



A Microfacet-based BRDF Generator

Michael Ashikhmin

Simon Premože

Peter Shirley

University of Utah

www.cs.utah.edu

Abstract

A method is presented that takes as an input a 2D microfacet orientation distribution and produces a 4D bidirectional reflectance distribution function (BRDF). This method differs from previous microfacet-based BRDF models in that it uses a simple shadowing term which allows it to handle very general microfacet distributions while maintaining reciprocity and energy conservation. The generator is shown on a variety of material types.

CR Categories: I.3.7 [Computing Methodologies]: Computer Graphics—3D Graphics

Keywords: Reflectance & Shading Models, Rendering

1 Introduction

Physically-based rendering systems describe reflection behavior using the *bidirectional reflectance distribution function* (BRDF) [7]. At a given point on a surface the BRDF is a function of two directions, one toward the light and one toward the viewer. The characteristics of the BRDF will determine what “type” of material the viewer thinks the displayed object is composed of, so the choice of BRDF model and its parameters is important. There are a variety of basic strategies for modeling BRDFs that we categorize as follows.

Direct measurement. BRDFs can be measured directly using *gonioreflectometers* which mechanically vary the direction to a small light source and a spectral sensor and thus collect a large number of point samples for the BRDF [7]. Simpler and less accurate devices can also be constructed using CCD imaging devices [26]. More complex CCD devices can also be used which gather data quickly with accuracy almost that of full gonioreflectometry [12]. If enough is known about the microstructure of a material, a BRDF can be simulated by using a *virtual gonioreflectometer*, where statistical ray tracing followed by density estimation is used to create BRDF data [3, 5, 27].

Empirical methods. There exist a variety of purely empirical reflection models, the most familiar being the models introduced by Gouraud [6] and Phong [15]. These two initial models were meant to be used with hand-chosen parameters, and thus these parameters are intuitive. A variety of more complex methods have been introduced to improve characteristics of the Phong model for efficiency [19], to include anisotropy [26], and enforce physical constraints such as reciprocity [9]. Other models have been developed to fit measurement data as opposed to being intuitive [10].




Figure 1: Images generated using the new BRDF model with unusual microfacet distributions. The BRDFs used to create these images are both reciprocal and energy-conserving. The only illumination is a small distant source, and the highlights will stay unchanged if the spheres rotate about the axes through their north and south poles.

Height correlation methods. In these methods a random rough surface is a realization of some Gaussian random process. Such a process can be described by its correlation function which is directly related to surface height correlations. This is the most complete surface representation used in computer graphics. Some of the most detailed descriptions of light scattering by a surface, including wave optics effects, were obtained using this approach [8, 22].

Microfacet methods. Somewhere between the height correlation methods and empirical methods lie models based on microfacet theory [2, 4]. Microfacet models assume the surface consists of a large number of small flat “micromirrors” (facets) each of which reflect light only in the specular direction. By computing the number of visible microfacets at the appropriate orientation to specularly reflect light from the source to the viewer, one can determine the BRDF.

All of these methods have their place. In applications where little is known about the low-level properties of the surface, measurement is essential. Where physical optics effects are important, height correlation methods should be used. Our interest is in visual computer graphics applications which do not have obvious physical optics effects (e.g. metal with relatively large scratches, fabric). The lesson from empirical models is that in many cases viewers are not particularly sensitive to the fine details of light scattering as long as the main character of the reflection is conveyed correctly. This paper uses this aspect of human sensitivity to suggest a new microfacet model specifically intended to capture the main character of reflection.

Microfacet models are able to capture the main character of reflection for surfaces whose appearance is dominated by surface scattering. Although microfacet models lack the precision of height correlation methods, they tend to be more intuitive with simpler expressions. However, to date there has been no microfacet model that is reasonably general in its assumptions, maintains a simple formulation, and conserves energy. In this paper we develop a model with all of these characteristics by introducing assumptions about surfaces that we believe are reasonable. These assumptions allow




Figure 2: *Geometry of reflection. Note that \mathbf{k}_1 , \mathbf{k}_2 , and \mathbf{h} share a plane, which usually does not include \mathbf{n} . On the left, the microfacet can “see” in directions \mathbf{k}_1 and \mathbf{k}_2 so it contributes to the BRDF. On the right, direction \mathbf{k}_2 is blocked and the microfacet does not contribute. Note that the microfacet distribution is not restricted to height fields.*

us to create a relatively simple formula for the probability that a microfacet at a certain orientation is visible to the light/viewer. The BRDF produced by this process is compact, reciprocal and energy-conserving with only mild restrictions on the distribution of microfacet orientation (e.g., the very general distributions in Figure 1).

Our assumptions and guiding principles in relation to microfacet theory are given in Section 2. Formalisms are developed in Section 3. The key development of the paper, a simplified shadowing term, is introduced in Section 4, and the resulting BRDF is derived. Section 5 shows that this BRDF model conserves energy, and derives a diffuse term to account for secondary and subsurface reflection. The model is applied to a variety of surfaces in Section 6. This last section serves as a set of case-studies which both show how the model can be applied, and that it is more general than previous microfacet approaches. We believe the only other method which is able to handle such a diverse set of surface microgeometries is the “virtual gonireflectometer” approach involving explicit modeling of the surface structure and statistical averaging the results of light scattering simulations.

2 Overview

The strategy behind our model is in balancing issues of practicality and accuracy to produce a simple formulation that is still expressive, reciprocal, and conserves energy. In this section we discuss the basic ideas of microfacet models, as well as our strategy for using this theory to produce BRDFs. Important symbols used in the paper are listed in Table 1.

Microfacet models assume that the surface consists of a large number of small flat “micromirrors” (facets) each of which reflects light only in the specular direction with respect to its own normal \mathbf{h} (Figure 2) and the overall appearance of the surface is governed by two assumptions:

- the microfacet normals have an underlying probability density function $p(\mathbf{h})$.
- a microfacet contributes to BRDF for a given pair of directions if and only if it is visible (not shadowed) relative to the lighting direction \mathbf{k}_1 and the viewing direction \mathbf{k}_2 .

The BRDF for a given direction pair $(\mathbf{k}_1, \mathbf{k}_2)$ is determined entirely by the Fresnel reflectance for that angle, the fraction of microfacets with normal vector \mathbf{h} exactly between \mathbf{k}_1 and \mathbf{k}_2 , and the *shadowing term*: the fraction of those microfacets which are visible to both eye and light (Figure 2). Microfacet theory’s only knowledge of the

($\mathbf{a}\mathbf{b}$)	scalar (dot) product of vectors \mathbf{a} and \mathbf{b}
\mathbf{k}_1	normalized vector to light
\mathbf{k}_2	normalized vector to viewer
\mathbf{n}	surface normal to macroscopic surface
$\rho(\mathbf{k}_1, \mathbf{k}_2)$	BRDF
\mathbf{h}	normalized half-vector between \mathbf{k}_1 and \mathbf{k}_2
$p(\mathbf{h})$	probability density function of microfacet normals
$F(\cos \theta)$	Fresnel reflectance for incident angle θ
$P(\mathbf{k}_1, \mathbf{k}_2, \mathbf{h})$	Probability that light from \mathbf{k}_1 reflecting in direction \mathbf{k}_2 is <i>not</i> shadowed
$\langle f \rangle$	average of function f over distribution $p(\mathbf{h})$ (see Equation 9)
$\Omega_+(\mathbf{k})$	set of directions \mathbf{h} where $(\mathbf{h}\mathbf{k}) > 0$ (see Figure 4)
$g(\mathbf{k})$	average of positive $(\mathbf{h}\mathbf{k})$ (see Equation 18)

Table 1: Important terms used in the paper

surface configuration is $p(\mathbf{h})$, and this alone does not uniquely determine the shadowing term. However, the shadowing term is still heavily constrained by energy conservation.

