[SQUEAKING]

[RUSTLING]

[CLICKING]

LAWRENCE GUTH: OK. Well, we can get started. So in the first class we talked about what is a projection estimate and what is projection theory. And I tried to give a little bit of motivation for it. And today we'll start getting our hands on things and looking at examples and proving theorems.

Now, the cleanest setting in which to get started is finite fields. And although most of our class will be working in RN, I'm going to start there because the proofs are very clean and the examples are very clean. And it is also interesting in itself.

OK.

OK. So here's our setup. So we'll have Fq, the finite field with q elements. And then theta, so then I'll have a map pi sub theta from Fq squared to Fq. These are the different projections from Fq squared to Fq. And it's defined by pi theta of X1, X2, is X1 times theta or X1 plus X2 times theta. So here theta is also in Fq.

So now our setup is that we'll have a set of points, X in the plane Fq squared. And we'll think about projecting this set of points in a bunch of different directions. So then we'll have a set of directions, which is called D. So D is a subset of Fq. And you think of it as a set of directions.

We're going to imagine taking this at X and projecting it in all of these different directions that are in D. And we'll see how big the different projections are. And the thing I'll measure, I'll call it S, which depends on X And D, is the size of the biggest projection. So it's the maximum theta in D of the size of pi theta of X.

OK. So one thing we study in projection theory is we'd like to understand for different sizes of X and D, what can we say about S. So let's look at a couple of examples. And then after that, I'll tell you some theorems, what we know, and some conjecture, what we think might be true based on the examples.

So the first example is a square grid or integer-- maybe it's good to say integer grid. So for this example, let's suppose that q is actually prime. You can do similar things if q is not prime. It's just easier to write it down. So our set X is going to be the set X1, X2 where X1 and X2 go between 1 and n. So even though we're over a finite field, you can think of this as a square integer grid.

And then we were saying last time that if you take a square grid, an integer grid, and you do a projection in a direction, that's kind of a rational angle, then you get a small projection. So that's what our set of directions is going to be. A set of directions is going to be the set of a1 over a2 where a1 and a2 go from 1 up to some A. So N and A parameters that we can play with. There are many square grids, many choices.

All right. So now let's figure out what is S of X, D. So what happens when I perform a projection of X at one of these angles? So if theta is in D, I'm going to take pi and if x is in X, so if I take pi theta of x1, x2, then I get x1 plus a1 over a2 x2. And I'm going to put it over a common denominator. So it's a2 x1 plus a1 x2 over a2.

All right. Now in this situation, we're imagining that theta is fixed and we're varying little x and big X. So we're going to try to see how big is pi theta of big X. So we're going to see if a is fixed, and I vary x1 and x2, how many different numbers do I have here? So the denominator isn't important because it's not changing. And I want to know how many different numerators I have. And the numerator is going to be an integer. And it has size a times n at most plus a times n. So it basically has size a times n.

So we see that the size of the projection of x is at most a times n. So this less than with a squiggle here, means there's a constant factor that I left out here. It's a factor of 2. But in general, a squiggle just means there's some constant factor left out. And S is the biggest of these. So S of X, D is around A times N.

OK, so then you could do a little algebra to see how S and X and D relate to each other. And if you do that, you see that D is around S squared over X. So that's how it works in an integer grid.

And there's one more example that I thought I would mention, which we also mentioned last time. It's an example that comes from subfields. So if q is not prime, then Fp is an interesting subset of Fq, and Fp squared is an interesting subset of Fq squared. So this is called the subfield example.

So let's say that q is p squared, or a similar phenomenon if q is a higher power of p, but the computations are a bit different. So if q is p squared, our set X is going to be Fp squared sitting in Fq squared. So the size of X is p squared, which is q. And our set of directions is going to be Fp which is sitting in Fq. So the size of our set of directions is p which is q to the one half. And then if theta is in D, and x1, x2 is in X, pi theta of x1, x2 is x1 plus x2 theta. All those things are in the subfield Fp. So this is in Fp.

So therefore, the size of pi theta of X is at most p. In fact exactly p. And so that's our S. And this example coming from subfields is more extreme than the example coming from grids. So if you were to take this value of X And this value of S, and plug them in here to make a fair comparison, in the grids we would have only one direction if we did that, but here we have many directions. So this thing has very small projections in many different directions, even more so than a square grid.

So those are a few examples. And based on that, I can tell you a conjecture that we stated also last time. And also, I can tell you some of the things we know about this conjecture, some of the main results.

