

THE KAKEYA CONJECTURE: WHERE DOES IT COME FROM AND WHY IS IT IMPORTANT?

JONATHAN HICKMAN

ABSTRACT. Roughly speaking, the Kakeya Conjecture asks to what extent lines which point in different directions can be packed together in a small space. In \mathbb{R}^2 , the problem is relatively straightforward and was settled in the 1970s. In \mathbb{R}^3 it is much more difficult and was only recently resolved in a monumental and groundbreaking work of Hong Wang and Joshua Zahl. This note describes the origins of the Kakeya Conjecture, with a particular focus on its classical connections to Fourier analysis, and concludes with a discussion of elements of the Wang–Zahl proof. The goal is to give a sense of why the problem is considered so central to mathematical analysis, and thereby underscore the importance of the Wang–Zahl result.

1. KAKEYA SETS

1.1. Kakeya’s question. A compact subset of \mathbb{R}^n is called a *Kakeya set* if it contains a line segment of length 1 in every possible direction. To avoid ambiguity, we describe this more formally as follows.

Definition 1.1. We say a compact set $K \subseteq \mathbb{R}^n$ is a *Kakeya set* if for all directions $\omega \in S^{n-1}$, there exists a position $a \in \mathbb{R}^n$ such that the line segment $\ell_{\omega,a} := \{a + t\omega : 0 \leq t \leq 1\}$ satisfies $\ell_{\omega,a} \subseteq K$.

In Figure 1 we illustrate three simple examples of planar Kakeya sets: a disc, an equilateral triangle and a deltoid. The most instructive is the equilateral triangle. If we imagine pinning one end of a line segment to the top corner, we can then sweep it along a 60° angle whilst staying within the triangle. This gives 60° worth of directions pointing upwards and 60° worth of directions pointing downwards, so a total of 120° of directions. We can do this for all three corners of the triangle to account for all 360° .

If we compare the areas of the three Kakeya sets in Figure 1, we see they get successively smaller. In fact, the deltoid has half the area of the disc. This means that, within the deltoid, we have managed to pack the line segments much more tightly. In the early twentieth century, Sōichi Kakeya was interested in these shapes and understanding different ways to pack lines into small spaces. From his work, a natural question arose: what is the smallest possible area of a planar Kakeya set?

1.2. Besicovitch’s surprising answer. Around the same time, Besicovitch was independently thinking about similar questions. He was able to devise new and

2020 *Mathematics Subject Classification.* Primary 28A78, 42B99.

Supported by New Investigator Award UKRI097. This article expands upon the Oberwolfach snapshot [12]. The author thanks Tony Carbery, Tom Leinster, Betsy Stovall and Joshua Zahl for helpful comments and feedback on an earlier draft.

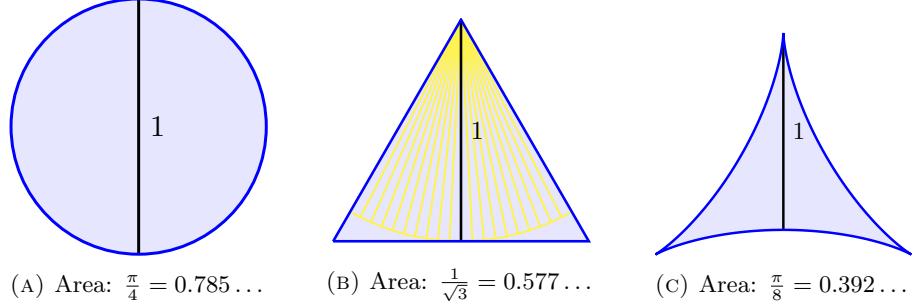


FIGURE 1. Three examples of Kakeya sets.

surprising ways to pack lines, which produced planar Kakeya sets with much, much smaller area than any example seen before. In fact, he was able to show that

there exist planar Kakeya sets of *arbitrarily small* area!

That is, for any $\varepsilon > 0$, there exists some Kakeya set $K \subseteq \mathbb{R}^2$ with $|K| < \varepsilon$. Here and below, given $E \subseteq \mathbb{R}^d$ for $d \in \mathbb{N}$, we let $|E|$ denote the d -dimensional measure of E .

How is this possible? How can we pack so many line segments into such a tiny space?¹ To understand this, we go back to the example of the equilateral triangle, which we denote by Δ . In Figure 2, we carry out the following ‘cut-and-shift’ procedure to arrive at a new shape:

- (a) Cut Δ into two ‘subtriangles’ Δ_1 and Δ_2 along the vertical midline.
- (b) Shift Δ_1 and Δ_2 horizontally in opposite directions so that they overlap.

The key observation is that, as we cut-and-shift, we do not lose any of the line segments swept out from the top corner. Furthermore, because of the overlap between the shifted subtriangles, the new shape has smaller area than Δ . The procedure therefore results in a shape with smaller area than the equilateral triangle Δ , but which still contains line segments in 120° of directions. This on its own is not very impressive. However, if we apply the procedure over and over many times, then we can amplify its effect and arrive at a startling conclusion.

Rather than cutting Δ into just two subtriangles, in Figure 3 we cut it into 2^m subtriangles for some large value of $m \in \mathbb{N}$. As illustrated, we then shift the subtriangles horizontally, first to form pairs, and then to form pairs of pairs, and then pairs of pairs of pairs and so on. By doing so, we ensure that there is a huge amount of overlap. At the end of this process, we arrive at a shape K called a *Perron tree*, represented in Figure 3d. The huge overlap means that the Perron tree K has much, much smaller area than the original triangle Δ . By choosing the number 2^m of subtriangles to be very large, and being careful how we shift, we can ensure that the area of K is as small as we wish. As before, K contains (shifted versions of) all the line segments spanned by the top corner of Δ , so lines in 120° of directions. We can then take three rotated copies of K to form a true Kakeya set containing line segments pointing in all directions.²

¹Here we sketch a simplification of Besicovitch’s argument. A detailed history of the problem, together with the precise description of the method, can be found in [9, Chapter 7].

²Taking three copies of K will increase the area by a factor of 3, but we can compensate for this by ensuring that the area of K is three times smaller than our target value.

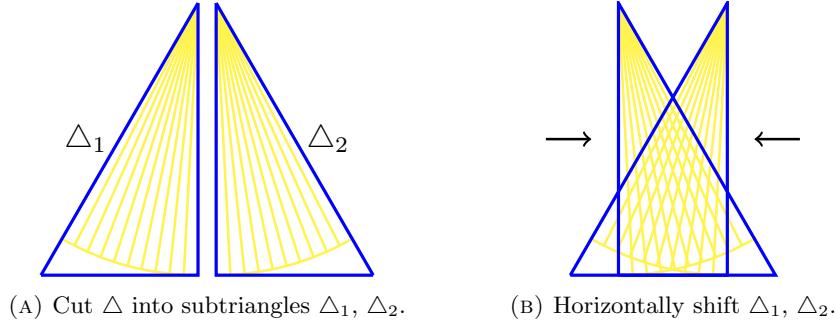


FIGURE 2. The cut-and-shift procedure.

The Perron tree construction produces a Kakeya set with area smaller than any prescribed value. In fact, combining these ideas with a limiting argument, Besicovitch was able to go one step further and show there exist Kakeya sets with *zero* area!

Theorem 1.2 (Besicovitch). *There exists a Kakeya set $K \subseteq \mathbb{R}^2$ with $|K| = 0$.*

Given a Kakeya set $K \subseteq \mathbb{R}^2$, the product $K \times [0, 1]^{n-2}$ is easily seen to be a Kakeya set in \mathbb{R}^n . Hence, Besicovitch's theorem also implies that for every $n \geq 2$ there exists a Kakeya set $K \subseteq \mathbb{R}^n$ with $|K| = 0$.

Besicovitch's construction is undoubtedly striking and beautiful. However, around the time of their discovery, it is perhaps fair to say that measure zero Kakeya sets were something of a curiosity. Besicovitch did apply his construction to solve a problem in integration theory, but beyond this it did not appear to have much use.

Mathematical ideas have a wonderful habit of reappearing in new and seemingly unrelated contexts. What was once thought a curiosity can later emerge as a profoundly important idea. Kakeya sets are a striking example of this. Fifty years after Besicovitch's work, startling discoveries uncovered their true significance.

2. CONVERGENCE OF FOURIER INTEGRALS

2.1. Fourier inversion. Fourier analysis is concerned with representing functions in terms of sine and cosine waves (or, equivalently, complex exponentials). Ostensibly, this has little to do with the kind of elementary geometric problems discussed above. However, as we shall see, it is within the world of Fourier analysis that the true significance of Kakeya sets and Besicovitch's construction comes to light.

