PROFESSOR: All right. Lecture 17 was about folding polygons into polyhedra, and today we will do it with real pieces of paper. But before we get there, I want to talk about some things we could make. In lecture, I demonstrated this perimeter having technique where you take any convex polygon and you pick any point on the perimeter. Then you measure out the perimeter halfway on each side. You get the antipodal point, y. And then just glue-- locally, you glue everything around x and around y.

So I mean, you just start. We call this zipping where you don't glue any extra material in here, you just start zipping these guys together. I guess these circles or when you happen to hit this vertex, well, then you just keep going-- zip, zip, zip, zip, zip. So that's one thing we could make with convex polygons.

Slightly more general is called a pita form. This is defined in the textbook. And the idea is instead of a convex polygon, you could take any convex body. It could have some corners or it could be smooth. So this is what we call a convex, or what I'm going to call today a convex 2D body, meaning it can be smooth in addition to polygonal. And then you pick any point on the boundary and you measure out the perimeter halfway, x and y, and you glue in exactly the same way.

Now, it does not follow from Alexandrov's theorem that this will make a convex thing, but it follows from a slightly more general form of Alexandrov's theorem called Alexandrov-Pogorelov theorem. So Pogorelov was a student of Alexandrov and he edited some things that Alexandrov wrote, but then also Pogorelov wrote his own papers and proved some stronger versions of Alexandrov's theorem that hold for smooth bodies, not just polygonal folding. So Alexandrov's theorem, you recall, if you have a convex polyhedral metric homeomorphic to a sphere, then it's realized by a unique convex polyhedron.

So the Pogorelov's extension is that if you have everything except the polyhedral part-- so if you have a convex metric homeomorphic to a sphere-- so topologically it's a sphere, but we omit the polyhedral part. So polyhedral was that you have finitely many vertices. In some sense, I have infinitely many vertices, infinitesimal curvature at every point here. Then it's realized by a unique convex 3D body.

So same concept of a convex object. Well, I should say this is really the surface of. So a convex body, I should probably define that. Convex body is one where you take any two points and you draw a straight line and you stay within the body for the entire straight line. So that's convex. You take any two points, like this one and this one, draw a line between them, and you stay inside.

The body is the piece of paper we're gluing up. Here we have a 3D body, but then we just look at the surface of it, the boundary. And that's going to be the surface we make. Technically, that surface is not convex. It's the interior that's convex because you can draw lines through it and you stay inside.

This is exactly the same thing, but instead of polyhedron here, we have body because we didn't have finitely many vertices. And I'm not going to prove this theorem. But existence is basically the same proof. You just take a limit. So you can polygonalize with lots of little edges here, add little points, and take the limit as that approximation gets very, very close to the actual curve.

At all times you have a polyhedron, it's easy to show you will converge to a convex body. The uniqueness is the harder part, I believe. I haven't read the proof, but it was more challenging. I think that took 20 years or so to settle uniqueness. Alexandrov's like 1950. This theorem I think is 1973 the final version. But the version without unique was proved back in 1950.

So that's something else we can do. Instead of cutting out a polygon, we can cut out a nice smooth thing or you can add some kinks and be smooth elsewhere. As long as it's convex, then you're guaranteed that-- we're only gluing together two convex points. It could be flat or it could be strictly convex, which means we'll have at most 360 degrees of material at any point because it's always doing two at most 180's.

And so we will have a convex metric, meaning always at most 360 of material everywhere. So then this theorem applies so we'll get a unique thing. Another thing we could make is called a D-form. And I have some examples of D-forms here. So the idea with the D-form is you take two convex 2D bodies-- so here they have some straight parts and some curve parts-- of the same perimeter.

I pick one point from one body and I pick another point from the other body. I attach them there. And then I zip. So it would be zip, zip, zip, zip, zip. Wherever my fingers were at the same time, those are points that get glued together. And this is with the convex body you get.

Again, Alexandrov-Pogorelov applies because we're still only gluing two convex points together at any point. As long as these guys have matching perimeter, we'll be OK. And so we will get a unique convex 3D body. In this case, we get this fun kind of shape.

So this is what I suggest you make. D-forms tend to look a little bit cooler somehow because they have two polygons. And the one that's particularly easy to build is two copies of the same shape, but then you don't want to glue corresponding points. Like if I glued this point to that point, I would just get a flat, doubly covered thing and that's boring. But you just pick some other point and glue it to non-matching points, then it acts like you have two very different convex bodies which is kind of cool.

