I could be bounded in a nut shell and count myself a king of infinite space, were it not that I have bad dreams.

William Shakespeare, *Hamlet*

1 Background and the first version of the Kakeya Conjecture

If you've ever carried a ladder, you know that it's pretty hard to turn corners; the ladder constantly bumps into things, and you need way more space than you might expect. This[1] is what inspired the mathematician Sōichi Kakeya to ask the following question in 1917:

Question (Kakeya). *What is the smallest area of a room in which you can turn a ladder of length* 1 *through a full* $360°$ *rotation?*

The first guess you might have as to the optimal shape is a circle of radius $1/2$, since the point of a circle is that you can turn things around in it. This gives an area of $\pi/4 \approx 0.785$. The next thing you might try is a quarter-circle of radius 1, which unfortunately does no better; it has area $\pi/4$ as well. If you think a bit more about this example, however, you might come up with the idea of using an equilateral triangle with height 1 (and thus side length $2\sqrt{3}/3$), which has area $\sqrt{3}/3 \approx .577$, beating our previous bounds. It was proved by Pál that in fact this equilateral triangle construction is the best we can do if we restrict ourselves to *convex* rooms. However, Kakeya himself found a better construction, the *three-pointed deltoid:*

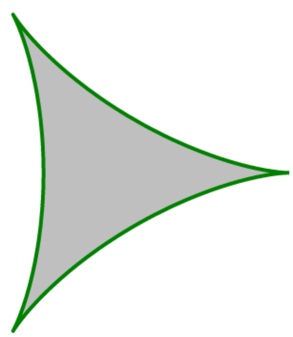

When chosen so that each altitude has length 1, this shape has area $\pi/8 \approx .393$, and Kakeya conjectured that this was the best. We can state this conjecture more formally after making the following definition:

Definition. A *Kakeya set* is a subset $K \subset \mathbb{R}^2$ with the property that a line segment of length 1 can be continuously rotated within K so that it returns to its original position after a rotation of $180°$.

With this definition, we can state the Kakeya Conjecture:

Conjecture (Kakeya, Version 1). *Every Kakeya set has area* $\geq \pi/8$*, and the deltoid is the only Kakeya set with area exactly* $\pi/8$*.*

2 Besicovitch's surprise

As it turns out, the Kakeya Conjecture is false, in a very strong sense. First off, the conjectured lower bound of $\pi/8$ is false; there are Kakeya sets with smaller area. Moreover, it turns out that the entire premise of the conjecture is false: there is *no* minimal area of a Kakeya set, and we can't hope to find some optimal shape like the deltoid!

[1] Strictly speaking, this backstory is totally false.

Theorem (Besicovitch, 1928)**.** *For every $\varepsilon > 0$, there is a Kakeya set $K_\varepsilon \subset \mathbb{R}^2$ with* $\text{area}(K_\varepsilon) \leq \varepsilon$.

We will present a construction due to Perron, which is simpler than Besicovitch's original construction (though most of the ideas are the same). We will do this in several steps.

First of all, let's forget about rotating a full 180°, and start by rotating 60°. If we can do this in an arbitrarily small area, then by taking three rotated copies of area $\leq \varepsilon/3$, we will be able to do the full 180° rotation. The basic idea of the construction is to start with an equilateral triangle of height 1, where we can definitely do a 60° rotation. We cut this triangle in two vertically, then translate the pieces so that they overlap:

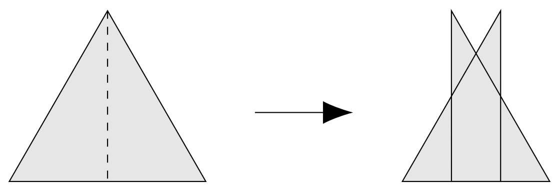

This new shape has a smaller area than the original triangle, and the goal is now to iterate this construction. However, there is one glaring problem with this trick: even though the new shape contains lines in all the directions that the original shape did, we can no longer continuously rotate from one to the other! It *almost* works, but it requires us to teleport from one vertical segment to the other (these are the two copies of the segment that we cut along; cutting and shifting does not preserve continuity).

To do this teleportation, we use a trick known as the *Pál join:*

Lemma (Pál)**.** *Let $\varepsilon > 0$. Let L_1, L_2 be two parallel lines, and let S be a unit segment on L_1. Then we can continuously move S onto L_2 using an area $\leq \varepsilon$.*

Proof. The first thing we might try is to just translate S diagonally:

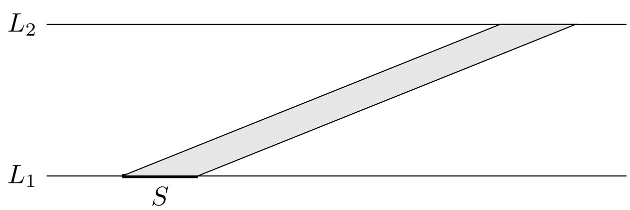

However, no matter how slanted we make this shift, we will always be tracing out a parallelogram whose base is S and whose height is the distance between L_1 and L_2, and this its area will always be the same.

So we need to be a bit more clever, and Pál's trick was to observe that we use up no area when we translate S in the direction it's pointing. So we will rotate S a tiny amount, then translate it until it hits L_2, and finally undo the rotation. We do this in such a way that the area of the sector it sweeps out is $\leq \varepsilon/2$, and thus the total area used is $\leq \varepsilon$:

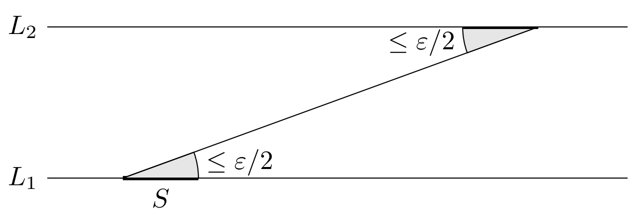

□

So to return to the construction of a small Kakeya set, by paying a very small amount for a bit of extra area (as small as we want, in fact), we can solve this teleportation problem: by possibly extending the vertical lines and adding such a Pál join, we can indeed rotate a segment 60° within our smaller chopped-up triangle.

