Using the Modified Phong Reflectance Model for Physically Based Rendering

Eric P. Lafortune
Yves D. Willems

Report CW 197, November 1994

Department of Computing Science, K.U. Leuven

Abstract

This text discusses a few aspects of reflectance models in physically based rendering:

- The first section presents the definition of the *bidirectional reflection distribution function* (*brdf*) of a surface and its physical properties.
- On a more practical level, the next section discusses models to represent *brdf*s and their desired properties in general for Monte Carlo algorithms.
- The third section goes into details about a specific reflectance model, the modified Phong *brdf*, with its definition, its properties and its use. We show how this model can be correctly integrated in importance sampling schemes for physically based Monte Carlo rendering algorithms.
- The fourth section is devoted to alternative parameter spaces in which reflectance models can be sampled, either deterministically or stochastically.
- The last section discusses an important implementational issue, more specifically the problem of verifying the implementation of a reflectance model.

Keywords : Reflectance models, BRDF, Monte Carlo rendering.
CR Subject Classification : I.3.7

Contents

1 The bidirectional reflection distribution function

1.1 Definition

The bidirectional reflection distribution function defines the local reflective properties of a surface. It describes in a physically and mathematically precise manner if a surface is reflective or not, diffuse or specular, etc. The *brdf* is formally defined as:

$$f_r(x, \Theta_i, \Theta_o) = \frac{dL_o(x, \Theta_o)}{L_i(x, \Theta_i)\ \cos\theta_i d\omega_i}$$

where:

- $x =$ the position on the surface,
- $\Theta_i =$ the direction from which the light comes in,
- $\Theta_o =$ the direction in which the light is reflected and measured,
- $dL_o(x, \Theta_o) =$ the differential amount of radiance that is reflected at point x in the outgoing direction Θ_o,
- $L_i(x, \Theta_i) =$ the amount of radiance coming in at point x along direction Θ_i through $d\omega_i$,
- $\cos\theta_i =$ the cosine of the angle between Θ_i and the surface normal at point x,
- $d\omega_i =$ a differential angle around Θ_i.

In this text we make abstraction of the wavelength dependency of the *brdf* and only consider monochromatic light.

1.2 Physical constraints

The set of all possible *brdfs* is constrained by some physical laws:

- Due to the reciprocal nature of the reflection of light the function of the incoming and the outgoing directions may be interchanged:

  $$f_r(x, \Theta_i, \Theta_o) = f_r(x, \Theta_o, \Theta_i)$$

  This property is called the Helmholtz-reciprocity.

- The fraction of light coming in from any direction that is reflected in the entire hemisphere should be smaller than 1 to ensure conservation of energy. This fraction is expressed by the total hemispherical reflectivity:

  $$\rho(x, \Theta_i) = \int_{\Omega_x} f_r(x, \Theta_i, \Theta_o)\ \cos\theta_o d\omega_o \leq 1, \quad \forall x, \Theta_i$$

Functions conforming to these constraints can be called *physically plausible* [1].

2 Reflectance models

2.1 Definition

A reflectance model is a practical representation for some set of *brdfs*. Any implementation of a global illumination algorithm uses a reflectance model. It may be empirical and simple, representing only a limited set of *brdfs* as in the classical radiosity method which assumes that all surfaces are Lambertian (i.e. perfectly diffuse). The model may also be more elaborate and flexible, allowing for specular and glossy reflections, anisotropy, etc. [2, 3, 4, 5, 6]. It should of course be physically plausible. If not, algorithms that use it may not converge to a correct solution or not converge at all.

2.2 Desired practical properties

Monte Carlo algorithms for physically based rendering typically impose the following requirements on the reflectance model:

- One must be able to evaluate the *brdf* for a given point on a surface and given incoming and outgoing directions.
- One should be able to sample the hemisphere at a given point and for a given outgoing (or incoming) direction, such that the probability density function (*pdf*) of the sampled direction follows the normalised *brdf* times the cosine factor as closely as possible. Ideally one should be able to sample according to the following *pdf*:

$$pdf(\Theta_i) = \frac{f_r(x, \Theta_i, \Theta_o)\ \cos\theta_i}{\int_{\Omega_x^{-1}} f_r(x, \Theta_i, \Theta_o)\ \cos\theta_i d\omega_i}$$

  In practice one often has to use an approximating *pdf* though. This type of importance sampling is commonly used in stochastic ray tracing where new directions have to be sampled recursively at each intersection point [7, 8, 9, 10, 11, 12, 13].

- Some algorithms require that *brdf* (times the cosine factor) can be integrated easily over the hemisphere to obtain the total hemispherical reflectivity ρ. Although the *brdf* can always be integrated numerically in principle, more efficient ways are much preferred. This problem is closely linked to the previous point because the reflectivity is the normalisation factor of the *pdf*. The reflectivity also determines whether the model is energy-conserving or not. One application is Russian roulette where the relative probabilities to reflect or absorb an incoming ray at a surface depend on it [14, 9, 10, 11, 12, 13]. Another application is the use of an ambient term times the *brdf* as a control variate for the rendering equation [15].

3 The modified Phong reflectance model

3.1 Definition

In 1975 Phong introduced a shading model which returns a colour for a point on a surface in the presence of some point light sources [16]. While the model does not have a physical basis and cannot be used as a *brdf* as it is, a physically plausible *brdf* which is very similar to the shading model is commonly used in rendering algorithms [1]. This modified Phong *brdf* can be written as the sum of a diffuse part and a specular part (Fig. 1):

$$\begin{aligned} f_r(x, \Theta_i, \Theta_o) &= f_{r,d}(x, \Theta_i, \Theta_o) + f_{r,s}(x, \Theta_i, \Theta_o) \\ &= k_d \frac{1}{\pi} + k_s \frac{n+2}{2\pi} cos^n \alpha \end{aligned}$$

where:

- α = the angle between the perfect specular reflective direction and the outgoing direction. Values larger than $\pi/2$ are clamped to $\pi/2$ in order not to get any negative cosine values.
- k_d = the diffuse reflectivity, ie. the fraction of the incoming energy that is reflected diffusely,
- k_s = the specular reflectivity, ie. the fraction of the *perpendicularly* incoming energy that is reflected specularly,
- n = the specular exponent. Higher values for n result in sharper specular reflections.

The additional factors in the model simplify the verification that a given *brdf* is energy-conserving, as we will demonstrate in the next paragraph.

