Preface

Overview

Why Mathematical Analysis?

Mathematical analysis is the foundation upon which nearly every area of applied mathematics is built. It is the language and intellectual framework for studying optimization, probability theory, stochastic processes, statistics, machine learning, differential equations, and control theory. It is also essential for rigorously describing the theoretical concepts of many quantitative fields, including computer science, economics, physics, and several areas of engineering.

Beyond its importance in these disciplines, mathematical analysis is also fundamental in the design, analysis, and optimization of algorithms. In addition to allowing us to make objectively true statements about the performance, complexity, and accuracy of algorithms, mathematical analysis has inspired many of the key insights needed to create, understand, and contextualize the fastest and most important algorithms discovered to date.

In recent years, the size, speed, and scale of computing has had a profound impact on nearly every area of science and technology. As future discoveries and innovations become more algorithmic, and therefore more computational, there will be tremendous opportunities for those who understand mathematical analysis. Those who can peer beyond the jargon-filled barriers of various quantitative disciplines and abstract out their fundamental algorithmic concepts will be able to move fluidly across quantitative disciplines and innovate at their crossroads. In short, mathematical analysis gives solutions to quantitative problems, and the future is promising for those who master this material.

To the Instructor

About this Text

This text modernizes and integrates a semester of advanced linear algebra with a semester of multivariable real analysis to give a new and redesigned year-long curriculum in linear and nonlinear analysis. The mathematical prerequisites are vector calculus, linear algebra, and a semester of undergraduate-level, single-variable real analysis.[1]

The content in this volume could be reasonably described as upper-division undergraduate or first-year graduate-level mathematics. It can be taught as a stand-alone two-semester sequence or in parallel with the second volume, *Foundations of Applied Mathematics, Volume 2: Algorithms, Approximation, and Optimization*, as part of a larger curriculum in applied and computational mathematics.

There is also a supplementary computer lab manual, containing over 25 computer labs that support this text. This text focuses on the theory, while the labs cover application and computation. Although we recommend that the manual be used in a computer lab setting with a teaching assistant, the material can be used without instruction. The concepts are developed slowly and thoroughly, with numerous examples and figures as pedagogical breadcrumbs, so that students can learn this material on their own, verifying their progress along the way. The labs and other classroom resources are open content and are available at

`http://www.siam.org/books/ot152`

The intent of this text and the computer labs is to attract and retain students into the mathematical sciences by modernizing the curriculum and connecting theory to application in a way that makes the students want to understand the theory, rather than just tolerate it. In short, a major goal of this text is to entice them to hunger for more.

Topics and Focus

In addition to standard material one would normally expect from linear and nonlinear analysis, this text also includes several key concepts of modern applied mathematical analysis which are not typically taught in a traditional applied math curriculum (see the Detailed Description, below, for more information).

We focus on both rigor and relevance to give the students mathematical maturity and an understanding of the most important ideas in mathematical analysis.

Detailed Description

Chapters 1–3 We give a rigorous treatment of the basics of linear algebra over both $\mathbb{R}$ and $\mathbb{C}$, including abstract vector spaces, linear transformations, matrices, the LU decomposition, inner product spaces, the QR decomposition, and least squares. As much as possible, we try to frame things in a way that does not require vector spaces to be finite dimensional, and we give many infinite-dimensional examples.

Chapter 4 We treat the spectral theory of matrices, including the spectral theorem for normal matrices. We give special attention to the singular value decomposition and its applications.

[1]Specifically, the reader should have had exposure to a rigorous treatment of continuity, convergence, differentiation, and Riemann–Darboux integration in one dimension, as covered, for example, in [Abb15].

Chapter 5 We present the basics of metric topology, including the ideas of completeness and compactness. We define and give many examples of Banach spaces. Throughout the rest of the text we formulate results in terms of Banach spaces, wherever possible. A highlight of this chapter is the continuous linear extension theorem (sometimes called the bounded linear transformation theorem), which we use to give a very slick construction of Riemann (or rather *regulated*) Banach-valued integration (single-variable in this chapter and multivariable in Chapter 8).

Chapters 6–7 We discuss calculus on Banach spaces, including Fréchet derivatives and Taylor's theorem. We then present the uniform contraction mapping theorem, which we use to prove convergence results for Newton's method and also to give nice proofs of the inverse and implicit function theorems.

Chapters 8–9 We use the same basic ideas to develop Lebesgue integration as we used in the development of the regulated integral in Chapter 5. This approach could be called the *Riesz* or *Daniell* approach. Instead of developing measure theory and creating integrals from simple functions, we define what it means for a set to have zero measure and create integrals from step functions. This is a very clean way to do integration, which has the additional benefit of reinforcing many of the functional-analytic ideas developed earlier in the text.