So first, I'll say the conjecture, so if q equals p which is prime and we're in this setup. And S is smaller than around X over 2, then we can't do better than the square grid.

So let's explain all the hypotheses in this conjecture. So why do we have this? Well, if you take any set D S of X, D is obviously smaller than the size of X. So if S was allowed to be the size of X then D could be everything. And there would be no meaningful upper bound on D. So we have to insist that it's somewhat smaller than that. You could ask about the range in between here and X, there's something to say but I think not super interesting. So we'll look at this version. OK.

We have this set up. And why do we also add here that q is a prime? Well, we have to add that because of example two. If q is not a prime, you have that example. And we just saw that this is false. And this conjecture is plausible in the sense that there's no counterexample, that no one has found an example that beats the square grid. And basically, no one has found an example that even is, well, OK, there are very few examples that are competitive with the square grid, and they're all small variations on it.

You also might wonder what happens if q is not a prime. And I haven't seen anyone write down a conjecture for that. I think, roughly speaking, what people believe is that all of the interesting examples should be of this type or that type, or maybe you can combine them a little bit by taking a grid of subfields. So it would be a reasonable little project if anyone is interested to try to write down the known examples and make a conjecture about what happens.

All right, now this conjecture is open. But let's start to talk about what we know. Well, I'm going to state some theorems for you and we'll prove the theorems. And maybe more importantly we'll introduce some of the basic techniques in projection theory. So we're going to talk about two techniques today. One of them is a double-counting argument that has a combinatorics flavor. And the other is a Fourier analytic argument that's based on orthogonality.

And this is what they give. So theorem one which is based on double counting, say if we're in this setup and S is smaller than X over 2, then the number of directions is smaller than around S.

Now, if S is around X over 2, this is actually sharp. Hold on. Let me look at my notes for a second. OK, actually I wrote that in my notes, but I'm not quite sure it's true. Let's just go on and I'll tell you theorem two.

So theorem 2 is based on what I'll call a Fourier method, or orthogonality, or something like that. It says, if we have this setup and this time if we suppose that's S is smaller than q over 2, then D is smaller than S times q over X.

So let's digest this statement a little bit. So the projection is always a subset of Fq. So S is always smaller than or equal to q. So this hypothesis says that it's significantly smaller than that. The projection is significantly smaller than the whole space OK. And then we get this bound, and this bound is the most interesting in the case S was actually close to q over 2. Because in that case this q And this S are about the same size. So this is around S squared over X, which matches the grid example. So this is actually sharp. if S is around q over 2.

OK, I find that quite striking and impressive. This is a great example. It has an interesting intricate structure. And we're wondering whether it was sharp and at least in this one regime, this theorem shows that it actually is sharp. On the other hand, S could be a lot smaller than q. Then we get some bound here and we don't have any example that matches that bound.

I guess I should also say this works over any field. So if you plug-in example two it's actually also sharp. And so it's also sharp in example two. All right. But there are a bunch of choices of the parameter. You could make where it's not clear that this is sharp.

OK, so the big goal of the class is to introduce these two techniques really. And we'll use them to prove these two theorems. But maybe it's a good place to pause and check if people have questions or comments. Yeah?

AUDIENCE: In the first theorem, are you assuming that it's prime?

LAWRENCE GUTH: We are not assuming that it's prime. Right, that's right. We're not assuming it's prime in either theorem. And so where it says setup it means we're assuming this. Yeah. That's right. And so that raises a point that it appears that the true answer to this question depends on whether q is prime or not prime. But neither of these techniques depend on that. They work for all q, and they're not able to notice whether q is prime or not prime. That will eventually be a big theme. We won't get to it today, but eventually we'll try to find some techniques that do notice whether q is prime or not prime Yeah?

AUDIENCE: For the premise, S is less than or equal to the size of X divided by 2, how easy is it to construct other examples that satisfy that element of the grid?

LAWRENCE GUTH: OK. So the question is, how many examples can we find which obey the hypotheses, so this theorem says anything about them? Well, I think it's not super hard because we're allowed to take D to be small. So if D is just one direction, then you could easily arrange the points so that their projection in one direction is quite small. And you have a lot more flexibility than a grid. And if you'd like the projection to be small in two directions, then you would want some kind of Cartesian product, but there's still a lot more Cartesian products than integer grids.

And then as you add more directions, it gets more rigid and we can start to think about it. We can think about it later on. It's a neat question how many examples we can come up with. Yeah. Yeah.

OK. Other questions?

AUDIENCE: In the example why you said q equals p prime. Can be p be a prime? [INAUDIBLE].