3.2 Physical constraints

The model should at least conform to the physical constraints discussed earlier:

- It can easily be verified that the modified Phong reflectance model obeys to the Helmholtz-reciprocity.
- The total hemispherical reflectivity for the Phong model becomes:

$$\begin{aligned} \rho(x, \Theta_i) &= \int_{\Omega_x} f_r(x, \Theta_i, \Theta_o) \cos\theta_o d\omega_o \\ &= \int_{\Omega_x} (k_d \frac{1}{\pi} + k_s \frac{n+2}{2\pi} cos^n \alpha) \cos\theta_o d\omega_o \\ &= k_d + k_s \frac{n+2}{2\pi} \int_{\Omega_x} cos^n \alpha \cos\theta_o d\omega_o \\ &= \rho_d + \rho_s(x, \Theta_i) \end{aligned}$$

The latter integral reaches its maximum $2\pi/(n+2)$ for a perpendicularly incoming direction (where α equals θ_o), so:

$$\rho_{max} = k_d + k_s$$

Conservation of energy is therefore guaranteed iff:

$$k_d + k_s \leq 1$$

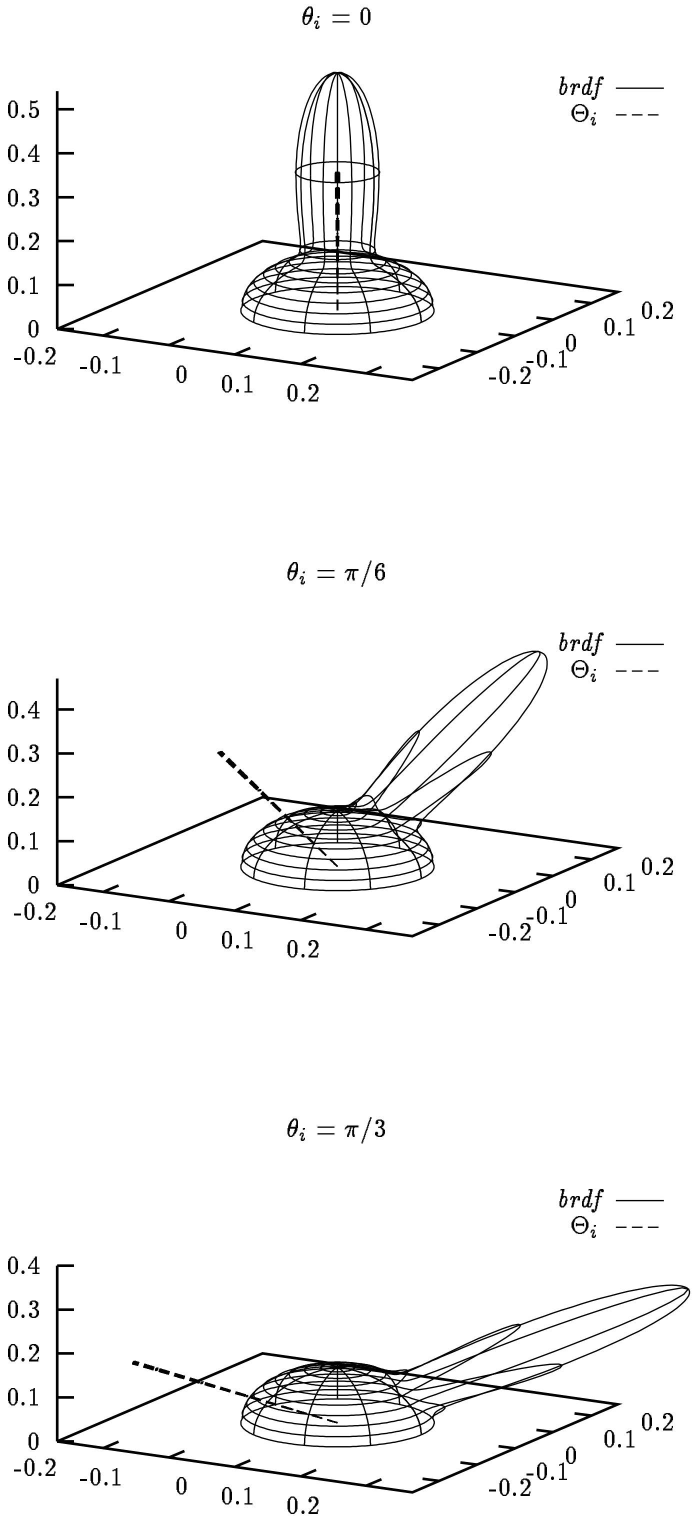

Figure 1: Polar diagrams of a modified Phong *brdf* with $k_d = 0.4$, $k_s = 0.05$ and $n = 50$ for various incoming angles.

3.3 Desired practical properties

We will now investigate how the modified Phong *brdf* meets the requirements in order to be of practical use:

- As the model is given in a closed form it can be evaluated directly for a specific point and an incoming and an outgoing direction, thereby satisfying the first condition.

- The *brdf* times the cosine factor cannot be used as a *pdf*. 'Hit or miss' sampling or rejection sampling will usually be inefficient, especially if there is a strongly peaked specular part. It is more effective to sample the diffuse part and the specular part of the *brdf* separately. Even this approach is difficult to implement correctly [7]. We propose a mathematically correct procedure in Section 3.3.2.

- The *brdf* times the cosine factor cannot be integrated analytically over the hemisphere because of the specular part. Firstly, the combination of the $\cos\alpha$ and $\cos\theta$ factors already precludes an exact analytical formula. Secondly, the $cos^n\alpha$ factor of the specular lobe may have a part below the surface and a part which is clamped to 0. Excluding these regions from the domain of integration causes some intricate integration boundaries, no matter which axes are used. We therefore estimate it by means of another Monte Carlo process, which is explained in Section 3.3.1.

3.3.1 Integrating the modified Phong *brdf*

Analytical integration of the specular part of the total hemispherical reflectivity is impossible. Monte Carlo integration which is optimised through importance sampling may therefore be a viable alternative. We propose to take one or more samples according to the following *pdf*:

$$pdf(\Theta_o) = \frac{n+1}{2\pi} cos^n\alpha$$

The primary estimate for the specular reflectivity is 0 if a direction below the surface is sampled. If a direction above the surface is selected the estimate becomes:

$$\begin{aligned} \langle \rho_s(x, \Theta_i) \rangle &= k_s \frac{n+2}{2\pi} \frac{2\pi}{n+1} \cos\theta_o \\ &= k_s \frac{n+2}{n+1} \cos\theta_o \end{aligned}$$

If $\frac{n+2}{n+1}\cos\theta_o > 1$ the estimate may conflict with the condition of energy conservation. In that case 1 will be a more accurate estimate.

