[SQUEAKING] [RUSTLING] [CLICKING]

BERTHOLD We're talking about representations for three-dimensional objects-- in particular, those that can't conveniently be
HORN: represented as polyhedra. And one representation is the extended Gaussian image, and for that, we needed to talk about Gaussian image and Gaussian curvature.

And the Gaussian image is a correspondence between the surface of an object and points on the unit sphere, simply based on the equality of surface normals. And we can extend that from points to areas. And if we do, then we can talk about curvature, in that, if the surface is highly curved, then the corresponding area on the sphere will be large. And if the surface is not curved very much, then things will be compressed. In the case of a planar surface, everything ends up in one spot and we have an impulse.

So the Gaussian curvature is just the ratio of those two areas. And we saw that in the case of a sphere that obviously becomes 1 over R squared because the area on the object is 4 pi R squared and the area of on the unit sphere is 4 pi.

OK, what do we do with that. Well, we're going to actually use the inverse of that quantity, and we will plot that as a function of position on the sphere. And that, in a crude way, can be thought of as defining how much of the surface has a normal that points in that direction.

Now, of course, for a convex, smoothly curved objects, there will typically be only one point that has exactly that normal. But if we take a small area around that, we can extend that idea and make that work.

One thing we can do with this quantity is take its integral either over the object or over the sphere. So for example, we can integrate the Gaussian curvature over some patch on the object and then change variables. And we get the area of the corresponding patch on the sphere.

And this is called the integral curvature. No surprise there. And one nice feature of the integral curvature is that it applies even when curvature itself can't be computed at a point because of discontinuities in surface orientation, like edges and corners.

But before we talk about that, let's go the other direction. Instead of integrating on the object, let's integrate on the sphere. 1 over K. And change variables.

And what do we get? It's the area on the object-- area on object-- that corresponds to that. And so that's a quantity, actually, that we're even more interested in. So it's saying that, if I take this quantity, which we'll call G, for Gauss, and we integrate it over a patch on the sphere, what we get is the area of the corresponding part of the object. So that's the generalization of that idea of-- it's the area with that surface normal. Well, here, we allow for the fact that the surface normal can have some variation and all of the points in this patch have surfaced normals that end up in that patch.

And so I'll just say something quickly about integral curvature. So suppose I have the corner of a cube. What's the curvature? Well, it's 0 there, 0 there, 0 there, and infinite on the edge, maybe-- infinite at the corner. So how can I talk about curvature?

Well, one thing I can do is take the integral of curvature of some area like that, and that captures the total change in orientation in that patch. And that's what this is computing. And so what does this look like on the sphere?

Well, we have these three distinct surface normals, and they correspond to three points on the sphere. They are 90 degrees apart in latitude or longitude. And what's the area of that? Well, to understand that, what we want to do is take a file to this cube and smooth off the edges so that, if we take a cross-section through a corner, it doesn't look like this but it looks like that.

And then, if you like, you can take the limit as you make the radius of this corner smaller and smaller. So in the ideal object with sharp corners, we only have these two surface orientations, and that's it. But if we think about it being a smooth transition of some sort, then we get all sorts of positive linear combinations of those two orientations. So that means that, if we think of this edge, for example, there's going to be some great circle on the sphere that corresponds to surface normals that-- very smoothly going from that one to that one.

And similarly, I can think of this edge as corresponding to that great circular arc and that edge here corresponding to this one. And then when I look at the corner, it's harder to draw, but I'll have, also, transitional directions, and they'll all be positive combinations of these three.

And so, at the corner, I'm actually dealing with surface orientations that are within this patch. So that means the integral curvature of that corner is the area of this patch. And that is-- it's one octant of a sphere. So the whole sphere has an area of 4 pi units here, and we have one octant.

So the integral curvature of the corner of a cube is pi over 2. And now you can imagine this could apply to other things, like, if it's a parallel pipette instead of a cube, you'll get a related but different result. And we can apply it to other objects, like cylinders and cones and so on. So it's a useful concept. We won't be doing much with it. We're mostly going with this integral. That's the one we're interested in.

So we're going to end up with some distribution on the sphere-- we will call G. And it'll depend on the orientation, which we can describe in various ways, like surface normal, unit normal. One question is, can we have any sort of distribution on the sphere, or are there some constraints?

So we saw earlier, when we were talking about polyhedra, that there might be a constraint. So with polyhedra, we saw that, if we create these vectors which have a length proportional to the area in the direction of the surface normal, that those all have to add up for it to be a closed object.

Well, something similar happens here. So the idea is that we have some object here-- let's take the discrete case first. We have facets. So here's a facet with some surface normal, and we imagine the whole thing is a mesh of facets like this. And then suppose we look at it from far away over here. Then the facet will appear foreshortened if the surface normal is not pointing directly at the observer.

So let's suppose we call this direction v. So this is the apparent area of-- suppose this is patch i. So that's the actual area of the patch. And that's the surface normal. And the way it appears to us is based on the cosine of the angle. And we can get that just using the product.

OK, well, now I'll see some part of this convex object, and I won't see other parts. So what am I going to see? Well, only the ones where that product is positive.

So that's the total cross-section that I see when I'm looking at the object from that direction. So I step through all of the triangles in this mesh, or whatever shape they are, and I pick the ones where the surface normal is within 90 degrees of the viewing direction, and I compute that sum. So-- what?

