

[SQUEAKING]

[RUSTLING]

[CLICKING]

BERTHOLD

HORN:

Let's have a quick review of what we learned about photometry. So there are a number of concepts, one of which was irradiance. And we use the symbol E for it. And it was power per unit area. And it's a way of talking about light falling on a surface. And it's what we measure in the image plane and convert to what's commonly called a gray level.

So the quantity of interest here is directly used when we're imaging. But it's also, of course, a measure of light falling on the objects that we're imaging. Then we talked about intensity, which applies to a point source. And it describes the power per unit solid angle. And so we had to define the solid angle. And it's the quantity that typically varies with the direction.

So if you have a good old incandescent light bulb, it's very low intensity in the direction of the base, because that's blocked by the base and some higher intensity in other directions. And that's a quantity that isn't of a whole lot of interest to us here. It's just interesting because, a, it's simple to define. And, b, it's the terminology incorrectly used to talk about the quantities that we are really wanting to talk about.

So the important one is radiance, which is basically a measure of how bright a surface appears. So again, we have a little facet of the surface. And we're looking at how much power is emitted per unit area and per unit solid angle.

And that's of interest to us, because that's what we actually measure with our instruments, cameras, and that's also obviously relevant to what we see. That small solid angle is perhaps the entrance pupil to your eye.

So we then looked at cameras and anything with a lens in it. And we came up with this relationship between the radiance of a surface that we're imaging, that's L , and the irradiance E of the corresponding part of the image. And so it gives us a direct relationship between something out there loosely called brightness and something inside the camera loosely called brightness. And the reason we can be loose about it is because they're proportional to each other.

So and then there's the π over 4, which is just a constant factor. And then there's this 1 over f -stop squared, which is kind of obvious, because we're limiting the solid angle-- the ω over there by opening or closing the aperture on the lens. And the area of that goes as the square of that ratio. And it's the area, of course, that we need when talking about the solid angle.

So then the next question is, OK, we're measuring E . And it's proportional to L . But where's L come from? What determines the radiance of a surface? And we already indicated that well, illumination, it's going to be directly proportional to the amount of illumination. And it's also going to depend on the geometry.

So how is the surface oriented? And it depends on the material. And that's where we-- bi-directional BDRF. And that's where we introduced the bidirectional reflectance distribution function, which is a function of the incident direction and the emitted direction.

So we have a light coming into a surface. And we have light re-emitted from that surface. And that's obviously the idea of reflectance. How much of that light going in is reflected.

Except it's not quite as simple as that. It's not simply a ratio of what percentage of the incoming light is reflected. But we're interested only in the light that's going to hit the camera or the eye. So we're actually using this terminology. So it's going to be ΔE -- let's see, ΔL of θ emitted.

So this is the radiance of the surface. And this is the irradiance. And so, that's what you imagine some definition of reflectance to be. And it's the detailed fine grain definition of reflectance, from which we can derive other quote reflectances. For example, albedo, which is the total output power divided by the total input power.

Well, in order to compute that, we just take this quantity, and we integrate over all possible output directions. Because in this case, we're interested in the total power going out, not just what's going to a particular light sensor. And in the process, we may need to think about spherical geometries.

Then we said that this quantity, this BRDF, has to satisfy a constraint. Which basically says, if you interchange the directions to the source and the direction to the viewer, the BRDF should come out the same. And that's because if it wasn't, then we'd be violating the second law of thermodynamics, which periodically people try and do, but generally don't have too much success with. And so, we can't just have any old function there.

By the way, in computer graphics obviously they use models of surface reflectance. And quite a number of those models violate this constraint. And yet we don't seem to care. We like the pictures, which suggest that this isn't something critical to a human or machine vision, other than it's kind of a shortcut. If you've measured one of them, then you've got the other one. So it cuts the number of measurements you need to take in two. Because you can just by symmetry find the other one.

So example. Well, we've been talking about Lambertian surfaces. And the Lambertian surface has the property that it appears equally bright from whatever viewing direction you have. And if it's an ideal Lambertian surface, it also reflects all the incident light.

So property number one, and this is the condition that's usually misstated in terms of emitted energy. So it's usually stated incorrectly as it's emitting light equally in all directions. So that's going to greatly simplify whatever formula we come up with for the Lambertian surface, because it's not going to depend on two of the four parameters.

And then the other condition is that if it's an ideal Lambertian surface, it reflects all light and doesn't generate any of its own. So as I indicated, a lot of work with the BRDF, the BRDF is the atomic thing. It's the low-level detail. And in many cases, we're interested in integrals of that.

So for example, if I don't have a point source, I have a distributed source like the lights in this room, how can I deal with that. Well, I can simply integrate over a hemisphere of incident directions. So I'd integrate over that quantity, taking into account how much light is coming from each direction.

So similarly, here we need to integrate over a hemisphere to get all of the energy that's coming off the surface. So we had this way of dividing up using-- so we use the polar angle. And then there was also an azimuth angle. So this is one way we can talk about the possible directions, two parameters. And if we perform the integral, we need to take into account the area of this patch, which is obviously going to involve $\Delta \theta$ and $\Delta \phi$.