3.3.2 Sampling the modified Phong *brdf*

We propose a sampling scheme for the modified Phong *brdf* which employs Russian roulette and importance sampling. It is suited for path tracing and similar Monte Carlo algorithms. The Russian roulette ensures that the recursion stops at some point. The importance sampling follows the *brdf* as closely as possible. Our first concern is correctness of the technique: the expected value of the estimate of the integral being computed should be the actual value. Our second concern is efficiency: importance sampling should be exploited maximally in order to minimise the variance.

We will demonstrate the sampling technique on a typical integral as it appears in the rendering equation:

$$\begin{aligned} L(x,\Theta_o) &= \int_{\Omega_x^{-1}} f_r(x,\Theta_i,\Theta_o) L(y,\Theta_i)\, \cos\theta_i d\omega_i \\ &= \int_{\Omega_x^{-1}} (k_d \frac{1}{\pi} \;+\; k_s \frac{n+2}{2\pi} cos^n\alpha) L(y,\Theta_i)\, \cos\theta_i d\omega_i \end{aligned}$$

Rewriting the integral already reveals the way in which it can be sampled:

$$\begin{aligned} L(x,\Theta_o) &= \rho_d \int_{\Omega_x^{-1}} [\frac{1}{\pi}\, \cos\theta_i]\, [\frac{k_d}{\rho_d} L(y,\Theta_i)]\, d\omega_i \\ &+ \rho_s \int_{\Omega_x^{-1}} [\frac{n+1}{2\pi} cos^n\alpha]\, [\frac{n+2}{n+1}\frac{k_s}{\rho_s} L(y,\Theta_i)\, \cos\theta_i]\, d\omega_i \\ &= \rho_d \int_{\Omega_x^{-1}} [\frac{1}{\pi}\, \cos\theta_i]\, L(y,\Theta_i)\, d\omega_i \\ &+ \rho_s \int_{\Omega_x^{-1}} [\frac{n+1}{2\pi} cos^n\alpha]\, [\frac{n+2}{n+1}\frac{k_s}{\rho_s} L(y,\Theta_i)\, \cos\theta_i]\, d\omega_i \end{aligned}$$

A first stochastic variable determines the type of reflection. It is sampled uniformly over the interval $[0,1]$ and the following actions are taken accordingly:

- $0 \le \xi < \rho_d$: take a diffuse sample and compute its contribution,
- $\rho_d \le \xi < \rho_d + \rho_s$: take a specular sample and compute its contribution,
- $\rho_d + \rho_s \le \xi \le 1$: the contribution is 0.

This type of importance sampling ensures that the most appropriate amount of work is put into each type of reflection (lacking any further information about $L(y,\Theta_i)$). Because the actual value of ρ_s cannot be computed analytically, an estimate $\langle\rho_s\rangle$ has to be used, as described above. The rewritten integral expression shows that any value can be used actually, so alternatively one can simply use a value of k_s. The expected value will remain correct, but the variance may change slightly as a result.

Diffuse and specular samples are then handled as follows:

- A diffuse sample can be sampled according to the cosine distribution:

  $$pdf(\Theta_i) = \frac{1}{\pi}\, \cos\theta_i$$

  The estimator for the integral then becomes:

  $$\langle L(x,\Theta_o)\rangle = L(y,\Theta_i)$$

- It is impossible to sample the complete specular part of the *brdf* directly because of the cosine factor. But it is feasible to sample according to the $cos^n\alpha$ factor, as we have done to compute the specular reflectivity:

  $$pdf(\Theta_i) = \frac{n+1}{2\pi} cos^n\alpha$$

  We then have to adapt the estimator accordingly. If a direction below the surface is sampled the weight is 0. One can see this as sampling a direction for which the incoming radiance is 0. If the direction points above the surface the estimator becomes:

  $$\langle L(x,\Theta_o)\rangle = \frac{n+2}{n+1}\frac{k_s}{\langle\rho_s\rangle}\, \cos\theta_i$$

The *pdf* ensures at least that samples are chosen in the specular lobe. This presents a major improvement over some sort of uniform sampling, because the lobe may be very sharp for highly specular surfaces which have a large specular exponent.

Sampling directional distributions

In the previous paragraphs two types of *pdfs* for directions were used. The space of directions in 3D-space is two-dimensional, so these *pdfs* can be sampled by selecting 2 uniform stochastic variables ξ_1 and ξ_2 over the interval $[0, 1]$ and transforming them appropriately:

- For $pdf(\Theta) = \frac{1}{\pi}\cos\theta$ the direction Θ can be specified in terms of (θ, ϕ), where θ is the polar angle with the surface normal and ϕ is the azimuthal angle:

  $$(\theta, \phi) = (\ \mathrm{acos}\sqrt{\xi_1}, 2\pi\xi_2)$$

  In terms of direction vector coordinates (x, y, z) this becomes:

  $$\begin{aligned}(x, y, z) &= (\ \sin\theta\ \cos\phi,\ \sin\theta\ \sin\phi,\ \cos\theta) \\ &= (\sqrt{1-\xi_1}\ \cos(2\pi\xi_2), \sqrt{1-\xi_1}\ \sin(2\pi\xi_2), \sqrt{\xi_1})\end{aligned}$$

- For $pdf(\Theta) = \frac{n+1}{2\pi}\cos^n\alpha$, which is actually a generalisation of the previous *pdf*, the direction Θ can be specified in terms of (α, ϕ), where:

  $$(\alpha, \phi) = (\ \mathrm{acos}\xi_1^{\frac{1}{n+1}}, 2\pi\xi_2)$$

  In terms of direction vector coordinates (x, y, z) this now becomes:

  $$\begin{aligned}(x, y, z) &= (\ \sin\alpha\ \cos\phi,\ \sin\alpha\ \sin\phi,\ \cos\alpha) \\ &= (\sqrt{1-\xi_1^{\frac{2}{n+1}}}\ \cos(2\pi\xi_2), \sqrt{1-\xi_1^{\frac{2}{n+1}}}\ \sin(2\pi\xi_2), \xi_1^{\frac{1}{n+1}})\end{aligned}$$