Well, now, if I look at it from the other side, I'll get a similar sum, except now this is reverse minus v. And the other thing that's changed is that now I'm only seeing things that have a positive product with minus v.

So we're looking at this object from one side. We see some area. It's like we were projecting it to the orthographic projection. And then we look at it from the other side, and, of course, there will be mirror image reversed-- that outline, the silhouette. But it should be the same area, right? And so these two are the same, right?

I'm looking at it from this side. Something appears foreshortened. I add all of those up. And presumably, there's no overlap in these sums. There's the borderline case, where this is equal to 0, but then that will contribute 0 to the sum. So we don't care.

Well, it turns out that now we've listed all of the facets for a convex object. We can see all of the facets from one side or the other. So each facet will appear either in this sum or in that sum. And then I bring this over to the other side, which flips the sign, and so that means that, actually, now the sum of all facets-- so I move this to the other side, and I end up with 0 over here. So the sum of all the facets is 0 when I take this dot product.

OK, now this is true for all possible viewing directions. Doesn't matter which viewing direction I look at it from. If I look at it from the other side, I don't have the same cross-sectional area. And so that means that this part better be 0, because, otherwise, if this wasn't 0, I could just pick a viewing direction in the direction of that vector, and I'd have a non-zero result.

So that means that sigma Aisi is 0. That's a vector equation, not just scalar. And so that means, really, that the centroid is at the origin.

So if I think of these areas as masses on the unit sphere, each of these facets gets mapped onto the unit sphere at a point that depends on this direction. And I put a mass down there that's proportional to this area. So I have this distribution on the sphere, and this is saying that the center of mass is at the origin, at the center of the sphere.

So these EGIs are distributions on the sphere, but there's one restriction. They can't be arbitrary. They have to have this property that the centroid is at the origin, and that corresponds to the surface being closed.

For example, I had that geometric object with conical and cylindrical parts. And if I plot that on the sphere, I have a small circle and a great circle. And if I stop there, then the centroid obviously is not at the center of the sphere because I have this great circle-- yes, centroid of that is at the center of the sphere. This one-- no, it's over here somewhere along this axis, and so the combination is off center.

So what's wrong? Well, what's wrong is that I forgot the plate at the back that closes it off, which would contribute to a big mass on the other side of the sphere, just enough to counteract that small circle of mass. And so the overall centroid is at the origin. So that's the discrete case.

I won't bother with going over the same argument in the continuous case. But if we take the integral of the mass on the sphere, the density on the sphere with the direction of point on the sphere-- that also is 0. So if I think of the EGI as a mass distribution on the sphere, its centroid is also going to be at the origin.

OK, now, we're going to look at discrete implementations. We have some surface data, perhaps from some machine vision method like photometric stereo. And we're going to map it onto the sphere in a discrete fashion.

But it's also useful to think of the continuous case. And one reason is that, if you have a model of an object, if it's a geometrically-defined object like the thing up there, you can actually exactly compute what its EGI is rather than having to approximate it. And so there's some advantage to doing that.

OK, examples. Now, we already know that, for a sphere, the EGI is very simple. It's just R squared because G is 1 over K, and k is 1 of R squared. And we got that from the ratio of the areas of the-- and this is symmetrical. So there's just one value I'm writing down. That's because G is the same everywhere around the unit sphere. So it gets more interesting when we have other objects, because, then, we have to have some coordinate system to refer to points on the sphere.

So what's the next most complex object to take as an example? Well, let's take ellipsoid since we know what those are since we talked about them in our discussion of critical surfaces. And we're not going to do the whole algebra, partly because we did the ellipse, and this just gets worse.

So here's our ellipsoid. And of course, we have the possibility of writing it in this form, as we saw, where the A, B, and C are the semi axes. And so we can get various shapes by changing the ratios of those axes. And of course, if A equals B equals C, then we have a sphere. And so this is an implicit equation for the surface.

It's not much use if you're trying to say, generate a visualization of that shape or if you're trying to say, integrate over its surface or something. So there are alternate ways of describing the same surface. And here's one of them.

And so that's a parametric description. And so I could generate points on the surface very easily just by sampling theta and phi and computing points on the surface. And these would correspond to latitude and longitude if it was a sphere. Phi and theta are a way of addressing points on that object.

OK, so I can always write a vector to any point on the surface by just listing those things. And what do I want? Well, a couple of things I need. One of them is surface normal, and the other one is curvature.

So let's start with the surface normal. So how can I get a surface normal? Well, the normal is perpendicular to any tangent. So if I had two tangents, I could just take their cross-products.

How do I get a tangent? Well, I just differentiate this with respect to the parameters. So I have two ways of doing that, and that gives me two tangents.

And let's see. So there's one. And there's another. And now I can take the cross-product to get a surface normal. And after some algebra, we get that expression.

This is not a unit normal, and we don't really care about the magnitude of this thing. We're going to normalize it anyway. So we can get rid of-- we can ignore that.

Then if we do that, then it actually looks similar to R itself. Just, we've replaced A with BC and B with AC and C with AB. So it's an interesting substitution. OK, so that gives us the surface normal. So that allows us-- for any point on the surface defined by theta and phi, we can compute the point on the unit sphere that corresponds to that.