The shadowing term is the most complex part of most microfacet-based models, even if additional $p(\mathbf{h})$ -specific information about the surface geometry is used. Because there are many possible surface geometries that are consistent with a given $p(\mathbf{h})$, it is the case that no specific shadowing function is “right”. We believe that in most cases the shape of $p(\mathbf{h})$ function itself has a much greater impact on the appearance than the shadowing. This suggests the key idea in this paper: the shadowing term should be made as simple as possible while remaining physically plausible. Such a shadowing term is developed in Section 4. This key extension of the standard microfacet theories allows us to construct a general procedure to create a BRDF for a statistical surface starting from $p(\mathbf{h})$.

Note that surface description in the language of $p(\mathbf{h})$ is less detailed than that of using height correlation functions. Nevertheless, we believe that the microfacet normal distribution is more intuitive to deal with than the correlation functions. As we will emphasize in Section 6, enough useful information about $p(\mathbf{h})$ can be obtained from general notion of surface structure obtained through visual examination of the surface and the specular reflection highlight. Moreover, attempting to obtain more detailed information about the distribution might not be worth the effort. As in any other model, we make simplification in our approach which affect the final result, but what we are trying to do is generate a physically plausible BRDF having the general character of the surface reflection while restricting the range of allowed surface microstructures as little as possible. This is in contrast to most other physics-based approaches which concentrate on a particular type of surface, usually Gaussian height field, and emphasize the need for precise knowledge of surface characteristics.

Some care should be exercised when specifying $p(\mathbf{h})$. In particular, because we do not make the common assumption of a surface being a height field, in this general case $p(\mathbf{h})$ should refer only to the distribution of “visually important” or “surface” part of the microfacets. For example, a homogeneous porous substance thought of as a collection of microfacets will have an overall “volume” distribution of microfacets $p_v(\mathbf{h}) = \text{const}$ over the whole sphere of directions. However, that most of these microfacets will be completely hidden and will not be of any significance for the scattering process which occurs on the surface. In this case it is rather difficult to separate surface from the rest of the substance and judge the exact shape of $p(\mathbf{h})$. Fortunately, because we are not trying to reproduce all the details of the reflection function, a reasonable guess for $p(\mathbf{h})$

is all we need and for this surface; it might be $p(\mathbf{h}) = \text{const}$ in the upper hemisphere and $p(\mathbf{h}) = 0$ in the lower one. Note, that by making this particular choice for $p(\mathbf{h})$ the surface is restricted to be a height field. The initial choice can be refined later if necessary but in this particular case it the surface will be mostly diffuse and small refinements will not dramatically change the appearance.

We are concerned with single-bounce reflections from the microfacets and stay within the limits of geometric optics and Fresnel reflection. The result is a new form of the specular component of the BRDF which constitutes the main contribution of the paper. The complete BRDF can also have a diffuse term which accounts for multiple bounces and subsurface scattering. This issue along with other important properties of the BRDFs produced with the generator are briefly discussed in Section 5. Our framework is modular and allows the user to choose the form of the final BRDF most appropriate for the particular application.

3 Microfacet Theory

We now review the main results of microfacet theory as developed by Torrance and Sparrow [23] and later introduced to computer graphics by Cook and Torrance [4]. We follow their approach of considering a collection of microfacets of small but finite size, and we derive the basic formula for BRDF in terms of quantities convenient for our model.

The quantity we wish to derive an expression for is the BRDF $\rho(\mathbf{k}_1, \mathbf{k}_2)$ which gives the ratio of radiance observed by a viewer in the direction \mathbf{k}_2 to irradiance from infinitesimal solid angle about \mathbf{k}_1 . Throughout the paper, all vectors are shown in **bold**. They are assumed to be normalized, and all quantities with subscript 1 refer to incident direction while those with subscript 2 belong to the outgoing direction. Both \mathbf{k}_1 and \mathbf{k}_2 and all normals point outward from the surface. If we expose the surface to a uniform radiance of L_1 coming from a small solid angle $\delta\omega_1$ around \mathbf{k}_1 , the outgoing radiance in direction \mathbf{k}_2 will be

$$L_2 = \rho(\mathbf{k}_1, \mathbf{k}_2) L_1 (\mathbf{k}_1 \mathbf{n}) \delta\omega_1, \quad (1)$$

where \mathbf{n} is the surface geometric normal and two vectors written next to each other in parenthesis denotes their scalar product, i.e., the cosine of the angle between them. The use of δ is not standard notation, but is used to make the algebra less cluttered without losing the gist of the argument. By the definition of radiance, if $(\mathbf{k}_2 \mathbf{n})A$ is the projected surface element area in the direction \mathbf{k}_2 and $\delta E(\mathbf{k}_1 \rightarrow \mathbf{k}_2)$ is the power reflected by the surface in the direction \mathbf{k}_2 , then

$$L_2 = \frac{\delta E(\mathbf{k}_1 \rightarrow \mathbf{k}_2)}{A(\mathbf{k}_2 \mathbf{n}) \delta\omega_2}, \quad (2)$$

and BRDF can be written as

$$\rho(\mathbf{k}_1, \mathbf{k}_2) = \frac{\delta E(\mathbf{k}_1 \rightarrow \mathbf{k}_2)}{A L_1 (\mathbf{k}_2 \mathbf{n}) (\mathbf{k}_1 \mathbf{n}) \delta\omega_1 \delta\omega_2}. \quad (3)$$

Only a fraction of all microfacets will participate in scattering the energy from \mathbf{k}_1 to \mathbf{k}_2 . If the number of these active microfacets is N_{active} and all microfacets have the same area A_{mf} , their total projected area in the direction of \mathbf{k}_1 is $N_{active} A_{mf} \langle (\mathbf{k}_1 \mathbf{n}) \rangle$ and the total scattered power is

$$\delta E(\mathbf{k}_1 \rightarrow \mathbf{k}_2) = L_1 \delta\omega_1 N_{active} A_{mf} \langle (\mathbf{k}_1 \mathbf{n}) \rangle F((\mathbf{k}_1 \mathbf{h})), \quad (4)$$

where \mathbf{h} is the normalized half-vector between \mathbf{k}_1 and \mathbf{k}_2 and $F((\mathbf{k}_1 \mathbf{h}))$ is Fresnel coefficient giving the fraction of incoming light which is specularly reflected by a microfacet. Note that we will drop subscripts in our notations if either of incoming and outgoing direction can be used in an expression (e.g., $\langle (\mathbf{k}_1 \mathbf{h}) \rangle$).

Out of the total of N surface microfacets, only $N p(\mathbf{h}) \delta\omega_h$ will have their normals oriented in the appropriate direction. The density $p(\mathbf{h})$ does not specify all surface properties uniquely, but in our simplified approach this is the only characteristic of the surface we will use in our analysis. Note that this function operates in the domain of microfacet normals which is different from the space of incoming and outgoing light directions. In particular, for the case of specularly reflecting microfacets, the relationship between elementary solid angles [23] can be shown to be

$$\delta\omega_2 = 4(\mathbf{k}_1 \mathbf{h}) \delta\omega_h. \quad (5)$$

Even if a microfacet has the required orientation, it might still not contribute to the single-bounce highlight if it is shadowed by other microfacets for either incoming or outgoing direction. Introducing the probability for a microfacet *not* to be shadowed in either incoming or outgoing directions as $0 \leq P(\mathbf{k}_1, \mathbf{k}_2, \mathbf{h}) \leq 1$ we will have $N_{active} = N p(\mathbf{h}) P(\mathbf{k}_1, \mathbf{k}_2, \mathbf{h}) \delta\omega_h$ and BRDF in the form

$$\rho(\mathbf{k}_1, \mathbf{k}_2) = \frac{N A_{mf} p(\mathbf{h}) P(\mathbf{k}_1, \mathbf{k}_2, \mathbf{h}) F((\mathbf{k}_1 \mathbf{h}))}{4 A (\mathbf{k}_1 \mathbf{n}) (\mathbf{k}_2 \mathbf{n})}. \quad (6)$$

Equation 6 is a somewhat modified version of the original result of Torrance and Sparrow who present its more detailed derivation [23].

