[SQUEAKING]

[RUSTLING]

[CLICKING]

LAWRENCE GUTH: So I decided to teach 18.156 this semester about something called projection theory. I think that projection theory, although it's not usually highlighted in graduate real analysis, is a fairly fundamental thing in analysis, a little bit like the isoperimetric inequality or the Sobolev inequality, maybe not quite that fundamental, but a fairly basic thing that comes up in a lot of different places.

And I was excited to teach about it this semester for a few reasons. So one is that some of the core basic questions-- one of the core basic questions in projection theory was just figured out in the last couple of years. And it's something that I'm excited about learning about. So I wanted to-- so that was one reason I was excited to teach about it. And the other reason is that I have gradually seen it appearing in more and more different places. And so I wanted to highlight the way it comes up in a bunch of different parts of math.

So today in the first class, I'm going to tell you-- I'm going to do an overview of projection theory. I'm going to tell you what it is, what the questions are, and about some of the different applications, and maybe a little bit about some of the different ways we think about it.

I guess it's a little crowded in the back. Are we-- I guess the seats are full. I'm going to start slowly. So if you want to grab a chair or anything, feel free.

So the setup is that we are in Rd, or maybe also a vector space over another field later. And then we have orthogonal projections. So V in Rd is a subspace. And pi sub v is the orthogonal projection.

So what we're going to do is we're going to consider some set, x in Rd. And then you can take the orthogonal projection of x in many different directions for many different v. And we'll study how that all fits together. Study x and this guy for many different V.

So a slightly corny metaphor for what we're doing is we're studying some set. And you can look at the set from many different directions. And we want to put together the information from these different points of view.

So let me make a picture. Red. So we'll come back to this picture a bunch. So maybe x is a set of red dots. And then we can take the projection of this set x onto different planes. So if I draw sort of a-- I'll draw one plane over here. And if I project x onto this plane, then I will get this. So there is pi v1 x.

But I have rigged this-- so if I project it onto a different plane, something a bit interesting will happen. Here is a different plane, v2. And if I project my set onto v2, all of those points will kind of line up on top of each other. And all of those points will kind of line up on top of each other. And all of these points kind of line up on top of each other.

Here are my observations about this situation. So x itself is 8 dots. And when I projected x onto v1, it was still 8 dots. But when I projected it onto v2, there were only 3 dots, because there were a lot of coincidences where dots got pushed on top of each other.

Now, intuitively, v2 is kind of special, and v1 is more typical. So most of the v's will be like v1. But there'll be a few exceptional ones that are like v2. So intuition is that a v2 type situation is rare. And one of our goals is to make that precise. Make a precise question, and say precisely how rare is that behavior.

So informally, if the set x is kind of spread out, most of its projections will also be kind of spread out. There could be a few projections where x gets kind of squished together. And a projection estimate means we'll try to quantify how often that happens and show that it's rare.

So we will make this precise in a bunch of different ways. And I think the most simple to ask question is when the set x is a finite set. So now let's suppose that x is a finite set. And I'm going to use a notation that absolute value of x is the cardinality of x.

So here's my setup. I'm going to take x in R2. And E sub s of x is going to be the set of v. Let me actually-- let me mention another piece of notation. So we're going to talk about all these different subspaces v. So there's a name for the family of all subspaces. It's called the Grassmannian. So another piece of Grassmannian notation is that the Grassmannian of k planes in Rd. This is the set of v and Rd, where v is a k-dimensional subspace.

So E sub s of x is the set of v. These are going to be one-dimensional subspaces in R2 with the property that the cardinality of pi v of x is at most s. So going back to our basic picture, if these red dots were just points, then this v2 would be in E sub 3 of x, because pi v2 is only 3 points.

So the first question, if you like, of projection theory would be estimate how many of these directions there are the size of Es of x in terms of s and the size of x.

So what would be an interesting set x to consider? What would be a set x of points or red dots that in many different directions things would get squished a lot? Do people have any thoughts? What would be an interesting x to look at?

AUDIENCE: A grid.

LAWRENCE GUTH: A grid, yeah. So as an example, let's look at a square grid. So x is going to be the set of pairs a, b, where a and b are integers, and they go from 1 up to n. Good. So this set has small projections in a bunch of different directions. Let me make it a little bit bigger.

So x has n squared points, 16 points in my picture. If I project it down, I only get 4 points. If I project it to the side, I only get 4 points. If I project it at 45 degrees, I also get rather few points. I project it at 45 degrees. The fibers of the projection will look like this. And they'll each have many points, and so the projection itself will be pretty small. Are there other angles as well where I could project it and it would be pretty small?

AUDIENCE: Some rational angles?