Chapters 10–11 We give an introduction to the fundamental tools of complex analysis, briefly covering first the main ideas of parametrized curves, surfaces, and manifolds as well as line integrals, and Green's theorem, to provide a solid foundation for contour integration and Cauchy's theorem.

Throughout both of these chapters, we express the main ideas and results in terms of Banach-valued functions whenever possible, so that we can use these powerful tools to study spectral theory of operators later in the book.

Chapters 12–13 One of the biggest innovations in the book is our treatment of spectral theory. We take the Dunford–Schwartz approach via resolvents. This approach is usually only developed from an advanced functional-analytic point of view, but we break it down to the level of an undergraduate math major, using the tools and ideas developed earlier in this text.

In this setting, we put a strong emphasis on eigenprojections, providing insights into the spectral resolution theorem. This allows for easy proofs of the spectral mapping theorem, the Perron–Frobenius theorem, the Cayley–Hamilton theorem, and convergence of the power method. This also allows for a nice presentation of the Drazin inverse and matrix perturbation theory. These ideas are used again in Volume 4 with dynamical systems, where we prove the stable and center manifold theorems using spectral projections and corresponding semigroup estimates.

Chapter 14 The pseudospectrum is a fundamental tool in modern linear algebra. We use the pseudospectrum to study sequences of the form $\|A^k\|$, their asymptotic and transient behavior, an understanding of which is important both for Markov chains and for the many iterative methods based on such sequences, such as successive overrelaxation.

Chapter 15 We conclude the book with a chapter on applied ring theory, focused on the algebraic structure of polynomials and matrices. A major focus of this chapter is the Chinese remainder theorem, which we use in many ways, including to prove results about partial fractions and Lagrange interpolation. The highlight of the chapter is Section 15.7.3, which describes a striking connection between Lagrange interpolation and the spectral decomposition of a matrix.

Teaching from the Text

In our courses we teach each section in a fifty-minute lecture. We require students read the section carefully before each class so that class time can be used to focus on the parts they find most confusing, rather than on just repeating to them the material already written in the book.

There are roughly five to seven exercises per section. We believe that students can realistically be expected to do all of the exercises in the text, but some are difficult and will require time, effort, and perhaps an occasional hint. Exercises that are unusually hard are marked with the symbol †. Some of the exercises are marked with * to indicate that they cover advanced material. Although these are valuable, they are not essential for understanding the rest of the text, so they may safely be skipped, if necessary.

Throughout this book the exercises, examples, and concepts are tightly integrated and build upon each other in a way that reinforces previous ideas and prepares students for upcoming ideas. We find this helps students better retain and understand the concepts learned, and helps achieve greater depth. Students are encouraged to do *all* of the exercises, as they reinforce new ideas and also revisit the core ideas taught earlier in the text.

Courses Taught from This Book

Full Sequence

At BYU we teach a year-long advanced undergraduate-level course from this book, proceeding straight through the book, skipping only the advanced sections and chapters (marked with *), and ending in Chapter 13. But this would also make a very good course at the beginning graduate level as well.

Graduate students who are well prepared could be further challenged either by covering the material more rapidly, so as to get to the very rewarding material at the end of the book, or by covering some or all of the advanced sections along the way.

Advanced Linear Algebra

Alternatively, Chapters 1–4 (linear analysis part I), Section 7.5 (conditioning), and Chapters 12–14 (linear analysis part II), as well as parts of Chapter 15, as time permits, make up a very good one-semester advanced linear algebra course for students who have already completed undergraduate-level courses in linear algebra, complex analysis, and multivariate real analysis.

Advanced Analysis

This book can also be used to teach a one-semester advanced analysis course for students who have already had a semester of basic undergraduate analysis (say, at the level of [Abb15]). One possible path through the book for this course would be to briefly review Chapter 1 (vector spaces), Sections 2.1–2.2 (basics of linear transformations), and Sections 3.1 and 3.5 (inner product spaces and norms), in order to set notation and to remind the students of necessary background from linear algebra, and then proceed through Part II (Chapters 5–7) and Part III (Chapters 8–11).

Figure 1 indicates the dependencies among the chapters.

Advanced Sections

A few problems, sections, and even chapters are marked with the symbol * to indicate that they cover more advanced topics. Although this material is valuable, it is not essential for understanding the rest of the text, so it may safely be skipped, if necessary.