LAWRENCE GUTH: Yes, right. OK. So the question was up in example one, we assume that q is prime. Are there examples similar to that if q is not prime? So if N and A are smaller, if q is p to the r and you choose N and A smaller than p, you can just use what's on the board verbatim. If you wanted to choose N and A bigger than p, it doesn't really make sense in Fq. So I think that there's a fix for that. But I didn't have time to think through the details before class. Yeah, so homework for one of us.

OK. So let's start proving some things. So first we'll do the double counting. So suppose that you just look at one direction where this projection is small. If S is a lot smaller than X, it will look something like this. So here is the space we're projecting onto and we'll have some points. And this is our projection map. So here, we have X, and this arrow is the projection map pi theta.

And then down here we have pi theta of X. OK. And you'll notice that for each point in the projection or at least for most points in the projection, we will have many points of X that are going on top of it. OK.

So we're going to pay attention to that. And the first thing we'll say is that if S is smaller than x over 2, then we have a lot of points going on top of each other. So for every theta in our set of directions, the number of pairs x1 and x2 in X that got projected to the same point. That should be at least around S times X over S squared.

OK, so let's digest that. First of all, I'll warn you that I changed notation. So x1 and x2 are each points in X. Those are not the components of anything. And so why is this? Well imagine first that each point in the projection has the same number of preimages. So each of these points has about X over S points of X sitting above it. And then the number of ways to choose two of them would be X over S squared. And that would be true for each of these S points. So that would give us this.

Now actually, let me add here that I'm interested in two different points that got projected to the same point. So this is still true. But because I put two different points here, that's where I needed to use this hypothesis. So that was some intuition why it might be true. Then you might wonder, well, maybe they're not evenly distributed. Maybe some of these points have more preimages than others. It turns out that that only helps you. I'm going to make it an exercise. It's kind of a standard combinatorics exercise. So it's exercise on your first problem set using Cauchy-Schwarz.

Cool. So there are a lot of points getting collapsed together in these different projections, and we're going to put it all together. So the key object that we're going to count is the number of directions theta and distinct points in X, so that they have the same projection.

And there have to be a lot of those because for every direction there are a lot of them. So the number of these is at least X squared S inverse number of directions. The X squared S inverse is just that simplified. And that was how many there are in each direction. And there are D directions.

So now the point is that it's hard to have this many collapses because if you take any two pairs of points in X, there is only one direction in which they get mapped on top of each other. So that gives us an upper bound on this star. But for any points x1 not equal to x2, the number of theta so that pi theta of x1 equals pi theta of x2, there's at most one. And that gives us an upper bound for star. Star is upper bounded by x squared. Once you choose x1 and x2, there's only one choice of theta.

So this key quantity is sandwiched in between here and here we have X squared S inverse D is less than star, is less than x squared. And you simplify this a little bit. And we get that D is less than S. And that's the proof.

Let me summarize it without the equations. The key object is this object, star. Let's call let's call this the set of coincidences. A coincidence is when you have two points in X that get mapped to the same point by pi theta. And if many projections are small, there must be many coincidences. But there can't be that many because if you take two points, x1 and x2, there's only one direction that makes a coincidence. And that's what we're using.

Any questions or comments about double counting? OK.

Let me show you the second approach, which is based on I called it a Fourier approach, or based on orthogonality, or based on cancelation. So this approach begins with an image that I want you to have in mind, and that we could try to work on this image in many ways. So it's the overlapping lines image. So suppose I have a bunch of points that have small projections in many directions. So if I project this down, it has a small projection. So maybe I'll call this pi theta 1.

If I look at the preimages of these, I'll get some lines in a certain direction that corresponds to theta 1. And so I've covered my set with not too many lines. So down here there are at most S points. So over here we have at most S lines. And I'll call this set of lines L theta 1.

So to summarize, if pi theta of X is less than or equal to S, then I get a set called L theta is a set of at most S lines that are I guess perpendicular to theta, and that they cover my set x.

So now if I have many small projections, I can do this in many directions. And so I'll have sets of lines in different directions. And I'll cover X many times. So I can do this for every theta in my set of directions. So I'll do one more. I was meaning to bring some colored chalk, but I didn't. I don't have colored chalk today, so I'll do my best. So let's say this is a dotted direction. And when I project in the theta 2 direction, we have this. So this would be L theta 2 in dotted lines.