LAWRENCE GUTH: Yeah, some rational angles. So if we project x at a rational angle, then pi v of x would be small. I'll draw one more, which I think is enough to get the idea. I remember when I first saw it, I could see the 45 degrees, and the vertical, and the horizontal. So if I project it perpendicular to these orange lines, where it has a slope like left 1 and up 2, all of those fibers will also be large. And so that projection will also be small.

Cool. So we will compute it tomorrow in the overview class. I'm not going to do computations with you. But if you do that, you will find the following thing, that Es of x is around s squared over the size of x for a big range of values of s. So in between x to the 1/2 and, say, 1/2 of x.

AUDIENCE: That's specifically for the grid?

LAWRENCE GUTH: Yeah, for the grid. Erdos did basically this computation in the 1960s. And he made a conjecture that the grid is the shape that makes these the biggest. These had a slightly different language, but his conjecture implied that the grid is the configuration that makes these sets the largest.

And that turned out to be true. And it was proven in the early '80s by Szemerédi and Trotter. So they proved that if x is in R2 and s is at most size of x/2, then the size of Es of x is at most 1 plus a constant times what you see for the grid.

So this can be considered, in a way, one of the fundamental theorems of projection theory. And I'll mention a couple interesting things about the proof. So one thing is the background of these people, Erdos, Szemerédi, and Trotter, these are people who come from combinatorics. So this theorem represents a connection between combinatorics and projection theory. And we'll talk about some other aspects as we go along.

However, this proof is not, at least at the time, was not a traditional combinatorial proof. So they first played around with a lot of combinatorial things, double counting and things like that, and they couldn't prove it. And then the proof eventually uses in a crucial way the topology of R2. So it's an interesting connection between topology and this stuff. So I'll just put here, proof based on topology.

So this story is one important story in projection theory. But it's a very discrete story about finite sets. And there's another more continuous, more analytic story, which will be important in our class, where instead of taking a finite set and projecting it, we take a function, which you can think of as like a density function.

Before I transition there, actually, let me pause here and see. Are there any questions or comments so far? Yeah.

AUDIENCE: So for the theorem that the grid kind of the maximizes Es of x, is that kind of the only example? Or is there a kind of big class of other examples that also reach that bound?

LAWRENCE GUTH: OK, cool. So the question is, is the grid the only example that's kind of tight for this theorem? Or are there lots of other examples? Yeah, so another-- actually, a third thing I like about projection theory is that it's full of open questions. So it is a long-standing open question whether the grid is basically the only example or whether there are lots of others.

In terms of other examples that we know about, there are a few, but I would describe them all as very minor variations on the grid. And we can meet them as we go along. But in terms of a theorem in the other direction, we know very little. So it may or may not be the case that there are many more examples out there. Nice. Other questions?

OK, great. So now, let me describe a much more analytic framework that also involves projections. And we'll see that it's sort of related-- that it's related to this.

Now, instead of having a set, suppose I have a function on R to the d, say a complex value. And I'd like you to visualize this as a density function. So actually, a good thing to picture is a CAT scan. So inside of your body, there's something that they'd like to scan, where are the bones or something. And so there's some density of that thing.

And then the CAT scan will shoot rays through your bodies in different directions. And it will compute the integral of that function along the ray. And so if you stand in front of a screen, maybe you just have an X-ray. X-rays go through. And what you see on the screen is the amount of stuff that that X-ray went through. So that's what we'd like to do with this function.

So suppose here is our subspace v. And if I have a point in it, y, then the orthogonal projection maps this fiber onto y. So that there is pi v inverse of y. So pi v inverse of y will be an affine subspace. And it has a volume measure, which I'm going to call d volume sub F. The letter F stands for fiber. It's just to distinguish it from d volume, which might mean the volume on the whole space.

So then the way I will define the projection of the function of y is I will integrate over the fiber f of x d volume of the fiber. So another way to think about this projection of the function is that if I were to integrate the projection of f on some set u in the fiber, so maybe over here I have some set u in my fiber.

AUDIENCE: [INAUDIBLE] dy?

LAWRENCE GUTH: Yeah, it's d volume of the fiber of y. Yeah, this thing you could call dy. Is that OK? Then this would be equal to the integral over the preimage of u of f of x dx. D-- oh, I'm sorry. You are right. Yes, that should not be d volume of the fiber. Thanks there.

So yeah, this is just dy. Or if you wanted to put d volume of something, it should be d volume in the space v of y, but not in the factor. That's absolutely right. What I wrote made no sense.

So if you integrate the projection of f over this region u. That's the integral of f over the whole preimage. Do you feel like you have an image of what this means? So now we can play the same game. We start with a function. And then you can consider many different projections of this function.

