Radiometry and Reflectance

Shree K. Nayar

Monograph: FPCV-3-1

Module: Reconstruction I

Series: First Principles of Computer Vision

Computer Science, Columbia University

February, 2025

FPCV Channel

FPCV Website

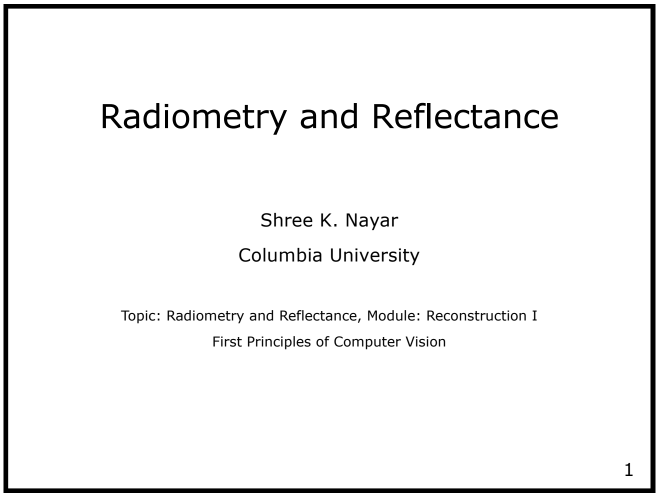

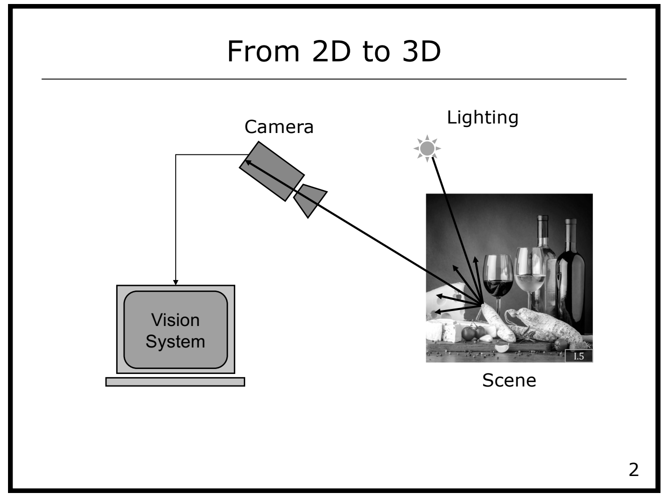

Shown on the right is a typical vision system, which consists of a three-dimensional scene and some form of lighting. While we have shown a simple point light source here, in general, we can have multiple point light sources, and even area sources, such as the sky. Each scene point receives light from the sources and reflects, or scatters, it in many different directions. The light scattered by the scene point in the direction of the camera is projected by the lens onto the image plane, where we have an image intensity value. Our goal is to explore what this measured intensity value reveals about the corresponding point in the scene. Before we can tackle this problem, we need to understand two things. First, we need to understand radiometric concepts related to the brightness of a light source, the illumination of a surface, and the brightness of a surface. Second, we need to understand reflectance, which is the ability of a surface, or a material, to take light from one direction and reflect it in another direction.

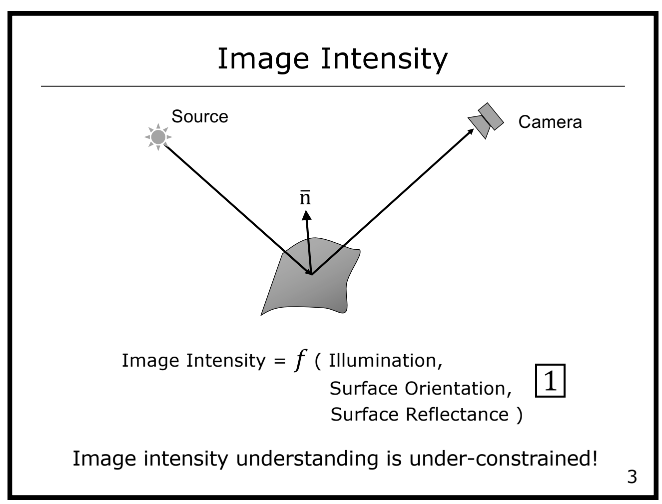

Let us look at the problem of image intensity understanding in greater detail. Shown here is a point on a surface that has an orientation given by the normal vector $\bar{n}$. In this case, we have a simple point light source illuminating the surface point. The surface point reflects some of the incident light in the direction of the camera. What factors influence the image intensity value of the surface point? The first is illumination. In general, the illumination could be arbitrarily complex and unknown—it could include multiple point sources and/or extended (area) sources. Whatever the case may be, we can imagine that we need at least a few parameters to define the illumination itself, and in general, these parameters are unknown. We also have the orientation of the surface at the point of interest, which includes an additional couple of unknown parameter. Finally, we have the ability of the surface to take light from one direction and reflect it in the direction of the camera. That is given by the reflectance of the material that the surface is made of, and that could have several parameters which are also unknown. In short, while we have one measurement, the image intensity of the scene point, we have many unknowns $\boxed{1}$. Therefore, image intensity understanding is a severely under-constrained problem. Fortunately, under certain settings and with certain assumptions, we can say quite a bit about the scene point from even a single measured image intensity.

In this lecture, we will start with radiometric concepts. These include the radiant intensity (brightness) of a source, the irradiance (illumination) of a surface, and the radiance (brightness) of the surface. Next, we will derive the relationship between the brightness of a scene point (surface radiance) and the intensity it produces in the image (image irradiance). This relationship is a fundamental one in computer vision.

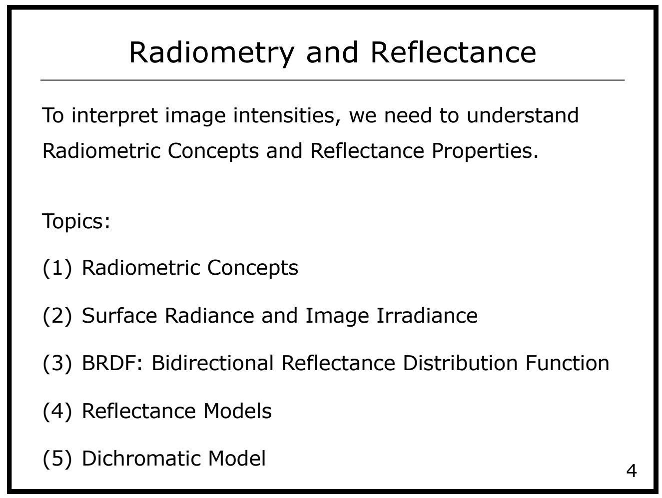

Next, we will describe the bidirectional reflectance distribution function (BRDF), which completely describes the ability of a surface to take light from any direction and reflect it in any another direction. We will present a few reflectance models that are widely used in computer vision and computer graphics. These include materials that are diffuse, or matte, in appearance, as well as surfaces that are mirror-like, or shiny, in appearance. We will also look at materials that exhibit both these phenomena at the same time. Finally, we will present the dichromatic reflectance model which describes how the appearance of a surface is impacted by the color of the illumination. We will also see how the dichromatic model can be used to take a single image and separate it into its reflection components.

Now, let us look at some radiometric concepts that are particularly useful for understanding image intensities.

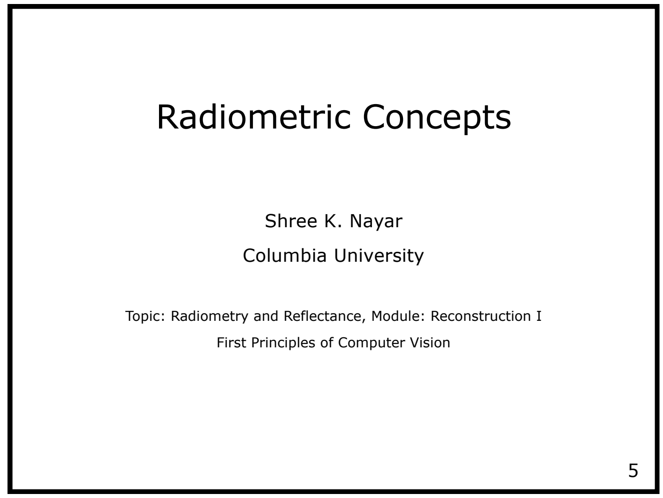

Let us begin with the definition of an angle in 2D. Shown here is a circle of radius r. The angle $d\theta$ subtended by the arc of the circle is the arc length dl divided by the radius r. The unit of $d\theta$ is radians, which is a dimensionless quantity because it is distance divided by distance.

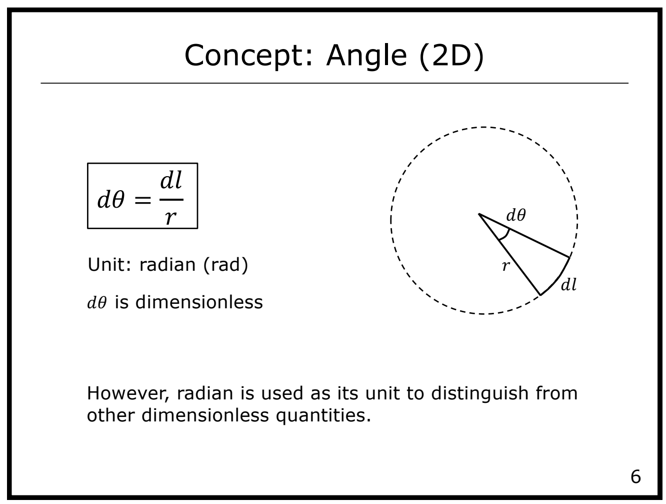

Now we can talk about an angle in 3D. In this case, we want to find the solid angle $d\omega$ subtended by the infinitesimal area dA from the point P. dA is at a distance r from P, and tilted at an angle θ with respect to it. We first compute the foreshortened area $dA' = dA\cos\theta$. The solid angle $d\omega$ is then dA' divided by r^2. The unit of $d\omega$ is again dimensionless and is called steradian. Using the above expression we can find the solid angle subtended by a hemisphere to be 2π, and by a sphere to be 4π.

