Preface

In this book we present the theory of *asymptotic geometric analysis*, a theory which stands at the midpoint between geometry and functional analysis. The theory originated from functional analysis, where one studied Banach spaces, usually of infinite dimensions. In the first few decades of its development it was called "local theory of normed spaces", which stood for investigating infinite dimensional Banach spaces via their finite dimensional features, for example subspaces or quotients. Soon, geometry started to become central. However, as we shall explain below in more detail, the study of "isometric" problems, a point of view typical for geometry, had to be substituted by an "isomorphic" point of view. This became possible with the introduction of an asymptotic approach to the study of high-dimensional spaces (asymptotic with respect to dimensions increasing to infinity). Finally, these finite but very high-dimensional questions and results became interesting in their own right, influential on other mathematical fields of mathematics, and independent of their original connection with infinite dimensional theory. Thus the name asymptotic geometric analysis nowadays describes an essentially new field.

Our primary object of study will be a finite dimensional normed space X; we may assume that X is $\mathbb{R}^n$ equipped with a norm $\|\cdot\|$. Such a space is determined by its unit ball $K_X = \{x \in \mathbb{R}^n : \|x\| \leqslant 1\}$, which is a compact convex set with non-empty interior (we call this type of set "a convex body"). Conversely, if K is a centrally symmetric convex body in $\mathbb{R}^n$, then it is the unit ball of a normed space $X_K = (\mathbb{R}^n, \|\cdot\|_K)$. Thus, the study of finite dimensional normed spaces is in fact a study of centrally symmetric convex bodies, but again, the low-dimensional type questions and the corresponding intuition are very different from what is needed when the emphasis is on high-dimensional asymptotic behaviour. An example that clarifies this difference is given by the following question: does there exist a universal constant $c > 0$ such that every convex body of volume one has a hyperplane section of volume more than c? In any fixed dimension n, simple compactness arguments show that the answer is affirmative (although the question to determine the optimal value of the corresponding constant c_n may remain interesting and challenging). However, this is certainly not enough to conclude that a constant $c > 0$ exists which applies to any body of volume one in any dimension. This is already an *asymptotic type question*. In fact, it is unresolved to this day and will be discussed in Chapter 10.

Classical geometry (in a fixed dimension) is usually an isometric theory. In the field of asymptotic geometric analysis, one naturally studies *isomorphic geometric objects* and derives *isomorphic geometric results*. By an "isomorphic" geometric object we mean a family of objects in different spaces of increasing dimension and by an "isomorphic" geometric property of such an "isomorphic" object we mean a property shared by the high-dimensional elements of this family. One is interested in the asymptotic behaviour with respect to some parameter (most often it is the dimension n) and in the control of how the geometric quantities involved depend on this parameter. The appearance of such an isomorphic geometric object is a new feature of asymptotic high-dimensional theory. Geometry and analysis meet here in a non-trivial way. We will encounter throughout the book many geometric inequalities in isomorphic form. Basic examples of such inequalities are the "isomorphic isoperimetric inequalities" that led to the discovery of the "concentration phenomenon", one of the most powerful tools of the theory, responsible for many counterintuitive results. Let us briefly describe it here, through the primary example of the sphere. A detailed account is given in Chapter 3. Consider the Euclidean unit sphere in $\mathbb{R}^n$, denoted S^{n-1}, equipped with the Lebesgue measure, normalized to have total measure 1. Let A be a subset of the sphere of measure $1/2$. Take an ε-extension of this set, with respect to Euclidean or geodesic distance, for some fixed but small ε; this is the set of all points which are at a distance of at most ε from the original set (usually denoted by A_ε). It turns out that the remaining set (that is, the set $S^{n-1} \setminus A_\varepsilon$ of all points in the sphere which are at a distance more than ε from A) has, in high dimensions, a very small measure, decreasing to zero exponentially fast as the dimension n grows. This type of statement has meaning only in asymptotic language, since in fact we are considering a sequence of spheres of increasing dimensions, and a sequence of subsets of these spheres, each of measure one half of its corresponding sphere, and the sequence of the measures of the ε-extensions (where ε is fixed for all n) is a sequence tending to 1 exponentially fast with dimension. We shall see how the above statement, which is proved very easily using the isoperimetric inequality on the sphere, plays a key role in some of the very basic theorems in this field.

We return to the question of changing intuition. The above paragraph shows that, for example, an ε-neighbourhood of the equator $x_1 = 0$ on S^{n-1} already contains an exponentially close to 1 part of the total measure of the sphere (since the sets $x_1 \leqslant 0$ and $x_1 \geqslant 0$ are both of measure $1/2$, and this set is the intersection of their ε-neighbourhoods). While this is again easy to prove (say, by integration) once it is observed, it does not correspond to our three-dimensional intuition. In particular, the far reaching consequences of these observations are hard to anticipate in advance. So, we see that in high dimension some of the intuition which we built for ourselves from what we know about three-dimensional space fails, and this "break" in intuition is the source of what one may call "surprising phenomena" in high dimensions. Of course, the surprise is there until intuition corrects itself, and the next surprise occurs only with the next break of intuition.

