying

- (i) $I_{n+1} \subseteq I_n$,
- (ii) $x_n \notin I_{n+1}$, and
- (iii) $I_n \cap P \neq \emptyset$.

To finish the proof, we let $K_n = I_n \cap P$. For each $n \in \mathbf{N}$, we have that K_n is closed because it is the intersection of closed sets, and bounded because it is contained in the bounded set I_n . Hence, K_n is compact. By construction, K_n is not empty and $K_{n+1} \subseteq K_n$. Thus, we can employ Theorem 3.3.5 to conclude that the intersection

$$\bigcap_{n=1}^{\infty} K_n \neq \emptyset.$$

But each K_n is a subset of P , and the fact that $x_n \notin I_{n+1}$ leads to the conclusion that $\bigcap_{n=1}^{\infty} K_n = \emptyset$, which is the sought-after contradiction. \square

Connected Sets

Although the two open intervals $(1, 2)$ and $(2, 5)$ have the limit point $x = 2$ in common, there is still some space between them in the sense that no limit point of one of these intervals is actually contained in the other. Said another way, the closure of $(1, 2)$ (see Definition 3.2.11) is disjoint from $(2, 5)$, and the closure of $(2, 5)$ does not intersect $(1, 2)$. Notice that this same observation cannot be made about $(1, 2]$ and $(2, 5)$, even though these latter sets are disjoint.

Definition 3.4.4. Two nonempty sets $A, B \subseteq \mathbf{R}$ are *separated* if $\bar{A} \cap B$ and $A \cap \bar{B}$ are both empty. A set $E \subseteq \mathbf{R}$ is *disconnected* if it can be written as $E = A \cup B$, where A and B are nonempty separated sets.

A set that is not disconnected is called a *connected set*.

Example 3.4.5. (i) If we let $A = (1, 2)$ and $B = (2, 5)$, then it is not difficult to verify that $E = (1, 2) \cup (2, 5)$ is disconnected. Notice that the sets $C = (1, 2]$ and $D = (2, 5)$ are not separated because $C \cap \bar{D} = \{2\}$ is not empty. This should be comforting. The union $C \cup D$ is equal to the interval $(1, 5)$, which better not qualify as a disconnected set. We will prove in a moment that every interval is a connected subset of \mathbf{R} and vice versa.

(ii) Let's show that the set of rational numbers is disconnected. If we let

$$A = \mathbf{Q} \cap (-\infty, \sqrt{2}) \quad \text{and} \quad B = \mathbf{Q} \cap (\sqrt{2}, \infty),$$

then we certainly have $\mathbf{Q} = A \cup B$. The fact that $A \subseteq (-\infty, \sqrt{2})$ implies (by the Order Limit Theorem) that any limit point of A will necessarily fall in $(-\infty, \sqrt{2}]$. Because this is disjoint from B , we get $\bar{A} \cap B = \emptyset$. We can similarly show that $A \cap \bar{B} = \emptyset$, which implies that A and B are separated.

The definition of connected is stated as the negation of disconnected, but a little care with the logical negation of the quantifiers in Definition 3.4.4 results in a positive characterization of connectedness. Essentially, a set E is connected

if, no matter how it is partitioned into two nonempty disjoint sets, it is always possible to show that at least one of the sets contains a limit point of the other.

Theorem 3.4.6. *A set $E \subseteq \mathbf{R}$ is connected if and only if, for all nonempty disjoint sets A and B satisfying $E = A \cup B$, there always exists a convergent sequence $(x_n) \rightarrow x$ with (x_n) contained in one of A or B , and x an element of the other.*

Proof. Exercise 3.4.6. □

The concept of connectedness is more relevant when working with subsets of the plane and other higher-dimensional spaces. This is because, in \mathbf{R} , the connected sets coincide precisely with the collection of intervals (with the understanding that unbounded intervals such as $(-\infty, 3)$ and $[0, \infty)$ are included).

Theorem 3.4.7. *A set $E \subseteq \mathbf{R}$ is connected if and only if whenever $a < c < b$ with $a, b \in E$, it follows that $c \in E$ as well.*

Proof. Assume E is connected, and let $a, b \in E$ and $a < c < b$. Set

$$A = (-\infty, c) \cap E \quad \text{and} \quad B = (c, \infty) \cap E.$$

Because $a \in A$ and $b \in B$, neither set is empty and, just as in Example 3.4.5 (ii), neither set contains a limit point of the other. If $E = A \cup B$, then we would have that E is disconnected, which it is not. It must then be that $A \cup B$ is missing some element of E , and c is the only possibility. Thus, $c \in E$.

Conversely, assume that E is an interval in the sense that whenever $a, b \in E$ satisfy $a < c < b$ for some c , then $c \in E$. Our intent is to use the characterization of connected sets in Theorem 3.4.6, so let $E = A \cup B$, where A and B are nonempty and disjoint. We need to show that one of these sets contains a limit point of the other. Pick $a_0 \in A$ and $b_0 \in B$, and, for the sake of the argument, assume $a_0 < b_0$. Because E is itself an interval, the interval $I_0 = [a_0, b_0]$ is contained in E . Now, bisect I_0 into two equal halves. The midpoint of I_0 must either be in A or B , and so choose $I_1 = [a_1, b_1]$ to be the half that allows us to have $a_1 \in A$ and $b_1 \in B$. Continuing this process yields a sequence of nested intervals $I_n = [a_n, b_n]$, where $a_n \in A$, $b_n \in B$, and the length $(b_n - a_n) \rightarrow 0$. The remainder of this argument should feel familiar. By the Nested Interval Property, there exists an

$$x \in \bigcap_{n=0}^{\infty} I_n,$$

and it is straightforward to show that the sequences of endpoints each satisfy $\lim a_n = x$ and $\lim b_n = x$. But now $x \in E$ must belong to either A or B , thus making it a limit point of the other. This completes the argument. □

Exercises

Exercise 3.4.1. If P is a perfect set and K is compact, is the intersection $P \cap K$ always compact? Always perfect?

Exercise 3.4.2. Does there exist a perfect set consisting of only rational numbers?

Exercise 3.4.3. Review the portion of the proof given for Theorem 3.4.2 and follow these steps to complete the argument.

(a) Because $x \in C_1$, argue that there exists an $x_1 \in C \cap C_1$ with $x_1 \neq x$ satisfying $|x - x_1| \leq 1/3$.

(b) Finish the proof by showing that for each $n \in \mathbf{N}$, there exists $x_n \in C \cap C_n$, different from x , satisfying $|x - x_n| \leq 1/3^n$.

Exercise 3.4.4. Repeat the Cantor construction from Section 3.1 starting with the open interval $[0, 1]$. This time, however, remove the open middle *fourth* from each component.

(a) Is the resulting set compact? Perfect?

(b) Using the algorithms from Section 3.1, compute the length and dimension of this Cantor-like set.

Exercise 3.4.5. Let A and B be subsets of \mathbf{R} . Show that if there exist disjoint open sets U and V with $A \subseteq U$ and $B \subseteq V$, then A and B are separated.

Exercise 3.4.6. Prove Theorem 3.4.6.

Exercise 3.4.7. (a) Find an example of a disconnected set whose closure is connected.

(b) If A is connected, is \bar{A} necessarily connected? If A is perfect, is \bar{A} necessarily perfect?

Exercise 3.4.8. A set E is *totally disconnected* if, given any two points $x, y \in E$, there exist separated sets A and B with $x \in A$, $y \in B$, and $E = A \cup B$.

(a) Show that \mathbf{Q} is totally disconnected.

(b) Is the set of irrational numbers totally disconnected?

Exercise 3.4.9. Follow these steps to show that the Cantor set $C = \bigcap_{n=0}^{\infty} C_n$ described in Section 3.1 is totally disconnected in the sense described in Exercise 3.4.8.

(a) Given $x, y \in C$, with $x < y$, set $\epsilon = y - x$. For each $n = 0, 1, 2, \dots$, the set C_n consists of a finite number of closed intervals. Explain why there must exist an N large enough so that it is impossible for x and y both to belong to the same closed interval of C_N .

(b) Argue that there exists a point $z \notin C$ such that $x < z < y$. Explain how this proves that there can be no interval of the form (a, b) with $a < b$ contained in C .

(c) Show that C is totally disconnected.

Exercise 3.4.10. Let $\{r_1, r_2, r_3, \dots\}$ be an enumeration of the rational numbers, and for each $n \in \mathbf{N}$ set $\epsilon_n = 1/2^n$. Define $O = \bigcup_{n=1}^{\infty} V_{\epsilon_n}(r_n)$, and let $F = O^c$.

(a) Argue that F is a closed, nonempty set consisting only of irrational numbers.

(b) Does F contain any nonempty open intervals? Is F totally disconnected? (See Exercise 3.4.8 for the definition.)

(c) Is it possible to know whether F is perfect? If not, can we modify this construction to produce a nonempty perfect set of irrational numbers?

3.5 Baire's Theorem

The nature of the real line can be deceptively elusive. The closer we look, the more intricate and enigmatic \mathbf{R} becomes, and the more we are reminded to proceed carefully (i.e., axiomatically) with all of our conclusions about properties of subsets of \mathbf{R} . The structure of open sets is fairly straightforward. Every open set is either a finite or countable union of open intervals. Standing in opposition to this tidy description of all open sets is the Cantor set. The Cantor set is a closed, uncountable set that contains no intervals of any kind. Thus, no such characterization of closed sets should be anticipated.

Recall that the arbitrary union of open sets is always an open set. Likewise, the arbitrary intersection of closed sets is closed. By taking unions of closed sets or intersections of open sets, however, it is possible to obtain a fairly broad selection of subsets of \mathbf{R} . In Exercises 3.2.14, we introduced the following two classes of sets.

Definition 3.5.1. A set $A \subseteq \mathbf{R}$ is called an F_σ set if it can be written as the countable union of closed sets. A set $B \subseteq \mathbf{R}$ is called a G_δ set if it can be written as the countable intersection of open sets.

Exercise 3.5.1. Argue that a set A is a G_δ set if and only if its complement is an F_σ set.

Exercise 3.5.2. Replace each _____ with the word *finite* or *countable*, depending on which is more appropriate.

- The _____ union of F_σ sets is an F_σ set.
- The _____ intersection of F_σ sets is an F_σ set.
- The _____ union of G_δ sets is a G_δ set.
- The _____ intersection of G_δ sets is a G_δ set.

Exercise 3.5.3. (This exercise has already appeared as Exercise 3.2.14.)

- Show that a closed interval $[a, b]$ is a G_δ set.
- Show that the half-open interval $(a, b]$ is both a G_δ and an F_σ set.
- Show that \mathbf{Q} is an F_σ set, and the set of irrationals \mathbf{I} forms a G_δ set.

It is not readily obvious that the class F_σ does not include every subset of \mathbf{R} , but we are now ready to argue that \mathbf{I} is not an F_σ set (and consequently \mathbf{Q} is not a G_δ set). This will follow from a theorem due to René Louis Baire (1874–1932).

Recall that a set $G \subseteq \mathbf{R}$ is *dense* in \mathbf{R} if, given any two real numbers $a < b$, it is possible to find a point $x \in G$ with $a < x < b$.

Theorem 3.5.2. *If $\{G_1, G_2, G_3, \dots\}$ is a countable collection of dense, open sets, then the intersection $\bigcap_{n=1}^{\infty} G_n$ is not empty.*

Proof. Before embarking on the proof, notice that we have seen a conclusion like this before. Theorem 3.3.5 asserts that a nested sequence of compact sets has a nontrivial intersection. In this theorem, we are dealing with dense, open sets, but as it turns out, we are going to use Theorem 3.3.5—and actually, just the Nested Interval Property—as the crucial step in the argument.

Exercise 3.5.4. (a) Starting with $n = 1$, inductively construct a nested sequence of *closed* intervals $I_1 \supseteq I_2 \supseteq I_3 \supseteq \dots$ satisfying $I_n \subseteq G_n$. Give special attention to the issue of the endpoints of each I_n .

(b) Now, use Theorem 3.3.5 or the Nested Interval Property to finish the proof.

□

Exercise 3.5.5. Show that it is impossible to write

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

where for each $n \in \mathbf{N}$, F_n is a closed set containing no nonempty open intervals.

Exercise 3.5.6. Show how the previous exercise implies that the set \mathbf{I} of irrationals cannot be an F_σ set, and \mathbf{Q} cannot be a G_δ set.

Exercise 3.5.7. Using Exercise 3.5.6 and versions of the statements in Exercise 3.5.2, construct a set that is neither in F_σ nor in G_δ .

Nowhere-Dense Sets

We have encountered several equivalent ways to assert that a particular set G is dense in \mathbf{R} . In Section 3.2, we observed that G is dense in \mathbf{R} if and only if every point of \mathbf{R} is a limit point of G . Because the closure of any set is obtained by taking the union of the set and its limit points, we have that

$$G \text{ is dense in } \mathbf{R} \text{ if and only if } \overline{G} = \mathbf{R}.$$

The set \mathbf{Q} is dense in \mathbf{R} ; the set \mathbf{Z} is clearly not. In fact, in the jargon of analysis, \mathbf{Z} is “nowhere-dense” in \mathbf{R} .

Definition 3.5.3. A set E is *nowhere-dense* if \overline{E} contains no nonempty open intervals.

Exercise 3.5.8. Show that a set E is nowhere-dense in \mathbf{R} if and only if the complement of \overline{E} is dense in \mathbf{R} .

Exercise 3.5.9. Decide whether the following sets are dense in \mathbf{R} , nowhere-dense in \mathbf{R} , or somewhere in between.

- (a) $A = \mathbf{Q} \cap [0, 5]$.
- (b) $B = \{1/n : n \in \mathbf{N}\}$.
- (c) the set of irrationals.
- (d) the Cantor set.

We can now restate Theorem 3.5.2 in a slightly more general form.

Theorem 3.5.4 (Baire's Theorem). *The set of real numbers \mathbf{R} cannot be written as the countable union of nowhere-dense sets.*

Proof. For contradiction, assume that E_1, E_2, E_3, \dots are each nowhere-dense and satisfy $\mathbf{R} = \bigcup_{n=1}^{\infty} E_n$.

Exercise 3.5.10. Finish the proof by finding a contradiction to the results in this section. □

3.6 Epilogue

Baire's Theorem is yet another statement about the size of \mathbf{R} . We have already encountered several ways to describe the sizes of infinite sets. In terms of cardinality, countable sets are relatively small whereas uncountable sets are large. We also briefly discussed the concept of “length,” or “measure,” in Section 3.1. Baire's Theorem offers a third perspective. From this point of view, nowhere-dense sets are considered to be “thin” sets. Any set that is the countable union—i.e., a not very large union—of these small sets is called a “meager” set or a set of “first category.” A set that is not of first category is of “second category.” Intuitively, sets of the second category are the “fat” subsets. The Baire Category Theorem, as it is often called, states that \mathbf{R} is of second category.

There is a significance to the Baire Category Theorem that is difficult to appreciate at the moment because we are only seeing a special case of this result. The real numbers are an example of a *complete metric space*. Metric spaces are discussed in some detail in Section 8.2, but here is the basic idea. Given a set of mathematical objects such as real numbers, points in the plane or continuous functions defined on $[0,1]$, a “metric” is a rule that assigns a “distance” between two elements in the set. In \mathbf{R} , we have been using $|x - y|$ as the distance between the real numbers x and y . The point is that if we can create a satisfactory notion of “distance” on these other spaces (we will need the triangle inequality to hold, for instance), then the concepts of convergence, Cauchy sequences, and open sets, for example, can be naturally transferred over. A complete metric space is any set with a suitably defined metric in which Cauchy sequences have limits. We have spent a good deal of time discussing the fact that \mathbf{R} is a complete metric space whereas \mathbf{Q} is not.

The Baire Category Theorem in its more general form states that *any* complete metric space must be too large to be the countable union of nowhere-dense subsets. One particularly interesting example of a complete metric space is the set of continuous functions defined on the interval $[0, 1]$. (The distance between two functions f and g in this space is defined to be $\sup |f(x) - g(x)|$, where $x \in [0, 1]$.) Now, in this space we will see that the collection of continuous functions that are differentiable at even one point *can* be written as the countable union of nowhere-dense sets. Thus, a fascinating consequence of Baire's Theorem in this setting is that *most continuous functions do not have derivatives at any point*. Chapter 5 concludes with a construction of one such function. This odd situation mirrors the roles of \mathbf{Q} and \mathbf{I} as subsets of \mathbf{R} . Just as the familiar rational numbers constitute a minute proportion of the real line, the differentiable functions of calculus are exceedingly atypical of continuous functions in general.

Chapter 4

Functional Limits and Continuity

4.1 Discussion: Examples of Dirichlet and Thomae

Although it is common practice in calculus courses to discuss continuity before differentiation, historically mathematicians' attention to the concept of continuity came long after the derivative was in wide use. Pierre de Fermat (1601–1665) was using tangent lines to solve optimization problems as early as 1629. On the other hand, it was not until around 1820 that Cauchy, Bolzano, Weierstrass, and others began to characterize continuity in terms more rigorous than prevailing intuitive notions such as “unbroken curves” or “functions which have no jumps or gaps.” The basic reason for this two-hundred year waiting period lies in the fact that, for most of this time, the very notion of *function* did not really permit discontinuities. Functions were entities such as polynomials, sines, and cosines, always smooth and continuous over their relevant domains. The gradual liberation of the term function to its modern understanding—a rule associating a unique output to a given input—was simultaneous with 19th century investigations into the behavior of infinite series. Extensions of the power of calculus were intimately tied to the ability to represent a function $f(x)$ as a limit of polynomials (called a *power series*) or as a limit of sums of sines and cosines (called a *trigonometric* or *Fourier series*). A typical question for Cauchy and his contemporaries was whether the continuity of the limiting polynomials or trigonometric functions necessarily implied that the limit f would also be continuous.

Sequences and series of functions are the topics of Chapter 6. What is relevant at this moment is that we realize why the issue of finding a rigorous definition for continuity finally made its way to the fore. Any significant progress on the question of whether the limit of continuous functions is continuous (for

Figure 4.1: DIRICHLET'S FUNCTION, $g(x)$.

Cauchy and for us) necessarily depends on a definition of continuity that does not rely on imprecise notions such as “no holes” or “gaps.” With a mathematically unambiguous definition for the limit of a sequence in hand, we are well on our way toward a rigorous understanding of continuity.

Given a function f with domain $A \subseteq \mathbf{R}$, we want to define continuity at a point $c \in A$ to mean that if $x \in A$ is chosen *near* c , then $f(x)$ will be *near* $f(c)$. Symbolically, we will say f is continuous at c if

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

The problem is that, at present, we only have a definition for the limit of a sequence, and it is not entirely clear what is meant by $\lim_{x \rightarrow c} f(x)$. The subtleties that arise as we try to fashion such a definition are well-illustrated via a family of examples, all based on an idea of the prominent German mathematician, Peter Lejeune Dirichlet. Dirichlet's idea was to define a function g in a piecewise manner based on whether or not the input variable x is rational or irrational. Specifically, let

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


The intricate way that \mathbf{Q} and \mathbf{I} fit inside of \mathbf{R} makes an accurate graph of g technically impossible to draw, but Figure 4.1 illustrates the basic idea.

Does it make sense to attach a value to the expression $\lim_{x \rightarrow 1/2} g(x)$? One idea is to consider a sequence $(x_n) \rightarrow 1/2$. Using our notion of the limit of a sequence, we might try to define $\lim_{x \rightarrow 1/2} g(x)$ as simply the limit of the sequence $g(x_n)$. But notice that this limit depends on how the sequence (x_n) is chosen. If each x_n is rational, then

$$\lim_{n \rightarrow \infty} g(x_n) = 1.$$

On the other hand, if x_n is irrational for each n , then

$$\lim_{n \rightarrow \infty} g(x_n) = 0.$$

Figure 4.2: MODIFIED DIRICHLET FUNCTION, $h(x)$.

This unacceptable situation demands that we work harder on our definition of functional limits. Generally speaking, we want the value of $\lim_{x \rightarrow c} g(x)$ to be independent of how we approach c . In this particular case, the definition of a functional limit that we agree on should lead to the conclusion that

$$\lim_{x \rightarrow 1/2} g(x) \quad \text{does not exist.}$$

Postponing the search for formal definitions for the moment, we should nonetheless realize that Dirichlet's function is not continuous at $c = 1/2$. In fact, the real significance of this function is that there is nothing unique about the point $c = 1/2$. Because both \mathbf{Q} and \mathbf{I} (the set of irrationals) are dense in the real line, it follows that for any $z \in \mathbf{R}$ we can find sequences $(x_n) \subseteq \mathbf{Q}$ and $(y_n) \subseteq \mathbf{I}$ such that

$$\lim x_n = \lim y_n = z.$$

(See Example 3.2.9 (iii).) Because

$$\lim g(x_n) \neq \lim g(y_n),$$

the same line of reasoning reveals that $g(x)$ is not continuous at z . In the jargon of analysis, Dirichlet's function is a *nowhere-continuous* function on \mathbf{R} .


What happens if we adjust the definition of $g(x)$ in the following way? Define a new function h (Fig. 4.2) on \mathbf{R} by setting

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

If we take c different from zero, then just as before we can construct sequences $(x_n) \rightarrow c$ of rationals and $(y_n) \rightarrow c$ of irrationals so that

$$\lim h(x_n) = c \quad \text{and} \quad \lim h(y_n) = 0.$$

Thus, h is not continuous at every point $c \neq 0$.

Figure 4.3: THOMAE'S FUNCTION, $t(x)$.

If $c = 0$, however, then these two limits are both equal to $h(0) = 0$. In fact, it appears as though no matter how we construct a sequence (z_n) converging to zero, it will always be the case that $\lim h(z_n) = 0$. This observation goes to the heart of what we want functional limits to entail. To assert that

$$\lim_{x \rightarrow c} h(x) = L$$

should imply that

$$h(z_n) \rightarrow L \quad \text{for all sequences } (z_n) \rightarrow c.$$

For reasons not yet apparent, it is beneficial to fashion the definition for functional limits in terms of neighborhoods constructed around c and L . We will quickly see, however, that this topological formulation is equivalent to the sequential characterization we have arrived at here.

To this point, we have been discussing continuity of a function at a particular point in its domain. This is a significant departure from thinking of continuous functions as curves that can be drawn without lifting the pen from the paper, and it leads to some fascinating questions. In 1875, K.J. Thomae discovered the function

$$t(x) = \begin{cases} 1 & \text{if } x = 0 \\ 1/n & \text{if } x = m/n \in \mathbf{Q} \setminus \{0\} \text{ is in lowest terms with } n > 0 \\ 0 & \text{if } x \notin \mathbf{Q}. \end{cases}$$

If $c \in \mathbf{Q}$, then $t(c) > 0$. Because the set of irrationals is dense in \mathbf{R} , we can find a sequence (y_n) in \mathbf{I} converging to c . The result is that

$$\lim t(y_n) = 0 \neq t(c),$$

and Thomae's function (Fig. 4.3) fails to be continuous at any rational point.

The twist comes when we try this argument on some irrational point in the domain such as $c = \sqrt{2}$. All irrational values get mapped to zero by t , so the natural thing would be to consider a sequence (x_n) of rational numbers that converges to $\sqrt{2}$. Now, $\sqrt{2} \approx 1.414213\dots$ so a good start on a particular

sequence of rational approximations for $\sqrt{2}$ might be

$$\left(1, \frac{14}{10}, \frac{141}{100}, \frac{1414}{1000}, \frac{14142}{10000}, \frac{141421}{100000}, \dots\right).$$

But notice that the denominators of these fractions are getting larger. In this case, the sequence $t(x_n)$ begins,

$$\left(1, \frac{1}{5}, \frac{1}{100}, \frac{1}{500}, \frac{1}{5000}, \frac{1}{100000}, \dots\right)$$

and is fast approaching $0 = t(\sqrt{2})$. We will see that this always happens. The closer a rational number is chosen to a fixed irrational number, the larger its denominator must necessarily be. As a consequence, Thomae's function has the bizarre property of being continuous at every irrational point on \mathbf{R} and discontinuous at every rational point.

Is there an example of a function with the opposite property? In other words, does there exist a function defined on all of \mathbf{R} that is continuous on \mathbf{Q} but fails to be continuous on \mathbf{I} ? Can the set of discontinuities of a particular function be arbitrary? If we are given some set $A \subseteq \mathbf{R}$, is it always possible to find a function that is continuous only on the set A^c ? In each of the examples in this section, the functions were defined to have erratic oscillations around points in the domain. What conclusions can we draw if we restrict our attention to functions that are somewhat less volatile? One such class is the set of so-called *monotone* functions, which are either increasing or decreasing on a given domain. What might we be able to say about the set of discontinuities of a monotone function on \mathbf{R} ?

4.2 Functional Limits

Consider a function $f : A \rightarrow \mathbf{R}$. Recall that a limit point c of A is a point with the property that every ϵ -neighborhood $V_\epsilon(c)$ intersects A in some point other than c . Equivalently, c is a limit point of A if and only if $c = \lim x_n$ for some sequence $(x_n) \subseteq A$ with $x_n \neq c$. It is important to remember that limit points of A do not necessarily belong to the set A unless A is closed.

If c is a limit point of the domain of f , then, intuitively, the statement

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

is intended to convey that values of $f(x)$ get arbitrarily close to L as x is chosen closer and closer to c . The issue of what happens when $x = c$ is irrelevant from the point of view of functional limits. In fact, c need not even be in the domain of f .

The structure of the definition of functional limits follows the “challenge–response” pattern established in the definition for the limit of a sequence. Recall




Figure 4.4: DEFINITION OF FUNCTIONAL LIMIT.

that given a sequence (a_n) , the assertion $\lim a_n = L$ implies that for every ϵ -neighborhood $V_\epsilon(L)$ centered at L , there is a point in the sequence—call it a_N —after which all of the terms a_n fall in $V_\epsilon(L)$. Each ϵ -neighborhood represents a particular challenge, and each N is the respective response. For functional limit statements such as $\lim_{x \rightarrow c} f(x) = L$, the challenges are still made in the form of an arbitrary ϵ -neighborhood around L , but the response this time is a δ -neighborhood centered at c .

Definition 4.2.1. Let $f : A \rightarrow \mathbf{R}$, and let c be a limit point of the domain A . We say that $\lim_{x \rightarrow c} f(x) = L$ provided that, for all $\epsilon > 0$, there exists a $\delta > 0$ such that whenever $0 < |x - c| < \delta$ (and $x \in A$) it follows that $|f(x) - L| < \epsilon$.

This is often referred to as the “ ϵ - δ version” of the definition for functional limits. Recall that the statement

$$|f(x) - L| < \epsilon \text{ is equivalent to } f(x) \in V_\epsilon(L).$$

Likewise, the statement

$$|x - c| < \delta \text{ is satisfied if and only if } x \in V_\delta(c).$$

The additional restriction $0 < |x - c|$ is just an economical way of saying $x \neq c$. Recasting Definition 4.2.1 in terms of neighborhoods—just as we did for the definition of convergence of a sequence in Section 2.2—amounts to little more than a change of notation, but it does help emphasize the geometrical nature of what is happening (Fig. 4.4).

Definition 4.2.1B (Topological Version). Let c be a limit point of the domain of $f : A \rightarrow \mathbf{R}$. We say $\lim_{x \rightarrow c} f(x) = L$ provided that, for every ϵ -neighborhood $V_\epsilon(L)$ of L , there exists a δ -neighborhood $V_\delta(c)$ around c with the property that for all $x \in V_\delta(c)$ different from c (with $x \in A$) it follows that $f(x) \in V_\epsilon(L)$.

The parenthetical reminder “ $(x \in A)$ ” present in both versions of the definition is included to ensure that x is an allowable input for the function in question. When no confusion is likely, we may omit this reminder, leaving it to the reader to understand that the appearance of $f(x)$ carries with it the implicit assumption that x is in the domain of f . On a similar note, there is no reason to discuss functional limits at isolated points of the domain. Thus, functional limits will only be considered as x tends toward a limit point of the function’s domain.

Example 4.2.2. (i) To familiarize ourselves with Definition 4.2.1, let’s prove that if $f(x) = 3x + 1$, then

$$\lim_{x \rightarrow 2} f(x) = 7.$$

Let $\epsilon > 0$. Definition 4.2.1 requires that we produce a $\delta > 0$ so that $0 < |x - 2| < \delta$ leads to the conclusion $|f(x) - 7| < \epsilon$. Notice that

$$|f(x) - 7| = |(3x + 1) - 7| = |3x - 6| = 3|x - 2|.$$

Thus, if we choose $\delta = \epsilon/3$, then $0 < |x - 2| < \delta$ implies $|f(x) - 7| < 3(\epsilon/3) = \epsilon$.

(ii) Let’s show

$$\lim_{x \rightarrow 2} g(x) = 4,$$

where $g(x) = x^2$. Given an arbitrary $\epsilon > 0$, our goal this time is to make $|g(x) - 4| < \epsilon$ by restricting $|x - 2|$ to be smaller than some carefully chosen δ . As in the previous problem, a little algebra reveals

$$|g(x) - 4| = |x^2 - 4| = |x + 2||x - 2|.$$

We can make $|x - 2|$ as small as we like, but we need an upper bound on $|x + 2|$ in order to know how small to choose δ . The presence of the variable x causes some initial confusion, but keep in mind that we are discussing the limit as x approaches 2. If we agree that our δ -neighborhood around $c = 2$ must have radius no bigger than $\delta = 1$, then we get the upper bound $|x + 2| \leq |3 + 2| = 5$ for all $x \in V_\delta(c)$.

Now, choose $\delta = \min\{1, \epsilon/5\}$. If $0 < |x - 2| < \delta$, then it follows that

$$|x^2 - 4| = |x + 2||x - 2| < (5)\frac{\epsilon}{5} = \epsilon,$$

and the limit is proved.

Sequential Criterion for Functional Limits

We worked very hard in Chapter 2 to derive an impressive list of properties enjoyed by sequential limits. In particular, the Algebraic Limit Theorem (Theorem 2.3.3) and the Order Limit Theorem (Theorem 2.3.4) proved invaluable in a large number of the arguments that followed. Not surprisingly, we are going to need analogous statements for functional limits. Although it is not difficult to

generate independent proofs for these statements, all of them will follow quite naturally from their sequential analogs once we derive the sequential criterion for functional limits motivated in the opening discussion of this chapter.

Theorem 4.2.3 (Sequential Criterion for Functional Limits). *Given a function $f : A \rightarrow \mathbf{R}$ and a limit point c of A , the following two statements are equivalent:*

- (i) $\lim_{x \rightarrow c} f(x) = L$.
- (ii) *For all sequences $(x_n) \subseteq A$ satisfying $x_n \neq c$ and $(x_n) \rightarrow c$, it follows that $f(x_n) \rightarrow L$.*

Proof. (\Rightarrow) Let's first assume that $\lim_{x \rightarrow c} f(x) = L$. To prove (ii), we consider an arbitrary sequence (x_n) , which converges to c and satisfies $x_n \neq c$. Our goal is to show that the image sequence $f(x_n)$ converges to L . This is most easily seen using the topological formulations of the definition.

Let $\epsilon > 0$. Because we are assuming (i), Definition 4.2.1B implies that there exists $V_\delta(c)$ with the property that all $x \in V_\delta(c)$ different from c satisfy $f(x) \in V_\epsilon(L)$. All we need to do then is argue that our particular sequence (x_n) is eventually in $V_\delta(c)$. But we are assuming that $(x_n) \rightarrow c$. This implies that there exists a point x_N after which $x_n \in V_\delta(c)$. It follows that $n \geq N$ implies $f(x_n) \in V_\epsilon(L)$, as desired.

(\Leftarrow) For this implication, we intend to argue by contradiction. Thus, we assume that statement (ii) is true, and (carefully) negate statement (i). To say that

$$\lim_{x \rightarrow c} f(x) \neq L$$

means that there exists at least one particular $\epsilon_0 > 0$ for which no δ is a suitable response. In other words, no matter what $\delta > 0$ we try, there will always be at least one point

$$x \in V_\delta(c) \quad \text{with} \quad x \neq c \quad \text{for which} \quad f(x) \notin V_{\epsilon_0}(L).$$

For each $n \in \mathbf{N}$, let $\delta_n = 1/n$. From the preceding discussion, it follows that for each $n \in \mathbf{N}$ we may pick an $x_n \in V_{\delta_n}(c)$ with $x_n \neq c$ and $f(x_n) \notin V_{\epsilon_0}(L)$. But now notice that the result of this is a sequence $(x_n) \rightarrow c$ with $x_n \neq c$, where the image sequence $f(x_n)$ certainly does *not* converge to L .

Because this contradicts (ii), which we are assuming is true for this argument, we may conclude that (i) must also hold. \square

Theorem 4.2.3 has several useful corollaries. In addition to the previously advertised benefit of granting us some short proofs of statements about how functional limits interact with algebraic combinations of functions, we also get an economical way of establishing that certain limits do not exist.

Corollary 4.2.4 (Algebraic Limit Theorem for Functional Limits). *Let f and g be functions defined on a domain $A \subseteq \mathbf{R}$, and assume $\lim_{x \rightarrow c} f(x) = L$ and $\lim_{x \rightarrow c} g(x) = M$ for some limit point c of A . Then,*

- (i) $\lim_{x \rightarrow c} kf(x) = kL$ for all $k \in \mathbf{R}$,
- (ii) $\lim_{x \rightarrow c} [f(x) + g(x)] = L + M$,
- (iii) $\lim_{x \rightarrow c} [f(x)g(x)] = LM$, and
- (iv) $\lim_{x \rightarrow c} f(x)/g(x) = L/M$, provided $M \neq 0$.

Proof. These follow from Theorem 4.2.3 and the Algebraic Limit Theorem for sequences. The details are requested in Exercise 4.2.5. \square

Corollary 4.2.5 (Divergence Criterion for Functional Limits). *Let f be a function defined on A , and let c be a limit point of A . If there exist two sequences (x_n) and (y_n) in A with $x_n \neq c$ and $y_n \neq c$ and*

$$\lim x_n = \lim y_n = c \quad \text{but} \quad \lim f(x_n) \neq \lim f(y_n),$$

then we can conclude that the functional limit $\lim_{x \rightarrow c} f(x)$ does not exist.

Example 4.2.6. Assuming the familiar properties of the sine function, let's show that $\lim_{x \rightarrow 0} \sin(1/x)$ does not exist (Fig. 4.5).

If $x_n = 1/2n\pi$ and $y_n = 1/(2n\pi + \pi/2)$, then $\lim(x_n) = \lim(y_n) = 0$. However, $\sin(1/x_n) = 0$ for all $n \in \mathbf{N}$ while $\sin(1/y_n) = 1$. Thus,

$$\lim \sin(1/x_n) \neq \lim \sin(1/y_n),$$

so by Corollary 4.2.5, $\lim_{x \rightarrow 0} \sin(1/x)$ does not exist.




Figure 4.5: THE FUNCTION $\sin(1/x)$ NEAR ZERO.

Exercises

Exercise 4.2.1. Use Definition 4.2.1 to supply a proof for the following limit statements.

(a) $\lim_{x \rightarrow 2} (2x + 4) = 8$.

(b) $\lim_{x \rightarrow 0} x^3 = 0$.

(c) $\lim_{x \rightarrow 2} x^3 = 8$.

(d) $\lim_{x \rightarrow \pi} \llbracket x \rrbracket = 3$, where $\llbracket x \rrbracket$ denotes the greatest integer less than or equal to x .

Exercise 4.2.2. Assume a particular $\delta > 0$ has been constructed as a suitable response to a particular ϵ challenge. Then, any larger/smaller (*pick one*) δ will also suffice.

Exercise 4.2.3. Use Corollary 4.2.5 to show that each of the following limits does not exist.

(a) $\lim_{x \rightarrow 0} |x|/x$

(b) $\lim_{x \rightarrow 1} g(x)$ where g is Dirichlet's function from Section 4.1.

Exercise 4.2.4. Review the definition of Thomae's function $t(x)$ from Section 4.1.

(a) Construct three different sequences (x_n) , (y_n) , and (z_n) , each of which converges to 1 without using the number 1 as a term in the sequence.

(b) Now, compute $\lim t(x_n)$, $\lim t(y_n)$, and $\lim t(z_n)$.

(c) Make an educated conjecture for $\lim_{x \rightarrow 1} t(x)$, and use Definition 4.2.1B to verify the claim. (Given $\epsilon > 0$, consider the set of points $\{x \in \mathbf{R} : t(x) \geq \epsilon\}$. Argue that all the points in this set are isolated.)

Exercise 4.2.5. (a) Supply the details for how Corollary 4.2.4 part (ii) follows from the sequential criterion for functional limits in Theorem 4.2.3 and the Algebraic Limit Theorem for sequences proved in Chapter 2.

(b) Now, write another proof of Corollary 4.2.4 part (ii) directly from Definition 4.2.1 without using the sequential criterion in Theorem 4.2.3.

(c) Repeat (a) and (b) for Corollary 4.2.4 part (iii).

Exercise 4.2.6. Let $g : A \rightarrow \mathbf{R}$ and assume that f is a bounded function on $A \subseteq \mathbf{R}$ (i.e., there exists $M > 0$ satisfying $|f(x)| \leq M$ for all $x \in A$). Show that if $\lim_{x \rightarrow c} g(x) = 0$, then $\lim_{x \rightarrow c} g(x)f(x) = 0$ as well.

Exercise 4.2.7. (a) The statement $\lim_{x \rightarrow 0} 1/x^2 = \infty$ certainly makes intuitive sense. Construct a rigorous definition in the "challenge-response" style of Definition 4.2.1 for a limit statement of the form $\lim_{x \rightarrow c} f(x) = \infty$ and use it to prove the previous statement.

(b) Now, construct a definition for the statement $\lim_{x \rightarrow \infty} f(x) = L$. Show $\lim_{x \rightarrow \infty} 1/x = 0$.

(c) What would a rigorous definition for $\lim_{x \rightarrow \infty} f(x) = \infty$ look like? Give an example of such a limit.

Exercise 4.2.8. Assume $f(x) \geq g(x)$ for all x in some set A on which f and g are defined. Show that for any limit point c of A we must have

$$\lim_{x \rightarrow c} f(x) \geq \lim_{x \rightarrow c} g(x).$$

Exercise 4.2.9 (Squeeze Theorem). Let f, g , and h satisfy $f(x) \leq g(x) \leq h(x)$ for all x in some common domain A . If $\lim_{x \rightarrow c} f(x) = L$ and $\lim_{x \rightarrow c} h(x) = L$ at some limit point c of A , show $\lim_{x \rightarrow c} g(x) = L$ as well.

4.3 Combinations of Continuous Functions

Definition 4.3.1. A function $f : A \rightarrow \mathbf{R}$ is *continuous at a point* $c \in A$ if, for all $\epsilon > 0$, there exists a $\delta > 0$ such that whenever $|x - c| < \delta$ (and $x \in A$) it follows that $|f(x) - f(c)| < \epsilon$.

If f is continuous at every point in the domain A , then we say that f is *continuous on* A .

The definition of continuity looks much like the definition for functional limits, with a few subtle differences. The most important is that we require the point c to be in the domain of f . The value $f(c)$ then becomes the value of $\lim_{x \rightarrow c} f(x)$. With this observation in mind, it is tempting to shorten Definition 4.3.1 to say that f is continuous at $c \in A$ if

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

This is fine except for the very minor fact that the definition of functional limits requires the point c to be a limit point of A , and this is not technically assumed in Definition 4.3.1. Now one consequence of Definition 4.3.1 is that any function is continuous at isolated points of its domain (Exercise 4.3.4), but this is hardly profound. As the name suggests, isolated points are too far from other points of the domain to contribute to any interesting phenomena.

With our attention firmly focused on limit points of the domain, we summarize several equivalent ways to characterize continuity.

Theorem 4.3.2 (Characterizations of Continuity). *Let $f : A \rightarrow \mathbf{R}$, and let $c \in A$ be a limit point of A . The function f is continuous at c if and only if any one of the following conditions is met:*

- (i) *For all $\epsilon > 0$, there exists a $\delta > 0$ such that $|x - c| < \delta$ (and $x \in A$) implies $|f(x) - f(c)| < \epsilon$;*
- (ii) $\lim_{x \rightarrow c} f(x) = f(c)$;
- (iii) *For all $V_\epsilon(f(c))$, there exists a $V_\delta(c)$ with the property that $x \in V_\delta(c)$ (and $x \in A$) implies $f(x) \in V_\epsilon(f(c))$;*
- (iv) *If $(x_n) \rightarrow c$ (with $x_n \in A$), then $f(x_n) \rightarrow f(c)$.*

Proof. Statement (i) is just Definition 4.3.1. Statement (ii) is seen to be equivalent to (i) by applying Definition 4.2.1 and observing that the case $x = c$ (which is excluded in the definition of functional limits) leads to the requirement $f(c) \in V_\epsilon(f(c))$, which is trivially true. Statement (iii) is the standard rewording of (i) using topological neighborhoods in place of the absolute value notation. Finally, statement (iv) follows using an argument nearly identical to that of Theorem 4.2.3 with some slight modifications for when $x_n = c$. \square

The length of this list is somewhat deceiving. Statements (i), (ii), and (iii) are closely related and essentially remind us that functional limits have an ϵ - δ formulation as well as a topological description. Statement (iv), however, is qualitatively different from the first three. As a general rule, the sequential characterization of continuity is usually the most useful for demonstrating that a function is *not* continuous at some point.

Corollary 4.3.3 (Criterion for Discontinuity). *Let $f : A \rightarrow \mathbf{R}$, and let $c \in A$ be a limit point of A . If there exists a sequence $(x_n) \subseteq A$ where $(x_n) \rightarrow c$ but such that $f(x_n)$ does not converge to $f(c)$, we may conclude that f is not continuous at c .*

The sequential characterization of continuity is also important for the other reasons that it was important for functional limits. In particular, it allows us to bring our catalog of results about the behavior of sequences to bear on the study of continuous functions. The next theorem should be compared to Corollary 4.2.3 as well as to Theorem 2.3.3.


Theorem 4.3.4 (Algebraic Continuity Theorem). *Assume $f : A \rightarrow \mathbf{R}$ and $g : A \rightarrow \mathbf{R}$ are continuous at a point $c \in A$. Then,*

- (i) $kf(x)$ is continuous at c for all $k \in \mathbf{R}$;
- (ii) $f(x) + g(x)$ is continuous at c ;
- (iii) $f(x)g(x)$ is continuous at c ; and
- (iv) $f(x)/g(x)$ is continuous at c , provided the quotient is defined.

Proof. All of these statements can be quickly derived from Corollary 4.2.4 and Theorem 4.3.2. \square

These results provide us with the tools we need to firm up our arguments in the opening section of this chapter about the behavior of Dirichlet's function and Thomae's function. The details are requested in Exercise 4.3.6. Later, we provide some more examples of arguments for and against continuity of some familiar functions.

Example 4.3.5. All polynomials are continuous on \mathbf{R} . In fact, rational functions (i.e., quotients of polynomials) are continuous wherever they are defined.

Figure 4.6: THE FUNCTION $x \sin(1/x)$ NEAR ZERO.

To see why this is so, we begin with the elementary claim that $g(x) = x$ and $f(x) = k$, where $k \in \mathbf{R}$, are continuous on \mathbf{R} (Exercise 4.3.3). Because an arbitrary polynomial

$$p(x) = a_0 + a_1x + a_2x^2 + \cdots + a_nx^n$$

consists of sums and products of $g(x)$ with different constant functions, we may conclude from Theorem 4.3.4 that $p(x)$ is continuous.

Likewise, Theorem 4.3.4 implies that quotients of polynomials are continuous as long as the denominator is not zero.

Example 4.3.6. In Example 4.2.6, we saw that the oscillations of $\sin(1/x)$ are so rapid near the origin that $\lim_{x \rightarrow 0} \sin(1/x)$ does not exist. Now, consider the function

$$g(x) = \begin{cases} x \sin(1/x) & \text{if } x \neq 0 \\ 0 & \text{if } x = 0. \end{cases}$$

To investigate the continuity of g at $c = 0$ (Fig. 4.6), we can estimate

$$|g(x) - g(0)| = |x \sin(1/x) - 0| \leq |x|.$$

Given $\epsilon > 0$, set $\delta = \epsilon$, so that whenever $|x - 0| = |x| < \delta$ it follows that $|g(x) - g(0)| < \epsilon$. Thus, g is continuous at the origin.

Example 4.3.7. The greatest integer function $[[x]]$ is defined for all $x \in \mathbf{R}$ by letting $[[x]]$ equal the largest integer $n \in \mathbf{Z}$ satisfying $n \leq x$. This familiar step function certainly has discontinuous “jumps” at each integer value of its domain, but it is a good exercise to try and articulate this observation in the language of analysis.

Given $m \in \mathbf{Z}$, define the sequence (x_n) by $x_n = m - 1/n$. It follows that $(x_n) \rightarrow m$, but

$$h(x_n) \rightarrow (m - 1),$$

which does not equal $m = h(m)$. By Corollary 4.3.3, we see that h fails to be continuous at each $m \in \mathbf{Z}$.

Now let's see why h is continuous at a point $c \notin \mathbf{Z}$. Given $\epsilon > 0$, we must find a δ -neighborhood $V_\delta(c)$ such that $x \in V_\delta(c)$ implies $h(x) \in V_\epsilon(h(c))$. We know that $c \in \mathbf{R}$ falls between consecutive integers $n < c < n + 1$ for some $n \in \mathbf{Z}$. If we take $\delta = \min\{c - n, (n + 1) - c\}$, then it follows from the definition of h that $h(x) = h(c)$ for all $x \in V_\delta(c)$. Thus, we certainly have

$$h(x) \in V_\epsilon(h(c))$$

whenever $x \in V_\delta(c)$.

