The endpoint case of the Bennett–Carbery–Tao multilinear Kakeya conjecture

by

Larry Guth

University of Toronto
Toronto, ON, Canada

1. Introduction

In [1], Bennett, Carbery and Tao formulated a multilinear Kakeya conjecture, and they proved the conjecture except for the endpoint case. In this paper, we slightly sharpen their result by proving the endpoint case of the conjecture.

Our method of proof is very different from the proof of Bennett, Carbery and Tao. The original proof was based on monotonicity estimates for heat flows. In 2007, Dvir [2] made a breakthrough on the Kakeya problem, proving the Kakeya conjecture over finite fields. His proof used polynomials in a crucial way. It was not clear whether Dvir's approach could be adapted to prove estimates in Euclidean space. Our proof of the multilinear Kakeya conjecture is based on Dvir's polynomial method. In my opinion, the method of proof is as interesting as the result.

The multilinear Kakeya conjecture concerns the overlap properties of cylindrical tubes in $\mathbb{R}^n$. Roughly, the (multilinear) Kakeya conjecture says that cylinders pointing in different directions cannot overlap too much.

Before coming to the Bennett–Carbery–Tao multilinear estimate, we want to state a weaker result, because it is easier to understand and easier to prove. To be clear about the notation, a cylinder of radius R around a line $L \subset \mathbb{R}^n$ is the set of all points $x \in \mathbb{R}^n$ within a distance R of the line L. We call the line L the *core* of the cylinder.

THEOREM 1.1. *Suppose that we have a finite collection of cylinders $T_{j,a} \subset \mathbb{R}^n$, where $1 \leqslant j \leqslant n$, and $1 \leqslant a \leqslant A$ for some integer A. Each cylinder has radius 1. Moreover, each cylinder $T_{j,a}$ runs nearly parallel to the x_j-axis. More precisely, we assume that the angle between the core of $T_{j,a}$ and the x_j-axis is at most $(100n)^{-1}$.*

Let I be the set of points that belong to at least one cylinder in each direction:

$$I := \bigcap_{j=1}^{n} \left(\bigcup_{a=1}^{A} T_{j,a} \right).$$

Then $\mathrm{Vol}(I) \leqslant C(n) A^{n/(n-1)}$.

As Bennett, Carbery and Tao point out in [1], this estimate can be viewed as a generalization of the Loomis–Whitney inequality.

THEOREM 1.2. (Special case of Loomis and Whitney [11]) *Let U be an open set in $\mathbb{R}^n$. Let π_j denote the projection from $\mathbb{R}^n$ onto the hyperplane perpendicular to the x_j-axis. Suppose that for each j, $\pi_j(U)$ has $(n-1)$-dimensional volume at most B.*
    Then $\mathrm{Vol}(U) \leqslant B^{n/(n-1)}$.

Suppose that each tube $T_{j,a}$ runs exactly parallel to the x_j-axis. It follows that $\pi_j(I)$ is contained in A unit balls and has volume at most $\omega_{n-1} A$. Applying the Loomis–Whitney inequality, we see that the volume of I is bounded by $\lesssim A^{n/(n-1)}$. Theorem 1.1 says that—up to a constant factor—this volume estimate continues to hold if we allow the tubes to tilt slightly.

The Loomis–Whitney inequality is sharp whenever the open set U is a cube. Similarly, Theorem 1.1 is essentially sharp whenever the tubes are arranged in a cubical lattice.

The proof of Theorem 1.1 uses the polynomial method of Dvir. The main new idea in the paper is a new approach for adapting Dvir's method to $\mathbb{R}^n$. The new approach uses algebraic topology. In particular, we will use a polynomial generalization of the ham-sandwich theorem, proven by Stone and Tukey [13] in the early 1940s.

Now we turn to the multilinear version of the Kakeya maximal conjecture, formulated by Bennett, Carbery and Tao.

THEOREM 1.3. (Multilinear Kakeya estimate) *For each $1 \leqslant j \leqslant n$, let $T_{j,a}$ be a collection of unit cylinders, where a runs from 1 to $A(j)$. We let $v_{j,a}$ be a unit vector parallel to the core of $T_{j,a}$. We assume that the cylinders from different classes are quantitatively transverse in the sense that any determinant of a matrix $(v_{1,a_1}, v_{2,a_2}, ..., v_{n,a_n})$ has norm at least $\theta > 0$.*

Under these hypotheses, the following inequality holds:

$$\int \left(\prod_{j=1}^{n} \left(\sum_{a=1}^{A(j)} \chi_{T_{j,a}} \right) \right)^{1/(n-1)} d\mathrm{vol} \leqslant C(n) \theta^{-1/(n-1)} \prod_{j=1}^{n} A(j)^{1/(n-1)}.$$

Theorem 1.3 generalizes Theorem 1.1. If each vector $v_{j,a}$ lies within a small angle of the x_j-axis, then the determinant condition is easy to check, and so Theorem 1.3 applies. Recall that I is the set of points lying in at least one cylinder with each value of j. At every point $x \in I$, the integrand in Theorem 1.3 is at least 1. Hence Theorem 1.3 gives an upper bound for the volume of I, recovering Theorem 1.1.

Theorem 1.3 improves on Theorem 1.1 in the following ways. First, we allow a more general condition on the angles of the tubes. Second, we allow the different classes to have different numbers of tubes: $A(j)$ depends on j. Third, and most importantly, we get an integral bound where the integrand is very large at points which lie in many tubes from each direction.

The paper [1] has a very nice introductory discussion of the multilinear Kakeya estimate. Some of the topics it describes are the original Kakeya conjecture, and linear and multilinear restriction estimates. Using their multilinear Kakeya estimates, Bennett, Carbery and Tao are able to prove nearly optimal multilinear restriction estimates. We refer to that paper for more context.

The proof of Theorem 1.3 is harder than the proof of Theorem 1.1. It uses more sophisticated tools from algebraic topology: cohomology classes, cup products, and the Lyusternik–Shnirel'man vanishing theorem. Theorem 1.3 is more important than Theorem 1.1 because Bennett, Carbery and Tao use Theorem 1.3 to prove L^p estimates for multilinear restriction operators. On the other hand, Theorem 1.1 contains the main ideas of this paper, and its proof is only three pages long.

This paper uses algebraic topology. We want it to be understandable to mathematicians who work in analysis and combinatorics, so we will try to introduce the algebraic topology in a friendly way. In particular, there is a short section introducing Lyusternik–Shnirel'man theory, and an appendix giving the proof of the Lyusternik–Shnirel'man vanishing lemma.

As a corollary of our method, we give a "planiness" estimate for unions of tubes in $\mathbb{R}^n$. An estimate of this kind can also be proven using the methods of [1], but the one below is slightly sharper. The phenomenon of "planiness" was discovered by Katz, Laba and Tao in [10], and the estimate below is similar to some estimates from that paper.

COROLLARY 1.4. (Box estimate) *There is a constant $C(n) > 0$ such that the following holds. Suppose that $X \subset \mathbb{R}^n$ is a union of cylinders with radius 1 and length $L \gg 1$. For each $x \in X$ we can choose a rectangular box $B(x)$ with the following properties*:

(1) The box $B(x)$ is centered at x. It may be oriented in any direction. It has volume at most $C(n) \operatorname{Vol}(X)$;

(2) For every cylinder $T \subset X$ of radius 1 and length L, if we pick a random point $x \in T$, then with probability at least $\frac{9}{10}$, the tube T lies in the box $B(x)$.

Acknowledgements. I would like to thank Nets Katz for showing me the multilinear Kakeya estimates in [1]. I showed him the proof of the box estimate, and he explained to me how that estimate is related to multilinear Kakeya estimates and the work of Bennett, Carbery and Tao. I would also like to thank Kannan Soundararajan for interesting conversations about combinatorial number theory. In particular, he pointed out to me Dvir's paper [2].

2. The polynomial ham-sandwich theorem

The main tool in our proof is a generalization of the ham-sandwich theorem to algebraic hypersurfaces. I learned about this result from Gromov's paper [3]. However, I recently learned that it was proven by Stone and Tukey [13] in 1941. In this section, we explain and prove this generalization of the ham-sandwich theorem, following Stone and Tukey.

First we recall the original ham-sandwich theorem.

THEOREM 2.1. (Ham-sandwich theorem) *Let $U_1, ..., U_n$ be finite-volume open sets in $\mathbb{R}^n$. Then there is a hyperplane H that bisects each set U_i.*

The 3-dimensional case of the ham-sandwich theorem was first proven in the 1930s by Stefan Banach, using the Borsuk–Ulam theorem. Stone and Tukey extended the method to the n-dimensional case. (There is a nice historical discussion on Wikipedia.) Stone and Tukey noticed that the same method can be used to prove many other bisection results. For example, they proved the following result.