And then we need curvature. And that's harder because we need to differentiate one more time. Now, we need to look at-- as you move on the surface, how fast does the surface normal move on the sphere? So anyway, it's in the paper and, I will not go through it. It's somewhat painful. I'll just give the result.

So I need to define a couple of other things. So on the unit sphere, the way of dramatizing that using latitude and longitude-- let's see. So which way do I want to do it?

So it's the same method we used over here to parameterize positions. OK, so this is a unit normal on the sphere. And obviously, I'm going to have to equate the normalized version of this thing with the terms of that. And then, in the process, it's convenient to have yet another vector, which is--

And what the significance of that is is not obvious. It's a little bit like when we're talking about coordinates on the sphere, which could be either geocentric or based on local surface normal. And but the answer now is-- and therefore, the thing we're interested in is--

So aside from computing the curvature, which involves seeing how fast the normal-- a major pain of this calculation is that we don't want the answer in theta phi coordinates. Those refer to points on the object. We want them in terms of coordinates on the unit sphere. And that's what we've done, because s is defined in terms of coordinates on the unit sphere. And so, if you give me the latitude and longitude, I just plug that in, and I get G.

OK, what does that look like? So that's some distribution on the sphere, and we'll use that for recognition and finding orientation. So what does it look like? Well, the first thing to notice is that there's some extreme-- no surprise. We'd expect that, where these semi axes hit the surface, those might be interesting places. And sure enough, we end up with maxima and minima. So those are the extreme values, and they occur at the places where the semi axes intersect the object.

And then, if I look at that on the sphere, that means I'm going to have three of these points. And one of them is a maximum, and one of them is a minimum, and they're 90 degrees apart. The surface normals here are pointing three orthogonal directions.

And let's see. What's the third one? Well, it can't be a maximum or minimum. It's a saddle point.

Now, when I go to the other side, I expect things to be symmetrical. So the curvature here is the same as the curvature here and so on. So the mirror images of these three points on the other side of the sphere-- so somewhere in the back here is another minimum. And down here, in the back, is another maximum, and here's another saddle point.

And there are some theorems that may seem intuitively obvious that tell you that you can't have maximum and minimum on the sphere without having a saddle point. So it's not too surprising that we have that. So that's the extended Gaussian image of an ellipsoid, and its maximum minimum and saddle points are lined up with its axes. We've chosen to line the axes up with x, y, and z-axis, but that'll be true in general.

If I rotate this object in space, what happens to the spherical image? It just rotates exactly the same way. And that's what we meant when we said we're not really looking for rotational invariance. We don't expect things to be constant when you rotate, but we want it to transform in some easily understandable way.

And I don't know-- people have some terminology for that. Rather than invariant, they say equivariants. There's no agreed terminology, as far as I'm concerned. But it's an important idea that we want the change to be easily understandable, unlike-- say if we've taken the perspective projection of the object. Then when the object rotates is a very complex transformation of the image, whereas here it's very straightforward.

And so we can now use this image with experimental data to both test whether we might be looking at an ellipsoid and to determine what its attitude in space is by taking the library version and trying to bring these two into alignment, and we'll talk about how to do that in a little while.

So the sphere is very simple. The ellipsoid is complicated because it has a full three-dimensional shape. Somewhere in between are things that are a little bit easier to handle. And for some purposes, those are of interest. There are certain objects that are in between in complexity.

And in particular, if we look at solids of revolution, we find that it's easier-- a lot easier-- to compute the EGI. And solids of revolution, of course, include cylinders, and cones, and spheres and hyperboloids of one sheet, hyperboloids of two sheets, assuming the parameters are chosen appropriately.

OK, so how to compute the EGI of a solid of revolution. So in the case of a revolution, there's a generator that we spin around some axes. And so let's suppose that this is a generator, and we're spinning it around this axis to produce some object.

And then we're going to map this object onto this sphere. So let's define a couple of things. So suppose we're here on the object. Then we're very interested in the radius, r.

There some other coordinates we might want to use. We might think of that as height. Then we might use the arc length along the generator. So we'll derive formulas for several cases because some are convenient in certain situations, more convenient than others.

OK, so surface normal. Here's our surface normal. And then angle with the equator. So the corresponding point on the sphere would be somewhere where, if you measure this angle at the equator, it would be the same angle. So that's the latitude on the sphere. OK, so that's how points correspond, but we need curvature. So how do we do that? Well, we go for that definition up there, and we look at some element that we can easily figure out the mapping of. So rather than just consider a point, let's consider this whole band. And let's say that the width of the band is delta s, the change in the arc length along the generator. And then the surface normal presumably will change a little bit as we go to the other edge of the band, so that band maps into this band.

So the nice feature of the solid of revolution is that it's symmetric both in the object and in the transform, in the EGI. And so it reduces it from 3D to 2D. Yeah.

So this really is the longitude on the Gaussian sphere. So in this direction, we have that angle, and then, in this direction, we have eta.

Well, it's the angle going around this way. And the great thing about the solid of revolution is that everything is constant in that direction. So we can cheat and forget about it.

So we need the area of this band. So the area of the object band is-- so it's 2 pi times r times the width of the band delta s. It may not be obvious. But we could take this ban, cut it, and lay it out in the plane and measure it, and this is what you'd get.

And correspondingly, then over here, we've got 2 pi. And what's r here? Well, it's depending on the latitude, the higher I go, the smaller r is. And in fact, it's the cosine of eta. So if I project this down here-- right angle-- and this is 1, then this is the cosine of eta.

