A Comparison of Four BRDF Models

Stephen H. Westin, Hongsong Li, and Kenneth E. Torrance
Cornell University Program of Computer Graphics

PCG-04-2 April, 2004

We compare four parametric reflectance models that are well-known in computer graphics: the Phong, Ward, Lafortune, and He-Torrance models. We compare the models with physical measurements on five representative sample surfaces. The surfaces span the domain of isotropic, homogeneous surfaces ranging from smooth to rough and including metal and dielectric surfaces. Since no one model was a clear winner in all cases, we draw conclusions about which of the models are best to represent various surfaces. We explain the differences in terms of the basic scattering phenomena involved.

[This report was submitted for publication at the 2004 Eurographics Symposium on Rendering.]

Eurographics Symposium on Rendering (2004), pp. 1–10
H. W. Jensen, A. Keller (Editors)

A Comparison of Four BRDF Models

Stephen H. Westin, Hongsong Li, and Kenneth E. Torrance

Cornell University Program of Computer Graphics

Abstract

We compare four parametric reflectance models that are well-known in computer graphics: the Phong, Ward, Lafortune, and He-Torrance models. We compare the models with physical measurements on five representative sample surfaces. The surfaces span the domain of isotropic, homogeneous surfaces ranging from smooth to rough and including metal and dielectric surfaces. Since no one model was a clear winner in all cases, we draw conclusions about which of the models are best to represent various surfaces. We explain the differences in terms of the basic scattering phenomena involved.
[This report was submitted for publication at the 2004 Eurographics Symposium on Rendering.]

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism I.6.5 [Simulation and Modeling]: Model Development

1. Introduction

In computer rendering, it is not always clear which of many available reflectance models should be used for a given surface; the relative strengths and weaknesses of the various models are not always clear. This paper seeks to provide a basis for selecting one model over another.

We have measured the BRDF of five carefully-chosen samples in the plane of incidence, fitting four different well-known models to each set of data. We then explain the phenomena found in each dataset in light of the physical processes involved, assess each model on its fidelity to the original data, and relate these phenomena and modeling differences to their visual effect in a rendered image.

We have chosen four models that are well-known in the computer graphics literature: Phong [Pho75], Ward [War92], Lafortune [LFTG97], and He-Torrance [HTSG91].

We will start the paper by reviewing the history of reflectance modeling for computer rendering, then elucidate some relevant reflectance phenomena, and finally present five sample surfaces, comparing how the four different models capture the reflectance behavior of each.

2. Background

Researchers have taken three main approaches to modeling reflectance for computer renderings: simple empirical models, physically-based models, and general mathematical fitting representations.

The earliest realistic computer renderings used a simple Lambertian model, limiting environments to purely diffuse surfaces. Phong [Pho75] introduced an empirical model, controlled by three simple parameters, that produced a glossy effect with minimal computation. Since then, Ward created a model almost as simple and efficient [War92], but designed to fit measured reflectance data and to incorporate effects of surface anisotropy. The Phong model itself has been modified to obey certain physical laws to allow use in simulations of global illumination [Lew93]. More recently, Westlund and Meyer [WM01] developed an empirical model based on traditional gloss and appearance measurements.

These models are compact and computationally efficient, and offer a small number of parameters that can be manipulated interactively to achieve a desired appearance. They may fall short in reproducing the exact reflectance of a surface, but there is controversy as to how visually important such errors may be.

Physically-based models were introduced to computer rendering by Blinn [Bli77] and Cook and Torrance [CT82], based on the work of [TS67]. The approach was generalized to arbitrary facet distributions by Ashikhmin et al. [APS00].

These were followed by several partial models of physical scattering processes including the anisotropic models of Kajiya [Kaj85] and Poulin and Fournier [PF90]. Oren and Nayar [ON94] presented a model for extremely rough surfaces such as unglazed ceramics. Finally, there are models of complete wave optics for certain classes of surfaces [HTSG91, Sta99].

These models offer the hope of accurate modeling of surface reflectance, but their parameters (e.g. index of refraction) often have no direct intuitive meaning and require special equipment to measure. This means that the models are unsuited to interactive use.

The final class of reflectance models is that of general mathematical fitting representations. Since the reflectance is a function of 3 or 4 geometric variables, it is possible to construct a representation for fitting based on orthogonal basis functions. Spherical harmonics [SAWG91, WAT92], Zernike polynomials [KvD98], and wavelets [SS95, LF97] have all been used as orthogonal bases to represent the BRDF. Kautz and McCool [KM99] represented the BRDF as a sum of separable functions. Each of these representations is theoretically capable of representing any BRDF exactly, but each requires many coefficients (dozens, hundreds, or even thousands) to achieve a good approximation, as compared to 3-4 coefficients for models in the other two categories. Lafortune et al. [LFTG97] struck a middle ground by generalizing the Phong model with multiple steerable lobes. A different approach was used by Matusik et al. [MPBM03], enumerating the entire space of BRDFs and using a relatively small number (15-30) of parameters to navigate the space.

These general representations seem to be used more rarely than models in the other two classes, perhaps because they must be fit to data from another source: either a physical model or measured data. It is telling that each of these papers either includes its own set of measurements or relies on some of the few BRDF datasets publicly available [DGNK97, Cor99]. In addition, the storage and computation required for many of these models is far greater than for models in the other two classes.

With such a variety of models, it is not at all clear which model to use in a particular case. Implementers of rendering systems must choose a model that combines efficiency, ease of use, and compatibility with a particular rendering or global illumination algorithm; users of such systems usually have a choice of models already implemented and must choose the most effective to reproduce the appearance of a given surface.