With this difficulty out of the way, let's think about how to iterate the chopped-up triangle construction in order to minimize the area we use. As an illustration, suppose chop up our equilateral triangle into 8 subtriangles:

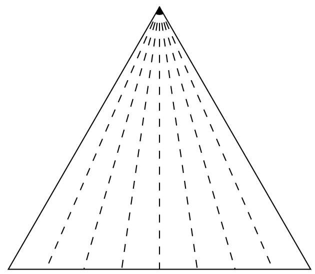

For each adjacent pair, we translate them to overlap, as before:

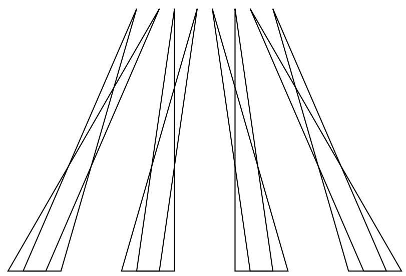

Now, we overlap pairs of these:

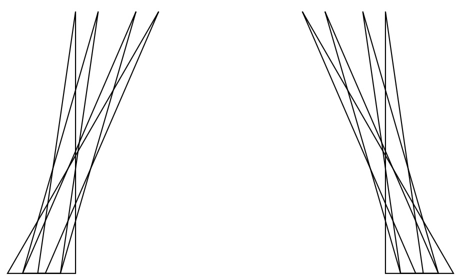

And finally, we overlap these two to get a single figure (with internal lines removed so that we can see what's going on):

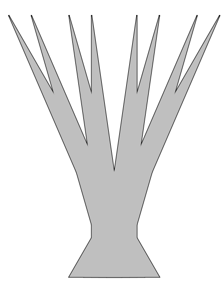

In this case, since we made seven cuts, we'll also need to add seven Pál joins to allow us to teleport between parallel segments. The general idea is that rather than cutting our triangle into 8 subtriangles, we will cut it into 2^k subtriangles for some large k, and shift them so that we can get the area arbitrarily small. For this, we need the following technical lemma:

Lemma. *Suppose T is a triangle, and we drop a median from the apex so that we end up with two subtriangles T_1, T_2, each of base b. Fix some parameter $\frac{1}{2} < \alpha < 1$, and suppose we translate T_2 to the left by a distance of $2b(1-\alpha)$, so that the resulting figure has a base of length*

$$2b - 2b(1-\alpha) = 2b\alpha$$

Then the resulting figure has area $(\alpha^2 + 2(1-\alpha)^2)\,\mathrm{area}(T)$.

Proof. We cut the new figure into shapes as follows:

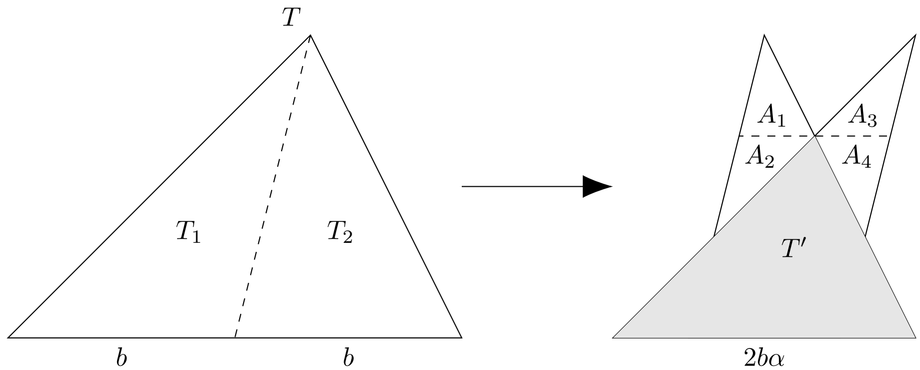

Then T' is similar to T, and the ratio of the side lengths is α, so $\mathrm{area}(T') = \alpha^2\,\mathrm{area}(T)$. Moreover, since the dashed line above is parallel to the base of T, we deduce that A_1 and A_4 are similar to T_2, while A_2, A_3 are similar to T_1. In all cases, the ratio of side lengths is $1-\alpha$, since the total length of the dashed segment is $2(1-\alpha)b$, the amount we translated T_2 by. Therefore,

$$\mathrm{area}(A_2) = \mathrm{area}(A_3) = (1-\alpha)^2\,\mathrm{area}(T_1) = \frac{1}{2}(1-\alpha)^2\,\mathrm{area}(T)$$
$$\mathrm{area}(A_1) = \mathrm{area}(A_4) = (1-\alpha)^2\,\mathrm{area}(T_2) = \frac{1}{2}(1-\alpha)^2\,\mathrm{area}(T)$$

Putting this all together, we see that the area of the new figure is

$$\alpha^2\,\mathrm{area}(T) + 4\left(\frac{1}{2}(1-\alpha)^2\,\mathrm{area}(T)\right) = (\alpha^2 + 2(1-\alpha)^2)\,\mathrm{area}(T)$$

□

With this lemma, we can prove Besicovitch's Theorem.

Proof. Fix some parameters $\frac{1}{2} < \alpha < 1$ and $k \in \mathbb{N}$, which we will pick later. Let T be an equilateral triangle, and divide it into 2^k subtriangles, $T_1, \ldots, T_{2^k}$. We split these into pairs T_{2i-1}, T_{2i} for $1 \le i \le 2^{k-1}$, and translate each T_{2i} towards T_{2i-1} by a fraction $1-\alpha$, as in the previous lemma. This yields a new shape, which we call $S_i^{(1)}$, consisting of a "heart" triangle $T_i^{(1)}$ that is similar to $T_{2i-1} \cup T_{2i}$, and two additional "ear" triangles $E_{2i-1}^{(1)}, E_{2i}^{(1)}$. By the previous lemma, we know that

$$\mathrm{area}(S_i^{(1)}) = (\alpha^2 + 2(1-\alpha)^2)\,\mathrm{area}(T_{2i-1} \cup T_{2i})$$

Now, we iterate this. For $1 \le j \le 2^{k-2}$, we translate $S_{2j}^{(1)}$ towards $S_{2j-1}^{(1)}$, and we apply the lemma to the hearts of $S_{2j-1}^{(1)}, S_{2j}^{(1)}$, which are triangles. Then the new figure we get, $S_j^{(2)}$, consists of the old ears, which had total area $2(1-\alpha)^2\,\mathrm{area}(T_{4j-3} \cup T_{4j-2} \cup T_{4j-1} \cup T_{4j})$, plus a new heart and new ears, whose total area is

$$(\alpha^2 + 2(1-\alpha)^2)\,\mathrm{area}(\heartsuit(S_{2j}^{(1)})) = (\alpha^4 + 2\alpha^2(1-\alpha)^2)\,\mathrm{area}(T_{4j-3} \cup T_{4j-2} \cup T_{4j-1} \cup T_{4j})$$

since $\heartsuit(S_{2j}^{(1)}) = \alpha^2 \operatorname{area}(T_{4j-3} \cup T_{4j-2} \cup T_{4j-1} \cup T_{4j})$. Putting this together and summing over all j, we find that

$$\sum_{j=1}^{2^{k-2}} \operatorname{area}(S_j^{(2)}) \leq (\alpha^4 + 2\alpha^2(1-\alpha)^2 + 2(1-\alpha)^2) \operatorname{area}(T)$$

where the $\leq$ comes from the fact that there might be some additional overlap between the old and the new ears that we're not taking into account.

