

A Practitioners' Assessment of Light Reflection Models

Peter Shirley

Computer Science Department
University of Utah
Salt Lake City, UT 84106 USA
shirley@cs.utah.edu

Brian Smits

Computer Science Department
University of Utah
Salt Lake City, UT 84106 USA
bes@cs.utah.edu

Helen Hu

Computer Science Department
University of Utah
Salt Lake City, UT 84106 USA
hhh@cs.utah.edu

Eric Lafortune

Program of Computer Graphics
Cornell University
Ithaca, NY 14853 USA
eric@graphics.cornell.edu

Abstract

We discuss the theory and practical issues behind creating reflection models to show the difficulty of the problem. We survey the current approaches towards reflection models for computer graphics to show that even for simple surfaces, the important issues are far from settled. We briefly discuss future directions for research. Finally, we present a case study of a particular type of light reflection that captures some important aspects of appearance for a limited class of materials with subsurface reflection.

1 Introduction

A goal of realistic image synthesis is to create images that evoke visual reactions similar to what a viewer would experience in the actual scenes, or more modestly, to create images of the scenes that might be captured by a camera or other sensor device. In either case, an accurate model of the scene is needed, describing salient aspects of shapes and materials that contribute to the final images. Typically, the material description is assumed to be the easiest part of creating an image [5]. However, it has been our experience that describing the materials can often be the most problematic part of image synthesis. In this paper we discuss why this is the case, and what issues need to be resolved to enable realistic images to be more easily generated.

Our goal in this paper is to provide a tutorial of reflection modeling terminology and methodology that is not made opaque by too much radiometric jargon, and to communicate the relative immaturity of reflection modeling technol-

ogy. We argue that this immaturity stems largely from barriers intrinsic to reflection modeling, rather than from a lack of previous work on the subject. We emphasize the issues that come up when using light reflection models to create realistic images. Dealing with these issues is vital when creating a rendering system, and we believe the community needs to simultaneously consider both theoretical and practical issues when developing reflection models. An application of this belief is shown in the case study reflection model developed in the paper.

In Section 2, we describe how light reflects from a smooth surface. We cover the complexities of reflection from rough surfaces in Section 3. We review the high-level issues related to the commonly-used reflection models and representations in Section 4. In Section 5 we provide a case study of a reflection model for smooth surfaces with subsurface scattering. This last section, where a relatively simple surface type is shown to be an unsolved problem, emphasizes the main point of the paper: even for simple surfaces, the important issues in reflection modeling are far from settled.

2 Properties of a smooth interface

The simplest case of reflection is for a smooth interface between two homogeneous materials. This situation is governed by the *Fresnel equations* which describe the ratio of reflected and transmitted energy as a function of incident direction, polarization, and properties of the two materials [37]. For simplicity we assume that one of the materials is vacuum or air, which have almost identical optical prop-

erties at a material interface.

A fundamental property of the Fresnel equations is that they predict different reflectance behavior for the two states of polarization. However, image synthesis practitioners almost never take polarization into account and assume random polarization everywhere. This should cause some unease, because in simple instances polarization can play an important role even when the emitted light is randomly polarized [40, 23]. To add to the unease, note that the light coming from the sky is partially polarized [19].

For engineering applications where accuracy is critical, polarization should be accounted for. For the rest of this paper we will assume that a world without polarized light would not look much different, so ignoring polarization will be acceptable for appearance-based applications. This is almost certainly not true in general, and should at some point be examined more carefully, however, ignoring polarization will probably not be the worst sin committed in the image synthesis process. We can ignore polarization by averaging the reflectance predicted by the two polarization cases. This will give us a reflectance that depends only on wavelength and incident angle: $R_f(\theta, \lambda)$ as well as the optical properties of the material.

The optical properties of the material that we are concerned with are its refractive index $n(\lambda)$ and its extinction coefficient $k(\lambda)$, both of which depend on wavelength λ . Materials fall into two broad categories: dielectrics, where $k(\lambda)$ is zero for λ in the visible range, and conductors (metals), where $k(\lambda)$ is non-zero. The equations are somewhat bulky, particularly in the case of a conductor, but are straightforward to implement.

It is occasionally possible to look up the optical properties for a material [27], however such data is hard to find for materials in typical scenes. Where experimental data exists, it often varies widely between sources. It is often easier to obtain, or at least estimate, reflectance data for normal incidence angle, $R_f(0, \lambda)$. Cook and Torrance used this to their advantage to approximate $R_f(\theta, \lambda)$ from $R_f(0, \lambda)$. This approximation was refined by Schlick [33]:

$$R_f(\theta, \lambda) \approx R_f(0, \lambda) + (1 - \cos \theta)^5 (1 - R_f(0, \lambda)). \quad (1)$$

Note that $R_f(90^\circ, \lambda) = 1$, which is important for achieving correct appearance at grazing angle. One might argue that the Fresnel equations are simple enough that it is unwise to approximate them. However, even when the data is available, ignoring the role of polarization introduces error likely to be much larger than that added by the Schlick approximation. Because the approximation is more efficient, stable, and easier to debug than the full Fresnel equations, we advocate its use without hesitation when ignoring polarization.