3. Surface Reflectance Phenomena

The models compared in this paper deal with a limited, but common, class of surfaces: those that

- can be treated as opaque (i.e. non-translucent), such that incident light is absorbed or scattered within a small region of the surface;
- can be described well by pure height fields;
- are rough in a random way, rather than with regular structure like a diffraction grating;
- are isotropic in roughness, so reflectance is invariant under a rotation about the surface normal;
- range from perfectly smooth to quite rough; and
- display pure Lambertian subsurface scattering.

All the selected models, as published, assume Lambertian subsurface scattering, if any. We include a more interesting metallic paint sample, which can be modeled effectively with a slight variation on the published models.

These limitations may seem to be quite restrictive, but they actually embrace a broad array of surfaces encountered in normal life. The examples in this paper include plastic and painted surfaces, metal surfaces, and even high-gloss metallic automotive paint.

3.1. Ideal Diffuse Reflectance

The simplest reflectance to model is Lambertian, with equal scattering in all directions. Strictly speaking, this behavior can never be realized physically, but near-Lambertian scattering can occur with near-normal illumination onto rough surfaces. The closest approach to Lambertian scattering is obtained with materials where subsurface scattering dominates.

3.2. Ideal Specular Reflectance

Ideal specular, or mirror-like, reflection is also relatively simple. All incident light is scattered in the specular direction with no spreading of the beam. Such reflectance comes from perfectly smooth surfaces. Unlike Lambertian reflection, ideal specular reflection can be realized. This is because of the wave nature of light; when the roughness of a surface is much smaller than the wavelength of light, it is perfectly smooth for optical purposes.

The magnitude of ideal specular reflectance depends on the material and the angle of incidence, and was quantified by Fresnel in the early 19th century. Figure 1 shows this relationship for several common materials.

The most important distinction between materials is between metals, which conduct electricity, and non-metals (or dielectrics), which do not. The mirror reflectance of dielectrics varies much more dramatically with angle, and its minimum is much lower, typically 4-5%. The distinction between metals and dielectrics is embodied in the refractive index n, which is a real number for dielectrics and a complex number for metals. Figure 1 applies for smooth surfaces, but the difference in reflectance between metals and dielectrics carries over to rough surfaces, as we shall see in Section 4.

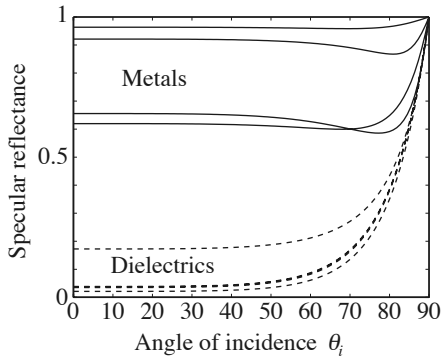

Figure 1: *Fresnel reflectance for smooth surfaces: metals vs. dielectrics*

3.3. Directional Diffuse Scattering

Most surfaces in our world are neither perfect mirrors nor Lambertian scatterers. These surfaces scatter light from a single source into many directions, but unequally. For the surfaces we are considering, the preference is for directions near the specular direction and consists of a smooth lobe much like that of the Phong model. Such surfaces have random roughness significantly larger than the light wavelength. Most reflectance models in computer graphics have been intended to operate in this regime, starting with Phong [Pho75] and including those of Ward [War92] and Cook and Torrance [CT82].

Extremely rough surfaces such as sandpaper or unglazed ceramic violate the assumptions of these models; because of their large mean surface slope, they are not matched well by these lobe-like models. Oren and Nayar [ON94] developed a model for these cases. This paper will not deal with large-slope surfaces.

The directional diffuse reflectance generally increases with increasing angle of incidence, for two reasons. Geometric factors cause a reflectance increase for all surfaces, and for dielectrics, the strong variation of Fresnel reflectance (evident in Figure 1) accentuates the increase.

3.4. Transitions Between Scattering Modes

If very smooth surfaces display mirror-like reflectance, but rougher surfaces show directional diffuse scattering, what happens when the roughness lies between the two domains?

Surfaces with very small roughness are similar in behavior to perfectly smooth surfaces, but their specular reflectance is attenuated by a well-known relationship [BS63]:

$$\rho \approx F e^{-\left(4\pi \frac{\sigma}{\lambda} \cos\theta_i\right)^2} \qquad (1)$$

where ρ is the resulting specular reflectance, F is the Fresnel reflectance, σ is the RMS roughness height, λ is the wavelength of incident light, and θ_i is the angle of incidence of that light. The fraction of light not reflected into the mirror direction, $F - \rho$, is essentially scattered to the rest of the hemisphere above the surface.

As the surface roughness increases, the ideal specular reflectance is attenuated, and a directional diffuse lobe emerges. At first, this lobe is centered on the specular direction, but can move to off-specular directions at large incidence angles [TS66]. Note that the transition between rough- and smooth-surface behavior is not a gradual change in shape. The two behaviors coexist when their domains overlap, in a proportion depending on surface roughness. Surfaces in this range of roughness can only be handled by a full physical-optics model like that of He et al. [HTSG91] or Stam [Sta99]. None of the simpler lobe-like models can accurately produce an ideal specular component, so the range of roughnesses they model is limited.

3.5. Effective Roughness

It may be surprising that the roughness attenuation of Equation 1 varies with θ_i; we might expect it to depend only on σ and λ. In fact, a particular surface will respond as though the roughness actually varies with incidence angle. This effect, illustrated in Figure 2, is called the effective roughness of a surface. With decreasing effective roughness, i. e. increasing θ_i, the directional diffuse lobe narrows and increases in magnitude, and an ideal specular component may emerge.