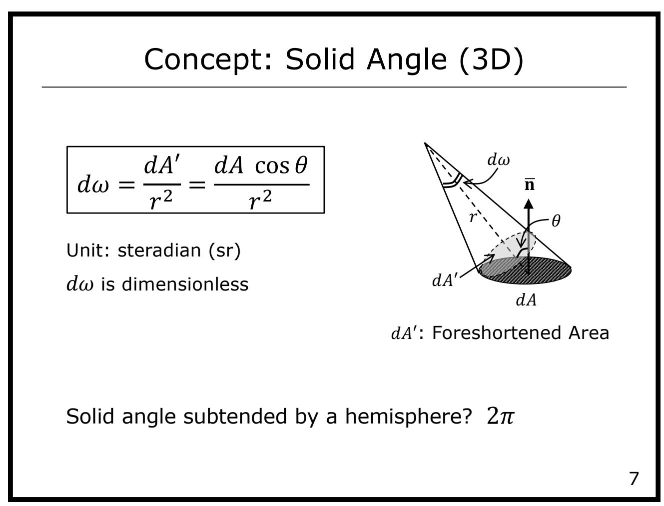

Now we can look at the radiometric concepts that are of interest to us. Let us start with light flux. Consider the point light source shown here. The light flux is the power emitted by the source within a solid angle. If you consider the solid angle $d\omega$ subtended by the surface patch dA, then we can say that the flux that is emitted within this solid angle and received by the surface patch is $d\Phi$. The unit of light flux is watts.

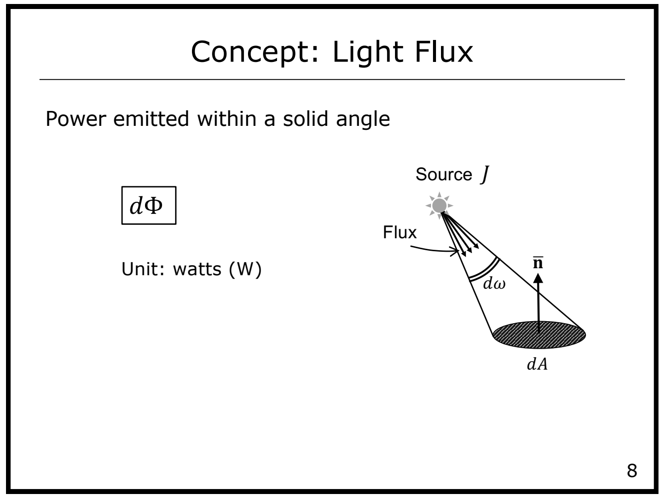

We can now define the brightness of a point light source, which is called its radiant intensity. The radiant intensity J is the flux emitted by the source per unit solid angle—it is $d\Phi$ divided by $d\omega$, and its unit is watts per steradian.

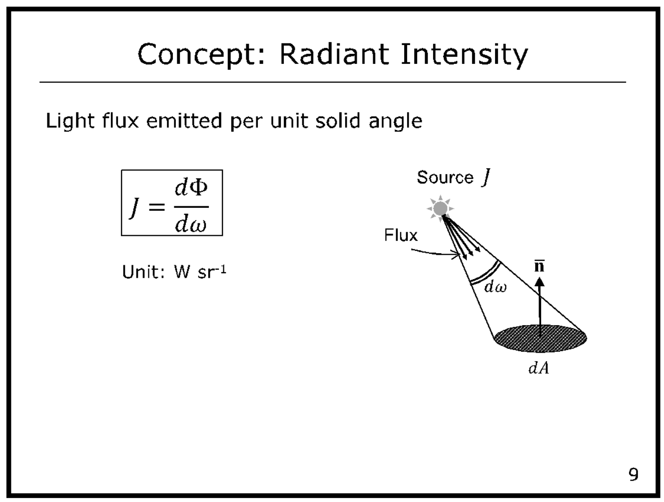

Let us now define the illumination of a surface, which is called surface irradiance. Surface irradiance E is the flux falling on the surface per unit surface area, that is, $d\Phi$ divided by dA. The unit of E is watts per meter squared. Thus, if we have a source with radiant intensity J, we can use the above definitions and find the surface irradiance E to be $J cos\,\theta$ divided by r^2. This expression implies that the irradiance of a surface is proportional to the radiant intensity of the source. It is inversely proportional to the square of the distance r between the source and the surface.

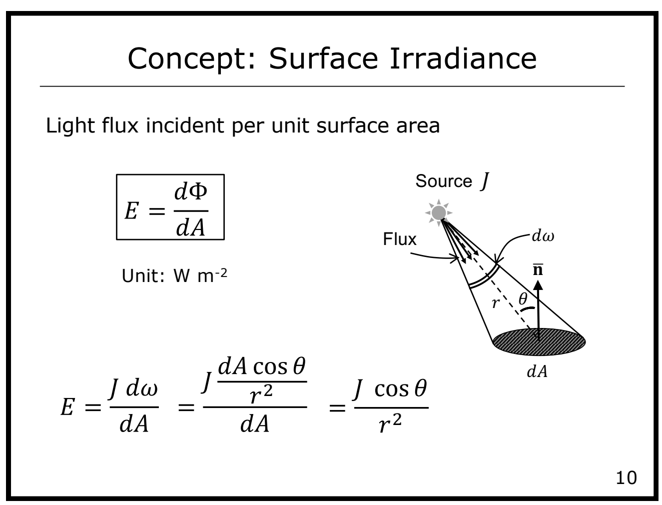

This is called the 1 over r^2 fall-off. Surface irradiance is also proportional to $cos\,\theta$, where θ is the tilt of the surface with respect to the direction of the light source. When the source is right above the surface (0^o) the irradiance is maximum, and it falls to zero when the source is at a grazing angle (90^o).

The brightness of a surface is called surface radiance. This is a bit more of a complicated concept, compared to the previous ones. How would you measure the brightness of a surface? Imagine that we have a surface patch, as shown here, and we want to measure its brightness using a sensor. Let us say the sensor has some area, shown as a disc in the figure $\boxed{1}$. We can see that if we move the sensor back, it is going to receive less light from the surface patch because the solid angle it subtends with respect to each point on the patch decreases. The second factor that influences this measurement is that if we increase the area of the patch, more light will enter the sensor. Therefore, the measured brightness is going to increase. So, when we define the brightness of a surface, we need to normalize with respect to both the area of the patch and the solid angle subtended by each point on the patch with respect to the sensor. Therefore, the radiance L of the surface is defined as the flux received by the sensor per unit foreshortened area of the surface, per unit solid angle subtended by the sensor. We use the foreshortened area in this case because we want to account for the fact that the area seen by the sensor is the foreshortened area. Note that radiance depends on the direction from which we look at the surface, which is given by θ_r. Finally, the radiance depends on the reflectance properties of the surface, that is, its material properties—the ability of the surface to take light from the source and reflect it in the direction of the sensor.

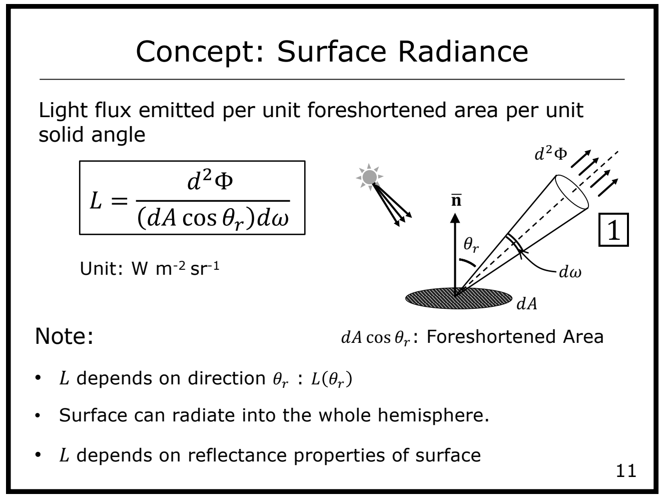

Now we will establish a relationship between the brightness of a point in the scene (scene radiance) and its brightness in the image (image irradiance).

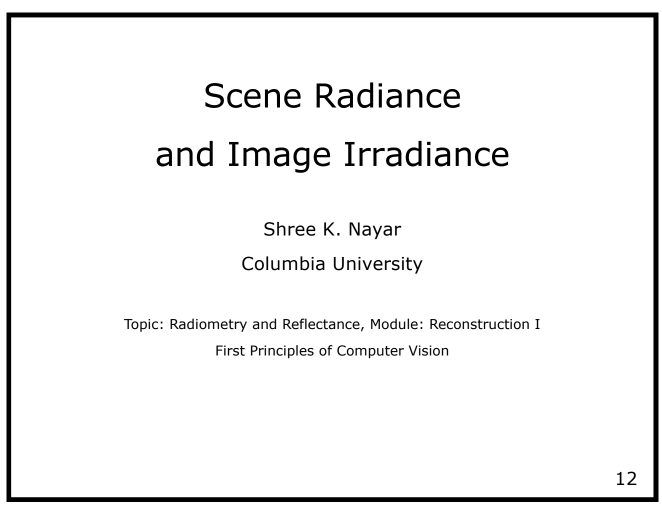

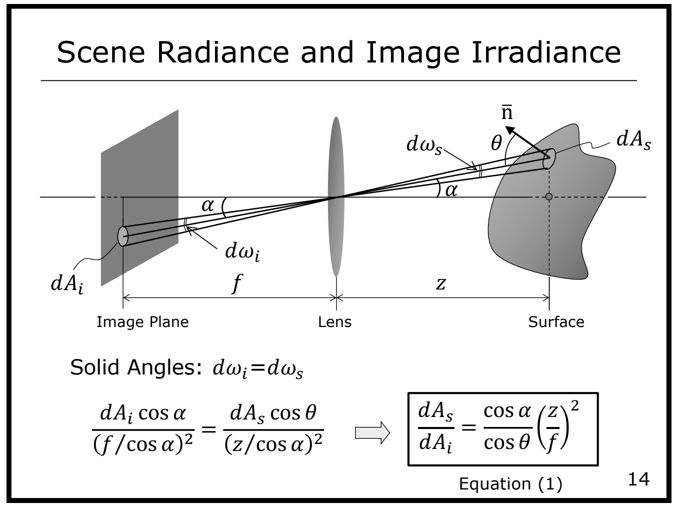

Here we have our lens camera used to image a surface $\boxed{1}$. We will assume that the distance between the lens and the image plane is f. It is important to note that this is not the focal length of the lens, but the effective focal length of the imaging system. Let's assume we are looking at the intensity captured by a single pixel of area dA_i. This pixel is observing a scene patch with area dA_s, orientation $\overline{n}$, and the depth of the patch from the lens is z.