The qualitative effects of the Fresnel equations are illustrated in Figure 1, where R_1 , R_2 , and R_3 are the fractions of




Figure 1. The fraction of energy specularly reflected (RE) and energy transmitted $((1-R)E)$ changes with incident angle, but at each angle the sum of the specularly reflected and transmitted energy equals the incident energy (E).

incident light specularly reflected, and the size of these fractions is shown by the length of the reflected vector. Note that the specular reflectance increases as the angle of incidence diverges from the surface normal. For a metal, the transmitted energy will be absorbed.

This is perhaps the only reflection model where the physics are well-understood and which applies to real situations. We can render an object with known parameters such as copper, and have some confidence in the result. However, the object will look like a unnaturally new copper coin. What is missing is corrosion, scratches, and patina. These can be simulated [8], but then the simple Fresnel equations do not apply. Also, uncorroded metal is likely to be an alloy, so we probably cannot easily get data for $k(\lambda)$ or $n(\lambda)$.

In typical image synthesis applications, surfaces are neither smooth nor clean. However, a case where they are smooth is when explicitly modeling the micro-geometry that creates a rough appearance; we will return to this topic in Section 5 where we create a picture using this technique. A case where surfaces might be both smooth *and* clean is glass and water. Here the data is readily available and we can get a very realistic appearance in practice¹.

In summary, the Fresnel equations are elegant and useful in some instances, but do not directly apply to most surfaces we need to model. Also, we advocate using Schlick’s approximation whenever the Fresnel equations are used without polarization.

¹When implementing a glass model, note that light is attenuated by impurities as it travels through glass, and this attenuation is important to model for realistic appearance. For homogeneous impurities as is found in typical glass, this attenuation is exponential:

$$I(t, \lambda) = I(0, \lambda) \exp(-a(\lambda)t),$$

where $I(t, \lambda)$ is the intensity of the beam at distance t from the interface, and $a(\lambda)$ is an attenuation coefficient that is typically hard to find. In practice we reverse-engineer $a(\lambda)$ by eye.




Figure 2. At a macroscopic scale the law of reflection is not followed but at point (u, v) we see that a single micro-facet is involved and the law of reflection does apply.




Figure 3. When light hits a set of micro-facets it creates an outgoing distribution. Once r is large enough, this distribution can be considered continuous.

3 Scattering functions

The reflection from the surface can be described by how light incident on the surface is scattered into a continuum of directions. This scattering function can be described mathematically, and there are several strategies for deriving a mathematical model for reflection.

3.1 The complexity of reflection

We now explore the complex nature of reflection functions by looking at the probabilistic behavior of a single photon-like light particle striking an opaque material. The particle is characterized by its incident direction (θ, ϕ) and a wavelength λ . We look only at geometric optics for this discussion, and gloss over many of the details in an attempt to arrive at a minimal set of parameters needed for describing a material.

A perfectly smooth, non-transparent surface absorbs or reflects the light striking the surface based on the properties of the material, the wavelength of the particle and the incident direction, as described in Section 2. This results in a scattering function with variables (θ, λ) . The outgoing direction is completely determined by the incoming direction (Figure 1).

Almost no real surface is completely smooth so we need a more descriptive representation for the surface. We could think of the surface as some sort of irregular height field. At each point the surface appears smooth, but the orientation at the point varies over the surface. Now we have a reflection model that is dependent upon surface position as well, giving us reflection model parameters $(u, v, \theta, \phi, \lambda)$ (Figure 2). This form is not very useful, as we are usually concerned with the average behavior over a small region of the surface. We can remove the (u, v) terms by looking at the aggregate behavior over a region of size r based on some statis-

tical model of the surface orientation. This statistical model is usually assumed to be *isotropic*, meaning the distribution of orientations is independent of the incident polar angle ϕ . (A surface with an oriented structure, such as typical brushed or milled metal, is called *anisotropic*.) The aggregate behavior of the surface results in light being scattered in many directions (Figure 3). This scattering distribution is known as *spread* reflection (sometimes called *directional diffuse* or *glossy* reflection). If we ignore the choice of statistical model and its parameters, we end up with the following independent variables for the model: $(r, \theta, \theta', \phi - \phi', \lambda)$, where (θ', ϕ') represents the outgoing direction. Here $\phi - \phi'$ is the polar angle relative to the incident direction. We can remove the restriction to isotropic surfaces by adding the ϕ term to completely describe the incident direction giving us $(r, \theta, \phi, \theta', \phi', \lambda)$.

We now have enough variables to describe the reflection properties of a uniform surface at some scale of interest at a given point. Very few real world materials are uniform, however. Properties such as roughness, color, composition, orientation of fine structures may all change as the point changes, requiring us to re-introduce the (u, v) terms into our set of variables. The final set of variables for describing a material are $(u, v, r, \theta, \phi, \theta', \phi', \lambda)$. Note that this discussion only describes reflection from the interface between the material and air. In dielectrics, light enters the material, propagates through it, possibly scattering as a result of impurities in the material, and is either transmitted out the far side of the material or reflected back out of the material. Light that escapes from the material does not necessarily do so at the same point where it went in, although for most materials the two points will be very close together. It is usually assumed that all sub-surface scattering and transmission can be treated as happening at (u, v) .

Even with these assumptions, a reflection model has eight degrees of freedom, making it difficult to formulate and represent. In order to simplify things, the graphics community

tends to remove the position (u, v) and treat this as a texture mapping operation. The size term r is pushed either into texture sampling or into anti-aliasing of the image, depending upon whether it is color or geometry that falls below the implicit r of the renderer. This simplification is useful for formally analyzing the properties of a reflection model at some point and scale, as some important properties are independent of position and scale issues. The simplification has been harmful because it has caused reflection models, texture mapping, and fine scale geometry to be thought of as independent issues when in fact they are very tightly linked.