Such behavior is visible in the real world. For example, rough pavement shows significant glare from headlights at night, since the incidence and viewing angles approach grazing. Rough paper, though normally highly diffuse, shows a sheen as viewing angles approach grazing. This effect is why a perfectly Lambertian surface can never be realized. Even a very rough surface that shows little directionality at normal incidence will exhibit a directional diffuse lobe near grazing incidence as the effective roughness becomes small. A surface of somewhat lesser, but still large, roughness that displays a directional diffuse lobe at angles near normal may resemble a perfect mirror as grazing incidence is approached.

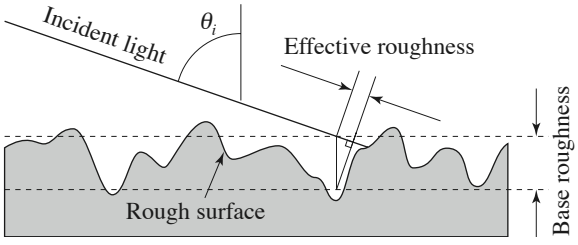

Figure 2: *Effective roughness decreases according to $\sigma \cos\theta_i$ as in Equation 1.*

3.6. Subsurface Scattering

Even the best model of surface scattering is insufficient to reproduce the appearance of most surfaces, since much of the appearance (e.g. color) of these surfaces comes from scattering beneath the surface, within the body of the material. The only exceptions are metallic surfaces and black dielectrics. A metallic surface prevents subsurface scattering, since energy is absorbed within a fraction of one wavelength of the surface. A black dielectric also absorbs all energy that enters, but over a longer length scale. All other dielectric surfaces reflect significant energy via subsurface scattering.

Each of the models compared here includes an empirical Lambertian term to model subsurface scattering. Strictly speaking, surface scattering will modify this pure behavior, as pointed out by Shirley et al. [SHSL97], but often this can be neglected, as we will see in Section 4. For more complex cases, an approach like that of Hanrahan and Krueger [HK93] may be useful.

4. Model Comparisons

In this section we take several examples of surfaces from within our domain and fit four different models to reflectance measurements of each, using the phenomena described in Section 3 to explain the results. The surfaces were chosen to include all of these phenomena; they include metal and dielectric surfaces and roughnesses ranging from the ideal specular domain to near-Lambertian behavior. Based on our experience, these samples represent the behavior of the entire class of surfaces described at the beginning of Section 3.

Surface reflectance is completely described by the BRDF, which is a function of three angles for these surfaces. We will examinine the function only in the plane of incidence, which is a well-accepted practice in surface optics [Sto95] that allows concise graphs that show the character, as well as the magnitude, of fitting errors. The profile in the incidence plane defines the major behavior with respect to two of the angles involved. We have fit the four models to full 3D data for two of the examples (those in Sections 4.3 and 4.5), without significant differences in the results.

We chose four models from the computer graphics literature for our comparision: the models of Phong [Pho75], Ward [War92], Lafortune [LFTG97], and He, Torrance, et al. [HTSG91]. The Phong model was the first model for glossy reflectance, and is still widely used. The Ward model is, like the Phong model, empirical, but gains sophistication from its verification with measured BRDF data. It is quite popular both because of the RADIANCE software package that depends on it and because of its convenience for Monte Carlo calculations. The Lafortune model represents the class of models that must be fit to some external source of data. He-Torrance is the most sophisticated of a series of physical models starting with Torrance-Sparrow [TS67] and including Blinn [Bli77] and Cook-Torrance [CT82]. The He-Torrance model gives results similar to those earlier models for the very rough surfaces where the latter are applicable (i.e. no ideal specular component).

All input data for this process were measured with the calibrated gonioreflectometer of the Cornell University Program of Computer Graphics [Foo96], except for the surface of Section 4.4, which was measured by Marx and Vorberger [MV90]. The BRDF was measured in the incidence plane, using a nonuniform sampling pattern to concentrate more samples near the specular direction where the sharpest features are seen. In each graph, we plot the original data points, showing the range of values as well as the angular range and resolution of the data. Measurements were taken at incidence angles of $15^\circ, 30^\circ, 45^\circ, 60^\circ$, and 75°.

Each model was fitted via a nonlinear optimization process in MATLAB, according to the L^2 norm of the BRDF error weighted by the cosine of the exitant angle. We used each model as originally published. For the Lafortune model, we used one or two lobes for each channel, as needed; for each of the other models, we used a single set of coefficients per channel. We chose to plot data for three angles of incidence: $15^\circ, 45^\circ$, and 75°. This shows both the "typical" behavior at near-normal incidence, the effects near grazing, and the transition between the two.

We have chosen to plot the data on a logarithmic scale because the range of values of the BRDF is large, and a linear scale would obliterate the effects seen at low angles of incidence. In addition, human perception of luminance is nonlinear, and the logarithmic scale follows Weber's law [WS82], a well-known approximation to this nonlinearity. This means that a given vertical distance on these graphs corresponds roughly to a similar perceptual difference.

Renderings were generated by a Monte Carlo path tracer, computing full global illumination for all transport paths. We present the five surfaces in order of increasing complexity.