But it's also going to get smaller the closer we get to the pole. And since we're measuring theta from the pole, this would be sine theta. It's kind of unfortunate that they didn't pick latitude. But they picked co-latitude. But whatever, we can do that.

So now we're dealing with-- we're trying to integrate our all emitted direction. So in this particular case, we're talking about those quantities. So azimuth, well that range is over a 2π range. So we're going to be integrating from minus pi to plus pi, for example.

The polar angle, well, we're not interested in points below the horizon, because the object itself is blocking. It's only emitting above the surface. So we only have to deal with 0 to pi over 2 for the polar angle.

And then we're going to have to include this term here. And that's obviously the Jacobian-- the determinant of the Jacobian of the coordinate transformation. But I find it easier just to draw the diagram.

And then what's in here. Well, there's F. And now, we've decided that F is a constant. So let's just do that, just write F. And then the light that's falling on the surface depends on the incoming radiation and the angle. And we're saying that all of that gets reflected.

So the lights coming in at a certain angle, there's foreshortening, so the power deposited on the surface is $E \cos \theta$ times the area of the surface. And we're saying that that's all going to be reflected. So when we integrate the reflected light, which is the BRDF times this quantity, then that should equal the incoming light. So we can just cancel it out conveniently.

So what we're looking for is what is this constant value of F for the Lambertian surface. Is it one, or some other convenient quantity. So first of all, the azimuth angle doesn't appear anywhere in the integrand. So we are going to evaluate this quantity and then just integrate over ϕ 2π .

So that's just 2π times that inner integral. So we're going to have $2\pi \int_0^{\pi/2} \sin \theta \cos \theta d\theta$ into-- oh, no. We've already dealt with that. F, well actually, we can take the F outside, so it's a constant.

Now this is a $\sin^2 \theta$ E. And so if we integrate that, we get $-\frac{1}{2} \cos^2 \theta$ E in the limits 0 to pi over 2. So we plug in pi over 2, we get cosine of pi, which is minus 1. Minus 1 times minus 1/2 is 1/2. And then we subtract what we get by plugging 0 in here, cosine of 0 is 1. And so we're going to subtract minus 1, which is like adding a 1/2. And so, this whole thing comes out to be 1. And so the result is that F is 1 over pi.

So that's it for Lambertian surface. That's the BRDF for the Lambertian surface. And that's as easy as it can get. And there's some question about why it's 1 over pi and not 1 over 2π . So let's think about that. So if you think about the hemisphere of possible directions. So here's our surface element. And it's radiating into all these directions.

And what is the solid angle that's occupied by that hemisphere? Well, 2π , of course. So the object is radiating into the hemisphere that's above its level, above the plane through the surface. And that's 2π .

And if we were radiating energy equally as isotopically into that hemisphere, then F should be 1 over 2π . And so those people who say it's radiating equally in all directions would end up with 1 over 2π for that. So what's wrong with that?

Well, what's wrong with that is that it appears equally bright from all directions doesn't mean it's radiating equally in all directions. So imagine that you're on the surface of the sea, and you're looking in at this object. There's going to be foreshortening. So if you're straight above it, you see its full area.

If you're off at an angle, you see an area that's reduced by the cosine of the polar angle. And so, what does that mean. Well, that means that if you emitted the same power, then the power per unit area would be growing. And when you get to be on the horizon, now the area of that surface element is pretty much shrunk to nothing. But you're still radiating the same power, supposedly.

Well, that means that the power per unit area is infinite, and it will fry your retina. So you don't want that. And that's not what it does. It is radiating this in this direction and in that direction. So but it's in proportion to the area. So that the power per area stays constant. So it appears equally bright. So that's condition number one. It appears equally bright.

And so, that means that actually it's radiating more up here and less down here. And in just such a way that we end up with $1/2$, we end up with 1 over π instead of 1 over 2π . So again, the idea that Lambertian surface radiates equally in all directions is wrong. And it'll give you the wrong answer here.

Now how do we use this? Well, let's-- simple case. Notice that there's no cosine theta ion here. So what's with that? We've made a big fuss about Lambert showing that it depends on the cosine of the incident angle. Well that's because that controls the foreshortening of the incoming radiation.

So suppose that we have a distant source of radiation. And that it has an irradiance perpendicular to its rays of E_0 watts per square meter. Now we are illuminating that surface. And of course, that surface has an area that's larger. So if we call this A , and we call this area A' , then $A' = A \cos \theta$.

So that means that this captures a smaller amount of the incident radiation than it would if it was oriented perpendicular to the surface. So we find that-- so if we measure our incoming light in terms of surface area, then L is 1 over π times that power per unit surface area.

Just as you'd expect, there's a 1 over-- there's the BRDF. If instead we measure it relative to the incoming radiation, perpendicular to the direction of that radiation, we have to take into account the foreshortening. And then we get the familiar expression for Lambert's law.