I have some images of D-forms here. There's actually a whole book on D-forms by John Sharp. They're invented by this guy Tony Wills who's an artist. At the top, the quote is, "there is no such thing as an ugly D-form." So we are guaranteed success here. And also it says, "they surprise, confuse, can be addictive. They're an intellectual virus. Beware." So you've been warned.

This is Tony Wills, the inventor. Now, from the artistic perspective, you don't have to start from convex polygons. We're going to start from convex bodies because they're guaranteed to work. If you're careful to satisfy Alexandrov-Pogorelov, you can have a convex metric even though you start from nonconvex bodies, although in this case, he's actually getting nonconvex shapes as well.

More typical D-form is something like this. This is made from two ellipses and he builds them out of metal. This is a simulation by this guy Kenneth Brakke who has this software called Surface Evolver. And it's usually used for computing minimal services, soap bundles and things like that. But here it's computing what D-forms look like, approximately.

And so this is an example of taking two ellipses and instead of-- if you glue at matching points, they would be flat-- but then you change the rotation. So here, they're rotated by 22.5 degrees, rotated by 45 degrees, rotated by 90 degrees, so you get this nice continuum of different things that you could zip together. So I propose we make some.

And so you have your scissors and tape. If you don't have two sheets of paper, get some from the front. And so everyone have scissors, tape, everything? We're going to have to share a bit.

So what we want to do is make two identical convex shapes. So you have some here. Two rectangles are valid convex shapes. You could cut not all. Everyone should do something different. I'm just going to make sure they're nicely aligned. Hold on tight because you don't want them to let go, and then just cut. Always turn left with your cutting.

You can be smooth. You can have some corners, what you want to do. Just always turn left. And once you have your two shapes, you pull them apart. Make sure you do not glue corresponding points. Pick some other points and attach them together with tape. Here's where it gets a little bit challenging.

With smooth curves, it's little harder to tape them together, so you're probably going to have to use a bunch of points of tape, so to speak, not literal points. And I just tape, tape, tape.

Before you get all the way, you want to make sure you can still pop it up and be convex. I think we'll be able to do that by [INAUDIBLE] inside the limit. Give me some scissors. All right. There's my D-form. Mostly convex, kind of fun.

So a funny thing you'll notice about D-forms, if you made yours smooth-- so if you don't have any kinks on the outer boundary-- then the surface looks smooth, which is kind of nice. The other thing is there's this noticeable seam where you taped things. There, you're not smooth, obviously. You've got a crease along the seam.

But also it looks like all that the surface is doing is kind of taking the convex hull, the envelope of that seam. So there's two properties we observe. One is that it looks smooth, except at the seam. And the other is that the hull surface is the convex hull of the 3D seam.

And I want to prove both of those things. Those are both true facts. But first I need to define some things. So we're going to prove that D-forms are smooth and other good things. This is a paper that was originally a class project in this class, I think in 2007. Sounds about right.

Submitted during class, it looks like, or just after semester with Greg Price. D-form have no spurious creases. So what's a D-form? You take two, for us, it's going to be two convex shapes of equal perimeter. You glue two points, two corresponding points P and Q together, and then you zip from there. We can actually talk about something even more general which we'll call a seam form.

A seam form, you can have multiple convex shapes. It's a little harder to imagine, but you can join multiple points together. Maybe you join all three of those corners together. As long as you have at most 360 of material everywhere, it'll be fine. Maybe do some zipping.

I'm not going to try to figure out exactly how this gluing pattern works, but you find a gluing pattern where you don't violate anything. So it looks like if I do this one, I better have these guys also a little bit sharp so that they don't add up to too much angle there. So these three points would also join together. This would be like three petals in like a seed pod or something, like those leaves that wrap around little fruits or something.

I won't try to draw it but-- I guess I will try to draw it. Something like this. So that's a thing. This could even fold not exactly a sphere, but something like a sphere. So seam forms, you could have any number of convex shapes and glue them together however you want, provided at all points you satisfy Alexandrov-Pogorelov, so you only have at most 360 degrees of material and you are topologically a sphere. Then you have a seam form.

