Convexity in High Dimension

Course Notes

A six-week course in asymptotic geometric analysis

Graduate course

Preface

These notes accompany the graduate course *Convexity in High Dimension*. The guiding question is

Where does Euclidean geometry hide inside high-dimensional normed spaces?

The answer depends on whether one studies a whole normed space, a subspace, or a quotient. Globally, John's theorem gives a universal Euclidean approximation whose distortion grows with the dimension, and the cube shows that this dependence cannot be improved in general. Locally, the same highly non-Euclidean spaces can contain large, nearly Euclidean subspaces. Allowing both sections and quotients leads to a still stronger structural statement.

The purpose of the course is to explain the mechanisms behind these phenomena. The first is the geometry of extremal ellipsoids and Banach–Mazur distance. The second is *concentration of measure*: a random direction typically behaves much more regularly than an arbitrary direction. Combined with nets, random matrices, Gaussian widths, and duality, concentration becomes a method for finding Euclidean structure in high dimension.

The manuscript is organized into *one chapter per week*. Within each chapter, definitions, examples, proofs, and complementary arguments are placed in their natural mathematical order, regardless of whether they are discussed during a lecture or in a problem session. We give complete proofs for the arguments developed in the course, including steps that may be omitted at the board. Results quoted from the research literature are identified rather than presented as proved here.

Three major inputs in the six-week programme are spherical isoperimetry, Gordon's Gaussian comparison theorem, and the ℓ-position (or MM^*) estimate. Other deep results, such as Gluskin's diameter theorem and the recent estimates for the Banach–Mazur radius of the cube, also appear as explicitly attributed background results. The roadmap below describes the *entire* course, while the table of contents records the chapters actually included in the present manuscript. The notes will grow as the course develops.

Mathematical roadmap

The six weeks form one sequence of questions, not a catalogue of independent theorems. We begin by measuring the failure of Euclidean geometry globally, learn why randomness produces regularity, and use that regularity to find increasingly strong forms of Euclidean structure.

	Geometric question	**Main mechanism**
Week 1	How far can a normed space be from Euclidean?	John's ellipsoid, contact decomposition, Banach–Mazur distance; first Euclidean sections
Week 2	Why are high-dimensional random measurements stable?	Spherical and Gaussian concentration; Johnson–Lindenstrauss
Week 3	How does concentration control an entire subspace?	Nets and quantitative Dvoretzky–Milman
Week 4	What dimensions can be guaranteed universally or in ℓ_1^n?	Dvoretzky–Rogers and Kashin's small-ball method
Week 5	How large can a well-controlled random section be?	Gaussian width, Gordon's theorem, low M^*
Week 6	What if we allow a section *and* a quotient?	Duality, ℓ-position, and quotient-of-subspace iteration

Week 1. Banach–Mazur geometry and John's theorem

The unit ball of a norm turns questions about finite-dimensional Banach spaces into questions about symmetric convex bodies. Banach–Mazur distance measures the best possible comparison after a linear change of coordinates. John's theorem gives

$$d_{\mathrm{BM}}(X, \ell_2^n) \leq \sqrt{n},$$

and the cube and cross-polytope show that this estimate is sharp. The contact-point decomposition of the identity is both an extremality certificate and a weighted Euclidean measurement system. Global distance, however, does not determine the geometry of sections: the cube is far from Euclidean, yet has nearly circular planar sections. That observation motivates the rest of the course.

Week 2. Concentration of measure and random projections

A random point of a high-dimensional sphere spends most of its time in a narrow range of values of any Lipschitz function. We derive Lévy's concentration lemma from spherical isoperimetry and establish Gaussian concentration by the Maurey–Pisier rotation argument. The Johnson–Lindenstrauss theorem then illustrates a fundamental principle: an exponentially small failure probability, combined with a union bound, can preserve many prescribed distances simultaneously. But a Euclidean *subspace* contains infinitely many directions; that is the next challenge.

Week 3. Quantitative Dvoretzky–Milman

For a norm on $\mathbb{R}^n$, write

$$M = \mathbb{E}\|\Theta\|_K, \qquad b = \max_{|\theta|_2=1} \|\theta\|_K, \qquad k_*(K) = n\left(\frac{M}{b}\right)^2.$$

Concentration controls the norm at one random direction, while a finite net allows control of all directions of a random k-dimensional subspace. This yields almost-Euclidean sections of dimension at least $c_\varepsilon k_*(K)$ for each fixed distortion parameter $\varepsilon > 0$. We keep track of the logarithmic loss in the distortion parameter caused by the elementary net argument. The resulting theorem depends on the linear position of the body, so a universal statement still needs another ingredient.

Week 4. Universal Dvoretzky and Kashin's theorem

The parameter $k_*(K)$ behaves very differently for different bodies: it is of order n for ℓ_1^n, but only of order $\log n$ for ℓ_∞^n. The Dvoretzky–Rogers argument, starting from John's contact decomposition, yields a universal logarithmic lower bound and hence an almost-Euclidean section in *every* finite-dimensional normed space. Kashin's theorem reveals an even stronger phenomenon in ℓ_1^n: for any fixed proportion $\lambda < 1$, one can find a subspace of dimension at least λn that is Euclidean up to a distortion depending on λ. Its proof replaces ordinary concentration by small-ball estimates strong enough to overcome a large net.

Week 5. Gaussian width, Gordon's theorem, and low M^*

Instead of asking that the norm be nearly constant on a section, we now ask whether a random subspace avoids a set of undesirable directions. Gaussian width measures the complexity of such a set. Gordon's comparison theorem, together with Gaussian concentration, controls the minimum of a Gaussian operator over it and yields the escape-through-a-mesh principle. Applied to suitable spherical slices of a convex body, this leads to the low-M^* estimate: a random section of sufficiently large codimension has controlled Euclidean radius. This one-sided conclusion is different from Dvoretzky's two-sided near-Euclidean estimate, and is precisely suited to the final week's duality arguments.

Week 6. Quotients of subspaces

Sections and quotients are dual geometric operations: the polar of a section corresponds to a projection of the polar body. Combining this dictionary with low M^* and an appropriately chosen linear position gives a one-step construction of a relatively Euclidean quotient of a subspace. The main admitted structural input is the ℓ-position/MM^* estimate. The first bound is logarithmic in the initial Banach–Mazur distance; an iteration improves it repeatedly while controlling the loss of dimension. The endpoint is Milman's quotient-of-subspace theorem: every n-dimensional normed space admits, for any fixed $0 < \lambda < 1$, a quotient of a subspace of dimension at least λn and Banach–Mazur distance from Euclidean space bounded by a constant depending only on λ.

The logical progression can be summarized as

$$\text{John's theorem} \longrightarrow \text{global geometry and its limits,}$$
$$\text{concentration} + \text{nets} \longrightarrow \text{Dvoretzky–Milman} \longrightarrow \text{universal Dvoretzky and Kashin,}$$
$$\text{Gaussian width} + \text{duality} + \ell\text{-position} \longrightarrow \text{quotients of subspaces.}$$

Contents

Preface **i**

Mathematical roadmap **ii**

1 Week 1: Banach–Mazur geometry and John's theorem **1**

1.1 Lecture 1: John's theorem, from the extremal ellipsoid to $\sqrt{n}$ 1

References **19**

Chapter 1

Week 1: Banach–Mazur geometry and John's theorem

1.1 Lecture 1: John's theorem, from the extremal ellipsoid to $\sqrt{n}$

1.1.1 Notation

Throughout, $\mathbb{R}^n$ has its standard scalar product $\langle x, y\rangle = \sum_{i=1}^n x_i y_i$ and Euclidean norm

$$|x|_2 = \langle x, x\rangle^{1/2}.$$

We write

$$B_2^n = \{x \in \mathbb{R}^n : |x|_2 \le 1\}, \quad S^{n-1} = \{x \in \mathbb{R}^n : |x|_2 = 1\}.$$

For $1 \le p < \infty$ and for $p = \infty$, respectively, let

$$\|x\|_p = \left(\sum_{i=1}^n |x_i|^p\right)^{1/p}, \quad \|x\|_\infty = \max_{1\le i\le n} |x_i|.$$

As usual, $\ell_p^n = (\mathbb{R}^n, \|\cdot\|_p)$ and $B_p^n = \{x \in \mathbb{R}^n : \|x\|_p \le 1\}$. In particular, B_∞^n is the cube, B_1^n is the cross-polytope, and B_2^n is the Euclidean ball. We write $GL_n(\mathbb{R})$ for the invertible $n \times n$ real matrices. Unless stated otherwise, all bodies and linear maps are considered in $\mathbb{R}^n$.

1.1.2 Why convex bodies represent norms

A *convex body* is a compact convex subset of $\mathbb{R}^n$ with nonempty interior. A body K is *centrally symmetric* (here simply *symmetric*) if $K = -K$. Suppose $K = -K$ and $0 \in \operatorname{int} K$.

