[SQUEAKING]

[RUSTLING]

[CLICKING]

LAWRENCE GUTH: OK, everyone. So today, we're going to learn some of the fundamental estimates of projection theory for balls in the plane for the setting of real analysis. Two goals of the class. The first goal is to adapt to the techniques that we learned last time in finite fields, so the double counting technique and the Fourier technique, to Euclidean space. And the second goal is to start to explore.

So there's a difference between Euclidean space and finite fields is that in Euclidean space, we have to pay attention to how our sets are clustered. And so the second goal of the class is to start to explore clustering--information about how the sets are clustered and how that comes into the story. OK.

So here's our setup. We'll have x will be a set of disjoint unit balls in the ball of radius R squared. OK. So that's our set of points or set of balls. And then we're going to have a set of directions called D. So D is a finite subset of S1. It's 1 over r separated. So the distance between any two points of D is at least 1 over R.

And then S which depends on X and D is the maximum size of the projection. So I look at all the directions in D and I project the set X in that direction, and I take its measure. And the maximum of that is called S. OK. So projection theory, we try to understand how these things are related to each other. And in this Euclidean setting, we also need to care about how this set X and how the set D are clustered.

So that's measured by this-- I guess call them clustering functions. So N sub X of little r is the measures-- the most amount of X in any ball of radius little r. So this is the maximum over C in R2 or of the size of x intersected with the ball around c of radius r.

So this measures how x is clustered. And we're also going to measure how D is clustered. So N sub D of rho is the maximum over any-- so sigma in S1 and arc of length rho of D intersect c. OK.

OK. So our big goal is in terms of this information and how big D is and how big X is, what can we say about S. OK. So we're first going to do the double counting approach. And the argument is, in spirit, the same as last time. But we'll see that when we try to do the double counting that we did last time, that the N sub X function and the N sub D function will come into the story, and it leads to the following bound.

So theorem one. I'll call it theorem 1R for the real case. Whereas last time we had something called theorem one, which, when I do it again, it'll be theorem 1F for the finite field case, and eventually we'll compare them. OK. So if we have this setup, then the number of directions is bounded by-- let's see.

OK. So you'll notice that this is not a formula that I have memorized. S over X times the sum over different scales. Sorry.

OK. So it's a sum over different scales, and we often in this class will sum over a dyadic sequence of scales, powers of two or whatever. OK. And so the formula doesn't yet look that similar, but the idea is the same as last time. And then eventually we'll compare the formulas and see that under certain assumptions on this it's actually very similar to last time.

OK. So the proof is to double count a certain quantity, which we called coincidences last time. So the key quantity called star is the number of B1 and B2 in X. These are unit balls. So remember, X is a set of unit balls. So these are two of the unit balls. And theta and D. I guess we'll say they're not equal to each other, although that won't matter as much this time.

So that pi theta of B1 and pi theta of B2 are similar to each other. And in this setting, we won't count-- they don't have to be exactly the same. I would just like them to overlap. So my condition will be that the intersection of these is non-empty. So that's my key quantity.

Actually, in this setup, I'm going to allow them to be equal to each other. That makes the double counting a little easier. And in a sense, we'll see, it's part of what's up here.

OK. So there's a lower bound, which is straightforward. The lower bound says that this star is at least the number of directions times in each direction X over S squared times S. So in each direction, these X different balls have been projected into only S slots.

So a typical slot has X over S balls in its fiber. And so we have this many squared choices to pick B1 and B2 in the same fiber. S times. So the lower bound is really the same as last time. Same as last time. Same as the homework.

OK. So now let's look at the upper bound. So I will-- for now I'll fix B1, B2 and I'll say that the distance between B1 and B2 is around r. So here is my picture B1 and B2 and the distance between them is around r.

OK. So now I'm interested in projections where the projection of B1 and the projection of B2 overlap. And so the angles for those projections are in a little angular range. The angle of the projection must be somewhere in this range here. So what I'll say is let's say that v is the direction from B1 to B2. So v is center of B2 minus center of center B1 over the norm of that.

And if pi theta of B1 intersect pi theta B2 is not empty, then theta must be pretty close to v. What I'll say is that the angle between theta and v is smaller than around 1 over r. OK. So now the number of theta so that pi theta of B1 intersect pi theta of B2 is not empty. That would be bounded by, well, the number of directions in this little arc. So that would be ND of 1 over r.

