

PROFESSOR: We spent the last few lectures developing the theory of graph limits. And one of the motivations I gave at the beginning of the lecture on graph limits was that there were certain graph inequalities. Specifically, if I tell you that your graph has edge density one half, what's the minimum possible C_4 density?

So for those kinds of problems, graph limits gives us a very nice language for describing what the answer is, and also sometimes for solving these problems. So today, I want to dive more into these types of problems. Specifically, we're going to be talking about homomorphism density inequalities. Homomorphism.

So trying to understand what is the relationship between possible subgraph densities or homomorphism densities within a large graph. We've seen these kind of problems in the past. So one of the very first theorems that we did in this course was Turan's theorem and Mantel's theorem.

So specifically, for Mantel's theorem, it tells us something about the possible edge versus triangle densities in the graph, which is something that I want to spend the first part of today's lecture focusing on. So what is the possible relationship? What are all the possible edge versus triangle densities in the graph?

Mantel's theorem tells us something-- namely, that if your edge density exceeds one half, then your triangle density cannot be zero. So that's what Mantel's theorem tells us. And let me write it down like this.

So the statement I just said, the one about Mantel's theorem, and more generally for Turan's theorem tells us that if the K_r plus 1 density in W is 0, then necessarily the edge density is at most $1 - \frac{1}{r}$.

So this is what it tells us. It gives us some information about what are the possible densities. But I would like to know more generally, or a good complete picture, of what is a set of edge versus triangle density inequalities. So let me draw a picture that captures what we're looking for.

So here on the x-axis, I have all the possible edge densities, and on the vertical axis, I have the triangle density. And I would like to know what is a set of feasible points in this box.

Mantel's theorem tells us already something-- namely, when can you-- so this region, the horizontal line at zero extends at most until the halfway point. Beyond this point, it's not a part of the feasible region.

So far that's the information that we know. Our discussion about graph limits, and in particular-- so let me first write down what is the question. So if you look at the set of possible edge versus triangle densities, so there is this region here. What is this region? It's a subset of this unit square.

We would like to understand what is the set of all possibilities. The compactness of the space of graphons tells us that this region is compact. So let me call this region D_{23} for edge versus triangle. So D_{23} is compact because the space of graphons is compact under the cut metric, and densities are continuous under cut distance.

So in particular, if you have some limit point of some sequence of graphs, that limit point's achieved by a corresponding limit graphon. So you really have a nice closed region over here. So we don't have to-- I should be able to tell you the answer. This is the region. There should not be any additional quantifiers. There's no optimizer zero, missing this point and missing that point. It's a closed region. So what is this closed region?

Equivalently, we can ask the following question. Suppose I give you the edge density. In other words, look at a particular horizontal place in this picture. What is the maximum and minimum possible triangle densities?

So I tell you that the edge density is 0.75. What is the upper and lower boundaries of this region? I want you to think about why this region is-- the vertical cross-section is a line segment. You cannot have any hulls. So that requires an argument, and I'll let you think about that.

So I want to complete this picture, and I'll show you some proofs. And at the end of-- well, by the middle of today's lecture, we'll see a picture of what this region looks like.

All right. First, let me do the easier direction, which is to find the upper boundary of this region. So what is the maximum possible triangle density for a given edge density? And the answer-- so it turns out to be what I will-- the result I will tell you is a special case of what's called Kruskal-Katona.

Think about it this way. Suppose I give you a very large number of vertices and I give you

some large number of edges, and I want you to put the edges into the graph in a way that generates as many triangles as possible. Intuitively, how should you put the edges in to try to make as many triangles as you can?

AUDIENCE: Clique.

PROFESSOR: In a clique. So you put all the edges as closely together as possible, try to form a clique. So maximize number of triangles by forming a clique. And that is indeed the answer. And this is what we'll prove, at least in the graph densities version.

So we will show that the upper boundary is given by the curve y equals to x to the $3/2$. So don't worry about the specific function. But what's important is that the upper bound is achieved by the following graphon. Namely, this graphon corresponding to a clique.

For this graphon here, the edge density is a squared, and the triangle density is a cubed. And it turns out this graphon is the best that you can do with a given edge density in order to generate as many triangles or the most triangle density possible.

In other words, what we'll prove is that the triangle density throughout W will be a graphon. So W is always a graphon. So values between 0 and 1. Then you have the following inequality.