And I'm not going to add any more directions to the picture, but in your imagination, imagine that this set of points, there are lines going many different ways that are covering the set of points. So in our second method, and maybe in a lot of things that we do, we're going to imagine these lines and work with these lines as a way of understanding the projection estimates.

So let me add a little vocabulary about this. So we've defined L theta. L, out whole set of lines is the union over all the different directions of the lines L theta. So that's a set of lines. And so if L is a line, I'm going to write L of x as an abbreviation for the characteristic function of that line 1 sub L of x

And we think about what happens when we add up all these characteristic functions. So f of x is the sum over all the lines in my set of lines, of L of x. So take a moment to visualize that. I have lines going in different directions. Imagine their characteristic functions and add them all up. So when I add them all up at any given point, it will tell me how many different lines to my set of lines went through that point.

And so if X is in my set of points, then f of x is the number of directions. If I take a point in my set, this point is going to lie in one line in every direction. If I take a point that's not in my set, it's not very clear what's going to happen. Obviously it will be in any lines, but it might be in some lines. We don't know much about what f is doing at those points that are not in my set X. OK, great.

So we're going to get at projection theory by trying to study this function, and in particular, trying to study the set where the function is big. So that's going to be our second approach. And we're going to study the set where this function is big using Fourier analysis and/or using orthogonality.

So orthogonality. So L of x are numbers, the characteristic function of a line. And I'm going to decompose it into pieces. So there are q points in a line. And there are q squared points in my whole plane. And so the average value of L of x is 1 over q I'm going to separate that out. That's 1 over q plus a piece that I'll call L sub high of x. That h is for high. We'll talk about that name more in a sec.

So this is the average value of our function. And therefore, the leftover piece has mean 0, it has average value 0, has mean 0. And I also will think of this as the zero frequency piece and the high frequency piece in terms of Fourier analysis. We can make that precise. We'll do that a little bit later. Depending on how much Fourier analysis you do, you may or may not like this intuition. But this is a constant part, and this is a higher frequency oscillating part OK. Those are just words.

So we decompose it like this. Now the key observation is that, oversimplifying just a little bit, these high frequency parts are orthogonal to each other. So here's a precise statement. So if L1 not equal to L2, are lines, we have two different lines. Then the sum over all of the points in the plane of L1 h of x, L2 h of x is less than or equal to 0.

Now when we prove this lemma in a little bit, we will see that it's almost always equal to 0. So these functions are almost always orthogonal. And once in a while they're not orthogonal. But then it's negative. That part isn't that important anyway. So roughly speaking, these functions L1h and L2h are orthogonal. Now how does that help us?

Well let's think about f of x. f of x, remember, is the sum over all the lines in our set of lines of L of x. So we can break it up this way also. When we add up the constant parts, we just see what we get. We get the number of lines divided by q. And then we have to add up all the high frequency parts. So plus the sum L in our set of lines of Lh of x. So this thing here is the high frequency part of f. I'll call it f sub h of x. For now you can think of that as just a name.

So this part we understand quite well. We know exactly what it is. And so understanding this function boils down to understanding this part. And for understanding this part, we can use this orthogonality. We've written it as a sum of pieces that are each pretty simple and they're orthogonal to each other. And so that lets us figure out or estimate quite well the L2 norm of this fh.

So after the key lemma, I guess there's the L2 estimate. It says this sum x in fq squared of fh of x squared, that's the sum L1 L2 in L of the sum x and Fq squared L1h of x L2 h of x.

And because of my orthogonality in the key lemma, almost all of these terms are less than or equal to 0. Namely, if L1 is different from L2, then this sum is less than or equal to 0. So then I'll just keep the terms where they're equal to each other.

Now this sum here is easy to do because I just have to think one line at a time, and all the lines will be equivalent to each other. So it's just the number of lines times whatever this is. This thing also is not very hard to compute, but let me say that this is actually smaller than if I did the whole L of x squared. I mean, you could just compute it and write down what happened. But anyway if you take a function and you subtract off its mean, then the L2 norm is smaller than for the original function.

So we could write this. And then we can compute this even in our heads. So L of x is 1 at q points, and otherwise it's 0. So this thing here is q. So we get the number of lines times q. So that's our L2 estimate.

OK, so now let's finish the proof of theorem two. There's theorem two up there. We have this hypothesis that S is a lot smaller than q. How does that matter? That controls this term. So if D was only as big as L over q, then this complicated thing wouldn't come into play. And f could be that big everywhere. So let's see how that works.

So the number of lines in my set of lines is S times D because I have D directions and I have S lines in each direction. So this L over q is S times D over q. And I have a hypothesis that's S at most q over 2. So that implies that when I take this ratio, this L over q, is less than half of D. Yeah?