The area A of the surface element can be written as a sum of the projected areas of all microfacets:

$$A = \sum_{facets} A_{mf} \langle (\mathbf{h} \mathbf{n}) \rangle P(\mathbf{n}, \mathbf{h}), \quad (7)$$

where we introduce probability $P(\mathbf{n}, \mathbf{h})$ for a microfacet not to be “shadowed” in the surface normal direction \mathbf{n} by other microfacets. If the surface is a height field, $P(\mathbf{n}, \mathbf{h}) = 1$ but in the general case some microfacets may not contribute to the area A of the projection. This question is related to the general shadowing term $P(\mathbf{k}_1, \mathbf{k}_2, \mathbf{h})$ and we postpone its discussion until the next section. The “ P ” is used with a variable number of arguments that depend on what assumptions are in play for that equation.

Given a large number of microfacets, Equation 7 can be rewritten using the average over the ensemble of microfacets as

$$A = N A_{mf} \langle (\mathbf{h} \mathbf{n}) \rangle P(\mathbf{n}, \mathbf{h})_{ens}, \quad (8)$$

where $\langle \dots \rangle_{ens}$ denotes the averaging procedure. One of the most fundamental results in statistics states that as the size of the ensemble increases, for a certain function f of a random variable its average over ensemble $\langle f \rangle_{ens}$ converges with probability one to its average $\langle f \rangle$ over the distribution of the random variable. In our case we can write for any quantity $f(\mathbf{h})$:

$$\langle f(\mathbf{h}) \rangle_{ens} = \langle f(\mathbf{h}) \rangle = \int_{\Omega} f(\mathbf{h}) p(\mathbf{h}) d\omega_h, \quad (9)$$

where the integration is done over the unit sphere Ω of microfacet normal directions (Gaussian sphere). So, for the BRDF we finally have

$$\rho(\mathbf{k}_1, \mathbf{k}_2) = \frac{p(\mathbf{h}) P(\mathbf{k}_1, \mathbf{k}_2, \mathbf{h}) F((\mathbf{k}_1 \mathbf{h}))}{4 (\mathbf{k}_1 \mathbf{n}) (\mathbf{k}_2 \mathbf{n}) \langle (\mathbf{h} \mathbf{n}) \rangle P(\mathbf{n}, \mathbf{h})}, \quad (10)$$

and in the important special case of surface being a height field,

$$\rho(\mathbf{k}_1, \mathbf{k}_2) = \frac{p(\mathbf{h}) P(\mathbf{k}_1, \mathbf{k}_2, \mathbf{h}) F((\mathbf{k}_1 \mathbf{h}))}{4 (\mathbf{k}_1 \mathbf{n}) (\mathbf{k}_2 \mathbf{n}) \langle (\mathbf{h} \mathbf{n}) \rangle}. \quad (11)$$

Although we have assumed that all microfacets have equal area A_{mf} the result does not change if there is an arbitrary distribution of microfacet areas so long as this distribution is not correlated with $p(\mathbf{h})$, the distribution of normals.

Given a density $p(\mathbf{h})$, all terms in Equation 11 are straightforward to compute except for the shadowing term $P(\mathbf{k}_1, \mathbf{k}_2, \mathbf{h})$. We now turn to the discussion of this shadowing term which is necessary to complete our formulation of the specular part of BRDF.

4 Shadowing Term

Most of the complexity of microfacet-based models arise from the shadowing function $P(\mathbf{k}_1, \mathbf{k}_2, \mathbf{h})$. In this section we describe how previous models deal with this term and introduce a new simplified shadowing term.

4.1 Previous Shadowing Terms

On any rough surface it is likely that some microfacets will either not receive light, or light reflected by them will be blocked by other microfacets. The first situation is referred to by many authors as *shadowing* and the second as *masking*. However, these events are symmetrical and for simplicity we will refer to both of them as *shadowing*. A rigorous derivation of the probability that a point on the surface is both visible and illuminated (also known as the bistatic shadowing function) leads to very complicated expressions and a set of approximations is made to make the problem tractable. Several forms of the shadowing term have been derived in different fields [1, 18, 21, 23, 25] and some of them (usually after further simplification) were later introduced to computer graphics reflection models [4, 8, 22].

The most popular shadowing functions currently used are modifications of those of Smith [21], Sancer [18] and the original Torrance and Sparrow shadowing term [23]. The first two formulations are rather complex and are designed only for Gaussian height fields. Smith, in addition, assumes an isotropic surface. The shadowing function by Torrance and Sparrow is simple, but assumes an inconsistent model of an isotropic surface exclusively made by very long V-cavities. None of the existing functions is flexible enough to accommodate a sufficiently general distribution of microfacets. Also, most of the formulations operate with height distributions, not the more intuitive normal distribution $p(\mathbf{h})$. In addition to space limitations, this is the reason we do not present the expressions of previously derived shadowing functions here.

The reason most authors deal with height distribution functions is that shadowing is clearly a non-local event intimately related to the height distribution of the surface and this information is necessary for rigorous treatment of shadowing. In the next subsection we will, however, make several assumptions which allows us to derive a very general form of the shadowing term $P(\mathbf{k}_1, \mathbf{k}_2, \mathbf{h})$ sufficient for our purposes.

4.2 New Shadowing Term

As indicated by the preceding discussion, we cannot treat shadowing rigorously if we assume a general form for the microfacet normal density function. Therefore, our generator is most appropriate in cases where the effects of shadowing are secondary compared with the influence of normal distribution shape. Even in these cases, however, we cannot ignore the shadowing term $P(\mathbf{k}_1, \mathbf{k}_2, \mathbf{h})$. As can be seen from Equation 10, at the very least shadowing should take care of the divergence at grazing angles where the denominator terms disappear: $(\mathbf{k}_1 \mathbf{n})(\mathbf{k}_2 \mathbf{n}) \rightarrow 0$.

The shadowing term can be written as

$$P(\mathbf{k}_1, \mathbf{k}_2, \mathbf{h}) = P(\mathbf{k}_1, \mathbf{h})P(\mathbf{k}_2, \mathbf{h} | \mathbf{k}_1), \quad (12)$$

where $P(\mathbf{k}_1, \mathbf{h})$ is the probability of not being shadowed in the direction \mathbf{k}_1 and $P(\mathbf{k}_2, \mathbf{h} | \mathbf{k}_1)$ is conditional probability of not being shadowed in the direction \mathbf{k}_2 given that the facet is not shadowed in direction \mathbf{k}_1 . In general, $P(\mathbf{k}_2, \mathbf{h} | \mathbf{k}_1) \neq P(\mathbf{k}_2, \mathbf{h})$. For example, it is easy to see that in the extreme case where $\mathbf{k}_1 = \mathbf{k}_2$ we have $P(\mathbf{k}_2, \mathbf{h} | \mathbf{k}_1) = 1$. This shows that visibilities in the incoming and outgoing directions are correlated. Most of shadowing functions, however, are derived under the assumption of uncorrelated

visibilities. Van Ginneken et al. [24] considered how this correlation affects Smith's shadowing function, and found that its effect can be accounted for by modifying the uncorrelated expression.