THEOREM 2.2. (Polynomial ham-sandwich theorem; Stone–Tukey [13]) *Let $N = \binom{n+d}{d} - 1$ and let $U_1, ..., U_N$ be finite-volume open sets in $\mathbb{R}^n$. Then there is a degree-d algebraic hypersurface Z which bisects each set U_i.*

We will prove the polynomial ham-sandwich theorem using the Borsuk–Ulam theorem, which we recall here.

THEOREM 2.3. (Borsuk–Ulam) *Let F be a continuous map from S^N to $\mathbb{R}^N$ obeying the antipodal condition*

$$F(-x) = -F(x) \quad \text{for every } x \in S^N.$$

Then the image of F contains 0.

For a proof of the Borsuk–Ulam theorem, the reader may look at Hatcher's book on algebraic topology [8, pp. 174–176]. Another reference is [12] by Matoušek. This book gives a proof of the theorem, and it also discusses interesting applications of the Borsuk–Ulam theorem, for example to Kneser's conjecture in combinatorics. We now turn to the proof of the polynomial ham-sandwich theorem.

Proof of Theorem 2.2. Let $V(d)$ denote the vector space of all real polynomials of degree at most d in n variables. The dimension of $V(d)$ is $\binom{n+d}{d}$. Let S^N denote the unit sphere in $V(d)$, where $N = \binom{n+d}{d} - 1$. For each set U_i, we define a function F_i from S^N to $\mathbb{R}$, by setting

$$F_i(P) = \text{Vol}(\{x \in U_i : P(x) > 0\}) - \text{Vol}(\{x \in U_i : P(x) < 0\}).$$

If we replace P by $-P$, then the two volumes trade places, so $F_i(-P) = -F_i(P)$. It is not hard to check that F_i is continuous (see below for the details). Combining all F_i into a vector-valued function, we get a continuous map $F: S^N \to \mathbb{R}^N$ obeying the antipodal condition. By the Borsuk–Ulam theorem, $F(P) = 0$ for some $P \in S^N \subset V(d)$. By the definition of F_i, we see that for each i,

$$\text{Vol}(\{x \in U_i : P(x) > 0\}) = \text{Vol}(\{x \in U_i : P(x) < 0\}).$$

Hence the hypersurface defined by $P(x) = 0$ bisects each set U_i. $\qquad\square$

For the sake of completeness, we include the proof that F_i is a continuous function.

LEMMA 2.4. (Continuity lemma) *If U is an open set of finite measure, then the measure of the set $\{x \in U : P(x) > 0\}$ depends continuously on $P \in V(d) \setminus \{0\}$.*

Proof. Suppose that P is a non-zero polynomial in $V(d)$ and $P_n \in V(d)$ with $P_n \to P$. Pick any $\varepsilon > 0$. We can find a subset $E \subset U$ so that $P_n \to P$ uniformly pointwise on $U \setminus E$, and $m(E) < \varepsilon$.

The set $\{x \in U : P(x) = 0\}$ has measure zero. Therefore, we can choose δ so that the set $\{x \in U : |P(x)| < \delta\}$ has measure less than ε.

Next we choose n large enough so that $|P_n(x) - P(x)| < \delta$ on $U \setminus E$. Then the measures of $\{x \in U : P_n(x) > 0\}$ and $\{x \in U : P(x) > 0\}$ differ by at most 2ε. But ε was arbitrary. $\qquad\square$

To make use of Theorem 2.2, we will use a standard volume estimate for hypersurfaces that bisect simple sets.

LEMMA 2.5. (Basic area estimate) *If a hypersurface S bisects a unit ball or a unit cube, then S has $(n-1)$-dimensional volume at least $c(n)$.*

3. Directed volume

The second tool in our paper is directed volume, which is a way of measuring the amount of volume of a hypersurface facing in different directions.

For a hypersurface $S \subset \mathbb{R}^n$, we define a directed volume function V_S by

$$V_S(v) := \int_S |v \cdot N| \, d\text{vol}_S. \tag{3.1}$$

In this formula, N denotes the normal vector to S, and $v \in \mathbb{R}^n$ is a fixed vector. Hence the directed volume is a non-negative function of $v \in \mathbb{R}^n$.

For a unit vector v, the directed volume $V_S(v)$ can be given a different, more geometric interpretation. Let $\pi_v \colon \mathbb{R}^n \to v^\perp$ be the orthogonal projection onto $v^\perp$. Then we can also think of $V_S(v)$ as the volume of $\pi_v(S)$, counted with geometric multiplicity. For each $y \in v^\perp$, we consider the intersection $S \cap \pi_v^{-1}(y)$. We let $|S \cap \pi_v^{-1}(y)|$ denote the number of points in $S \cap \pi_v^{-1}(y)$. For a compact smooth hypersurface S (possibly with boundary), this number of points is finite for almost every y. If v is a unit vector, then $V_S(v)$ is given by

$$V_S(v) = \int_{v^\perp} |S \cap \pi_v^{-1}(y)| \, dy. \tag{3.2}$$

Equations (3.1) and (3.2) will both be useful to us. Using equation (3.2), we can prove a key estimate about the directed volumes of algebraic hypersurfaces in cylinders.

Lemma 3.1. (Cylinder estimate) *If T is a cylinder of radius r, v is a unit vector parallel to the core of T and Z is an algebraic hypersurface of degree d, then the directed volume $V_{Z \cap T}(v)$ is bounded as follows*:

$$V_{Z \cap T}(v) \leqslant \omega_{n-1} r^{n-1} d.$$

Proof. The projection $\pi_v(T)$ is an $(n-1)$-dimensional disk of radius r. The function $y \mapsto |Z \cap T \cap \pi_v^{-1}(y)|$ is supported in this disk. But since Z is a degree-d algebraic hypersurface, Z intersects a line in at most d points, unless Z contains the entire line. Hence $|Z \cap \pi_v^{-1}(y)| \leqslant d$ for almost every y. $\square$

Equation (3.1) is also useful. For example, it allows us to see that a surface of volume 1 must have a fairly large directed volume in some direction.

Lemma 3.2. *Suppose that $v_1, \ldots, v_n$ are unit vectors. Let e_j denote the coordinate unit vectors, and suppose that $|e_j - v_j| < (100n)^{-1}$. Let S be any hypersurface in $\mathbb{R}^n$. Then*

$$\text{Vol}(S) \leqslant 2 \sum_{j=1}^n V_S(v_j).$$

Proof. For each x in S, let $N(x)$ denote the unit normal vector to S at x. Because of the angle condition on v_j, we know that $|v_j \cdot N(x)| \geqslant |e_j \cdot N(x)| - (100n)^{-1}$. Hence

$$\sum_{j=1}^n |v_j \cdot N(x)| \geqslant \left(\sum_{j=1}^n |e_j \cdot N(x)| \right) - \frac{1}{100} \geqslant \frac{99}{100}.$$

Integrating this inequality over S, we see that

$$\sum_{j=1}^{n} V_S(v_j) = \int_S \sum_{j=1}^{n} |v_j \cdot N(x)| \, d\mathrm{vol}_S(x) \geqslant \int_S \frac{99}{100} \, d\mathrm{vol}_S(x) = \frac{99}{100} \, \mathrm{Vol}(S). \qquad \square$$

Estimates for directional volumes appeared in [6] and [7], where quantitative estimates for certain homotopy invariants of a map are given, in terms of its Lipschitz constant.

4. The proof of Theorem 1.1

Proof of Theorem 1.1. Look at the standard unit lattice in $\mathbb{R}^n$. Let $Q_1, ..., Q_V$ be the set of n-cubes in the lattice which intersect I. Here V is the number of cubes that intersect I. It suffices to prove the estimate $V \lesssim A^{n/(n-1)}$.

According to the polynomial ham-sandwich theorem, we may find a degree-d algebraic hypersurface Z which bisects Q_k for every k, with degree $d \lesssim V^{1/n}$. Because of the bisection property, the volume of $Q_k \cap Z$ is $\gtrsim 1$ for each Q_k.

For each Q_k, we pick a tube in each direction that goes through Q_k. So we have labels $a_1(k), ..., a_n(k)$ such that $T_{j,a_j(k)}$ intersects Q_k. By assumption, the unit vector $v_{j,a_j(k)}$ parallel to the core of $T_{j,a_j(k)}$ is within $(100n)^{-1}$ of the coordinate vector e_j. Applying Lemma 3.2, we get the estimate

$$\sum_{j=1}^{n} V_{Z \cap Q_k}(v_{j,a_j(k)}) \gtrsim \mathrm{Vol}(Z \cap Q_k) \gtrsim 1.$$

So for each k, we can choose a tube $T_{j(k),a(k)}$ which meets Q_k and is such that

$$V_{Z \cap Q_k}(v_{j(k),a(k)}) \gtrsim 1.$$

We have just associated a tube with each cube. There are in total only nA tubes. By the pigeonhole principle, there is a tube associated with $\gtrsim V/A$ different cubes. Let $T_{j,a}$ be such a tube. Then we have $\gtrsim V/A$ different cubes Q_k which intersect $T_{j,a}$ and with $V_{Z \cap Q_k}(v_{j,a}) \gtrsim 1$.