So the quality L equals S times D, does that assume that every set of lines has the same-- every L is the same size? Because I thought that they could be different sizes.

Yes, right. OK. That's a good technical catch. So in each direction I have at most S lines. But it might be OK to work with that. But I find it slightly confusing to keep track of. And you can always add some extra lines. And the lines will still cover the set X.

So I don't know that this is crucial, but what I like to do is to say that there are equal to S lines and they cover X, and we can say that, and then it makes this computation clean. OK.

So this L over q is less than half of D. Now remember that at each point in X, f of x is D. And f of x breaks up as L over q plus the high frequency part. And what we just saw is that L over q is less than half of D. So the high frequency part is the majority of f at those points.

So we conclude that for every x in X, the high frequency part of X is at least half of D. I'll just write greater than D. And now we can compare this with our L2 bound there. So we can say that sum over all the points of f h of x squared, well, that should be at least as big as the points in X times the number of directions squared. So in this sum, some of these points are in my set x. And each of those points, this has size around D. So that contributes this much. Plus there are other points which I really don't know what they're doing.

And then we have the upper bound the L2 estimate that we just proved over there. That is upper bounded by the number of lines times q, which is S times D times q. And now we just rearrange this and it will give us some information about S And D And X and q. So I'm going to get the D by itself. So I divide by that and I divide by that. And I'm left with D is bounded by S times q over x. There it is. Yeah?

AUDIENCE: I just am going back and forth between that [INAUDIBLE]. The mean value here is taken over like we're talking about the mean over the points x?

LAWRENCE GUTH: Yeah, when I say the word mean, so the question was when we said mean value the mean over what. So let me put that up here. So if I have a function g Fq squared to C, then when I say the mean value of g, what I mean is 1 over q squared sum x in Fq squared of g of x. So it's the mean over the whole vector space.

Other questions or comments about the proof? OK.

I am also working on digesting the proof. So the proof is all about breaking up f as this constant part plus this high frequency part. And the proof is saying that in a certain sense the high frequency part is actually small compared to the constant part. So I wanted to draw that out because I feel like it helps my intuition about what the proof is saying. So let's look at this L2 estimate.

Our L2 estimate was that the sum of h of x squared x in the whole plane was bounded by the number of lines times q. So I want to think about what this means. And it helps me a little bit to replace the sum by an average. So I'm going to divide both sides by q squared. So if I take 1 over q squared, sum x and Fq squared of h of x squared, then that's bounded by the number of lines divided by q.

So this is an L2 version of an average. And by Cauchy-Schwarz, that controls-- let's see. So if I took a square root to get an L2 norm, I would have a square root over here. And by Cauchy-Schwarz, I would get the L1 average has a bound. And the bound is now the square root of number of lines over q. Does that look right to everybody?

So a bit informally, there might be some points x that are different. But on average for a typical point x, the high frequency part only has the size square root of L over q. Whereas the low frequency part, the constant part has size L over q at every point. Here let's call this constant part f0, that 0 frequency part.

So for comparison, the size of f0 of x is always number of lines over q. So the conclusion is that if the number of lines is much bigger than q, then this is way bigger than that. So then the high frequency part is much smaller than the constant part, on average.

So it's telling us that whenever I add up a lot of lines, it has to be significantly more than a q lines, then that sum is pretty close to being a constant function. At a typical point, the difference from being constant is extremely small, and there may be a few points where the difference from being constant is bigger, but there's a pretty strong bound on that.

Any questions or comments about this orthogonality proof? Yeah?

AUDIENCE: Could you go over again where the inequality h of x is greater than 1/2 D? So how do you get that from the previous slide?

LAWRENCE GUTH: Yeah. So to recap a little bit, we have these lines going in each direction. And we have D directions. And we add up all the characteristic functions of the lines. So that's called f. And each point of x is in all is in a line in every direction. So f of x is D. Now we broke up f into two pieces-- the constant piece, which was number of lines over q, and then the mean 0 piece or high frequency piece.

And I claim that the constant piece is small. The constant piece is less than half of D, and that would mean that the high frequency part should be more than half of D. And why is the constant piece small? Well, this L is the number of directions times S. And we plugged it in. And we used the fact, that's S at most q over 2. It's one of our hypotheses. That's where it fell out. OK.