After we've done this r times, then we need to apply the lemma to the heart of the $(r-1)$st iteration, whose area is $\alpha^{2r-2} \operatorname{area}(T_{j2^r+1} \cup \cdots \cup T_{(j+1)2^r})$. Thus, we get a new heart of area $\alpha^{2r} \operatorname{area}(T_{j2^r+1} \cup \cdots \cup T_{(j+1)2^r})$, and new ears of total area $2\alpha^{2r-2}(1-\alpha)^2 \operatorname{area}(T_{j2^r+1} \cup \cdots \cup T_{(j+1)2^r})$, and therefore

$$\sum_{j=1}^{2^{k-r}} \operatorname{area}(S_j^{(r)}) \leq \left(\alpha^{2r} + 2(1-\alpha)^2 \sum_{m=0}^{r-1} \alpha^{2m}\right) \operatorname{area}(T)$$

Thus, at the end of the day, after we do this a total of k times, we find that the final figure S satisfies

$$\begin{aligned}
\operatorname{area}(S) &\leq \left(\alpha^{2k} + 2(1-\alpha)^2 \sum_{m=0}^{k-1} \alpha^{2m}\right) \operatorname{area}(T) \\
&\leq \left(\alpha^{2k} + 2(1-\alpha)^2 \sum_{m=0}^{\infty} \alpha^{2m}\right) \operatorname{area}(T) \\
&= \left(\alpha^{2k} + \frac{2(1-\alpha)^2}{1-\alpha^2}\right) \operatorname{area}(T) \\
&= \left(\alpha^{2k} + \frac{2(1-\alpha)}{1+\alpha}\right) \operatorname{area}(T) \\
&\leq (\alpha^{2k} + 2(1-\alpha)) \operatorname{area}(T)
\end{aligned}$$

Now, we first choose α sufficiently close to 1 so that $2(1-\alpha)\operatorname{area}(T) \leq \varepsilon/16$. Then, we pick k large enough that $\alpha^{2k} \operatorname{area}(T) \leq \varepsilon/16$. So we get that $\operatorname{area}(S) \leq \varepsilon/8$. By adding $2^k - 1$ Pál joins, each of area $\leq \varepsilon/2^{k+3}$, we can turn S into a set S' of area $\leq \varepsilon/4$ where we can rotate a segment $60°$. Finally, we form two more rotated copies of S' and add three new Pál joins each of area $\leq \varepsilon/12$ to connect these rotated copies, and we end up with a Kakeya set K_ε of area $\leq \varepsilon$. □

3 Can we do better?

We found that there are Kakeya sets of arbitrarily small area. But can we do better?

First of all, we should understand what it means to do "better" than a set of arbitrarily small area.

Definition. A set $Z \subset \mathbb{R}^2$ is said to have *zero area* if for every $\varepsilon > 0$, there is an open set U_ε with $Z \subset U_\varepsilon$ and $\operatorname{area}(U_\varepsilon) \leq \varepsilon$.

(Recall that U is an open set if for every $x \in U$ and any y sufficiently close to x, y is also in U.)

So we can rephrase the question above as follows: do there exist Kakeya sets with zero area?

The answer, perhaps disappointingly, is no. A fully formal proof is actually a bit tricky to produce, but the intuition is as follows: when we rotate the needle by some infinitesimal amount inside the set K, some tiny sector of a circle is swept out, so K contains a tiny sector. Since this sector has positive (though tiny) area, K must have positive area as well.

However, we *can* do better if we slightly weaken our notion of a Kakeya set.

Definition. A subset $B \subset \mathbb{R}^2$ is called a *Besicovitch set* if B contains a unit line segment in every direction (with no assumption about being able to continuously rotate).

Since every Kakeya set is a Besicovitch set, we already know that there exist Besicovitch sets of arbitrarily small area. In fact, the proof is even simpler: if you think back to what we did above, you'll realize that we were just constructing a Besicovitch set of arbitrarily small area, and then adding to it a bunch of Pál joins of small area to turn it into a Kakeya set. Moreover, with this new notion, we *can* do better:

Theorem (Besicovitch, 1919)**.** *There is a Besicovitch set $B \subset \mathbb{R}^2$ with zero area.*

To prove this, we first need a simple lemma.

Lemma. *If T is a triangle and $U \supset T$ is an open set, then we can cut T into subtriangles and translate them to form a set S of arbitrarily small area so that $S \subset U$.*

Proof. This lemma simply says that the above Perron tree construction allows us to stay pretty close to the original triangle T. To prove this, let T have base b. Divide T into subtriangles and fix T_1 so that it never moves during the whole process. Then every other triangle will move a distance at most b.

Now, pick some $\varepsilon > 0$ so that every point at distance $\leq \varepsilon$ from T is contained in U; such an ε exists since U is open. Then divide T into $\lceil 1/\varepsilon \rceil$ subtriangles, each of base $\leq \varepsilon$. Now, perform the Perron tree construction independently for each of these subtriangles. By the previous paragraph, in doing so, no point will move more than ε away from T. Thus, in the end, we get a set S all of whose points are distance $\leq \varepsilon$ from T, so $S \subset U$. □

With this lemma, we can prove Besicovitch's Theorem:

Proof. Fix an equilateral triangle T, and an open set $U_1 \supset T$ with $\operatorname{area}(U_1) \leq 2 \operatorname{area}(T)$. By chopping up and translating T, we form a new set S_1 of area $\leq \frac{1}{2}$, and by the lemma we can ensure that $S_1 \subset U_1$. Now pick a new open set $U_2 \supset S_1$ such that $\operatorname{area}(U_2) \leq 2 \operatorname{area}(S_1)$. Since S_1 is a union of triangles, we can apply the previous lemma to each such triangle and get a new set S_2 of area $\leq \frac{1}{4}$, such that $S_2 \subset U_2$. Again, we pick a new open set $U_3 \supset S_2$ with $\operatorname{area}(U_3) \leq 2 \operatorname{area}(S_2)$, and iterate this.