Definition 1.1.1 (Gauge, or Minkowski functional)**.**

$$\|x\|_K := \inf\{t > 0 : x \in tK\}.$$

Proposition 1.1.2. *The function $\|\cdot\|_K$ is a norm, and its closed unit ball is exactly K.*

Proof. Positivity. Since $0 \in \operatorname{int}(K)$ and K is bounded, there exist $0 < r \le R$ such that

$$rB_2^n \subset K \subset RB_2^n.$$

By the definition of the gauge, this implies

$$\frac{|x|_2}{R} \le \|x\|_K \le \frac{|x|_2}{r}.$$

Thus $\|x\|_K$ is finite, nonnegative, and vanishes only at 0.
Homogeneity. For $\lambda > 0$, a change of variable in the infimum gives

$$\|\lambda x\|_K = \inf\{t > 0 : \lambda x \in tK\} = \lambda \|x\|_K.$$

Since $K = -K$, we also have $\|-x\|_K = \|x\|_K$. Hence $\|\lambda x\|_K = |\lambda|\|x\|_K$ for all $\lambda \in \mathbb{R}$.
Triangle inequality. Let $s > \|x\|_K$ and $t > \|y\|_K$. Since K is convex and contains 0, we have $x/s, y/t \in K$. By convexity,

$$\frac{x+y}{s+t} = \frac{s}{s+t}\frac{x}{s} + \frac{t}{s+t}\frac{y}{t} \in K.$$

Therefore $\|x+y\|_K \le s+t$. Letting $s \downarrow \|x\|_K$ and $t \downarrow \|y\|_K$ yields

$$\|x+y\|_K \le \|x\|_K + \|y\|_K.$$

Thus $\|\cdot\|_K$ is a norm.
Unit ball. If $x \in K$, then clearly $\|x\|_K \le 1$.
Conversely, if $\|x\|_K \le 1$, then for every $\varepsilon > 0$,

$$\frac{x}{1+\varepsilon} \in K.$$

Letting $\varepsilon \to 0$ and using the closedness of K, we obtain $x \in K$. Consequently,

$$K = \{x \in \mathbb{R}^n : \|x\|_K \le 1\}.$$

□

Basic identities. We record several useful properties of gauges and ellipsoids.

1. *Reversal of inclusions.* If $K \subset L$ are symmetric convex bodies, then $tK \subset tL$ for every $t > 0$. Thus $x \in tK$ implies $x \in tL$, and taking the infimum over admissible t gives

   $$\|x\|_L \le \|x\|_K \quad (x \in \mathbb{R}^n).$$

   Hence inclusions of bodies run in the *opposite direction* to inequalities between their gauges: a larger unit ball defines a smaller norm.

2. *Invariance under linear transformations.* If $T \in GL_n(\mathbb{R})$, then by linearity and invertibility,

   $$Tx \in t(TK) \iff Tx \in T(tK) \iff x \in tK.$$

   Taking the infimum over $t > 0$ yields

   $$\|Tx\|_{TK} = \|x\|_K.$$

   Thus T is a linear isometry from $(\mathbb{R}^n, \|\cdot\|_K)$ to $(\mathbb{R}^n, \|\cdot\|_{TK})$. In particular, invertible linear changes of coordinates preserve the intrinsic normed-space geometry.

3. *Representation of ellipsoids.* Every ellipsoid centered at 0 can be written as $E = AB_2^n$ for some invertible A. By polar decomposition, $A = UP$, where U is orthogonal and P is symmetric positive definite. Since $UB_2^n = B_2^n$,

   $$AB_2^n = UPB_2^n = (UPU^T)B_2^n.$$

   The matrix $Q = UPU^T = (AA^T)^{1/2}$ is symmetric positive definite. Hence every centered ellipsoid can also be written as QB_2^n with $Q = Q^T > 0$. Its eigenvectors give the principal directions, and its eigenvalues give the corresponding semiaxis lengths.

4. *Volume of an ellipsoid.* For $E = AB_2^n$, the linear change of variables $y = Ax$ has constant Jacobian $|\det A|$. Therefore

   $$\mathrm{vol}_n(AB_2^n) = \int_{B_2^n} |\det A| dx = |\det A| \,\mathrm{vol}_n(B_2^n).$$

   If $\lambda_1, \dots, \lambda_n > 0$ are the singular values of A, they are the lengths of the principal semiaxes, and

   $$\mathrm{vol}_n(E) = \lambda_1 \cdots \lambda_n \,\mathrm{vol}_n(B_2^n),$$

   since $|\det A| = \lambda_1 \cdots \lambda_n$. Consequently, finding a maximum-volume centered ellipsoid contained in a symmetric convex body K amounts to solving

   $$\max\{|\det A| : A \in GL_n(\mathbb{R}), AB_2^n \subset K\}.$$

   This explains why *determinant maximization* is the natural optimization problem.

5. *Adapted Euclidean coordinates.* For an ellipsoid $E = AB_2^n$, the definition of the gauge gives

$$\|x\|_E = \inf\{t > 0 : A^{-1}x \in tB_2^n\} = |A^{-1}x|_2.$$

Thus, after the change of coordinates $z = A^{-1}x$, the ellipsoidal norm becomes the standard Euclidean norm. More generally, choosing adapted linear coordinates allows us to normalize a convex body and remove anisotropy arising solely from the original choice of coordinates.

These identities show why convex bodies should be compared up to invertible linear transformations. The Banach–Mazur distance introduced below measures the smallest multiplicative distortion that remains after optimizing over all such transformations.

1.1.3 Banach–Mazur distance and its geometric formulation

For normed spaces X, Y of dimension n and a linear map $T : X \to Y$, its *operator norm* is

$$\|T\|_{X\to Y} = \sup_{\|x\|_X \le 1} \|Tx\|_Y.$$

For a Euclidean operator $A : \mathbb{R}^n \to \mathbb{R}^n$, we also write

$$\|A\|_{\mathrm{op}} = \sup_{|x|_2=1} |Ax|_2.$$

The latter notation will be used for the symmetric-matrix perturbations in John's proof.

Definition 1.1.3 (Banach–Mazur distance)**.** For two n-dimensional normed spaces X, Y, define

$$d_{\mathrm{BM}}(X, Y) = \inf_{T:X\to Y \text{ invertible}} \|T\|_{X\to Y}\|T^{-1}\|_{Y\to X}.$$

In particular, $d_{\mathrm{BM}}(X, Y) \ge 1$. The product is unchanged if T is multiplied by a nonzero scalar.

Proposition 1.1.4 (Geometric formulation)**.** *Let $K = B_X$ and $L = B_Y$ be the unit balls of X and Y. Then*

$$d_{\mathrm{BM}}(X, Y) = \inf\{d \ge 1 : \exists S \in GL_n(\mathbb{R}), L \subset SK \subset dL\}.$$

Proof. Let K and L be the unit balls of X and Y, respectively, so that

$$K = \{x \in X : \|x\|_X \le 1\}, \quad L = \{y \in Y : \|y\|_Y \le 1\}.$$

Let $T : X \to Y$ be invertible and put $a = \|T\|_{X\to Y}$ and $b = \|T^{-1}\|_{Y\to X}$. By the definition of the operator norm, for $x \in K$ and $y \in L$,

$$\|Tx\|_Y \le a\|x\|_X \le a, \quad \|T^{-1}y\|_X \le b\|y\|_Y \le b.$$

Thus $Tx \in aL$ and $T^{-1}y \in bK$, or equivalently $y \in bTK$. Therefore

$$TK \subset aL, \quad L \subset bTK.$$

Setting $S = bT$, we obtain

$$L \subset SK \subset abL.$$

Thus the geometric infimum is at most $ab = \|T\|_{X\to Y}\|T^{-1}\|_{Y\to X}$. Taking the infimum over all invertible T gives one inequality.
Conversely, suppose that $S : X \to Y$ is invertible and

$$L \subset SK \subset dL.$$

The inclusion $SK \subset dL$ means that $\|Sx\|_Y \le d$ for every $x \in K$, while $L \subset SK$ implies $\|S^{-1}y\|_X \le 1$ for every $y \in L$. By the definition of the operator norm,

$$\|S\|_{X\to Y} \le d, \quad \|S^{-1}\|_{Y\to X} \le 1.$$

Hence

$$\|S\|_{X\to Y}\|S^{-1}\|_{Y\to X} \le d.$$

Taking the infimum over all admissible S and d gives the reverse inequality, proving that the two definitions of $d_{\mathrm{BM}}(X, Y)$ coincide. □

For symmetric convex bodies K, L, we also use the shorthand

$$d_{\mathrm{BM}}(K, L) = d_{\mathrm{BM}}((\mathbb{R}^n, \|\cdot\|_K), (\mathbb{R}^n, \|\cdot\|_L)).$$

1.1.4 Basic properties of Banach–Mazur distance

Proposition 1.1.5. *For n-dimensional normed spaces X, Y, Z the following hold:*

$$1 \le d_{\mathrm{BM}}(X,Y) = d_{\mathrm{BM}}(Y,X),$$
$$d_{\mathrm{BM}}(X,Z) \le d_{\mathrm{BM}}(X,Y) d_{\mathrm{BM}}(Y,Z),$$
$$d_{\mathrm{BM}}(X,Y) = 1 \quad \Longleftrightarrow \quad X \text{ and } Y \text{ are linearly isometric.}$$

Consequently $\log d_{\mathrm{BM}}$ *is a metric on the set of linear-isometry classes of n-dimensional normed spaces.*

