Preface

Convexity is very easy to define, to visualize and to get an intuition about. A set is called convex if for every two points a and b in the set, the straight line interval $[a, b]$ is also in the set. Thus the main building block of convexity theory is a straight line interval.

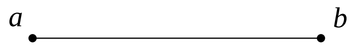

Convexity is more intuitive than, say, linear algebra. In linear algebra, the interval is replaced by the whole straight line. We have some difficulty visualizing a straight line because it runs unchecked in both directions.

On the other hand, the structure of convexity is richer than that of linear algebra. It is already evident in the fact that all points on the line are alike whereas the interval has two points, a and b, which clearly stand out.

Indeed, convexity has an immensely rich structure and numerous applications. On the other hand, almost every "convex" idea can be explained by a two-dimensional picture. There must be some reason for that apart from the tautological one that all our pictures are two-dimensional. One possible explanation is that since the definition of a convex set involves only three points (the two points a and b and a typical point x of the interval) and every three points lie in some plane, whenever we invoke a convexity argument in our reasoning, it can be properly pictured (moreover, since our three points a, b and x lie on the same line, we have room for a fourth point which often plays the role of the origin). Simplicity, intuitive appeal and universality of applications make teaching convexity (and writing a book on convexity) a rather gratifying experience.

About this book. This book grew out of sets of lecture notes for graduate courses that I taught at the University of Michigan in Ann Arbor since 1994. Consequently, this is a graduate textbook. The textbook covers several directions, which, although not independent, provide enough material for several one-semester three-credit courses.

One possibility is to follow discrete and combinatorial aspects of convexity: combinatorial properties of convex sets (Chapter I) – the structure of some interesting polytopes and polyhedra (the first part of Chapter II, some results of Chapter IV and Chapter VI) – lattice points and convex bodies (Chapter VII) – lattice points and polyhedra (Chapter VIII).

Another possibility is to follow the analytic line: basic properties of convex sets (Chapter I) – the structure of some interesting non-polyhedral convex sets, such as the moment cone, the cone of non-negative polynomials and the cone of positive semidefinite matrices (Chapter II and some results of Chapter IV) – metric properties of convex bodies (Chapter V).

Yet another possibility is to follow infinite-dimensional and dimension-free applications of convexity: basic properties of convex sets in a vector space (Chapter I) – separation theorems and the structure of some interesting infinite-dimensional convex sets (Chapter III) – linear inequalities and linear programming in an abstract setting (Chapter IV).

The main focus of the book is on applications of convexity rather than on studying convexity for its own sake. Consequently, mathematical applications range from analysis and probability to algebra to combinatorics to number theory. Finite- and infinite-dimensional optimization problems, such as the Transportation Problem, the Diet Problem, problems of optimal control, statistics and approximation are discussed as well.

The choice of topics covered in the book is entirely subjective. It is probably impossible to write a textbook that covers "all" convexity just as it is impossible to write a textbook that covers all mathematics. I don't even presume to claim to cover all "essential" or "important" aspects of convexity, although I believe that many of the topics discussed in the book belong to both categories.

The audience. The book is intended for graduate students in mathematics and other fields such as operations research, electrical engineering and computer science. That was the typical audience for the courses that I taught. This is, of course, reflected in the selection of topics covered in the book. Also, a significant portion of the material is suitable for undergraduates.

Prerequisites. The main prerequisite is linear algebra, especially the coordinate-free linear algebra. Knowledge of basic linear algebra should be sufficient for understanding the main convexity results (called "Theorems") and solving problems which address convex properties per se.

In many places, knowledge of some basic analysis and topology is needed. In most cases, some general understanding coupled with basic computational skills will be sufficient. For example, when it comes to the topology of Euclidean space, it suffices to know that a set in Euclidean space is compact if and only if it is closed and bounded and that a linear functional attains its maximum and minimum on such a set. Whenever the book says "Lebesgue integral" or "Borel set", it does so for the sake of brevity and means, roughly, "the integral makes sense" and "the set is nice and behaves predictably". For the most part, the only properties of the integral that the book uses are linearity (the integral of a linear combination of two functions is the linear combination of the integrals of the functions) and monotonicity (the integral of a non-negative function is non-negative). The relative abundance of integrals in a textbook on convexity is explained by the fact that the most natural way to define a linear functional is by using an integral of some sort. A few exercises openly require some additional skills (knowledge of functional analysis or representation theory).