So that's the radius times 2 pi. And then I still need to multiply by that. So that gives me K is-- so 2 pi cancels, and I'm left with that.

Now, actually, I'm mostly interested in 1 over K. So let me flip this over. Yeah.

So what are these things? So this is the rate of change of direction of the surface normal as I move along the arc. So that's a curvature. That's a rate of turning as I move along the arc. So that's actually the 2D curvature. It's the curvature of the generator.

So that's interesting. That means that I've been able to reduce things to-- well, let me stick with this form-- from 3D to 2D. So I've got K is cosine eta over r times KG. And so if I can get the curvature of the generator, then I'm done.

The important thing to see is that we used the same idea all along, which is that the Gaussian curvature is the ratio of those two areas. We just need to find corresponding patches and measure the areas.

Now, that's one way of expressing it if we know KG. But as I mentioned, this curve may be given in various forms. We may give it in an implicit form, or we may give it as r as a function of s or r as a function of height z. So it's convenient to have different versions of that formula.

This one is ds, so the object and the sphere. So if we blow up that place where the narrow band hits the solid of revolution, here we got delta s. That's the step in the arc length. And here's a delta z, the change in height. And here's a change in radius, which, for positive eta, is negative. If eta is a positive quantity, then the curve is coming in towards the origin. So the change in radius is negative.

And then where is eta? Well, it's here, because the surface normal is there, and this is eta. And so that better be eta. And then I can read off trigonometric terms from that diagram. For example, I can get sine ADA is minus the minus dr over ds, which is minus rs. The subscript now denotes differentiation. So if I have r as a function of s, I can use this method and just differentiate with respect to this.

But what I need in the formula is cosine of eta. So one thing I can do is differentiate with respect to s. And of course, the sine becomes a cosine.

And there we have the second derivative, which shouldn't be a surprise because we expect curvature to have to do with a second derivative. And so here we have a very convenient formula for the curvature of a solid of revolution if we're given r as a function of s.

And so right away, we can do an example. In the case of the sphere, we can write that this small radius, little r here, which is the same as this thing, is the big R times cosine of that angle. And that angle, of course, is just the arc length divided by the radius, our usual formula for and defining angles and radians.

OK, so that's r. The other thing I need is our differentiated twice with respect to s And of course, if you differentiate cosine twice, you get minus cosine, and then there's an r squared. So we should get 1 over r minus sine.

And so then, when we put it together, hopefully we get that. So we already know that. But it's a good way to check the result.

So that's if we're given the generator as r as a function of arc length. Well, it's one of the 12 most common ways of specifying a curve, but it's not in the top three. So let's go a little bit further.

So the other thing we can look at is z. It's more likely we're given r as a function of z, because, if we turn this sideways, that would be the normal way you'd specify y as a function of x when you're defining a curve.

So let's see. That one also comes out pretty simply. So if we look at that diagram over here again, we can relate z and s to trigonometric term in eta, so we have 10 over eta.

So we can get 10 of eta from r given as a function of z, just by differentiating with respect to z And again, I need secant or cosine or something. So I can differentiate this with respect to z. And I'm going to get the secant squared eta d eta dz. Oh, ds, sorry. And then that's going to be d ds of minus rz. And that's minus-- the chain formula times dz ds.

And then, from the same diagram, I can read off-- so I'm reading off all of the possible trigonometric things. So cosine of eta is just dz over ds. Well, no. OK.

Putting it all together, I get rzz cosine to the fourth of eta. And so K ends up being-- this formula is slightly messier than the other one. I left out a couple of steps there, just to avoid the monotony.

One thing you'll need is secant squared eta is 1 plus tan squared. And so in our case, that's 1 plus rz squared. And when you put all that together, you get that formula. So that's a second way of getting the Gaussian curvature of a solid of revolution. If you generate a curve, it's given as r as a function of z instead of r as a function of s. And again, we can apply this to an example just to have a sanity check. So in the case of the sphere, we have r is the square root of r squared minus z squared, assuming that z starts at 0 at this point. And so rz is minus z of a square root of r squared minus z squared. But we need the second derivative. So we differentiate again, and we end up with--

And let's see. The other thing we need is 1 plus rz squared. And if you put it all together, we get the correct result. OK, so this gives several methods for generating extended Gaussian images of solids over evolution.

One reason we did this is because we're going to look at a particular solid of revolution and study its EGI and talk about how you would use it in alignment and recognition. And that's a donut.

So here's a cross-section. And basically, we take a circle as a generator, and we rotate it around this axis. So we generate a torus. And in terms of specifying how big this thing is, there are two things we need.

Let's call this rho. We need the small radius, and then we need the large radius. And those two define the shape.

So is this going to make sense? Can we compute an extended Gaussian image for that? So it's a solid of revolution, so we should be able to just use one of these formulas we derived over there. What might be a potential problem?

AUDIENCE: It's not convex.

BERTHOLD Right, it's not convex. Right. So, so far, we've focused on convex objects. In the case of a convex object, the
HORN: Gaussian image is convertible-- that is, you can go from the object to the sphere, and there's a unique place to go back to on the object because there's only one point that has that surface orientation. And we know that there are some powerful properties in that case, such-- will improve it. But it's unique.