This latter proof is quite different from the typical situation in that the value of δ does not actually depend on the choice of ϵ . Usually, smaller ϵ 's require smaller δ 's in response, but here the same value of δ works no matter how small ϵ is chosen.

Example 4.3.8. Consider $f(x) = \sqrt{x}$ defined on $A = \{x \in \mathbf{R} : x \geq 0\}$. Exercise 2.3.2 outlines a sequential proof that f is continuous on A . Here, we give an ϵ - δ proof of the same fact.

Let $\epsilon > 0$. We need to argue that $|f(x) - f(c)|$ can be made less than ϵ for all values of x in some δ neighborhood around c . If $c = 0$, this reduces to the statement $\sqrt{x} < \epsilon$, which happens as long as $x < \epsilon^2$. Thus, if we choose $\delta = \epsilon^2$, we see that $|x - 0| < \delta$ implies $|f(x) - 0| < \epsilon$.

For a point $c \in A$ different from zero, we need to estimate $|\sqrt{x} - \sqrt{c}|$. This time, write

$$|\sqrt{x} - \sqrt{c}| = |\sqrt{x} - \sqrt{c}| \left(\frac{\sqrt{x} + \sqrt{c}}{\sqrt{x} + \sqrt{c}} \right) = \frac{|x - c|}{\sqrt{x} + \sqrt{c}} \leq \frac{|x - c|}{\sqrt{c}}.$$

In order to make this quantity less than ϵ , it suffices to pick $\delta = \epsilon\sqrt{c}$. Then, $|x - c| < \delta$ implies

$$|\sqrt{x} - \sqrt{c}| < \frac{\epsilon\sqrt{c}}{\sqrt{c}} = \epsilon,$$

as desired.

Although we have now shown that both polynomials and the square root function are continuous, the Algebraic Continuity Theorem does not provide the justification needed to conclude that a function such as $h(x) = \sqrt{3x^2 + 5}$ is continuous. For this, we must prove that *compositions* of continuous functions are continuous.

Theorem 4.3.9 (Composition of Continuous Functions). *Given $f : A \rightarrow \mathbf{R}$ and $g : B \rightarrow \mathbf{R}$, assume that the range $f(A) = \{f(x) : x \in A\}$ is contained in the domain B so that the composition $g \circ f(x) = g(f(x))$ is well-defined on A .*

If f is continuous at $c \in A$, and if g is continuous at $f(c) \in B$, then $g \circ f$ is continuous at c .

Proof. Exercise 4.3.2. □

Exercises

Exercise 4.3.1. Let $g(x) = \sqrt[3]{x}$.

- (a) Prove that g is continuous at $c = 0$.
 (b) Prove that g is continuous at a point $c \neq 0$. (The identity $a^3 - b^3 = (a - b)(a^2 + ab + b^2)$ will be helpful.)

Exercise 4.3.2. (a) Supply a proof for Theorem 4.3.9 using the ϵ - δ characterization of continuity.

- (b) Give another proof of this theorem using the sequential characterization of continuity (from Theorem 4.3.2 (iv)).

Exercise 4.3.3. Using the ϵ - δ characterization of continuity (and thus using no previous results about sequences), show that the linear function $f(x) = ax + b$ is continuous at every point of \mathbf{R} .

Exercise 4.3.4. (a) Show using Definition 4.3.1 that any function f with domain \mathbf{Z} will necessarily be continuous at every point in its domain.

- (b) Show in general that if c is an isolated point of $A \subseteq \mathbf{R}$, then $f : A \rightarrow \mathbf{R}$ is continuous at c .

Exercise 4.3.5. In Theorem 4.3.4, statement (iv) says that $f(x)/g(x)$ is continuous at c if both f and g are, provided that the quotient is defined. Show that if g is continuous at c and $g(c) \neq 0$, then there exists an open interval containing c on which $f(x)/g(x)$ is always defined.

Exercise 4.3.6. (a) Referring to the proper theorems, give a formal argument that Dirichlet's function from Section 4.1 is nowhere-continuous on \mathbf{R} .

- (b) Review the definition of Thomae's function in Section 4.1 and demonstrate that it fails to be continuous at every rational point.

(c) Use the characterization of continuity in Theorem 4.3.2 (iii) to show that Thomae's function is continuous at every irrational point in \mathbf{R} . (Given $\epsilon > 0$, consider the set of points $\{x \in \mathbf{R} : t(x) \geq \epsilon\}$. Argue that all the points in this set are isolated.)

Exercise 4.3.7. Assume $h : \mathbf{R} \rightarrow \mathbf{R}$ is continuous on \mathbf{R} and let $K = \{x : h(x) = 0\}$. Show that K is a closed set.

Exercise 4.3.8. (a) Show that if a function is continuous on all of \mathbf{R} and equal to 0 at every rational point, then it must be identically 0 on all of \mathbf{R} .

- (b) If f and g are defined on all of \mathbf{R} and $f(r) = g(r)$ at every rational point, must f and g be the same function?

Exercise 4.3.9 (Contraction Mapping Theorem). Let f be a function defined on all of \mathbf{R} , and assume there is a constant c such that $0 < c < 1$ and

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

for all $x, y \in \mathbf{R}$.

- (a) Show that f is continuous on \mathbf{R} .

(b) Pick some point $y_1 \in \mathbf{R}$ and construct the sequence

$$(y_1, f(y_1), f(f(y_1)), \dots).$$

In general, if $y_{n+1} = f(y_n)$, show that the resulting sequence (y_n) is a Cauchy sequence. Hence we may let $y = \lim y_n$.

(c) Prove that y is a fixed point of f (i.e., $f(y) = y$) and that it is unique in this regard.

(d) Finally, prove that if x is *any* arbitrary point in \mathbf{R} then the sequence $(x, f(x), f(f(x)), \dots)$ converges to y defined in (b).

Exercise 4.3.10. Let f be a function defined on all of \mathbf{R} that satisfies the additive condition $f(x + y) = f(x) + f(y)$ for all $x, y \in \mathbf{R}$.

(a) Show that $f(0) = 0$ and that $f(-x) = -f(x)$ for all $x \in \mathbf{R}$.

(b) Show that if f is continuous at $x = 0$, then f is continuous at every point in \mathbf{R} .

(c) Let $k = f(1)$. Show that $f(n) = kn$ for all $n \in \mathbf{N}$, and then prove that $f(z) = kz$ for all $z \in \mathbf{Z}$. Now, prove that $f(r) = kr$ for any rational number r .

(d) Use (b) and (c) to conclude that $f(x) = kx$ for all $x \in \mathbf{R}$. Thus, any additive function that is continuous at $x = 0$ must necessarily be a linear function through the origin.

Exercise 4.3.11. For each of the following choices of A , construct a function $f : \mathbf{R} \rightarrow \mathbf{R}$ that has discontinuities at every point x in A and is continuous on A^c .

(a) $A = \mathbf{Z}$.

(b) $A = \{x : 0 < x < 1\}$.

(c) $A = \{x : 0 \leq x \leq 1\}$.

(d) $A = \{\frac{1}{n} : n \in \mathbf{N}\}$.

Exercise 4.3.12. Let C be the Cantor set constructed in Section 3.1. Define $g : [0, 1] \rightarrow \mathbf{R}$ by

$$g(x) = \begin{cases} 1 & \text{if } x \in C \\ 0 & \text{if } x \notin C. \end{cases}$$

(a) Show that g fails to be continuous at any point $c \in C$.

(b) Prove that g is continuous at every point $c \notin C$.

4.4 Continuous Functions on Compact Sets

Definition 4.4.1. Given a function $f : A \rightarrow \mathbf{R}$ and a subset $B \subseteq A$, let $f(B)$ represent the range of f over the set B ; that is, $f(B) = \{f(x) : x \in B\}$. We say f is *bounded* if $f(A)$ is bounded in the sense of Definition 2.3.1. For a given subset $B \subseteq A$, we say f is *bounded on B* if $f(B)$ is bounded.

The adjectives open, closed, bounded, compact, perfect, and connected are all used to describe subsets of the real line. An interesting question is to sort out which, if any, of these properties are preserved when a particular set $A \subseteq \mathbf{R}$

is mapped to $f(A)$ via a continuous function. For instance, if A is open and f is continuous, is $f(A)$ necessarily open? The answer to this question is no. If $f(x) = x^2$ and A is the open interval $(-1, 1)$, then $f(A)$ is the interval $[0, 1)$, which is not open.

The corresponding conjecture for closed sets also turns out to be false, although constructing a counterexample requires a little more thought. Consider the function

$$g(x) = \frac{1}{1+x^2}$$

and the closed set $A = [0, \infty) = \{x : x \geq 0\}$. Because $g(A) = (0, 1]$ is not closed, we must conclude that continuous functions do not, in general, map closed sets to closed sets. Notice, however, that our particular counterexample required using an *unbounded* closed set A . This is not incidental. Sets that are closed and bounded—that is, compact sets—always get mapped to closed and bounded subsets by continuous functions.

Theorem 4.4.2 (Preservation of Compact Sets). *Let $f : A \rightarrow \mathbf{R}$ be continuous on A . If $K \subseteq A$ is compact, then $f(K)$ is compact as well.*

Proof. Let (y_n) be an arbitrary sequence contained in the range set $f(K)$. To prove this result, we must find a subsequence (y_{n_k}) , which converges to a limit also in $f(K)$. The strategy is to take advantage of the assumption that the domain set K is compact by translating the sequence (y_n) —which is in the range of f —back to a sequence in the domain K .

To assert that $(y_n) \subseteq f(K)$ means that, for each $n \in \mathbf{N}$, we can find (at least one) $x_n \in K$ with $f(x_n) = y_n$. This yields a sequence $(x_n) \subseteq K$. Because K is compact, there exists a convergent subsequence (x_{n_k}) whose limit $x = \lim x_{n_k}$ is also in K . Finally, we make use of the fact that f is assumed to be continuous on A and so is continuous at x in particular. Given that $(x_{n_k}) \rightarrow x$, we conclude that $(y_{n_k}) \rightarrow f(x)$. Because $x \in K$, we have that $f(x) \in f(K)$, and hence $f(K)$ is compact. \square

An extremely important corollary is obtained by combining this result with the observation that compact sets are bounded and contain their supremums and infimums (Exercise 3.3.1).

Theorem 4.4.3 (Extreme Value Theorem). *If $f : K \rightarrow \mathbf{R}$ is continuous on a compact set $K \subseteq \mathbf{R}$, then f attains a maximum and minimum value. In other words, there exists $x_0, x_1 \in K$ such that $f(x_0) \leq f(x) \leq f(x_1)$ for all $x \in K$.*

Proof. Exercise 4.4.3. \square

Uniform Continuity

Although we have proved that polynomials are always continuous on \mathbf{R} , there is an important lesson to be learned by constructing direct proofs that the functions $f(x) = 3x + 1$ and $g(x) = x^2$ (previously studied in Example 4.2.2) are everywhere continuous.

Example 4.4.4. (i) To show directly that $f(x) = 3x + 1$ is continuous at an arbitrary point $c \in \mathbf{R}$, we must argue that $|f(x) - f(c)|$ can be made arbitrarily small for values of x near c . Now,

$$|f(x) - f(c)| = |(3x + 1) - (3c + 1)| = 3|x - c|,$$

so, given $\epsilon > 0$, we choose $\delta = \epsilon/3$. Then, $|x - c| < \delta$ implies

$$|f(x) - f(c)| = 3|x - c| < 3\left(\frac{\epsilon}{3}\right) = \epsilon.$$

Of particular importance for this discussion is the fact that the choice of δ is the same regardless of which point $c \in \mathbf{R}$ we are considering.

(ii) Let's contrast this with what happens when we prove $g(x) = x^2$ is continuous on \mathbf{R} . Given $c \in \mathbf{R}$, we have

$$|g(x) - g(c)| = |x^2 - c^2| = |x - c||x + c|.$$

As discussed in Example 4.2.2, we need an upper bound on $|x + c|$, which is obtained by insisting that our choice of δ not exceed 1. This guarantees that all values of x under consideration will necessarily fall in the interval $(c - 1, c + 1)$. It follows that

$$|x + c| \leq |x| + |c| \leq (|c| + 1) + |c| = 2|c| + 1.$$

Now, let $\epsilon > 0$. If we choose $\delta = \min\{1, \epsilon/(2|c| + 1)\}$, then $|x - c| < \delta$ implies

$$|f(x) - f(c)| = |x - c||x + c| < \left(\frac{\epsilon}{2|c| + 1}\right)(2|c| + 1) = \epsilon.$$

Now, there is nothing deficient about this argument, but it is important to notice that, in the second proof, the algorithm for choosing the response δ depends on the value of c . The statement

$$\delta = \frac{\epsilon}{2|c| + 1}$$

means that larger values of c are going to require smaller values of δ , a fact that should be evident from a consideration of the graph of $g(x) = x^2$ (Fig. 4.7). Given, say, $\epsilon = 1$, a response of $\delta = 1/3$ is sufficient for $c = 1$ because $2/3 < x < 4/3$ certainly implies $0 < x^2 < 2$. However, if $c = 10$, then the steepness of the graph of $g(x)$ means that a much smaller δ is required— $\delta = 1/21$ by our rule—to force $99 < x^2 < 101$.

The next definition is meant to distinguish between these two examples.

Definition 4.4.5. A function $f : A \rightarrow \mathbf{R}$ is *uniformly continuous on A* if for every $\epsilon > 0$ there exists a $\delta > 0$ such that $|x - y| < \delta$ implies $|f(x) - f(y)| < \epsilon$.




Figure 4.7: $g(x) = x^2$; A LARGER c REQUIRES A SMALLER δ .

Recall that to say that “ f is continuous on A ” means that f is continuous at each individual point $c \in A$. In other words, given $\epsilon > 0$ and $c \in A$, we can find a $\delta > 0$ *perhaps depending on c* such that if $|x - c| < \delta$ then $|f(x) - f(c)| < \epsilon$. Uniform continuity is a strictly stronger property. The key distinction between asserting that f is “uniformly continuous on A ” versus simply “continuous on A ” is that, given an $\epsilon > 0$, a single $\delta > 0$ can be chosen that works simultaneously for all points c in A . To say that a function is *not* uniformly continuous on a set A , then, does not necessarily mean it is not continuous at some point. Rather, it means that there is some $\epsilon_0 > 0$ for which no single $\delta > 0$ is a suitable response for all $c \in A$.

Theorem 4.4.6 (Sequential Criterion for Nonuniform Continuity). *A function $f : A \rightarrow \mathbf{R}$ fails to be uniformly continuous on A if there exists a particular $\epsilon_0 > 0$ and two sequences (x_n) and (y_n) in A satisfying*

$$|x_n - y_n| \rightarrow 0 \quad \text{but} \quad |f(x_n) - f(y_n)| \geq \epsilon_0.$$

Proof. Take the logical negation of Definition 4.4.5, and consider the particular values $\delta_n = 1/n$. The details are requested in Exercise 4.4.5. \square

Example 4.4.7. The function $h(x) = \sin(1/x)$ (Fig. 4.5) is continuous at every point in the open interval $(0, 1)$ but is not uniformly continuous on this interval. The problem arises near zero, where the increasingly rapid oscillations take domain values that are quite close together to range values a distance 2

apart. To illustrate Theorem 4.4.6, take $\epsilon_0 = 2$ and set

$$x_n = \frac{1}{\pi/2 + 2n\pi} \quad \text{and} \quad y_n = \frac{1}{3\pi/2 + 2n\pi}.$$

Because each of these sequences tends to zero, we have $|x_n - y_n| \rightarrow 0$, and a short calculation reveals $|h(x_n) - h(y_n)| = 2$ for all $n \in \mathbf{N}$.

Whereas continuity is defined at a single point, uniform continuity is always discussed in reference to a particular domain. In Example 4.4.4, we were not able to prove that $g(x) = x^2$ is uniformly continuous on \mathbf{R} because larger values of x require smaller and smaller values of δ . (As another illustration of Theorem 4.4.6, take $x_n = n$ and $y_n = n + 1/n$.) It is true, however, that $g(x)$ is uniformly continuous on the bounded set $[-10, 10]$. Returning to the argument set forth in Example 4.4.4 (ii), notice that if we restrict our attention to the domain $[-10, 10]$ then $|x + y| \leq 20$ for all x and y . Given $\epsilon > 0$, we can now choose $\delta = \epsilon/20$, and verify that if $x, y \in [-10, 10]$ satisfy $|x - y| < \delta$, then

$$|f(x) - f(y)| = |x^2 - y^2| = |x - y||x + y| < \left(\frac{\epsilon}{20}\right) 20 = \epsilon.$$

In fact, it is not difficult to see how to modify this argument to show that $g(x)$ is uniformly continuous on any bounded set A in \mathbf{R} .

Now, Example 4.4.7 is included to keep us from jumping to the erroneous conclusion that functions that are continuous on bounded domains are necessarily uniformly continuous. A general result does follow, however, if we assume that the domain is compact.

Theorem 4.4.8. *A function that is continuous on a compact set K is uniformly continuous on K .*

Proof. Assume $f : K \rightarrow \mathbf{R}$ is continuous at every point of a compact set $K \subseteq \mathbf{R}$. To prove that f is uniformly continuous on K we argue by contradiction.

By the criterion in Theorem 4.4.6, if f is not uniformly continuous on K , then there exist two sequences (x_n) and (y_n) in K such that

$$\lim |x_n - y_n| = 0 \quad \text{while} \quad |f(x_n) - f(y_n)| \geq \epsilon_0$$

for some particular $\epsilon_0 > 0$. Because K is compact, the sequence (x_n) has a convergent subsequence (x_{n_k}) with $x = \lim x_{n_k}$ also in K .

We could use the compactness of K again to produce a convergent subsequence of (y_n) , but notice what happens when we consider the particular subsequence (y_{n_k}) consisting of those terms in (y_n) that correspond to the terms in the convergent subsequence (x_{n_k}) . By the Algebraic Limit Theorem,

$$\lim(y_{n_k}) = \lim((y_{n_k} - x_{n_k}) + x_{n_k}) = 0 + x.$$

The conclusion is that both (x_{n_k}) and (y_{n_k}) converge to $x \in K$. Because f is assumed to be continuous at x , we have $\lim f(x_{n_k}) = f(x)$ and $\lim f(y_{n_k}) = f(x)$, which implies

$$\lim(f(x_{n_k}) - f(y_{n_k})) = 0.$$

A contradiction arises when we recall that (x_n) and (y_n) were chosen to satisfy

$$|f(x_n) - f(y_n)| \geq \epsilon_0$$

for all $n \in \mathbf{N}$. We conclude, then, that f is indeed uniformly continuous on K . \square

Exercises

Exercise 4.4.1. (a) Show that $f(x) = x^3$ is continuous on all of \mathbf{R} .

(b) Argue, using Theorem 4.4.6, that f is not uniformly continuous on \mathbf{R} .

(c) Show that f is uniformly continuous on any bounded subset of \mathbf{R} .

Exercise 4.4.2. Show that $f(x) = 1/x^2$ is uniformly continuous on the set $[1, \infty)$ but not on the set $(0, 1]$.

Exercise 4.4.3. Furnish the details (including an argument for Exercise 3.3.1 if it is not already done) for the proof of the Extreme Value Theorem (Theorem 4.4.3).

Exercise 4.4.4. Show that if f is continuous on $[a, b]$ with $f(x) > 0$ for all $a \leq x \leq b$, then $1/f$ is bounded on $[a, b]$.

Exercise 4.4.5. Using the advice that follows Theorem 4.4.6, provide a complete proof for this criterion for nonuniform continuity.

Exercise 4.4.6. Give an example of each of the following, or state that such a request is impossible. For any that are impossible, supply a short explanation (perhaps referencing the appropriate theorem(s)) for why this is the case.

(a) a continuous function $f : (0, 1) \rightarrow \mathbf{R}$ and a Cauchy sequence (x_n) such that $f(x_n)$ is not a Cauchy sequence;

(b) a continuous function $f : [0, 1] \rightarrow \mathbf{R}$ and a Cauchy sequence (x_n) such that $f(x_n)$ is not a Cauchy sequence;

(c) a continuous function $f : [0, \infty) \rightarrow \mathbf{R}$ and a Cauchy sequence (x_n) such that $f(x_n)$ is not a Cauchy sequence;

(d) a continuous bounded function f on $(0, 1)$ that attains a maximum value on this open interval but not a minimum value.

Exercise 4.4.7. Assume that g is defined on an open interval (a, c) and it is known to be uniformly continuous on $(a, b]$ and $[b, c)$, where $a < b < c$. Prove that g is uniformly continuous on (a, c) .

Exercise 4.4.8. (a) Assume that $f : [0, \infty) \rightarrow \mathbf{R}$ is continuous at every point in its domain. Show that if there exists $b > 0$ such that f is uniformly continuous on the set $[b, \infty)$, then f is uniformly continuous on $[0, \infty)$.

(b) Prove that $f(x) = \sqrt{x}$ is uniformly continuous on $[0, \infty)$.

Exercise 4.4.9. A function $f : A \rightarrow \mathbf{R}$ is called *Lipschitz* if there exists a bound $M > 0$ such that

$$\left| \frac{f(x) - f(y)}{x - y} \right| \leq M$$

for all $x, y \in A$. Geometrically speaking, a function f is Lipschitz if there is a uniform bound on the magnitude of the slopes of lines drawn through any two points on the graph of f .

(a) Show that if $f : A \rightarrow \mathbf{R}$ is Lipschitz, then it is uniformly continuous on A .

(b) Is the converse statement true? Are all uniformly continuous functions necessarily Lipschitz?

Exercise 4.4.10. Do uniformly continuous functions preserve boundedness? If f is uniformly continuous on a bounded set A , is $f(A)$ necessarily bounded?

Exercise 4.4.11 (Topological Characterization of Continuity). Let g be defined on all of \mathbf{R} . If A is a subset of \mathbf{R} , define the set $g^{-1}(A)$ by

$$g^{-1}(A) = \{x \in \mathbf{R} : g(x) \in A\}.$$

Show that g is continuous if and only if $g^{-1}(O)$ is open whenever $O \subseteq \mathbf{R}$ is an open set.

Exercise 4.4.12. Construct an alternate proof of Theorem 4.4.8 using the open cover characterization of compactness from Theorem 3.3.8 (iii).

Exercise 4.4.13 (Continuous Extension Theorem). (a) Show that a uniformly continuous function preserves Cauchy sequences; that is, if $f : A \rightarrow \mathbf{R}$ is uniformly continuous and $(x_n) \subseteq A$ is a Cauchy sequence, then show $f(x_n)$ is a Cauchy sequence.

(b) Let g be a continuous function on the open interval (a, b) . Prove that g is uniformly continuous on (a, b) if and only if it is possible to define values $g(a)$ and $g(b)$ at the endpoints so that the extended function g is continuous on $[a, b]$. (In the forward direction, first produce candidates for $g(a)$ and $g(b)$, and then show the extended g is continuous.)

4.5 The Intermediate Value Theorem

The Intermediate Value Theorem (IVT) is the name given to the very intuitive observation that a continuous function f on a closed interval $[a, b]$ attains every value that falls between the range values $f(a)$ and $f(b)$ (Fig. 4.8).

Here is this observation in the language of analysis.

Theorem 4.5.1 (Intermediate Value Theorem). *If $f : [a, b] \rightarrow \mathbf{R}$ is continuous, and if L is a real number satisfying $f(a) < L < f(b)$ or $f(a) > L > f(b)$, then there exists a point $c \in (a, b)$ where $f(c) = L$.*




Figure 4.8: INTERMEDIATE VALUE THEOREM.

This theorem was freely used by mathematicians of the 18th century (including Euler and Gauss) without any consideration of its validity. In fact, the first analytical proof was not offered until 1817 by Bolzano in a paper that also contains the first appearance of a somewhat modern definition of continuity. This emphasizes the significance of this result. As discussed in Section 4.1, Bolzano and his contemporaries had arrived at a point in the evolution of mathematics where it was becoming increasingly important to firm up the foundations of the subject. Doing so, however, was not simply a matter of going back and supplying the missing proofs. The real battle lay in first obtaining a thorough and mutually agreed-upon understanding of the relevant concepts. The importance of the Intermediate Value Theorem for us is similar in that our understanding of continuity and the nature of the real line is now mature enough for a proof *to be possible*. Indeed, there are several satisfying arguments for this simple result, each one isolating, in a slightly different way, the interplay between continuity and completeness.

Preservation of Connected Sets

The most potentially useful way to understand the Intermediate Value Theorem (IVT) is as a special case of the fact that continuous functions map connected sets to connected sets. In Theorem 4.4.2, we saw that if f is a continuous function on a compact set K , then the range set $f(K)$ is also compact. The analogous observation holds for connected sets.

Theorem 4.5.2 (Preservation of Connectedness). *Let $f : A \rightarrow \mathbf{R}$ be continuous. If $E \subseteq A$ is connected, then $f(E)$ is connected as well.*

Proof. Intending to use the characterization of connected sets in Theorem 3.4.6,

let $f(E) = A \cup B$ where A and B are disjoint and nonempty. Our goal is to produce a sequence contained in one of these sets that converges to a limit in the other.

Let

$$C = \{x \in E : f(x) \in A\} \quad \text{and} \quad D = \{x \in E : f(x) \in B\}.$$

The sets C and D are called the *preimages* of A and B , respectively. Using the properties of A and B , it is straightforward to check that C and D are nonempty and disjoint and satisfy $E = C \cup D$. Now, we are assuming E is a connected set, so by Theorem 3.4.6, there exists a sequence (x_n) contained in one of C or D with $x = \lim x_n$ contained in the other. Finally, because f is continuous at x , we get $f(x) = \lim f(x_n)$. Thus, it follows that $f(x_n)$ is a convergent sequence contained in either A or B while the limit $f(x)$ is an element of the other. With another nod to Theorem 3.4.6, the proof is complete. \square

In \mathbf{R} , a set is connected if and only if it is a (possibly unbounded) interval. This fact, together with Theorem 4.5.2, leads to a short proof of the Intermediate Value Theorem (Exercise 4.5.1). We should point out that the proof of Theorem 4.5.2 does not make use of the equivalence between connected sets and intervals in \mathbf{R} but relies only on the general definitions. The previous comment that this is the most *useful* way to approach IVT stems from the fact that, although it is not discussed here, the definitions of continuity and connectedness can be easily adapted to higher-dimensional settings. Theorem 4.5.2, then, remains a valid conclusion in higher dimensions, whereas the Intermediate Value Theorem is essentially a one-dimensional result.

Completeness

A typical way the Intermediate Value Theorem is applied is to prove the existence of roots. Given $f(x) = x^2 - 2$, for instance, we see that $f(1) = -1$ and $f(2) = 2$. Therefore, there exists a point $c \in (0, 1)$ where $f(c) = 0$.


In this case, we can easily compute $c = \sqrt{2}$, meaning that we really did not need IVT to show that f has a root. We spent a good deal of time in Chapter 1 proving that $\sqrt{2}$ exists, which was only possible once we insisted on the Axiom of Completeness as part of our assumptions about the real numbers. The fact that the Intermediate Value Theorem has just asserted that $\sqrt{2}$ exists suggests that another way to understand this result is in terms of the relationship between the continuity of f and the completeness of \mathbf{R} .

The Axiom of Completeness (AoC) from the first chapter states that “Sets which are bounded above have least upper bounds.” Later, we saw that the Nested Interval Property (NIP) is an equivalent way to assert that the real numbers have no “gaps.” Either of these characterizations of completeness can be used as the cornerstone for an alternate proof of Theorem 4.5.1.

Proof. I. (First approach using AoC.) To simplify matters a bit, let’s consider the special case where f is a continuous function satisfying $f(a) < 0 < f(b)$ and

show that $f(c) = 0$ for some $c \in (a, b)$. First let

$$K = \{x \in [a, b] : f(x) \leq 0\}.$$



Notice that K is bounded above by b , and $a \in K$ so K is not empty. Thus we may appeal to the Axiom of Completeness to assert that $c = \sup K$ exists.

There are three cases to consider:


$$f(c) > 0, f(c) < 0, \text{ and } f(c) = 0.$$

The fact that c is the least upper bound of K can be used to rule out the first two cases, resulting in the desired conclusion that $f(c) = 0$. The details are requested in Exercise 4.5.5.

II. (*Second approach using NIP.*) Again, consider the special case where $L = 0$ and $f(a) < 0 < f(b)$. Let $I_0 = [a, b]$, and consider the midpoint

$$z = (a + b)/2.$$

If $f(z) \geq 0$, then set $a_1 = a$ and $b_1 = z$. If $f(z) < 0$, then set $a_1 = z$ and $b_1 = b$. In either case, the interval $I_1 = [a_1, b_1]$ has the property that f is negative at the left endpoint and nonnegative at the right.



This procedure can be inductively repeated, setting the stage for an application of the Nested Interval Property. The remainder of the argument is left as Exercise 4.5.6.

□

The Intermediate Value Property

Does the Intermediate Value Theorem have a converse?

Definition 4.5.3. A function f has the *intermediate value property* on an interval $[a, b]$ if for all $x < y$ in $[a, b]$ and all L between $f(x)$ and $f(y)$, it is always possible to find a point $c \in (x, y)$ where $f(c) = L$.

Another way to summarize the Intermediate Value Theorem is to say that every continuous function on $[a, b]$ has the intermediate value property. There is an understandable temptation to suspect that any function that has the intermediate value property must necessarily be continuous, but that is not the case. We have seen that

$$g(x) = \begin{cases} \sin(1/x) & \text{if } x \neq 0 \\ 0 & \text{if } x = 0 \end{cases}$$

is not continuous at zero (Example 4.2.6), but it does have the intermediate value property on $[0, 1]$.

The intermediate value property *does* imply continuity if we insist that our function is monotone (Exercise 4.5.4).

Exercises

Exercise 4.5.1. Show how the Intermediate Value Theorem follows as a corollary to Theorem 4.5.2.

Exercise 4.5.2. Decide on the validity of the following conjectures.

(a) Continuous functions take bounded open intervals to bounded open intervals.

(b) Continuous functions take bounded open intervals to open sets.

(c) Continuous functions take bounded closed intervals to bounded closed intervals.

Exercise 4.5.3. Is there a continuous function on all of \mathbf{R} with range $f(\mathbf{R})$ equal to \mathbf{Q} ?

Exercise 4.5.4. A function f is *increasing* on A if $f(x) \leq f(y)$ for all $x < y$ in A . Show that the Intermediate Value Theorem does have a converse if we assume f is increasing on $[a, b]$.

Exercise 4.5.5. Finish the proof of the Intermediate Value Theorem using the Axiom of Completeness started previously.

Exercise 4.5.6. Finish the proof of the Intermediate Value Theorem using the Nested Interval Property started previously.

Exercise 4.5.7. Let f be a continuous function on the closed interval $[0, 1]$ with range also contained in $[0, 1]$. Prove that f must have a fixed point; that is, show $f(x) = x$ for at least one value of $x \in [0, 1]$.

Exercise 4.5.8. Imagine a clock where the hour hand and the minute hand are indistinguishable from each other. Assuming the hands move continuously around the face of the clock, and assuming their positions can be measured with perfect accuracy, is it always possible to determine the time?

4.6 Sets of Discontinuity

Given a function $f : \mathbf{R} \rightarrow \mathbf{R}$, define $D_f \subseteq \mathbf{R}$ to be the set of points where the function f fails to be continuous. In Section 4.1, we saw that Dirichlet's function $g(x)$ had $D_g = \mathbf{R}$. The modification $h(x)$ of Dirichlet's function had $D_h = \mathbf{R} \setminus \{0\}$, zero being the only point of continuity. Finally, for Thomae's function $t(x)$, we saw that $D_t = \mathbf{Q}$.

Exercise 4.6.1. Using modifications of these functions, construct a function $f : \mathbf{R} \rightarrow \mathbf{R}$ so that

- (a) $D_f = \mathbf{Z}$.
- (b) $D_f = \{x : 0 < x \leq 1\}$.

We concluded the introduction with a question about whether D_f could take the form of *any* arbitrary subset of the real line. As it turns out, this is not the case. The set of discontinuities of a real-valued function on \mathbf{R} has a specific topological structure that is not possessed by every subset of \mathbf{R} . Specifically, D_f , no matter how f is chosen, can always be written as the countable union of closed sets. In the case where f is *monotone*, these closed sets can be taken to be single points.

Monotone Functions

Classifying D_f for an arbitrary f is somewhat involved, so it is interesting that describing D_f is fairly straightforward for the class of monotone functions.

Definition 4.6.1. A function $f : A \rightarrow \mathbf{R}$ is *increasing on A* if $f(x) \leq f(y)$ whenever $x < y$ and *decreasing* if $f(x) \geq f(y)$ whenever $x < y$ in A . A *monotone* function is one that is either increasing or decreasing.

Continuity of f at a point c means that $\lim_{x \rightarrow c} f(x) = f(c)$. One particular way for a discontinuity to occur is if the limit from the right at c is different from the limit from the left at c . As always with new terminology, we need to be precise about what we mean by “from the left” and “from the right.”

Definition 4.6.2 (Right-hand limit). Given a limit point c of a set A and a function $f : A \rightarrow \mathbf{R}$, we write

$$\lim_{x \rightarrow c^+} f(x) = L$$

if for all $\epsilon > 0$ there exists a $\delta > 0$ such that $|f(x) - L| < \epsilon$ whenever $0 < x - c < \delta$.

Equivalently, in terms of sequences, $\lim_{x \rightarrow c^+} f(x) = L$ if $\lim f(x_n) = L$ for all sequences (x_n) satisfying $x_n > c$ and $\lim(x_n) = c$.

Exercise 4.6.2. State a similar definition for the left-hand limit

$$\lim_{x \rightarrow c^-} f(x) = L.$$

Theorem 4.6.3. Given $f : A \rightarrow \mathbf{R}$ and a limit point c of A , $\lim_{x \rightarrow c} f(x) = L$ if and only if

$$\lim_{x \rightarrow c^+} f(x) = L \quad \text{and} \quad \lim_{x \rightarrow c^-} f(x) = L.$$

Exercise 4.6.3. Supply a proof for this proposition.

Generally speaking, discontinuities can be divided into three categories:

- (i) If $\lim_{x \rightarrow c} f(x)$ exists but has a value different from $f(c)$, the discontinuity at c is called *removable*.
- (ii) If $\lim_{x \rightarrow c^+} f(x) \neq \lim_{x \rightarrow c^-} f(x)$, then f has a *jump* discontinuity at c .
- (iii) If $\lim_{x \rightarrow c} f(x)$ does not exist for some other reason, then the discontinuity at c is called an *essential* discontinuity.

We are now equipped to characterize the set D_f for an arbitrary monotone function f .

Exercise 4.6.4. Let $f : \mathbf{R} \rightarrow \mathbf{R}$ be increasing. Prove that $\lim_{x \rightarrow c^+} f(x)$ and $\lim_{x \rightarrow c^-} f(x)$ must each exist at every point $c \in \mathbf{R}$. Argue that the only type of discontinuity a monotone function can have is a jump discontinuity.

Exercise 4.6.5. Construct a bijection between the set of jump discontinuities of a monotone function f and a subset of \mathbf{Q} . Conclude that D_f for a monotone function f must either be finite or countable, but not uncountable.

D_f for an Arbitrary Function

Recall that the intersection of an infinite collection of closed sets is closed, but for unions we must restrict ourselves to *finite* collections of closed sets in order to ensure the union is closed. For open sets the situation is reversed. The arbitrary union of open sets is open, but only finite intersections of open sets are necessarily open.

Definition 4.6.4. A set that can be written as the countable union of closed sets is in the class F_σ . (This definition also appeared in Section 3.5.)

To this point, we have constructed functions where the set of discontinuity has been \mathbf{R} (Dirichlet's function), $\mathbf{R} \setminus \{0\}$ (modified Dirichlet function), \mathbf{Q} (Thomae's function), \mathbf{Z} , and $(0, 1]$ (Exercise 4.6.1).

Exercise 4.6.6. Show that in each case we get an F_σ set as the set where each function is discontinuous.

The upcoming argument depends on a concept called α -continuity.

Definition 4.6.5. Let f be defined on \mathbf{R} , and let $\alpha > 0$. The function f is α -continuous at $x \in \mathbf{R}$ if there exists a $\delta > 0$ such that for all $y, z \in (x - \delta, x + \delta)$ it follows that $|f(y) - f(z)| < \alpha$.

The most important thing to note about this definition is that there is no “for all” in front of the $\alpha > 0$. As we will investigate, adding this quantifier would make this definition equivalent to our definition of continuity. In a sense, α -continuity is a measure of the variation of the function in the neighborhood of a particular point. A function is α -continuous at a point c if there is some interval centered at c in which the variation of the function never exceeds the value $\alpha > 0$.

Given a function f on \mathbf{R} , define D_α to be the set of points where the function f fails to be α -continuous. In other words,

$$D_\alpha = \{x \in \mathbf{R} : f \text{ is not } \alpha\text{-continuous at } x\}.$$

Exercise 4.6.7. Prove that, for a fixed $\alpha > 0$, the set D_α is closed.

The stage is set. It is time to characterize the set of discontinuity for an arbitrary function f on \mathbf{R} .

Theorem 4.6.6. Let $f : \mathbf{R} \rightarrow \mathbf{R}$ be an arbitrary function. Then, D_f is an F_σ set.

Proof. Recall that

$$D_f = \{x \in \mathbf{R} : f \text{ is not continuous at } x\}.$$

Exercise 4.6.8. If $\alpha_1 < \alpha_2$, show that $D_{\alpha_2} \subseteq D_{\alpha_1}$.

Exercise 4.6.9. Let $\alpha > 0$ be given. Show that if f is continuous at x , then it is α -continuous at x as well. Explain how it follows that $D_\alpha \subseteq D_f$.

Exercise 4.6.10. Show that if f is not continuous at x , then f is not α -continuous for some $\alpha > 0$. Now explain why this guarantees that

$$D_f = \bigcup_{n=1}^{\infty} D_{\frac{1}{n}}.$$

Because each $D_{\frac{1}{n}}$ is closed, the proof is complete. □

4.7 Epilogue

Theorem 4.6.6 is only interesting if we can demonstrate that not every subset of \mathbf{R} is in an F_σ set. This takes some effort and was included as an exercise in Section 3.5 on the Baire Category Theorem. Baire’s Theorem states that if \mathbf{R} is written as the countable union of closed sets, then at least one of these sets must contain a nonempty open interval. Now \mathbf{Q} is the countable union of singleton points, and we can view each point as a closed set that obviously contains no

intervals. If the set of irrationals \mathbf{I} were a countable union of closed sets, it would have to be that none of these closed sets contained any open intervals or else they would then contain some rational numbers. But this leads to a contradiction to Baire's Theorem. Thus, \mathbf{I} is not the countable union of closed sets, and consequently it is not an F_σ set. We may therefore conclude that there is no function f that is continuous at every rational point and discontinuous at every irrational point. This should be compared with Thomae's function discussed earlier.

The converse question is interesting as well. Given an arbitrary F_σ set, W.H. Young showed in 1903 that it is always possible to construct a function that has discontinuities precisely on this set. His construction involves the same Dirichlet-type definitions we have seen but is understandably more intricate. By contrast, a function demonstrating the converse for the monotone case is not too difficult to describe. Let

$$D = \{x_1, x_2, x_3, x_4, \dots\}$$

be an arbitrary countable set of real numbers. In order to construct a monotone function that has discontinuities precisely on D , intuitively attach a "weight" of $1/2^n$ to each point $x_n \in D$. Now, define

$$f(x) = \sum_{n: x_n < x} \frac{1}{2^n}$$

where for each $x \in \mathbf{R}$ the sum is intended to be taken over all of those weights corresponding to points to the left of x . (If there are no points in D to the left of x , then set $f(x) = 0$.) Any worries about the order of the sum can be alleviated by observing that the convergence is absolute. It is not too hard to show that the resulting function f is monotone and has jump discontinuities of size $1/2^n$ at each point x_n in D , as desired (Exercise 6.4.8).

Chapter 5

The Derivative

5.1 Discussion: Are Derivatives Continuous?

The geometric motivation for the derivative is most likely familiar territory. Given a function $g(x)$, the derivative $g'(x)$ is understood to be the slope of the graph of g at each point x in the domain. A graphical picture (Fig. 5.1) reveals the impetus behind the mathematical definition

$$g'(c) = \lim_{x \rightarrow c} \frac{g(x) - g(c)}{x - c}.$$

The difference quotient $(g(x) - g(c))/(x - c)$ represents the slope of the line through the two points $(x, g(x))$ and $(c, g(c))$. By taking the limit as x approaches c , we arrive at a well-defined mathematical meaning for the slope of the tangent line at $x = c$.

The myriad applications of the derivative function are the topic of much of the calculus sequence, as well as several other upper-level courses in mathematics. None of these applied questions are pursued here in any length, but it should be pointed out that the rigorous underpinnings for differentiation worked




Figure 5.1: DEFINITION OF $g'(c)$.

out in this chapter are an essential foundation for any applied study. Eventually, as the derivative is subjected to more and more complex manipulations, it becomes crucial to know precisely how differentiation is defined and how it interacts with other mathematical operations.

Although physical applications are not explicitly discussed, we will encounter several questions of a more abstract quality as we develop the theory. Most of these are concerned with the relationship between differentiation and continuity. Are continuous functions always differentiable? If not, how nondifferentiable can a continuous function be? Are differentiable functions continuous? Given that a function f has a derivative at every point in its domain, what can we say about the function f' ? Is f' continuous? How accurately can we describe the set of all possible derivatives, or are there no restrictions? Put another way, if we are given an arbitrary function g , is it always possible to find a differentiable function f such that $f' = g$, or are there some properties that g must possess for this to occur? In our study of continuity, we saw that restricting our attention to monotone functions had a significant impact on the answers to questions about sets of discontinuity. What effect, if any, does this same restriction have on our questions about potential sets of nondifferentiable points? Some of these issues are harder to resolve than others, and some remain unanswered in any satisfactory way.

A particularly useful class of examples for this discussion are functions of the form

$$g_n(x) = \begin{cases} x^n \sin(1/x) & \text{if } x \neq 0 \\ 0 & \text{if } x = 0. \end{cases}$$

When $n = 0$, we have seen (Example 4.2.6) that the oscillations of $\sin(1/x)$ prevent $g_0(x)$ from being continuous at $x = 0$. When $n = 1$, these oscillations are squeezed between $|x|$ and $-|x|$, the result being that g_1 is continuous at $x = 0$ (Example 4.3.6). Is $g'_1(0)$ defined? Using the preceding definition, we get

$$g'_1(0) = \lim_{x \rightarrow 0} \frac{g_1(x)}{x} = \lim_{x \rightarrow 0} \sin(1/x),$$

which, as we now know, does not exist. Thus, g_1 is not differentiable at $x = 0$. On the other hand, the same calculation shows that g_2 is differentiable at zero. In fact, we have

$$g'_2(0) = \lim_{x \rightarrow 0} x \sin(1/x) = 0.$$

At points different from zero, we can use the familiar rules of differentiation (soon to be justified) to conclude that g_2 is differentiable everywhere in \mathbf{R} with

$$g'_2(x) = \begin{cases} -\cos(1/x) + 2x \sin(1/x) & \text{if } x \neq 0 \\ 0 & \text{if } x = 0. \end{cases}$$

But now consider

$$\lim_{x \rightarrow 0} g'_2(x).$$

Because the $\cos(1/x)$ term is not preceded by a factor of x , we must conclude that this limit does not exist and that, consequently, the derivative function is




Figure 5.2: THE FUNCTION $g_2(x) = x^2 \sin(1/x)$ NEAR ZERO.

not continuous. To summarize, the function $g_2(x)$ is continuous and differentiable everywhere on \mathbf{R} (Fig. 5.2), the derivative function g_2' is thus defined everywhere on \mathbf{R} , but g_2' has a discontinuity at zero. The conclusion is that derivatives need not, in general, be continuous!

The discontinuity in g_2' is *essential*, meaning $\lim_{x \rightarrow 0} g_2'(x)$ does not exist as a one-sided limit. But, what about a function with a simple jump discontinuity? For example, does there exist a function h such that

$$h'(x) = \begin{cases} -1 & \text{if } x \leq 0 \\ 1 & \text{if } x > 0. \end{cases}$$

A first impression may bring to mind the absolute value function, which has slopes of -1 at points to the left of zero and slopes of 1 to the right. However, the absolute value function is *not* differentiable at zero. We are seeking a function that is differentiable everywhere, including the point zero, where we are insisting that the slope of the graph be -1 . The degree of difficulty of this request should start to become apparent. Without sacrificing differentiability at any point, we are demanding that the slopes jump from -1 to 1 and not attain any value in between.

Although we have seen that continuity is not a required property of derivatives, the intermediate value property will prove a more stubborn quality to ignore.

5.2 Derivatives and the Intermediate Value Property

Although the definition would technically make sense for more complicated domains, all of the interesting results about the relationship between a function and its derivative require that the domain of the given function be an interval.

Thinking geometrically of the derivative as a rate of change, it should not be too surprising that we would want to confine the independent variable to move about a connected domain.

The theory of functional limits from Section 4.2 is all that is needed to supply a rigorous definition for the derivative.

Definition 5.2.1. Let $g : A \rightarrow \mathbf{R}$ be a function defined on an interval A . Given $c \in A$, the *derivative of g at c* is defined by

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

provided this limit exists. If g' exists for all points $c \in A$, we say that g is *differentiable on A* .