We are going to derive a set of equations that we will use to find the relationship between scene radiance and image irradiance. The first equation relates the solid angle $d\omega_s$ subtended by the surface patch with respect to the lens, and the solid angle $d\omega_i$ subtended by the pixel with respect to the lens. We can see from the figure that these two solid angles are equal. $d\omega_i$ is equal to the foreshortened area of the pixel as seen from the center of the lens, $dA_i cos\ \alpha$, divided by the square of the distance from the lens to the pixel, which is $f\ /\ cos\ \alpha$. Similarly, $d\omega_s$ is equal to $dA_s cos\ \theta$ divided by the square of the distance from the lens to the patch, which is $z\ /\ cos\ \alpha$. The end result is equation (1).

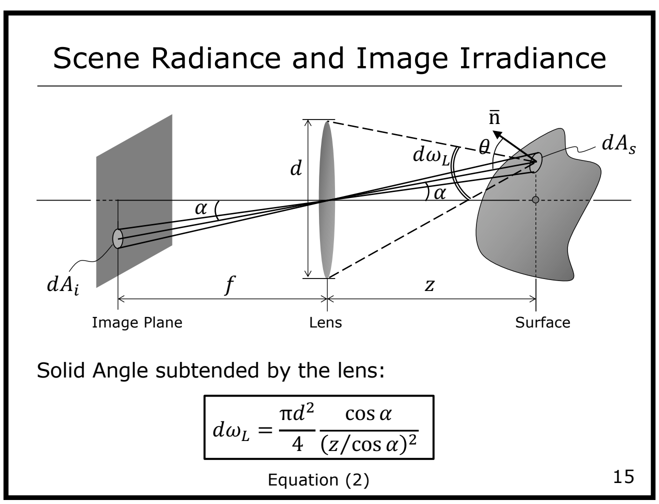

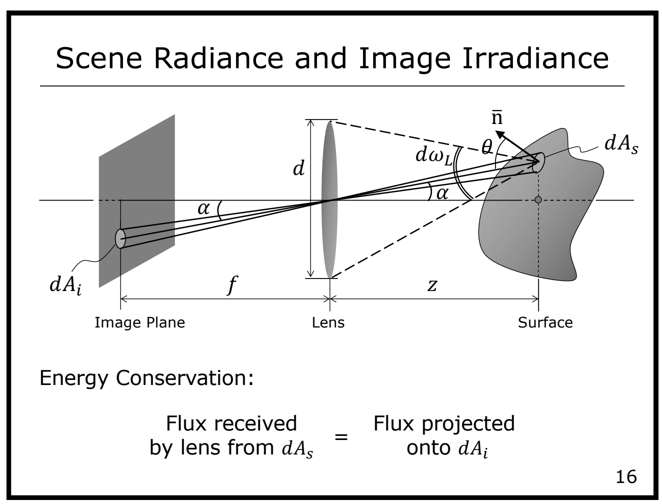

Next, consider the solid angle $d\omega_L$ subtended by the lens from the surface patch. If the lens has a diameter d, $d\omega_L$ is the foreshortened area of the lens, $\frac{\pi d^2}{4} cos\ \alpha$, divided by the square of the distance $z\ /cos\ \alpha$ of the lens from the patch. This gives us equation (2).

We know that all the light flux received by the lens from the scene patch of area dA_s, is projected by the lens onto the pixel of area dA_i.

The flux received by the lens can be determined from the radiance L of the scene patch, using the definition of surface radiance. That gives us equation (3).

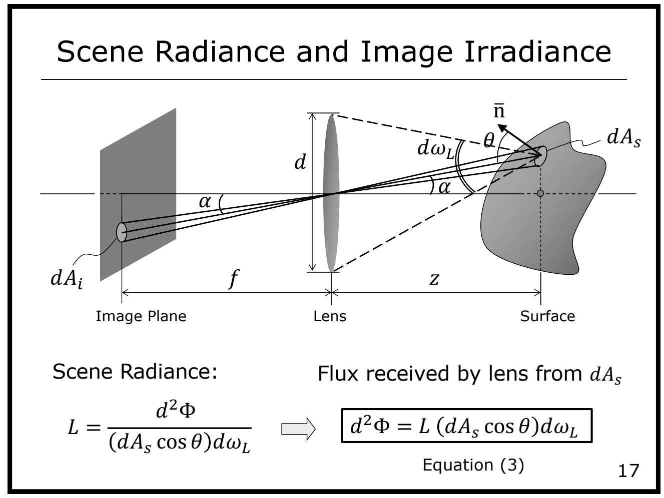

Since all the flux received by the lens eventually arrives at the pixel, the irradiance of the pixel is simply the flux received by it divided by its area, dA_i. This is equation (4).

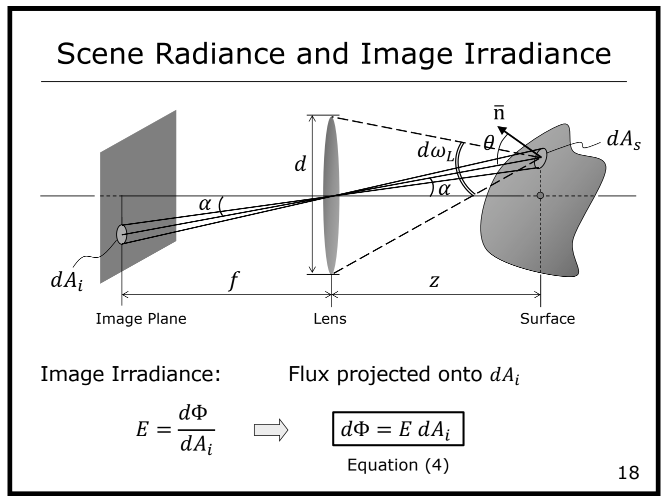

Using the four equations we obtained above, we can derive the relationship between scene radiance L and image irradiance E.

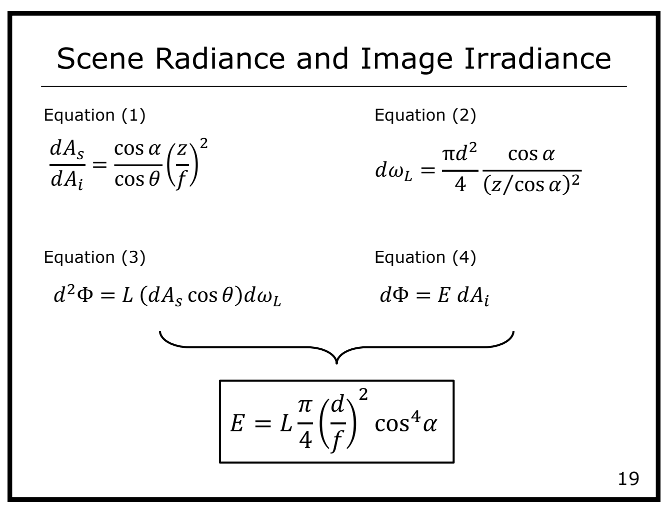

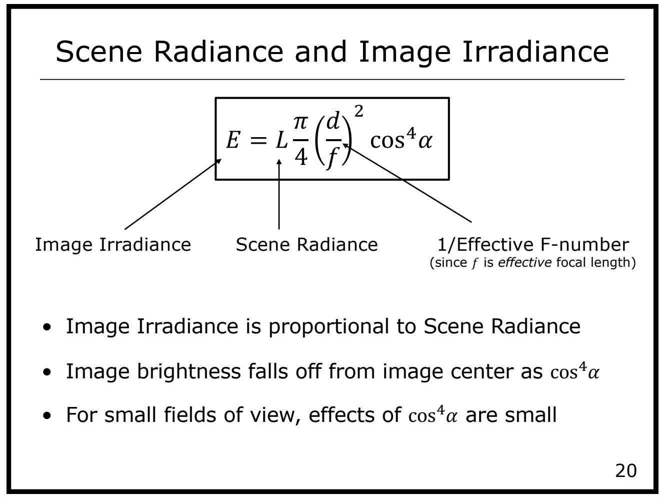

We can make a few observations with respect to the relation between scene radiance and image irradiance. First, image irradiance is proportional to scene radiance. If we double the radiance of the scene patch, the image irradiance will double as well. Second, image irradiance falls off from the center of the image as the fourth power of the cosine of the angle between the direction of the scene patch and the optical axis. As a result, if we move the scene patch away from the center of the image (the optical axis), its brightness in the image will decrease. This fall-off will, however, be small if the field of view of the camera is small. The fall-off is a consequence of using a single lens to image the scene. Compound lenses, which include a series of lenses, are designed to minimize the fall-off.

