Flavors of Geometry
MSRI Publications
Volume **31**, 1997

An Elementary Introduction to Modern Convex Geometry

KEITH BALL

CONTENTS

Preface

These notes are based, somewhat loosely, on three series of lectures given by myself, J. Lindenstrauss and G. Schechtman, during the Introductory Workshop in Convex Geometry held at the Mathematical Sciences Research Institute in Berkeley, early in 1996. A fourth series was given by B. Bollobás, on rapid mixing and random volume algorithms; they are found elsewhere in this book.

The material discussed in these notes is not, for the most part, very new, but the presentation has been strongly influenced by recent developments: among other things, it has been possible to simplify many of the arguments in the light of later discoveries. Instead of giving a comprehensive description of the state of the art, I have tried to describe two or three of the more important ideas that have shaped the modern view of convex geometry, and to make them as accessible as possible to a broad audience. In most places, I have adopted an informal style that I hope retains some of the spontaneity of the original lectures. Needless to say, my fellow lecturers cannot be held responsible for any shortcomings of this presentation.

I should mention that there are large areas of research that fall under the very general name of convex geometry, but that will barely be touched upon in these notes. The most obvious such area is the classical or "Brunn–Minkowski" theory, which is well covered in [Schneider 1993]. Another noticeable omission is the combinatorial theory of polytopes: a standard reference here is [Brøndsted 1983].

Lecture 1. Basic Notions

The topic of these notes is convex geometry. The objects of study are convex bodies: compact, convex subsets of Euclidean spaces, that have nonempty interior. Convex sets occur naturally in many areas of mathematics: linear programming, probability theory, functional analysis, partial differential equations, information theory, and the geometry of numbers, to name a few.

Although convexity is a simple property to formulate, convex bodies possess a surprisingly rich structure. There are several themes running through these notes: perhaps the most obvious one can be summed up in the sentence: "All convex bodies behave a bit like Euclidean balls." Before we look at some ways in which this is true it is a good idea to point out ways in which it definitely is not. This lecture will be devoted to the introduction of a few basic examples that we need to keep at the backs of our minds, and one or two well known principles.

The only notational conventions that are worth specifying at this point are the following. We will use $|\cdot|$ to denote the standard Euclidean norm on $\mathbb{R}^n$. For a body K, $\mathrm{vol}(K)$ will mean the volume measure of the appropriate dimension.

The most fundamental principle in convexity is the *Hahn–Banach separation theorem*, which guarantees that each convex body is an intersection of half-spaces, and that at each point of the boundary of a convex body, there is at least one supporting hyperplane. More generally, if K and L are disjoint, compact, convex subsets of $\mathbb{R}^n$, then there is a linear functional $\phi : \mathbb{R}^n \to \mathbb{R}$ for which $\phi(x) < \phi(y)$ whenever $x \in K$ and $y \in L$.

The simplest example of a convex body in $\mathbb{R}^n$ is the cube, $[-1, 1]^n$. This does not look much like the Euclidean ball. The largest ball inside the cube has radius 1, while the smallest ball containing it has radius $\sqrt{n}$, since the corners of the cube are this far from the origin. So, as the dimension grows, the cube resembles a ball less and less.

The second example to which we shall refer is the n-dimensional regular solid simplex: the convex hull of $n+1$ equally spaced points. For this body, the ratio of the radii of inscribed and circumscribed balls is n: even worse than for the cube. The two-dimensional case is shown in Figure 1. In Lecture 3 we shall see

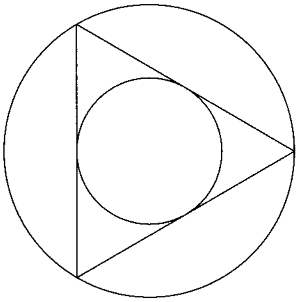

Figure 1. Inscribed and circumscribed spheres for an n-simplex.

that these ratios are extremal in a certain well-defined sense.

