BRDFs

Lecture #1b: Wednesday, 28 September 2000
Lecturer: Pat Hanrahan, Szymon Rusinkiewicz, Steve Marschner
Scribe: Ravi Ramamoorthi
Reviewer: Yelena Vileshina

1 Introduction

Bidirectional Reflectance Distribution Functions (BRDFs) are the fundamental quantity describing reflectance. After a brief introduction (by Pat), this lecture discusses common features of BRDFs and possible representations (by Szymon), a few general constructions (by Steve), and BRDF reparameterizations (by Szymon and Steve).

First, we consider L, the radiance. The radiance is roughly the power per unit direction per unit area and is more formally defined as the power per unit projected area perpendicular to the ray per unit solid angle in the direction of the ray.

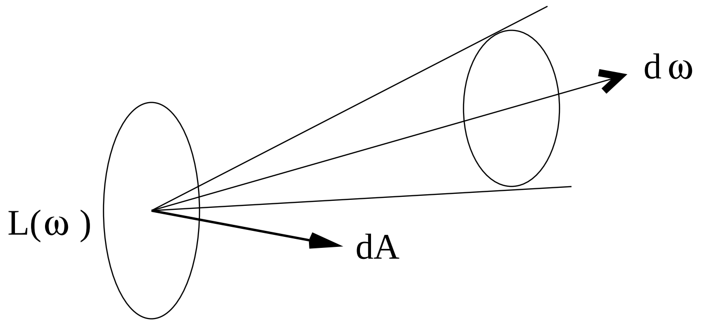

Figure 1: Radiance is the power per unit projected area per unit solid angle

We can also write the radiance in terms of the *throughput* as

$$L = \frac{d\Phi}{dT} \tag{1}$$

where Φ is the power and T is the throughput defined as

$$dT = d\vec{\omega} \cdot d\vec{A} \tag{2}$$

where ω is the solid angle, A is the area, and we consider $\vec{\omega}$ and $\vec{A}$ as vectors—so their dot product will include the cosine of the angle between them. The differential throughput

dT is a measure of the number of lines in the beam; the size of the beam depends on the throughput $T = \int dT$. For the same beam, the throughput remains constant. Therefore, in geometric optics, both the throughput and the power (or energy) are conserved.

Thus, the limit for the radiance

$$L = \lim_{\triangle T \to 0} \frac{\triangle \Phi}{\triangle T} \tag{3}$$

will be well defined since both the numerator and denominator are invariant quantities. Therefore, the radiance is a fundamental quantity, and the radiance along a ray remains constant as it propagates, assuming there are no losses because of absorption in a medium. Also, the response of a sensor is proportional to the radiance of the surface visible to the sensor; the constant of proportionality is the throughput, a function only of the geometry of the sensor.

As far as reflection is concerned, we want some measure of the ratio of outgoing to incoming energy. However, it is important to define the limit properly in order for it to make sense. Consider as a first attempt,

$$r = \frac{\triangle \Phi_o}{\triangle \Phi_i} = \frac{\triangle L_o(\vec{\omega_o}) \cos(\theta_o) \triangle \omega_o \triangle A}{L_i(\vec{\omega_i}) \cos(\theta_i) \triangle \omega_i \triangle A}$$

Taking the limit of this quantity as $\triangle A$, $\triangle \omega_i$ and $\triangle \omega_o$ tend to 0 will lead to an inconsistency. Basically, there are three differential quantities in the numerator and only two in the denominator. Also, $\triangle L_o$ varies linearly and is related to $\triangle \omega_i$ which leaves $\triangle \omega_o$ not cancelled by any term in the denominator. Similarly, an attempt to define $r = \triangle L_o / L_i$ will lead to an inconsistency.

In order to have a well-defined limit for the BRDF, we must first define the irradiance, E as the radiant power per unit area incident on a surface. The differential irradiance because of a small solid angle $\triangle \omega_i$ is then

$$\triangle E = L_i(\vec{\omega_i}) \cos(\theta_i) \triangle \omega_i \tag{4}$$

Note that if we were to multiply by the differential area, this would simply be the incoming radiance multiplied by the differential throughput i.e. the differential response of a sensor.

Finally, we define the BRDF f_r as the ratio of the outgoing (differential) radiance to the incoming differential irradiance that produces it.

$$f_r(\vec{\omega_i}, \vec{\omega_o}) = \lim_{\triangle E, \triangle L_o \to 0} \frac{\triangle L_o(\vec{\omega_o})}{L_i(\vec{\omega_i}) \cos\theta_i \triangle \omega_i} = \lim_{\triangle \omega_i, \triangle \omega_o \to 0} \frac{\triangle L_o(\vec{\omega_o})}{L_i(\vec{\omega_i}) \cos(\theta_i) \triangle \omega_i} \tag{5}$$

Upon integrating the outgoing radiance over the the incoming hemisphere, we obtain the *reflection equation*:

$$L_o(\vec{\omega_o}) = \int_{\Omega_i} f_r(\vec{\omega_i}, \vec{\omega_o}) L_i(\vec{\omega_i}) \cos(\theta_i) d\,\omega_i \tag{6}$$