Here is a very simple example: The volume of the Euclidean ball B_2^n of radius one seems to be increasing with dimension. Indeed, denote this by κ_n and compute:

$$\kappa_1 = 2 < \kappa_2 = \pi < \kappa_3 = \frac{4\pi}{3} < \kappa_4 < \kappa_5 < \kappa_6.$$

However, a simple computation which is usually performed in Calculus III classes shows that

$$\mathrm{Vol}_n(B_2^n) = \kappa_n = \frac{\pi^{n/2}}{\Gamma(\frac{n}{2}+1)} = \left(c_n/\sqrt{n}\right)^n$$

where $c_n \to \sqrt{2\pi e}$. We thus see that in fact the volume of the Euclidean unit ball decreases like $n^{-n/2}$ with dimension (and one has the recursion formula $\kappa_n = \frac{2\pi}{n}\kappa_{n-2}$). So, for example, if one throws a point into the cube circumscribing the ball, at random, the chance that it will fall inside the ball, even in dimension 20, say, is practically zero. One cannot find this ball inside the cube.

Let us try to develop an intuition of high-dimensional spaces. We illustrate, with another example, how changing the intuition can help us understand, and anticipate, results. To begin, we should understand how to draw "high-dimensional" pictures, or, in other words, to try and imagine what do high-dimensional convex bodies "look like". The first non-intuitive fact is that the volume of parallel hypersections of a convex body decays exponentially after passing the median level (this is a consequence of the Brunn-Minkowski inequality, see Section 3.5). If we want to capture this property, it means that our two- or three-dimensional pictures of a high-dimensional convex body should have a "hyperbolic" form! Thus, K is a convex set but, as the rate of volume decay has a crucial influence on the geometry, we should find a way to visualize it in our pictures. For example, one may draw the convex set K as follows:

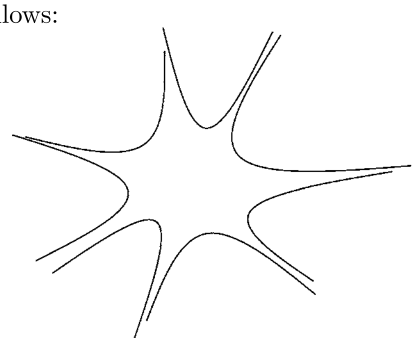

The convexity is no longer seen in the picture, but the volumetric properties are apparent. Next, with such a picture in mind, we may intuitively understand the following fact (it is a special case of Theorem 5.5.4 in Section 5.5): Consider the convex body $K = \sqrt{n}B_1^n := \mathrm{conv}(\pm\sqrt{n}e_i)$ (also called the unit ball of L_1^n). Take a random rotation UK of K and intersect it with the original body.

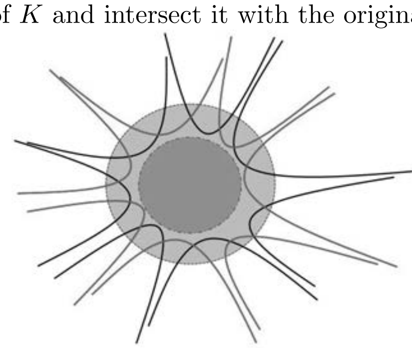

The resulting body, $K \cap UK$ is, with high probability over the choice of the random rotation U, contained in a Euclidean ball of radius C, where C is a universal constant independent of the dimension. Note that the original body, which contains a Euclidean ball of radius 1 (as does the intersection), has points in distance $\sqrt{n}$ from the origin. That is, the smallest Euclidean ball containing K is $\sqrt{n}B_2^n$. However, the simple (random) procedure of rotation and intersection, with high probability cuts out all these "remote regions" and regularizes the body completely so that it becomes an isomorphic Euclidean ball.

This was an example of a very concrete body, but it turns out the same property holds for a large class of bodies (called "finite volume ratio" bodies, see Section 5.5). Actually, if one allows more rotations, $\frac{n}{\log n}$ of them in dimension n, one may always regularize any body by the same process to become an isomorphic Euclidean ball. This last claim needs a small correction to be completely true: we have not explained how one chooses a random rotation. To this end one considers the Haar probability measure on the space of orthogonal rotations. To consider orthogonal rotations one must first fix a Euclidean structure, and the above statement is true after fixing the "right" structure corresponding to the body in question. The story of choosing a Euclidean structure, which is the same as choosing a "position" for the body, is an important topic, and for different goals different structures should be chosen. This topic is covered in Chapter 2.