4 Parametrisations of the hemisphere

Consider the integral expressing the total hemispherical reflectivity:

$$\rho(x, \Theta_i) = \int_\Omega f_r(x, \Theta_i, \Theta_o)\ \cos\theta_o d\omega_o$$

Related integrals and integral equations such as the rendering equation are usually very similar, just containing an additional –mostly unknown– factor such as the radiance function. There are a number of alternative ways to parametrise the direction Θ_o in this integral. The parametrisation has a large influence on the efficiency of any deterministic or stochastic sampling scheme. In fact, one can easily show that in Monte Carlo integration performing importance sampling in one parameter space is equivalent to performing uniform sampling in another parameter space. Transformation to a parameter space where the integrand is smooth is preferable, as smoother functions are easier to integrate than wildly varying functions.

- In the classical (θ, ϕ) parameter space (Fig. 2) this integral can be written out as:

  $$\rho(x, \Theta_i) = \int_0^{2\pi} \int_0^{\pi/2} f_r(x, \Theta_i, \Theta_o)\ \cos\theta_o\ \sin\theta_o d\theta_o d\phi$$

  Even without considering the specular peak of the *brdf* one can see that the function has a strong relative variation over the domain, descending to 0 for $\theta = 0$ and $\theta = \pi/2$.

- In the (c, ϕ) parameter space (Fig. 3) where $(c, \phi) = (1 - \cos\theta, \phi)$ the integral is written out as:

  $$\rho(x, \Theta_i) = \int_0^{2\pi} \int_0^{1} f_r(x, \Theta_i, \Theta_o)(1 - c)\, dc\, d\phi$$

  Equally sized elements in the parameter space correspond to equally sized spatial angles. The integrand no longer goes to 0 for $\theta = 0$ $(c = 0)$, but still descends to 0 for $\theta = \pi/2$ $(c = 1)$.

- In the (ξ_1, ξ_2) parameter space (Fig. 4) where $(\theta, \phi) = (\mathrm{acos}\sqrt{1-\xi_1}, 2\pi\xi_2)$ $= (\mathrm{asin}\sqrt{\xi_1}, 2\pi\xi_2)$ the integral becomes:

  $$\rho(x, \Theta_i) = \pi \int_0^1 \int_0^1 f_r(x, \Theta_i, \Theta_o) d\xi_1 d\xi_2$$

  This parametrisation has the advantage that equally sized elements correspond to equally sized spatial angles weighed by the cosine term. As a result a perfectly diffuse part of a *brdf* will appear as a constant part in the integrand which can be sampled perfectly by a single sample.

- Yet another alternative is the (x, y) parameter space (Fig. 5) which is also used in the Nusselt analog. Directions are determined by their projections in the (x, y) plane: $(x, y) = (\ \sin\theta\ \cos\phi,\ \sin\theta\ \sin\phi)$, or the inverse transformation: $(\theta, \phi) = (\ \mathrm{asin}\sqrt{x^2 + y^2},\ \mathrm{atan2}(y, x))$. This yields an integral over the unit disc:

  $$\rho(x, \Theta_i) = \int\int_{\bigcirc} f_r(x, \Theta_i, \Theta_o) dx dy$$

As in the (ξ_1, ξ_2) parametrisation equally sized elements correspond to equally sized spatial angles weighed by the cosine term. Although the integration domain is more complex, this parametrisation has the interesting property that it is contiguous, in the sense that neighbouring points in parameter space map to neighbouring directions.

The above parametrisations do not depend on the *brdf*. For more specular *brdf*s though the specular peak may easily become very large and thin. Sampling then becomes ineffective in these schemes because of the quick variations of the integrand. Alternative parametrisations such as the one used implicitly in the importance sampling technique for the modified Phong reflectance model explained in Section 3.3.2 are then much preferred.