Solid simplices are particular examples of cones. By a *cone* in $\mathbb{R}^n$ we just mean the convex hull of a single point and some convex body of dimension $n-1$ (Figure 2). In $\mathbb{R}^n$, the volume of a cone of height h over a base of $(n-1)$-dimensional volume B is Bh/n.

The third example, which we shall investigate more closely in Lecture 4, is the n-dimensional "octahedron", or *cross-polytope:* the convex hull of the $2n$ points $(\pm 1, 0, 0, \ldots, 0)$, $(0, \pm 1, 0, \ldots, 0)$, $\ldots$, $(0, 0, \ldots, 0, \pm 1)$. Since this is the unit ball of the ℓ_1 norm on $\mathbb{R}^n$, we shall denote it B_1^n. The circumscribing sphere of B_1^n has radius 1, the inscribed sphere has radius $1/\sqrt{n}$; so, as for the cube, the ratio is $\sqrt{n}$: see Figure 3, left.

B_1^n is made up of 2^n pieces similar to the piece whose points have nonnegative coordinates, illustrated in Figure 3, right. This piece is a cone of height 1 over a base, which is the analogous piece in $\mathbb{R}^{n-1}$. By induction, its volume is

$$\frac{1}{n} \cdot \frac{1}{n-1} \cdot \cdots \cdot \frac{1}{2} \cdot 1 = \frac{1}{n!},$$

and hence the volume of B_1^n is $2^n/n!$.

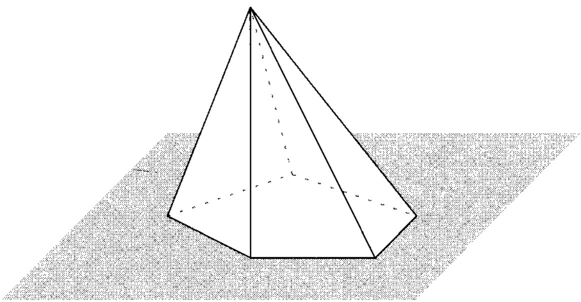

Figure 2. A cone.

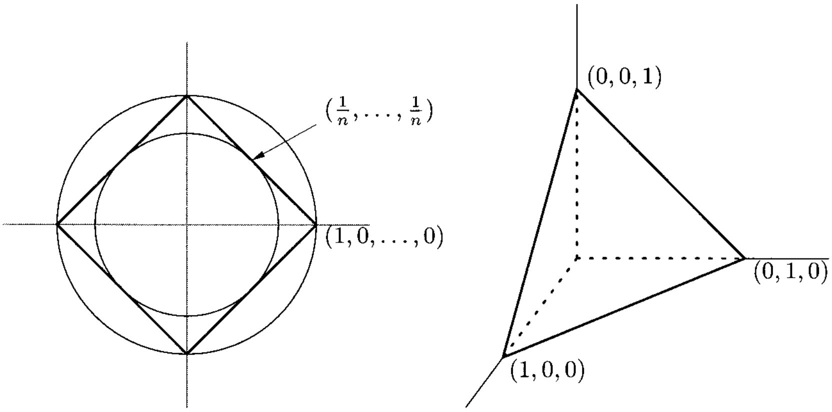

Figure 3. The cross-polytope (left) and one orthant thereof (right).

The final example is the Euclidean ball itself,

$$B_2^n = \left\{ x \in \mathbb{R}^n : \sum_1^n x_i^2 \le 1 \right\}.$$

We shall need to know the volume of the ball: call it v_n. We can calculate the surface "area" of B_2^n very easily in terms of v_n: the argument goes back to the ancients. We think of the ball as being built of thin cones of height 1: see Figure 4, left. Since the volume of each of these cones is $1/n$ times its base area, the surface of the ball has area nv_n. The sphere of radius 1, which is the surface of the ball, we shall denote S^{n-1}.