Let us emphasize that while the geometric picture is what helps us understand which phenomena may occur, the picture is of course not a proof, and in each case a proof should be developed and is usually non-trivial.

This last example brings us to another important point which will be a central theme in this book, and this is the way in which, in this theory, *randomness* and *patterns* appear together. A perceived random nature of high dimensions is at the root of the reasons for the patterns produced and the unusual phenomena observed in high dimensions. In the dictionary, "randomness" is the exact opposite of "pattern". Randomness means "no pattern". But, in fact, objects created by independent identically distributed random processes, while being different from one another, are many times indistinguishable and similar in the statistical sense. Consider for example the unit cube, $[0,1]^n$. Choosing a random point inside it with respect to the uniform distribution means simply picking the n coordinates independently and uniformly at random in $[0,1]$. We know that such a point has some very special statistical properties (the simplest of which is the law of large numbers and the central limit theorem regarding the behaviour of the sum of these coordinates). It turns out that similar phenomena occur when the unit cube is replaced by a general convex body (again, a position should be specified). It is a challenge to uncover these similarities, a pattern, in very different looking objects. When we discover very similar patterns in arbitrary, and apparently very diverse convex bodies or normed spaces, we interpret them as a manifestation of the randomness principle mentioned above.

On the one hand, high dimension means many variables and many "possibilities", so one may expect an increase in the diversity and complexity as dimension increases. However, the concentration of measure and similar effects caused by the convexity assumption imply in fact a reduction of the diversity with increasing dimension, and the collapse of many different possibilities into one, or, in some cases, a few possibilities only. We quote yet another simple example which is a version of the "global Dvoretzky-type theorem". For details see Section 5.6. (The Minkowski sum of two sets is defined by $A + B = \{a + b : a \in A, b \in B\}$.)

> Let $n \in \mathbb{N}$ and let $K \subset \mathbb{R}^n$ be a convex body such that the Euclidean ball B_2^n is the ellipsoid of maximal volume inside K. Then, for $N = Cn/\log n$ random orthogonal transformations $U_i \in O(n)$, with probability at least $1 - e^{-cn}$ we have that

$$B_2^n \subset \frac{1}{N}\big(U_1K + U_2K + \cdots + U_NK\big) \subset C'B_2^n,$$

where $0 < c, C, C' < \infty$ are universal constants (independent of K and of n).

One way in which diversity is compensated and order is created in the mixture caused by high dimensionality, is the concentration of measure phenomenon. As the dimension n increases, the covering numbers of a generic body of the same volume as the unit Euclidean ball, say, by the Euclidean ball itself (this means the number of translates of the ball needed to cover the body, see Sections 4.1 and 4.2) become large, usually exponentially so, meaning e^{cn} for some constant $c > 0$, and so seem impossible to handle. The concentration of measure is, however, of exponential order too (this time $e^{-c'n}$ for some constant $c' > 0$), so that in the end proofs become a matter of comparison of different constants in the various exponents (this is, of course, a very simplistic description of what is going on).

Let us quote from the preface of P. Lévy to the second edition of his book of 1951 [**380**]:

> "It is quite paradoxical, that an increase in the number of variables might cause simplifications. In reality, any law of large numbers presupposes the existence of some rule governing the influence of sequential variables; starting with such a rule, we often obtain simple asymptotic results. Without such a rule, complete chaos ensues, and since we are unable to describe, for instance, an infinite sequence of numbers, without resorting to an exact rule, we are unable to find order in the chaos, where, as we know, one can find mysterious non-measurable sets, which we can never truly comprehend, but which nevertheless will not cease to exist."

As we shall see below, the above facts reflect the probabilistic nature of high dimensions. We mean by this more than just the fact that we are using probabilistic techniques in many steps of the proofs. Let us mention one more very concrete example to illustrate this "probabilistic nature": Assume you are given a body $K \subset \mathbb{R}^n$, and you know that there exist 3 orthogonal transformations $U_1, U_2, U_3 \in O(n)$ such that the intersection of U_1K, U_2K and U_3K is, up to constant 2, say, a Euclidean ball. Then, for a *random* choice of 10 rotations, $\{V_i\}_{i=1}^{10} \subset O(n)$, with high probability on their choice, one has that $\bigcap_{i=1}^{10} V_iK$ is up to constant C (which depends on the numbers 2, 3 and 10, not on the dimension n, and may be computed) a Euclidean ball. This is a manifestation of a principle which is sometimes called "random is the best", namely that in various situations the results obtained by a random method cannot be substantially improved if the random choice is replaced by the best choice for the specific goal.