3.2 Formal representations

To do any simulation using spread reflection, we need to develop a convention for describing it mathematically to be able to determine its properties. As described above, we ignore position and size for this analysis. We could look at a concept similar to the reflectance of the last section. The simplest way to do this is to describe the ratio of incident light power at a given angle to outgoing light power integrated over all outgoing angles:

$$R(\theta, \phi, \lambda) = \frac{\text{power incident at } \lambda \text{ from } (\theta, \phi)}{\text{total outgoing power at } \lambda}.$$

This quantity is often called *directional hemispherical reflectance*. But to describe the directional behavior of spread reflectance, we need some function over the outgoing directions (θ', ϕ') . One way to do this is to imagine a hypothetical photon-like light particle of wavelength λ that is incident to the surface from direction (θ, ϕ) . This particle will be absorbed by the surface with probability $1 - R(\theta, \phi, \lambda)$. If it is not absorbed it will be scattered in some outgoing direction (θ', ϕ') with probability density function $p(\theta, \phi, \theta', \phi', \lambda)$. We call $p(\theta, \phi, \theta', \phi', \lambda)$ the *scattering probability density function*, or *SPDF* for short. Note that there is no standard term for the SPDF. Also note that the SPDF refers to a potentially different probability density function for each (θ, ϕ, λ) . Readers that prefer to think deterministically can think of the product of R and p as an energy density function without losing any information.

Given the directional hemispherical reflectance and the SPDF for a particular isotropic metal, we have completely specified its reflectance behavior. We could change the definition in several ways without changing the completeness. The most common way to do this is to multiply the directional hemispherical reflectance, the SPDF, and reciprocal of $\cos \theta'$ into a single product:

$$\rho(\theta, \phi, \theta', \phi', \lambda) = \frac{R(\theta, \phi, \lambda)p(\theta, \phi, \theta', \phi', \lambda)}{\cos \theta'}.$$

There is no gain or loss of information in this transformation, but the form of ρ is convenient for many image synthesis and

measurement applications as it relates incoming irradiance to outgoing radiance, quantities that are often computed and measured. The standard name for ρ is the *bidirectional reflectance distribution function*, or *BRDF* for short. One fundamental property of a BRDF is that

$$\rho(\theta, \phi, \theta', \phi', \lambda) = \rho(\theta', \phi', \theta, \phi, \lambda).$$

This property is known as *reciprocity*; incident and outgoing directions can be interchanged. Although almost all of the reflection literature is written in terms of BRDF, it is often convenient to think in terms of the SPDF (or the outgoing energy density function) when intuition can play a role.

The directional hemispherical reflectance can be written in terms of the BRDF:

$$R(\theta, \phi, \lambda) = \int_0^{2\pi} \int_0^{\frac{\pi}{2}} \rho(\theta, \phi, \theta', \phi', \lambda) \cos \theta' \sin \theta' d\theta' d\phi'.$$

This equation is equivalent to saying the SPDF must have unit volume, which is required for any probability density function. For a physically valid material model, it is necessary that $R(\theta, \phi, \lambda) \leq 1$; it must be *energy conserving*.

Many rendering algorithms rely on these properties, so it is important when using these renderers that BRDF models both conserve energy and be reciprocal. Lewis calls BRDF models that are faithful to these constraints *physically plausible* [22]. As an example we see that the Lambertian BRDF:

$$\rho_L(\theta, \phi, \theta', \phi', \lambda) = \frac{R_d(\lambda)}{\pi},$$

is both reciprocal and energy conserving² provided the reflectance $R_d(\lambda)$ is less than one for all wavelengths.

4 Reflection models and representation

In the last section we reviewed the definition of the BRDF, and the main constraints on the BRDF: energy conservation for every incident (θ, ϕ) , and reciprocity. To implement rendering software, some actual BRDF model or models must be deployed. In this section we discuss some of the options with which the implementor is presented. Readers can find more extensive reviews of reflection models on the world-wide-web [9, 32], and in Glassner's book [12].

4.1 Acquiring BRDF data

To acquire BRDF data for a specific material we can measure the BRDF directly a large number of angles and wavelengths. Such measurements are done with a device called

²The division by π is necessary because the BRDF is integrated over the cosine-weighted hemisphere.

gonioreflectometers. These are challenging devices to build and calibrate [10]. Assuming we can get the measurements, the next issue is a good representation for the data, as will be discussed in Section 4.2.

Measuring BRDFs is a slow and expensive process that requires having a sample of the material on hand. This makes measurement infeasible in many cases. Sometimes the physical properties and micro-geometry of the desired surface can be determined and modeled. In this case it is possible to compute a BRDF through simulating the scattering of light by the surface, in effect, using the computer to simulate a gonioreflectometer. This is usually done by casting rays at the surface and computing the distribution of the resulting scattered rays [4, 42]. These methods are limited by the difficulty of getting the data and their restriction to geometric optics. They are capable of simulating very complex BRDFs due to highly structured and complicated geometry.