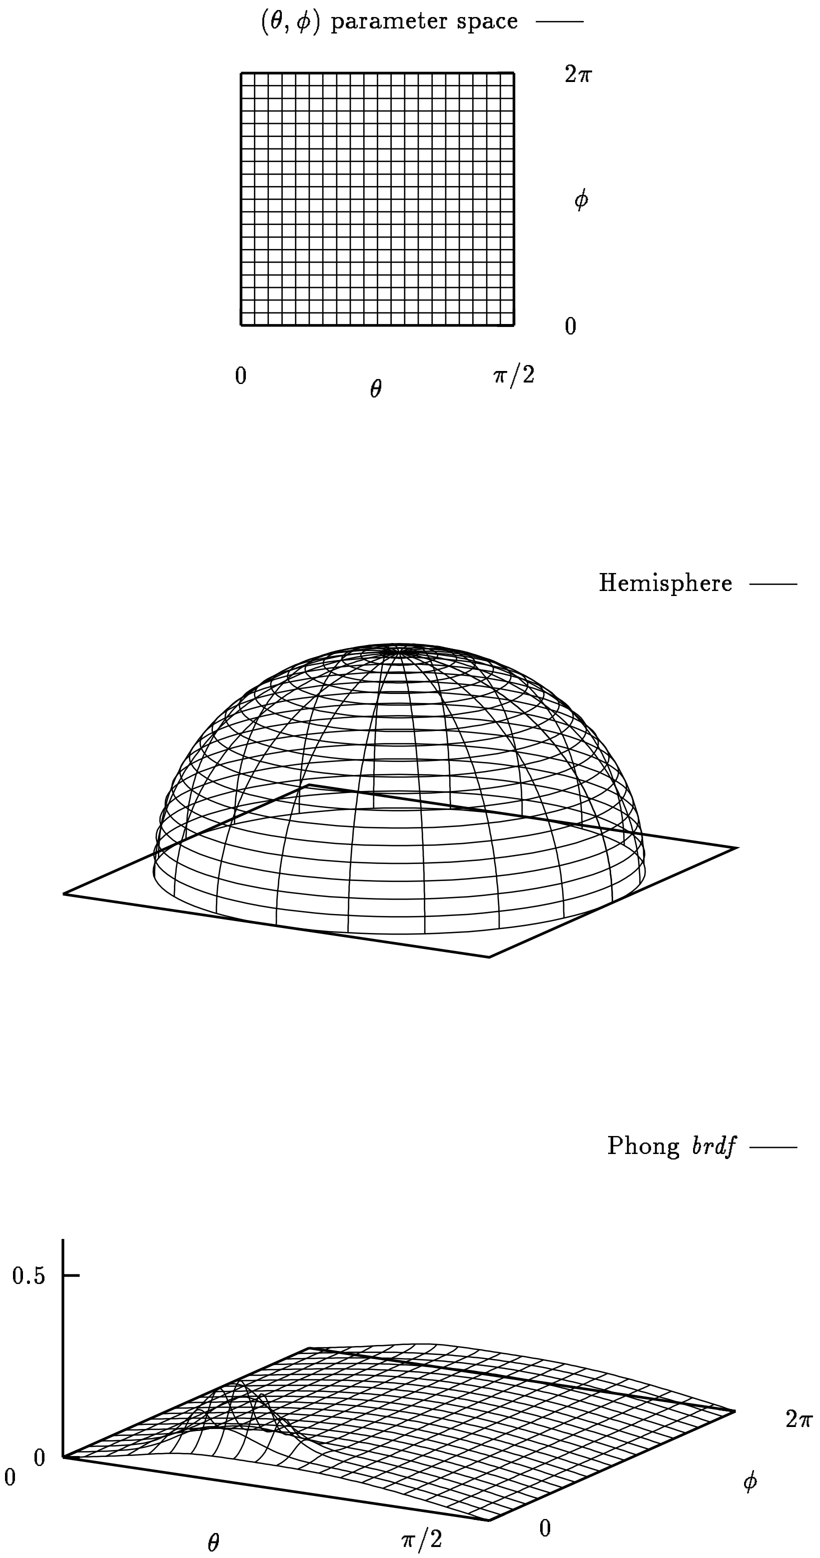

Figure 2: The (θ, ϕ) parametrisation with its parameter space, its mapping onto a hemisphere in a polar diagram and the resulting integrand with a Phong *brdf* ($k_d = 0.4$, $k_s = 0.05$, $n = 50$ and $\theta_i = \pi/6$) plotted as a function of θ and ϕ.

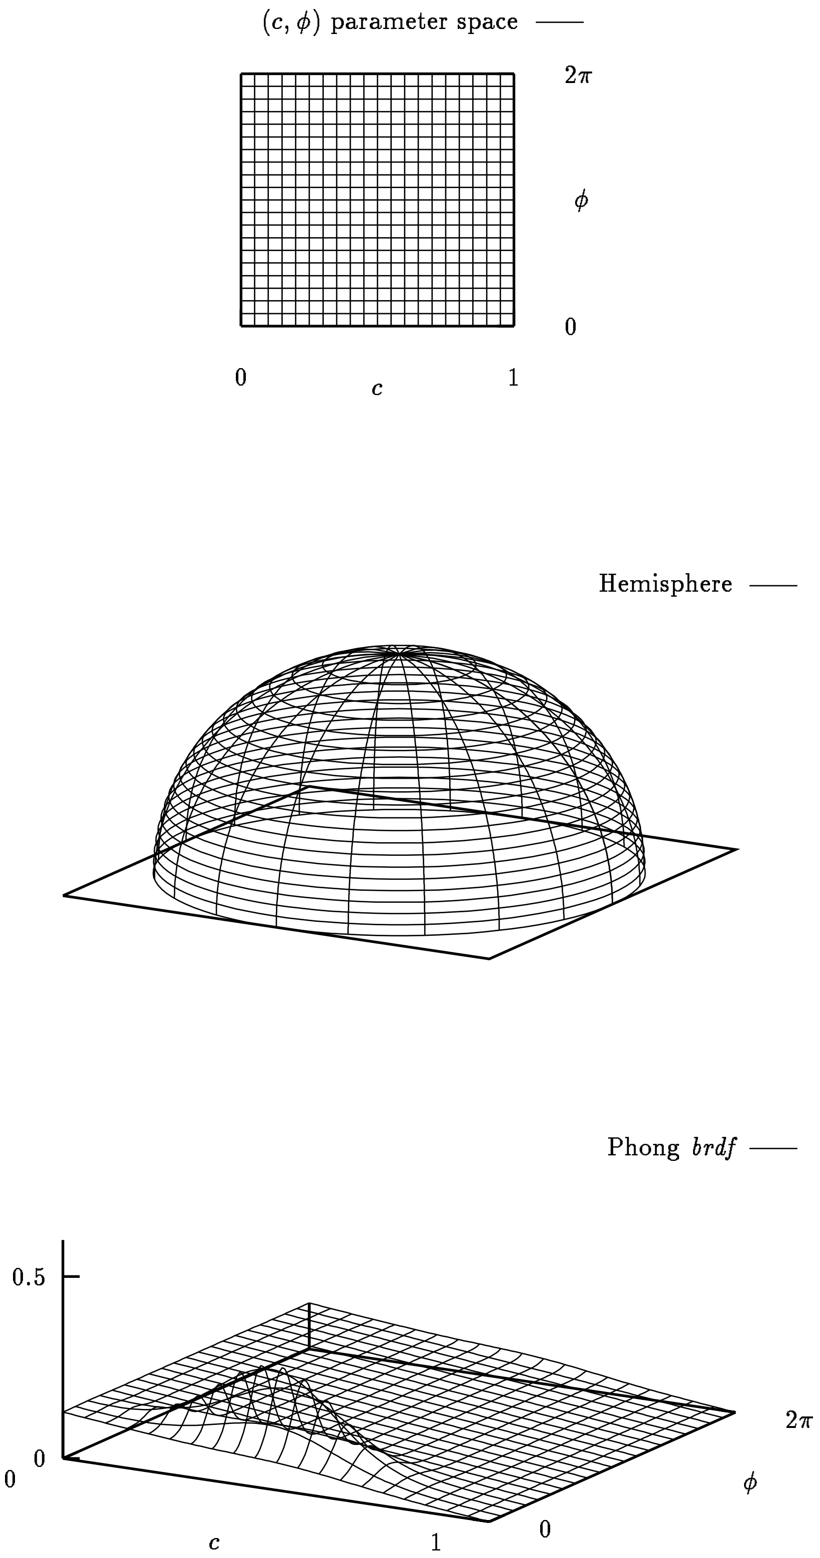

Figure 3: The (c, ϕ) parametrisation with its parameter space, its mapping onto a hemisphere in a polar diagram and the resulting integrand with a Phong *brdf* plotted as a function of c and ϕ.