So most of the theorems I'm going to talk about apply to general seam forms, but we're interested in D-forms and pita forms in particular. So let me tell you some theorems about seam forms. Theorem 1 is that a seam form equals convex hull of its seams. So the seams in general are these parts where you did gluing, wherever you taped stuff.

The boundaries of the convex polygons, those are the seams. They map to some curves in 3D. If you take the convex hull, the claim is you get the entire seam form. That's actually pretty easy to prove.

Then the second claim is there aren't many creases other than the seams of course. The seams are usually going to be creased. And these creases are going to be line segments. And not just any line segments, but their endpoints are pretty special.

So the claim is that the creases have to be line segments basically connecting strict vertices. What do I mean by strict vertex? Something like these three points which come together, provided the total angle here is strictly less than 360. I call that a strict vertex. Whereas these points connecting from here to here, those don't count as strict vertices.

They do have curvature, but only infinitesimal curvature. Because this is essentially 180 degrees of material, just slightly less because it's curved. So only where I have kinks can I potentially have strictly less than 180. And if I join them together to be strictly less than 360, that's a strict vertex.

So potentially, I have a seam coming from here, also at the other endpoint. So in this picture, I might have a crease like this, I might have a crease like this, I might have a crease like this, potentially. But in something like a D-form, there are no strict vertices.

If these are smooth-- so I'll assume here these are smooth. We could call it a smooth D-form if you like. Then there's no vertices, no strict vertices, only these sort of barely vertices. So there could be no creases.

There's one other situation which is creases could be tangent to seams. This should seem impossible because the way I've defined things it is impossible. You're tangent to a seam, that would be you're like this. How am I supposed to have a crease inside that's tangent to the seam? The answer is you can't. If these regions are convex, you can't have them.

This statement is actually about a more general form of seam forms where you can have a nonconvex piece. So in general seam form, you have a bunch of flat pieces and you somehow search for them together so that you satisfy Alexandrov-Pogorelov. And now you could have a tangent. I could, for example, have a crease that emanates from there or have a crease that emanates from there.

But right now, it has nowhere to stop. The only way for these creases to actually exist is if there's a kink here, then this could go from there to there. So this is starting at a vertex and ending tangent. You could, of course, also start tangent and end tangent if there's another bend.

But the claim is, all crease look like that, which means if you carefully designed your thing or vaguely carefully such as a smooth D-form, then you're guaranteed there are no creases which is what we observe. So this is justifying our intuition from these examples. One other case is the pita form.

So pita form does actually make vertices. This point x is going to be a strict vertex because it has less than 180 of material, so definitely strictly less than 360. Same with y. So pita form is potentially going to have a seam from x to y, but that's it. Sorry, a crease from x to y. It has at most one crease. I think most examples we've made do have such a crease from x to y.

Let's prove these theorems, or at least sketch the proofs. Don't want to get too technical here. So to prove the first part that the seam form is convex hull of its seams, it's helpful to have a tool here which is a theorem by Minkowski. We'll call it Minkowski's theorem, although he has a bunch. Hermann Minkowski.

It relates convex things to convex hulls. It says, any convex body is the convex hull of its extreme points. What are extreme points? Extreme points are points on the surface where you can touch just that point with a tangent plane.

So if you think of polyhedra, these are vertices of the polyhedra. But I want to handle things that are smooth, so they may might not really have any vertices. In general, I have some convex body. If I can draw-- imagine this is in two dimensions-- if I can draw a tangent plane that just touches at a single point, then I call that point extreme.

Let's look at an example here. Here I believe every seam point is going to be an extreme point because I can put a tangent plane. It just touches at that point. What am I distinguishing from?

Well, for example, this point, this surface is developable. We know it's ruled. So there's a straight line here. You can see it in the shadow, in the silhouette. If I said, is this point extreme? I try to put a tangent plane on, the only way to get a tangent plane is to include this entire line. So it's impossible to just hit this point or any point along that line.

In fact, every interior point, I claim, you cannot just hit that point. You've got to hit an entire line. Whereas at the seam, I can do it. I can angle between this and this. There's a tangent plane, because this is curving, that only hits at one point. If I had a straight segment, then the endpoints of the straight segment of the seam would be extreme, but the points in the middle of the segment would not be extreme.

So that's the meaning of extreme. And we're going to use this because I claim that these interior points can never be extreme. It can only be the seam points that are extreme. And therefore, we are the convex hull of the seam points because we are the convex hull of the extreme points.

