[SQUEAKING] [RUSTLING] [CLICKING]

BERTHOLD We're starting to talk about what determines brightness in an image and how we can exploit that. And weHORN: introduced the idea of a gradient space.

Why? Well, because brightness is going to depend on the illumination, obviously, and it's going to depend on the geometry of the situation, including surface orientation. Obviously, the amount of light falling on the surface will, per unit area, will depend on its orientation. And then different types of surfaces will reflect that light in different ways.

In any case, we expect that the brightness we observe in an image is going to depend on the surface orientation of the corresponding patch on an object. And so we need to talk about the orientation of the patch. And we had unit normals, and then we had p and q, but these are just convenient shorthands for those slopes in the image, slopes on the surface. So those are derivatives of height rather than derivatives of brightness.

And then since we were busy with photometric stereo and we just talked a bit about Lambertian surfaces, which have a property that their brightness depends on the cosine of the incident angle and it does not depend on the viewing direction. So a surface of that type will appear pretty much equally bright from all viewing points that you might have, which is fairly common for material on our human scale.

So what determines how surfaces reflect light? Well, we're getting into that, but largely it's microstructure. You know, photons get into the fibers in my paper, they bounce around, they come back out again in a different direction. And that's what determines how bright the surface will appear from a certain direction. And so a lot depends on the imaging situation.

If I'm looking at the moon, what constitute microstructure are craters, not fibers of paper. So as we'll discuss later, Lambertian surfaces or near Lambertian surfaces are fairly common in our world. A lot of matte surfaces are pretty good approximations, snow and whatever. But they don't necessarily apply when we go to microscopic scale or to a cosmic scale.

Anyway, so Lambertian is a handy approximation for some surfaces. And we can address it in different ways. And if we expand out that end on s, you may remember we got something like this, which is linear in p and q. So that's the good part.

But then unfortunately, we divide by this term here, which is not linear in p and q. And there's also this term, although we don't worry too much about that because that's a constant. If the light source is somewhere as defined by ps qs, then that's a constant. And I'm going to get tired of writing that, so I'm just going to introduce a constant rs to represent it. And then we went through the business of what are the isophotes. Well, that's where this expression is constant. And we can get rid of the square root by squaring. And if we do that, we end up with an expression, which is second order in p and q.

So it's got p's, p squared, q squared, pq, p and q, and constant. And plotting those gives us conic sections. And in particular, we end up with a diagram like this.

And what is this? Well, this says if you tell me what the surface orientation is, I can tell you how bright it's supposed to look. So again, imagine a terrain built above the floor. And I don't know-- x-axis to the right and y-axis forward. Then I can take the derivative in the x direction. That's p. And I can take the derivative in the y direction. That's q.

And those define the orientation locally of that surface. Of course, it might be a different p and q somewhere else. And that defines a point in this plane. And I can go to that point and say, what is the value of this function there? And that will be the brightness. And this is obviously handy in graphics because we can have surface models.

We can determine not just the z position, the depth, but also these derivatives-- surface orientation. And then we go to this diagram which could be perhaps a lookup table in the computer. And we just look up the appropriate gray level to paint at that point in the image. That's the forward problem. And we are dealing with the inverse problem.

So our problem is, OK, I have measured the brightness of E equals E1. What can you tell me about the surface orientation? Well, it's confined to that curve. It's unfortunately not going to tell me uniquely what the surface orientation is. But it's restricted now to a great deal.

And so I need additional information. And there are various ways of getting additional information. One is to say, well, most objects in the world aren't haphazard collections of blobs in three space, but they hang together. Objects are solid. They have surfaces. Neighboring points tend to often have similar properties. Different parts of this table have pretty much the same surface orientation until you get to the edge.

But that's a hard constraint to implement. We'll get to that later. A much easier idea is, well, if I illuminate the surface differently, I'll get a different map. And I'll get a constraint on that map. So for example, suppose I move a light source from there to over here, and then I draw the same diagram. And now I measure the brightness in that second image under different illumination.

And let's suppose it comes out to be, I don't know, this one, here. Well, then I know that the surface orientation is where those two curves intersect. So that's photometric stereo done in the graphical way. We've done it before in an algebraic way.

Of course, that algebraic way only worked if we had Lambertian surfaces. This is going to work for any surface as long as you can draw this diagram, which is called a reflectance map. So the reflectance map basically is a diagram that shows you for every orientation how bright that surface will look for that orientation.

Now we've talked a lot about Lambertian surfaces, and that's mostly because they allow us to write down the equations and solve them. But real surfaces aren't perfectly Lambertian. And some are dramatically non-Lambertian. So what do we do in that case?

Well, it's pretty easy to imagine we can create a diagram like this and use that. And we can perhaps think about building a lookup table. Now this is going the wrong way. Here, if you've got p and q, you can look up what the brightness should be. We need a table that goes in the opposite direction. So let's think about how to do that.

So one way is to-- let's just say a normal inversion. One way is to take a surface element and look at it under different lighting conditions and record what you see and then repeat that for different orientations. And you can do that. It's going to get pretty tedious because you have to basically explore this whole space.

And every cell in this space in your lookup table is going to require that you reorient that piece. Of course, you can automate that and have some robot do the calibration for you. But it's an alternative that's often easier is to use a calibration object of known shape. And what better than a sphere, partly because it's very easy to make a sphere? And what do we do?

Well, we take an image. First, we'll have it lit up from all sides. And the sphere, of course, we'll image as a circle. And presumably the brightness inside the circle is going to be much larger than outside. So we should be able to distinguish the two and fit a circle to this.