Proof. Let $T : X \to Y$ be an isomorphism. Since $T^{-1}T = I_X$, submultiplicativity of the operator norm gives

$$1 = \|I_X\| = \|T^{-1}T\| \le \|T^{-1}\|\|T\|.$$

Taking the infimum over all invertible T yields $d_{\mathrm{BM}}(X,Y) \ge 1$. Moreover, since every isomorphism $T : X \to Y$ has an inverse $T^{-1} : Y \to X$, and the product $\|T\|\|T^{-1}\|$ is unchanged when exchanging T and T^{-1}, we obtain

$$d_{\mathrm{BM}}(X,Y) = d_{\mathrm{BM}}(Y,X).$$

Next, let $T : X \to Y$ and $S : Y \to Z$ be isomorphisms. Their composition $ST : X \to Z$ is also an isomorphism, with $(ST)^{-1} = T^{-1}S^{-1}$. By submultiplicativity,

$$\|ST\| \le \|S\|\|T\|, \quad \|(ST)^{-1}\| \le \|T^{-1}\|\|S^{-1}\|.$$

Multiplying these inequalities gives

$$\|ST\|\|(ST)^{-1}\| \le (\|S\|\|S^{-1}\|)(\|T\|\|T^{-1}\|).$$

Since ST is admissible in the definition of $d_{\mathrm{BM}}(X,Z)$, taking the infimum over all isomorphisms S and T yields the multiplicative triangle inequality

$$d_{\mathrm{BM}}(X,Z) \le d_{\mathrm{BM}}(X,Y) d_{\mathrm{BM}}(Y,Z).$$

We now show that $d_{\mathrm{BM}}(X,Y) = 1$ if and only if X and Y are linearly isometric. If $T : X \to Y$ is a linear isometry, then $\|T\| = \|T^{-1}\| = 1$, so $d_{\mathrm{BM}}(X,Y) = 1$.
Conversely, suppose $d_{\mathrm{BM}}(X,Y) = 1$. By the definition of the infimum, there exist isomorphisms $T_j : X \to Y$ such that

$$\|T_j\|\|T_j^{-1}\| \to 1.$$

Since multiplying T_j by a positive scalar does not change this product, we may normalize so that $\|T_j\| = 1$. It follows that $\|T_j^{-1}\| \to 1$.
After identifying X and Y linearly with $\mathbb{R}^n$, the sequence (T_j) is bounded in the finite-dimensional space of matrices, since all norms on this space are equivalent. We may therefore pass to a convergent subsequence $T_j \to T$ for some linear map $T : X \to Y$.
For every $x \in X$, the definition of the operator norm gives

$$\|T_j x\|_Y \le \|T_j\|\|x\|_X = \|x\|_X.$$

On the other hand, since $x = T_j^{-1}T_j x$,

$$\|x\|_X \le \|T_j^{-1}\|\|T_j x\|_Y,$$

and therefore

$$\frac{\|x\|_X}{\|T_j^{-1}\|} \le \|T_j x\|_Y \le \|x\|_X.$$

Passing to the limit, using $T_j \to T$ and $\|T_j^{-1}\| \to 1$, yields

$$\|Tx\|_Y = \|x\|_X$$

for every $x \in X$. Thus T is injective and, since X and Y have the same finite dimension, invertible. Hence T is a linear isometry.
Finally, define $d(X,Y) = \log d_{\mathrm{BM}}(X,Y)$. The preceding properties give nonnegativity, symmetry, and $d(X,Y) = 0$ exactly when X and Y are linearly isometric. Moreover, taking logarithms in the multiplicative triangle inequality gives

$$d(X,Z) \le d(X,Y) + d(Y,Z).$$

Thus d defines a metric on the classes of finite-dimensional normed spaces of fixed dimension, up to linear isometry. □

1.1.5 Elementary examples and the ℓ_p^n spaces

Example 1.1.6 (Ellipsoids)**.** For $K = AB_2^n$, $A \in GL_n(\mathbb{R})$, one has

$$\|x\|_K = |A^{-1}x|_2, \quad d_{\mathrm{BM}}\big((\mathbb{R}^n, \|\cdot\|_K), \ell_2^n\big) = 1.$$

Thus a linear change of coordinates makes every ellipsoid exactly Euclidean.

Proposition 1.1.7 (Comparison of ℓ_p norms)**.** *If $1 \le p \le q \le \infty$, then for every $x \in \mathbb{R}^n$,*

$$\|x\|_q \le \|x\|_p \le n^{1/p-1/q}\|x\|_q,$$

where $1/\infty = 0$. In particular,

$$d_{\mathrm{BM}}(\ell_p^n, \ell_q^n) \le n^{|1/p-1/q|} \quad (1 \le p, q \le \infty).$$

Proof. We first establish the inequalities between the norms, then use them to bound the Banach–Mazur distance. The cases $p = q$ and $x = 0$ are immediate, so assume $1 \le p < q$ and $x \ne 0$.
Step 1: Prove $\|x\|_q \le \|x\|_p$ for $q < \infty$.
Set $a = \|x\|_p$. Since $|x_i| \le a$ for every i, we have $0 \le |x_i|/a \le 1$. As $p < q$, raising these quantities to the power q gives smaller values than raising them to the power p. Therefore

$$\sum_{i=1}^n \left(\frac{|x_i|}{a}\right)^q \le \sum_{i=1}^n \left(\frac{|x_i|}{a}\right)^p = 1.$$

Multiplying by a^q and taking the q-th root gives

$$\|x\|_q \le \|x\|_p.$$

Step 2: Prove $\|x\|_p \le n^{1/p-1/q}\|x\|_q$ for $q < \infty$.
Apply Hölder's inequality with conjugate exponents $r = q/p$ and $s = q/(q-p)$, which satisfy $1/r+1/s = 1$. We obtain

$$\sum_{i=1}^n |x_i|^p = \sum_{i=1}^n |x_i|^p \cdot 1 \le \left(\sum_{i=1}^n |x_i|^q\right)^{p/q} n^{1-p/q}.$$

Taking the p-th root yields

$$\|x\|_p \le n^{1/p-1/q}\|x\|_q.$$

Combining Steps 1 and 2 gives

$$\|x\|_q \le \|x\|_p \le n^{1/p-1/q}\|x\|_q.$$

Step 3: Consider the case $q = \infty$.
By the definition of the maximum norm, $|x_i| \le \|x\|_\infty$ for every i. Moreover, the largest term in the sum defining $\|x\|_p^p$ cannot exceed the sum itself. Hence

$$\max_i |x_i| \le \left(\sum_{i=1}^n |x_i|^p\right)^{1/p} \le n^{1/p} \max_i |x_i|.$$

Thus the same inequalities hold for $q = \infty$, with the convention $1/\infty = 0$.
Step 4: Deduce the Banach–Mazur bound.
Consider the identity map $I : \ell_p^n \to \ell_q^n$. By the norm inequalities established above,

$$\|I\|_{\ell_p^n \to \ell_q^n} = \sup_{x \ne 0} \frac{\|x\|_q}{\|x\|_p} \le 1$$

and

$$\|I^{-1}\|_{\ell_q^n \to \ell_p^n} = \sup_{x \ne 0} \frac{\|x\|_p}{\|x\|_q} \le n^{1/p-1/q}.$$

Since I is an admissible isomorphism in the definition of the Banach–Mazur distance,

$$d_{\mathrm{BM}}(\ell_p^n, \ell_q^n) \le \|I\|\|I^{-1}\| \le n^{1/p-1/q}.$$

Finally, if $p > q$, exchanging their roles and using the symmetry of d_{BM} gives

$$d_{\mathrm{BM}}(\ell_p^n, \ell_q^n) \le n^{|1/p-1/q|}.$$

□

Example 1.1.8 (The square and the diamond)**.** In dimension two, the unit balls of ℓ_1^2 and ℓ_∞^2 are the diamond and the square, respectively. Although these bodies are different, they are related by an invertible linear transformation.

Consider the map $T : \ell_1^2 \to \ell_\infty^2$ defined by

$$T(x, y) = (x + y, x - y).$$

Its matrix has determinant $-2 \neq 0$, so T is invertible. Moreover,

$$\|T(x, y)\|_\infty = \max\{|x + y|, |x - y|\} = |x| + |y| = \|(x, y)\|_1.$$

Indeed, if x and y have the same sign, then $|x + y| = |x| + |y|$, while if they have opposite signs, then $|x - y| = |x| + |y|$. Thus T preserves the norm and is a linear isometry between ℓ_1^2 and ℓ_∞^2.

Consequently, $\|T\| = \|T^{-1}\| = 1$, and by the definition of the Banach–Mazur distance,

$$d_{\mathrm{BM}}(\ell_1^2, \ell_\infty^2) = 1.$$

This shows that the elementary comparison bound $d_{\mathrm{BM}}(\ell_1^2, \ell_\infty^2) \leq 2$, obtained using the identity map, is not sharp. In general, the identity map need not minimize the distortion.

This equality is special to dimensions one and two. For $n \geq 3$, the unit balls of ℓ_1^n and ℓ_∞^n have $2n$ and 2^n vertices, respectively, so they cannot be linearly isometric. In fact, their Banach–Mazur distance grows asymptotically as $\sqrt{n}$.

Remark 1.1.9 (Some exact values)**.** The norm comparison inequalities give the general upper bound

$$d_{\mathrm{BM}}(\ell_p^n, \ell_q^n) \leq n^{|1/p-1/q|}.$$

However, this bound is not always sharp, since the identity map need not minimize the distortion. There are two important cases in which equality holds.