Example 5.2.2. (a) Consider $f(x) = x^n$, where $n \in \mathbf{N}$, and let c be any arbitrary point in \mathbf{R} . Using the algebraic identity

$$x^n - c^n = (x - c)(x^{n-1} + cx^{n-2} + c^2x^{n-3} + \cdots + c^{n-1}),$$

we can calculate the familiar formula

$$\begin{aligned} f'(c) &= \lim_{x \rightarrow c} \frac{x^n - c^n}{x - c} = \lim_{x \rightarrow c} (x^{n-1} + cx^{n-2} + c^2x^{n-3} + \cdots + c^{n-1}) \\ &= c^{n-1} + c^{n-1} + \cdots + c^{n-1} = nc^{n-1}. \end{aligned}$$

(b) If $g(x) = |x|$, then attempting to compute the derivative at $c = 0$ produces the limit

$$g'(0) = \lim_{x \rightarrow 0} \frac{|x|}{x},$$

which is $+1$ or -1 depending on whether x approaches zero from the right or left. Consequently, this limit does not exist, and we conclude that g is not differentiable at zero.

Example 5.2.2 (b) is a reminder that continuity of g does not imply that g is necessarily differentiable. On the other hand, if g is differentiable at a point, then it is true that g must be continuous at this point.

Theorem 5.2.3. *If $g : A \rightarrow \mathbf{R}$ is differentiable at a point $c \in A$, then g is continuous at c as well.*

Proof. We are assuming that

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

exists, and we want to prove that $\lim_{x \rightarrow c} g(x) = g(c)$. But notice that the Algebraic Limit Theorem for functional limits allows us to write

$$\lim_{x \rightarrow c} (g(x) - g(c)) = \lim_{x \rightarrow c} \left(\frac{g(x) - g(c)}{x - c} \right) (x - c) = g'(c) \cdot 0 = 0.$$

It follows that $\lim_{x \rightarrow c} g(x) = g(c)$. □

Combinations of Differentiable Functions

The Algebraic Limit Theorem (Theorem 2.3.3) led easily to the conclusion that algebraic combinations of continuous functions are continuous. With only slightly more work, we arrive at a similar conclusion for sums, products, and quotients of differentiable functions.

Theorem 5.2.4. *Let f and g be functions defined on an interval A , and assume both are differentiable at some point $c \in A$. Then,*

- (i) $(f + g)'(c) = f'(c) + g'(c)$,
- (ii) $(kf)'(c) = kf'(c)$, for all $k \in \mathbf{R}$,
- (iii) $(fg)'(c) = f'(c)g(c) + f(c)g'(c)$, and
- (iv) $(f/g)'(c) = \frac{g(c)f'(c) - f(c)g'(c)}{[g(c)]^2}$, provided that $g(c) \neq 0$.

Proof. Statements (i) and (ii) are left as exercises. To prove (iii), we rewrite the difference quotient as

$$\begin{aligned} \frac{(fg)(x) - (fg)(c)}{x - c} &= \frac{f(x)g(x) - f(x)g(c) + f(x)g(c) - f(c)g(c)}{x - c} \\ &= f(x) \left[\frac{g(x) - g(c)}{x - c} \right] + g(c) \left[\frac{f(x) - f(c)}{x - c} \right]. \end{aligned}$$

Because f is differentiable at c , it is continuous there and thus $\lim_{x \rightarrow c} f(x) = f(c)$. This fact, together with the functional-limit version of the Algebraic Limit Theorem (Theorem 4.2.4) justifies the conclusion

$$\lim_{x \rightarrow c} \frac{(fg)(x) - (fg)(c)}{x - c} = f(c)g'(c) + f'(c)g(c).$$

A similar proof of (iv) is possible, or we can use an argument based on the next result. Each of these options is each discussed in Exercise 5.2.2. \square

The composition of two differentiable functions also fortunately results in another differentiable function. This fact is referred to as the *chain rule*. To discover the proper formula for the derivative of the composition $g \circ f$, we can write

$$\begin{aligned} (g \circ f)'(c) &= \lim_{x \rightarrow c} \frac{g(f(x)) - g(f(c))}{x - c} = \lim_{x \rightarrow c} \frac{g(f(x)) - g(f(c))}{f(x) - f(c)} \cdot \frac{f(x) - f(c)}{x - c} \\ &= g'(f(c)) \cdot f'(c). \end{aligned}$$

With a little polish, this string of equations could qualify as a proof except for the pesky fact that the $f(x) - f(c)$ expression causes problems in the denominator if $f(x) = f(c)$ for x values in arbitrarily small neighborhoods of c . (The function $g_2(x)$ discussed in Section 5.1 exhibits this behavior near $c = 0$.) The upcoming proof of the chain rule manages to finesse this problem but in content is essentially the argument just given.

Theorem 5.2.5 (Chain Rule). *Let $f : A \rightarrow \mathbf{R}$ and $g : B \rightarrow \mathbf{R}$ satisfy $f(A) \subseteq B$ so that the composition $g \circ f$ is well-defined. If f is differentiable at $c \in A$ and if g is differentiable at $f(c) \in B$, then $g \circ f$ is differentiable at c with $(g \circ f)'(c) = g'(f(c)) \cdot f'(c)$.*

Proof. Because g is differentiable at $f(c)$, we know that

$$g'(f(c)) = \lim_{y \rightarrow f(c)} \frac{g(y) - g(f(c))}{y - f(c)}.$$

Another way to assert this same fact is to let $d(y)$ be the difference

$$(1) \quad d(y) = \frac{g(y) - g(f(c))}{y - f(c)} - g'(f(c)),$$

and observe that $\lim_{y \rightarrow f(c)} d(y) = 0$. At the moment, $d(y)$ is not defined when $y = f(c)$, but it should seem natural to declare that $d(f(c)) = 0$, so that d is continuous at $f(c)$.

Now, we come to the finesse. Equation (1) can be rewritten as

$$(2) \quad g(y) - g(f(c)) = [g'(f(c)) + d(y)](y - f(c)).$$

Observe that this equation holds for all $y \in B$ including $y = f(c)$. Thus, we are free to substitute $y = f(t)$ for any arbitrary $t \in A$. If $t \neq c$, we can divide equation (2) by $(t - c)$ to get

$$\frac{g(f(t)) - g(f(c))}{t - c} = [g'(f(c)) + d(f(t))] \frac{(f(t) - f(c))}{t - c}$$

for all $t \neq c$. Finally, taking the limit as $t \rightarrow c$ and applying the Algebraic Limit Theorem yields the desired formula. \square

Darboux's Theorem

One conclusion from this chapter's introduction is that although continuity is necessary for the derivative to exist, it is not the case that the derivative function itself will always be continuous. Our specific example was $g_2(x) = x^2 \sin(1/x)$, where we set $g_2(0) = 0$. By tinkering with the exponent of the leading x^2 factor, it is possible to construct examples of differentiable functions with derivatives that are unbounded, or twice-differentiable functions that have discontinuous second derivatives (Exercise 5.2.5). The underlying principle in all of these examples is that by controlling the size of the oscillations of the original function, we can make the corresponding oscillations of the slopes volatile enough to prevent the existence of the relevant limits.

It is significant that for this class of examples, the discontinuities that arise are never simple jump discontinuities. (A precise definition of "jump discontinuity" is presented in Section 4.6.) We are now ready to confirm our earlier suspicions that, although derivatives do not in general have to be continuous,




Figure 5.3: THE INTERIOR EXTREMUM THEOREM.

they do possess the intermediate value property. (See Definition 4.5.3.) This surprising observation is a fairly straightforward corollary to the more obvious observation that differentiable functions attain maximums and minimums at points where the derivative is equal to zero (Fig. 5.3).

Theorem 5.2.6 (Interior Extremum Theorem). *Let f be differentiable on an open interval (a, b) . If f attains a maximum value at some point $c \in (a, b)$ (i.e., $f(c) \geq f(x)$ for all $x \in (a, b)$), then $f'(c) = 0$. The same is true if $f(c)$ is a minimum value.*

Proof. Because c is in the open interval (a, b) , we can construct two sequences (x_n) and (y_n) , which converge to c and satisfy $x_n < c < y_n$ for all $n \in \mathbf{N}$. The fact that $f(c)$ is a maximum implies that $f(y_n) - f(c) \leq 0$ for all n , and thus

$$f'(c) = \lim_{n \rightarrow \infty} \frac{f(y_n) - f(c)}{y_n - c} \leq 0$$

by the Order Limit Theorem (Theorem 2.3.4). In a similar way,

$$\frac{f(x_n) - f(c)}{x_n - c} \geq 0$$

for each x_n because both numerator and denominator are negative. This implies that

$$f'(c) = \lim_{n \rightarrow \infty} \frac{f(x_n) - f(c)}{x_n - c} \geq 0,$$

and therefore $f'(c) = 0$, as desired. \square

The Interior Extremum Theorem is the fundamental fact behind the use of the derivative as a tool for solving applied optimization problems. This idea, discovered and exploited by Pierre de Fermat, is as old as the derivative itself. In a sense, finding maximums and minimums is arguably why Fermat invented his method of finding slopes of tangent lines. It was 200 years later that the French mathematician Gaston Darboux (1842–1917) pointed out that Fermat's method of finding maximums and minimums carries with it the implication that

if a derivative function attains two distinct values $f'(a)$ and $f'(b)$, then it must also attain every value in between. The noticeably long delay between these discoveries is indicative of the difference between the kinds of mathematical questions that were relevant during these two eras. Whereas Fermat was creating a tool for solving a computational problem, by the middle of the 19th century mathematics had become more introspective. Mathematicians were directing their energies toward understanding their subject for its own sake.

Theorem 5.2.7 (Darboux's Theorem). *If f is differentiable on an interval $[a, b]$, and if α satisfies $f'(a) < \alpha < f'(b)$ (or $f'(a) > \alpha > f'(b)$), then there exists a point $c \in (a, b)$ where $f'(c) = \alpha$.*

Proof. We first simplify matters by defining a new function $g(x) = f(x) - \alpha x$ on $[a, b]$. Notice that g is differentiable on $[a, b]$ with $g'(x) = f'(x) - \alpha$. In terms of g , our hypothesis states that $g'(a) < 0 < g'(b)$, and we hope to show that $g'(c) = 0$ for some $c \in (a, b)$.

The remainder of the argument is outlined in Exercise 5.2.6. □

Exercises

Exercise 5.2.1. Supply proofs for parts (i) and (ii) of Theorem 5.2.4.

Exercise 5.2.2. (a) Use Definition 5.2.1 to produce the proper formula for the derivative of $f(x) = 1/x$.

(b) Combine the result in part (a) with the chain rule (Theorem 5.2.5) to supply a proof for part (iv) of Theorem 5.2.4.

(c) Supply a direct proof of Theorem 5.2.4 (iv) by algebraically manipulating the difference quotient for (f/g) in a style similar to the proof of Theorem 5.2.4 (iii).

Exercise 5.2.3. By imitating the Dirichlet constructions in Section 4.1, construct a function on \mathbf{R} that is differentiable at a single point.

Exercise 5.2.4. Let $f_a(x) = \begin{cases} x^a & \text{if } x \geq 0 \\ 0 & \text{if } x < 0. \end{cases}$

(a) For which values of a is f continuous at zero?

(b) For which values of a is f differentiable at zero? In this case, is the derivative function continuous?

(c) For which values of a is f twice-differentiable?

Exercise 5.2.5. Let

$$g_a(x) = \begin{cases} x^a \sin(1/x) & \text{if } x \neq 0 \\ 0 & \text{if } x = 0. \end{cases}$$

Find a particular (potentially noninteger) value for a so that

(a) g_a is differentiable on \mathbf{R} but such that g'_a is unbounded on $[0, 1]$.

(b) g_a is differentiable on \mathbf{R} with g'_a continuous but not differentiable at zero.

(c) g_a is differentiable on \mathbf{R} and g'_a is differentiable on \mathbf{R} , but such that g''_a is not continuous at zero.

Exercise 5.2.6. (a) Assume that g is differentiable on $[a, b]$ and satisfies $g'(a) < 0 < g'(b)$. Show that there exists a point $x \in (a, b)$ where $g(a) > g(x)$, and a point $y \in (a, b)$ where $g(y) < g(b)$.

(b) Now complete the proof of Darboux's Theorem started earlier.

Exercise 5.2.7. Review the definition of uniform continuity (Definition 4.4.5) and also the content of Theorem 4.4.8, which states that continuous functions on compact sets are uniformly continuous.

(a) Propose a definition for what it should mean to say that $f : A \rightarrow \mathbf{R}$ is *uniformly differentiable* on A .

(b) Give an example of a uniformly differentiable function on $[0, 1]$.

(c) Is there a theorem analogous to Theorem 4.4.8 for differentiation? Are functions that are differentiable on a closed interval $[a, b]$ necessarily uniformly differentiable? The class of examples discussed in Section 5.1 may be useful.

Exercise 5.2.8. Decide whether each conjecture is true or false. Provide an argument for those that are true and a counterexample for each one that is false.

(a) If a derivative function is not constant, then the derivative must take on some irrational values.

(b) If f' exists on an open interval, and there is some point c where $f'(c) > 0$, then there exists a δ -neighborhood $V_\delta(c)$ around c in which $f'(x) > 0$ for all $x \in V_\delta(c)$.

(c) If f is differentiable on an interval containing zero and if $\lim_{x \rightarrow 0} f'(x) = L$, then it must be that $L = f'(0)$.

(d) Repeat conjecture (c) but drop the assumption that $f'(0)$ necessarily exists. If $f'(x)$ exists for all $x \neq 0$ and if $\lim_{x \rightarrow 0} f'(x) = L$, then $f'(0)$ exists and equals L .

5.3 The Mean Value Theorem

The Mean Value Theorem (Fig. 5.4) makes the geometrically plausible assertion that a differentiable function f on an interval $[a, b]$ will, at some point, attain a slope equal to the slope of the line through the endpoints $(a, f(a))$ and $(b, f(b))$. More tersely put,

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

for at least one point $c \in (a, b)$.

On the surface, there does not appear to be anything especially remarkable about this observation. Its validity appears undeniable—much like the Intermediate Value Theorem for continuous functions—and its proof is rather short. The ease of the proof, however, is misleading, as it is built on top of some hard-fought accomplishments from the study of limits and continuity. In this regard, the Mean Value Theorem is a kind of reward for a job well done. As we will see, it is a prize of exceptional value. Although the result itself is geometrically obvious, the Mean Value Theorem is the cornerstone for almost every major theorem pertaining to differentiation. We will use it to prove L'Hospital's rules




Figure 5.4: THE MEAN VALUE THEOREM.

regarding limits of quotients of differentiable functions. A rigorous analysis of how infinite series of functions behave when differentiated requires the Mean Value Theorem (Theorem 6.4.3), and it is the crucial step in the proof of the Fundamental Theorem of Calculus (Theorem 7.5.1). It is also the fundamental concept underlying Lagrange's Remainder Theorem (Theorem 6.6.1) which approximates the error between a Taylor polynomial and the function that generates it.

The Mean Value Theorem can be stated in various degrees of generality, each one important enough to be given its own special designation. Recall that the Extreme Value Theorem (Theorem 4.4.3) states that continuous functions on compact sets always attain maximum and minimum values. Combining this observation with the Interior Extremum Theorem for differentiable functions (Theorem 5.2.6) yields a special case of the Mean Value Theorem first noted by the mathematician Michel Rolle (1652–1719) (Fig. 5.5).

Theorem 5.3.1 (Rolle's Theorem). *Let $f : [a, b] \rightarrow \mathbf{R}$ be continuous on $[a, b]$ and differentiable on (a, b) . If $f(a) = f(b)$, then there exists a point $c \in (a, b)$ where $f'(c) = 0$.*

Proof. Because f is continuous on a compact set, f attains a maximum and a minimum. If both the maximum and minimum occur at the endpoints, then f is necessarily a constant function and $f'(x) = 0$ on all of (a, b) . In this case, we can choose c to be any point we like. On the other hand, if either the maximum or minimum occurs at some point c in the interior (a, b) , then it follows from the Interior Extremum Theorem (Theorem 5.2.6) that $f'(c) = 0$. \square

Theorem 5.3.2 (Mean Value Theorem). *If $f : [a, b] \rightarrow \mathbf{R}$ is continuous on $[a, b]$ and differentiable on (a, b) , then there exists a point $c \in (a, b)$ where*

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





Figure 5.5: ROLLE'S THEOREM.

Proof. Notice that the Mean Value Theorem reduces to Rolle's Theorem in the case where $f(a) = f(b)$. The strategy of the proof is to reduce the more general statement to this special case.

The equation of the line through $(a, f(a))$ and $(b, f(b))$ is

$$y = \left(\frac{f(b) - f(a)}{b - a} \right) (x - a) + f(a).$$



We want to consider the difference between this line and the function $f(x)$. To this end, let

$$d(x) = f(x) - \left[\left(\frac{f(b) - f(a)}{b - a} \right) (x - a) + f(a) \right],$$

and observe that d is continuous on $[a, b]$, differentiable on (a, b) , and satisfies $d(a) = 0 = d(b)$. Thus, by Rolle's Theorem, there exists a point $c \in (a, b)$ where $d'(c) = 0$. Because

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

we get

$$0 = f'(c) - \frac{f(b) - f(a)}{b - a},$$

which completes the proof. \square

The point has been made that the Mean Value Theorem manages to find its way into nearly every proof of any statement related to the geometrical nature of the derivative. As a simple example, if f is a constant function $f(x) = k$ on some interval A , then a straightforward calculation of f' using Definition 5.2.1 shows that $f'(x) = 0$ for all $x \in A$. But how do we prove the converse statement? If we know that a differentiable function g satisfies $g'(x) = 0$ everywhere on A , our intuition suggests that we should be able to prove $g(x)$ is constant. It is the Mean Value Theorem that provides us with a way to articulate rigorously what seems geometrically valid.

Corollary 5.3.3. *If $g : A \rightarrow \mathbf{R}$ is differentiable on an interval A and satisfies $g'(x) = 0$ for all $x \in A$, then $g(x) = k$ for some constant $k \in \mathbf{R}$.*

Proof. Take $x, y \in A$ and assume $x < y$. Applying the Mean Value Theorem to g on the interval $[x, y]$, we see that

$$g'(c) = \frac{g(y) - g(x)}{y - x}$$

for some $c \in A$. Now, $g'(c) = 0$, so we conclude that $g(y) = g(x)$. Set k equal to this common value. Because x and y are arbitrary, it follows that $g(x) = k$ for all $x \in A$. \square

Corollary 5.3.4. *If f and g are differentiable functions on an interval A and satisfy $f'(x) = g'(x)$ for all $x \in A$, then $f(x) = g(x) + k$ for some constant $k \in \mathbf{R}$.*

Proof. Let $h(x) = f(x) - g(x)$ and apply Corollary 5.3.3 to the differentiable function h . \square

The Mean Value Theorem has a more general form due to Cauchy. It is this generalized version of the theorem that is needed to analyze L'Hospital's rules and Lagrange's Remainder Theorem.

Theorem 5.3.5 (Generalized Mean Value Theorem). *If f and g are continuous on the closed interval $[a, b]$ and differentiable on the open interval (a, b) , then there exists a point $c \in (a, b)$ where*

$$[f(b) - f(a)]g'(c) = [g(b) - g(a)]f'(c).$$

If g' is never zero on (a, b) , then the conclusion can be stated as

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

Proof. This result follows by applying the Mean Value Theorem to the function $h(x) = [f(b) - f(a)]g(x) - [g(b) - g(a)]f(x)$. The details are requested in Exercise 5.3.4. \square

L'Hospital's Rules

The Algebraic Limit Theorem asserts that when taking a limit of a quotient of functions we can write

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

provided that each individual limit exists and $\lim_{x \rightarrow c} g(x)$ is not zero. If the denominator does converge to zero and the numerator does not, then it is not difficult to argue that the quotient $f(x)/g(x)$ grows in absolute value without bound as x approaches c (Exercise 5.3.9). L'Hospital's rules are named for the Marquis de L'Hospital (1661–1704), who learned the results from his tutor, Johann Bernoulli (1667–1748), and published them in 1696 in what is regarded as the first calculus text. Stated in different levels of generality, they are a favorite tool for handling the indeterminate cases when either numerator and denominator both tend to zero or both tend simultaneously to infinity.

Theorem 5.3.6 (L'Hospital's Rule: 0/0 case). *Assume f and g are continuous functions defined on an interval containing a , and assume that f and g are differentiable on this interval, with the possible exception of the point a . If $f(a) = 0$ and $g(a) = 0$, then*

$$\lim_{x \rightarrow a} \frac{f'(x)}{g'(x)} = L \quad \text{implies} \quad \lim_{x \rightarrow a} \frac{f(x)}{g(x)} = L.$$

Proof. This argument follows from a straightforward application of the Generalized Mean Value Theorem. It is requested as Exercise 5.3.11. \square

L'Hospital's rule remains true if we replace the assumption that $f(a) = g(a) = 0$ with the hypothesis that $\lim_{x \rightarrow a} g(x) = \infty$. To this point we have not been explicit about what it means to say that a limit equals ∞ . The logical structure of such a definition is precisely the same as it is for finite functional limits. The difference is that rather than trying to force the function to take on values in some small ϵ -neighborhood around a proposed limit, we must show that $g(x)$ eventually *exceeds* any proposed upper bound. The arbitrarily small $\epsilon > 0$ is replaced by an arbitrarily large $M > 0$.

Definition 5.3.7. Given $g : A \rightarrow \mathbf{R}$ and a limit point c of A , we say that $\lim_{x \rightarrow c} g(x) = \infty$ if, for every $M > 0$, there exists a $\delta > 0$ such that whenever $0 < |x - c| < \delta$ it follows that $g(x) \geq M$.

We can define $\lim_{x \rightarrow c} g(x) = -\infty$ in a similar way.

The following version of L'Hospital's rule is referred to as the ∞/∞ case, although the hypothesis only requires that the function in the denominator tend to infinity. If the numerator is bounded, then it is a straightforward exercise to prove that the resulting quotient tends to zero (Exercise 5.3.10). The argument for the general case is relatively involved when compared to the 0/0 case. To simplify the notation of the proof, we state the result using a one-sided limit.

Theorem 5.3.8 (L'Hospital's Rule: ∞/∞ case). Assume f and g are differentiable on (a, b) , and that $\lim_{x \rightarrow a} g(x) = \infty$ (or $-\infty$). Then

$$\lim_{x \rightarrow a} \frac{f'(x)}{g'(x)} = L \quad \text{implies} \quad \lim_{x \rightarrow a} \frac{f(x)}{g(x)} = L.$$

Proof. Because $\lim_{x \rightarrow a} \frac{f'(x)}{g'(x)} = L$, there exists a $\delta_1 > 0$ such that

$$\left| \frac{f'(x)}{g'(x)} - L \right| < \frac{\epsilon}{2}$$

for all $a < x < a + \delta_1$. For convenience of notation, let $t = a + \delta_1$ and mentally note that t is fixed for the remainder of the argument.

Our functions are not defined at a , but for any $a < x < t$ we can apply the Generalized Mean Value Theorem on the interval $[x, t]$ to get

$$\frac{f(x) - f(t)}{g(x) - g(t)} = \frac{f'(c)}{g'(c)}$$

for some $c \in (x, t)$. Our choice of t then implies

$$(1) \quad L - \frac{\epsilon}{2} < \frac{f(x) - f(t)}{g(x) - g(t)} < L + \frac{\epsilon}{2}$$

for all x in (a, t) .

In an effort to isolate the fraction $\frac{f(x)}{g(x)}$, the strategy is to multiply equation (1) by $(g(x) - g(t))/g(x)$. We need to be sure, however, that this quantity is positive, which amounts to insisting that $1 \geq g(t)/g(x)$. Because t is fixed and $\lim_{x \rightarrow a} g(x) = \infty$, we can choose $\delta_2 > 0$ so that $g(x) \geq g(t)$ for all $a < x < a + \delta_2$. Carrying out the desired multiplication results in

$$\left(L - \frac{\epsilon}{2} \right) \left(1 - \frac{g(t)}{g(x)} \right) < \frac{f(x) - f(t)}{g(x)} < \left(L + \frac{\epsilon}{2} \right) \left(1 - \frac{g(t)}{g(x)} \right),$$

which after some algebraic manipulations yields

$$L - \frac{\epsilon}{2} + \frac{-Lg(t) + \frac{\epsilon}{2}g(t) + f(t)}{g(x)} < \frac{f(x)}{g(x)} < L + \frac{\epsilon}{2} + \frac{Lg(t) - \frac{\epsilon}{2}g(t) + f(t)}{g(x)}.$$

Again, let's remind ourselves that t is fixed and that $\lim_{x \rightarrow a} g(x) = \infty$. Thus, we can choose a δ_3 such that $a < x < a + \delta_3$ implies that $g(x)$ is large enough to ensure that both

$$\frac{-Lg(t) + \frac{\epsilon}{2}g(t) + f(t)}{g(x)} \quad \text{and} \quad \frac{Lg(t) - \frac{\epsilon}{2}g(t) + f(t)}{g(x)}$$

are less than $\epsilon/2$. Putting this all together and choosing $\delta = \min\{\delta_1, \delta_2, \delta_3\}$ guarantees that

$$\left| \frac{f(x)}{g(x)} - L \right| < \epsilon$$

for all $a < x < a + \delta$. □

Exercises

Exercise 5.3.1. Recall from Exercise 4.4.9 that a function $f : A \rightarrow \mathbf{R}$ is “Lip-schitz on A ” if there exists an $M > 0$ such that

$$\left| \frac{f(x) - f(y)}{x - y} \right| \leq M$$

for all $x, y \in A$. Show that if f is differentiable on a closed interval $[a, b]$ and if f' is continuous on $[a, b]$, then f is Lipschitz on $[a, b]$.

Exercise 5.3.2. Recall from Exercise 4.3.9 that a function f is contractive on a set A if there exists a constant $0 < s < 1$ such that

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

for all $x, y \in A$. Show that if f is differentiable and f' is continuous and satisfies $|f'(x)| < 1$ on a closed interval, then f is contractive on this set.

Exercise 5.3.3. Let h be a differentiable function defined on the interval $[0, 3]$, and assume that $h(0) = 1$, $h(1) = 2$, and $h(3) = 2$.

- Argue that there exists a point $d \in [0, 3]$ where $h(d) = d$.
- Argue that at some point c we have $h'(c) = 1/3$.
- Argue that $h'(x) = 1/4$ at some point in the domain.

Exercise 5.3.4. (a) Supply the details for the proof of Cauchy’s Generalized Mean Value Theorem (Theorem 5.3.5).

(b) Give a graphical interpretation of the Generalized Mean Value Theorem analogous to the one given for the Mean Value Theorem at the beginning of Section 5.3. (Consider f and g as parametric equations for a curve.)

Exercise 5.3.5. A *fixed point* of a function f is a value x where $f(x) = x$. Show that if f is differentiable on an interval with $f'(x) \neq 1$, then f can have at most one fixed point.

Exercise 5.3.6. Let $g : [0, 1] \rightarrow \mathbf{R}$ be twice-differentiable (i.e., both g and g' are differentiable functions) with $g''(x) > 0$ for all $x \in [0, 1]$. If $g(0) > 0$ and $g(1) = 1$, show that $g(d) = d$ for some point $d \in (0, 1)$ if and only if $g'(1) > 1$. (This geometrically plausible fact is used in the introductory discussion to Chapter 6.)

Exercise 5.3.7. (a) Recall that a function $f : (a, b) \rightarrow \mathbf{R}$ is *increasing* on (a, b) if $f(x) \leq f(y)$ whenever $x < y$ in (a, b) . Assume f is differentiable on (a, b) . Show that f is increasing on (a, b) if and only if $f'(x) \geq 0$ for all $x \in (a, b)$.

(b) Show that the function

$$g(x) = \begin{cases} x/2 + x^2 \sin(1/x) & \text{if } x \neq 0 \\ 0 & \text{if } x = 0 \end{cases}$$

is differentiable on \mathbf{R} and satisfies $g'(0) > 0$. Now, prove that g is *not* increasing over any open interval containing 0.

Exercise 5.3.8. Assume $g : (a, b) \rightarrow \mathbf{R}$ is differentiable at some point $c \in (a, b)$. If $g'(c) \neq 0$, show that there exists a δ -neighborhood $V_\delta(c) \subseteq (a, b)$ for which $g(x) \neq g(c)$ for all $x \in V_\delta(c)$. Compare this result with Exercise 5.3.7.

Exercise 5.3.9. Assume that $\lim_{x \rightarrow c} f(x) = L$, where $L \neq 0$, and assume $\lim_{x \rightarrow c} g(x) = 0$. Show that $\lim_{x \rightarrow c} |f(x)/g(x)| = \infty$.

Exercise 5.3.10. Let f be a bounded function and assume $\lim_{x \rightarrow c} g(x) = \infty$. Show that $\lim_{x \rightarrow c} f(x)/g(x) = 0$.

Exercise 5.3.11. Use the Generalized Mean Value Theorem to furnish a proof of the $0/0$ case of L'Hospital's rule (Theorem 5.3.6).


Exercise 5.3.12. Assume f and g are as described in Theorem 5.3.6, but now add the assumption that f and g are differentiable at a and f' and g' are continuous at a . Find a short proof for the $0/0$ case of L'Hospital's rule under this stronger hypothesis.

Exercise 5.3.13. Review the hypothesis of Theorem 5.3.6. What happens if we do not assume that $f(a) = g(a) = 0$, but assume only that $\lim_{x \rightarrow a} f(x) = 0$ and $\lim_{x \rightarrow a} g(x) = 0$? Assuming we have a proof for Theorem 5.3.6 as it is written, explain how to construct a valid proof under this slightly weaker hypothesis.

5.4 A Continuous Nowhere-Differentiable Function

Exploring the relationship between continuity and differentiability has led to both fruitful results and pathological counterexamples. The bulk of discussion to this point has focused on the continuity of derivatives, but historically a significant amount of debate revolved around the question of whether continuous functions were necessarily differentiable. Early in the chapter, we saw that continuity was a requirement for differentiability, but, as the absolute value function demonstrates, the converse of this proposition is not true. A function can be continuous but not differentiable at some point. But just how nondifferentiable can a continuous function be? Given a *finite* set of points, it is not difficult to imagine how to construct a graph with corners at each of these points, so that the corresponding function fails to be differentiable on this finite set. The trick gets more difficult, however, when the set becomes infinite. For instance, is it possible to construct a function that is continuous on all of \mathbf{R} but fails to be differentiable at every rational point? Not only is this possible, but the situation is even more delightful. In 1872, Karl Weierstrass presented an example of a continuous function that was not differentiable at *any* point. (It seems to be the case that Bernhard Bolzano had his own example of such a beast as early as 1830, but it was not published until much later.) Weierstrass actually discovered a class of nowhere-differentiable functions of the form

$$f(x) = \sum_{n=0}^{\infty} a^n \cos(b^n x)$$

Figure 5.6: THE FUNCTION $h(x)$.

where the values of a and b are carefully chosen. Such functions are specific examples of Fourier series discussed in Section 8.3. The details of Weierstrass' argument are simplified if we replace the cosine function with a piecewise linear function that has oscillations qualitatively like $\cos(x)$.

Define

$$h(x) = |x|$$

on the interval $[-1, 1]$ and extend the definition of h to all of \mathbf{R} by requiring that $h(x + 2) = h(x)$. The result is a periodic "sawtooth" function (Fig. 5.6).

Exercise 5.4.1. Sketch a graph of $(1/2)h(2x)$ on $[-2, 3]$. Give a qualitative description of the functions

$$h_n(x) = \frac{1}{2^n} h(2^n x)$$

as n gets larger.

Now, define

$$g(x) = \sum_{n=0}^{\infty} h_n(x) = \sum_{n=0}^{\infty} \frac{1}{2^n} h(2^n x).$$

The claim is that $g(x)$ is continuous on all of \mathbf{R} but fails to be differentiable at any point.

Infinite Series of Functions and Continuity

The definition of $g(x)$ is a significant departure from the way we usually define functions. For each $x \in \mathbf{R}$, $g(x)$ is defined to be the value of an infinite series.

Exercise 5.4.2. Fix $x \in \mathbf{R}$. Argue that the series

$$\sum_{n=0}^{\infty} \frac{1}{2^n} h(2^n x)$$

converges absolutely and thus $g(x)$ is properly defined.




Figure 5.7: A SKETCH OF $g(x) = \sum_{n=0}^{\infty} (1/2^n)h(2^n x)$.

Exercise 5.4.3. Taking the continuity of $h(x)$ as given, reference the proper theorems from Chapter 4 that imply that the *finite* sum

$$g_m(x) = \sum_{n=0}^m \frac{1}{2^n} h(2^n x)$$

is continuous on \mathbf{R} .

This brings us to an archetypical question in analysis: When do conclusions that are valid in finite settings extend to infinite ones? A finite sum of continuous functions is certainly continuous, but does this necessarily hold for an infinite sum of continuous functions? In general, we will see that this is *not* always the case. For this particular sum, however, the continuity of the limit function $g(x)$ can be proved. Deciphering when results about finite sums of functions extend to infinite sums is one of the fundamental themes of Chapter 6. Although a self-contained argument for the continuity of g is not beyond our means at this point, we will nevertheless postpone the proof (Exercise 6.4.4), leaving it as an enticement for the upcoming study of uniform convergence (or as an exercise for those who have already covered it).

Nondifferentiability

With the proper tools in place, the proof that g is continuous is quite straightforward. The more difficult task is to show that g is not differentiable at any point in \mathbf{R} .

Let's first look at the point $x = 0$. Our function g does not appear to be differentiable here (Fig. 5.7), and a rigorous proof is not too difficult. Consider the sequence $x_m = 1/2^m$, where $m = 0, 1, 2, \dots$

Exercise 5.4.4. Show that

$$\frac{g(x_m) - g(0)}{x_m - 0} = m + 1,$$

and use this to prove that $g'(0)$ does not exist.

Any temptation to say something like $g'(0) = \infty$ should be resisted. Setting $x_m = -(1/2^m)$ in the previous argument produces difference quotients heading toward $-\infty$. The geometric manifestation of this is the “cusp” that appears at $x = 0$ in the graph of g .

Exercise 5.4.5. (a) Modify the previous argument to show that $g'(1)$ does not exist. Show that $g'(1/2)$ does not exist.

(b) Show that $g'(x)$ does not exist for any rational number of the form $x = p/2^k$ where $p \in \mathbf{Z}$ and $k \in \mathbf{N} \cup \{0\}$.

The points described in Exercise 5.4.5 (b) are called “dyadic” points. If $x = p/2^k$ is a dyadic rational number, then the function h_n has a corner at x as long as $n \geq k$. Thus, it should not be too surprising that g fails to be differentiable at points of this form. The argument is more intricate at points between the dyadic points.

Assume x is *not* a dyadic number. For a fixed value of $m \in \mathbf{N} \cup \{0\}$, x falls between two adjacent dyadic points,

$$\frac{p}{2^m} < x < \frac{p+1}{2^m}.$$

Set $x_m = p/2^m$ and $y = (p+1)/2^m$. Repeating this for each m yields two sequences (x_m) and (y_m) satisfying

$$\lim x_m = \lim y_m = x \quad \text{and} \quad x_m < x < y_m.$$

Exercise 5.4.6. (a) Without working too hard, explain why the partial sum $g_m = h_0 + h_1 + \cdots + h_m$ is differentiable at x . Now, prove that, for every value of m , we have

$$|g'_{m+1}(x) - g'_m(x)| = 1.$$

(b) Prove the two inequalities

$$\frac{g(y_m) - g(x)}{y_m - x} < g'_m(x) < \frac{g(x_m) - g(x)}{x_m - x}.$$

(c) Use parts (a) and (b) to show that $g'(x)$ does not exist.

Weierstrass' original 1872 paper contained a demonstration that the infinite sum

$$f(x) = \sum_{n=0}^{\infty} a^n \cos(b^n x)$$

defined a continuous nowhere-differentiable function provided $0 < a < 1$ and b was an odd integer satisfying $ab > 1 + 3\pi/2$. The condition on a is easy to understand. If $0 < a < 1$, then $\sum_{n=0}^{\infty} a^n$ is a convergent geometric series, and the forthcoming Weierstrass M-Test (Theorem 6.4.5) can be used to conclude that f is continuous. The restriction on b is more mysterious. In 1916, G.H. Hardy extended Weierstrass' result to include any value of b for which $ab \geq 1$. Without looking at the details of either of these arguments, we nevertheless get a sense that the lack of a derivative is intricately tied to the relationship between the compression factor (the parameter a) and the rate at which the frequency of the oscillations increases (the parameter b).

Exercise 5.4.7. Review the argument for the nondifferentiability of $g(x)$ at nondyadic points. Does the argument still work if we replace $g(x)$ with the summation $\sum_{n=0}^{\infty} (1/2^n)h(3^n x)$? Does the argument work for the function $\sum_{n=0}^{\infty} (1/3^n)h(2^n x)$?

5.5 Epilogue

Far from being an anomaly to be relegated to the margins of our understanding of continuous functions, Weierstrass' example and those like it should actually serve as a guide to our intuition. The image of continuity as a smooth curve in our mind's eye severely misrepresents the situation and is the result of a bias stemming from an overexposure to the much smaller class of differentiable functions. The lesson here is that continuity is a strictly weaker notion than differentiability. In Section 3.6, we alluded to a corollary of the Baire Category Theorem, which asserts that Weierstrass' construction is actually typical of continuous functions. We will see that most continuous functions are nowhere-differentiable, so that it is really the differentiable functions that are the exceptions rather than the rule. The details of how to phrase this observation more rigorously are spelled out in Section 8.2.

To say that the nowhere-differentiable function g constructed in the previous section has “corners” at every point of its domain slightly misses the mark. Weierstrass' original class of nowhere-differentiable functions was constructed from infinite sums of *smooth* trigonometric functions. It is the densely nested oscillating structure that makes the definition of a tangent line impossible. So what happens when we restrict our attention to monotone functions? How nondifferentiable can an increasing function be? Given a finite set of points, it is not difficult to piece together a monotone function which has actual corners—and thus is not differentiable—at each point in the given set. A natural question is whether there exists a continuous, monotone function that is nowhere-differentiable. Weierstrass suspected that such a function existed but only managed to produce an example of a continuous, increasing function which failed to be differentiable on a countable dense set (Exercise 7.5.11). In 1903, the French mathematician Henri Lebesgue (1875–1941) demonstrated that Weierstrass' intuition had failed on this account. Lebesgue proved that a continuous, monotone function would have to be differentiable at “almost” every point in

its domain. To be specific, Lebesgue showed that, for every $\epsilon > 0$, the set of points where such a function fails to be differentiable can be covered by a countable union of intervals whose lengths sum to a quantity less than ϵ . This notion of “zero length,” or “measure zero” as it is called, was encountered in our discussion of the Cantor set and is explored more fully in Section 7.6, where Lebesgue’s substantial contribution to the theory of integration is discussed.

With the relationship between the continuity of f and the existence of f' somewhat in hand, we once more return to the question of characterizing the set of all derivatives. Not every function is a derivative. Darboux’s Theorem forces us to conclude that there are some functions—those with jump discontinuities in particular—that cannot appear as the derivative of some other function. Another way to phrase Darboux’s Theorem is to say that all derivatives must satisfy the intermediate value property. Continuous functions do possess the intermediate value property, and it is natural to ask whether every continuous function is necessarily a derivative. For this smaller class of functions, the answer is yes. The Fundamental Theorem of Calculus, treated in Chapter 7, states that, given a continuous function f , the function $F(x) = \int_a^x f$ satisfies $F' = f$. This does the trick. The collection of derivatives at least contains the continuous functions. The search for a concise characterization of *all* possible derivatives, however, remains largely unsuccessful.

As a final remark, we will see that by cleverly choosing f , this technique of defining F via $F(x) = \int_a^x f$ can be used to produce examples of continuous functions which fail to be differentiable on interesting sets, *provided we can show that $\int_a^x f$ is defined*. The question of just how to define integration became a central theme in analysis in the latter half of the 19th century and has continued on to the present. Much of this story is discussed in detail in Chapter 7 and Section 8.1.

Chapter 6

Sequences and Series of Functions

6.1 Discussion: Branching Processes

The fact that polynomial functions are so ubiquitous in both pure and applied analysis can be attributed to any number of reasons. They are continuous, infinitely differentiable, and defined on all of \mathbf{R} . They are easy to evaluate and easy to manipulate, both from the points of view of algebra (adding, multiplying, factoring) and calculus (integrating, differentiating). It should be no surprise, then, that even in the earliest stages of the development of calculus, mathematicians experimented with the idea of extending the notion of polynomials to functions that are essentially polynomials of infinite degree. Such objects are called *power series*, and are formally denoted by

$$\sum_{n=0}^{\infty} a_n x^n = a_0 + a_1 x + a_2 x^2 + a_3 x^3 + \cdots .$$

The basic dilemma from the point of view of analysis is deciphering when the desirable qualities of the limiting functions (the polynomials in this case) are passed on to the limit (the power series). To put the discussion in a more concrete context, let's look at a particular problem from the theory of probability.

In 1873, Francis Galton asked the London Mathematical Society to consider the problem of the survival of surnames (which at that time were passed to succeeding generations exclusively by adult male children). "Assume," Galton said, "that the law of population is such that, in each generation, p_0 percent of the adult males have no male children who reach adult life; p_1 have one such male child; p_2 percent have two; and so on... Find [the probability that] the surname will become extinct after r generations." We should add (or make explicit) the assumption that the lives of each offspring, and the descendants thereof, proceed independently of the fortunes of the rest of the family.

Galton asks for the probability of extinction after r generations, which we will call d_r . If we begin with one parent, then $d_1 = p_0$. If $p_0 = 0$, then d_r will clearly equal 0 for all generations r . To keep the problem interesting, we will insist that from here on $p_0 > 0$. Now, d_2 , whatever it equals, will certainly satisfy $d_1 \leq d_2$ because if the population is extinct after one generation it will remain so after two. By this reasoning, we have a monotone sequence

$$d_1 \leq d_2 \leq d_3 \leq d_4 \cdots,$$

which, because we are dealing with probabilities, is bounded above by 1. By the Monotone Convergence Theorem, the sequence converges, and we can let

$$d = \lim_{r \rightarrow \infty} d_r$$

be the probability that the surname eventually goes extinct at any time in the future. Knowing it exists, our task is to find d .

The truly clever step in the solution is to define the function

$$G(x) = p_0 + p_1x + p_2x^2 + p_3x^3 + \cdots.$$

In the case of producing male offspring, it seems safe to assume that this sum terminates after five or six terms, because nature would have it that $p_n = 0$ for all values of n beyond this point. However, if we were studying neutrons in a nuclear reactor, or heterozygotes carrying a mutant gene (as is often the case with the theory of branching processes), then the notion of an infinite sum becomes a more attractive model. The point is this: We will proceed with reckless abandon and treat the function $G(x)$ as though it were a familiar polynomial of finite degree. At the end of the computations, however, we will have to again become well-trained analysts and be prepared to justify the manipulations we have made under the hypothesis that $G(x)$ represents an infinite sum for each value of x .

The critical observation is that

$$G(d_r) = d_{r+1}.$$

The way to understand this is to view the expression

$$G(d_r) = p_0 + p_1d_r + p_2d_r^2 + p_3d_r^3 + \cdots$$


as a sum of the probabilities for different distinct ways extinction could occur in $r + 1$ generations based on what happens after the first generational step. Specifically, p_0 is the probability that the initial parent has no offspring and so still has none after $r + 1$ generations. The term p_1d_r is the probability that the initial parent has one male child times the probability that this child's own lineage dies out after r generations. Thus, the probability p_1d_r is another contribution toward the probability of extinction in $r + 1$ steps. The third term represents the probability that the initial parent has two children and that the surnames of each of these two children die out within r generations. Continuing in this way, we see that every possible scenario for extinction in $r + 1$ steps is

accounted for exactly once within the sum $G(d_r)$. By the definition of d_{r+1} , we get $G(d_r) = d_{r+1}$.

Now for some analysis. If we take the limit as $r \rightarrow \infty$ on each side of the equation $G(d_r) = d_{r+1}$, then on the right-hand side we get $\lim d_{r+1} = d$. Assuming G is continuous, we have

$$d = \lim_{r \rightarrow \infty} d_{r+1} = \lim_{r \rightarrow \infty} G(d_r) = G(d).$$

The conclusion that $d = G(d)$ means that the point d is a *fixed point* of G . It can be located graphically by finding where the graph of G intersects the line $y = x$.



(i) $G'(1) \leq 1$

(ii) $G'(1) > 1$

It is always the case that

$$G(1) = p_0 + p_1 + p_2 + p_3 + \cdots = 1$$

because the probabilities (p_k) form a complete distribution. But $d = 1$ is not necessarily the only candidate for a solution to $G(d) = d$. Graph (ii) illustrates a scenario in which G has another fixed point in the interval $(0, 1)$ in addition to $x = 1$.

Treating G as though it were a polynomial, we differentiate term-by-term to get

$$G'(x) = p_1 + 2p_2x + 3p_3x^2 + 4p_4x^3 + \cdots$$

and

$$G''(x) = 2p_2 + 6p_3x + 12p_4x^2 + \cdots$$

On the interval $[0, 1]$, every term in G' and G'' is nonnegative which means G is an increasing, convex function from $G(0) = p_0 > 0$ up to $G(1) = 1$. This suggests that the two preceding graphs form a rather complete picture of the possibilities for the behavior of G with regard to fixed points. Of particular interest is graph (ii), where the graph of $y = x$ intersects G twice in $[0, 1]$. Using the Mean Value Theorem, we can prove (Exercise 5.3.6) that $G(d) = d$ for some other point $d \in (0, 1)$ if and only if $G'(1) > 1$.