In most of this paper we will use the uncorrelated form of the shadowing term written as a product of the two independent factors for each of the two directions:

$$P(\mathbf{k}_1, \mathbf{k}_2, \mathbf{h}) = P(\mathbf{k}_1, \mathbf{h})P(\mathbf{k}_2, \mathbf{h}). \quad (13)$$

This leads to some underestimation of the BRDF if directions \mathbf{k}_1 and \mathbf{k}_2 are close to each other. If the viewing conditions are such that this arrangement is of particular importance (in a night driving simulator, for example) or if retroreflection is one of the pronounced features of surface appearance (see Section 6.4) we propose using a different form of the shadowing term:

$$P(\mathbf{k}_1, \mathbf{k}_2, \mathbf{h}) = (1 - t(\phi))P(\mathbf{k}_1, \mathbf{h})P(\mathbf{k}_2, \mathbf{h}) + t(\phi)\min(P(\mathbf{k}_1, \mathbf{h}), P(\mathbf{k}_2, \mathbf{h})), \quad (14)$$

where $-\pi < \phi < \pi$ is the angle between the projections of vectors \mathbf{k}_1 and \mathbf{k}_2 onto the tangent plane and $t(\phi)$ is a correlation factor with values between 0 and 1. The case $t(\phi) = 0$ corresponds to the completely uncorrelated case. This form of correlated shadow term was chosen because it is simple and the resulting BRDF will still conserve energy with arbitrary $t(\phi)$, as will be shown in Section 5.3. We have not done extensive experimentation with the particular form of $t(\phi)$ but we do not believe it makes a large difference as long as $t(0) = 1$ and $t(\phi)$ monotonically decreases to almost zero as $|\phi|$ increases. The range of correlation effects was found in [24] to be on the order of 15-25 degrees, so we use a Gaussian in ϕ with the width of 15 degrees.

All we need now is an expression for $P(\mathbf{k}, \mathbf{h})$, the probability for a microfacet to be visible in a given direction \mathbf{k} . Note that $P(\mathbf{n}, \mathbf{h})$ in Equations 7, 8 and 10 of the previous section is just a special case of this probability with $\mathbf{k} = \mathbf{n}$. The key assumption we make is that probability for a microfacet to be visible in direction \mathbf{k} does not depend on the microfacet's orientation \mathbf{h} as long as it is not turned away from \mathbf{k} (not self-shadowed), namely

$$P(\mathbf{k}, \mathbf{h}) = \begin{cases} P(\mathbf{k}) & \text{if } (\mathbf{k}\mathbf{h}) > 0 \\ 0 & \text{if } (\mathbf{k}\mathbf{h}) \leq 0 \end{cases} \quad (15)$$

This assumption is equivalent to the absence of correlation between the microfacet orientation and its position. This "distant shadower" assumption has been invoked before to simplify complicated shadowing expressions obtained in other fields [1, 21, 25] but we will use it in a different way - as a basis for deriving a simple and general shadowing function. Intuitively, it corresponds to rather rough surfaces and does not hold if the microfacets with certain orientation are more likely to be found at a certain height. For example, a surface made of cylinders as shown in Figure 3a will not obey this assumption while a very similar surface in Figure 3b might. In general, the more correlated the surface microfacets are, the less likely $P(\mathbf{k}, \mathbf{h})$ is to obey Equation 15.

The two surfaces in Figure 3 may still have the same distribution $p(\mathbf{h})$ and there is no way for us to distinguish between the two cases. Similarly, we will not be able to distinguish, for example, between "positive" and "negative" cylinders of Poulin and Fournier [16] but from their images it is clear that the differences in appearance due to microfacet visibility issues and not to the distribution of microfacets are minor in this case. If finer details of microfacet arrangement not captured by $p(\mathbf{h})$ are expected to substantially affect the appearance, some different framework should be used (see also Section 6.4).

The total projected area of a surface element onto direction \mathbf{k} is $A(\mathbf{kn})$. It can also be written in a way similar to Equation 7:

$$A(\mathbf{kn}) = \sum_{\text{facets}} A_{mf}(\mathbf{hk})_+ P(\mathbf{k}). \quad (16)$$




Figure 3: *Examples of surface microgeometry. Top: microfacets with almost vertical orientation are more likely to be found near the “bottom” of the surface and, therefore, are more likely to be shadowed. Bottom: orientation and height are largely uncorrelated.*

Here the subscript ‘+’ refers to the fact that the summation is performed only over microfacets turned towards \mathbf{k} , namely the ones with $\langle \mathbf{h}\mathbf{k} \rangle > 0$. Introducing averaging over microfacets and, as before, replacing it by averaging over distribution, we get

$$A(\mathbf{k}\mathbf{n}) = NA_{mf}P(\mathbf{k})\langle \langle \mathbf{h}\mathbf{k} \rangle_+ \rangle. \quad (17)$$

We are able to take $P(\mathbf{k})$ out of the averaging integral because of our assumption that it does not depend on \mathbf{h} . Because of the great importance of quantity $\langle \langle \mathbf{h}\mathbf{k} \rangle_+ \rangle$ we introduce a new notation

$$g(\mathbf{k}) = \langle \langle \mathbf{h}\mathbf{k} \rangle_+ \rangle = \int_{\Omega_+(\mathbf{k})} \langle \mathbf{h}\mathbf{k} \rangle_+ p(\mathbf{h}) d\omega_{\mathbf{h}}, \quad (18)$$

where the integration is done in \mathbf{h} -space over the hemisphere $\Omega_+(\mathbf{k})$ of directions $\langle \mathbf{h}\mathbf{k} \rangle > 0$ (Figure 4). Note that if the surface is a height field, $P(\mathbf{n}) = 1$ and Equations 8 and 17 immediately give a useful expression for $P(\mathbf{k})$:

$$P(\mathbf{k}) = \frac{\langle \mathbf{k}\mathbf{n} \rangle g(\mathbf{n})}{g(\mathbf{k})}. \quad (19)$$

In this special case $p(\mathbf{h}) = 0$ in the lower hemisphere and the averaging in $g(\mathbf{n})$ is effectively done over the complete distribution.

To handle a more general case, we note that each microfacet turned away from the direction \mathbf{k} will have a shadow with area $A_{mf}(\mathbf{h}\mathbf{k})$. This area must be subtracted from the contribution of microfacets turned towards \mathbf{k} . Again replacing sums by averages over ensemble and then over distribution, we write the projected area on the right-hand side of Equation 17 as

$$NA_{mf}P(\mathbf{k})\langle \langle \mathbf{h}\mathbf{k} \rangle_+ \rangle = NA_{mf}\langle \langle \mathbf{h}\mathbf{k} \rangle_+ \rangle + NA_{mf}\langle \langle \mathbf{h}\mathbf{k} \rangle_- \rangle, \quad (20)$$

or

$$P(\mathbf{k}) = 1 + \frac{\langle \langle \mathbf{h}\mathbf{k} \rangle_- \rangle}{g(\mathbf{k})}. \quad (21)$$

The second term is negative and the integration in it is done over the part $\Omega_-(\mathbf{k})$ of distribution complimentary to $\Omega_+(\mathbf{k})$ (Figure 4). It is clear from this equation that $P(\mathbf{k}) \leq 1$ as it should be. For a distribution of microfacet normals $p(\mathbf{h})$ to represent a valid surface, at the very least the average normal vector over the entire distribution must lie in the direction of the geometric normal \mathbf{n} of the surface:

$$\int_{\Omega_+(\mathbf{k})} \mathbf{h}p(\mathbf{h})d\omega_{\mathbf{h}} + \int_{\Omega_-(\mathbf{k})} \mathbf{h}p(\mathbf{h})d\omega_{\mathbf{h}} = \int_{\Omega} \mathbf{h}p(\mathbf{h})d\omega_{\mathbf{h}} = \mathbf{n}\langle \langle \mathbf{h}\mathbf{n} \rangle \rangle \quad (22)$$




Figure 4: *Integration domain for $g(\mathbf{k})$*

Multiplying both sides of this equation by scalar \mathbf{k} we have

$$\langle \langle \mathbf{h}\mathbf{k} \rangle_+ \rangle + \langle \langle \mathbf{h}\mathbf{k} \rangle_- \rangle = \langle \mathbf{k}\mathbf{n} \rangle \langle \langle \mathbf{h}\mathbf{n} \rangle \rangle, \quad (23)$$

or

$$\langle \langle \mathbf{h}\mathbf{k} \rangle_- \rangle = \langle \mathbf{k}\mathbf{n} \rangle \langle \langle \mathbf{h}\mathbf{n} \rangle \rangle - g(\mathbf{k}). \quad (24)$$

Substituting this into Equation 21 we obtain an expression for $P(\mathbf{k})$:

$$P(\mathbf{k}) = \frac{\langle \mathbf{k}\mathbf{n} \rangle \langle \langle \mathbf{h}\mathbf{n} \rangle \rangle}{g(\mathbf{k})}. \quad (25)$$

Averaging in the numerator is done over the complete sphere Ω of directions. Note that Equation 19 is now just a special case of Equation 25 and that Equations 21 and 25 show that for any physically valid distribution $p(\mathbf{h})$ our probability of being visible will indeed lie between 0 and 1.