And so that's a little thing that you have to keep in mind and avoid confusion. Here is an example of how you might get confused. Helmholtz reciprocity. You look at this formula, and you say, oops, there's no cosine theta E . So it doesn't satisfy Helmholtz reciprocity. So it's not a physically possible surface. But the Helmholtz reciprocity applies to the BRDF, not that.

And here, this is obviously if I interchange θ_I and θ_E , it's the same. It's one over π . So we have to be a little bit careful when we ask questions about Helmholtz reciprocity, for example. This is a perfectly valid formula, but that's not the one that you want to apply Helmholtz reciprocity to. It's instead the BRDF, underlying BRDF.

So that's Lambertian, which is really simple. And we should have some other examples. So let's see. So let's try this. So this is another example. So I'm not picking this totally at random. We're going to use this particular type of surface quite a bit. So I might as well introduce it at this point.

So for this one, the BRDF, let's see, isn't the constant. It's something like that. It's $1/\sqrt{\cos \theta_i}$. And in that form, we can immediately answer the question, does this type of surface satisfy Helmholtz reciprocity. Well, yeah. If you interchange θ_i with θ_e , you get the same answer.

Now when we use this model in practice, we're adding illumination and looking at how bright the surface will appear under certain illumination. And so, we do what we did over there.

So the radiance is going to be the irradiance times the BRDF. And it's the irradiance in terms of power per unit area on the surface patch. And that's going to be affected by the foreshortening, because when it's tilted, it's going to receive less power. So that's interesting, because that's now going to be-- so here we have a surface that acts quite differently from a Lambertian surface. Instead of having $\cos \theta_i$, we have this funny ratio.

And so, it turns out that this type of behavior is what we find on the lunar surface. Well, more specifically, the area, the dark areas, the maria, where volcanic eruptions have occurred to fill in the basins-- but actually, rocky planets in general, and asteroids-- some asteroids. And it's not a bad model for them. And it's significantly different from Lambertian. And by the way, this was the basis of the first methods for recovering shape from variations in brightness.

Now let's see if we can learn something about this type of surface. So one question is, if we look at the moon, what are the isophotes. Now of course, we know that the lunar surface has some texture on it and rays ejecting from craters that are brighter than the background and so on. But let's pretend that the lunar surface was pretty much uniform in its reflecting properties. And what we'd like to know is what is the contour map of brightness.

Now if it was Lambertian, we know that all of the points that are the same angle from the sun, where the surface normally has the same angle with respect to the sun have the same brightness, $\cos \theta_i$. And so, if we were to look at the isophotes on the sphere, they'd all be nested circles.

And then if I project them into the image plane, those circles are at an angle. So the circle gets turned into an ellipse. And eccentricity depending on just how much of an angle. So this is what I'd expect to see. And this is what I would see if say I took that calibration object, that sphere painted white in the lab. If you plot its isophotes, they pretty much look like that. So that's for Lambertian.

So what about this other material? Well that's a little bit more tricky. So let's see how we can do that. Now with an a Lambertian, we could just set $\cos \theta_i$ as the constant. And that means θ_i is a constant. And then you find all the places where the angle between the direction to the light source and the local surface normal is the same.

And you just spin that around to get a cone, and you're done. So this one's a little bit harder. So we're doing it for this one. And so here we have L is constant for-- so we're now looking for all of the points on the surface that have a certain ratio of $\cos \theta_i$ over $\cos \theta_e$.

Well we can write this in terms of unit normals that might make it easier to see what's going on. So $\cos \theta_i$ is the unit normal dot product with a light source direction. And $\cos \theta_e$ is the unit vector with respect to the viewing direction. So we're now looking for all values of n that make this a constant.

And so, for the constant C , that's what we get. And now-- and so we have a dark product that's equal to 0. That means we have two vectors that are perpendicular to each other. And so, let's fix the constant C for the moment. Then this is some fixed vector.

And this is saying that all of the ends that satisfy that, all of the ends that have the same brightness must be perpendicular to that vector. So what is the set of vectors that are perpendicular to a particular vector. What does that look like. What's the locus of the endpoints of those unit vectors.

So we have some vector. And then we're saying that $\mathbf{n}$ is perpendicular to that. That's one $\mathbf{n}$. Well, here's another one. So we actually get a plane. So that if we think of the unit vectors of all the points on the surface that have the same brightness, they all lie in a plane.

So that's already useful information, which isn't the case here. Because the unit vectors-- this unit vector points up towards the pole a little bit. And this one points in a different direction. So that's already very non-Lambertian.

Well we're not going to do all the details, but we can benefit somewhat from thinking about spherical coordinates to figure this one out. Now again, it's sad that they picked a polar angle instead of latitude. So it looks a little different from your usual formula. But of course, it's really just the same thing with subtract from 90.

So I can always write a unit vector in this form. A unit vector has only two degrees of freedom. And I've picked the polar angle and the azimuth here as those two convenient parameters. So what I'm going to try and do is make some advance in understanding this by substituting for all of them in that form.

Now I'm going to end up with an algebraic mess unless I pick my coordinate system as well. And so I happen to know that here's a good way to pick the coordinate system. So here's the direction to the sun. Here's the direction to the Earth, where the viewer is.