First, for every $1 \leq p \leq \infty$, the exact distance to Euclidean space is

$$d_{\mathrm{BM}}(\ell_p^n, \ell_2^n) = n^{|1/p-1/2|}.$$

More generally, if p and q lie on the same side of 2, meaning either $1 \leq p, q \leq 2$ or $2 \leq p, q \leq \infty$, then

$$d_{\mathrm{BM}}(\ell_p^n, \ell_q^n) = n^{|1/p-1/q|}.$$

In these cases, the identity map is optimal. Indeed, the norm comparison gives the upper bound, while the reverse inequality requires showing that every invertible linear transformation has distortion at least $n^{|1/p-1/q|}$. These general lower bounds are more involved and will not be proved here. Lecture 2 establishes the endpoint cases $p = 1, \infty$ against 2.

When p and q lie on opposite sides of 2, the situation is more subtle. The identity-map bound may be strictly larger than the actual Banach–Mazur distance, as illustrated by

$$d_{\mathrm{BM}}(\ell_1^2, \ell_\infty^2) = 1 < 2.$$

Thus, determining the exact distance requires optimizing over all invertible linear transformations, rather than relying only on inequalities between the norms.

1.1.6 John's theorem: the statement and the central mechanism

The natural question is now precise: how large can $d_{\mathrm{BM}}(X, \ell_2^n)$ be when $\dim X = n$?

Theorem 1.1.10 (John's theorem, symmetric case)**.** *For every symmetric convex body $K \subset \mathbb{R}^n$ there is $T \in GL_n(\mathbb{R})$ such that*

$$B_2^n \subset TK \subset \sqrt{n} B_2^n.$$

Equivalently, every n-dimensional normed space satisfies

$$d_{\mathrm{BM}}(X, \ell_2^n) \leq \sqrt{n}.$$

The geometric theorem rests on a more structural result.

Theorem 1.1.11 (John's contact-point decomposition)**.** *Assume that $K = -K$, that $B_2^n \subset K$, and that B_2^n has maximal volume among all ellipsoids contained in K. There exist contact points $u_i \in \partial K \cap S^{n-1}$ and positive weights c_i such that*

$$I_n = \sum_{i=1}^{m} c_i u_i u_i^T, \quad m \le \frac{n(n+1)}{2}.$$

The proof has two parts: maximality of the ellipsoid forces the contact-point identity (by separation and a volume-increasing perturbation), and that identity immediately implies the $\sqrt{n}$ inclusion. We establish the necessary lemmas and then prove both theorems.

1.1.7 Existence of a maximal-volume ellipsoid

Lemma 1.1.12 (Centering)**.** *If $K = -K$ and $a + AB_2^n \subset K$, then $AB_2^n \subset K$ as well. In particular, maximal-volume inscribed ellipsoids may be sought among centered ellipsoids.*

Proof. Let $E = a + AB_2^n \subset K$ be an ellipsoid centered at a. Since K is centrally symmetric, $K = -K$, and since AB_2^n is symmetric about the origin, we have

$$-(a + AB_2^n) = -a + AB_2^n \subset K.$$

Thus both $a + AB_2^n$ and $-a + AB_2^n$ are contained in K.
Now let $z \in AB_2^n$. The two inclusions imply that $a + z \in K$ and $-a + z \in K$. By convexity of K, their midpoint also belongs to K, so

$$z = \frac{(a+z) + (-a+z)}{2} \in K.$$

Since z was arbitrary, we conclude that

$$AB_2^n \subset K.$$

Finally, translations preserve volume, hence

$$\mathrm{vol}_n(a + AB_2^n) = \mathrm{vol}_n(AB_2^n) = |\det A| \, \mathrm{vol}_n(B_2^n).$$

Therefore, every ellipsoid contained in K can be replaced by a centered ellipsoid of the same volume, still contained in K. □

Lemma 1.1.13 (Existence)**.** *Every symmetric convex body has a maximal-volume inscribed ellipsoid.*

Proof. Since $0 \in \mathrm{int}(K)$ and K is bounded, there exist $0 < r < R$ such that

$$rB_2^n \subset K \subset RB_2^n.$$

By Lemma 1.1.12, it suffices to consider centered ellipsoids. Using the representation by symmetric positive-definite matrices and the volume formula, the problem reduces to showing that the supremum

$$M = \sup\{\det A : A = A^T > 0, AB_2^n \subset K\}$$

is attained.
For every admissible A, the inclusion $AB_2^n \subset RB_2^n$ implies

$$\|A\|_{\mathrm{op}} = \sup_{|v|_2 \le 1} |Av|_2 \le R.$$

Since A is symmetric positive definite, all its eigenvalues lie in $(0, R]$, so $\det A \le R^n$. Moreover, rI is admissible because $rB_2^n \subset K$, and $\det(rI) = r^n > 0$. Thus

$$r^n \le M \le R^n.$$

Let (A_j) be a maximizing sequence, so that $\det A_j \to M$. Since $M \ge r^n$, we may assume, after discarding finitely many terms, that $\det A_j \ge r^n/2$. Denoting the eigenvalues of A_j by

$$0 < \lambda_1^{(j)} \le \cdots \le \lambda_n^{(j)} \le R,$$

we obtain

$$\frac{r^n}{2} \le \det A_j = \prod_{i=1}^{n} \lambda_i^{(j)} \le \lambda_1^{(j)} R^{n-1},$$

and hence

$$\lambda_1^{(j)} \ge \frac{r^n}{2R^{n-1}} > 0.$$

Thus the eigenvalues are uniformly bounded above and away from zero. Consequently, (A_j) lies in a compact subset of the symmetric positive-definite matrices. Passing to a subsequence, we obtain $A_j \to A$, where A remains symmetric positive definite.
It remains to show that A is admissible and attains the supremum. For every $v \in B_2^n$, we have $A_j v \in K$. Since $A_j v \to Av$ and K is closed, it follows that $Av \in K$. Therefore $AB_2^n \subset K$. Finally, by continuity of the determinant,

$$\det A = \lim_{j\to\infty} \det A_j = M.$$

Hence the supremum is attained, and AB_2^n is a maximum-volume ellipsoid contained in K. □

Lemma 1.1.14 (First variation of ellipsoid volume)**.** *Let $H = H^T$ and $A_t = I_n + tH$. For t sufficiently close to 0, A_t is positive definite and*

$$\left.\frac{d}{dt}\right|_{t=0} \log\det(I_n + tH) = \operatorname{tr} H, \quad \det(I_n + tH) = 1 + t\operatorname{tr} H + O(t^2).$$

In particular, $\operatorname{tr} H > 0$ makes the ellipsoid $A_t B_2^n$ larger in volume for all sufficiently small positive t.

Proof. Since H is symmetric, the spectral theorem gives $H = QDQ^T$, where Q is orthogonal and $D = \operatorname{diag}(\lambda_1, \dots, \lambda_n)$. Thus

$$\det(I_n + tH) = \det(Q(I_n + tD)Q^T) = \prod_{i=1}^{n}(1 + t\lambda_i).$$

For t sufficiently close to 0, all factors are positive, so taking logarithms gives

$$\log\det(I_n + tH) = \sum_{i=1}^{n} \log(1 + t\lambda_i).$$

Differentiating with respect to t and evaluating at $t = 0$, we obtain

$$\left.\frac{d}{dt} \log\det(I_n + tH)\right|_{t=0} = \sum_{i=1}^{n} \lambda_i = \operatorname{tr} H.$$

To obtain the determinant expansion, set $f(t) = \det(I_n + tH)$. Since $f(0) = 1$ and

$$\frac{f'(0)}{f(0)} = \left.\frac{d}{dt} \log f(t)\right|_{t=0} = \operatorname{tr} H,$$

we have $f'(0) = \operatorname{tr} H$. As f is a polynomial in t, its expansion at 0 gives

$$\det(I_n + tH) = 1 + t\operatorname{tr} H + O(t^2).$$

□

1.1.8 Normalization

Lemma 1.1.15 (John normalization)**.** *Let $K \subset \mathbb{R}^n$ be a convex body and let $E = AB_2^n \subset K$ be a maximum-volume ellipsoid, where A is invertible. Then B_2^n is a maximum-volume ellipsoid contained in $A^{-1}K$.*

Proof. Clearly,

$$B_2^n \subset A^{-1}K.$$

Suppose, for a contradiction, that there exists an ellipsoid $F \subset A^{-1}K$ such that

$$\mathrm{vol}_n(F) > \mathrm{vol}_n(B_2^n).$$

Since A is invertible, AF is an ellipsoid contained in K. Moreover,

$$\mathrm{vol}_n(AF) = |\det A|\,\mathrm{vol}_n(F) > |\det A|\,\mathrm{vol}_n(B_2^n) = \mathrm{vol}_n(E),$$

contradicting the maximality of E. □

Definition 1.1.16 (John position)**.** A convex body K is said to be in *John position* if B_2^n is a maximum-volume ellipsoid contained in K.

By the preceding lemma, every convex body can be brought into John position by an invertible affine transformation. For centrally symmetric bodies, a linear transformation suffices.