We begin by recalling some of the basic principles of the Fourier theory. Given $f \in L^1(\mathbb{R}^n)$, the *Fourier transform* is defined by³

$$(2.1) \quad \widehat{f}(\xi) := \int_{\mathbb{R}^n} e^{-2\pi i x \cdot \xi} f(x) dx, \quad \xi \in \widehat{\mathbb{R}}^n.$$

Here we can think of the $\widehat{\mathbb{R}}^n$ notation in (2.1) as simply denoting a copy of \mathbb{R}^n . We include the hat to remind ourselves this is the domain of the Fourier transform, since later it will be useful to distinguish the domain of \widehat{f} and the domain of f .

³Recall that $L^p(\mathbb{R}^n)$ denotes the space of all measurable functions $f: \mathbb{R}^n \rightarrow \mathbb{C}$ satisfying $\|f\|_{L^p(\mathbb{R}^n)} := \left(\int_{\mathbb{R}^n} |f|^p \right)^{1/p} < \infty$.

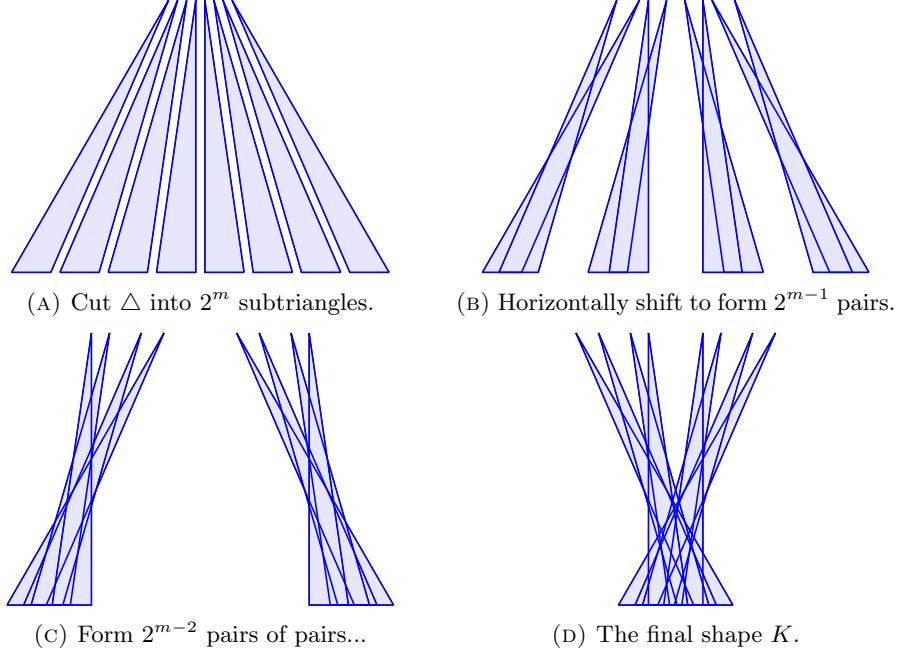


FIGURE 3. Perron tree construction of a Kakeya set.

Roughly speaking, $\widehat{f}(\xi)$ corresponds to the amplitude of the component of f which oscillates with fixed frequency vector $\xi \in \widehat{\mathbb{R}}^n$. This leads to the fundamental *Fourier inversion formula*.

Theorem 2.1 (Fourier inversion formula). *If $f \in L^1(\mathbb{R}^n)$ and $\widehat{f} \in L^1(\widehat{\mathbb{R}}^n)$, then*

$$(2.2) \quad f(x) = \int_{\mathbb{R}^n} e^{2\pi i x \cdot \xi} \widehat{f}(\xi) d\xi \quad \text{for almost every } x \in \mathbb{R}^n.$$

The inversion formula is one realisation of the fundamental principle of Fourier analysis. In particular, under the above hypotheses, (2.2) represents f as a superposition of simple waves

$$x \mapsto e^{2\pi i x \cdot \xi} \widehat{f}(\xi).$$

Each simple wave oscillates with a fixed frequency $|\xi|$ and has fixed amplitude $\widehat{f}(\xi)$. For dimensions $n \geq 2$, each wave also has an orientation: the peaks and troughs lie orthogonal to ξ . See Figure 4.

2.2. Shortcomings of the inversion formula. The hypotheses of Theorem 2.1 require both $f \in L^1(\mathbb{R}^n)$ and $\widehat{f} \in L^1(\widehat{\mathbb{R}}^n)$. These conditions are natural: we need $f \in L^1(\mathbb{R}^n)$ to define the integral in (2.1) and we need $\widehat{f} \in L^1(\widehat{\mathbb{R}}^n)$ to define the integral in (2.2). Despite this, the condition $\widehat{f} \in L^1(\widehat{\mathbb{R}}^n)$ is, in some sense, restrictive and a severe limitation on the theory. It requires

$$\int_{\widehat{\mathbb{R}}^n} |\widehat{f}(\xi)| d\xi < \infty,$$

which means $\widehat{f}(\xi)$ has fast decay as $|\xi| \rightarrow \infty$. In other words, the high frequency components of f must have very small amplitude. Intuitively, the graph of such

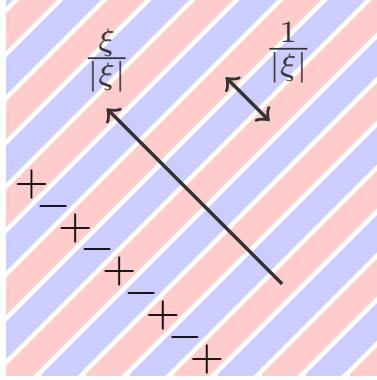


FIGURE 4. The level sets of the plane wave $x \mapsto \cos(2\pi x \cdot \xi)$, corresponding to the real part of $e^{2\pi i x \cdot \xi}$. The alternating bands represent regions where the function is either positive or negative. The wave has frequency $|\xi|$, which means that the wavelength (distance between consecutive peaks) is $1/|\xi|$. Each wave also has an orientation: the peaks and troughs lie orthogonal to ξ .

a function cannot have sudden jumps, since jumps require high frequencies. This forces f to be rather regular, which drastically limits the applicability of Theorem 2.1. Moreover, often in applications of Fourier analysis, it is critical to understand precisely *high* frequency behaviour. Since the inversion formula does not apply to functions with significant high frequency components, it is poorly suited to many of these applications.

2.3. Convergence of Fourier integrals: one dimensional case. Without the strong hypothesis $\widehat{f} \in L^1(\widehat{\mathbb{R}}^n)$, the integral formula in (2.2) is not well defined. Thus, on the face of it, it is unclear how to make sense of Fourier inversion without requiring $\widehat{f} \in L^1(\widehat{\mathbb{R}}^n)$. However, restricting to the $n = 1$ case, we can still make sense of the *partial Fourier integrals*

$$(2.3) \quad S_R f(x) := \int_{-R}^R e^{2\pi i x \xi} \widehat{f}(\xi) d\xi, \quad R > 0.$$

Since the integration now takes place over a bounded interval, the integral in (2.3) always converges. Moreover, with a small amount of work, (2.3) can be rewritten as the convolution

$$(2.4) \quad S_R f(x) = D_R * f(x) \quad \text{where} \quad D_R(x) := \frac{\sin(2\pi R x)}{\pi x}.$$

The precise formula for D_R is not too important for us. However, what is important is that (2.4) makes no reference to the Fourier transform of f and, consequently, one can easily show $S_R f$ is well defined for any $f \in L^p(\mathbb{R})$ for $1 \leq p < \infty$ (in particular, we no longer require $f \in L^1(\mathbb{R})$). It is natural to ask whether, at this level of generality, f can be reconstructed from the partial Fourier integrals $S_R f$.

Question 2.2. Does $S_R f \rightarrow f$ as $R \rightarrow \infty$?

A positive answer to Question 2.2 acts as a surrogate for the inversion formula (2.2). However, as stated, the question is somewhat vague. Since we are dealing

with limits of families of *functions*, there are many different modes of convergence to consider. The following theorem, which combines two foundational results in Fourier analysis, treats the two most natural modes.

Theorem 2.3 (Riesz [18], Carleson–Hunt [5, 13]). *Let $1 < p < \infty$ and $f \in L^p(\mathbb{R})$. Then*

- i) $\|S_R f - f\|_{L^p(\mathbb{R})} \rightarrow 0$ as $R \rightarrow \infty$;
- ii) $S_R f(x) \rightarrow f(x)$ as $R \rightarrow \infty$ for almost every $x \in \mathbb{R}$.