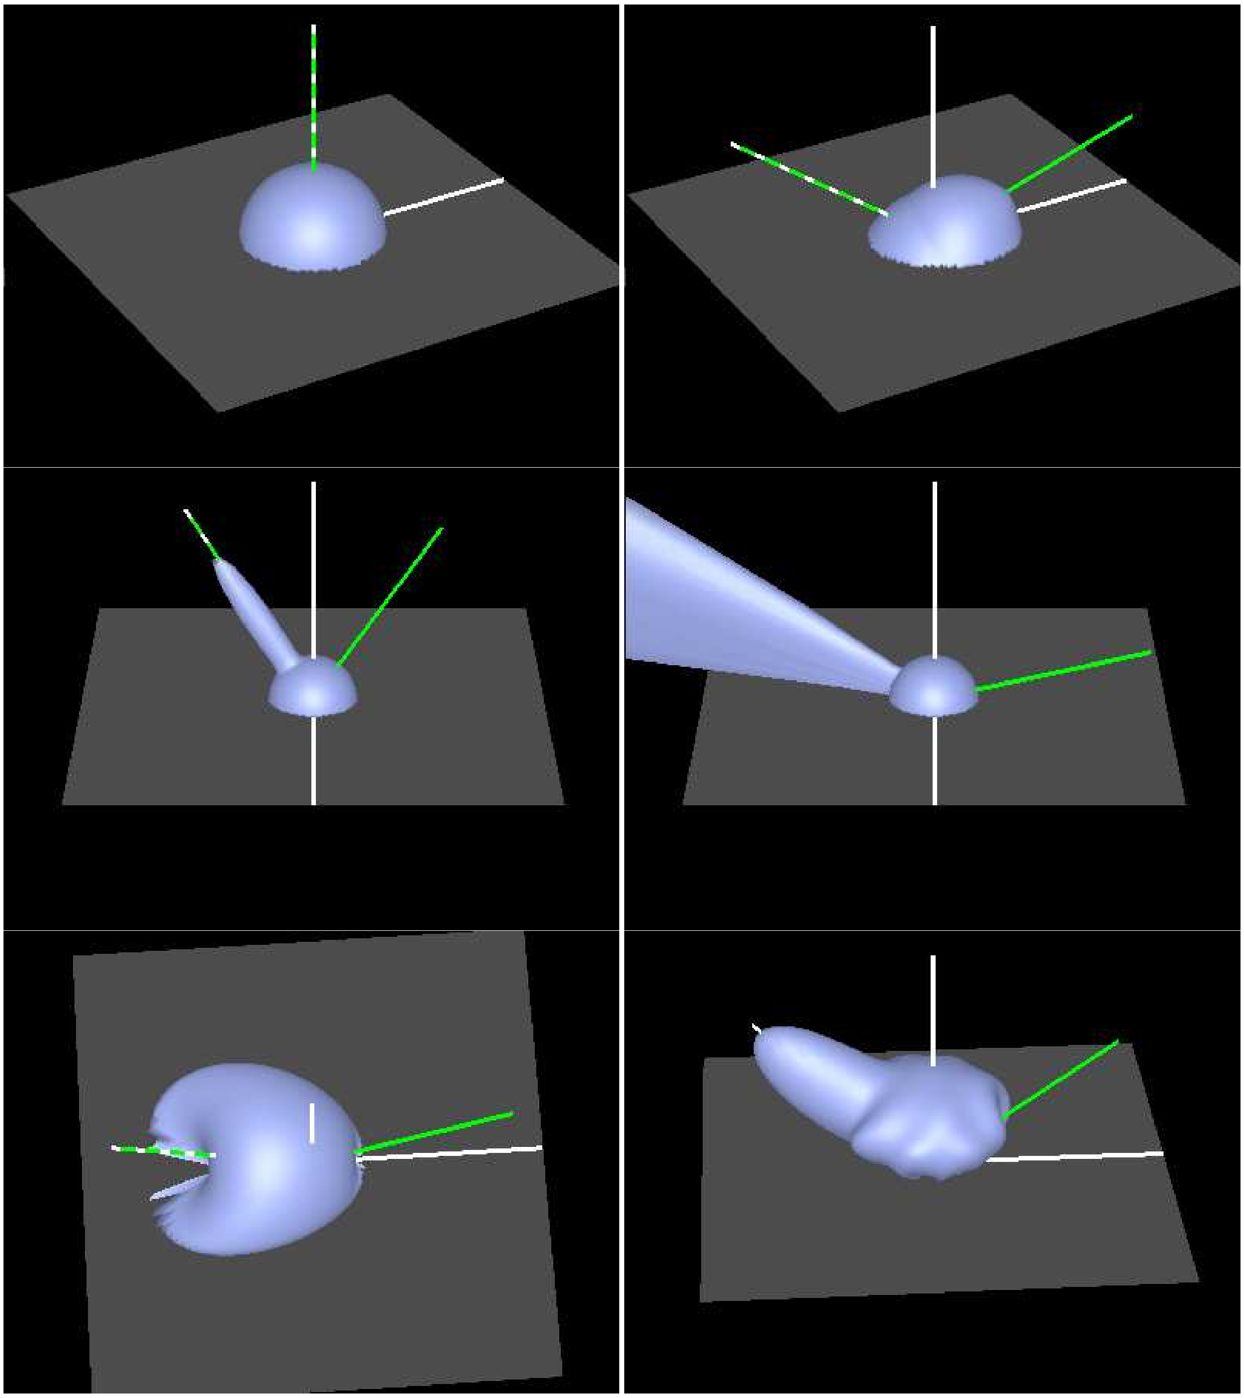

Figure 2: Visualization of some common BRDFs. *Top:* Diffuse (left) and the Hapke/Lommel-Seeliger lunar BRDF with retroreflection (right). *Middle:* Torrance-Sparrow (left) and at grazing angles (right) with increased Fresnel reflection. This also involves a significant off-specular component. *Right:* Anisotropic BRDF (left) and using spherical harmonics to approximate (right) demonstrating ringing.

2 Common Features of the BRDF

Visualizations of some common BRDFs discussed in this section are found in figure 2. We briefly go thru some common BRDFs, discussing their key features.

Diffuse: The diffuse or Lambertian BRDF is just a constant and is the simplest BRDF. Note that while the BRDF is constant, there is a cosine falloff of intensity toward grazing incident angles with respect to the light source. While not physically realizable, the Lambertian surface is often a good first-order approximation and is a widely used phenomenological model.

Torrance-Sparrow: The Torrance-Sparrow BRDF includes a peak in the direction of specular reflection. For a perfect mirror, there would be a single spike in the mirror direction. For general somewhat rough microgeometry, we instead get a lobe around the specular direction, the width of which depends on how rough the surface is. Other features are a sharp increase in the intensity of this specular reflection at grazing angles—because of the Fresnel effect to be discussed below—and off-specular peaks in the BRDF distribution because of similar effects.

Below, we present a derivation of the *Fresnel* term. This describes the amount of light reflected in the specular direction by a surface and is responsible for increase reflection as we approach grazing angles.

