[SQUEAKING]

[RUSTLING]

[CLICKING]

LAWRENCE GUTH: OK, today is our second and probably last day of homogeneous dynamics and projection theory. Let's remember where we are, and then we're going to try to use the things we learned about last time to say something about dynamics.

So we'll have a Lie group G, which so far is SL2R. And we have a discrete subgroup, which so far has been SL2Z. And we're interested in G mod gamma. We'll have a metric m, which is a right invariant metric on G. And then it leads to a metric that I'll also call m on G mod gamma.

So now G also has a left action on G mod gamma. And this left action does not preserve the metric, because the metric is not left invariant. So this left action distorts the metric. And we started to study that last time, and it's going to continue to be crucial how it distorts the metric. But it preserves the volume.

Now, inside of this group, we have the unitary group. Sorry, the unipotent group. Let's say UT is 1 1 t. And our goal is to understand how the unipotent orbits in G mod gamma look. So we're going to study u dot x in G mod gamma. So x is in G mod gamma.

And there's an old theorem of Hedlund from the '30s, which is the kind of thing about this orbit that we're interested in that u dot x is either periodic or dense. We won't exactly prove that, but we'll prove some things that are a little bit related.

Now, the unipotent group is special. And the one way of getting our hands on something special about it is to see how it relates to diagonal matrices. So let's say this a sub r is that matrix, that diagonal matrix in SL2R. And then acute observation is that a sub r UT a sub r inverse is u e to the 2rt. It's just a computation with matrices. I've done it a couple of times, and I start to-- anyway, doing it in front of you, I think, is not very useful. But if you are interested in this, if you do it yourself, it starts to make sense.

So now let's say that u0t of x, that means the set of ut of x t is in 0t. So if capital T is large, this is a big piece of an orbit. And that's more or less what we'd like to study. What is the shape of a big piece of the orbit? And this observation tells us that u0t of x is ar u01 ar inverse of x.

And this is helpful, because this u01 is much less intimidating than u0t. So this here is just a little piece of the orbit. And so to understand a big orbit, we have to understand how this a sub r acts on a little orbit.

And another simple observation is that ar is like a little r to the j. So if big R is big J times little r, then we have this. Oh, in this equation, by the way e to the big r is t So this ar, we can think of it as a little r applied many times. So therefore, it makes sense that we could just try to understand how a little r acts on unipotent orbits.

So we made a picture of this at the end of class last time with models and props and everything. And this picture is the main tool that we're going to think about for homogeneous dynamics. So I'm going to draw this picture again, and then we're going to talk about why it is the way we drew it.

So over here, a shape that I'll draw as being kind of a cylinder. And this is a fundamental domain. So this is either the whole fundamental domain or maybe a tube in f. So f is sitting in the group G, and it's a fundamental domain for G mod gamma. And so this length scale is 1, and this length scale is also 1. Or we could also make it smaller if we'd like to have a thinner tube.

And now inside of this tube are a bunch of pieces of unipotent orbit that I'll draw like this. So each of those is a piece of unipotent orbit. Here is t equals 0 and here is t equals 1.

And now I understand what is going to happen to this picture when we apply a sub r. So we're going to apply, I guess strictly speaking, we're going to apply it on the left. We're going to multiply by a sub r on the left.

So what's going to happen? So first of all, the unipotent orbits are going to stretch that way. And then second of all, so if you just think about what happens to the cylinder, there's going to be a direction where it stretches, there's going to be a direction that compresses, and there's going to be a direction that stays the same. That would be true at every point.

But what's important is that the direction that compresses is kind of not the same everywhere in this picture. And so to illustrate that, I'm going to draw what happens at a couple of different heights. So at t equals 1, I have this.

And there will be some direction that compresses. Maybe it's that direction. And so when we apply lar inverse, this disk is going to be squished vertically. And when it's squished vertically, then we'll have a triangle that looks sort of like so. And also at t equals 0, we have the same disk with the same triangle in it. But there's a compression direction, which is different.

So maybe down here it's this direction. So when we apply la sub r inverse, then we will get a pancake and get squished this way. And that will affect these points differently. After squishing, the green point and the orange point will be quite close together. So that's the picture that we ended with last time.

And what is important here, key point, is that the compression direction is twisting relative to the unipotent orbits. Meaning that if at a given height, two unipotent orbits are being compressed towards each other, then at most other heights, they won't be. So this is where we left off last time. Are there questions or comments about this picture? Yeah.