4.1. Rough Metal

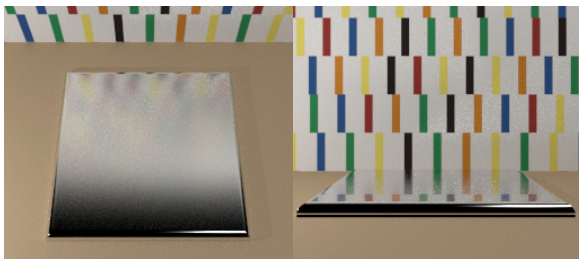

Figure 3: *Renderings of rough aluminum surface at two viewing angles*

Our first sample is a rough aluminum surface created by coating a sheet of anti-glare picture frame glass with pure

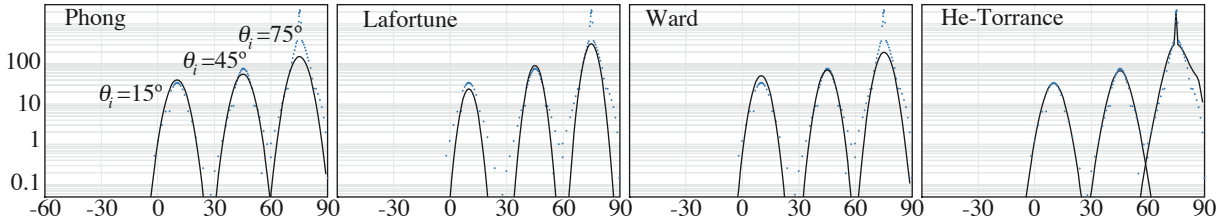

Figure 4: *Fitting results for rough aluminum surface: BRDF vs. angle of reflection for three incidence angles, $\lambda = 550nm$. Blue dots show actual measurement points; solid lines show results from each model.*

aluminum; its visual appearance is shown in computer renderings in Figure 3. The measured BRDF data are shown in Figure 4; they show a directional diffuse lobe with a small ideal specular (mirror-like) component emerging at an incidence angle of $\theta_i = 75°$. This ideal specular component corresponds approximately to the mirror reflection relation, Equation 1, but its magnitude is relatively small. Gloss and some mirror reflection are suggested in the right-hand rendering at the larger viewing angle in Figure 3.

For this surface, there is no striking difference in the performance of the four models for the directional diffuse lobes. The Phong model shows the least increase toward grazing incidence, but not much less than the other three models. It may be surprising to see any increase at all in the Phong model; the model was originally formulated in terms of radiance, but the BRDF is scaled by irradiance, which includes an extra term $\cos\theta_i$. Dividing by this term gives the increase shown in the directional diffuse lobe.

4.2. Rough Plastic

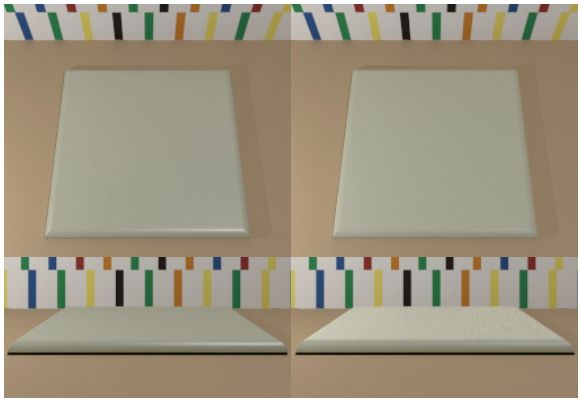

Figure 5: *Comparison of rendered images of off-white rough plastic for two reflectance models.*

We now move from a metallic surface to a dielectric: a section of an off-white computer case, which shows a finish with a slight sheen. An example rendering appears in Figure 5. The BRDF data are shown in Figure 6. Both this surface and the rough metal in Figure 4 show a directional diffuse lobe, but several differences are apparent. First, the minimum BRDF value in Figure 6 is increased by a constant Lambertian component due to subsurface scattering, which gives the surface its off-white color. Second, the magnitude of the directional diffuse lobe varies much more with incidence angle for the dielectric surface, since it is a dielectric material and its Fresnel reflectance varies dramatically with angle, as shown in Figure 1. Third, the magnitude of the directional diffuse lobe is roughly a full order of magnitude less for the dielectric surface. The two surfaces show similar brightness, however, since the wider directional diffuse lobe and the constant Lambertian subsurface component of the plastic surface spread a similar amount of energy over a broader angular range. Finally, the wider directional diffuse lobe indicates a greater mean roughness slope on the plastic surface. This surface also shows no emergence of an ideal specular peak, indicating that the surface roughness σ is larger than that for the metal surface.

This material shows dramatic differences between the four models. The He-Torrance model fits fairly well, but the Phong model cannot model the increase in the directional diffuse lobe. The Ward model also shows an inadequate increase toward grazing. The rate of increase toward grazing is fixed in the Ward model, whereas the He-Torrance model includes an explicit factor for the Fresnel reflectance and so models the increase. The Lafortune fitting model, as we might expect, has sufficient flexibility to reproduce the effect accurately. To see the visual effects of this difference, we have rendered a simple environment using two of the models and show the results in Figure 5. Notice how the Phong model is slightly brighter than the Lafortune model at viewing angles near normal, but much dimmer near grazing, where both the Fresnel reflectance and the reduced effective roughness increase the directional reflectance.

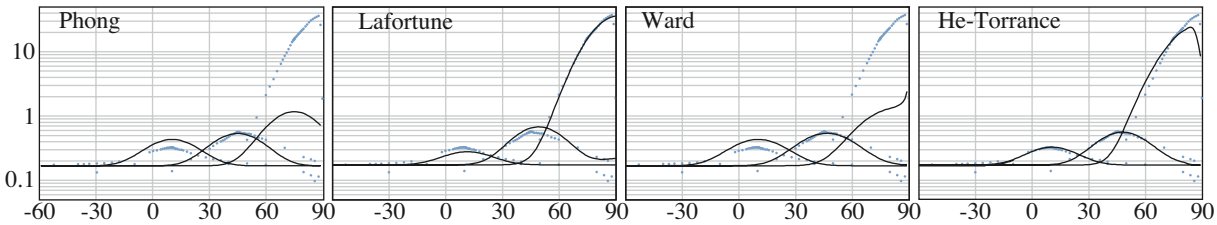

Figure 6: *Fitting results for rough plastic surface: BRDF vs. angle of reflection for three incidence angles,* $\lambda = 550nm$.

