22. Sharp Projection Theorems I: Introduction and Beck's Theorem

May 6

In the last two years, Orponen-Shmerkin and Ren-Wang proved the Furstenberg set conjecture. As a special case, this gives a sharp projection theorem in $\mathbb{R}^2$, completely answering the questions in projection theory first raised by Kaufman in the 1960s. It can also be viewed as a harmonic analysis cousin of the Szemeredi-Trotter theorem. I think it is a remarkable result, and this work was one of my main motivations to teach this class.

The full proof of the Furstenberg set conjecture spans several long papers. It is too long and too technical to give the full proof in these lectures. But in the last three lectures we will discuss some of the ideas of the proof.

We begin this section by restating the Szemeredi Trotter theorem, an important sharp theorem in projection theory.

Theorem 22.1 (1982)**.** *Let E be a set of points in $\mathbb{R}^2$. Pick some integer $S > 1$. For every x in E, let L_x be a set of S lines passing through x. Define $L = \bigcup_{x \in E} L_x$. Then*

$$|L| \gtrsim \min(|E| \cdot S, |E|^{1/2} S^{3/2}) \tag{46}$$

This theorem is discussed in more detail earlier in these notes. It is one of the earlier examples of a sharp theorem in projection theory. The Furstenberg set conjecture is a continuous analogue of this theorem, in which points are replaced by δ-balls and lines are replaced by δ-tubes. To state it, we first recall the (δ, s, C) spacing condition.

Definition 22.2. *A set $E \subset \mathbb{R}^d$ contained in the unit ball centered at the origin is (δ, s, C) if*

$$|E \cap B_x(r)|_\delta \leq C r^s |E|_\delta \tag{47}$$

for all balls of radius r with $r \geq \delta$ centered at arbitrary points x, where $|\cdot|_\delta$ is the δ-ball covering number.

The following theorem by Orponen, Shmerkin, Ren, Wang (OSRW) gives an analogous statement to the Szemeredi Trotter theorem for (δ, s, C) sets (this statement is also known as the Furstenberg conjecture (or FC)) :

Theorem 22.3. *[Furstenberg Conjecture, OSRW (2024)] Let $E \subset \mathbb{R}^2$ be a (δ, t, C) set. Define a δ tube as a $1 \times \delta$ rectangle in the plane. For every $x \in E$, let $\mathbb{T}_x$ be a set of δ tubes in $\mathbb{R}^2$ passing through x. For each x, let $\operatorname{Dir} \mathbb{T}_x \subset S^1$ be the set of* *directions of the δ tubes in $\mathbb{T}_x$. Assume that for all $x \in E$, $\operatorname{Dir} \mathbb{T}_x$ is a (δ, s, C) subset of S^1. Define*

$$\mathbb{T} = \bigcup_x \mathbb{T}_x$$

Then for every $\epsilon > 0$

$$|\mathbb{T}| \geq c_\epsilon C^{-O(1)} \delta^\epsilon \min(\delta^{-t-s}, \delta^{-t/2-3s/2}, \delta^{-1-s}) \tag{48}$$

where c_ϵ is a constant depending on ϵ.

Denote the three cases of the minimum value, A, B, and C, in order. The first two cases of the minimum value are analogous to the cases of the Szemeredi Trotter theorem. In case A, each point has many lines passing through it, and each line passes through only one point. The second case corresponds to a grid of points with lines corresponding to rational angles.

The third case is new in the setup with δ-balls and δ-tubes. Notice that if we randomly pick a δ-ball in $B^2(1)$ and a δ-tube in $B^2(1)$, the the probability that the δ-ball intersects the δ-tube is $\sim \delta$. To get an example in this third case, we randomly pick a set E consisting of δ^{-t} δ-balls and a set $\mathbb{T}$ consisting of δ^{-1-s} δ-tubes. For each $x \in E$, we define $\mathbb{T}_x$ to be the set of $T \in \mathbb{T}$ so that $x \in T$. With high probability, for every $x \in E$, $|\mathbb{T}_x| \approx \delta^{-s}$. Moreover, for any $\eta > 0$, with high probability the set E will be $(\delta, t, \delta^{-\eta})$ and each $\operatorname{Dir} \mathbb{T}_x$ will be $(\delta, s, \delta^{-\eta})$.

22.1. **History of the Furstenberg conjecture.** The first set of methods that were applied are classical methods, due to Kaufman, Falconer, and Wolff. These consisted of double counting arguments and Fourier methods.