And there's a phenomenon, which is similar to this phenomenon, but the way it plays out in this setting of functions is that most of the projections are smoother than the original function. And I'll illustrate that with a proposition.

So proposition, if f is just an L2 function in R7, then for almost every v, which is a one-dimensional subspace in R7, if you look at the projection of f onto v, it is a C2 function. So this function on R7 is an L2. It doesn't have to be continuous anywhere. It certainly doesn't have to be differentiable. You project it onto various one-dimensional subspaces.

Now, not every one. There are. There are some bad ones, where it's only L2, or maybe even worse. But for most of them, the projection is C2. And if you play around with these numbers, if you make this bigger, you can make it even smoother. So I want to try to-- yeah.

AUDIENCE: What's special about R7 here that makes this true?

LAWRENCE GUTH: 7 is not that special. But increasing this number makes it smoother. I mean, there are similar theorems in any dimension.

So let me try to give some intuition about why that is and to try to convince you that that is a closely related phenomenon to the phenomenon that we started with that appears in the Szemerédi-Trotter theorem. So what's the intuition about this? So I have this high dimensional function, which is quite irregular. And because it's quite irregular, you could picture it having many peaks and valleys and being very jagged.

But then when you integrate it over the fibers, if you integrate both a peak and a valley, they cancel each other out. So a lot of those peaks and valleys cancel each other out. And you could hope to be left with something smoother.

Let's try to make a picture. So I'm going to pick-- I'm going to erase this, but I'm going to draw a very similar picture. So in my picture, my f is going to have-- it's going to have peaks on the red dots, just to help me visualize. So this function doesn't exactly have values. It's a picture of flat landscape, but with lots of little peaks. And so here are some red dots to make them a little bit fatter.

So I think I'm not going to try to draw the function, but what I want you to visualize is there's the plane of the blackboard, and then there's a peak coming out over each of these things. Now, let's try to visualize what happens when I project onto this space. And in yellow, I'm going to draw the function pi v of f.

So to help myself draw, I'll put some little guidelines. So out here, pi v of f is just 0, because there were no peaks. And then it's going to go up as I see this peak. And then as I get here, there are two peaks. So that's even taller. And then here, a little shorter, here a little taller, a little shorter. It looks something like that. You might notice already that this is not as jagged as the original thing. Actually, probably, it would be better if I did something like this.

And because it's the law of averages, if I stick these around, or if I stick them around randomly, then because of the law of averages, this will be pretty steady. And even if I don't pick them randomly, if I choose v randomly, there's a certain amount of the law of averages. And this is going to look relatively steady.

Cool. So qualitatively, this is kind of similar to the Szemerédi-Trotter theorem. It says that there are only a few bad directions. And except for those bad directions, when I project this down, I will see things kind of averaging out. And this proposition is one way to make it precise, but actually, we'll spend some time thinking together about various ways to make this precise. Questions or comments? Yeah.

AUDIENCE: Is it possible sometimes to prove the discrete statements like this based on the continuous versions?

LAWRENCE GUTH: Yeah, good question. So the question is, can we go back and forth between these two worlds? Can we take, say, the Szemerédi-Trotter theorem or something like that and use it to prove some continuous estimates? Or vice versa, can we take some continuous estimates and use it to prove things like this?

So both of those things are sometimes possible. But both of those things have some difficulties. And in particular, there's a big story in the field in the class that people spent a lot of energy trying to translate this into the continuous setting. And the original proof doesn't work. And we had to think a lot. So maybe we'll come back to that a little later in the class. But there's a story I want to share with you about how it's difficult to get from this discrete setting to the continuous set.

Actually, let me say a little bit more. So this example is really pretty typical. So you can suppose that your function is a sum of bumps on different little balls. And so it is just a question about how these different balls, small balls are projecting down. And it sounds, at first, very similar to have a bunch of delta balls and project them down to have a bunch of points and project them down. But the thickness of the balls is really important. It makes things different in ways that we'll talk about. Yeah.

AUDIENCE: Sorry. Would you mind saying a bit more about-- so if I understand correctly, you're saying in the typical case of a continuous setting, you have some [INAUDIBLE]? In the typical case, in the discrete setting, you recover all the points. Could you say a bit more about [INAUDIBLE]?

LAWRENCE GUTH: Yeah. OK, great. So the question is-- so the observation is that in the discrete setting, we had a theorem that says for most directions of projections, the images of the points are different from each other. They stay spread out.

And in the continuous case, we said something else, that if you have a bunch of these peaks, then in most directions when you project, things kind average out and are spaced evenly. And how are those two things connected to each other? Yeah, that's a great question.