So let's prove it. First let me draw you what this shape looks like. Because of the relationship between graphs and graph limits, any of these inequalities about graph limits, about graphons, it's sufficient to prove the corresponding inequality for graphs because the set of graphs is dense within the space of graphons according to the topology-- namely, the cut metric that we discussed.

So it suffices to show the corresponding inequality about graphs-- namely, that the K_3 density in a graph is at most the K_2 density in a graph raised to the power $3/2$.

So let me belabor this point just a little bit more. This inequality is a subset of those inequalities up there because graphs sit inside a space of graphons. But because they sit inside in a dense subset, if you know this inequality and everything is continuous, then you know that inequality. So these two are equivalent to each other.

Now, with graphs-- and specifically, these counts here, so triangle densities and edge densities-- they correspond to counting closed walks in the graph. So in particular, if we're interested in the number of K_3 homomorphisms in a graph, this is the same as counting

closed walks of length 3.

And there was important identity we used earlier when we were discussing the proof of quasi-random graphs, that for counting closed walks you should look at the spectral moment. So that's a very important tool to look at the spectral moment-- namely, the third power of the eigenvalues of the adjacency matrix of this graph. This is the eigenvalues of the adjacency matrix of G .

I claim that this sum here is upper bounded by a corresponding sum of squares raised to the power that normalizes. The first time I saw this I was a bit confused because I remember, power means inequality. Shouldn't go the other way. But actually, no, this is the correct direction. So let me remind you why.

So if you have a positive t , then claim that-- and you have a bunch of non-negative reals, then the claim is that this t -th power sum is less than or equal to the t -th power of the sum.

Now, there are several ways to see why this is true. You can do induction. But let me show you one way which is quite neat. Because it's homogeneous in the variables, I can assume that the sum is 1, in which case the left-hand side is equal to this sum of t -th powers.

So because I assumed that everything is non-negative, all these a 's are between 0 and 1. So now, this sum is less than or equal to the same sum without the t 's because you're using it like that. And that's equal to 1, which is the right-hand side.

So this is true. And now we have the sum of the squares of the eigenvalues, which is also a moment of the eigenvalues-- namely, corresponding to K_2 . So the same inequality is true for graph homomorphisms. And to get to the inequality for densities, we just divide by the number of vertices raised to the third power from both sides, and we get the inequality that we're looking for.

So that's the proof of the upper bound. Any questions?

There is something that bothers me slightly about this proof. Look, it's a correct proof. So there is nothing wrong with this proof. Everything is kosher. Everything is correct. You might ask, is there a way to do this spectral argument in graphons without passing to graphs? And yes, you can, because for graphons you can also talk about spectrum.

It turns out to be a compact operator, so that spectrum makes sense. You have to develop a

little bit more theory about the spectrum of compact operators, but everything, more or less, works exactly the same way. It's just easier to talk about graphs.

But what bothers me about this proof is that we started with what I would call a physical inequality, meaning that it only has to do with the actual edges and subgraph densities. But the proof involved going to the spectrum. And that bothers me a little bit.

There's nothing incorrect about it, but somehow in my mind a physical inequality deserves a physical proof. So I use the word physical in contrast to frequency, which is coming from Fourier analysis. And that's the next thing we'll do in this course.

But this proof goes to the spectrum. It goes to something beyond the physical domain. OK. It's neat. But I want to show you a different proof that stays within the physical domain. And this other proof-- I mean, it's always nice to see some different proofs because you can use it to apply to different situations.

And there are some situations where you might not be able to use this spectral characterization. For example, what if your K_3 is now K_4 ? A similar inequality is true, but this proof doesn't show it, at least not directly. You have to do a little bit of extra work.

So let me show you a different proof of the upper bound. And we'll prove a slightly stronger statement. Namely, that for all not just graphons-- it's not so important but for all symmetric measurable functions, from the unit square to \mathbb{R} , one has the following inequality-- namely, that the K_3 density in W is upper bounded by the K_2 density of W squared raised to power $3/2$. Here, the square is meant to be a point-wise square.

So a couple of things. If your graph is a graphon or a graph, then-- if it's a graph, then it's 0 comma 1 value. So taking this point by square doesn't do anything. If you're a graphon, you can always put one more inequality that replaces it by the thing that we're looking for because W is always bounded between 0 and 1.