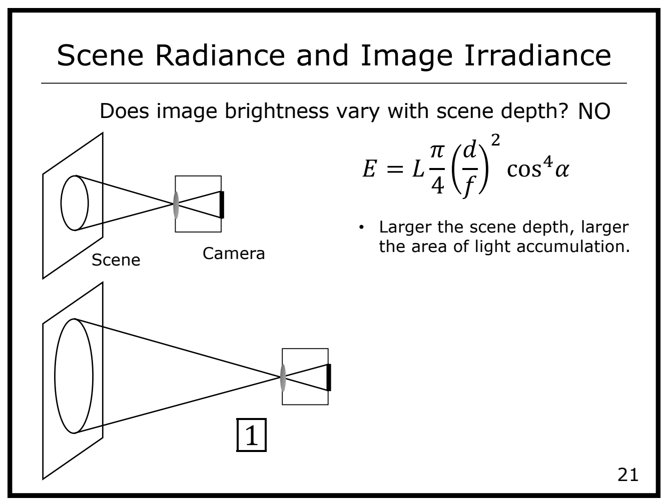

Here is a question that is critical in the context of computer vision: does image brightness vary with scene depth? That is, if we increase the distance of the surface patch, does its image irradiance change? Intuitively, we might think the image irradiance of the patch should decrease as its distance increases. But if we examine the image irradiance equation we derived, we see that the distance between the scene patch and the lens, which is z, does not appear in equation. Therefore, image brightness does not vary with scene depth.

Let us examine why image brightness does not vary with scene depth. Shown on the left is a planar scene imaged with a camera with a single pixel. If we pull this camera back, away from the scene plane $\boxed{1}$, the pixel is going to receive light from a larger area in the scene plane. This suggests that the brightness of the pixel will increase with distance of the camera from the scene!

However, we need to also consider the solid angle subtended by the lens from each point on the scene patch, which determines how much light is being collected by each point in the scene patch. Note that this solid angle decreases as the distance of the camera from the scene increases. Hence, while moving away from the scene causes the pixel to accumulate light from a larger area of the scene, the pixel simultaneously collects less light from each point within the area. These two effects cancel each other, and, as a result, image irradiance is independent of scene depth.

Scene Radiance and Image Irradiance

Does image brightness vary with scene depth? NO

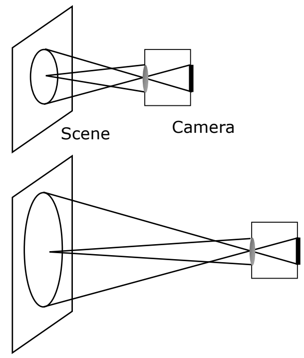

$$E = L\frac{\pi}{4}\left(\frac{d}{f}\right)^2 \cos^4\alpha$$

- Larger the scene depth, larger the area of light accumulation.
- Larger the scene depth, smaller the solid angle subtended by each point onto the lens, and hence less light from each point.

22

We know the relationship between image irradiance and scene radiance. Now, we will describe how surface radiance is related to surface irradiance (illumination). This relation is clearly driven by the physical properties of the material the surface is made of. We refer to it as the reflectance of the surface and it is concisely represented by the bidirectional reflectance distribution function, or BRDF.

BRDF: Bidirectional Reflectance Distribution Function

Shree K. Nayar
Columbia University

Topic: Radiometry and Reflectance, Module: Reconstruction I
First Principles of Computer Vision

23

Here are several spheres in a scene. The light sources are distant from the scene and hence every sphere receives the same illumination. Furthermore, the spheres are observed by a camera that is far away, so the viewing direction for every sphere is also the same. The only reason then that the spheres look different from each other is that they are made of different materials and thus have different reflectance properties.

Surface Appearance

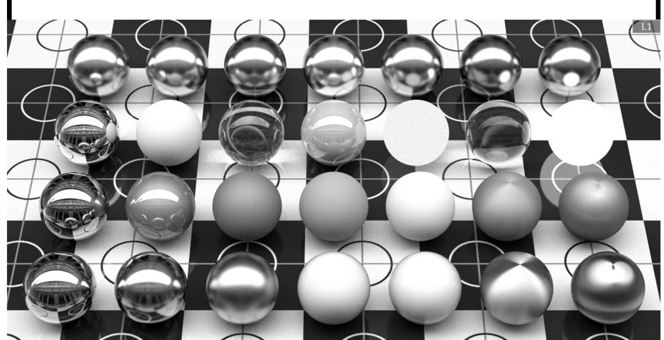

24

That brings us to surface reflectance. We see that two directions are important while defining reflectance. The first is the direction from which the light arrives, and second is the direction from which light is viewed.

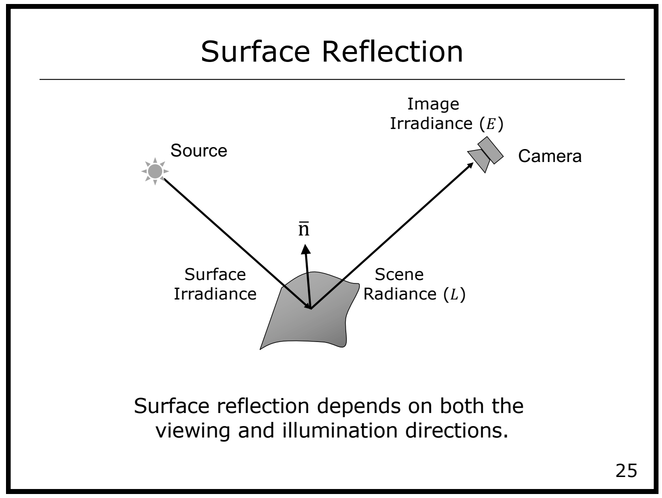

The reflectance of a material can be fully described by the bidirectional reflectance distribution function (BRDF). Let us represent direction using two angles: the zenith angle θ and azimuth angle ϕ, as shown on the left. Let the illumination direction be (θ_i, ϕ_i) and the viewing (reflection) direction be (θ_r, ϕ_r). The BRDF of a surface is then the ratio of its radiance $L(\theta_r, \phi_r)$ and its irradiance $E(\theta_i, \phi_i)$. It is therefore a four-dimensional function that fully describes the reflectance properties of a scene point. The unit of the BRDF is 1 over steradian.

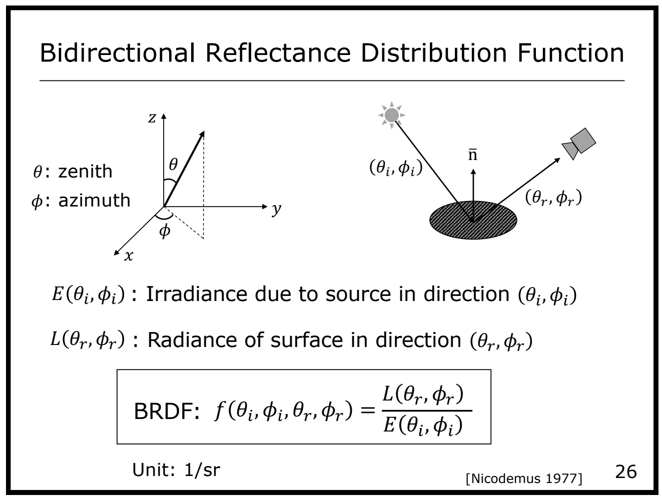

Let us take a look at some important properties of the BRDF. The first is that it is always greater than, or equal to, zero. Since irradiance and radiance cannot be negative, their ratio f is always positive. A more interesting property is Helmholtz reciprocity, which says that we get the same value of the BRDF if we flip the illumination and reflection directions. That is, if we swap the camera and the light source in the figure, the value of the BRDF remains the same. This property is based on the law of conservation of energy.

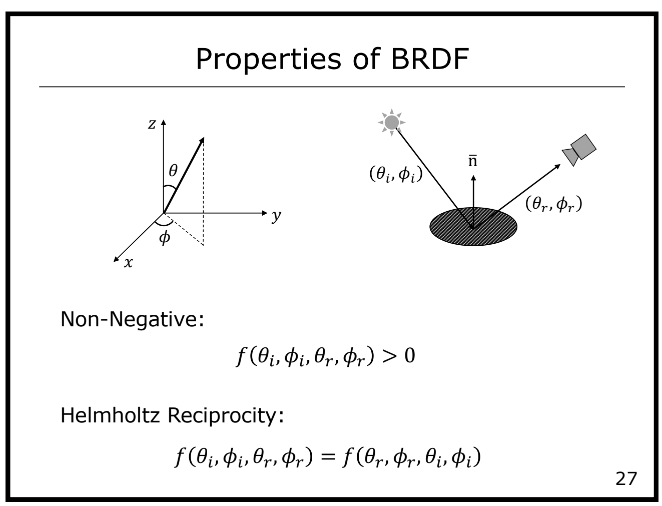

While the BRDF is, in general, a four-dimensional function, for a large class of materials, it reduces to a three-dimensional function. These are referred to as isotropic surfaces. Such a surface has reflectance properties that are rotationally symmetric—that is, in any setting of the illumination and the camera, if the surface is rotated about its surface normal, its brightness as seen by the camera will not change. Surfaces that are not isotropic are called anisotropic, and their BRDFs remain four-dimensional.

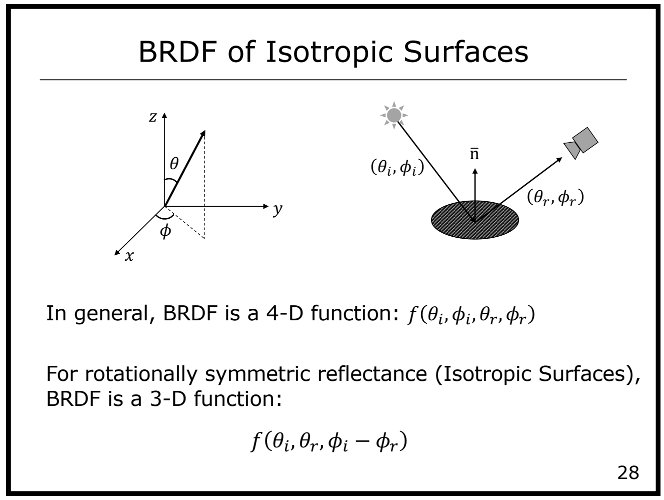

Shown on the left is a sphere with an isotropic BRDF and on the right is a sphere with an anisotropic BRDF. If we rotate the isotropic sphere about its center, its appearance will not change. Whereas, if we rotate the anisotropic sphere about its center, its appearance will change, possibly in a dramatic way. Examples of anisotropic surfaces include ones that are machined or sandpapered. In such cases, the surface can end up with a micro-structure that is directional (grooves, for instance), causing its BRDF to vary as a function of the azimuth angle ϕ.