To calculate v_n, we use integration in spherical polar coordinates. To specify a point x we use two coordinates: r, its distance from 0, and θ, a point on the sphere, which specifies the direction of x. The point θ plays the role of $n-1$ real coordinates. Clearly, in this representation, $x = r\theta$: see Figure 4, right. We can

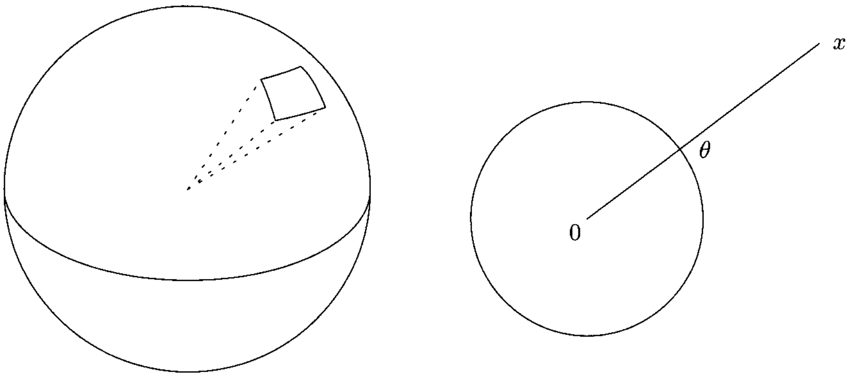

Figure 4. Computing the volume of the Euclidean ball.

write the integral of a function on $\mathbb{R}^n$ as

$$\int_{\mathbb{R}^n} f = \int_{r=0}^{\infty} \int_{S^{n-1}} f(r\theta) \text{ “}d\theta\text{” } r^{n-1}\, dr. \tag{1.1}$$

The factor r^{n-1} appears because the sphere of radius r has area r^{n-1} times that of S^{n-1}. The notation “$d\theta$” stands for “area” measure on the sphere: its total mass is the surface area nv_n. The most important feature of this measure is its rotational invariance: if A is a subset of the sphere and U is an orthogonal transformation of $\mathbb{R}^n$, then UA has the same measure as A. Throughout these lectures we shall normalise integrals like that in (1.1) by pulling out the factor nv_n, and write

$$\int_{\mathbb{R}^n} f = nv_n \int_0^{\infty} \int_{S^{n-1}} f(r\theta) r^{n-1}\, d\sigma(\theta)\, dr$$

where $\sigma = \sigma_{n-1}$ is the rotation-invariant measure on S^{n-1} of total mass 1. To find v_n, we integrate the function

$$x \mapsto \exp\left(-\tfrac{1}{2} \sum_1^n x_i^2\right)$$

both ways. This function is at once invariant under rotations and a product of functions depending upon separate coordinates; this is what makes the method work. The integral is

$$\int_{\mathbb{R}^n} f = \int_{\mathbb{R}^n} \prod_1^n e^{-x_i^2/2}\, dx = \prod_1^n \left(\int_{-\infty}^{\infty} e^{-x_i^2/2}\, dx_i \right) = \left(\sqrt{2\pi}\,\right)^n.$$

But this equals

$$nv_n \int_0^{\infty} \int_{S^{n-1}} e^{-r^2/2} r^{n-1}\, d\sigma\, dr = nv_n \int_0^{\infty} e^{-r^2/2} r^{n-1}\, dr = v_n 2^{n/2} \Gamma\left(\frac{n}{2}+1\right).$$

Hence

$$v_n = \frac{\pi^{n/2}}{\Gamma\left(\frac{n}{2}+1\right)}.$$

This is extremely small if n is large. From Stirling's formula we know that

$$\Gamma\left(\frac{n}{2}+1\right) \sim \sqrt{2\pi}\, e^{-n/2} \left(\frac{n}{2}\right)^{(n+1)/2},$$

so that v_n is roughly

$$\left(\sqrt{\frac{2\pi e}{n}}\right)^n.$$

To put it another way, the Euclidean ball of *volume* 1 has *radius* about

$$\sqrt{\frac{n}{2\pi e}},$$

which is pretty big.