AUDIENCE: So in the picture you had last week, these [INAUDIBLE] things are all kind of stretched out and compressed and then mapped back onto the fundamental domain. But if they've been twisted around, are they kind of mapped back onto f without twisting in some way to make it all perfectly line up? Or is there some other twisting that happens again as they get back into f?

LAWRENCE GUTH: Yeah, yeah. OK, that's a good question. So the question is about how this gets mapped back into the fundamental domain. Let's try to draw that also. So here's our fundamental domain. First we apply this map just in our group, and we get something which is tall and flat. And perhaps, it is twisted a little bit between the bottom and the top.

Anyway, then that is going to get mapped back into our domain. So pi is the projection from G to G mod gamma. And G mod gamma will identify with our fundamental domain. So this thing is going to get back into the fundamental domain. What does that look like?

So the fundamental domain is kind of a unit cube. But this thing is way taller than a unit cube. So imagine cutting it into pieces whose height is only around 1. And this was only ever around 1. So each of those pieces now is small enough to fit in the fundamental domain. And let me call it piece one, piece two, piece three, and piece four.

Now, those are going to fold back into the fundamental domain. And at least roughly, I think it's good to imagine it like this. So each one of these guys corresponds to one of those guys. But they are not in order. So it's not like they go 1, 2, 3, 4. They're in some order, 1, 2, 3. So I don't know much about this order. I personally don't know anything about this order.

OK, so now I think the question was, how would this picture look if this thing were twisted a little bit? So first of all, when we say that this thing is twisted a little bit, it's something like a change of 90 degrees or 180 degrees between the bottom and the top. And there are many steps between the bottom and the top. So that means if you just look at one of these guys, you don't notice or barely notice any twisting. This thing is basically a planar slab.

And then these planar slabs are at slightly different angles. But this map pi, it doesn't exactly preserve-- I mean, those angles aren't exactly well defined, because this is not living in R3. So it's actually not clear exactly what we mean when we say that this slab and this slab, are they parallel or are they twisted compared to each other? Anyway, when they get mapped back in, I think that imagining this is reasonable.

Now, that raises a bit of a question, coming back to this picture that we were trying to digest last time. What exactly do we mean when we say that the compression angle is twisting? Because the compression angle here lives at one height. It lives over here, and the compression angle there lives over there. And how do we compare an angle over here with an angle over here given that this is not R3, this is living in a curved thing?

So I actually rethought how to state it, and I added some words that I don't think I mentioned last time. I think when we talk about this twisting, we need to include these words relative to the unipotent orbits.

So let me try to make a more precise statement about what this twisting means, and then we'll see how we can do some computations in the Lie group with matrices and see that it happens. All right. OK. So precise statement about what is twisting. What is twisting.

So suppose I have a piece of unipotent orbit in my fundamental domain. It starts at G0. That's over there, some group element, and our fundamental domain. And then over here I have UT of G0. And t goes, say, from 0 to 1.

Now, at each point here, there's going to be a vector that compresses when I perform ar. So let's say v compression of t. This is a vector, and this is the direction that compresses when we apply a left ar inverse. So it's a vector. It lives in the tangent space of our group at the given point. And it's the smallest singular vector for the derivative of this map. So that's what we mean by the compression direction.

Now, what does it mean relative to the unipotent orbits? So suppose that over here I have another unipotent orbit. And this is like a nearby unipotent orbit. So this might be G0 plus epsilon times v0. Epsilon times v0 is a little vector that goes between these starting points of these two nearby unipotent orbits.

And up here, or maybe here, is UT of G0 plus epsilon v0. And here there's a vector that goes between them, I guess, epsilon times v orbit of t. So v orbit of t is going in the direction from the red orbit to the orange orbit.

So the thing that is twisting is v compression of t relative to v orbit of t is twisting. The angle between those two vectors is changing. That's the precise statement of what is twisting. And we can check what it's happening by computing these two vectors. Yeah?

AUDIENCE: [INAUDIBLE]

LAWRENCE GUTH: So the question has to do with keeping track of things, given that they live in different tangent spaces. So the compression direction is a vector in the tangent space at the given point along the orbit. Is that OK so far?

AUDIENCE: Yeah. Yeah, yeah, yeah.

LAWRENCE GUTH: OK. But then the comment is if I have a family of vectors in one vector space, then I can say that they're twisted. But here I have a family of vectors in a family of vector spaces. And so it's not so clear what it means that they're twisting.