When it comes to applications (often called "Propositions"), the reader is expected to have some knowledge in the general area which concerns the application.

Style. The numbering in each chapter is consecutive: for example, Theorem 2.1 is followed by Definition 2.2 which is followed by Theorem 2.3. When a reference is made to another chapter, a roman numeral is included: for example, if Theorem 2.1 of Chapter I is referenced in Chapter III, it will be referred to as Theorem I.2.1. Definitions, theorems and other numbered objects in the text (except figures) are usually followed by a set of problems (exercises). For example, Problem 5 following Definition 2.6 in Chapter II will be referred to as Problem 5 of Section 2.6 from within Chapter II and as Problem 5 of Section II.2.6 from everywhere else in the book. Figures are numbered consecutively throughout the book. There is a certain difference between "Theorems" and "Propositions". Theorems state some general and fundamental convex properties or, in some cases, are called "Theorems" historically. Propositions describe properties of particular convex sets or refer to an application.

Problems. There are three kinds of problems in the text. The problems marked by * are deemed difficult (they may be so marked simply because the author is unaware of an easy solution). Problems with straightforward solutions are marked by $^\circ$. Solving a problem marked by $^\circ$ is essential for understanding the material and its result may be used in the future. Some problems are not marked at all. There are no solutions at the end of the book and there is no accompanying solution manual (that I am aware of) , which, in my opinion, makes the book rather convenient for use in courses where grades are given. On the other hand, many of the difficult and some of the easy problems used later in the text are supplied with a hint to a solution.

Acknowledgment. My greatest intellectual debt is to my teacher A.M. Vershik. As a student, I took his courses on convexity and linear programming. Later, we discussed various topics in convex analysis and geometry and he shared his notes on the subject with me. We planned to write a book on convexity together and actually started to write one (in Russian), but the project was effectively terminated by my relocation to the United States. The overall plan, structure and scope of the book have changed since then, although much has remained the same. All unfortunate choices, mistakes, typos, blunders and other slips in the text are my own. A.M. Vershik always insisted on a "dimension-free"approach to convexity, whenever possible, which simplifies and makes transparent many fundamental facts, and on stressing the idea of duality in the broadest sense. In particular, I learned the algebraic approach to the Hahn-Banach Theorem (Sections II.1, III.1–3) and

the general view of infinite-dimensional linear programming (Chapter IV) from him. This approach makes the exposition rather simple and elegant. It makes it possible to deduce a variety of strong duality results from a single simple theorem (Theorem IV.7.2). My interest in quadratic convexity (Section II.14) and other "hidden convexity" results (Section III.7) was inspired by him. He also encouraged my preoccupation with lattice points (Chapter VIII) and various peculiar polytopes (Sections II.5–7).

On various stages of the project I received encouragement from A. Björner, L. Billera, R. Pollack, V. Klee, J.E. Goodman, G. Kalai, A. Frieze, L. Lovász, W.T. Gowers and I. Bárány.

I am grateful to my colleagues in the Department of Mathematics at the University of Michigan in Ann Arbor, especially to P. Hanlon, B.A. Taylor, J. Stembridge and S. Fomin with whose blessings I promoted convexity within the Michigan combinatorics curriculum. I thank the students who took Math 669 convexity classes in 1994–2001. Special thanks to G. Blekherman who contributed some of his interesting results on the metric structure of the set of non-negative multivariate polynomials (Problems 8 and 9 of Section V.2.4).

Since the draft of this book was posted on the web, I received very useful and detailed comments from R. Connelly, N. Ivanov, J. Lawrence, L. Lovász, G. Ziegler and A.M. Vershik. I am particularly grateful to J. Lawrence who suggested a number of essential improvements, among them are the greater generality of the "polarity as a valuation" theorem (Theorem IV.1.5), a simplified proof of the Euler-Poincaré Formula (Corollary VI.3.2) and an elegant proof of Gram's relations (Problem 1 of Section VIII.4.4) and many mathematical, stylistic and bibliographical corrections.

I thank A. Yong for reading the whole manuscript carefully and suggesting numerous mathematical and stylistic corrections. I thank M. Wendt for catching a mistake and alerting me by e-mail.

I thank S. Gelfand (AMS) for insisting over a number of years that I write the book and for believing that I was able to finish it.

I am grateful to the National Science Foundation for its support.

Ann Arbor, 2002
Alexander Barvinok