The combination of Equations 10, 13 (or 14) and 25 completely describes the specular part of BRDF. Using the uncorrelated form of shadowing term of Equation 13, we get

$$\rho(\mathbf{k}_1, \mathbf{k}_2) = \frac{p(\mathbf{h})\langle \langle \mathbf{h}\mathbf{n} \rangle \rangle F(\langle \mathbf{k}\mathbf{h} \rangle)}{4g(\mathbf{k}_1)g(\mathbf{k}_2)}. \quad (26)$$

Note the interesting fact that p.d.f. $p(\mathbf{h})$ does not even have to be normalized to be used in this equation. The above formula is well-suited to evaluation. Given $p(\mathbf{h})$, it is straightforward to evaluate the BRDF. Equation 26 is the main contribution of this paper. For the rest of the paper we will discuss implications and applications of this formula.

5 Extensions and Discussion

In this section we discuss several issues related to the specular-only single bounce BRDF model derived in the last section. In particular, we discuss an energy-conserving diffuse term, implementation issues, extension to non-Fresnel microfacets, and prove energy conservation.

5.1 Diffuse Term

Equation 26 describes the part of scattering process due to single-bounce reflections from microfacets. In addition to this specular part there will be other scattering events, such as multiple bounces and subsurface scattering. A complete description of these processes is rarely attempted in a general-purpose BRDF model and

their combined contribution is usually represented by adding a diffuse component to the specular BRDF. The most common form of the diffuse term is Lambertian:

$$\rho(\mathbf{k}_1, \mathbf{k}_2) = \frac{k_d \rho_d}{\pi} + k_s \rho_s(\mathbf{k}_1, \mathbf{k}_2), \quad (27)$$

where $0 \leq \rho_d \leq 1$ is diffuse albedo of the surface while k_d and k_s are user-specified constants controlling the relative importance of specular and diffuse reflections. This is a perfectly valid option in our case as well. We can simply use Equation 26 for ρ_s and ensure that $k_d + k_s \leq 1$ to preserve the energy conservation achieved for the specular part (Section 5.3).

However, this simple form of diffuse term has problems. First of all, it is not obvious how to choose weights k_d and k_s . Second, it is clear that as more light is being reflected specularly, less of it is available for diffuse scattering, so the relative weights k_d and k_s of diffuse and specular reflections should not be constants. If Fresnel effects can cause k_s to approach one for grazing angles, k_d must be set to zero for all angles (since it is a constant). To take this effect into account in a way preserving reciprocity, we use a method of Shirley et al. [20] and write for k_d

$$k_d(\mathbf{k}_1, \mathbf{k}_2) = c(1 - R(\mathbf{k}_1))(1 - R(\mathbf{k}_2)), \quad (28)$$

where

$$R(\mathbf{k}) = \int \rho_s(\mathbf{k}, \mathbf{k}')(\mathbf{k}' \cdot \mathbf{n}) d\omega_{\mathbf{k}'} \quad (29)$$

is the directional hemispherical reflectance of the specular term, where \mathbf{k}' is the mirrored direction of \mathbf{k} . We also completely dispose of k_s by allowing the specular reflection to “have its way” and adjust the diffuse term so that it consistently follows the specular reflection. The normalization constant c is computed such that for $\rho_d = 1$ the total incident and reflected energies are the same. A complete BRDF will have the form

$$\rho(\mathbf{k}_1, \mathbf{k}_2) = c(1 - R(\mathbf{k}_1))(1 - R(\mathbf{k}_2))\rho_d + \rho_s(\mathbf{k}_1, \mathbf{k}_2). \quad (30)$$

This form of diffuse term implicitly assumes that there is no absorption on the surface and all the energy which is not reflected specularly is available for diffuse scattering. The situation is different in case of metals. First, if f_0 is the normal reflectance of the metal, only approximately f_0 fraction of incoming light is not absorbed by a flat metal surface. Second, diffuse scattering here is exclusively due to multiple bounces and thus the diffusely scattered light has a more saturated color of the metal than the primary reflection does. We attempt to take both of these effects into account by replacing ones in Equation 30 by f_0 and assigning ρ_d for a metal (which otherwise does not have any physical sense) to be f_0 . Because the true fraction of non-absorbed light is greater than f_0 , factor $(f_0 - R(\mathbf{k}))$ can become negative for some surfaces due to our approximation. We simply set the diffuse term to zero in such cases.

5.2 Implementation Issues

Implementation of our model in a rendering system is straightforward. For the Fresnel coefficient we use Schlick’s approximate formula [19]

$$F((\mathbf{k}\mathbf{n})) = f_0 + (1 - f_0)(1 - (\mathbf{k}\mathbf{n}))^5 \quad (31)$$

where again f_0 is the Fresnel factor at normal incidence. Note that we could also use the full Fresnel equations, but we use Schlick’s formula only for convenience. This should not lead to significant accuracy problems as for the error introduced by Schlick’s formula is smaller than one percent compared with the full Fresnel expression [19]. To generate a BRDF for a new distribution $p(\mathbf{h})$ all we

need, in addition to the implementation of $p(\mathbf{h})$ itself, are values for $g(\mathbf{k})$ and $R(\mathbf{k})$. Unfortunately, because of the non-standard integration domain of $g(\mathbf{k})$, analytical expressions for this function can be obtained only for the most trivial $p(\mathbf{h})$ ’s and we need to resort to numerical integration.

However, the integrals are well-behaved and the results are smooth functions for non-singular $p(\mathbf{h})$. This allows us to compute values of both $g(\mathbf{k})$ and $R(\mathbf{k})$ on a very coarse grid using available numerical packages, store the results in a table and use bilinear interpolation during the rendering process. We have used a total of 200 grid points (for many distributions an even coarser grid should be sufficient). Integration was done using both Matlab and a simple home-built Monte Carlo routine. Two sets of computed $R(\mathbf{k})$ (one with $f_0 = 1$ and one with $f_0 = 0$) are sufficient to compute $R(\mathbf{k})$ for a material with arbitrary f_0 for a given microfacet distribution.

In the BRDF generation phase we start from $p(\mathbf{h})$ and output a compact numerical representation of three two-dimensional functions: $g(\mathbf{k})$, $R(\mathbf{k})$ with $f_0 = 0$ and $R(\mathbf{k})$ with $f_0 = 1$. The last two functions are only used for the diffuse term and are not required for its simpler form in Equation 27. During rendering we use these data to compute the full four dimensional BRDF for arbitrary \mathbf{k}_1 and \mathbf{k}_2 . At this stage we also use data for normal reflectance f_0 and diffuse albedo ρ_d . Wavelength dependence of these quantities controls the color of the surface. We have not done a careful performance analysis but from our experience for a non-trivial $p(\mathbf{h})$ most of the BRDF computation time is due to evaluating this normal distribution function.

Note that most distributions have some symmetry which can be exploited to further reduce the amount of data and/or generation time. Data for an anisotropic Gaussian distribution of normals, for example, need be computed only over a quarter of the hemisphere and for any isotropic distribution functions $g(\mathbf{k})$ and $R(\mathbf{k})$ become one dimensional.

Finally, if a particular type of parameterized distribution (Gaussian, for example) is used often it should be possible to approximate $g(\mathbf{k})$ with a simple function of \mathbf{k} and distribution parameters as is commonly done to increase the efficiency of reflection models. The same is true for $R(\mathbf{k})$ but these functions usually have more complex shapes.

