



HAL
open science

A Physically-Based Reflectance Model Combining Reflection and Diffraction

Nicolas Holzschuch, Romain Pacanowski

► **To cite this version:**

Nicolas Holzschuch, Romain Pacanowski. A Physically-Based Reflectance Model Combining Reflection and Diffraction. [Research Report] RR-8964, INRIA. 2016. hal-01386157

HAL Id: hal-01386157

<https://inria.hal.science/hal-01386157>

Submitted on 7 Nov 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



A Physically-Based Reflectance Model Combining Reflection and Diffraction

Nicolas Holzschuch , Romain Pacanowski

**RESEARCH
REPORT**

N° 8964

October 2016

Project-Teams Maverick and
Manao

ISSN INRIA/RR--8964--FR+ENG

ISSN 0249-6399



A Physically-Based Reflectance Model Combining Reflection and Diffraction

Nicolas Holzschuch ^{*}, Romain Pacanowski [†]

Project-Teams Maverick and Manao

Research Report n° 8964 — October 2016 — 23 pages

Abstract: Reflectance properties express how objects in a virtual scene interact with light; they control the *appearance* of the object: whether it looks shiny or not, whether it has a metallic or plastic appearance. Having a good reflectance model is essential for the production of photorealistic pictures. Measured reflectance functions provide high realism at the expense of memory cost. Parametric models are compact, but finding the right parameters to approximate measured reflectance can be difficult. Most parametric models use a model of the surface micro-geometry to predict the reflectance at the macroscopic level. In this paper, we show that this micro-geometry causes *two* different physical phenomena: reflection and diffraction. Their relative importance is connected to the surface roughness. Taking both phenomena into account, we develop a new reflectance model that is compact, based on physical properties and provides a good approximation of measured reflectance.

Key-words: Appearance, material, Reflectance, BRDF, Measured materials, Diffraction

^{*} Inria ; Univ. Grenoble Alpes, LJK ; CNRS, LJK

[†] CNRS-LP2N, Univ. Bordeaux, IOGS; Inria

Un modèle de réflectance combinant diffraction et réflexion

Résumé : Les propriétés de réflexion expriment la façon dont les objets interagissent avec la lumière dans les scènes virtuelles. Elles contrôlent l'apparence de l'objet: s'il apparaît brillant ou non; son aspect métallique ou plastique ? Le modèle de réflexion (BRDF) est primordial afin d'obtenir des images photo-réalistes. Les modèles de réflexion mesurés fournissent un haut degré de réalisme au détriment du coût mémoire. Les modèles paramétrés sont compacts mais il reste difficile de trouver les bonnes valeurs des paramètres à partir des réflectances mesurées. Dans ce rapport de recherche, nous montrons que deux phénomènes physiques sont présents dans les mesures de réflexion : la réflexion à proprement parler auquel s'ajoute le phénomène de diffraction. En prenant les deux en compte nous présentons un modèle de BRDF fournit une très bonne approximation des mesures tout en restant compact en mémoire. Les artistes peuvent modifier les paramètres du modèle qui sont reliés aux propriétés de la surface afin de créer de nouveaux matériaux.

Mots-clés : Modèles de matériaux, simulation de l'éclairage, réflectance, diffraction

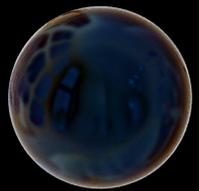
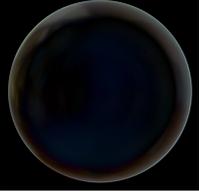
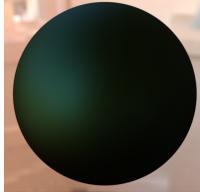
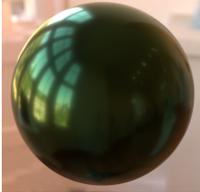
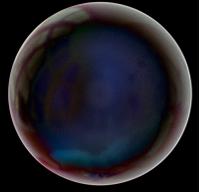
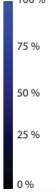
| | Our model | | | Reference | Difference (sMAPE) | Scale |
|-----------------------|---|---|---|--|---|---|
| nickel |  | +  | =  |  |  |  |
| alum-bronze |  | +  | =  |  |  |  |
| green-metallic-paint2 |  | +  | =  |  |  |  |
| | Diffraction (GHS) | + Cook-Torrance | = Together | | | |

Figure 1: Surface micro-geometry contributes to its visible aspect (material reflectance). Two physical phenomena are acting together: reflection on micro-facets and diffraction. Our reflectance model combines them, with the proper energy repartition between them. The importance of diffraction depends on the roughness of the material. Even when it is relatively small, as for green-metallic-paint2, it has a significant impact of the aspect of the material. Our model explains even a very difficult material like alum-bronze as a single material.