Cambridge University Press
978-0-521-62962-1 - Flavors of Geometry
Edited by Silvio Levy
Excerpt
More information

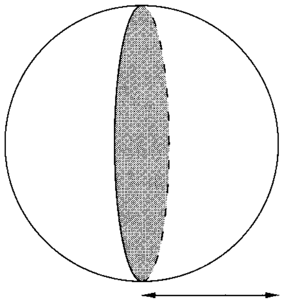

Figure 5. Comparing the volume of a ball with that of its central slice.

This rather surprising property of the ball in high-dimensional spaces is perhaps the first hint that our intuition might lead us astray. The next hint is provided by an answer to the following rather vague question: how is the mass of the ball distributed? To begin with, let's estimate the $(n-1)$-dimensional volume of a slice through the centre of the ball of volume 1. The ball has radius

$$r = v_n^{-1/n}$$

(Figure 5). The slice is an $(n-1)$-dimensional ball of this radius, so its volume is

$$v_{n-1} r^{n-1} = v_{n-1} \left(\frac{1}{v_n} \right)^{(n-1)/n}.$$

By Stirling's formula again, we find that the slice has volume about $\sqrt{e}$ when n is large. What are the $(n-1)$-dimensional volumes of parallel slices? The slice at distance x from the centre is an $(n-1)$-dimensional ball whose radius is $\sqrt{r^2 - x^2}$ (whereas the central slice had radius r), so the volume of the smaller slice is about

$$\sqrt{e} \left(\frac{\sqrt{r^2 - x^2}}{r} \right)^{n-1} = \sqrt{e} \left(1 - \frac{x^2}{r^2} \right)^{(n-1)/2}.$$

Since r is roughly $\sqrt{n/(2\pi e)}$, this is about

$$\sqrt{e} \left(1 - \frac{2\pi e x^2}{n} \right)^{(n-1)/2} \approx \sqrt{e} \exp(-\pi e x^2).$$

Thus, if we project the mass distribution of the ball of volume 1 onto a single direction, we get a distribution that is approximately Gaussian (normal) with variance $1/(2\pi e)$. What is remarkable about this is that the variance does not depend upon n. Despite the fact that the radius of the ball of volume 1 grows like $\sqrt{n/(2\pi e)}$, almost all of this volume stays within a slab of fixed width: for example, about 96% of the volume lies in the slab

$$\{x \in \mathbb{R}^n : -\tfrac{1}{2} \le x_1 \le \tfrac{1}{2}\}.$$

See Figure 6.

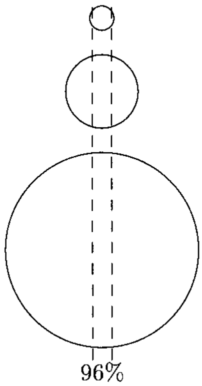

Figure 6. Balls in various dimensions, and the slab that contains about 96% of each of them.

So the volume of the ball concentrates close to *any* subspace of dimension $n-1$. This would seem to suggest that the volume concentrates near the centre of the ball, where the subspaces all meet. But, on the contrary, it is easy to see that, if n is large, most of the volume of the ball lies near its surface. In objects of high dimension, measure tends to concentrate in places that our low-dimensional intuition considers small. A considerable extension of this curious phenomenon will be exploited in Lectures 8 and 9.

To finish this lecture, let's write down a formula for the volume of a general body in spherical polar coordinates. Let K be such a body with 0 in its interior, and for each direction $\theta \in S^{n-1}$ let $r(\theta)$ be the radius of K in this direction. Then the volume of K is

$$nv_n \int_{S^{n-1}} \int_0^{r(\theta)} s^{n-1}\, ds\, d\sigma = v_n \int_{S^{n-1}} r(\theta)^n\, d\sigma(\theta).$$

This tells us a bit about particular bodies. For example, if K is the cube $[-1,1]^n$, whose volume is 2^n, the radius satisfies

$$\int_{S^{n-1}} r(\theta)^n = \frac{2^n}{v_n} \approx \left(\sqrt{\frac{2n}{\pi e}}\right)^n.$$

So the "average" radius of the cube is about

$$\sqrt{\frac{2n}{\pi e}}.$$

This indicates that the volume of the cube tends to lie in its corners, where the radius is close to $\sqrt{n}$, not in the middle of its facets, where the radius is close to 1. In Lecture 4 we shall see that the reverse happens for B_1^n and that this has a surprising consequence.

