

Recent progress on the Kakeya conjecture


Terence Tao

University of California, Los Angeles

Mahler Lecture Series


The Kakeya problem

- In 1917, Soichi Kakeya proposed the following problem:
What is the smallest amount of area required to continuously rotate a unit line segment in the plane by a full rotation?
- Clearly, one can rotate the needle within a circle of radius $1/2$ (which has area $\pi/4$). Using a **deltoid** and performing a “three-point U-turn”, one can rotate using area $\pi/8$.




Besicovitch's solution

- In 1928, Besicovitch observed that one could in fact rotate a needle in arbitrary small amounts of area. The proof relied on two observations.
- First, one could translate a needle by any distance using as little area as one pleased:



- Secondly, one could find sets of arbitrarily small area that contained line segments (or thin triangles or rectangles) in every direction:



(Image from Wikipedia)

- Inspired by this, define a **Keakeya set** in \mathbf{R}^n to be a set which contained a line segment in every direction. Besicovitch's construction showed that Keakeya sets in \mathbf{R}^2 could have arbitrarily small measure; in fact, one can construct Keakeya sets which have Lebesgue measure zero.
- But not all sets of measure zero are created equal. What about the Hausdorff dimension of Keakeya sets? (One can also ask questions about other types of dimension, such as Minkowski dimension.)

- It is conjectured that
- **Keakeya set conjecture** A Keakeya set in \mathbf{R}^n has Hausdorff and Minkowski dimension n .
- This conjecture was solved for $n = 2$ by Davies in 1971, but remains open for $n \geq 3$, but there are many partial results.

- There are quantitative versions of this conjecture that are closely related. For instance, the Minkowski version of the Kakeya set conjecture is equivalent to
- **Minkowski Kakeya set conjecture, equivalent version**
For each direction $\omega \in S^{n-1}$, let T_ω be a $1 \times \delta$ tube oriented in the direction ω , where $0 < \delta \ll 1$ is small. Then the set $\bigcup_\omega T_\omega$ has volume at least $c_{n,d} \delta^{n-d}$ for every $0 < d < n$, where $c_{n,d} > 0$ depends only on n and d .
- Establishing this conjecture for a fixed d is equivalent to showing that the (lower) Minkowski dimension of Kakeya sets is at least d . For instance, the case $d = 1$ is trivial, since every tube T_ω has volume roughly δ^{n-1} .
- There is also a more technical **Kakeya maximal function conjecture** which is stronger than the above conjectures, which we will not discuss in detail here.

The Kakeya conjecture is not just a curiosity; it has turned out to have consequences for fields as diverse as

- Multidimensional Fourier summation (Fefferman 1972, Córdoba 1982, Bourgain 1995, ...);
- Wave equations (Wolff 1999);
- Behaviour of Dirichlet series in analytic number theory (Bourgain 2001); and
- Random number generation (Dvir-Wigderson 2008).

- Very roughly speaking, the connections to multidimensional Fourier summation and wave equations arise due to the existence of **wave packets** that concentrate on tubes T_ω ;
- the connection to Dirichlet series arises from a relationship between tubes T_ω and arithmetic progressions $a, a + r, a + 2r, \dots$ (which Dirichlet series can concentrate on);
- and the connection to random number generation arises by noting that a random point selected from a random tube T_ω is quite uniformly distributed even if the tubes T_ω are not.

In addition to the diverse applications of the Kakeya conjecture, the conjecture is also remarkable for the breadth of mathematical fields used in obtaining partial results on the problem, including

- Incidence geometry;
- Additive combinatorics;
- Multiscale analysis;
- Heat flows;
- Algebraic geometry;
- Algebraic topology.

In this talk I would like to discuss each of these methods in turn.

To simplify the discussion, I would like to focus on an **finite field model** for the Kakeya conjecture - a simpler “toy” version of the problem which is easier to analyse (and which, in fact, has been completely solved). The conjecture (introduced by Wolff in 1995) is as follows:

- **Finite field Kakeya conjecture** Let F be a finite field, and let E be a subset of F^n that contains a line ℓ in every direction. Then the cardinality of E is at least $c_{n,d}|F|^d$ for every $0 \leq d \leq n$, where $c_{n,d} > 0$ depends only on n, d .
- The bound is trivial for $d \leq 1$; the difficulty increases with d .