And usually, we think of the North Pole as being up above this plane, right angles above that plane. And this isn't going to be perfect, because the plane in which the moon orbits around the Earth is not exactly the same plane as the plane that the Earth moves around the sun. But let's pretend it's the same.

So we're going to pick this preferred coordinate system where this is the z direction. And so, the sun and the Earth is at Z equal 0. And that means that when we write the vectors for them, we can leave out the third part.

So their position depends only on the azimuth, since we've picked this convenient coordinate system. So it's a little bit simpler. There's only that one unknown. And the third component is 0. And so now, we go back to our expression for these normals. And I guess I can write it here. And I suppose I do need the other board.

So before we do that, by the way, there's something you can already ascertain right now, which is what happens at full moon. Well, at full moon, from Earth you're looking at the moon in the same direction as the incident light from the sun. So that means that θ_I is the same as θ_E . So that means that it's constant.

What does that mean? Well, that means that the disk you see should be uniformly bright aside from the surface markings that we discussed. And that's completely normal Lambertian. If we had a Lambertian surface sphere illuminated from the same direction as the viewer, you're holding the flashlight next to your camera, then we expect to see isophotes that look like this.

Because the incident angle here at 0 and then it increases to 90 degrees. So the cosine of the incident angle goes down. And so if we are looking at a sphere, and it has isophotes like this, we recognize it has a three-dimensional shape, and it's kind of spherical. And all of a sudden here, that's not the case.

And so, that's pretty interesting. Next time you look at a full moon, you realize that it doesn't look like a ball. I mean, you know intellectually, that it's, well, unless you belong to the Flat Earth Society, you probably believe the moon is flat as well. But leaving that out, it doesn't really look round.

You can see that it must be sort of round, because the outline of it is a circle. But it doesn't look quite right. And this is why. Because it in fact, in opposition at that time, is pretty much uniformly bright. And this is it. It's because it's not Lambertian. It's a different microstructure. And the Hapke model is a pretty good one for predicting that kind of thing.

Well, let me kind of jump from that. So we're going to have n_s -- which way around did I have it? n_s over n_v is a constant. And now I can plug in the spherical coordinate versions of those vectors and leave out a couple of steps. What I'm going to get is a bunch of terms

Canceled, and we end up with that. Now unless, of course, if θ_S is the same as θ_V , i.e. opposition, then we just get 1. But suppose that we're not, then this can only be true if that is true.

So what does that mean? Well, it means that all of the points on the surface that have the same brightness have the same azimuth. So in our coordinate system, that means that here's a great circle, not drawn very well, but that has a fixed azimuth, fixed angle.

If you go into the center, and we look at this direction, that's the same for all of them. So that's one isophote. And here's another one. So the lines of constant longitude are isophotes. And that's again very different from Lambertian and it makes the moon look odd.

I'm just thinking of something. That when I was a little kid, we went for a walk in the Black Forest in Germany, and the moon was just rising. And the adults thought they'd have a little joke at my part, and they said, well, how far away do you think the moon is? And I'm like, OK, if they ask me that, it must be much further away than I think. So I said, I don't know. 100 meters. And they all laughed.

So anyway. So it's hard to estimate properties of celestial bodies. For example, I already mentioned that we would be surprised to know-- that most people are surprised to know that the reflectance, the albedo of the moon is about 0.1, which is the albedo of coal. And yet it looks so bright in the sky.

And that's because we don't have any comparison. We don't have anything near it. And all we can measure is the product of incident light times reflectance. And then we try to separate that into those two components.

Now in our own world, often we have that the illumination is more or less constant of an area. And so we can separate changes in reflectance from changes in brightness, particularly if we have some calibration objects like a piece of white paper. You recognize that, and you say, OK that's one. And everything else can be measured in relation to that.

But if we only see the product of illumination and reflectance, it is totally ambiguous. We don't know whether it's dark because the illumination is weak, or because the reflectance is low, and so on.

So we can go a little bit further with this. For example, we might say, well, suppose we take a picture of the surface under two different illumination conditions. Can we find the orientation, the surface orientation locally. And yes, we can. And it's just photometric stereo the way we've done it before.

But then the remaining question is, can we get the shape of the surface. So let's talk about that a little bit. So this Hapke model, we've looked at it here in terms of the angles. But we also know-- and the unit vector-- we also know that we can use the gradient as a way of talking about surface orientation. So let's look at what this is like in terms of surface orientation.

By the way, you may wonder why there's a square root. I mean, it would still be-- well, it's partly because we want to make sure that we satisfy Helmholtz. Because if it wasn't the square root, then when you divide by cosine theta I, you get, I don't know, $1/\cos\theta_i$, which is not symmetrical.

So that wouldn't work. And the other reason is that you want the integral of all the outgoing radiation to equal the incoming radiation. And that doesn't work if you don't have the square root. It becomes infinite.

So now we can plug in for the various unit vectors. So this was our way of converting from gradient to unit vector. And then, we can use the same notation to talk about the position of the light source.

