High-dimensional tennis balls

W. T. Gowers

University of Cambridge

May 30, 2020

Joint work with Kasia Wyczesany

Joint work with Kasia Wyczesany

What do we mean by a high-dimensional tennis ball?

Joint work with Kasia Wyczesany

What do we mean by a high-dimensional tennis ball?

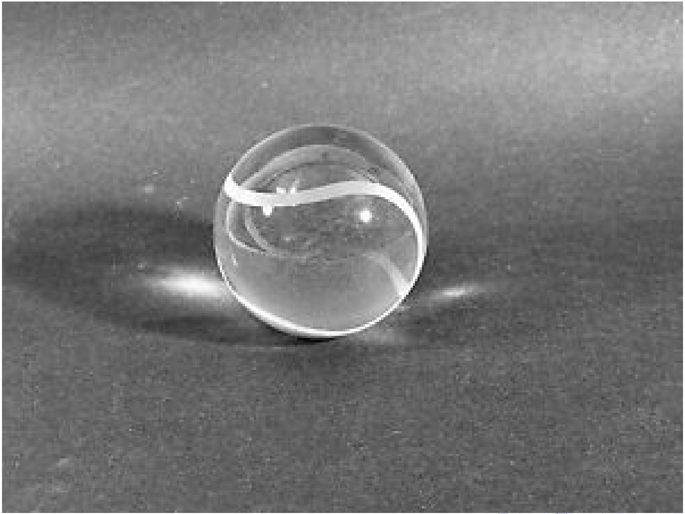

Adding interesting conditions to extremal problems

Example 1. Turán's theorem states that the graph with the largest number of edges and no K_r is a complete $(r-1)$-partite graph with vertex classes of sizes as equal as possible. But such graphs have huge independent sets. What happens if in addition the largest independent set has $o(n)$ vertices?

Adding interesting conditions to extremal problems

Example 1. Turán's theorem states that the graph with the largest number of edges and no K_r is a complete $(r-1)$-partite graph with vertex classes of sizes as equal as possible. But such graphs have huge independent sets. What happens if in addition the largest independent set has $o(n)$ vertices?

Example 2. The Erdős-Ko-Rado theorem states that if $k < n/2$, then the largest intersecting families of k-sets from $[n]$ are those of the form $\{A \subset [n] : |A| = k, x \in A\}$. What happens if in addition we insist that the family is invariant under a transitive group of permutations of $[n]$?

An extra condition with a slight twist

How large can a subset of the n-sphere be if its ϵ-expansion does not contain a linear subsphere of dimension k?

$$\downarrow$$

How 'high-dimensional' can a subset of the n-sphere be if its ϵ-expansion does not contain the unit sphere of a 2-dimensional subspace?

Some definitions

An m-dimensional *topological subsphere* of S^n is the image of a continuous odd function $f : S^m \to S^n$. It is a *linear subsphere* if it is the intersection of S^n with an $(m+1)$-dimensional linear subspace of $\mathbb{R}^{n+1}$. A 1-dimensional linear subsphere is a *great circle*.

Some definitions

An m-dimensional *topological subsphere* of S^n is the image of a continuous odd function $f : S^m \to S^n$. It is a *linear subsphere* if it is the intersection of S^n with an $(m+1)$-dimensional linear subspace of $\mathbb{R}^{n+1}$. A 1-dimensional linear subsphere is a *great circle*.

Remark. An m-dimensional topological subspace must intersect every $(n-m)$-dimensional linear subspace.

Proof. WLOG the subspace is $\{x \in S^n : x_1 = \cdots = x_m = 0\}$. But if $f : S^m \to S^n$ is a continuous function and P_m is the projection to the first m coordinates, then by Borsuk-Ulam we can find $x \in S^m$ such that $P_m f(x) = P_m f(-x)$, and since f is odd it follows that $P_m f(x) = 0$. $\square$

(Note that the conclusion is completely false for a union of a spherical cap and minus that cap.)

Main result

There are absolute constants $c, \epsilon > 0$ such that for every n, S^n contains a topological subspace X of dimension at least cn such that X_ϵ contains no great circle.

Informally, tennis balls exist with linear-dimensional seams!

Motivation

Question. *(Milman) Let $k \in \mathbb{N}, C \in \mathbb{R}, \epsilon > 0$. Does there exist $n_0(k, C, \epsilon)$ such that if $n \geq n_0$, then every n-dimensional space X with $d(X, \ell_2^n) \leq C$ has a k-dimensional $(1+\epsilon)$-complemented subspace Y such that $d(Y, \ell_2^k) \leq 1 + \epsilon$?*

Not known even for $k = 2$.

A natural weakening

Let $X = (\mathbb{R}^n, \|.\|)$ be a normed space. Say that a subspace $Y \subset X$ is *strongly α-Euclidean* if there exists t such that $t\|y\|_2 \leq \|y\| \leq \alpha t\|y\|_2$ for every $y \in Y$.

Say that Y is *strongly α-complemented* if the orthogonal projection to Y has norm at most α.

'Strong Milman question.' *Let $k \in \mathbb{N}, C \in \mathbb{R}, \epsilon > 0$. Does there exist $n_0(k, C, \epsilon)$ such that if $n \geq n_0$, then every n-dimensional space X that is strongly C-Euclidean has a k-dimensional subspace Y that is strongly $(1 + \epsilon)$-Euclidean and strongly $(1 + \epsilon)$-complemented?*

A reformulation

Let $x \in X$, $x \neq 0$, and let P_x be the orthogonal projection to the 1D subspace generated by x. Say that x is *ϵ-good* if $\|P_x\| \leq 1 + \epsilon$.

Geometrical interpretation: if $\|x\| = 1$, then x is ϵ-good if all of the unit ball of X lies on the same side of the hyperplane that is orthogonal to x and that goes through the point $(1+\epsilon)x$.