1.1.9 Supporting hyperplanes at contact points

Define the support function

$$h_K(v) = \max_{x \in K} \langle x, v \rangle, \quad v \in \mathbb{R}^n.$$

If $B_2^n \subset K$, then $h_K(u) \geq 1$ for every $u \in S^{n-1}$. The contact set is

$$U \coloneqq \partial K \cap S^{n-1}.$$

Lemma 1.1.17 (Three facts about support functions)**.** *Let K, L be nonempty compact convex sets and let A be an invertible matrix.*

1. *$K \subset L$ if and only if $h_K(v) \leq h_L(v)$ for every $v \in S^{n-1}$.*
2. *$h_{AB_2^n}(v) = |A^T v|_2$.*
3. *If $K \subset RB_2^n$, then $|h_K(v) - h_K(w)| \leq R|v - w|_2$ for every $v, w \in \mathbb{R}^n$.*

Proof. (1) Inclusion of convex bodies. Support functions are positively homogeneous and vanish at 0, so comparison on S^{n-1} is equivalent to comparison on all of $\mathbb{R}^n$. Suppose first that $K \subset L$. By the definition of the support function, for every $v \in \mathbb{R}^n$,

$$h_K(v) = \max_{x \in K} \langle x, v \rangle \leq \max_{x \in L} \langle x, v \rangle = h_L(v).$$

Conversely, suppose that $h_K(v) \leq h_L(v)$ for every $v \in \mathbb{R}^n$. If $K \not\subset L$, there exists $x \in K \setminus L$. Since L is closed and convex, the strict separation theorem gives a vector $v \in \mathbb{R}^n$ such that

$$\langle x, v \rangle > h_L(v).$$

But $x \in K$ implies $\langle x, v \rangle \leq h_K(v)$, hence

$$h_K(v) \geq \langle x, v \rangle > h_L(v),$$

contradicting our assumption. Therefore $K \subset L$.
(2) Support function of an ellipsoid. Let $E = AB_2^n$. Every point of E has the form Az with $|z|_2 \leq 1$. Using the definition of the support function and the identity $\langle Az, v \rangle = \langle z, A^T v \rangle$, we obtain

$$h_E(v) = \max_{|z|_2 \leq 1} \langle Az, v \rangle = \max_{|z|_2 \leq 1} \langle z, A^T v \rangle.$$

By Cauchy–Schwarz,

$$\langle z, A^T v \rangle \leq |z|_2 |A^T v|_2 \leq |A^T v|_2.$$

Equality is attained by choosing $z = A^T v / |A^T v|_2$ when $A^T v \neq 0$, while the case $A^T v = 0$ is immediate. Therefore

$$h_{AB_2^n}(v) = |A^T v|_2.$$

(3) Lipschitz continuity. Suppose that $K \subset RB_2^n$. For every $x \in K$ and $v, w \in \mathbb{R}^n$, Cauchy–Schwarz gives

$$\langle x, v \rangle = \langle x, w \rangle + \langle x, v - w \rangle \leq \langle x, w \rangle + R|v - w|_2.$$

Taking the maximum over $x \in K$ yields

$$h_K(v) \leq h_K(w) + R|v - w|_2.$$

Interchanging v and w, we also obtain

$$h_K(w) \leq h_K(v) + R|v - w|_2.$$

Combining both inequalities gives

$$|h_K(v) - h_K(w)| \leq R|v - w|_2.$$

Thus h_K is Lipschitz continuous. □

Lemma 1.1.18 (Contact points and their normals)**.** *If $B_2^n \subset K$, then for $u \in S^{n-1}$,*

$$u \in U \quad \Longleftrightarrow \quad h_K(u) = 1.$$

At such a point, the hyperplane $\{x : \langle x, u\rangle = 1\}$ supports K. In particular, if $K = -K$, then

$$|\langle x, u\rangle| \leq 1 \quad \textit{for every } x \in K.$$

Proof. Let $U = \partial K \cap S^{n-1}$ and suppose that $u \in U$. Since $u \in \partial K$, the supporting-hyperplane theorem gives a unit vector v such that

$$\langle x, v\rangle \leq \langle u, v\rangle \quad (x \in K).$$

Thus $h_K(v) = \langle u, v\rangle$. Since $B_2^n \subset K$, the support function of K is at least that of B_2^n, so

$$h_K(v) \geq h_{B_2^n}(v) = |v|_2 = 1.$$

On the other hand, since $|u|_2 = |v|_2 = 1$, Cauchy–Schwarz gives

$$1 \leq h_K(v) = \langle u, v\rangle \leq |u|_2|v|_2 = 1.$$

Hence $\langle u, v\rangle = 1$, and equality in Cauchy–Schwarz forces $v = u$. Consequently,

$$h_K(u) = \langle u, u\rangle = 1.$$

Conversely, suppose that $u \in S^{n-1}$ satisfies $h_K(u) = 1$. Since $B_2^n \subset K$, we have $u \in K$. Moreover, by the definition of the support function,

$$\langle x, u\rangle \leq h_K(u) = 1 = \langle u, u\rangle \quad (x \in K).$$

Thus the hyperplane $\{x : \langle x, u\rangle = 1\}$ supports K at u, so $u \in \partial K$. Since $u \in S^{n-1}$, we conclude that $u \in U$.
Finally, if K is symmetric, then $x \in K$ implies $-x \in K$. Applying the preceding inequality to both x and $-x$ gives

$$-1 \leq \langle x, u\rangle \leq 1 \quad (x \in K).$$

Therefore $|\langle x, u\rangle| \leq 1$ for every $x \in K$ and $u \in U$. □

Geometric interpretation. Suppose $B_2^n \subset K$ and let $u \in U = \partial K \cap S^{n-1}$ be a contact point. The Euclidean ball and K touch at u. The ball has a unique supporting hyperplane at this point, given by

$$\{x \in \mathbb{R}^n : \langle x, u\rangle = 1\},$$

whose outward unit normal is u. Since $B_2^n \subset K$, any supporting hyperplane of K at u must also support the ball. Its normal is therefore necessarily u, and we obtain

$$\langle x, u\rangle \leq 1 \quad (x \in K).$$

If K is symmetric, applying the same inequality to $-x \in K$ gives

$$|\langle x, u\rangle| \leq 1 \quad (x \in K).$$

Although this inequality comes from a single contact point, it holds for every $x \in K$, because a supporting hyperplane bounds the entire convex body. Thus each contact point provides a global constraint on the projection of points of K in its direction.

Question. A single contact direction does not control the Euclidean norm of every point of K. Why do the *active constraints* at the maximal ellipsoid collectively contain enough information to control points far from the contact points? John's identity decomposition will show that certain contact directions can be combined to recover the Euclidean norm, allowing us to add the squares of these inequalities and obtain a global bound on K.

1.1.10 Closed cone of contact matrices

The ambient vector space is

$$\mathrm{Sym}_n = \{H \in \mathbb{R}^{n\times n} : H^T = H\},$$

with inner product $\langle A, B\rangle_{\mathrm{HS}} = \mathrm{tr}(AB)$ and associated norm $\|A\|_{\mathrm{HS}} = \sqrt{\mathrm{tr}(A^2)}$. For $u \in S^{n-1}$, write $u \otimes u = uu^T$. Note

$$\mathrm{tr}(uu^T) = 1, \quad \langle H, uu^T\rangle_{\mathrm{HS}} = u^T H u.$$

Lemma 1.1.19 (Closed cone of contact matrices)**.** *If $U \subset S^{n-1}$ is nonempty and compact, then*

$$\mathcal{C} \coloneqq \mathrm{cone}\{uu^T : u \in U\} = \left\{ \sum_{j=1}^{m} t_j u_j u_j^T : t_j \geq 0, u_j \in U \right\}$$

is a closed convex cone in Sym_n.

Proof. By definition, $\mathcal{C}$ consists of all finite nonnegative linear combinations of matrices uu^T with $u \in U$. It is therefore closed under addition and multiplication by nonnegative scalars, hence a convex cone.
To prove closedness, consider

$$Q = \mathrm{conv}\{uu^T : u \in U\}.$$

Since U is compact and the map $u \mapsto uu^T$ is continuous, the set $\{uu^T : u \in U\}$ is compact. Its convex hull Q is also compact, since the space of symmetric matrices is finite-dimensional, by Carathéodory's theorem.
For every $u \in U \subset S^{n-1}$, we have

$$\mathrm{tr}(uu^T) = |u|_2^2 = 1.$$

By linearity of the trace, every $V \in Q$ satisfies $\mathrm{tr}\, V = 1$. Moreover, every nonzero $M \in \mathcal{C}$ can be written as $M = sV$ with $s > 0$ and $V \in Q$. Taking traces gives

$$\mathrm{tr}\, M = s\,\mathrm{tr}\, V = s.$$

Thus the scaling factor is uniquely determined by $s = \mathrm{tr}\, M$.
Now let (M_j) be a sequence in $\mathcal{C}$ such that $M_j \to M$. Write $M_j = s_j V_j$ with $s_j \geq 0$ and $V_j \in Q$. By continuity of the trace,

$$s_j = \mathrm{tr}\, M_j \to \mathrm{tr}\, M = s \geq 0.$$

Since Q is compact, we may pass to a subsequence such that $V_j \to V$ for some $V \in Q$. Therefore

$$M = \lim_{j\to\infty} s_j V_j = sV.$$

Since $s \geq 0$ and $V \in Q$, we conclude that $M \in \mathcal{C}$. In particular, if $s = 0$, then $M = 0 \in \mathcal{C}$. Hence $\mathcal{C}$ is closed. □