In nature, there are a lot of materials that are anisotropic. We see two examples here—butterfly wings and peacock feathers. When we move around, their appearances can change dramatically, producing striking visual effects.

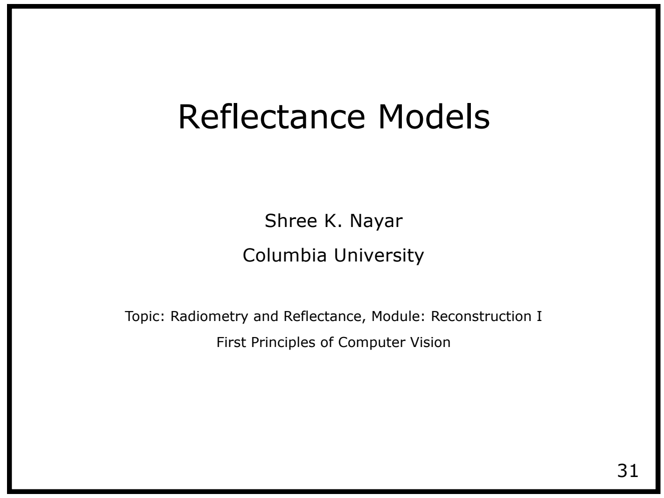

Using the BRDF, we can express the reflectance properties of any material found in the real world. Here, we will focus on a couple of reflectance models that are commonly found in the real world, and widely used in computer vision and computer graphics.

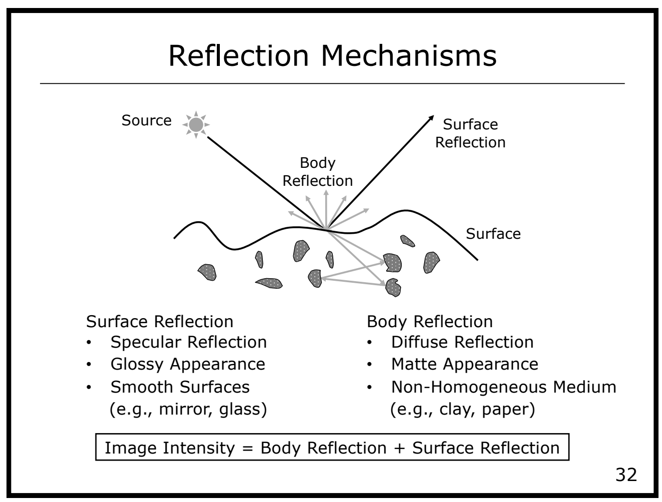

Let us first take a closer look at the main mechanisms that produce reflection. The first one is called surface reflection, where light falling on the surface is reflected at the interface (the surface itself). It is also called specular reflection and gives objects a shiny appearance. Surface reflection is strong in the case of smooth surfaces like mirrors, glass and polished metal.

The second mechanism is called body reflection, where a portion of the light enters the body of the surface. Most materials are non-homogenous, in that, they are composed of particles that have different physical properties, and hence different refractive indices. The light that enters the material therefore gets refracted and reflected multiple times. Some of this light remerges at the surface and leaves the material in a wide range of directions. This body reflection is also called diffuse reflection, as it gives the material a matte appearance. It dominates in the case of non-homogeneous materials such as clay, plaster and paper.

In general, when we look at the image intensity due to a scene point, it is a combination of two components—a body reflection component and a surface reflection component.

Here are a couple of examples of objects that exhibit the reflection mechanisms discussed above. The clay vase on the top has primarily body reflection and hence has a matte appearance. In contrast, surface reflection dominates in the case of the mirror sphere. In general, however, surfaces can have a combination of these two mechanisms, which we refer to as hybrid reflection. While the wooden tiles and the painted can at the bottom have body (diffuse) components, they also have highlights which are due to surface reflection.

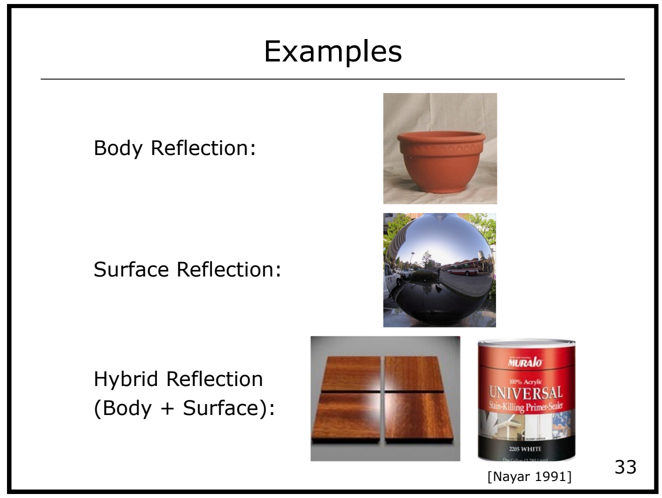

Now let us look at some commonly used models for body and surface reflection. We start with the Lambertian model for body reflection, which was introduced by Lambert in 1760. This model states that a surface appears equally bright in all directions. So, its BRDF is a constant, which is given by ρ_d/π. ρ_d is called the albedo and it ranges from 0 to 1. It is 0 for a perfectly black material that absorbs all the light falling on it, and 1 for a perfectly white material which reflects all the light falling on it. This model is widely used in computer vision and computer graphics, not just because of its simplicity, but also because it describes a variety of surfaces found in the real world.

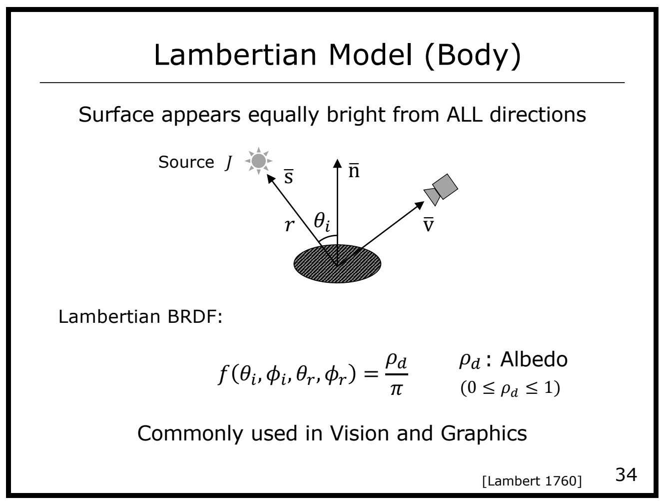

Now, let us look at the relationship between the irradiance E and the radiance L of a Lambertian surface. Since its BRDF is ρ_d/π, we have $L = (\rho_d/\pi)\,E$. Using the definition of irradiance we derived in slide 10, we can express E in terms of the radiant intensity J of the source, the distance r of the source, and the angle θ_i of the source with respect to the surface normal. We also know that $\cos\theta_i$ can be written as the dot product of the surface normal vector $\bar{n}$ and the source direction vector $\bar{s}$. The resulting expression for the radiance L of a Lambertian surface is given at the bottom.

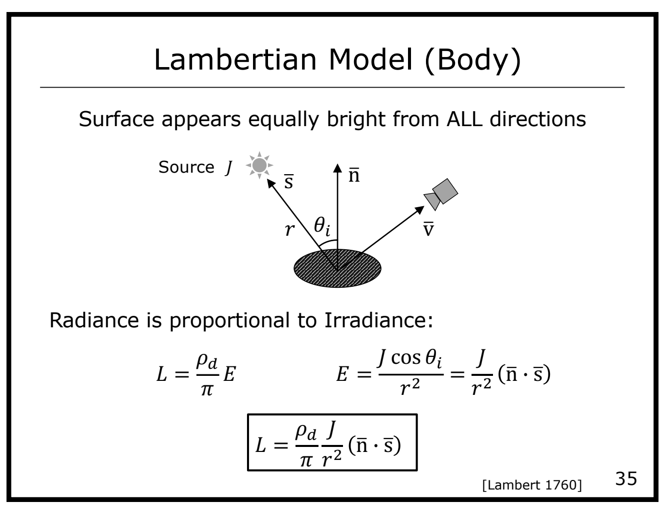

In the above equation for the radiance of a Lambertian surface, let us assume that the distance of the light source, r, is a constant, and the direction θ_i of the source with respect to the surface is varied. When the surface is lit from the top; θ_i is equal to 0, and $\overline{n} \cdot \overline{s}$ is equal to 1. The radiance is maximum in this case and equal in all directions. As we increase θ_i, the radiance drops, while still being equal in all directions. Eventually, the radiance drops to 0 at θ_i = $90°$.