Theorem 2.3 represents a sweeping and powerful extension of the inversion formula in the one-dimensional setting. Since the hypotheses only require $f \in L^p(\mathbb{R})$ (and no additional conditions on the Fourier transform), the results apply in great generality, including to very rough functions. The first part of Theorem 2.3, due to Riesz [18], shows that $S_R f$ converges to f with respect to the metric structure on $L^p(\mathbb{R})$. Since the L^p -norms involve an integral, this can be interpreted as showing convergence holds ‘on average’. The second result, on almost everywhere convergence, involves no averaging and is far more subtle. This is the celebrated Carleson–Hunt theorem [5, 13], a landmark in twentieth century analysis.

2.4. Convergence of Fourier integrals: higher dimensions. As soon as we move to higher dimensions, convergence of Fourier integrals becomes a much more difficult and subtle problem. To begin with, there is no longer a ‘canonical’ choice of truncation for defining the partial Fourier integrals S_R as in (2.3). For instance, two natural choices are given by the *square* and *ball* partial integrals⁴

$$S_R^{\text{square}} f(x) := \int_{[-R, R]^n} e^{2\pi i x \cdot \xi} \widehat{f}(\xi) d\xi \quad \text{and} \quad S_R^{\text{ball}} f(x) := \int_{B(0, R)} e^{2\pi i x \cdot \xi} \widehat{f}(\xi) d\xi,$$

respectively. The above definitions coincide for $n = 1$, but are different beasts in higher dimensions. Taking the limit of the partial Fourier integrals $S_R^{\text{square}} f$ and $S_R^{\text{ball}} f$ as $R \rightarrow \infty$ corresponds to two different *summation methods* for f : that is, different ways to attempt to reconstruct f from its Fourier transform.

Square summation. The summation method for the square partial integrals is relatively easy to understand. In this setting, the one-dimensional theory from Theorem 2.3 extends directly to higher dimensions.

Theorem 2.4. *For $n \geq 1$, let $1 < p < \infty$ and $f \in L^p(\mathbb{R}^n)$. Then*

- i) $\|S_R^{\text{square}} f - f\|_{L^p(\mathbb{R}^n)} \rightarrow 0$ as $R \rightarrow \infty$;
- ii) $S_R^{\text{square}} f(x) \rightarrow f(x)$ as $R \rightarrow \infty$ for almost every $x \in \mathbb{R}^n$.

Since the square $[-R, R]^n$ is formed by a product of intervals $[-R, R]$, the square partial integral operators S_R^{square} have an underlying tensor product structure. Using this, the norm convergence problem for S_R^{square} immediately reduces to the $n = 1$ case. A similar, but this time somewhat non-trivial, reduction applies to almost everywhere convergence [11]. In this way, Theorem 2.4 can be derived as a corollary of Theorem 2.3.

Ball summation. In contrast with the previous case, the ball partial integral operators S_R^{ball} do not possess a tensor structure. Consequently, there is no obvious or easy way to reduce the convergence problem for the ball summation method to the $n = 1$ case. This situation is ‘genuinely’ higher dimensional.

⁴Here $B(0, R)$ denotes the Euclidean ball centred at the origin of radius R .

Understanding the norm convergence problem for S_R^{ball} was a major open problem in the 1960s. Using Plancherel's theorem, it is a trivial matter to see that norm convergence holds for $p = 2$; that is,

$$\|S_R^{\text{ball}} f - f\|_{L^2(\mathbb{R}^n)} \rightarrow 0 \quad \text{as } R \rightarrow \infty \quad \text{for all } f \in L^2(\mathbb{R}^n).$$

However, the behaviour for $p \neq 2$ is much less clear. It came as a great surprise when, in 1971, Charles Fefferman [10] produced a counterexample showing that norm convergence holds *only* in the trivial $p = 2$ case!

Theorem 2.5 (Fefferman [10]). *For $n \geq 2$ and $1 \leq p \leq \infty$ with $p \neq 2$, there exists some $f \in L^p(\mathbb{R}^n)$ such that*

$$\|S_R^{\text{ball}} f - f\|_{L^p(\mathbb{R}^n)} \not\rightarrow 0 \quad \text{as } R \rightarrow \infty.$$

*In particular, norm convergence holds **only** for $p = 2$.*

Fefferman's theorem also rules out almost everywhere convergence for $p \neq 2$. These results illustrate a remarkable difference between the square and ball summation methods. Whilst the square Fourier integrals have a robust convergence theory and behave just about as well as one could hope for, the ball Fourier integrals behave just about as badly as one could dread!

Fefferman's theorem is celebrated not only for its startling conclusion, but also for the powerful ideas used in the proof. In particular, his counterexample relies heavily on Besicovitch's construction of a measure zero Kakeya set! Moreover, as we shall see, the proof of Theorem 2.5 revealed a profound relationship between the geometry of Kakeya sets and the theory of the Fourier transform in two and higher dimensions. This marked a tremendous shift in perspective, and the ramifications of Fefferman's insight remain a central and intensely active topic of contemporary research in the modern Fourier analysis and beyond.

3. FEFFERMAN'S COUNTEREXAMPLE: WHY KAKEYA SETS ARE IMPORTANT

Here we provide a non-rigorous sketch of the proof of Theorem 2.5, focusing on how Kakeya sets arise in Fefferman's counterexample. We shall focus on the planar case $n = 2$, which already contains all the essential ideas. Moreover, once the $n = 2$ case is known, standard arguments easily lift the result to higher dimensions.

3.1. A low pass filter. Consider the partial integral operator $S := S_1^{\text{ball}}$, given by

$$Sf(x) = \int_{B(0,1)} e^{2\pi i x \cdot \xi} \widehat{f}(\xi) d\xi.$$

This has a natural interpretation as a *low pass filter*. It acts on an input function f by removing all high frequency components, which corresponds to multiplying \widehat{f} against the characteristic function $\chi_{B(0,1)}$ of the unit ball. A basic question is to understand the continuity of this filtering operation. Given $1 \leq p \leq \infty$, the continuity of S on the metric space $L^p(\mathbb{R}^2)$ is equivalent to the existence of a constant $C > 0$ such that

$$(3.1) \quad \|Sf\|_{L^p(\mathbb{R}^2)} \leq C\|f\|_{L^p(\mathbb{R}^2)} \quad \text{for all } f \in L^p(\mathbb{R}^2).$$

The question of whether S is continuous on $L^p(\mathbb{R}^2)$ is closely related to the question of whether L^p convergence

$$(3.2) \quad \|S_R^{\text{ball}} f - f\|_{L^p(\mathbb{R}^2)} \rightarrow 0 \quad \text{as } R \rightarrow \infty \quad \text{for all } f \in L^p(\mathbb{R}^2)$$

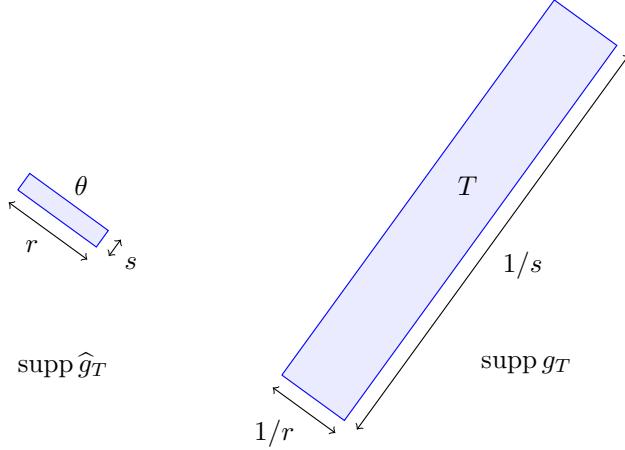


FIGURE 5. The uncertainty principle in \mathbb{R}^2 . The Fourier transform \widehat{g}_T is a bump adapted to the $r \times s$ rectangle θ . The function g_T is a normalised bump adapted to the $1/r \times 1/s$ dual rectangle T .

holds. To understand this relationship, note:

- i) L^p convergence always holds for smooth, rapidly decaying functions g . Indeed, in this case g and \widehat{g} are both well behaved and one may establish convergence via elementary tools such as the inversion formula (Theorem 2.1).
- ii) Any $f \in L^p(\mathbb{R}^2)$ can be approximated to arbitrary accuracy in $L^p(\mathbb{R}^2)$ by some smooth, rapidly decaying function $g \in L^p(\mathbb{R}^2)$.