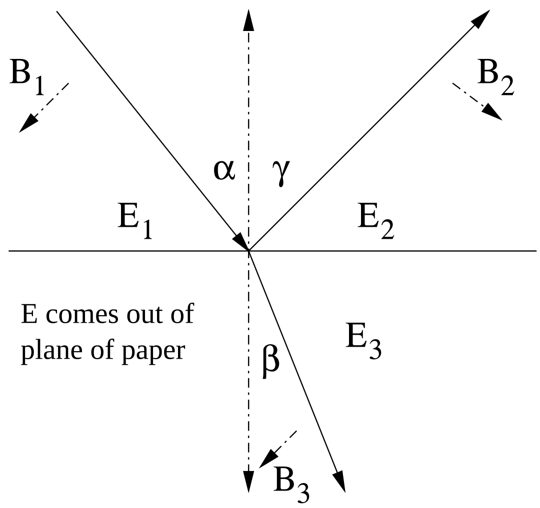

Figure 3: Diagram for Fresnel reflection, E perpendicular to plane of incidence

We derive the equations from a number of continuity conditions. Here, we have only shown the polarization with E coming out of the plane of the paper. We first define $r = E_2/E_1$ and $t = E_3/E_1$. r and t are the amplitude coefficients for reflection and transmission, respectively.

Our first condition of continuity is that the perpendicular or tangential component of E is continuous at the interface. This implies that $E_1 + E_2 = E_3$. Note that this is because of our sign convention. We can then write (dividing by E_1)

$$1 + r = t \tag{7}$$

Our second condition of continuity is on the normal component of the magnetic field i.e. $B_1 \sin(\alpha) + B_2 \sin(\gamma) = B_3 \sin(\beta)$. Accounting for the index of refraction n of the medium, and dividing by B_1, this becomes

$$\sin(\alpha) + r \sin(\gamma) = nt \sin(\beta) \tag{8}$$

It should be noted that $B = E/v$ where v is the speed of light in the medium, and $v = c/n$ where c is the speed of light in free space. This accounts for the n on the right-hand side.

Next, we consider the continuity of the tangential component of the magnetic field. Here, we must actually consider continuity of B/μ, but we assume as is common that the magnetic permeability of the medium $\mu = 1$. In that case, the equation is similar to the one above viz. $B_1 \cos(\alpha) - B_3 \cos(\gamma) = B_2 \cos(\beta)$ and becomes

$$\cos(\alpha) - r \cos(\gamma) = nt \cos(\beta) \tag{9}$$

Finally, we consider energy conservation viz. that energy does not build up at the interface. Mathematically this is $E_1 B_1 \cos(\alpha) - E_3 B_3 \cos(\gamma) = E_2 B_2 \cos(\beta)$. Again, dividing by $E_1 B_1$ and accounting for the index of refraction, we obtain

$$\cos(\alpha) - r^2 \cos(\gamma) = nt^2 \cos(\beta) \tag{10}$$

We can write the right-hand side of the equation above as $t[nt \cos(\beta)]$. But the bracketed term is the right-hand side of equation 9 and from equation 7, $t = 1 + r$. Making the substitutions, we get

$$\cos(\alpha) - r^2 \cos(\gamma) = (1 + r)[\cos(\alpha) - r \cos(\gamma)]$$

Upon simplifying and dividing by r, this reduces to

$$\cos(\alpha) = \cos(\gamma) \tag{11}$$

This is the *law of reflection* in wave optics stating that the angle of incidence is equal to the angle of reflection.

Using the above result to substitute for γ in equation 8 and using $t = 1 + r$,

$$\sin(\alpha)(1 + r) = n(1 + r) \sin(\beta)$$

from which we derive *Snell's Law*

$$\sin(\alpha) = n \sin(\beta) \tag{12}$$

Finally, we can use $t = 1 + r$ in equation 9 to obtain

$$(1 - r)\cos(\alpha) = n(1 + r)\cos(\beta)$$

from which we can derive

$$r_\perp = \frac{\cos(\alpha) - n\cos(\beta)}{\cos(\alpha) + n\cos(\beta)} \tag{13}$$

A similar derivation for when the electric field is polarized parallel to the plane of incidence yields

$$r_\parallel = \frac{n\cos(\alpha) - \cos(\beta)}{n\cos(\alpha) + \cos(\beta)} \tag{14}$$

The reflectance, a measure of the ratio of the reflected power to the incident power is given by the square of these quantities.

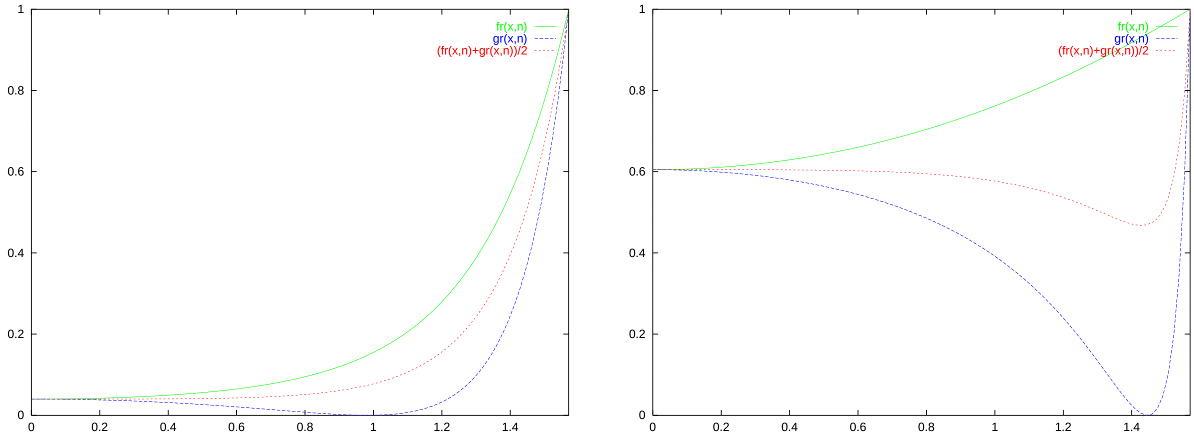

Figure 4: Plots of Fresnel reflectance as a function of incident angle for a dielectric (left) and a metal (right). Note Brewster's angle when reflectance for one of the polarizations vanishes.