So we take this image, and then we fit a circle. And what does that mean? That means we find estimates for the center and the radius such that all points within that are bright, and all points out are not. And you can do some subtle, clever things to make that accurate to subpixels and so on, but we won't talk about that now.

So what good is that? Well, for a sphere, we have a very convenient relationship. If we draw a point from the center of the sphere to the surface, and then we, at that point on the surface, draw a unit vector, guess what? Those are parallel. That's a unique property of the sphere. So it makes it really easy to know what the surface orientation is because we just connect the center of the sphere to the point on the surface you're interested in.

And of course, this doesn't quite apply to the Earth, because it's not a sphere. So need to modify that slightly. And that's why there are several different definitions of latitude, depending on whether you're talking about the local surface normal or the vector from the center of the Earth. And usually we take the local surface normal as the definition for the angle that's latitude.

So now what? Well, we could be a little bit more precise here. So we have-- for every point in this image now-let's call this point x, y-- we can calculate what the surface orientation is. And how do we do this? Well, let's start with a cross-section. Here, we're looking down on the sphere.

Now we're looking sideways across it. And I don't know. The camera is way up there. So the surface normal is parallel to this. It's just-- well, this is the point x0, y0, z0, and this is the point x, y, z. And so that's the surface normal. So all we really need is a formula for z minus z0. We're measuring x and y in the image plane, so we know those.

We can't measure the depth, so we have to calculate that. And of course, it's just going to be r0 squared. Where's that come from? Well, it just comes from the formula for a sphere, which is that the radius vector has a fixed size. x minus x0 squared plus y minus y0 squared plus z minus z0 squared is r squared.

So having that, we can then compute p and q. That comes from the formula we had last time. Last time, we showed how to compute n from p and q. And we also said you can go the other way. If you're given a p and q, you can compute n. And so that's where that comes from. And you'll notice that-- oh, sorry. I think this is plus. Let's see. If we go to the right of the center, the vector tilts to the right, so it's plus. Sorry.

So for every pixel, we can-- we know the orientation of the surface, p and q. And then we take a picture. And we get a map under different lighting conditions. So under one lighting condition, we get one-- E1. Under another one, we get E2.

And let's suppose we take three. So we develop this numerical mapping from surface orientation to brightness in three pictures taken under different lighting conditions. So we're treating pixels completely independently. Just think of a single point that we're imaging and a single pixel that's imaging because we repeat that at all the others.

So now that's going the wrong way. We really want to go the other way. We want to use this information so that when we then later take three images of an object under the three lighting conditions, we can go to some sort of table or some calculation that gives us p and q. So we want to go that way.

And in terms of the implementation, we left Lambertian behind. Lambertian's nice and analytic. You can invert it and do all sorts of stuff. But once we've decided that we want to be completely general, then we can't depend on analytically inverting things. So let's use a numerical table. And so this is a three-dimensional array in the computer.

And each little voxel here is one entry. And what is in that box? Well, each of these little boxes has a p and q. So what I would do is I would make the measurement in three image situations-- use them-- quantize them to the discrete intervals of this lookup table. Go to that place, and it tells me what the surface orientation is. So that's what I'd like. But how do I build that?

So is that clear? I mean, you can see how computationally trivial this is. It costs you nothing just about to interpret these images. There's no complicated iteration or anything. It's just a table look-up. It's hard to imagine anything cheaper than that. So if I can build this table, I'll have a very efficient method for getting the shape. Well, not quite. We'll get local surface orientation. We still have to talk about how to patch that together into shape. But for the moment, we're just going to get surface orientation.

Now what I'm doing then is I'm running over this calibration image. And at every pixel, I'm looking-- I'm computing p and q. And I'm measuring E1, E2, and E3. And I use that to put something into this table. So I can quantize my E1, E2, E3. That gives me an index into this 3D array. And I write the p and q there.

And I do this for every pixel. And there might be some overlap because I can't represent the table with infinite precision. I'll have to make some sort of compromise where I say, OK, I'm going to quantize-- let's suppose I quantize 200 different values. That means the lookup table will be what? A million entries.

It's a lot in terms of cache size and so on. But it's a reasonable value, whereas, let's say, a petabyte lookup table probably wouldn't be very satisfactory. But in any case, that means we need to quantize. And we might need to quantize fairly costly, like 1 and 100. And that means that some of these points may produce the same-- they happen to produce the same rounded-off values. And so they'll be writing on top of each other. So that's one issue, and, well, how do you deal with that? Well, one way is to average them, because presumably, they're slightly different p and q's. And by combining them in some weighted fashion, we can get higher accuracy.

But a more serious problem is there could be cells here that never get touched, that never get filled in. And there's a couple of reasons for that. One is just the nature of these numerical quantization effects. But there's a more basic one which is p and q. Gradient space is a two-dimensional space.

And we're mapping it here into a three-dimensional space. So we're not actually filling that space at all. So what do we get? Well, we get a surface in that space. So if we ignore the quantization for the moment, we're not filling that space. We're getting some surface. And if you like, you can address points on that surface using p and q. So that's point number one.

So that means that later on, if I find some combination of E1, E2, and E3, it's quite possible that it won't happen to be on that surface. Then that means the lookup table is not giving me an answer. So what's with that? Well, if you remember, when we went from two images to three, we introduced the albedo.

So that was one case where making a problem more complicated made it easier. Before, we only had two unknowns, p and q. And we ended up with a quadratic, so we had two solutions. And we said, ambiguity. Let's add a third unknown. And then suddenly, we get linear equations, unique answer. Well, this would apply here as well.