If K is (*centrally*) *symmetric*, that is, if $-x \in K$ whenever $x \in K$, then K is the unit ball of some norm $\|\cdot\|_K$ on $\mathbb{R}^n$:

$$K = \{x : \|x\|_K \le 1\}.$$

This was already mentioned for the octahedron, which is the unit ball of the ℓ_1 norm

$$\|x\| = \sum_1^n |x_i|.$$

The norm and radius are easily seen to be related by

$$r(\theta) = \frac{1}{\|\theta\|}, \quad \text{for } \theta \in S^{n-1},$$

since $r(\theta)$ is the largest number r for which $r\theta \in K$. Thus, for a general symmetric body K with associated norm $\|\cdot\|$, we have this formula for the volume:

$$\text{vol}(K) = v_n \int_{S^{n-1}} \|\theta\|^{-n}\, d\sigma(\theta).$$

Lecture 2. Spherical Sections of the Cube

In the first lecture it was explained that the cube is rather unlike a Euclidean ball in $\mathbb{R}^n$: the cube $[-1,1]^n$ includes a ball of radius 1 and no more, and is included in a ball of radius $\sqrt{n}$ and no less. The cube is a bad approximation to the Euclidean ball. In this lecture we shall take this point a bit further. A body like the cube, which is bounded by a finite number of flat facets, is called a *polytope*. Among symmetric polytopes, the cube has the fewest possible facets, namely $2n$. The question we shall address here is this:

If K is a polytope in $\mathbb{R}^n$ with m facets, how well can K approximate the Euclidean ball?

Let's begin by clarifying the notion of approximation. To simplify matters we shall only consider symmetric bodies. By analogy with the remarks above, we could define the distance between two convex bodies K and L to be the smallest d for which there is a scaled copy of L inside K and another copy of L, d times as large, containing K. However, for most purposes, it will be more convenient to have an affine-invariant notion of distance: for example we want to regard all parallelograms as the same. Therefore:

DEFINITION. The distance $d(K, L)$ between symmetric convex bodies K and L is the least positive d for which there is a linear image $\tilde{L}$ of L such that $\tilde{L} \subset K \subset d\tilde{L}$. (See Figure 7.)

Note that this distance is multiplicative, not additive: in order to get a metric (on the set of linear equivalence classes of symmetric convex bodies) we would need to take $\log d$ instead of d. In particular, if K and L are identical then $d(K, L) = 1$.

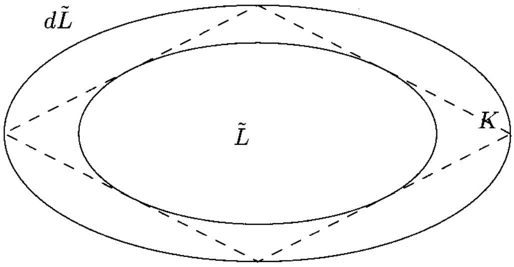

Figure 7. Defining the distance between K and L.

Our observations of the last lecture show that the distance between the cube and the Euclidean ball in $\mathbb{R}^n$ is *at most* $\sqrt{n}$. It is intuitively clear that it really is $\sqrt{n}$, i.e., that we cannot find a linear image of the ball that sandwiches the cube any better than the obvious one. A formal proof will be immediate after the next lecture.

The main result of this lecture will imply that, if a polytope is to have small distance from the Euclidean ball, it must have very many facets: exponentially many in the dimension n.

THEOREM 2.1. *Let K be a (symmetric) polytope in $\mathbb{R}^n$ with $d(K, B_2^n) = d$. Then K has at least $e^{n/(2d^2)}$ facets. On the other hand, for each n, there is a polytope with 4^n facets whose distance from the ball is at most 2.*

The arguments in this lecture, including the result just stated, go back to the early days of packing and covering problems. A classical reference for the subject is [Rogers 1964].