- The earliest positive results on the Kakeya problem, when interpreted from a modern perspective, were based primarily on exploiting elementary quantitative results in incidence geometry - the study of how points, lines, planes, and the like (or more precisely, δ -thickened versions of such concepts, such as balls, tubes, and slabs) intersect each other.
- These results mostly predate the finite field model, but they are easier to explain in the context of that model, so we will use the finite field setting throughout.

- Very roughly speaking, every incidence geometry fact in classical Euclidean geometry can be converted (at least in principle) via elementary combinatorial arguments into a lower bound on Kakeya sets.
- Consider for instance the Euclidean axiom that any two lines intersect in at most one point. It turns out that this can be converted to the fact that Kakeya sets in finite fields have cardinality $\gg |F|^2$ for $n \geq 2$, thus solving the Kakeya conjecture in two dimensions. (The Euclidean versions of these statements were established by Davies and by Córdoba.)


Here is a heuristic sketch of the argument.

- The set E contains about $|F|^{n-1}$ lines, one in each direction. Each contains about q points, thus leading to $|F|^n$ points in E (counting repetitions). We write $|E| = |F|^d$ for some d , then we expect each point in E to be counted about $|F|^{n-d}$ times, i.e. every point in E is incident to about $|F|^{n-d}$ lines ℓ .
- Now, let us count the number of configurations consisting of two distinct lines ℓ, ℓ' in E intersecting at a single point p in E .

- On the one hand, since there are about $|F|^{n-1}$ choices for ℓ and ℓ' , and any two lines intersect in at most one point, the number of such configurations is at most $|F|^{2(n-1)}$ or so.
- On the other hand, since there are about $|F|^d$ points, and each point is incident to about $|F|^{n-d}$ lines, then the number of configurations is at least $|F|^d \times |F|^{n-d} \times |F|^{n-d}$.
- Comparing the lower and upper bounds gives the desired bound $d \geq 2$.
- This argument can be made rigorous by exploiting the Cauchy-Schwarz inequality to handle the case when the $|F|^n$ points are not spread uniformly in E .

In a similar vein, other incidence geometry facts can be converted into bounds on this problem:

- The axiom that two distinct points determine a line can be used to give the bound $d \geq \frac{n+1}{2}$ in any dimension n . (Drury 1983, Christ, Duandikoextea, Rubio de Francia 1986).
- The fact that any three lines in general position determine a **regulus** (ruled surface), which contains a one-dimensional family of lines incident to all three of the original lines, can be used to give the bound $d \geq \frac{7}{3}$ in the three-dimensional case $n = 3$ (Schlag 1998; the bound was also obtained in Bourgain 1991).




A regulus is the union of all lines that intersect three fixed lines in general position.

- The axiom that any two intersecting lines determine a plane, which contains a one-dimensional family of possible directions of lines, was used to give the bound $d \geq \frac{n+2}{2}$ in all dimensions $n \geq 2$ (Wolff 1995). This argument relied crucially on the fact that the lines pointed in different directions; note that if one allows lines to be parallel, then a plane contains a two-dimensional family of lines rather than a one-dimensional one.
- In four dimensions (and for the finite field set problem), the Wolff bound of $d \geq 3$ in four dimensions $n = 4$ was improved slightly to $d \geq 3 + \frac{1}{16}$, by a more complicated axiom from incidence geometry, namely that any three reguli in general position determine an algebraic hypersurface, the lines in which only point in a two-dimensional family of directions (Tao 2005).

- From this sequence, it seems that one needs to imply increasingly “high-degree” facts from incidence geometry to make deeper progress on the Kakeya conjecture, in particular one begins to transition from incidence geometry to algebraic geometry.
- However, there is some evidence that the incidence geometry approach, by itself, is not sufficient to establish the full conjecture; in particular it appears difficult to raise d much above $n/2 + O(1)$ with these methods.