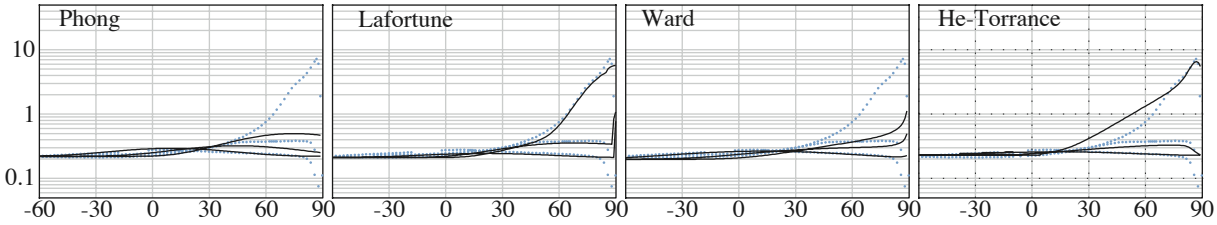

Figure 8: *Fitting results for plain white paper: BRDF vs. angle of reflection for three incidence angles,* $\lambda = 550nm$.

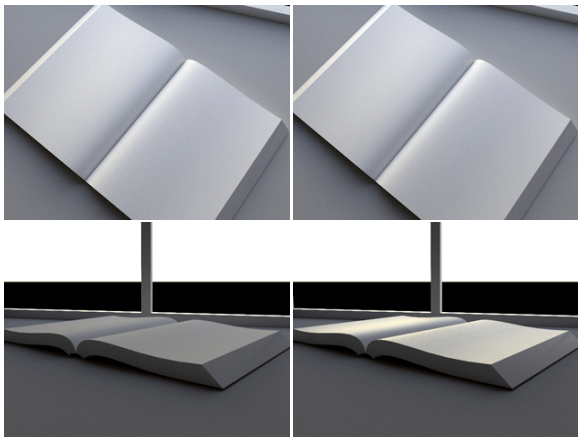

Figure 7: *Comparison of rendered images of paper for two reflectance models*

4.3. White Paper

This example is an ordinary sheet of office paper, shown rendered in Figure 7. The reflectance is plotted in Figure 8. Like the rough plastic surface, it has a strong subsurface scattering component. Unlike the plastic, there is no significant directional diffuse lobe at incidence angles near normal, so the paper has no visible gloss. Because of this, this surface might normally be modeled as purely Lambertian. However, its reflectance in Figure 8 varies dramatically near grazing angles due partly to the decrease in effective roughness predicted by Equation 1.

The He-Torrance model matches the reflectance increase fairly well with its effective roughness term, but falls somewhat short. The differerence may be attributed to the subsurface scattering, which is not truly Lambertian; as the angle of incidence increases, the subsurface scattering is biased in the forward direction, as has been observed with paper samples [GJM03]. The Lafortune model is flexible enough to model the large-angle behavior better.

The Phong and Ward models show dramatic departures from the true behavior. This is borne out in the renderings shown in Figure 7; near-normal viewing shows virtually no difference between the models, but viewing from near grazing shows a dramatic difference, with the Lafortune model looking much more natural. The golden color in the images is introduced by the illumination, which simulates a sunset sky.

4.4. Smoother Metal

Our next example is a metal surface in the transition region between rough and smooth surface behavior. An example rendering is shown in Figure 9; the BRDF data are plotted in Figure 10. The surface, made of stainless steel and hand-lapped to near-mirror smoothness, was studied, and its reflectance measured, by Marx and Vorberger [MV90]. We used the surface roughness statistics given in the paper: RMS roughness $\sigma = 0.08\mu m$, horizontal roughness $\tau = 1.52\mu m$. Since we were unable to obtain refractive index information for the stainless steel, we used the index for iron, the major component, but the behavior is relatively insensitive to this approximation. We have normalized both the measurement data from the paper and our BRDF results, as the