So we see the spacing function of the set D comes into play here. OK. So that's what happens when we fixed B1 and B2. And the case B1 equals B2 is allowed. And I'll make this slightly funny convention that then r is one unit ball, so the distance between them is one.

And this is true, then. I mean, v is not well defined, but this condition includes the whole range of angles, so it's true. So remark if B1 equals B2, take little r equals 1. It's slightly exceptional but all the bounds are true.

OK. So if we want to upper bound star, now we need to count the number of choices for B1 and B2. So now I'm interested in the number of pairs, B1, B2, and x with the property that the distance from B1 to B2 is at most r. OK. How big is that? Well, I have norm of X choices for B1. But once I choose B1, I don't have all that many choices for B2 because B2 has to be close to B1. It has to be in a little r ball.

So the number of choices for B2 is at most NX of r. So at this moment, the spacing function of X comes into play. OK. So now we can combine these to upper bound star. So star is now upper bounded by X times NX of r times ND of 1 over r. Ah, sorry. So star is upper bounded by the sum over the dyadic r of this quantity.

OK. So let me say why. So star is this set of coincidences. It's a B1 and B2 and a theta, and for each coincidence, I sort these coincidences by the distance from B1 to B2, which I call r. So I take the coincidences where the distance is around r and I sum over the possible r's. Once I fixed r, I have this many choices for the pair B1, B2. And once I've fixed B1 and B2, I have this many choices for theta. So that's the point.

OK. The rest is algebra. I'll sketch it quickly on the board, and that'll be the end of the proof. And then we'll pause to digest and check it.

OK. So all together, we have X squared S inverse D is bounded by star which is bounded by the sum over r dyadic of X and X of r and D of 1 over r. And then I just get the D by itself. If you move this stuff over to the other side, you get the formula on the board. OK. So let's pause there for questions and comments.

OK. All right. So I wanted to get a sense of start to explore this function NX of r by looking at a few different examples of sets X and writing down what is NX of r. It's not difficult, but having a few different examples in our mind will help to orient us as we keep going. And actually will help also to digest this formula. So this proof was clean. This formula is a little bit messy. So after looking at some examples, we'll will digest this formula and think about a situation where it works out a little more cleanly. All right.

All right. So examples for N sub X of r. OK. So number one, X is the one neighborhood of a curve, like a circle or parabola or something. So here's the ball of radius r. And inside of there, I have a nice curve. And all along the curve I have unit balls max X. So if I have this set, how many unit balls do I have in a ball of radius little r? That would be around r.

OK. Do a second example. Second example would be very spread out. If you imagine that these points are positive charges that repel each other, and they push as far away from each other as possible, they would arrange themselves in something like a lattice inside of here. OK. So I'll call this the well spaced example, the most spread out example. So it looks like this.

So it's let's say it's a well spaced n by n grid. And so the distance between neighboring balls would be like r over N. So now what would NX of r be like. Well, it would be one for quite a while. If our radius is less than r over N, then this would be 1. But then eventually, as the radius gets bigger, it would get bigger than 1. And at that point, it would be its fair share of the total number or total size of X.

So once r is bigger than this threshold, every little r ball contains the same amount of X. And so the density of X is like N squared over r squared. The area of the ball is like this. OK. And I'll do a third example, which is the points could be very clustered. Cluster. So we could have N squared unit balls in an N ball. So the picture would be like this.

We have our ball of radius r. Inside of there we have a smaller ball of radius N, and we pack the smaller ball with unit balls of X. So now what is N sub X of r? Well, it grows like r squared until r is as big as this N ball. And after that, there are only N squared balls at all so you can make it a little bigger. It doesn't grow anymore. So it caps off it N squared if r is bigger than N.

OK. So now I will draw some pictures of these three functions. You can compare them. OK. We'll do some pictures. I'm going to normalize in my pictures, I have to choose an N. I'm going to choose N equals r to the 1/2. That's helpful because it says that the total size of X is around r in all three examples. It makes them a little easier to compare with each other.

OK. So now I'll make a graph. Basically, I'd like to put on this axis little r and on this axis N sub X of little r. But this graph, like many graphs in math, looks nicer if we use a log log plot. So on this axis, I'm going to put the log of little r. And the nicest base for that is big R. This goes from 0 to 1. And on this axis, I'm going to put the log of N sub X of r. And it's natural to use the same base.