Plots of the Fresnel coefficient as a function of the incident angle are given in figure 4. It will be noted that for dielectrics with low n, the reflectance is substantially higher at grazing angles. For metals with higher n, this effect also holds, but the difference between normal incidence and grazing angles is less marked. Note that there is an angle, *Brewsters angle*, given by $tan^{-1}(n)$ at which light is completely polarized, since the reflectance for the other direction of polarization falls to 0.

Retroreflection: As an example, consider images of the moon. Interesting features include no large falloff in intensity towards grazing angles as in a diffuse BRDF, and that the full moon is significantly brighter. The Hapke/Lommel-Seeliger model of the BRDF for the moon has been postulated to take these features into account. This BRDF includes a strong retroreflective peak. Since the surface is very rough, only areas with

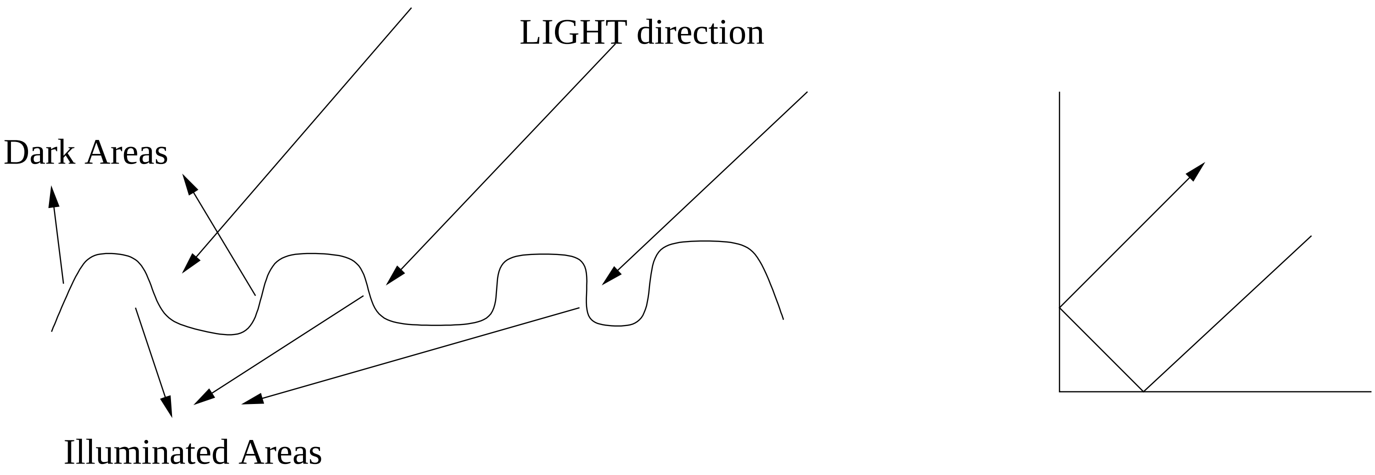

Figure 5: Showing retroreflection from a very rough surface (left). Only areas with normals close to the light direction are well lit, so there is a strong retroreflective peak. On the right, we see a corner reflector (the inside corner of 3 planes is the 3D analog) which produces the same effect.

normals close to the light direction are well lit. These in turn are visible only when the viewing direction is close to the direction of the light source, which explains the retroreflective peak. A similar effect is obtained from a corner reflector, shown on the right of figure 5, and this effect is utilized in bicycle reflectors and highway paint.

Dusty Surfaces: These types of surfaces appear brighter towards grazing angles. This is a similar phenomenological effect as what we saw earlier for the Fresnel effect, although the cause is different. At a normal viewing angle, we see the surface directly, with not much scattering off the dust. However, at grazing angles, we see increased scattering because of the dust and this makes the surface appear brighter. For the same reason, the earth when viewed from space appears brighter near the edges, because of increased scattering of the atmosphere.

Anisotropy: For the most part, BRDFs are isotropic. This means that if we rotate the surface about its normal, leaving the light and viewing vectors unchanged, there is no change in the BRDF. Alternatively, if we rotate both incident and exitant directions about the surface normal by the same amount, we get the same result i.e. there is no preferred tangential direction. This is important for representation because it allows us to use 3 parameters instead of 4 (2 for each of the incident and outgoing angles), with one parameter made redundant because of isotropy. However, BRDFs may in certain cases be anisotropic. This happens, for instance, in machined aluminium, where the scratches lie in particular directions. Anisotropic BRDFs must be represented considering all 4 dimensions.

3 BRDF representations

There are three principal ways to obtain a BRDF. We can derive an analytic formula using physical principles, use simulation given an assumed or measured model of the surface microgeometry, or measure the BRDF based on empirical observation. In any case, we need a method to represent and store the BRDF. This section considers the merits and demerits of various BRDF representations, followed by a discussion of some general representation schemes. We must keep in mind the features discussed already when choosing a representation.

Tabular Representation: The advantage of a table is simplicity, However, general BRDFs are 4 dimensional quantities which can make the table occupy a large amount of storage. Further, quadralinear interpolation of samples must be performed to find an intermediate value. Another troubling aspect of this representation is that BRDFs often have peaks which concentrate a large part of the energy of the BRDF in a very small region, while remaining relatively constant elsewhere. In order to capture this correctly, we will need a very high resolution. It is possible to build schemes with adaptive resolution; see the paragraph on wavelets.

Splines: The benefit of splines over a tabular representation is that it should be possible to represent the smooth parts with fewer control points. However, splines are relatively complicated and can be difficult to sample correctly.

Analytical Formulae: A number of analytic formulae have been proposed for BRDFs. These may be purely ad-hoc or phenomenological, as in the Phong model where we use the specular term $(R_e \cdot L)^n$ with R_e being the reflection of the eye or viewing vector about the surface normal, L being the light vector, and n being an exponent which controls the width of the highlight. Perhaps the best known phenomenological model is the diffuse or Lambertian model with a constant BRDF. Alternatively, formulae may be derived from physical principles. A number of models such as the Torrance-Sparrow model have been derived from the distribution of microfacet orientations. Finally, one may postulate a certain microgeometry for the surface, and simulate the resulting BRDF by using a software raytracer. The disadvantage of all these approaches is lack of generality; features of only a particular class of BRDFs are captured. However, this representation can be very compact—only a few parameters need be stored, although the parameters are often not inutuitive.