Now,

$$G'(1) = p_1 + 2p_2 + 3p_3 + 4p_4 + \cdots$$

has a very interesting interpretation within the language of probability. The sum is a weighted average, where in each term we have multiplied the number of male children by the probability of actually producing this particular number. The result is a value for the *expected number of male offspring* from a given parent. Said another way, $G'(1)$ is the average number of male children produced by the parents in this particular family tree.

It is not difficult to argue that (d_r) will converge to the smallest solution to $G(d) = d$ on $[0, 1]$ (Exercise 6.5.12), and so we arrive at the following conclusion. If each parent produces, on average, more than one male child, then there is a positive probability that the surname will survive. The equation $G(d) = d$ will have a unique solution in $(0, 1)$, and $1 - d$ represents the probability that the surname does not become extinct. On the other hand, if the expected number of male offspring per parent is one or less than one, then extinction occurs with probability one.

The implications of these results on nuclear reactions and the spread of cancer are fascinating topics for another time. What is of concern to us here is whether our manipulations of $G(x)$ are justified. The assumption that $\sum_{n=0}^{\infty} p_n = 1$ guarantees that G is at least defined at $x = 1$. The point $x = 0$ poses no problem, but is G necessarily well-defined for $0 < x < 1$? If so, how might we prove that G is continuous on this set? Differentiable? Twice-differentiable? If G is differentiable, can we compute the derivative by naively differentiating each term of the series? Our initial attack on these questions will require us to focus attention on the interval $[0, 1)$. Some interesting subtleties arise when we try to extend our results to include the endpoint $x = 1$.

6.2 Uniform Convergence of a Sequence of Functions

Just as in chapter two, we will initially concern ourselves with the behavior and properties of converging *sequences* of functions. Because convergence of infinite sums is defined in terms of the associated sequence of partial sums, the results from our study of sequences will be immediately applicable to the questions we have raised about power series and about infinite series of functions in general.

Pointwise Convergence

Definition 6.2.1. For each $n \in \mathbf{N}$, let f_n be a function defined on a set $A \subseteq \mathbf{R}$. The sequence (f_n) of functions *converges pointwise on A* to a function $f : A \rightarrow \mathbf{R}$ if, for all $x \in A$, the sequence of real numbers $f_n(x)$ converges to $f(x)$.

In this case, we write $f_n \rightarrow f$, $\lim f_n = f$, or $\lim_{n \rightarrow \infty} f_n(x) = f(x)$. This last expression is helpful if there is any confusion as to whether x or n is the limiting variable.

Example 6.2.2. (i) Consider

$$f_n(x) = (x^2 + nx)/n$$




Figure 6.1: f_1, f_5, f_{10} , AND f_{20} WHERE $f_n = (x^2 + nx)/n$.

on all of \mathbf{R} . Graphs of f_1, f_5, f_{10} , and f_{20} (Fig. 6.1) give an indication of what is happening as n gets larger. Algebraically, we can compute

$$\lim_{n \rightarrow \infty} f_n(x) = \lim_{n \rightarrow \infty} \frac{x^2 + nx}{n} = \lim_{n \rightarrow \infty} \frac{x^2}{n} + x = x.$$

Thus, (f_n) converges pointwise to $f(x) = x$ on \mathbf{R} .

(ii) Let $g_n(x) = x^n$ on the set $[0, 1]$, and consider what happens as n tends to infinity (Fig 6.2). If $0 \leq x < 1$, then we have seen that $x^n \rightarrow 0$. On the other hand, if $x = 1$, then $x^n \rightarrow 1$. It follows that $g_n \rightarrow g$ pointwise on $[0, 1]$, where

$$g(x) = \begin{cases} 0 & \text{for } 0 \leq x < 1 \\ 1 & \text{for } x = 1. \end{cases}$$




Figure 6.2: $g(x) = \lim_{n \rightarrow \infty} x^n$ IS NOT CONTINUOUS ON $[0, 1]$.




Figure 6.3: $h_n \rightarrow |x|$ ON $[-1, 1]$; LIMIT IS NOT DIFFERENTIABLE.

(iii) Consider $h_n(x) = x^{1+\frac{1}{2^n-1}}$ on the set $[-1, 1]$ (Fig. 6.3). For a fixed $x \in [-1, 1]$ we have

$$\lim_{n \rightarrow \infty} h_n(x) = x \lim_{n \rightarrow \infty} x^{\frac{1}{2^n-1}} = |x|.$$

Examples 6.2.2 (ii) and (iii) are our first indication that there is some difficult work ahead of us. The central theme of this chapter is analyzing which properties the limit function inherits from the approximating sequence. In Example 6.2.2 (iii) we have a sequence of differentiable functions converging pointwise to a limit that is not differentiable at the origin. In Example 6.2.2 (ii), we see an even more fundamental problem of a sequence of continuous functions converging to a limit that is not continuous.

Continuity of the Limit Function

With Example 6.2.2 (ii) firmly in mind, we begin this discussion with a doomed attempt to prove that the pointwise limit of continuous functions is continuous. Upon discovering the problem in the argument, we will be in a better position to understand the need for a stronger notion of convergence for sequences of functions.

Assume (f_n) is a sequence of continuous functions on a set $A \subseteq \mathbf{R}$, and assume (f_n) converges pointwise to a limit f . To argue that f is continuous, fix a point $c \in A$, and let $\epsilon > 0$. We need to find a $\delta > 0$ such that

$$|x - c| < \delta \quad \text{implies} \quad |f(x) - f(c)| < \delta.$$

By the triangle inequality,

$$\begin{aligned} |f(x) - f(c)| &= |f(x) - f_n(x) + f_n(x) - f_n(c) + f_n(c) - f(c)| \\ &\leq |f(x) - f_n(x)| + |f_n(x) - f_n(c)| + |f_n(c) - f(c)|. \end{aligned}$$

(We should really call this the “quadrilateral inequality” because we are using three joined “sides” as an overestimate for the length of the fourth.) Our first, optimistic impression is that each term in the sum on the right-hand side can be made small—the first and third by the fact that $f_n \rightarrow f$, and the middle term by the continuity of f_n . In order to use the continuity of f_n , we must first establish which particular f_n we are talking about. Because $c \in A$ is fixed, choose $N \in \mathbf{N}$ so that

$$|f_N(c) - f(c)| < \frac{\epsilon}{3}.$$

Now that N is chosen, the continuity of f_N implies that there exists a $\delta > 0$ such that

$$|f_N(x) - f_N(c)| < \frac{\epsilon}{3}$$

for all x satisfying $|x - c| < \delta$.

But here is the problem. We also need

$$|f_N(x) - f(x)| < \frac{\epsilon}{3} \quad \text{for all } |x - c| < \delta.$$

The values of x depend on δ , which depends on the choice of N . Thus, we cannot go back and simply choose a different N . More to the point, the variable x is not fixed the way c is in this discussion but represents any point in the interval $(c - \delta, c + \delta)$. Pointwise convergence implies that we can make $|f_n(x) - f(x)| < \epsilon/3$ for large enough values of n , but *the value of n depends on the point x* . It is possible that different values for x will result in the need for different—larger—choices for n . This phenomenon is apparent in Example 6.2.2 (ii). To achieve the inequality

$$|g_n(1/2) - g(1/2)| < \frac{1}{3},$$

we need $n \geq 2$, whereas

$$|g_n(9/10) - g(9/10)| < \frac{1}{3}$$

is true only after $n \geq 11$.

Uniform Convergence

To resolve this dilemma, we define a new, stronger notion of convergence of functions.

Definition 6.2.3. Let f_n be a sequence of functions defined on a set $A \subseteq \mathbf{R}$. Then, (f_n) converges uniformly on A to a limit function f defined on A if, for every $\epsilon > 0$, there exists an $N \in \mathbf{N}$ such that $|f_n(x) - f(x)| < \epsilon$ whenever $n \geq N$ and $x \in A$.

To emphasize the difference between uniform convergence and pointwise convergence, we restate Definition 6.2.1, being more explicit about the relationship

between ϵ , N , and x . In particular, notice where the domain point x is referenced in each definition and consequently how the choice of N then does or does not depend on this value.

Definition 6.2.1B. Let f_n be a sequence of functions defined on a set $A \subseteq \mathbf{R}$. Then, (f_n) converges pointwise on A to a limit f defined on A if, for every $\epsilon > 0$ and $x \in A$, there exists an $N \in \mathbf{N}$ (perhaps dependent on x) such that $|f_n(x) - f(x)| < \epsilon$ whenever $n \geq N$.

The use of the adverb *uniformly* here should be reminiscent of its use in the phrase “uniformly continuous” from Chapter 4. In both cases, the term “uniformly” is employed to express the fact that the response (δ or N) to a prescribed ϵ can be chosen to work simultaneously for all values of x in the relevant domain.

Example 6.2.4. (i) Let

$$g_n(x) = \frac{1}{n(1+x^2)}.$$

For any fixed $x \in \mathbf{R}$, we can see that $\lim g_n(x) = 0$ so that $g(x) = 0$ is the pointwise limit of the sequence (g_n) on \mathbf{R} . Is this convergence uniform? The observation that $1/(1+x^2) \leq 1$ for all $x \in \mathbf{R}$ implies that

$$|g_n(x) - g(x)| = \left| \frac{1}{n(1+x^2)} - 0 \right| \leq \frac{1}{n}.$$

Thus, given $\epsilon > 0$, we can choose $N > 1/\epsilon$ (which does not depend on x), and it follows that

$$n \geq N \quad \text{implies} \quad |g_n(x) - g(x)| < \epsilon$$

for all $x \in \mathbf{R}$. By Definition 6.2.3, $g_n \rightarrow 0$ uniformly on \mathbf{R} .

(ii) Look back at Example 6.2.2 (i), where we saw that $f_n(x) = (x^2 + nx)/n$ converges pointwise on \mathbf{R} to $f(x) = x$. On \mathbf{R} , the convergence is not uniform. To see this write

$$|f_n(x) - f(x)| = \left| \frac{x^2 + nx}{n} - x \right| = \frac{x^2}{n},$$



and notice that in order to force $|f_n(x) - f(x)| < \epsilon$, we are going to have to choose

$$N > \frac{x^2}{\epsilon}.$$

Although this is possible to do for each $x \in \mathbf{R}$, there is no way to choose a single value of N that will work for all values of x at the same time.

On the other hand, we can show that $f_n \rightarrow f$ uniformly on the set $[-b, b]$. By restricting our attention to a bounded interval, we may now assert that

$$\frac{x^2}{n} \leq \frac{b^2}{n}.$$

Figure 6.4: $f_n \rightarrow f$ UNIFORMLY ON A .Figure 6.5: $g_n \rightarrow g$ POINTWISE, BUT NOT UNIFORMLY.

Given $\epsilon > 0$, then, we can choose

$$N > \frac{b^2}{\epsilon}$$

independently of $x \in [-b, b]$.

Graphically speaking, the uniform convergence of f_n to a limit f on a set A can be visualized by constructing a band of radius $\pm\epsilon$ around the limit function f . If $f_n \rightarrow f$ uniformly, then there exists a point in the sequence after which each f_n is *completely* contained in this ϵ -strip (Fig. 6.4). This image should be compared with the graphs in Figures 6.1–6.2 from Example 6.2.2 and the one in Figure 6.5.

Cauchy Criterion

Recall that the Cauchy Criterion for convergent sequences of real numbers was an equivalent characterization of convergence which, unlike the definition, did not make explicit mention of the limit. The usefulness of the Cauchy Criterion suggests the need for an analogous characterization of uniformly convergent sequences of functions. As with all statements about uniformity, pay attention to where the quantifying phrase “for all $x \in A$ ” appears in the statement.

Theorem 6.2.5 (Cauchy Criterion for Uniform Convergence). *A sequence of functions (f_n) defined on a set $A \subseteq \mathbf{R}$ converges uniformly on A if and only if for every $\epsilon > 0$ there exists an $N \in \mathbf{N}$ such that $|f_n(x) - f_m(x)| < \epsilon$ for all $m, n \geq N$ and all $x \in A$.*

Proof. Exercise 6.2.6. □

Continuity Revisited

The stronger assumption of uniform convergence is precisely what is required to remove the flaws from our attempted proof that the limit of continuous functions is continuous.

Theorem 6.2.6. *Let (f_n) be a sequence of functions defined on $A \subseteq \mathbf{R}$ that converges uniformly on A to a function f . If each f_n is continuous at $c \in A$, then f is continuous at c .*

Proof. Fix $c \in A$ and let $\epsilon > 0$. Choose N so that

$$|f_N(x) - f(x)| < \frac{\epsilon}{3}$$

for all $x \in A$. Because f_N is continuous, there exists a $\delta > 0$ for which

$$|f_N(x) - f_N(c)| < \frac{\epsilon}{3}$$

is true whenever $|x - c| < \delta$. But this implies

$$\begin{aligned} |f(x) - f(c)| &= |f(x) - f_N(x) + f_N(x) - f_N(c) + f_N(c) - f(c)| \\ &\leq |f(x) - f_N(x)| + |f_N(x) - f_N(c)| + |f_N(c) - f(c)| \\ &< \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} = \epsilon. \end{aligned}$$

Thus, f is continuous at $c \in A$. □

Exercises

Exercise 6.2.1. Let

$$f_n(x) = \frac{nx}{1 + nx^2}.$$

- Find the pointwise limit of (f_n) for all $x \in (0, \infty)$.
- Is the convergence uniform on $(0, \infty)$?
- Is the convergence uniform on $(0, 1)$?
- Is the convergence uniform on $(1, \infty)$?

Exercise 6.2.2. Let

$$g_n(x) = \frac{nx + \sin(nx)}{2n}.$$

Find the pointwise limit of (g_n) on \mathbf{R} . Is the convergence uniform on $[-10, 10]$? Is the convergence uniform on all of \mathbf{R} ?

Exercise 6.2.3. Consider the sequence of functions

$$h_n(x) = \frac{x}{1+x^n}$$

over the domain $[0, \infty)$.

- Find the pointwise limit of (h_n) on $[0, \infty)$.
- Explain how we know that the convergence *cannot* be uniform on $[0, \infty)$.
- Choose a smaller set over which the convergence is uniform and supply an argument to show that this is indeed the case.

Exercise 6.2.4. For each $n \in \mathbf{N}$, find the points on \mathbf{R} where the function $f_n(x) = x/(1+nx^2)$ attains its maximum and minimum values. Use this to prove (f_n) converges uniformly on \mathbf{R} . What is the limit function?

Exercise 6.2.5. For each $n \in \mathbf{N}$, define f_n on \mathbf{R} by

$$f_n(x) = \begin{cases} 1 & \text{if } |x| \geq 1/n \\ n|x| & \text{if } |x| < 1/n. \end{cases}$$

- Find the pointwise limit of (f_n) on \mathbf{R} and decide whether or not the convergence is uniform.
- Construct an example of a pointwise limit of continuous functions that converges everywhere on the compact set $[-5, 5]$ to a limit function that is unbounded on this set.

Exercise 6.2.6. Using the Cauchy Criterion for convergent sequences of real numbers (Theorem 2.6.4), supply a proof for Theorem 6.2.5. (First, define a candidate for $f(x)$, and then argue that $f_n \rightarrow f$ uniformly.)

Exercise 6.2.7. Assume that (f_n) converges uniformly to f on A and that each f_n is uniformly continuous on A . Prove that f is uniformly continuous on A .

Exercise 6.2.8. Decide which of the following conjectures are true and which are false. Supply a proof for those that are valid and a counterexample for each one that is not.

- If $f_n \rightarrow f$ pointwise on a compact set K , then $f_n \rightarrow f$ uniformly on K .
- If $f_n \rightarrow f$ uniformly on A and g is a bounded function on A , then $f_n g \rightarrow f g$ uniformly on A .
- If $f_n \rightarrow f$ uniformly on A , and if each f_n is bounded on A , then f must also be bounded.
- If $f_n \rightarrow f$ uniformly on a set A , and if $f_n \rightarrow f$ uniformly on a set B , then $f_n \rightarrow f$ uniformly on $A \cup B$.
- If $f_n \rightarrow f$ uniformly on an interval, and if each f_n is increasing, then f is also increasing.
- Repeat conjecture (e) assuming only pointwise convergence.

Exercise 6.2.9. Assume (f_n) converges uniformly to f on a compact set K , and let g be a continuous function on K satisfying $g(x) \neq 0$. Show (f_n/g) converges uniformly on K to f/g .

Exercise 6.2.10. Let f be uniformly continuous on all of \mathbf{R} , and define a sequence of functions by $f_n(x) = f(x + \frac{1}{n})$. Show that $f_n \rightarrow f$ uniformly. Give an example to show that this proposition fails if f is only assumed to be continuous and not uniformly continuous on \mathbf{R} .

Exercise 6.2.11. Assume (f_n) and (g_n) are uniformly convergent sequences of functions.

- (a) Show that $(f_n + g_n)$ is a uniformly convergent sequence of functions.
- (b) Give an example to show that the product $(f_n g_n)$ may not converge uniformly.
- (c) Prove that if there exists an $M > 0$ such that $|f_n| \leq M$ and $|g_n| \leq M$ for all $n \in \mathbf{N}$, then $(f_n g_n)$ does converge uniformly.

Exercise 6.2.12. Theorem 6.2.6 has a partial converse. Assume $f_n \rightarrow f$ pointwise on a compact set K and assume that for each $x \in K$ the sequence $f_n(x)$ is increasing. Follow these steps to show that if f_n and f are continuous on K , then the convergence is uniform.

- (a) Set $g_n = f - f_n$ and translate the preceding hypothesis into statements about the sequence (g_n) .
- (b) Let $\epsilon > 0$ be arbitrary, and define $K_n = \{x \in K : g_n(x) \geq \epsilon\}$. Argue that $K_1 \supseteq K_2 \supseteq K_3 \supseteq \dots$ is a nested sequence of compact sets, and use this observation to finish the argument.

Exercise 6.2.13 (Cantor Function). Review the construction of the Cantor set $C \subseteq [0, 1]$ from Section 3.1. This exercise makes use of results and notation from this discussion.

- (a) Define $f_0(x) = x$ for all $x \in [0, 1]$. Now, let

$$f_1(x) = \begin{cases} (3/2)x & \text{for } 0 \leq x \leq 1/3 \\ 1/2 & \text{for } 1/3 < x < 2/3 \\ (3/2)x - 1/2 & \text{for } 2/3 \leq x \leq 1. \end{cases}$$

Sketch f_0 and f_1 over $[0, 1]$ and observe that f_1 is continuous, increasing, and constant on the middle third $(1/3, 2/3) = [0, 1] \setminus C_1$.

- (b) Construct f_2 by imitating this process of flattening out the middle third of each nonconstant segment of f_1 . Specifically, let

$$f_2(x) = \begin{cases} (1/2)f_1(3x) & \text{for } 0 \leq x \leq 1/3 \\ f_1(x) & \text{for } 1/3 < x < 2/3 \\ (1/2)f_1(3x - 2) + 1/2 & \text{for } 2/3 \leq x \leq 1. \end{cases}$$

If we continue this process, show that the resulting sequence (f_n) converges uniformly on $[0, 1]$.

- (c) Let $f = \lim f_n$. Prove that f is a continuous, increasing function on $[0, 1]$ with $f(0) = 0$ and $f(1) = 1$ that satisfies $f'(x) = 0$ for all x in the open set $[0, 1] \setminus C$. Recall that the “length” of the Cantor set C is 0. Somehow, f manages to increase from 0 to 1 while remaining constant on a set of “length 1.”

Exercise 6.2.14. Recall that the Bolzano–Weierstrass Theorem (Theorem 2.5.5) states that every bounded sequence of real numbers has a convergent subsequence. An analogous statement for bounded sequences of functions is not true in general, but under stronger hypotheses several different conclusions are possible. One avenue is to assume the common domain for all of the functions in the sequence is countable. (Another is explored in the next two exercises.)

Let $A = \{x_1, x_2, x_3, \dots\}$ be a countable set. For each $n \in \mathbf{N}$, let f_n be defined on A and assume there exists an $M > 0$ such that $|f_n(x)| \leq M$ for all $n \in \mathbf{N}$ and $x \in A$. Follow these steps to show that there exists a subsequence of (f_n) that converges pointwise on A .

(a) Why does the sequence of real numbers $f_n(x_1)$ necessarily contain a convergent subsequence (f_{n_k}) ? To indicate that the subsequence of functions (f_{n_k}) is generated by considering the values of the functions at x_1 , we will use the notation $f_{n_k} = f_{1,k}$.

(b) Now, explain why the sequence $f_{1,k}(x_2)$ contains a bounded subsequence.

(c) Carefully construct a nested family of subsequences $(f_{m,k})$, and use Cantor’s diagonalization technique (from Theorem 1.5.1) to produce a single subsequence of (f_n) that converges at every point of A .

Exercise 6.2.15. A sequence of functions (f_n) defined on a set $E \subseteq \mathbf{R}$ is called *equicontinuous* if for every $\epsilon > 0$ there exists a $\delta > 0$ such that $|f_n(x) - f_n(y)| < \epsilon$ for all $n \in \mathbf{N}$ and $|x - y| < \delta$ in E .

(a) What is the difference between saying that a sequence of functions (f_n) is equicontinuous and just asserting that each f_n in the sequence is individually uniformly continuous?

(b) Give a qualitative explanation for why the sequence $g_n(x) = x^n$ is not equicontinuous on $[0, 1]$. Is each g_n uniformly continuous on $[0, 1]$?

Exercise 6.2.16 (Arzela–Ascoli Theorem). For each $n \in \mathbf{N}$, let f_n be a function defined on $[0, 1]$. If (f_n) is bounded on $[0, 1]$ —that is, there exists an $M > 0$ such that $|f_n(x)| \leq M$ for all $n \in \mathbf{N}$ and $x \in [0, 1]$ —and if the collection of functions (f_n) is equicontinuous (Exercise 6.2.15), follow these steps to show that (f_n) contains a uniformly convergent subsequence.

(a) Use Exercise 6.2.14 to produce a subsequence (f_{n_k}) that converges at every rational point in $[0, 1]$. To simplify the notation, set $g_k = f_{n_k}$. It remains to show that (g_k) converges uniformly on all of $[0, 1]$.

(b) Let $\epsilon > 0$. By equicontinuity, there exists a $\delta > 0$ such that

$$|g_k(x) - g_k(y)| < \frac{\epsilon}{3}$$

for all $|x - y| < \delta$ and $k \in \mathbf{N}$. Using this δ , let r_1, r_2, \dots, r_m be a *finite* collection of rational points with the property that the union of the neighborhoods $V_\delta(r_i)$ contains $[0, 1]$.

Explain why there must exist an $N \in \mathbf{N}$ such that

$$|g_s(r_i) - g_t(r_i)| < \frac{\epsilon}{3}$$

for all $s, t \geq N$ and r_i in the finite subset of $[0, 1]$ just described. Why does having the set $\{r_1, r_2, \dots, r_m\}$ be finite matter?

(c) Finish the argument by showing that, for an arbitrary $x \in [0, 1]$,

$$|g_s(x) - g_t(x)| < \epsilon$$

for all $s, t \geq N$.

6.3 Uniform Convergence and Differentiation

Example 6.2.2 (iii) imposes some significant restrictions on what we might hope to be true regarding differentiation and uniform convergence. If $h_n \rightarrow h$ uniformly and each h_n is differentiable, we should not anticipate that $h'_n \rightarrow h'$ because in this example $h'(x)$ does not even exist at $x = 0$.

The key assumption necessary to be able to prove any facts about the derivative of the limit function is that the *sequence of derivatives* be uniformly convergent. This may sound as though we are assuming what it is we would like to prove, and there is some validity to this complaint. The more hypotheses a proposition has, the more difficult it is to apply. The content of the next theorem is that if we are given a pointwise convergent sequence of differentiable functions, and if we know that the sequence of derivatives converges uniformly to *something*, then the limit of the derivatives is indeed the derivative of the limit.

Theorem 6.3.1. *Let $f_n \rightarrow f$ pointwise on the closed interval $[a, b]$, and assume that each f_n is differentiable. If (f'_n) converges uniformly on $[a, b]$ to a function g , then the function f is differentiable and $f' = g$.*

Proof. Let $\epsilon > 0$ and fix $c \in [a, b]$. We want to argue that $f'(c)$ exists and equals $g(c)$. Because f' is defined by the limit

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

our task is to produce a $\delta > 0$ so that

$$\left| \frac{f(x) - f(c)}{x - c} - g(c) \right| < \epsilon$$

whenever $0 < |x - c| < \delta$. Using the triangle inequality, we can estimate

$$\begin{aligned} \left| \frac{f(x) - f(c)}{x - c} - g(c) \right| &\leq \left| \frac{f(x) - f(c)}{x - c} - \frac{f_n(x) - f_n(c)}{x - c} \right| \\ &\quad + \left| \frac{f_n(x) - f_n(c)}{x - c} - f'_n(c) \right| + |f'_n(c) - g(c)|. \end{aligned}$$

Our intent is to force each of the three terms on the right-hand side to be less than $\epsilon/3$. This will not be too difficult in the case of the third term (because

$f'_n \rightarrow g$ uniformly) nor in the case of the second (because f'_n is differentiable). Handling the first term requires the most delicate touch, and we tend to this task first.

Apply the Mean Value Theorem to the function $f_m - f_n$ on the interval $[c, x]$. (If $x < c$, the argument is the same.) By MVT, there exists an $\alpha \in (c, x)$ such that

$$f'_m(\alpha) - f'_n(\alpha) = \frac{(f_m(x) - f_n(x)) - (f_m(c) - f_n(c))}{x - c}.$$

Now, by the Cauchy Criterion for Uniform Convergence (Theorem 6.2.5), there exists an $N_1 \in \mathbf{N}$ such that $m, n \geq N_1$ implies

$$|f'_m(\alpha) - f'_n(\alpha)| < \frac{\epsilon}{3}.$$

We should point out that α depends on the choice of m and n , so it is crucial to have uniform convergence of (f'_n) at this point in the argument. Putting these last two statements together leads to the conclusion that

$$\left| \frac{f_m(x) - f_m(c)}{x - c} - \frac{f_n(x) - f_n(c)}{x - c} \right| < \frac{\epsilon}{3}$$

for all $m, n \geq N_1$ and all $x \in [a, b]$. Because $f_m \rightarrow f$, we can take the limit as $m \rightarrow \infty$ and use the Order Limit Theorem (Theorem 2.3.4) to assert that

$$\left| \frac{f(x) - f(c)}{x - c} - \frac{f_n(x) - f_n(c)}{x - c} \right| \leq \frac{\epsilon}{3}$$

for all $n \geq N_1$.

To complete the proof, choose N_2 large enough so that

$$|f'_m(c) - g(c)| < \frac{\epsilon}{3}$$

for all $m \geq N_2$, and then let $N = \max\{N_1, N_2\}$. Having settled on a choice of N , we use the fact that f_N is differentiable to produce a $\delta > 0$ for which

$$\left| \frac{f_N(x) - f_N(c)}{x - c} - f'_N(c) \right| < \frac{\epsilon}{3}$$

whenever $0 < |x - c| < \delta$. Finally, we observe that for these values of x ,

$$\begin{aligned} \left| \frac{f(x) - f(c)}{x - c} - g(c) \right| &\leq \left| \frac{f(x) - f(c)}{x - c} - \frac{f_N(x) - f_N(c)}{x - c} \right| \\ &\quad + \left| \frac{f_N(x) - f_N(c)}{x - c} - f'_N(c) \right| + |f'_N(c) - g(c)| \\ &< \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} = \epsilon \end{aligned}$$

□

The hypothesis in Theorem 6.3.1 is unnecessarily strong. We actually do not need to assume that $f_n(x) \rightarrow f(x)$ at each point in the domain because the assumption that the sequence of derivatives (f'_n) converges uniformly is nearly strong enough to *prove* that (f_n) converges, uniformly in fact. Two functions with the same derivative may differ by a constant, so we must assume that there is at least one point x_0 where $f_n(x_0) \rightarrow f(x_0)$.

Theorem 6.3.2. *Let (f_n) be a sequence of differentiable functions defined on the closed interval $[a, b]$, and assume (f'_n) converges uniformly on $[a, b]$. If there exists a point $x_0 \in [a, b]$ where $f_n(x_0)$ is convergent, then (f_n) converges uniformly on $[a, b]$.*

Proof. Exercise 6.3.5. □

Combining the last two results produces a stronger version of Theorem 6.3.1.

Theorem 6.3.3. *Let (f_n) be a sequence of differentiable functions defined on the closed interval $[a, b]$, and assume (f'_n) converges uniformly to a function g on $[a, b]$. If there exists a point $x_0 \in [a, b]$ for which $f_n(x_0)$ is convergent, then (f_n) converges uniformly. Moreover, the limit function $f = \lim f_n$ is differentiable and satisfies $f' = g$.*

Exercises

Exercise 6.3.1. (a) Let

$$h_n(x) = \frac{\sin(nx)}{n}.$$

Show that $h_n \rightarrow 0$ uniformly on \mathbf{R} . At what points does the sequence of derivatives h'_n converge?

(b) Modify this example to show that it is possible for a sequence (f_n) to converge uniformly but for (f'_n) to be unbounded.

Exercise 6.3.2. Consider the sequence of functions defined by

$$g_n(x) = \frac{x^n}{n}.$$

(a) Show (g_n) converges uniformly on $[0, 1]$ and find $g = \lim g_n$. Show that g is differentiable and compute $g'(x)$ for all $x \in [0, 1]$.

(b) Now, show that (g'_n) converges on $[0, 1]$. Is the convergence uniform? Set $h = \lim g'_n$ and compare h and g' . Are they the same?

Exercise 6.3.3. Consider the sequence of functions

$$f_n(x) = \frac{x}{1 + nx^2}.$$

Exercise 6.2.4 contains some advice for how to show that (f_n) converges uniformly on \mathbf{R} . Review or complete this exercise.

Now, let $f = \lim f_n$. Compute $f'_n(x)$ and find all the values of x for which $f'(x) = \lim f'_n(x)$.

Exercise 6.3.4. Let

$$g_n(x) = \frac{nx + x^2}{2n},$$

and set $g(x) = \lim g_n(x)$. Show that g is differentiable in two ways:

(a) Compute $g(x)$ by algebraically taking the limit as $n \rightarrow \infty$ and then find $g'(x)$.

(b) Compute $g'_n(x)$ for each $n \in \mathbf{N}$ and show that the sequence of derivatives (g'_n) converges uniformly on every interval $[-M, M]$. Use Theorem 6.3.3 to conclude $g'(x) = \lim g'_n(x)$.

(c) Repeat parts (a) and (b) for the sequence $f_n(x) = (nx^2 + 1)/(2n + x)$.

Exercise 6.3.5. Use the following advice to supply a proof for Theorem 6.3.2. To get started, observe that the triangle inequality implies that, for any $x \in [a, b]$,

$$|f_n(x) - f_m(x)| \leq |(f_n(x) - f_m(x)) - (f_n(x_0) - f_m(x_0))| + |f_n(x_0) - f_m(x_0)|.$$

Now, apply the Mean Value Theorem to $f_n - f_m$.

6.4 Series of Functions

Definition 6.4.1. For each $n \in \mathbf{N}$, let f_n and f be functions defined on a set $A \subseteq \mathbf{R}$. The infinite series

$$\sum_{n=1}^{\infty} f_n(x) = f_1(x) + f_2(x) + f_3(x) + \cdots$$

converges pointwise on A to $f(x)$ if the sequence $s_k(x)$ of partial sums defined by

$$s_k(x) = f_1(x) + f_2(x) + \cdots + f_k(x)$$

converges pointwise to $f(x)$. The series converges uniformly on A to f if the sequence $s_k(x)$ converges uniformly on A to $f(x)$.

In either case, we write $f = \sum_{n=1}^{\infty} f_n$ or $f(x) = \sum_{n=1}^{\infty} f_n(x)$, always being explicit about the type of convergence involved.

If we have a series $\sum_{n=1}^{\infty} f_n$ where the functions f_n are continuous, then the Algebraic Continuity Theorem (Theorem 4.3.4) guarantees that the partial sums—because they are finite sums—will be continuous as well. A corresponding observation is true if we are dealing with differentiable functions. As a consequence, we can immediately translate the results for sequences in the previous sections into statements about the behavior of infinite series of functions.

Theorem 6.4.2. Let f_n be continuous functions defined on a set $A \subseteq \mathbf{R}$, and assume $\sum_{n=1}^{\infty} f_n$ converges uniformly on A to a function f . Then, f is continuous on A .

Proof. Apply Theorem 6.2.6 to the partial sums $s_k = f_1 + f_2 + \cdots + f_k$. \square

Theorem 6.4.3. Let f_n be differentiable functions defined on an interval $[a, b]$, and assume $\sum_{n=1}^{\infty} f'_n(x)$ converges uniformly to a limit $g(x)$ on A . If there exists a point $x_0 \in [a, b]$ where $\sum_{n=1}^{\infty} f_n(x_0)$ converges, then the series $\sum_{n=1}^{\infty} f_n(x)$ converges uniformly to a differentiable function $f(x)$ satisfying $f'(x) = g(x)$ on $[a, b]$. In other words,

$$f(x) = \sum_{n=1}^{\infty} f_n(x) \quad \text{and} \quad f'(x) = \sum_{n=1}^{\infty} f'_n(x).$$

Proof. Apply Theorem 6.3.3 to the partial sums $s_k = f_1 + f_2 + \cdots + f_k$. Observe that Theorem 5.2.4 implies that $s'_k = f'_1 + f'_2 + \cdots + f'_k$. \square

In the vocabulary of infinite series, the Cauchy Criterion takes the following form.

Theorem 6.4.4 (Cauchy Criterion for Uniform Convergence of Series). A series $\sum_{n=1}^{\infty} f_n$ converges uniformly on $A \subseteq \mathbf{R}$ if and only if for every $\epsilon > 0$ there exists an $N \in \mathbf{N}$ such that for all $n > m \geq N$,

$$|f_{m+1}(x) + f_{m+2}(x) + f_{m+3}(x) + \cdots + f_n(x)| < \epsilon$$

for all $x \in A$.

The benefits of uniform convergence over pointwise convergence suggest the need for some ways of determining when a series converges uniformly. The following corollary to the Cauchy Criterion is the most common such tool. In particular, it will be quite useful in our upcoming investigations of power series.

Corollary 6.4.5 (Weierstrass M-Test). For each $n \in \mathbf{N}$, let f_n be a function defined on a set $A \subseteq \mathbf{R}$, and let $M_n > 0$ be a real number satisfying

$$|f_n(x)| \leq M_n$$

for all $x \in A$. If $\sum_{n=1}^{\infty} M_n$ converges, then $\sum_{n=1}^{\infty} f_n$ converges uniformly on A .

Proof. Exercise 6.4.2. \square

Exercises

Exercise 6.4.1. Prove that if $\sum_{n=1}^{\infty} g_n$ converges uniformly, then (g_n) converges uniformly to zero.

Exercise 6.4.2. Supply the details for the proof of the Weierstrass M-Test (Corollary 6.4.5).

Exercise 6.4.3. (a) Show that $g(x) = \sum_{n=1}^{\infty} \cos(2^n x)/2^n$ is continuous on all of \mathbf{R} .

(b) Prove that $h(x) = \sum_{n=1}^{\infty} x^n/n^2$ is continuous on $[-1, 1]$.

Exercise 6.4.4. In Section 5.4, we postponed the argument that the nowhere-differentiable function

$$g(x) = \sum_{n=0}^{\infty} \frac{1}{2^n} h(2^n x)$$

is continuous on \mathbf{R} . Use the Weierstrass M-Test to supply the missing proof.

Exercise 6.4.5. Let

$$f(x) = \sum_{k=1}^{\infty} \frac{\sin(kx)}{k^3}.$$

(a) Show that $f(x)$ is differentiable and that the derivative $f'(x)$ is continuous.

(b) Can we determine if f is twice-differentiable?

Exercise 6.4.6. Observe that the series

$$f(x) = \sum_{n=1}^{\infty} \frac{x^n}{n} = x + \frac{x^2}{2} + \frac{x^3}{3} + \frac{x^4}{4} + \cdots$$

converges for every x in the half-open interval $[0, 1)$ but does not converge when $x = 1$. For a fixed $x_0 \in (0, 1)$, explain how we can still use the Weierstrass M-Test to prove that f is continuous at x_0 .

Exercise 6.4.7. Let

$$h(x) = \sum_{n=1}^{\infty} \frac{1}{x^2 + n^2}.$$

(a) Show that h is a continuous function defined on all of \mathbf{R} .

(b) Is h differentiable? If so, is the derivative function h' continuous?

Exercise 6.4.8. Let $\{r_1, r_2, r_3, \dots\}$ be an enumeration of the set of rational numbers. For each $r_n \in \mathbf{Q}$, define

$$u_n(x) = \begin{cases} 1/2^n & \text{for } x > r_n \\ 0 & \text{for } x \leq r_n. \end{cases}$$

Now, let $h(x) = \sum_{n=1}^{\infty} u_n(x)$. Prove that h is a monotone function defined on all of \mathbf{R} that is continuous at every irrational point.

6.5 Power Series

It is time to put some mathematical teeth into our understanding of functions expressed in the form of a power series; that is, functions of the form

$$f(x) = \sum_{n=0}^{\infty} a_n x^n = a_0 + a_1 x + a_2 x^2 + a_3 x^3 + \cdots.$$

The first order of business is to determine the points $x \in \mathbf{R}$ for which the resulting series on the right-hand side converges. This set certainly contains $x = 0$, and, as the next result demonstrates, it takes a very predictable form.

Theorem 6.5.1. *If a power series $\sum_{n=0}^{\infty} a_n x^n$ converges at some point $x_0 \in \mathbf{R}$, then it converges absolutely for any x satisfying $|x| < |x_0|$.*

Proof. If $\sum_{n=0}^{\infty} a_n x_0^n$ converges, then the sequence of terms $(a_n x_0^n)$ is bounded. (In fact, it converges to 0.) Let $M > 0$ satisfy $|a_n x_0^n| \leq M$ for all $n \in \mathbf{N}$. If $x \in \mathbf{R}$ satisfies $|x| < |x_0|$, then

$$|a_n x^n| = |a_n x_0^n| \left| \frac{x}{x_0} \right|^n \leq M \left| \frac{x}{x_0} \right|^n.$$

But notice that

$$\sum_{n=0}^{\infty} M \left| \frac{x}{x_0} \right|^n$$

is a geometric series with ratio $|x/x_0| < 1$ and so converges. By the Comparison Test, $\sum_{n=0}^{\infty} a_n x^n$ converges absolutely. \square

The main implication of Theorem 6.5.1 is that the set of points for which a given power series converges must necessarily be $\{0\}$, \mathbf{R} , or a bounded interval centered around $x = 0$. Because of the strict inequality in Theorem 6.5.1, there is some ambiguity about the endpoints of the interval, and it is possible that the set of convergent points may be of the form $(-R, R)$, $[-R, R)$, $(-R, R]$, or $[-R, R]$.

The value of R is referred to as the *radius of convergence* of a power series, and it is customary to assign R the value 0 or ∞ to represent the set $\{0\}$ or \mathbf{R} , respectively. Some of the standard devices for computing the radius of convergence for a power series are explored in the exercises. Of more interest to us here is the investigation of the properties of functions defined in this way. Are they continuous? Are they differentiable? If so, can we differentiate the series term-by-term? What happens at the endpoints?

Establishing Uniform Convergence

The positive answers to the preceding questions, and the usefulness of power series in general, are largely due to the fact that they converge uniformly on compact sets contained in their domain of convergent points. As we are about to see, a complete proof of this fact requires a fairly delicate argument attributed to the Norwegian mathematician Niels Abel. A significant amount of progress, however, can be made with the Weierstrass M-Test (Corollary 6.4.5).

Theorem 6.5.2. *If a power series $\sum_{n=0}^{\infty} a_n x^n$ converges absolutely at a point x_0 , then it converges uniformly on the closed interval $[-c, c]$, where $c = |x_0|$.*

Proof. This proof requires a straightforward application of the Weierstrass M-Test. The details are requested in Exercise 6.5.5. \square

For many applications, Theorem 6.5.2 is good enough. For instance, combining our results about uniform convergence and power series to this point, we

can now argue that a power series that converges on an open interval $(-R, R)$ is necessarily continuous on this interval (Exercise 6.5.4).

But what happens if we know that a series converges at an endpoint of its interval of convergence? Does the good behavior of the series on $(-R, R)$ necessarily extend to the endpoint $x = R$? If the convergence of the series at $x = R$ is absolute convergence, then we can again rely on Theorem 6.5.2 to conclude that the series converges uniformly on the set $[-R, R]$. The remaining interesting open question is what happens if a series converges *conditionally* at a point $x = R$. We may still use Theorem 6.5.1 to conclude that we have pointwise convergence on the interval $(-R, R]$, but more work is needed to establish uniform convergence on compact sets containing $x = R$.

Abel's Theorem

We should remark that if the power series $g(x) = \sum_{n=0}^{\infty} a_n x^n$ converges conditionally at $x = R$, then it is possible for it to diverge when $x = -R$. The series

$$\sum_{n=1}^{\infty} \frac{(-1)^n x^n}{n}$$

with $R = 1$ is an example. To keep our attention fixed on the convergent endpoint, we will prove uniform convergence on the set $[0, R]$.

The argument we need is very similar to the proof of Abel's Test, which is included in the exercises of Section 2.7. The first step in the proof of Abel's Test is an estimate, sometimes called Abel's Lemma, which we will presently need again.

Lemma 6.5.3 (Abel's Lemma). *Let b_n satisfy $b_1 \geq b_2 \geq b_3 \geq \cdots \geq 0$, and let $\sum_{n=1}^{\infty} a_n$ be a series for which the partial sums are bounded. In other words, assume there exists $A > 0$ such that*

$$|a_1 + a_2 + \cdots + a_n| \leq A$$

for all $n \in \mathbf{N}$. Then, for all $n \in \mathbf{N}$,

$$|a_1 b_1 + a_2 b_2 + a_3 b_3 + \cdots + a_n b_n| \leq 2A b_1.$$

Proof. This inequality follows from the so-called summation-by-parts formula. Exercises 2.7.12 and 2.7.14(b) from Section 2.7 contain the relevant definitions and advice required to complete the argument. \square

It is worth observing that if A were an upper bound on the partial sums of $\sum |a_n|$ (note the absolute value bars), then the proof of Lemma 6.5.3 would be a simple exercise in the triangle inequality. (Also, we would not need the factor of 2, and, actually, it is not needed in general except that our particular method of proof requires it.) The point of the matter is that because we are only assuming conditional convergence, the triangle inequality is not going to be of any use in proving Abel's Theorem, but we are now in possession of an inequality that we can use in its place.

Theorem 6.5.4 (Abel's Theorem). *Let $g(x) = \sum_{n=0}^{\infty} a_n x^n$ be a power series that converges at the point $x = R > 0$. Then the series converges uniformly on the interval $[0, R]$. A similar result holds if the series converges at $x = -R$.*

Proof. To set the stage for an application of Lemma 6.5.3, we first write

$$g(x) = \sum_{n=0}^{\infty} a_n x^n = \sum_{n=0}^{\infty} (a_n R^n) \left(\frac{x}{R}\right)^n.$$

Let $\epsilon > 0$. By the Cauchy Criterion for Uniform Convergence of Series (Theorem 6.4.4), we will be done if we can produce an N such that $n > m \geq N$ implies

$$(1) \quad \left| (a_{m+1} R^{m+1}) \left(\frac{x}{R}\right)^{m+1} + (a_{m+2} R^{m+2}) \left(\frac{x}{R}\right)^{m+2} + \cdots \right. \\ \left. + (a_n R^n) \left(\frac{x}{R}\right)^n \right| < \epsilon.$$

Because we are assuming that $\sum_{n=1}^{\infty} a_n R^n$ converges, the Cauchy Criterion for convergent series of real numbers guarantees that there exists an N such that

$$|a_{m+1} R^{m+1} + a_{m+2} R^{m+2} + \cdots + a_n R^n| < \frac{\epsilon}{2}$$

whenever $n > m \geq N$. But now, for any fixed $m \in N$, we can apply Lemma 6.5.3 to the sequences obtained by omitting the first m terms. Using $\epsilon/2$ as a bound on the partial sums of $\sum_{j=1}^{\infty} a_{m+j} R^{m+j}$ and observing that $(x/R)^{m+j}$ is monotone decreasing, an application of Lemma 6.5.3 to equation (1) yields

$$\left| (a_{m+1} R^{m+1}) \left(\frac{x}{R}\right)^{m+1} + (a_{m+2} R^{m+2}) \left(\frac{x}{R}\right)^{m+2} + \cdots \right. \\ \left. + (a_n R^n) \left(\frac{x}{R}\right)^n \right| \leq 2 \left(\frac{\epsilon}{2}\right) \left(\frac{x}{R}\right)^{m+1} \leq \epsilon.$$

The fact that the inequality is not strict (as the Cauchy Criterion technically requires) is a distraction but not a real deficiency. We leave it as a point for discussion. \square

The Success of Power Series

An economical way to summarize the conclusions of Theorem 6.5.2 and Abel's Theorem is with the following statement.

Theorem 6.5.5. *If a power series converges pointwise on the set $A \subseteq \mathbf{R}$, then it converges uniformly on any compact set $K \subseteq A$.*

Proof. Exercise 6.5.6. \square

This fact leads to the desirable conclusion that a power series is continuous at every point at which it converges. To make an argument for differentiability, we would like to appeal to Theorem 6.4.3; however, this result has a slightly more involved set of hypotheses. In order to conclude that a power series $\sum_{n=0}^{\infty} a_n x^n$ is differentiable, and that term-by-term differentiation is allowed, we need to know *beforehand* that the differentiated series $\sum_{n=1}^{\infty} n a_n x^{n-1}$ converges uniformly.