- In 1998, Bourgain introduced a somewhat different approach to the Kakeya problem, which relied on elementary facts of **arithmetic** (and in particular, addition and subtraction), rather than that of **geometry**, to obtain new bounds on the Kakeya problem.
- This additive-combinatorics method fared better than the geometric method in higher dimensions, basically because the additive structure of high-dimensional spaces was much the same as for low-dimensional spaces, even if the geometric structure becomes much different (and less intuitive).

- Let $E \subset F^n$ be a finite field Kakeya set, and suppose it has cardinality about $|F|^d$.
- We then write F^n as $F \times F^{n-1}$ and consider the three “slices” A, B, C of E , defined as the intersection of E with the horizontal hyperplanes $\{0\} \times F^{n-1}$, $\{1\} \times F^{n-1}$, and $\{1/2\} \times F^{n-1}$ respectively (let us assume $|F|$ is odd for the sake of discussion, so that $1/2$ is well-defined).
- Heuristically, we expect A, B, C to all have cardinality about q^{d-1} .



A Kakeya set E and three of its slices A, B, C . This generates a graph G of pairs (a, b) with few sums but many differences.

- Now, E contains about $|F|^{n-1}$ lines, each of which connect a point $a \in A$ to a point $b \in B$, with the midpoint $\frac{a+b}{2}$ lying in C .
- Since two points uniquely determine a line, we thus have a set G of about $|F|^{n-1}$ many pairs (a, b) in $A \times B$ whose sums $a + b$ are contained in a small set, namely $2 \cdot C$. (This fact alone already leads to the lower bound $d \geq \frac{n+1}{2}$ mentioned earlier.)
- On the other hand, since the lines all point in different directions, the differences $a - b$ of all these pairs (a, b) are distinct.

- The additive combinatorial strategy is then to play off the compressed nature of the sums on one hand, and the dispersed nature of the differences on the other. One of the main ingredients here are elementary additive identities, such as


$$a + b = a' + b' \implies a - b' = a' - b,$$

which suggests that collisions of sums should imply collisions of differences.

- By combining these sorts of identities with tools from combinatorics, various new lower bounds on the dimension d of Kakeya sets have been made; for instance, the current record in high dimension is $d \geq \frac{n-1}{\alpha} + 1$, where $\alpha = 1.675\dots$ is the largest root of $\alpha^3 - 4\alpha + 2 = 0$ (Katz-Tao 2002).
- While in principle the additive combinatorics method could potentially solve the whole conjecture, in practice this seems to be out of reach of current technology.

- The finite field model is considered simpler than the Euclidean model for a number of reasons, but one of the main ones is that the finite field model does not have the infinite number of scales that are present in the Euclidean setting.
- But one can reverse this viewpoint, and instead look for ways to exploit the multitude of scales available in the Euclidean case.

- One promising strategy in this direction is the **induction on scales strategy**, introduced by Bourgain for the closely related restriction and Bochner-Riesz problems, and developed further by Wolff.
- The basic idea here is to deduce a Keakeya estimate at scale δ from the corresponding the Keakeya estimate at a coarser scale, such as $\sqrt{\delta}$, by organising “thin” $\delta \times 1$ tubes into “fat” $\sqrt{\delta} \times 1$ tubes.
- A key point is that if one looks at all the thin tubes inside a single fat tube, and rescales the fat tube to be of unit size, then the thin tubes themselves turn into fat tubes.



Analysing thin tubes by placing them inside fat tubes, then blowing up the fat tubes so that the thin tubes themselves turn into fat tubes.

- Actually formalising this strategy turns out to be remarkably difficult. Using this approach (combined with the additive combinatorial approach), the dimension bound $d \geq 5/2$ of Wolff in three dimensions was improved very slightly to $d \geq 5/2 + 10^{-10}$ (Katz-Laba-Tao 2000), which is still the best bound known in the Euclidean setting in three dimensions.
- Nevertheless, the method still has further potential. In particular, it led to the **heat flow** method, which we discuss next.