So this goes from 0 to 1. This could go above one potentially, but because the size of X happens to be r, we're going to stop there. So now what do these three examples look like. The first example is a straight line. So that's example one.

The second example grows more slowly than the straight line for a while. It doesn't grow at all for a while. And then it grows quickly to catch up. So that's the second example. It lies below the straight line.

And then the third example grows more quickly than the first example, and then it reaches a certain point and it stops growing. So it looks like this. And it lies above the straight line. OK. So in this theory, there turns out to be that things that grow-- the straight line case is an important case. And then there's an important distinction between things that are below the straight line, things that are above the straight line.

So the straight line case would be that N-- that the size of X would be r to the alpha for some alpha, and then N sub little r of X or N sub X of little r would be little r to the alpha. That would be the straight line case. And the below the straight line case, X would still be R to the alpha. So the end point would be the same, but now in the middle we're going to go below what we had before. So N sub X of r would be at most r to the alpha.

OK. So this thing here is called a regular alpha dimensional spacing. And this thing is called alpha dimensional spacing. I'd like to make one more definition. I'll say that X has Hausdorff spacing if it has alpha dimensional spacing for some alpha.

But another way of saying that is if you look at how many points of X are in a ball of radius R to the beta, it's bounded by the total number of points of X to the beta. And this is equivalent. This is still below the straight line.

OK. So over here we have example one. And in general, if you have any-- if you have a submanifold in any dimensional space, it will behave like this. And alpha is the dimension of the submanifold.

Also some self-similar Cantor set kind of fractals, like the Cantor set, have this kind of spacing where alpha would be the dimension of your Cantor set. That's why it's called alpha dimensional. And these things, this thing called Hausdorff spacing, I called it that because it appears naturally in the theory of Hausdorff dimensions.

Yeah. So it's maybe worth saying briefly that in the geometric measure theory literature, a lot of the stuff we're talking about today was originally described in the language of Hausdorff dimensions. But my impression is that that's not really that helpful. Like, if we defined Hausdorff dimensions, and then we wanted to prove things about Hausdorff dimensions, we would add an extra layer of definition. And then we'd have to check the extra layer of definition.

But then in all the applications that we'll talk about, we would never use those definitions. We would go back to this theorem the way it's stated now and use that. So I have-- a bit tentatively, but I've come to the conclusion that Hausdorff dimension is not really a helpful notion for this subject. And so I'm going to leave it out. OK. But any questions or comments about these different ways that a set can be spaced. Yeah?

AUDIENCE: I'm wondering, like in the theorem one for the lower bound, basically, if we try to calculate lower bound for V1, V2, is this somehow r from R1 to R2, like that, and compare this with the upper bound would be the distance from V1 and V2 that gives more information than the inequality?

LAWRENCE GUTH: Yeah. OK. So the question is, when we were processing the upper bound for this, we sorted it according to the distance. So we really have a set star sub r where we add the condition that the distance from V1 to V2 is around r. And the question is if we looked at the lower bound we could also sort it by r, and would that be helpful. Yeah.

Yeah. That is a good thing to look into. Off the top of my head, I don't know if it would be helpful. Our lower bound with this S here depends on taking all of them, but it would be interesting. Yeah, it would be interesting to think about. Other questions or comments?

OK. OK. So let me show you. So I claim-- OK. It seems that this spacing condition, say Hausdorff spacing, is helpful in the theory. And one way that comes about is that a lot of the theorem statements become a lot cleaner if you put in Hausdorff space. So let's do that here. Let's suppose that X and D have Hausdorff spacing. We'll use that to simplify this thing.

OK. So corollary one, real, is that if we have the setup and X and D are Hausdorff spacing, then we get a simpler inequality, which matches quite well what happens in finite fields. Then D is bounded by S plus X. So let me check.

OK. So I'll prove this and then I'll try to argue. Actually, so OK. So this looks a little nicer than before. It might still not look that nice to you. But I would summarize it like this. Either S is pretty close to X or D is smaller than S. And that's what we had in finite fields. So why is it? If S is a lot smaller than X, this is much less than 1, so we can rearrange that. And we get that D is less than S.

OK. So what's the proof? Well, we're going to have an expression like NX of r ND of 1 over r. And so now let's suppose that little r is big R to the beta. And then being Hausdorff. It tells us that NX of r to the beta ND of r to the minus beta there, that is bounded by X to the beta times D to the 1 minus beta. That's what being Hausdorff tells us.