Let $\widetilde{T}_{j,a}$ be the $\sqrt{n}$-neighborhood of $T_{j,a}$. The set $\widetilde{T}_{j,a}$ is itself a cylinder of radius $1 + \sqrt{n}$, with core parallel to $v_{j,a}$, and it contains all the cubes Q_k which overlap $T_{j,a}$.

Therefore, the directed volume $V_{Z \cap \widetilde{T}_{j,a}}(v_{j,a}) \gtrsim V/A$.

On the other hand, by Lemma 3.1, the same directed volume is $\lesssim V^{1/n}$.

Hence $V/A \lesssim V^{1/n}$. Rearranging, we get $V \lesssim A^{n/(n-1)}$. $\qquad \square$

5. The Lyusternik–Shnirel'man vanishing lemma

In order to prove Theorem 1.3, we use some more sophisticated algebraic topology: the Lyusternik–Shnirel'man vanishing lemma. In this section, we will introduce it and try to explain what it is good for. The basic message is that the vanishing lemma is similar to the ham-sandwich theorem, but it is more flexible.

The vanishing lemma is about cup-products of cohomology classes.

LEMMA 5.1. (Vanishing lemma) *Let X be a CW complex (for example a manifold). Let a_1 and a_2 be cohomology classes in $H^*(X, R)$, where R may be any ring of coefficients, such as $\mathbb{R}$, $\mathbb{Z}$ or $\mathbb{Z}_2$. Suppose that a_1 vanishes on some open set $S_1 \subset X$ and that a_2 vanishes on some open set $S_2 \subset X$. Then the cup product $a_1 \cup a_2$ vanishes on the union $S_1 \cup S_2$.*

The vanishing lemma is one of the fundamental topological facts about cup products. I believe that it was first proven by Lyusternik and Shnirel'man in the 1930s, as part of their project for proving the existence of closed geodesics. The proofs I have seen in the literature are a little more abstract than I would like, so I wrote an appendix giving the proof.

Here is the basic intuition behind the vanishing lemma. Suppose that f_1 and f_2 are functions on X. If f_1 vanishes on S_1 and f_2 vanishes on S_2, then clearly the product $f_1 f_2$ vanishes on the union $S_1 \cup S_2$. The vanishing lemma holds because cohomology classes are not so different from functions. A cohomology class can be represented by either a differential form or a singular cocycle, and these objects have enough in common with functions to make the vanishing lemma hold. For details, see the appendix.

To apply the vanishing lemma, we need to know something about the cup products of cohomology classes. For this paper, the key example is the cohomology ring of real projective space.

THEOREM 5.2. (Cohomology ring of $\mathbb{RP}^N$) *The cohomology group $H^i(\mathbb{RP}^N, \mathbb{Z}_2)$ is isomorphic to $\mathbb{Z}_2$ for $0 \leqslant i \leqslant N$, and is trivial otherwise. Let a denote the non-zero element in $H^1(\mathbb{RP}^N, \mathbb{Z}_2)$. Then, for $1 \leqslant i \leqslant N$, a^i is the non-zero element of $H^i(\mathbb{RP}^N, \mathbb{Z}_2)$.*

This theorem may be found in Hatcher's topology book [8, p. 212].

Using the vanishing lemma, we can give a different proof of the polynomial ham-sandwich theorem.

As before, we let $V(d)$ denote the vector space of all real polynomials of degree at most d in n variables. The dimension of $V(d)$ is $\binom{n+d}{d}$. For each non-zero polynomial P in $V(d)$, there is an associated variety, the zero-set of P. If we replace P by some multiple λP, the zero-set remains unchanged, and so the real algebraic hypersurfaces of degree at

most d are parametrized by the projectivization of $V(d)$, which is a real projective space $\mathbb{RP}^N$, where $N = \binom{n+d}{d} - 1$.

We are interested in hypersurfaces that bisect open sets. Given a finite-volume open set $U \subset \mathbb{R}^n$, we let $\mathrm{Bi}(U) \subset \mathbb{RP}^N$ consist of the algebraic hypersurfaces that bisect the set U. If Z is a real algebraic hypersurface given by the equation $P = 0$, then we say that Z bisects U if

$$\mathrm{Vol}(\{x \in U : P(x) > 0\}) = \mathrm{Vol}(\{x \in U : P(x) < 0\}).$$

By the continuity lemma (Lemma 2.4), these volumes change continuously with P, and so $\mathrm{Bi}(U)$ is a closed subset of $\mathbb{RP}^N$. The key topological result about $\mathrm{Bi}(U)$ is the following lemma.

LEMMA 5.3. (Bisection lemma) *Let a denote the non-trivial cohomology class in $H^1(\mathbb{RP}^N, \mathbb{Z}_2)$. Let U be a finite-volume subset of $\mathbb{R}^n$. Then the cohomology class a vanishes on the complement $\mathbb{RP}^N \setminus \mathrm{Bi}(U)$.*

Proof. Suppose that a does not vanish on $\mathbb{RP}^N \setminus \mathrm{Bi}(U)$. Then the class a is detected by a 1-cycle c in $\mathbb{RP}^N \setminus \mathrm{Bi}(U)$. Without loss of generality, we may assume that c has only one component, and so c is topologically a circle. Pick a point in c and look at the corresponding hypersurface Z. We can assume that Z does not bisect U, so the complement of Z has a big half and a little half. Now we pick a polynomial P_Z representing Z, and we choose it so that the big half is where the polynomial P_Z is positive. We can lift our 1-parameter family of hypersurfaces to a 1-parameter family of polynomials that goes from P_Z to $-P_Z$. The part of U where P_Z is positive has more than half measure. The part of U where $-P_Z$ is positive has less than half measure. According to Lemma 2.4, the measure changes continuously as the polynomial changes. By continuity, there is a polynomial in the family that bisects U. $\square$

Combining the bisection lemma and the vanishing lemma, we can say something about hypersurfaces that bisect multiple sets. Suppose that $U_1, ..., U_r \subset \mathbb{R}^n$ are finite-volume open sets, where r is any positive integer. Let $\mathrm{Bi}(U_1, ..., U_r) \subset \mathbb{RP}^N$ denote the set of algebraic hypersurfaces which bisect all the open sets $U_1, ..., U_r$. The set $\mathrm{Bi}(U_1, ..., U_r)$ is just the intersection of $\mathrm{Bi}(U_i)$, $1 \leqslant i \leqslant r$. In particular, $\mathrm{Bi}(U_1, ..., U_r) \subset \mathbb{RP}^N$ is a closed set.

LEMMA 5.4. (Multiple bisection lemma) *Let $\mathrm{Bi}(U_1, ..., U_r)$ be as above. Then the cohomology class a^r vanishes on $\mathbb{RP}^N \setminus \mathrm{Bi}(U_1, ..., U_r)$.*

Proof. By the bisection lemma, the cohomology class a vanishes on $\mathbb{RP}^N \setminus \mathrm{Bi}(U_i)$ for each i. Each of these sets is open. Therefore, the vanishing lemma tells us that a^r vanishes on their union. But the union $\bigcup_{i=1}^r (\mathbb{RP}^N \setminus \mathrm{Bi}(U_i))$ is exactly $\mathbb{RP}^N \setminus \mathrm{Bi}(U_1, ..., U_r)$. $\square$

Combining the multiple bisection lemma and the cohomology ring of $\mathbb{RP}^N$, we can reprove the polynomial ham-sandwich theorem. This proof was given by Gromov in [3].

Proof of Theorem 2.2. Recall that a is the non-zero element in $H^1(\mathbb{RP}^N, \mathbb{Z}_2)$. By the multiple bisection lemma, a^N vanishes on $\mathbb{RP}^N \setminus \mathrm{Bi}(U_1, ..., U_N)$. But in the cohomology ring of $\mathbb{RP}^N$, a^N does not vanish on $\mathbb{RP}^N$. Hence $\mathrm{Bi}(U_1, ..., U_N)$ must be non-empty. In other words, there is a degree-d hypersurface Z that bisects each open set U_i. $\square$

The vanishing lemma has other applications besides the ham-sandwich theorem. One classical application is to give covering estimates.

PROPOSITION 5.5. (Covering estimate; Lyusternik–Shnirel'man) *Suppose that* $\mathbb{RP}^N$ *is covered by some contractible open sets* $S_1, ..., S_r$. *Then* $r \geqslant N+1$.

