Virtual Reality
Prof. Steve Lavalle
Department of Multidisciplinary
Indian Institute of Technology, Madras

Lecture - 15-1
Visual Rendering (shading models)

(Refer Slide Time: 00:16)

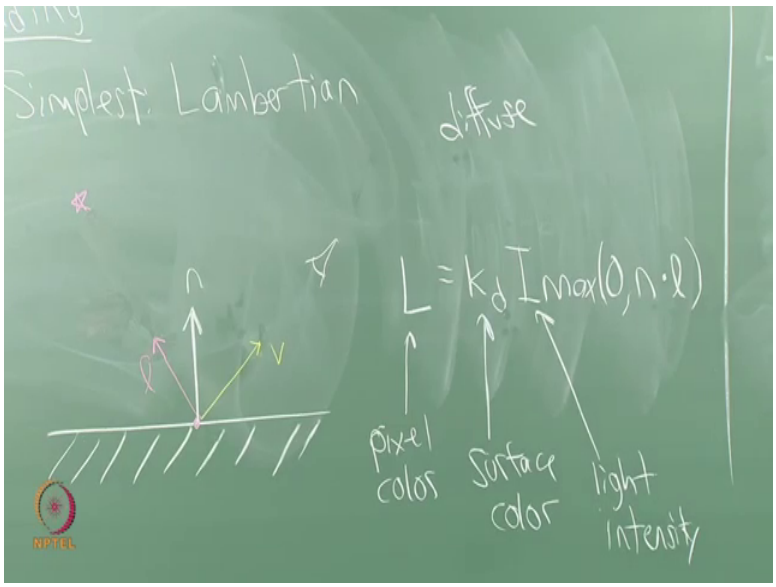

So, under shading with simplest case is called Lambertian shading, which looks like the following. Let me draw the surface. So, this would correspond to the part of the triangle that was struck by the visibility segment and we have a light source. Let me draw, let us see a light source up here and we have, let us see fixate on this particular point. So, we have one direction vector corresponding to the direction to the light source. I will call this, L for light.

We have the normal to the surface n. I will keep everything normalized. So, let me. In fact, shorten this l vector here. So, we will deal with unit vectors. So, I do not want to make it too vector that goes all the way to the light source. So, let us just put it down like this and. So, normalize that, that should appear to be the same length, I guess, I did not quite get it right. And then there is a direction that you are looking at the surface from. So, this would be, the direction of this ray as it hits, the surface over here. So, the direction of where is the eye? Let us say the virtual eye that is looking at the surface.

So, I will draw and a virtual eye here, looking at the surface. So, that goes off in this

direction again. I want it to be normalized, I will put, v label here for viewing direction and now to determine the color of a pixel that corresponds to this situation. So, L is meant to be pixel color, it is equal to some constant for, for this model, which is called a surface color and perfectly fine for this to be a vector, for example, it could correspond to R G B values that capture, the surface color or it could be H S V or some other representation. So, it could be a vector of parameters here and then there is another constant that is the intensity light, intensity for the source and then the max of 0 comma n dot l we do not want negative.

Numbers there. So, if n and l have an angle between them greater than 90 degrees, then this whole thing just collapses to 0 rather than trying to subtract light away in some strange way. So, this is the way the model works and. So, one interesting thing about this is that I did not use v. So, in the Lambertian model, the particular direction that you are viewing the surface at, does not matter, all that matters is the direction of the incident light and how that varies with respect to the normal right. So, the surface is most illuminated when the light source is perpendicular to it and when it is off close to tangential then that particular patch is less illuminated right.

So, you can sort of imagine, it is like having a giant dish in the rain and you are trying to collect as much raindrops as you as you as you can and. So, if you oriented perpendicular to the rain of something that rain is coming straight down, you get the maximum amount of collection and has the angle varies towards being 90 degrees off you collect very little right. So, that is the Lambertian shading model, it is also often referred to as diffuse shading. So, this is what you might call a flat or a matte surface very much like a carpet floor for example. So, this is the least amount of what we would perceive to be shininess in the surface. So, it is the opposite, it is a very flat surface, it is nice called diffuse. That is why there is a d here, for the, for the constant case of d questions about that right.