Algebraic interpretation: x is ϵ-good if the inequality

$$\langle x, y\rangle \leq (1+\epsilon)\frac{\|y\|}{\|x\|}\|x\|_2^2$$

holds for every $y \in X$.

Remark. If $\|x\|/\|x\|_2$ is near to either its maximum or its minimum, then x is ϵ-good for a small ϵ.

A useful equivalence

Y is strongly $(1+\epsilon)$-Euclidean and strongly $(1+\epsilon)$-complemented if and only if every point in Y is δ-good.

Are almost all points ϵ-good?

Are almost all points ϵ-good?

Counterexample. $\|x\|^2 = x_1^2 + \cdots + x_{n/2}^2 + 2(x_{n/2+1}^2 + \cdots + x_n^2)$.

The ϵ-good points are close to the eigenspaces, so their total measure is exponentially small.

More general ellipsoidal norms also work for similar reasons.

A more interesting example

Let $\|x\|^2 = \max_{|A|=n/2} \sum_{i \in A} x_i^2$.

This has a symmetric basis, so is in 'the right position', but ϵ-good points are still quite special.

Moral. We can't hope to use concentration of measure to obtain a positive answer.

Observation

The set of ϵ-good points appears to be 'high dimensional'. Also, a small perturbation of an ϵ-good point is an ϵ'-good point.

Observation

The set of ϵ-good points appears to be ‘high dimensional’. Also, a small perturbation of an ϵ-good point is an ϵ'-good point.

The connection. Must a small expansion of a high-dimensional set contain a k-dimensional linear subsphere?

Observation

The set of ϵ-good points appears to be 'high dimensional'. Also, a small perturbation of an ϵ-good point is an ϵ'-good point.

The connection. Must a small expansion of a high-dimensional set contain a k-dimensional linear subsphere?

(**Question.** Is it true that the set of ϵ-good points has to be high dimensional?)

Crude random constructions don't give counterexamples to strong Milman question

1. Let N be exponential in n, let $x_1, \dots, x_N$ be random unit vectors, and let $\|x\| = \max |\langle x, x_i \rangle|$.

2. Let $x_1, \dots, x_N$ be as above and let the unit ball of a space X be the convex hull of the $\pm x_i$.

Our tennis ball construction

Basic idea: define a continuous 'distorting' map $\psi : S^n \to S^n$ and take the image of an m-dimensional linear subsphere.

Our tennis ball construction

Basic idea: define a continuous 'distorting' map $\psi : S^n \to S^n$ and take the image of an m-dimensional linear subsphere.

More detail: define a continuous map $\phi : \mathbb{R} \to \mathbb{R}$, apply it pointwise, and normalize to get ψ; then apply ψ to a random m-dimensional subsphere.

Yet more detail

Let $\sigma : \mathbb{R} \to \mathbb{R}$ be the function that takes t to 2^{2k} when $2^{2k-1} < t \leq 2^{2k+1}$ and let ϕ be a smooth (except at 0) strictly increasing odd function that approximates σ when $t > 0$, such that ϕ' is bounded away from 0 and ∞. For convenience we assume that $\phi(4t) = 4\phi(t)$ for every t.

If $x \in \mathbb{R}^{n+1}$, write $\phi(x)$ for $(\phi(x_1), \ldots, \phi(x_{n+1}))$, and let $\psi(x) = \frac{\phi(x)}{\|\phi(x)\|_2}$.

ψ is a bi-Lipschitz bijection from S^n to S^n.

What we need to prove

Let $c, \epsilon > 0$ be small enough.

We would like

If X is a random linear subsphere of dimension cn, then $\psi(X)_\epsilon$ contains no great circle.

for which it is sufficient to prove

There is a set $\Gamma \subset S^n$ such that $S^n \setminus \Gamma$ has exponentially small measure and every great circle contains a point x such that $d(x, \psi(\Gamma_\epsilon)) > \epsilon$.

Why might we expect that to be true?

There is a set $\Gamma \subset S^n$ such that $S^n \setminus \Gamma$ has exponentially small measure and every great circle contains a point x such that $d(x, \psi(\Gamma_\epsilon)) > \epsilon$.

Oversimplified argument

Choose $x = (x_1, \dots, x_n)$ at random. Then with high probability most values of $\phi(x_i)$ are close to $\pm$ a power of 4. So most ratios $\psi(x_i)/\psi(x_j)$ are close to a power of 4.

But an averaging argument shows that in a great circle, most vectors have many coordinate ratios that are not close to a power of 4.

What does 'most' mean?

Let E be a subset of $\{1, 2, \ldots, n+1\}$ and let $x \in S^n$. Define $\mathbb{P}_x(E)$ to be $\sum_{i \in E} x_i^2$.

Similarly for pairs of coordinates.

So the statement 'Most ratios $|u_i|/|u_j|$ are close to a power of 4' means that

$$\sum \{u_i^2 u_j^2 : |u_i|/|u_j| \text{ is close to a power of } 4\}$$

is close to 1.

Concluding remarks

Would guess that Milman's question has a negative answer.

Would also guess that there are *cn*-dimensional tennis balls for every $c \in (0, 1)$ – possible barrier at $c = 1/2$.

Our example rules out natural proof strategies. Currently looking for a smoother example to give a counterexample to the (strong at first) Milman question.

There is a 2-colouring $S^n = R \cup B$ such that $R_{\epsilon/2}$ does not contain the sphere of a $(1-c)n$-dimensional subspace and $B_{\epsilon/2}$ does not contain the sphere of a 2-dimensional subspace.

(Take B to be $X_{\epsilon/2}$, where X is a *cn*-dimensional tennis ball.)