5.3 Energy Conservation

By inspection of the formulas, it is clear that generated BRDFs are reciprocal. We now prove now that they also conserve energy for any physically plausible $p(\mathbf{h})$. To do this, we assume the worst-case scenario of $F((\mathbf{k}\mathbf{h})) = 1$ and shadowing term in Equation 14 with $t(\phi) = 1$ (because $P(\mathbf{k}) \leq 1$ this corresponds to the largest possible shadowing term for our model). The BRDF in this case will be

$$\rho(\mathbf{k}_1, \mathbf{k}_2) = \frac{p(\mathbf{h}) \min(P(\mathbf{k}_1), P(\mathbf{k}_2))}{4\langle(\mathbf{h}\mathbf{n})\rangle(\mathbf{k}_1\mathbf{n})(\mathbf{k}_2\mathbf{n})} \leq \frac{p(\mathbf{h})P(\mathbf{k}_1)}{4\langle(\mathbf{h}\mathbf{n})\rangle(\mathbf{k}_1\mathbf{n})(\mathbf{k}_2\mathbf{n})}$$

Hemispherical reflectance for a given incoming direction is

$$R(\mathbf{k}_1) = \int \rho_s(\mathbf{k}_1, \mathbf{k}_2)(\mathbf{k}_2\mathbf{n}) d\omega_2 \leq \frac{P(\mathbf{k}_1)}{4\langle(\mathbf{h}\mathbf{n})\rangle(\mathbf{k}_1\mathbf{n})} \int p(\mathbf{h}) d\omega_2 = \frac{P(\mathbf{k}_1)}{4\langle(\mathbf{h}\mathbf{n})\rangle(\mathbf{k}_1\mathbf{n})} \int p(\mathbf{h}) 4(\mathbf{k}_1\mathbf{h}) d\omega_h$$

The last transition is done using Equation 5. The integration is done over a complex region of h -space which is in any case contained

in the hemisphere $\Omega_+(\mathbf{k}_1)$. Extending the integral over the whole $\Omega_+(\mathbf{k}_1)$ and using definitions 18 of $g(\mathbf{k})$ and 25 of $P(\mathbf{k})$ we complete the proof:

$$R(\mathbf{k}_1) \leq \frac{P(\mathbf{k}_1)}{\langle(\mathbf{h}\mathbf{n})\rangle(\mathbf{k}_1\mathbf{n})} \int_{\Omega_+(\mathbf{k}_1)} (\mathbf{k}_1\mathbf{h})p(\mathbf{h})d\omega_h = \frac{P(\mathbf{k}_1)g(\mathbf{k}_1)}{\langle(\mathbf{h}\mathbf{n})\rangle(\mathbf{k}_1\mathbf{n})} = 1 \quad (32)$$

The only fact we used in our proof is that $P(\mathbf{k}) \leq 1$ for any \mathbf{k} . In Section 4, in turn, this was shown to be the case for any $p(\mathbf{h})$ whose average normal vector $\langle\mathbf{h}\rangle$ is parallel to the geometric normal of the surface. This is the only restriction on microfacet distribution $p(\mathbf{h})$. If it is satisfied, the generated BRDF will conserve energy.

5.4 Non-Fresnel Microfacets

Our model is not restricted to perfectly specular microfacets. In general, microfacets with many orientations will contribute to surface BRDF for given incoming and outgoing directions and integration of their contribution is necessary.

Let all microfacets have elementary BRDF β . Then we can repeat with some modifications the derivation from Sections 3 and 4 to arrive at the result

$$\rho(\mathbf{k}_1, \mathbf{k}_2) = \frac{P(\mathbf{k}_1, \mathbf{k}_2)}{(\mathbf{k}_1\mathbf{n})(\mathbf{k}_2\mathbf{n})\langle(\mathbf{h}\mathbf{n})\rangle} \int \beta(\mathbf{k}_1, \mathbf{k}_2)(\mathbf{k}_1\mathbf{h})_+(\mathbf{k}_2\mathbf{h})_+p(\mathbf{h})d\omega_h \quad (33)$$

The integration is done over the sector where both $(\mathbf{k}_1\mathbf{h})$ and $(\mathbf{k}_2\mathbf{h})$ are positive and any of shadowing terms $P(\mathbf{k}_1, \mathbf{k}_2)$ from Section 4 can be used. Note that $\beta(\mathbf{k}_1, \mathbf{k}_2)$ is usually specified with respect to microfacet's local coordinate system and a coordinate transformation is necessary to obtain its value for the integral in Equation 33.

Although this extension considerably broadens the range of surfaces our model is applicable to, we also lose one of the main advantages of our approach: compactness. Before, we could represent a general four dimensional BRDF using only two dimensional functions. The integral in Equation 33, however, is a four dimensional function by itself and does not, in general, allow lower dimensional representation. For some special cases, such as Lambertian elementary BRDF coupled with isotropic $p(\mathbf{h})$ the integral becomes three dimensional and, therefore, feasible to compute, store and use in a way similar to that described in Section 5.2. For an isotropic Gaussian distribution of Lambertian microfacets the general behavior of the generated BRDF is similar to that of Oren-Nayar's model [14], namely, retroreflection is increased compared to a Lambertian surface (Figure 8).

6 Applications

In this section we apply our model to a variety of surface types. Although we have implemented our model in a Monte Carlo ray tracer capable of handling complex geometries and illumination effects, our images in this section intentionally show very simple objects and lighting conditions. In particular, illumination is coming from a single small light source far from the scene and indirect lighting is not included. This is done to emphasize effects due to BRDF of the material and to make the comparison with previous results easier.

Reflectance data of gold are used as f_0 (see Section 5.2) for all metal objects while for non-metals f_0 is set to 5% across the visible spectrum.




Figure 5: *Anisotropic Gaussian golden spheres with $\sigma_x = 0.1$, $\sigma_y = 1.0$. Left: Ward. Right: new model.*




Figure 6: *Directional hemispherical reflectance as a function of incoming angle for perfectly reflecting microfacets with Gaussian distribution $\sigma_x = 0.1$, $\sigma_y = 0.2$. For an ideal flat surface R should be 1.0 everywhere. Left: Ward. Right: new model.*




Figure 7: *Anisotropic Gaussian golden painted plastic spheres with $\sigma_x = 0.1$, $\sigma_y = 0.2$. Left: Ward. Right: new model.*




Figure 8: *Gaussian spheres with Lambertian microfacets. Right: new model with $\sigma_x = \sigma_y = 1.0$. Left: Oren-Nayar with compatible parameters.*

6.1 Gaussian Surfaces

By far the most popular distribution used in BRDF research literature is Gaussian. This is due to both its practical importance and nice mathematical properties. Gaussians are used in all four major categories of BRDF models outlined in the introduction. While some of this work is closer to our approach in its theoretical foundations, we feel that from the practical point of view our model is closest to that of Ward [26]. Ward’s BRDF is simple, handles anisotropic distributions and seeks to reproduce the main character of the material’s reflectance behavior without attempting an overly detailed description. Other previous models do not simultaneously possess all these properties.

To create an anisotropic Gaussian BRDF, we use the distribution

$$p(\mathbf{h}) = c * \exp(-\tan^2 \theta (\cos^2 \phi / \sigma_x^2 + \sin^2 \phi / \sigma_y^2)) \quad (34)$$

where θ is the angle between the half vector \mathbf{h} and the surface normal, ϕ is the azimuth angle of \mathbf{h} and c is a normalization constant.

Two side-by-side comparisons of our model with Ward’s are shown in Figures 5 and 7. Note that the shape of highlight is nearly identical while there are some differences in the diffuse part of images which is due to Ward effectively using a simpler form (Equation 27) of the diffuse component. In particular, for our metal sphere on Figure 5 the diffuse component appears automatically when there is enough energy left after single-bounce scattering. To achieve the same effect in Ward’s model (and any other using the popular Lambertian diffuse term) it would be necessary to manually adjust the diffuse reflectance parameter.