So, we define a light source, it has, it is into it is own intensity. I am assuming a white light source here. So, it has the entire color spectrum and then the reflected back colors come from the surface color here and. So, based on that the overall amount of intensity you get depends after you get these constants out of the way, just on the angle that the light source is coming in. So, we would like to now, make a model that depends on the particular direction that you are viewing the surface. This will be the next level of complexity. So, this is the diffuse case does not really matter, which angle you are

looking at the surface from, we would like to make another one that takes into account shininess of a surface, as well start to get interesting and complex for, for virtual reality, head mounted displays.

(Refer Slide Time: 06:05)

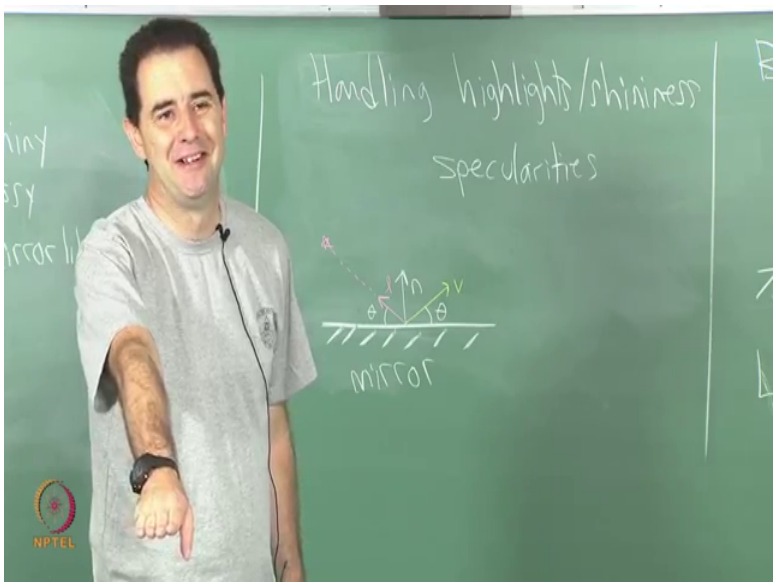

So, handling what are called highlights or shininess well here is a very simple case, what if I am looking at a mirror? What should happen? Let me draw that a little more level here. So, I have the surface normal, I have a light source. So, just like I had before. So, L for light source and this has some, let us say angle theta. So, where should I be looking over here. So, that the light source appears directly on my retina right. what angle should i be looking in at, if this were a perfect mirror, this is a mirror. So, it should be the case that the angle that I am looking in.

Here is theta with respect to the surface right, because for a perfect reflection, the angle of reflection is equal to the angle of incidence of the light. So, it should be a perfect reflection model. So, this should come back off like this, with these two angles. These are both theta and that is the viewing direction ? One thing to also pay attention to which is going to be important.

In the next model that I give is, you may note that it does not quite look perfect here, the bisector of this angle between that, that is formed by v and l is the normal to the surface right. So, notice that the normal is exactly should be exactly at the midpoint. I think my normal, the drawings, a little bit slanted, but if it were straight up, it should look exactly

like the bisector of this angle. So, if you want to make surfaces that are shiny and say there is somewhat mirror like, but not exactly a mirror like.