One way to do it is to pick a reference vector to compare them with. And so that's the job of this v orbit vector, which goes from the red vector to this-- from the red orbit to this orange orbit. And that's like a reference vector. And then we'll see that-- so at each t, there's an orbit vector and a compression vector, and the angle between them is changing.

So there was a question last class that basically was like, how are we sure that this is happening in this picture? And the one way to be sure is to compute these two vectors and see that they're not the same.

So let's compute them. Actually, one of them we can compute over here. So UT is a matrix. So we can just multiply this through. This is UT G0 plus epsilon UT of v0. And if we compare sides, we see that v orbit of t is u sub t of v. Now, let's figure out what v compression of t is.

So say we have some G0 and we're going to apply the left action of ar inverse. And that will move us to a sub r of G0. And we want to understand how this map is distorting the geometry of G0. The geometry near G0. And in particular, figure out which direction is being compressed.

So I think the best way I could figure out to understand this is to put it in a picture connecting this thing that I don't understand yet to some maps that are easier to understand. So here's the identity. And from here to here, I could apply a right multiplication by G0. And right multiplication is much easier to understand, because that's an isometry. And the identity also is kind of nicer than other matrices.

And then I'm going to go back to the identity. So this is right multiplication by ar g0 all inverse, which is also an isometry. So if you put this all together, I'm going to have a map from the tangent space of the identity to itself. And that map is an isometry times the map I want to understand times an isometry so that all of the metric distortion is going to come from here.

Now, this map is not so hard to understand. So if I apply-- what am I doing? Sorry. I start with some matrix h. Right multiply by G0. Then I left multiply by ar. Then I right multiply by the inverse of ar G0. What do I get? Well, I have h. I multiply on the right by G0. Now I multiply on the left by ar. I multiply on the right by this.

When I invert this, I switch those. I get a G0 inverse and they cancel. So this operation is none other than conjugating by ar. So this is the conjugation by arh. So conjugation by ar is a map from our new group to itself, takes the identity to the identity. And therefore, this map, or really its derivative, takes the tangent space of the identity to itself.

All right. And now you can write this down. This ar is an explicit matrix. A Tangent vector here is a matrix ABCD. So this is just ABCD goes to DCAR of this is er minus r, ABCD e or minus r. So this is a thing you can write down. It's not that scary. And we worked out last time what it does.

So recall from last time that we have-- first of all, we have an orthonormal basis of Te of G, which are-- there's a basis vector n, 0, 1, 0, 0. Last time I called this u, but there were so many u's on the blackboard that I was worried it was not a great idea, so I'm going to call it n. n for nilpotent, which is also true. There's n tilde, which is 0, 1, 0, 0. And there's a diagonal guy, which looks like this. OK.

So What we figured out is that the conjugation makes n way bigger, it makes n tilde way smaller, and it leaves the diagonal fixed. OK. So, if I start at the identity and I go all the way around, I'm doing this conjugation, and n tilde is the vector that's going to get compressed.

But now, if I want to know what vector over here is going to get compressed, Well, I start with n tilde, I map it over here. It hasn't gotten any shorter yet because this is an isometry. That's the vector that's going to be compressed in this step. OK.

So, it's the right action by g0 of n tilde. This is the singular vector of dLar inverse with singular value e to the minus 2r. OK, so that's basically the answer to our question. So the compression vector at T, that's r at-- so now I'm not necessarily at the point g0, but I'm at the point ut g0, so it's right by ut g0 of n tilde. So that's n tilde ut g.

OK. So now, if you compare how these two things are changing with T, you see, it is not the same because our group is not commutative. Because when you're changing the orbit of T, you're multiplying by ut on the left. But when you're changing the compression direction, you're multiplying by ut over here in the middle. And it's a non-commutative group, so those two operations are not the same as each other.

OK. So I don't necessarily think it will be helpful to do more computations, but I did some at home to be sure, if you imagined that at time 0, these two directions coincided, and then you see what happens as T changes, you'll see that they stop to coincide. Yeah?

AUDIENCE: So it can be the inner product of the vectors, does it make sense as an operation to compare the angles in the space?

LAWRENCE GUTH: Yeah, that's right. So these are vectors in a Riemannian manifold, and so they have a well-defined angle, and we can compute it by taking inner products. And then-- correct. So I guess we could say that the cosine of the angle between v of t and w of t would be v dot w over norm of v norm of w.