- The induction on scales method uses a discrete induction on a scale parameter δ . It is possible to get sharper results by using a *continuous* scale parameter instead, which we now re-interpret as a time variable.
- To do this, one needs to replace a tube $1 \times \delta$ tube by a gaussian function, concentrated near a $1 \times \sqrt{t}$ tube, and consider various L^p -type expressions involving combinations of such functions. As t increases, these combinations of the gaussians evolve by a version of the heat equation. (Informally, one is “melting” a Kakeya set into a shapeless blob.)

- The heat flow obeys many monotonicity properties. (A well-known example is that the L^p norm of a heat flow contracts as time increases; but several other monotonicity formulae exist.)
- Using these monotonicity properties, one can establish (up to endpoints) a variant of the Kakeya conjecture, known as the **multilinear Kakeya conjecture**, in which “coplanar intersections” have been removed. (Bennett-Carbery-Tao 2006)
- Unfortunately, the method encounters technical difficulties when the tubes are allowed to intersect in a coplanar fashion; this issue has not been surmounted yet.

- A striking breakthrough on the finite field side of the Kakeya problem was achieved recently by Dvir (2008), by introducing a high-degree analogue of the low-degree algebraic geometry used in the incidence geometry approaches.
- In fact, Dvir's argument solves the finite field Kakeya conjecture completely, using two new observations from algebraic geometry.

The first idea is that low-degree hypersurfaces cannot contain many lines:

- **Lemma 1.** If S is a degree k hypersurface in F^d with $k < |F|$, then the number of possible directions of lines in S is at most $O(k|F|^{d-2})$.
- The reason for this is that the intersection of S with the hyperplane at infinity also has degree k and dimension $d - 2$, and the claim follows from the Schwartz-Zippel lemma (a higher-dimensional generalisation of the fact that a degree k polynomial of one variable has at most k zeroes).

The second idea is that small sets are contained in low-degree hypersurfaces:

- **Lemma 2.** If E is a set of points in F^d , then E is contained inside a hypersurface of degree $O(|E|^{1/d})$.
- The reason for this is that the dimension of the space of polynomials of degree at most $O(|E|^{1/d})$ is larger than $|E|$, so one can use linear algebra to find a non-trivial polynomial of this degree that vanishes at every point of E .

- Putting together the two ingredients, one obtains the Kakeya conjecture in finite fields.
- Further refinements by Dvir, Kopparty, Saraf, Sudan, and Wigderson.

- It was thought that the high-degree algebraic geometry methods of Dvir and others were strongly dependent on the discrete nature of the finite field setting, and would not extend to continuous settings such as that of Euclidean space.
- Very recently, however, Larry Guth (2008) managed to partially extend Dvir's results to this case.

- The key observation is to use a variant of a classical result in algebraic topology, namely the **Ham Sandwich Theorem**: given n bodies in \mathbf{R}^n , there is a hyperplane that bisects all of them.



- Actually, one needs a polynomial version of this theorem, first discovered by Stone and Tukey in 1941 (and rediscovered by Gromov in 2003):
- **Polynomial Ham Sandwich theorem** Given E bodies in \mathbf{R}^n , there exists an algebraic hypersurface of degree $O(E^{1/d})$ that bisects all of them.
- Compare this to “Lemma 2” from the algebraic geometry approach: given a set E in F^n , there exists an algebraic hypersurface of degree $O(|E|^{1/d})$ that passes through every point of E .

- By using the polynomial ham sandwich theorem as a substitute for Lemma 2, together with some additional arguments, Guth established the endpoint multilinear Kakeya conjecture in all dimensions. Further work by Guth-Katz, Elekes-Kaplan-Sharir, and Quilodrán extended the method to other Euclidean problems (in particular, the “joints” problem of Sharir), but the full Kakeya conjecture remains out of reach.
- Nevertheless, with so many techniques being developed, I am optimistic that the conjecture will be solved in the near future.