1 Introduction

Reflectance properties express how materials interact with light. They are responsible for the overall aspect of objects in virtual scenes: whether they look shiny, metallic, plastic. . . Measured reflectance functions [MPBM03] reproduce the appearance of the material well, but have storage issues (33 MB for a single isotropic material). They also make it more difficult to edit materials or to use importance sampling.

Parametric reflectance models, such as Phong, Lafortune or Cook-Torrance, are more compact and easier to edit. They are also easier to combine with importance sampling in Monte-Carlo integration. On the other hand, it is difficult for users to find the parameters that correspond to the visual characteristic of a specific material. To combine the advantages of both approaches, we want to compute the parameters for one BRDF model so it produces the same appearance as a given measured material. This would provide a compact representation and allow importance sampling. It would also make it easier to edit material properties, to make them smoother or rougher.

Many parametric models are based on a physical representation of the surface micro-geometry and how it interacts with incoming light. The Cook-Torrance model [CT82] is very popular in Computer Graphics. Based on the assumption that surface micro-geometry acts as a set of specular mirrors, it connects the *normal* distribution of the surface at the microscopic level with its visual appearance at the macroscopic level. Another model, widely used in the Optics community, computes diffraction effects caused by differences in *height* in the surface micro-geometry, and predicts visual appearance from the frequency content of the height distribution. In this model, the width of material response intrinsically depends on the light wavelength.

It is interesting to compare measured reflectance functions with what the models predict. Butler et al. observed that measured material response is well approximated with a micro-facet model [BNM15b] but that the width of material response appears to vary with the wavelength [BNM15c], in contradiction with the model. Löw et al. [LKYU12] observed that the behaviour of many measured reflectance corresponds to the predictions of the diffraction model for low intensities. Several researchers, e.g. [NDM05, Bur12, LKYU12], have observed that some measured materials appear to be the sum of two different behaviors.

In this paper, we provide a physical explanation to all these observations, even where they appear to contradict the physical models: when modelling the influence of surface micro-geometry on material reflectance, diffraction and reflection are not

mutually exclusive. The Harvey-Shack diffraction model [Har75] explicitly predicts the energy repartition between specular reflection and diffraction effects. Integrating this energy repartition, we propose a new reflectance model that combines micro-facet reflection and height-based diffraction.

With our model, reflection effects, governed by the micro-geometry normals, dominate near the specular direction. Diffraction effects, governed by the micro-geometry height distribution, are dominant for wide-angle scattering, away from the specular direction. The relative importance of these two phenomena is governed by the *variance* of the height distribution. All these effects come directly from the full application of Generalized Harvey-Shack theory. For smooth surfaces, the two effects are well separated and appear as two different lobes. For rougher surfaces, the two effects overlap and result in a single lobe, whose width changes with the wavelength.

For conductors, our model express reflectance properties *without* a diffuse term. We introduce two variations of our model for materials that explicitly cause a diffuse effects: sub-surface scattering and multi-layered plastic materials.

All parameters in our model are connected to physical quantities. The color of the material comes solely from its index of refraction. The shape of the reflectance lobe comes from the normal distribution function, the variance of the surface height and the frequency spectrum of the height distribution. We validate our model by using it to approximate measured material reflectance functions. It is a good predictor of material behavior for smooth and moderately rough surfaces.

In the next section, we review some of the previous work on material reflectance, both in the Computer Graphics and Optical Engineering communities. We then describe in Section 3 the scientific background: the micro-facet reflectance models and the Generalized Harvey-Shack diffraction theory. In Section 4, we describe our model, combining reflection and diffraction, and fully describe each component. In Section 5, we present our method to extract model parameters from measured reflectance functions, and discuss the results. In Section 6, we conclude and present avenues for future work.

2 Previous Work

2.1 Micro-Facet Model

The micro-facet model [TS67, CT82] assumes that surface micro-geometry is made of specular micro-facets. It predicts the overall surface appearance from the probability distribution of these micro-facet normals. Existing models include Gaussian [CT82], rational fraction [TR75, WMLT07], fraction to the power p [LKYU12, Bur12], Shifted-Gamma Distribution [BSH12], exponential of a power function [BLPW14].

The shadowing and masking term is essential for energy conservation in the micro-facet model. Assuming that micro-facets position and orientation are independent, Smith [Smi67] computes the shadowing/masking term from two successive integrations of the slope distribution. Heitz [Hei14] shows that this is the most physically consistent method to compute the shadowing and masking term, and provides an improved shadowing term taking into account correlation between input and output directions.