Before we embark upon a proof of Theorem 2.1, let's look at a reformulation that will motivate several more sophisticated results later on. A symmetric convex body in $\mathbb{R}^n$ with m pairs of facets can be realised as an n-dimensional slice (through the centre) of the cube in $\mathbb{R}^m$. This is because such a body is the intersection of m slabs in $\mathbb{R}^n$, each of the form $\{x : |\langle x, v\rangle| \le 1\}$, for some nonzero vector v in $\mathbb{R}^n$. This is shown in Figure 8.

Thus K is the set $\{x : |\langle x, v_i\rangle| \le 1 \text{ for } 1 \le i \le m\}$, for some sequence $(v_i)_1^m$ of vectors in $\mathbb{R}^n$. The linear map

$$T : x \mapsto (\langle x, v_1\rangle, \dots, \langle x, v_m\rangle)$$

embeds $\mathbb{R}^n$ as a subspace H of $\mathbb{R}^m$, and the intersection of H with the cube $[-1, 1]^m$ is the set of points y in H for which $|y_i| \le 1$ for each coordinate i. So this intersection is the image of K under T. Conversely, any n-dimensional slice of $[-1, 1]^m$ is a body with at most m pairs of faces. Thus, the result we are aiming at can be formulated as follows:

> *The cube in $\mathbb{R}^m$ has almost spherical sections whose dimension n is roughly* $\log m$ *and not more.*

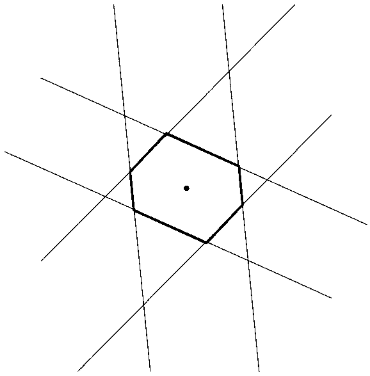

Figure 8. Any symmetric polytope is a section of a cube.

In Lecture 9 we shall see that all symmetric m-dimensional convex bodies have almost spherical sections of dimension about $\log m$. As one might expect, this is a great deal more difficult to prove for general bodies than just for the cube.

For the proof of Theorem 2.1, let's forget the symmetry assumption again and just ask for a polytope

$$K = \{x : \langle x, v_i \rangle \leq 1 \text{ for } 1 \leq i \leq m\}$$

with m facets for which

$$B_2^n \subset K \subset dB_2^n.$$

What do these inclusions say about the vectors (v_i)? The first implies that each v_i has length at most 1, since, if not, $v_i/|v_i|$ would be a vector in B_2^n but not in K. The second inclusion says that if x does not belong to dB_2^n then x does not belong to K: that is, if $|x| > d$, there is an i for which $\langle x, v_i \rangle > 1$. This is equivalent to the assertion that for every unit vector θ there is an i for which

$$\langle \theta, v_i \rangle \geq \frac{1}{d}.$$

Thus our problem is to look for as few vectors as possible, $v_1, v_2, \ldots, v_m$, of length at most 1, with the property that for every unit vector θ there is some v_i with $\langle \theta, v_i \rangle \geq 1/d$. It is clear that we cannot do better than to look for vectors of length exactly 1: that is, that we may assume that all facets of our polytope touch the ball. Henceforth we shall discuss only such vectors.

For a fixed unit vector v and $\varepsilon \in [0, 1)$, the set $C(\varepsilon, v)$ of $\theta \in S^{n-1}$ for which $\langle \theta, v \rangle \geq \varepsilon$ is called a *spherical cap* (or just a *cap*); when we want to be precise, we will call it the *ε-cap about v*. (Note that ε does not refer to the radius!) See Figure 9, left.

We want every $\theta \in S^{n-1}$ to belong to at least one of the $(1/d)$-caps determined by the (v_i). So our task is to estimate the number of caps of a given size needed to cover the entire sphere. The principal tool for doing this will be upper and lower estimates for the area of a spherical cap. As in the last lecture, we shall