Together, i) and ii) tell us that the limit identity in (3.2) holds for all f belonging to a dense subspace of $L^p(\mathbb{R}^2)$. If the operators $(S_R^{\text{ball}})_{R>0}$ were (equi)continuous, then one could take limits and pass from the dense subspace to *all* $L^p(\mathbb{R}^2)$ functions. More precisely, suppose we have an inequality of the form

$$(3.3) \quad \|S_R^{\text{ball}} f\|_{L^p(\mathbb{R}^2)} \leq C \|f\|_{L^p(\mathbb{R}^2)} \quad \text{for all } f \in L^p(\mathbb{R}^2),$$

where the constant $C > 0$ is uniform in R . This is a statement about the continuity of the family of linear operators $(S_R^{\text{ball}})_{R>0}$ on $L^p(\mathbb{R}^2)$. Following the above discussion, the inequality (3.3) implies the L^p convergence result (3.2).

It turns out that the converse of the above observation is also true. That is, if L^p convergence result (3.2) holds, then the inequality (3.3) must also be true. This implication is deeper, relying on the *principle of uniform boundedness*; we do not discuss the details here. The key takeaway is that, in order to prove Fefferman's theorem that the L^p convergence result (3.2) fails, it suffices to show that the inequality (3.3) fails for all $p \neq 2$. We shall demonstrate this failure for $R = 1$, in which case (3.3) agrees with (3.1). Thus, matters are reduced to showing that

$$\text{the inequality (3.1) fails for all } p \neq 2.$$

This is the goal for the remainder of the section.

3.2. Elementary properties of the Fourier transform. Fefferman's argument relies on two elementary properties of the Fourier transform.

The uncertainty principle. Rather than a single result or set of results, the uncertainty principle is a guiding heuristic which permeates all aspects of Fourier analysis. Loosely speaking, it says that:

Localising \hat{f} in the frequency domain forces f to spread out in the spatial domain.

To illustrate this idea on the real line, for a small number $r > 0$, consider the scaled Gaussian $f_r(x) := re^{-\pi r^2 x^2}$ with variance $1/r^2$. Then $\hat{f}_r(\xi) = e^{-\pi x^2/r^2}$ is a Gaussian with variance r^2 . The variance of f_r and the variance of \hat{f}_r satisfy a reciprocal relationship. By shrinking r , the distribution of \hat{f}_r becomes narrower (more certain) whilst the distribution of f_r becomes wider (more uncertain). This behaviour is a manifestation of the uncertainty principle.

We are interested in a two-dimensional variant of the above example. For small numbers $r, s > 0$, let $\theta \subseteq \mathbb{R}^2$ be a rectangle of dimensions $r \times s$, with arbitrary centre and orientation. We define the *dual rectangle* $T \subseteq \mathbb{R}^2$ to be the $1/r \times 1/s$ rectangle centred at the origin with the same orientation as θ . We illustrate this definition in Figure 5.

Suppose g_T is a function on \mathbb{R}^2 whose Fourier transform \hat{g}_T is a *bump adapted to the rectangle* θ . By this, we mean that \hat{g}_T looks like⁵ a smoothed out version of the characteristic function χ_θ . Then the uncertainty principle tells us that the support of g_T spreads out over the dual rectangle T . This is a two-dimensional analogue of the Gaussian example above. In particular, if we think of \hat{g}_T as a smoothed out version of the characteristic function χ_θ of θ , then, at least morally, $|g_T|$ should be a smoothed out version of the normalised characteristic function $|T|^{-1} \chi_T$.^{6,7}

Translation vs modulation symmetry. Given $y \in \mathbb{R}^2$, the Fourier transform intertwines the *translation* and *modulation* operations

$$(\tau_y f)(x) := f(x - y) \quad \text{and} \quad \text{Mod}_y h(\xi) := e^{-2\pi y \cdot \xi} h(\xi).$$

In particular, the integral formula (2.1), together with a change of variables, implies

$$(\tau_y f)^\wedge(\xi) = \int_{\mathbb{R}^n} e^{-2\pi i x \cdot \xi} f(x - y) dx = e^{-2\pi i y \cdot \xi} \int_{\mathbb{R}^n} e^{-2\pi i x \cdot \xi} f(x) dx = \text{Mod}_y \hat{f}(\xi).$$

Thus, translating the function f in the physical domain corresponds to modulating \hat{f} in the frequency domain.

3.3. The counterexample. We now have all the ingredients we need to sketch Fefferman's argument. We break the proof into steps.

Step 0: Recap. We begin by recalling what it is we wish to prove. As above, S denotes the low pass filter

$$Sf(x) := \int_{B(0,1)} e^{2\pi i x \cdot \xi} \hat{f}(\xi) d\xi.$$

⁵We intentionally keep the discussion vague here to avoid technicalities. For the purposes of this exposition, one could simply think of $\hat{g}_T = \chi_\theta$. However, in practice, it is useful to take a smoothed out version of the characteristic function, since the Fourier transform is poorly behaved in the presence of discontinuities.

⁶The normalising factor $|T|^{-1}$ naturally appears owing to the scaling properties of the Fourier transform. It plays the same role as the r factor in the definition of f_r in the Gaussian example.

⁷This discussion applies to the *absolute value* $|g_T|$. The function g_T itself carries some oscillation, which is dictated by the location of θ in the frequency domain.

This removes components of f oscillating with frequency $|\xi| > 1$. By our earlier reductions, our goal is to show

$$(3.4) \quad \|Sf\|_{L^p(\mathbb{R}^2)} \leq C\|f\|_{L^p(\mathbb{R}^2)} \quad \text{for all } f \in L^p(\mathbb{R}^2)$$

fails whenever $p \neq 2$. Suppose, to the contrary, that the inequality (3.4) holds. What does this tell us?

The L^p norms are sensitive to the distributional properties (that is, the shape of the graph) of a given function. Since L^∞ functions are by definition bounded, their distributions are typically broad and flat. On the other hand, L^1 functions are defined in terms of the convergence of an integral, which requires some decay at ∞ , so their distributions are typically tall and narrow. A function in L^p for some $1 < p < \infty$ lies in a spectrum between these two extreme cases.

If the inequality (3.4) were to hold, say for large $p > 2$, then this would tell us that the filter S preserves certain distributional properties of the input. In particular, if the input function f is broad and flat, then the output function Sf should also be broad and flat. Therefore, to disprove the inequality, we shall construct an input function f whose distributional properties change significantly under filtering:

- The input f will have a broad, flat distribution;
- The output Sf will have a tall, narrow distribution.

This kind of behaviour is incompatible with (3.4) for $p > 2$ and therefore provides a counterexample for this exponent regime. A similar dual construction can be used to rule out (3.4) for $p < 2$.

Step 1: Arrange frequency support. We now turn to constructing the counterexample. That is, we set about finding a function f with the above properties. Fix a small parameter $0 < r < 1$. As illustrated in Figure 6a, we place $r \times r^2$ rectangles θ tangentially to the unit circle $\partial B(0, 1)$ in the frequency domain $\widehat{\mathbb{R}}^2$. The choice of side lengths $r \times r^2$ reflects the curvature of the circle: with this choice, the θ can be arranged to fit neatly around the circle and form an efficient covering.

Step 2: Initial physical support. Given θ , let T denote the corresponding $1/r \times 1/r^2$ dual rectangle, which is centred at the origin. The collection of all such dual rectangles is illustrated in Figure 6b. By the **uncertainty principle**, we can find a function g_T such that \widehat{g}_T is a bump adapted to θ and g_T is a normalised bump adapted to T . We initially consider the function $g := \sum_T g_T$. The two sides of Figure 6 then provide a schematic of the *frequency support* of g (the support of \widehat{g} in $\widehat{\mathbb{R}}^2$) and *physical support* of g (the support of g itself in \mathbb{R}^2).

Step 3: Rearrange physical support. Using **translation vs modulation symmetry**, by independently modulating each \widehat{g}_T , we translate the g_T in physical space. This corresponds to translating the underlying supporting rectangles T . We rearrange the T so, rather than all passing through the origin, they are disjoint. Moreover, as illustrated in Figure 7a, we strategically place the T below the base of a Perron tree.⁸ We let f_T denote the corresponding translate of g_T and $f = \sum_T f_T$. Figure 7a then provides a schematic of the physical support of f .

Recall that \widehat{f}_T is obtained by modulating \widehat{g}_T . Modulation does not alter the support of a function so \widehat{f}_T and \widehat{g}_T share the same same support. Thus, Figure 6a provides a schematic of the frequency support of both $g = \sum_T g_T$ and $f = \sum_T f_T$.

⁸Since a single Perron tree only accounts for lines in 120° of directions, in practice we use multiple rotated Perron trees. However, for the sake of simplicity we ignore this minor detail.