How do we argue that? So I claim an extreme point can't be locally flat and at the same time be convex, because convex in 3D, that's what we need. If you had a point and it just meets a tangent plane at that single point, that means locally you're kind of going down from that plane, if you think of that plane as vertical. So this is a situation. I want to prove this.

To prove it, I need to introduce another tool which is a generally good thing for you guys to know about, so kind of an excuse to tell you about the Gauss sphere, which I don't think we actually cover in lectures, but might come up again. Gauss sphere is a simple idea. For every tangent plane-- let's just think of this single point. I want to look at the-- it's called the tangent space. Look at all the tangent planes that touch this point.

Because this particular tangent plan I drew only touches that one point, it has some wiggle room. I can pick any direction and just wiggle and rotate the plane around this point and it won't immediately hit the surface. I've got a little bit of time before it hits the surface. So there's a two-dimensional space of tangent planes that touch just this point. Because there's at least one, I've got to have some wiggle room.

So I want to draw that space. And a clean way to draw the space is for every such plane to draw a normal vector perpendicular to the plane, take the direction of that normal vector, and draw it on a sphere. It's a unit sphere. So this is a sphere.

This one looks pretty vertical so that would correspond to the north pole of the sphere. So that direction becomes a point on the sphere. Think of it as this vector from the center of the sphere to the north pole. But we'll just draw it as a point. And because I've got a two-dimensional space of maneuverability or rotatability of this plane, I'm going to get some region-- which maybe I should draw like this-- of the sphere. Those are all the possible normal vectors of tangent planes at that point.

So this is called the Gauss map where you map all these normals to the Gauss sphere, which Gauss sphere is just a unit sphere. Fun fact. The area of that thing, which is the map of all the tangent normal directions, equals the curvature at that point. You may recall at some point curvature is actually Gaussian curvature, so that's why Gauss is all over the place here.

Now because I claim this is a two-dimensional space, this thing will have positive area. Therefore, at this point, that vertex has positive curvature on the surface. And yet, it was supposed to be in the middle of one of these flat shapes. That's supposed to have zero curvature. Contradiction. Or stated more positively, if I have a point that is an extreme point, it's only touched by one of these planes, I have this flexibility. Therefore, I get area. Therefore, it is not locally flat, so it must be a seam point.

It could be one like this where I'm barely non-flat, but I am non-flat. There's kind of infinitesimal curvature here. Or it could be one of these points where I'm very non-flat, or I guess X and Y here are other examples of here I'm very not flat. I've only got 180 degrees of curvature roughly.

So it's got to be one of those if you're going to be an extreme point. Therefore, all the extreme points are seam points. Therefore, convex hull of the extreme points is the convex hull of the seams, and that is your seam form. So that's part one. It's pretty easy.

Let me tell you about part two. Part two builds on part one. It's one reason why we care about. So part two is that there are no spurious creases. Unless you have strict vertices, then you could connect those up.

So let's look at a locally flat crease point. Actually, at this point, I should probably mention. This paper, I think of as a forerunner to the "How Paper Folds Between Creases" paper which is that hyperbolic parabolas don't exist. You may recall there were theorems about what creases look like, what ruled surfaces look like, and so on between creases.

In that setting, we've proved things like straight creases stay straight and all these good things. Here, it's a little bit different because we have kind of two levels of creases. There's the seam which is special. And then we're imagining hypothetical creases between the seam. And could you have a curved crease? Claim is no.

Claim is all the creases in between the seams have to be straight. And in fact, they can't look like this. They have to be at corners somehow. So that's what we're going to prove. But a lot of the same techniques, most of the same definitions. This paper was kind of a warm up, I feel like, for the nonexistence of hyperbolic parabolas.

Although they use a lot of the same tools. They have two of the same authors, but there's no theorem that really is shared between the two. They're not identical.

So we're looking inside a flat region. We're imagining, let's look at a crease point. It lies on some crease locally. And it's a fold crease. So I'm going to wave my hands a little bit.

But imagine a folded crease. Maybe it's a curve. We don't know. It's folded by some angle, some nonzero angle. Otherwise, it's not a crease. So locally, it looks like two kind of planes, at least to the first order they're going to be planes. Maybe something like this.