Proof. Since each S_i is contractible, the cohomology class a vanishes on each S_i. Applying the vanishing lemma once, we see that a^2 vanishes on $S_1 \cup S_2$. Proceeding inductively, we see that a^r vanishes on the union of all S_r, which is $\mathbb{RP}^N$. But in the cohomology ring of $\mathbb{RP}^N$, a^i is non-zero for all $i \leqslant N$. Hence $r \geqslant N+1$. $\square$

Our proof of the multilinear Kakeya estimate combines the polynomial ham-sandwich theorem with some covering estimates similar to the one above. The Lyusternik–Shnirel'man vanishing lemma allows us to combine these two techniques, making it a little more flexible than the Borsuk–Ulam theorem.

6. The visibility lemma

In the proof of Theorem 1.1, we found an algebraic hypersurface whose intersection with many unit cubes has volume $\gtrsim 1$. The total volume is not as important as the directional volumes $V_{Z \cap Q}(v)$ in various directions. (Here Z is the hypersurface, Q is a cube and v is a direction.) In this section, we build a hypersurface which has large directional volumes in many directions.

We now define the visibility of a hypersurface, which measures whether the surface has a large directional volume in many directions. Roughly, a hypersurface has large visibility if either $V_S(v)$ is large for every unit vector v, or else $V_S(v)$ is extremely large for some vectors v.

We define the *visibility* of a surface S to be

$$\mathrm{Vis}[S] := \frac{1}{\mathrm{Vol}(\{v : |v| \leqslant 1 \text{ and } V_S(v) \leqslant 1\})}.$$

This definition is a little long, so we make some comments about it. The reader may wonder, why not just look at the average directional volume in all unit directions v:

$\int_{S^{n-1}} V_S(v) \, d\mathrm{vol}(v) / \mathrm{Vol}(S^{n-1})$? For our arguments, it is crucial to know whether $V_S(v)$ is small in some directions even if the set of such directions has a small measure. The average directional volume above will not detect small values of $V_S(v)$, but the definition of visibility is quite sensitive to small values of $V_S(v)$.

We now compute the visibility in two examples. First, suppose that S is a unit $(n-1)$-disk in the hyperplane $x_n = 0$. Then the function $V_S(v) = \omega_{n-1}|v_n|$, where v_n is the nth component of v. Therefore, the set $\{v : V_s(v) \leqslant 1\}$ is an infinite slab of the form $|v_n| \leqslant C$. The set of v with $V_S(v) \leqslant 1$ and $|v| \leqslant 1$ is roughly the unit ball, and so the visibility of S is on the order of 1. We had to include the condition $|v| \leqslant 1$ in the definition, or else the visibility of a disk would be zero. Including $|v| \leqslant 1$ in the definition has the unpleasant effect that the visibility of the empty set is also around 1. In practice, we will speak of the visibility of S for surfaces S contained in a unit cube and with volume at least 1, and in this range, the visibility behaves reasonably.

A second important example is a union of unit disks with N_j disks perpendicular to the x_j-axis. In this case, $V_S(v)$ is roughly $\sum_{j=1}^{n} N_j |v_j|$. Hence the region where $V_S(v) \leqslant 1$ is roughly $\{v \in \mathbb{R}^n : |v_j| \leqslant N_j^{-1}\}$. The volume of this region is roughly $N_1^{-1} \dots N_n^{-1}$, and so the visibility of this surface is roughly $N_1 \dots N_n$. This is the best example to keep in mind to understand what visibility means.

Our next goal is to find algebraic hypersurfaces which have large visibility in many cubes. We recall from §5 that the space of degree-d hypersurfaces in $\mathbb{R}^n$ is parametrized by $\mathbb{RP}^N$ for $N = \binom{n+d}{d} - 1$. We will slightly abuse notation by identifying a degree-d hypersurface Z and the corresponding point in $\mathbb{RP}^N$—we will speak of $Z \in \mathbb{RP}^N$.

If we fix some cube $Q \subset \mathbb{R}^n$, we want to study $\mathrm{Vis}[Z \cap Q]$ as a function of $Z \in \mathbb{RP}^N$. Unfortunately, this function is not continuous in Z. Even the $(n-1)$-dimensional volume of $Z \cap Q$ is not continuous in Z. Because we make topological arguments using the Lyusternik–Shnirel'man vanishing lemma, this discontinuity leads to some technical problems. To deal with these, we define mollified continuous versions of the directed volume and the visibility. We mollify these functions by averaging over small balls in $\mathbb{RP}^N$. We use the standard metric on $\mathbb{RP}^N$, and let $B(Z, \varepsilon)$ denote the ball around $Z \in \mathbb{RP}^N$ of radius ε.

For any open set U, we define a *mollified* version of $V_{Z \cap U}(v)$ as follows:

$$\overline{V}_{Z \cap U}(v) := \frac{1}{|B(Z, \varepsilon)|} \int_{B(Z, \varepsilon)} V_{Z' \cap U}(v) \, dZ'.$$

We define a mollified visibility function using the mollified directional volumes:

$$\overline{\mathrm{Vis}}[Z \cap U] := \frac{1}{\mathrm{Vol}(\{v : |v| \leqslant 1 \text{ and } \overline{\overline{V}}_{Z \cap U}(v) \leqslant 1\})}.$$

We will choose ε extremely small compared to all other constants in the paper. In practice, the mollified directional volumes and visibilities maintain all the useful properties of the unmollified versions, and they are also continuous. Therefore, on an early reading of the paper, it makes sense to ignore the mollification and just pretend that the visibility is continuous in Z.

In the following lemma, we collect the properties of the mollified volumes and visibilities which we will use.

LEMMA 6.1. *Let U be a bounded open set in $\mathbb{R}^n$. The mollified directed volume $\overline{V}_{Z \cap U}(v)$ and the mollified visibility $\overline{\mathrm{Vis}}[Z \cap U]$ obey the following properties*:

(i) (Scaling) *For any constant λ, $\overline{V}_{Z \cap U}(\lambda v) = |\lambda| \overline{V}_{Z \cap U}(v)$*;

(ii) (Convexity) *The function $\overline{V}_{Z \cap U}(v)$ is convex in v*;

(iii) (Disjoint unions) *If U is a disjoint union of U_1 and U_2, then*

$$\overline{V}_{Z \cap U}(v) = \overline{V}_{Z \cap U_1}(v) + \overline{V}_{Z \cap U_2}(v);$$

(iv) (Cylinder estimate) *If T is a cylinder of radius r with core vector v, and if Z is a degree-d hypersurface, then $\overline{V}_{Z \cap T}(v) \leqslant \omega_{n-1} r^{n-1} d$*;

(v) (Bisection) *If Z bisects a unit ball B, and if ε is small enough, then $\overline{V}_{Z \cap B}(v) \gtrsim 1$ for some unit vector v*;

(vi) (Continuity) *The functions $\overline{V}_{Z \cap U}(v)$ and $\overline{\mathrm{Vis}}[Z \cap U]$ depend continuously on $Z \in \mathbb{RP}^N$.*

Proof. (i) This follows by plugging in the formulas.

(ii) For each vector N, the function $|N \cdot v|$ is a convex function of v. Since a positive combination of convex functions is convex, $V_{Z \cap U}(v)$ is convex in v. As an average of convex functions is convex, $\overline{V}_{Z \cap U}$ is also convex.

(iii) This also follows by plugging in the formulas.

(iv) Lemma 3.1 tells us that $V_{Z' \cap T}(v) \leqslant \omega_{n-1} r^{n-1} d$ for every degree-d hypersurface Z'. Taking an appropriate average, we see that $\overline{V}_{Z \cap T}(v) \leqslant \omega_{n-1} r^{n-1} d$.

(v) Suppose that Z bisects B. By Lemma 2.4, we can choose ε small enough so that each Z' in $B(Z, \varepsilon)$ nearly bisects B. Hence the volume of $Z' \cap B$ is $\gtrsim 1$. We let $e_1, ..., e_n$ be the standard orthonormal basis of $\mathbb{R}^n$. By Lemma 3.2,

$$\sum_{j=1}^{n} V_{Z' \cap B}(e_i) \geqslant \frac{\mathrm{Vol}(Z' \cap B)}{2} \gtrsim 1.$$

Taking an average over Z' in $B(Z, \varepsilon)$, we see that $\sum_{j=1}^{n} \overline{V}_{Z \cap B}(e_i) \gtrsim 1$.

(vi) The function $V_{Z \cap U}(v)$ is a bounded measurable function on $\mathbb{RP}^N$. Hence its averages over ε-balls form a continuous function. So $\overline{V}_{Z \cap U}(v)$ depends continuously on Z.