So it's a slightly stronger inequality. Let me show it to you by writing down a series of inequalities and applying the Cauchy-Schwarz inequality repeatedly. So it's, again, an exercise in using Cauchy-Schwarz. And we will apply three applications of Cauchy-Schwarz.

Essentially, three applications-- one corresponding to every edge of this triangle. So let me begin by writing down the expression in graphons corresponding to the K_3 density.

I'm going to apply Cauchy-Schwarz by-- so I'm going to apply Cauchy-Schwarz to the variable x , holding all the other variables constant. So hold dy and dz constant. Going to apply to dz . You see there are two factors that involve the variable x . So apply Cauchy-Schwarz to them, you split each of them into an L_2 .

So one of these factors become that. By the way, all of these are definite integrals. I'm just omitting the domain of integrations. All the integrals are integrated over from 0 to 1. So the second application-- sorry, the second factor becomes like that. And the third factor is left intact.

So that's the first application of Cauchy-Schwarz. You apply it with respect to dx to these two factors. Split them like that.

AUDIENCE: There's a normalization missing.

PROFESSOR: Thank you. There is a normalization missing. OK. Guess what the second step is? Going to apply Cauchy-Schwarz again, but now to dy , to one more variable. Cauchy-Schwarz with respect to dy .

There are two factors now that involve the letter y . So I apply Cauchy-Schwarz and I get the following. The first factor now just becomes the L_2 norm of W . The second factor does not involve y , so it is left intact. And the third factor is again integrated with respect to y after taking the square.

And there's now dz that remains. Last step. You can guess, you integrate with respect to dz and apply Cauchy-Schwarz. Apply Cauchy-Schwarz to the last two factors. And there, actually, the outside integral goes away.

OK. So you get this product. And you see every single term is just the L_2 norm of W . So you have that, which is the same as what I wrote over here. Any questions? Yeah.

AUDIENCE: Where do you use the fact that W is symmetric?

PROFESSOR: Great question. So where do I use the fact that W is symmetric? So let's see. In some sense, we're not using the fact that W is symmetric because there is a slightly more general inequality you can write down. And actually, the question gives me a good chance to do a slight diversion into how this inequality is related to Holder's inequality.

So this is actually one of my favorite inequalities for this kind of combinatorial inequalities on graphons. So many of you may be familiar with Holder's inequality in the following form. If I have three functions, if I integrate them, then you can upper bound this integral by the product of the L^3 norms.

And likewise, if you have more functions. So if you apply just this inequality directly, you get a weaker estimate. So you don't get anything that's quite as strong as what you're looking for over there.

So what happens is that if you know-- so if f , g , and h each depends only on a subset of the coordinates in the following way that f depends on only x and y , g depends only on x and z , and h depends only on y and z , then if you repeat that proof verbatim with three different functions, you will find that you can upper bound this product, this integral, by the product of the L^2 norms.

So L^2 norms are in general less than or equal to the L^3 norms. So here we're inside a probability measure space. So the entire space has volume 1. So this is a stronger inequality, and this is the inequality that comes up over there. Yeah.

AUDIENCE: Is there an entirely graph theoretic proof of this-- say, for graphs instead of graphons-- that doesn't involve going to spectrum?

PROFESSOR: Great. So the question is, is there entirely graph theoretic proof of this? So the reason why I mentioned that this result is a special case of Kruskal-Katona-- so Kruskal-Katona actually is a stronger result, which tells you precisely how you should construct a graph. So given exactly m edges, what's the maximum number of triangles?

And the statement that there is actually-- it's a very precise result. It tells you, for example, if you have K choose 2 edges, you have at most K choose 3 triangles. It's not just at the density level but exactly. And even if the number of edges is not in the form of K choose 2, it tells you what to do.

And actually, the answer is pretty easy to describe. It's almost intuitive so if I give you a bunch of matchsticks and ask you to construct a graph with as many triangles as you can, what should you do? You start with one, two, filling a triangle. Start filling a triangle. Another vertex. 1, 2, 3, 4. You keep going. And that's the best way to do it.

And that's what Kruskal-Katona tells you. So that's a more precise version of this inequality.

And the Kruskal-Katona, the combinatorial version, is proved via a combinatorial shifting argument, also known as a compression argument. Namely, if you start with a given graph, there are some transformations you do to that graph to push your edges in one direction that saves the number of edges exactly the same but increases the number of triangles at each step.