So our calibration object, of course, is made with one particular albedo. Perhaps we painted it white, and so its albedo is basically 1. But suppose we also want to deal with similar objects where not all the light is reflected? And so how do we do that? Well, it's very easy, because rho albedo linearly scales E1, E2, E3. And so that means that anything on array out here is connected.

And rho is 0 on this end and 1 on that end. So suppose we've painted our sphere nice and white, and we made these measurements. Then we've defined that shape, that surface in the space. And now that's for rho equal to 1. Because it's perfectly linear for other rhos, we can just generate those. So when we place an entry in the lookup table up here, we can just say, OK, now we scan along this line.

And we fill in all of those cubes. And now instead of just writing a p and q, we write p, q, and rho in there. So it's a 3 to 3 lookup table now-- three dimensions to three dimensions. And in a real situation, maybe you know that the object is supposed to be a certain color.

And so having a value, different value of rho, is not really acceptable. It indicates there's something wrong. And that can be pretty useful, like there's a smudge on the surface, or it's not actually the surface you were told it was. So it's like a check-- an error condition that you can check for.

Or, for example, suppose that there's something blocking one of the three light sources, casting a shadow. Well, that means that one of these E1, E2, E3s is going to be relatively small. And so you'll be away from the surface. And if you're using this method, you can pick that up and say, oh, I don't know what p and q is here, but there's something weird going on. Or I don't know if you remember the slides, but one of the other things that can happen is that if highly reflective surfaces are close to each other, there will be interreflection. And we will have brightnesses that are abnormally high. And we'll actually be outside the surface. And again, we can say, OK, my method doesn't tell me what the surface orientation is.

But there's something going on here. And so I'm not going to make that part of the surface. I'm going to use that to break up the image into parts that hang together and parts that don't. So in the case of those overlapping donuts, this is a way to segment it into-- you start off, the image is one thing. It's just e of x and y.

But you know that it's an image of multiple objects. And so segmentation is a big problem. And if we use these methods, we can segment on cast shadows where one donut casts a shadow on another. And we can segment on areas of high interreflection which is where they touch. And so what looks like a drawback can actually be helpful.

Now this still doesn't guarantee that we'll have filled in all of the voxels in this three-dimensional space. So actually, there's a customary way of solving this kind of problem, which is to go the other direction. So what we did was we went to the image. And we said, OK, for that pixel we get p and q. And we measure E1, E2, 3, then we put that in this table.

The other way is to systematically step through the table rather than step through the pixels, and then for each of these voxels, find the corresponding thing over there. And that way, you can be assured that you filled in the table. This way, it's you're projecting in a nonlinear way from a lattice, a cubic lattice, onto this curvilinear space. And there's no guarantee that you won't have overlaps, and yet you will fill everything in. But that's fairly boring, so I won't talk about that.

Now I was going to talk about Lambert's lunch paper and explain exactly why when you cannot see the fatty spot, the illumination from the left and the right is equal. But I'm hoping you'll believe me. It's just a page of fairly boring algebra. So maybe we'll turn it into a homework problem or something. Anyway, that was Lambert's instrument for discovering a lot of things about photometry.

So let's get a little more serious about photometry. We use these terms like "brightness" and "intensity" rather loosely. And it's good to make them precise. So the first term is "irradiance." And it's the most trivial concept you can imagine. Here's a patch of the surface of area delta A. And here's some light source.

And there's a bit of the power emitted by the light source, which is intercepted by the surface. And the irradiance is just a power per unit area. And in terms of-- it's watts per square meter, if you like. And just for reference, noonday sun in Washington, DC is supposedly 1 kilowatt per square meter.

And what does that mean? Well, it's not a very precise value, because it will depend on the state of the atmosphere and the time of year. And what are you measuring? Are you measuring only visible light? Are you measuring near infrared as well, and so on? But roughly speaking, that's a useful number to know, because from that, we can calculate things like-- suppose you have an image sensor that's 4 microns by 4 microns.

What energy falls on that? And then you put that image sensor behind the lens, which attenuates it even more. How much comes out? So it's a very simple concept. And unfortunately, it's not terribly useful for us, because we have an imaging system. We're not exposing the sensor directly to the illumination. Well, if we did, we wouldn't learn anything about the environment. Now you might say, well, what we're interested in is how much light comes off. So we could have perhaps a complementary idea where we have-- give a name to this quantity, where we take the power that's emitted divided by the area. Trouble is, that could be going anywhere. And there is terminology for this quantity, but since it's useless to us, we're not going to bother with it.

So, I mean, this is obviously useful if you're worrying about heat exchange, like how do I keep my satellite cool enough given that the sun is illuminating it on one side, and there's heat going off into black space in another direction, and so on? But it's not useful for us because we're not intercepting all of this. We're over here, intercepting a tiny part of that. And therefore, it matters that this radiation is not isotropic. It's not going in all directions equally. So we don't want that.

So then many textbooks will use this term "intensity." And we often talk about-- well, I try not to talk about "image intensity," because it's really wrong. Intensity has a technical-- it's a technical term that has a meaning. And what is it? Well, it's useful for a point source to talk about how much radiation is going off into a certain direction.

So here's a point source-- I don't know, a star, a light bulb. And we're measuring how much power is going in a certain direction. And, well, we need to normalize that, right? We need to take this cone of possible directions and somehow measure how big it is.

So we have to define this measure, which is called the solid angle, and the units of which are steradians. And in case you haven't come across it before, it's very simple. So in the plane in 2D, the preferred way of talking about angles is in terms of radians. And it's what you get by cutting that circle with that angle and looking at the length of the arc of the circle and dividing by the radius. So we all know how to do that.