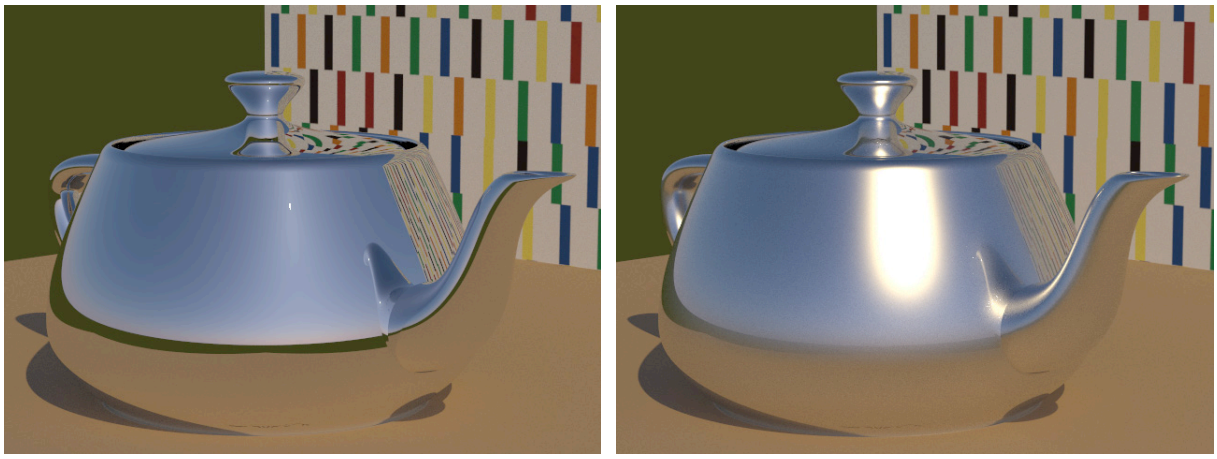

Simple Mirror He-Torrance Model

Figure 9: *Renderings of nearly-smooth metal teapot*

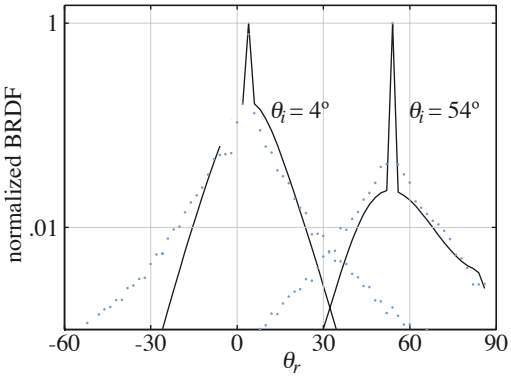

Figure 10: *Fitting results for nearly-smooth metal surface: data from Marx et al., He-Torrance model,* $\lambda = 633nm$.

Figure 11: *Rendering using Lafortune model of metallic paint*

data in the paper were not absolute; the result is shown in Figure 10. The specular peak is prominent even at the near-normal incidence angle of 4°, in contrast to the surface in Section 4.1, where it emerged only at large angles. $\theta_i = 54°$ is the largest incidence angle measured by Marx et al., but Equation 1 makes it clear that the specular component will grow still farther at greater incidence angles. We show only the results from the He-Torrance model, as none of the other models is capable of reproducing the ideal specular component. While the He-Torrance model does not exactly match the measurements, there is qualitative agreement, with the directional diffuse lobes similar in width and position over a range of two orders of magnitude.

Since the specular peak is so strong, we might avoid the computational complexity of the He-Torrance model by modelling this surface as a mirror, either by using the Fresnel reflectance or with a single reflectance value for all angles. Figure 9 shows the result of this approach, compared to a rendering with the He-Torrance model. The dramatic highlight from the sun essentially disappears, since its brightness is limited to a few pixels whose values are severely clipped. Whitted [Whi80] avoided this by using a non-physical model that provides a directional diffuse lobe such as Phong for light sources. Unfortunately, such an approach would still miss subtler effects that are still important. Notice how the horizon is blurred on the front of the teapot in the right-hand image, but becomes perfectly distinct at the edge of the teapot, where angles approach grazing. Similarly, the reflection of the spout in the body of the teapot on the right is blurred in a way that is distinctive of slightly rough metal, such as the outside of an unpainted airliner.

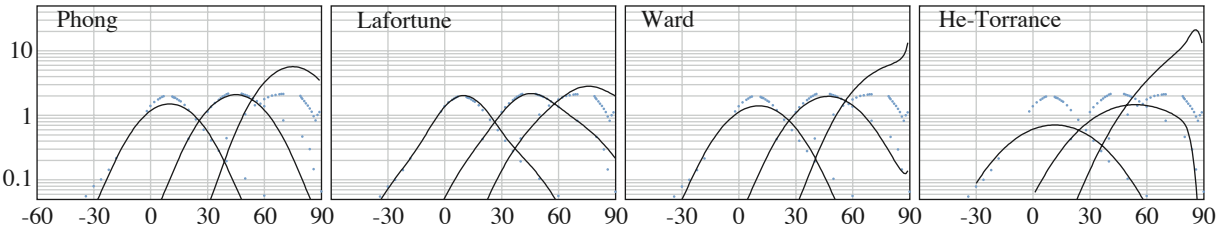

Figure 12: *Fitting results for metallic paint with clearcoat: BRDF vs. angle of reflection for three incidence angles,* $\lambda = 550nm$.

4.5. Metallic Paint

Our final example is an automotive metallic paint, showing how a more complex surface can be modeled with a combination of commonly-available models. An example rendering is shown in Figure 11; the BRDF data are plotted in Figure 12. The near-specular lobe comes from subsurface scattering; the surface itself is smooth and mirror-like. To display the subsurface scattering, these graphs omit the ideal specular surface reflection; for rendering, the specular reflection can be modeled simply via the Fresnel equation.

This paint has a "metallic" visual appearance, but its reflectance is quite different from that of a real metal. We see a directional diffuse lobe in Figure 12 that resembles the lobe from the rough aluminum sample in Figure 4, but shows very little variation with angle of incidence. This demonstrates a difficulty faced by the makers of the ill-fated Delorean automobile: it was impossible to paint the plastic body parts to match the stainless-steel body panels, since there is no paint that will match from all angles.

The Phong model reproduces this behavior well due its lack of variation with incidence angle. The Lafortune fitting model again shows the flexibility to match this performance, and the more sophisticated Ward and He-Torrance models are actually less accurate because this surface violates one of their assumptions.

The subsurface scattering comes from a collection of polished metal flakes suspended at random angles in a clear binder, much like the idealized microfacets of Cook and Torrance [CT82]. As angles approach grazing, the presence of a mirror-like clear overcoat modifies behavior for two reasons, both arising from the Fresnel reflectance. First, as the incident light approaches grazing, the Fresnel reflectance sends more of the total energy into the ideal specular component (not shown in the BRDF graph), reducing the magnitude of the directional lobe. Second, light scattered by the aluminum must traverse the smooth boundary between clearcoat and air; since the subsurface internal reflection increases with angle, light scattered into a near-grazing direction will be preferentially reflected back to the aluminum flakes and will be multiply scattered.