Next we address continuity in v. The function $V_{Z\cap U}(v)$ is Lipschitz in v with a Lipschitz constant $C(d, U, n)$ independent of Z. To see this, we expand

$$|V_{Z\cap U}(v_1) - V_{Z\cap U}(v_2)|$$

as an integral:

$$\left| \int_{Z\cap U} |N\cdot(v_1-v_2)|\, d\mathrm{vol} \right| \leqslant |v_1-v_2|\, \mathrm{Vol}(Z\cap U).$$

Now U is a bounded domain, so it fits in a ball of some radius $R(U)$, and standard algebraic geometry shows that $\mathrm{Vol}(Z\cap U) \leqslant C_n R^{n-1} d$. (For more details on this Crofton estimate, see [5, p. 58].)

Hence $\overline{V}_{Z\cap U}(v)$ is also Lipschitz in v with a Lipschitz constant $C(d, U, n)$. Therefore, $\overline{V}_{Z\cap U}(v)$ is jointly continuous as a function of $(Z, v) \in \mathbb{RP}^N \times \mathbb{R}^n$. Hence $\overline{\mathrm{Vis}}[Z\cap U]$ is continuous in Z. $\qquad\square$

The next lemma allows us to find algebraic hypersurfaces with large visibility. It is analogous to the bisection lemma, but instead of producing surfaces that bisect a ball, it produces surfaces with large visibility in a ball.

LEMMA 6.2. (Visibility lemma) *There is an integer constant $C_n > 1$ such that the following holds. Fix any degree d and any unit ball $B(p, 1) \subset \mathbb{R}^n$. Consider the space of degree-d algebraic hypersurfaces in $\mathbb{R}^n$, parametrized by $\mathbb{RP}^N$. Let L_M denote the subset of algebraic surfaces Z with $\overline{\mathrm{Vis}}[Z\cap B(p, 1)] \leqslant M$, where $M \geqslant 1$ is an integer. Let a denote the non-zero cohomology class in $H^1(\mathbb{RP}^N, \mathbb{Z}_2)$. Then the cohomology class $a^{C_n M}$ vanishes on a neighborhood of L_M.*

Proof. Let E be an ellipsoid contained in the unit ball in $\mathbb{R}^n$, with the volume of E at least $1/M$. Let $L(E)$ denote the set of degree-d hypersurfaces Z such that

$$\overline{V}_{Z\cap B(p,1)}(v) \leqslant 1 \quad \text{for all } v \in E.$$

Notice that if $\overline{\mathrm{Vis}}[Z\cap B(p, 1)] \leqslant M$, then Z is in $L(E)$ for some ellipsoid E of volume $\gtrsim 1/M$. We will first deal with the different ellipsoids individually and then see how to deal with all of them simultaneously.

LEMMA 6.3. (Weak visibility lemma) *If E is an ellipsoid contained in the unit ball with volume at least $1/M$, then the cohomology class $a^{C_n M}$ vanishes on a neighborhood of $L(E)$.*

Proof. Let $A(n)$ be a large number we will choose later.

We let E' be a rescaling of E by a factor $1/A(n)$ (so that E' is smaller than E). We let $U_1, ..., U_k$ denote disjoint parallel copies of E' contained in $B(p, 1)$. We take a

maximal family of parallel copies of E' in $B(p, 1)$—meaning that there is no room to add an additional parallel copy of E'. From the maximality, we see that $\mathrm{Vol}(E')k \sim 1$, where k is the number of parallel copies. Since the volume of E' is at least $1/A(n)^n M$, we also know that $k \lesssim A(n)^n M$.

Now suppose that a^k does not vanish on $L(E)$. Using Lemma 5.4, we can pick a cycle Z in $L(E)$ such that Z bisects each set U_i. Next we investigate the directional volumes of a surface bisecting a copy of E'.

Suppose that Z bisects E'. Let $E'_1, ..., E'_n$ be the lengths of the principal axes of E'. Let $e_1, ..., e_n$ be unit length vectors with e_j lying on the jth principal axis of E'. (To check the notation, each point $\pm E'_j e_j$ lies on the boundary of E'.) The vectors $e_1, ..., e_n$ form an orthonormal basis of $\mathbb{R}^n$.

LEMMA 6.4. *Under the hypotheses in the last paragraph, the following estimate holds for some $1 \leqslant j \leqslant n$:*

$$\overline{V}_{Z \cap E'}(e_j) \gtrsim \frac{\mathrm{Vol}(E')}{E'_j}.$$

Proof. Let L be a linear map taking E' diffeomorphically to the unit ball. The map L is diagonal with respect to the basis $\{e_j\}$: in this basis, it scales the jth coordinate by $1/E'_j$. Then $L(Z)$ bisects the unit ball. According to the bisection clause in Lemma 6.1, $\overline{V}_{L(Z) \cap B}(e_j) \gtrsim 1$ for some j. When we change coordinates back and interpret this inequality in E', it gives the lemma. We now explain the coordinate change in detail. We let π_j denote the orthogonal projection from $\mathbb{R}^n$ to $e_j^\perp$. Next we use equation (3.2) to write directional volumes in terms of π_j:

$$V_{L(Z') \cap B}(e_j) = \int_{e_j^\perp} |L(Z') \cap B \cap \pi_j^{-1}(y)| \, dy,$$

$$V_{Z' \cap E}(e_j) = \int_{e_j^\perp} |Z' \cap E \cap \pi_j^{-1}(y)| \, dy.$$

Comparing the right-hand sides, we get

$$V_{Z' \cap E}(e_j) = \left(\prod_{i=1}^{n} E'_i \right) \frac{1}{E'_j} V_{L(Z') \cap B}(e_j).$$

Averaging over Z', we get the inequality for the mollified directional volumes

$$\overline{V}_{Z \cap E}(e_j) = \left(\prod_{i=1}^{n} E'_i \right) \frac{1}{E'_j} \overline{V}_{L(Z) \cap B}(e_j) \gtrsim \frac{\mathrm{Vol}(E')}{E'_j}. \qquad \square$$

Since U_i is a translation of E', we get the estimate

for each i there is some coordinate j such that $\overline{V}_{Z \cap U_i}(e_j) \gtrsim \dfrac{\mathrm{Vol}(E')}{[E'_j]}$.

The number of translated ellipsoids U_i is k, where $\mathrm{Vol}(E')k \sim 1$. Combining our last estimate over all these ellipsoids, we see that for a popular coordinate j,

$$\overline{V}_{Z \cap B(p,1)}(e_j) \gtrsim \frac{1}{E'_j} = \frac{A(n)}{E_j}.$$

Now we choose $A(n)$ sufficiently large compared to our dimensional constants, and we conclude that $\overline{V}_{Z \cap B(p,1)}(e_j) > 1/E_j$, and so $\overline{V}_{Z \cap B(p,1)}(E_j e_j) > 1$. But the vector $E_j e_j$ is contained in E. By the definition of $L(E)$, we should have $\overline{V}_{Z \cap B(p,1)}(v) \leqslant 1$ for every $v \in E$. This contradiction shows that our assumption was wrong, and a^k vanishes on $L(E)$. But $k \lesssim A(n)^n M$, and so $a^{C(n)M}$ vanishes on $L(E)$ for an appropriate dimensional constant $C(n)$.

Reinspecting the argument we see that $a^{C_n M}$ vanishes on the union of $\mathbb{RP}^n \setminus \mathrm{Bi}(U_i)$. This latter set is open and we have shown that it contains $L(E)$, and so $a^{C_n M}$ vanishes on a neighborhood of $L(E)$. $\qquad\square$

Next we explain how to upgrade this weak visibility lemma to get the visibility lemma we originally stated. For each sufficiently large ellipsoid E, we have seen that $a^{C_n M}$ vanishes on $L(E)$. Remarkably, $a^{C_n M}$ vanishes on the union $\bigcup_E L(E)$ as E varies over all ellipsoids with volume at least $1/M$. We can use the vanishing lemma to show that $a^{pC_n M}$ vanishes on the union of any p sets $L(E_k)$, but we do not have any good control of the size of p. The situation is analogous to the following proposition, which is used in Gromov's paper [4].

PROPOSITION 6.5. (Gromov) *Let X be a manifold and let $f \colon X \to \mathbb{R}^m$ be a map. Suppose that for each unit ball $B(y, 1)$ in $\mathbb{R}^m$, the cohomology class $\alpha \in H^*(X)$ vanishes on $f^{-1}(B(y, 1))$. Then α^{m+1} vanishes on all of X.*

Proof. Triangulate $\mathbb{R}^m$ so that each simplex has diameter at most $\frac{1}{4}$. Let U_i be an open cover on $\mathbb{R}^m$, indexed by the simplices of the triangulation (including simplices of all dimensions). It is possible to choose U_i in such a way that U_i intersects U_j only if one of the corresponding simplices contains the other one. In particular, the open sets corresponding to two simplices of the same dimension never intersect. Also, each U_i is contained in a $\frac{1}{10}$-neighborhood of the corresponding simplex. Since each U_i has diameter at most $\frac{1}{2}$, each U_i is contained in some unit ball, and so α vanishes on $f^{-1}(U_i)$. Now, for $0 \leqslant l \leqslant m$, let V_l denote the union of U_i as i varies among all the l-dimensional simplices of

our triangulation. For any two l-dimensional simplices, i_1 and i_2, the corresponding sets U_{i_1} and U_{i_2} are disjoint, and hence their preimages $f^{-1}(U_{i_1})$ and $f^{-1}(U_{i_2})$ are disjoint open subsets of X. Therefore, α vanishes on $f^{-1}(V_l)$. Finally, by the Lyusternik–Shnirel'man vanishing lemma, α^{m+1} vanishes on all of X. □

(The clever covering for $\mathbb{R}^m$ that appears here originated in dimension theory. See the book [9] for more information.)