Notice that when we do this, we get a sequence of sets S_i with $\operatorname{area}(S_i) \leq 2^{-i}$, and a nested sequence of set $U_1 \supset U_2 \supset U_3 \supset \cdots$, with

$$\operatorname{area}(U_i) \leq 2 \operatorname{area}(S_{i-1}) \leq 2^{-i}$$

Therefore, if we set $B = \bigcap_{i=1}^{\infty} U_i$, then B will automatically have measure zero, since it is covered by open sets of arbitrarily small area. Moreover, since each U_i contains S_i, which contains a segment in every direction, one can prove that B will also contain a line segment in every direction (proving this fully formally is a bit subtle and requires some notions like compactness that I don't want to get into, though hopefully the intuition is clear). □

Finally, let's end this section with a more general notion of Besicovitch set:

Definition. Let $n \in \mathbb{N}$. A set $B \subset \mathbb{R}^n$ is called a Besicovitch set if it contains a line segment in every direction.

4 Can we do even better?

We now know that there is a Besicovitch set of zero area in $\mathbb{R}^2$. But can we do even better?

Again, we need to make clear what we mean by this question. We certainly can't do any better when it comes to area—zero is the smallest it can get. However, it turns out that not all zero-area sets are created equal. As an illustrative example, think about the difference between a point and a line segment. They both have zero area, but we still want to think of the segment as "bigger" than the point. If we now work inside $\mathbb{R}^3$, then a point, a line segment, and a flat square all have zero volume, but again, there is a sense in which the square is the "biggest" of the three.

As you may have guessed, the difference characterizing the above examples is *dimension*, which is the most important way of differentiating sets of zero area (or volume).

The notion of dimension that we will be considering is called *Minkowski dimension.* It is not the "best" notion of dimension (that honor goes to the so-called Hausdorff dimension), but it is the easiest to define.

Let S be a bounded set in $\mathbb{R}^n$, and let $N(S, \varepsilon)$ be the minimal number of balls of radius ε needed to cover S. For instance, if S is a line segment of length L, then $N(S, \varepsilon) \approx L/2\varepsilon$, whereas if T is a square of area A, then $N(S, \varepsilon) \approx A/\pi\varepsilon^2$. Moreover, both of these calculations continue to hold even if S is a curve or T is a curved surface, so long as ε is small enough. These examples suggest that

$$N(S, \varepsilon) \approx \frac{C}{\varepsilon^d}$$

where C is some constant and d is the dimension of S. Taking logarithms gives that $\log N(S, \varepsilon) \approx \log C - d \log \varepsilon$, and rearranging shows that

$$d \approx \frac{\log C}{\log \varepsilon} - \frac{\log N(S, \varepsilon)}{\log \varepsilon}$$

We expect this approximation to get better and better as ε gets smaller. Moreover, as ε becomes small, the first term above tends to zero, since C is a constant. All of this motivates the following definition:

Definition. For a bounded set $S \subset \mathbb{R}^n$, its *Minkowski dimension* is defined by

$$\dim_M S = \lim_{\varepsilon \to 0} \frac{\log N(S, \varepsilon)}{-\log \varepsilon}$$

assuming this limit exists.

Example. The Sierpinski gasket is a fractal gotten by starting with an equilateral triangle, removing its middle triangle, and then iterating this on every smaller subtriangle.

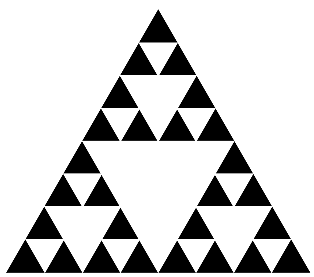

For each n, we can cover the nth stage of the construction with 3^n balls (or triangles), each of radius $2^{-n}/\sqrt{3}$, since each triangle in the nth stage has side length 2^{-n}. So we find that

$$N(S, 2^{-n}/\sqrt{3}) = 3^n$$

and thus, assuming that the limit exists (which it does, though it's not so easy to prove), we find that

$$\dim_M S = \lim_{n \to \infty} \frac{\log(3^n)}{-\log(2^{-n}/\sqrt{3})} = \lim_{n \to \infty} \frac{n \log 3}{n \log 2 + \log \sqrt{3}} = \frac{\log 3}{\log 2} = \log_2 3$$

In particular, the Sierpinski gasket, like many fractals, has a non-integer dimension.

With this in hand, we can state the modern version of the Kakeya Conjecture:

Conjecture (Kakeya, Version 2)**.** *Any Besicovitch set $B \subset \mathbb{R}^n$ has $\dim_M B = n$.*

An equivalent definition of Minkowski dimension is as follows. Suppose we partition $\mathbb{R}^n$ into cubes of side length ε by drawing a very fine lattice; let $N'(S, \varepsilon)$ be the number of these cubes that intersect S.

Proposition.

$$\dim_M S = \lim_{\varepsilon \to 0} \frac{\log N'(S, \varepsilon)}{-\log \varepsilon}$$

In fact, many similar things you could think of will also give the same notion of dimension. You proved this proposition on the homework.

Why might we expect the Kakeya Conjecture to be true? My intuition for it is that the prototypical example of a set that does not have dimension n is a hyperplane, which does not contain lines in a large number of directions; so since a Besicovitch set can't "look" like a hyperplane, it must have full dimension. Of course, this intuition is at the very least sketchy, and at worst might be actively misleading: there are many other low-dimensional sets that look really weird, and in particular look nothing like a hyperplane. For instance, on the homework you saw the Cantor dust, a subset of $\mathbb{R}^2$ with dimension 1 that is very dissimilar from a line.

The Kakeya Conjecture is considered the most important open problem in the field of Geometric Measure Theory, and is one of the biggest open problems in all of analysis. Relatively few partial results are known. The conjecture is proven in the case $n = 1$ (which is simple, since a Besicovitch set in $\mathbb{R}^1$ must contain a line segment, and thus automatically has dimension 1) and in the case $n = 2$, due to Davies in 1971. For all $n \geq 3$, it is unknown; Bourgain proved that every Besicovitch set has dimension at least $(n+1)/2$, Wolff improved this to $(n+2)/2$, and Katz-Tao improved this to $(2-\sqrt{2})(n-4)+3$, which is a better bound when $n \geq 5$. Apart from a few other special cases, this is basically the state of the art.

5 Finite Fields

Recall that in the integers mod n, we can add, subtract, and multiply, and many of the properties we expect these operations to have indeed occur. However, when n is not prime, some strange things can occur; for instance, we can have two non-zero elements that multiply to give zero. But if p is a prime, then the integers mod p form a *field,* which is an algebraic structure where we can add, subtract, multiply, and divide, and everything works the way we expect. Fields you are more familiar with include $\mathbb{Q}, \mathbb{R}$, and $\mathbb{C}$; as it turns out, there are also other finite fields in addition to the integers mod p. If you want to learn more about what these fields are, take Aaron's class next week, but for now, we will simply denote by $\mathbb{F}_q$ a finite field of order q; you're welcome to think of q as a prime and $\mathbb{F}_q$ simply being the integers mod q.