As mentioned earlier, dielectrics transmit part of the incident light, which is then scattered and absorbed by impurities. The scattered light may eventually leave the surface again. Computing the resulting distribution is even harder than for metallic surfaces, because it not only depends on the properties of the interface, but on the properties of the inclusions in the material as well. A typical example is latex paint, where colored particles are suspended in a transparent substrate. The scattered component is usually approximated by a Lambertian term. Hanrahan and Krueger [14], Gondek et al. [13] and Schramm et al. [34] have simulated subsurface scattering using Monte Carlo simulations, obtaining BRDFs of materials such as leaves, skin and paints. A problem with these simulation approaches is that the substrate geometric data is very difficult to acquire, but their generality and simplicity are appealing.

4.2 BRDF representations

Given the sampled data for a BRDF, we need to fit it to a particular model or representation. An example of a model is the Cook-Torrance model [6], where low-level physics is used to infer a natural representation for the BRDF. An example of a representation is a piece-wise linear table. A model is of course also a representation, so we will call all BRDF storage strategies representations for the remainder of this discussion.

For the isotropic BRDF case at a particular point and sample size, a representation is parameterized by four variables: $(\theta, \theta', \phi - \phi', \lambda)$. Typically this fit is done separately at each sample for λ , so it becomes a three dimensional problem. Many different representations have been examined. The simplest is to create an explicit three-dimensional table for each wavelength. This strategy is memory intensive and how the interpolation should be done in the ta-

ble is not clear. Because $(\theta', \phi - \phi')$ is a coordinate system on the hemisphere, spherical harmonics have been employed [4, 38, 42]. Zernike polynomials, defined on the hemisphere rather than the sphere, have also been used [7]. The danger of using these orthogonal bases is ringing, which can cause visual artifacts. Spherical wavelets have been employed to attempt to get the cleanliness of tables and the compactness of polynomial representation [35]. The practical details of this approach have been developed by researchers at the University of British Columbia, as is detailed in Lalonde's forthcoming dissertation [21]. This is a promising approach, although somewhat daunting to an implementer. Two flaws limit the above approaches. They require expensive data acquisition and storage, and are not able to allow natural user modification. Rather than using a general-purpose representation, more compactness and efficiency can be achieved by using a special-purpose low-dimensional representation. This representation can be derived from low-level physics and fit with either direct micro-geometric parameters, or fit with measured BRDF data. Examples of this approach for isotropic surface-reflection are the models of Blinn [3], Cook and Torrance [6], and He et al. [17]. For the anisotropic case there are the model of Kajiyama [18], and Poulin and Fournier [31]. For rough surfaces with subsurface reflection there is the model of Oren and Nayar [26]. The resulting representations are defined by a limited number of parameters, which in some cases can be measured directly. However, they are often restricted by their assumptions and approximations. For example, most do not account for secondary reflection (reflection from two or more micro-facets for a single light particle), or the vertical blurring that occurs in reflections (see Figure 4) for example.

A final approach to representation is top-down, where intuition and experience is used to develop an expressive and convenient representation with a small number of free parameters. This approach was first employed in graphics by Phong [30], and was generalized by Ward [41] and Schlick [33]. A more general approach has been implemented by Lafortune [20], who has achieved good data fits for the isotropic case. We think it likely that this avenue is likely to continue to be fruitful.

4.3 Practical issues

In computer graphics applications, reflectance functions are usually part of a larger illumination simulation framework. For that purpose, the functions should be efficient to evaluate. For instance, the He model is sophisticated but computationally expensive, because its evaluation includes a slowly converging series. Researchers therefore often look for approximations that are more efficient [16, 33]. Reflectance models for which entire illumination compu-

tations can be carried out analytically, rather than numerically, present an interesting option, as Arvo [1] showed for the Phong BRDF. In Monte Carlo lighting simulations it is often advantageous if one can sample according to the BRDF, when tracing random walks of light particles through a scene. An example of this is discussed in Section 5.1.1.

Another aspect is memory use. Most analytic BRDF models are compact, as they are defined by only a few surface parameters. More general representations like tables or spherical harmonics usually lead to large memory requirements, to avoid sampling problems and ringing. Specular BRDFs are difficult to represent because of their narrow specular peaks. Wavelets may offer a solution, as they allow refinement of the representation in highly varying regions of the function.

The scale at which BRDFs should be modeled is still largely an open problem. Scratches on a surface may be visually important when viewed from arms length, but they become invisible from larger distances. Westin et al. [42] recognized the need for different geometric models and reflectance models at different scales. Becker and Max [2] presented techniques to smoothly transit between different representations. Pharr et al. [28] discussed efficient techniques to store and render the geometric detail.

Measuring how the BRDF varies over a surface (as a function of (u, v) in our terminology) is largely an unaddressed problem. Dana et al. have begun a database of measurements of surfaces with variable BRDF, and their description of the task makes it clear that very concerted efforts are needed to make any progress in this area [7].

5 Case study: matte/specular reflection

Suppose we were to create a realistic renderer that modeled micro-geometry explicitly, and used very simple reflection models for the smooth facets. This type of renderer has been shown to be practical for at least some scenes by Pharr et al. [28, 29]. For metal and glass surfaces, we would use the Fresnel equations directly. For dull surfaces we could use the Lambertian model because in many applications the errors that result from using the model to approximate matte reflection are not visually significant [24]. But what model should we use for *matte/specular* (sometimes called *polished*) materials that are smooth but have significant subsurface scattering? In this section we argue that current “general” models do not model this matte/specular case well, and develop a special-purpose model that we call the *coupled* model.