I used to teach this a lot in 1802. I think that's right. So we have formulas for v and w, so we can compute this. And it'll be a little messy, but we'll see that it is not constant in T. OK. Other questions or comments? OK.

So the first big goal-- that was the first chunk of the class. The first big goal of the class was to check that this picture happens. And then the second part of the class, we will discuss how to use this picture to figure out stuff about along unipotent orbit. OK. So let's switch gears. OK.

OK. So let's call this tracking the spread of an order. All right. So remember that we had this, U0,T of x is aR U0, 1 aR inverse x. Let's call this guy x tilde. OK. Now this-- there was a discussion last time, what happens if this orbit here is periodic?

If this orbit here is periodic, the special thing that happens is that this point, x tilde, will be deep in the cusp. And the deep in the cusp is a special case because we were imagining that this was a regular nice orbit of length 1, but if it's deep in the cusp, it's different.

So, let's assume that x tilde is not deep in the cusp. And that's equivalent to saying that Ux is not close to periodic. OK. And let's break up R as capital J times little r. And so then I'm going to say that Uj is a little r to the j of U0,1 of x tilde. OK.

And I want to track how this thing is spreading out through my homogeneous space, GY gamma. And we're going to track it by noticing how many delta balls it takes to cover it. So recall that X delta is the minimum number of delta balls needed to cover X.

And so our goal is going to be to estimate. Uj delta in terms of j, delta, and r. Oh. And for context, for reference, we have a three-dimensional space. So for the whole space, the number of delta balls that we need to cover it is delta to the minus 3.

OK. Cool. All right, so next of all, our Uj is like a bunch of these unipotent orbits. Looks like that. And so you might notice that it only really matters what the top cross-section looks like. All the other cross-sections are similar. So let's take that top part down. Inside of there, we see this.

So let's call that Xj. That's the top cross-section of Uj. And then it's not hard to check that Uj delta is like delta inverse times Xj delta. I have to cut it into slabs of thickness delta. For each thickness, the number of delta balls that we need to cover that slab is Xj delta, and this is the number of slats. OK.

So basically, we want to track how Xj delta depends on delta. So, fully spread. Or maybe I'll say very spread. That means that Uj delta is like delta to the minus 3, and that would be like saying Xj delta, so delta to the minus 2 at Xj. Really fills a lot of the disc.

So now, how does this picture help us to study Xj delta? So using the key picture. All right. So let me label these maps pi 1 and pi minus 0, and somewhere in the middle we would have pi t. Pi t is what our map aR is doing to the slice at height t. OK. Using the key picture, have a lemma.

So if e to the minus 2r is delta, then Xj plus 1 delta is around a sum over t of pi t of xj delta. And which t's do we sum over? The t's are in delta Z, and t goes from 0 to 1. This is all-- we can also write this as delta inverse times the average t between 0 and 1 of pi t height of Xj delta.

So, proof sketch. OK. So here is our fundamental domain. Cut it into slices of thickness delta. So here is a slice at height t with thickness delta. OK. Now, when we apply Lar, we stretch it up vertically. And we compress it in some other direction, we get something like this. And then we divide this into pieces at different heights. And let me color this particular slab in red. So that red one is one of them over here. So these are each of height 1. OK. Right. Right.

So, in this new picture, there will be-- so in the old picture, there were some unipotent orbits that went through our slab. So, X describes where the unipotent orbits hit the top of our slab. And now, we apply this map, and now there will be some points over here in the top of this slab. And who will they be? Those points over there would be pi t of X.

OK. Now we want to cover things with delta balls. So this distance here is e to the minus 2r, so that's delta. So we want to know how many delta balls does it take to cover that? So, that number there is delta covering number of pi t of Xj. So over here, we started with xj.

Now, what happens when we apply pi? Each of these slabs gets folded back in. And the red one went somewhere. I don't really know where, but maybe it went over here. And here in the top of it is pi t of Xj and its delta-covering number. So the number of delta balls in this strip, that would be pi t of Xj.

OK. And now, this thing-- these dots up here are our Xj plus 1. So the set of dots on the top face. And so how many delta balls do we need to cover them? Well, for these ones, we needed pi t Xj delta-covering number, and these other guys, those corresponded to different slabs here, so there are different terms in this set. OK.

So now, we can try to use projection theory to study this thing, and it will tell us how the covering number Xj plus 1 delta evolves as we increase J. Any questions-- yeah.