Double counting methods give sharp bounds when the first term in Theorem 22.3 dominates, which happens when $s \geq t$.

Fourier methods give sharp bounds when the third term in Theorem 22.3 dominates, which happens when $s + t \geq 2$.

When the second term dominates, classical methods are not sharp. One of their key deficiencies that they cannot distinguish between $\mathbb{R}$ and $\mathbb{C}$. As the Furstenberg conjecture is false in $\mathbb{C}$, it is essential to use methods that distinguish the two spaces.

The second set of methods, from 2000 to 2022 was ϵ improvements. These methods began with Bourgain's projection thoerem in 2000, and showed bounds that were ϵ better than the trivial or classical bounds. For instance, these methods were used to show that if $t = 1$ and $s = 1/2$, then

$$|\mathbb{T}| \gtrsim C^{-O(1)} \delta^{-1-\epsilon}$$

for some tiny but explicit $\epsilon > 0$. Without the ϵ, this bound is trivial, but the ϵ improvement was a large step forward. In particular it was the first result to distinguish $\mathbb{R}$ from $\mathbb{C}$.

In 2021, Orponen and Shmerkin pushed these methods further, proving a very general ϵ-improvement result.

Theorem 22.4. *[Orponen-Shmerkin (2021)] Under the same hypotheses as Theorem 22.3, for every $0 < s < t$ there is $\epsilon > 0$ so that*

$$|\mathbb{T}| \gtrsim \delta^{-2s-\epsilon} \tag{49}$$

In this situation, the classical method gives the lower boundf $|\mathbb{T}| \gtrapprox \delta^{-2s}$ and this theorem improves the classical bound by ϵ. The proof uses Bourgain's projection thoerem, as well as other ideas.

The main progress in the second stage consisted of proving ϵ improvements in more and more general situations. This last theorem of Orponen-Shmerkin was an important step in that direction. However, thoughtout this second phase, the value of ϵ remained quite small.

The third phase is based on repeatedly applying the ϵ improvement results to reach a sharp result. The result by OSRW is a key example of these methods.

It is striking and surprising that it is possible to bootstrap the ϵ-improvement theorems to get sharp bounds, and I think this is one of the main ideas to take away from the recent work in projection theory. In this class and the next class, we will try to explain how it works.

We begin in this class with the simplest example I know in which an ϵ-improvement can be bootstrapped to get a sharp bound. The result is an analogue of Beck's theorem from combinatorial geometry, and we begin by stating Beck's theorem.

Theorem 22.5 (Beck)**.** *Let E be a set of points in $\mathbb{R}^2$ and for any line ℓ, assume that ℓ intersects at most half of the points in E. i.e.*

$$|\ell \cap E| \leq \frac{1}{2}|E|$$

For every $x \in E$, let $L_{x,E}$ be the set of lines passing through x that also pass through an additional point in E, i.e.

$$L_{x,e} = \{\ell : \ell \text{ is a line passing through } x \text{ such that } |\ell \cap E| \geq 2\}$$

Then for every x

$$|L_{x,E}| \gtrsim |E|$$

Proof sketch. We assume that the lines in each $L_{x,E}$ are uniform, that is each $L_{x,E}$ contains approximately the same number of lines, and each point has approximately the same number of lines passing through it. This implies that

$$|E \cap \ell| \sim \frac{|E|}{|L_{x,E}|}$$

for each ℓ and each x. Let $S \sim |L_{x,E}|$ be the number of lines through each point. Also let $L = \cup_{x \in E} L_{x,E}$.

Double counting shows that

$$|E| \cdot S \sim |L| \cdot |E|/S$$

The left hand side is the number of points multiplied by the number of lines per point, so is the total number of lines multiplied by the number of points per line. The right hand side is number of lines multiplied by $|E|/S$, which is the number of points per line. The two sides are therefore equal. By manipulating the equation, we get

$$|L| \sim S^2$$

On the other hand, the Szemeredi Trotter theorem tells us that

$$|L| \gtrsim \min(S|E|, S^{3/2}|E|^{1/2}).$$

Since $|L| \sim S^2$, this implies that

$$S \gtrsim |E|$$

the desired conclusion. □

Note that to obtain this conclusion, a weaker version of Szemeredi Trotter is sufficient. We only need to know that if

$$|E| \gg S$$

then

$$|L| \gg S^2$$

This weaker version of Szemeredi-Trotter is only an ϵ improvement of a double counting bound. This bound is analogous to the bound in Theorem 22.4. Using Theorem 22.4, Orponen, Shmerkin, and Wang were able to prove a continuum analogue of Beck's theorem. Here is the statement.