And eventually, you push everything into a clique. So it's something you can read about. It's a very nice result. Other questions?

So we've solved the upper bound. So from examples and from this upper bound proof, we see that it's the upper bound. Now let me tell you a fairly general result that says something about graph theoretic inequalities but for a specific kind of linear inequalities. So here's a theorem due to Bollobas.

I'm interested in an inequality of the form like-- so I'm interested in inequality of this form, where I have a bunch of real coefficients, and I'm looking at a linear combination of the clique densities. I would like to know if this inequality is true.

So somebody gives us this inequality, whatever the numbers may be. You can also have a constant term. The constant term corresponds to r equals to 1. So the point density. That's the constant term. And asks you to decide is this inequality true. And if so, prove it. If not, find a counterexample.

So the theorem tells you that this is actually not hard to do. So this inequality holds for all G if and only if it holds whenever G is a clique. Maybe somebody gives you this inequality about-- it's a linear inequality about clique densities.

Then, to check this inequality, you only have to check over all cliques G , which is much easier than checking for all graphs. For each clique G this is just some specific expression you can write down, and you can check.

So I want to show you the proof of Bollobas' theorem. It's a quite nice result. But before that, any questions about the statement. All right. So the reason I say that this is very easy to check if I actually give you what the numbers are is because this inequality for cliques-- so the inequality is equivalent to just the statement of inequality that I'm writing down now, where I tell you precisely what the r clique density is in an n clique. Because that's just some combinatorial expression.

So to check whether this inequality is true for all graphs, I just have to check the specific inequality for all integers n , which is straightforward. All right. So let's see how to prove that inequality up there.

And here we're-- I mean, we're not going to exactly use the theorems about graphons, but it's useful to think about graphons. So if and only if one of the directions is trivial-- so let's get that out of the way first. But also-- so the only if is clear. So for the if direction, first note that this is true for all graphs if and only if it is true for all graphons and where I replaced G by W . By the general theory of graph limits and whatnot, this is true.

So in particular, there is one class of class that I would like to look at-- namely, I want to consider the set of node weighted-- so I want to consider the set of node weighted simple graphs. So node weighted simple graphs, by this I mean a graph where some of the edges are present and I have a node weight-- a weight for each node.

And to normalize things properly, I'm going to assume that the nodes' weights add up to 1. Now, you see that each graph like that, you can be represented by a graphon where-- so you can have a graphon. So they're not meant to be the same picture, but you have some graphon like this, which corresponds to a node weighted graph.

And the set of such node weighted graphs is dense in the space of graphons. In particular, as far as graph densities are concerned, they include all the simple graphs. So it suffices-- I mean, it's equivalent to-- the inequality is equivalent to it being true for all node weighted simple graphs.

But for this space of graphs, suppose that the inequality fails. Suppose that inequality is false. Then there exists a node weighted simple graph. I'm going to actually drop the word simple from now on. So node weighted graph H , such that f of H being the above sum is less than zero.

And there could be many possibilities for such an H . But let me choose. So among all the possible H 's, let's choose one that is minimal in the sense that it has the smallest possible number of nodes. So with this minimum-- has a minimum number of nodes.

And furthermore, among all H with this number of nodes, choose the node weights, which we'll denote by a_1 through a_n , summing to 1. Chooses node weights so that this expression, the

sum, is minimized.

And by compactness-- and now we're not even talking about compact in the space of graphons. You have a finite number of parameters. It's a continuous function. So just by compactness, there exists such an H for which the minimum is achieved.

This is minimizing over integers. And here, minimizing over a finite set of bounded real numbers. So the name of the game now is we have this H , which is minimizing. And I want to show that H has certain properties. If it doesn't have these properties, I can decrease those values.

So let's see what properties this H must have if it has the minimum number of nodes and f of H is minimum possible. So first I claim that all the node weights are positive. If not, I can delete that node and decrease the number of nodes.

I would like to claim that H must be a complete graph because if some ij is not edge of H -- so here i is different from j . I do not allow loops. It's just simple. Then let's think about what this expression f of H should be. So I don't want to write this down, but I want you to imagine in your head.

So you have this graphon H . I'm Looking at the clique density. It's some polynomial. In fact, it's some multilinear-- it's some polynomial in these node weights.

So I want to understand what is the shape of this polynomial as a function of the node weights. And I observe that it has to be multilinear in-- has to be multilinear in particular in α_i and α_j . It's a polynomial. That should be clear.