AUDIENCE: Why do-- like, we have delta balls start to overlap when j is large enough and the compression is more than delta because delta is the size of one round of compression. If you do multiple, why do they overlap?

LAWRENCE GUTH: So I think the question is, suppose at the beginning that there happened to be two guys here in X whose distance was less than delta, which might happen. And then, what would happen to them? Well--

AUDIENCE: Not that they're less than delta. If slightly more than delta, but less than e to the minus 2r or e to the 2r delta, so that in f, they're in the same delta ball. Then applying the transformation, the [INAUDIBLE] or transformation method, then they become [INAUDIBLE] delta ball. But before they're not in-- before, there were two separate delta balls; after the transformation, there's one delta ball.

LAWRENCE GUTH: Right, right. OK, so it could happen that we have two guys-- like maybe these two guys are pretty far apart, but after we perform pi t, they got squished together and they're now in the same delta ball over here. Yeah, that could happen. And the main thing we have to study is how often that happens.

But this lemma is still true. This might be a lot smaller than Xj delta because many points that were far apart in Xj may have gotten squished together in pi t of Xj. But this is just the number of delta balls that it takes to cover the image there. Does that answer your question?

AUDIENCE: Yeah. Also, in the third image, why is each delta ball containing only the size of one of those strips?

LAWRENCE GUTH: Ah, OK. So we tuned delta and r to play together. So that we picked r so that the thickness of this strip is delta.

AUDIENCE: I thought if the transformation is applied multiple times, then you get multiple e to the minus 2rj scale, right?

LAWRENCE GUTH: That's not wrong, but when we apply it multiple times, we're not changing little r, we're just applying a sub r-- a sub little r multiple times. So this picture is describing what happens in one time, and then you could use this in many rounds. So you could use this lemma over and over again. Which is what we'll do. OK.

Let me change my notation a little bit. So I was using the letter pi for this maybe a little bit prematurely. What's clear so far is that this is a map, so I'm going to call it f1, ft, and f0. And so then I'll change this. OK. OK. So now, the next remark is that these maps look a lot like linear projections. So remark, ft is kind of approximately an orthogonal projection pi t. So let me make a picture.

So ft does something like this, and this thickness is delta. This is 1. And then pi t would just be an actual orthogonal projection. OK. So ft is not perfectly linear, but it is smooth. OK.

So now, here is a proposition-- or exercise about projections. It says that if X is in the unit ball in a plane, then you take the average over all directions, theta and S1, of pi theta of X delta that is at least as big as X delta to the 1/2. Here's an example that shows where the X delta to the 1/2 comes from.

So here's the unit ball. And then inside of it, I take a smaller ball and I call that X. So, X is the ball of some small radius rho inside of the ball of radius 1. OK. Then when I perform an orthogonal projection-- there is pi theta of X, in any direction I'll get a line of length rho.

So I see that X delta is like rho over delta squared, and pi theta of X delta-covering number is like rho over delta, so this is the square root of that OK. And this is the best thing to do. You could devise an example where in one particular direction, the compression is much smaller than the square root. Here's an example like that. B1. Here's X. Here's theta. OK.

In this example, pi theta of X delta is much smaller-- I'll call this theta 0-- much smaller than X delta to the 1/2 for one theta 0. But at least in this picture, most of the projections are way bigger. OK. So it's not super difficult to check this proposition, and I'm going to make it an exercise. I think we're out of problem sets, but if we weren't, this is a nice one to review the class.

OK. Another comment is even though these ft's are not exactly linear maps, so that we cannot just literally apply this proposition, they are close enough that by using the proof of this, you can get it for ft, too. So remark, the proof applies to ft. OK. So if we put together this lemma and this proposition, we'll get an estimate about the delta-covering number of Xj.

OK. So corollary, if we put these together, Xj plus 1 delta-covering number is at least around delta inverse times Xj delta to the 1/2. Proof. So Xj plus 1 delta is at least, by our lemma, delta inverse times the average over t of Xj-- of ft of Xj delta. And then by our proposition, that's at least delta inverse times Xj delta to the 1/2, and that's the proof.

OK. So then let's just track what this means. So X0 delta would be 1. Think about what X0 delta means, it's one unipotent orbit. So therefore, X1 delta would be at least around delta inverse, X2 delta would be at least around delta to the minus 3/2, X3 delta be at least around delta to the minus 1 and 3/4, and so on, and you'd get pretty close to delta to the minus 2. OK. And delta to the minus 2 would be like a definite fraction of the whole space.