Well, this is very similar in that we take this cone of possible directions in three space now instead of in two space, and we cut it with a circle. And we imagine we're at the center of a sphere. We're at the center of the sphere. We're cutting it, and we get a certain area. And now to normalize it, we need to divide by r squared, because this area will grow with radius squared rather than radius.

And so that's the definition of solid angle. So this allows us to talk about a set of direction. Doesn't have to be a right circular cone. Could be any shape. I could have something like this. All that matters is what this area is on the sphere.

So radians-- we go from 0 to 2 pi. So what's the corresponding thing for steradians? So if I want to talk about all possible directions around me, how many steradians is that? What's the surface area of a sphere? Yeah?

STUDENT: 4 pi r squared.

BERTHOLD 4 pi r squared. Thank you. And so it's 4 pi, right? Here, we go 0 to 2 pi. And this goes up to 4 pi. So the set of all possible directions around me-- so if I'm radiating energy, it could go into a solid angle of 4 pi. If I'm, for example, only worrying about light coming from the sky, that's a hemisphere.

So that's obviously 2 pi steradians. So there's just one more little subtle thing that's kind of handy, which is if that surface area is inclined relative to the direction to the center of the sphere, in that case-- so this is just another manifestation of that foreshortening phenomenon. And this is sometimes handy because we will have cases where there is an inclination. For example, now we're going to be talking about cameras. And the lens is going to be tilted relative to a subject that's off-center. And so we'll need to account for that. And that's obviously-- and why is that? Well, because this area at an angle is equivalent to an area that's at right angle to the axis.

And the ratio of these two lines to each other is cosine theta. And that also is the ratio of the two areas. So now we know how to calculate solid angles-- very handy concept.

We go back to the intensity. And intensity, I, is defined as the power per solid angle. And you can see that it's independent of distance. So if I go further out, it will cover a larger area. But we're assuming that there's no loss.

So the power going into that cone is the same the further out I go. And so that's a useful quantity. And it's a way of describing how a distant point source might act. And usually in machine vision, that's not the case we're dealing with. We're dealing mostly with continuous surfaces. And we could imagine breaking them up into an array of point sources, but that isn't usually done, and it's not very helpful.

So, intensity-- so if you use that word in your homework problems, you will flunk the course. No, just kidding. Since all the textbooks use intensities, unfortunately, we have to accept that that is an alternate term for something else, which we'll talk about next. So the true meaning of intensity is this. And it's never something that people in machine vision consider, unless they're doing things like trying to reconstruct the center of our galaxy and finding what that black hole looks like.

So what is it that we want? Well, we kind of had it here. We don't want to know the whole power coming off the surface. We're only interested in what reaches the observer or the camera. And so we introduced this idea of radiance, which is power per unit area per unit solid angle. So we have a little part of the surface here with an area A. And we have a person or a camera over here.

And this is a solid angle delta omega. And there's some power delta p going that way. And that's obviously much more what we want because that's what we're measuring. In the image, this power will be projected onto a certain area. And our sensors are measuring power, basically, which, by the way, is unfortunate.

It would be great if they measured the electric field. But they're instead measuring the absolute value squared of the electric field. And so amongst other things, we know that quantity can never be negative. And so we can't do any phase imaging with our usual cameras. There would be some real advantages to that, but we can't do that. So that's radiance. And it's power per unit area per unit solid angle. Or in terms of units, it's watts per square meter per solid angle-- steradian.

Now if you're a mechanical engineer, and you do dimensional analysis, that's kind of a nuisance, because steradians, like radians, doesn't have units, like meters per kilogram per second or something. They're just ratios. And so the clever trick mechanical engineers use to guess at the answer by just matching dimensions-- it doesn't quite work, because you have these quantities that are dimensionless.

So what to do next? Well, what I want to do next is to relate brightness out there to brightness in the camera. And so I'm using this term "brightness" in a very loose way. And we'll justify in a little while why that's acceptable. So in a way, we can talk about a brightness of a surface in terms of radiance. This is how bright it's going to appear. Then in the image plane, we can talk about the brightness we measure as irradiance. So notice that before, I talked about light falling on an object in the world. But the same concept, of course, applies in the image plane. I've got these little areas that are light-sensitive. And what are they measuring? They're measuring irradiance.

Well, they're actually measuring energy of a certain time. But of course, energy of a certain time is the same as power, so they are measuring power. So I want to relate those two. And what we're going to end up showing is that they're proportional to each other. And therefore, we can be sloppy. We can call both of them "brightness," as people actually usually do. And so once we've defined the-- being careful about the terminology-- we can justify being sloppy about the terminology.

So for this to be meaningful, we need a finite aperture. So, so far, we've talked about pinhole model. And the pinhole model gives us perspective projection, and it's very useful for that. But we know that in terms of image measurement is problematic.

And if we make the pinhole too large, the image is going to be blurred. And if we make the pinhole too small, we don't have enough photons to count. And also, there'll be diffraction. So if you make it small enough, it'll appear just like a isotropic source that's radiating into 2 pi steradians.

So I don't know if anyone else thought about this. But in notation of infinitesimals, it is customary to do that, because you're dividing by two infinitesimal. So this should be, I don't know, two infinitesimals-- product of two infinitesimals.

Now fortunately, we're not going to get too perturbed by that, but that's a good observation. And I guess in textbooks, that'll be the notation. And then you have someone else saying, oh, what do you mean? It's del squared p. That's the power going there. What is that? Why is it squared? So I didn't want to go there.