For example, this particular case here, has some shininess to it, but it does not look like a mirror. I do not see my face inside, when I looking at it says, shininess to it, then I want to have an increased amount of light reflected as we get as our viewing direction gets very close to this situation, this mirror situation, but it might not be perfectly lined up with the angle of incidence being the angle of reflection, but as we get close to it. We would like there to be a greater amount of light reflected have it be an intense sort of difference. Let us say. So, this intuition or this idea is called the Blin

(Refer Slide Time: 09:03)

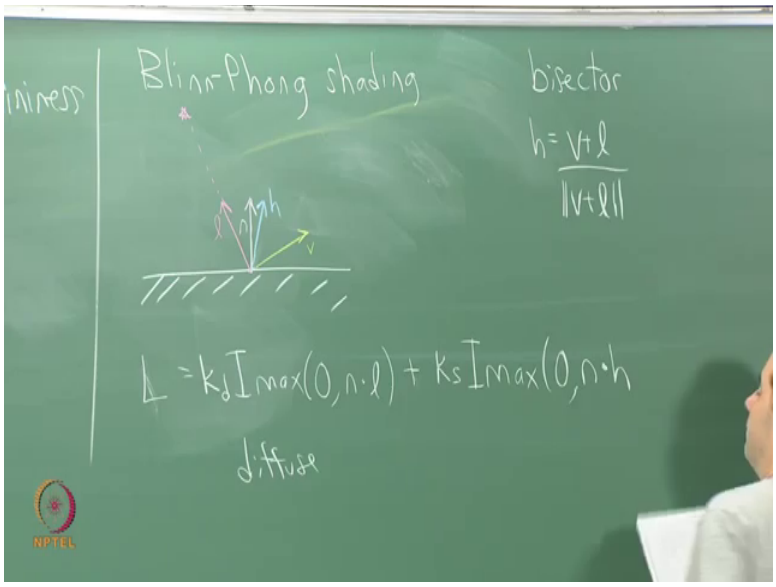

Phong shading model, which is very widely used in computer graphics and also shows up in G P U's Graphical Processing Units; so, in this particular case, let me draw my picture again, we have the light source. So, this is the l vector, we have the normal and we have a viewing vector V and now, I want to take the bisector between l and v, which I've tried to draw this picture. So, that it will look like the bisector is not the normal, but off a little bit, let us see here. So, let us draw it off that is gonna be tough here, to do. So, let us say this bisector is h.

So, I can calculate this bisector as h equals v plus l and normalize it. So, that gives me this bisector vector and now, if I want to take into account, these kind of highlights or shininess, which is also called specularities. I just introduce another term in addition to

having this diffuse component, I make what is called a specular component.

And it will compare the difference between n and h and the closer one gets, the closer h is to n, the more amount of light that is reflected. So, if I do this, I get the following formula. So, the total amount of light again, this can be decomposed into r g and b values, if that is what you are using for your color representation, I still have this diffuse component. So, this is defuse, I am just copying what is over there and now, I add this new part, I will have k s for specular light intensity again and I have max again, I just want to avoid the negative case.

So, I put max 0 comma n dot h. So, instead of subtracting the 2, I am just using the dot product to compare them. As they become very far apart, when they become orthogonal, this goes to 0 and if they go further away than orthogonal or I have to go the other way. So, that the angle between them is greater than 90 degrees, this max keeps this term to be 0 anyway So, when they are very close, these are all normalized. So, this should become 1, when they become n equals h.

Now, it is not exactly the formula that is used yet, because people want to control exactly the amount of drop-off as the distance between n and h increases. So, they want to control the rate of drop off and. So, what is done is an extra exponent parameter is put in here, if you are obsessed about physical realism, this part may look disgusting to you all right. So, it is just an extra parameter, it gives you some control, you tune it. So, that it looks like the way, you want the pictures to look and people are happy with it and there is been this way since the 1970s.

So,, you can do more physically realistic rendering and I will talk a little bit about that, but this was a kind of hack that continues to get used and works very well for most people's purposes I thought it is great to, quote a particular table that is in the, the surely a tall textbook, which gives different values of p and their interpretations. So, if So, for different values of p, we do not want to run out of space here. Let me go over to this side.