So being Hausdorff allows us to relate NX of r for any r to the size of X, which is NX or big R. And algebraically it works out like this, and for D, it works out like that. OK. And then by [? AM-GM, ?] this is smaller than X plus D. So every one of these complicated things, they're all bounded by X plus D. And the number of them is only log of r. So OK, sorry. So I should put a log of r here.

OK. So now we're just going to plug this in up there. I'll have an X plus D up there and a log r factor. The X and the X plus D gives me this, and the D and the X plus D gives me that. That's just algebra.

OK. Now let's compare this with the theorem we proved last time in finite fields. So theorem one in finite fields said that if we had X in Fq2 and D in Fq and S is the maximum theta in D size of pi theta of X then either S is around X or D is bounded by S. And the numerology is now the same in finite fields, or if you have balls, with the Hausdorff space.

OK. So in general, the numerology in a lot of things is the same between finite fields and balls with Hausdorff spacing. And so now I'll-- so this was theorem one from last time. I'll track through also theorem two from last time and the conjecture from last time. And we'll see that they all match up.

OK. So we have projection theory in Fq2 versus unit balls in B2R and B2R with Hausdorff space. So here was theorem two last time in finite fields. So if we have our finite field setup, which is over here, and S is smaller than q, then the conclusion is that the number of directions is smaller than Sq over X.

And here's what we'll prove in the real case corollary 2R. So if we have our setup of the day and X and D are Hausdorff and S is a good bit smaller than R, then maybe this needs a little bit more. Then D is smaller than SR over X. OK. So I use this symbol with the extra tilde. That's hiding something like a log factor, similar to the log factor that appeared over here.

Like, if you compare the real case theorem one and the finite field case, we lost an extra log r factor. That will happen sometimes. So notation. A less than some B means for every epsilon bigger than 0, there exists C epsilon. So A is less than or equal to C epsilon R to the epsilon.

But if you compare these, q corresponds to R. That makes sense. There are q squared points here and there are r squared unit balls here. And these basically match up except that there's a minor loss here. OK. And so this is the Fourier method that we'll turn to in a little bit. There is also a conjecture about the sharp behavior.

So conjecture. If p is prime and we have our finite field setup, and S is smaller than the minimum of q and X. So either of these are trivial. S is always at most q and it's always at most X. So we're make it less than that, so we're in a non-trivial range. Then D should be bounded by S squared over X, which would be sharp if it's true, and it matches the behavior in an integer grid.

OK. This conjecture has an analog over over-- in the real case. And the real analog is due to Furstenberg. It's called the Furstenberg set conjecture. So it says if we have the setup and X and D are Hausdorff, and S is significantly less than the minimum of R and X. So again, S is always smaller than R. It's always smaller than X. And we're saying it's significantly less than that, so it's non-trivial.

Then the number of directions would be less than tilde tilde S squared over X. And you see that these conjectures again match up. Cool. OK. Now I mentioned on the first day that one of the fundamental problems of projection theory recently had been solved. This is the problem. So this conjecture has been proven. So this is a theorem from 2024 of Orponen, Shmerkin Ren, and Wang. Conjecture is true.

AUDIENCE: In the first, the plan, it is still open?

LAWRENCE GUTH: Yes. Thank you. Yeah. Conjecture R is true. Finite field is still wide open. Yes. Yes. Cool. So actually, another reason I was excited to do projection theory this semester is that a bunch of the architects of the recent developments will be visiting. Pablo Shmerkin is going to be here for about a month. He's going to talk in the analysis seminar. I might rope him into talking in this class.

Kevin Ren is going to visit and give a talk. Hong will be here at some point, although maybe not during classes. Yeah. Let me say something very briefly. So these two conjectures, I presented them as being parallel to each other. I think that's fair. And they're both quite difficult. This one people have managed to prove, which is exciting. And this one is still open.

So the difference is that in this problem, there are many different scales. We've seen the many different scales in what we just did with double counting. We had to think about all these different scales. And we're going to do it again with the Fourier method. So far, you might have the impression that these different scales are just a nuisance. They make simple arguments more complicated and the expressions longer. But actually, having these many different scales will turn out to be a help.

So they play a key role in the proof of this that, somehow, because of all these scales, if something goes wrong, something has to go wrong at many different scales. The argument also has different cases, depending upon whether the sets look like-- set have spacing like number one or spacing like number two. There are different things to play with. They're both helpful in different ways. And none of that structure is available here, so probably completely different ideas would be needed.