So one part of it, I think, is kind of clear. So one part of it I can explain easily, and then this won't be everything. But if my balls were arranged in such a way that when I projected they went on top of each other, so then the projection of my function would be really big here, and then it would be 0, and it would be really big here, and it would be 0, and then it would be really big there. So that would not be smooth. It would be oscillating violently.

So if you have a direction, which is bad in the sense that these balls are landing on top of each other, then it will be bad for this analysis problem. But a set could be bad for this analysis problem in another way. There could be two balls over here, and one ball over there, and two balls over here, and one ball over there, and then the density would still be oscillating a lot, but the image wouldn't be particularly small. So that will require some additional ideas. But it still turns out to be the case that trying to prove stuff about this, the ideas, the ideas from ruling out this case are related. Yeah.

AUDIENCE: It's probably the next-- exactly the problem relating these two, but is there an analog for these distributions?

LAWRENCE GUTH: Yeah. Yeah. There's an analog for distributions. So the question is, what if f was a distribution instead of a function? Maybe it was something messier than being an L2. So yeah, you could still say things. I wasn't planning-- so I think all of those questions you can reduce to something about how balls get projected. And then if you want to, you can say something about distributions.

So the next character in the story I wanted to mention to you is the Fourier transform, that in this continuous setup, there's a very close connection between this projection operation and the Fourier transform. And that brings Fourier analysis into projection theory.

So the connection with the Fourier transform is the following lemma. Well, actually let me make a-- so here's our space Rd. And in there, we have our subspace v. So our function f lives on R to the d. And the projection of f lives on v.

Now, how does the Fourier transform work? So here's Fourier space. So Fourier space is also R to the d. So f hat will live on R to the d. And the Fourier transform of something on v lives on v as well. So in this R to the d, I'll draw here. I'll call it v hat. I guess if we wanted to, we could call this R to the d hat.

So if I have this function, Rd to c, then f hat is a function R to the d hat goes to c. The projection of f is a function from v to c. And the projection of f hat is a function from v hat to c. And then the lemma about how they are related is that if you take the projection Fourier transform of c, it's just f hat c for any c and v.

So let me write down the sketch, the proof of this lemma. It has a simple explanation. So pi v f hat of c is the integral over v pi v f of y e to the minus i xi dot y dy. Now, what is pi v of f? Well, it's the integral over the fiber. So integral over v, integral over pi inverse of y f of x d volume of the fiber e to the minus i xi dot y dy.

And now notice that we have a double integral. We're integrating over the fiber and over y. So if you put those together, you just get a dx. Now, what happens to this e to the minus i xi dot y? So remark that since x is in pi v inverse of y, that tells us that x minus y is perpendicular to v.

AUDIENCE: The inverse is on the pi.

LAWRENCE GUTH: It goes on the pi. Thanks. Over here is y. Here's the fiber. And in our fiber, we have x. So x minus y is perpendicular to v.

On the other hand, we also know that the frequency xi is in v. xi is in this copy of v over here. So x minus y is perpendicular to xi. x minus y, x minus y dot xi is 0. So that means I can change this xi dot y for xi dot x. And I put these volume forms together. I have the integral of f of x e to the minus i xi dot x dx, which is f hat of x.

OK, cool. So that's an interesting connection between projections and Fourier analysis. And it was used classically to solve the problem of recovering a density from the CAT scan. So mathematically, the problem would be you're given pi v of f. For every v, you want to recover f.

And the way you can do it is because of this formula, you can recover-- if you're given-- if you know all the pi v of f's, then you can compute their Fourier transform. So you can compute f hat of xi for all the xi in v hat. Well, you have all the different v's, so you put it together. Now you have computed f hat. So you can find f.

Cool. But this connection with the Fourier transform also can be used to study this continuous projection theory problem. And it's used to prove this proposition here. Yeah.

AUDIENCE: What is v hat? What's v?

LAWRENCE GUTH: Yeah, thank you. So the question is, what's the difference between v hat and v? I'm not sure it was helpful to put the hats. It might be clearer if we remove the hats. So here's another copy of R to the d, where the variable xi lives. And here's the same space v. And maybe it's clearer to remove all these hats.

So we mentioned before that this continuous problem has something to do with what happens when you take balls and you project the balls instead of projecting points. And it's a little tricky to formulate exactly what the question should be about projecting balls. I want to take you through the first steps and just see why it's a little bit tricky and what direction we're going to go in.

So I have projection theory for unit balls, maybe first steps or something like that. So let's say that x is a set of disjoint unit balls in a big ball of radius R in R2. So we could consider the projection in different directions. And we'll have sets of unit intervals, more or less.