(Refer Slide Time: 13:55)

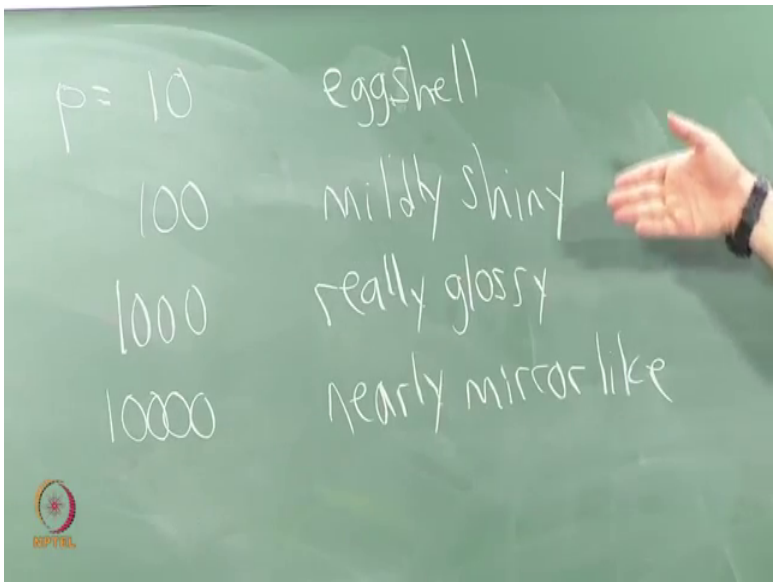

So, if p equals 10, you get what is called eggshell (Refer time: 14:00). This may also remind you of, if you get photographs developed, people do not do that as much anymore, but if you get photographs develop, there is different choices about the shininess of the cover. You can get a matte surface, which corresponds to diffuse, you can get eggshell right, which is the little bit of shininess to it 100 is named again, this is in the textbook, but I found this amusing to give some intuition, mildly shiny a 1000 is called really glossy and 10000 thousand is nearly mirror like.

So, I find this somewhat amusing these are; obviously, not very technical terms and it gives you the idea that, this p parameter is a kind of hack, but it, let us, you have a knob that you can adjust to control the amount of shininess of the material and in one extreme in the limit as p goes to infinity, this formula should boil down to saying that the reflection looks exactly like a mirror and as it gets smaller then you might as well just make a diffuse term only. Any questions about that?

Student: Sir.

Yes.

Student: In the two models, we do not consider the color of the light source.

That is correct, that is correct. I would have to, I would have to then So, I am assuming that the light source is white. So, we will have the full spectrum, if I want to put colored

lights in, then I am gonna have to in these coefficients here, consider the combination of them right. I am gonna have to look at this, you might remember in the lecture, where we talked about, lens aberrations. I talked about chromatic aberration and I talked about the frequency spectrum of, the light and the frequency spectrum of, the reflectance due to a particular material.

So, you will have to take both of those into account to find the resulting color, that is right. So, if you shine a, a blue light on a, on a, on a, on a surface that only reflects red then it will look black right and. So, that is right. That needs to be taken into account. So, I am covering the easy case, where you have a full uniform spectral representation. Let us say the full energy spectrum or power spectrum is, is represented by the white light. other questions. So, far if we just use this model, what happens to objects that have their surface normals pointing away from the source of the light. How will they be illuminated, what will they look like all right.

So, I have one source of light, let us say I have some bright white light and I have an object and I want to know, what is the backside gonna look like, what should it look like, if I just start doing these models, I do, I do both diffuse and specular did not write it here, specular. What will these two terms give ok. what will this one give all right. So, the light source is back here, now. So, the surface is up the light source is back here. So, what should this give?

Student: 0.