Matte/specular materials such as plastics or polished woods have reflection governed by Fresnel equations at the surface, and scattering within the subsurface. An example of this reflection can be seen in the tiles in the photographs




Figure 4. Photographs illustrating appearance of polished tiles. Specularity increases and matte behavior fades as viewing angle approaches grazing.

in Figure 4. Also note that the blurring in the specular reflection is mostly vertical. This is due to the compression of apparent bump spacing in the view direction. This effect causes the vertically-streaked reflections seen on lakes on windy days[25], and can either be modeled using explicit micro-geometry and a simple smooth-surface reflection model or by an as yet unrealized more general model that accounts for this asymmetry.

We could use the traditional Lambertian-specular model for the tiles, which uses two constants to modulate a constant and specular component of the BRDF [30]. In standard radiometric terms, this idea is expressed as:

$$\rho(\theta, \phi, \theta', \phi', \lambda) = \frac{R_d(\lambda)}{\pi} + R_s \rho_s(\theta, \phi, \theta', \phi'),$$

where $R_d(\lambda)$ is the hemispherical reflectance of the matte term, R_s is the specular reflectance, and ρ_s is the normalized specular BRDF (a weighted Dirac delta function on the sphere). This equation is a simplified version of the BRDF where R_s is independent of wavelength. This independence causes a highlight that is the color of the luminaire, so a polished rather than a metal appearance will be achieved [6]. Ward suggests that in order to conserve energy, $R_d(\lambda) + R_s \leq 1$ [41]. However, such models with constant R_s fail to show the increase in specularity for steep viewing angles.

This is the key point: in the real world the relative proportions of matte and specular appearance change with viewing angle.

He et al. suggest using the Fresnel equation for the coefficient of the specular term [17], but do not address the subsurface term’s angular behavior because this model is intended primarily to simulate surface physics. Since the Fresnel term of the He model goes to one for $\theta = 90^\circ$, the Lambertian term would have to be set to zero to enforce energy conservation for all incident (θ, ϕ) . Because in the case of smooth polished surfaces energy conservation is important to us, and the explicit spread reflection is not (we assume that we will model it with micro-geometry for this discussion), the He model is not appropriate for our purposes.

Shirley attempted to simulate the change in the matte appearance with angle by explicitly dampening $R_d(\lambda)$ as R_s increases [36]:

$$\rho(\theta, \phi, \theta', \phi', \lambda) = \frac{R_f(\theta)\rho_s(\theta, \phi, \theta', \phi') + R_d(\lambda)(1 - R_f(\theta))}{\pi},$$

where $R_f(\theta)$ is the Fresnel reflectance for a polish-air interface. The problem with this equation is that it is not reciprocal, as can be seen by exchanging θ and θ' which changes the value of the matte dampening factor because of the multiplication by $(1 - R_f(\theta))$. The specular term, a scaled Dirac delta function, is reciprocal, but this does not make up for the non-reciprocity of the matte term. Because Shirley’s BRDF is not physically plausible, it will cause some rendering methods to have ill-defined solutions.

Schlick proposed a general reflectance model tuned for efficiency. In his model a matte/specular surface could have constant Lambertian and specular coefficients, or the Fresnel reflectance could be used. In the latter case, which he calls a *double* surface the BRDF becomes:

$$\rho(\theta, \phi, \theta', \phi', \lambda) = \frac{R_f(\theta)\rho_s(\theta, \phi, \theta', \phi') + R_d(\lambda)(1 - R_f(\alpha))}{\pi},$$

where $R_d(\lambda)$ is a matte coefficient and α is half the angle between incident and outgoing directions. However, this form does not conserve energy for all incident angles: for example, at $\theta = 90^\circ$ the specular reflectivity goes to one, and the fraction of the hemispherical reflectance is still above zero (e.g., plug in $\theta' = 0$). So that part of the Schlick model is not appropriate for our purposes.

In review of our attempt to model a smooth matte-specular surface, none of the commonly used general models is appropriate for our purposes. The Lambertian-specular and Ward models do not have the appropriate angular trade-off between the matte and specular term. For smooth surfaces, the He model has a constant subsurface

term that must be set to zero if energy is to be conserved. The Schlick model either defaults to the Lambertian-specular model, or it accounts for the Fresnel equation effects but does not conserve energy. One reason these models fail for this case is that they are all intended to model spread reflection for a variety of material types. In our case-study we do not need this generality, so we can develop a simple model that is customized for this narrow class of materials and captures the angular-dependent relationship of the matte and specular coefficients. This model uses a physically-based specular coefficient derived from the Fresnel equations, and a heuristic matte component of the BRDF. To our knowledge, it is the first model that produces the matte/specular tradeoff while remaining reciprocal and energy conserving. Because the key feature of the new model is that it couples the matte and specular scaling coefficients, we will hereafter refer to it as the *coupled* model.

5.1 A coupled model

Surfaces which have a glossy appearance are often a clear dielectric, such as polyurethane or oil, with some subsurface structure. The specular (mirror-like) component of the reflection is caused by the smooth dielectric surface and is independent of the structure below this surface. The magnitude of this specular term is governed by the Fresnel equations.