There are a number of reasons for this observed ordered behaviour. One may mention "repetition", which creates order, as statistics demonstrates. What we explain here and shall see throughout the book is that very high dimensions, or more generally, high parametric families, are another source of order.

We mention at this point that historically we observe the study of finite, but very high-dimensional spaces and their asymptotic properties as dimension increases already in Minkowski's work, who for the purposes of analytic number theory considered n-dimensional space from a geometric point of view. Before him, as well as long after him, geometry had to be two- or three-dimensional, see, e.g., the works of Blaschke. A paper of von Neumann from 1942 also portrays the same asymptotic point of view. We quote below from Sections 4 and 5 of the introduction of [**601**]. Here E_n denotes n-dimensional Euclidean space and M_n denotes the space of all $n \times n$ matrices. Whatever is in brackets is the present authors' addition.

> "Our interest will be concentrated in this note on the conditions in E_n and M_n - mainly M_n - when n is *finite*, but *very great*. This is an approach to the study of the infinite dimensional, which differs essentially from the usual one. The usual approach consists in studying an actually infinite dimensional unitary space, i.e. the Hilbert space E_∞. We wish to investigate instead the *asymptotic* behaviour of E_n and M_n for finite n, when $n \to \infty$.
>
> We think that the latter approach has been unjustifiably neglected, as compared with the former one. It is certainly not contained in it, since it permits the use of the notions $\|A\|$ and $t(A)$ (normalized Hilbert Schmidt norm, and trace) which, owing to the factors $1/n$ appearing in (their definitions) possess no analogues in E_∞.
>
> Since Hilbert space E_∞ was conceived as a limiting case of the E_n for $n \to \infty$, we feel that such a study is necessary in order to clarify to what extent E_∞ is or is not the only possible limiting case. Indeed we think that it is not, and that investigations on operator rings by F. J. Murray and the author show that other limiting cases exist, which under many aspects are more natural ones.
>
> Our present investigations originated in fact mainly from the desire to solve certain questions... We hope, however, that the reader will find that they also have an interest of their own, mainly in the sense indicated above: as a study of the asymptotic behaviour of E_n and M_n for finite n, when $n \to \infty$.
>
> From the point of view described (above) it seems natural to ask this question: How much does the character of E_n and M_n change when n increases - especially if n has already assumed very great values?"

Let us turn to a short description of the various chapters of the book; this will give us the opportunity to comment on additional fundamental ideas of the theory.

In Chapter 1 we recall basic notions from classical convexity. In fact, a relatively large portion of this book is dedicated to convexity theory, since a large part of the development of asymptotic geometric analysis is connected strongly with the classical theory. We present several proofs of the Brunn-Minkowski inequality and some of its fundamental applications. We have chosen to discuss in detail those proofs as they allow us to introduce fruitful ideas which we shall revisit throughtout the book. In the appendices we provide a more detailed exposition of basic facts from elementary convexity, convex analysis and the theory of mixed volumes. In particular, we describe the proof of Minkowski's theorem on the polynomiality of the volume of the sum of compact convex sets, and of the Alexandrov-Fenchel inequality, one of the most beautiful, non-trivial and profound theorems in convexity, which is linked with algebraic geometry and number theory. We emphasize the functional analytic point of view into classical convexity. This point of view opened a new field which is sometimes called "*functionalization of geometry*" or "*geometrization of probability*": It turns out that almost any notion or inequality connected with convex bodies has an analogous notion or inequality in the world of convex *functions*. This analogy between bodies and functions is fruitful in many different ways. On the one hand, it allows to predict functional inequalities which then are interesting in their own right. On the other hand, the generalization into the larger world of convex functions enables one to see the bigger picture and better understand what is going on. Finally, the results for functions may sometimes have implications back in the convex bodies world. This general idea is considered in parallel with the classical theory throughout the book.

In Chapter 2 we introduce the most basic and classical positions of convex bodies: Given a convex body K in $\mathbb{R}^n$, the family of its *positions* is the family of its affine images $\{x_0 + T(K)\}$ where $x_0 \in \mathbb{R}^n$ and $T \in GL_n$. In the context of functional analysis, one is given a norm (whose unit ball is K) and the choice of a position reflects a choice of a Euclidean structure for the linear space $\mathbb{R}^n$. Note that the choice of a Euclidean structure specifies a unit ball of the Euclidean norm, which is an ellipsoid. Thus, we may equivalently see a "position" as a choice of a special ellipsoid. The different ellipsoids connected with a convex body (or the different positions, corresponding to different choices of a Euclidean structure) that we consider in this chapter reflect different traces of symmetries which the convex body has. We introduce John position (also called maximal volume ellipsoid position), minimal surface area position and minimal mean width position. It turns out that when a position is extremal then some differential must vanish, and its vanishing is connected with isotropicity of some connected measure.