1.1.11 Proof of John's contact-point decomposition [Theorem 1.1.11]

Proof. Recall that K is in John position, so $B_2^n \subset K$ is a maximum-volume ellipsoid. The contact set $U = \partial K \cap S^{n-1}$ is compact and nonempty. Indeed, if U were empty, the compact sphere S^{n-1} would lie in the interior of K, so K would contain a slightly larger Euclidean ball, contradicting maximality.
Let $\mathcal{C}$ be the closed convex cone generated by the matrices uu^T with $u \in U$, as established in

Lemma 1.1.19. We want to prove that $I_n \in \mathcal{C}$.
Step 1: Separation of the identity from the cone.
Suppose, for a contradiction, that $I_n \notin \mathcal{C}$. Since $\mathcal{C}$ is closed and convex in the finite-dimensional Euclidean space Sym_n, there exists a nearest point $Q \in \mathcal{C}$ to I_n. Set $H_0 = I_n - Q$. By the nearest-point variational inequality,

$$\langle H_0, Z - Q\rangle_{\mathrm{HS}} \le 0 \quad (Z \in \mathcal{C}).$$

Since $\mathcal{C}$ is a cone, both 0 and $2Q$ belong to $\mathcal{C}$. Taking $Z = 0$ and $Z = 2Q$ gives

$$\langle H_0, Q\rangle_{\mathrm{HS}} = 0.$$

Consequently,

$$\langle H_0, Z\rangle_{\mathrm{HS}} \le 0 \quad (Z \in \mathcal{C}).$$

Moreover, since $I_n = Q + H_0$ and $H_0 \neq 0$,

$$\operatorname{tr} H_0 = \langle H_0, I_n\rangle_{\mathrm{HS}} = \langle H_0, Q\rangle_{\mathrm{HS}} + \|H_0\|_{\mathrm{HS}}^2 = \|H_0\|_{\mathrm{HS}}^2 > 0.$$

For every $u \in U$, the matrix uu^T belongs to $\mathcal{C}$, so

$$u^T H_0 u = \langle H_0, uu^T\rangle_{\mathrm{HS}} \le 0.$$

To obtain strict negativity, choose $0 < \eta < \operatorname{tr} H_0/n$ and set $H = H_0 - \eta I_n$. Since $|u|_2 = 1$ for every $u \in U$, we have

$$\operatorname{tr} H = \operatorname{tr} H_0 - n\eta > 0, \quad u^T H u = u^T H_0 u - \eta \le -\eta.$$

Thus H has positive trace, while its quadratic form is strictly negative on every contact direction.
Step 2: Construction of a larger-volume ellipsoid.
For $t > 0$, consider the symmetric matrix $A_t = I_n + tH$ and the ellipsoid $A_t B_2^n$. Since H is symmetric, the eigenvalues of A_t are $1 + t\lambda_i(H)$. If $t\|H\|_{\mathrm{op}} < 1$, they are all positive, so A_t is positive definite.
By Lemma 1.1.14,

$$\det A_t = \det(I_n + tH) = 1 + t \operatorname{tr} H + O(t^2).$$

Since $\operatorname{tr} H > 0$, for every sufficiently small $t > 0$,

$$\det A_t > 1.$$

Therefore,

$$\mathrm{vol}_n(A_t B_2^n) = \det A_t \, \mathrm{vol}_n(B_2^n) > \mathrm{vol}_n(B_2^n).$$

To contradict maximality, it remains to show that $A_t B_2^n \subset K$ for some such t.
Step 3: Containment of the perturbed ellipsoid.
By Lemma 1.1.17, the inclusion $A_t B_2^n \subset K$ is equivalent to

$$h_{A_t B_2^n}(v) = |A_t v|_2 \le h_K(v) \quad (v \in S^{n-1}),$$

where we used the symmetry of A_t. We verify this inequality separately near the contact set and away from it.
First, consider the continuous function $q(v) = v^T H v$ on S^{n-1}. Since $q(u) \le -\eta$ for every $u \in U$ and U is compact, there exists an open neighborhood V of U in S^{n-1} such that

$$v^T H v \le -\frac{\eta}{2} \quad (v \in V).$$

For every $v \in V$, using $|v|_2 = 1$, we obtain

$$|A_t v|_2^2 = |v + tHv|_2^2 = 1 + 2tv^T H v + t^2 |Hv|_2^2 \le 1 - \eta t + t^2 \|H\|_{\mathrm{op}}^2.$$

For sufficiently small $t > 0$, the right-hand side is strictly less than 1. Since $B_2^n \subset K$ implies $h_K(v) \ge 1$, we conclude that

$$|A_t v|_2 < 1 \le h_K(v) \quad (v \in V).$$

Next, consider the compact set $S^{n-1} \setminus V$, which contains no contact points. By Lemma 1.1.18, $h_K(v) > 1$ on this set. Since h_K is continuous, if the complement is nonempty, there exists $\alpha > 0$ such that

$$\alpha = \min_{v \in S^{n-1} \setminus V} (h_K(v) - 1) > 0.$$

On the other hand, for every $v \in S^{n-1}$,

$$|A_t v|_2 = |v + tHv|_2 \leq 1 + t\|H\|_{\mathrm{op}}.$$

Thus, choosing $t > 0$ sufficiently small that $t\|H\|_{\mathrm{op}} \leq \alpha$, we obtain

$$|A_t v|_2 \leq 1 + \alpha \leq h_K(v) \quad (v \in S^{n-1} \setminus V).$$

If the complement is empty, this second estimate is unnecessary.
We may therefore choose a single sufficiently small $t > 0$ such that A_t is positive definite, $\det A_t > 1$, and

$$|A_t v|_2 \leq h_K(v) \quad (v \in S^{n-1}).$$

By the support-function criterion, this implies $A_t B_2^n \subset K$. But

$$\mathrm{vol}_n(A_t B_2^n) > \mathrm{vol}_n(B_2^n),$$

contradicting the maximality of B_2^n. Hence our assumption was false, and

$$I_n \in \mathcal{C}.$$

By the definition of $\mathcal{C}$, there exist contact points $u_1, \ldots, u_m \in U$ and coefficients $c_i > 0$ such that

$$I_n = \sum_{i=1}^{m} c_i u_i u_i^T.$$

It remains to bound the number of terms. Let $d = \dim \mathrm{Sym}_n = n(n+1)/2$. If $m > d$, the matrices $u_i u_i^T$ are linearly dependent, so there exist scalars a_i, not all zero, satisfying

$$\sum_{i=1}^{m} a_i u_i u_i^T = 0.$$

Taking traces and using $\mathrm{tr}(u_i u_i^T) = 1$ gives $\sum_i a_i = 0$. Thus some a_i are positive and others negative. Set

$$s = \min_{a_i > 0} \frac{c_i}{a_i}.$$

Replacing c_i by $c_i - s a_i$ preserves nonnegativity and the identity decomposition, since

$$\sum_{i=1}^{m} (c_i - s a_i) u_i u_i^T = I_n - s \sum_{i=1}^{m} a_i u_i u_i^T = I_n.$$

At least one coefficient becomes zero. Repeating this reduction and discarding zero coefficients yields a decomposition with

$$m \leq \frac{n(n+1)}{2}.$$

Finally, taking traces in the identity decomposition gives $\sum_i c_i = n$. □

Remark 1.1.20 (Why the proof has three logically distinct ingredients)**.** The argument combines three essential ideas:

1. *Optimization.* The maximality of the Euclidean ball provides the starting point: no admissible perturbation can increase its volume.

2. *Separation.* If the contact matrices do not positively generate I_n, separation produces a symmetric matrix H with $\mathrm{tr}\, H > 0$ and $u^T H u < 0$ for every contact point $u \in U$. This gives a volume-increasing perturbation that moves inward at the contact directions.

3. *Uniformity.* Strict negativity near the contact set and positive slack away from it allow us to choose one sufficiently small perturbation that remains inside K in every direction.

The last step is essential: checking the perturbation only at the contact points would not guarantee containment in K. Compactness of the sphere provides the uniform estimates needed to conclude.