That is, if you have a particular extended Gaussian image, there's only one convex object that corresponds to that. And I didn't mention, but Alexandroff proved that one. And again, it's an unconstructed proof, but it means that there's no confusion. There's only one.

So we're going to lose some of that. We'll see that we can take the extended Gaussian image of this. But some of the nice properties that we talked about won't apply here. So then-- well, we'll wait until we get there.

So right away, in terms of the issue of inverting the mapping, we see that, in this case, instead of there being a unique point on the object that has a certain surface orientation, there are two places. So that means that the mapping is not inevitable.

And those two places differ in an important way, which is that the object is convex here. If we move in the blackboard direction, you can see it's curving that way. And if we move out of the blackboard direction, it's curving in a similar way. And so that part is convex, whereas, over here, again, in the blackboard direction we have this convex shape, the circle. But then when we go out of the blackboard, it's curving this way.

So this is actually a saddle point, and so the curvature here will be negative. So we'll have to deal with that. And so that's, again, just a reflection of the fact that it's not a convex object. If it was a convex object, the curvature would be everywhere non-negative, and in most cases, just plain positive. OK, so, keeping all that in mind, let's just blindly apply our formulas for computing its extended Gaussian image. So what we're going to need is the radius, the little radius, r. So that's this quantity. And so little r is big R plus rho times cosine eta.

To apply the formula we need r as either a function of s or z. So let's do it in terms of s tho, because this is s, and so this angle is divided by rho. Cosine is divided by rho.

And then I take the second derivative, and, of course, the big odd drops out. And I'm going to get a negative sign because of the cosine turning into a minus cosine. And I have to divide by r squared. OK, so that's my second derivative.

And so, by the way, a number of things become apparent, that-- so that's curvature of the generator-- that, as I go from this orientation up to the top, that is going down all the time because cosine from 0 to pi over 2 goes down until it hits 0.

So something interesting happens up there. And then, if I go further, it's going to go negative. And that's the area we're talking about here, where the surface curvature is actually negative while it's positive here. So that divides the torus in to two parts in a way-- one that's inside this cylinder, where everything is negative curvature, and the other part, which is outside.

OK, so combining those using the formula we had over there for the case where r is given as a function of the arc length, this is what we get, and it's for this part.

So now, what happens on the other part? Well, we can do the same calculation there. And now r r minus rho cosine eta. And so r is r minus rho cosine s over rho. And so the second derivative is going to be-- and so we get a contribution. Minus, minus.

So the contribution there is a different sign and also a different magnitude. So what do we do? Well, we have two obvious choices. One of them is to just add them up. So if we have a non-convex object, one thing we might do is just compute the Gaussian curvature at all of the points that have the same surface orientation and then add those up, and it seems like a reasonable thing to do.

And actually, of course, we're interested in the inverse. We're interested G. So if we think of those as K plus and K minus, it turns out that that's nice because stuff cancels out, and we just get that.

So that's not very satisfactory, because it says that, if we define the extended Gaussian image this way, then we get a constant for donut, which is what you get for a sphere. So this is like a sphere of radius square root of 2 times rho. And it has no orientation. It's constant all the way around, so it's not going to help us in determining the attitude of an object. So no, we don't want to do that.

Another way to think about it is that, when we do this in practice, we simply project out the normals wherever they come from without taking into account local curvature. So we could, for example, carve this up into lots of little facets. And for each facet, we compute the area on the surface normal, and then we put a mass equal to the area on the unit sphere at the corresponding place. And there's no account taken of whether the surface curls in or out. And so we don't want to do this. We actually want to do this so that the second term, which is negative, makes a positive contribution. So if we do that, we get this. And so that's going to be our EGI for the torus.

And let's look at that. So that's pretty interesting. So it's not constant. That's good. And actually, it has a singularity at the pole. So when we were dealing with that object up top right-hand corner, on the EGI, at the back, there's the mass concentration, which, in the limit, is an impulse. So that's certainly one form of singularity.

This also has a singularity, but it's not an impulse. It's just that, as you approach the pole, this keeps on going up, and it's infinite at the pole. And actually, if you want to know how it varies, it's pretty easy because-- imagine that we embed our unit sphere in a unit cylinder. Let's see. Cosine is opposite of adjacent, so secant is the inverse. And so the height of this thing gives us the secant.

We start at the equator with a non-zero value, but then it monotonically increases until, when eta approaches 90 degrees, we're off at infinity. So this is a way of visualizing how that varies. It's symmetric in this direction, which is appropriate for solid of revolution.

And you can now imagine that we could use this for alignment, because, if we have a model of this object and we have data from machine vision, then we have these two spheres with the distribution, which is non-zero everywhere, veering smoothly, but it has this rapid growth towards the poles. So we can bring them together to line up those singularities, or, in practice just very, very large values.

Notice that that doesn't give us, completely, the attitude because we can still spin this around this axis without anything changing. And that's also appropriate because we're dealing with a solid of revolution. So of course, it's ambiguous, what angle it's turning around its axis. Yeah.

We briefly mentioned this last time, and somebody sent me email about how to trace that down. So first of all, the proofs are non-constructive. Both Minkowski's proof for the discrete-- the polyhedral case and Alexandroff's for the smooth case, the smoothly curved surface.