We also discuss some applications, mainly of John position, and introduce a main tool, which is useful in many other results in the theory, called the Brascamp-Lieb inequality. We state and prove one of its most useful forms, which is the so-called "normalized form" put forward by K. Ball, together with its reverse form, using F. Barthe's transportation of measure argument. In the second volume of this book we shall discuss the general form of the Brascamp-Lieb inequality, its various versions, proofs, and reverse form, as well as further applications to convex geometric analysis.

In Chapter 3 we discuss the concentration of measure phenomenon, first put forward in V. Milman's version of Dvoretzky theorem. Concentration is the central phenomenon that is responsible for the main results in this book. We present a number of approaches, all leading to the same type of behaviour: in high parametric families, under very weak assumptions of various types, a function tends to concentrate around its mean or median. Classical isoperimetric inequalities for metric probability spaces, such as the sphere, Gauss space and the discrete cube, are at the origin of measure concentration, and we start our exposition with these examples. Once the extremal sets (the solutions of the isoperimetric problem) are known, concentration inequalities come as a consequence of a simple computation. However, in very few examples are the exteremal sets known. We therefore do not focus on extremal sets but mainly on different ways to get concentration inequalities. We explore various such ways, and determine the different sources for concentration. In the second volume of this book we shall come back to this subject and study its functional aspects: Sobolev and logarithmic Sobolev inequalities, tensorization techniques, semi-group approaches, Laplace transform and infimum convolutions, and investigate in more detail the subject of transportation of measure.

In Chapter 4 we introduce the covering numbers $N(A, B)$ and the entropy numbers $e_k(A, B)$ as a way of measuring the "size" of a set A in terms of another set B. As we will see in the next chapters, they are a very useful tool and play an important role in the theory. Here, we explain some of their properties, derive relations and duality between these numbers, and estimate them in terms of other parameters of the sets involved — estimates which shall be useful in the sequel.

Chapter 5 is the starting point for our exposition of the asymptotic theory of convex bodies. It is devoted to the Dvoretzky-Milman theorem and to the main developments around it. In geometric language the theorem states that every high-dimensional centrally symmetric convex body has central sections of high dimension which are almost ellipsoidal. The dependence of the dimension k of these sections on the dimension n of the body is as follows: for every n-dimensional normed space $X = (\mathbb{R}^n, \|\cdot\|)$ and every $\varepsilon \in (0, 1)$ there exist an integer $k \geqslant c\varepsilon^2 \log n$ and a k-dimensional subspace F of X which satisfies $d_{BM}(F, \ell_2^k) \leqslant 1+\varepsilon$, where d_{BM} denotes Banach-Mazur distance, a natural geometric distance between two normed spaces, and c is some absolute constant. The proof of the Dvoretzky-Milman theorem exploits the concentration of measure phenomenon for the Euclidean sphere S^{n-1}, in the form of a deviation inequality for Lipschitz functions $f : S^{n-1} \to \mathbb{R}$, which implies that the values of $\|\cdot\|$ on S^{n-1} concentrate near their average

$$M = \int_{S^{n-1}} \|x\|\, d\sigma(x).$$

A remarkable fact is that in Milman's proof, a formula for such a k is given in terms of n, M and the Lipschitz constant (usually called b) of the norm, and that this formula turns out to be sharp (up to a universal constant) in full generality. This gives us the opportunity to introduce one more new idea of the theory, which is *universality.* In different fields, and also in the origins of asymptotic geometric analysis, for a long time one knew how to write very precise estimates reflecting different asymptotic behaviour of certain specific high-dimensional (or high parametric) objects (say, for the spaces ℓ_p^n). Usually, one could show that these estimates are sharp, in an isomorphic sense at least. However, an accumulation of results indicates that, in fact, available estimates are exact for *every* sequence of spaces in increasing dimension (and thus one is tempted to say "for every space"). These kinds of estimates are called "asymptotic formulae". Let us demonstrate another such formula, concerning the diameter of a random projection of a convex body. All constants appearing in the statement below (C, c_1, C_2, c') are universal and do not depend on the body or on the dimension. Let $K \subset \mathbb{R}^n$ be a centrally symmetric convex body. One denotes by $h_K(u)$ the support function of K in direction u, defined as half the width of the minimal slab orthogonal to u which includes K, that is,

$$h_K(u) = \max\{\langle x, u\rangle : x \in K\}.$$

Denote by $d = d(K)$ the smallest constant such that $K \subset dB_2^n$, that is, half of the diameter of K, and actually $d = \max_{u\in S^{n-1}} h_K(u)$. Denote by $M^* = M^*(K)$ the average of h_K over S^{n-1}, that is,