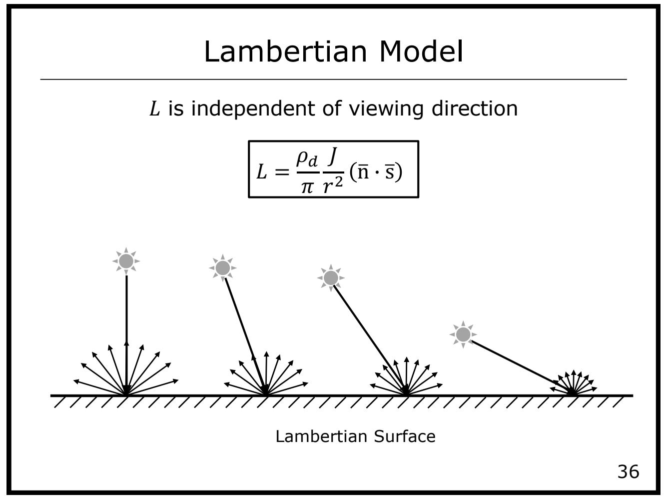

Now let us go to the other end of the spectrum and look at a perfect mirror. In this case, we have pure surface reflection, or specular reflection, and no body reflection. Since it is a perfect mirror, the light incident from a single direction is reflected in a single direction. Consider the setup shown here, where a mirror with normal $\overline{n}$ is illuminated by a source in direction $\overline{s}$. In this case, all the incident light is reflected in the direction $\overline{r}$, which is called the specular direction. The vector $\overline{r}$ is a reflection of $\overline{s}$ about $\overline{n}$, such that the angle of reflection equals the angle of incidence. So, a viewer (camera) receives light only when the viewing direction is equal to the specular direction $\overline{r}$.

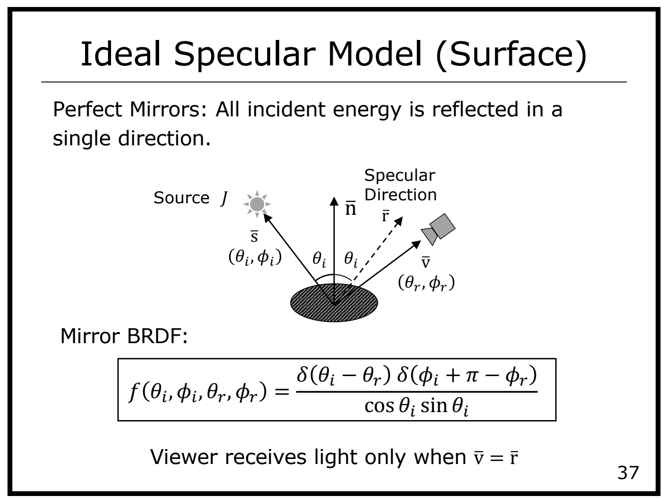

In this case, the BRDF is expressed as a product of two delta functions. The first delta function ensures that the two zenith angles are the same. The second delta function ensures that the two azimuthal angles are the opposite of each other. The $cos\ \theta_i\ \ sin\ \theta_i$ in the denominator ensures that the law of conservation of energy is satisfied—the light energy that is reflected by the surface exactly equals the light energy received by it.

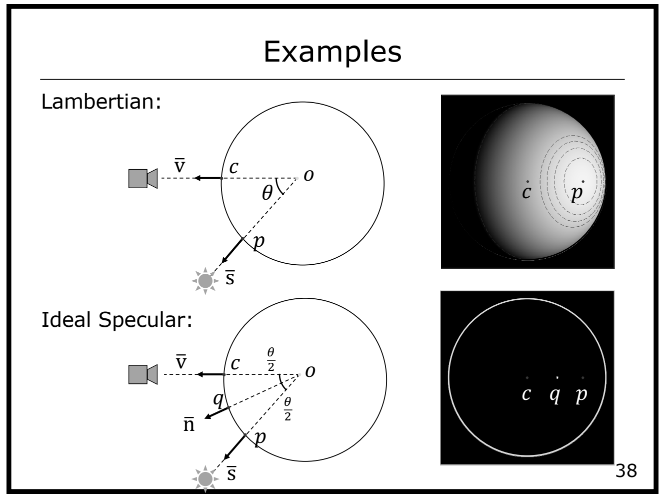

Let us look at how the Lambertian and specular models manifest in terms of object appearance. Let's first take a look at the Lambertian model. Let us assume that we are looking at a Lambertian sphere, where the viewing direction is $\bar{v}$ and the source direction is $\bar{s}$. Let us assume that the source and camera are far away (at infinity) and the angle between them is θ. Shown on the right is the image that the camera captures. The brightest point is p, for which the angle of incidence is equal to 0. Contours of equal brightness in the image correspond to points on the sphere that have the same angle of incidence.

Now, let us look at the other extreme, which is a specular sphere. We will assume the same setup of the camera and the light source. In this case, the camera observes a reflection, a non-zero value, only for one point, q, on the sphere. The surface normal of q is the bisector of $\bar{v}$ and $\bar{s}$.

The Lambertian and specular models are simple but are widely used because we have many mirror-like objects and matte objects in the real world.

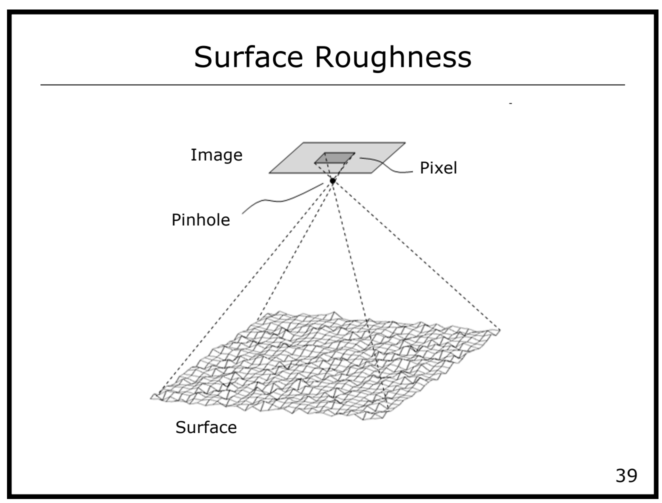

Next, we discuss the effect of surface roughness on the BRDF. Let us assume that we are looking at one pixel in the image shown here. This pixel projects onto a patch on the surface. Within this patch, the surface is likely to have undulations (roughness). While each point within the patch may be Lambertian or specular, the roughness changes the aggregate BRDF of the patch. We would like to devise a way to model the roughness of a surface and then determine the BRDF of the patch, for any given local BRDF (Lambertian or specular).

A simple way to model surface roughness is by assuming the surface to be a collection of microfacets that are tilted in different directions. The surface has a mean orientation $\bar{n}$, but each microfacet has its own orientation angle, α. We can describe the roughness of the surface using a distribution for α. For instance, a Gaussian distribution can be used where the standard deviation σ is then a measure of surface roughness.

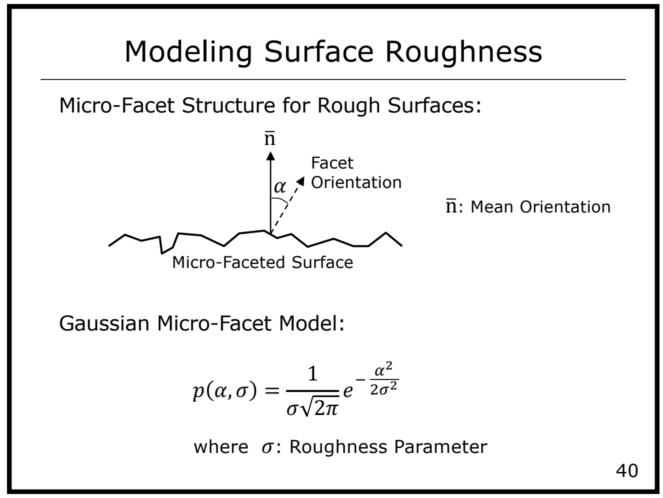

Shown here is the effect of varying the σ of the Gaussian microfacet model. When σ is equal to 0, we get a perfectly flat surface. The surface is slightly rough for σ equal to 0.1, and appears more rough as σ increases.

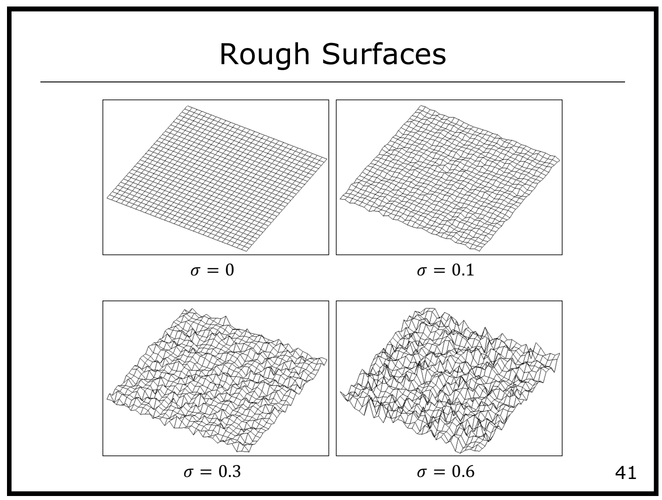

The Torrance-Sparrow BRDF model shown here predicts reflection from a rough surface where each microfacet is a perfect mirror. The BRDF is calculated for the entire ensemble, or collection, of facets. Here, ρ_s represents the specular reflectivity of each facet. The Gaussian roughness model is given by the function p. The geometric attenuation factor G accounts for the fact that adjacent facets can cast shadows on each other, as well as mask each from the observer (camera). An interesting feature of this BRDF is that when σ (roughness) is set to zero, we get the BRDF of a smooth mirror, which is the ideal specular BRDF model.