Basis Functions: Just as the Fourier basis can be used to represent functions over the real line, Spherical Harmonics or Zernike Polynomials—these will be discussed in more detail under general constructions—can be used to represent BRDFs. These bases are designed to work well for smooth functions. The individual basis functions can generally be assigned intuitive meanings; for instance, the lowest order Zernike polynomial or spherical harmonic merely gives the Lambertian or diffuse BRDF. However, BRDFs have sharp specular peaks which require many basis functions to represent. Further, these peaks are compact in angular space but not compact in frequency space, and representing them by basis functions compact in frequency space, but not angular space, has repercussions over the entire BRDF, leading to undesirable ringing effects.

Wavelets: Wavelets are an alternative multiresolution construction which are compact in both angular and frequency space. As such, they are a good basis for approximating BRDFs. The disadvantage of wavelets as compared to more traditional basis functions is that they are harder to analyze as functional objects in order to derive analytic formulae, or to assign intuitive meaning to. However, for many representational tasks, they appear to be the basis of choice.

General Constructions

We now discuss a few general constructions in terms of basis functions that can be used to represent arbitrary BRDFs.

Spherical Harmonics: Spherical harmonics are an orthogonal basis over the sphere, denoted by $Y_{lm}(\theta, \phi)$, where $l \geq 0$ is the degree, and m with $-l \leq m \leq l$ is the order of the spherical harmonic. An important property is that if a spherical harmonic is rotated, the rotated version can be expressed as a linear combination of spherical harmonics of the same degree l i.e. the degree is invariant to rotation. Spherical harmonics are used in other contexts, for instance, in describing the wave functions of electron orbitals. The functions can be written as $Y_{lm}(\theta, \phi) = P_{lm}(\cos\theta)e^{Im\phi}$ where P_{lm} is an appropriately normalized associated Legendre function, and $I = \sqrt{-1}$.

For the purpose of representing BRDFs, we collapse the indices, defining a single index $u = l^2 + l + m$. We also consider the product of two spherical harmonic basis functions to give a mapping $S^2 \times S^2 \rightarrow R$ from the product of spheres to real numbers. We can define a combined basis function:

$$\tilde{Y}_{uv}(\vec{\omega_i}, \vec{\omega_o}) = Y_u(\vec{\omega_i})Y_v(\vec{\omega_o})$$

To preserve reciprocity, this expression should also be symmetrized with respect to incident and outgoing angles. We can do this by making

$$\tilde{Y}_{uv}(\vec{\omega_i}, \vec{\omega_o}) = N(Y_u(\vec{\omega_i})Y_v(\vec{\omega_o}) + Y_v(\vec{\omega_i})Y_u(\vec{\omega_o})) \tag{15}$$

where N is an appropriate normalizing constant.

A potentially difficult issue is that $\vec{\omega_i}$ and $\vec{\omega_o}$ are defined over only the upper hemisphere while the spherical harmonics are defined over the whole sphere. Further, Fresnel effects cause BRDFs to increase toward grazing angles, causing a spherical harmonic approximation to exhibit ringing throughout the domain.

To work around these issues, we may first multiply the BRDF by $\cos(\theta_i)\cos(\theta_o)$ and represent this new function instead. This goes to 0 at grazing angles, reducing the effects of the Fresnel term, and ensuring continuity at the edges of the hemisphere. To get C^1 continuity everywhere, we may also extend the BRDF to the lower hemisphere, so that it has the same value as the antipodal point in the upper hemisphere. Specifically, $f_r(-\vec{\omega_i}, \vec{\omega_o}) = f_r(\vec{\omega_i}, \vec{\omega_o})$. This makes the BRDF an even function of angle. Since the multiplying cosine term is an odd function, the cosine-multiplied BRDF becomes an odd function of angle, giving us C^1 continuity. A related benefit is that we need only use odd spherical harmonics to express this new function, which means we need not store coefficients of the even spherical harmonics. Isotropy can be maintained as discussed below for Zernike polynomials.

Zernike Polynomials: While spherical harmonics can be used, we need several tricks to convert them from a basis on the sphere to one over the hemisphere on which BRDFs are defined. This motivates us to look for a basis defined on the hemisphere. Zernike polynomials have been used to express spherical aberration in optics. They are functions over the unit disk $D^2 \rightarrow R$.

Since, the hemisphere has the same topology as the disk, they can be extended to the hemisphere $H^2 \rightarrow R$. The Zernike functions on the hemisphere can be written as $K_n^m(\theta, \phi) = R_n^m(\sqrt{2}\sin(\theta/2)) az_m(\phi)$ where R is an appropriately normalized Zernike polynomial of degree at least m and az_m is a real version of the complex exponential, being given by $cos(m\phi)$ when $m > 0$, $\sqrt{2}$ when $m = 0$ and $-\sin(m\phi)$ when $m < 0$.

The extension to a product of hemispheres is very similar to that for spherical harmonics.

$$H_{mn}^{lk}(\vec{\omega_i}, \vec{\omega_o}) = N(K_m^l(\vec{\omega_i})K_n^k(\vec{\omega_o}) + K_n^k(\vec{\omega_o})K_m^l(\vec{\omega_i})) \tag{16}$$

Again, N is a normalization constant.

To maintain isotropy, we require that adding a constant phase $\triangle\phi$ to both incident and outgoing angles does not change the function. The ϕ dependence of the above basis function is the real part of $e^{Il\phi_i}e^{Ik\phi_o}$ plus the symmetrized version with the subscripts swapped. It can be seen that this should depend on $\phi_i - \phi_o$ for the additive phase to not play a role. Thus, $l + k = 0$. For isotropic BRDFs, only basis functions with $l + k = 0$ need be taken into account, so there are only three relevant indices instead of four.