$$M^*(K) = \int_{S^{n-1}} h_K(u) d\sigma(u)$$

where σ is the Haar probability measure on S^{n-1}. It turns out that for dimensions larger than $k^* = C(M^*/d)^2 n$, the diameter of the projection of K onto a random k-dimensional subspace is, with high probability, approximately $d\sqrt{k/n}$. That is, between $c_1 d\sqrt{k/n}$ and $C_2 d\sqrt{k/n}$. Around the critical dimension $k^* = k^*(K)$, the projection becomes already (with high probability on the choice of a subspace) a Euclidean ball of radius approximately $M^*(K)$, and this will be, again up to constants, the diameter (and the inner-radius) of a random projection onto dimension $c'k^*$ and less. In this result the isomorphic nature of the result is very apparent. Indeed, the diameter need not be ε-isometrically close to $d\sqrt{k/n}$ for k in the range between k^* and n, but only isomorphically. Isometric results are known in the regime $k \leqslant c'k^*$ when the projection is already with high probability a Euclidean ball (this is actually the Dvoretzky-Milman theorem). We describe this result in detail in Section 5.7.1. Another property of this last example is a threshold behaviour of the function $f(t)$ giving the average diameter of a projection into dimension tn. The function, which is monotone, attains its maximum, d, at $t = 1$, behaves like $d\sqrt{t}$ in the range $[C(M^*/d)^2, 1]$, and like a constant, close to M^*, in the range $[0, c'(M^*/d)^2]$. Threshold phenomena have been known for a long time in many areas of mathematics, for example, in mathematical physics. Here we see that these occur in complete generality (for *any* convex body, the same type of threshold). More examples of threshold behaviour in asymptotic geometric analysis shall be demonstrated in the book.

Before moving to the description of Chapters 6–10 we mention another point of view one should keep in mind when reading the book: the comparison between *local* and *global* type results. The careful readers may have already noted the similarity of two of the statements given so far in this preface: a part of the statement about decrease of diameter in fact said that after some critical dimension, a random projection of a convex body is with high probability close to a Euclidean ball (this also follows from the Dvoretzky-Milman theorem by duality of projections and sections). This is called a “local” statement. Two other theorems quoted above regarded what happens when one intersects random rotations of a convex body (for example, B_1^n), or when one takes the Minkowski sum (average) of random rotations of a convex body (for example, the cube). Again the results were that after a suitable (and not very large) number of such rotations, the resulting body is an isomorphic Euclidean ball. These types of results, pertaining to the body as a whole and not its sections or projections, are called “global” results. At the heart of the global results presented in this book, which have convex geometric flavor, stand methods which come from functional analysis (considering norms, their averages, etc). Again, by global properties we refer to properties of the original body or norm in question, while the local properties pertain to the structure of lower dimensional sections and projections of the body or normed space. From the beginning of the 1970’s the need for geometric functional analysis led to a deep investigation of the linear structure of finite dimensional normed spaces (starting with Dvoretzky theorem). However, it had to develop a long way before this structure was understood well enough to be used for the study of the global properties of a space. The culmination of this study was an understanding of the fact that subspaces (and quotient spaces) of proportional dimension behave very predictably. An example is the theorem quoted above regarding the decay of diameter. This understanding formed a bridge between the problems of functional analysis and the global asymptotic properties of convex sets, and is the reason the two fields of convexity and of functional analysis work together nowadays.

In Chapter 6 we discuss upper bounds for the parameter $M(K)M^*(K)$, or equivalently, the product of the mean width of K and the mean width of its polar, the main goal being to minimize this parameter over all positions of the convex body. (The polar of a convex body K is the closed convex set generating the norm given by h_K, and is denoted K°.) We will see that the quantity MM^* can be bounded from above by a parameter of the space $(X, \|\cdot\|_K)$ which is called its K-convexity constant, and which in turn can be bounded from above, for X of dimension n, by $c[\log(d_{BM}(X, \ell_2^n)) + 1] \leqslant c' \log n$ for universal c, c'. This estimate for the K-convexity constant is due to G. Pisier and as we will see it is one of the fundamental facts in the asymptotic theory. The estimate for $M(K)M^*(K)$ brings us to one more main point, which concerns *duality*, or *polarity*. In many situations two dual operations performed one after the other already imply complete regularization. That is, one operation cancels a certain type of "bad behaviour", and the dual operation cancels the "opposite" bad behaviour. Other examples include the quotient of a subspace theorem (see Chapter 7) or its corresponding global theorem: if one takes the sum of a body and a random (in the right coordinate system) rotation of it, then considers the polar of this set, to which again one applies a random rotation and takes the sum, the resulting body will be with high probability on the choice of rotations, an isomorphic Euclidean ball. If one uses just one of these two operations, there may be a need for $n/\log n$ such operations.