All right. So this proves that a unipotent orbit that's not close to periodic, it proves that it fills up a lot of the homogeneous space. And it's not Hedlund's theorem. Hold on a second. So we didn't prove yet that the orbit is dense. The orbit could be missing an entire ball that's like half the size of the space, and this could still be true. But we did prove that the orbit goes a lot of places. Yeah?

AUDIENCE: In this proof, does-- since applying it each time loses some of implicit constant, how do we know that it's not becoming dense and just, like, an infinitesimally small region of this whole space?

LAWRENCE GUTH: Right. So I think the question is like, we didn't actually get to a constant times delta to the minus 2, and if I tried to do this infinitely many times, I would be losing a constant each time. Yeah, that's right. So imagine doing this 10 times and we get delta to the minus 1.99. And that's what the method proves so far. So stronger things are true, but this method doesn't prove them yet.

So remark, this doesn't show that the orbit is actually dense, and so maybe question, could the orbit do this? So here's our homogeneous space, there's a cusp, but it's not that important. And then the orbit goes around a good bit, but it only ever goes around over here, and it never goes over there. So this proof doesn't say anything about that.

OK. So, as far as I know, to finish Hedlund's theorem requires another step that is based on different ideas from these ideas. So there's more to homogeneous dynamics than this story, but we did still learn something interesting. Cool. Yeah?

AUDIENCE: Do how close as possible to get to the full theorem by trying to optimize all the constants in this proof?

LAWRENCE GUTH: OK. So there are two things that we might like to improve here. So the question is, how close can we get to the full theorem by optimizing the argument that we've been talking about?

So one thing that you might like is you might like to actually get to delta to the minus 2. That still wouldn't necessarily rule out this possibility, but it would be interesting. And if you look at this proposition, one thing that you notice, that this proposition is sharp in this very clumped case. And if you remember back to the early months of our class, if x was not so compressed, then there would be even better estimates.

So if you knew as an input that x is not too clustered, then you could get all the way to delta to the minus 2 in one step or two steps or something. And so you might try to track that. You might try to track not just the delta-covering number, but also the rho-covering number for different delta and rho, and then you could take advantage of this clustering.

And I am not positive, but I think it is a reasonable 18-156 project to try to use that to get all the way to delta to the minus 2. Oh. Even if you got a constant times delta to the minus 2, it would still not be obvious to me that you filled the entire manifold and that you're not missing a ball.

And I thought about it yesterday and today, and I don't have a good understanding yet of what you need to get everywhere. So I would be interested to learn if people have thoughts about it or know about it. OK. Cool.

So I think that there are actually are a bunch of ways to prove Hedlund's theorem, and this is not one that I have seen in the literature. And all of these problems become more difficult and maybe more interesting when you go to higher dimensions.

So the projection theory also becomes more difficult and more interesting in higher dimensions, and I wanted, in the last part of the class, to just show you what happens. And in particular, we will see a new question about projection theory that we haven't mentioned so far in this class, but which I think is a really cool question. OK.

OK. So let me put it this way. What about SLn R? OK. So, I'm going to take a somewhat specific example, although I think most of what we'll say then would work for any n. So for example, we could have SL3 R for our G, gamma could be SL3 z, and U could be the unipotent group of matrices that look like this.

So from what I understand, there are several different one-parameter unipotent subgroups of SL3 R. They're slightly different from each other. And this one is particularly well-studied in the homogeneous dynamics community because it's the one that comes up in the Oppenheim conjecture about quadratic forms. But anyway, so we have this one-parameter subgroup. And so let's call this u sub t.

OK. So our basic setup with conjugating by a diagonal matrix, that works. And the relevant diagonal matrix is this matrix. OK. And then you can check that a sub r u sub t a sub r inverse is u sub e to the r t. Turns out to be e to the r t, not e to the 2r t, but that's not too big a deal.

OK. So now, we have a bunch of unipotent orbits like before. It was a bunch of unipotent orbits. But the dimension of our group is 8 instead of 3. And so we have the direction along the orbits, t equals 0 up to t equals 1, that's still one-dimensional. So this thing here is a seven-dimensional ball instead of a disc. OK.

So now, we apply-- we study the action of Lar inverse. And we'll have a picture like this. f1. f0. And, OK, something here will happen, and something different here will happen. OK. So you can study the singular values-- the singular values-- of Lar inverse. And those are equal to the singular values of conjugation by ar.