Generalized Cosine Lobes: The previous models have started with a general well-known set of basis functions, and have manipulated them to be able to represent the BRDFs reasonably well. However, these linear models are often ill-equipped to deal with the special features of BRDFs, a point noted in our earlier discussion of using basis functions to represent BRDFs. The generalized cosine lobe model starts from a phenomenological model—the Phong model—and generalizes it to represent a fairly large class of common BRDFs. The phong model specular lobe can be viewed as

$$(\vec{\omega_i}^T R \vec{\omega_e})^n$$

where R is a 3x3 matrix that reflects the eye-vector ω_e about the surface normal. The idea is to generalize R to M, a general matrix. However, to preserve reciprocity, we require M to be symmetric. This still leaves a 6 parameter family of functions, but 3 parameters serve merely to rotate the entire lobe, leaving 3 interesting parameters. We can decompose $M = Q^T DQ$ where D is a diagonal matrix to write

$$f_r(\vec{\omega_i}, \vec{\omega_e}) = (\vec{\omega_i}^T Q^T DQ\vec{\omega_e})^n \tag{17}$$

The Q matrices merely serve to set the coordinate frame while the diagonal D matrix with parameters c_x,c_y,c_z actually has the interesting parameters. Together with n, this gives a phong lobe parameterized with 4 parameters. If z is the normal direction, the original phong model has $c_x = -1, c_z = 1$. A positive c_x represents retroreflection. Off-specular peaks and increased reflection at grazing angles are obtained by setting $c_x < -1, c_z = 1$. Thus, many features of real BRDFs can be represented using only a few parameters. By summing several lobes, more complex BRDFs can be represented.

Reparameterization

The BRDF is defined over a product of two hemispheres $H^2 \times H^2$. Therefore, four parameters are required, and the traditional parameterization is by incident angle (θ_i, ϕ_i) and exitant angle (θ_e, ϕ_e). Thus, the BRDF is a function of $(\theta_i, \phi_i, \theta_e, \phi_e)$. This parameterization is good for measurement since all variables correspond directly to physical quantities. However, it can miss several symmetries of the BRDF. For instance, there is the phenomenon of off-specular reflection at near-grazing angles, which seems rather strange in the conventional parameterization; refer to figure 6. Also, reciprocity—the symmetry between incident and reflected angles— is not explicit.

The reason for reparameterization is to provide a better understanding of the BRDF, to make the symmetries more explicit, and also to help in creating a more compact representation. A problem of the standard representation with respect to compactness is that there are peaks all over the distribution, since the specular direction varies with incident angle. We seek reparameterizations where variation occurs in a smaller number of dimensions, ideally only 1, and where the BRDF varies slowly with respect to the other parameters. The rest of the lecture briefly discusses various parameterizations.

Difference from Mirror Direction: If we make one of the variables the angle to the ideal mirror direction, we allow the specular peak to remain stationary at 0, thus reducing somewhat the complexity or dimensionality of the space. However, for anisotropic BRDFs, the specular lobes will spin or rotate as we move the incident angle. Retroreflective peaks are not localized, either. Off-specular peaks will remain, since the distribution has only been shifted.

Half-Angle Parameterization: We may also base our reparameterization on H, the *half-angle* midway between the incident and exitant angles. If the exitant angle is in the

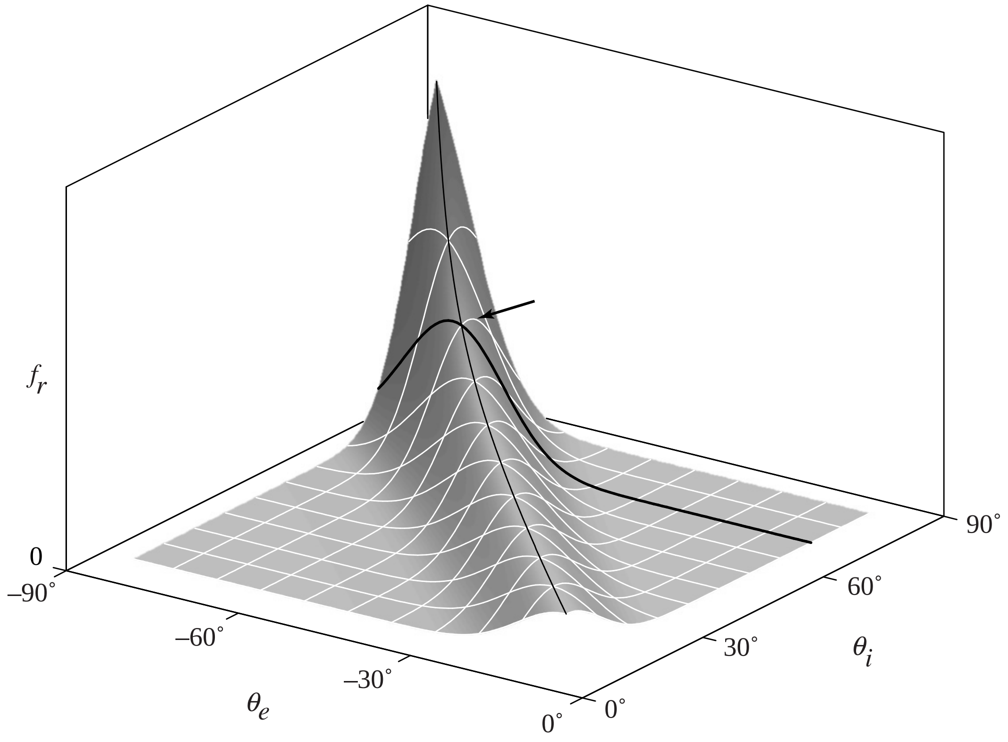

Figure 6: A depiction of off-specular reflection. There is a ridge where incident and reflected angles are equal, but the ridge is steeper toward grazing angles. For a given incident angle, this creates off-specular peaks. A half-angle reparameterization as discussed below is more natural and does not have these counter-intuitive results.