Our argument is a variation on the proof of this proposition. The role of the space $\mathbb{R}^m$ is played by the space of all ellipsoids in $\mathbb{R}^n$.

Let Ell denote the set of all closed ellipsoids in $\mathbb{R}^n$ centered at the origin. We put a distance function on Ell by saying that $\mathrm{dist}_{\mathrm{Ell}}(E_1, E_2) \leqslant \log D$ if and only if

$$\frac{1}{D} E_1 \subset E_2 \subset D E_1.$$

We let Ell[M] denote the set of all ellipsoids contained in the unit ball with volume at least $1/M$. Then we choose a maximal 1-separated subset of Ell[M], given by finitely many ellipsoids $\mathrm{Ell}_1, ..., \mathrm{Ell}_s$. The number of ellipsoids is finite, but it grows exponentially with M.

Recall that L_M is the set of hypersurfaces $Z \in \mathbb{RP}^N$ so that $\overline{\mathrm{Vis}}[Z \cap B(p,1)] \leqslant M$. Next we divide the set L_M into classes. For any hypersurface Z, we let $K[Z]$ be the convex set $\{v : |v| \leqslant 1 \text{ and } \overline{V}_{Z \cap B(p,1)}(v) \leqslant 1\}$. We say that a hypersurface Z lies in A_k if and only if $K[Z]$ resembles Ell_k in the sense that $(10n)^{-1/2} \mathrm{Ell}_k \subset K[Z] \subset (10n)^{1/2} \mathrm{Ell}_k$. Because our mollified function $\overline{V}_{Z \cap B(p,1)}$ is continuous, the sets $A_k \subset \mathbb{RP}^N$ are closed.

According to a lemma of Fritz John, any symmetric convex set K can be approximated by an ellipsoid E in the sense that $n^{-1/2} E \subset K \subset n^{1/2} E$. From this estimate, it follows that the sets A_k cover L_M.

On the other hand, $A_k \subset L[(10n)^{-1/2} \mathrm{Ell}_k]$. By the weak visibility lemma, we see that $a^{C_n M}$ vanishes on a neighborhood of each A_k.

Two sets A_k and A_l overlap only if the corresponding ellipsoids Ell_k and Ell_l lie within a distance $C(n)$ of each other, using our metric on Ell. We want to bound the multiplicity of the cover of L_M by the sets A_k. It suffices to bound the number of ellipsoids Ell_k inside a ball of radius $C(n)$ in the space Ell. Let Ell_0 denote the unit ball. The closed ball of radius $C(n)$ around Ell_0 is a compact subset of Ell. (The space Ell is a finite-dimensional manifold, and our metric defines the usual topology on the manifold.) By compactness, any set of 1-separated ellipsoids Ell_i inside this ball has cardinality bounded by some $C'(n)$. But there is nothing special about the unit ball Ell_0. In fact, the space Ell is extremely symmetrical. The group $\mathrm{GL}(n, \mathbb{R})$ acts on Ell in the following way. Given a linear map $M \in \mathrm{GL}(n, \mathbb{R})$ and an ellipsoid $E \in \mathrm{Ell}$, we define $M(E)$ to be the

THE ENDPOINT CASE OF THE BENNETT–CARBERY–TAO CONJECTURE

image of E under the map M. This group action is an isometry using our metric on Ell. It is also transitive because of the spectral theorem. Therefore, any ball of radius $C(n)$ contains at most $C'(n)$ 1-separated points. Hence the multiplicity of the cover $\{A_k\}_k$ is bounded by $C'(n)$.

Let B_k be tiny open neighborhoods of A_k such that $a^{C_n M}$ vanishes on B_k. Since the sets A_k are closed, we can arrange that B_k and B_l intersect only if A_k and A_l intersect. The B_k form an open cover of a neighborhood of L_M with multiplicity at most $C'(n)$. We color the sets B_k using $C'(n)$ colors so that overlapping sets have distinct colors. For each color α from 1 to $C'(n)$, we let C_α denote the union of all sets B_k with the color α. Because these sets are disjoint, $a^{C_n M}$ vanishes on C_α for each α. Now by the Lyusternik–Shnirel'man vanishing lemma, $a^{C_n'' M}$ vanishes on the union of C_α, which includes a neighborhood of L_M. $\qquad\square$

Combining the visibility lemma and the Lyusternik–Shnirel'man vanishing lemma, we can find a degree-d algebraic hypersurface with large visibility on various cubes. The following lemma is the main result of this section.

LEMMA 6.6. (Large visibility on many cubes) *Consider the standard unit lattice in* $\mathbb{R}^n$. *Let M be a function from the set of n-cubes in the unit lattice to the non-negative integers. Then we can find an algebraic hypersurface of degree d such that*

$$\overline{\mathrm{Vis}}[Z \cap Q_k] \geqslant M(Q_k)$$

for every cube Q_k, where the degree d is bounded by $C(n)\left(\sum_k M(Q_k)\right)^{1/n}$.

Proof. The space of degree-d hypersurfaces is parametrized by $\mathbb{RP}^N$, where

$$N = \binom{n+d}{d} - 1 \geqslant c(n)d^n.$$

Let a denote the fundamental cohomology class of $\mathbb{RP}^N$. Let $S[Q_k]$ denote the set of surfaces Z where $\overline{\mathrm{Vis}}[Z \cap Q_k] < M(Q_k)$. According to the visibility lemma, the cohomology class $a^{C(n)M(Q_k)}$ vanishes on a neighborhood of $S[Q_k]$. By the Lyusternik–Shnirel'man vanishing lemma, the cohomology class $a^{C(n)\sum_k M(Q_k)}$ vanishes on a neighborhood of $\bigcup_k S[Q_k]$. But a^N does not vanish on $\mathbb{RP}^N$. As long as

$$C(n) \sum_k M(Q_k) < c(n)d^n \leqslant N,$$

there is a variety Z which does not lie in any $S(Q_k)$. Unwinding the definition, we see that $\overline{\mathrm{Vis}}[Z \cap Q_k] \geqslant M(Q_k)$ for every k. Our condition on d is $C(n)\sum_k M(Q_k) < c(n)d^n$, which holds for any $d > C'(n)\left(\sum_k M(Q_k)\right)^{1/n}$. $\qquad\square$

7. Multilinear Kakeya estimates

We now give a proof of the multilinear Kakeya estimate.

Proof of Theorem 1.3. We consider the standard unit cube lattice in $\mathbb{R}^n$. For each cube Q_k in this lattice, we define the following functions, measuring how many tubes of different types go through Q_k. We let $M_j(Q_k)$ denote the number of tubes $T_{j,a}$ which go through Q_k. Then we let $F(Q_k)$ be the product of these:

$$F(Q_k) := \prod_{j=1}^{n} M_j(Q_k).$$

It suffices to prove the following estimate for $F(Q_k)$:

$$\sum_k F(Q_k)^{1/(n-1)} < C(n) \theta^{-1/(n-1)} \prod_{j=1}^{n} A(j)^{1/(n-1)}. \tag{7.1}$$

Since we have only finitely many tubes, the function $F(Q_k)$ vanishes outside of finitely many cubes. We fix a large cube of side length S containing all of the relevant cubes Q_k. Next we apply Lemma 6.6, which guarantees that we can find a hypersurface Z_0 of degree $d \lesssim S$ obeying the following visibility estimates: for every cube Q_k,

$$\overline{\mathrm{Vis}}[Z_0 \cap Q_k] \geqslant S^n F(Q_k)^{1/(n-1)} \left(\sum_k F(Q_k)^{1/(n-1)} \right)^{-1}. \tag{7.2}$$

Adding $C(n)S$ hyperplanes to Z_0, we produce a variety Z of degree $d \lesssim S$ that still obeys the visibility estimate above, and also $\overline{V}_{Z \cap Q_k}(v) \geqslant |v|$ for each cube Q_k where $F(Q_k) > 0$. Equation (7.2) gives a strong lower bound for $V_{Z \cap Q_k}(v)$ in some directions, and this last estimate gives a weak lower bound in all directions.