The proof does not include reconstruct from this one and that one, and they're the same. That's not the way the proof works. And so people have tried to find some way of iteratively reconstructing. And for the discrete case, there are two very slow solutions.

So imagine that we have a polyhedran and we know the orientation of each of its faces and we know the area. What do we do? Well, one thing we can do is we can construct a plane. So if we know a normal, we can construct the plane. And a piece of that plane is going to be part of that object.

And then we can move that plane in and out. And as it intersects other planes, its area will change. And so we can set up some big search or optimization problem where the knobs we can turn are the distances of all of the planes from the origin. And the objective is to match as close as possible the areas that somebody told us that each of the facets is.

And it's a nasty problem because computing those intersections is a lot of work, and some of those faces may not exist. You may have pushed it so far out that it no longer intersects with the rest of the object. And in many cases, if you think about it, you have some complicated object with many faces, and then you push a face outward. It's likely its area is going to decrease because it's moving out. So here's the thing we've reconstructed. Here's the surface that we're playing with. If we move it out, its area will increase. Well, unfortunately, for some objects, that's not the case. It actually goes the other way around. So if you have some iterative negative feedback scheme, it'll suddenly be positive feedback and blow up.

So it's been done, and Katsuya Akiyoshi started that game, and I guess Jim Little is another reference. So short answer, no. But you can approximate it if you're willing to do a lot of computing.

So now, what's important for us, though, is that we don't need that because what we're doing is we're working entirely in this space. We don't go back to the object space. So we're doing the recognition by comparing the distributions on the sphere, and we're doing the alignment by trying to rotate one sphere relative to the other until we get a good match.

It's intellectually very intriguing. Why can't you just compute the object? But it's actually something, fortunately, we don't need to do-- fortunately, since it's not easily done.

It's interesting distribution on the sphere. And there's another way of understanding it that's quite useful, which is the same argument we made about bands on the surface and bands on the sphere, except, this time, in a different direction.

So here's a donut. And now imagine that we divide it up into bands. We're cutting it based on an axis, and we have a plane that goes through the axis. And we rotate it slightly to generate this band, and then you look at the sphere, and we look at the corresponding band. Now this goes through a full rotation, so we'll have a full rotation on the sphere.

But that same plane I was talking about that we used to slice this when we sliced the sphere-- we get a crescent shape, like a slice of lemon, not this bad. Now, this band is actually not as wide on this side as it is on that side. But if the radius of the donut is large enough, the difference will be small, whereas, in this one, it's obviously a function of latitude, and the width of it goes as the cosine of eta. So the area on the sphere goes as cosine of eta.

And so the curvature goes as cosine of eta, and so G goes secant of eta. You see that? So we're really, as usual, taking the ratio of the area here to the area there. And what's happened is that this has gotten squashed near the poles, and so we have that cosine of eta dependence.

Now, you might say, but this band is in constant width. It's slightly narrower on the other side. Well, combine it with the other band that has the same orientations. So when we cut it with a plane, in this case, the object actually gets cut in two places. And so we need to add contributions from both of them, and this one happens to be asymmetrical in exactly the opposite way as that one is. So the end result is-- but between the two of them they're perfectly even all the way around if you add them up, whereas this one is perfectly cosine of eta.

So what's next? What's the area of a torus? Anyone have that off the top of their head?

OK, it's that. 4 pi squared. And why is that?

Well, one way we can think about it is that we take a circle of circumference 2 pi rho. So that's the circumference of the circle. Then we spin it with an axis around that circle, which has a length of 2 pi r. And the product gives us the area. So we have that generator that we're sweeping along an axis, and that's a hand-waving argument, but it can be made rigorous. In any case, the area of the torus is that.

Now, if we look at the equation, the formula here for the Gaussian image, you see that r and rho don't appear separately. Only the product appears.

And so what does that mean? Well, it means that two donuts of different shapes but the same area have the same EGI. So that's the price we pay for going to allowing non-convex objects. We lost uniqueness.

And so before, there was only one convex object that corresponded to a valid EGI. Now we have a bunch of them. And so if, for example, if you had a bicycle tire with a big R and a small rho, you might have the same EGI as a scooter tire with a large rho and the small big R, if the product happens to be the same.

So that's a shortcoming of this representation, and it may or may not matter in an application. If you're dealing with a car repair shop and you have trucks and scooters coming in, then this might be an issue. You might need to use some other method to distinguish them.

If you're in the world where these donuts are in with other objects, like spheres and cubes and bricks and tabletops and so on And this is the only torus you're going to run into, then it doesn't make any difference. But it shows that, when we extend this to non-convex objects, things aren't quite as nice.

And there are other issues with this. For example, if we image this donut-- if we image a convex object from opposite sides, we get the whole surface. There's nothing hidden. Here, there are little pieces that are missing because they're hidden by--

So normally, you will see all surface elements where the surface normal is not more than 90 degrees away from your viewing direction. If it's more than 90 degrees from your viewing direction, it's in the back. It's self-shadowed. You won't see it anyway.

But here, the small parts of the surface hidden back in here, where the surface normal is pointing towards you, so they should be counted in constructing the EGI-- and they are in the mathematical version. But if I take this from, say, photometric stereo data, there will be some small error introduced because I'm missing part of the surface. And again, that's because the object is not convex.