Theorem 22.6. *[Continuum Beck's Theorem, OSW (2023)] Choose* $\eta > 0$ *and let* E *be a* (δ, u, C) *set in the plane such that for all* $\rho \times 1$ *rectangle* R*,*

$$|E \cap R|_\delta \leq C\rho^\eta |E|_\delta$$

Then for most $x \in E$*,*

$$|L_{x,E}|_\delta \gtrsim \delta^\epsilon \min(\delta^{-u}, \delta^{-1})$$

(Here $L_{x,E}$ is a set of lines through the point x. We define the distance between two such lines as the angle between them, and so we can define $|L_{x,E}|_\delta$.)

This theorem is sharp. And the result is false over $\mathbb{C}$. It is one of the first sharp theorems in projection theory which distinguishes $\mathbb{R}$ from $\mathbb{C}$.

The proof is based on the proof of Beck's theorem, but there is a new issue in this setting, and a new idea to deal with the issue. Here we give only a proof sketch, explaining the new issue and the new idea.

Suppose we try to imitate the proof of Beck's theorem using Theorem 22.4 in place of the Szemeredi-Trotter theorem. In order to apply Theorem 22.4, we need to assume that each $L_{x,E}$ is a (δ, s, C) set for some s. By doing some uniformization, we can reduce to the case that all the sets $L_{x,E}$ are similar to each other: $|L_{x,E}|$ is roughly constant in x and every $L_{x,E}$ is a (δ, s, C) set for the same s, C.

As above, we let $L = \cup_{x \in E} L_{x,E}$. We let $\mathbb{T}$ be the set of δ-tubes formed by thickening the line segments of L. Several lines may thicken to essentially the same δ-tube $T \in \mathbb{T}$. We let $\mathbb{T}_x$ be the set of tubes of $\mathbb{T}$ passing through x. So we have $|\mathbb{T}_x| \sim |L_{x,E}|_\delta$.

A version of the same double counting argument as above shows that

$$|\mathbb{T}| \sim |\mathbb{T}_x|^2 \sim |L_{x,E}|_\delta^2.$$

On the other hand, Theorem 22.4 tells us that if $0 < s < \min(u, 1)$ then

$$|\mathbb{T}| \gtrsim \delta^{-2s-\epsilon}$$

for some small $\epsilon = \epsilon(s, u)$. Comparing the last two equations, we see that

$$|L_{x,E}|_\delta \gtrsim \delta^{-s-\epsilon}$$

We state what we have learned as a lemma.

Lemma 22.7. *If* $0 < s < \min(u, 1)$*, and a typical set* $L_{x,E}$ *is* (δ, s, C)*, then*

$$|L_{x,E}|_\delta \gtrsim \delta^{-s-\epsilon}$$

Let us reflect on the lemma. If $L_{x,E}$ is (δ, s, C), then it follows that $|L_{x,E}|_\delta \gtrsim \delta^{-s}$. This lemma improves on that trivial bound by an ϵ. However, it looks far from the sharp bound in Theorem 22.6.

Orponen, Shmerkin, and Wang proved Theorem 22.6 by a bootstrapping argument, where Theorem 22.4 is used not just once but many times at many different scales.

We now sketch this bootstrapping argumen. Suppose that $L_{x,E}$ is uniform and (δ, s, c) for s, where $0, s < \min(u, 1)$. Then $L_{x,E}$ is (ρ, s, C) for $\rho \geq \delta$. The lemma then implies that for every ρ,

$$|L_{x,E}|_\rho \gtrsim \rho^{-s-\epsilon}$$

for every $\rho \geq \delta$. From the assumption that $L_{x,E}$ is a uniform set, $L_{x,E}$ is therefore a $(\delta, s+\epsilon, C)$ set. To summarize, we now have a stronger lemma:

Lemma 22.8. *If $0 < s < \min(u, 1)$, and if a typical set $L_{x,E}$ is uniform and (δ, s, C), then a typical $L_{x,E}$ is $(\delta, s+\epsilon, C')$ where $\epsilon = \epsilon(s, u) > 0$.*

The hypothesis in Theorem 22.6 that E does not concentrate too much in rectangles shows that each $L_{x,E}$ is a (δ, η, C) set. Starting with this assumption, we can then apply Lemma 22.8 repeatedly. As we keep iterating, the value of s will approach $\min(u, 1)$.