OK. Good. So the second half of the class, the goal is to do the Fourier method in the real case and build up towards doing this. Let me remind you first of the main thing we proved using Fourier analysis over finite fields. And then we'll set up a version over the reals. OK. So part two of the class, Fourier method.

OK. So the main lemma that we proved over finite fields is like this. So let's say fancy L is a set of lines in fq squared. And f is a sum over all these lines of the characteristic function of the line. OK. If that's true, then we can break up f into two pieces, a zero frequency piece and a non-zero frequency piece. So then f is equal to F0 plus f-- we called it f high. Zero frequency. High frequency.

So the support of f0 hat is just zero. In other words, f0 is constant. And we can find out exactly what the constant is. And based on that, we can say that the L2 norm squared of this is the number of lines squared. Then the support of the high frequency part is the complement of zero. And the high frequency part has an L2 bound. L2 squared is L times q.

OK. I remember teaching intro to real analysis class once, and so I tried to answer the question on the first day, what is real analysis? So there are several. There are lots of different things you might say, something about convergence and infinity or whatever. But another thing I proposed is analysis is complicated equals simple plus small.

Study something complicated. The real analysis point of view is to try to write it as something simple plus something small. So this is our function that maybe we feel is complicated. We don't understand it that well, yet. We wrote it as something simple. This is a constant function, which makes it quite simpler than a general function. And then this bit-- and the point here is that this is smaller than f0 or f.

Now that only becomes true, and this only becomes at all interesting, when the number of lines is much bigger than q. When that happens, this L2 norm is much smaller than that L2 norm. So complicated equals simple plus small. OK. So our goal is I will describe to you the version of this lemma over the real numbers.

OK. So instead of a line, so an infinitely thin line, what we're going to talk about is a rectangle, or I usually call it a tube. So capital T for tube is a 1 by R rectangle in the plane in any direction. OK. And so basically, we're going to add up characteristic functions of rectangles. But when you do Fourier analysis, characteristic functions of rectangles-- characteristic functions in general are not super nice.

Because of the discontinuous sharp cutoffs, their Fourier transforms are not very well behaved. So what we often do is we work with a smooth approximation of the characteristic function. So c sub t is a smooth approximation of the characteristic function of t. OK. So another thing about this step of doing Fourier analysis on R2 instead of finite fields is that certain things are a little bit messy.

For example, it's a little bit messy to write down the formula for c sub t. So I'm going to try the approach of not doing the messy details in class. So in class I'll give you an overview. And then on the homework. I'll do all of the details, and you can check the things and make sure you understand how it works. And I do that partly because I feel like sometimes the messy details obscure the main trajectory of what we want to do.

So it's good to first learn what you want to do, and then, once you know where you're trying to go, fill in all the little steps to get there. So for this class, I won't say anything more precise than c sub t is a smooth approximation of the characteristic function of t. OK. So now here is our main lemma in the real case.

OK. So instead of L is a set of lines, t is a set of 1 by R rectangles in the plane R2. And then f is the sum over all the rectangles of this smoothed characteristic function, and that plays the role of the function f up there. OK. And then we're going to break this function up into pieces. There are going to be low frequency pieces that have the advantage of being kind of constant. And there are going to be some high frequency pieces that have the advantage of being small.

But the difference between the real world and fq is that fq, there are only two kinds of frequencies. There was zero and not zero. But in the reals, there are a lot of frequencies. There's large, medium, small, very small, and all of those, they're each going to be there. OK. So instead of having two pieces, we will have many pieces.

So f is going to be the sum of f little r. Little r goes from 1 up to capital R and is dyadic, just like in theorem one. So we'll have log of r many pieces. And these are the different scales of the frequencies. OK. We'll do this so that the support of Fr hat is contained in the ball of radius 1 over little r. So that's the frequency information. And we will say something about the L2 norm.

OK. Now, our bounds on the L2 norm before only depended on how many lines there were, not about how the lines were arranged. But that won't be quite true anymore. Like we've seen in other places, it will matter how these tubes are clumped. So if you have many tubes that are clumped together inside of a-- yeah. OK. So let me put the other part of the definition.

So N sub r of T is the maximum over T tilde, which is like a two little r by two big R rectangle of the number of T in our set of T so that T is contained in T tilde. So here's a picture. So here is T tilde. This axis is little r basically, and this axis is big R, basically. Maybe it's good to put two big R. And inside of here there may be a bunch of R2's.