Instructors New to the Material

We've taken a tactical approach that combines professional development for faculty with instruction for the students. Specifically, the class instruction is where the theory lies and the supporting media (labs, etc.) are provided so that faculty need not be computer experts nor familiar with the applications in order to run the course.

The professor can teach the theoretical material in the text and use teaching assistants, who may be better versed in the latest technology, to cover the applications and computation in the labs, where the "hands-on" part of the course takes place. In this way the professor can gradually become acquainted with the applications and technology over time, by working through the labs on his or her own time without the pressures of staying ahead of the students.

A more technologically experienced applied mathematician could flip the class if she wanted to, or change it in other ways. But we feel the current format is most versatile and allows instructors of all backgrounds to gracefully learn and adapt to the program. Over time, instructors will become familiar enough with the content that they can experiment with various pedagogical approaches and make the program theirs.

To the Student

Examples

Although some of the topics in this book may seem familiar to you, especially many of the linear algebra topics, we have taken a very different approach to these topics by integrating many different topics together in our presentation, so examples treated in a discussion of vector spaces will appear again in sections on nonlinear analysis and other places throughout the text. Also, notation introduced in the examples is often used again later in the text.

Because of this, we recommend that you *read all the examples* in each section, even if the definitions, theorems, and other results look familiar.

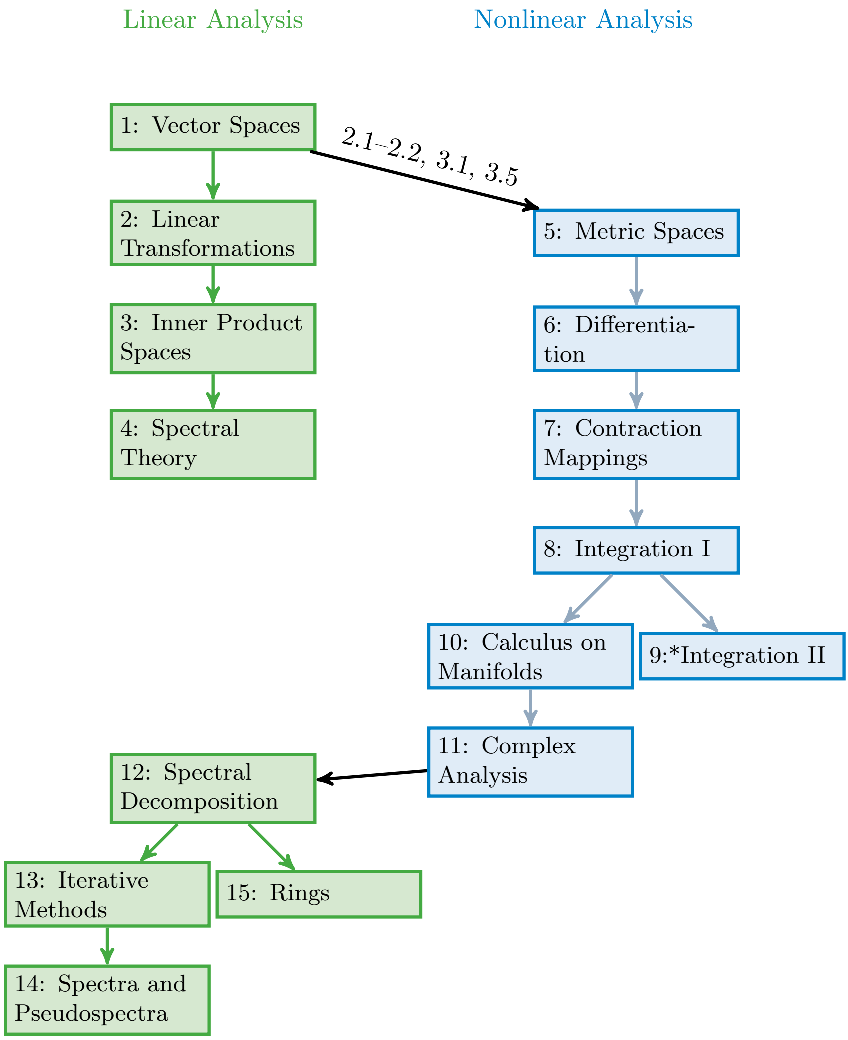

Figure 1. *Diagram of the dependencies among the chapters of this book. Although we usually proceed straight through the book in order, it could also be used for either an advanced linear algebra course or a course in real and complex analysis. The linear analysis half (left side) provides a course in advanced linear algebra for students who have had complex analysis and multivariate real analysis. The nonlinear analysis half (right side) could be used for a course in real and complex analysis for students who have already had linear algebra and a semester of real analysis. For that track, we recommend briefly reviewing the material of Chapter 1 and Sections 2.1–2.2, 3.1, and 3.5, before proceeding to the nonlinear material, in order to fix notation and ensure students remember the necessary background.*

Exercises

Each section of the book has several exercises, all collected at the end of each chapter. Horizontal lines separate the exercises for each section from the exercises for the other sections. We have carefully selected these exercises. You should work them all (but your instructor may choose to let you skip some of the advanced exercises marked with *)—each is important for your ability to understand subsequent material.