Chapter 7 is devoted to results about proportional subspaces and quotients of an n-dimensional normed space, i.e., of dimension λn, where the "proportion" $\lambda \in (0,1)$ can sometimes be very close to 1. The first step in this direction is Milman's M^*-estimate. In a geometric language, it says that there exists a function $f : (0,1) \to \mathbb{R}^+$ such that, for every centrally symmetric convex body K in $\mathbb{R}^n$ and every $\lambda \in (0,1)$, a random $\lfloor \lambda n \rfloor$-dimensional section $K \cap F$ of K satisfies the inclusion

$$K \cap F \subseteq \frac{M^*(K)}{f(\lambda)} B_2^n \cap F.$$

In other words, the diameter of a random "proportional section" of a high-dimensional centrally symmetric convex body K is controlled by the mean width $M^*(K)$ of the body. We present several proofs of the M^*-estimate; based on these, we will be able to say more about the best possible function f for which the theorem holds true and about the corresponding estimate for the probability of subspaces in which this occurs. As an application of the M^* estimate we obtain Milman's quotient of a subspace theorem. We also complement the M^* estimate by a lower bound for the outer-radius of sections of K, which holds for *all* subspaces, we compare "best" sections with "random" ones of slightly lower dimension, and we provide a linear relation between the outer-radius of a section of K and the outer-radius of a section of K°.

In Chapter 8 we present one of the deepest results in asymptotic geometric analysis: the existence of an M-position for every convex body K. This position can be described "isometrically" (if, say, K has volume 1) as minimizing the volume of $T(K) + B_2^n$ over all $T \in SL_n$. However, such a characterization hides its main properties and advantages that are in fact of an "isomorphic" nature. The isomorphic formulation of the result states that there exists an ellipsoid of the same volume as the body K, which can replace K, in many computations, up to universal constants. This result, which was discovered by V. Milman, leads to the reverse Santaló inequality and the reverse Brunn-Minkowski inequality. The reverse Santaló inequality concerns the volume product, sometimes called the Mahler product, of K which is defined by

$$s(K) := \mathrm{Vol}_n(K)\mathrm{Vol}_n(K^\circ).$$

The classical Blaschke-Santaló inequality states that, given a centrally symmetric convex body K in $\mathbb{R}^n$, the volume product $s(K)$ is less than or equal to the volume product $s(B_2^n) = \kappa_n^2$, and that equality holds if and only if K is an ellipsoid. In the opposite direction, a well-known conjecture of Mahler states that $s(K) \geqslant 4^n/n!$ for every centrally symmetric convex body K (i.e., the cube is a minimizer for $s(K)$ among centrally symmetric convex bodies) and that $s(K) \geqslant (n+1)^{n+1}/(n!)^2$ in the not necessarily symmetric case, meaning that in this case the simplex is a minimizer. The reverse Santaló inequality of Bourgain and Milman verifies this conjecture in the asymptotic sense: there exists an absolute constant $c > 0$ such that

$$\left(\frac{s(K)}{s(B_2^n)}\right)^{1/n} \geqslant c$$

for every centrally symmetric convex body K in $\mathbb{R}^n$. Milman's reverse Brunn-Minkowski inequality states that for any pair of convex bodies K and T that are in M-position, one has

$$\mathrm{Vol}_n(K+T)^{1/n} \leqslant C\left[\mathrm{Vol}_n(K)^{1/n} + \mathrm{Vol}_n(T)^{1/n}\right].$$

(The reverse inequality, with constant 1, is simply the Brunn-Minkowski inequality of Chapter 1.)

Another way to define the M-position of a convex body is through covering numbers, as was presented in Milman's proof. Pisier has proposed a different approach to these results, which allows one to find a whole family of special M-ellipsoids satisfying stronger entropy estimates. We describe his approach in the last part of Chapter 8.

In Chapter 9 we introduce a "Gaussian approach" to some of the main results which were presented in previous chapters, including sharp versions of the Dvoretzky-Milman theorem and of the M^*-estimate. The proof of these results is based on comparison principles for Gaussian processes, due to Gordon, which extend a theorem of Slepian. The geometric study of random processes, and especially of Gaussian processes, has strong connections with asymptotic geometric analysis. The tools presented in this chapter will appear again in the second volume of the book.