And we could ask for the number of directions where something or another happens. Actually, let's step back a little bit. So OK, so we have a set of unit balls. And in this context, I'm going to use the absolute value of x to denote the area of x, or volume in whatever dimension. And hopefully, that will be helpful and not confusing compared with before. If you have a bunch of unit balls, its area is more or less the same as the number of unit balls. So then if we look at pi v of x, that would be the length of pi v of x.

So now I want to make a question like the Szemerédi-Trotter question. So I could look at the set of directions where the projection has length less than s maybe. Maybe I'll do that. So Es of x might be set of v Gr 1, 2, so that pi v of x has length at most s.

But it would now be silly to ask for the number of elements of this thing. This thing is very likely infinite, if it's not empty, because if you make a tiny perturbation in v, it will have a tiny effect on pi v of x. So there's a natural discretization.

So we know that if the angle between v1 and v2 is less than 1/R, then pi v1 of x and pi v2 of x are basically the same as each other. And so it doesn't make a lot of sense to count them separately.

So if we want to count something, one trick we could make is we could have this fancy v be a set of R v's in Grassmann 1, 2 that are evenly spaced. So I just draw-- oh, sorry-- r different lines through the origin evenly spaced at an angle of about 1/R.

And then my new question will be, if x is a set of disjoint unit balls in the ball of radius R and R2, then estimate-- yeah, so actually, let's change this definition. We're not going to count all of the v's, but we're going to count the ones in this set. And that's a more fair comparison to before.

So we're going to estimate the size of Es of x in terms of s and the size of x. And we can compare it to the Szemerédi-Trotter theorem. So we could ask, is it still always true that this is bounded by that expression there? I'll write that down.

Is it true that Es of x is bounded by 1 plus S squared over x? And this is not true. There's a configuration of unit balls where Es of x is much bigger than s. What can we do with unit balls that sort of wasn't relevant before when we had perfect points? What other example might we look at? Yeah, what do you think?

AUDIENCE: You could integrate functions over them?

LAWRENCE GUTH: We could. Yeah, so we can integrate functions over them. So here's a configuration of unit balls. It kind of doesn't have an analog when you're just looking at points. But what we're going to do is we're just going to pack them into a larger ball.

So here's the example. So here's our big ball. And then inside of it, I have a smaller ball. And I pack the smaller ball full of unit balls. So that's x. So the area of x or the number of unit balls is around n squared. And no matter what direction I project this thing in, it's going to fill in an interval of length n, about n. So for every v, I'm going to have pi v of x around n, which is x to the 1/2.

So if I make s equal to x to the 1/2, or a bit less, then Es of x is the whole thing, every direction. And so the size of Es of x is R. It's all of them. It's as big as it could possibly be. Whereas if you plug in s is the square root of x here, you would get 1. So it's actually as bad as possible. The prediction of Szemerédi-Trotter is 1, and the truth is everything.

So these small balls are not really behaving like points because of the fact that they've been packed so densely into a big ball. And so if we-- it is perfectly reasonable to analyze this situation and answer this question. We will do that in this class. There is an answer to it, and it's worth knowing.

However, as you insist that these small balls are spread apart farther and farther, then they start to behave more and more like ideal points. And eventually, this is true. So a goal of this part is to add hypotheses that the set of balls is spread out, and then get stronger conclusions. And probably next class, we will play around with how to state some precise hypotheses and what kind of conclusions we might expect.

But the big progress that I referred to at the beginning of the class in the field is that in the last couple of years, people proved, let's say, the main conjecture about this question for unit balls in R2. And we'll try to flesh it out, because you could imagine many different hypotheses. But I think the theory gives a fairly complete picture of how this works for balls in R2. Different hypotheses about how spread out they are, and what you can conclude about how often the projections have things piled on top of each other.

Now, there are many variations on these questions. And because this one is a little bit difficult to state, I thought I'd mention one other, which is a beautiful open problem. And it's pretty clean to state. And it will let me explain one of the key difficulties.

So these questions make sense over any field. I started with the real numbers, because I was kind of used to it, because I'm an analyst. But you could ask about any field, any vector space. And actually, a very interesting case is finite fields.

So the cousin problem over finite fields. So let's say Fq is the finite field with q elements. And if I write Fp, that will mean p is a prime. So it's Fp of a field with a prime number of elements. And Fq is a general finite field. q will be some prime power.

So here's our question. So now I should say, what is an orthogonal projection? So it's not hard to say what a subspace is. Over finite fields, we don't have, obviously, a notion of orthogonality, but orthogonality isn't really the crucial thing. The crucial thing is just that we have a lot of linear maps from our vector space to a lower dimensional vector space. So we're really looking at all the possible linear maps.