direction of mirror specular reflection, the half-angle coincides with the surface normal. Thus, the deviation of the half-angle from the normal direction, indicates the strength of the specular reflection, and is a fundamental quantity in many analytic BRDF models, which are exponential in $N \cdot H$ where N is the normal direction, or include a phong-like lobe $(N \cdot H)^n$. Furthermore, incident and exitant angles are explicitly symmetric with respect to the halfway vector. We write the BRDF as a function of (H, D) where D is a difference vector. The BRDF is often factorizable as a function of H times a function of D, where the D dependence takes care of Fresnel-effects at grazing angles, and where the H dependence takes care of the specular peak. There are many symmetries exposed by this parameterization. For instance, different specular peaks correspond to the same H for different values of D, and remain stationary in H as D is varied. The effects of the Fresnel term depend only on D, not H, and so off-specular reflection is folded into the model, and does not appear explicitly. For a given D, the peak is always at $H = 0$. The

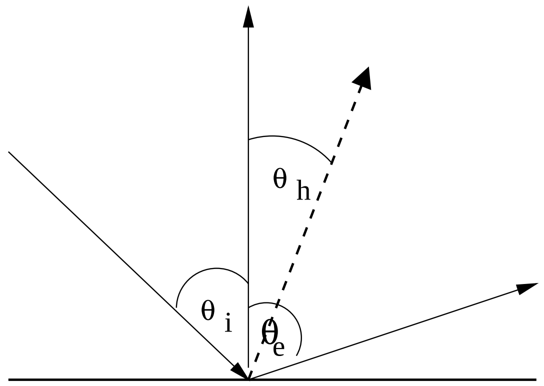

Figure 7: The half-angle, halfway between incident and exitant angles is a key quantity in BRDF representation

orientations of anisotropic peaks also remain fixed. Retroreflection for different angles corresponds to changing H while keeping D fixed, and is therefore also easily factorizable, and the retroreflective peaks remain stationary. In fact, retroreflection can be seen as essentially the dual to ordinary reflection, under an exchange of D and H.

Factorized Representations: Many BRDF models can be reparameterized so they can be written as the product of several terms, each of which depends on different parameters. Each of these product terms can be stored as a texture, with texture co-ordinates being the appropriate parameters. These parameters are usually computable in terms of dot-products of vectors generated by the graphics system e.g. $N \cdot H$. Therefore, such a representation is very efficient for hardware rendering. Simple dot-products index into a texture map, and the various terms are multiplied together to compute the appropriate products in hardware. This allows a wide variety of BRDFs to be efficiently rendered with standard graphics hardware.

Isotropic BRDFs: This paragraph introduces a more intuitive parameterization for isotropic BRDFs. In general, BRDFs are defined over a product of hemispheres $H^2 \times H^2$. However, isotropic BRDFs have an equivalence relation in that rotations of the surface about the normal should result in the same BRDF and so the topology is $H^2 \times H^2 \backslash E$ where E is the equivalence class. The usual parameterization is with $(\theta_i, \theta_e, \triangle\phi)$ where $\triangle\phi = \phi_r - \phi_i$. The goal is to map these 3 parameters into 3-space so as to preserve the topology of the BRDF space.

A first attempt is to use cylindrical coordinates with

$$\begin{aligned} r &= \theta_i \\ x &= r \cos \triangle\phi \\ y &= r \sin \triangle\phi \\ z &= \theta_e \end{aligned}$$

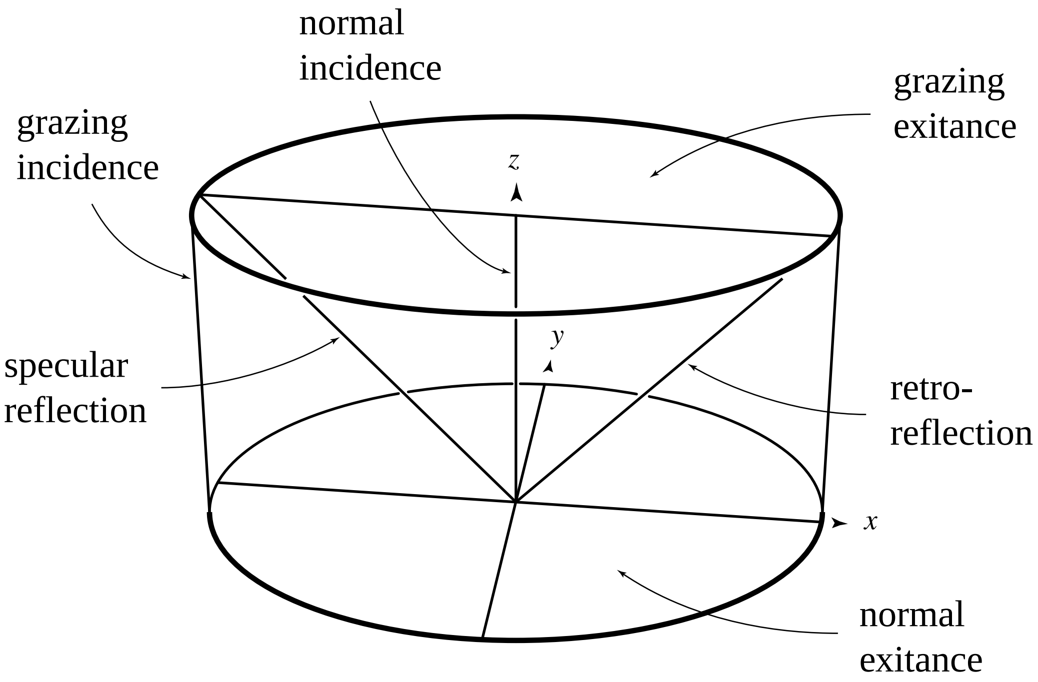

Figure 8: A first attempt at a 3-space mapping of isotropic BRDF parameters into cylindrical coordinates.

Some good properties of the above parameterization, depected in figure 8 are that normal incidence corresponds to the z axis and is described by only 1 parameter , θ_e, since only the angle of exitance matters. However, the reciprocity condition is a reflection about a cone and is not simple. Normal exitance maps to an entire plane (the xy plane), instead of a line as for normal incidence. Finally, for a constant incident direction, the parameter space is a cylinder, not a hemisphere.

A second, more sophisticated attempt is to use

$$\begin{aligned} r &= \sin\theta_i \sin\theta_e \\ x &= r\cos\triangle\phi \\ y &= r\sin\triangle\phi \\ z &= \cos\theta_e - \cos\theta_i \end{aligned} \tag{18}$$

Reciprocity now translates into a simple reflection about the origin. Normal incidence and exitance both correspond to lines since r vanishes, and so too do retro and specular reflection for which z vanishes.