Now, I called this early on the Fourier method. We didn't actually mention the Fourier transform or Fourier series of anything. We didn't have to. We did use orthogonality crucially, which is an idea that I first saw in Fourier analysis. But it is possible to rewrite this proof using the Fourier transform. And I think it's worth doing even though we didn't need the Fourier transform. For this, as we keep going and do more sophisticated things in this spirit, it helps to have the Fourier transform. OK.

I guess the other thing is we didn't prove the key lemma, right? So I'm going to prove the key lemma. And I'll show you two proofs. One of them is direct and the other one is based on the Fourier transform.

So I think we can erase this board here. We're just going to talk about the key lemma. So I'll call this the direct proof of the key lemma, and then we'll also have a Fourier proof. So actually two cases in the lemma. So the main case is that L1 and L2 are not parallel.

And in this case, the high frequency parts are orthogonal. So sum of x L1 h of x L2 h of x, in this case, it's exactly zero. The less than 0 comes from the parallel case. OK, now there's a lot of symmetry in this problem. And so if we have the two lines that intersect, we can make a linear change of variables in a translation. So that one of them is the x1 axis and the other one is the x2 axis. So after change of variables, L1 is actually I'm going to say L1 is the x2 axis and L2 is the x1 axis. I know that looks a little silly, but somewhere there's going to be a-- it's not important.

So over here is L2 and over here is L1. All right. Now when we do that, if I take L1 high of a point x1, x2, it only depends. So the function L1 of x only depends upon x1. If x1 is 0, L1 is 1. If x1 is not 0, we get 0. Then we're subtracting a constant function from that. So it still only depends on x1.

So this is a function g of x1 and its mean 0, so the sum over x1 and Fq of g of x1 is 0. Similarly, L2 high of x1, x2 only depends upon x2 and it's g of x2.

This is the trickiest bit. Let me make sure that looks OK. So remember L1 is 1 if x0 is 0, and it's 0 otherwise. And then we subtract. So that's a function that only depends on x1. We subtract off its mean. We get a new function that only depends on x1. I guess it's 1 minus 1 over q if x1 is 0. And it's minus 1 over q if x1 is not 0. That's what g is.

So now when we sum L1 h of x1, x2, L2 h of x1, x2. Well, it's the sum over x1, x2 of g of x1, g of x2. And now we can just break up this sum. So it's the sum on x1, of g of x1, times the sum on x2 of g of x2. It's 0 times 0.

Does that look OK to people? OK, then there's the minor case that they're parallel. I don't have anything super instructive to say about this case. They might as well both be parallel to the x2 axis. It's not hard to compute it by hand, and you'll see that it's negative. So I'm going to not do it on the board because I think it's not very illuminating.

So that's one way to prove the key lemma. But the whole setup also has a nice relationship to Fourier analysis. And we can also use Fourier analysis to prove the key lemma. And if we do it that way, we can justify in a precise way saying that something is the zero frequency part or the high frequency part. OK.

So let me remind everybody how Fourier analysis works over finite fields. So I have a Fourier analysis based on having some group homomorphism from your group into C star. So I'm going to call it E. So E is a map from Fq to C star. Think of it as Fq with addition. So this is a group, an abelian group. That's an abelian group, complex numbers with multiplication. This is a homomorphism. And I want it to be non-trivial, to not just send everything to 1

So for example, if q equals p is prime, then it would be e of x is e to the 2 PI i x over p. And if q is not prime, it takes a little bit more doing to write down a homomorphism. But there is one that's not true.

So now suppose I have a function on Fq to the D. It has a Fourier transform. So the Fourier transform of C is defined to be sum over q and Fq to the d, f of x, e of negative x dot C. So there's x. C. This is a dot product. It's just x1c1 plus, x2c2. So that's the Fourier transform. It looks more or less like how it looks in Euclidean space.

So this C also lives in fq. There are two main really fundamental theorems about Fourier transform. One of them is Fourier inversion and the other one is Plancherel theorem with orthogonality. And in the finite field setting, they work out in the following way. So theorem one, Fourier inversion f of x is the sum xi in Fq to the d. And here we need to take an average. f hat of xi e of x dot xi.

So it looks like Fourier inversion in Euclidean space, although there's this average here. Maybe worth saying that there are several reasonable conventions that you might put about it. Somewhere in this setup, we need to divide something by some powers of q And there are several reasonable conventions, and you see different conventions in different places. But anyway, this is the one that I picked. And I never remember it. So I always take the function 1 and compute its Fourier transform and check that this works. And that tells me I have this factor right.