0, good and what does this give? 0, so; that means, that the object should look black right. There is no light being reflected back does that correspond to the real world look like, we did some reasonable stuff here right. It does not correspond to the real world, does it and I can look, I look up under the desk here. I do not see any light sources well ok. Maybe there is a little bit there, but there is not many and I can certainly go to places, where if I were to put my eye there, I could not see any light sources directly, yet there is still some light there. Why is that is there inter-reflections right. So, there is light propagating throughout this room. I have not taken that into account at all here.

So, So, one thing I can do after I go through and I make a, a some kind of virtual world and I consider a light source, I make diffuse component, a specular component. I light up my objects or model one thing by the way that comes up very quickly here and I talked

about this in the beginning, when I talked about making geometric models, I talked about their coherency and getting all the surface normals in the right direction. If one of your triangles has a surface normal that is accidentally pointing in the wrong direction. It will just show up and look black all of a sudden. So, you might start having black triangles on a surface, if you have done computer graphics projects before, you may have seen this happening right from imperfections in the way of assigned the surface normals.

So, but even if you have done it correctly, the backsides of objects may become black, because of the reason that I just mentioned, you are not taking into account, any into reflections, anything else like that. Well here is one simple hack, you could add, an ambient light component, which basically, just adds to, l some extra term. So, add, let us say case of a, for ambient, I sub ambient to l. So, I just put it over here. For example, if I want to combine everything together, I just get some extra component here, k a i a, some ambient component.

So, I just can add an extra term to, every one of my materials and I still can make the color of the material be dependent here, with this case of a, but I have some ambient intensity, i sub a. So, that I cause everything in the environment to glow a little bit does that, seem physically reasonable, probably not I think it, only accomplishes the goal of making everything glow a little bit. So, that nothing turns completely black. So, this is yet another hack.

But nevertheless, it is commonly used and if it looks reasonable, people do it, it is almost like the interpretation is that all of my triangles are glowing a little bit on their own If you have multiple light sources, then you just extend this by just adding up the diffuse and specular components for each one of your light sources. So, you just do a linear combination, it is a simple principle of superposition, you just add up the contributions from all of the light sources.

So, if I have seventeen sources of light in the environment. I have 17 diffused terms, 17 specular terms and I still may want to add this ambient term, at the end all right. So, that is the easy way to handle more light sources, if you are not very happy with these kind of strange hacks of taking into account specularities and this, ambient part here, for illuminating surfaces that have their normals pointing away from all the light sources, then you may want to solve the problem of what is called global illumination.

(Refer Slide Time: 21:38)

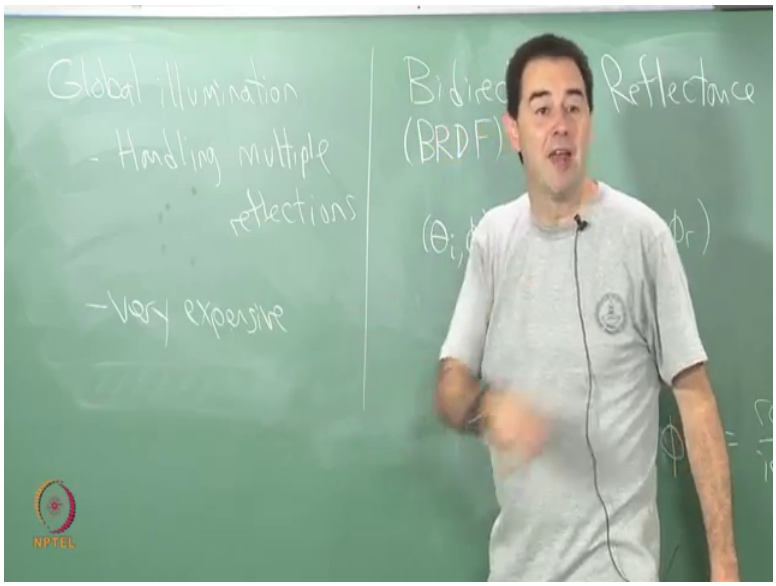

In this case, you are handling, multiple reflections or inter reflections back and forth between the objects. So, light will bounce back and forth, how far do we have to go right, if you are in a hall of mirrors then there may be quiet, you may have to go out, you may have to consider a hundred bounces or more perhaps right of the array of light, if you are in an environment, such as this one we are. There are some shiny surfaces, but, it may be a fair assumption that after three or four reflections, we are done.