Lenses-- so what we're going to do is invent this device that has the property that it provides the same projection as the pinhole, but has the huge advantage that it actually gives you a finite number of photons. And there'll be a penalty, which is it only works at a certain focal length. It only works at a certain distance. So let's talk about lenses-- ideal thin lenses.

So Gauss already showed that there's no such thing. You cannot make a perfect lens, but you can come incredibly close to making this ideal object. So let me-- there are three rules. So of course, you know that lenses are made of glass or transparent material. They have a refractive index that's different from air.

And so light rays are deflected when they hit the surface. And simple lenses are made with spherical surfaces. They're particularly easy to grind. You can plunk epoxy-- 100 of them onto a sphere and grind them all at once, because they all have spherical surfaces. Fancier lenses are aspherical, and they're obviously much harder to make and more expensive. But the simple ones are-- but for us, what's important is the following.

So rule number one is a central ray undeflected-- un. So what does that mean? It means that if you have a ray of light going through the center of the lens, it comes out traveling in the same direction on the other side, and not just that ray, but any ray that goes through the center, and since time is reversible, it works the other way around as well. So we'll just show rays going in one direction.

And so that's a pretty remarkable property. And why is that important? Well, because that gives us perspective projection. That means that in terms of projection, this is acting just like our pinhole.

Then number two is a ray from focal center emerges parallel to optical access. So what is that? Well, lenses have a focal length. And let's call that f0. So I can define a point here that's 1 focal length away from the lens. And what this rule is saying-- that if I take any ray coming from that focal center, and I see what it does after it comes through the lens is going to be parallel to the optical axis-- and again, of course, not just that ray, but any ray.

And then rule number 3 is-- well, it's really the same thing. Parallel ray goes through focal center. So if I have a parallel coming in from the right, it's going to go through-- that's starting to look a bit messy. So you can see that if you have a bundle of parallel rays coming in from the right, they're all going to go through that point. That's where they are in focus.

And from that diagram, we can use similar triangles to get the lens formula. And I'm not going to go through-- it's kind of boring algebra, again, possibly suitable for a homework problem. But at least I'll draw the diagram. So we have these two special points, one of which is 1 focal length ahead. And one is 1 focal length behind.

And we have a central ray. And let's see. So then that's parallel, and it's going to go through here. And notice how that didn't hit the lens. That's because this simple model doesn't have an idea of the diameter of the lens. It's what happens in this plane. It doesn't-- I can get a larger lens to do that if I want to.

And so what this lens does is kind of remarkable. Namely, it takes all of the rays that come from this point up here. And it magically brings them back together again over here. So that's the good news. The bad news is that if I move this in depth, if I go over here, it will no longer be focused in this plane. It'll be focused in a different plane.

We didn't have that with a pinhole. We didn't say anything about focus or length or whatever. So the upside is that we are now bringing a significant number of photons together in order to make a good measurement. The downside is we have this penalty that things can be in and out of focus.

Then what you can do with this is draw similar triangles. And you get three equations. And if you put them together just the right way-- and again, it's a page of fairly boring algebra. I don't really want to do that. And also, you should know that from physics. What else?

So as I think I mentioned before, we can think of the lens as this amazing analog computer, because rays come in from the side, going in every which direction. And it magically figures out where to send them to. I don't think we really appreciate how amazing this is because we all use lenses since we were very small. And we just accept that you can do things with them that depend on this property.

Now this is a planar diagram, but this applies in 3D as well, which makes it even more amazing, because now the rays coming from this point occupy a whole solid angle. One of them is going to come out of the board, and it's going to hit the lens over here. And magically, that part of the lens will redirect it to be over there.

And I mentioned that actually, you can't make a perfect ideal lens. And that's connected to this property-- that each part of the surface of the lens has to somehow deal with rays coming from a whole solid angle of possible directions. And so you can't optimize it for one particular direction. Nevertheless, by combining different lenses, you can get very good performance, very accurate performance. And there are trade-offs between different kinds of defects, one of which is radial distortion. So if you think of an image and particularly think of a fisheye lens image, where more than 90 degrees of the world is image-perhaps as much as 180 degrees. Well, in order to bring that into the plane, you have to squash the outer parts.

And so in terms of the image, what's happening is that the distance from the center of the image is distorted and is made smaller than it would be with perspective projection. And that's the only way you can get a large angle in a wide-angle lens or fisheye lens. And that type of distortion is actually inherent in lenses.

And we're not very sensitive to it if it's relatively small. And so oftentimes, lenses are designed such that they will suppress certain defects at the cost of increasing radial distortion. And that's why when we talked about calibration, I mentioned that, unfortunately, with real lenses, we need to not just find the principal distance and the principal point, but we may also need to talk about radial distortion.

So now we're ready to put it together to see what the irradiance in the image is, given a object radiance out in the world. So I'll draw a diagram of a simple imaging system that includes a lens. Now biological vision systems don't have flat image planes. But all of our cameras do.

People have built cameras with curved retinas, but it's hard. And it doesn't seem to serve any particular purpose. So let's assume as usual that we have a planar image plane. So then there's a lens, and there's an object. Let's pick some.

So I'm calling the distance from the lens to the image plane f as opposed to f0. Why is that? Well, I'm using f0 to denote the focal length, which is a fixed property of the lens. And because of this formula over here, we know that for things to be in focus, they have to actually be further away from the lens than f. And so-- than f0. And so this is going to be a little bit larger than f0, depending on the magnification. Well, the way I've drawn it, actually, it's going to be a lot larger than f0.