It is multilinear because, well, you have-- why is it multilinear? Why do I not have α_i squared? Either of you.

AUDIENCE: It says the 0 is not [INAUDIBLE].

PROFESSOR: So we're forbidding-- so here's α_1 , α_2 , α_3 , α_4 , α_1 , α_2 , α_3 , α_4 . So understand what the triangle density-- if you write down the triangle density as an expression in terms of the parameters, think about what comes out, what it looks like. And they essentially consist of you choosing a subgraph, which you cannot have repeats.

So it's multilinear. So it's multilinear in particular in α_i and α_j . So no term has the

product $\alpha_i \alpha_j$ in it because ij is not an edge. So here's where we're really using that we're only considering clique densities.

So the theorem is completely false without the assumption of clique densities. If we have a general inequality, general linear inequality, then the statement is completely false. So it's multilinear. So if we now fix all the other variables and just think about how to optimize, how to minimize f of H by tweaking α_i and α_j , well, it's linear, so you should minimize it by setting one of them to be zero.

And that would then decrease the number of nodes. So can shift α_i and α_j while preserving $\alpha_i + \alpha_j$ and not changing-- so not increasing f of H . And then we get either α_i to go to zero or α_j to go to zero, in which case we decrease the number of nodes, thereby contradicting the minimality assumption.

So this argument here then tells you that H must be a clique. So hence, H is complete. And if H is complete, then as a polynomial in these alphas, what should f look like? Well, it has to be symmetric with respect to all these alphas.

So in particular, it has to be-- so since H is complete, we see that, in fact, now you can write down exactly what f of H is in terms of the parameters described in the problem. Namely, it's C_r times r factorial times S_r , where S_r is a symmetric polynomial where you look at-- you choose r of the terms, r of these alphas for each term in this sum. It's just elementary symmetric polynomial.

And I would like to know, given such a polynomial, how to minimize this number by choosing the alphas. But if you think about what happens if you fix again everything but two of the alphas-- so by fixing all of, let's say, α_3 to α_n , we find that-- so as a function in just α_1 and α_2 , f of H has the following form.

And because it's symmetric, these two B 's are actually the same. So if we now vary α_1 and α_2 but fixing everything else, because $\alpha_1 + \alpha_2$ is constant, I can even get rid of this linear part. So that linear part is fixed as a constant.

I want to minimize this expression with $\alpha_1 + \alpha_2$, how it's fixed. So there are two possibilities depending on whether C is positive or negative or, I guess, 0. So now you're here. So depending if C is positive or negative, it's minimized by either the two alphas equal to each other or one of the two of alphas should be zero.

The latter cannot occur because we assume minimality. So the first must occur. And hence, by symmetry, if you apply the same argument to all the other alphas, all the alphas are equal to each other, which means that H is a simple clique. It's basically an unweighted clique.

So in other words, if this inequality fails for some H , some node weighted H , then it must fail for a simple clique H . And that's the claim above. Yeah?

AUDIENCE: So in the statement, there are two n 's, are those two n 's different n 's then?

PROFESSOR: OK. Question. There are two n 's. Yeah. Thank you. So these are two-- yeah. So these are two different n 's. Great. yeah.

AUDIENCE: I have a question. Which is the node weight such that f of H [INAUDIBLE]?

PROFESSOR: Question is, why can we assume that you can choose H so that f of H is minimized? Its because once-- OK. So you agreed the first thing you can minimize because the number of nodes is a positive integer. So if there's a counterexample, choose the minimum counterexample.

Now, you fixed that number of vertices, and the number of-- then this is an optimization problem. It's minimizing continuous function with a finite number of variables. So it has a minimum just by compactness of a continuous function. So I choose that minimizer.

Any more questions? So we have this rather general looking theorem. So in the second part of today's lecture, after taking a short break, I want to discuss what are some of the consequences and also variations of that statement up there. And I want to also show you what the rest of this picture looks like.

So let's continue to deduce some consequences of this theorem up there that tells us that it is pretty easy to decide linear inequalities between clique densities. Namely, to decide it, just check the inequalities on cliques.

So as a corollary for each n -- yes, for each n , the extremal points-- so the extremal points of the convex hull of this set where I record the clique densities overall graphons W . So think about this set as the higher dimensional generalization of that picture I drew up there.