In the last chapter of this volume, Chapter 10, we discuss more recent discoveries on the distribution of volume in high-dimensional convex bodies, together with the unresolved "slicing problem" which was mentioned briefly at the beginning of this preface, with some of its equivalent formulations. A natural framework for this study is the *isotropic position* of a convex body: a convex body $K \subset \mathbb{R}^n$ is called isotropic if $\mathrm{Vol}_n(K) = 1$, its barycenter (center of mass) is at the origin and its inertia matrix is a multiple of the identity, that is, there exists a constant $L_K > 0$

such that

$$\int_K \langle x,\theta\rangle^2 dx = L_K^2$$

for every θ in the Euclidean unit sphere S^{n-1}. The number L_K is then called the isotropic constant of K. The isotropic position arose from classical mechanics back in the 19$^{\text{th}}$ century. It has a useful characterization as a solution of an extremal problem: the isotropic position $\tilde{K} = T(K)$ of K minimizes the quantity

$$\int_{\tilde{K}} |x|^2 dx$$

over all $T \in GL_n$ such that $\mathrm{Vol}_n(\tilde{K}) = 1$ and $\int_{\tilde{K}} x dx = 0$.

The central theme in Chapter 10 is the *hyperplane conjecture* (or *slicing problem*): it asks whether there exists an absolute constant $c > 0$ such that $\max_{\theta\in S^{n-1}} \mathrm{Vol}_{n-1}(K\cap\theta^{\perp}) \geqslant c$ for every n and every convex body K of volume 1 in $\mathbb{R}^n$ with barycenter at the origin. We will see that an affirmative answer to this question is equivalent to the fact that there exists an absolute constant $C > 0$ such that

$$L_n := \max\{L_K : K \text{ is an isotropic convex body in } \mathbb{R}^n\} \leqslant C.$$

We shall work in the more general setting of a finite log-concave measure μ, where a corresponding notion of isotropicity is defined via the covariance matrix $\mathrm{Cov}(\mu)$ of μ. We present the best known upper bounds for L_n. Around 1985-86, Bourgain obtained the upper bound $L_n \leqslant c\sqrt[4]{n}\log n$ and, in 2006, this estimate was improved by Klartag to $L_n \leqslant c\sqrt[4]{n}$. In fact, Klartag obtained a solution to an isomorphic version of the hyperplane conjecture, the "isomorphic slicing problem", by showing that, for every convex body K in $\mathbb{R}^n$ and any $\varepsilon \in (0,1)$, one can find a centered convex body $T \subset \mathbb{R}^n$ and a point $x \in \mathbb{R}^n$ such that $(1+\varepsilon)^{-1}T \subseteq K + x \subseteq (1+\varepsilon)T$ and $L_T \leqslant C/\sqrt{\varepsilon}$ for some absolute constant $C > 0$. An additional essential ingredient in Klartag's proof of the bound $L_n \leqslant c\sqrt[4]{n}$, which is a beautiful and important result in its own right, is the following very useful deviation inequality of Paouris: if μ is an isotropic log-concave probability measure on $\mathbb{R}^n$, then

$$\mu(\{x \in \mathbb{R}^n : |x| \geqslant ct\sqrt{n}\}) \leqslant \exp\left(-t\sqrt{n}\right)$$

for every $t \geqslant 1$, where $c > 0$ is an absolute constant. The proof is presented in Section 10.4 along with the basic theory of the L_q-centroid bodies of an isotropic log-concave measure. Another important result regarding isotropic log-concave measures is the *central limit theorem* of Klartag, which states that the 1-dimensional marginals of high-dimensional isotropic log-concave measures μ are approximately Gaussian with high probability. We will come back to this result and related ones in the second volume of the book and we will see that precise quantitative relations exist between the hyperplane conjecture, the optimal answer to the central limit problem, and other conjectures regarding volume distribution in high dimensions.

Acknowledgements. This book is based on material gathered over a long period of time with the aid of many people. We would like to mention two ongoing working seminars in which many of the ideas and results were presented and discussed: these are the Asymptotic Geometric Analysis seminars at the University of Athens and at Tel Aviv University. The active participation of faculty members, students and visitors in these seminars, including many discussions and collaborations, have made a large contribution to the possibility of this book. We would like to mention the names of some people whose contribution was especially important, whether in offering us mathematical and technical advice, in reading specific chapters of the book, in allowing us to make use of their research notes and material, and for sending us to correct and less known references and sources. We thank S. Alesker, S. Bobkov, D. Faifman, B. Klartag, H. König, A. Litvak, G. Pisier, R. Schneider, B. Slomka, S. Sodin and B. Vritsiou. Finally, we would like to thank S. Gelfand and the AMS team for offering their publishing house as a home for this manuscript, and for encouraging us to complete this project.

The second named author would like to acknowledge partial support from the ARISTEIA II programme of the General Secretariat of Research and Technology of Greece during the final stage of this project. The first and third named authors would like to acknowledge partial support from the Israel Science Foundation.

November 2014