We are not going to consider how these rays are propagating anymore after that. So, taking into account multiple reflections in order to solve the global illumination problem, we will fix this simple model. So, that it is not a hack anymore, but it becomes extremely expensive computationally. So, it is very expensive, but solving the full global illumination problem is the way to do the most photorealistic rendering.

That is if you wanted to have the best rendering possible, take a look at global illumination techniques, and the problems go on and on after this. I mean also in considering global illumination, what if I have windows right to have transparency or um? So, I or I may have objects that are translucent in some way, they are just I can see the light coming through, but it is not transparent enough. So, that I do not see an image through it or maybe I just see, some kind of shadows on the surface of it.

So, it can get more and more complex as you consider different kinds of materials, especially light propagating. All the way through materials of course, I guess, I should be

able to make lenses and optics inside of my virtual world and make all of that work correctly a lot of interesting questions of whether or not you need to do that in video games. For example, if you have a game that involves shooting and you are looking through the scope, do you have optics?

I need to be taken into account there or not and if. So, is it better to do that or is it worse in terms of discomfort or simulator sickness? So, there are interesting questions, I hope to come back to those a little bit later, but So, the problems go on and on depending on the types of materials, if you have materials that act like lenses and glass things like this, it becomes even more complex. In the most general case, if we are considering global illumination and we do not like doing this kind of strange separation into these components, we can model the reflectance properties of a surface using, what is called a BRDF. I just want to give you the general formulation here just; so, you know what people are doing, when they care more about physical realism.

(Refer Slide Time: 24:23)

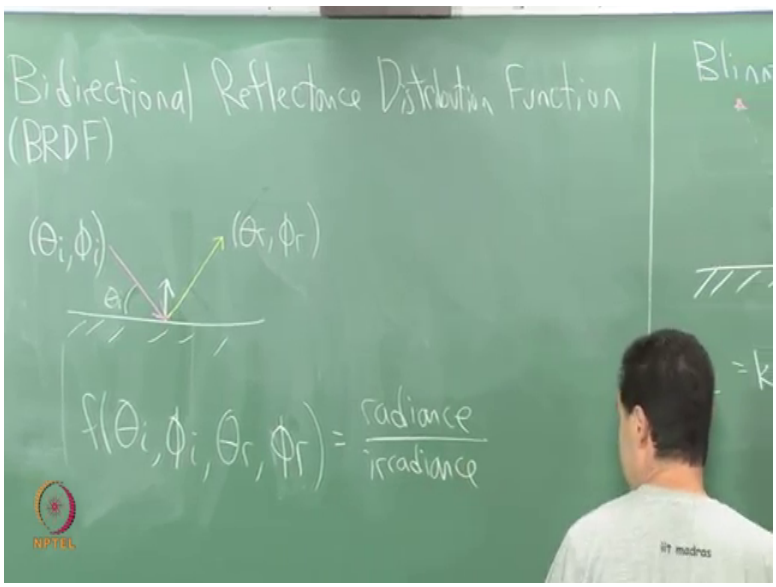

So, this is called the bidirectional reflectance distribution function or as I said B R D F. In this particular case, we have the surface, we have, it is normal. We once again, have the light source and we have the viewing direction, draw it the arrow coming in the other way here to emphasize something in this particular model. So, we have the direction that light is coming in and we have the, the direction from which we are interested in how it reflects and.

So, for the general B R D F model not only do we look at the angles, in this two-dimensional picture, the angle of incidence and this angle of reflection here, but we also consider the direction that cannot be seen very well, which would be if I drew a three-dimensional picture, it would be going around in this direction right. So, coming in laterally, protruding outward or going into the board. So, that is another angle to consider.