OK, well, let's talk about how we would do this using numerical data rather than nice, mathematically-defined shapes. So this is a little bit like in our discussion of patents, where the ultimate application is where you have a real image-- training image, and you fit edges and do all that stuff.

But there's some utility to being able to also deal with CAD data, where you have an analytic description of an object and now you don't need a perfectly made one because you've got the perfect thing there in the CAD. So similarly here, in practice, we'll be looking at real objects which are imperfect. But if it's possible to put something in the library based on the true shape that this thing is supposed to be, that's valuable.

Hard to do it numerically. And for example, if we have photometric stereo data or if we have a mesh model as people use in graphics. So then we have patches on the surface. And in the case of photometric stereo, those patches will typically correspond to pixels. So there's a small, quadrilateral patch on the surface that maps into a pixel, and we know its orientation. But whatever it is we have the same job. We know that facet by, say, its corners. Let's say it's triangular. And we need two things. The one is the surface normal. And then the other one is the area.

So let's start with the surface normal. Well, that's easy to get because we can take any two edges and take the cross-product-- for example, that. And that looks asymmetrical because we have b appearing twice. Why should b be appearing twice and the other's only once? But we can easily show that it's actually that, which is symmetrical. OK, so it's easy to compute the surface normal.

And then the other thing is we need the area. And so there, the area of the triangle is 1/2-- where does that come from. Well, the dot product is proportional to the length of the two vectors times the cosine of the angle in between them. And that's exactly what we need for the area.

So if we have a parallelogram, the area of that is the product of those two vectors. And we don't have a parallelogram. We only have half of a parallelogram. So we get that. And again, this is asymmetrical, so that's not-- that seems odd.

And of course, you can do this three ways. You can either have two copies of A in the formula or two copies of B and so on. And if you add them all up just for fun-- there, that's symmetrical. How will I fit this?

Oh, no, there are, because we added three of them, and each of them is a half.

So anyway, easy to compute. So this is normal and the area. And now what do we do? Well, we put a mass on the sphere at the point. Based on the surface normal, the mass will be proportional to this area.

And then we repeat this for the other facets of the object, and that way we get a mass distribution on the sphere. And the density of that is our G. So this is another way of understanding why, when we add the two contributions from the two sides of the donut, we want to add them the way we do. We don't want to subtract stuff and have it cancel out, because, here, we're not taking into account anything about curvature directly.

It's interesting that that's the effect we get, that if the curvature is high, these guys will be all spread out, but it doesn't matter whether the curvature is positive or negative. They'll be spread out.

So how to represent this? So basically, what we're building is the direction histogram. So you can imagine that we would somehow divide the sphere up into boxes, just as we do when we're computing histograms. And then we just count everything that falls into each of those cells.

And direction histograms are used in other contexts. They're pretty interesting. So for example, if you look at the fine structure of, say, muscle you'll find that most of the fibers are parallel. And so how do you express that--which direction are they going, and how many are not in that category?

Well, you plot them on the sphere and build this orientation histogram, and then it'll become apparent that there's a strong concentration, a particular point on the sphere that corresponds to the longitudinal axis of these fibers and not so much elsewhere, but there will be some. And this has an application in, for example, neuroimaging. So as with MRI, you can find out the flow directions of water in your brain and thereby determine the dominant axon directions, and then you can plot these connecting cables that go from one part to another. And you can study them by this method of plotting these directions histograms.

They do the same with blood vessels. So one method for trying to distinguish tumor from other tissue is to make note of the fact that the tumor needs blood supply to grow. So it puts out stuff that attracts blood vessels.

But it's disorganized It's not built the way it is when you're growing from a small cell to many cells. And so it's a mess. It's leaky. But more importantly, from our point of view, blood vessels go every which way.

So if you image a tumor and you plot the orientation histogram, you'll get a uniform distribution around the sphere. That's bad. If you image real tissue, you'll find that, yes, there are vessels going every which way, but there's typically a dominant direction or multiple dominant directions. And when you plot the directions on an orientation histogram, they'll be strong blobs, whereas, in this disorganized tissue, it'll all be spread around.

So orientation histograms aren't really a new thing. Most people don't know about them. But they're used in other areas-- cryomicroscopy, and what have you.

Now, ordinary histograms are pretty straightforward. In 1D, you just divide things up, and you just count how many things end up in each slot. And then maybe, based on that, you create some estimate of a probability distribution. It's slightly more difficult in 2D. Well, OK, so we just divided up into cells. And same thing-- we count.

Squares? Not so good. Well, one reason is that they aren't round if we could somehow fill the plane with disks, it would be better, but we can't without overlap or leaving gaps.

Why is this bad? Well, take a more extreme case. Suppose your tessellation is triangles. We could certainly use that way of dividing up the plane as well. But you see that, in the case of a triangle, you're always combining slightly different things. But in the case of a triangle, you're combining things that are pretty far away from the center compared-- for the same area triangle-- than if you had a square, or, even better, if you had a more rounded shape like a hexagon.

So in the case of the hexagon, the ratio of the largest radius to the smallest radius is very small. In the case of this triangle, it's quite large. And the square is intermediate.

So in the case of 2D, depending on your application-- this is what people generally just use because it's trivial. It separates the problem of dealing with x from dealing with y This you don't want to use, and this would be better. But it's extra work, so people typically don't do it. And the improvement is not huge. It's not like it's twice as good as square. So that's one issue, how to pick the cells.