Theorem 6.5.6. *If $\sum_{n=0}^{\infty} a_n x^n$ converges for all $x \in (-R, R)$, then the differentiated series $\sum_{n=1}^{\infty} n a_n x^{n-1}$ converges at each $x \in (-R, R)$ as well. Consequently, the convergence is uniform on compact sets contained in $(-R, R)$.*

Proof. Exercise 6.5.7. □

We should point out that it is possible for a series to converge at an endpoint $x = R$ but for the differentiated series to diverge at this point. The series $\sum_{n=1}^{\infty} (-x)^n/n$ has this property when $x = 1$. On the other hand, if the differentiated series does converge at the point $x = R$, then Abel's Theorem applies and the convergence of the differentiated series is uniform on compact sets that contain R .

With all the pieces in place, we summarize the impressive conclusions of this section.

Theorem 6.5.7. *Assume*

$$g(x) = \sum_{n=0}^{\infty} a_n x^n$$

converges on an interval $A \subseteq \mathbf{R}$. The function g is continuous on A and differentiable on any open interval $(-R, R) \subseteq A$. The derivative is given by

$$g'(x) = \sum_{n=1}^{\infty} n a_n x^{n-1}.$$

Moreover, g is infinitely differentiable on $(-R, R)$, and the successive derivatives can be obtained via term-by-term differentiation of the appropriate series.

Proof. The details for why g is continuous are requested in the exercises. Theorem 6.5.6 justifies the application of Theorem 6.4.3, which verifies the formula for g' .

A differentiated power series is a power series in its own right, and Theorem 6.5.6 implies that, although the series may no longer converge at a particular endpoint, the radius of convergence does not change. By induction, then, power series are differentiable an infinite number of times. □

Exercises

Exercise 6.5.1. Consider the function g defined by the power series

$$g(x) = x - \frac{x^2}{2} + \frac{x^3}{3} - \frac{x^4}{4} + \frac{x^5}{5} - \cdots.$$

(a) Is g defined on $(-1, 1)$? Is it continuous on this set? Is g defined on $(-1, 1]$? Is it continuous on this set? What happens on $[-1, 1]$? Can the power series for $g(x)$ possibly converge for any other points $|x| > 1$? Explain.

(b) For what values of x is $g'(x)$ defined? Find a formula for g' .

Exercise 6.5.2. Find suitable coefficients (a_n) so that the resulting power series $\sum a_n x^n$

(a) converges absolutely for all $x \in [-1, 1]$ and diverges off of this set;

(b) converges conditionally at $x = -1$ and diverges at $x = 1$;

(c) converges conditionally at both $x = -1$ and $x = 1$.

(d) Is it possible to find an example of a power series that converges conditionally at $x = -1$ and converges absolutely at $x = 1$?

Exercise 6.5.3. Explain why a power series can converge conditionally for at most two points.

Exercise 6.5.4. (a) By referencing the proper theorems, produce a detailed argument that a power series that converges on the interval $(-R, R)$ necessarily represents a continuous function at each point $x \in (-R, R)$.

(b) If the series converges at an endpoint $x = R$, point out how we know continuity extends to the set $(R, R]$.

Exercise 6.5.5. Use the Weierstrass M-Test to prove Theorem 6.5.2.

Exercise 6.5.6. Show how Theorem 6.5.1, Theorem 6.5.2, and Abel's Theorem together imply that if a power series converges pointwise on a compact set, then the convergence is actually uniform on this set.

Exercise 6.5.7. (a) The Ratio Test (from Exercise 2.7.9) states that if (b_n) is a sequence of nonzero terms satisfying $\lim |b_{n+1}/b_n| = r < 1$, then the series $\sum b_n$ converges. Use this to argue that if s satisfies $0 < s < 1$, then ns^{n-1} is bounded for all $n \geq 1$.

(b) Given an arbitrary $x \in (-R, R)$, pick t to satisfy $|x| < t < R$. Use the observation

$$|na_n x^{n-1}| = \frac{1}{t} \left(n \left| \frac{x^{n-1}}{t^{n-1}} \right| \right) |a_n t^n|$$

to construct a proof for Theorem 6.5.6.

Exercise 6.5.8. Let $\sum a_n x^n$ be a power series with $a_n \neq 0$, and assume

$$L = \lim_{n \rightarrow \infty} \left| \frac{a_{n+1}}{a_n} \right|$$

exists.

(a) Show that if $L \neq 0$, then the series converges for all x in $(-1/L, 1/L)$. (The advice in Exercise 2.7.9 may be helpful.)

(b) Show that if $L = 0$, then the series converges for all $x \in \mathbf{R}$.

(c) Show that (a) and (b) continue to hold if L is replaced by the limit

$$L' = \lim_{n \rightarrow \infty} s_n \quad \text{where} \quad s_n = \sup \left\{ \left| \frac{a_{k+1}}{a_k} \right| : k \geq n \right\}.$$

The value L' is called the “limit superior” or “lim sup” of the sequence $|a_{n+1}/a_n|$. It exists if and only if the sequence is bounded (Exercise 2.4.6).

(d) Show that if $|a_{n+1}/a_n|$ is unbounded, then the original series $\sum a_n x^n$ converges only when $x = 0$.

Exercise 6.5.9. Use Theorem 6.5.7 to argue that power series are unique. If we have

$$\sum_{n=0}^{\infty} a_n x^n = \sum_{n=0}^{\infty} b_n x^n$$

for all x in an interval $(-R, R)$, prove that $a_n = b_n$ for all $n = 0, 1, 2, \dots$. (Start by showing that $a_0 = b_0$.)

Exercise 6.5.10. Review the definitions and results from Section 2.8 concerning products of series and Cauchy products in particular. At the end of Section 2.9, we mentioned the following result: If both $\sum a_n$ and $\sum b_n$ converge conditionally to A and B respectively, then it is possible for the Cauchy product,

$$\sum d_n \quad \text{where} \quad d_n = a_0 b_n + a_1 b_{n-1} + \cdots + a_n b_0,$$

to diverge. However, if $\sum d_n$ does converge, then it must converge to AB . To prove this, set

$$f(x) = \sum a_n x^n, \quad g(x) = \sum b_n x^n, \quad \text{and} \quad h(x) = \sum d_n x^n.$$

Use Abel’s Theorem and the result in Exercise 2.8.8 to establish this result.

Exercise 6.5.11. A series $\sum_{n=0}^{\infty} a_n$ is said to be *Abel-summable to L* if the power series

$$f(x) = \sum_{n=0}^{\infty} a_n x^n$$

converges for all $x \in [0, 1)$ and $L = \lim_{x \rightarrow 1^-} f(x)$.

(a) Show that any series that converges to a limit L is also Abel-summable to L .

(b) Show that $\sum_{n=0}^{\infty} (-1)^n$ is Abel-summable and find the sum.

Exercise 6.5.12. Consider the function G from the opening discussion on branching processes, and recall that the increasing monotone sequence of probabilities (d_r) has a limit $d = \lim d_r$ that satisfies $G(d) = d$. Assume we are in the situation where there are two fixed points: $G(1) = 1$ and some other value $0 < d_0 < 1$ satisfying $G(d_0) = d_0$. Formulate an argument for why the sequence (d_r) necessarily converges to the value $d = d_0$ and not to $d = 1$.

6.6 Taylor Series

Our study of power series has led to some enthusiastic conclusions about the nature of functions of the form

$$f(x) = a_0 + a_1x + a_2x^2 + a_3x^3 + a_4x^4 + \cdots .$$

Despite their infinite character, power series can be manipulated more or less as though they are polynomials. On its interval of convergence, a power series is continuous and infinitely differentiable, and successive derivatives can be computed by performing the desired operation on each individual term in the series—just as it is done for polynomials. As we will see in the next chapter, the situation regarding integrals of power series is just as pleasant. As students of introductory calculus know well, the processes of integration and differentiation, as well as basic algebraic manipulations, are rather straightforward when applied only to polynomials. It is the introduction of functions such as $\sin(x)$, $\ln(x)$, and $\arctan(x)$ that necessitates much of the symbolic acrobatics taught in first and second semester calculus courses. This phenomenon leads us to ask whether functions such as $\arctan(x)$ and the rest have representations as power series.

In the examples in this section, we will assume all of the familiar properties of the trigonometric, inverse trigonometric, exponential, and logarithmic functions. Rigorously defining these functions is an interesting exercise in analysis. In fact, one of the most common methods for providing definitions is through power series. The point of this discussion, however, is to come at this question from the other direction. Assuming we are in possession of an infinitely differentiable function such as $\arctan(x)$, can we find suitable coefficients a_n so that

$$\arctan(x) = a_0 + a_1x + a_2x^2 + a_3x^3 + a_4x^4 + \cdots$$

for at least some nonzero values of x ?

Manipulating Series

We already have one example of a power series expansion of a familiar function. In the material on infinite series in Example 2.7.5 from Chapter 2, we proved that

$$(1) \quad \frac{1}{1-t} = 1 + t + t^2 + t^3 + t^4 + \cdots$$

for all $t \in (-1, 1)$. Substituting $-t^2$ for t yields

$$\frac{1}{1+t^2} = 1 - t^2 + t^4 - t^6 + t^8 - \cdots .$$

But now we can use the fact that

$$\arctan(x) = \int_0^x \frac{1}{1+t^2} dt.$$

Although integration has not been rigorously studied at this point, we will see in Chapter 7 that if $f_n \rightarrow f$ uniformly on an interval $[a, b]$, then $\int_a^b f_n \rightarrow \int_a^b f$. This observation, together with the Fundamental Theorem of Calculus, leads to the formula

$$\arctan(x) = x - \frac{1}{3}x^3 + \frac{1}{5}x^5 - \frac{1}{7}x^7 + \cdots .$$

Exercise 6.6.1. Upcoming results in Chapter 7 will justify this equation for all $x \in (-1, 1)$, but notice that this series actually converges when $x = 1$. Assuming that $\arctan(x)$ is continuous, explain why the value of the series at $x = 1$ must necessarily be $\arctan(1)$. What interesting identity do we get in this case?

Exercise 6.6.2. Starting from the identity in equation (1) of this section, find a power series representation for $\ln(1+x)$. For which values of x is this expression valid?

Taylor's Formula for the Coefficients

Manipulating old series to produce new ones was a well-honed craft in the 17th and 18th centuries when the usefulness of infinite series was first being realized. But there also emerged a formula for producing the coefficients from “scratch” — a recipe for generating a power series representation using only the function in question and its derivatives. The technique is named after the mathematician Brook Taylor (1685–1731) who published it in 1715, although it was certainly known previous to this date.

Given an infinitely differentiable function f defined on some interval centered at zero, the idea is to assume that f has a power series expansion and deduce what the coefficients must be; that is, write

$$(2) \quad f(x) = a_0 + a_1x + a_2x^2 + a_3x^3 + a_4x^4 + a_5x^5 + \cdots .$$

Setting $x = 0$ in this expression gives $f(0) = a_0$.

Exercise 6.6.3. (a) Take derivatives of each side of equation (2), and deduce that $f'(0) = a_1$. In general, prove that if f has a power series expansion, then the coefficients must be given by the formula

$$a_n = \frac{f^{(n)}(0)}{n!} .$$

Supply references to the theorem(s) that justify the manipulations carried out on the series in equation (2).

Exercise 6.6.4. Use Taylor's formula for a_n from the preceding exercise to produce/verify the so-called *Taylor series* for $\sin(x)$ given by

$$\sin(x) \sim x - \frac{x^3}{3!} + \frac{x^5}{5!} - \frac{x^7}{7!} + \cdots .$$

Lagrange's Remainder Theorem

We need to be very clear about what we have proved to this point. To derive Taylor's formula, *we assumed that f actually had a power series representation*. The conclusion is that if f is infinitely differentiable on an interval centered at zero, and if f can be expressed in the form

$$f(x) = \sum_{n=0}^{\infty} a_n x^n,$$

then it must be that

$$a_n = \frac{f^{(n)}(0)}{n!}.$$

But what about the converse question? Assume f is infinitely differentiable in a neighborhood of zero. If we let

$$a_n = \frac{f^{(n)}(0)}{n!},$$

does the resulting series

$$\sum_{n=0}^{\infty} a_n x^n$$

converge to $f(x)$ on some nontrivial set of points? Does it converge at all? If it does converge, we know that the limit function is a well-behaved, infinitely differentiable function whose derivatives at zero are exactly the same as the derivatives of f . Is it possible for this limit to be different from f ? In other words, might the Taylor series of a function f converge to the wrong thing?

Let

$$S_N(x) = a_0 + a_1x + a_2x^2 + \cdots + a_Nx^N.$$

The polynomial $S_N(x)$ is a partial sum of the Taylor series expansion for the function $f(x)$. Thus, we are interested in whether or not

$$\lim_{N \rightarrow \infty} S_N(x) = f(x)$$

for some values of x besides zero. A powerful tool for analyzing this question was provided by Joseph Louis Lagrange (1736–1813). The idea is to consider the difference

$$E_N(x) = f(x) - S_N(x),$$

which represents the error between f and the partial sum S_N .

Theorem 6.6.1 (Lagrange's Remainder Theorem). *Let f be infinitely differentiable on $(-R, R)$, define $a_n = f^{(n)}(0)/n!$, and let*

$$S_N = a_0 + a_1x + a_2x^2 + \cdots + a_Nx^N.$$

Given $x \neq 0$, there exists a point c satisfying $|c| < |x|$ where the error function $E_N(x) = f(x) - S_N(x)$ satisfies

$$E_N(x) = \frac{f^{(N+1)}(c)}{(N+1)!} x^{N+1}.$$

Before embarking on a proof, let's examine the significance of this result. Proving $S_N(x) \rightarrow f(x)$ is equivalent to showing $E_N(x) \rightarrow 0$. There are three components to the expression for $E_N(x)$. In the denominator, we have $(N+1)!$, which helps to make E_N small as N tends to infinity. In the numerator, we have x^{N+1} , which potentially grows depending on the size of x . Thus, we should expect that a Taylor series is less likely to converge the farther x is chosen from the origin. Finally, we have $f^{(N+1)}(c)$, which is a bit of a mystery. For functions with straightforward derivatives, this term can often be handled using a suitable upper bound.

Example 6.6.2. Consider the Taylor series for $\sin(x)$ generated earlier. How well does

$$S_5(x) = x - \frac{1}{3!}x^3 + \frac{1}{5!}x^5$$

approximate $\sin(x)$ on the interval $[-2, 2]$? Lagrange's Remainder Theorem asserts that the difference between these two functions is

$$E_5(x) = \sin(x) - S_5(x) = \frac{-\cos(c)}{6!}x^6$$

for some c in the interval $(-|x|, |x|)$. Without knowing the value of c , we can be quite certain that $|\cos(c)| \leq 1$. Because $x \in [-2, 2]$, we have that

$$|E_5(x)| \leq \frac{2^6}{6!} \approx .089.$$

Exercise 6.6.5. Prove that $S_N(x)$ converges uniformly to $\sin(x)$ on $[-2, 2]$. Generalize this proof to show that the convergence is uniform on any interval of the form $[-R, R]$.

Exercise 6.6.6. (a) Generate the Taylor coefficients for the exponential function $f(x) = e^x$, and then prove that the corresponding Taylor series converges uniformly to e^x on any interval of the form $[-R, R]$.

(b) Verify the formula $f'(x) = e^x$.

(c) Use a substitution to generate the series for e^{-x} , and then calculate $e^x \cdot e^{-x}$ by multiplying together the two series and collecting common powers of x .

Proof of Lagrange's Remainder Theorem: Review the Generalized Mean Value Theorem (Theorem 5.3.5) from Chapter 5.

Exercise 6.6.7. Explain why the error function $E_N(x) = f_N(x) - S_N(x)$ satisfies

$$E_N^{(n)}(0) = 0 \quad \text{for all } n = 0, 1, 2, \dots, N.$$

To simplify notation, let's assume $x > 0$ and apply the Generalized Mean Value Theorem to the functions $E_N(x)$ and x^{N+1} on the interval $[0, x]$. Thus, there exists a point $x_1 \in (0, x)$ such that

$$\frac{E_N(x)}{x^{N+1}} = \frac{E'_N(x_1)}{(N+1)x_1^N}.$$

Exercise 6.6.8. Finish the proof of Lagrange's Remainder Theorem.

A Counterexample

Lagrange's Remainder Theorem is extremely useful for determining how well the partial sums of the Taylor series approximate the original function, but it leaves unresolved the central question of whether or not the Taylor series necessarily converges to the function that generated it. The appearance of the n th derivative $f^{(n)}(c)$ in the error formula makes any general statement impossible. There are, in fact, several other ways to represent the error between the partial sum $S_N(x)$ and the function $f(x)$, but none lend themselves to a proof that $S_N \rightarrow f$. This is because no such proof exists! Let

$$g(x) = \begin{cases} e^{-1/x^2} & \text{for } x \neq 0, \\ 0 & \text{for } x = 0. \end{cases}$$

In the exercises that follow we will need the familiar formula $\frac{d}{dx}e^x = e^x$ and the property that $e^{-x} = 1/e^x$. (Note that we could use the series generated in Exercise 6.6.6 as the *definition* of the exponential function e^x . Parts (b) and (c) of this exercise verify that this series possesses these properties.) Although we have proved all of the standard rules of differentiation, none of these rules can be used to directly compute the derivatives of g at $x = 0$.

Exercise 6.6.9. Use the ∞/∞ version of L'Hospital's rule (Theorem 5.3.8) to prove that $g'(0) = 0$.

Exercise 6.6.10. Compute $g'(x)$ for $x \neq 0$. Compute $g''(x)$ and $g'''(x)$ for $x \neq 0$. Use these observations and invent whatever notation is needed to give a general description for the n th derivative $g^{(n)}(x)$ at points different from zero.

Now,

$$g''(0) = \lim_{x \rightarrow 0} \frac{g'(x) - g'(0)}{x - 0} = \lim_{x \rightarrow 0} \frac{g'(x)}{x}.$$

Exercise 6.6.11. Compute $g''(0)$. From this example, produce a general argument for how to compute $g^{(n)}(0)$.

Exercise 6.6.12. Discuss the consequences of this example. Is g infinitely differentiable? What does its Taylor series look like? At what points does this series converge? To what? What are the implications of this example on the conjecture that every infinitely differentiable function can be represented by its Taylor series expansion?

6.7 Epilogue

The fact that power series behave so impeccably well under the operations of calculus makes the search for Taylor series expansions a worthwhile enterprise. As it turns out, the traditional list of functions from calculus— $\sin(x)$, $\ln(x)$, $\arccos(x)$, $\sqrt{1+x}$ —all have Taylor series representations that converge on some nontrivial interval to the function from which they were derived. This fact played a major role in the expanding achievements of calculus in the 17th and 18th centuries and understandably led to speculation that every function could be represented in such a fashion. (The term “function” at this time implicitly referred to functions that were infinitely differentiable.) This point of view effectively ended with Cauchy’s discovery in 1821 of the counterexample presented at the end of the previous section. So under what conditions does the Taylor series necessarily converge to the generating function? Lagrange’s Remainder Theorem states that the difference between the Taylor polynomial $S_N(x)$ and the function $f(x)$ is given by

$$E_N(x) = \frac{f^{(N+1)}(c)}{(N+1)!} x^{N+1}.$$

The Ratio Test shows that the $(N+1)!$ term in the denominator grows more rapidly than the x^{N+1} term in the numerator. Thus, if we knew for instance that

$$|f^{(N+1)}(c)| \leq M$$

for all $c \in (-R, R)$ and $N \in \mathbf{N}$, we could be sure that $E_N(x) \rightarrow 0$ and hence that $S_N(x) \rightarrow f(x)$. This is the case for $\sin(x)$, $\cos(x)$, and e^x , whose derivatives do not grow at all as $N \rightarrow \infty$. It is also possible to formulate weaker conditions on the rate of growth of $f^{(N+1)}$ that guarantee convergence.

It is not altogether clear whether Cauchy’s counterexample should come as a surprise. The fact that every previous search for a Taylor series ended in success certainly gives the impression that a power series representation is an intrinsic property of infinitely differentiable functions. But notice what we are saying here. A Taylor series for a function f is constructed from the values of f and its derivatives at the origin. If the Taylor series converges to f on some interval $(-R, R)$, then the behavior of f near zero completely determines its behavior at every point in $(-R, R)$. One implication of this would be that if two functions with Taylor series agree on some small neighborhood $(-\epsilon, \epsilon)$, then these two functions would have to be the same everywhere. When it is put this way, we probably should not expect a Taylor series to always converge back to the function from which it was derived. As we have seen, this is not the case for real-valued functions. What is fascinating, however, is that results of this nature *do* hold for functions of a complex variable. The definition of the derivative looks symbolically the same when the real numbers are replaced by complex numbers, but the implications are profoundly different. In this setting, a function that is differentiable at every point in some open disc must necessarily be infinitely differentiable on this set. This supplies the ingredients to construct

the Taylor series that in every instance converges uniformly on compact sets to the function that generated it.

Chapter 7

The Riemann Integral

7.1 Discussion: How Should Integration be Defined?

The Fundamental Theorem of Calculus is a statement about the inverse relationship between differentiation and integration. It comes in two parts, depending on whether we are differentiating an integral or integrating a derivative. Under suitable hypotheses on the functions f and F , the Fundamental Theorem of Calculus states that

$$(i) \int_a^b F'(x) dx = F(b) - F(a) \text{ and}$$

$$(ii) \text{ if } G(x) = \int_a^x f(t) dt, \text{ then } G'(x) = f(x).$$

Before we can undertake any type of rigorous investigation of these statements, we need to settle on a definition for $\int_a^b f$. Historically, the concept of integration was defined as the inverse process of differentiation. In other words, the integral of a function f was understood to be a function F that satisfied $F' = f$. Newton, Leibniz, Fermat, and the other founders of calculus then went on to explore the relationship between antiderivatives and the problem of computing areas. This approach is ultimately unsatisfying from the point of view of analysis because it results in a very limited number of functions that can be integrated. Recall that every derivative satisfies the intermediate value property (Darboux's Theorem, Theorem 5.2.7). This means that any function with a jump discontinuity cannot be a derivative. If we want to define integration via antidifferentiation, then we must accept the consequence that a function as simple as

$$h(x) = \begin{cases} 1 & \text{for } 0 \leq x < 1 \\ 2 & \text{for } 1 \leq x \leq 2 \end{cases}$$

is not integrable on the interval $[0, 2]$.




Figure 7.1: A RIEMANN SUM.

A very interesting shift in emphasis occurred around 1850 in the work of Cauchy, and soon after in the work of Bernhard Riemann. The idea was to completely divorce integration from the derivative and instead use the notion of “area under the curve” as a starting point for building a rigorous definition of the integral. The reasons for this were complicated. As we have mentioned earlier (Section 1.2), the concept of *function* was undergoing a transformation. The traditional understanding of a function as a holistic formula such as $f(x) = x^2$ was being replaced with a more liberal interpretation, which included such bizarre constructions as Dirichlet’s function discussed in Section 4.1. Serving as a catalyst to this evolution was the budding theory of Fourier series (discussed in Section 8.3), which required, among other things, the need to be able to integrate these more unruly objects.

The Riemann integral, as it is called today, is the one usually discussed in introductory calculus. Starting with a function f on $[a, b]$, we partition the domain into small subintervals. On each subinterval $[x_{k-1}, x_k]$, we pick some point $c_k \in [x_{k-1}, x_k]$ and use the y -value $f(c_k)$ as an approximation for f on $[x_{k-1}, x_k]$. Graphically speaking, the result is a row of thin rectangles constructed to approximate the area between f and the x -axis. The area of each rectangle is $f(c_k)(x_k - x_{k-1})$, and so the total area of all of the rectangles is given by the *Riemann sum* (Fig. 7.1)

$$\sum_{k=1}^n f(c_k)(x_k - x_{k-1}).$$

Note that “area” here comes with the understanding that areas below the x -axis are assigned a negative value.

What should be evident from the graph is that the accuracy of the Riemann-sum approximation seems to improve as the rectangles get thinner. In some

sense, we take the *limit* of these approximating Riemann sums as the width of the individual subintervals of the partitions tends to zero. This limit, if it exists, is Riemann's definition of $\int_a^b f$.

This brings us to a handful of questions. Creating a rigorous meaning for the limit just referred to is not too difficult. What will be of most interest to us—and was also to Riemann—is deciding what types of functions can be integrated using this procedure. Specifically, what conditions on f guarantee that this limit exists?

The theory of the Riemann integral turns on the observation that smaller subintervals produce better approximations to the function f . On each subinterval $[x_{k-1}, x_k]$, the function f is approximated by its value at some point $c_k \in [x_{k-1}, x_k]$. The quality of the approximation is directly related to the difference

$$|f(x) - f(c_k)|$$

as x ranges over the subinterval. Because the subintervals can be chosen to have arbitrarily small width, this means that we want $f(x)$ to be close to $f(c_k)$ whenever x is close to c_k . But this sounds like a discussion of continuity! We will soon see that the continuity of f is intimately related to the existence of the Riemann integral $\int_a^b f$.

Is continuity sufficient to prove that the Riemann sums converge to a well-defined limit? Is it necessary, or can the Riemann integral handle a discontinuous function such as $h(x)$ mentioned earlier? Relying on the intuitive notion of area, it would seem that $\int_0^2 h = 3$, but does the Riemann integral reach this conclusion? If so, how discontinuous can a function be before it fails to be integrable? Can the Riemann integral make sense out of something as pathological as Dirichlet's function on the interval $[0, 1]$?

A function such as

$$g(x) = \begin{cases} x^2 \sin(\frac{1}{x}) & \text{for } x \neq 0 \\ 0 & \text{for } x = 0 \end{cases}$$

raises another interesting question. Here is an example of a differentiable function, studied in Section 5.1, where the derivative $g'(x)$ is *not* continuous. As we explore the class of integrable functions, some attempt must be made to reunite the integral with the derivative. Having defined integration independently of differentiation, we would like to come back and investigate the conditions under which equations (i) and (ii) from the Fundamental Theorem of Calculus stated earlier hold. If we are making a wish list for the types of functions that we want to be integrable, then in light of equation (i) it seems desirable to expect this set to at least contain the set of derivatives. The fact that derivatives are not always continuous is further motivation not to content ourselves with an integral that cannot handle some discontinuities.




Figure 7.2: UPPER AND LOWER SUMS.

7.2 The Definition of the Riemann Integral

Although it has the benefit of some modern polish, the development of the integral presented in this chapter is closely related to the procedure just discussed. In place of Riemann sums, we will construct *upper sums* and *lower sums* (Fig. 7.2), and in place of a limit we will use a supremum and an infimum.

Throughout this section, it is assumed that we are working with a *bounded* function f on a closed interval $[a, b]$, meaning that there exists an $M > 0$ such that $|f(x)| \leq M$ for all $x \in [a, b]$.

Partitions, Upper Sums, and Lower Sums

Definition 7.2.1. A *partition* P of $[a, b]$ is a finite, ordered set

$$P = \{a = x_0 < x_1 < x_2 < \cdots < x_n = b\}.$$

For each subinterval $[x_{k-1}, x_k]$ of P , let

$$m_k = \inf\{f(x) : x \in [x_{k-1}, x_k]\} \quad \text{and} \quad M_k = \sup\{f(x) : x \in [x_{k-1}, x_k]\}.$$

The *lower sum* of f with respect to P is given by

$$L(f, P) = \sum_{k=1}^n m_k(x_k - x_{k-1}).$$

Likewise, we define the *upper sum* of f with respect to P by


$$U(f, P) = \sum_{k=1}^n M_k(x_k - x_{k-1}).$$

For a particular partition P , it is clear that $U(f, P) \geq L(f, P)$. The fact that this same inequality holds if the upper and lower sums are computed with respect to different partitions is the content of the next two lemmas.

Definition 7.2.2. A partition Q is a *refinement* of a partition P if Q contains all of the points of P . In this case, we write $P \subseteq Q$.

Lemma 7.2.3. If $P \subseteq Q$, then $L(f, P) \leq L(f, Q)$, and $U(f, P) \geq U(f, Q)$.

Proof. Consider what happens when we refine P by adding a single point z to some subinterval $[x_{k-1}, x_k]$ of P .



Focusing on the lower sum for a moment, we have

$$\begin{aligned} m_k(x_k - x_{k-1}) &= m_k(x_k - z) + m_k(z - x_{k-1}) \\ &\leq m'_k(x_k - z) + m''_k(z - x_{k-1}), \end{aligned}$$

where

$$m'_k = \inf \{f(x) : x \in [z, x_k]\} \quad \text{and} \quad m''_k = \inf \{f(x) : x \in [x_{k-1}, z]\}$$

are each necessarily as large or larger than m_k .

By induction, we have $L(f, P) \leq L(f, Q)$, and an analogous argument holds for the upper sums. \square

Lemma 7.2.4. If P_1 and P_2 are any two partitions of $[a, b]$, then $L(f, P_1) \leq U(f, P_2)$.

Proof. Let $Q = P_1 \cup P_2$ be the so-called *common refinement* of P_1 and P_2 . Because $Q \subseteq P_1$ and $Q \subseteq P_2$, it follows that

$$L(f, P_1) \leq L(f, Q) \leq U(f, Q) \leq U(f, P_2).$$

\square

Integrability

Intuitively, it helps to visualize a particular upper sum as an overestimate for the value of the integral and a lower sum as an underestimate. As the partitions get more refined, the upper sums get potentially smaller while the lower sums get potentially larger. A function is *integrable* if the upper and lower sums “meet” at some common value in the middle.

Rather than taking a limit of these sums, we will instead make use of the Axiom of Completeness and consider the *infimum* of the upper sums and the *supremum* of the lower sums.

Definition 7.2.5. Let \mathcal{P} be the collection of all possible partitions of the interval $[a, b]$. The *upper integral* of f is defined to be

$$U(f) = \inf\{U(f, P) : P \in \mathcal{P}\}.$$

In a similar way, define the *lower integral* of f by

$$L(f) = \sup\{L(f, P) : P \in \mathcal{P}\}.$$

The following fact is not surprising.

Lemma 7.2.6. For any bounded function f on $[a, b]$, it is always the case that $U(f) \geq L(f)$.

Proof. Exercise 7.2.1. □

Definition 7.2.7 (Riemann Integrability). A bounded function f defined on the interval $[a, b]$ is *Riemann-integrable* if $U(f) = L(f)$. In this case, we define $\int_a^b f$ or $\int_a^b f(x) dx$ to be this common value; namely,

$$\int_a^b f = U(f) = L(f).$$

The modifier “Riemann” in front of “integrable” accurately suggests that there are other ways to define the integral. In fact, our work in this chapter will expose the need for a different approach, one of which is discussed in Section 8.1. In this chapter, the Riemann integral is the only method under consideration, so it will usually be convenient to drop the modifier “Riemann” and simply refer to a function as being “integrable.”

Criteria for Integrability

To summarize the situation thus far, it is always the case for a bounded function f on $[a, b]$ that

$$\sup\{L(f, P) : P \in \mathcal{P}\} = L(f) \leq U(f) = \inf\{U(f, P) : P \in \mathcal{P}\}.$$

The function f is integrable if the inequality is an equality. The major thrust of our investigation of the integral is to describe, as best we can, the class

of integrable functions. The preceding inequality reveals that integrability is really equivalent to the existence of partitions whose upper and lower sums are arbitrarily close together.

Theorem 7.2.8. *A bounded function f is integrable on $[a, b]$ if and only if, for every $\epsilon > 0$, there exists a partition P_ϵ of $[a, b]$ such that*

$$U(f, P_\epsilon) - L(f, P_\epsilon) < \epsilon.$$

Proof. Let $\epsilon > 0$. If such a partition P_ϵ exists, then

$$U(f) - L(f) \leq U(f, P_\epsilon) - L(f, P_\epsilon) < \epsilon.$$

Because ϵ is arbitrary, it must be that $U(f) = L(f)$, so f is integrable. (To be absolutely precise here, we could throw in a reference to Theorem 1.2.6.)

The proof of the converse statement is a familiar triangle inequality argument with parentheses in place of absolute value bars because, in each case, we know which quantity is larger. Because $U(f)$ is the greatest lower bound of the upper sums, we know that, given some $\epsilon > 0$, there must exist a partition P_1 such that

$$U(f, P_1) < U(f) + \frac{\epsilon}{2}.$$

Likewise, there exists a partition P_2 satisfying

$$L(f, P_2) > L(f) - \frac{\epsilon}{2}.$$

Now, let $P_\epsilon = P_1 \cup P_2$ be the common refinement. Keeping in mind that the integrability of f means $U(f) = L(f)$, we can write

$$\begin{aligned} U(f, P_\epsilon) - L(f, P_\epsilon) &\leq U(f, P_1) - L(f, P_2) \\ &= (U(f, P_1) - U(f)) + (L(f) - L(f, P_2)) \\ &< \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon. \end{aligned}$$

□

In the discussion at the beginning of this chapter, it became clear that integrability is closely tied to the concept of continuity. To make this observation more precise, let $P = \{a = x_0 < x_1 < x_2 < \cdots < x_n = b\}$ be an arbitrary partition of $[a, b]$, and define $\Delta x_k = x_k - x_{k-1}$. Then,

$$U(f, P) - L(f, P) = \sum_{k=1}^n (M_k - m_k) \Delta x_k,$$


where M_k and m_k are the supremum and infimum of the function on the interval $[x_{k-1}, x_k]$ respectively. Our ability to control the size of $U(f, P) - L(f, P)$ hinges on the differences $M_k - m_k$, which we can interpret as the variation in the range of the function over the interval $[x_{k-1}, x_k]$. Restricting the variation of f over arbitrarily small intervals in $[a, b]$ is *precisely* what it means to say that f is uniformly continuous on this set.

Theorem 7.2.9. *If f is continuous on $[a, b]$, then it is integrable.*

Proof. The first crucial observation is that because f is continuous on a compact set, it is uniformly continuous. This means that, given $\epsilon > 0$, there exists a $\delta > 0$ so that $|x - y| < \delta$ guarantees

$$|f(x) - f(y)| < \frac{\epsilon}{b - a}.$$

Now, let P be a partition of $[a, b]$ where $\Delta x_k = x_k - x_{k-1}$ is less than δ for every subinterval of P .



Given a particular subinterval $[x_{k-1}, x_k]$ of P , we know from the Extreme Value Theorem (Theorem 4.4.3) that the supremum $M_k = f(z_k)$ for some $z_k \in [x_{k-1}, x_k]$. In addition, the infimum m_k is attained at some point y_k also in the interval $[x_{k-1}, x_k]$. But this means $|z_k - y_k| < \delta$, so

$$M_k - m_k = f(z_k) - f(y_k) < \frac{\epsilon}{b - a}.$$

Finally,

$$U(f, P) - L(f, P) = \sum_{k=1}^n (M_k - m_k) \Delta x_k < \frac{\epsilon}{b - a} \sum_{k=1}^n \Delta x_k = \epsilon,$$

and f is integrable by the criterion given in Theorem 7.2.8. \square

Exercises

Exercise 7.2.1. Let f be a bounded function on $[a, b]$, and let P be an arbitrary partition of $[a, b]$. First, explain why $U(f) \geq L(f, P)$. Now, prove Lemma 7.2.6.

Exercise 7.2.2. Consider $f(x) = 2x + 1$ over the interval $[1, 3]$. Let P be the partition consisting of the points $\{1, 3/2, 2, 3\}$.

(a) Compute $L(f, P)$, $U(f, P)$, and $U(f, P) - L(f, P)$.

(b) What happens to the value of $U(f, P) - L(f, P)$ when we add the point $5/2$ to the partition?

(c) Find a partition P' of $[1, 3]$ for which $U(f, P') - L(f, P') < 2$.

Exercise 7.2.3. Show directly (without appealing to Theorem 7.2) that the constant function $f(x) = k$ is integrable over any closed interval $[a, b]$. What is $\int_a^b f$?

Exercise 7.2.4. (a) Prove that a bounded function f is integrable on $[a, b]$ if and only if there exists a sequence of partitions $(P_n)_{n=1}^{\infty}$ satisfying

$$\lim_{n \rightarrow \infty} [U(f, P_n) - L(f, P_n)] = 0.$$

(b) For each n , let P_n be the partition of $[0, 1]$ into n equal subintervals. Find formulas for $U(f, P_n)$ and $L(f, P_n)$ if $f(x) = x$. The formula $1 + 2 + 3 + \cdots + n = n(n + 1)/2$ will be useful.

(c) Use the sequential criterion for integrability from (a) to show directly that $f(x) = x$ is integrable on $[0, 1]$.

Exercise 7.2.5. Assume that, for each n , f_n is an integrable function on $[a, b]$. If $(f_n) \rightarrow f$ uniformly on $[a, b]$, prove that f is also integrable on this set. (We will see that this conclusion does not necessarily follow if the convergence is pointwise.)

Exercise 7.2.6. Let $f : [a, b] \rightarrow \mathbf{R}$ be increasing on the set $[a, b]$ (i.e., $f(x) \leq f(y)$ whenever $x < y$). Show that f is integrable on $[a, b]$.

7.3 Integrating Functions with Discontinuities

The fact that continuous functions are integrable is not so much a fortunate discovery as it is evidence for a well-designed integral. Riemann's integral is a modification of Cauchy's definition of the integral, which was crafted specifically to work on continuous functions. The interesting issue is discovering just how dependent the Riemann integral is on the continuity of the integrand.

Example 7.3.1. Consider the function

$$f(x) = \begin{cases} 1 & \text{for } x \neq 1 \\ 0 & \text{for } x = 1 \end{cases}$$

on the interval $[0, 2]$. If P is any partition of $[0, 2]$, a quick calculation reveals that $U(f, P) = 2$. The lower sum $L(f, P)$ will be less than 2 because any subinterval of P that contains $x = 1$ will contribute zero to the value of the lower sum. The way to show that f is integrable is to construct a partition that minimizes the effect of the discontinuity by embedding $x = 1$ into a very small subinterval.

Let $\epsilon > 0$, and consider the partition $P_\epsilon = \{0, 1 - \epsilon/3, 1 + \epsilon/3, 2\}$. Then,

$$\begin{aligned} L(f, P_\epsilon) &= 1 \left(1 - \frac{\epsilon}{3}\right) + 0(\epsilon) + 1 \left(1 - \frac{\epsilon}{3}\right) \\ &= 2 - \frac{2}{3}\epsilon. \end{aligned}$$

Because $U(f, P_\epsilon) = 2$, we have

$$U(f, P_\epsilon) - L(f, P_\epsilon) = \frac{2}{3}\epsilon < \epsilon.$$

We can now use Theorem 7.2.8 to conclude that f is integrable.

Although the function in Example 7.3.1 is extremely simple, the method used to show it is integrable is really the same one used to prove that any bounded function with a single discontinuity is integrable. The notation in the following proof is more cumbersome, but the essence of the argument is that the misbehavior of the function at its discontinuity is isolated inside a particularly small subinterval of the partition.

Theorem 7.3.2. *If $f : [a, b] \rightarrow \mathbf{R}$ is bounded, and f is integrable on $[c, b]$ for all $c \in (a, b)$, then f is integrable on $[a, b]$. An analogous result holds at the other endpoint.*

Proof. Let M be a bound for f so that $|f(x)| \leq M$ for all $x \in [a, b]$. If

$$P = \{a = x_0 < x_1 < x_2 < \cdots < x_n = b\}$$

is a partition of $[a, b]$, then

$$\begin{aligned} U(f, P) - L(f, P) &= \sum_{k=1}^n (M_k - m_k) \Delta x_k \\ &= (M_1 - m_1)(x_1 - a) + \sum_{k=2}^n (M_k - m_k) \Delta x_k \\ &= (M_1 - m_1)(x_1 - a) + (U(f, P_{[x_1, b]}) - L(f, P_{[x_1, b]})), \end{aligned}$$

where $P_{[x_1, b]} = \{x_1 < x_2 < \cdots < x_n = b\}$ is the partition of $[x_1, b]$ obtained by deleting a from P .

Given $\epsilon > 0$, the first step is to choose x_1 close enough to a so that

$$(M_1 - m_1)(x_1 - a) < \frac{\epsilon}{2}.$$

This is not too difficult. Because $M_1 - m_1 \leq 2M$, we can pick x_1 so that

$$x_1 - a \leq \frac{\epsilon}{4M}.$$

Now, by hypothesis, f is integrable on $[x_1, b]$ so there exists a partition P_1 of $[x_1, b]$ for which

$$U(f, P_1) - L(f, P_1) < \frac{\epsilon}{2}.$$

Finally, we let $P_2 = \{a\} \cup P_1$ be a partition of $[a, b]$, from which it follows that

$$\begin{aligned} U(f, P_2) - L(f, P_2) &\leq (2M)(x_1 - a) + (U(f, P_1) - L(f, P_1)) \\ &< \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon. \end{aligned}$$

□

Theorem 7.3.2 only allows for a discontinuity at the endpoint of an interval, but that is easily remedied. In the next section, we will prove that integrability on the intervals $[a, b]$ and $[b, d]$ is equivalent to integrability on $[a, d]$. This property, together with an induction argument, leads to the conclusion that any function with a *finite* number of discontinuities is still integrable. What if the number of discontinuities is infinite?

Example 7.3.3. Recall Dirichlet's function

$$g(x) = \begin{cases} 1 & \text{for } x \text{ rational} \\ 0 & \text{for } x \text{ irrational} \end{cases}$$

from Section 4.1. If P is some partition of $[0, 1]$, then the density of the rationals in \mathbf{R} implies that every subinterval of P will contain a point where $g(x) = 1$. It follows that $U(g, P) = 1$. On the other hand, $L(g, P) = 0$ because the irrationals are also dense in \mathbf{R} . Because this is the case for every partition P , we see that the upper integral $U(f) = 1$ and the lower integral $L(f) = 0$. The two are not equal, so we conclude that Dirichlet's function is *not* integrable.

How discontinuous can a function be before it fails to be integrable? Before jumping to the hasty (and incorrect) conclusion that the Riemann integral fails for functions with more than a finite number of discontinuities, we should realize that Dirichlet's function is discontinuous at *every* point in $[0, 1]$. It would be useful to investigate a function where the discontinuities are infinite in number but do not necessarily make up all of $[0, 1]$. Thomae's function, also defined in Section 4.1, is one such example. The discontinuous points of this function are precisely the rational numbers in $[0, 1]$. In Section 7.6, we will see that Thomae's function *is* Riemann-integrable, raising the bar for allowable discontinuous points to include potentially infinite sets.

The conclusion of this story is contained in the doctoral dissertation of Henri Lebesgue, who presented his work in 1901. Lebesgue's elegant criterion for Riemann integrability is explored in great detail in Section 7.6. For the moment, though, we will take a short detour from questions of integrability and construct a proof of the celebrated Fundamental Theorem of Calculus.

Exercises

Exercise 7.3.1. Consider the function

$$h(x) = \begin{cases} 1 & \text{for } 0 \leq x < 1 \\ 2 & \text{for } x = 1 \end{cases}$$

over the interval $[0, 1]$.

- Show that $L(f, P) = 1$ for every partition P of $[0, 1]$.
- Construct a partition P for which $U(f, P) < 1 + 1/10$.
- Given $\epsilon > 0$, construct a partition P_ϵ for which $U(f, P_\epsilon) < 1 + \epsilon$.

Exercise 7.3.2. In Example 7.3.3, we learned that Dirichlet's function $g(x)$ is not Riemann-integrable. Construct a sequence $g_n(x)$ of integrable functions with $g_n \rightarrow g$ pointwise on $[0, 1]$. This demonstrates that the pointwise limit of integrable functions need not be integrable. Compare this example to the result requested in Exercise 7.2.5.

Exercise 7.3.3. Here is an alternate explanation for why a function f on $[a, b]$ with a finite number of discontinuities is integrable. Supply the missing details.

Embed each discontinuity in a sufficiently small open interval and let O be the union of these intervals. Explain why f is uniformly continuous on $[a, b] \setminus O$, and use this to finish the argument.

Exercise 7.3.4. Assume $f : [a, b] \rightarrow \mathbf{R}$ is integrable.

(a) Show that if one value of $f(x)$ is changed at some point $x \in [a, b]$, then f is still integrable and integrates to the same value as before.

(b) Show that the observation in (a) holds if a *finite* number of values of f are changed.

(c) Find an example to show that by altering a countable number of values, f may fail to be integrable.

Exercise 7.3.5. Let

$$f(x) = \begin{cases} 1 & \text{if } x = 1/n \text{ for some } n \in \mathbf{N} \\ 0 & \text{otherwise.} \end{cases}$$

Show that f is integrable on $[0, 1]$ and compute $\int_0^1 f$.

Exercise 7.3.6. A set $A \subseteq [a, b]$ has *content zero* if for every $\epsilon > 0$ there exists a finite collection of open intervals $\{O_1, O_2, \dots, O_N\}$ that contain A in their union and whose lengths sum to ϵ or less. Using $|O_n|$ to refer to the length of each interval, we have

$$A \subseteq \bigcup_{n=1}^N O_n \quad \text{and} \quad \sum_{k=1}^N |O_n| \leq \epsilon.$$

(a) Let f be bounded on $[a, b]$. Show that if the set of discontinuous points of f has content zero, then f is integrable.

(b) Show that any finite set has content zero.

(c) Content zero sets do not have to be finite. They do not have to be countable. Show that the Cantor set C defined in Section 3.1 has content zero.

(d) Prove that

$$h(x) = \begin{cases} 1 & \text{if } x \in C \\ 0 & \text{if } x \notin C. \end{cases}$$

is integrable, and find the value of the integral.

7.4 Properties of the Integral

Before embarking on the proof of the Fundamental Theorem of Calculus, we need to verify what are probably some very familiar properties of the integral. The discussion in the previous section has already made use of the following fact.