And we've usually chosen the coordinate system so that Z is the direction that's coming straight up at me. So it's along the optical axis. So Z is the viewing direction. So V is just Z. So now, I'm going to take those dot products.

And this was the messy part of Lambertian. We had this nonlinear term. So if we're trying to plot isophotes and so on, this would create a second order component. So we ended up with conic sections. But if we now take the ratio of these two, we get something that's linear.

And r_s is just a shorthand for this thing. So r_s is square root of $1 + p^2 + q^2$. And it's constant if the light source is in a fixed position. And it's just a nuisance to have to write that out all the time. So we're not quite done, because actually, we want the square root of this thing.

But what do the isophotes look like in gradient space? So well if the square root of this is a constant, then this quantity itself is a constant squared. So we can look at isophotes in terms of this formula. So what are the isophotes? Well it's when $1 + p^2 + q^2$ is a constant.

And that's a linear equation in p and q. So it's-- what is the curb in pq space? If it's a linear equation in p and q? It's a line. Right. So I can-- which is going to be great compared to Lambertian, which had these conic sections.

So that's one line. Now suppose I plot another isophote. Well, it's going to be a line again, just with a different constant, because this will be different. But the $p^2 + q^2$ will be the same. So it'll have the same orientation. So the other isophote might look like that. And so there's a whole bunch of parallel lines that are the isophotes. Now, they're not equally spaced, because I'm taking the square root of this thing.

But other than that, so I'm going to have, I don't know, something like that. So this is my plot in gradient space. And there's one particular line, which is where brightness is 0, where I've turned 90 degrees away from the light source. And just as with our Lambertian cosine theta, we need to be aware of the fact that brightness can't be negative. So actually, this part of the diagram is 0.

So why is this exciting? Well, because it's linear. It's going to make it very easy to solve all sorts of problems. And so first of all, since we are coming from photometric stereo, suppose that we have one lighting condition. And then we have a different lighting condition. Well, then we'll get straight lines, but different straight lines. So I don't know, maybe like this.

And then obviously, if I have the two measurements, I can find the intersection of the corresponding lines. So suppose that the measurement in the one lighting condition was that. And on the other lighting condition it was this line. Then there's the intersection. So that's the answer. That's the surface orientation.

So photometric stereo is very easy. And of course, I can-- this is geometrically. I can do this with equations. Just we have two linear equations of this form, $1 + p^2 + q^2 = \text{something}$. And of course, we know how to solve linear equations. So that's kind of neat. And there's no ambiguity.

With Lambertian, we had to conic sections intersecting. And they could intersect in up to two places. No, actually, by Bézout's theorem, we got two second order equations. And by Bézout's theorem, there might be as many as four solutions.

And so it turns out that, well, that's for arbitrary second order equations. But what if you have the particular equations we have. Well it turns out that in that case, there can be only two. Anyway, here there's only one. So that's an advantage.

Then another thing that we can read right off this diagram is that-- so we don't know from one measurement, we can't determine the surface orientation, as usual. But we do have something pretty powerful, which is the surface orientation in a particular direction. And so when we-- oops.

So this is for a particular orientation of the coordinate system. I've picked some x, y, z coordinate system. And this is what I get in the diagram. What if I pick a different coordinate system. What if I turn this by some angle. I don't know, call it α . Well, it turns out that then I turn this. And if I turn it the correct way, I get a pretty neat result.

I'm just stating this. I'm not proving it. But you can prove it, keeping in mind that p is dc, dx , and q is dc, dy . And p prime is dz, dx prime. And q prime is dz, dy prime. And then use chain rule, and you basically get a rotation through an angle α . So it's sort of surprising that the first derivatives rotate the same way as the coordinate system. But that they do.

Why is that great? Well because, too bad I messed up that diagram. Well, let me do it again. So I'm just copying that first diagram. And now, suppose that I pick this as α . Well then, when I measure a particular brightness, let's say this.

As usual, I don't know the surface orientation. But I know one component of it. I know that the component in this x prime direction has a certain value. Because all of the points on this line have that same slope in that direction. I don't know anything about the slope in the direction at right angles. But I know the slope in that direction.

And that's different from Lambertian. Because in a Lambertian case, we had this curve. And so the orientation was different along the curve. Here we've got a line. And all points on that line have the same distance from this origin from here. So you can find p prime.

And then, what is this angle alpha? Well, it's obviously some function of p_s and q_s . $\tan \alpha$ is, let's see-- and how do I know that? Well because I want this straight line, the $1 + p_s$, plus $p_s p$ plus $q_s q$, I want that to become the vertical axis. And so I have to find the angle that will make one of the two terms disappear after the rotation. So anyway, that it's not a particularly important point. But that's the angle I actually want to use.

So this is more exciting than you might think. Because what this means is, I can look at the surface, if I pick the coordinate system right, and measure the brightness. And I immediately know how steep the surface is in a certain direction. So I could then traverse. I could say, it's going up by 1 meter and 10. So I'm going to take a 1 meter step. And it's a tenth of a meter up in z .