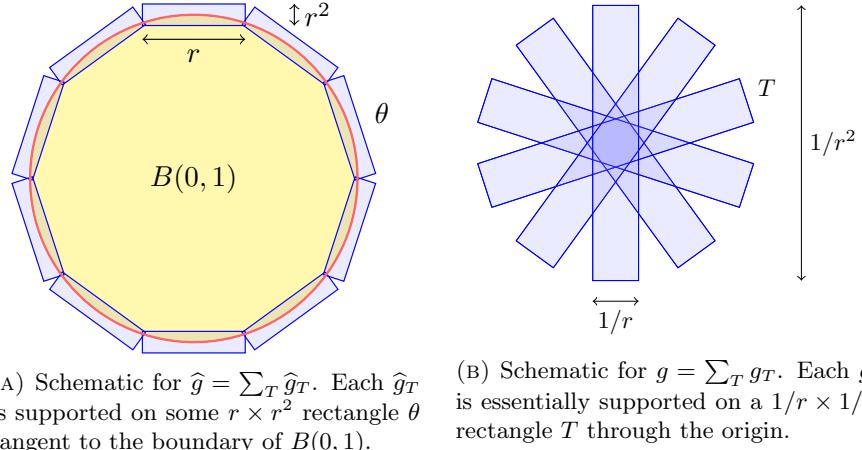


FIGURE 6. The function g from Step 2 of Fefferman's construction, viewed from both the frequency and physical domain. Each spatial rectangle T in Figure 6b is dual to some frequency rectangle θ in Figure 6a.

Step 4: Apply the operator. We now consider what happens when we apply the filter S to the function f constructed above. By linearity, $Sf = \sum_T Sf_T$ and so it suffices to study the individual f_T .

We first describe the action of S from the perspective of the frequency domain. Recall that \hat{f}_T is supported on a frequency rectangle θ , arranged along the boundary of the ball as in Figure 6a. The rectangle θ is tangential to the unit circle, and the filter S acts by cutting off frequencies to the unit ball. Consequently, S effectively cuts the support θ of \hat{f}_T in half.

We now consider the action of S from the perspective of the spatial domain. As discussed above, S effectively halves the frequency support of f_T . By the **uncertainty principle**, halving the frequency support causes the physical support to double. Consequently, in the spatial domain, each Sf_T spreads out over a double of T . This phenomenon is illustrated in Figure 7.

Step 5: Conclusion. As we can see in Figure 7b, the strategic placement of the rectangles T causes the doubles to pile up on top of one another. Moreover, since the Perron tree packs lines in different directions into a tiny space, the pile up between the doubles is enormous. This pile up is key.

From the previous steps we know:

- The disjoint rectangles in Figure 7a correspond to the supports of the functions f_T forming $f = \sum_T f_T$;
- The piled-up doubles in Figure 7b correspond to the supports of the functions Sf_T forming $Sf = \sum_T Sf_T$.

Comparing both sides of Figure 7, we pass from an input function f with a broad, flat distribution to an output function Sf with a tall, narrow distribution. Thus, S fails to preserve the distributional properties of the input, and therefore cannot be bounded on L^p .

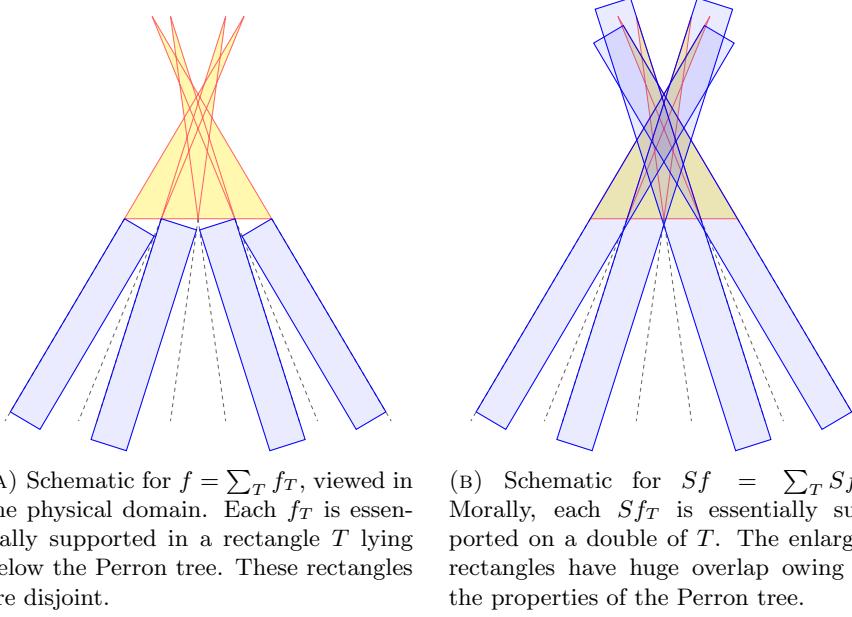


FIGURE 7. Applying the operator S to the function f . The rectangles on the left represent the support of the input f . The rectangles on the right represent the support of the output Sf .

The above argument is a rough sketch, and for the sake of brevity omits some details.⁹ Nevertheless, it serves to highlight the essential role of the geometry of the Kakeya sets in Fefferman's argument. It is only by arranging the rectangles in relation to (an approximation of) a measure zero Kakeya set that the pile up effect is sufficient to prove the theorem.

Fefferman's profound discovery showed that the geometry of Kakeya sets underpins Fourier analysis in two and higher dimensions. Since Fourier analysis is so central to the study of many physical processes, it soon became apparent that Kakeya sets were also deeply connected to the theory of many important equations governing the laws of nature. This includes the wave equation, which describes a diverse range of phenomena such as water waves, seismic waves and sound waves, and also Schrödinger's famous equation from quantum mechanics [19, 25]. Kakeya sets, which were once was a curiosity, suddenly became one of the most important objects in mathematical analysis!

4. THE NEXT STEPS IN FOURIER ANALYSIS

4.1. A possible fix: the Bochner–Riesz Conjecture. Fefferman's example demonstrated a dramatic and highly unexpected failure of the ball Fourier summation method. In response to this, the hunt was on for a possible fix. A natural candidate is to consider slightly smoothed out versions of the ball summation method.

⁹A significant omission is the effect of the oscillation of the individual terms of the sum $Sf = \sum_T Sf_T$. If there were a lot of destructive interference between the Sf_T , then this could, in principle, cancel out the effect of the pile up of the supports. However, it is reasonably straightforward to understand and control the oscillation in the proof.

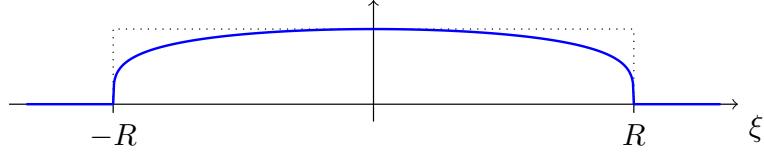


FIGURE 8. A radial slice of $m_R^\alpha(\xi) := (1 - |\xi|^2/R^2)^\alpha \chi_{B(0,R)}(\xi)$ from the definition of the Bochner–Riesz means. Here $\alpha = 1/4$. In general, for small $\alpha > 0$, the function is continuous but fails to be smooth around the boundary of the $B(0, R)$.

Given $\alpha \geq 0$, we define the *Bochner–Riesz means*

$$B_R^\alpha f(x) := \int_{B(0,R)} e^{2\pi i x \cdot \xi} (1 - |\xi|^2/R^2)^\alpha \widehat{f}(\xi) d\xi, \quad R > 0.$$

When $\alpha = 0$, this definition agrees with $S_R^{\text{ball}} f$. However, for $\alpha > 0$, the Bochner–Riesz mean $B_R^\alpha f$ replaces the rough frequency cutoff $\chi_{B(0,R)}$ in the definition of $S_R^{\text{ball}} f$ with the smoothed out version $m_R^\alpha(\xi) := \chi_{B(0,R)}(\xi)(1 - |\xi|^2/R^2)^\alpha$. The graph of this function is sketched in Figure 8. For $\alpha > 0$ small, m_R^α is continuous, but fails to be differentiable at the boundary of $B(0, R)$. As α increases, it becomes more and more regular along the boundary of the ball.

The Bochner–Riesz means provide a natural mollification of the ball summation method. Given $1 \leq p \leq \infty$, this prompts the question of whether

$$(4.1) \quad \|B_R^\alpha f - f\|_{L^p(\mathbb{R}^n)} \rightarrow 0 \quad \text{as } R \rightarrow \infty \quad \text{holds for all } f \in L^p(\mathbb{R}^n).$$

As for the ball summation method, Plancherel's theorem shows (4.1) is trivially true for $p = 2$. The *Bochner–Riesz Conjecture* asserts that, as soon as $\alpha > 0$, there is a range¹⁰ of exponents p for which (4.1) holds, beyond the trivial $p = 2$ case.