Theorem 7.4.1. *Assume $f : [a, b] \rightarrow \mathbf{R}$ is bounded, and let $c \in (a, b)$. Then, f is integrable on $[a, b]$ if and only if f is integrable on $[a, c]$ and $[c, b]$. In this case, we have*

$$\int_a^b f = \int_a^c f + \int_c^b f.$$

Proof. If f is integrable on $[a, b]$, then for $\epsilon > 0$ there exists a partition P such that $U(f, P) - L(f, P) < \epsilon$. Because refining a partition can only potentially bring the upper and lower sums closer together, we can simply add c to P if it is not already there. Then, let $P_1 = P \cap [a, c]$ be a partition of $[a, c]$, and $P_2 = P \cap [c, b]$ be a partition of $[c, b]$. It follows that

$$U(f, P_1) - L(f, P_1) < \epsilon \quad \text{and} \quad U(f, P_2) - L(f, P_2) < \epsilon,$$

implying that f is integrable on $[a, c]$ and $[c, b]$.

Conversely, if we are given that f is integrable on the two smaller intervals $[a, c]$ and $[c, b]$, then given an $\epsilon > 0$ we can produce partitions P_1 and P_2 of $[a, c]$ and $[c, b]$, respectively, such that

$$U(f, P_1) - L(f, P_1) < \frac{\epsilon}{2} \quad \text{and} \quad U(f, P_2) - L(f, P_2) < \frac{\epsilon}{2}.$$

Letting $P = P_1 \cup P_2$ produces a partition of $[a, b]$ for which

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

Thus, f is integrable on $[a, b]$.

Continuing to let $P = P_1 \cup P_2$ as earlier, we have

$$\begin{aligned} \int_a^b f &\leq U(f, P) < L(f, P) + \epsilon \\ &= L(f, P_1) + L(f, P_2) + \epsilon \\ &\leq \int_a^c f + \int_c^b f + \epsilon, \end{aligned}$$

which implies $\int_a^b f \leq \int_a^c f + \int_c^b f$. To get the other inequality, observe that

$$\begin{aligned} \int_a^c f + \int_c^b f &\leq U(f, P_1) + U(f, P_2) \\ &< L(f, P_1) + L(f, P_2) + \epsilon \\ &= L(f, P) + \epsilon \\ &\leq \int_a^b f + \epsilon. \end{aligned}$$

Because $\epsilon > 0$ is arbitrary, we must have $\int_a^c f + \int_c^b f \leq \int_a^b f$, so

$$\int_a^c f + \int_c^b f = \int_a^b f,$$

as desired. □

The proof of Theorem 7.4.1 demonstrates some of the standard techniques involved for proving facts about the Riemann integral. Admittedly, manipulating partitions does not lend itself to a great deal of elegance. The next result catalogs the remainder of the basic properties of the integral that we will need in our upcoming arguments.

Theorem 7.4.2. *Assume f and g are integrable functions on the interval $[a, b]$.*

- (i) *The function $f + g$ is integrable on $[a, b]$ with $\int_a^b (f + g) = \int_a^b f + \int_a^b g$.*
- (ii) *For $k \in \mathbf{R}$, the function kf is integrable with $\int_a^b kf = k \int_a^b f$.*
- (iii) *If $m \leq f \leq M$, then $m(b - a) \leq \int_a^b f \leq M(b - a)$.*
- (iv) *If $f \leq g$, then $\int_a^b f \leq \int_a^b g$.*
- (v) *The function $|f|$ is integrable and $|\int_a^b f| \leq \int_a^b |f|$.*

Proof. Properties (i) and (ii) are reminiscent of the Algebraic Limit Theorem and its many descendants (Theorems 2.3.3, 2.7.1, 4.2.4, and 5.2.4). In fact, there is a way to use the Algebraic Limit Theorem for this argument as well. An immediate corollary to Theorem 7.2.8 is that a function f is integrable on $[a, b]$ if and only if there exists a sequence of partitions (P_n) satisfying

$$(1) \quad \lim_{n \rightarrow \infty} [U(f, P_n) - L(f, P_n)] = 0,$$

and in this case $\int_a^b f = \lim U(f, P_n) = \lim L(f, P_n)$. (A proof for this was requested as Exercise 7.2.4.)

To prove (ii) for the case $k \geq 0$, first verify that for any partition P we have

$$U(kf, P) = kU(f, P) \quad \text{and} \quad L(kf, P) = kL(f, P).$$

Exercise 1.3.5 is used here. Because f is integrable, there exist partitions (P_n) satisfying (1). Turning our attention to the function (kf) , we see that

$$\lim_{n \rightarrow \infty} [U(kf, P_n) - L(kf, P_n)] = \lim_{n \rightarrow \infty} k [U(f, P_n) - L(f, P_n)] = 0,$$

and the formula in (ii) follows. The case where $k < 0$ is similar except that we have

$$U(kf, P_n) = kL(f, P_n) \quad \text{and} \quad L(kf, P_n) = kU(f, P_n).$$

A proof for (i) can be constructed using similar methods and is requested in Exercise 7.4.5.

To prove (iii), observe that

$$U(f, P) \geq \int_a^b f \geq L(f, P)$$

for any partition P . Statement (iii) follows if we take P to be the trivial partition consisting of only the endpoints a and b .

For (iv), let $h = g - f \geq 0$ and use (i) and (iii).

Because $-|f| \leq f \leq |f|$, statement (v) will follow from (iv) provided that we can show that $|f|$ is actually integrable. The proof of this fact is outlined in Exercise 7.4.1. \square

To this point, the quantity $\int_a^b f$ is only defined in the case where $a < b$.

Definition 7.4.3. If f is integrable on the interval $[a, b]$, define

$$\int_b^a f = - \int_a^b f.$$

Also, define

$$\int_c^c f = 0.$$

Definition 7.4.3 is a natural convention to simplify the algebra of integrals. If f is an integrable function on some interval I , then it is straightforward to verify that the equation

$$\int_a^b f = \int_a^c f + \int_c^b f$$

from Theorem 7.4.1 remains valid for *any* three points a, b , and c chosen in any order from I .

Uniform Convergence and Integration

If (f_n) is a sequence of integrable functions on $[a, b]$, and if $f_n \rightarrow f$, then we are inevitably going to want to know whether

$$(2) \quad \int_a^b f_n \rightarrow \int_a^b f.$$

This is an archetypical instance of one of the major themes of analysis: When does a mathematical manipulation such as integration respect the limiting process?

If the convergence is pointwise, then any number of things can go wrong. It is possible for each f_n to be integrable but for the limit f not to be integrable

(Exercise 7.3.2). Even if the limit function f is integrable, equation (2) may fail to hold. As an example of this, let

$$f_n(x) = \begin{cases} n & \text{if } 0 < x < 1/n \\ 0 & \text{if } x = 0 \text{ or } x \geq 1/n. \end{cases}$$

Each f_n has two discontinuities on $[0, 1]$ and so is integrable with $\int_0^1 f_n = 1$. For each $x \in [0, 1]$, we have $\lim f_n(x) = 0$ so that $f_n \rightarrow 0$ pointwise on $[0, 1]$. But now observe that the limit function $f = 0$ certainly integrates to 0, and

$$0 \neq \lim_{n \rightarrow \infty} \int_0^1 f_n.$$

As a final remark on what can go wrong in (2), we should point out that it is possible to modify this example to produce a situation where $\lim \int_0^1 f_n$ does not even exist.

One way to resolve all of these problems is to add the assumption of uniform convergence.

Theorem 7.4.4. *Assume that $f_n \rightarrow f$ uniformly on $[a, b]$ and that each f_n is integrable. Then, f is integrable and*

$$\lim_{n \rightarrow \infty} \int_a^b f_n = \int_a^b f.$$

Proof. The proof that f is integrable was requested as Exercise 7.2.5. The remainder of this argument is asked for in Exercise 7.4.3. \square

Exercises

Exercise 7.4.1. (a) Let f be a bounded function on a set A , and set

$$M = \sup\{f(x) : x \in A\}, \quad m = \inf\{f(x) : x \in A\},$$

$$M' = \sup\{|f(x)| : x \in A\}, \quad \text{and} \quad m' = \inf\{|f(x)| : x \in A\}.$$

Show that $M - m \geq M' - m'$.

(b) Show that if f is integrable on the interval $[a, b]$, then $|f|$ is also integrable on this interval.

(c) Provide the details for the argument that in this case we have $|\int_a^b f| \leq \int_a^b |f|$.

Exercise 7.4.2. Review Definition 7.4.3. Show that if $c \leq a \leq b$ and f is integrable on the interval $[c, b]$, then it is still the case that $\int_a^b f = \int_a^c f + \int_c^b f$.

Exercise 7.4.3. Prove Theorem 7.4.4 including an argument for Exercise 7.2.5 if it is not already done.

Exercise 7.4.4. Decide which of the following conjectures is true and supply a short proof. For those that are not true, give a counterexample.

- (a) If $|f|$ is integrable on $[a, b]$ then f is also integrable on this set.
- (b) Assume g is integrable and $g \geq 0$ on $[a, b]$. If $g(x) > 0$ for an infinite number of points $x \in [a, b]$, then $\int g > 0$.
- (c) If g is continuous on $[a, b]$ and $g \geq 0$ with $g(x_0) > 0$ for at least one point $x_0 \in [a, b]$, then $\int_a^b g > 0$.
- (d) If $\int_a^b f > 0$, there is an interval $[c, d] \subseteq [a, b]$ and a $\delta > 0$ such that $f(x) \geq \delta$ for all $x \in [c, d]$.

Exercise 7.4.5. Let f and g be integrable functions on $[a, b]$.

- (a) Show that if P is any partition of $[a, b]$, then

$$U(f + g, P) \leq U(f, P) + U(g, P).$$

Provide a specific example where the inequality is strict. What does the corresponding inequality for lower sums look like?

- (b) Review the proof of Theorem 7.4.2 (ii), and provide an argument for part (i) of this theorem.

Exercise 7.4.6. Review the discussion immediately preceding Theorem 7.4.4.

- (a) Produce an example of a sequence $f_n \rightarrow 0$ pointwise on $[0, 1]$ where $\lim_{n \rightarrow \infty} \int_0^1 f_n$ does not exist.
- (b) Produce another example (if necessary) where $f_n \rightarrow 0$ and the sequence $\int_0^1 f_n$ is unbounded.
- (c) Is it possible to construct each f_n to be continuous in the examples of parts (a) and (b)?
- (d) Does it seem possible to construct the sequence (f_n) to be uniformly bounded? (Uniformly bounded means that there exists a single $M > 0$ satisfying $|f_n| \leq M$ for all $n \in \mathbf{N}$.)

Exercise 7.4.7. Assume that g_n and g are bounded integrable functions with $g_n \rightarrow g$ on $[0, 1]$. The convergence is not uniform; however, the convergence is uniform on any set of the form $[\delta, 1]$ where $0 < \delta < 1$. Show that $\lim_{n \rightarrow \infty} \int_0^1 g_n = \int_0^1 g$.

7.5 The Fundamental Theorem of Calculus

The derivative and the integral have been independently defined, each in its own rigorous mathematical terms. The definition of the derivative is motivated by the problem of finding tangent lines and is given in terms of functional limits of difference quotients. The definition of the integral grows out of the desire to describe areas under nonconstant functions and is given in terms of supremums and infimums of finite sums. The Fundamental Theorem of Calculus reveals the remarkable inverse relationship between the two processes.

The result is stated in two parts. The first is a computational statement that describes how an antiderivative can be used to evaluate an integral over a particular interval. The second statement is more theoretical in nature, expressing the fact that every continuous function is the derivative of its indefinite integral.

Theorem 7.5.1 (Fundamental Theorem of Calculus). (i) If $f : [a, b] \rightarrow \mathbf{R}$ is integrable, and $F : [a, b] \rightarrow \mathbf{R}$ satisfies $F'(x) = f(x)$ for all $x \in [a, b]$, then

$$\int_a^b f = F(b) - F(a).$$

(ii) Let $g : [a, b] \rightarrow \mathbf{R}$ be integrable, and define

$$G(x) = \int_a^x g$$

for all $x \in [a, b]$. Then, G is continuous on $[a, b]$. If g is continuous at some point $c \in [a, b]$, then G is differentiable at c and $G'(c) = g(c)$.

Proof. (i) Let P be a partition of $[a, b]$ and apply the Mean Value Theorem to F on a typical subinterval $[x_{k-1}, x_k]$ of P . This yields a point $t_k \in (x_{k-1}, x_k)$ where

$$\begin{aligned} F(x_k) - F(x_{k-1}) &= F'(t_k)(x_k - x_{k-1}) \\ &= f(t_k)(x_k - x_{k-1}). \end{aligned}$$

Now, consider the upper and lower sums $U(f, P)$ and $L(f, P)$. Because $m_k \leq f(t_k) \leq M_k$ (where m_k is the infimum on $[x_{k-1}, x_k]$ and M_k is the supremum), it follows that

$$L(f, P) \leq \sum_{k=1}^n [F(x_k) - F(x_{k-1})] \leq U(f, P).$$

But notice that the sum in the middle telescopes so that

$$\sum_{k=1}^n [F(x_k) - F(x_{k-1})] = F(b) - F(a),$$

which is *independent* of the partition P . Thus we have

$$L(f) \leq F(b) - F(a) \leq U(f).$$

Because $L(f) = U(f) = \int_a^b f$, we conclude that $\int_a^b f = F(b) - F(a)$.

(ii) To prove the second statement, take $x, y \in [a, b]$ and observe that

$$\begin{aligned} |G(x) - G(y)| &= \left| \int_a^x g - \int_a^y g \right| = \left| \int_y^x g \right| \\ &\leq \int_y^x |g| \\ &\leq M|x - y|, \end{aligned}$$

where $M > 0$ is a bound on $|g|$. This shows that G is Lipschitz and so is uniformly continuous on $[a, b]$ (Exercise 4.4.9).

Now, let's assume that g is continuous at $c \in [a, b]$. In order to show that $G'(c) = g(c)$, we rewrite the limit for $G'(c)$ as

$$\begin{aligned} \lim_{x \rightarrow c} \frac{G(x) - G(c)}{x - c} &= \lim_{x \rightarrow c} \frac{1}{x - c} \left(\int_a^x g(t) dt - \int_a^c g(t) dt \right) \\ &= \lim_{x \rightarrow c} \frac{1}{x - c} \left(\int_c^x g(t) dt \right). \end{aligned}$$

We would like to show that this limit equals $g(c)$. Thus, given an $\epsilon > 0$, we must produce a $\delta > 0$ such that if $|x - c| < \delta$ then

$$(1) \quad \left| \frac{1}{x - c} \left(\int_c^x g(t) dt \right) - g(c) \right| < \epsilon.$$

The assumption of continuity of g gives us control over the difference $|g(t) - g(c)|$. In particular, we know that there exists a $\delta > 0$ such that

$$|t - c| < \delta \text{ implies } |g(t) - g(c)| < \epsilon.$$

To take advantage of this, we cleverly write the constant $g(c)$ as

$$g(c) = \frac{1}{x - c} \int_c^x g(c) dt$$

and combine the two terms in equation (1) into a single integral. Keeping in mind that $|x - c| \geq |t - c|$, we have that for all $|x - c| < \delta$,

$$\begin{aligned} \left| \frac{1}{x - c} \left(\int_c^x g(t) dt \right) - g(c) \right| &= \left| \frac{1}{x - c} \int_c^x [g(t) - g(c)] dt \right| \\ &\leq \frac{1}{(x - c)} \int_c^x |g(t) - g(c)| dt \\ &< \frac{1}{(x - c)} \int_c^x \epsilon dt = \epsilon. \end{aligned}$$

□

Exercises

Exercise 7.5.1. We have seen that not every derivative is continuous, but explain how we at least know that every continuous function is a derivative.

Exercise 7.5.2. (a) Let $f(x) = |x|$ and define $F(x) = \int_{-1}^x f$. Find a formula for $F(x)$ for all x . Where is F continuous? Where is F differentiable? Where does $F'(x) = f(x)$?

(b) Repeat part (a) for the function

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

Exercise 7.5.3. The hypothesis in Theorem 7.5.1 (i) that $F'(x) = f(x)$ for all $x \in [a, b]$ is slightly stronger than it needs to be. Carefully read the proof and state exactly what needs to be assumed with regard to the relationship between f and F for the proof to be valid.

Exercise 7.5.4 (Natural Logarithm). Let

$$H(x) = \int_1^x \frac{1}{t} dt,$$

where we consider only $x > 0$.

(a) What is $H(1)$? Find $H'(x)$.

(b) Show that H is strictly increasing; that is, show that if $0 < x < y$, then $H(x) < H(y)$.

(c) Show that $H(cx) = H(c) + H(x)$. (Think of c as a constant and differentiate $g(x) = H(cx)$.)

Exercise 7.5.5. The Fundamental Theorem of Calculus can be used to supply a shorter argument for Theorem 6.3.1 under the additional assumption that the sequence of derivatives is continuous.

Assume $f_n \rightarrow f$ pointwise and $f'_n \rightarrow g$ uniformly on $[a, b]$. Assuming each f'_n is continuous, we can apply Theorem 7.5.1 (i) to get

$$\int_a^x f'_n = f_n(x) - f_n(a)$$

for all $x \in [a, b]$. Show that $g(x) = f'(x)$.

Exercise 7.5.6. Use part (ii) of Theorem 7.5.1 to construct another proof of part (i) of Theorem 7.5.1 using the following strategy. Given f and F as in part (i), set $G(x) = \int_a^x f$. What is the relationship between F and G ?

Exercise 7.5.7 (Average Value). If g is continuous on $[a, b]$, show that there exists a point $c \in (a, b)$ where

$$g(c) = \frac{1}{b-a} \int_a^b g.$$

Exercise 7.5.8. Given a function f on $[a, b]$, define the *total variation* of f to be

$$Vf = \sup \left\{ \sum_{k=1}^n |f(x_k) - f(x_{k-1})| \right\},$$

where the supremum is taken over all partitions P of $[a, b]$.

(a) If f is continuously differentiable (f' exists as a continuous function), use the Fundamental Theorem of Calculus to show $Vf \leq \int_a^b |f'|$.

(b) Use the Mean Value Theorem to establish the reverse inequality and conclude that $Vf = \int_a^b |f'|$.

Exercise 7.5.9. Let

$$h(x) = \begin{cases} 1 & \text{if } x < 1 \text{ or } x > 1 \\ 0 & \text{if } x = 1, \end{cases}$$

and define $H(x) = \int_0^x h$. Show that even though h is not continuous at $x = 1$, $H(x)$ is still differentiable at $x = 1$.

Exercise 7.5.10. Assume f is integrable on $[a, b]$ and has a “jump discontinuity” at $c \in (a, b)$. This means that both one-sided limits exist as x approaches c from the left and from the right, but that

$$\lim_{x \rightarrow c^-} f(x) \neq \lim_{x \rightarrow c^+} f(x).$$

(This phenomenon is discussed in more detail in Section 4.6.)

Show that $F(x) = \int_a^x f$ is not differentiable at $x = c$.

Exercise 7.5.11. The Epilogue to Chapter 5 mentions the existence of a continuous monotone function that fails to be differentiable on a dense subset of \mathbf{R} . Combine the results of Exercise 7.5.10 and Exercise 6.4.8 to show how to construct such a function.

7.6 Lebesgue's Criterion for Riemann Integrability

We now return to our investigation of the relationship between continuity and the Riemann integral. We have proved that continuous functions are integrable and that the integral also exists for functions with only a finite number of discontinuities. At the opposite end of the spectrum, we saw that Dirichlet's function, which is discontinuous at every point on $[0, 1]$, fails to be Riemann-integrable. The next examples show that the set of discontinuities of an integrable function can be infinite and even uncountable.

Riemann-integrable Functions with Infinite Discontinuities

Recall from Section 4.1 that Thomae's function

$$t(x) = \begin{cases} 1 & \text{if } x = 0 \\ 1/n & \text{if } x = m/n \in \mathbf{Q} \setminus \{0\} \text{ is in lowest terms with } n > 0 \\ 0 & \text{if } x \notin \mathbf{Q} \end{cases}$$

is continuous on the set of irrationals and has discontinuities at every rational point. Let's prove that Thomae's function is integrable on $[0, 1]$ with $\int_0^1 t = 0$.

Let $\epsilon > 0$. The strategy, as usual, is to construct a partition P_ϵ of $[0, 1]$ for which $U(t, P_\epsilon) - L(t, P_\epsilon) < \epsilon$.

- Exercise 7.6.1.** a) First, argue that $L(t, P) = 0$ for any partition P of $[0, 1]$.
 b) Consider the set of points $D_{\epsilon/2} = \{x : t(x) \geq \epsilon/2\}$. How big is $D_{\epsilon/2}$?
 c) To complete the argument, explain how to construct a partition P_ϵ of $[0, 1]$ so that $U(t, P_\epsilon) < \epsilon$.

We first met the Cantor set C in Section 3.1. We have since learned that C is a compact, uncountable subset of the interval $[0, 1]$. The request of Exercise 4.3.12 is to prove that the function

$$g(x) = \begin{cases} 1 & \text{if } x \in C \\ 0 & \text{if } x \notin C \end{cases}$$

is continuous at every point of the complement of C and has discontinuities at each point of C . Thus, g is not continuous on an uncountably infinite set.

Exercise 7.6.2. Using the fact that $C = \bigcap_{n=0}^{\infty} C_n$, where each C_n consists of a finite union of closed intervals, argue that g is Riemann-integrable on $[0, 1]$.

Sets of Measure Zero

Thomae's function fails to be continuous at each rational number in $[0, 1]$. Although this set is infinite, we have seen that any subset of \mathbf{Q} is countable. Countably infinite sets are the smallest type of infinite set. The Cantor set is uncountable, but it is also small in a sense that we are now ready to make precise. In the introduction to Chapter 3, we presented an argument that the Cantor set has zero "length." The term "length" is awkward here because it really should only be applied to intervals or unions of intervals, which the Cantor set is not. There is a generalization of the concept of length to more general sets called the *measure* of a set. Of interest to our discussion are subsets that have *measure zero*.

Definition 7.6.1. A set $A \subseteq \mathbf{R}$ has *measure zero* if, for all $\epsilon > 0$, there exists a countable collection of open intervals O_n with the property that A is contained in the union of all of the intervals O_n and the sum of the lengths of all of the intervals is less than or equal to ϵ . More precisely, if $|O_n|$ refers to the length of the interval O_n , then we have

$$A \subseteq \bigcup_{n=1}^{\infty} O_n \quad \text{and} \quad \sum_{n=1}^{\infty} |O_n| \leq \epsilon.$$

Example 7.6.2. Consider a finite set $A = \{a_1, a_2, \dots, a_N\}$. To show that A has measure zero, let $\epsilon > 0$ be arbitrary. For each $1 \leq n \leq N$, construct the interval

$$G_n = \left(a_n - \frac{\epsilon}{2N}, a_n + \frac{\epsilon}{2N} \right).$$

Clearly, A is contained in the union of these intervals, and

$$\sum_{n=1}^N |G_n| = \sum_{n=1}^N \frac{\epsilon}{N} = \epsilon.$$

Exercise 7.6.3. Show that any countable set has measure zero.

Exercise 7.6.4. Prove that the Cantor set (which is uncountable) has measure zero.

Exercise 7.6.5. Show that if two sets A and B each have measure zero, then $A \cup B$ has measure zero as well. In addition, discuss the proof of the stronger statement that the countable union of sets of measure zero also has measure zero. (This second statement is true, but a completely rigorous proof requires a result about double summations discussed in Section 2.8.)

α -Continuity

Definition 7.6.3. Let f be defined on $[a, b]$, and let $\alpha > 0$. The function f is α -continuous at $x \in [a, b]$ if there exists $\delta > 0$ such that for all $y, z \in (x - \delta, x + \delta)$ it follows that $|f(y) - f(z)| < \alpha$.

Let f be a bounded function on $[a, b]$. For each $\alpha > 0$, define D_α to be the set of points in $[a, b]$ where the function f fails to be α -continuous; that is,

$$(1) \quad D_\alpha = \{x \in [a, b] : f \text{ is not } \alpha\text{-continuous at } x\}.$$

The concept of α -continuity was previously introduced in Section 4.6. Several of the ensuing exercises appeared as exercises in this section as well.

Exercise 7.6.6. If $\alpha_1 < \alpha_2$, show that $D_{\alpha_2} \subseteq D_{\alpha_1}$.

Now, let

$$(2) \quad D = \{x \in [a, b] : f \text{ is not continuous at } x\}.$$

Exercise 7.6.7. (a) Let $\alpha > 0$ be given. Show that if f is continuous at $x \in [a, b]$, then it is α -continuous at x as well. Explain how it follows that $D_\alpha \subseteq D$.

(b) Show that if f is not continuous at x , then f is not α -continuous for some $\alpha > 0$. Now, explain why this guarantees that

$$D = \bigcup_{n=1}^{\infty} D_{1/n}.$$

Exercise 7.6.8. Prove that for a fixed $\alpha > 0$, the set D_α is closed.

Exercise 7.6.9. By imitating the proof of Theorem 4.4.8, show that if, for some fixed $\alpha > 0$, f is α -continuous at every point on some compact set K , then f is *uniformly* α -continuous on K . By uniformly α -continuous, we mean that there exists a $\delta > 0$ such that whenever x and y are points in K satisfying $|x - y| < \delta$, it follows that $|f(x) - f(y)| < \alpha$.

Compactness Revisited

Compactness of subsets of the real line can be described in three equivalent ways. The following theorem appears toward the end of Section 3.3.

Theorem 7.6.4. *Let $K \subseteq \mathbf{R}$. The following three statements are all equivalent, in the sense that if any one is true, then so are the two others.*

- (i) *Every sequence contained in K has a convergent subsequence that converges to a limit in K .*
- (ii) *K is closed and bounded.*
- (iii) *Given a collection of open intervals $\{G_\alpha : \alpha \in \Lambda\}$ that covers K ; that is, $K \subseteq \bigcup_{\alpha \in \Lambda} G_\alpha$, there exists a finite subcollection $\{G_{\alpha_1}, G_{\alpha_2}, G_{\alpha_3}, \dots, G_{\alpha_N}\}$ of the original set that also covers K .*

The equivalence of (i) and (ii) has been used throughout the core material in the text. Characterization (iii) has been less central but is essential to the upcoming argument. So that the material in this section is self-contained, we quickly outline a proof that (i) and (ii) imply (iii). (This also appears as Exercise 3.3.8.)

Proof. Assume K satisfies (i) and (ii), and let $\{G_\alpha : \alpha \in \Lambda\}$ be an open cover of K . For contradiction, let's assume that no finite subcover exists.

Let I_0 be a closed interval containing K , and then bisect I_0 into two closed intervals A_1 and B_1 . It must be that either $A_1 \cap K$ or $B_1 \cap K$ (or both) has no finite subcover consisting of sets from $\{G_\alpha : \alpha \in \Lambda\}$. Let I_1 be a half of I_0 containing a part of K that cannot be finitely covered. Repeating this construction results in a nested sequence of closed intervals $I_0 \supseteq I_1 \supseteq I_2 \supseteq \dots$ with the property that, for any n , $I_n \cap K$ cannot be finitely covered and $\lim_n |I_n| = 0$.

Exercise 7.6.10. (a) Show that there exists an $x \in K$ such that $x \in I_n$ for all n .

(b) Because $x \in K$, there must exist an open set G_{α_0} from the original collection that contains x as an element. Explain why this furnishes us with the desired contradiction.

□

Lebesgue's Theorem

We are now prepared to completely categorize the collection of Riemann-integrable functions in terms of continuity.

Theorem 7.6.5 (Lebesgue's Theorem). *Let f be a bounded function defined on the interval $[a, b]$. Then, f is Riemann-integrable if and only if the set of points where f is not continuous has measure zero.*

Proof. Let $M > 0$ satisfy $|f(x)| \leq M$ for all $x \in [a, b]$, and let D and D_α be defined as in the preceding equations (1) and (2). Let's first assume that D has measure zero and prove that our function is integrable.

(\Leftarrow) Set

$$\alpha = \frac{\epsilon}{2(b-a)}.$$

Exercise 7.6.11. Show that there exists a *finite* collection of disjoint open intervals $\{G_1, G_2, \dots, G_N\}$ whose union contains D_α and that satisfies

$$\sum_{n=1}^N |G_n| < \frac{\epsilon}{4M}.$$

Exercise 7.6.12. Let K be what remains of the interval $[a, b]$ after the open intervals G_n are all removed; that is, $K = [a, b] \setminus \bigcup_{n=1}^N G_n$. Argue that f is uniformly α -continuous on K .

Exercise 7.6.13. Finish the proof in this direction by explaining how to construct a partition P_ϵ of $[a, b]$ such that $U(f, P_\epsilon) - L(f, P_\epsilon) \leq \epsilon$. It will be helpful to break the sum

$$U(f, P_\epsilon) - L(f, P_\epsilon) = \sum_{k=1}^n (M_k - m_k) \Delta x_k$$

into two parts, one over those subintervals that contain points of D_α and the other over subintervals that do not.

(\Rightarrow) For the other direction, assume f is Riemann-integrable. We must argue that the set D of discontinuities of f has measure zero.

Fix $\alpha > 0$, and let $\epsilon > 0$ be arbitrary. Because f is Riemann-integrable, there exists a partition P_ϵ of $[a, b]$ such that $U(f, P_\epsilon) - L(f, P_\epsilon) < \alpha\epsilon$.

Exercise 7.6.14. (a) Use the subintervals of the partition P_ϵ to prove that D_α has measure zero. Point out that it is possible to choose a cover for D_α that consists of a *finite* number of open intervals. (Sets for which this is possible are sometimes called *content zero*. See Exercise 7.3.6.)

(b) Show how this implies that D has measure zero.

□

A Nonintegrable Derivative

To this point, our one example of a nonintegrable function is Dirichlet's nowhere-continuous function. We close this section with another example that has special significance. The content of the Fundamental Theorem of Calculus is that integration and differentiation are inverse processes of each other. This led us to ask (in the final paragraph of the discussion in Section 7.1) whether we could integrate every derivative. For the Riemann integral, the answer is a resounding




Figure 7.3: A PRELIMINARY SKETCH OF $f_1(x)$.

no. What follows is the construction of a differentiable function whose derivative cannot be integrated with the Riemann integral.

We will once again be interested in the Cantor set

$$C = \bigcap_{n=0}^{\infty} C_n,$$

defined in Section 3.1. As an initial step, let's create a function $f(x)$ that is differentiable on $[0, 1]$ and whose derivative $f'(x)$ has discontinuities at every point of C . The key ingredient for this construction is the function

$$g(x) = \begin{cases} x^2 \sin(1/x) & \text{if } x > 0 \\ 0 & \text{if } x \leq 0. \end{cases}$$

Exercise 7.6.15. (a) Find $g'(0)$.

(b) Use the standard rules of differentiation to compute $g'(x)$ for $x \neq 0$.


(c) Explain why, for every $\delta > 0$, $g'(x)$ attains every value between 1 and -1 as x ranges over the set $(-\delta, \delta)$. Conclude that g' is not continuous at $x = 0$.

Now, we want to transport the behavior of g around zero to each of the endpoints of the closed intervals that make up the sets C_n used in the definition of the Cantor set. The formulas are awkward but the basic idea is straightforward. Start by setting

$$f_0(x) = 0 \quad \text{on} \quad C_0 = [0, 1].$$

To define f_1 on $[0, 1]$, first assign

$$f_1(x) = 0 \quad \text{for all} \quad x \in C_1 = \left[0, \frac{1}{3}\right] \cup \left[\frac{2}{3}, 1\right].$$

Figure 7.4: A GRAPH OF $f_2(x)$.

In the remaining open middle third, put translated “copies” of g oscillating toward the two endpoints (Fig. 7.3). In terms of a formula, we have

$$f_1(x) = \begin{cases} 0 & \text{if } x \in [0, 1/3] \\ g(x - 1/3) & \text{if } x \text{ is just to the right of } 1/3 \\ g(-x + 1/3) & \text{if } x \text{ is just to the left of } 2/3 \\ 0 & \text{if } x \in [2/3, 1]. \end{cases}$$

Finally, we splice the two oscillating pieces of f_1 together in such a way that makes f_1 differentiable. This is no great feat, and we will skip the details so as to keep our attention focused on the two endpoints $1/3$ and $2/3$. These are the points where $f_1'(x)$ fails to be continuous.

To define $f_2(x)$, we start with $f_1(x)$ and do the same trick as before, this time in the two open intervals $(1/9, 2/9)$ and $(7/9, 8/9)$. The result (Fig. 7.4) is a differentiable function that is zero on C_2 and has a derivative that is not continuous on the set

$$\{1/9, 2/9, 1/3, 2/3, 7/9, 8/9\}.$$

Continuing in this fashion yields a sequence of functions f_0, f_1, f_2, \dots defined on $[0, 1]$.

Exercise 7.6.16. (a) If $c \in C$, what is $\lim_{n \rightarrow \infty} f_n(c)$?

(b) Why does $\lim_{n \rightarrow \infty} f_n(x)$ exist for $x \notin C$?

Now, set

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

Exercise 7.6.17. (a) Explain why $f'(x)$ exists for all $x \notin C$.

(b) If $c \in C$, argue that $|f(x)| \leq (x - c)^2$ for all $x \in [0, 1]$. Show how this implies $f'(c) = 0$.

(c) Give a careful argument for why $f'(x)$ fails to be continuous on C . Remember that C contains many points besides the endpoints of the intervals that make up C_1, C_2, C_3, \dots .

Let's take inventory of the situation. Our goal is to create a nonintegrable derivative. Our function $f(x)$ is differentiable, and f' fails to be continuous on C . We are not quite done.

Exercise 7.6.18. Why is $f'(x)$ Riemann-integrable on $[0, 1]$?

The reason the Cantor set has measure zero is that, at each stage, 2^{n-1} open intervals of length $1/3^n$ are removed from C_{n-1} . The resulting sum

$$\sum_{n=1}^{\infty} 2^{n-1} \left(\frac{1}{3^n} \right)$$

converges to one, which means that the approximating sets C_1, C_2, C_3, \dots have total lengths tending to zero. Instead of removing open intervals of length $1/3^n$ at each stage, let's see what happens when we remove intervals of length $1/3^{n+1}$.

Exercise 7.6.19. Show that, under these circumstances, the sum of the lengths of the intervals making up each C_n no longer tends to zero as $n \rightarrow \infty$. What is this limit?

If we again take the intersection $\bigcap_{n=0}^{\infty} C_n$, the result is a Cantor-type set with the same topological properties—it is closed, compact and perfect. But a consequence of the previous exercise is that it no longer has measure zero. This is just what we need to define our desired function. By repeating the preceding construction of $f(x)$ on this new Cantor-type set of *positive* measure, we get a differentiable function whose derivative has too many points of discontinuity. By Lebesgue's Theorem, this derivative cannot be integrated using the Riemann integral.

7.7 Epilogue

Riemann's definition of the integral was a modification of Cauchy's integral, which was originally designed for the purpose of integrating continuous functions. In this goal, the Riemann integral was a complete success. For continuous functions at least, the process of integration now stood on its own rigorous footing, defined independently of differentiation. As analysis progressed, however, the dependence of integrability on continuity became problematic. The last example of Section 7.6 highlights one type of weakness: not every derivative can be integrated. Another limitation of the Riemann integral arises in association with limits of sequences of functions. To get a sense of this, let's once again consider Dirichlet's function $g(x)$ introduced in Section 4.1. Recall that $g(x) = 1$ whenever x is rational, and $g(x) = 0$ at every irrational point. Focusing on the interval $[0, 1]$ for a moment, let

$$\{r_1, r_2, r_3, r_4 \dots\}$$

be an enumeration of the countable number of rational points in this interval. Now, let $g_1(x) = 1$ if $x = r_1$ and define $g_1(x) = 0$ otherwise. Next, define $g_2(x) = 1$ if x is either r_1 or r_2 , and let $g_2(x) = 0$ at all other points. In general, for each $n \in \mathbf{N}$, define

$$g_n(x) = \begin{cases} 1 & \text{if } x \in \{r_1, r_2, \dots, r_n\} \\ 0 & \text{otherwise.} \end{cases}$$

Notice that each g_n has only a finite number of discontinuities and so is Riemann-integrable with $\int_0^1 g_n = 0$. But we also have $g_n \rightarrow g$ pointwise on the interval $[0, 1]$. The problem arises when we remember that Dirichlet's nowhere-continuous function is not Riemann-integrable. Thus, the equation

$$(1) \quad \lim_{n \rightarrow \infty} \int_0^1 g_n = \int_0^1 g$$

fails to hold, not because the values on each side of the equal sign are different but because the value on the right-hand side does not exist. The content of Theorem 7.4.4 is that this equation does hold whenever we have $g_n \rightarrow g$ *uniformly*. This is a reasonable way to resolve the situation, but it is a bit unsatisfying because the deficiency in this case is not entirely with the type of convergence but lies in the strength of the Riemann integral. If we could make sense of the right-hand side via some other definition of integration, then maybe equation (1) would actually be true.

Such a definition was introduced by Henri Lebesgue in 1901. Generally speaking, Lebesgue's integral is constructed using a generalization of length called the *measure* of a set. In the previous section, we studied sets of *measure zero*. In particular, we showed that the rational numbers in $[0, 1]$ (because they are countable) have measure zero. The irrational numbers in $[0, 1]$ have measure one. This should not be too surprising because we now have that the measures of these two disjoint sets add up to the length of the interval $[0, 1]$. Rather than chopping up the x -axis to approximate the area under the curve, Lebesgue suggested partitioning the y -axis. In the case of Dirichlet's function g , there are only two range values—zero and one. The integral, according to Lebesgue, could be defined via

$$\begin{aligned} \int_0^1 g &= 1 \cdot [\text{measure of set where } g = 1] + 0 \cdot [\text{measure of set where } g = 0] \\ &= 1 \cdot 0 + 0 \cdot 1 = 0. \end{aligned}$$

With this interpretation of $\int_0^1 g$, equation (1) is now valid!

The Lebesgue integral is presently the standard integral in advanced mathematics. The theory is taught to all graduate students, as well as to many advanced undergraduates, and it is the integral used in most research papers where integration is required. The Lebesgue integral generalizes the Riemann integral in the sense that any function that is Riemann-integrable is Lebesgue-integrable and integrates to the same value. The real strength of the Lebesgue

integral is that the class of integrable functions is much larger. Most importantly, this class includes the limits of different types of Cauchy sequences of integrable functions. This leads to a group of extremely important convergence theorems related to equation (1) with hypotheses much weaker than the uniform convergence assumed in Theorem 7.4.4.

Despite its prevalence, the Lebesgue integral does have a few drawbacks. There are functions whose *improper* Riemann integrals exist but that are not Lebesgue-integrable. Another disappointment arises from the relationship between integration and differentiation. Even with the Lebesgue integral, it is still not possible to prove

$$\int_a^b f' = f(b) - f(a)$$

without some additional assumptions on f . Around 1960, a new integral was proposed that can integrate a larger class of functions than either the Riemann integral or the Lebesgue integral and suffers from neither of the preceding weaknesses. Remarkably, this integral is actually a return to Riemann's original technique for defining integration, with some small modifications in how we describe the "finess" of the partitions. An introduction to the generalized Riemann integral is the topic of Section 8.1.

Chapter 8

Additional Topics

The foundation in analysis provided by the first seven chapters is sufficient background for the exploration of some advanced and historically important topics. The writing in this chapter is similar to that in the concluding project sections of each individual chapter. Exercises are included within the exposition and are designed to make each section a narrative investigation into a significant achievement in the field of analysis.

8.1 The Generalized Riemann Integral

Chapter 7 concluded with Henri Lebesgue's elegant result that a bounded function is Riemann-integrable if and only if its points of discontinuity form a set of measure zero. To eliminate the dependence of integrability on continuity, Lebesgue proposed a new method of integration that has become the standard integral in mathematics. In the Epilogue to Chapter 7, we briefly outlined some of the strengths and weaknesses of the Lebesgue integral, concluding with a look back to the Fundamental Theorem of Calculus (Theorem 7.5.1). (Lebesgue's measure-zero criterion is not a prerequisite for understanding the material in this section, but the discussion in Section 7.7 provides some useful context for what follows.)

If F is a differentiable function on $[a, b]$, then in a perfect world we might hope to prove that

$$(1) \quad \int_a^b F' = F(b) - F(a).$$

Notice that although this is the conclusion of part (i) of Theorem 7.5.1, there we needed the additional requirement that F' be Riemann-integrable. To drive this point home, Section 7.6 concluded with an example of a function that has a derivative that the Riemann integral cannot handle. The Lebesgue integral alluded to earlier is a significant improvement. It can integrate our example

from Section 7.6, but ultimately it too suffers from the same setback. Not every derivative is integrable, no matter which integral is used.

What follows is a short introduction to the generalized Riemann integral, discovered independently around 1960 by Jaroslav Kurzweil and Ralph Henstock. As mentioned in Section 7.7, this lesser-known modification of the Riemann integral can actually integrate a larger class of functions than Lebesgue's ubiquitous integral and yields a surprisingly simple proof of equation (1) above with no additional hypotheses.

The Riemann Integral as a Limit

Let

$$P = \{a = x_0 < x_1 < x_2 < \cdots < x_n = b\}$$

be a partition of $[a, b]$. A *tagged partition* is one where in addition to P we have chosen points c_k in each of the subintervals $[x_{k-1}, x_k]$. This sets the stage for the concept of a Riemann sum. Given a function $f : [a, b] \rightarrow \mathbf{R}$, and a tagged partition $(P, \{c_k\}_{k=1}^n)$, the *Riemann sum* generated by this partition is given by

$$R(f, P) = \sum_{k=1}^n f(c_k)(x_k - x_{k-1}).$$

Looking back at the definition of the upper sum

$$U(f, P) = \sum_{k=1}^n M_k(x_k - x_{k-1}) \quad \text{where} \quad M_k = \sup\{f(x) : x \in [x_{k-1}, x_k]\},$$

and the lower sum

$$L(f, P) = \sum_{k=1}^n m_k(x_k - x_{k-1}) \quad \text{where} \quad m_k = \inf\{f(x) : x \in [x_{k-1}, x_k]\},$$

it should be clear that

$$L(f, P) \leq R(f, P) \leq U(f, P)$$

for any bounded function f . In Definition 7.2.7, we characterized integrability by insisting that the infimum of the upper sums equal the supremum of the lower sums. Any Riemann sum is going to fall between a particular upper and lower sum. If the upper and lower sums are converging on some common value, then the Riemann sums are also eventually close to this value as well. The next theorem shows that it is possible to characterize Riemann integrability in a way equivalent to Definition 7.2.7 using an ϵ - δ -type definition applied to Riemann sums.

Definition 8.1.1. Let $\delta > 0$. A partition P is δ -*fine* if every subinterval $[x_{k-1}, x_k]$ satisfies $x_k - x_{k-1} < \delta$. In other words, every subinterval has width less than δ .

Theorem 8.1.2. A bounded function $f : [a, b] \rightarrow \mathbf{R}$ is Riemann-integrable with

$$\int_a^b f = A$$

if and only if, for every $\epsilon > 0$, there exists a $\delta > 0$ such that, for any tagged partition $(P, \{c_k\})$ that is δ -fine, it follows that

$$|R(f, P) - A| < \epsilon.$$

Before attempting the proof, we should point out that, in some treatments, the criterion in Theorem 8.1.2 is actually taken as the *definition* of Riemann integrability. In fact, this is how Riemann originally defined the concept. The spirit of this theorem is close to what is taught in most introductory calculus courses. To approximate the area under the curve, Riemann sums are constructed. The hope is that as the partitions become finer, the corresponding approximations get closer to the value of the integral. The content of Theorem 8.1.2 is that if the function is integrable, then these approximations do indeed converge to the value of the integral, regardless of how the tags are chosen. Conversely, if the approximating Riemann sums for finer and finer partitions collect around some value A , then the function is integrable and integrates to A .

Proof. (\Rightarrow) For the forward direction, we begin with the assumption that f is integrable on $[a, b]$. Given an $\epsilon > 0$, we must produce a $\delta > 0$ such that if $(P, \{c_k\})$ is any tagged partition that is δ -fine, then $|R(f, P) - \int_a^b f| < \epsilon$.

Because f is integrable, we know there exists a partition P_ϵ such that

$$U(f, P_\epsilon) - L(f, P_\epsilon) < \frac{\epsilon}{3}.$$

Let $M > 0$ be a bound on $|f|$, and let n be the number of subintervals of P_ϵ (so that P_ϵ really consists of $n + 1$ points in $[a, b]$). We will argue that choosing

$$\delta = \epsilon/9nM$$

has the desired property.

Here is the idea. Let $(P, \{c_k\})$ be an arbitrary tagged partition of $[a, b]$ that is δ -fine, and let $P' = P \cup P_\epsilon$. The key is to establish the string of inequalities

$$L(f, P') - \frac{\epsilon}{3} < L(f, P) \leq U(f, P) < U(f, P') + \frac{\epsilon}{3}.$$

Exercise 8.1.1. (a) Explain why both the Riemann sum $R(f, P)$ and $\int_a^b f$ fall between $L(f, P)$ and $U(f, P)$.

(b) Explain why $U(f, P') - L(f, P') < \epsilon/3$.

By the previous exercise, if we can show $U(f, P) < U(f, P') + \epsilon/3$ (and similarly $L(f, P') - \epsilon/3 < L(f, P)$), then it will follow that

$$\left| R(f, P) - \int_a^b f \right| < \epsilon$$

and the proof will be done. Thus, we turn our attention toward estimating the distance between $U(f, P)$ and $U(f, P')$.