So just like-- oh, and I guess to be consistent with before, I should actually do this. So N sub X of r measured whether many of the balls in X were clumped together, and N sub T of r measures whether many of the tubes of T are clumped together. OK. So the bound will depend on that thing. Formula is-- OK. N sub T of little r times the number of tubes times little r inverse times big R.

OK. So you can see that I don't remember this formula that well. But one thing to notice about it is that as this number gets bigger, then our bound for this gets bigger. I also claim that this tracks with the finite field case. So if you plug in little r equals 1, that's the highest frequency part. N sub T of 1 is going to be just 1. If the tubes are disjoint, then-- or the tubes are distinct, then.

Anyway. So then if you plug in little r equals 1, this number will match with the high frequency case. And if you plug in little r equals big R, then this will match with the zero frequency case. But now we also have all this in between stuff. OK. Any questions or comments about this statement?

So I'm going to do a proof sketch. Proof sketch of main lemma are details on homework two. How's homework one going? If you're interested, there are office hours after class today and after class on Thursday. OK.

OK. There's a Fourier analysis proof. And an important character is, well, we'll take the Fourier transform of each of these CT hat. And the support of CT hat is going to be contained in T star, where T star is defined as follows. It's the set of ksi in R2 so that the dot product X1 minus X2 dot ksi has norm at most 1 for any X1 and X2 in T. Here's a picture.

So in X space, here is our tube T. And I take a couple of points in it, X1 and X2. So our ksi is going to have a small dot product with any vector like that. And so the dimensions of this T are-- this is 1 and this is R. OK. Now in ksi space, T star is going to be perpendicular to T. There's T star. And it has length 1 and width like 1 over R.

So geometrically, the idea is that if I take my vector here, my vector there, the long component of this one has size R and it gets paired with the short component of that one which has size 1 over R, which gives me around 1. And the two short components-- the short component here has size 1, which gets paired with the long component here, which has size 1.

OK. And this is analogous to the lemma we had in the finite field case that the characteristic function of a line hat was supported in the line perp, which is like star. OK.

OK. So that's the first important ingredient where the Fourier transforms live. And I guess you're doing the homework now where you're checking what line hat looks like. And then in the next homework, you'll check what CT hat looks like, and you'll see it supported like this.

OK. So now I'll tell you who f sub r is. So let's next do-- so the f sub r comes from Littlewood-Paley theory, where you can take any function and you can break its frequencies into dyadic pieces, the Littlewood-Paley pieces. That's what we're doing here.

So Littlewood-Paley decomposition. So I'm going to write 1 as a sum on my dyadic little r's of eta r of ksi. So this is a partition of unity in Fourier space. It's a partition of unity, so these are all positive. And the support of eta r is contained in an annulus where the norm of ksi is at most 1 over r and at least 1 over 10 r. And that's true for r strictly between 1 and r. And at the endpoints, we need to have not just an annulus but something a little bit bigger to cover all the frequencies.

So the support of eta capital R is contained in the ball of radius 1 over R, and the support of eta 1 is contained in the set of ksi, where the norm of ksi is bigger than 1/10. All right. All right. Oh, I guess maybe I should mention. One thing that we see from this is that the support of f hat is contained in the unit ball, because every T star is in the unit ball.

OK. So now I can say who fr is. So fr is just taking f and multiplying on the frequency side by eta r. So I take f hat. I multiply it by eta r and I take the Fourier transform back. So this is also equal to f convolved with eta r inverse Fourier transform. OK. And because this is a partition of unity, that tells me that f is, indeed, the sum over r of fr. So we can check off our first conclusion.

And from this formula, this here is the Fourier transform of fr hat. fr hat is eta r times f hat. And that tells us that the support of fr hat is contained in the ball of radius 1 over r. So we can check off that part of the conclusion. And then the difficult and interesting part will be to figure out the L2 norm of each of these pieces fr.

OK. Now this operation of localizing frequencies to an annulus around 1 over r, you can do it with any function, not just with f. So you can also do it with each of the psi T one at a time. So psi T comma r would be I take psi T hat, I multiply it by eta r, and I take the Fourier transform back. And so FR is the sum over all the T in my set of tubes of psi Tr.

OK. So the next thing I'd like to do is to visualize together this function eta r check and this function psi Tr. Do a visual of eta r check and psi Tr. OK. So the function eta r itself, eta r of ksi, that's around 1 on the annulus of size 1 over little r.