Figure 11 shows a car model rendered using the Lafortune model, with a Fresnel term added to describe the specular reflection from the clearcoat.

5. Summary and Conclusions

We have presented the first comparison of major alternative reflectance models by fitting to a common set of measured reflectance data. The data include smooth and rough surfaces, as well as metals and dielectrics. We have compared predicted and measured reflectances, and rendered images.

We have considered only surfaces that exhibit a smooth directional-diffuse reflectance lobe at or near the specular direction, giving the classic computer graphics highlight, with possibly an additional mirror-like component.

Our main results are that

- the differences between various reflectance models are visible in rendered images;
- no one model performs well on every surface;
- sometimes the most sophisticated model is not the best choice for a given surface; and
- the differences can be understood in light of a few basic principles.

Of the four models compared, the Lafortune model performed well for all the surfaces studied, since it was designed to fit almost arbitrary BRDF data. The physically-based He-Torrance model also performed well, except when its assumptions were violated. The Ward and Phong models each performed well for certain surfaces, generally those where the dependence of reflectance on incidence angle is relatively weak. We believe that our incidence-plane studies will be substantially validated by future work over the entire hemisphere.

We believe that this work will be of help

- to implementers of rendering systems, who need to choose models that are efficient, accurate, and compatible with the algorithms used;
- to users of rendering systems, who must choose from a sometimes bewildering array of models implemented in any given system; and
- to researchers creating new reflectance models of greater accuracy, generality, and efficiency.

References

[APS00] ASHIKHMIN M., PREMOZE S., SHIRLEY P. S.: A microfacet-based BRDF generator. In *Proceedings of ACM SIGGRAPH 2000* (July 2000), Computer Graphics Proceedings, Annual Conference Series, pp. 65–74.

[Bli77] BLINN J. F.: Models of light reflection for computer synthesized pictures. In *Computer Graphics (Proceedings of SIGGRAPH 77)* (July 1977), vol. 11, pp. 192–198.

[BS63] BECKMANN P., SPIZZICHINO A.: *The Scattering of Electromagnetic Waves from Rough Surfaces*. Pergamon Press, 1963, p. 81.

[Cor99] CORNELL UNIVERSITY PROGRAM OF COMPUTER GRAPHICS: Measurement data. http://www.graphics.cornell.edu/online/measurements/, 1999.

[CT82] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. *ACM Transactions on Graphics 1*, 1 (Jan. 1982), 7–24.

[DGNK97] DANA K. J., GINNEKEN B. V., NAYAR S. K., KOENDERINK J. J.: Columbia-Utrecht Reflectance and Texture Database. http://www1.cs.columbia.edu/CAVE/curet/, 1997.

[Foo96] FOO S.-C.: *A gonioreflectometer for measuring the bidirectional reflectance of material for use in illumination computation*. Master's thesis, Cornell University, Ithaca, New York, 1996.

[GJM03] GRANBERG H., JENSEN J., MATTSSON L.: Forward scattering of fiber-containing surfaces studied by 3-D reflectance distribution simulations and measurements. *Opt. Eng. 42* (Aug. 2003), 2384–2390.

[HK93] HANRAHAN P., KRUEGER W.: Reflection from layered surfaces due to subsurface scattering. In *Proceedings of SIGGRAPH 93* (Aug. 1993), Computer Graphics Proceedings, Annual Conference Series, pp. 165–174.

[HTSG91] HE X. D., TORRANCE K. E., SILLION F. X., GREENBERG D. P.: A comprehensive physical model for light reflection. In *Computer Graphics (Proceedings of SIGGRAPH 91)* (July 1991), vol. 25, pp. 175–186.

[Kaj85] KAJIYA J. T.: Anisotropic reflection models. In *Computer Graphics (Proceedings of SIGGRAPH 85)* (July 1985), vol. 19, pp. 15–21.

[KM99] KAUTZ J., MCCOOL M. D.: Interactive rendering with arbitrary BRDFs using separable approximations. In *Eurographics Rendering Workshop 1999* (June 1999).

[KvD98] KOENDERINK J., VAN DOORN A.: Phenomenological description of bidirectional surface reflection. *J. Opt. Soc. Am. A 15*, 11 (Nov. 1998), 2903–2912.

[Lew93] LEWIS R.: Making shaders more physically plausible. In *Fourth Eurographics Workshop on Rendering* (June 1993), pp. 47–62.

[LF97] LALONDE P., FOURNIER A.: A wavelet representation of reflectance functions. *IEEE Transactions on Visualization and Computer Graphics 3*, 4 (October–December 1997), 329–336.

[LFTG97] LAFORTUNE E. P. F., FOO S.-C., TORRANCE K. E., GREENBERG D. P.: Non-linear approximation of reflectance functions. In *Proceedings of SIGGRAPH 97* (Aug. 1997), Computer Graphics Proceedings, Annual Conference Series, pp. 117–126.

[MPBM03] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: A data-driven reflectance model. *ACM Transactions on Graphics 22*, 3 (July 2003), 759–769.

[MV90] MARX E., VORBERGER T. V.: Direct and inverse problems for light scattered by rough surfaces. *Applied Optics 29*, 25 (Sept. 1990), 3613–3626.

[ON94] OREN M., NAYAR S. K.: Generalization of Lambert's reflectance model. In *Proceedings of SIGGRAPH 94* (July 1994), Computer Graphics Proceedings, Annual Conference Series, pp. 239–246.