And then I look at the brightness there. And now I again calculate the slope. And it's whatever the slope is, it allows me to calculate how much height I gain or lose in the next step. And so following that idea, I can actually get a profile of the surface. I can actually keep on going and measuring brightness, and computing the slope, and taking a small step.

So the idea is, I'm here. And I measure the brightness. And it gives me a slope. So now I can take a small step in that direction. If I go at right angles, I could fall off a cliff. I know nothing about the slope in the direction at right angles.

So now I'm at this point. I measure the brightness. It gives me a different slope. And I take another small step, and I'm at that point, and so on. So you can see how I can get a profile on the surface. Now of course, there are issues of accuracy, because we already mentioned that it's hard to measure brightness accurately.

And also, the surface may not be perfectly uniform. There may be some variation in reflecting properties and so on. But conceptually, I can do this. And by the way, I can go in the other direction as well. Because if at this point, the slope has a certain value, I can go in the other direction by minus the slope times the step size. And so I can continue this profile on this side.

Well that requires that I need some sort of initial condition. So I need to know z at my starting point so that I can incrementally change z . Do I know z ? No. When I measure brightness, brightness gives me information about orientation not about absolute depth. And remember the formula L equals π over 4 E , blah, blah, blah? z doesn't appear in there. That's very important. That when I walk towards the wall, it has the same brightness.

And we went through that argument, two changes that are proportional to r squared, which cancel each other. So z doesn't appear in there. And conversely that means that, I mean, in a way it's nice. It means that things don't-- you don't burn your eyes when you get too close to somebody-- well, most people. And so it's pretty much the same brightness. So that's the good part, that you can recognize things because their color doesn't change as you move around.

The bad part is you can't get distance that way. You can't invert that process to get distance. So we don't know the distance. So that's an important consideration here. I can get this profile if I have the initial value, but so what happens if I don't know the initial value? Well the profile might be up here. It'd be the same shape. So I can get the shape of the profile. I just can't get its absolute vertical position.

So that's pretty exciting, because it means that, for example, I can look at the moon other than at full moon when everything is independent of surface orientation, and I can run a profile like this. And it so happens that in the case of the-- we had the coordinate system up there. And not too surprisingly, the direction is going to be parallel to the equator. So of course, you typically do this on a very small scale.

But suppose I start here. Well, I can get a profile that way. Well, why don't I start somewhere else. Why don't I start here? Well I can get a profile there. Get a profile there. So you can see how I can explore the whole surface. I can get lots of profiles. And they may not be very accurate and whatever. But do not worry about that for the moment.

And there I've got the shape. And as I said, typically you'd be doing this in a small area like in some crater where you might think of landing or something. You wouldn't do it for the whole moon. But this gives the direction of the profiles. It's parallel to the equator. It's along lines of latitude.

So that's the good news. The bad news is, I don't know how these relate to each other. Because when I'm standing here, I have no idea what the slope is perpendicular to this profile. And the same with all the other profiles. So the good news is, I can get the profiles. The bad news is, they're all independent.

And now you can start imagining various heuristics like saying, oh, well, there aren't any gigantic cliffs. So typically, neighboring profiles will be similar. And maybe the average height along one profile should be the same as the average height along a profile next to it, and so on, and then stitch them together into a 3D surface-- attempt to stitch them together into a 3D surface.

But that's going to depend on prior information, like what are the topographic properties of the lunar surface. Because if you had a surface with those reflecting light properties, and these were all-- you could shift each of these independently in the vertical direction, you get the same image. So you don't know which of those it is.

So another idea is, if I have a crater that has some sort of rotational symmetry, not perfect perhaps, then I'm scanning across like this. Well if I'm lucky, this cross-section will be very similar to this central cross-section. So once I've got this horizontal cross-section, I can pretend that I know the vertical cross-section. And that will tie them all together. And of course, that makes an assumption about the symmetry of the crater and so on. Anyway, that's what people did.

Now for a moment, I want to have a complete change of topic. We'll get back to this later. It's the very beginnings of shape from shading. And this was the first shape from shading problem solved, because it's so easy. And at the time, there was a strong incentive to solve it.

So I want to get back a little bit to lenses. And there's a reason. Because we'll be switching to orthographic projection. And I'm going to try and justify that. So we talked about thin lenses. So a thin lens has the property that it has exactly the same projection as a pinhole perspective projection. And the advantage is that it actually gathers a certain amount of light.

Now real lenses aren't thin. And so, if you actually look at a catalog of fancy expensive high quality lenses, there will be all sorts of diagrams of many different elements. I don't know, just keep on going. Lots of individual elements symmetrically arranged around some optical axis. And so how do those work, and why do they do that?

Well, as I already mentioned, it's impossible to build a perfect lens. There will always be trade-offs between different kinds of aberrations. But by compounding, by adding different lenses, you can compensate. So for example, glass has a refractive index that varies with wavelength.

And so, that means that the focal length will depend on wavelength. And that means that red light will be brought into focus at a slightly different place from blue light. And so you get chromatic aberrations. You get fringes, color fringes around things.