So then we said that the central rays are undeflected. So the rays going into the lens here come out in the same direction. And therefore, a couple of things. One of them is that suppose there--- it's a very narrow cone of directions. I can assign an angle to this that it makes with respect to the optical axis.

And that angle is going to be the same on the two sides. And the other thing I can say is that these cones of directions have the same solid angle, because they're the same rays, just turned around. There's a small patch in the image that's being illuminated. I'll call that delta I. And there's a small patch on the object that's being illuminated. And we're going to see how much power coming off that patch ends up in this patch. That's the kind of thing we'll be looking at.

But first, let's equate those two solid angles. I'm trying to relate these two areas. Then in order to do that, I need to take into account the foreshortening. So there's a unit vector on the surface of the object. And, well, this also is not immune to that effect, because the light is not coming in perpendicular to the image sensor. The light is coming in at an angle, right?

If I draw a surface normal to the image plane, then it's not the center of this cone of rays. And what is the angle between the surface normal to the image plane and the incoming rays here? So remember that diagram, where these angles are the same, and these angles are the same? So over here, this angle and that angle are the same so that alpha affects the incident light onto the image sensor as well. So now I can write down-- I can use that formula I had for a solid angle. So on this side, the area is delta I. And then there's a foreshortening effect, cosine alpha. And then I have to divide by the distance squared. So that's f squared or maybe not.

It's bigger than f squared, right? So it's f secant alpha squared. So this part is f. This is alpha. And I'm measuring this length. So you can imagine, if I draw this further out, then that's going to become even more obvious.

And now that, since we said that the central rays are undeflected, this cone of directions is the same as that cone of directions. So I can just equate that to-- and it's the same thing on the other side. And fortunately, those secants cancel out. So let's see. Do I want delta o over delta I, or do I want the other way around?

So that's half of the story. We're making good progress here. And that's important because the total energy coming off that patch depends on how big it is. And then that energy gets concentrated into that smaller patch in the image. And so the irradiance in the image is going to be whatever power ends up here divided by this area, delta I. And so we'll be able to relate the radiance over there to the irradiance over here. And this is what we measure. It's the irradiant state.

So now we have to think about how much of the light from that patch on the surface actually is going to be concentrated into that image. So the good news is that because the lens focuses the rays, all I need to know really is the solid angle that the lens occupies when viewed from the object. And so I need to know what the solid angle is because the rays that come out the other side all get concentrated into the corresponding patch.

If things are in focus, the light that comes off this patch and goes through the lens is all concentrated into this area in the image. And conversely, light coming from anywhere else has no effect on this. It's imaged somewhere else. So there's a direct match between the power coming off here that makes it through the lens and the power that's delivered to that small area in the image.

So what is this? Well, it's the area of the lens. Let's suppose the lens has a diameter d. So it's 4 pi squared. And then there's a foreshortening effect. There's an angle which we drew as alpha. And then we have to divide by the distance squared. And that, again, is f secant alpha squared. So that's the solid angle.

And typically, that's actually quite small. And so we typically only gather up a relatively small fraction of the-- oh, sorry, not f, z. And now this time, the cosines and secants don't cancel out, unfortunately. And we get that. So we're almost done.

So the power delivered to that small area in the image is the radiance of the surface times its area times the solid angle times cosine theta. And that's going to be L delta o pi over 4 d over z squared. So that's the total power we're delivering. And it's concentrated into an area of delta I. So the power per unit area is we just divided through by delta I. And that's what we actually measure. So that's the conclusion of that. So let's study that a little bit carefully.

Let's look at this d over f. What is that? Or maybe f over d looks more familiar. I guess there are no amateur photographers here. That's the F stop. That tells you how open your aperture is, is f over d. And so typically, SLRs will have a maximum opening of maybe, I don't know, 1.8 for that ratio.

And you can stop it down to, I don't know, 22, let's say. And so obviously, the square of that controls how much light you get. And that's one way of controlling the exposure. The other one is time. And so often, there are trade-offs between using the aperture opening versus using time. For example, if you want a lot of depth of field, then you can achieve that by making the aperture very small, approaching a pinhole.

But the cost is you need a longer exposure. So if things are moving, that's not going to work. Conversely, if you want to have a great portrait, and you want to wash out the background out of focus, then you go the other direction. You open the lens very wide so that it only has a narrow field of depth and then use a very short exposure.

So anyway, that quantity is one that's well-known to people working with cameras. And it's intuitive that the image irradiance goes as the inverse square of the F stop. And that's why the F stops are usually done in square roots of 2. So it goes 1, 1.4, 2, 2.8, 4, 5.6, 8, 11, 16, 22, 32.

And those are the steps of square root of 2 in size of the aperture, which give you steps of 2 in exposure. So that's that. So that's kind of intuitive and not particularly interesting. The pi over 4-- who cares? It's just a constant. We're not too excited about that.

What's really exciting, though, is that the thing we measure, E, our image irradiance, is proportional to L, the thing we're interested in out in the world-- the radiance. So that's why we can be sloppy about talking about brightness, because the brightness of the surface, radiance, is proportional to the brightness in the image, irradiance. So we're measuring brightness in the image. And that has a meaning beyond just power per unit area in the image plane. But it has a meaning out there in terms of how much that object is radiating.

And the remaining part is this annoying thing here, cosine to the 4th of alpha. What does that do? Well, it means that the brightness is dropping off as we go off-axis. So if you have a part of the image that's way out in the corner here, it's going to receive less power per unit area than something in the middle of the image. And that means that you need to take that into account.