Exercise 8.1.2. Explain why $U(f, P) - U(f, P') \geq 0$.

A typical term in either $U(f, P)$ or $U(f, P')$ has the form $M_k(x_k - x_{k-1})$, where M_k is the supremum of f over $[x_{k-1}, x_k]$. A good number of these terms appear in both upper sums and so cancel out.

Exercise 8.1.3. (a) In terms of n , what is the largest number of terms of the form $M_k(x_k - x_{k-1})$ that could appear in one of $U(f, P)$ or $U(f, P')$ but not the other?

(b) Finish the proof in this direction by arguing that

$$U(f, P) - U(f, P') < \epsilon/3.$$

(\Leftarrow) For this direction, we assume that the ϵ - δ criterion in Theorem 8.1.2 holds and argue that f is integrable. Integrability, as we have defined it, depends on our ability to choose partitions for which the upper sums are close to the lower sums. We have remarked that given any partition P , it is always the case that

$$L(f, P) \leq R(f, P) \leq U(f, P)$$

no matter which tags are chosen to compute $R(f, P)$.

Exercise 8.1.4. (a) Show that if f is continuous, then it is possible to pick tags $\{c_k\}_{k=1}^n$ so that

$$R(f, P) = U(f, P).$$

Similarly, there are tags for which $R(f, P) = L(f, P)$ as well.

(b) If f is not continuous, it may not be possible to find tags for which $R(f, P) = U(f, P)$. Show, however, that given an arbitrary $\epsilon > 0$, it is possible to pick tags for P so that

$$U(f, P) - R(f, P) < \epsilon.$$

The analogous statement holds for lower sums.

Exercise 8.1.5. Use the results of the previous exercise to finish the proof of Theorem 8.1.2. It may be easier to first argue that f is integrable using the criterion in Theorem 7.2.8 and then argue that $\int_a^b f = A$.

□

Gauges and $\delta(x)$ -fine Partitions

The key to the generalized Riemann integral is to allow the δ in Theorem 8.1.2 to be a *function of x* .

Definition 8.1.3. A function $\delta : [a, b] \rightarrow \mathbf{R}$ is called a *gauge* on $[a, b]$ if $\delta(x) > 0$ for all $x \in [a, b]$.

Definition 8.1.4. Given a particular gauge $\delta(x)$, a tagged partition $(P, \{c_k\}_{k=1}^n)$ is $\delta(x)$ -fine if every subinterval $[x_{k-1}, x_k]$ satisfies $x_k - x_{k-1} < \delta(c_k)$. In other words, each subinterval $[x_{k-1}, x_k]$ has width less than $\delta(c_k)$.

It is important to see that if $\delta(x)$ is a constant function, then Definition 8.1.4 says precisely the same thing as Definition 8.1.1. In the case where $\delta(x)$ is not a constant, Definition 8.1.4 describes a way of measuring the fineness of partitions that is quite different.

Exercise 8.1.6. Consider the interval $[0, 1]$.

(a) If $\delta(x) = 1/9$, find a $\delta(x)$ -fine tagged partition of $[0, 1]$. Does the choice of tags matter in this case?

(b) Let

$$\delta(x) = \begin{cases} 1/4 & \text{if } x = 0 \\ x/3 & \text{if } 0 < x \leq 1. \end{cases}$$

Construct a $\delta(x)$ -fine tagged partition of $[0, 1]$.

The tinkering required in Exercise 8.1.6 (b) may cast doubt on whether an arbitrary gauge always admits a $\delta(x)$ -fine partition. However, it is not too difficult to show that this is indeed the case.

Theorem 8.1.5. *Given a gauge $\delta(x)$ on an interval $[a, b]$, there exists a tagged partition $(P, \{c_k\}_{k=1}^n)$ that is $\delta(x)$ -fine.*

Proof. Let $I_0 = [a, b]$. It may be possible to find a tag such that the trivial partition $P = \{a = x_0 < x_1 = b\}$ works. Specifically, if $b - a < \delta(x)$ for some $x \in [a, b]$, then we can set c_1 equal to such an x and notice that $(P, \{c_1\})$ is $\delta(x)$ -fine. If no such x exists, then bisect $[a, b]$ into two equal halves.

Exercise 8.1.7. Apply the previous algorithm to each half and then explain why this procedure must eventually terminate after some finite number of steps.

□

Generalized Riemann Integrability

Keeping in mind that Theorem 8.1.2 offers an equivalent way to define Riemann integrability, we now propose a new method for defining the value of the integral.

Definition 8.1.6 (Generalized Riemann Integrability). A function f on $[a, b]$ has *generalized Riemann integral* A if, for every $\epsilon > 0$, there exists a gauge $\delta(x)$ on $[a, b]$ such that for each tagged partition $(P, \{c_k\}_{k=1}^n)$ that is $\delta(x)$ -fine, it is true that

$$|R(f, P) - A| < \epsilon.$$

In this case, we write $A = \int_a^b f$.

Theorem 8.1.7. *If a function has a generalized Riemann integral, then the value of the integral is unique.*

Proof. Assume that a function f has generalized Riemann integral A_1 and that it also has generalized Riemann integral A_2 . We must prove $A_1 = A_2$.

Let $\epsilon > 0$. Definition 8.1.6 assures us that there exists a gauge $\delta_1(x)$ such that

$$|R(f, P) - A_1| < \frac{\epsilon}{2}$$

for all tagged partitions that are $\delta_1(x)$ -fine. Likewise, there exists another gauge $\delta_2(x)$ such that

$$|R(f, P) - A_2| < \frac{\epsilon}{2}$$

for all $\delta_2(x)$ -fine tagged partitions.

Exercise 8.1.8. Finish the argument. □

The implications of Definition 8.1.6 on the resulting class of integrable functions are far reaching. This is somewhat surprising given that the criteria for integrability in Definition 8.1.6 and Theorem 8.1.2 differ in such a small way. One observation that should be immediately evident is the following.

Exercise 8.1.9. Explain why every function that is Riemann-integrable with $\int_a^b f = A$ must also have generalized Riemann integral A .

The converse statement is not true, and that is the important point. One example that we have of a non-Riemann-integrable function is Dirichlet's function

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

which has discontinuities at every point of \mathbf{R} .

Theorem 8.1.8. *Dirichlet's function $g(x)$ is generalized Riemann-integrable on $[0, 1]$ with $\int_0^1 g = 0$.*

Proof. Let $\epsilon > 0$. By Definition 8.1.6, we must construct a gauge $\delta(x)$ on $[0, 1]$ such that whenever $(P, \{c_k\}_{k=1}^n)$ is a $\delta(x)$ -fine tagged partition, it follows that

$$0 \leq \sum_{k=1}^n g(c_k)(x_k - x_{k-1}) < \epsilon.$$

The gauge represents a restriction on the size of $\Delta x_k = x_k - x_{k-1}$ in the sense that $\Delta x_k < \delta(c_k)$. The Riemann sum consists of products of the form $g(c_k)\Delta x_k$. Thus, for irrational tags, there is nothing to worry about because $g(c_k) = 0$ in this case. Our task is to make sure that any time a tag c_k is rational, it comes from a suitably thin subinterval.

Let $\{r_1, r_2, r_3, \dots\}$ be an enumeration of the countable set of rational numbers contained in $[0, 1]$. For each r_k , set $\delta(r_k) = \epsilon/2^{k+1}$. For x irrational, set $\delta(x) = 1$.

Exercise 8.1.10. Show that if $(P, \{c_k\}_{k=1}^n)$ is a $\delta(x)$ -fine tagged partition, then $R(f, P) < \epsilon$. Keep in mind that each rational number r_k can show up as a tag in at most two subintervals of P .

□

Dirichlet's function fails to be Riemann-integrable because, given any (un-tagged) partition, it is possible to make $R(f, P) = 1$ or $R(f, P) = 0$ by choosing the tags to be either all rational or all irrational. For the generalized Riemann integral, choosing all rational tags results in a tagged partition that is not $\delta(x)$ -fine (when $\delta(x)$ is small on rational points) and so does not have to be considered. In general, allowing for nonconstant gauges allows us to be more discriminating about which tagged partitions qualify as $\delta(x)$ -fine. The result, as we have just seen, is that it may be easier to achieve the inequality

$$|R(f, P) - A| < \epsilon$$

for the often smaller and more carefully selected set of tagged partitions that remain.

The Fundamental Theorem of Calculus

We conclude this brief introduction to the generalized Riemann integral with a proof of the Fundamental Theorem of Calculus. As was alluded to earlier, the most notable distinction between the following theorem and part (i) of Theorem 7.5.1 is that here we do not need to assume that the derivative function is integrable. Using the generalized Riemann integral, every derivative is integrable, and the integral can be evaluated using the antiderivative in the familiar way. It is also interesting to note that in Theorem 7.5.1 the Mean Value Theorem played the crucial role in the argument, but it is not needed here.

Theorem 8.1.9. *Assume $F : [a, b] \rightarrow \mathbf{R}$ is differentiable at each point in $[a, b]$ and set $f(x) = F'(x)$. Then, f has the generalized Riemann integral*

$$\int_a^b f = F(b) - F(a).$$

Proof. Let $P = \{a = x_0 < x_1 < x_2 < \cdots < x_n = b\}$ be a partition of $[a, b]$. Both this proof and the proof of Theorem 7.5.1 make use of the following fact.

Exercise 8.1.11. Show that

$$F(b) - F(a) = \sum_{k=1}^n [F(x_k) - F(x_{k-1})].$$

If $\{c_k\}_{k=1}^n$ is a set of tags for P , then we can estimate the difference between the Riemann sum $R(f, P)$ and $F(b) - F(a)$ by

$$\begin{aligned} |F(b) - F(a) - R(f, P)| &= \left| \sum_{k=1}^n [F(x_k) - F(x_{k-1}) - f(c_k)(x_k - x_{k-1})] \right| \\ &\leq \sum_{k=1}^n |F(x_k) - F(x_{k-1}) - f(c_k)(x_k - x_{k-1})|. \end{aligned}$$

Let $\epsilon > 0$. To prove the theorem, we must construct a gauge $\delta(c)$ such that

$$(2) \quad |F(b) - F(a) - R(f, P)| < \epsilon$$

for all $(P, \{c_k\})$ that are $\delta(c)$ -fine. (Using the variable c in the gauge function is more convenient than x in this case.)

Exercise 8.1.12. For each $c \in [a, b]$, explain why there exists a $\delta(c) > 0$ (a $\delta > 0$ depending on c) such that

$$\left| \frac{F(x) - F(c)}{x - c} - f(c) \right| < \epsilon$$

for all $0 < |x - c| < \delta(c)$.

This $\delta(c)$ is the desired gauge on $[a, b]$. Let $(P, \{c_k\}_{k=1}^n)$ be a $\delta(c)$ -fine partition of $[a, b]$. It just remains to show that equation (2) is satisfied for this tagged partition.

Exercise 8.1.13. (a) For a particular $c_k \in [x_{k-1}, x_k]$ of P , show that

$$|F(x_k) - F(c_k) - f(c_k)(x_k - c_k)| < \epsilon(x_k - c_k)$$

and

$$|F(c_k) - F(x_{k-1}) - f(c_k)(c_k - x_{k-1})| < \epsilon(c_k - x_{k-1}).$$

(b) Now, argue that

$$|F(x_k) - F(x_{k-1}) - f(c_k)(x_k - x_{k-1})| < \epsilon(x_k - x_{k-1}),$$

and use this fact to complete the proof of the theorem. □

If we consider the function

$$F(x) = \begin{cases} x^{(3/2)} \sin(1/x) & \text{if } x \neq 0 \\ 0 & \text{if } x = 0 \end{cases}$$

then it is not too difficult to show that F is differentiable everywhere, including $x = 0$, with

$$F'(x) = \begin{cases} (3/2)\sqrt{x} \sin(1/x) - (1/\sqrt{x}) \cos(1/x) & \text{if } x \neq 0 \\ 0 & \text{if } x = 0. \end{cases}$$

What is notable here is that the derivative is unbounded near the origin. The theory of the ordinary Riemann integral begins with the assumption that we only consider bounded functions on closed intervals, but there is no such restriction for the generalized Riemann integral. Theorem 8.1.9 proves that F' has a generalized integral. Now, *improper* Riemann integrals have been created to extend Riemann integration to some unbounded functions, but it is another interesting fact about the generalized Riemann integral that any function having an improper integral must already be integrable in the sense described in Definition 8.1.6.

As a parting gesture, let's show how Theorem 8.1.9 yields a short verification of the change-of-variables technique from calculus.

Theorem 8.1.10 (Substitution Formula). *Let $g : [a, b] \rightarrow \mathbf{R}$ be differentiable at each point of $[a, b]$, and assume F is differentiable on the set $g([a, b])$. If $f(x) = F'(x)$ for all $x \in g([a, b])$, then*

$$\int_a^b (f \circ g) \cdot g' = \int_{g(a)}^{g(b)} f.$$

Proof. The hypothesis of the theorem guarantees that the function $(F \circ g)(x)$ is differentiable for all $x \in [a, b]$.

Exercise 8.1.14. (a) Why are we sure that $(F \circ g)'(x)$ has a generalized Riemann integral?

(b) Use the chain rule (Theorem 5.2.5) and Theorem 8.1.9 to prove that

$$\int_a^b (f \circ g) \cdot g' = F(g(b)) - F(g(a)).$$

(c) Finish the proof by showing that

$$\int_{g(a)}^{g(b)} f = F(g(b)) - F(g(a)).$$

□

The impressive properties of the generalized Riemann integral do not end here. The central source for the material in this section is Robert Bartle's excellent article "Return to the Riemann Integral," which appeared in the *American Mathematical Monthly*, October, 1996. This article goes on to outline convergence theorems in the spirit of Theorem 7.4.4 for the generalized Riemann integral as well as its relationship to the theory of the Lebesgue integral. A more detailed development can be found in the recently published book *Integral: An Easy Approach after Kurzweil and Henstock* by Rudolph V'ýborn'ý and Lee Peng Yee or in a forthcoming book by Robert Bartle to be published by the American Mathematical Society.

8.2 Metric Spaces and the Baire Category Theorem

A natural question to ask is whether the theorems we have proved about sequences, series, and functions in \mathbf{R} have analogs in the plane \mathbf{R}^2 or in even higher dimensions. Looking back over the proofs, one crucial observation is that most of the arguments depend on just a few basic properties of the absolute value function. Interpreting the statement “ $|x - y|$ ” to mean the “distance from x to y in \mathbf{R} ,” our aim is to experiment with other ways of measuring “distance” on other sets such as \mathbf{R}^2 and $C[0, 1]$, the space of continuous functions on $[0, 1]$.

Definition 8.2.1. Given a set X , a function $d : X \times X \rightarrow \mathbf{R}$ is a *metric* on X if for all $x, y \in X$:

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

A *metric space* is a set X together with a metric d .

Property (iii) in the previous definition is the “triangle inequality.” The next two exercises illustrate the point that the same set X can be home to several different metrics. When referring to a metric space, we must specify the set *and* the particular distance function d .

Exercise 8.2.1. Decide which of the following are metrics on $X = \mathbf{R}^2$. For each, we let $x = (x_1, x_2)$ and $y = (y_1, y_2)$ be points in the plane.

- (a) $d(x, y) = \sqrt{(x_1 - y_1)^2 + (x_2 - y_2)^2}$.
- (b) $d(x, y) = 1$ if $x \neq y$; and $d(x, x) = 0$.
- (c) $d(x, y) = \max\{|x_1 - y_1|, |x_2 - y_2|\}$.
- (d) $d(x, y) = |x_1x_2 + y_1y_2|$.

Exercise 8.2.2. Let $C[0, 1]$ be the collection of continuous functions on the closed interval $[0, 1]$. Decide which of the following are metrics on $C[0, 1]$.

- (a) $d(f, g) = \sup\{|f(x) - g(x)| : x \in [0, 1]\}$.
- (b) $d(f, g) = |f(1) - g(1)|$.
- (c) $d(f, g) = \int_0^1 |f - g|$.

The following distance function is called the *discrete metric* and can be defined on any set X . For any $x, y \in X$, let

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

Exercise 8.2.3. Verify that the discrete metric is actually a metric.

Basic Definitions

Definition 8.2.2. Let (X, d) be a metric space. A sequence $(x_n) \subseteq X$ converges to an element $x \in X$ if for all $\epsilon > 0$ there exists an $N \in \mathbf{N}$ such that $d(x_n, x) < \epsilon$ whenever $n \geq N$.

Definition 8.2.3. A sequence (x_n) in a metric space (X, d) is a *Cauchy sequence* if for all $\epsilon > 0$ there exists an $N \in \mathbf{N}$ such that $d(x_m, x_n) < \epsilon$ whenever $m, n \geq N$.

Exercise 8.2.4. Show that a convergent sequence is Cauchy.

The Cauchy Criterion, as it is called in \mathbf{R} , was an “if and only if” statement. In the general metric space setting, however, the converse statement does not always hold. Recall that, in \mathbf{R} , the assertion that “Cauchy sequences converge” was shown to be equivalent to the Axiom of Completeness. In order to transport the Axiom of Completeness into a metric space, we would need to have an ordering on our space so that we could discuss such things as upper bounds. It is an interesting observation that not every set can be ordered in a satisfying way (the points in \mathbf{R}^2 for example). Even without an ordering, we are still going to want completeness. For metric spaces, the convergence of Cauchy sequences is taken to be the definition of completeness.

Definition 8.2.4. A metric space (X, d) is *complete* if every Cauchy sequence in X converges to an element of X .

Exercise 8.2.5. (a) Consider \mathbf{R}^2 with the metric defined in Exercise 8.2.1 (b). What do Cauchy sequences look like in this space? Is \mathbf{R}^2 complete with respect to this metric?

(b) Show that $C[0, 1]$ is complete with respect to the metric in Exercise 8.2.2 (a).

(c) Define $C^1[0, 1]$ to be the collection of differentiable functions on $[0, 1]$ whose derivatives are also continuous. Is $C^1[0, 1]$ complete with respect to the metric defined in Exercise 8.2.2 (a)?

(d) What does a convergent sequence in \mathbf{R} look like when we consider the discrete metric $\rho(x, y)$ examined in Exercise 8.2.3?

The metric on $C[0, 1]$ in Exercise 8.2.2 (a) is important enough to have earned the nickname “sup norm” and is denoted by

$$d(f, g) = \|f - g\|_\infty = \sup\{|f(x) - g(x)| : x \in [0, 1]\}.$$

In all upcoming discussions, it is assumed that the space $C[0, 1]$ is endowed with this metric unless otherwise specified.

Definition 8.2.5. Let (X, d) be a metric space. A function $f : X \rightarrow \mathbf{R}$ is *continuous at a point* $x \in X$ if for all $\epsilon > 0$ there exists a $\delta > 0$ such that $|f(x) - f(y)| < \epsilon$ whenever $d(x, y) < \delta$.

Exercise 8.2.6. Which of these functions on $C[0, 1]$ are continuous?

- (a) $g(f) = \int_0^1 fk$, where k is some fixed function in $C[0, 1]$.
 - (b) $g(f) = f(1/2)$.
 - (c) $g(f) = f(1/2)$, but this time with respect to the metric in Exercise 8.2.2
- (c).

Topology on Metric Spaces

Definition 8.2.6. Given $\epsilon > 0$ and an element x in the metric space (X, d) , the ϵ -neighborhood of x is the set $V_\epsilon(x) = \{y \in X : d(x, y) < \epsilon\}$.

Exercise 8.2.7. (a) Describe the ϵ -neighborhoods in \mathbf{R}^2 for each of the different metrics described in Exercise 8.2.1. How about for the discrete metric?

- (b) What do ϵ -neighborhoods in \mathbf{R} look like using the discrete metric $\rho(x, y)$?

With the definition of an ϵ -neighborhood, we can now define *open sets*, *limit points*, and *closed sets* exactly as we did before. A set $O \subseteq X$ is *open* if for every $x \in O$ we can find a neighborhood $V_\epsilon(x) \subseteq O$. A point x is a *limit point* of a set A if every $V_\epsilon(x)$ intersects A in some point other than x . A set C is *closed* if it contains its limit points.

Exercise 8.2.8. (a) Let (X, d) be a metric space, and pick $x \in X$. Verify that an ϵ -neighborhood $V_\epsilon(x)$ is an open set. Is the set

$$C_\epsilon(x) = \{y \in X : d(x, y) \leq \epsilon\}$$

a closed set?

- (b) Show that the set $Y = \{f \in C[0, 1] : \|f\|_\infty \leq 1\}$ is closed in $C[0, 1]$.
- (c) Is the set $T = \{f \in C[0, 1] : f(0) = 0\}$ open, closed, or neither in $C[0, 1]$?

We define compactness in metric spaces just as we did for \mathbf{R} .

Definition 8.2.7. A subset K of a metric space (X, d) is *compact* if every sequence in K has a convergent subsequence that converges to a limit in K .

An extremely useful characterization of compactness in \mathbf{R} is the proposition that a set is compact if and only if it is closed and bounded. For abstract metric spaces, this proposition only holds in the forward direction.

Exercise 8.2.9. (a) Supply a definition for *bounded* subsets of a metric space (X, d) .

(b) Show that if K is a compact subset of the metric space (X, d) , then K is closed and bounded.

(c) Show that $Y \subseteq C[0, 1]$ from Exercise 8.2.8 (b) is closed and bounded but not compact.

A good hint for part (c) of the previous exercise can be found in Exercise 6.2.15 from Chapter 6. This exercise defines the concept of an *equicontinuous* family of functions, which is a key ingredient in the Arzela–Ascoli Theorem (Exercise 6.2.16). The Arzela–Ascoli Theorem states that any bounded,

equicontinuous collection of functions in $C[0, 1]$ must have a uniformly convergent subsequence. One way to summarize this famous result—which we did not have the language for in Chapter 6—is as a statement describing a particular class of compact subsets in $C[0, 1]$. Looking at the definition of compactness, and remembering that the uniform limit of continuous functions is continuous, the Arzela–Ascoli Theorem states that any closed, bounded, equicontinuous collection of functions is a compact subset of $C[0, 1]$.

Definition 8.2.8. Given a subset E of a metric space (X, d) , the *closure* \overline{E} is the union of E together with its limit points. The *interior* of E is denoted by E° and is defined as

$$E^\circ = \{x \in E : \text{there exists } V_\epsilon(x) \subseteq E\}.$$

Closure and interior are dual concepts. Results about these concepts come in pairs and exhibit an elegant and useful symmetry.

Exercise 8.2.10. (a) Show that E is closed if and only if $\overline{E} = E$. Show that E is open if and only if $E^\circ = E$.

(b) Show that $\overline{E^c} = (E^\circ)^c$, and similarly that $(E^\circ)^c = \overline{E^c}$.

A good hint for the previous exercise is to review the proofs from Chapter 3, where closure at least is discussed. Thinking of all of these concepts as they relate to \mathbf{R} or \mathbf{R}^2 with the usual metric is not a bad idea. However, it is important to remember also that rigorous proofs must be constructed purely from the relevant definitions.

Exercise 8.2.11. To keep things from sounding too familiar, find an example of a metric space (from somewhere in the preceding discussion) where

$$\overline{V_\epsilon(x)} \neq \{y \in X : d(x, y) \leq \epsilon\}$$

for some ϵ -neighborhood in a metric space (X, d) .

We are on our way toward the Baire Category Theorem. The next definitions provide the final bit of vocabulary needed to state the result.

Definition 8.2.9. A set $A \subseteq X$ is *dense* in the metric space (X, d) if $\overline{A} = X$. A subset E of a metric space (X, d) is *nowhere-dense* in X if \overline{E}° is empty.

Exercise 8.2.12. If E is a subspace of a metric space (X, d) , show that E is nowhere-dense in X if and only if $\overline{E^c}$ is dense in X .

The Baire Category Theorem

In Section 3.5, we proved Baire’s Theorem, which states that it is impossible to write the real numbers \mathbf{R} as the countable union of nowhere-dense sets. Previous to this, we knew that \mathbf{R} was too big to be written as the countable union of single points (\mathbf{R} is uncountable), but Baire’s Theorem improves on this by asserting

that the only way to make \mathbf{R} from a countable union of arbitrary sets is for the closure of at least one of these sets to contain an interval. The keystone to the proof of Baire's Theorem is the completeness of \mathbf{R} . The idea now is to replace \mathbf{R} with an arbitrary complete metric space and prove the theorem in this more general setting. This leads to a statement that can be used to discuss the size and structure of other spaces such as \mathbf{R}^2 and $C[0, 1]$. At the end of Chapter 3, we mentioned one particularly fascinating implication of this result for $C[0, 1]$, which is that—despite the substantial difficulty required to produce an example of one—most continuous functions are nowhere-differentiable. It would be a good idea at this point to reread Sections 3.6 and 5.5. We are now equipped to carry out the details promised in these discussions.

Theorem 8.2.10. *Let (X, d) be a complete metric space, and let $\{O_n\}$ be a countable collection of dense, open subsets of X . Then, $\bigcap_{n=1}^{\infty} O_n$ is not empty.*

Proof. When we proved this theorem on \mathbf{R} , completeness manifested itself in the form of the Nested Interval Property. We could derive something akin to NIP in the metric space setting, but instead let's take an approach that uses the convergence of Cauchy sequences (because this is how we have defined completeness).

Pick $x_1 \in O_1$. Because O_1 is open, there exists an $\epsilon_1 > 0$ such that $V_{\epsilon_1}(x_1) \subseteq O_1$.

Exercise 8.2.13. (a) Give the details for why we know there exists a point $x_2 \in V_{\epsilon_1}(x_1) \cap O_2$ and an $\epsilon_2 > 0$ satisfying $\epsilon_2 < \epsilon_1/2$ with $V_{\epsilon_2}(x_2)$ contained in O_2 and

$$\overline{V_{\epsilon_2}(x_2)} \subseteq V_{\epsilon_1}(x_1).$$

(b) Proceed along this line and use the completeness of (X, d) to produce a single point $x \in O_n$ for every $n \in \mathbf{N}$. □

Theorem 8.2.11 (Baire Category Theorem). *A complete metric space is not the union of a countable collection of nowhere-dense sets.*

Proof. Let (X, d) be a complete metric space.

Exercise 8.2.14. If E is nowhere-dense in X , then what can we say about $(\overline{E})^c$? Now, complete the proof of the theorem. □

This result is called the Baire Category Theorem because it creates two categories of size for subsets in a metric space. A set of “first category” is one that *can* be written as a countable union of nowhere-dense sets. These are the small, intuitively thin subsets of a metric space. We now see that if our metric space is complete, then it is necessarily of “second category,” meaning it cannot be written as a countable union of nowhere-dense sets. Given a subset A of a complete metric space X , showing that A is of first category is a mathematically

precise way of proving that A constitutes a very minor portion of the set X . The term “meager” is often used to mean a set of first category.

With the stage set, we now outline the argument that continuous functions that are differentiable at even one point of $[0,1]$ form a meager subset of the metric space $C[0,1]$.

Theorem 8.2.12. *The set*

$$D = \{f \in C[0,1] : f'(x) \text{ exists for some } x \in [0,1]\}$$

is a set of first category in $C[0,1]$.

Proof. For each pair of natural numbers m, n , define

$$A_{m,n} = \{f \in C[0,1] : \text{there exists } x \in [0,1] \text{ where} \\ \left| \frac{f(x) - f(t)}{x - t} \right| \leq n \text{ whenever } 0 < |x - t| < \frac{1}{m}\}.$$

This definition takes some time to digest. Think of $1/m$ as defining a δ -neighborhood around the point x , and view n as an upper bound on the magnitude of the slopes of lines through the two points $(x, f(x))$ and $(t, f(t))$. The set $A_{m,n}$ contains any function in $C[0,1]$ for which it is possible to find at least one point x where the slopes through $(x, f(x))$ and points on the function nearby—within $1/m$ to be precise—are bounded by n .

Exercise 8.2.15. Show that if $f \in C[0,1]$ is differentiable at a point $x \in [0,1]$, then $f \in A_{m,n}$ for some pair $m, n \in \mathbf{N}$.

The collection of subsets $\{A_{m,n} : m, n \in \mathbf{N}\}$ is countable, and we have just seen that the union of these sets contains our set D . Because it is not difficult to see that a subset of a set of first category is first category, the final hurdle in the argument is to prove that each $A_{m,n}$ is nowhere-dense in $C[0,1]$.

Fix m and n . The first order of business is to prove that $A_{m,n}$ is a closed set. To this end, let (f_k) be a sequence in $A_{m,n}$ and assume $f_k \rightarrow f$ in $C[0,1]$. We need to show $f \in A_{m,n}$.

Because $f_k \in A_{m,n}$, then for each $k \in \mathbf{N}$ there exists a point $x_k \in [0,1]$ where

$$\left| \frac{f_k(x_k) - f_k(t)}{x_k - t} \right| \leq n$$

for all $0 < |x - t| < 1/m$.

Exercise 8.2.16. (a) The sequence (x_k) does not necessarily converge, but explain why there exists a subsequence (x_{k_i}) that is convergent. Let $x = \lim(x_{k_i})$.

(b) Prove that $f_{k_i}(x_{k_i}) \rightarrow f(x)$.

(c) Let $t \in [0,1]$ satisfy $0 < |x - t| < 1/m$. Show that

$$\left| \frac{f(x) - f(t)}{x - t} \right| \leq n$$

and conclude that $A_{m,n}$ is closed.

Because $A_{m,n}$ is closed, $\overline{A_{m,n}} = A_{m,n}$. In order to prove that $A_{m,n}$ is nowhere-dense, we just have to show that it contains no ϵ -neighborhoods, so pick an arbitrary $f \in A_{m,n}$, let $\epsilon > 0$, and consider the ϵ -neighborhood $V_\epsilon(f)$ in $C[0, 1]$. To show that this set is not contained in $A_{m,n}$, we must produce a function $g \in C[0, 1]$ that satisfies $\|f - g\|_\infty < \epsilon$ and has the property that there is no point $x \in [0, 1]$ where

$$\left| \frac{g(x) - g(t)}{x - t} \right| \leq n \quad \text{for all} \quad 0 < |x - t| < 1/m.$$

Exercise 8.2.17. A function is called *piecewise linear* if its graph consists of a finite number of line segments.

(a) Show that there exists a piecewise linear function $p \in C[0, 1]$ satisfying $\|f - p\|_\infty < \epsilon/2$.

(b) Show that if h is any function in $C[0, 1]$ that is bounded by 1, then the function

$$g(x) = p(x) + \frac{\epsilon}{2}h(x)$$

satisfies $g \in V_\epsilon(f)$.

(c) Construct a piecewise linear function $h(x)$ in $C[0, 1]$ that is bounded by 1 and leads to the conclusion $g \notin A_{m,n}$, where g is defined as in (b). Explain how this completes the argument for Theorem 8.2.12.

□

8.3 Fourier Series

In his famous treatise, *Theorie Analytique de la Chaleur* (The Analytical Theory of Heat), 1822, Joseph Fourier (1768–1830) boldly asserts, “Thus there is no function $f(x)$, or part of a function, which cannot be expressed by a trigonometric series.”¹

It is difficult to exaggerate the mathematical richness of this idea. It has been convincingly argued by mathematical historians that the ensuing investigation into the validity of Fourier’s conjecture was the fundamental catalyst for the pursuit of rigor that characterizes 19th century mathematics. Power series had been in wide use in the 150 years leading up to Fourier’s work, largely because they behaved so well under the operations of calculus. A function expressed as a power series is continuous, differentiable an infinite number of times, and can be integrated and differentiated as though it were a polynomial. In the presence of such agreeable behavior, there was no compelling reason for mathematicians to formulate a more precise understanding of “limit” or “convergence” because there were no arguments to resolve. Fourier’s successful implementation of trigonometric series to the study of heat flow changed all of this. To understand what the fuss was really about, we need to look more closely at what

¹Quotes in this section are taken from the article by W.A. Coppel, “J.B. Fourier—On the Occasion of his Two Hundredth Birthday,” *American Mathematical Monthly*, 76, 1969.

Fourier was asserting, focusing individually on the terms “function,” “express,” and “trigonometric series.”

Trigonometric Series

The basic principle behind any series representations is to express a given function $f(x)$ as a sum of simpler functions. For power series, the component functions are $\{1, x, x^2, x^3, \dots\}$, so that the series takes the form

$$f(x) = \sum_{n=0}^{\infty} a_n x^n = a_0 + a_1 x + a_2 x^2 + a_3 x^3 + \dots$$

A *trigonometric series* is a very different type of infinite series where the functions

$$\{1, \cos(x), \sin(x), \cos(2x), \sin(2x), \cos(3x), \sin(3x), \dots\}$$

serve as the components. Thus, a trigonometric series has the form

$$\begin{aligned} f(x) &= a_0 + a_1 \cos(x) + b_1 \sin(x) + a_2 \cos(2x) + b_2 \sin(2x) + a_3 \cos(3x) + \dots \\ &= a_0 + \sum_{n=1}^{\infty} a_n \cos(nx) + b_n \sin(nx). \end{aligned}$$

The idea of representing a function in this way was not completely new when Fourier first publicly proposed it in 1807. About 50 years earlier, Jean Le Rond d’Alembert (1717–1783) published the partial differential equation

$$(1) \quad \frac{\partial^2 u}{\partial x^2} = \frac{\partial^2 u}{\partial t^2}$$

as a means of describing the motion of a vibrating string. In this model, the function $u(x, t)$ represents the displacement of the string at time $t \geq 0$ and at some point x , which we will take to be in the interval $[0, \pi]$. Because the string is understood to be attached at each end of this interval, we have

$$(2) \quad u(0, t) = 0 \quad \text{and} \quad u(\pi, t) = 0$$

for all values of $t \geq 0$. Now, at $t = 0$, the string is displaced some initial amount, and at the moment it is released we assume

$$(3) \quad \frac{\partial u}{\partial t}(x, 0) = 0,$$

meaning that, although the string immediately starts to move, it is given no initial velocity at any point. Finding a function $u(x, t)$ that satisfies equations (1), (2), and (3) is not too difficult.

Exercise 8.3.1. (a) Verify that

$$u(x, t) = b_n \sin(nx) \cos(nt)$$

satisfies equations (1), (2), and (3) for any choice of $n \in \mathbf{N}$ and $b_n \in \mathbf{R}$. What goes wrong if $n \notin \mathbf{N}$?

(b) Explain why any finite sum of functions of the form given in part (a) would also satisfy (1), (2), and (3). (Incidentally, it is possible to hear the different solutions in (a) for values of n up to 4 or 5 by isolating the harmonics on a well-made stringed instrument.)

Now, we come to the truly interesting issue. We have just seen that any function of the form

$$(4) \quad u(x, t) = \sum_{n=1}^N b_n \sin(nx) \cos(nt)$$

solves d'Alembert's *wave equation*, as it is called, but the particular solution we want depends on how the string is originally "plucked." At time $t = 0$, we will assume that the string is given some initial displacement $f(x) = u(x, 0)$. Setting $t = 0$ in our family of solutions in (4), the hope is that the initial displacement function $f(x)$ can be expressed as

$$(5) \quad f(x) = \sum_{n=1}^N b_n \sin(nx).$$

What this means is that *if* there exist suitable coefficients b_1, b_2, \dots, b_N so that $f(x)$ can be written as a sum of sine functions as in (5), then the vibrating-string problem is completely solved by the function $u(x, t)$ given in (4). The obvious question to ask, then, is just what types of functions can be constructed as linear combinations of the functions $\{\sin(x), \sin(2x), \sin(3x), \dots\}$. How general can $f(x)$ be? Daniel Bernoulli (1700–1782) is usually credited with proposing the idea that by taking an *infinite* sum in equation (5), it may be possible to represent *any* initial position $f(x)$ over the interval $[0, \pi]$.

Fourier was studying the propagation of heat when trigonometric series resurfaced in his work in a very similar way. For Fourier, $f(x)$ represented an initial temperature applied to the boundary of some heat-conducting material. The differential equations describing heat flow are slightly different from d'Alembert's wave equation, but they still involve the second derivatives that make expressing $f(x)$ as a sum of trigonometric functions the crucial step in finding a solution.

Periodic Functions

In the early stages of his work, Fourier focused his attention on even functions (i.e., functions satisfying $f(x) = f(-x)$) and sought out ways to represent them as series of the form $\sum a_n \cos(nx)$. Eventually, he arrived at the more general formulation of the problem, which is to find suitable coefficients (a_n) and (b_n) to express a function $f(x)$ as

$$(6) \quad f(x) = a_0 + \sum_{n=1}^{\infty} a_n \cos(nx) + b_n \sin(nx).$$




Figure 8.1: $f(x) = x^2$ OVER $(-\pi, \pi]$, EXTENDED TO BE 2π PERIODIC.

As we begin to explore how arbitrary $f(x)$ can be, it is important to notice that every component of the series in equation (6) is periodic with period 2π . Turning our attention to the term “function,” it now follows that any function we hope to represent by a trigonometric series will necessarily be periodic as well. We will give primary attention to the interval $(-\pi, \pi]$. What this means is that, given a function such as $f(x) = x^2$, we will restrict our attention to f over the domain $(-\pi, \pi]$ and then extend f periodically to all of \mathbf{R} via the rule $f(x) = f(x + 2k\pi)$ for all $k \in \mathbf{Z}$ (Fig. 8.1).

This convention of focusing on just the part of $f(x)$ over the interval $(-\pi, \pi]$ hardly seems controversial, but it did generate some confusion in Fourier’s time. In Sections 1.2 and 4.1, we alluded to the fact that in the early 1800s the term “function” was used to mean something more like “formula.” It was generally believed that a function’s behavior over the interval $(-\pi, \pi]$ determined its behavior everywhere else, a point of view that follows naturally from an overly zealous faith in Taylor series. The modern definition of function given in Definition 1.2.3 is attributed to Dirichlet from the 1830s, although the idea had been suggested earlier by others. In *Theorie Analytique de la Chaleur*, Fourier clarifies his own use of the term by stating that a “function $f(x)$ represents a succession of values or ordinates, each of which is arbitrary... We do not suppose these ordinates to be subject to a common law; they succeed each other in any matter whatever, and each of them is given as if it were a single quantity.”

In the end, we will need to make a few assumptions about the nature of our functions, but the requirements we will need are quite mild, especially when compared with restrictions such as “infinitely differentiable,” which are necessary—but not sufficient—for the existence of a Taylor series representation.

Types of Convergence

This brings us to a discussion of the word “expressed.” The assumptions we must ultimately place on our function depend on the kind of convergence we aim to demonstrate. How are we to understand the equal sign in equation (6)?

Our usual course of action with infinite series is first to define the partial sum

$$(7) \quad S_N(x) = a_0 + \sum_{n=1}^N a_n \cos(nx) + b_n \sin(nx).$$

To “express $f(x)$ as a trigonometric series” then means finding coefficients $(a_n)_{n=0}^{\infty}$ and $(b_n)_{n=1}^{\infty}$ so that

$$(8) \quad f(x) = \lim_{N \rightarrow \infty} S_N(x).$$

The question remains as to what kind of limit this is. Fourier probably imagined something akin to a *pointwise* limit because the concept of uniform convergence had not yet been formulated. In addition to pointwise convergence and uniform convergence, there are still other ways to interpret the limit in equation (8). Although it won't be discussed here, it turns out that proving

$$\int_{-\pi}^{\pi} |S_N(x) - f(x)|^2 dx \rightarrow 0$$

is a natural way to understand equation (8) for a particular class of functions. This is referred to as L^2 convergence. An alternate type of convergence that we will discuss, called *Cesaro mean convergence*, relies on demonstrating that the *averages* of the partial sums converge, in our case uniformly, to $f(x)$.

Fourier Coefficients

In the discussion that follows, we are going to need a few calculus facts.

Exercise 8.3.2. Using trigonometric identities when necessary, verify the following integrals.

(a) For all $n \in \mathbf{N}$,

$$\int_{-\pi}^{\pi} \cos(nx) dx = 0 \quad \text{and} \quad \int_{-\pi}^{\pi} \sin(nx) dx = 0.$$

(b) For all $n \in \mathbf{N}$,

$$\int_{-\pi}^{\pi} \cos^2(nx) dx = \pi \quad \text{and} \quad \int_{-\pi}^{\pi} \sin^2(nx) dx = \pi.$$

(c) For all $m, n \in \mathbf{N}$,

$$\int_{-\pi}^{\pi} \cos(mx) \sin(nx) dx = 0.$$

For $m \neq n$,

$$\int_{-\pi}^{\pi} \cos(mx) \cos(nx) dx = 0 \quad \text{and} \quad \int_{-\pi}^{\pi} \sin(mx) \sin(nx) dx = 0.$$

The consequences of these results are much more interesting than their proofs. The intuition from inner-product spaces is useful. Interpreting the integral as a kind of dot product, this exercise can be summarized by saying that the functions

$$\{1, \cos(x), \sin(x), \cos(2x), \sin(2x), \cos(3x), \dots\}$$

are all *orthogonal* to each other. The content of what follows is that they in fact form a *basis* for a large class of functions.

The first order of business is to deduce some reasonable candidates for the coefficients (a_n) and (b_n) in equation (6). Given a function $f(x)$, the trick is to *assume* we are in possession of a representation described in (6) and then manipulate this equation in a way that leads to formulas for (a_n) and (b_n) . This is exactly how we proceeded with Taylor series expansions in Section 6.6. Taylor's formula for the coefficients was produced by repeatedly differentiating each side of the desired representation equation. Here, we integrate.

To compute a_0 , integrate each side of equation (6) from $-\pi$ to π , brazenly take the integral inside the infinite sum, and use Exercise 8.3.2 to get

$$\begin{aligned} \int_{-\pi}^{\pi} f(x) dx &= \int_{-\pi}^{\pi} \left[a_0 + \sum_{n=1}^{\infty} a_n \cos(nx) + b_n \sin(nx) \right] dx \\ &= \int_{-\pi}^{\pi} a_0 dx + \sum_{n=1}^{\infty} \int_{-\pi}^{\pi} [a_n \cos(nx) + b_n \sin(nx)] dx \\ &= a_0(2\pi) + \sum_{n=1}^{\infty} a_n 0 + b_n 0 = a_0(2\pi). \end{aligned}$$

Thus,

$$(9) \quad a_0 = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x) dx.$$

The switching of the sum and the integral sign in the second step of the previous calculation should rightly raise some eyebrows, but keep in mind that we are really working backward from a hypothetical representation for $f(x)$ to get a *proposal* for what a_0 should be. The point is not to justify the derivation of the formula but rather to show that using this value for a_0 ultimately gives us the representation we want. That hard work lies ahead.

Now, consider a fixed $m \geq 1$. To compute a_m , we first multiply each side of equation (6) by $\cos(mx)$ and again integrate over the interval $[-\pi, \pi]$.

Exercise 8.3.3. Derive the formulas

$$(10) \quad a_m = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) \cos(mx) dx \quad \text{and} \quad b_m = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) \sin(mx) dx$$

for all $m \geq 1$.

Let's take a short break and empirically test our recipes for (a_m) and (b_m) on a few simple functions.

Example 8.3.1. Let

$$f(x) = \begin{cases} 1 & \text{if } 0 < x < \pi \\ 0 & \text{if } x = 0 \text{ or } x = \pi \\ -1 & \text{if } -\pi < x < 0. \end{cases}$$

The fact that f is an odd function (i.e., $f(-x) = -f(x)$) means we can avoid doing any integrals for the moment and just appeal to a symmetry argument to conclude

$$a_0 = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x) dx = 0 \quad \text{and} \quad a_n = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) \cos(nx) = 0$$

for all $n \geq 1$. We can also simplify the integral for b_n by writing

$$\begin{aligned} b_n &= \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) \sin(nx) = \frac{2}{\pi} \int_0^{\pi} \sin(nx) dx \\ &= \frac{2}{\pi} \left(\frac{-1}{n} \cos(nx) \Big|_0^{\pi} \right) \\ &= \begin{cases} 4/n\pi & \text{if } n \text{ is odd} \\ 0 & \text{if } n \text{ is even.} \end{cases} \end{aligned}$$


Proceeding on blind faith, we plug these results into equation (6) to get the representation

$$f(x) = \frac{4}{\pi} \sum_{n=0}^{\infty} \frac{1}{2n+1} \sin((2n+1)x).$$

A graph of a few of the partial sums of this series (Fig. 8.2) should generate some optimism about the legitimacy of what is happening.

Exercise 8.3.4. (a) Referring to the previous example, explain why we can be sure that the convergence of the partial sums to $f(x)$ is *not* uniform on any interval containing 0. (b) Repeat the computations of Example 8.3.1 for the function $g(x) = |x|$ and examine graphs for some partial sums. This time, make use of the fact that g is even ($g(x) = g(-x)$) to simplify the calculations. By just looking at the coefficients, how do we know this series converges uniformly to something?

(c) Use graphs to collect some empirical evidence regarding the question of term-by-term differentiation in our two examples to this point. Is it possible to conclude convergence or divergence of either differentiated series by looking at the resulting coefficients? In Chapter 6, we have a theorem about the legitimacy of term-by-term differentiation. Can it be applied to either of these examples?

Figure 8.2: f , S_4 , AND S_{20} ON $[-\pi, \pi]$.

The Riemann–Lebesgue Lemma

In the examples we have seen to this point, the sequences of Fourier coefficients (a_n) and (b_n) all tend to 0 as $n \rightarrow \infty$. This is always the case. Understanding why this happens is crucial to our upcoming convergence proof.