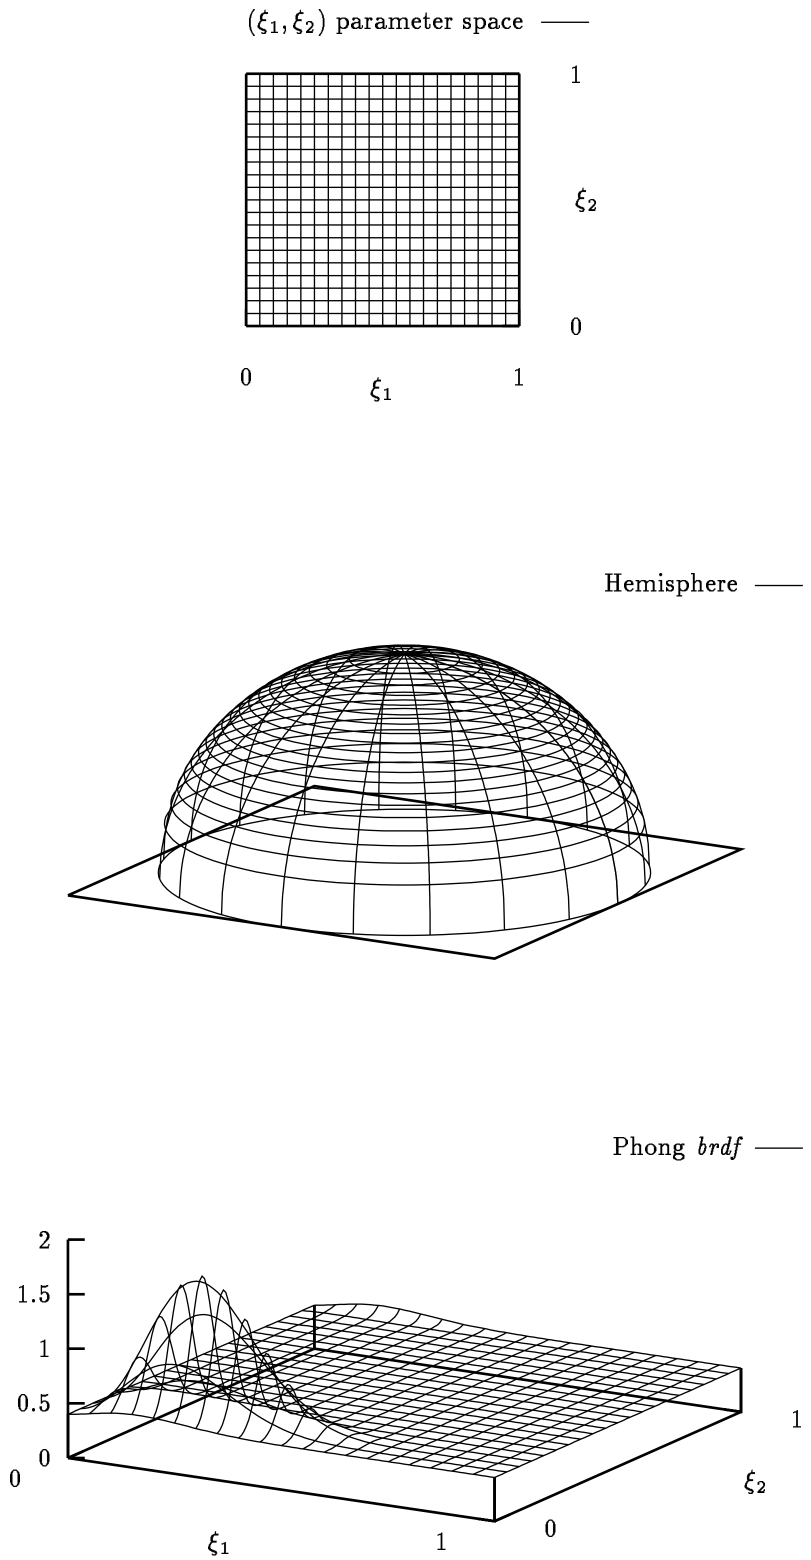

Figure 4: The (ξ_1, ξ_2) parametrisation with its parameter space, its mapping onto a hemisphere in a polar diagram and the resulting integrand with a Phong *brdf* plotted as a function of ξ_1 and ξ_2.