Now fortunately, for a small alpha, cosine to the 4th alpha is as close to 1 as you can get. And so as long as alpha is small, you can completely ignore this. So when does alpha get big? Well, it only really gets big when you have a wide-angle lens-- when you're taking in a large, solid angle of the world. And then alpha may be 20, 30, 40 degrees. And then it starts to matter because cosine of 40 degrees is not any more close to 1.

Fortunately, it's a fixed thing, so you can compensate for it. And in fact, if you buy a DSLR, part of the magic that happens that you're not allowed to see or understand is to compensate for this. Also, we're not very sensitive to this effect. So if you have an image that slowly gets darker towards the edge, you really have to focus on it to notice it.

So when I was much younger, I was very keen to have a telephoto lens. But I couldn't afford a telephoto lens. And then I saw an advertisement for a Russian catadioptric telephoto lens. And having worked in telescopes, that rang a bell. And I could afford that one, so I got it. And it was quite nice.

And then I looked at-- in those days, we used slides and projected them on the wall. So I'm looking at this slide of a predatory bird that I took looking up in the tree and the sky background behind it. And the sky behind the bird is white. And I'm like, OK, this is odd. And strangely, in the corners, the sky is blue. So what's going on there? Well, what's going on is that this lens was cheap in part because it had a rapid drop-off in brightness with the angle, perhaps even worse than this. And so the corners were not illuminated as well as the center. And in the center, there was enough light so that these-- the colors were oversaturated. So not only did I get the green and the red channels overexposed, but I even got the blue channel overexposed.

So it just looked white. It had more than the maximum intensity that the film could handle in all three channels. Anyway, so this comes up a lot in other situations. In X-ray imaging, it's slightly different. It's cosine cubed.

But it's something that people often forget about. And as I said, you can kind of forget about it, because it's something you can compensate for just in your image-processing chain. It doesn't change from image to image unless you change the lens and the-- yeah, OK.

So this formula is central, not because we're going to play a lot with it, but just because it justifies this whole idea of talking about brightness and measuring it using gray levels in the image and thinking that that has something to do with what's out in the real world. That's our description of what the camera does in terms of brightness. So this is the counterpart to what it does in terms of position, which was the perspective projection equation, which, of course, was trivial in comparison. But those are the two key cornerstones of understanding of cameras.

So now that we know what it is we're measuring-- well, we already knew it was the power per unit area we're measuring. But now we also know that that corresponds to radiance in the world. We need to try and understand what determines the radiance in the world.

And we already mentioned it depends on the illumination. It depends on the material. It depends on orientation. So let's try and make that a little bit clearer and talk about the bidirectional reflectance distribution function.

So let me tell you that until, I don't know, the '80s, this field was a complete mess. And there were dozens of different terms, all competing for attention, some of whom had famous people's names on them, and some didn't. And then the National Bureau of Standards stepped in. And a brilliant man named Nicodemus cleaned it up. And who knows? Maybe the last thing the National Bureau of Standards did that was very interesting. But it was a very powerful effect on imaging, not just optical, but X-ray as well.

So what is this? Well, the idea is that we have light coming in, and we have light going out. And crudely speaking, when we talk about reflectance, the ratio of those two things is reflectance, right?

So crudely speaking, something that's white reflects all of the light coming in. Something that's black reflects none of it. But where's that light going? So it's not as simple as just saying, oh, reflectance is 0.3. It's more nuanced. And that's what this is about.

So this basically is a quantity that depends on the incident and the emitted direction. And it tells you how bright the surface will appear. So let's look at it. So we have a patch of the surface. I'll use a little diagram-- should have a rubber stamp for that. And we have light coming in, and we have light going out.

And how bright the surface will appear, its radiance, will depend on those angles, but more, because this is in 3D. I've drawn it on a plane, but it's, of course, really in 3D. So I really needed to talk about the directions of these two rays in more detail than just saying they make an angle with the surface normal. And so how do I talk about directions? Well, we've already been through this-- unit vectors, points on a sphere, latitude, longitude. So in this field, it's customary to use these angles. So let's suppose that this line-- this curve down here is in the surface. And I've constructed the hemisphere above it. And this is called the polar angle, also called colatitude. Why? Because it's 90 degrees minus latitude. And this is called azimuth-- azimuth angle.

So to specify the direction of light coming in or light going out, I need two angles, polar angle and azimuth. And so here, I've only shown the polar angle. And to draw the azimuth, I'd have to project this down into the plane of the surface and then look at the direction of those lines.

And then since the brightness is going to depend on all of those, I can write it as some function like this. And the official terminology is that's the bidirectional reflectance distribution function-- theta E. And as we said, reflectance should be power going out divided by power going in. And so that's what this is, clean and formalized.

So what we've got is it's the radiance, how bright the object will appear when viewed from this position, divided by how much energy I'm putting in from the source direction. So this is finally a definition of reflectance that actually works, because the other definition of just saying, oh, it's 0.3 or something, doesn't. Unfortunately, this is much more complicated. But any other definition of reflectance can be based on this. Basically, it's an integral of this.

Well, there are lots of questions that immediately come up. One of them is, how on Earth do you measure this thing? Well, one way you can do it is put the light source in a certain position. Put your camera in a certain position. Take a measurement.

Then move your camera, blah, blah, blah. Move your light source. But you're exploring a four-dimensional space, so that's pretty expensive. Nevertheless, it's done. And one method is using goniometers. That's just a fancy way of saying angle measurement devices.

So what I'm trying to draw here is something that will rotate about this vertical axis and then has an arc along which the apparatus can move. So the first rotation corresponds to the azimuth angle, and the second movement corresponds to the polar angle. So the pole would be up here.