4.2. A counterexample to the Bochner–Riesz Conjecture? One lesson of Fefferman's work is that it is healthy to have a good dose of scepticism. Rather than try to prove a positive result on the Bochner–Riesz Conjecture, we could first try to construct a counterexample.

As it stands, Fefferman's argument, based on the Perron tree construction of a Kakeya set of measure zero, is too weak to contradict the Bochner–Riesz Conjecture. Any tiny additional degree of smoothing (that is, any choice of $\alpha > 0$) is enough to dampen and overcome the pile up effect described in §3. But what would happen if there was a way to construct even smaller Kakeya sets?

Despite having zero measure, the planar Kakeya set arising from Perron tree construction is still relatively large: it has *dimension* 2. Here *dimension* can either refer to the *Minkowski* or the *Hausdorff dimension*. We shall (briefly and informally) review these concepts in §5.3 below. For our present purposes, however, it suffices to intuitively think of a 2-dimensional subset of \mathbb{R}^2 as large and filling out space (whilst still possibly having zero measure).

Suppose there exists a Kakeya set $K \subseteq \mathbb{R}^n$ which not only has measure zero, but is lower dimensional; say the Minkowski dimension of K is less than n . This pathological object would pack all the line segments into a truly tiny space. We

¹⁰The precise statement of the conjecture posits that this should hold for $\frac{2n}{n+1} < p < \frac{2n}{n-1}$, but we shall not concern ourselves with the exact numerology here.

could then apply Fefferman's argument to *this* Kakeya set. The pile up effect for our new K would be monstrous: far worse than the Perron tree construction. In particular, it would overwhelm the effect of smoothing, at least for small values of $\alpha > 0$, and give a counterexample to the Bochner–Riesz Conjecture.

5. DIMENSIONS OF KAKEYA SETS

5.1. The Kakeya Conjecture. The existence of low dimensional Kakeya sets would have profound consequences for harmonic analysis, providing counterexamples to many important problems, including the Bochner–Riesz Conjecture discussed above. However, it seems very difficult to construct a Kakeya set which is significantly smaller than that coming from the Perron tree construction. This led to the widely held belief that low dimensional Kakeya sets do not exist.

Conjecture 5.1. Any Kakeya set $K \subseteq \mathbb{R}^n$ has dimension n .

Here, as before, *dimension* refers either to the *Minkowski* or the *Hausdorff* dimension. At its heart the meaning of the conjecture is simple: although Kakeya sets can have measure zero, they cannot be *too* small.

If the Kakeya Conjecture failed, then there would exist a Kakeya set in \mathbb{R}^n of small dimension (less than n). Following the discussion of the previous section, this set could be used to produce a Fefferman-style counterexample to the Bochner–Riesz Conjecture. In this way, we see that

$$(5.1) \quad \text{Bochner–Riesz Conjecture} \implies \text{Kakeya Conjecture.}$$

Many other central conjectures in Fourier analysis are known to imply the Kakeya Conjecture, following a similar scheme. But the significance of the Kakeya Conjecture stretches far beyond Fourier analysis! There are important problems in PDE, combinatorics and number theory which are also known to either imply the Kakeya Conjecture or some variant.

The Kakeya Conjecture in the plane is not too difficult. In the 1970s, Davies [7] and later Córdoba [6] gave two very different proofs showing that any planar Kakeya set has dimension 2.¹¹ Correspondingly, the Bochner–Riesz Conjecture is also known to be true in two dimensions, and in general we have a relatively good understanding of Fourier analysis in the planar setting. By contrast, the Kakeya Conjecture in three (and higher) dimensional space is much, much harder. The additional directions in three dimensions complicate matters enormously. Correspondingly, our understanding of many central questions in Fourier analysis in \mathbb{R}^3 is severely lacking. The Kakeya Conjecture in higher dimensions has been a major bottleneck in harmonic analysis over the last 50 years. For this reason, it is considered one of the most important open problems in the area.

5.2. Spectacular progress. Many mathematicians have studied the Kakeya Conjecture in three (and also higher) dimensions over the last 50 years. By their efforts, deep connections were discovered between the problem and questions in PDE, arithmetic combinatorics, number theory, and beyond [14, 25]. Partial results on the

¹¹See, for instance, [25] and [17] for expositions of these results.

Kakeya Conjecture have led to dramatic breakthroughs on other problems.¹² However, despite remarkable successes, a full solution, even for $n = 3$, always appeared far out of reach. The Kakeya Conjecture achieved a certain notoriety and became a byword for a difficult, near impossible challenge. This was the state of affairs until February 2025, when Hong Wang and Joshua Zahl [22] announced the final piece of a spectacular full solution of the three dimensional problem!

Theorem 5.2 (Wang–Zahl [21, 22]). *Any Kakeya set $K \subseteq \mathbb{R}^3$ has dimension 3.*

The monumental proof of Theorem 5.2 is hundreds of pages long and is spread over multiple papers. It combines a rich array of ideas, and relies on contributions from many different mathematicians over the last 30 years. However, their arguments also incorporate a wealth of new and deep insights and mark a major leap in our understanding. The work of Wang and Zahl is a great and historic achievement in modern mathematics.

5.3. Elements of the proof. We conclude with a discussion of some of the features of the Wang–Zahl proof. Given the length and complexity of their arguments, we only provide the roughest of sketches.

Discretisation. The first step is a standard reformulation of the problem. Suppose $K \subseteq \mathbb{R}^3$ is a Kakeya set and, for $0 < \delta < 1$, let $N_\delta K$ denote the δ -neighbourhood of K . We may assume K has measure zero (since otherwise it automatically has dimension 3). It then follows that

$$(5.2) \quad |N_\delta K| \rightarrow 0 \quad \text{as } \delta \rightarrow 0_+,$$

where here the notation $|\cdot|$ is used to denote 3-dimensional volume.

The Minkowski dimension quantifies the rate of convergence in (5.2). In order to prove our Kakeya set K has Minkowski dimension 3, given $\varepsilon > 0$, we need to show that there exists some constant $c_\varepsilon > 0$ such that

$$(5.3) \quad |N_\delta K| \geq c_\varepsilon \delta^\varepsilon \quad \text{for all } \delta > 0.$$

In particular, although $|N_\delta K| \rightarrow 0$ as $\delta \rightarrow 0_+$, the rate of convergence is much slower than any power of δ . In this sense, K is ‘large’. We remark that Hausdorff dimension provides a more nuanced notion of size, which relies on analysing K simultaneously at many scales, rather than just one fixed scale δ . For the purposes of this sketch, however, we stick to the simpler notion of Minkowski dimension.

When we form the δ -neighbourhood $N_\delta K$ of K , each line segment in K fattens up into a solid shape which is roughly a cylinder of length 1 and radius δ . We call these cylinders δ -tubes. We expect many of the δ -tubes formed from lines in K to essentially coincide. For this reason, we sample from K a finite family of line segments \mathcal{L} with directions forming a maximal, δ -separated set. The corresponding δ -tubes are then *essentially distinct* (they differ significantly from one another). To succinctly describe this setup, we make the following definition.

Definition 5.3. For $0 < \delta < 1$, we say that a set \mathbb{T} of δ -tubes is *direction-separated* if the directions of the core lines of the tubes in \mathbb{T} form a δ -separated set.

In light of the above, the proof of (5.3) reduces to showing the following.

¹²To highlight one particularly striking instance, the *multilinear* Kakeya theory of Bennett–Carbery–Tao [1] led to the resolution of the longstanding main conjecture in Vinogradov’s mean value theorem [4] in analytic number theory.

Theorem 5.4 (Kakeya Conjecture, discretised version [21, 22]). *Given any $\varepsilon > 0$, there exists some $c_\varepsilon > 0$ such that for all $0 < \delta < 1$ the inequality*

$$(5.4) \quad \left| \bigcup_{T \in \mathbb{T}} T \right| \geq c_\varepsilon \delta^\varepsilon \sum_{T \in \mathbb{T}} |T|$$

holds whenever \mathbb{T} is a direction-separated set of δ -tubes in \mathbb{R}^3 .

If the tubes in \mathbb{T} were disjoint, then we would have $\left| \bigcup_{T \in \mathbb{T}} T \right| = \sum_{T \in \mathbb{T}} |T|$. Thus, (5.4) can be interpreted as saying the tubes in \mathbb{T} are ‘almost’ disjoint.