Next we apply the cylinder estimate to control the directed volumes of Z in cubes along a given tube $T_{j,a}$. (The estimate we need is the cylinder clause of Lemma 6.1.) For each tube $T_{j,a}$, we have the estimate

$$\sum_{Q_k \text{ that intersect } T_{j,a}} \overline{V}_{Z \cap Q_k}(v_{j,a}) \lesssim S.$$

We would like to sum this inequality over all a from 1 to $A(j)$, but the vectors $v_{j,a}$ are changing.

For each cube Q_k and each j, we pick a vector $v_{j,k}$ from among $v_{j,a}$ such that

$$\overline{V}_{Z \cap Q_k}(v_{j,k}) = \min\{\overline{V}_{Z \cap Q_k}(v_{j,a}) : a = 1, ..., A(j)\}.$$

Substituting $v_{j,k}$ for $v_{j,a}$ in the last inequality and summing over a yields

$$\sum_{Q_k} M_j(Q_k) \overline{V}_{Z \cap Q_k}(v_{j,k}) \lesssim SA(j). \tag{7.3}$$

Next we need a lemma relating $\overline{\mathrm{Vis}}$ and $\overline{V}$.

LEMMA 7.1. *For each cube Q_k, the following inequality holds*:

$$\overline{\text{Vis}}[Z \cap Q_k] < \frac{C(n)}{\theta} \prod_{j=1}^{n} \overline{V}_{Z \cap Q_k}(v_{j,k}).$$

Proof. Let

$$v'_{j,k} = \frac{v_{j,k}}{\overline{V}_{Z \cap Q_k}(v_{j,k})}.$$

Because we added the hyperplanes to Z_0, we know that $\overline{V}_{Z \cap Q_k}(v) \geqslant 1$ for all unit vectors v. Hence $|v'_{j,k}| \leqslant 1$. We know that $\overline{V}_{Z \cap Q_k}(\pm v'_{j,k}) = 1$ for each j. Since the directed volume is a convex function of v, $\overline{V}_{Z \cap Q_k}(v) \leqslant 1$ for every v in the convex hull of the $2n$ points $\pm v'_{j,k}$. This convex hull is contained in the unit ball, and its volume is

$$c(n) \det(v'_{1,k}, ..., v'_{n,k}) = c(n) \left(\prod_{j=1}^{n} \overline{V}_{Z \cap Q_k}(v_{j,k}) \right)^{-1} \det(v_{1,k}, ..., v_{n,k}).$$

Because of our transverality assumption, the volume is

$$\geqslant c(n) \theta \left(\prod_{j=1}^{n} \overline{V}_{Z \cap Q_k}(v_{j,k}) \right)^{-1}.$$

Hence the set of all v with $|v| \leqslant 1$ and $\overline{V}_{Z \cap Q_k}(v) \leqslant 1$ has volume at least

$$\gtrsim \left(\prod_{j=1}^{n} \overline{V}_{Z \cap Q_k}(v_{j,k}) \right)^{-1} \theta.$$

The visibility $\overline{\text{Vis}}[Z \cap Q_k]$ is the inverse of this volume, which is at most

$$\frac{C(n)}{\theta} \prod_{j=1}^{n} \overline{V}_{Z \cap Q_k}(v_{j,k}). \qquad \square$$

Now we follow a string of inequalities powered by the visibility estimate in equation (7.2) and the cylinder estimate in equation (7.3):

$$S \left(\sum_{k} F(Q_k)^{1/(n-1)} \right)^{(n-1)/n} = \sum_{k} \left(S F(Q_k)^{1/(n-1)} \Big/ \left(\sum_{l} F(Q_l)^{1/(n-1)} \right)^{1/n} \right).$$

Using the visibility estimate in equation (7.2), we get

$$S \left(\sum_{k} F(Q_k)^{1/(n-1)} \right)^{(n-1)/n} \leqslant \sum_{k} F(Q_k)^{1/n} \overline{\text{Vis}}[Z \cap Q_k]^{1/n}.$$

We now apply Hölder's inequality to the products

$$F(Q_k) = \prod_{j=1}^{n} M_j(Q_k) \quad \text{and} \quad \overline{\mathrm{Vis}}[Z \cap Q_k] \lesssim \frac{1}{\theta} \prod_{j=1}^{n} \overline{V}_{Z \cap Q_k}(v_{j,k}),$$

obtaining

$$S \left(\sum_k F(Q_k)^{1/(n-1)} \right)^{(n-1)/n} \lesssim \theta^{-1/n} \prod_{j=1}^{n} \left(\sum_k M_j(Q_k) \overline{V}_{Z \cap Q_k}(v_{j,k}) \right)^{1/n}.$$

Using the cylinder estimate in equation (7.3),

$$S \left(\sum_k F(Q_k)^{1/(n-1)} \right)^{(n-1)/n} \lesssim \theta^{-1/n} \prod_{j=1}^{n} (SA(j))^{1/n} = S\theta^{-1/n} \prod_{j=1}^{n} A(j)^{1/n}.$$

Finally, we cancel the S on each side and raise the equation to the power $n/(n-1)$:

$$\sum_k F(Q_k)^{1/(n-1)} \lesssim \theta^{-1/(n-1)} \prod_{j=1}^{n} A(j)^{1/(n-1)}.$$

This establishes the inequality (7.1) and hence the theorem. $\qquad\square$

8. Box estimates for unions of tubes

The multilinear Kakeya estimate of Bennett, Carbery and Tao implies that Kakeya sets must be rather "plany". Here we give a quantitative estimate of planiness.

LEMMA 8.1. (Box estimate) *There is a constant $C(n) > 0$ such that the following holds. Suppose that $X \subset \mathbb{R}^n$ is a union of cylinders with radius 1 and length $L \gg 1$. For each $x \in X$ we can choose a convex set $B(x)$ with the following properties:*

(1) The set $B(x)$ contains x. In fact, $B(x)$ is a symmetric convex body translated so that the center is x;

(2) The set $B(x)$ has volume at most $\mathrm{Vol}(X)$;

(3) For every cylinder $T \subset X$ of radius 1 and length L, if we pick a random point $x \in T$, then the tube T lies in the rescaled set $\sigma B(x)$ with probability at least $1 - C(n)/\sigma$. (This probability estimate holds for every $\sigma > 1/C(n)$.)

Proof. We pick a collection of disjoint balls B_i of radius $\frac{1}{10}$ so that the union of $3B_i$ covers X. The number of balls is $\lesssim \mathrm{Vol}(X)$.

We can assume that $\mathrm{Vol}(X)$ is significantly less than L^n, because otherwise we just take each $B(x)$ to be a cube with side length L. By Lemma 6.6, we can choose an

algebraic hypersurface Z such that $\overline{\mathrm{Vis}}[Z \cap B_i] \geqslant L^n / \mathrm{Vol}(X)$ for each ball in our cover, with degree $\lesssim L$.

We use the hypersurface Z to define our box function $B(x)$. First take the set

$$\{v : |v| \leqslant 1 \text{ and } \overline{V}_{Z \cap B(x,1)}(v) \leqslant 1\}.$$

Let $B_0(x)$ be the translate of this set so that it is centered at x instead of at the origin. Then let $B(x)$ be the rescaling of $B_0(x)$ by a factor L, keeping it centered at x. For each $x \in X$, the unit ball $B(x,1)$ contains at least one ball B_i from our set of balls, and so $\overline{\mathrm{Vis}}[Z \cap B(x,1)] \geqslant L^n / \mathrm{Vol}(X)$. Therefore, the convex set $B_0(x)$ has volume at most $\mathrm{Vol}(X)/L^n$, and so the box $B(x)$ has volume at most $\mathrm{Vol}(X)$.

Now fix a number $\sigma > 1$ and a tube $T \subset X$ with radius 1 and length L. Let v be a unit vector pointing parallel to the core of T. If $x \in T$, then T lies in $\sigma B(x)$ unless $\overline{V}_{Z \cap B(x,1)}(v) \geqslant \frac{1}{2}\sigma$.

On the other hand, we will estimate the average value of $\overline{V}_{Z \cap B(x,1)}(v)$ as x varies in T.

LEMMA 8.2. *Let Z' denote any algebraic hypersurface of degree $\lesssim L$. Then the average value of $V_{Z' \cap B(x,1)}(v)$ over x in T is bounded as follows:*

$$\frac{1}{|T|} \int_T V_{Z' \cap B(x,1)}(v) \, dx \lesssim 1.$$

This lemma is essentially the cylinder estimate (Lemma 3.1), as we will see below. Given the lemma, we can finish the proof of the box estimate. Applying the lemma to averages over appropriate Z', we get the following estimate for the mollified volume $\overline{V}$:

$$\frac{1}{|T|} \int_T \overline{V}_{Z \cap B(x,1)}(v) \, dx \lesssim 1.$$

Let $B \subset T$ be the set of bad points where T is not contained in $\sigma B(x)$. At each bad point, $\overline{V}_{Z \cap B(x,1)}(v) \geqslant \frac{1}{2}\sigma$. Since the average value of $\overline{V}_{Z \cap B(x,1)}(v)$ is at most $C(n)$, it follows that the volume of B is at most $2C(n)|T|/\sigma$.   □

Now we turn to the proof of Lemma 8.2.