But no previously we had n equals to 3, and we're still interested in n equals to 3. But in general, you have this set sitting in this box. And so it's some set. And if I take the convex hull

of the set, what that theorem tells us-- and it requires maybe one bit of extra computation. But what it tells us is that the extremal points are precisely the points given by W equals to K_m for all m equal to 1.

So evaluate, find what this point is for each m , and you have a bunch of points. And those are the convex hull. So I'll illustrate by drawing what the points are for the picture over there. But it essentially follows from Bollobas' theorem with one extra bit of computation to make sure that all of these are actually extremal points of the convex hull. None of them is contained in the convex hull of the other points.

So for example, we can also deduce very easily Turan's theorem. So what does Turan's theorem tell us? It tells us that if the r plus 1 clique density is zero, then the K_2 density is at most $1 - r$. So why does Turan's theorem follow from the above claims?

It should follow because all the data here has to do with clique densities. And everything we saw so far says that if you just want to understand linear inequalities between clique densities, it's super easy. Maybe I'll draw the picture for triangles, and then you'll see what it's like.

So the corollary tells us for this picture, corresponding to n equals to 3, what the points, the extreme point of the convex hull are. So let me let me draw these points for you. So one of these points is this $1/2$ comma 0. So that corresponds to Mantel's theorem.

Now, if you go to the other values of m , you find that those points-- so the extreme points-- they are of the form $m - 1$ divided by m , $m - 1$ $m - 2$ divided by m squared for positive integers m .

So for m equals to 2, that's the point that we just drew. And the next point-- so next two points, one third and one fourth, they are at, if you plug it in-- thank you. $2/3$ and $3/4$. They correspond to $2/9$ and $3/8$.

So let me show you where these points are. So they are at here and over there. And you have this sequence of points going up. So this is the convex hull. And from that information, you should already be able to deduce Mantel's theorem because this right half is not part of this convex hull. So that's what Mantel's theorem. And similarly, the deduction to Turan's theorem also follows by a similar logic.

OK. So you have this sequence of points. Now, it happens that all of these points lie on a curve. So let me try to draw what this extra curve is. So there is some curve, like that. So

there's some curve like that. The equation of this curve happens to be $x^2 - x$. And because the region is contained in the convex hull, the yellow points, it certainly lies above this convex red curve.

You've seen this red curve before. From where? So what is that saying? It's saying that if your edge density is beyond above one half, then you have some lower bound on the triangle density. Where have we seen this before? Problem set one. There was a problem on problem set one that says exactly this inequality. So go back and compare what it did.

But of course, the convex hull result tells you even a little bit more-- namely, that you can draw line segments between these convex hull points. So you have some polygonal region that lower bounds the actual region.

So what is the actual region? So leaving you in suspense. So let me tell you what the actual region is now. So it turns out to be actually-- it's beautiful and it's quite deep, that the region is now completely understood. And it's a fairly recent result. It's only about 10 years ago roughly that there are some concave curves. The sequence of scallops going up to the top right corner.

And this is now understood to be the complete region between these lower and upper curves. So this is the complete set of feasible regions for edge versus triangle densities. So this lower curve is a difficult result due to Razborov.

And I want to give you a statement what this curve is. And Razborov came up with this machinery, this technique, known by the name of flag algebra. So actually, he came up with this name. So I won't really tell you what flag algebra is, but it's kind of a computerized way of doing Cauchy-Schwarz inequalities.

So many of our proofs for this graph through inequalities, they go through some kind of Cauchy-Schwarz or sum of squares equivalently. But there are some very large or difficult inequalities you can also prove this way. But it may be difficult to find exactly what is the actual inequality-- the chain of Cauchy-Schwarz or the sum of squares that you should write down.

So this machinery, flag algebra, is a language, is a framework for setting up those sum of squares inequalities in the context of proving graph theoretic inequalities. So it can be used in many different ways. And notably, a lot of people have used serious computer computations. If I want to prove something is true, I plug it into what's called a semidefinite program that allows

me to decide what kinds of Cauchy-Schwarz inequalities I should be applying to derive the result I want to prove.

So that's what flag algebra roughly is. So what Razborov proved is the following. So Razborov's theorem, which is drawn up there-- that's the lower curve-- is that for fixed-- so for a fixed value of edge densities, if it lies between two specific points, drawn above, the minimum value of triangle density with a fixed value of edge density is attained via the following construction.