Because $\mathbb{F}_q$ is a field, we can do geometry over it. Specifically, let $\mathbb{F}_q^n$ denote the collection of n-tuples of elements of $\mathbb{F}_q$. Then, for instance, a *line* in $\mathbb{F}_q^n$, in the direction of some $m \in \mathbb{F}_q^n$ and going through some $b \in \mathbb{F}_q^n$, is simply the set

$$\ell_{m,b} = \{b + t \cdot m : t \in \mathbb{F}_q\} \subseteq \mathbb{F}_q^n$$

where the multiplication $t \cdot m$ means that we multiply each coordinate of m by the scalar t, and the addition $b + t \cdot m$ is also done componentwise. Because $\mathbb{F}_q$ is a field, lines in $\mathbb{F}_q^n$ work basically like lines in $\mathbb{R}^n$; for instance, any two points define a line, and any two lines intersect in at most one point. We can also make the following definition:

Definition. A set $B \subseteq \mathbb{F}_q^n$ is called a *Besicovitch set* if it contains a full line in every direction. In other words, B is a Besicovitch set if for all $m \in \mathbb{F}_q^n$, there is some $b \in \mathbb{F}_q^n$ so that $\ell_{m,b} \subseteq B$.

There are (at least) two important things to note about this definition. First, we can't hope to define a Kakeya set in $\mathbb{F}_q^n$, since there is no notion of "continuously rotating" when we are working with a discrete space like $\mathbb{F}_q^n$. Second, observe that we require a full line in every direction, rather than a unit segment; this too follows from the fact that there is no good notion of "length" or "scaling" over $\mathbb{F}_q$.

Suppose we partition the interval $[0, 1]$ into q subintervals of equal length, and use this to define a partition of $[0, 1]^n$ into q^n boxes. If we put a point in the center of each box, then we can pretend that $\mathbb{F}_q^n$ is a discrete approximation to $[0,1]^n$, and we can guess that this approximation gets better and better as q gets larger. In particular, we might hope that as q gets larger, a Besicovitch set in $\mathbb{F}_q^n$ looks more and more like a Besicovitch set in $\mathbb{R}^n$. Recall that $N'(S,\varepsilon)$ denotes the number of boxes of side-length ε that intersect a set S, and that the Minkowski dimension of S is given by

$$\dim_M S = \lim_{\varepsilon\to 0} \frac{\log N'(S,\varepsilon)}{\log(1/\varepsilon)}$$

If we think of S as coming from a subset of $\mathbb{F}_q^n$, again imagined as living inside $[0,1]^n$, then we have that $N'(S,1/q) = |S|$, since the number of boxes intersecting S is precisely the number of elements of $\mathbb{F}_q^n$ in S. Therefore, if we believe the Kakeya conjecture, that $\dim_M B = n$ for any Besicovitch set $B \subseteq \mathbb{R}^n$, we might hope that something like the following holds:

$$n = \dim_M B = \lim_{q\to\infty} \frac{\log |B_q|}{\log q}$$

where $B_q \subseteq \mathbb{F}_q^n$ is a Besicovitch set in $\mathbb{F}_q^n$. Rearranging this gives us the following guess:

Conjecture (Finite Field Kakeya Conjecture)**.** *For every n, there is some constant C_n so that for any q and any Besicovitch set $B_q \subseteq \mathbb{F}_q^n$, we have*

$$|B_q| \geq C_n q^n$$

Note that if this conjecture is true, then we indeed have that

$$\lim_{q\to\infty} \frac{\log |B_q|}{\log q} \geq \lim_{q\to\infty} \left(\frac{\log C_n}{\log q} + \frac{n \log q}{\log q} \right) = n$$

as our heuristic argument above suggested. This Finite Field Kakeya Conjecture was first conjectured by Wolff in 1999, and his idea was that it might serve as another regime where ideas for the real Kakeya conjecture could be tested. There is no formal reduction from one conjecture to the other (all our arguments above were purely heuristic, and cannot be turned into proofs), but the hope was that understanding one problem would help us understand the other.

Indeed, for many years this worked. Wolff himself extended Davies' ideas to prove the Finite Field Kakeya Conjecture in dimension $n = 2$, along with several partial results for higher dimensions. For about a decade, any time someone made an advance towards solving either the Kakeya Conjecture or the Finite Field Kakeya Conjecture, some work very quickly followed that got the same result for the other conjecture. Moreover, the relationship between these two conjectures allowed new ideas to come into play.

However, this all changed in 2008, when Zeev Dvir shocked everyone and proved the full Finite Field Kakeya Conjecture. More precisely, he showed

Theorem (Dvir)**.** *For every n, q, and every Besicovitch set $B \subseteq \mathbb{F}_q^n$, we have*

$$|B| \geq \binom{q+n-1}{n} \geq \frac{1}{n!} q^n$$

Thus, the Finite Field Kakeya Conjecture is true with $C_n = 1/n!$.

This result was later improved by Dvir, Kopparty, Saraf, and Sudan, who showed that in fact we can take $C_n \approx 2^{-n}$.

6 A digression on polynomials

Dvir's remarkably simple proof uses the so-called "polynomial method," which is not really even a method; it is simply the observation that polynomials are weirdly useful for proving lots of difficult-seeming results. The basic idea that went into Dvir's proof is the following heuristic: "a set is small if and only if there is a polynomial of low degree that is identically zero on it."

One way that this heuristic plays out is in the following well-known theorem:

Theorem (Factor Theorem)**.** *Let $\mathbb{F}$ be any field, and let $P(x)$ be a non-zero polynomial over $\mathbb{F}$ with degree d. Then P has at most d roots, i.e. there are at most d values of $a \in \mathbb{F}$ so that $P(a) = 0$.*

Proof. We prove this by induction on d. The base case is when $d = 0$, which means that P must be a constant polynomial. Since we assumed that P was a non-zero polynomial, this constant value cannot be 0, so P has no roots, and the base case is true.