I'm going to draw the creases straight because to the first order, it is straight. But it might be slightly curved. So we're looking at a point here. And my point is, it's bent by some angle. So I'm going to draw the two tangent planes. There's a tangent plane on the right.

Whatever the surface is doing on the right, there's a tangent plane over there. Whatever the surface is doing on the left of the crease, there's a tangent plane there. They're not the same plane because we are creased. Whatever the crease angle is, that's the angle between these two planes.

I want to look at this tangent space again a little bit. You can think of it, the tangent space at the least it has this tangent plane. For this point, you have at least this tangent plane and at least this other tangent plane. And so in particular, you have to have the sweep between them. So you get at least a one-dimensional set there.

Now, this can happen. So we're not going to get a contradiction. We're not going to suddenly discover there's a two-dimensional area and get that it wasn't flat. But at least we've got a one-dimensional arc. On the Gaussian sphere here, we've got something like this so far.

This is one extreme tangent plane on the left. This is extreme tangent plane on the right. These two points correspond to those two. And then we've got the arc of-- you can sweep the tangent plane around and still stay outside the surface,

Well, that's interesting. It's interesting because all of those tangent planes share a line, which is going to be this line that I drew here. I mean, you look at these two planes, there's a line that they share. And if you rotate the plane through that line, you'll still pass through this point and you won't hit the surface. So in fact, all those tangent planes share this line.

I think we actually just need these two. So if you don't see that, don't worry. There are these two planes. There's a line there. They're tangent, which means that line is on or outside the surface because tangent planes don't hit the surface. I mean, they touch the surface barely, but they don't go interior.

Well, I claim that in fact the surface must touch this line. I claim the surface, at least locally, has got to follow along that intersection line of the left and right tangent planes. Why? It comes down to this Gaussian sphere thing.

We know that the surface lies inside these two planes. There's this wedge here. The surface must be below. We also know the surface comes and touches it at this point. Could it, from here, kind of dip down away from this line? I claim no.

If it dipped down, then there'd be a third tangent plane. The tangent plane could follow and dip down as well. And then we've got not just a one-dimensional arc, but we're actually going to get a whole triangle in the Gauss sphere. Once you have a whole triangle, that means you weren't flat. You have positive curvature.

So we're looking at a locally flat crease point somewhere in the middle of one of these polygons. So in fact, you can't afford to dip down because then you'd have at least a little bit of curvature there. Because you have zero curvature, you've got to stay straight. So that means that the surface locally has to exist along the segment in both directions until something happens.

Now, we assumed that our point was locally flat and a crease point. So the surface must continue in this direction until we reach a point that is either not locally flat or not a crease point. I claim it's got to remain a crease point. Because look, you're following along this intersection line of these two tangent planes. These remain tangent planes.

These are not only tangent planes at this point. They're also tangent planes at this point and potentially any point along this segment. As long as the surface is here, they are tangent planes. Therefore, you are creased here. I mean, you can't be smooth if you're butting up against these two tangent planes.

So crease must be preserved, so the only issue is locally flat. It could be at some point you become not locally flat. And that indeed happens. That's when you hit the seam. When you hit the seam, then you're no longer locally flat.

So this proves the creases must be line segments between two points on the seam. There's a little bit more to prove which is it either hits a strict vertex or it's tangent to seams. And this basically follows the same argument that if you hit something other than a strict vertex or a seam, you would have positive curvature where you shouldn't, basically.

So I'll just leave it at that. I think that's enough for the proof. You get the idea why creases can't exist. Any questions about that?

I have few more questions. One question is, can we see more examples of rolling belts? And in some sense, this example is an example of a rolling belt. So D-forms, you take two, say, ellipses. And the seam that connects them is a rolling belt. And this is an illustration of the rolling action.

As you change, which point glues to which point? Let's say I fix this point. I can glue to this one or to this one or to this one or to this one. And then I zip around. That's the rolling of the belt. More generally, I'll draw a picture.

More generally, a rolling belt from the viewpoint of a gluing tree is that you have basically it's like a conveyor belt like at the airport or on a tank or something. You've got treads here and it's going to roll. So if I mark a point, let's say I mark this point. And I roll it around this way a little bit, this point maybe stays stationary and I end up with something like this.

And what that means is whatever this guy glued to over here is now over here. So we're changing the gluing. This guy's now getting glued to some point like this because it kind of rolls around. So in general, you're rolling this, but then you're always just gluing across this path.