The light that is not reflected specularly at the surface is transmitted through the surface. There either it is absorbed by the subsurface or it is reflected from a pigment or a subsurface and transmitted back through the surface of the polish. This transmitted light forms the matte component of reflection. Since the matte component can only consist of as much light as is transmitted, it will naturally decrease in total magnitude for increasing angle (Figure 5).

Figures 1 and 5 illustrate why the magnitudes of the matte and specular behaviors are coupled; because they partition the incident energy, one must decrease as the other increases. This inverse relationship is the fundamental behavior that the traditional Lambertian-specular model does not capture.

To avoid choosing between physically plausible models and models with good qualitative behavior over a range of incident angles, we note that the Fresnel equations that account for the specular term, $R_f(\theta)$, are derived directly from the physics of the dielectric-air interface. Therefore the problem must lie in the matte term. We could use a full-blown simulation of subsurface scattering as implemented by Hanrahan and Krueger [14], but this technique is both costly and requires detailed knowledge of subsurface structure, which is usually neither known nor easily measurable. Instead, we can modify the matte term to be a simple approximation that captures the important qualitative angular behavior shown in Figure 4. Let us assume that the matte




Figure 5. Viewing along the same incident angles as in figure 2, the fraction of energy in the matte term is at most what is transmitted and thus varies with incident angle. The energy of the matte component is spread over many outgoing directions, as indicated by the grey lines. The length of each grey line indicates the relative energy near that direction.

term is not Lambertian, but instead is some other function that depends only on θ , θ' and λ : $\rho_m(\theta, \theta', \lambda)$. We discard behavior that depends on ϕ or ϕ' in the interest of simplicity. We try to keep the formulas reasonably simple because the physics of the matte term is complicated and sometimes requires unknown parameters. We expect the matte term to be close to constant, and roughly rotationally symmetric, as is argued in He’s dissertation [15].

An obvious candidate for the matte component $\rho_m(\theta, \theta', \lambda)$ that will be reciprocal is the *separable* form $k R_m(\lambda) f(\theta) f(\theta')$ for some constant k and matte reflectance parameter $R_m(\lambda)$. We could merge k and $R_m(\lambda)$ into a single term, but we choose to keep them separated because this makes it more intuitive to set $R_m(\lambda)$ which must be between 0 and 1 for all wavelengths. Separable BRDFs have been shown to have several computational advantages [11], which suggests the separable model:

$$\rho(\theta, \phi, \theta', \phi', \lambda) = R_f(\theta) \rho_s(\theta, \phi, \theta', \phi') + k R_m(\lambda) f(\theta) f(\theta'). \quad (2)$$

We know that the matte component can only contain energy not reflected in the surface (specular) component. This means that for $R_m(\lambda) = 1$, the incident and reflected energy are the same, which suggests the following constraint on the BRDF for each incident θ and λ :

$$R_f(\theta) + 2\pi k f(\theta) \int_0^{\frac{\pi}{2}} f(\theta') \cos \theta' \sin \theta' d\theta' = 1. \quad (3)$$

We can see that $f(\theta)$ must be proportional to $(1 - R_f(\theta))$. If we assume matte components that absorb some energy have the same directional pattern as this ideal, we get a BRDF of the form:

$$\rho(\theta, \phi, \theta', \phi', \lambda) = R_f(\theta) \rho_s(\theta, \phi, \theta', \phi') + k R_m(\lambda) [1 - R_f(\theta)] [1 - R_f(\theta')].$$

This is similar to a BRDF model used in the sensor community, although the constants used in that model do not have

the normalization properties we desire [9]. We could now insert the full form of the Fresnel equations to get $R_f(\theta)$ and then use energy conservation to solve for constraints on k . Instead we will use the approximation discussed in Section 2. This implies that

$$f(\theta) \propto (1 - (1 - \cos \theta)^5).$$

Applying Equation 3 gives:

$$k = \frac{21}{20\pi(1 - R_0)}. \quad (4)$$

The full coupled BRDF is then:

$$\rho(\theta, \phi, \theta', \phi', \lambda) = [R_0 + (1 - \cos \theta)^5 (1 - R_0)] \rho_s(\theta, \phi, \theta', \phi') + k R_m(\lambda) [1 - (1 - \cos \theta)^5] [1 - (1 - \cos \theta')^5]. \quad (5)$$

To test this new coupled model we created an approximate geometric model of the scene in Figure 4, where each tile has a displacement map that roughly corresponds to the subtle roughness of the real tiles. We took two photographs of the test scenes from different view angles, resulting in two different values of θ .

The results of running the coupled model is shown in Figure 6. Note that for the high viewpoint the specular reflection is almost invisible, but is clearly visible in the low-angle photograph image, while the matte behavior is less obvious.

For reasonable values of refractive indices, the R_0 is limited to approximately the range 0.03 to 0.06 (the value $R_0 = 0.05$ was used for the figures). The value of R_s in a traditional Phong model is harder to choose because it must typically be tuned for viewpoint in static images, and tuned for a particular camera sequence for animations. Thus, the coupled model is easier to use in a “hands-off” mode.

We did not attempt to mimic all subtleties of geometry exactly, so the reader should concentrate on the gross appearance features in the rendered and photographic images of the model. These images were produced using a Monte Carlo path tracer.

5.1.1 Monte Carlo sampling

It is not enough just to come up with an expression for the BRDF. We also need some computational tools if it is to be useful for statistical simulations. Given an incident angle θ and a BRDF, we often want to scatter a reflected ray whose direction follows the outgoing energy distribution indicated by the BRDF. In the case of the coupled BRDF presented in this section, one can first choose probabilistically between a reflection governed by the matte term and the specular term.