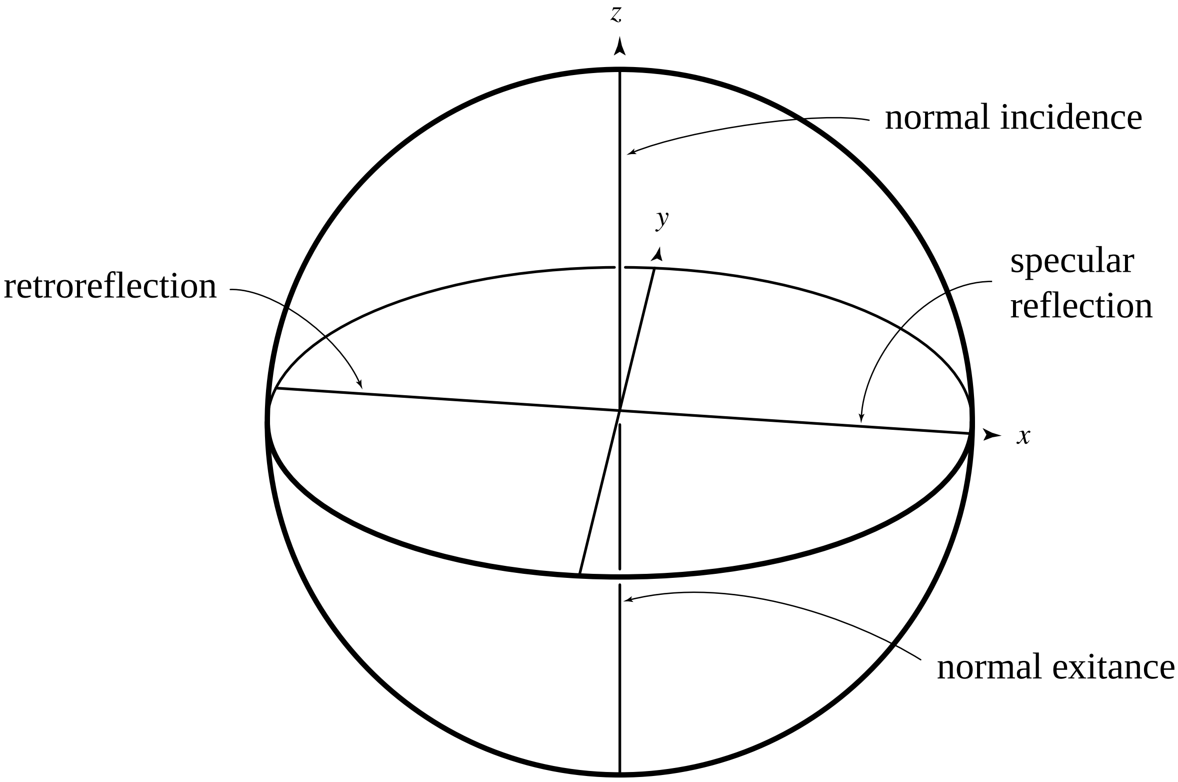

Figure 9: A second, more sophisticated attempt at a 3-space mapping of isotropic BRDF parameters into cylindrical coordinates.

4 Further Reading

This lecture has skimmed a lot of material fairly briefly. Introductory material on Radiometry and BRDFs can be found in the chapter on Rendering Concepts by Pat Hanrahan in Cohen and Wallace [1]. Nicodemus' article [8] first introduces the BRDF as a unified description of reflection. The derivation of Fresnel reflectance and a discussion of spherical harmonics can be found in a standard Physics textbook, for instance, [3]. BRDF modeling and representation has a long history. The original paper by Torrance-Sparrow [10] first introduces that model. Spherical harmonics have been used to simulate reflectance from microgeometry by Westin, Arvo and Torrance [11]. Koenderink and van Doorn [5] introduce the phenomenological model based on Zernike polynomials. The generalized cosine lobe model is introduced in a paper by LaFortune et al. [6]. The later papers also have pointers to some of the earlier literature. The half-angle parameterization is due to Rusinkiewicz [9]. BRDF factorizations have been explored by Heidrich and Seidel [2], and by Kautz and McCool [4]. Marschner's [7] thesis introduces the mappings we discuss for isotropic BRDFs. The BRDF plots shown during this lecture were done using a BRDF viewer written by Szymon Rusinkiewicz available at *http://graphics.stanford.edu/~smr/brdf/bv/*

References

[1] M. F. Cohen and J. R. Wallace. *Radiosity and Realistic Image Synthesis*. Academic Press, 1993.

[2] Wolfgang Heidrich and Hans-Peter Seidel. Realistic, hardware-accelerated shading and lighting. In *Proceedings of SIGGRAPH 99*, pages 171–178, August 1999.

[3] J.D. Jackson. *Classical Electrodynamics*. John Wiley, 1975.

[4] Jan Kautz and Michael D. McCool. Interactive rendering with arbitrary brdfs using separable approximations. In *Eurographics Rendering Workshop 1999*, June 1999.

[5] Jan J. Koenderink and Andrea J. van Doorn. Phenomenological description of bidirectional surface r eflection. *Journal of the Optical Society of America A*, 15(11):2903–2912, 1998.

[6] E. P. F. Lafortune, S. Foo, K. E. Torrance, and D. P . Greenberg. Non-linear approximation of reflectance functions. In *SIGGRAPH 97*, pages 117–126, 1997.

[7] Stephen R. Marschner. *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell University, 1998.

[8] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. Geometric considerations and nomenclature for reflectance. *National Bureau of Standards (US)*, October 1977.

[9] Szymon M. Rusinkiewicz. A new change of variables for efficient brdf representation. In *Eurographics Rendering Workshop 1998*, pages 11–22, June 1998.

[10] K. Torrance and E. Sparrow. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America*, 57:1105–1114, 1967.

[11] Stephen H. Westin, James R. Arvo, and Kenneth E. Torrance. Predicting reflectance functions from complex surfaces. In *SIGGRAPH 92*, 1992.