For this to work, every angle out here must be less than or equal to 180, like the material that's out on the outside. Remember, the gluing tree has the outside of the polygon in here. The inside of polygon's out here. So all the material's on the outside here.

All of these things must be less than or equal to 180. If you can find a path in your gluing tree where you've got less than 180 material all around you, then you can do this rolling. And no matter how you roll, and then just glue across, you will still have a valid gluing, a valid Alexandrov gluing. So that's what rolling belts look like in general.

We see them all over the place with these D-forms. The simple example I showed in class was when you take a rectangle, you glue into a cylinder. And now up here, I've got a rolling belt. Because if you look at the sides here, I've got actually exactly 180 degrees of material all around the belt.

And so I could collapse it like this or I could collapse it-- let me tape it. A little hard to hold everything. If you look at this belt, I could collapse it like this. Or I could roll it a little bit and classify like that. Or I could roll it a little bit, collapse it like that. Roll it a little more. Collapse it like that.

All of these are valid gluing because they're always gluing 180 to 180. And that's what rolling belts are. Does the broken applet work now? I don't know why it wasn't working. Probably some conflict with suspension or just the Java installation was bad.

But here it is. It's freely available online. You draw your graph. It's a little hard to just use. It's hard to draw an exciting example. Here I'll a draw a tetrahedron because that's simple.

In this case, it's abstract, so all of these edges are unit lengths. It's kind of like TreeMaker. You separately specify the lengths. Then you say Compute. And it will find, in this case, it's a tetrahedron, an equilateral. A regular tetrahedron is a little hard to see because we've got either a weird field of view or a lot of perspective here.

But this is a regular tetrahedron. You can verify that by say making some of these lengths longer. I'll make three of them five. So then this is like a pyramid. Of course, we're just essentially applying Cauchy's rigidity theorem here and say, oh, this uniquely assembles into this spiky tetrahedron.

You could, in principle, drawn an across here with the appropriate gluing. It's a little tricky to draw in because you have to also triangulate. It'd be hard to do from scratch to compute gluings. But in theory, especially if you use software to compute the gluings and compute shortest path in a glued polygon-- which we'll be getting more to next class. Next lecture is about algorithms for finding gluings. Then computing shortest paths.

Then you could use that as input to this software which will then compute what the 3D polyhedron is. So that's the Alexandrov implementation. This is by, I think, a student of Bobenko and Izmestiev.

One quick question about-- running out of time-- about the nonconvex case. So we're doing all this work about making convex surfaces. I said nonconvex case is trivial. In a sense, the decision problem is trivial. It's not an easy theorem to prove.

And aha, taped it up. It follows from what's called the Burago-Zalgaller theorem. And it says that if you tap any polyhedral metric-- so that means you take a polygon, you glue it however you want. You could even add boundary. Pick any polyhedron metrics, so you take any polygon and glue it however you want.

It could have handles. It could be a doughnut. It could be whatever. Then it has an isometric polyhedral realization in 3D. So there is a way-- so we're doing it abstractly. We're gluing stuff together. It's kind of topologically. We don't know what we're doing.

But then there is a way to actually embed it as a polyhedral surface in 3D without stretching any of the lengths. So it's a folding of the piece of paper. It could have creases in it, will have actually tons of creases in it. And it has exactly the right-- it connects all the things you want it to connect together.

It's kind of crazy. Furthermore, it could be self-crossing. But if your surface is either orientable-- so it doesn't have any cross caps or Mobius strips in it-- or it has boundary, then there's even a way to do without crossings. So that's basically all the cases we care about.

You could glue this together into a sphere or into a disk or into a torus or two-handled torus or whatever. No matter what you do, there is a way to embed it in 3D as a presumably nonconvex very creased surface. This is a hard to prove. If you're familiar with Nash's embedding theorem about embedding Riemannian surfaces, it uses a bunch of techniques from there called spiraling perturbations, which I don't know.

And one description I read calls the resulting surfaces "strongly corrugated," which I think means tons of creases. It is finitely many creases. That's the theorem. There's no bound on them. As far as I know, there's no algorithm to compute this thing. I've never seen a picture of them, otherwise, I'd show it to you.

So lots of open problems in making this real. But at least the decision problem is kind of boring. You do any gluing. It will make a nonconvex thing. To find that thing is still pretty interesting question though. That's it.