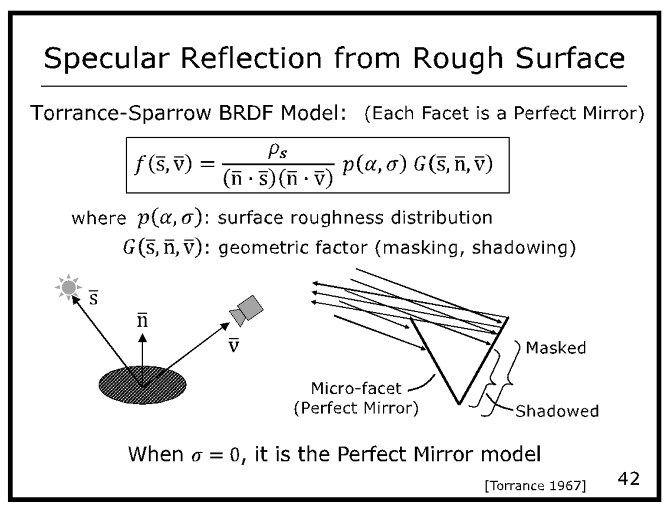

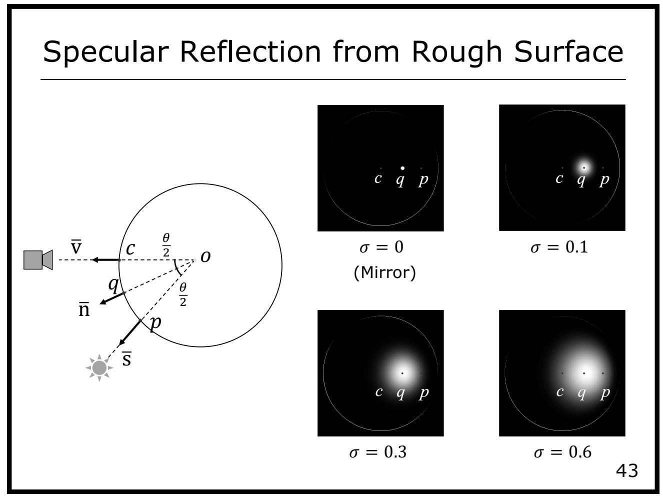

Let us look at how specular reflection from a rough surface manifests in images. Once again, $\bar{v}$ is the viewing direction, $\bar{s}$ is the source direction, and we assume that the camera and source are far from the sphere. For σ equal to zero, we have a perfect mirror (slide 38) and hence get a single bright spot at q, the point whose surface normal is the bisector of $\bar{v}$ and $\bar{s}$. Increasing σ results in a lobe, or a highlight, which gets broader with σ. An interesting feature of the Torrance-Sparrow model is that, for a fairly rough surface, the brightest point is not q; it is a point that lies between q and p. This is due to the geometric effects of masking and shadowing we discussed earlier.

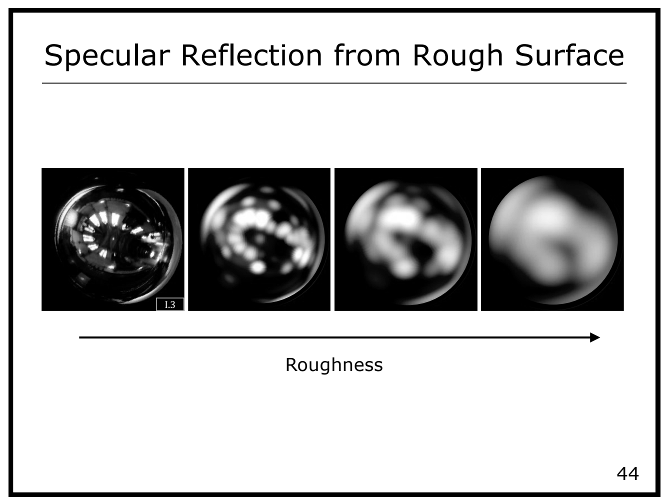

We will now see how a specular rough sphere might appear in the real world. Shown here are spheres with increasing roughness, from left to right, lit by a complex environmental illumination. As roughness increases, as expected, the scene reflected by the sphere looks more and more blurred.

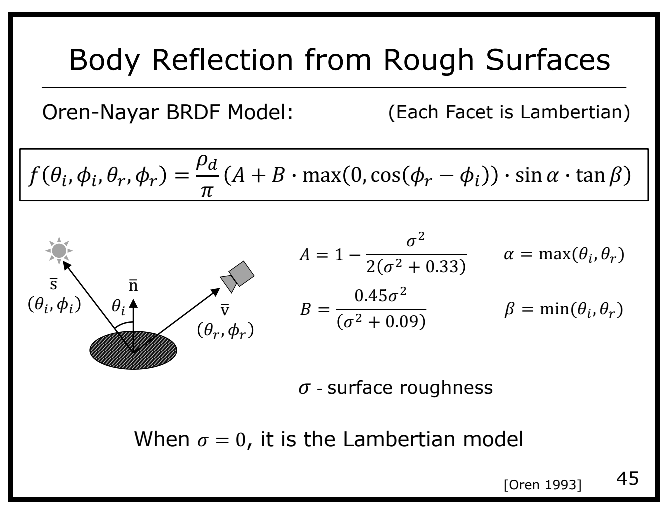

At the other end of the spectrum, we can consider the same microfacet model for surface roughness, but assume that each facet is Lambertian instead of specular. This is the Oren-Nayar BRDF model, and it can be used to represent materials such as plaster, clay and bricks. In this model, ρ_d is the Lambertian albedo of each microfacet. When σ is equal to zero, the model reduces to the perfect Lambertian model.

Here we use the Oren-Nayar model to render the appearance of rough spheres, where each microfacet is Lambertian. The image configuration is the same as that used previously. When σ is equal to zero, we get Lambertian reflectance. As σ increases, we see that the sphere begins to appear flatter.

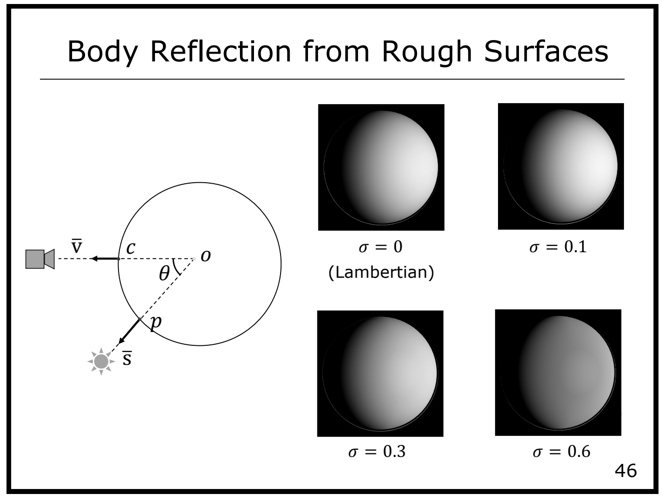

This flattening effect is particularly pronounced when the viewing and source directions are the same, which is the case for the three spheres on the left. When σ is equal to zero, we get a perfectly Lambertian sphere where brightness falls as we go from the center to the edge. As we increase σ, the sphere begins to appear flatter. For the largest σ (right most sphere), we can see that the sphere appears more or less like a flat disc. Indeed, this is the phenomenon we observe in the case of the full moon. The full moon does not look like a shaded sphere but more like a flat disc. It is hypothesized that this is because of the roughness of the surface of the moon.

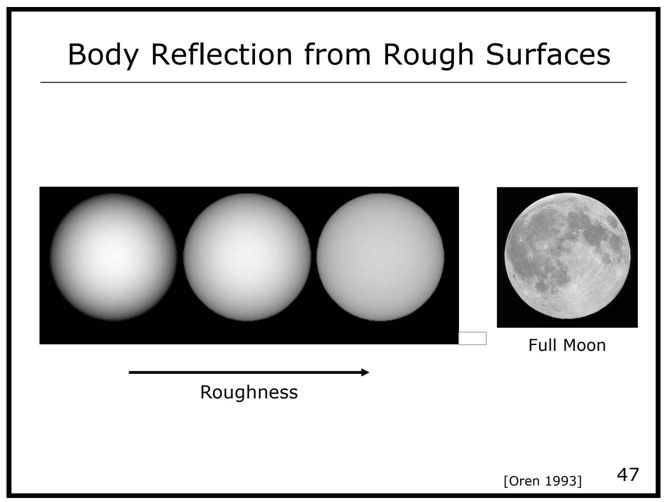

We have discussed two basic mechanisms of reflection: surface reflection and body reflection. We have not yet discussed the impact of these mechanisms on the color of incident light. This brings us to the dichromatic reflectance model.

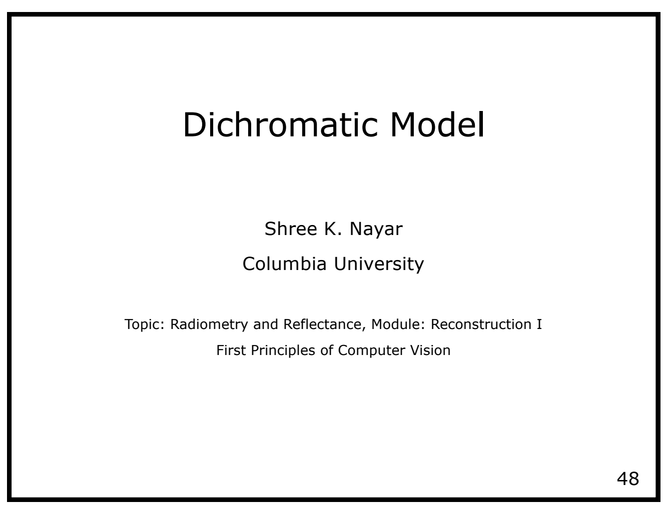

Consider our two main mechanisms of reflectance once again. The first is surface reflection, which happens at the interface. It turns out that for most materials, the color of surface reflection is the same as the color of the incident light. That is, surface reflection preserves the color of the light source. When it comes to body reflection, the light enters the material and gets bounced around by non-homogeneities in the material. In this process, it is expected that the material will absorb some wavelengths of light more than others. Thus, body reflection is going to have a different color from the incident light. Loosely speaking, we can say that the color of the body reflection is the product of the color of the illumination and the color of the object.