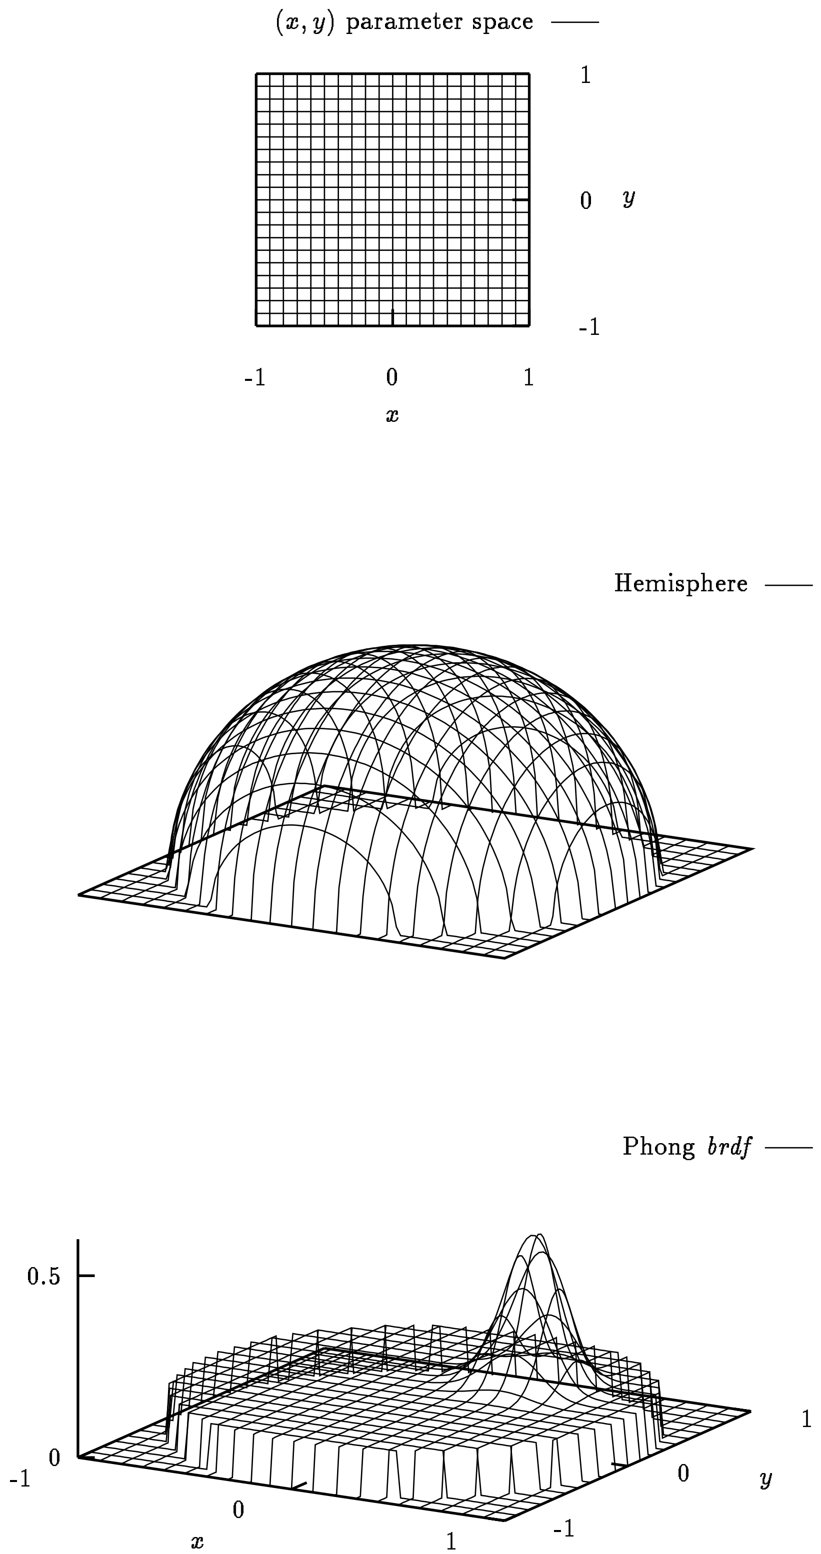

Figure 5: The (x, y) parametrisation with its parameter space, its mapping onto a hemisphere in a polar diagram and the resulting integrand with a Phong *brdf* plotted as a function of x and y.

5 Verifying the implementation of a reflectance model

An implementation of a *brdf* model will basically contain the three functions which have been outlined in Section 2.2. They respectively evaluate, sample and integrate the model for a given set of input parameters (the parameters describing the surface properties at a given position and the incoming direction). Experience has shown that is is very easy to make mistakes while implementing even the simplest model. A slip of attention or a typing error may have a substantial influence on the results even though images created using the implementation may look acceptable. The physical constraints, Helmholtz-reciprocity and conservation of energy, can be tested easily with some random input parameters. We suggest an additional test which helps finding contradictions between the three basic functions.

The idea is to compute the fractions of the power that are reflected through a set of solid angles, for some incoming direction. These partial reflectivities have to be the same when computed using different techniques. The partial reflectivity through a solid angle Ω_k is:

$$\Delta\rho_k(x, \Theta_i) = \int_{\Delta\Omega_k} f_r(x, \Theta_i, \Theta_o) \cos\theta_o d\omega_o$$

We have chosen to work in the (x, y) parameter space, because diagrams in this space are easy to interpret and conversion between an (x, y, z) direction vector and the (x, y) parameters is straightforward. The outgoing solid angles $\Delta\Omega_k$ are chosen to correspond to the squares of size $\Delta x \times \Delta y$ in a uniform grid in the (x, y) parametrisation, as in Fig. 5. For these outgoing solid angles $\Delta\Omega_k$ the partial reflectivities $\Delta\rho_k$ can be evaluated using two alternative Monte Carlo techniques. One technique uses the sampling function of the implementation and the other technique is based on the evaluating function. In order for the implementation to be consistent the results have to match within the accuracy of the computations. They can easily be visualised in a three-dimensional graph on the grid.

- A first primary estimator for the fluxes uses the sampling function of the implementation. All fluxes are computed at the same time. The *subcritical pdf* implemented in the function samples a direction over the entire hemisphere:

  $$spdf(\Theta_o) = f_r(x, \Theta_i, \Theta_o) \cos\theta_o$$

  The correct estimator is then assigned to the appropriate solid angle:

  $$\langle \Delta\rho_k(x, \Theta_i) \rangle = 2\pi$$

  Averaging several primary estimators yields more accurate secondary estimators for all the partial reflectivities (Fig. 6). The variance of this estimator is high because each partial reflectivity on its own is effectively computed using 'hit or miss' sampling. For more specular *brdf*s it may be increasingly efficient however, as alternative sampling approaches will usually fail to capture the specular peak adequately.

- The alternative primary estimator uses the evaluating function of the implementation. All partial reflectivities are computed in turn. The uniform *pdf* over the appropriate grid square in the (x, y) parameter space for outgoing solid angle $\Delta\Omega_k$ is:

$$pdf(p_o) = \frac{1}{\Delta x \Delta y}$$

The sampled point p_o is projected back onto the hemisphere to obtain the outgoing direction Θ_o. If there is no corresponding direction the estimator is 0. Otherwise the corresponding estimator can be evaluated as:

$$\langle \Delta\rho_k(x, \Theta_i) \rangle = \Delta x \Delta y f_r(x, \Theta_i, \Theta_o)$$

For perfectly diffuse surfaces the *brdf* is a constant and the estimator is perfect. Alternatively several primary estimators will yield more accurate secondary estimators (Fig. 7).

- After these computations the result of function which returns the total reflectivity $\rho(x, \Theta_i)$ can be checked by computing the estimate:

  $$\langle \rho(x, \Theta_i) \rangle = \sum_k \langle \Delta\rho_k(x, \Theta_i) \rangle$$

  where the estimates of the partial reflectivities $\langle \Delta\rho_k \rangle$ have been computed using one of the two algorithms above.

Of course these verification techniques can never *prove* the correctness of the implementation because they can only verify the results for a few random incoming directions. Nevertheless, they have already proven to be helpful in practice for demonstrating implementational errors.