To pass from (5.4) to (5.3), we simply take \mathbb{T} to be the set of tubes formed from the line segments belonging to \mathcal{L} as defined earlier. On the one hand, $N_\delta K \supseteq \bigcup_{T \in \mathbb{T}} T$. On the other hand, maximality ensures that $\sum_{T \in \mathbb{T}} |T| \sim 1$. These observations combine with (5.4) to give (5.3). The advantage of the reformulation in Theorem 5.4 is that, rather than studying a *continuum* of line segments in K , here we study a *finite* set of δ -tubes. This is a form of *discretisation*, a common tool in Fourier analysis and geometric measure theory.¹³

Exploiting distinct directions. It is surprisingly difficult to directly leverage the *direction-separated* hypothesis from the Kakeya set definition. In practice, one often works with weaker, and more malleable, conditions. We noted earlier that direction-separated tubes are *essentially distinct*. This is formally defined as follows.

Definition 5.5. Let $0 < \delta < 1$ and \mathbb{T} be a set of δ -tubes. We say the tubes in \mathbb{T} are *essentially distinct* if $|T_1 \cap T_2| \leq |T_1|/2$ for all $T_1, T_2 \in \mathbb{T}$ with $T_1 \neq T_2$.

It is natural to ask whether the conclusion of Theorem 5.4 continues to hold under the weaker hypothesis that the tubes in \mathbb{T} are essentially distinct. However, there is an easy counterexample. Assume $0 < \delta < 1$ satisfies $\delta^{-1} \in \mathbb{N}$ and let

$$A := \{(\delta j, 0, 0) : 1 \leq j \leq \delta^{-1}\} \quad \text{and} \quad B := \{(\delta j, 1, 0) : 1 \leq j \leq \delta^{-1}\}$$

be sets of equally spaced points along a pair of parallel lines in $S := [0, 1]^2 \times \{0\}$. Form a family \mathcal{L} of δ^{-2} line segments by joining every point in A to a point in B . If \mathbb{T} denotes the corresponding set of δ -tubes, then $\#\mathbb{T} = \delta^{-2}$ and the tubes in \mathbb{T} are essentially distinct. However, $T \subseteq N_\delta S$ for all $T \in \mathbb{T}$ and, moreover, $|\bigcup_{T \in \mathbb{T}} T| \sim \delta$. Thus, (5.4) fails in this case.

In the above example, the δ -tubes concentrate in $N_\delta S$, which is essentially a $1 \times 1 \times \delta$ rectangular prism. What happens if we rule out this behaviour?

Definition 5.6 ([24]). A set \mathbb{T} of δ -tubes in \mathbb{R}^3 satisfies the *Wolff axiom* if

$$(5.5) \quad \#\{T \in \mathbb{T} : T \subseteq R\} \leq \delta^{-2}|R|$$

holds whenever $R \subseteq \mathbb{R}^3$ is a rectangular prism.

If \mathbb{T} is direction-separated, then it is not difficult to see that \mathbb{T} automatically satisfies the Wolff axiom. It is natural to ask whether the Wolff axiom is sufficient to ensure the conclusion of Theorem 5.4. It turns out that this is indeed the case.

Theorem 5.7 (Strong Kakeya Conjecture [21, 23, 22]). *Given any $\varepsilon > 0$, there exists some $c_\varepsilon > 0$ such that for all $0 < \delta < 1$ the inequality*

$$(5.6) \quad \left| \bigcup_{T \in \mathbb{T}} T \right| \geq c_\varepsilon \delta^\varepsilon \sum_{T \in \mathbb{T}} |T|$$

¹³We also saw an instance of discretisation in the construction of the Fefferman counterexample.

holds whenever \mathbb{T} is a collection of δ -tubes in \mathbb{R}^3 which satisfy the Wolff axiom.

At this point we encounter the first major indication of the subtlety of the Kakeya problem: Theorem 5.7 *fails* over the complex field! In [15, §13], authors considered the compact piece of the *Heisenberg group* in \mathbb{C}^3 , defined by

$$H := \{(z_1, z_2, z_3) \in \mathbb{C}^3 : \Im(z_1) = \Im(z_2\bar{z}_3), |z_1|, |z_2|, |z_3| \leq 2\}.$$

It is easy to see that H contains a 4 (real) parameter family of complex line segments

$$\ell_{a,b,w} := \{(z, w + az, z\bar{w} + b) : z \in \mathbb{C}, |z| \leq 1/2\},$$

where $a, b \in \mathbb{R}$ and $w \in \mathbb{C}$ with $|a|, |b|, |w| \leq 1$. From this, one may construct a set \mathbb{T} of complex δ -tubes with $\#\mathbb{T} \sim \delta^{-4}$ which satisfies the natural complex analogue of the Wolff axiom. However, $T \subseteq N_\delta H$ for all $T \in \mathbb{T}$ and, moreover, $|\bigcup_{T \in \mathbb{T}} T| \sim \delta$. Thus, \mathbb{T} fails (5.6).

In light of the above example, any proof of Theorem 5.7 must somehow distinguish between the real and complex field. The definition of the Heisenberg group relies on the notion of the real and imaginary parts of a complex number, and the accompanying conjugation operator $z \mapsto \bar{z}$. These definitions in turn rest on the existence of the half-dimensional subfield $\mathbb{R} \subseteq \mathbb{C}$. There is no apparent analogue of this structure in the real line. Indeed, a famous result of Edgar–Miller [8] shows that any Borel subring of \mathbb{R} either has Hausdorff dimension 0 or is the whole of \mathbb{R} . This fact will play an important role later in the proof.

Reducing to a symmetric case. To prove Theorem 5.7, we can attempt to isolate configurations of tubes which are extremal, or optimal, in the sense that $|\bigcup_{T \in \mathbb{T}} T|$ is minimal over all sets of tubes \mathbb{T} of a fixed cardinality satisfying the Wolff axiom. If we can prove that the Strong Kakeya Conjecture holds for such extremal configurations, it will automatically hold in all cases.

It is natural to expect an optimal configuration of tubes should saturate the Wolff axiom hypothesis of Theorem 5.7. Indeed, if there was any leeway in our hypotheses, we could hope to exploit it to construct a tighter packing of tubes. This leads to the definition of a *sticky* set of tubes.

Definition 5.8 (Sticky, informal). For each intermediate scale $\delta \leq \rho \leq 1$, we consider fattening the δ -tubes T into ρ -tubes T_ρ . We expect many of the ρ -tubes formed in this way to essentially coincide. We therefore extract a maximal family of essentially distinct ρ -tubes \mathbb{T}_ρ . We say the collection of δ -tubes is *sticky* if

$$(5.7) \quad \#\{T \in \mathbb{T} : T \subseteq T_\rho\} \sim (\rho/\delta)^2 \quad \text{for all } T_\rho \in \mathbb{T}_\rho \text{ and all } \delta \leq \rho \leq 1.$$

In other words, each ρ -tube in \mathbb{T}_ρ contains the maximal number of δ -tubes permitted by the Wolff axiom condition (5.5).

Sticky sets of tubes are highly symmetric from the perspective of metric dimension. The condition (5.7) relates the distribution of the tubes \mathbb{T} at different scales $\delta \leq \rho \leq 1$ and is an example of *statistical self-similarity*: while the set is not literally self-similar (in the sense of the middle third Cantor set, for example), certain quantitative properties of the set are identical at many different scales. Our experience and intuition suggest optimal configurations should be highly symmetric. Moreover, the Kakeya Conjecture is both a quantitative and a multi-scale problem: we are interested in the size, and not the exact shape, of Kakeya sets and we need to prove results *for all* scales δ . The relevant notions of symmetry should reflect these fundamental features. This inexorably leads to statistical self-similarity. Thus,

sticky sets of tubes arise as a natural and important special case in the study of the Kakeya Conjecture.

Theorem 5.9 (Sticky reduction [22]). *In order to prove Theorem 5.7, it suffices to assume, in addition to the existing hypotheses, that \mathbb{T} is sticky.*

This is the most significant innovation in the work of Wang–Zahl. Loosely speaking, Theorem 5.9 formalises our intuition that optimal configurations of tubes for Theorem 5.7 correspond to the highly symmetric arrangements for which the Wolff axiom is saturated. However, making this intuition precise is monumentally challenging. Reduction to the sticky case was previously explored by Wolff and Katz–Laba–Tao [15] in the early 2000s, but with limited success. Since then very little progress had been made on this problem. Indeed, prior to Wang–Zahl [22], there was no hint of a viable mechanism for reducing to the sticky case; to the contrary, evidence was mounting [16] that such a reduction was infeasible. The proof of Theorem 5.9 forms the entirety of the 127 page preprint [22] and introduces very important new ideas.