So in terms of what are the possible projections, I can parameterize them just like this. So if t is in Fq, pi t will be a map from Fq squared to Fq. And pi t of x1, x2 is just x1 plus t times x2. So those are a bunch of linear maps.

And actually, you might say, why don't I write all the linear maps? I could have t1 x1 plus t2 x2. Those are all the linear maps. But if you multiply t1 and t2 by the same factor, then you've just taken a linear map, and then composed with a linear map from Fp to itself. So those are all going to be equivalent for us. And so that's why we just set this to be 1, and let this be t.

So now let's say x is in Fq squared and Es of x is the set of t. So the t just live in Fq with the property that pi t of x is at most s. And now we can ask the same question. Question-- estimate the size of Es of x in terms of s and x, ideally, and I guess maybe q.

Now, an integer grid still makes sense over a finite field. And so it works the same way as in Szemerédi and Trotter. And over the prime fields, that's the best example that we know. So there is a conjecture, if x is in Fp squared-- and p is always prime, but it's important here, so I will add it-- then the size of Es of x is bounded by 1 plus s squared over x. Matching that. But this is not true in Fq squared. q is not prime.

So let me show you an example. And this example plays a big role in the theory. So here's the example. Suppose that q is p squared. There are similar examples for higher powers, but this is just simple concrete. So Fp is a subset of Fq. It's a subfield of Fq. And our set x is going to be Fp squared, which is a subset of Fq squared. Cardinality of x is p squared, which is q.

So now notice that if t is also in Fp and x is in x, then pi t of x, which is x1 plus t times x2, well, all of these characters, I've assumed to be in Fp. And Fp is a field. So it's closed under addition, closed under multiplication. So this is in Fp.

So that tells me that if t is in Fp, pi T of x is Fp-- is contained in Fp. It's actually all of it. And so the size of pi t of x is at most p, which is q to the 1/2, which is much smaller than the size of x.

So we conclude that the size of Ep of x is at least p. So there are p traces t that all lie in Ep of x. And that's much bigger than the bound in the conjecture for what should happen over prime fields.

So if we were to plug in these numbers up here, s would be p, which is the square root of q, and x would be q. So this would be y. Over prime fields, this should be order of 1. But over this not prime field, it's much bigger.

So this shows that projection theory is also connected to some-- I don't know-- algebra, algebraic structure questions of whether fields have subfields. And this issue, which is-- so this is an open conjecture. And it's hard for two reasons. The first reason is, well, you look at it. You say, OK, it's like Szemerédi-Trotter, but over a prime field. So I should try to take the proof of Szemerédi-Trotter.

And that proof completely breaks down. That proof was based on the topology of R. We'll go through it, but the topology of R and the topology of Fp are very different. And it just doesn't work at all.

The second reason is that the statement of the conjecture really requires this to be prime. And so it's not true for any field. It's only true for prime fields. And so you have to wonder, how are we going to use that in the proof? And a lot of methods, a lot of methods don't distinguish between prime fields and non-prime fields. And therefore, they can't prove the whole thing.

In the real analysis setting, actually, there's the same issue. So these theorems that I wasn't able to state precisely for you today, they are true over R2, but they are false in C2, because the complex numbers have a subfield R. So the complex numbers play the role of Fq, where q is p squared. And so over here, also, we need to do something to distinguish between the complex numbers and the real numbers. And many techniques don't do that.

So a bit surprisingly, these real analysis problems that eventually get used in estimates and Fourier analysis with integral operators that really look like analysis, they take something from these questions over finite fields. These are more closely related to each other than either one of them is to that first Szemerédi-Trotter theorem.

So at the beginning of the class, I will teach you some classical techniques, some double counting techniques, some Fourier analysis techniques. Hold on a second. And they prove interesting things. And they work over here. And they work over here. But they don't distinguish between the prime fields and the not-prime fields. Yeah. Was there a question?

AUDIENCE: Does the fact that the reals have the rationals as a subfield end up mattering for this? Or just kind of the topology argument [INAUDIBLE]?

LAWRENCE GUTH: Ah, OK. So the question is-- so the observation is, I said that the situation for Fq, when q is p squared, is kind of complicated, because Fq has a subfield, namely Fp. And then I tried to say there's something similar going on with the complex numbers and the real numbers.

Now, the observation is the real numbers have many proper subfields. One of them is the rational numbers. There are many more number fields. Doesn't that matter? Yeah, that's a good thing for us to reflect on. It turns out not to matter that much. And it's because those subfields are so small. The real numbers are an infinite order extension of the rational numbers. I think even an uncountable order extension of the rational numbers, whereas Fq is a degree 2 extension of Fp.