1.1.12 From contact points to John's $\sqrt{n}$ theorem

Lemma 1.1.21 (The trace identity)**.** *If $I_n = \sum_i c_i u_i u_i^T$ with $|u_i|_2 = 1$, then*

$$\sum_i c_i = n, \quad |x|_2^2 = \sum_i c_i \langle x, u_i \rangle^2 \quad (x \in \mathbb{R}^n).$$

Proof. Taking traces in the identity decomposition and using $\operatorname{tr}(u_i u_i^T) = |u_i|_2^2 = 1$, we obtain

$$n = \operatorname{tr} I_n = \sum_i c_i \operatorname{tr}(u_i u_i^T) = \sum_i c_i.$$

For the second identity, let $x \in \mathbb{R}^n$. Multiplying the matrix identity on the left by x^T and on the right by x gives

$$|x|_2^2 = x^T I_n x = \sum_i c_i x^T u_i u_i^T x = \sum_i c_i \langle x, u_i \rangle^2.$$

Thus the squared Euclidean norm can be expressed as a weighted sum of squared projections onto the contact directions. □

1.1.13 Proof of Theorem 1.1.10

Proof. By Lemmas 1.1.12 and 1.1.13, a maximum-volume ellipsoid exists and can be chosen centered at the origin. After an invertible linear change of coordinates, we may put K in John position, so that B_2^n is a maximum-volume ellipsoid contained in K. In particular,

$$B_2^n \subset K.$$

It remains to prove the reverse inclusion $K \subset \sqrt{n} B_2^n$.
By Theorem 1.1.11, there exist contact points $u_1, \dots, u_m \in \partial K \cap S^{n-1}$ and positive coefficients $c_1, \dots, c_m$ such that

$$I_n = \sum_{i=1}^m c_i u_i u_i^T.$$

Since K is symmetric, Lemma 1.1.18 gives

$$|\langle x, u_i \rangle| \leq 1 \quad (x \in K).$$

Thus, for every $x \in K$, we have $\langle x, u_i \rangle^2 \leq 1$. Applying Lemma 1.1.21 and using $c_i > 0$, we obtain

$$|x|_2^2 = \sum_{i=1}^m c_i \langle x, u_i \rangle^2 \leq \sum_{i=1}^m c_i = n.$$

Hence $|x|_2 \leq \sqrt{n}$ for every $x \in K$, which means

$$K \subset \sqrt{n} B_2^n.$$

Combining both inclusions, we conclude that

$$B_2^n \subset K \subset \sqrt{n} B_2^n.$$

□

Norm formulation. Recall that inclusions of convex bodies reverse the inequalities between their gauges. Since

$$B_2^n \subset K \subset \sqrt{n} B_2^n,$$

we obtain

$$\|x\|_{\sqrt{n} B_2^n} \leq \|x\|_K \leq \|x\|_{B_2^n}.$$

Using $\|x\|_{B_2^n} = |x|_2$ and $\|x\|_{\sqrt{n}B_2^n} = |x|_2/\sqrt{n}$, this becomes

$$\frac{1}{\sqrt{n}}|x|_2 \le \|x\|_K \le |x|_2.$$

Now consider the identity map $I : X \to \ell_2^n$, where $X = (\mathbb{R}^n, \|\cdot\|_K)$. By the definition of the operator norm, these inequalities give

$$\|I\|_{X\to\ell_2^n} \le \sqrt{n}, \quad \|I^{-1}\|_{\ell_2^n\to X} \le 1.$$

Therefore, by Proposition 1.1.4,

$$d_{\mathrm{BM}}(X, \ell_2^n) \le \|I\|\|I^{-1}\| \le \sqrt{n}.$$

Since the initial change of coordinates is a linear isometry between the corresponding normed spaces, the same bound holds for the original space X.

1.1.14 Interlude: John's decomposition as an exact Euclidean measurement system

John's identity decomposition has an interpretation beyond convex geometry: it allows us to recover the Euclidean norm of any vector from finitely many weighted projections onto contact directions.

Proposition 1.1.22 (A weighted tight frame from John contact points)**.** *Suppose that $u_1, \dots, u_m \in S^{n-1}$ and $c_1, \dots, c_m > 0$ satisfy*

$$I_n = \sum_{i=1}^m c_i u_i u_i^T.$$

Then, for every $x \in \mathbb{R}^n$,

$$|x|_2^2 = \sum_{i=1}^m c_i \langle x, u_i\rangle^2.$$

Equivalently, the linear map $\Phi : \mathbb{R}^n \to \mathbb{R}^m$ defined by

$$\Phi x = (\sqrt{c_1}\langle x, u_1\rangle, \dots, \sqrt{c_m}\langle x, u_m\rangle)$$

is an isometric embedding, meaning that $|\Phi x|_2 = |x|_2$ for every $x \in \mathbb{R}^n$.

Proof. Multiplying the identity decomposition on the left by x^T and on the right by x, we obtain

$$|x|_2^2 = x^T I_n x = \sum_{i=1}^m c_i x^T u_i u_i^T x = \sum_{i=1}^m c_i \langle x, u_i\rangle^2.$$

On the other hand, by the definition of Φ,

$$|\Phi x|_2^2 = \sum_{i=1}^m (\sqrt{c_i}\langle x, u_i\rangle)^2 = \sum_{i=1}^m c_i \langle x, u_i\rangle^2 = |x|_2^2.$$

Thus Φ preserves the Euclidean norm. Since it is linear, it also preserves distances, and is therefore a linear isometric embedding. □

The weighted vectors $\sqrt{c_i}u_i$ form a *Parseval frame*, also called a *tight frame with frame bound* 1. Unlike an orthonormal basis, these vectors need not be orthogonal, and their number may exceed the dimension n. Nevertheless, their weighted projections recover the Euclidean norm exactly.

In fact, applying the matrix identity directly to x gives the reconstruction formula

$$x = \sum_{i=1}^m c_i \langle x, u_i\rangle u_i.$$

Thus the directional measurements determine not only the norm of x, but the vector itself. This interpretation is useful beyond John's theorem, particularly in the study of Euclidean sections and random embeddings.

Example 1.1.23 (Two easy frames in the plane)**.** First, consider the standard coordinate directions $u_1 = e_1$ and $u_2 = e_2$, with weights $c_1 = c_2 = 1$. Since

$$I_2 = e_1 e_1^T + e_2 e_2^T,$$

the frame identity becomes

$$|x|_2^2 = \langle x, e_1\rangle^2 + \langle x, e_2\rangle^2 = x_1^2 + x_2^2.$$

This is the usual Parseval identity for an orthonormal basis. Geometrically, these directions correspond to the contact points $\pm e_1, \pm e_2$ between the Euclidean unit disk and the square B_∞^2.

A less obvious example uses three nonorthogonal unit vectors

$$u_j = (\cos(j\pi/3), \sin(j\pi/3)), \quad j = 0, 1, 2.$$

Explicitly,

$$u_0 = (1, 0), \quad u_1 = (1/2, \sqrt{3}/2), \quad u_2 = (-1/2, \sqrt{3}/2).$$

Computing their outer products gives

$$\sum_{j=0}^{2} u_j u_j^T = \begin{pmatrix} 1 + 1/4 + 1/4 & \sqrt{3}/4 - \sqrt{3}/4 \\ \sqrt{3}/4 - \sqrt{3}/4 & 3/4 + 3/4 \end{pmatrix} = \frac{3}{2} I_2.$$

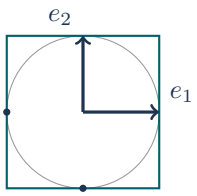
Thus, choosing equal weights $c_j = 2/3$, we obtain

$$I_2 = \frac{2}{3} \sum_{j=0}^{2} u_j u_j^T,$$

and therefore

$$|x|_2^2 = \frac{2}{3} \sum_{j=0}^{2} \langle x, u_j\rangle^2.$$

Hence three nonorthogonal directions recover the Euclidean norm exactly when their squared projections are weighted appropriately. This decomposition will reappear in Lecture 2 when constructing a hexagonal section of the three-dimensional cube.



The unit disk touches the square at $\pm e_1, \pm e_2$

1.1.15 The converse: why a decomposition certifies optimality

Theorem 1.1.24 (Converse contact criterion)**.** *Suppose $K = -K$ and $B_2^n \subset K$. If there are contact points $u_i \in \partial K \cap S^{n-1}$ and $c_i > 0$ with*

$$I_n = \sum_i c_i u_i u_i^T,$$

then B_2^n is a maximal-volume inscribed ellipsoid.

Proof. By Lemma 1.1.12, it suffices to test centered ellipsoids. Write such an ellipsoid as $AB_2^n \subset K$ with A symmetric positive definite. For each contact point, using the support function of the ellipsoid,

$$|Au_i|_2 = h_{AB_2^n}(u_i) \le h_K(u_i) = 1.$$

Multiply the decomposition of the identity by A^2 and take traces:

$$\operatorname{tr}(A^2) = \operatorname{tr}(A^2 I_n) = \sum_i c_i \operatorname{tr}(A^2 u_i u_i^T) = \sum_i c_i |Au_i|_2^2 \le \sum_i c_i = n.$$

If $\lambda_1, \dots, \lambda_n > 0$ are the eigenvalues of A, the arithmetic–geometric mean inequality gives

$$(\det A)^{2/n} = \left(\prod_i \lambda_i^2 \right)^{1/n} \le \frac{1}{n} \sum_i \lambda_i^2 = \frac{\operatorname{tr}(A^2)}{n} \le 1.$$

Thus $\det A \le 1$, and no centered inscribed ellipsoid has larger volume than B_2^n. Centering takes care of translated ellipsoids. □

1.1.16 Concrete contact decompositions

Example 1.1.25 (The cube)**.** For $K = B_\infty^n$, $B_2^n \subset K$. Each e_i is a contact point, since $|e_i|_2 = 1$ and $\|e_i\|_\infty = 1$. The identity

$$I_n = \sum_{i=1}^n e_i e_i^T$$

certifies, by Theorem 1.1.24, that B_2^n is the John ellipsoid of the cube. The vertex $(1, \dots, 1)$ has Euclidean norm $\sqrt{n}$, so the outer inclusion is sharp in this position.