[PF90] POULIN P., FOURNIER A.: A model for anisotropic reflection. In *Computer Graphics (Proceedings of SIGGRAPH 90)* (Aug. 1990), vol. 24, pp. 273–282.

[Pho75] PHONG B.-T.: Illumination for computer generated pictures. *Communications of the ACM 18*, 6 (June 1975), 311–317.

[SAWG91] SILLION F. X., ARVO J. R., WESTIN S. H., GREENBERG D. P.: A global illumination solution for general reflectance distributions. In *Computer Graphics (Proceedings of SIGGRAPH 91)* (July 1991), vol. 25, pp. 187–196.

[SHSL97] SHIRLEY P., HU H., SMITS B., LAFORTUNE E. P.: A practitioners' assessment of light reflection models. In *Pacific Graphics '97* (Oct. 1997).

[SS95] SCHRÖDER P., SWELDENS W.: Spherical wavelets: Efficiently representing functions on the sphere. In *Proceedings of SIGGRAPH 95* (Aug. 1995), Computer Graphics Proceedings, Annual Conference Series, pp. 161–172.

[Sta99] STAM J.: Diffraction shaders. In *Proceedings of SIGGRAPH 99* (Aug. 1999), Computer Graphics Proceedings, Annual Conference Series, pp. 101–110.

[Sto95] STOVER J. C.: *Optical Scattering, Measurement and Analysis*, 2nd ed. ed. SPIE Optical Engineering Press, Bellingham, WA, 1995.

[TS66] TORRANCE K. E., SPARROW E. M.: Off-specular peaks in the directional distribution of reflected thermal distribution. *Journal of Heat Transfer — Transactions of the ASME Series C* (May 1966), 223–230.

[TS67] TORRANCE K. E., SPARROW E. M.: Theory for off-specular reflection from roughened surfaces. *Journal of Optical Society of America 57*, 9 (Sept. 1967), 1105–1114.

[War92] WARD G. J.: Measuring and modeling anisotropic reflection. *Computer Graphics 26*, 2 (July 1992), 265–272. Proceedings of SIGGRAPH '92.

[WAT92] WESTIN S. H., ARVO J. R., TORRANCE K. E.: Predicting reflectance functions from complex surfaces. In *Computer Graphics (Proceedings of SIGGRAPH 92)* (July 1992), vol. 26, pp. 255–264.

[Whi80] WHITTED T.: An improved illumination model for shaded display. *Communications of the ACM 23*, 6 (June 1980), 343–349.

[WM01] WESTLUND H. B., MEYER G. W.: Applying appearance standards to light reflection models. In *Proceedings of ACM SIGGRAPH 2001* (Aug. 2001), Computer Graphics Proceedings, Annual Conference Series, pp. 501–510.

[WS82] WYSZECKI G., STILES W. S.: *Color Science: Concepts and Methods, Quantitative Data and Formulae*, second edition ed. John Wiley and Sons, 1982, ch. 6, p. 490.

Appendix: Parameter Values Used in This Paper

We list here all the parameter values used for BRDF graphs and renderings in this paper. Each graph uses data measured at the single wavelength $\lambda = 550nm$, which is the peak of human luminance sensitivity, except Figure 10, where data were only available at $\lambda = 633nm$.

For most renderings, we used this set of coefficients as a monochrome BRDF; for Figure 5, we converted spectral data (gathered at $10nm$ intervals) to RGB at each data point, then fit each model to the individual color channels. We give all sets of coefficients in the tables below.

Sample	Channel	Phong			Ward		
		K_d	K_s	n	ρ_d	ρ_s	α
Rough metal (§4.1)	$550nm$	0.0	1.1	224	0.0	1.07	0.0413
White Plastic (§4.2)	$550nm$	0.663	0.186	67.8	0.526	0.0904	0.167
	R	0.525	0.0792	19.8			
	G	0.541	0.0792	19.9			
	B	0.396	0.0834	19.7			
White Paper (§4.3)	$550nm$	0.686	0.0999	3.48	0.664	0.212	0.486
Metallic Paint (§4.5)	$550nm$	0.0562	0.529	16.2	0.0456	0.612	0.188

Table 1: *Parameters for Phong and Ward models*

Sample	Channel	C_{xy}	C_z	n	K_d
Rough metal	$550nm$	-1.01	1.01	405.5	0
White Plastic	$550nm$	-1.22	-0.142	21.4	0.542
		-1.04	0.893	21.1	
	R	-1.17	-0.210	20.2	0.585
		-0.989	0.854	21.5	
	G	-1.16	-0.164	20.9	0.604
		-0.988	0.848	20.8	
	B	-1.17	-0.171	20.8	0.446
		-0.989	0.858	21.9	
White Paper	$550nm$	-1.11	0.0052	20.4919	0.653
		-0.701	0.314	3.0265	
Metallic Paint	$550nm$	-0.992	1.01	34.4	0.021
		-1.11	0.914	7.2	

Table 2: *Parameter values used for the Lafortune model*

Sample	$\lambda(nm)$	σ	τ	n	ρ_{ud}
Rough metal	550	0.200	8.17	$0.958 + 6.69i$	0
White Plastic	550	0.512	6.97	1.5	0.534
	650 (R)	0.512	6.97	1.5	0.532
	550 (G)	0.512	6.97	1.5	0.548
	450 (B)	0.512	6.97	1.5	0.413
White Paper	550	0.798	4.87	1.5	0.718
Smoother Metal (§4.4)	633	0.08	1.52	$1.3730 + 7.6193i$	0
Metallic Paint	550	0.687	6.00	5.0103	0.0239

Table 3: *Parameter values used for the He-Torrance model*