And so we can mount the light sensor on one of these. Then we get a second one, and we mount the light source on it as well. And then we make measurements. Well, you'll get tired of it pretty soon, because you're exploring this four-dimensional space.

And even if you sample it pretty coarsely, like 1 in 10, you're not talking about 10,000 measurements. But of course, you can automate it. You can build a robotic device that mechanically moves this and go away for the weekend, and it'll do these measurements for you.

If you're in a hurry, as people in movie-making are, you can have many light sources that you can turn on and off. So you could construct maybe a whole sphere or maybe a hemisphere. And you distribute light sources all over the surface that can be individually controlled. And then, if you want-- so that takes care of two of the angles.

So by picking-- by turning on one of these light sources, you've controlled one of the goniometers to position the light source. And then you can have light sensors or perhaps even cameras interspersed. And now you can do all of your measurements very quickly because you just flash one source at a time and take pictures and then process the images you get to-- and why would you want to do this? Well, suppose you want to realistically model how someone's skin reflects light. Well, you could try and build some mathematical model. But who knows how good that is? The better way is just to measure it. And so that's something people do. You don't want to approximate it with, I don't know, Lambertian or some other well-known model. You want the real thing. And so you can do this.

So that's part of the story of this four-dimensional space. The next part is that, well, in most cases, it's not fourdimensional. And I think if you look at this diagram, you can see that what really matters is the angle between these two azimuth lines. [INAUDIBLE]. So this is a general formula that applies to any surface.

But for many surfaces, what really only matters is the difference between these two angles. And what type of surfaces are those? Well, those are surfaces where if you rotate them in the plane, if you rotate them about their surface normal, they don't change brightness. So it's pretty much true of this thing. And it's pretty much true of lots of things like wood and this floor and so on. So that's a dramatic improvement because now you've got a 3D lookup table to fill in instead of 4D.

So what kind of materials are there that do not satisfy this? What materials require that we take the full four dimensions into account? And think about it. Is there surface where, if you look at it, and you rotate it in the plane of the surface, it changes appearance? Can you think of something like that?

STUDENT: [INAUDIBLE] iridescent?

BERTHOLD Something that's iridescent. Why? Because an iridescent material has microstructure that's oriented, and like a
HORN: hummingbird's neck, the feathers are lined up. It produces the color not by pigment but by interference. And so if the microstructure is my fingers, and then I rotate that microstructure, it'll diffract light differently.

Our ruby hummingbirds, if you look at the neck of the male from the wrong angle, it looks black. Why? Because it's reflecting all of the light in one direction, and it's only the red. So that's an example of something where unfortunately, you need the full model. You can't reduce and use this.

And the other examples, like the semi-precious stones called tiger eye or various other things-- they're basically very fancy forms of asbestos. And people don't call them that, because as soon as you say "asbestos," lawyers come and so on. But it's basically-- just as asbestos, it has a microstructure that's very linear and very tightly packed on the scale of the wavelength of light.

And so it has a very different appearance as you rotate it which gives it its appeal as a piece of jewelry. Then some people have very straight black hair that's very parallel. And that will have the same effect, where because they're all parallel, they'll reflect light in a certain way. And so as the head rotates, the sheen on that hair will move. And so you need the full four-dimensional model for that, whereas for a lot of things, like paper and snow and strawberries, the three-dimensional model is-- so what else do we know about this bidirectional reflectance function?

Well, there's an important property due to Helmholtz called the Helmholtz reciprocity. Now Helmholtz, of course, lived a long, long time ago, long before Nicodemus of the National Bureau of Standards. So how could he have come up with a property of the bidirectional reflectance distribution function? Well, he didn't call it that, but he had the basic idea, which was basically the second law of thermodynamics. So suppose we have two objects at different temperatures. And there's an object patch down here. And so there's radiation coming from one and going to the other. And reciprocally, there'll be radiation coming from the object at temperature T2 and arriving at T1. And it takes a little bit of handwaving, but basically what it's saying is if it's not a reciprocal, then there will be energy transfer from the colder object to the hotter object, which we know doesn't happen.

So what this is saying-- if you interchange incident and emitted, you should get the same value for your bidirectional-- so there's a symmetry, which, in a way, helps you in the data collection, because there's half of the data you don't have to collect. By the way, it reminds me of when I was a student here, there was a famous professor called Lettvin. And he, Jerome Lettvin, came up with a paper called, "What the Frog's Eye Tells the Frog's Brain," which was one of the early attempts to try and understand how neurons work and how image processing might work and so on.

And everyone wanted to hear his talk. And he was talking about color. And I was way back in the room. And as a student, I was really intimidated by these people. But I just had to ask. And so I stuck up my hand, and I asked him some question. I at this point don't even remember what it was. And he stared at me for about five seconds.

And then he said, well, if you had read the original book by Helmholtz-- the book by Helmholtz in the original German version-- you would know blah, blah, blah. And I'm like, oh, my god, I really killed my career here, because-- and then years later, I was thinking about that same problem. And I realized that he was just using a wonderful debating technique because he had no idea what the answer was.

But he certainly put me in my place. And some of these people went to schools where they taught you how to debate. I didn't. So I was just flabbergasted-- anyway, Helmholtz. So that was Helmholtz there.

So next time what we're going to do is apply this to look at different types of surface material models, of course, using Lambertian again, and also some new ones that apply to the moon and rocky planets in our solar system and that enable us to determine their surface shape. And I guess there was an extension, right?

STUDENT: Yeah, till tomorrow.

BERTHOLDSo make sure you keep up to date on what's on Piazza because there was an extension on the homework**HORN:**problem. And other good stuff is happening there, so.