And theorem two is Plancherel. So that says that if I take the sum over Fq to the d of f of x, g of x bar, then that's equivalent to a sum on the Fourier side. And there's also a normalization. So that's 1 over q to the d sum x and sum xi and Fq to the d, f hat of xi g hat of xi bar. So again, if you compute a simple example, you'll see that this is the right normalization.

OK, so now if you take the characteristic function of the line, you can take its Fourier transform. And something interesting happens that represents an important connection between lines in the Fourier transform and therefore between projection theory and the Fourier transform. So here is what happens. So suppose L is an affine line in Fq squared, or really the characteristic function of the affine line

So then it has a perp, so L perp is the set of xi so that xi dot x1 minus x2 is 0 for any x1 and x2 in the line L. So in Euclidean space this would look like this. To here is my line L. This is x space. And then over here in xi space I would have L perp. Notice that L perp always goes through 0. So here is my x1 and x2. So this direction here is a tangent direction to L. And then these are the perpendicular directions.

All right. So proposition, if you take L hat of xi, we take its absolute value. There are two cases. One of them is that you get q And that's if xi is in L perp. And the other one is you get 0. And that's if xi is not in L pip. So the Fourier transform of a line, first of all has is very sparse. It's 0 almost everywhere. The only places where it's not 0 are on this perpendicular line.

So I think it is a good experience to compute this out and check that it happens. When you write down this sum, this sum will reduce to a geometric series. And then you can sum the geometric series and you can see exactly what happens. So this is an exercise on the homework.

OK, so we'll also do this thing of subtracting off the mean. And that has a very natural interpretation in Fourier analysis. Let's erase this.

So if you have a function g from Fq to the d goes to C, then we can write g as 1 over q to the d sum over xi g hat of xi, E of x dot xi. And one of these frequencies is often very special, namely the frequency zero. So I'm just going to break up this sum. It's 1 over q to the d g hat of 0 plus the sum over the non-zero xi 1 over q to the d g hat xi of x dot xi.

All right. So if you think about what this is, g hat of 0 is just the sum of the values of g, and then I divide by that. I get this formula here. That's the mean value of g. So this is g0 and it's the mean value of g. And this thing here we'll call g high. g high is g minus its mean value. So it has mean 0.

So that's what we did with our lines. This is the zero frequency part of L. And this is the nonzero frequency. Or I called it high frequency part of L. So from the Fourier analysis point of view, that's what these things precisely mean.

Now we can put this together to give a new explanation of the key orthogonality lemma. Where should I put it? Yeah, so we gave one proof of the key lemma over here. We can do now another one.

All right. So let's give the Fourier proof of the key lemma. We'll still do the main case, which is that L1 and L2 are not parallel. And that's where we get literal orthogonality. Then we get that the sum of L1 h of x and L2 h of x is 0. How will we do it? We'll evaluate it with Plancherel.

So that's Plancherel, theorem two. And now I have the high part here. What does the high part mean? The high part means I crossed out the 0 frequency. So this is 1 over q squared times the sum on non-zero xi, L1 hat of xi, L2 hat xi. So if you compare the Fourier transform of L1 and L1 high, the only difference is the zero frequency. I take the zero frequency here and I throw it out. All the other ones are the same. So this is the same as this.

OK, now I'll make a picture. So here in x space I have a line L1 and I have a line L2. What happens in Fourier space, xi space. Well there's a L1 perp and then there's L2 perp. And L1 and L2 perp intersect only at 0. So here I'm summing over non-zero frequencies.

So if xi is non-zero then either xi is not in L1 perp or xi is not in L2 perp, because those two perps, they only intersect at the origin. But this function is only supported when xi is in L1 perp, and this one is only supported when xi is in L2 perp. So at every xi in this sum, one of these guys is zero. Or to put it a little bit differently, these two functions have disjoint Fourier support. So they are orthogonal.

So these are the two main classical techniques of projection theory, the double counting argument, and the Fourier orthogonality argument. They're each not that difficult, but it actually is really pretty hard to go beyond them, even though there's a lot of situations where they're not sharp. So that's a theme that we'll get to. But before we get to that theme, we'll explore these for a while and we'll see that a bunch of people have used them, in fact, in different areas of math to do some different things, and it's things that I think are interesting.

So the next thing that we'll do is the more analytic setting of working in R2 instead of over a finite field, and having balls instead of points. And I waved my hands about it last class about exactly what we would try to prove. And so in our last 10 minutes, I'll introduce the setup with balls. And I'll tell you what are the analog of theorem one and theorem two and this conjecture.