So, I can draw two angles here, for the incident part or incoming, which is theta i and phi i. Let us say right. So, theta can correspond exactly to this part and then phi i is what I would be angled, I would get if I am looking straight down at this picture and it is the angle coming in. And then the outgoing part also has two parameters, the reflected part theta r and phi r and then I can make a very general function of 4 variables. Let us say f of theta i phi i theta r phi r, which is going to loosely be defined as and I would not go into the exact definitions here, I just want to give you the general idea of it, it is the radiance divided by the irradiance, which in other words means the, the upper part is the reflected energy.

(Refer Slide Time: 27:34)

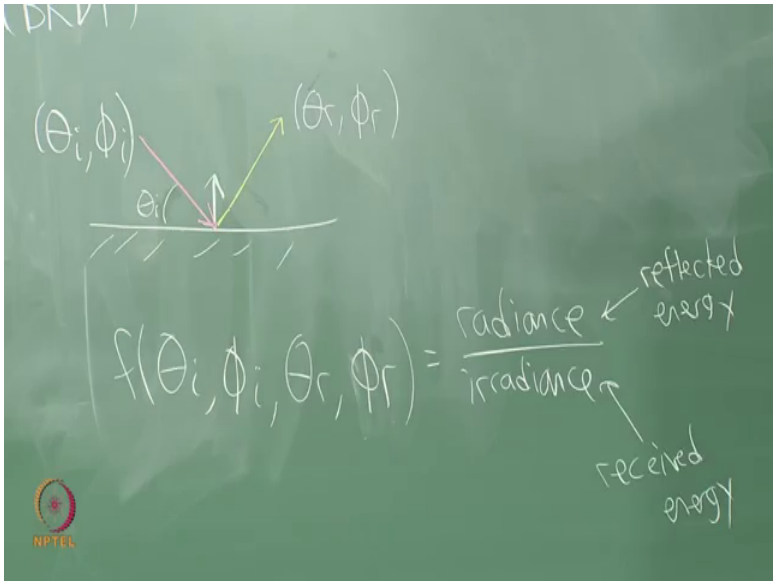

And the lower part is the received or incident energy. So, this really helps us to nail down exactly, what happens from every possible direction. Now, if we wanted to make this model even more precise, what else could we add to this? We could also make it. So, that, if this is some kind of rough surface, then we might even have two more parameters

that correspond to, where exactly this point is along the surface. So, we could add position parameters, what else might we add it could be the case that this reflectance depends on the wavelength right, very natural, we certainly have seen this with lenses right.

The, the in this case, in the case of lenses, we were talking about the transmission part or propagation through the lens, but in general, the reflectance can certainly depend on frequency. So, I could add three more parameters easily here I could add, let us say x y position along the surface and then I could add a, a lambda wavelength parameter I may also make simplifications it may be the case that the, that all I care about is the difference between a phi sub I and phi.

Sub r and there may be a kind of rotational symmetry with respect to that. So, there is some simplifications, you could make there is some additional complications, but I just want to say that in general people do make models. I take into account the complete reflectance and, and parameterize every possible aspect or. So, it seems here and take this into account. So, that is the other alternative and of course, this gets again very complicated, if I want to use a bidirectional reflectance distribution function and use that for global illumination and take into account multiple reflections and also let us handle non white, light sources. Let us put it all together in the pot, then you end up with very complex calculations involving effectively.

The propagation of these light waves throughout this continuous medium and considering all of the reflections that go back and forth over effectively great distances right. So, you do not, you no longer get these kind of simple combinatorial expressions. Let us say where you just simply do very simple visibility calculations and then, you know adjust a few parameters and, and, and render a scene. You have to do, very complex kinds of operations that correspond to let us say numerical high resolution propagations of essentially wave fronts propagating through the environment questions about that.