Well you don't see that in your camera. That's because they have then put in a second lens of a different material that has a different wavelength dependence carefully designed to compensate for that. And depending on how fancy they are, it compensates exactly at two wavelengths. Or if you're more fuzzy at three wavelengths.

Anyway, so there is a need for compound lenses. And they then have different properties. But those properties can be approximated very well as follows. So I don't know if you remember, you should remember, that for the thin lens, we had this notion that the central ray was undeflected. So a ray coming into the center of the lens at an angle α would be emitted at an angle α .

Well the thick lens can be approximated this way, which is very similar. So these points are called nodal points. And anything coming into the front nodal point at a certain direction will leave the back nodal point in the same direction. The planes through those points are often called principal planes.

Now in the thin lens, the two nodal points are on top of each other, and the two principal planes are on top of each other. And not too surprisingly, the distance between them is called thickness. And usually the notation is T .

So that makes it actually quite simple to deal with thick lenses, because it doesn't change things a whole lot. I mean, it does mess a little bit with our lens formula, because now A and B are not measured relative to one place, but they're measured relative to those points. So it's-- and just how do you compound the lenses to create this effect? Well, that's not our job. The people at Zeiss and such know how to do that.

So why are we even talking about this? Well, because now there's a neat trick you can do. It turns out that T doesn't have to be positive. That is the nodal point, front nodal point can be actually behind the real nodal point. So who cares? Well if you make this pretty large, you can make a short telephoto lens.

So normally, a telephoto lens is one that has a long focal length obviously. And small field of view. And the lens with a long focal length means you need a tube, a long tube. Well if you make T negative, you can compress it. And you can get a significant reduction.

If you typically buy a telephoto lens from Nikon or Canon, and you look at how many millimeters its focal length is, if you actually go measure the lens, you'll find that in many cases the lens is shorter than the focal length. And so that's one trick to play with this. Then another one is to move one of these points far away off to infinity, in fact. So and this is used quite a bit in machine vision.

So why? Well there are a couple of reasons. One is that when we have perspective projection, the magnification changes with distance. So if you're, say, looking down at a conveyor belt and reading labels and whatever, or trying to make a precise measurement of some dimension, well the image size will depend on focal length and the distance to the object.

And if there's any variation in the distance to the object, that's problematic. Or say you're doing, I don't know, printed circuit board inspection, something like that. And you want to be insensitive to small changes in distance, well if you can get rid of perspective projection, that would be good.

So how do you do that? Well what you need is a very far distant center of projection. If you move the center of projection far away, then that effect of varying magnification with distance gets less and less. Because that cone of directions gets more and more parallel. And in fact, if you could move the nodal point to minus infinity, then there would be no change in magnification.

And it's amazing. But yeah, by building a compound lens, you can do that. So that's object space, stella centricity. And as I mentioned, a lot of machine vision systems commercially used in the industry use this. They're not cheap. Partly because they need a lot of glass. And the reason they need a lot of glass is that normally, a lens images a cone of the world.

A telecentric lens, because the center of projection is way back, actually images a cylinder. So if you think about it, here's the center of projection. Here's the lens. Now normally, you're imaging this whole area with a magnification that changes with distance that gets smaller and smaller-- the image of an object gets smaller and smaller the further the object is back.

Well now, imagine that you move this point way back there. Then this cone becomes shallower and shallower until eventually you have a cylindrical volume. So an object space telecentric lens will image a cylindrical volume. And that means that the lens has to be as big as the object, otherwise it won't be imaged.

And actually, it has to be a little bit bigger. So that means that if you're trying to read a circuit board or something, you may need a lens with a substantial bit of glass. And accurately made, so and that gets expensive. But it's done.

Now that's moving one of these nodes. Now you can actually move the other one as well. So let's keep this one in the same place, but move the other one, so image space. So we have the same kind of diagram with a cone of rays on the other side hitting the image. So we have-- here's the image plane, and here is our center of projection.

And we know a number of things. One of them is that if your image plane isn't in exactly the right place, the magnification changes. So if I move my image plane there, the magnification is different. Now in order to achieve a sharp image, I'm going to focus the lens, which ta-da, that means I'm either changing the focal length of the glass, which is not possible, or I'm moving the lens relative to the image plane, and therefore I'm changing the magnification. Maybe by a small amount, but if you're making accurate measurements, that's a drawback.

So that's one issue. And the other one is that cosine to the fourth law, we really don't want that. Now if we move this center of projection off to plus infinity, then this cone becomes more and more like a cylinder. And first of all, that means that as I move-- if the image plane is moved, if I got it in the wrong place, it doesn't change the size of the image.

It may make it more blurry. That's another issue. But and so, that's very useful in the metric situation, where you actually trying to measure something. So it turns out that-- what's the terminology for this-- there are lenses which are telecentric on both sides. Hmm. Double telecentric. Now that makes sense-- double telecentric.

So why does this cosine to the fourth go away? Well because that came from the inclination of the rays coming into the sensors. And so by moving the nodal point way out, now the radiation reaching a particular sensor is coming in perpendicular to the sensor.