Figure 6. Renderings of polished tiles using coupled model.

If the matte term of Equation 5 is chosen, then we use a probability density function p that is proportional to the cosine-weighted BRDF for the given incident θ :

$$p(\theta', \phi') \propto [1 - (1 - \cos \theta')^5] \cos \theta'.$$

Because p does not depend on ϕ' we can choose ϕ' uniformly given a canonical (uniform on $[0,1)$) random number ξ' : $\phi' = 2\pi\xi'$. Because the differential measure on the sphere is $\sin \theta' d\theta' d\phi'$, we can write down a one-dimensional probability density function $q(\theta')$:

$$q(\theta') \propto [1 - (1 - \cos \theta')^5] \cos \theta' \sin \theta'.$$

Using standard techniques we can transform a second canonical random number ξ into a θ' ([39], pp.184-185), but this requires solving for a root of a high-degree polynomial in $\cos \theta'$:

$$\xi = \frac{7}{2} \cos^3 \theta' - \frac{21}{4} \cos^4 \theta' + \frac{21}{5} \cos^5 \theta' - \frac{7}{4} \cos^6 \theta' + \frac{3}{10} \cos^7 \theta'. \quad (6)$$

The solution to this equation that gives a $\xi \in [0, 1)$ can be approximated by the following rational function:

$$\cos \theta' = \frac{38295.9\xi^3 + 13002.4\xi^2 + 154.4\xi + 0.0114813}{16520.2\xi^3 + 33596.4\xi^2 + 1483.57\xi + 1}.$$

This equation is accurate to within 1.2% for all ξ . If more accuracy is needed, the equation provides an excellent initial guess for a Newton iteration to approximate $\cos \theta'$ to machine precision in only a few iterations.

6 Conclusion

We hope to have communicated three things in this paper. First, reflectance modeling is a hard problem that has many competing constraints. Second, that there is currently no general-purpose reflectance model that is suitable for all or even most situations. This was emphasized by the case study of ideal smooth/subsurface reflection, where the popular models all broke down. Third, that where reflection modeling is headed is unclear. It is possible that special purpose reflection modeling such as our case-study is the wave of the future. It is also possible that there is an as yet undiscovered general purpose model that will suffice. The possibly orthogonal issue of how scale should be handled is also an open question.

Acknowledgments

Thanks to Ken Torrance, Don Greenberg, Jon Blockson, Sing Foo, Bruce Walter, Steve Marschner, Dan Kartch, Dave Weinstein, Dean Brederson, and Lisa Durbeck for many helpful comments. This work was supported by the NSF Science and Technology Center for Computer Graphics and Scientific Visualization (ASC-8920219) and by NSF CCR-9401961.

References

- [1] J. Arvo. Applications of irradiance tensors to the simulation of non-lambertian phenomena. In *Proceedings of SIGGRAPH '95*, pages 437–438, August 1995.
- [2] B. G. Becker and N. L. Max. Smooth transitions between bump rendering algorithms. *Computer Graphics*, pages 183–190, August 1993. ACM Siggraph '93 Conference Proceedings.
- [3] J. Blinn. Simulation of wrinkled surfaces. *Computer Graphics*, 12(3):286–292, August 1978. ACM Siggraph '78 Conference Proceedings.
- [4] B. Cabral, N. Max, and R. Springmeyer. Bidirectional reflectance functions from surface bump maps. *Computer Graphics*, 21(4):273–282, July 1987. ACM Siggraph '87 Conference Proceedings.
- [5] M. F. Cohen. Is image synthesis a solved problem? In *Proceedings of the Third Eurographics Workshop on Rendering*, pages 161–167, 1992.
- [6] R. L. Cook and K. E. Torrance. A reflectance model for computer graphics. *Computer Graphics*, 15(3):307–316, August 1981. ACM Siggraph '81 Conference Proceedings.