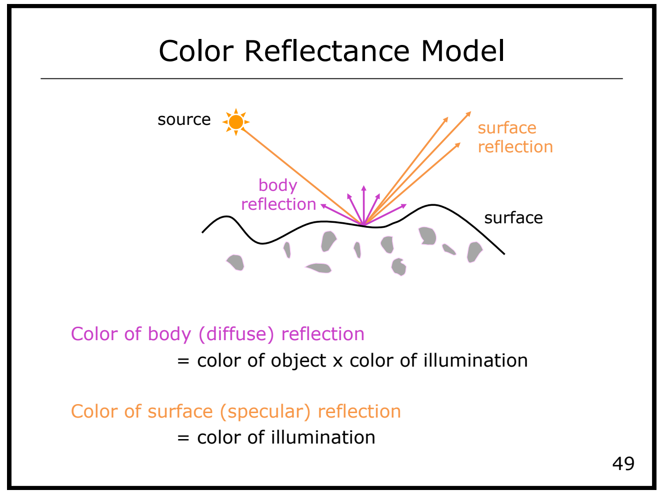

The color at an image pixel is therefore a linear combination of the color of body reflection and the color of surface reflection. On the right, the colors of the body and surface reflections are shown as vectors in RGB color space. These two vectors define what is referred to as the dichromatic plane. Let us assume that the object is uniform in terms of its material properties. Then, any image pixel that lies on the object will measure a color that is a linear combination of the two vectors and hence must lie on the dichromatic plane.

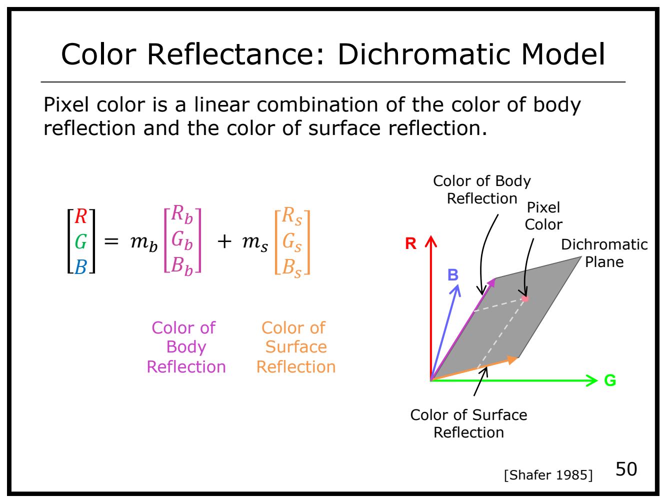

Shown here is a reddish sphere that is illuminated with blueish light. When we map the colors of all the pixels on the sphere to the RGB color space, we get the distribution seen on the right. In this distribution, the lower edge represents the color of body reflection, and the right edge represents the color of surface reflection. We also see an artifact at the top of the distribution. These are "clipped" colors where one or more of the color channels are saturated. If we ignore the clipping artifact, the distribution of colors appears like a "skewed-T".

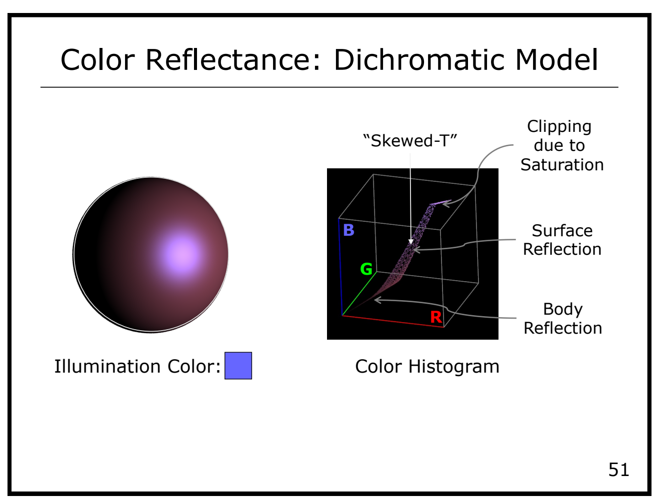

On the left is an image with three plastic objects of different colors. If we map every pixel in this image to color space, we get the three distributions on the right. Each one of these skewed-T distributions is associated with one of the three objects. Consider the skewed-T of a single object shown at the bottom. If we take the points on the left edge of the skewed-T, we can get the color of the body reflection. Similarly, the points at the top of the skewed-T reveal the color of surface reflection. Once these two vectors have been computed, we can decompose the color of each point on the object into its two components: its body reflection and surface reflection.

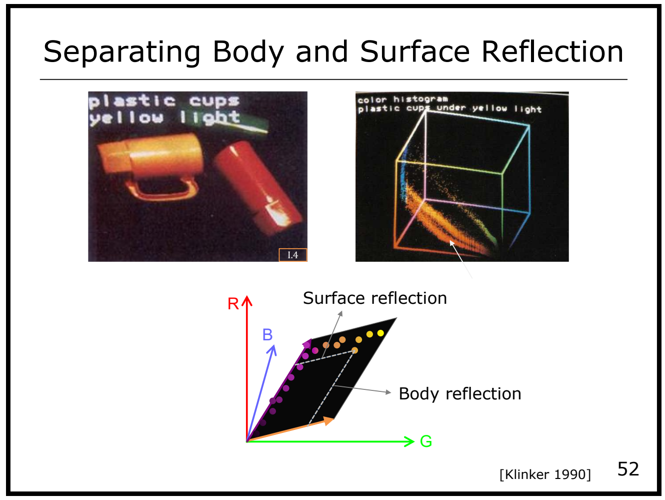

By applying the above separation to every pixel in the input image (top-left), we can decompose it into a body reflection image (bottom-left) and a surface reflection image (bottom-right). This is useful because highlights, or specular reflections ,are often a nuisance in computer vision as they float on top of the objects and move over the object when the illumination direction changes. On the other hand, the body reflection gives us the true shading of the object, which, as we will see in a subsequent lecture, can be used to compute the three-dimensional shape of the object.

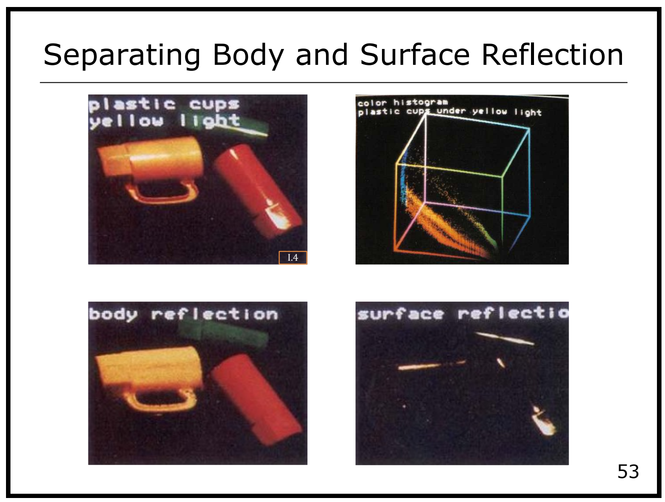

References and Credits

Shree K. Nayar

Columbia University

Topic: Radiometry and Reflectance, Module: Reconstruction I

First Principles of Computer Vision

54

References: Papers

[Klinker 1990] G. J. Klinker, S. A. Shafer, and T. Kanade. *A Physical Approach to Color Image Understanding*. IJCV, vol.4, pp.7-38. 1990

[Lambert 1760] Johann Heinrich Lambert. *Photometria*. 1760.

[Nayar 1991] S. K. Nayar, K. Ikeuchi, and T. Kanade. *Surface Reflection: Physical and Geometrical Perspectives*. PAMI 1991.

[Nicodemus 1977] F. E. Nicodemus, J. C. Ridmond, I. W. Ginsberg, and T. Limperis. *Geometrical considerations and nomenclature for reflectance*. National Bureau of Standards Monograph, 160:201-231. October, 1977.

[Oren 1993] M. Oren and S.K. Nayar. *Generalization of Lambert's Reflectance Model*. SIGGRAPH. pp.239-246, Jul, 1994

[Shafer 1985] S. A. Shafer. *Using color to separate reflection components*. Color Research and Application, 10:210-218, 1985.

[Torrance 1967] K. E. Torrance and E. M. Sparrow. *Theory for off-specular reflection from roughened surfaces*. J. Opt. Soc. Am. 57, 9 (Sep 1967).

[Westin 1992] S. H. Westin, J. R. Arvo, and K. E. Torrance. *Predicting Reflectance Functions from Complex Surfaces*. SIGGRAPH 1992

55

Image Credits

I.1 Kevin Beason. Used with permission.
I.2 Stephen Westin. Used with permission.
I.3 Fredo Durand. Used with permission.
I.4 G.K. Klinker. Used with permission.
I.5 Purchased from iStock by Getty Images.

56

Acknowledgements: Thanks to Joel Salzman, Tracy Cui, Nikhil Nanda and Ayush Sharma for their help with transcription, editing and proofreading.

References

[Klinker 1990] G. J. Klinker, S. A. Shafer, and T. Kanade, A Physical Approach to Color Image Understanding, IJCV, vol.4, pp.7-38, 1990.

[Lambert 1760] Johann Heinrich Lambert, Photometria, 1760.

[Nayar 1991] S. K. Nayar, K. Ikeuchi, and T. Kanade, Surface Reflection: Physical and Geometrical Perspectives, PAMI 1991.

[Nicodemus 1977] F. E. Nicodemus, J. C. Ridmond, I. W. Ginsberg, and T. Limperis, Geometrical considerations and nomenclature for reflectance, National Bureau of Standards Monograph, 160:201-231. October, 1977.

[Oren 1993] M. Oren and S.K. Nayar, Generalization of Lambert's Reflectance Model, SIGGRAPH, pp.239-246, Jul, 1994.

[Shafer 1985] S. A. Shafer, Using color to separate reflection components, Color Research and Application, 10:210-218, 1985.

[Torrance 1967] K. E. Torrance and E. M. Sparrow, Theory for off-specular reflection from roughened surfaces, J. Opt. Soc. Am., 57, 9 (Sep 1967).

[Westin 1992] S. H. Westin, J. R. Arvo, and K. E. Torrance, Predicting Reflectance Functions from Complex Surfaces, SIGGRAPH, 1992.