And so that actually has other effects. So here's our little sensing element. And the radiation is coming in this way. And before it might come in at an angle. So particularly near the edge of the sensor, the light is coming from the center of the lens, and it's coming in at an angle. And so we get that effect.

Well there's another reason not to want that. And here's one, which is that often the sensor has right in front of it a set of little lenslets, which concentrate the incoming light into a smaller area. So they don't create your image or focus your image or something. What they're doing is they're taking the light that is covering a certain area and concentrating it into smaller area.

And this is very common. And why? Well because their circuitry. So the surface of the sensor isn't all sensitive to light. If we look down on it from above, it might look like this. And then there's a lot of switching circuitry and stuff around it.

So here's the area that's actually sensitive to light. And then there's this other thing. And there's something called the fill factor. So of course, there are many different designs. But that's something that happens in many designs.

So there are different issues. One of them is if you image without the lenses, you're throwing away light, not measuring it. What's worse, you have to protect the circuitry from the light, because light goes into the semiconductor, creates electron hole pairs, and oh, if that's right in the middle of your MOSFET, that's not such a good thing. So that means you have to put a metal layer on top of it.

So that's one reason why people add the little lenslet arrays. The other reason that is aliasing. So those of you may remember in 6003 or some other equivalent signal and system course, that when you sample discretely, you have to be sure that there aren't high frequency components in the signal.

And in our case, we could have sharp transitions in brightness from one area to another. And those will create effects, where some high-frequency component is it looks like a low-frequency component. Maybe alias down to a lower frequency. And how do you avoid that? Well you low-pass filter first.

And then there's a wonderful theorem that says, if you have a low-pass signal, you only need to sample it twice the bandwidth of the signal, and you can reconstruct it perfectly. And what's the relevance here. Well, if we are not-- if we're measuring over the whole area, we're performing a crude form of low-pass filtering, we're block averaging, which isn't-- you know the real low-pass filter is a sinc function, but we can't build that.

So if we have the large pixel, we get a certain amount of low-pass filtering that's advantageous. And fancy cameras have additional mechanisms for this. But if we have these smaller areas, it's more like point sampling. It's more like we didn't low-pass filter, we're just sampling. And that has very bad aliasing effects.

So by using this lens array, we're actually using the light from that whole area and measuring it, because we're projecting it onto the sensitive area. And so we reduce the aliasing problem. Now this works great if light is coming in more or less perpendicular to the surface.

It's not so good if the light is coming in at an angle. Because then you'd have to somehow change the scale of the lenslet array. And even then, you can't make it work correctly, because there's a spread. The lens has areas that are in different directions. So anyway, cut a long story short, there are several reasons why people like light to come in perpendicular to the sensor, starting with the cosine to the fourth.

And so, in high quality digital SLRs, the lenses tend to be image space telecentric, or at least partially. I mean, they don't actually move the center of projection all the way to infinity, but they move it far enough out that that cosine to the fourth becomes negligible, and we have those effects.

So that's telecentric lenses. And these used to not be available. And it took a while for people to figure out how to design them. But now they are all the rage. So double telecentric.

So where are we going with this? Orthographic projection. So we said that we no longer have a dependence on distance. That in the object space telecentric device, an object of a certain size will be imaged the same size independent of its distance. The sharpness of the image will change, just as in a normal lens. But the size will be the same. And similarly for the image space, telecentric.

So what we're really doing is taking perspective projection equation and making the focal length huge, so that our center of projection is far away. And in effect, we can then pretend that we're dealing with orthographic projection rather than a perspective projection.

So we saw that perspective projection was quite useful in a way. So this is where we started. Because we had this dependence on depth, and particularly in terms of motion that was helpful. But now, suppose that the changes in depth in the scene are much smaller than the depth itself. Then we can write x is F over z_0 , times x .

So if z is approximately-- so this is another way to get to orthographic projection. We can make z pretty much constant. And how can we do that? Well one way is to add a very large number to z . And that's essentially what happens when we move the center of projection. We're adding a very large number to z .

And so, some small variations in z aren't going to make any difference. The projection is pretty much independent of the position. And so, we have a linear relationship between x and y in the world and x and y in the image. And amongst other things, this means we can measure distances, sizes of objects, independent of how far away they are.

Now in many cases, it's convenient to just pretend that that scale factor is 1. And often we'll just use that version of it.

Hmm. I guess it's 12:30. So that's where we're going to go. Orthographic projection is useful in practice with telecentric lenses. And it's also going to be greatly simplifying some of the problems we're going to work on. That's not-- so this is a little bit like the Lambertian thing. A lot of people say, oh, these methods only work for Lambertian. No they don't. They work for everything. It's just that for anything but Lambertian, the equations get messy.

And it's the same thing here. The kind of reconstruction we're going to address next can be done under perspective projection. It's just complicated and not very insightful. If the math gets very complicated, you lose track of what you're doing. When we change to orthographic projection, it turns out that many of these problems become quite clear. So there'll be a new homework problem as usual on Thursday. So.