This figure also shows that the highlight is brighter for our BRDF. The general reason for this is clear from Figure 6 where the hemispherical reflectance R is plotted versus the incoming light direction. To make the plots directly comparable, we show data for most reflecting specular BRDF in both cases ($f_0 = 1$ for our model and $\rho_s = 1$ in Equation 5 of Ward’s paper [26]) and do not include the diffuse term. For the values of parameters shown, the surface is quite close to being flat, so one would expect that R should be close to that of flat surface, 1.0 in this case. One can see from the plots that our model behaves as expected while Ward’s does not. Note also that the true value for R at the grazing angle ($(\mathbf{kn}) = 0$) is infinite for Ward’s model [13] and we simply extrapolate previous behavior to get the data point at the grazing angle.

While our approach does require an extra generation step, computation time during the rendering process of our BRDF is close to that of Ward’s and our model is a viable alternative where energy conservation is of great importance for a particular application.

Figure 8 compares a BRDF generated for an isotropic Gaussian distribution of Lambertian microfacets with an extension of our process (Section 5.4) and Oren-Nayar model with compatible parameters. Both BRDFs have the tendency to make objects appear “flatter” than the Lambertian BRDF due to increased retroreflection.

6.2 Grooved Surface

A surface consisting of ideal V-grooves all running in a given direction will have its $p(\mathbf{h})$ proportional to the sum of two delta functions, each accounting for microfacets forming one side of a groove. Replacing these delta functions with narrow Gaussians ($\sigma = 0.1$) to account for imperfections and going through our generation process, we create a BRDF which correctly shows the main feature of a grooved surface’s reflectance, double reflections. Figure 9 shows a piece of grooved metal illuminated by a *single* light source. The orientation of the grooves on the left is perpendicular to the viewing direction while on the right they are parallel.




Figure 9: *Double highlights from a single light source for the same metallic grooved surface at two orientations of the grooves. Grooves are symmetrical with the angle of 40 degrees*

6.3 Satin

The microstructure of woven cloth is usually thought of as a symmetric pattern of interwoven cylindrical fibers running in perpendicular directions. While it would be possible to generate a BRDF corresponding to this structure with our approach, the surface of particular fabric we studied had a different microstructure shown in Figure 10. It is created almost exclusively by fibers running in one direction with about 70% of the fiber length lying in the relatively flat part of the fiber while the other 30% corresponding to the bent parts at the ends. We model the distribution of microfacets as a linear combination of two terms corresponding to these flat and bent parts of the cylindrical fiber: $p(\mathbf{h}) = 0.7 * p_{flats}(\mathbf{h}) + 0.3 * p_{ends}(\mathbf{h})$. The coefficients reflect mutual area contributions of the two parts to the complete distribution. Both $p_{flats}(\mathbf{h})$ and $p_{ends}(\mathbf{h})$ were chosen to be “cylindrical” Gaussian heightfields ($\sigma_y = \infty$, $p(\mathbf{h}) = 0$ for $(\mathbf{hn}) < 0$) with different widths. Values $\sigma_x = 0.1$ for $p_{flats}(\mathbf{h})$ and $\sigma_x = 0.3$ for $p_{ends}(\mathbf{h})$ were used. Strictly speaking, the shape of real $p_{ends}(\mathbf{h})$ would probably be more accurately modeled by a distribution with flatter top and faster drop-offs than that of a Gaussian. This was attempted but the results were almost identical visually, so a simpler Gaussian distribution was used for the final image. This is consistent with our belief that the very precise characterization of the microfacet distribution is not needed for visual applications. Note that because $g(\mathbf{k})$ is linear in $p(\mathbf{h})$, no new integration is necessary to compute $g(\mathbf{k})$ if g ’s corresponding to $p_{flats}(\mathbf{h})$ and $p_{ends}(\mathbf{h})$ are already computed. This suggests an efficient way of creating new distributions as a linear combination of ones for which $g(\mathbf{k})$ has been previously computed. For example, small contribution due to perpendicular fibers can be added in this manner if necessary.

Because the appearance of real cloth is dramatically affected by the presence of characteristic wrinkles, we used a dynamic simulation method [17] to create cloth geometry. The left side of Figure 11 shows a satin tablecloth rendered with generated BRDF. It is interesting to contrast this image with the image on the right using the same geometric model with the BRDF described in the next section.

6.4 Velvet

Velvet is another example of a material with interesting reflectance properties not easily conveyed by conventional BRDFs. In their virtual gonioreflectometer, Westin et al. [27], model velvet microstructure as a forest of narrow cylinders (fibers) with the orientation of each cylinder perturbed randomly. While it is difficult to write an exact $p(\mathbf{h})$ corresponding to such “surface” for the reasons outlined in Section 2, a simple intuitive form of this function written as an “inverse Gaussian” heightfield is enough to capture the main char-




Figure 10: *Microgeometry of our sample of satin.*



Figure 11: *Synthetic satin (left) and velvet (right) tablecloths. The geometries are identical.*




Figure 12: *Microgeometry of velvet (left) and $p(\mathbf{h})$ used to model it (right).*



Figure 13: *A tablecloth made of two different colors of slanted fiber velvets.*

acter of the distribution:

$$p(\mathbf{h}) = c * \exp(-\cot^2 \theta / \sigma^2), \quad (35)$$

with $\sigma = 0.5$ for the image on the right of Figure 11 which shows a material with distinct velvet-like reflectance properties. Because retroreflection is one of the most pronounced reflection properties of velvet [11], we used the correlated form of shadowing term (Equation 14) to generate both this and slanted fiber (see below) velvet BRDFs. Contrary to Westin et al. we ignore the tips of the fibers due to their very small area. If there were any specular highlights due to the tips, their contribution can be easily added by forming a linear combination of an inverse Gaussian with a regular Gaussian distribution.

Although this approach produced good results, a symmetric forest of fibers was not what we saw when we examined a piece of real velvet. More realistic structure is shown on the left of Figure 12. The fabric consists of rows of tightly woven bundles of filament. Each bundle is slanted with the angle of about 40 degrees with respect to the geometric normal of the cloth surface. We can call this arrangement millscale geometry in contrast with microgeometry formed by the thin fibers themselves. Similar geometry was credited as the major reason for velvet anisotropic reflection behavior by Lu et al. [11]. Strictly speaking, our model does not take into account visibility issues due to this higher-order arrangement of microfacets. The most consistent approach therefore would be to model this structure explicitly, for example as a collection of slanted cylinders applying two different BRDFs (both of which can be generated by our process) to the tops and to the sides of these cylinders. An easier alternative would be an attempt to create a simple distribution of microfacets $p(\mathbf{h})$ which, although potentially non-physical, can account for the millscale visibility and produce a BRDF with necessary reflection properties.

Looking carefully at the velvet highlight structure we saw that it is the sides of the bundles and not the tops which contribute the most to the reflection. This suggests that we can try to reproduce most of the behavior with a specular BRDF based exclusively on the $p(\mathbf{h})$ accounting for the microfacets on the sides of the bundles. A “slanted” version of cylindrical Gaussian distribution ($\sigma_y = \infty$, $\sigma_x = 0.5$) schematically shown on the right of Figure 12 was used. The only place where we used the part of distribution due to the tops of the bundles is the computation of $\langle (\mathbf{h}\mathbf{n}) \rangle$ when we double this value due to the tops contribution. Note two facts about this distribution: it is not a height field and its average vector $\langle \mathbf{h} \rangle$ does not point in the direction of geometric normal. While the first feature does not present any problem in our approach, the second one shows that this distribution is not physically realizable and, as a result, the energy conservation of the generated BRDF is not guaranteed. Computations of $R(\mathbf{k})$ show that this quantity indeed exceeds one for 14 out of our 200 directional data points in the hypothetical case of perfectly reflecting ($f_0 = 1.0$) fiber material but was never a problem for our $f_0 = 0.05$ synthetic fibers.