Proof of Lemma 8.2. We want to understand the average

$$\frac{1}{|T|} \int_T V_{Z' \cap B(x,1)}(v) \, dx.$$

The directional volume is itself an integral. We expand that integral and apply Fubini's theorem:

$$
\begin{aligned}
\frac{1}{|T|} \int_T V_{Z' \cap B(x,1)}(v)\, dx &= \frac{1}{|T|} \int_T \left(\int_{Z' \cap B(x,1)} |N(y) \cdot v|\, dy \right) dx \\
&\leqslant \frac{1}{|T|} \int_{Z' \cap 3T} |N(y) \cdot v| \left(\int_{B(y,1)} dx \right) dy \qquad (8.1) \\
&= \frac{C(n)}{L} V_{Z' \cap 3T}(v).
\end{aligned}
$$

But according to the cylinder estimate (Lemma 3.1), $V_{Z' \cap 3T}(v) \lesssim L$. Plugging this into (8.1), we see that our average is $\lesssim 1$.   □

Appendix A. The Lyusternik–Shnirel'man vanishing lemma

In this section, we give a proof of the vanishing lemma. There are proofs in the literature, but we will try to write the proof here in a way that is accessible with a minimum of background in algebraic topology.

First, we will prove the lemma in the special case of de Rham cohomology on a manifold. This setting is probably familiar to more readers, and the proof in this setting is clearest. In the paper, we have to apply the vanishing lemma to mod-2 cohomology, so we do the general case afterwards.

LEMMA A.1. (Vanishing lemma for de Rham cohomology; non-optimal version) *Let M be a smooth manifold. Let a_1 and a_2 be cohomology classes in $H^*(M, \mathbb{R})$. Suppose that a_1 vanishes on some open set $S_1 \subset X$ and that a_2 vanishes on some open set $S_2 \subset X$. Let $K \subset S_1 \cup S_2$ be a compact set. Then the cup product $a_1 \cup a_2$ vanishes on K.*

In fact, $a_1 \cup a_2$ vanishes on all of $S_1 \cup S_2$, not just on the compact subsets. But we have chosen to prove the weaker statement above because it makes the proof shorter and clearer.

Proof. Because we are using cohomology with real coefficients and working on a manifold, we may use de Rham cohomology. Let α_1 be a differential form that represents the cohomology class a_1. The first point of the proof is that we can choose α_1 to vanish on almost all of S_1. Let us see how to do this. We know that the restriction of a_1 to $S_1 \subset M$ is zero. In other words, the restriction of α_1 to S_1 is exact, that is there is a form β on S_1 such that $d\beta = \alpha_1$ on S_1. The form β is only defined on S_1. Now let $K_1 \subset S_1$ be any compact subset—the reader should imagine that K_1 is almost all of S_1. We can find a form β' on all of M so that β' restricted to K_1 agrees with β. Hence $d\beta'$ is an exact

form on all of M. Also $d\beta'=\alpha_1$ on K_1. Since $d\beta'$ is exact, $\alpha_1-d\beta'$ still represents the cohomology class a_1. But $\alpha_1-d\beta'$ vanishes pointwise on K_1.

By the previous paragraph, we may pick a differential form α_1 on M which represents a_1 and vanishes pointwise on K_1. By the same argument, for any compact $K_2\subset S_2$, we can pick a differential form α_2 on M which represents a_2 and vanishes pointwise on K_2. Now the wedge product of forms $\alpha_1\wedge\alpha_2$ represents the cup product $a_1\cup a_2$. On the other hand, the wedge product $\alpha_1\wedge\alpha_2$ vanishes pointwise on $K_1\cup K_2$. Hence $a_1\cup a_2$ vanishes on $K_1\cup K_2$. Therefore, $a_1\cup a_2$ vanishes on any compact subset $K\subset S_1\cup S_2$. $\square$

We now prove the vanishing lemma in general.

Proof of Lemma 5.1. This time we work with singular cohomology. Singular cohomology and cup products are well explained in Hatcher's book on algebraic topology, [8, §3.1 and §3.2]. Let α_1 be a singular cocycle representing a_1. We know that a_1 restricted to S_1 is zero. Therefore, we can choose a singular cochain β on S_1 such that $\partial\beta$ is equal to the restriction of α_1 to S_1. We can automatically extend β to a singular cochain β' on all of X. Then we look at the cocycle $\alpha_1-\partial\beta'$. Since $\partial\beta'$ is exact, this cocycle still represents the cohomology class a_1. The cocycle $\alpha_1-\partial\beta'$ vanishes on any chain supported in S_1.

By the previous paragraph, we may pick a singular cocycle α_1 representing a_1 such that α_1 vanishes on S_1. Similarly, we may pick a cocycle α_2 representing a_2 such that α_2 vanishes on S_2. We now look at the cup product $\alpha_1\cup\alpha_2$, which represents $a_1\cup a_2$. The product $\alpha_1\cup\alpha_2$ vanishes on any singular simplex supported in S_1 or supported in S_2.

Now let f denote a singular simplex supported in $S_1\cup S_2$. We will subdivide S into small pieces so that each piece lies in either S_1 or S_2. Recall that f is a continuous map from the simplex Δ to X. We subdivide the simplex into many small simplices. Restricting f to each small simplex, we get various maps g_i from Δ to X. The sum $\sum_i g_i$ is a singular chain that parametrizes the image of f. Now it is not true that $f=\sum_i g_i$ as singular chains. But it is true that $f-\sum_i g_i$ is a boundary. Since $\alpha_1\cup\alpha_2$ is a cocycle, $\alpha_1\cup\alpha_2(f)=\sum_i \alpha_1\cup\alpha_2(g_i)$. If we subdivide finely enough, then each g_i is contained in either S_1 or S_2. (At this step, we use the fact that S_1 and S_2 are open.) So each term $\alpha_1\cup\alpha_2(g_i)$ vanishes. Hence $\alpha_1\cup\alpha_2(f)=0$ for any singular simplex f in $S_1\cup S_2$. In other words, $\alpha_1\cup\alpha_2$ vanishes on $S_1\cup S_2$. Hence the cup product $a_1\cup a_2$ vanishes on $S_1\cup S_2$. $\square$

References

[1] BENNETT, J., CARBERY, A. & TAO, T., On the multilinear restriction and Kakeya conjectures. *Acta Math.*, 196 (2006), 261–302.

[2] DVIR, Z., On the size of Kakeya sets in finite fields. *J. Amer. Math. Soc.*, 22 (2009), 1093–1097.

[3] GROMOV, M., Isoperimetry of waists and concentration of maps. *Geom. Funct. Anal.*, 13 (2003), 178–215.

[4] — Singularities, expanders and topology of maps. I. Homology versus volume in the spaces of cycles. *Geom. Funct. Anal.*, 19 (2009), 743–841.

[5] GUTH, L., Minimax problems related to cup powers and Steenrod squares. *Geom. Funct. Anal.*, 18 (2009), 1917–1987.

[6] — Directional isoperimetric inequalities and rational homotopy invariants. Preprint, 2008. `arxiv:0802.3549v1 [math.DG]`.

[7] — Isoperimetric inequalities and rational homotopy invariants. Preprint, 2008. `arxiv:0802.3550v1 [math.DG]`.

[8] HATCHER, A., *Algebraic Topology*. Cambridge University Press, Cambridge, 2002.

[9] HUREWICZ, W. & WALLMAN, H., *Dimension Theory*. Princeton Mathematical Series, 4. Princeton University Press, Princeton, NJ, 1941.

[10] KATZ, N. H., ŁABA, I. & TAO, T., An improved bound on the Minkowski dimension of Besicovitch sets in $\mathbf{R}^3$. *Ann. of Math.*, 152 (2000), 383–446.

[11] LOOMIS, L. H. & WHITNEY, H., An inequality related to the isoperimetric inequality. *Bull. Amer. Math. Soc*, 55 (1949), 961–962.

[12] MATOUŠEK, J., *Using the Borsuk–Ulam Theorem*. Universitext. Springer, Berlin–Heidelberg, 2003.

[13] STONE, A. H. & TUKEY, J. W., Generalized "sandwich" theorems. *Duke Math. J.*, 9 (1942), 356–359.

LARRY GUTH
Department of Mathematics
University of Toronto
40 St. George St.
Toronto, ON M5S 2E4
Canada
lguth@math.toronto.edu