It is unfortunately outside the modest scope of this article to delve into the proof of Theorem 5.9. However, we mention that it relies on an intricate induction or self-improving mechanism. For this reason, it is essential to work with the Wolff axiom formulation of the problem as in the Strong Kakeya Conjecture. Indeed, the hypotheses must be sufficiently general in order to close the inductive arguments.

Forbidden symmetries. It remains to prove the Strong Kakeya Conjecture under the additional hypothesis that the set of tubes \mathbb{T} is sticky. To treat this case, Wang–Zahl [21, 23] were inspired by an approach originally proposed by Katz and Tao, and shared with the public in 2014 in a talk and blogpost [20] by Tao. At that time, Katz–Tao were unable to rigorously implement their proposal. Indeed, the implementation in Wang–Zahl [21] involves substantial new ideas and the use of mathematical tools and techniques which were unavailable back in 2014.

Katz–Tao proposed a proof by contradiction. We assume there exists a sticky set of tubes \mathbb{T} which satisfies the hypotheses of the Strong Kakeya Conjecture, but for which $K := \bigcup_{T \in \mathbb{T}} T$ has very small volume. Moreover, we may assume the volume of K is minimal over all such sets of tubes \mathbb{T} of a fixed cardinality.

The stickiness hypothesis means that K is highly symmetric. Under the minimal volume condition, the symmetries are further amplified. In particular, following landmark ideas introduced by Katz–Laba–Tao [15], one can show the tubes in \mathbb{T} must have extremely rich structural properties. Katz–Tao [20] proposed using these structural properties to encode highly symmetric *algebraic* structures inside the set K . This is in part motivated by the earlier observation that any proof of Theorem 5.7 must distinguish between the real and complex field (and, consequently, it is natural to expect algebra to play a role at some level).

One of the most symmetric objects in mathematics is a commutative ring: it simultaneously admits *two* symmetry groups! By analysing the interplay between the structure of the Kakeya set at two different scales, Katz–Tao [20] sketched a method for encoding a ring-like object R into K . More precisely, they proposed encoding an *approximate subring* R of \mathbb{R} into K . This is a set which is approximately closed under the addition and multiplication operations, in the sense that sum set $R + R := \{x + y : x, y \in R\}$ and product set $R \cdot R := \{x \cdot y : x, y \in R\}$ have

size comparable to that of R . Furthermore, this approximate subring would have Hausdorff dimension $0 < \dim_H R < 1$.

At this point, the proof sketch is nearing its endgame. We have already mentioned that the Edgar–Miller theorem [8] rules out the existence of subrings of intermediate Hausdorff dimension. This would give the desired contradiction, except we are dealing with *approximate* subrings rather than *bona fide* subrings. However, there is a more quantitative version of the Edgar–Miller theorem called *Bourgain’s sum-product theorem* [2, 3], which does apply in the approximate setting. Using this, Katz–Tao hoped to obtain a contradiction and thereby conclude the proof of the Strong Kakeya Conjecture in the sticky case.

As mentioned earlier, the rigorous proof in [21] differs significantly from that originally envisioned by Katz–Tao [20]. Most notably, the connection with approximate subrings is far less direct. Rather than appeal to Bourgain’s sum-product theorem [3], Wang–Zahl [21] invoke various recent results in *projection theory*, a subfield of geometric measure theory, to arrive at a contradiction. These tools ultimately do rely on Bourgain’s sum-product theorem [3] but, taken as a whole, the scheme of the proof is far more complex than that sketched in [20].

REFERENCES

1. Jonathan Bennett, Anthony Carbery, and Terence Tao, *On the multilinear restriction and Kakeya conjectures*, Acta Math. **196** (2006), no. 2, 261–302. MR 2275834
2. J. Bourgain, *On the Erdős–Volkmann and Katz–Tao ring conjectures*, Geom. Funct. Anal. **13** (2003), no. 2, 334–365. MR 1982147
3. Jean Bourgain, *The discretized sum-product and projection theorems*, J. Anal. Math. **112** (2010), 193–236. MR 2763000
4. Jean Bourgain, Ciprian Demeter, and Larry Guth, *Proof of the main conjecture in Vinogradov’s mean value theorem for degrees higher than three*, Ann. of Math. (2) **184** (2016), no. 2, 633–682. MR 3548534
5. Lennart Carleson, *On convergence and growth of partial sums of Fourier series*, Acta Math. **116** (1966), 135–157. MR 199631
6. Antonio Cordoba, *The Kakeya maximal function and the spherical summation multipliers*, Amer. J. Math. **99** (1977), no. 1, 1–22. MR 447949
7. Roy O. Davies, *Some remarks on the Kakeya problem*, Proc. Cambridge Philos. Soc. **69** (1971), 417–421.
8. G. A. Edgar and Chris Miller, *Borel subrings of the reals*, Proc. Amer. Math. Soc. **131** (2003), no. 4, 1121–1129. MR 1948103
9. K. J. Falconer, *The geometry of fractal sets*, Cambridge Tracts in Mathematics, vol. 85, Cambridge University Press, Cambridge, 1986. MR 867284
10. Charles Fefferman, *The multiplier problem for the ball*, Ann. of Math. (2) **94** (1971), 330–336. MR 296602
11. ———, *On the convergence of multiple Fourier series*, Bull. Amer. Math. Soc. **77** (1971), 744–745. MR 435724
12. Jonathan Hickman, *The Kakeya conjecture*, Snapshots of modern mathematics from Oberwolfach, to appear.
13. Richard A. Hunt, *On the convergence of Fourier series*, Orthogonal Expansions and their Continuous Analogues (Proc. Conf., Edwardsville, Ill., 1967), Southern Illinois Univ. Press, Carbondale, IL, 1968, pp. 235–255. MR 238019
14. Nets Katz and Terence Tao, *Recent progress on the Kakeya conjecture*, Proceedings of the 6th International Conference on Harmonic Analysis and Partial Differential Equations (El Escorial, 2000), 2002, pp. 161–179. MR 1964819
15. Nets Hawk Katz, Izabella Łaba, and Terence Tao, *An improved bound on the Minkowski dimension of Besicovitch sets in \mathbb{R}^3* , Ann. of Math. (2) **152** (2000), no. 2, 383–446. MR 1804528
16. Nets Hawk Katz and Joshua Zahl, *An improved bound on the Hausdorff dimension of Besicovitch sets in \mathbb{R}^3* , J. Amer. Math. Soc. **32** (2019), no. 1, 195–259. MR 3868003

17. Pertti Mattila, *Fourier analysis and Hausdorff dimension*, Cambridge Studies in Advanced Mathematics, vol. 150, Cambridge University Press, Cambridge, 2015. MR 3617376
18. Marcel Riesz, *Sur les fonctions conjuguées*, Math. Z. **27** (1928), no. 1, 218–244. MR 1544909
19. Terence Tao, *From rotating needles to stability of waves: emerging connections between combinatorics, analysis, and PDE*, Notices Amer. Math. Soc. **48** (2001), no. 3, 294–303. MR 1820041
20. ———, *Stickiness, graininess, planiness, and a sum-product approach to the Kakeya problem*, <https://terrytao.wordpress.com/2014/05/07/>, 2014, accessed 10 September 2025.
21. Hong Wang and Joshua Zahl, *Sticky Kakeya sets and the sticky Kakeya conjecture*, Preprint: [arXiv:2210.09581](https://arxiv.org/abs/2210.09581).
22. ———, *Volume estimates for unions of convex sets, and the Kakeya set conjecture in three dimensions*, Preprint: [arXiv:2502.17655](https://arxiv.org/abs/2502.17655).
23. ———, *The Assouad dimension of Kakeya sets in \mathbb{R}^3* , Invent. Math. **241** (2025), no. 1, 153–206. MR 4921573
24. Thomas Wolff, *An improved bound for Kakeya type maximal functions*, Rev. Mat. Iberoamericana **11** (1995), no. 3, 651–674. MR 1363209
25. ———, *Recent work connected with the Kakeya problem*, Prospects in mathematics (Princeton, NJ, 1996), Amer. Math. Soc., Providence, RI, 1999, pp. 129–162. MR 1660476

SCHOOL OF MATHEMATICS AND MAXWELL INSTITUTE FOR MATHEMATICAL SCIENCES, JAMES CLERK MAXWELL BUILDING, THE KING'S BUILDINGS, PETER GUTHRIE TAIT ROAD, EDINBURGH, EH9 3FD, UK.

Email address: jonathan.hickman@ed.ac.uk