Figure 13 shows the results of this process. The illumination and viewing directions are almost parallel but due to the slant of the fibers the left side of the tablecloth is substantially brighter than the right one. This is in good agreement with the behavior of real velvet we observed. The right image of Figure 13 shows some limitations of our approach. Because we do not handle the details of multiple-bounce scattering and simply introduce a diffuse term to account for them, the right side of the red tablecloth does not look as it does for the real velvet. In the real material, light experiences multiple bounces among the red fibers for this viewing geometry acquiring a deep dark (almost black) color in the process. This is not captured by our simple diffuse term.

6.5 Unusual Distributions

We can take to extreme the use of the desired reflection properties as the only guidance in creating the distribution $p(\mathbf{h})$ regardless of whether a material described by this function exists or is even physically possible. For example, we can modulate a Gaussian $p(\mathbf{h})$ with an arbitrary function or even an image to create the unusual highlights shown in Figure 1. As long as the modulation is symmetric enough to keep the average vector $\langle \mathbf{h} \rangle$ in the normal direction (such as the distribution used for the image on the left of Figure 1), the BRDF will be energy conserving. A more general modulation may result in $\langle \mathbf{h} \rangle$ no longer parallel to \mathbf{n} but in practice we notice that as long as this effect is not very strong, the energy conservation is not affected. For example, image on the right of Figure 1 was created with an energy conserving BRDF. While such unusual distributions are not of great value in realistic image synthesis, they clearly demonstrate the generality of our approach and can potentially find applications in the special effects industry.

7 Conclusion

The new BRDF model presented in this paper is well-suited to surfaces whose primary characteristic is the shape of the specular highlight. We have found it reasonably straightforward to design new BRDFs for surfaces because the diffuse term and energy conservation are handled in a natural manner that does not require substantial user intervention, and the parameters used in the model are intuitive. However, for surfaces whose appearance is not dominated by the specular highlight, our model is not well-suited.

We have found that using our model does not require much hand-tuning of parameters; the images in the last section were generated with very few iterations on parameter values. We speculate that a model for subsurface effects in a similar spirit to our model is possible. The user would specify some simple parameters analogous to $p(\mathbf{h})$ and a BRDF would be generated. We also believe that there should ultimately be separate terms for the components of the BRDF accounted for by primary specular reflection, multiple-bound specular reflection, and subsurface scattering.

Acknowledgements

This work was supported by NSF grants 96-23614, 97-96136, 97-31859, and 98-18344. Thanks to Robert McDermott for his help with production issues. The tablecloth models were done using Maya software generously donated by Alias/Wavefront.

References

- [1] BECKMANN, P. Shadowing of random rough surfaces. *IEEE Transactions on Antennas and Propagation* 13 (1965), 384–388.
- [2] BLINN, J. F. Models of light reflection for computer synthesized pictures. *Computer Graphics (Proceedings of SIGGRAPH 77)* 11, 2 (July 1977), 192–198.
- [3] CABRAL, B., MAX, N., AND SPRINGMEYER, R. Bidirectional reflectance functions from surface bump maps. *Computer Graphics* 21, 4 (July 1987), 273–282. ACM Siggraph '87 Conference Proceedings.
- [4] COOK, R. L., AND TORRANCE, K. E. A reflectance model for computer graphics. *Computer Graphics* 15, 3 (August 1981), 307–316. ACM Siggraph '81 Conference Proceedings.
- [5] GONDEK, J. S., MEYER, G. W., AND NEWMAN, J. G. Wavelength dependent reflectance functions. In *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)* (July 1994), A. Glassner, Ed., Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH, ACM Press, pp. 213–220.
- [6] GOURAUD, H. Continuous shading of curved surfaces. *Communications of the ACM* 18, 6 (June 1971), 623–629.
- [7] GREENBERG, D. P., TORRANCE, K. E., SHIRLEY, P., ARVO, J., FERWERDA, J. A., PATTANAIAK, S., LAFORTUNE, E. P. F., WALTER, B., FOO, S.-C., AND TRUMBORE, B. A framework for realistic image synthesis. *Proceedings of SIGGRAPH 97* (August 1997), 477–494.
- [8] HE, X. D., TORRENCE, K. E., SILLION, F. X., AND GREENBERG, D. P. A comprehensive physical model for light reflection. *Computer Graphics* 25, 4 (July 1991), 175–186. ACM Siggraph '91 Conference Proceedings.
- [9] LAFORTUNE, E. P., AND WILLEMS, Y. D. Using the modified phong BRDF for physically based rendering. Tech. Rep. CW197, Computer Science Department, K.U.Leuven, November 1994.
- [10] LAFORTUNE, E. P. F., FOO, S.-C., TORRANCE, K. E., AND GREENBERG, D. P. Non-linear approximation of reflectance functions. *Proceedings of SIGGRAPH 97* (August 1997), 117–126.
- [11] LU, R., KOENDERINK, J. J., AND KAPPERS, A. M. L. Optical properties (bidirectional reflection distribution functions) of velvet. *Applied Optics* 37, 25 (1998), 5974–5984.
- [12] MARSCHNER, S. R., WESTIN, S. H., LAFORTUNE, E. P. F., TORRANCE, K. E., AND GREENBERG, D. P. Image-based BRDF measurement including human skin. *Eurographics Rendering Workshop 1999* (June 1999).
- [13] NEUMANN, L., NEUMANN, A., AND SZIRMAY-KALOS, L. Compact metallic reflectance models. *Computer Graphics Forum* 18, 13 (1999).
- [14] OREN, M., AND NAYAR, S. K. Generalization of lambert's reflectance model. In *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)* (July 1994), A. Glassner, Ed., Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH, ACM Press, pp. 239–246.
- [15] PHONG, B.-T. Illumination for computer generated images. *Communications of the ACM* 18, 6 (June 1975), 311–317.
- [16] POULIN, P., AND FOURNIER, A. A model for anisotropic reflection. *Computer Graphics* 24, 3 (August 1990), 267–282. ACM Siggraph '90 Conference Proceedings.
- [17] PROVOT, X. Deformation constraints in a mass-spring model to describe rigid cloth behavior. In *Proceedings of Graphics Interface '95* (1995), pp. 147–154.
- [18] SANCER, M. I. Shadow corrected electromagnetic scattering from randomly rough surfaces. *IEEE Transactions on Antennas and Propagation AP-17*, 5 (September 1969), 577–585.
- [19] SCHLICK, C. An inexpensive BRDF model for physically-based rendering. *Computer Graphics Forum* 13, 3 (1994), 233–246.
- [20] SHIRLEY, P., HU, H., SMITS, B., AND LAFORTUNE, E. A practitioners' assessment of light reflection models. In *Pacific Graphics* (October 1997), pp. 40–49.
- [21] SMITH, B. G. Geometrical shadowing of a random rough surface. *IEEE Transactions on Antennas and Propagation* 15 (1967), 668–671.
- [22] STAM, J. Diffraction shaders. *Proceedings of SIGGRAPH 99* (August 1999), 101–110.
- [23] TORRANCE, K. E., AND SPARROW, E. M. Theory for off-specular reflection from roughened surfaces. *Journal of Optical Society of America* 57, 9 (1967).
- [24] VAN GINNEKEN, B., STAVRIDIS, M., AND KOENDERINK, J. J. Diffuse and specular reflectance from rough surfaces. *Applied Optics* 37, 1 (1998), 130–139.
- [25] WAGNER, R. J. Shadowing of randomly rough surfaces. *Journal of Acoustic Society of America* 41 (1967), 138–147.
- [26] WARD, G. J. Measuring and modeling anisotropic reflection. *Computer Graphics* 26, 4 (July 1992), 265–272. ACM Siggraph '92 Conference Proceedings.
- [27] WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. Predicting reflectance functions from complex surfaces. *Computer Graphics* 26, 2 (July 1992), 255–264. ACM Siggraph '92 Conference Proceedings.