Although the exercises are gathered together at the end of the chapter, we strongly recommend that you do the exercises for each section as soon as you have completed the section, rather than saving them until you have finished the entire chapter. Learning mathematics is like developing physical strength. It is much easier to improve, and improvement is greater, when exercises are done daily, in measured amounts, rather than doing long, intense bouts of exercise separated by long rests.

Origins

This curriculum evolved as an outgrowth of lecture notes and computer labs that were developed for a 6-credit summer course in computational mathematics and statistics. This was designed to introduce groups of undergraduate researchers to a number of core concepts in mathematics, statistics, and computation as part of a National Science Foundation (NSF) funded mentoring program called CSUMS: Computational Science Training for Undergraduates in the Mathematical Sciences.

This NSF program sought out new undergraduate mentoring models in the mathematical sciences, with particular attention paid to computational science training through genuine research experiences. Our answer was the Interdisciplinary Mentoring Program in Analysis, Computation, and Theory (IMPACT), which took cohorts of mathematics and statistics undergraduates and inserted them into an intense summer "boot camp" program designed to prepare them for interdisciplinary research during the school year. This effort required a great deal of experimentation, and when the dust finally settled, the list of topics that we wanted to teach blossomed into 8 semesters of material—essentially an entire curriculum.

After we explained the boot camp concept to one visitor, he quipped, "It's the minimum number of instructions needed to create an applied mathematician." Our goal, however, is much broader than this. We don't want to train or create a specific type of applied mathematician; we want a curriculum that supports all types, simultaneously. In other words, our goal is to take in students with diverse and evolving interests and backgrounds and provide them with a common corpus of mathematical, statistical, and computational content so that they can emerge well prepared to work in *their own* chosen areas of specialization. We also want to draw their attention to the core ideas that are ubiquitous across various applications so that they can navigate fluidly across fields.

Acknowledgments

We thank the National Science Foundation for their support through the TUES Phase II grant DUE-1323785. We especially thank Ron Buckmire at the National

Science Foundation for taking a chance on us and providing much-needed advice and guidance along the way. Without the NSF, this book would not have been possible. We also thank the Department of Mathematics at Brigham Young University for their generous support and for providing a stimulating environment in which to work.

Many colleagues and friends have helped shape the ideas that led to this text, especially Randy Beard, Rick Evans, Shane Reese, Dennis Tolley, and Sean Warnick, as well as Bryant Angelos, Jonathan Baker, Blake Barker, Mylan Cook, Casey Dougal, Abe Frandsen, Ryan Grout, McKay Heasley, Amelia Henricksen, Ian Henricksen, Brent Kerby, Steven Lutz, Shane McQuarrie, Ryan Murray, Spencer Patty, Jared Webb, Matthew Webb, Jeremy West, and Alexander Zaitzeff, who were all instrumental in helping to organize this material.

We also thank the students of the BYU Applied and Computational Mathematics Emphasis (ACME) cohorts of 2013–2015, 2014–2016, 2015–2017, and 2016–2018, who suffered through our mistakes, corrected many errors, and never hesitated to tell us what they thought of our work.

We are deeply grateful to Chris Grant, Todd Kapitula, Zach Boyd, Rachel Webb, Jared McBride, and M. A. Averill, who read various drafts of this volume very carefully, corrected many errors, and gave us a tremendous amount of helpful feedback. Of course, all remaining errors are entirely our fault. We also thank Amelia Henricksen, Sierra Horst, and Michael Hansen for their help illustrating the text and Sarah Kay Miller for her outstanding graphic design work, including her beautifully designed book covers. We also appreciate the patience, support, and expert editorial work of Elizabeth Greenspan and the other editors and staff at SIAM.

Finally, we thank the folks at Savvysherpa, Inc., for corporate sponsorship that greatly helped make the transition from IMPACT to ACME and their help nourishing and strengthening the ACME development team.