6 Conclusion

In this text we have discussed some physical and practical properties of *bidirectional reflection distribution functions* and their models. We have demonstrated how the modified Phong reflectance model can be applied in a mathematically correct way for physically based rendering using Monte Carlo techniques. An overview of some alternative parametrisations of the directions in the hemisphere shows why some parametrisations and importance sampling techniques may be preferable over others. We have presented a verification technique in one of these parameter spaces which may be helpful for finding implementational errors.

7 Acknowledgement

The first author would like to acknowledge the financial support by a grant from the Belgian "Instituut tot Aanmoediging van Wetenschappelijk Onderzoek in Nijverheid en Landbouw" (I.W.O.N.L. #910184).

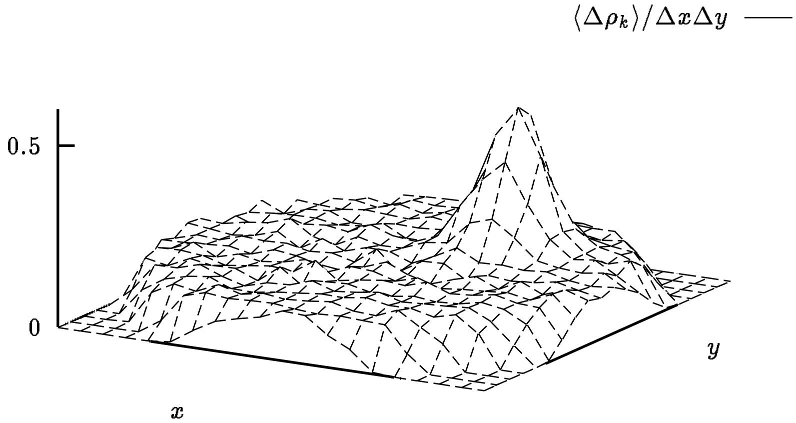

Figure 6: The partial reflectivities $\Delta\rho_k$ computed by sampling according to the Phong *brdf* with $k_d = 0.4$, $k_s = 0.05$ and $n = 50$ for an incoming angle $\theta_i = \pi/6$. A total of 100.000 samples was taken.

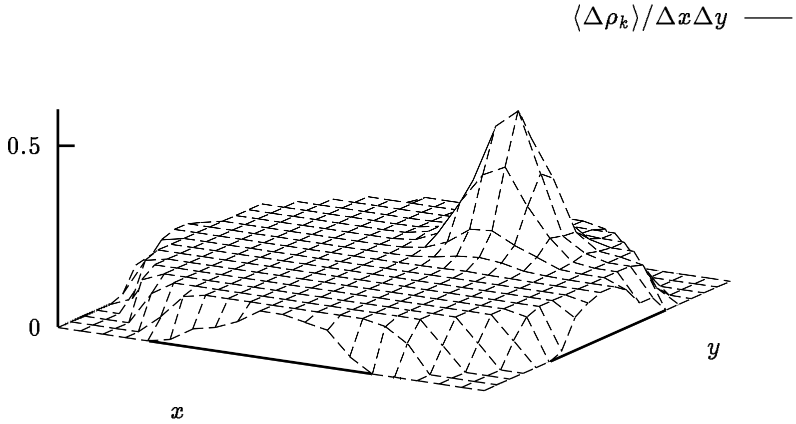

Figure 7: The partial reflectivities $\Delta\rho_k$ computed by sampling the grid squares in the (x, y) parameter space uniformly. 10 samples per grid square were taken.

References

[1] R. Lewis, "Making shaders more physically plausible," *Computer Graphics Forum*, vol. 13, pp. 109–120, June 1994.

[2] R. Cook and K. Torrance, "A reflectance model for computer graphics," *ACM Transactions on Graphics*, vol. 1, pp. 7–24, Jan. 1982.

[3] X. He, K. Torrance, F. Sillion, and D. Greenberg, "A comprehensive physical model for light reflection," *Computer Graphics*, vol. 25, pp. 175–186, July 1991.

[4] G. Ward, "Measuring and modeling anisotropic reflection," *Computer Graphics*, vol. 26, pp. 265–272, July 1992.

[5] C. Schlick, "A customizable reflectance model for everyday rendering," in *Proceedings of the Fourth Eurographics Workshop on Rendering*, (Paris, France), pp. 73–83, June 1993.

[6] C. Schlick, "A survey of shading and reflectance models," *Computer Graphics Forum*, vol. 13, pp. 121–131, June 1994.

[7] P. Shirley, *Physically Based Lighting Calculations for Computer Graphics.* PhD thesis, University of Illinois, Nov. 1990.

[8] P. Shirley and C. Wang, "Distribution ray tracing: Theory and practice," in *Proceedings of the Third Eurographics Workshop on Rendering*, (Bristol, UK), pp. 33–43, May 1992.

[9] S. Pattanaik and S. Mudur, "Computation of global illumination by monte carlo simulation of the particle model of light," in *Proceedings of the Third Eurographics Workshop on Rendering*, (Bristol, UK), pp. 71–83, May 1992.

[10] S. Pattanaik, *Computational Methods for Global Illumination and Visualisation of Complex 3D Environments.* PhD thesis, Birla Institute of Technology & Science, Pilani, India, Feb. 1993.

[11] P. Dutré, E. Lafortune, and Y. Willems, "Monte carlo light tracing with direct computation of pixel intensities," in *Proceedings of CompuGraphics*, (Alvor, Portugal), pp. 128–137, Dec. 1993.

[12] E. Lafortune and Y. Willems, "Bi-directional path tracing," in *Proceedings of CompuGraphics*, (Alvor, Portugal), pp. 145–153, Dec. 1993.

[13] E. Lafortune and Y. Willems, "A theoretical framework for physically based rendering," *Computer Graphics Forum*, vol. 13, pp. 97–107, June 1994.

[14] J. Arvo and D. Kirk, "Particle transport and image synthesis," *Computer Graphics*, vol. 24, no. 4, pp. 63–66, 1990.

[15] E. Lafortune and Y. Willems, "The ambient term as a variance reducing technique for monte carlo ray tracing," in *Proceedings of the Fifth Eurographics Workshop on Rendering*, (Darmstadt, Germany), pp. 163–171, June 1994.

[16] B. Phong, "Illumination for computer generated pictures," *Communications of the ACM*, vol. 18, no. 6, pp. 311–317, 1975.