So conjugation by ar is a very nice operation on matrices, you can write it down and see what happens. And you can compute these singular values. And they are as follows. e to the minus 2r, e to the minus r, e to the minus r, 1, 1, e to the r, e to the r, e to the 2. I won't do it on the board, but I promise it's not hard.

OK. Now one of these singular values corresponds to the U-direction. So, little n is the vector that's the tangent vector to the group u at the identity. And so little n is-- it's not super important what the formula is, but if you differentiate that, you'll see that the tangent vector is that. And that's one of these guys.

So as before, the direction of the orbits is expanding. And then on the B 7, we have the other seven directions. So this thing over here is an ellipsoid with axes e to the minus 2r, e to the minus r, e to the minus r, 1, 1, e to the r, e to the 2. OK.

So, as we go from here to here, these maps are not exactly linear, but they're close to linear, they're close to linear maps, and they stretch some directions and they compress other directions by these amounts. OK. So, what is different from before? OK. One difference is that this thing is not exactly a projection. But that's not the most important difference.

So differences from last time. Right. So if we had an approximate projection that went from the 7 ball to some vector space V of dimension three, say-- so four dimensions are getting squished and three dimensions are staying as they are, that would be approximately a linear map with singular values e to the minus r, e to the minus r, e to the minus r, e to the minus r, 1, 1, 1. If I had a linear map like that, these guys are almost 0. So these are the four directions that are being squished, these are the three directions that are staying the same.

The list of singular values here is not quite like that. Some of them are bigger than 1, we didn't have that before. And the ones that are-- some of them are equal to 1. The ones that are smaller than 1 have two different scales. So that's a little bit more complicated, but I don't want to focus on it. So, our singular values are a little more complicated.

OK. Let's ignore that. Suppose each of these linear maps actually had those singular values, so it was more like what we were used to in projection theory. There would be a different-- there would be another difference from before, which is more important. So number 1. Number 2, we still have only a one parameter family of linear maps. OK.

So before, when we talked about projections from B 7 to V-- to the three-dimensional vector space, we talked about all of them. If you look at all of the projections from B 7 to the three-dimensional vector spaces, then something nice happens most of the time.

But we don't have all of them. We just have a very thin subfamily. So this is not all of the projections. Cool. OK. So there's also a cosmetic complication that 7 is a lot of dimensions and there are a lot of numbers, but that's not that important, it's just a pain. But these are this to some extent, and this is a really fundamental conceptual difference. So I wanted to show you not this problem, but the simplest problem of this type. Which is called the restricted projection problem. OK.

All right. So this question, which I think is a cool question, was raised by Fassler, Katrin Fassler, and Tuomas Orponen in 2013. And the question is like this. So here's the 2 sphere. So we're going to be doing projection theory in three dimensions, in R 3, so the different directions are parametrized by the 2 sphere. But we're not going to project in all of them, we're just going to pick some curve gamma and project in the directions of gamma. So gamma contained in S2 is a curve.

The question is like this. If X is contained in the unit 3 ball, and let's say x is a delta, s, C set-- we'll put-- to make it interesting, we put some kind of non-concentration condition, and this is the one that they had in their paper, basically. OK. Then estimate the average over all the directions in gamma of the pi theta of X delta-covering number.

OK, so pi-- so theta is a direction, an element of S 2, and pi theta is the orthogonal projection from R 3 to theta perp. So from three dimensions to two dimensions. So estimate the average of that. So we still have that-- I erased the wrong thing.

OK. On this very board, not two minutes ago, there was a proposition that said if you had a set in R 2, and you projected it in all the different angles, then on average, the delta-covering number of the projection was at least something or another. So this is like an analog of that in three dimensions for a set with the spacing condition, but with the key difference that we're not averaging over all of them, we're only averaging over the ones in a particular curve gamma. OK.

Now, an interesting thing about this question is that it matters what the curve gamma is. And in particular, there's a big difference between an equator, which is like a straight line on the sphere, and a non-degenerate, or curved gamma. So the answer depends on whether gamma-- gamma is the equator-- is one extreme or gamma non-degenerate is another extreme.

Instead of writing the definition of non-degenerate, let me just give you one example of a non-degenerate curve. So, an example of a non-degenerate curve is gamma is equal to the set of theta in S 2 so that the third component of theta is 1/2. OK. So it looks like the boundary of the spherical cap.