And then we'll see the proofs of all that next class. And we'll see the proofs are morally the same as these proofs in finite fields, but with a little bit of extra technical issues from working in Euclidean space. OK.

All right. So let's do a setup with balls, setup in R2 with balls. All right. So our set X is going to be a set of disjoint unit balls in a large ball B2R. And then D, our set of directions, I'll think of it as a subset of S1. So this is a set of directions that we'll do projections in. And they will also be disjoint. And what I'm going to put here is that they're 1 over R separated. So you can think of D as a set of angles. And any two angles are at least 1 over R apart from each other. The reason for that is if they are closer together than that, the two projections would basically be the same in terms of what they do to a set of unit balls in a ball of radius R.

OK, so that's X and D. And then our S, S of XD is the maximum over any direction in our set of directions of the size of the projection of X in that direction. So this is a one-dimensional set. It's basically a set of unit intervals. And this absolute value is its length.

All right. So we could ask just in terms of the size of X and the size of D, what can you say about S and how does it compare to the finite field case? And we saw last time that it's rather different from the finite field case because of an example where X is very clumped together. So let me put up here a couple of examples in our new geometric setup.

So example one is a square grid that's spaced out. OK, so it looks like this. It's going to be an n by n grid. And it's inside of an R ball. But I'm going to make it as spread out as possible. So the distance here is about R over N. So that's my set X. And D could be the same as before. It could be the directions with slope a1 over a2, where a1 and a2 are integers going from 1 up to A.

And if you compute how big the projection is, it's the same as before, A times N. That's the same thing we got before at the beginning of class, except that the projection is never bigger than R. If you direct a whole ball of radius R, you get an interval of length R. So it's never bigger than that. So I'll put the minimum this thing.

So there's one scenario. This is an integer grid, which was an important example before. But another thing you can do is you can take these points and clump them way together. So example two is a clump. So I take a ball of radius N And I fill it with unit balls. So we still have the number of balls is around N squared. But they're much more clumped together.

And when you do that, pi theta of X is only around N. You project this set X in any direction, it will be an interval of length N. And so that's true for everything. And if you compare these examples, you'll see that this one has much smaller S's than this one. This example is much more extreme than the grid.

OK, so it turns out that just asking the question, if you know the cardinality of X, and you know the cardinality of D, what can you say about S? We'll answer that question. That's a somewhat interesting question. And we'll answer it. But it's not the most interesting question. There's a phenomenon that's more interesting, which is that as the points get spread out, you get stronger estimates.

So what does it mean that the points get spread out? You need to have some rules about it. How do we make that precise that they're spread out? So being clumped together means there are a lot of points in a small ball. And we can quantify that in the following way. N sub x of r is the maximum over any choice of center of the size of X in the ball around center c with radius r. So you can see that this quantity would be quite large in this example, especially if you took little r to be N. And it would be a lot smaller in this example. So that point is more spread out.

And it turns out it doesn't just matter how our set of points is spaced. It also matters how our set of directions is spaced. So we'll talk about that also. So N sub D of R is the maximum over any-- actually let me give this another letter. I'll call it rho, any arc of length rho in S1. So we have the directions as a subset of S1 of the number of the directions in the arc. So if you pick a particular parameter, a size rho, and you ask how many directions are there in an arc of size of rho. So that's this thing.

So you could describe our goal in projection theory in Euclidean space in the following way. Given some information about this thing and that thing, try to estimate S And D And X.

You can phrase it that way. And then there's a whole world of questions because there's a lot of different things you could assume about N, X, and D. There is one particular scenario, one particular set of assumptions that has been investigated a lot. And it's not the only interesting one, but it seems to be important and a good thing to focus on. So you could say that x is an alpha dimensional set of balls if the total number of balls is around R to the alpha and N sub r of x is less than little r to the alpha.

This numerology is consistent with how fractals behave. And then you can make a similar definition for a D. OK.

People might feel it's not so clear why to focus on this particular set of conditions. We can talk about it. Anyway, so the final thing is that if you take this set of conditions, then we can do a version in Euclidean space of the conjecture in theorem one and theorem two, and all the numbers match up, and we can prove a version of theorem one that looks just like that with this double counting, and we can prove a version of theorem two that looks just like that with Fourier analysis. And we have a conjecture that looks just like that.

And the big recent result that I hinted at in the first class is the analog of that conjecture is proven. So at the beginning of class on Tuesday, I will write all those things down and then we'll do analysis in R2 and we'll prove theorem one prime and theorem two prime. OK, great. Let's stop there. Have a great weekend and I will see you next week.