We should note that this proof sketch was not a complete proof. The technical work that is missing is to make precise what we mean when we say that $L_{x,E}$ is typical. This requires some careful uniformizing and pigeonholing.

I was very impressed when Theorem 22.6 was proven, because it gives the sharp answer to a natural question in projection theory and distinguishes $\mathbb{R}$ from $\mathbb{C}$. On the other hand, it was not at all clear to me whether these ideas would lead to sharp answers to more difficult problems like the Furstenberg set conjecture. Here is one issue. In the combinatorial geometry world, it was well known that an ϵ-improvement to Szemeredi-Trotter implies Beck's theorem, and that Beck's theorem is sharp. The proof of the continuum Beck's theorem builds on this observation. But on the other hand, no one knows how to bootstrap an ϵ-improvement to Szemeredi-Trotter in order to prove the full Szemeredi-Trotter theorem. So it was not all clear whether to expect that we could bootstrap the ϵ-improvement estimate in Theorem 22.4 in order to prove the Furstenberg set conjecture. In fact, it would be fair to stay that this strategy sounded very doubtful to me.

As we will see, OSRW did prove the Fustenberg set conjecture, and bootstrapping theorem 22.4 played a key role.

22.2. **Outline of OSRW proof of Furstenberg conjecture.** We now begin to discuss the proof of the Furstenberg set conjecture, just at the level of a very broad outline.

The proof is split into cases based on the spacing of the set E. I believe this division into cases is a second major takeaway from the recent work. Until recently, most proofs in projection theory applies for all (δ, s, C) sets. But different (δ, s, C)

sets can have quite different spacing properties. And it turns out that depending on the way a set E is spaced, different tools are helpful to bound the projection theory of E.

The best language for describing the spacing of E is the language of branching functions and uniform sets. First by a pigeonholing argument, we can reduce to the case where E is a uniform set with $\delta = \Delta^m$ for some large m. Recall that the uniform condition on E means that for any dyadic Δ^j cube Q with j an integer between 1 and m, then

$$|E \cap Q|_{\delta^{j+1}} \sim R_j$$

where $1 \leq R_j \leq \Delta^{-2}$ is a branching number that determines the spacing of E.

The sequence of branching numbers R_j gives very precise information on "the way E is spaced". Notice that recording the sequence of branching numbers contains a lot more information than a single number s that would appear if we said that E is a (δ, s, C) set.

To build intuition, it is well worth a little time to draw sets with a few different branching functions. Here are two different cases that turn out to play an important role in the story.

AD regular case. For every j, $R_j \sim \Delta^{-t}$. In this case, the set E is a (δ, t, C) set. But not all (δ, t, C) sets are AD regular.

Well spaced case. In this case, $R_j = \Delta^{-2}$ for $j \leq m$ and $R_j = 1$ for $j > m$. The number of points in the set E is Δ^{-2m}, and these points are as well-separated as possible. If we choose t so that $\delta^{-t} = \Delta^{-2m}$, then the set E is a (δ, t, C) set. But it looks very different from an AD regular set.

Among (δ, t, C) sets, the AD regular set is the most compressed (the distances between points are as small as possible). And the well spaced case is the most spread out.

There is also a continuum of cases in between.

One important feature of the proof is that there are different tools for the AD regular case and the well spaced case.

In 2024, Orponen-Shmerkin proved the AD regular case of the Furstenberg set conjecture. Their proof uses a bootstrapping argument and uses the continuum Beck theorem. It could be described as an elaborate bootstrapping argument which uses the ϵ-improvement in Theorem 22.4 many times. (Recall that there is no known bootstrapping argument to deduce Szemeredi-Trotter from a weaker ϵ-improvement version of Szemeredi-Trotter. But Orponen and Shmerkin showed that the AD regular case has a lot of special structure, and in this case the sharp estimate does ultimately follow from an ϵ-improvment version.)

Somewhat earlier, Guth-Solomon-Wang proved that well spaced case of the Furstenberg set conjecture. The proof is based on Fourier methods.

At this point, the Furstenberg conjecture had been proven in two extreme cases by very different methods. But there were many other cases in between these.

A little later in 2024, Ren and Wang proved the full Furstenberg set conjecture. They used a multiscale argument which breaks the problem into several different scales. And they showed that, if the sequence of scales is picked carefully, then each scale can be controlled using either the AD regular case or a Fourier method generalizing the GSW method.

In the next two lectures, we will survey these developments, spending one lecture on the AD regular case, and one lecture on the rest of the proof.