So when I take its inverse Fourier transform, eta r check of x will be around-- it will be around 1 over r squared on most of x less than r. So the biggest this could possibly be is integrating this would be the area of this annulus, which is 1 over r squared. And then, because this is smooth, once I look at x's that are significantly bigger than r, I start to get a lot of cancelation in the integral.

Let me write down the formula for this. So eta r check of x is the integral e to the ix ksi eta r of ksi. So if x is smaller than r, there is no significant cancelation in this integral because the ksi is around 1 over r. But then when x is bigger than r, there starts to be cancelation in this integral. And because this is smooth, the cancelation is quite dramatic. So this is rapidly decaying. x bigger.

OK, so that's what the size of this function is like. But also there's some cancelation in it. So if you integrate this function. Well, if you integrate that, you get eta r of 0, which is 0 because eta r is supported on an annulus and not a ball. So this function has some cancelation in it. It averages to 0. OK.

So I can graph it. Ksi sub r check is radial. So here is the radius norm of x and here's eta r check of x. And it would look something like this. And you have to do this carefully, but the negative part is going to balance the positive part to integrate to 0. So imagine rotating this. And this scale here, that's r.

OK. Now these different exponents here work out nicely and a different way-- a way to describe it is that if I take the integral of the norm of eta r check of x dx, well, it's around 1 over r squared times the area of the ball of radius r, which is 1. That integral is around 1. So if I were to take some function and convolve it with the absolute value of eta r check of x, that would roughly be the average value of f on the ball around x of radius r. The weighted average is not exactly the standard average, but that's roughly what it is.

AUDIENCE: So is it the case that convolution with a function on the annu-- like, if eta r is on the annulus, basically behaves the same as if it were on the ball of radius.

LAWRENCE GUTH: So the question is, how would we compare an eta r that's supported on an annulus with a similar function that was supported on the whole ball? Yeah. It's mostly similar. Because of the annulus, there's more cancelation. So if we did it on the whole ball, we wouldn't have this integral property. There may or may not be a little negative part if we did it on the whole ball. But the negative part would probably would be dominated by the positive part.

OK. So that's what the function eta r looks like. And now that we understand that, we can understand the function psi Tr because psi Tr is equal to psi T convolved with eta r check. OK. So the norm of psi Tr of x is around r inverse times the characteristic function of a thickened tube. So if I take the r neighborhood of T.

So if this is T I could take a fat neighborhood of that that looks like this. So that's the r neighborhood of T. So psi T itself is basically supported on T, although it's smoothed out a little bit. And then when I take the low frequency part, the 1 over r frequency part, then it gets spread out and it gets smaller.

OK. And it also has some cancelation in it. So if the graph of this function mostly depends on the transverse direction, and then it looks a lot like this. So this would be the transverse direction and this would be the height. So here I'll put that in the picture. So this transverse direction maybe I'll call y. And then we could make a graph where here we put y, and here we put psi Tr of y.

And we would see the same thing. Positive. Negative. And this length scale here is r. The negative part would cancel the positive part, because if you convolve something who integrates to 0 with something else, then the whole thing will integrate to 0. OK. So that's what these functions look like. OK.

All right. So before I erase this, maybe I'll say. So we mentioned earlier that if we have a smooth function supported on this tube, we take its Fourier transform. It's essentially supported on this dual tube. That proof is a slightly fancier version of what we were just talking about.

So we were just saying that if you have a smooth function on a certain annulus or a certain ball, then its Fourier transform is essentially supported on a dual ball and then its rapidly decaying outside. So this first statement is exactly the same as that, except that we have a tube instead of a ball. Having a tube instead of a ball makes it a little-- there's a little bit more writing or thinking involved, but it's the same idea. The details will be on the homework. OK.

OK. So in the finite field case, the heart of the argument, the key thing called the key lemma was an orthogonality lemma that said that the high frequency parts two different lines are orthogonal. And something like that is true here. So lemma on orthogonality. This is like the key lemma. So it says the following. T1 and T2 are tubes, 1 by R tubes.

Then the psi Tr's are usually orthogonal. So if I take the inner product of psi T1r with psi T2r, it's usually really small. R to the minus 1,000, something that's really small compared to everything in our proof, unless the two tubes are quite similar. And the way we'll say this is there exists T tilde, which is just basically a little r by big R tube. I'm going to magnify it just a little bit that. That contains them both. T tilde and inside of here. And they have these T1 and T2.