So one way people talk about it is in terms of Hausdorff dimension. So the real numbers have subfields, but their Hausdorff dimension is either 0 or 1 if it's the whole real numbers and not in the middle. Anyway, we'll have a look at it. Yeah.

AUDIENCE: Do you have any intuition about why things unfold for the small fields?

LAWRENCE GUTH: Why you can't do something bad with small fields?

AUDIENCE: Yeah, if I understand the intuition correctly, the real numbers are fine, because they're large enough extension over the rationals. So I don't know. I'm just wondering if you have something about the size of the space.

LAWRENCE GUTH: So here's an analogy for you. So what if we didn't have p squared? What if we had p to the 17th? Now, Fq still has a subfield Fp, but it's kind of a small subfield. We could do the same computation. We would see that this doesn't fully hold, but for sets that are pretty large, for sets that are a lot larger than that, we would still-- we might still hope for pretty good bounce. And I think that the situation for the small subfields of R is similar to that.

AUDIENCE: So is this conjecture known for any prime number or is it for every prime? So I want to prove this conjecture for R. [INAUDIBLE]

LAWRENCE GUTH: So we stick this tilde here, which means there's a missing constant factor.

AUDIENCE: [INAUDIBLE]

LAWRENCE GUTH: We do that to be safe, I guess. It probably is true with a small constant factor. But strictly speaking, you can't falsify it with one prime number. But that's not the real reason.

Then you might say, if you take a medium prime number, like 101, could you check it numerically? Maybe for 100-- even for 101, I don't have any wisdom about how to do this. And if you literally did it by brute force and checked all the possibilities x, you really couldn't go very far.

And I'm not sure if we'll address this, but there is a computational question, like, maybe I can't prove this conjecture, but do I understand it well enough that I could prove it for a particular p faster than brute force? I don't. And so we're not able to have much numerical evidence.

Yeah, it could be worth looking for a counterexample on the computer. The evidence is not that strong in favor of this conjecture. OK, cool.

In the last 15 minutes, I wanted to mention to you some other parts of math that this connects to. So one connection is to a part of combinatorial number theory called sum product theory, which we'll see it fits in naturally with this issue.

So in this example, this set that we're worried about is closely connected to the fact that Fq has a interesting subfield, an interesting subset that's closed under addition and multiplication. And the field-- the sum product theory and combinatorial number theory is about what happens to a set under addition and multiplication. So initially, it's over R over the reals. But then you can ask in any field.

So let's say that R is a ring. So we can define sums and products. And A is a subset of R. So A plus A is all of the sums you can make, two elements of A. So it's the set of all the A1 plus A2, where A1 and A2 are in A. And the product set A times A is similar to the set of all the products.

So actually, motivating this from before, what was so bad about this set Fp as a subset of Fq? The problem was just that, well, if these three guys are in Fp, then x1 plus t times x2 is still in Fp. So the problem was that Fp is closed under sums and under products. So it raises the question, are there other sets that are maybe not exactly closed, but roughly closed under sums and products?

So the idea of sum product theory is take a subset, A in R. Could it happen that both A plus A and A times A are pretty small? So let me mention a few examples. So if A is a subset of the real numbers, which is generic, then A plus A and A times A would both be around the size of A squared. So generically, there would be no coincidences in the sums or in the products.

But it could happen that A plus A is not very big. So if A is an arithmetic progression, then the size of A plus A is only around 2 times the size of A. But in this case, A times A would be very big. On the other hand, if A is a geometric progression, then the size of A times A would be around 2 times the size of A. But A is a geometric progression, then A plus A would be very big.

So based on these few examples, Erdos and Szemerédi conjectured that always one of these guys is big. So here's their conjecture, as it's still an open problem. If A is a finite subset of the reals, then the maximum of A plus A and A times A is at least C epsilon A to the 2 minus epsilon.

So for every-- so at the beginning, I should put for every epsilon bigger than 0. So it's not quite a squared. That example comes from taking the numbers 1 to n. Then A plus A is clearly quite small. A times A is many of the integers from 1 to n squared, but not all of them. Some of them have a prime factor that's bigger than n. So they're not in the list. And Erdos computed that carefully. So he said you need the epsilon.

So this type of problem is connected to progression-- to projection theory because of this type of example. And the connection goes in both directions. So in one direction, there's a theorem of Elekes about the sum product theorem. So we can't prove this. But any exponent here that's bigger than 1 would be non-trivial. And we can prove it with exponents that are bigger than 1. And historically, there was a big jump in the quality of the exponent when Elekes proved this theorem.