Then there's another one, which even occurs up here, which is-- suppose that something falls there, and with a little bit of noise, it would have fallen there. And so when you look at the histogram, you have to take that into account, that there's some sort of randomness going on here and that when you compare two histograms, you want to be careful to take that into account.

And then how do you do that? Well, one way is to have a second histogram that's shifted. And so, in that one, these fall into the same cell. Problem solved. Except now you have to compare a shifted and an unshifted histogram.

Well, in 2D, of course, you can do the same thing. But you have to be careful. You end up having to do it four times. And so as you go up in dimensions, this, quote, "solution" gets more and more expensive. You have to shift it to half an x, a half a y, and then shift both x and y. And then together with the original grid, you've got four grids. So that's a common solution for the 2D binning problem.

There's another way, which is to say, well, when I deposit my contribution here, I put some of it over there and some of it over there. Basically, you're convolving your distribution with some spread function, and, depending on the implementation, this may be cheaper to do than that. In the case of 2D, you'll have to put it into four places. And again, this is like doing it at the time you enter the data versus doing it at the time you read it out. So there are those issues.

But we actually have a worse problem. We have a sphere. And so how do we divide up the sphere? And we already talked about longitude, latitude not being a great way to divide it up.

So let's, before we do it, summarize what of the desired properties we want of a tessellation. And in the planar case, people don't even think about it. It's just obvious. But when we have a curved surface, it's more complicated.

We would like the cells to all have the same area. And again, here, it's trivial to arrange for that, hard to do on the sphere. Then we might want them to be equal shapes. And again, these are all the same shape. No problem.

And then we might want the shapes to be rounded. And that refers to the discussion we had about triangular grids and hexagonal grids. And again, on the sphere, it's very easy to build triangular grids, but they're not particularly good. We want something that's more rounded, like hexagons, pentagons, dodecagons, and so on.

OK, what else do we want? Equal area, equal shapes. We'd like to have a regular pattern. We want it to be easy to do the binning.

So over here, how do we do the binning? Well, we just do an integer division and throw away the fraction. Or we round off to an integer-- both in x and y, If necessary. But that's not so obvious to do here, particularly if we have some interesting pattern with lots of hexagons and pentagons. So if I have a unit vector, where does it go? Now, over here, you even think about it. It's so trivial. You just divide by the interval, and there's an integer part, and that's what you use.

If you have a bunch of facets on the sphere and you have a unit normal vector, what do you do? Well, you can do something brute force. You can just take the dot product of that unit vector with the unit vector of each of the cells, and then you pick the cell that has the largest value because it's cosine of the angle and the largest value means the smallest angle. But obviously, that's not practical because that means that, every time you access the orientation histogram, you need to step through all of the cells.

OK, easy to bin. Let's see, one, two, three, four, five. I thought I had eight. Let's see what else we need. We want to have alignment on rotation.

So what's that all about? Well, in doing the matching of these edges, we will need to bring one object into alignment with the other. And again, in the planar case, it's just translation. You just shift things around. It's very straightforward.

And in particular, you can shift it around by discrete increments equal to the size of the cell and, each time, test in that full match. And there's no loss of quality because you just take the numbers as they are.

Well, that's not going to be so easy in the case of rotation. So let's think about how that might work. So suppose that we have divided up the sphere into-- I don't know. Let's make it a dodecahedron. Dodecahedron.

OK, so here's our sphere. We've taken the dodecahedron and centrally projected it out onto the surface. So we have these pentagons. Well, they have curved edges, but they're the result of projecting a pentagon up.

And so this is one of the cells. And there's another one. And there are 12 of them, right? Dodecahedron.

And so what's my data representation? Well, it's just 12 numbers. Now, of course, this is not really a good example because that's too few, but just to illustrate the point. So my orientation histogram is 12 numbers.

Now, if I rotate the sphere so as to bring the facets of the dodecahedron back into alignment with itself, what happens to these numbers? Well, a facet goes to some other facet. So maybe A1 goes over here and A7 comes back here, and A9, A3.

So all that happens is that they're permuted and there's no loss in quality, and it's easy to compute. So that's what happens over here. If I shift this whole thing, the entries in the data are permuted, but I don't even worry about that, because I just have an array, and I just imagine that the array starts somewhere else. So that's the advantage of alignment on rotation. So this means that, here, for any rotation in the group of rotations of the pentagon, my data changes in a very systematic way that does not involve any loss of quality.

What am I talking about-- loss of quality? Well, suppose that, after rotation, these cells didn't line up but overlapped in some way. Suppose, after my rotation, It looks like that. Well, that means then I'm going to have to redistribute whatever weight was in here, a little bit in there, and this red cell would pick up a bit of the weight from there, a bit of the weight from there.

So I'd be doing some interpolation convolution operation, and the result could be useful. But then maybe I'd have to do it again. And then I'm going to be in the Xerox of the Xerox of the Xerox problem, where each step, I lose a little bit of quality, and, after a while, it's not really useful anymore.

So let's see. Regular pattern, regular shape.

So then the question is, what patterns can we use? And so the reason we talked about platonic and Archimedean solids is because those are the starting points for these orientation histograms. And unfortunately, we've run out of time. So we'll talk about that next time.

So there is a quiz. And I think we've covered everything that you need to do the quiz. Otherwise, we'll finish on next Tuesday.