It's attained by the step function of the graphon corresponding to a K clique. So a complete graph on K vertices with node weights α_1 through α_K summing to 1, and such that the first $K-1$ of the node weights are equal. And the last one is smaller

All right. And the point here is that if you are given a specific edge weight, edge density, then there is a unique choice of these alphas that achieve that edge density. And that is the graphon you should use that minimizes the triangle density-- describes the lower curve.

So you can write down specific equations for the lower curve, but it's not so important. This is a more important description. These are the graphs that come out. And what is something that is actually quite-- I mean, why you should suspect this theorem is difficult is that unlike Turan's theorem-- so Turan's theorem, which corresponds to all those discrete points. In Turan's theorem, the minimizer is unique.

I tell you the number-- I tell you that the edge density is $2/3$, and I want you to minimize the number of triangles. Not from Turan's theorem, but it turns out that this extremal point is unique. Essentially corresponds to a complete three partite graph.

But for the intermediate values, the constructions are not unique. So unless the K_2 density is exactly of this form, the minimizer is not unique. And the reason why it is not unique is that you can replace-- so what's going on here? So you have this graphon. $\alpha_1, \alpha_2, \alpha_3$.

I can replace this graphon here by any triangle free graphon of the same edge density. And there are lots and lots of them. And the non-uniqueness of the minimizer makes this minimization problem much more difficult.

So Razborov proved this result for edge versus triangle densities. And this program was later completed to K_4 , and more generally, to K_r So K_4 is due to a result of Nikiforov, and the K_r

result of Reihher So a similar picture. It's more or less that picture up there but with the actual numbers shifted. Instead of edge versus triangle, it is now edge versus Kr.

I should say that it's worth-- so this is a picture that I drew up there, and this is roughly the picture that you see in textbooks-- how they draw these scallops. I once plotted what this picture looks like in Mathematica, just to see for myself where the actual graph is. And it doesn't actually look like that.

The concaveness is very subtle. If you draw it on a computer, they look like straight lines. So in some sense, that's a cartoon. So the concaveness is caricatured. So it's not actually as concave as it is drawn, but I think it's a good illustration of what's happening in reality.

Questions?

So on one hand, every polynomial graph inequality-- so what do I mean by a polynomial graph inequality? So something like-- suppose I have some inequality of this form. And I want to know, is this true? It turns out that I don't actually need these squares in some sense because I can always replace them by what happens if you take disjoint unions.

So all I'm trying to say is that every polynomial graph inequality can be written as a linear graph inequality of densities. But nevertheless, this still captures a very large class of graph inequalities. And if I just give you some arbitrary one that is not of that form, it can be often very difficult to decide whether it is true or not.

So over here it's not so hard. You just plug it in, and then you can decide whether it is true. I mean, it turns out to decide whether this inequality is true, it's really a polynomial. And then you just check. It's not too hard to do.

But in general, suppose I give you an inequality of this form. So some generalized version of a linear inequality, like that. It's even decidable if the inequality holds. Decidable in the sense of Turing halting problem. So is there some computer program give you this inequality is true? I wonder, can you write a computer program that decides the truthfulness?

It turns out-- OK. So before telling you what the answer is, let me just put it in some context. What about more classical questions before we jump into graph theory? If I give you some polynomial p over the real numbers and I want to check is that true-- so this is not too hard. So this is not too hard.

But what if you have multivariate for all real? Does anyone know the answer? Is this

decidable? So as you can imagine, these things were studied pretty classically. And so it turns out that every first order theory over the real numbers is decidable. So this is a result of Tarski.

In particular, such questions are decidable. And in fact, there is a very nice characterization of - so there's a result called Artin's theorem that tells you that every such polynomial, if it is non-negative, then if and only if, it can be written as a sum of squares of rational functions. So there's a very nice characterization of positiveness of polynomials over the reals.

But now I change the question and I ask, what about over the integers? So if I give you a polynomial, is it always non-negative if I have integer entries? Is this decidable? So turns out, this is not decidable.

And this is related. So it's more or less the same as the undecidability of Diophantine equations, which is also known as Hilbert's tenth problem. So there is no computer program where we give you a Diophantine equation and solves the question or even tells you whether the equation has a solution.

And this is part of what makes number theory, makes Diophantine equations interesting. So it's undecidable, but we talk about it. So undecidability is a famous result due to Matiyasevich.