For the inductive step, suppose that any polynomial of degree $d-1$ has at most $d-1$ roots, and let P be a non-zero polynomial of degree d. If P has no roots, then we are done. If not, then there is some $a \in \mathbb{F}$ with $P(a) = 0$. By polynomial long division, we may write

$$P(x) = (x-a)Q(x) + R(x)$$

where $\deg R \leq \deg(x-a) = 1$. This means that R is a degree-zero polynomial, so it's just a constant. On the other hand, plugging in $x = a$ to both sides shows us that

$$0 = P(a) = (a-a)Q(a) + R(a) = R(a)$$

So $R(a) = 0$, so in fact R must be identically zero. Thus, we've written $P(x) = (x-a)Q(x)$, where $\deg Q = d-1$. So by the inductive hypothesis, Q has at most $d-1$ roots; adding back to this the root a of P, we find that P has at most d roots, as desired. □

We also have a sort of converse to the factor theorem:

Theorem. *Let $\mathbb{F}$ be any field, and let $S \subseteq \mathbb{F}$ have $|S| = d$. Then there is a non-zero polynomial of degree d that vanishes identically on S.*

Proof. We can simply define the polynomial to be

$$P(x) = \prod_{a \in S} (x-a)$$

Then since $|S| = d$, we find that P is a product of d linear factors, so $\deg P = d$. Moreover, P is a non-zero polynomial, since it has a non-zero leading coefficient (namely, its first term is just x^d). Finally, P indeed vanishes on S, since if we plug in some $a \in S$, then the right-hand side will be zero, so $P(a) = 0$. □

Thus, we find that the heuristic mentioned above is quite precise when we are dealing with subsets of a field $\mathbb{F}$: a set has size at most d if and only if there is a non-zero degree-d polynomial that vanishes on it. However, we are going to be working with subsets of $\mathbb{F}_q^n$, so we first need to generalize these results to work in this higher-dimensional setting. This presents some difficulties—for instance, we can't hope that a polynomial in many variables will have some bounded number of roots, since e.g. the polynomial $P(x, y) = x + y$ has infinitely many roots in $\mathbb{R}^2$, despite having degree 1.

Theorem (Schwartz-Zippel)**.** *Let $\mathbb{F}$ be any feild, let $P(x_1, \ldots, x_n)$ be a non-zero polynomial in n variables and total degree d, and let $S \subseteq \mathbb{F}$ be a set with $|S| > d$. Then there are some $a_1, \ldots, a_n \in S$ so that $P(a_1, \ldots, a_n) \neq 0$. In other words, even though P can have many roots, it can't have too many in any product set $S \times \cdots \times S \subseteq \mathbb{F}^n$.*

Proof. We proceed by induction on n. The base case is $n = 1$, where we need to show that P has some non-root in S. But since $|S| > \deg P$, the Factor Theorem we proved above guarantees exactly that. For the inductive case, we expand P as a polynomial in x_n, writing

$$P(x_1, \ldots, x_n) = \sum_{j=0}^{d} x_n^j P_j(x_1, \ldots, x_{n-1})$$

where each P_j is some polynomial in the variables $x_1, \ldots, x_{n-1}$. We can always write down such a decomposition, by simply writing down all the monomials of P, arranging them by the exponent of x_n in each monomial, and then factoring out that power of x_n. Since P is a non-zero polynomial, then there is some i for which P_i is a non-zero polynomial. Since $\deg P_i \leq \deg P = d$, the inductive hypothesis guarantees that there are some $a_1, \ldots, a_{n-1} \in S$ so that $P_i(a_1, \ldots, a_n) \neq 0$. Now, let $Q(x)$ be the single-variate polynomial given by plugging $a_1, \ldots, a_{n-1}$ into P, namely

$$Q(x) = P(a_1, \ldots, a_{n-1}, x)$$

If we write $b_j = P_j(a_1, \ldots, a_{n-1})$, then we have that

$$Q(x) = \sum_{j=0}^{d} b_j x^j$$

Since we know that $b_i \neq 0$, we find that Q is a non-zero polynomial. So by the Factor Theorem, we know that there is some $a_n \in S$ so that $Q(a_n) \neq 0$. But that precisely means that $P(a_1, \ldots, a_n) \neq 0$, as desired. □

The next subtlety in passing to more variables is in producing the lower bound, namely proving that every "small" set has a low-degree polynomial that vanishes on it. The precise statement is as follows.

Proposition. *Let $\mathbb{F}$ be any field and $d, n \in \mathbb{N}$, and let $T \subseteq \mathbb{F}^n$ be any set so that $|T| < \binom{d+n}{n}$. Then there is a non-zero polynomial $P(x_1, \ldots, x_n)$ of degree at most d such that P vanishes on T, namely $P(a_1, \ldots, a_n) = 0$ for any $(a_1, \ldots, a_n) \in T$.*

Proof. First, note that we may write any polynomial P of degree at most d as a sum of monomials, namely

$$P(x_1, \ldots, x_n) = \sum_{\substack{e_1, \ldots, e_n \geq 0 \\ e_1 + \cdots + e_n \leq d}} c_{e_1, \ldots, e_n} x_1^{e_1} \cdots x_n^{e_n}$$

How many such monomials are there? That is the same as asking for the number of n-tuples $(e_1, \ldots, e_n)$ with $e_i \geq 0$ and $\sum e_i \leq d$. The standard technique for counting such things is the "balls and bins" method: imagine we have n bins, corresponding to $e_1, \ldots, e_n$, and we are tossing into them at most d balls, corresponding to the value of e_i. For convenience, we add one more bin, which will be the trash can: we are now throwing in *exactly* d balls, and all the ones that land in the trash can will not be assigned to any e_i.

To count how many ways we can throw d balls into $n + 1$ bins, we draw a picture like this (for five balls and four bins):

$$\circ \quad | \quad \circ \quad | \quad \circ \quad \circ \quad \circ \quad |$$

The vertical bars denote the dividers between the bins, and the circles are the balls. In this case, the first two bins each have one ball, the third has three, and the fourth has zero. As this example shows, the number of ways of putting d balls into $n+1$ bins is the same as the number of ways of ordering d circles and n bars, which is precisely $\binom{d+n}{n}$.

So what we find is that there are $\binom{d+n}{n}$ possible monomials in a polynomial of degree at most d. We will find a polynomial P that vanishes on T as follows. For each point $(a_1, \ldots, a_n) \in T$, we want $P(a_1, \ldots, a_n) = 0$, which gives us

$$\sum_{e_1, \ldots, e_n} c_{e_1, \ldots, e_n} a_1^{e_1} \cdots a_n^{e_n} = 0$$

For each fixed $(a_1, \dots, a_n)$, this is a linear equation in the unknown quantities $c_{e_1,\dots,e_n}$. We get such a linear equation for each point of T, thus giving us a homogeneous system of $|T|$ linear equations in $\binom{d+n}{n}$ unknowns. Since we assumed that $|T| < \binom{d+n}{n}$, this system is underdetermined, so it has a non-zero solution. This solution precisely defines the coefficients of a polynomial that vanishes on T. □

So as we see, the situation in higher dimensions is a bit more complicated than the one-dimensional case, but the basic heuristic is indeed true (when interpreted correctly): a set in $\mathbb{F}^n$ is "small" if and only if there is a low-degree polynomial that vanishes on it.