OK. So let me show you how these are different from each other. OK. So example. Say gamma is the equator and S is 2. So we have a, quote unquote, "two-dimensional set," delta 2 set. OK. So our set X is going to be a delta-by-1-by-1 slab. So it looks-- and so imagine that the two big directions are in the plane of the blackboard, and the slab is sticking out of the blackboard by a little bit delta.

OK. Now, there's an equator of directions that we're going to project in, and think of that as the equator that's like the sphere with center at the blackboard, intersected the sphere with the blackboard, you would get a circle. So here's gamma.

So now, if theta is in gamma, then the projection of X is a delta-by-1 rectangle. So there's a circle of directions where if you project this slab, you just get a delta-by-1 rectangle. OK. So then we see that the average pi theta of x delta is like delta inverse. OK. OK.

Now there's no obvious version of this example where gamma is that curve that's up near the top. And the difference is-- so if you're projecting in a certain direction, then the fibers are like lines in that direction. And so what's special here is that all of these lines are coplanar, and so you can have a planar object that has a lot that-- where these lines intersected a lot.

And instead, when we have those directions, then all of our lines make an angle of 45 degrees with the horizontal. They're like light cones. And they just don't fit together in the same nice way that coplanar lines fit together. So it makes a difference.

And Fassler and Orponen, they proved that it makes a difference. So they proved some bounds for that thing that are strictly bigger than this thing. But it took a while to figure out what is the sharpest bound. And that was figured out some time later in a paper of six authors, Shengwen Gan Shaoming Guo me, Terence Harris, Dominique Maldague, and Hong Wang, many of whom were here working on it at various times. OK.

And we proved-- I'll just do a special case. So if XCB 3 is a delta, 2, C set and gamma is non-degenerate, like this example above, then the average theta in gamma of pi theta of X delta-covering number is bigger than C epsilon delta to the negative 2 plus epsilon. OK.

And the best thing you could hope for here is delta to the minus 2 because x itself probably has delta-covering number delta to the minus 2. So up to small issues that-- this is the Sharpe theorem. OK. This proof is based on some version of the Fourier analysis method, but it's a more difficult version that involves decoupling. OK. Cool.

So this connection between projection theory and-- between homogeneous dynamics and projection theory was discovered by Elon Lindenstrauss and Amir Mohammadi, who were working on quantitative estimates for how unipotent orbits become dense in this setup. And they realized this connection with projection theory, and they found the literature, they found the paper by Fassler and Orponen and all the follow-up papers. And they used those tools to prove quantitative estimates about how these unipotent orbits mix. Yeah?

AUDIENCE: There's a question of projection theory where each unipotent element induces a different projection. Is that set of projections dense in the whole-space projections, or is it some non-dense subset?

LAWRENCE GUTH: No, it's some non-dense subset. Yeah, so here, t is going from 0 to 1. For each t, we have more or less a projection. And so we have a smooth curve of length 1 in the space of-- or not exactly projections, but the space of linear maps.

And so it is not, in any sense, becoming dense in the space of linear maps. It's really just a few special ones. And this is a reasonable model to picture in your mind. That sphere is like all of the projections or all of the linear maps, and that orange circle is the ones that come up in homogeneous dynamics. Yeah.

OK. So in the last two minutes, I will write down a theorem that they proved. At least approximately. OK, so this is Lindenstrauss, Mohammadi, Wang, and Yang. So setup is on that board. So if G is SL 3, R, gamma is SL 3, Z, U as above. OK. We take X and we suppose U dot X is not close to a proper homogeneous subspace. So this is some analog of being not close to periodic. OK.

Then the conclusion is that U 0, T of X is delta dense in the homogeneous space, but so the homogeneous space has infinite diameter, so this can't quite be true. But in the-- if you cut off the cusp-- so you cut off the cusp, so you have a compact piece of this.

And delta is T to the minus c. c greater than 0 is a universal constant. OK. And a key step in the proof is to first show an estimate of the kind that we were talking about, that the delta-covering number of this orbit is pretty big. So the biggest it could possibly be would be delta to the minus the dimension of the group, that's the delta-covering number of the whole space. And it's quite close to that. That's a key step. It's an analogous to the proof that we sketched for the simpler SL 2, R case.

And there is an additional argument, but it seems like the most novel and important part of this paper was using projection theory, using those kind of theorems, plus some other smart stuff to show that the orbit covers a lot of ground. OK.

OK, so homogeneous dynamics is a big subject, and there are different techniques, but in this course, I wanted to show that there's a connection to projection theory, which is one set of techniques, and which helps to say some interesting things.