Example 1.1.26 (The cross-polytope)**.** The inequality $\|x\|_1 \le \sqrt{n}|x|_2$ shows that $(1/\sqrt{n})B_2^n \subset B_1^n$. Put $K = \sqrt{n}B_1^n$, so $B_2^n \subset K$. For each sign vector $\varepsilon \in \{-1, +1\}^n$, set

$$u_\varepsilon = \frac{1}{\sqrt{n}}(\varepsilon_1, \dots, \varepsilon_n).$$

Then $|u_\varepsilon|_2 = 1$ and $\|u_\varepsilon\|_1 = \sqrt{n}$, so $u_\varepsilon \in \partial K \cap S^{n-1}$. Furthermore,

$$\frac{1}{2^n} \sum_\varepsilon u_\varepsilon u_\varepsilon^T = \frac{1}{n} I_n :$$

the diagonal entries are $1/n$, and the off-diagonal entries vanish because the signs are independent and symmetric. Hence

$$I_n = \frac{n}{2^n} \sum_\varepsilon u_\varepsilon u_\varepsilon^T.$$

Theorem 1.1.24 proves that B_2^n is the John ellipsoid of $\sqrt{n}B_1^n$.

1.1.17 Extremal questions in Banach–Mazur geometry

John's theorem shows that every n-dimensional normed space is at Banach–Mazur distance at most $\sqrt{n}$ from Euclidean space. Combined with the multiplicative triangle inequality, this gives a universal bound for the distance between any two normed spaces of the same dimension.

Corollary 1.1.27 (A universal bound for two normed spaces)**.** *If* $\dim X = \dim Y = n$*, then*

$$d_{\mathrm{BM}}(X, Y) \le n.$$

Proof. By the multiplicative triangle inequality, using ℓ_2^n as an intermediate space,

$$d_{\mathrm{BM}}(X, Y) \le d_{\mathrm{BM}}(X, \ell_2^n) d_{\mathrm{BM}}(\ell_2^n, Y).$$

John's theorem gives $d_{\mathrm{BM}}(X, \ell_2^n) \le \sqrt{n}$ and, by symmetry, $d_{\mathrm{BM}}(\ell_2^n, Y) = d_{\mathrm{BM}}(Y, \ell_2^n) \le \sqrt{n}$. Therefore

$$d_{\mathrm{BM}}(X, Y) \le \sqrt{n} \cdot \sqrt{n} = n.$$

□

Definition 1.1.28 (Banach–Mazur compactum)**.** Let BM_n denote the set of linear-isometry classes of real n-dimensional normed spaces. By Proposition 1.1.5,

$$\rho(X, Y) = \log d_{\mathrm{BM}}(X, Y)$$

defines a metric on BM_n. The space (BM_n, ρ) is called the *Banach–Mazur compactum*. We also define its *multiplicative diameter* by

$$\operatorname{diam}_{\mathrm{mult}}(BM_n) = \sup_{X, Y \in BM_n} d_{\mathrm{BM}}(X, Y).$$

The compactness of BM_n follows from John's theorem and the Blaschke selection theorem. Indeed, every isometry class admits a representative whose unit ball K is in John position, so

$$B_2^n \subset K \subset \sqrt{n} B_2^n.$$

Thus the normalized unit balls form a uniformly bounded family of convex bodies. By the Blaschke selection theorem, every sequence of such bodies admits a subsequence converging in the Hausdorff metric. The limit remains a symmetric convex body containing B_2^n. Moreover, Hausdorff convergence $K_j \to K$ in this normalized family implies that, for every $\varepsilon > 0$ and all sufficiently large j, $(1-\varepsilon)K \subset K_j \subset (1+\varepsilon)K$. Thus the corresponding Banach–Mazur distances tend to 1, and BM_n is compact.

The universal bound from the preceding corollary immediately gives

$$\operatorname{diam}_{\mathrm{mult}}(BM_n) \le n.$$

This leads to two distinct extremal questions.

1. *Distance to Euclidean space.* John's theorem gives

$$\sup_{\dim X = n} d_{\mathrm{BM}}(X, \ell_2^n) \le \sqrt{n}.$$

   Is this bound sharp? A natural candidate is ℓ_∞^n, whose unit ball is the cube B_∞^n. The elementary inclusions

$$B_2^n \subset B_\infty^n \subset \sqrt{n} B_2^n$$

   give $d_{\mathrm{BM}}(\ell_\infty^n, \ell_2^n) \le \sqrt{n}$. However, these inclusions only use the standard coordinates and do not exclude the possibility of a better invertible linear transformation.

   Lecture 2 proves the matching lower bound

$$d_{\mathrm{BM}}(\ell_\infty^n, \ell_2^n) = \sqrt{n}.$$

   Consequently,

$$\sup_{\dim X = n} d_{\mathrm{BM}}(X, \ell_2^n) = \sqrt{n}.$$

   Thus the cube is extremal for the distance to Euclidean space, and the factor $\sqrt{n}$ in John's theorem cannot be improved.

2. *Diameter of the entire compactum.* We now allow both spaces to vary and consider

$$\operatorname{diam}_{\mathrm{mult}}(BM_n) = \sup_{X,Y \in BM_n} d_{\mathrm{BM}}(X, Y).$$

   The universal bound gives an upper estimate of n, but can its dependence on the dimension be improved?

   A theorem of Gluskin shows that there exists an absolute constant $c > 0$ such that

$$cn \le \operatorname{diam}_{\mathrm{mult}}(BM_n) \le n.$$

   Thus the multiplicative diameter is of order n. In particular, there exist pairs of n-dimensional normed spaces whose Banach–Mazur distance is at least cn. The lower bound relies on probabilistic constructions of convex bodies and will not be proved here.

   In the usual metric $\rho = \log d_{\mathrm{BM}}$, taking logarithms gives

$$\log(cn) \le \operatorname{diam}_\rho(BM_n) \le \log n.$$

   Hence the metric diameter of BM_n grows asymptotically as $\log n$.

The distinction between these two problems is important. When one space is fixed to be Euclidean, the maximal distance is exactly $\sqrt{n}$. When both spaces are allowed to vary, the maximal distance is of order n. Thus the universal estimate obtained from John's theorem and the multiplicative triangle inequality is optimal in its dependence on the dimension, up to an absolute constant.

The first extremal problem is the subject of Lecture 2, where we prove that no linear change of coordinates can improve the $\sqrt{n}$ bound for the cube.

References

- F. John, *Extremum problems with inequalities as subsidiary conditions*, in *Studies and Essays Presented to R. Courant*, 1948.
- K. Ball, *An elementary introduction to modern convex geometry*, in *Flavors of Geometry*, MSRI Publications 31, 1997.
- S. Artstein-Avidan, A. Giannopoulos and V. D. Milman, *Asymptotic Geometric Analysis, Part I*, AMS, 2015.
- A. Giannopoulos and V. D. Milman, *Euclidean structure in finite dimensional normed spaces*, in *Handbook of the Geometry of Banach Spaces*, Vol. 1, 2001 (survey; includes Banach–Mazur geometry and Gluskin's diameter theorem).
- B. Grünbaum, *Projection constants*, Trans. Amer. Math. Soc. 95 (1960), 451–465.
- A. Pełczyński, *Structural theory of Banach spaces and its interplay with analysis and probability*, ICM Warsaw 1983 proceedings (1984), 237–269.
- J. Bourgain and S. J. Szarek, *The Banach–Mazur distance to the cube and the Dvoretzky–Rogers factorization*, Israel J. Math. 62 (1988), 169–180.
- S. J. Szarek and M. Talagrand, *An "isomorphic" version of the Sauer–Shelah lemma and the Banach–Mazur distance to the cube*, Lecture Notes in Math. 1376 (1989), 105–112.
- S. J. Szarek, *Spaces with large distance to ℓ_∞^n and random matrices*, Amer. J. Math. 112 (1990), 899–942.
- A. A. Giannopoulos, *A note on the Banach–Mazur distance to the cube*, Oper. Theory Adv. Appl. 77 (1995), 67–73.
- P. Youssef, *Restricted invertibility and the Banach–Mazur distance to the cube*, Mathematika 60 (2014), 201–218; doi:10.1112/S0025579313000144.
- K. Tikhomirov, *On the Banach–Mazur distance to cross-polytope*, Adv. Math. 345 (2019), 598–617; arXiv:1804.08212.
- A. Naor and P. Youssef, *Restricted invertibility revisited*, in *A Journey Through Discrete Mathematics*, Springer, 2017, 657–691.
- O. Friedland, *An improved lower bound on the Banach–Mazur distance to the cross-polytope*, 2026 preprint; arXiv:2602.10665.
- O. Friedland, *A 5/8 lower bound on the Banach–Mazur distance to the cross-polytope*, 2026 preprint; arXiv:2608.17743.
- A. Hmadi, *Gaussian polytopes with large Banach–Mazur distance to the cross-polytope*, 2026 preprint; arXiv:2609.27796.
- O. Friedland, A. Naor and P. Youssef, *On the Banach–Mazur radius of ℓ_∞^n*, 2026 preprint; arXiv:2608.21112.
- O. Friedland, *The sharp exponent for the Banach–Mazur distance to the cube*, 2026 preprint; arXiv:2610.36917.