7 Back to Kakeya

With these technical tools out of the way, we are able to prove the Finite Field Kakeya Conjecture. Recall the statement we want to prove:

Theorem (Dvir)**.** *For every n, q, and every Besicovitch set $B \subseteq \mathbb{F}_q^n$, we have*

$$|B| \geq \binom{q+n-1}{n}$$

Proof. Suppose for contradiction that we had a Besicovitch set $B \subseteq \mathbb{F}_q^n$ with

$$|B| < \binom{q+n-1}{n}$$

By the last proposition we proved, this implies that there is some non-zero polynomial $P(x_1, \dots, x_n)$ with coefficients in $\mathbb{F}_q$ and degree at most $q-1$ with the property that P vanishes on B. Suppose $\deg P = d \leq q-1$, and write

$$P = \sum_{i=0}^{d} P_i$$

where each P_i is a homogeneous polynomial of degree i; in other words, we simply group together the monomials of P by their degree. Since $\deg P = d$, we know that P_d is not the zero polynomial (for otherwise the degree would be strictly smaller).

For any $0 \neq m \in \mathbb{F}_q^n$, we know that B contains a line in the direction of m, namely there is some $b \in \mathbb{F}_q^n$ so that $\ell_{m,b} \subseteq B$, where

$$\ell_{m,b} = \{b + t \cdot m : t \in \mathbb{F}_q\}$$

Define a new single-variate polynomial $Q_m(t)$ by

$$Q_m(t) = P(b + t \cdot m)$$

Since we are just plugging in values to P, we find that $\deg Q_m \leq \deg P \leq q-1$. On the other hand, for any value of t, we have that $b + t \cdot m \in B$, so $P(b + t \cdot m) = 0$. Thus, $Q_m(t) = 0$ for every $t \in \mathbb{F}_q$, so Q_m has at least q roots (this is where we use that $m \neq 0$, so that distinct values of t give us distinct values of $b + t \cdot m$). Since $\deg Q_m \leq q-1$, by the Factor Theorem, this implies that Q_m is the zero polynomial. Thus, in particular, the coefficient of t^d in $Q_m(t)$ is zero. However, the coefficient of t^d in $Q_m(t)$ is precisely the value of $P_d(m)$. So we find that $P_d(m) = 0$ for every $0 \neq m \in \mathbb{F}_q^n$. Moreover, since P_d is homogeneous of degree d, this implies that in fact, P_d vanishes on all of $\mathbb{F}_q^n$. Finally, since $d < q = |\mathbb{F}_q|$, Schwartz-Zippel implies that P_d must in fact be the zero polynomial. This is a contradiction. □

8 Kakeya in dimension 2

Today, we will prove the Kakeya Conjecture in dimension 2. This was originally proved by Davies in 1971; his proof is not really the one we will do today, though the ideas are more or less the same. This proof is basically due to Bourgain.

First, we need yet another definition of Minkowski dimension.

Definition. Let $S \subset \mathbb{R}^n$ be a bounded subset, and let $\varepsilon > 0$. By $N_\varepsilon(S)$ we denote the ε-neighborhood of S, namely all points that are at distance $\leq \varepsilon$ away from S. Formally,

$$N_\varepsilon(S) = \{x \in \mathbb{R}^n : \exists y \in S \text{ such that } d(x, y) \leq \varepsilon\}$$

where $d(x, y)$ is the Euclidean distance between x and y. Finally, let

$$V(S, \varepsilon) = \mathrm{vol}(N_\varepsilon(S))$$

Proposition. *The Minkowski dimension of S is given by*

$$\dim_M S = \lim_{\varepsilon \to 0} \left(n + \frac{\log V(S, \varepsilon)}{\log \frac{1}{\varepsilon}} \right)$$

Proof-ish. If we cover S by balls of radius ε, then we can also cover the ε-neighborhood by balls of radius 2ε, centered at the same points. This implies that

$$V(S, \varepsilon) \approx \mathrm{vol}(B_{2\varepsilon}) \cdot N(S, \varepsilon)$$

where $B_{2\varepsilon}$ is a ball of radius 2ε. This implies that

$$\begin{aligned}
\dim_M S &= \lim_{\varepsilon \to 0} \frac{\log N(S, \varepsilon)}{-\log \varepsilon} \\
&\approx \lim_{\varepsilon \to 0} \frac{\log V(S, \varepsilon) - \log \mathrm{vol}(B_{2\varepsilon})}{-\log \varepsilon} \\
&= \lim_{\varepsilon \to 0} \left(\frac{\log(\nu 2^n \varepsilon^n)}{\log \varepsilon} - \frac{\log V(S, \varepsilon)}{\log \varepsilon} \right) \\
&= \lim_{\varepsilon \to 0} \left(\frac{\log(\nu 2^n)}{\log \varepsilon} + \frac{n \log \varepsilon}{\log \varepsilon} - \frac{\log V(S, \varepsilon)}{\log \varepsilon} \right) \\
&= \lim_{\varepsilon \to 0} \left(n + \frac{\log V(S, \varepsilon)}{\log \frac{1}{\varepsilon}} \right)
\end{aligned}$$

where ν is the volume of a ball of radius 1 in $\mathbb{R}^n$, which is just some constant (depending on n). □

What this proposition implies is that in order to prove that every Besicovitch set in $\mathbb{R}^2$ has dimension 2, it suffices to prove that $V(B, \varepsilon)$ is big for each Besicovitch set B. In fact, we will prove the following theorem:

Theorem (Davies, basically)**.**

$$V(B, \varepsilon) \geq \frac{s}{\log \frac{1}{\varepsilon}}$$

for some constant s.