We start with a simple observation. The reason

$$\int_{-\pi}^{\pi} \sin(x) dx = 0$$

is that the positive and negative portions of the sine curve cancel each other out. The same is true of

$$\int_{-\pi}^{\pi} \sin(nx) dx = 0.$$

Now, when n is large, the period of the oscillations of $\sin(nx)$ becomes very short— $2\pi/n$ to be precise. If $h(x)$ is a continuous function, then the values of h do not vary too much as $\sin(nx)$ ranges over each short period. The result is that the successive positive and negative oscillations of the product $h(x)\sin(nx)$ (Fig. 8.3) are nearly the same size so that the cancellation leads to a small value for

$$\int_{-\pi}^{\pi} h(x) \sin(nx) dx.$$

Theorem 8.3.2 (Riemann–Lebesgue Lemma). *Assume $h(x)$ is continuous on $(-\pi, \pi]$. Then,*

$$\int_{-\pi}^{\pi} h(x) \sin(nx) dx \rightarrow 0 \quad \text{and} \quad \int_{-\pi}^{\pi} h(x) \cos(nx) dx \rightarrow 0$$

as $n \rightarrow \infty$.




Figure 8.3: $h(x)$ AND $h(x) \sin(nx)$ FOR LARGE n .

Proof. Remember that, like all of our functions from here on, we are mentally extending h to be 2π -periodic. Thus, while our attention is generally focused on the interval $(-\pi, \pi]$, the assumption of continuity is intended to mean that the periodically extended h is continuous on all of \mathbf{R} . Note that in addition to continuity on $(-\pi, \pi]$, this amounts to insisting that $\lim_{x \rightarrow -\pi^+} h(x) = h(\pi)$.

Exercise 8.3.5. Explain why h is uniformly continuous on \mathbf{R} .

Given $\epsilon > 0$, choose $\delta > 0$ such that $|x - y| < \delta$ implies $|h(x) - h(y)| < \epsilon/2$. The period of $\sin(nx)$ is $2\pi/n$, so choose N large enough so that $2\pi/n < \delta$ whenever $n \geq N$. Now, consider a particular interval $[a, b]$ of length $2\pi/n$ over which $\sin(nx)$ moves through one complete oscillation.

Exercise 8.3.6. Show that $\int_a^b h(x) \sin(nx) dx < \epsilon/n$, and use this fact to complete the proof.

□

Applications of Fourier series are not restricted to continuous functions (Example 8.3.1). Even though our particular proof makes use of continuity, the Riemann–Lebesgue lemma holds under much weaker hypotheses. It is true, however, that any proof of this fact ultimately takes advantage of the cancellation of positive and negative components. Recall from Chapter 2 that this type of cancellation is the mechanism that distinguishes conditional convergence from absolute convergence. In the end, what we discover is that, unlike power series, Fourier series can converge conditionally. This makes them less robust, perhaps, but more versatile and capable of more interesting behavior.

A Pointwise Convergence Proof

Let's return once more to Fourier's claim that every "function" can be "expressed" as a trigonometric series. Our recipe for the Fourier coefficients in equations (9) and (10) implicitly requires that our function be integrable. This is the major motivation for Riemann's modification of Cauchy's definition of the integral. Because integrability is a prerequisite for producing a Fourier series, we would like the class of integrable functions to be as large as possible. The natural question to ask now is whether Riemann integrability is enough or whether we need to make some additional assumptions about f in order to guarantee that the Fourier series converges back to f . The answer depends on the type of convergence we hope to establish.

$$f(x) = a_0 + \sum_{n=1}^{\infty} a_n \cos(nx) + b_n \sin(nx)$$

bounded	pointwise convergence
integrable	uniform convergence
continuous	L^2 convergence
differentiable	Cesaro mean convergence
f' continuous	

There is no tidy way to summarize the situation. For pointwise convergence, integrability is not enough. At present, "integrable" for us means Riemann-integrable, which we have only rigorously defined for bounded functions. In 1966, Lennart Carleson proved (via an extremely complicated argument) that the Fourier series for such a function converges pointwise at every point in the domain excluding possibly a set of *measure zero*. This term surfaced in our discussion of the Cantor set (Section 3.1) and is defined rigorously in Section 7.6. Sets of measure zero are small in one sense, but they can be uncountable, and there are examples of continuous functions with Fourier series that diverge at uncountably many points. Lebesgue's modification of Riemann's integral in 1901 proved to be a much more natural setting for Fourier analysis. Carleson's proof is really about Lebesgue-integrable functions which are allowed to be unbounded but for which $\int_{-\pi}^{\pi} |f|^2$ is finite. One of the cleanest theorems in this area states that, for this class of square Lebesgue-integrable functions, the Fourier series always converges to the function from which it was derived if we interpret convergence in the L^2 sense described earlier. As a final warning about how fragile the situation is, there is an example due to A. Kolmogorov (1903–1987) of a Lebesgue-integrable function where the Fourier series fails to converge at any point.

Although all of these results require significantly more background to pursue in any rigorous way, we are in a position to prove some important theorems that

require a few extra assumptions about the function in question. We will content ourselves with two interesting results in this area.

Theorem 8.3.3. *Let $f(x)$ be continuous on $(-\pi, \pi]$, and let $S_N(x)$ be the N th partial sum of the Fourier series described in equation (7), where the coefficients (a_n) and (b_n) are given by equations (9) and (10). It follows that*

$$\lim_{N \rightarrow \infty} S_N(x) = f(x)$$

pointwise at any $x \in (-\pi, \pi]$ where $f'(x)$ exists.

Proof. Cataloging a few preliminary facts makes for a smoother argument.

Fact 1: (a) $\cos(\alpha - \theta) = \cos(\alpha)\cos(\theta) + \sin(\alpha)\sin(\theta)$.

(b) $\sin(\alpha + \theta) = \sin(\alpha)\cos(\theta) + \cos(\alpha)\sin(\theta)$.

Fact 2: $\frac{1}{2} + \cos(\theta) + \cos(2\theta) + \cos(3\theta) + \cdots + \cos(N\theta) = \frac{\sin((N + 1/2)\theta)}{2 \sin(\theta/2)}$ for any $\theta \neq 2n\pi$.

Facts 1(a) and 1(b) are familiar trigonometric identities. Fact 2 is not as familiar. Its proof (which we omit) is most easily derived by taking the real part of a geometric sum of complex exponentials. The function in Fact 2 is called the *Dirichlet kernel* in honor of the mathematician responsible for the first rigorous convergence proof of this kind. Integrating both sides of this identity leads to our next important fact.

Fact 3: Setting

$$D_N(\theta) = \begin{cases} \frac{\sin((N+1/2)\theta)}{2 \sin(\theta/2)}, & \text{if } \theta \neq 2n\pi \\ 1/2 + N, & \text{if } \theta = 2n\pi \end{cases}$$

from Fact 2, we see that

$$\int_{-\pi}^{\pi} D_N(\theta) d\theta = \pi.$$

Although we will not restate it, the last fact we will use is the Riemann–Lebesgue lemma.

Fix a point $x \in (-\pi, \pi]$. The first step is to simplify the expression for $S_N(x)$. Now x is a fixed constant at the moment, so we will write the integrals in equations (9) and (10) using t as the variable of integration. Keeping an eye

on Facts 1(a) and (2), we get that

$$\begin{aligned}
 S_N(x) &= a_0 + \sum_{n=1}^N a_n \cos(nx) + b_n \sin(nx) \\
 &= \left[\frac{1}{2\pi} \int_{-\pi}^{\pi} f(t) dt \right] + \sum_{n=1}^N \left[\frac{1}{\pi} \int_{-\pi}^{\pi} f(t) \cos(nt) dt \right] \cos(nx) \\
 &\quad + \sum_{n=1}^N \left[\frac{1}{\pi} \int_{-\pi}^{\pi} f(t) \sin(nt) dt \right] \sin(nx) \\
 &= \frac{1}{\pi} \int_{-\pi}^{\pi} f(t) \left[\frac{1}{2} + \sum_{n=1}^N \cos(nt) \cos(nx) + \sin(nt) \sin(nx) \right] dt \\
 &= \frac{1}{\pi} \int_{-\pi}^{\pi} f(t) \left[\frac{1}{2} + \sum_{n=1}^N \cos(nt - nx) \right] dt \\
 &= \frac{1}{\pi} \int_{-\pi}^{\pi} f(t) D_N(t - x) dt.
 \end{aligned}$$

As one final simplification, let $u = t - x$. Then,

$$S_N(x) = \frac{1}{\pi} \int_{-\pi-x}^{\pi-x} f(u+x) D_N(u) du = \frac{1}{\pi} \int_{-\pi}^{\pi} f(u+x) D_N(u) du.$$

The last equality is a result of our agreement to extend f to be 2π -periodic. Because D_N is also periodic (it is the sum of cosine functions), it does not matter over what interval we compute the integral as long as we cover exactly one full period.


To prove $S_N(x) \rightarrow f(x)$, we must show that $|S_N(x) - f(x)|$ gets arbitrarily small when N gets large. Having expressed $S_N(x)$ as an integral involving $D_N(u)$, we are motivated to do a similar thing for $f(x)$. By Fact 3,

$$f(x) = f(x) \frac{1}{\pi} \int_{-\pi}^{\pi} D_N(u) du = \frac{1}{\pi} \int_{-\pi}^{\pi} f(x) D_N(u) du,$$

and it follows that

$$(11) \quad S_N(x) - f(x) = \frac{1}{\pi} \int_{-\pi}^{\pi} (f(u+x) - f(x)) D_N(u) du.$$

Our goal is to show this quantity tends to zero as $N \rightarrow \infty$. A sketch of $D_N(u)$ (Fig. 8.4) for a few values of N reveals why this might happen. For large N , the Dirichlet kernel $D_N(u)$ has a tall, thin spike around $u = 0$, but this is precisely where $f(u+x) - f(x)$ is small (because f is continuous). Away from zero, $D_N(u)$ exhibits the fast oscillations that harken back to the Riemann–Lebesgue lemma (Theorem 8.3.2). Let's see how to use this theorem to finish the argument.

Figure 8.4: $D_6(u)$ AND $D_{16}(u)$.

Using Fact 1(b), we can rewrite the Dirichlet kernel as

$$D_N(u) = \frac{\sin((N + 1/2)u)}{2 \sin(u/2)} = \frac{1}{2} \left[\frac{\sin(Nu) \cos(u/2)}{\sin(u/2)} + \cos(Nu) \right].$$

Then, equation (11) becomes

$$\begin{aligned} S_N(x) - f(x) &= \frac{1}{2\pi} \int_{-\pi}^{\pi} (f(u+x) - f(x)) \left[\frac{\sin(Nu) \cos(u/2)}{\sin(u/2)} + \cos(Nu) \right] du \\ &= \frac{1}{2\pi} \int_{-\pi}^{\pi} (f(u+x) - f(x)) \left(\frac{\sin(Nu) \cos(u/2)}{\sin(u/2)} \right) \\ &\quad + (f(u+x) - f(x)) \cos(Nu) du \\ &= \frac{1}{2\pi} \int_{-\pi}^{\pi} p_x(u) \sin(Nu) du + \frac{1}{2\pi} \int_{-\pi}^{\pi} q_x(u) \cos(Nu) du, \end{aligned}$$

where in the last step we have set

$$p_x(u) = \frac{(f(u+x) - f(x)) \cos(u/2)}{\sin(u/2)} \quad \text{and} \quad q_x(u) = f(u+x) - f(x).$$

Exercise 8.3.7. (a) First, argue why the integral involving $q_x(u)$ tends to zero as $N \rightarrow \infty$.

(b) The first integral is a little more subtle because the function $p_x(u)$ has the $\sin(u/2)$ term in the denominator. Use the fact that f is differentiable at x (and a familiar limit from calculus) to prove that the first integral goes to zero as well.

□

This completes the argument that $S_N(x) \rightarrow f(x)$ at any point x where f is differentiable. If the derivative exists everywhere, then we obviously get

$S_N \rightarrow f$ pointwise. If we add the assumption that f' is continuous, then it is not too difficult to show that the convergence is uniform. In fact, there is a very strong relationship between the speed of convergence of the Fourier series and the smoothness of f . The more derivatives f possesses, the faster the partial sums S_N converge to f .

Cesaro Mean Convergence

Rather than pursue the proofs in this interesting direction, we will finish this very brief introduction to Fourier series with a look at a different type of convergence called Cesaro mean convergence.

Exercise 8.3.8. Prove that if a sequence of real numbers (x_n) converges, then the arithmetic means

$$y_n = \frac{x_1 + x_2 + x_3 + \cdots + x_n}{n}$$

also converge to the same limit. Give an example to show that it is possible for the sequence of means (y_n) to converge even if the original sequence (x_n) does not.

The discussion preceding Theorem 8.3.3 is intended to create a kind of reverence for the difficulties inherent in deciphering the behavior of Fourier series, especially in the case where the function in question is not differentiable. It is from this humble frame of mind that the following elegant result due to L. Fejér in 1904 can best be appreciated.

Theorem 8.3.4 (Fejér's Theorem). Let $S_N(x)$ be the N th partial sum of the Fourier series for a function f on $(-\pi, \pi]$. Define

$$\sigma_N(x) = \frac{1}{N+1} \sum_{n=0}^N S_n(x).$$

If f is continuous on $(-\pi, \pi]$, then $\sigma_N(x) \rightarrow f(x)$ uniformly.

Proof. This argument is patterned after the proof of Theorem 8.3.3 but is actually much simpler. In addition to the trigonometric formulas listed in Facts 1 and 2, we are going to need a version of Fact 2 for the sine function, which looks like

$$\sin(\theta) + \sin(2\theta) + \sin(3\theta) + \cdots + \sin(N\theta) = \frac{\sin\left(\frac{N\theta}{2}\right) \sin\left(\frac{(N+1)\theta}{2}\right)}{\sin\left(\frac{\theta}{2}\right)}.$$

Exercise 8.3.9. Use the previous identity to show that

$$\frac{1/2 + D_1(\theta) + D_2(\theta) + \cdots + D_N(\theta)}{N+1} = \frac{1}{2(N+1)} \left[\frac{\sin\left(\frac{(N+1)\theta}{2}\right)}{\sin\left(\frac{\theta}{2}\right)} \right]^2.$$

The expression in Exercise 8.3.9 is called the *Fejér kernel* and will be denoted by $F_N(\theta)$. Analogous to the Dirichlet kernel $D_N(\theta)$ from the proof of Theorem 8.3.3, F_N is used to greatly simplify the formula for $\sigma_N(x)$.

Exercise 8.3.10. (a) Show that

$$\sigma_N(x) = \frac{1}{\pi} \int_{-\pi}^{\pi} f(u+x)F_N(u) du.$$

(b) Graph the function $F_N(u)$ for several values of N . Where is F_N large, and where is it close to zero? Compare this function to the Dirichlet kernel $D_N(u)$. Now, prove that $F_N \rightarrow 0$ uniformly on any set of the form $\{u : |u| \geq \delta\}$, where $\delta > 0$ is fixed (and u is restricted to the interval $(-\pi, \pi]$).

(c) Prove that $\int_{-\pi}^{\pi} F_N(u) du = \pi$.

(d) To finish the proof of Fejér's theorem, first choose a $\delta > 0$ so that

$$|u| < \delta \quad \text{implies} \quad |f(x+u) - f(x)| < \epsilon.$$

Set up a single integral that represents the difference $\sigma_N f(x) - f(x)$ and divide this integral into sets where $|u| \leq \delta$ and $|u| \geq \delta$. Explain why it is possible to make each of these integrals sufficiently small, independently of the choice of x .

□

Weierstrass Approximation Theorem

The hard work of proving Fejér's theorem has many rewards, one of which is access to a relatively short argument for a profoundly important theorem discovered by Weierstrass in 1885.

Theorem 8.3.5 (Weierstrass Approximation Theorem). *If f is a continuous function on a closed interval $[a, b]$, then there exists a sequence of polynomials that converges uniformly to $f(x)$ on $[a, b]$.*

Proof. We have actually seen a special case of this result once before in Section 6.6 on Taylor series. Even if this material has not been covered, the formula

$$(12) \quad \sin(x) = x - \frac{x^3}{3!} + \frac{x^5}{5!} - \frac{x^7}{7!} + \frac{x^9}{9!} - \dots$$

is probably familiar from a course in calculus. The content of Section 6.6 is that, by using Lagrange's Remainder Theorem, we can prove that the Taylor series in (12) converges uniformly to $\sin(x)$ on any bounded subset of \mathbf{R} . Uniform convergence of a series means the partial sums converge uniformly, and the partial sums in this case are polynomials. Notice that this is precisely what Theorem 8.3.5 asks us to prove, only we must do it for an arbitrary, continuous function in place of $\sin(x)$.

Using Taylor series will not work in general. The major problem is that to construct a Taylor series we need the function to be infinitely differentiable, and

even in this case we might get a series that either does not converge or converges to the wrong thing. We do plan to use Taylor series, however. The important point about them for this conversation is that the Taylor series for $\sin(x)$ and $\cos(x)$ *do* converge uniformly to the proper limit on any bounded set.

Exercise 8.3.11. (a) Make use of the previous comments and use Fejér's Theorem to complete the proof of Theorem 8.3.5 under the added assumption that the interval $[a, b]$ is $[0, \pi]$.

(b) Show how the case for an arbitrary interval $[a, b]$ follows from this one.

□

It is interesting to juxtapose this result of Weierstrass with his demonstration of a continuous nowhere-differentiable function. Although there exist continuous functions that oscillate so wildly that they fail to have a derivative at any point, these unruly functions are always uniformly within ϵ of an infinitely differentiable polynomial.

Approximation as a Unifying Theme

Viewing the last section of this chapter as a kind of appendix (included to clear up some loose ends from Chapter 1 regarding the definition of the real numbers), Weierstrass' Approximation Theorem makes for a fitting close to our introductory survey of some of the gems of analysis.

The idea of approximation permeates the entire subject. Every real number can be approximated with rational ones. The value of an infinite sum is approximated with partial sums, and the value of a continuous function can be approximated with its values nearby. A function is differentiable when a straight line is a good approximation to the curve, and it is integrable when finite sums of rectangles are a good approximation to the area under the curve. Now, we learn that every continuous function can be approximated arbitrarily well with a polynomial. In every case, the approximating objects are tangible and well-understood, and the issue is how well these properties survive the limiting process. By viewing the different infinities of mathematics through pathways crafted out of finite objects, Weierstrass and the other founders of analysis created a paradigm for how to extend the scope of mathematical exploration deep into territory previously unattainable. Although our journey ends here, the road is long and continues to be written.

8.4 A Construction of \mathbf{R} From \mathbf{Q}

This entire section is devoted to constructing a proof for the following theorem:

Theorem 8.4.1. *There exists an ordered field in which every nonempty set that is bounded above has a least upper bound. In addition, this field contains \mathbf{Q} as a subfield.*

There are a few terms to define before this statement can be properly understood and proved, but it can essentially be paraphrased as “the real numbers exist.” In Section 1.1, we encountered a major failing of the rational number system as a place to do analysis. Without the square root of 2 (and uncountably many other irrational numbers) we cannot confidently move from a Cauchy sequence to its limit because in \mathbf{Q} there is no guarantee that such a number exists. (A review of Sections 1.1 and 1.3 is highly recommended at this point.) The resolution we proposed in Chapter 1 came in the form of the Axiom of Completeness, which we restate.

Axiom of Completeness. *Every nonempty set of real numbers that is bounded above has a least upper bound.*

Now let’s be clear about how we actually proceeded in Chapter 1. This is the property that distinguishes \mathbf{Q} from \mathbf{R} , but by referring to this property as an *axiom* we were making the point that it was not something to be proved. The real numbers were defined simply as an extension of the rational numbers in which bounded sets have least upper bounds, but no attempt was made to demonstrate that such an extension is actually possible. Now, the time has finally come. By explicitly building the real numbers from the rational ones, we will be able to demonstrate that the Axiom of Completeness does not need to be an axiom at all; it is a theorem!

There is something ironic about having the final section of this book be a construction of the number system that has been the underlying subject of every preceding page, but there is something perfectly apt about it as well. Through eight chapters stretching from Cantor’s Theorem to the Baire Category Theorem, we have come to see how profoundly the addition of completeness changes the landscape. We all grow up believing in the existence of real numbers, but it is only through a study of classical analysis that we become aware of their elusive and enigmatic nature. It is because completeness matters so much, and because it is responsible for such perplexing phenomena, that we should now feel obliged—compelled really—to go back to the beginning and verify that such a thing really exists.

As we mentioned in Chapter 1, proceeding in this order puts us in good historical company. The pioneering work of Cauchy, Bolzano, Abel, Dirichlet, Weierstrass, and Riemann preceded—and in a very real sense led to—the host of rigorous definitions for \mathbf{R} that were proposed in 1872. Georg Cantor is a familiar name responsible for one of these definitions, but alternate constructions of the real number system also came from Charles Meray (1835–1911), Heinrich Heine (1821–1881), and Richard Dedekind (1831–1916). The formulation that follows is the one due to Dedekind. In a sense it is the most abstract of the approaches, but it is the most appropriate for us because the verification of completeness is done in terms of least upper bounds.

Dedekind Cuts

We begin this discussion by assuming that the rational numbers and all of the familiar properties of addition, multiplication, and order are available to us. At the moment, there is no such thing as a real number.

Definition 8.4.2. A subset A of the rational numbers is called a *cut* if it possesses the following three properties:

- (c1) $A \neq \emptyset$ and $A \neq \mathbf{Q}$.
- (c2) If $r \in A$, then A also contains every rational $q < r$.
- (c3) A does not have a maximum; that is, if $r \in A$, then there exists $s \in A$ with $r < s$.

Exercise 8.4.1. (a) Fix $r \in \mathbf{Q}$. Show that the set $C_r = \{t \in \mathbf{Q} : t < r\}$ is a cut.

The temptation to think of all cuts as being of this form should be avoided. Which of the following subsets of \mathbf{Q} are cuts?

- (b) $S = \{t \in \mathbf{Q} : t \leq 2\}$
- (c) $T = \{t \in \mathbf{Q} : t^2 < 2 \text{ or } t < 0\}$
- (d) $U = \{t \in \mathbf{Q} : t^2 \leq 2 \text{ or } t < 0\}$

Exercise 8.4.2. Let A be a cut. Show that if $r \in A$ and $s \notin A$, then $r < s$.

To dispel any suspense, let's get right to the point.

Definition 8.4.3. Define the *real numbers* \mathbf{R} to be the set of all cuts in \mathbf{Q} .

This may feel awkward at first—real numbers should be numbers, not sets of rational numbers. The counterargument here is that when working on the foundations of mathematics, sets are about the most basic building blocks we have. We have defined a set \mathbf{R} whose elements are subsets of \mathbf{Q} . We now must set about the task of imposing some algebraic structure on \mathbf{R} that behaves in a way familiar to us. What exactly does this entail? If we are serious about constructing a proof for Theorem 8.4.1, we need to be more specific about what we mean by an “ordered field.”

Field and Order Properties

Given a set F and two elements $x, y \in F$, an *operation* on F is a function that takes the ordered pair (x, y) to a third element $z \in F$. Writing $x + y$ or xy to represent different operations reminds us of the two operations that we are trying to emulate.

Definition 8.4.4. A set F is a *field* if there exist two operations—addition $(x + y)$ and multiplication (xy) —that satisfy the following list of conditions:

- (f1) (commutativity) $x + y = y + x$ and $xy = yx$ for all $x, y \in F$.

- (f2) (associativity) $(x+y)+z = x+(y+z)$ and $(xy)z = x(yz)$ for all $x, y, z \in F$.
- (f3) (identities exist) There exist two special elements 0 and 1 with $0 \neq 1$ such that $x+0 = 0$ and $x1 = x$ for all $x \in F$.
- (f4) (inverses exist) Given $x \in F$, there exists an element $-x \in F$ such that $x+(-x) = 0$. If $x \neq 0$, there exists an element x^{-1} such that $xx^{-1} = 1$.
- (f5) (distributive property) $x(y+z) = xy+xz$ for all $x, y, z \in F$.

Exercise 8.4.3. Using the usual definitions of addition and multiplication, determine which of these properties are possessed by \mathbf{N} , \mathbf{Z} , and \mathbf{Q} , respectively.

Although we will not pursue this here in any depth, all of the familiar algebraic manipulations in \mathbf{Q} (e.g., $x+y = x+z$ implies $y = z$) can be derived from this short list of properties.

Definition 8.4.5. An *ordering* on a set F is a relation, represented by \leq , with the following three properties:

- (o1) For arbitrary $x, y \in F$, at least one of the statements $x \leq y$ or $y \leq x$ is true.
- (o2) If $x \leq y$ and $y \leq x$, then $x = y$.
- (o3) If $x \leq y$ and $y \leq z$, then $x \leq z$.

We will sometimes write $y \geq x$ in place of $x \leq y$. The strict inequality $x < y$ is used to mean $x \leq y$ but $x \neq y$.

A field F is called an *ordered field* if F is endowed with an ordering \leq that satisfies

- (o4) If $y \leq z$, then $x+y \leq x+z$.
- (o5) If $x \geq 0$ and $y \geq 0$, then $xy \geq 0$.

Let's take stock of where we are. To prove Theorem 8.4.1, we are accepting as given that the rational numbers are an ordered field. We have defined the real numbers \mathbf{R} to be the collection of cuts in \mathbf{Q} , and the challenge now is to invent addition, multiplication, and an ordering so that each possesses the properties outlined in the preceding two definitions. The easiest of these is the ordering. Let A and B be two arbitrary elements of \mathbf{R} .

Define $A \leq B$ to mean $A \subseteq B$.

Exercise 8.4.4. Show that this defines an ordering on \mathbf{R} by verifying properties (o1), (o2), and (o3) from Definition 8.4.5.

Algebra in \mathbf{R}

Given A and B in \mathbf{R} , define

$$A + B = \{a + b : a \in A \text{ and } b \in B\}.$$

Before checking properties (f1)–(f4) for addition, we must first verify that our definition really defines an operation. Is $A + B$ actually a cut? To get the flavor of how these arguments look, let's verify property (c2) of Definition 8.4.2 for the set $A + B$.

Let $a + b \in A + B$ be arbitrary and let $s \in \mathbf{Q}$ satisfy $s < a + b$. Then, $s - b < a$, which implies that $s - b \in A$ because A is a cut. But then

$$s = (s - b) + b \in A + B,$$

and (c2) is proved.

Exercise 8.4.5. (a) Show that (c1) and (c3) also hold for $A + B$. Conclude that $A + B$ is a cut.

(b) Check that addition in \mathbf{R} is commutative (f1) and associative (f2).

(c) Show that the cut


$$O = \{p \in \mathbf{Q} : p < 0\}$$

successfully plays the role of the additive identity (f3). (Showing $A + O = O$ amounts to proving that these two sets are the same. The standard way to prove such a thing is to show two inclusions: $A + O \subseteq O$ and $O \subseteq A + O$.)

What about additive inverses? Given $A \in \mathbf{R}$, we must produce a cut $-A$ with the property that $A + (-A) = O$. This is a bit more difficult than it sounds. Conceptually, the cut $-A$ consists of all rational numbers less than $-\sup A$. The problem is how to define this set without using suprema, which are strictly off limits at the moment. (We are building the field in which they exist!)

Given $A \in \mathbf{R}$, define

$$-A = \{r \in \mathbf{Q} : \text{there exists } t \notin A \text{ with } t < -r\}.$$



Exercise 8.4.6. (a) Prove that $-A$ defines a cut.

(b) What goes wrong if we set $-A = \{r \in \mathbf{Q} : -r \notin A\}$?

(c) If $a \in A$ and $r \in -A$, show $a + r \in O$. This shows $A + (-A) \subseteq O$. Now, finish the proof of property (f4) for addition in Definition 8.4.4.

(d) Show that property (o3) holds.

Although the ideas are similar, the technical difficulties increase when we try to create a definition for multiplication in \mathbf{R} . This is largely due to the fact that the product of two negative numbers is positive. The standard method of attack is first to define multiplication on the positive cuts.

Given $A \geq O$ and $B \geq O$ in \mathbf{R} , define the product

$$AB = \{ab : a \in A, b \in B \text{ with } a, b \geq 0\} \cup \{q \in \mathbf{Q} : q < 0\}.$$

Exercise 8.4.7. (a) Show that AB is a cut.

(b) Prove property (o4) from Definition 8.4.5.

(c) Propose a good candidate for the multiplicative identity (1) on \mathbf{R} and show that this works for all cuts $A \geq O$.

(d) Show that $AO = O$ for all cuts $A \geq O$.

Products involving at least one negative factor can be defined in terms of the product of two positive cuts by observing that $-A \geq 0$ whenever $A \leq O$. (Given $A \leq O$, property (c4) implies $A + (-A) \leq O + (-A)$, which yields $O \leq -A$.)

For any A and B in \mathbf{R} , define

$$AB = \begin{cases} \text{as given} & \text{if } A \geq O \text{ and } B \geq O \\ -[A(-B)] & \text{if } A \geq O \text{ and } B < O \\ -[(-A)B] & \text{if } A < O \text{ and } B \geq O \\ (-A)(-B) & \text{if } A < O \text{ and } B < O. \end{cases}$$

Verifying that multiplication defined in this way satisfies all the required field properties is important but uneventful. The proofs generally fall into cases for when terms are positive or negative and follow a pattern similar to those for addition. We will leave them as an unofficial exercise and move on to the punch line.

Least Upper Bounds

Having proved that \mathbf{R} is an ordered field, we now set our sights on showing that this field is complete. We defined completeness in Chapter 1 in terms of least upper bounds. Here is a summary of the relevant definitions from that discussion.

Definition 8.4.6. A set $\mathcal{A} \subseteq \mathbf{R}$ is *bounded above* if there exists a $B \in \mathbf{R}$ such that $A \leq B$ for all $A \in \mathcal{A}$. The number B is called an *upper bound* for \mathcal{A} .

A real number $S \in \mathbf{R}$ is the *least upper bound* for a set $\mathcal{A} \subseteq \mathbf{R}$ if it meets the following two criteria:

- (i) S is an upper bound for \mathcal{A} and
- (ii) if B is any upper bound for \mathcal{A} , then $S \leq B$.

Exercise 8.4.8. Let $\mathcal{A} \subseteq \mathbf{R}$ be nonempty and bounded above, and let S be the *union* of all $A \in \mathcal{A}$.

- (a) First, prove that $S \in \mathbf{R}$ by showing that it is a cut.
- (b) Now, show that S is the least upper bound for \mathcal{A} .

This finishes the proof that \mathbf{R} is complete. Notice that we could have proved that least upper bounds exist immediately after defining the ordering on \mathbf{R} , but saving it for last gives it the privileged place in the argument it deserves. There is, however, still one loose end to sew up. The statement of Theorem 8.4.1 mentions that our complete ordered field contains \mathbf{Q} as a subfield. This is a slight abuse of language. What it should say is that \mathbf{R} contains a subfield that looks and acts exactly like \mathbf{Q} .

Exercise 8.4.9. Consider the collection of so-called “rational” cuts of the form

$$C_r = \{t \in \mathbf{Q} : t < r\}$$

where $r \in \mathbf{Q}$. (See Exercise 8.4.1.)

(a) Show that $C_r + C_s = C_{r+s}$ for all $r, s \in \mathbf{Q}$. Verify $C_r C_s = C_{rs}$ for the case when $r, s \geq 0$.

(b) Show that $C_r \leq C_s$ if and only if $r \leq s$ in \mathbf{Q} .

Cantor’s Approach

As a way of giving Georg Cantor the last word, let’s briefly look at his very different approach to constructing \mathbf{R} out of \mathbf{Q} . One of the many equivalent ways to characterize completeness is with the assertion that “Cauchy sequences converge.” Given a Cauchy sequence of rational numbers, we are now well aware that this sequence may converge to a value not in \mathbf{Q} . Just as before, the goal is to create something, which we will call a *real number*, that can serve as the limit of this sequence. Cantor’s idea was essentially to define a real number to be the entire Cauchy sequence. The first problem one encounters with this approach is the realization that two different Cauchy sequences can converge to the same real number. For this reason, the elements in \mathbf{R} are more appropriately defined as *equivalence classes* of Cauchy sequences where two sequences (x_n) and (y_n) are in the same equivalence class if and only if $(x_n - y_n) \rightarrow 0$.

As with Dedekind’s approach, it can be momentarily disorienting to supplant our relatively simple notion of a real number as a decimal expansion with something as unruly as an equivalence class of Cauchy sequences. But what exactly do we mean by a decimal expansion? And how are we to understand the number $1/2$ as both $.5000\dots$ and $.4999\dots$? We leave it as an exercise.

Bibliography

- [1] Robert G. Bartle, *The Elements of Real Analysis*. Second Edition. John Wiley and Sons, New York, 1964.
- [2] Robert G. Bartle, “Return to the Riemann Integral,” *American Mathematical Monthly*, October, 1996.
- [3] E. T. Bell, *Men of Mathematics*. Simon and Schuster, New York, 1937.
- [4] Carl B. Boyer, *A History of Mathematics*. Princeton University Press, Princeton, New Jersey, 1969.
- [5] David Bressoud, *A Radical Approach to Real Analysis*. The Mathematical Association of America, Washington D.C., 1994.
- [6] Soo Bong Chae, *Lebesgue Integration*. Monographs and Textbooks in Pure and Applied Mathematics, Marcel Dekker, New York, 1980.
- [7] W.A. Coppel, “J.B. Fourier—On the Occasion of his Two Hundredth Birthday,” *American Mathematical Monthly*, 76, 1969, pp. 468–483.
- [8] Roger Cooke, “Uniqueness of Trigonometric Series and Descriptive Set Theory,” Offprint from *Archive for History of Exact Sciences*, Volume 45, number 4, Springer–Verlag, New York, 1993, pp. 281–334.
- [9] William Dunham, *Journey Through Genius*. John Wiley and Sons, New York, 1990.
- [10] H. Dym and H.P. McKean, *Fourier Series and Integrals*. Academic Press, Inc., New York, 1972.
- [11] E. Hairer and G. Wanner, *Analysis by Its History*. Undergraduate Texts in Mathematics, Springer–Verlag, New York, 1996.
- [12] Paul R. Halmos, *Naive Set Theory*. Undergraduate Texts in Mathematics, Springer–Verlag, New York, 1974.
- [13] G.H. Hardy, *A Mathematician’s Apology*. Cambridge University Press (Canto Edition), Cambridge, 1992.

- [14] E.W. Hobson, *The Theory of Functions of a Real Variable and the Theory of Fourier's Series*. Volume 1, Third Edition. Harren Press, Washington D.C., 1950.
- [15] Cameron Parks, *Branching Processes*. Senior thesis, Middlebury College, 1998.
- [16] Walter Rudin, *Principles of Mathematical Analysis*. International Series in Pure and Applied Mathematics, McGraw–Hill, New York, 1964.
- [17] George Simmons, *Calculus Gems: Brief Lives and Memorable Mathematics*. McGraw–Hill, New York, 1992.
- [18] Rudolph Vyborný and Lee Peng Yee, *Integral: An Easy Approach after Kurzweil and Henstock*. Cambridge University Press, Cambridge, 2000.

Index

- Abel, Niels Henrik, 13, 170
- Abel's Lemma, 171
- Abel's Test, 69, 171
- Abel's Theorem, 172
- Abel-summable, 175
- absolute convergence, 65
 - of power series, 170
- Absolute Convergence Test, 65
- accumulation point, 79
- additive inverse, 17, 246
- Algebraic Limit Theorem
 - for continuous functions, 110
 - for derivatives, 133
 - for functional limits, 106
 - for sequences, 45
 - for series, 62
- algebraic number, 28
- Alternating Series Test, 65
- antichain, 32
- Archimedean property, 19
- Arzela–Ascoli Theorem, 34, 163, 224
- average value, 202
- Axiom of Completeness, 14, 16, 17, 60, 62, 122, 244

- Baire, René Louis, 94
- Baire Category Theorem, 96, 127, 148, 226
- Baire's Theorem, 96, 127
- Bernoulli, Daniel, 230
- Bernoulli, Johann, 141
- Bernstein, Felix, 33
- Bolzano, Bernhard, 13, 60, 99, 121, 144
- Bolzano–Weierstrass Theorem, 56, 60, 62

- bounded
 - function, 114
 - sequence, 44
 - set, 84, 224
- branching processes, 151

- Cantor, Georg, 22, 75, 244, 249
- Cantor diagonalization method, 29, 34
- Cantor function, 162
- Cantor set, 75, 89, 149, 204, 205, 208
- Cantor's Theorem, 32
- cardinal number, 33
- cardinality, 22, 33
- Carleson, Lennart, 237
- category, first or second, 96, 226
- Cauchy, Augustin Louis, 13, 60, 99, 181
- Cauchy Condensation Test, 53
- Cauchy Criterion, 59, 61, 62
 - for series, 63
 - for uniform convergence, 160
 - for uniform convergence of series, 168
 - in a metric space, 223
- Cauchy product, 73, 74, 175
- Cauchy sequence, 58
 - as a real number, 249
 - converges, 59
 - in a metric space, 223
- Cesaro means, 50, 241
- chain rule, 134
- closed set, 80
 - in a metric space, 224
- closure, 80, 81
 - in a metric space, 225
- cluster point, 79
- Cohen, Paul, 34
- compact, 84
 - in a metric space, 224
 - subsets of $C[0, 1]$, 225
- Comparison Test, 63
- complement, 6, 82
- complete metric space, 223, 226
- conditional convergence, 65
- connected set, 91, 121
- continuity, 109

- α , 127, 205
- and integrability, 190
- and Riemann integrability, 185
- and uniform convergence, 160
- characterizations of, 109
- in a metric space, 223
- nowhere, 101
- of compositions, 112
- on a set, 109
- on compact sets, 115
- uniform, 116
- Continuous Extension Theorem, 120
- continuum hypothesis, 34
- Contraction Mapping Theorem, 113
- convergence
 - L^2 , 232, 237
 - absolute, 65, 70
 - Cesaro mean, 232, 241
 - conditional, 65
 - metric space, 223
 - of a Cauchy sequence, 59
 - of a sequence, 39, 58
 - of a subsequence, 55
 - of p -series, 54
 - of series, 51, 62
 - pointwise, 154, 158
 - pointwise for series, 167
 - to infinity, 44
 - uniform, 157
 - uniform for series, 167
- countable, 24
 - \mathbf{Q} is, 24
 - \mathbf{R} is not, 24
 - Cantor set is not, 77, 90
 - subsets, 26
 - unions, 26
- cut, 245
- d'Alembert, Jean Le Rond, 229
- Darboux, Gaston, 135
- Darboux's Theorem, 136, 149, 183
- De Morgan's Laws, 6, 11, 83
- decimal expansion, 30, 249
- decreasing
 - function, 125
 - sequence, 51
- Dedekind, Richard, 244
- Dedekind cut, 245
- density
 - in a metric space, 225
 - nowhere, 95, 225
 - of \mathbf{Q} , 20, 81
- derivative, 129, 132, 199
- dimension, of Cantor set, 77
- Dirichlet, Peter Lejeune, 7, 13, 100
- Dirichlet kernel, 238, 239, 242
- Dirichlet's function, 7, 100, 125, 184, 193, 210, 218
- Dirichlet's Test, 69
- disconnected
 - set, 91
 - totally, 93
- discontinuity
 - all types, 126
 - essential, 131
- divergence
 - of a sequence, 42, 56
 - of functional limits, 107
- domain, 7
- double summation, 37, 69
- empty set, 5
- equivalence classes
 - of Cauchy sequences, 249
 - of sets, 33
- equivalence relation, 28, 33
- eventually, 44, 49, 64
- Extreme Value Theorem, 115
- Fejér, Lipót, 241
- Fejér kernel, 242
- Fejér's Theorem, 241
- Fermat, Pierre de, 99, 135
- field, 3, 14, 17, 245
- fixed point, 143, 153
- Fourier, Joseph, 228
- Fourier coefficients, 232
 - converge to zero, 235
- Fourier series, 99, 145, 184, 228
 - Cesaro mean convergence of, 241
 - pointwise convergence of, 238
- fractal, 78
- frequently, 44
- function, 7
- functional limit, 104
- Fundamental Theorem of Calculus, 138, 149, 177, 183, 200, 202, 213, 219
- Galton, Francis, 151

- gauge, 216
- Generalized Mean Value Theorem, 140
- generalized Riemann integral, 217
- Gödel, Kurt, 34
- greatest lower bound, 14, 17

- halting problem, 34
- Hardy, Godfrey Harold, 1, 148
- harmonic series, 52
 - alternating, 73
- Heine, Heinrich, 244
- Heine–Borel Theorem, 84

- increasing
 - function, 125
 - sequence, 51
- infimum, 14, 17
- infinite series, 51, 62
 - associative property, 57
 - comparison test, 63
 - converges, 62
 - double summations, 69
 - of functions, 167
 - partial sum, 62
 - products, 63, 72
 - ratio test, 68
- integer, 3
 - countable, 24
- integral
 - generalized Riemann, 212, 213, 217
 - improper, 221
 - Lebesgue, 211, 213, 237
 - lower, 188
 - Riemann, 188
 - substitution formula, 221
 - upper, 188
- interior, 225
- Interior Extremum Theorem, 135
- intermediate value property, 124, 131
 - of derivatives, 136
- Intermediate Value Theorem, 120
- irrational number, 1, 4
- isolated point, 80

- Kolmogorov, Andrey, 237
- Kronecker, Leopold, 2, 11

- Lagrange, Joseph Louis, 178
- Lagrange’s Remainder Theorem, 178
- least upper bound, 14, 16, 243, 248

- Lebesgue, Henri, 148, 193, 213
- Lebesgue integral, 211
- Lebesgue’s Theorem, 206
- length, of Cantor set, 76, 204
- L’Hospital, Guillaume François Antoine de, 141
- L’Hospital’s rule, 141
- lim inf, 55
- lim sup, 55, 175
- limit
 - of a function, 104
 - of a sequence, 39
 - of Riemann sums, 185, 215
 - right-hand, 125
 - superior, 55, 175
- limit point, 79
 - in metric space, 224
- Lipschitz function, 120
- lower integral, 188
- lower sum, 186

- Mandelbrot, Benoit, 78
- maximum, 15
 - attained on compact sets, 115
- Mean Value Theorem, 137, 138, 219
 - generalized, 140
- measure zero, 204, 213, 237
- Meray, Charles, 244
- metric, 222
 - discrete, 222
- metric space, 96, 222
 - complete, 223
- minimum, 15
- monotone
 - function, 103, 125, 128, 148
 - sequence, 51
- Monotone Convergence Theorem, 51, 60, 62
- multiplicative inverse, 17, 246

- natural logarithm, 202
- natural number, 2
- neighborhood, 39, 78
 - in a metric space, 224
- Nested Interval Property, 18, 60, 122
- nowhere
 - continuous, 101
 - dense, 95, 225, 226
 - differentiable, 144, 226, 227

- one-to-one, 23
- onto, 23
- open cover, 86, 206
- open set, 78
 - in a metric space, 224
- Order Limit Theorem, 48
- ordered field, 3, 14, 246
- ordering, 246

- partial sum, 51
- partition, 186
 - δ -fine, 214
 - $\delta(x)$ -fine, 217
 - refinement, 187
 - tagged, 214
- perfect set, 89
- piecewise linear, 228
- pointwise convergence, 154
 - for series, 167
 - of Fourier series, 237, 238
- power series, 74, 99, 151, 169, 229
 - differentiation of, 173
 - uniform convergence of, 172
- power set, 31
- preimage, 12, 122
- proof
 - by contradiction, 9
 - by induction, 10
 - contrapositive, 9
 - of convergence, 41

- quantifiers, 41

- radius of convergence, 170
- range, 7
- Ratio Test, 68
- rational number, 1, 3, 249
 - countable, 24
- real number, 4, 13
 - as a Cauchy sequence, 249
 - as a cut, 245
 - uncountable, 24
- rearrangement, 36, 66
- refinement, 187
 - common, 187
- Riemann, Bernhard, 13, 184
- Riemann integral, 184, 188
 - and continuity, 190, 210
 - and discontinuity, 192, 193, 203, 206
 - criterion for existence, 189
 - improper, 212
 - properties of, 195
- Riemann sum, 184, 214
- Riemann–Lebesgue lemma, 235
- Rolle, Michel, 138
- Rolle’s Theorem, 138

- Schröder, Ernst, 33
- Schröder–Bernstein Theorem, 29, 33
- separated sets, 91
- sequence, 38
- sequential criterion
 - for continuity, 102, 109
 - for functional limits, 106
- set, 5
 - F_σ , 94, 126
 - G_δ , 94
 - closed, 80
 - compact, 84
 - complement of, 6
 - connected, 91
 - disjoint, 5
 - empty, 5
 - fat, 96
 - first or second category, 96, 226
 - inclusion, 5
 - meager, 96, 227
 - of measure zero, 204
 - open, 78
 - perfect, 89
 - subset, 5
- square roots, 1, 54
- Squeeze Theorem, 49, 109
- subsequence, 55
- substitution formula, 221
- summation by parts, 68
- sup norm, 223
- supremum, 14

- Taylor, Brook, 177
- Taylor series, 176, 231, 242
 - formula for coefficients, 177
 - remainder formula, 178
- Thomae, K.J., 102
- Thomae’s function, 102, 125, 203
- total variation, 202
- transcendental number, 28
- triangle inequality, 8, 11, 47, 222
- trigonometric series, 229

- uniform convergence, 157
 - and continuity, 160, 167
 - and differentiation, 164, 166, 168
 - and integration, 197, 212
 - of power series, 170, 172
 - of series, 167
- uniformly
 - α -continuous, 205
 - continuous, 116
- uniqueness
 - of generalized Riemann integral, 217
 - of limits, 49
- upper bound, 14, 248
- upper integral, 188
- upper sum, 186
- wave equation, 230
- Weierstrass, Karl, 13, 60, 99, 144, 148
- Weierstrass Approximation Theorem, 242
- Weierstrass M-Test, 168