- [7] K. Dana, B. van Ginneken, S. Nayar, and J. Koenderink. Reflectance and texture of real-world surfaces. Technical Report CUCS-046-96, Computer Science Department, Columbia University, 1996.
- [8] J. Dorsey and P. Hanrahan. Modeling and rendering of metallic patinas. In *SIGGRAPH 96 Conference Proceedings*, pages 387–396, 1996.
- [9] K. Ellis. Reflectance phenomenology and modeling tutorial. <http://www.erim.org/on-line-docs/GUIDE/guide.frm.html>, 1994. Report of the Environmental Research Institute of Michigan.
- [10] S.-C. Foo. A gonioreflectometer for measuring the bidirectional reflectance of material for use in illumination computation. Master's thesis, Cornell University, Ithaca, New York, June 1997.
- [11] A. Fournier. Separating reflection functions for linear radiosity. In *Proceedings of the Sixth Eurographics Workshop on Rendering*, pages 296–305, June 1995.
- [12] A. S. Glassner. *Principles of Digital Image Synthesis*. Morgan-Kaufman, San Francisco, 1995.
- [13] J. S. Gondek, G. W. Meyer, and J. G. Newman. Wavelength dependent reflectance functions. In *Computer Graphics*, pages 213–220, 1994. ACM Siggraph '94 Conference Proceedings.
- [14] P. Hanrahan and W. Krueger. Reflection from layered surfaces due to subsurface scattering. *Computer Graphics*, pages 165–174, 1993. ACM Siggraph '93 Conference Proceedings.
- [15] X. D. He. *Physically-Based Models for the Reflection, Transmission and Subsurface Scattering of Light by Smooth and Rough Surfaces, with Applications to Realistic Image Synthesis*. PhD thesis, Cornell University, Ithaca, New York, Jan. 1993.
- [16] X. D. He, P. O. Heynen, R. L. Phillips, K. E. Torrance, D. H. Salesin, and D. P. Greenberg. A fast and accurate light reflection model. In *Computer Graphics (SIGGRAPH '92 Proceedings)*, pages 253–254, 1992.
- [17] X. D. He, K. E. Torrance, F. X. Sillion, and D. P. Greenberg. A comprehensive physical model for light reflection. *Computer Graphics*, 25(4):175–186, July 1991. ACM Siggraph '91 Conference Proceedings.
- [18] J. T. Kajiya. Anisotropic reflection models. *Computer Graphics*, 19(3):15–22, July 1985. ACM Siggraph '85 Conference Proceedings.
- [19] R. V. Klassen. Modeling the effect of the atmosphere on light. *ACM Transactions on Graphics*, 6(3):215–237, July 1987.
- [20] E. Lafortune, S. Foo, K. Torrance, and D. Greenberg. Non-linear approximation of reflectance functions. In *SIGGRAPH 97 Conference Proceedings*, Los Angeles, California, Aug. 1997.
- [21] P. Lalonde. (*In preparation*). PhD thesis, University of British Columbia, 1997.
- [22] R. Lewis. Making shaders more physically plausible. In *Fourth Eurographics Workshop on Rendering*, pages 47–62, June 1993.
- [23] C. Lo, B. Palmer, M. Drost, and J. Welty. Incorporation of polarization effects in Monte Carlo simulations of radiative heat transfer. *Numerical Heat Transfer, Part A*, 27:129–142, 1995.
- [24] G. W. Meyer, H. E. Rushmeyer, M. F. Cohen, D. P. Greenberg, and K. E. Torrance. An experimental evaluation of computer graphics imagery. *ACM Transactions on Graphics*, 5(1):30–50, January 1986.
- [25] M. Minnaert. *Light and Color in the Outdoors*. Springer-Verlag, New York, N.Y., 1974.
- [26] M. Oren and S. K. Nayar. Generalization of lambert's reflectance model. In *Proceedings of SIGGRAPH '94*, pages 239–246, 1994.
- [27] E. D. Palik. *Handbook of Optical Constants of Solids*. Academic Press, New York, N.Y., 1985.
- [28] M. Pharr and P. Hanrahan. Geometry caching for ray-tracing displacement maps. In *Eurographics Rendering Workshop*, June 1996.
- [29] M. Pharr, C. Kolb, R. Gershbein, and P. Hanrahan. Rendering complex scenes with memory-coherent ray tracing. In *Computer Graphics*, August 1997. ACM Siggraph '97 Conference Proceedings.
- [30] B.-T. Phong. Illumination for computer generated images. *Communications of the ACM*, 18(6):311–317, June 1975.
- [31] P. Poulin and A. Fournier. A model for anisotropic reflection. *Computer Graphics*, 24(3):267–282, August 1990. ACM Siggraph '90 Conference Proceedings.
- [32] S. Rusinkiewicz. A survey of brdf representation for computer graphics. <http://www-graphics.stanford.edu/smr/cs348c/paper.html>, 1997.
- [33] C. Schlick. A customizable reflectance model for everyday rendering. In *Proceedings of the Fourth Eurographics Workshop on Rendering*, pages 73–84, June 1993.
- [34] M. Schramm, J. Gondek, and G. Meyer. Light scattering simulations using complex subsurface models. In *Proceedings of Graphics Interface '97*, pages 56–67, May 1997.
- [35] P. Schroeder and W. Sweldens. Spherical wavelets: texture processing. In *Eurographics Rendering Workshop 1995*. Eurographics, June 1995.
- [36] P. Shirley. *Physically Based Lighting Calculations for Computer Graphics*. PhD thesis, University of Illinois at Urbana-Champaign, January 1991.
- [37] R. Siegel and J. R. Howell. *Thermal Radiation Heat Transfer*. Hemisphere, Washington, third edition, 1992.
- [38] F. X. Sillion, J. Arvo, S. Westin, and D. Greenberg. A global illumination algorithm for general reflection distributions. *Computer Graphics*, 25(4):187–196, July 1991. ACM Siggraph '91 Conference Proceedings.
- [39] F. X. Sillion and C. Puech. *Radiosity & Global Illumination*. Morgan-Kaufmann, San Francisco, 1994.
- [40] D. C. Tannenbaum, P. Tannenbaum, and M. J. Wozney. Polarization and birefringency considerations in rendering. *Computer Graphics*, pages 221–222, July 1994. ACM Siggraph '94 Conference Proceedings.
- [41] G. J. Ward. Measuring and modeling anisotropic reflection. *Computer Graphics*, 26(4):265–272, July 1992. ACM Siggraph '92 Conference Proceedings.
- [42] S. H. Westin, J. R. Arvo, and K. E. Torrance. Measuring and modeling anisotropic reflection. *Computer Graphics*, 26(2):255–264, July 1992. ACM Siggraph '92 Conference Proceedings.