Indeed, this suffices. For if this happens, then

$$\dim_M B = \lim_{\varepsilon \to 0} \left(2 + \frac{\log V(B, \varepsilon)}{\log \frac{1}{\varepsilon}} \right) \geq \lim_{\varepsilon \to 0} \left(2 + \frac{\log s}{\log \frac{1}{\varepsilon}} - \frac{\log \log \frac{1}{\varepsilon}}{\log \frac{1}{\varepsilon}} \right) = 2$$

In order to do this, we will actually prove that a subset of $N_\varepsilon(B)$ has such a large area; this large subset will consist of rectangles around lines in various directions. Assume for convenience that $1/\varepsilon = 2^\ell$ is a power of two, and set

$$\Omega = \left\{0, \frac{\pi\varepsilon}{2}, \pi\varepsilon, \ldots, \frac{\pi}{2} - \varepsilon, \frac{\pi}{2}\right\}$$

For $\omega \in \Omega$, let L_ω be a unit line segment with angle ω contained in B. Also, let R_ω be a $1 \times \varepsilon$ rectangle whose central axis is L_ω. Then since $L_\omega \subseteq B$, we find that

$$T := \bigcup_{\omega \in \Omega} R_\omega \subseteq N_\varepsilon(B)$$

So it suffices to prove that T has a big area. The basic intuition here is that each R_ω has a pretty big area (namely ε), and we have a bunch of them (namely $1/\varepsilon$). Moreover, since they go in different directions, they have a pretty small intersection (they don't overlap a lot), so their union should have a pretty big area. The precise lemma we need is as follows:

Lemma. *Let $A_1, \ldots, A_n$ be subsets of $\mathbb{R}^2$, and let $A = \bigcup_{i=1}^n A_i$. Then*

$$\left(\sum_{i=1}^n \operatorname{area}(A_i)\right)^2 \leq \operatorname{area}(A) \sum_{i=1}^n \sum_{i'=1}^n \operatorname{area}(A_i \cap A_{i'})$$

Proof. Overlay all of the A_is, and consider all subsets of the plane that are formed in this way, namely all possible intersections of various A_is. Call these regions $B_1, \ldots, B_m$. For each $1 \leq j \leq m$, let $w(j)$ denote the number of A_is that contain the region B_j, namely the number of A_is that overlap on this B_j. Then we have that

$$\sum_{i=1}^n \operatorname{area}(A_i) = \sum_{j=1}^m \operatorname{area}(B_j) w(j)$$

This follows immediately from the definition of $w(j)$, when we split up the sum as a sum over all the B_js. Similarly, observe that

$$\sum_{j=1}^m \operatorname{area}(B_j) w(j)^2 = \sum_{i=1}^n \sum_{i'=1}^n \operatorname{area}(A_i \cap A_{i'})$$

This is because the area of each B_j is counted on the right-hand side multiple times, where the number of times is precisely the number of pairs i, i' with $B_j \subseteq A_i \cap A_{i'}$, which is precisely $w(j)^2$.

Now, recall the Cauchy-Schwarz inequality, which says that for any two sequences $c_1, \ldots, c_m, d_1, \ldots, d_m$, we have that

$$\left(\sum_{j=1}^m c_j d_j\right)^2 \leq \left(\sum_{j=1}^m c_j^2\right)\left(\sum_{j=1}^m d_j^2\right)$$

We apply this with $c_j = \sqrt{\operatorname{area}(B_j)}$ and $d_j = w(j)\sqrt{\operatorname{area}(B_j)}$. Then it tells us that

$$\left(\sum_{j=1}^m \operatorname{area}(B_j) w(j)\right)^2 \leq \left(\sum_{j=1}^m \operatorname{area}(B_j)\right)\left(\sum_{j=1}^m \operatorname{area}(B_j) w(j)^2\right) = \operatorname{area}(A) \sum_{i=1}^n \sum_{i'=1}^n \operatorname{area}(A_i \cap A_{i'})$$

□

With this lemma, we can now prove Davies' theorem.

Proof. We apply this lemma where our A_is our just the sets $\{R_\omega\}_{\omega\in\Omega}$, and $A = \bigcup A_i = T$. Notice that since each R_ω is a $1\times\varepsilon$ rectangle, it has area ε. Since there are $1/\varepsilon$ of them total, we have that

$$\sum_{\omega\in\Omega} \text{area}(R_\omega) = 1$$

and thus this inequality can be rearranged to

$$\text{area}(T) \geq \frac{1}{\sum_{\omega_1\in\Omega}\sum_{\omega_2\in\Omega}\text{area}(R_{\omega_1}\cap R_{\omega_2})}$$

So, in order to prove that $\text{area}(T) \geq s/\log\frac{1}{\varepsilon}$, it suffices to prove that

$$\sum_{\omega_1\in\Omega}\sum_{\omega_2\in\Omega}\text{area}(R_{\omega_1}\cap R_{\omega_2}) \leq S\log\frac{1}{\varepsilon}$$

for some new constant S. For this, we use a fact you proved on the homework, namely that when two $1\times\varepsilon$ rectangles meet at an angle θ, then we have that the area of their intersection is at most $C\varepsilon^2/\theta$, for some constant C (in fact $C = \pi/2$ suffices).

Recall that $\ell = \log_2\frac{1}{\varepsilon}$. For some fixed $\omega_1\in\Omega$, and some $1\leq k\leq \ell$, let

$$D_k(\omega_1) = \{\omega_2\in\Omega : \varepsilon 2^k \leq |\omega_2-\omega_1| < \varepsilon 2^{k+1}\}$$

Then for each $\omega_2\in D_k(\omega_1)$, we have that

$$\text{area}(R_{\omega_1}\cap R_{\omega_2}) \leq \frac{C\varepsilon^2}{|\omega_1-\omega_2|} \leq \frac{C\varepsilon^2}{\varepsilon 2^k} = \frac{C\varepsilon}{2^k}$$

Additionally, $|D_k(\omega_1)| \leq 2^k$. Therefore, for each fixed ω_1,

$$\begin{aligned}\sum_{\omega_2\in\Omega}\text{area}(R_{\omega_1}\cap R_{\omega_2}) &= \sum_{k=1}^{\ell}\sum_{\omega_2\in D_k(\omega_1)}\text{area}(R_{\omega_1}\cap R_{\omega_2})\\ &\leq \sum_{k=1}^{\ell} 2^k\cdot\frac{C\varepsilon}{2^k}\\ &= C\varepsilon\sum_{k=1}^{\ell}1\\ &= C\varepsilon\ell\\ &= S\varepsilon\log\frac{1}{\varepsilon}\end{aligned}$$

for a new constant $S = C\log 2$. Finally, we sum over all ω_1, and find that

$$\sum_{\omega_1\in\Omega}\sum_{\omega_2\in\Omega}\text{area}(R_{\omega_1}\cap R_{\omega_2}) \leq \sum_{\omega_1\in\Omega} S\varepsilon\log\frac{1}{\varepsilon} = S\log\frac{1}{\varepsilon}$$

since $|\Omega| = \frac{1}{\varepsilon}$. □