OK. So we could have some intuition about this from our visualization of psi Tr. If I were to take T1 and T2 in the same T tilde-- so this thickened neighborhood, notice, is basically the same as T tilde. If I pick T1 and T2 corresponding to the same T tilde, then when I thicken them both, I get basically the same rectangle.

And in fact, then the functions psi T1r and psi T2r are basically the same function, so they are not at all orthogonal. And that's the only thing that can go wrong if the thickened neighborhoods are disjoint or certainly orthogonal. And the more interesting part is that if the thickened neighborhoods are transverse to each other, they are orthogonal. And we can check that by Fourier analysis.

So proof sketch. If the angle between T1 and T2 is bigger than maybe r to the epsilon 1 over R-- or little r over big R, then the support of psi Tr hat intersected the support of T1 of psi T2r hat is empty. So the support of psi T1-- so this is the support of psi T1 hat intersected with the annulus of size 1 over r, intersected with the support of psi T2 hat. And these are two rectangles that are transverse to each other. And so you can check the intersection is empty.

AUDIENCE: Why does it rule out the translation for the rectangles if they're [INAUDIBLE]?

LAWRENCE GUTH: Yeah. And so that's case one. So the other case is if the R neighborhood of T1 is disjoint from the R neighborhood of T2 or a little bit more than r, then we use the formula psi Tr is psi T1. So then we have here integral psi T1r, psi T2r, complex conjugate, is the integral of psi T1 convolved with eta 1 and psi T2 convolved eta r check. And these two guys have disjoint support.

OK. So now we are ready to estimate the L2 norm. So our orthogonality lemma says that these different psi Tr's are usually orthogonal with a few exceptions. And we're going to use that orthogonality to estimate the L2 norm, but we have to keep track of how many exceptions there are.

All right. So L2 estimate. So the last thing we need to check is fr L2 squared. So fr is the sum over all of our T of psi Tr. L2 squared. So that's this thing inner product with itself. So if I expand it out, it's the sum T1, T2 of the inner product of psi T1r, psi T2r.

OK. And the key point is orthogonality lemma tells us that most of the inner products in this sum are negligible. There's only a few that are not. Let me give a bit of a name to the ones that are not. So this scenario here, I'm going to abbreviate it as saying that T1 is related to T2. T1 is close to T2. I guess we could put a little r here.

So this inner product is negligible unless T1 is related to T2. OK. So let's look at our big sum of inner products. And I'll break it up into two parts. It's the sum where T1 is related to T2 inner product, plus the rest of them. And the rest of them are really tiny by the orthogonality lemma, so the rest of them are negligible.

OK. Now each of these inner products, all arithmetic, geometric mean it. So that's less than the sum T1 related to T2 of psi Tr L2 squared psi T1r L2 squared plus psi T2r L2 squared.

OK. So I want to know how many times does each of these functions psi T1 appear in this sum. And the number of times is the number of different T2's that are related to T1. And that number is basically the number that appears in our main lemma. This thing that we were counting here, how many tubes lie in a little r by big R big tube? That's exactly how many tubes T2 are related to T1.

So this thing. You can say the number of T2 so that T2 is related to T1. This is bounded by Nr of T. So this whole sum is bounded by Nr of T times just the sum on all the T and T of psi Tr L2 squared. Cool. And then finally, each of these functions we've already described. So hopefully I didn't erase it.

Yeah. So we described it. It has height at r inverse, and it's supported on a little r by big R rectangle. So we can compute its L2 norm. So this is N sub T of r times the number of tubes times the height was r inverse. So that's r to the minus 2 times the-- so this is the amplitude squared. And this is the area for this psi Tr. And that was the thing that was on the right hand side.

OK. Let me summarize the proof for the last minute. Let's look back at our statement here. So if we look at this sum of tubes, we're going to break it into different frequency pieces. Littlewood-Paley theory. And then the high frequency pieces are usually orthogonal to each other. If you look at two tubes and you look at the frequency 1 over r part, they're orthogonal to each other unless those two tubes are in a common fat tube, little r by big R. And so you take advantage of that orthogonality where it's available and it gives this bound.

OK. Good. Let's stop there for today. We will wrap this up on Thursday and then discuss how similar ideas apply in a different context in sieve theory. In number theory.