So what about graph theoretic inequalities? So is a graph homomorphism inequality decidable? I mean, the question you should ask yourself is, which one is it closer to? Is it closer to deciding the positiveness of polynomials over reals or over integers?

On one hand, you might think that it is more similar to the question of polynomials over real. So first of all, why it's similar to polynomials, I hope that's at least intuitively-- nothing's a proof, but intuitively it feels somewhat similar to polynomials. And all of these guys you can write down as polynomial-like quantities. And we saw this earlier in the proof of Bollobas' theorem.

So you might think it's similar to reals because, well, for graphons, you can take arbitrary real weights. So it feels like the reals. So it turns out, due to a theorem of Hatami and Norine, that the answer is no. It is not decidable.

And roughly the reason has to do with this picture. Even though the space of graphons is not discrete, it's a very continuous object, even if you just look at this picture here, you have a bunch of discrete points along this scallop. So here's a potential strategy for proving the

undecidability of graph homomorphism inequalities.

I start by just restricting myself to this curve. I restrict myself to the red curve. If you restrict yourself to the red curve, then the set of possibilities-- it's now a discrete set, which is like the positive integers. And now I start with-- I can reduce the problem to the problem of decidability of integer inequalities.

I start with an integer inequality. I convert it to an inequality about points on this red curve. And that turns into a corresponding graph inequality, which must then be undecidable. So this undecidability result is related to the discreteness of points on this red curve.

So general undecidability results are interesting. But often, we're interested in specific problems. So I give you some specific inequality and ask, is it true? And there are a lot of interesting open problems of that type. My favorite one, and also a very important problem in extremal graph theory, is known as Sidorenko's conjecture.

So the main cause conjecture-- it's a conjecture-- says that if H is bipartite, then the H density in G or W is at least the edge density raised to the power of the number of edges of H . So we saw one example of this inequality when H is the fourth cycle. So when we discussed quasi-randomness we saw that this is true.

And in the homework, you'll have a few more-- so in the next problem homework, you'll have a few more examples where you're asked to show this inequality. It is open. We don't know any counterexamples. And the first open example, it's known as something called a Mobius strip.

So the Mobius strip graph, which is a fancy name for the graph consisting of taking a $K_{5,5}$ and removing a 10 cycle. So that's the graph. It is open whether this inequality holds for that graph there. And this is something of great interest. So if you can make progress on this problem, people will be very excited.

Now, why is this called a Mobius strip? This took me a while to figure out. So there are many different interpretations. I think the reason why it's called a Mobius strip is that if you think about the usual simplicial complex for a Mobius strip. And then this is the face vertex incidence bipartite graph.

So five vertices, one for each face. Five vertices, one for each vertex. And if you draw the incident structure, that's the graph. I'm not sure if this topological formulation will help you improving Sidorenko's conjecture or disprove it, but certainly that that's why it's called a

Möbius strip. And there are some people believe who believe that it may be false. So it's still open. It's still open.

The one last thing I want to mention is that even though the inequality written up there in general is undecidable, if you only want to know whether this inequality is true up to an epsilon error, then it has decidable. In fact, there is an algorithm that I can tell you.

So there exists an algorithm that decides, for every epsilon, that decides-- so I just want to know whether that inequality is true. But I allow an epsilon error, meaning it decides correctly this inequality is true up to an epsilon error for all G or outputs a G such that the sum here is negative.

So up to an epsilon of error, I can give you an algorithm. And the algorithm follows-- I mean, it's not too hard to describe. Basically, the idea is that if I take an epsilon regular partition, then all the data about edge densities can be encoded in the epsilon regular partition. So apply even the weak regularity lemma is enough.

And then we can test the bounded number of possibilities with some fixed number of parts. And by the counting lemma, you lose some epsilon error if I check over all weighted graphs on some bounded number of parts whose edge weights are multiples of epsilon, let's say, whether this is true. If it's true, then it is true with this epsilon. If it is false, then I can already output a counterexample.

So there is only finitely many possibilities as a result of weak regularity lemma. And therefore, this version here is decidable. So today, we saw many different graph theoretic inequalities and some general results. And there are lots of open problems about graph homomorphism inequalities.

So this concludes roughly the extremal graph theory section of this course. So starting from next lecture, we'll be looking at Roth's theorem. So looking at the Fourier analytic proof of Roth's theorem.