So he proved that A is a subset of R, which is finite, then the maximum of A plus A and A times A is at least A to the 5/4. And the key ingredient in the proof was Szemerédi-Trotter. That's not the whole proof, but the proof uses Szemerédi-Trotter, and not much else. It's quite short.

And so if you know something about projection theory, you can say something about sum-product theory. But later things went in the other direction. So there's a theorem of Bourgain, Katz, and Tao, which says that not only a prime field does not have any subfields, which we knew, but it doesn't have anything that's even close to being a subring.

So if p is prime and A is a subset of Fp and A is not too big. So it's not super important what this number is, but A is significantly smaller than the whole field. The maximum of A plus A and A times A is at least the size of A to the 1 plus epsilon, where epsilon is a small, but explicit number. Think something like 1/20.

This theorem is actually more accessible than projection theory in general. Having these sums and products gives you a little bit more structure. It turns out to be important. But it shows that in Fp, over Fp there is no example which is too much like this example. You can't just choose a subset of Fp, which is almost closed under sums and products, and then mimic this argument. And it's not immediate, but this is used to improve the projection theory in Fp squared. And something similar will also happen for problems about balls in R2.

I wanted to mention some other applications, but I don't want to cram things in too much in the last 5 minutes. I'll just say a few words, and then I'll pause for final questions. So one application or story is from sieve theory in number theory. So in sieve theory, you have a set of integers instead of a set of points in a vector space. And in sieve theory, we study what happens if you take the set of integers and you reduce it mod p or mod q for many different numbers q?

And there are some questions about general sets, like, could it happen that you have a set of integers? And whenever you reduce it mod q for any q, you always get something small? What could you say about your set of integers?

So an interesting example of that are the square numbers. If you take the squares and you reduce it mod p for any prime-- you start with primes-- you get the quadratic residues, which is only half of the residue classes. So whenever you reduce it mod p, you only get half of the residue classes.

How many other sets of integers can you think of that whenever you reduce the mod p, you only get half of the residue classes? How big could such a set of integers be? What fraction of all the integers? And so on. So there's one set of questions about this.

And there's another set of questions where you take a particular set of integers that you might be interested in, of which the most famous is the primes. And you think about what happens when you reduce the set the primes mod q for many different q. So it appears to be the case that the prime numbers are evenly distributed mod q for every q.

Actually, depending on how you quantify it, that's a theorem. So if you fix a q and you take the primes mod q, and you take primes up to x, where x is much, much, much bigger than q. They will be evenly distributed among the residue classes mod q that are possible, the ones that have no common factor with q.

But you could have-- you might also want to know if you take the primes up to some fixed x and you reduce them mod q for many different q, it still appears-- it appears that those are evenly distributed for quite big q, even q's that are only slightly less than x.

We can't prove anything like that. But there's a line of thought that maybe we could relax what we're trying to prove a little bit. And instead of trying to prove it for every q, we could try to prove it for most of the q. Now we're in this setting where we have one fixed set, and we're looking at how it reduces mod q for many different q. And that problem is very similar to the projection theory problem.

So anyway, so we'll teach some core techniques of projection theory that were developed by people in geometric measure theory in the '60s, '70s, and '80s. And those are almost-- those are very closely parallel to the classic techniques of sieve theory that were developed by people in the '40s, '50s, '60s. So interesting connections there. In sieve theory, there are lots of open problems. So that's one set of applications, if you like, or connections.

And the other one I wanted to mention is dynamical systems. So Elon Lindenstrauss last year gave the Simons lectures. And one of the things was about taking orbits in homogeneous spaces, unipotent-- so I don't have time to say what all the things are. But you take unipotent orbits in homogeneous spaces, then they distribute very evenly.

So there are some classical important theorems about that. And the recent thing he and Amir Mohammadi were trying to make it more quantitative. And a new idea in their proof is they brought in projection theory. So anyway, so I will tell you the story, slowly. First of all, what is the question? And then, what does it have to do with projection theory? So those are some applications I was thinking that I'd like to talk about with you.

OK. Any last questions or comments? Yeah.

AUDIENCE: I was wondering if you'll tell us about the [INAUDIBLE]?

LAWRENCE GUTH: So the question is that the Szemerédi-Trotter theorem is usually set in a different way, actually a more general way, where you have a set of lines. And you look at the points that have many lines going through them. And we'll see that this is a corollary or special case of that.

AUDIENCE: That's like the version of this [INAUDIBLE] problem in Fp. They are quite different. For example, if you just estimate the number of average points, we know that the bulk in our tree does not hold.

LAWRENCE GUTH: Yeah. Let's talk about that after class.

AUDIENCE: Yeah.

LAWRENCE GUTH: OK, good. Let's stop there for today. I will see you on Thursday.