2.2 Diffraction Models

Diffraction models focus on diffraction effects caused by variations in surface height. They are widely used in the Optical Engineering Community, for example to measure imperfections on smooth surfaces, such as mirrors. There are several models expressing the relationship between surface height distribution and reflectance properties: the Rayleigh-Rice vector perturbation theory is designed for smooth surfaces and wide-angle scattering; the Beckmann-Kirchhoff [BS87] and Modified Beckmann-Kirchhoff [VH98] theories express the BRDF as an infinite sum; they are valid for a large range of surfaces, but for small angles. For very smooth or very rough surfaces, the infinite sum is approximated by a single term. The Harvey-Shack model [Har75] express the diffraction effects using a surface transfer function. For smooth surfaces, this transfer function is the Power Spectral Distribution of surface height. Krywonos [Kry06] extended this with the Generalized Harvey-Shack model, where the surface transfer function is recomputed for each incoming direction. This model is valid for all surfaces, rough and smooth, and for all angles. Krywonos [Kry06] compared models with measured reflectance and showed that the Generalized Harvey-Shack model provides a better approximation than other diffraction models.

He *et al.* [HTSG91] derived the most comprehensive reflectance model for isotropic surfaces, based on Beckmann-Kirchhoff theory; it takes into account all scales in the micro-geometry in a single model. Stam [Sta99] introduced an anisotropic reflectance model, also based on the Beckmann-Kirchhoff model. His model reverts to He's for isotropic surfaces. These models remain impractical for rendering and fitting since their number of terms grows quickly with the surface roughness.

In this paper, we build on the Generalized Harvey-Shack model, as it models diffraction effects using a single term instead of an infinite sum and has been shown to agree better with actual reflectance measurements.

2.3 Comparison with Measured Materials

Butler et al. studied the connection between these two models. They found that measured reflectance functions are well approximated by the micro-facet model in the incident plane [BNM15b], but that the normal distribution appears to change with the wavelength [BNM15c]. They studied the connections between the two models and proposed an approximation for the ratio of the Fresnel term for reflection and the color term for diffraction [BNM15a].

The two models have different parameters and thus predict different shapes for reflectance lobes outside of the incident plane. Löw et al. [LKYU12] showed that the shape of the reflectance lobe in many measured materials is consistent with the predictions of diffraction theory. They also showed that the diffraction model provides a good fit with measured reflectance. To conduct their fitting, they removed the wavelength dependency, in contradiction with the diffraction model. They also replaced the color term for diffraction Q with the Fresnel term for reflection F .

Holzschuch and Pacanowski [HP15b], through an empirical study, showed that measured reflectance appear to be a combination of the two phenomena: diffraction for wide-angle scattering, micro-facet reflection close to the specular direction. In this paper, we provide a physical explanation for their empirical observation. In later work [HP15a], they showed that summing the two components provides a good approximation to all measured reflectance. Compared to their work, we remove the need for a diffuse lobe, explicitly model the energy repartition between the two lobes and predict the color of each lobe using only the material index of refraction, reducing the number of parameters strongly.

Dong et al [DWMG15] compared measured material reflectance and measured micro-geometry; they found that measured normals are a good predictor of the specular peak of the reflectance lobe, if the surface features are filtered to remove sub-wavelength details. They also compared measured reflectance with Kirchhoff diffraction theory, using a single term.

Matusik et al. [MPBM03] measured and released reflectance properties for a large range of materials. We use their database in our tests. Ngan et al. [NDM05] tried fitting parametric BRDF models to this measured data. They found the best fits for the Cook-Torrance [CT82] and Lafortune [LFTG97] models. They also reported that many specular materials are poorly approximated with a single lobe. The quality of the fit improves with multiple lobes, but the fitting process becomes unstable.

Ashikhmin and Premože [AP07] approximated measured BRDFs using back-scattering: if input and output directions are equal, the entire BRDF can be expressed as a function of the half-vector. By storing this function, they get a compact BRDF model, that fits measured data very well. We use their intuition in our fitting process to provide a first approximation. Brady *et al.* [BLPW14] applied genetic programming to discover new analytical BRDF models from measured BRDFs. We use their distribution as the basis of our reflectance lobe. Our contributions include the normalization constant, the Smith shadowing term and the importance sampling method for this distribution.

3 Scientific Background

3.1 Parameters

We consider a given material; the surface of this material is described by a height field h , as well as a normal distribution function $D(\mathbf{m})$. We are interested in the reflectance properties and whether they can be predicted from the surface properties, h and D .

The BRDF $\rho(\mathbf{i}, \mathbf{o})$ encodes the way the material reflects light. It is defined as the ratio of outgoing radiance in direction \mathbf{o} by incoming irradiance along \mathbf{i} .

The two competing models for reflectance properties, micro-facet model and diffraction models, use different parameters, making direct comparison difficult. Micro-facet models depend mostly on the half-vector \mathbf{h} , defined as:

$$\mathbf{h} = (\mathbf{i} + \mathbf{o}) / \|\mathbf{i} + \mathbf{o}\|, \quad (1)$$

while diffraction models depend on a vector \mathbf{f} , defined as the projection on the tangent plane of the difference between the reflected incoming direction and the outgoing direction, divided by the wavelength (see Figure 2):

$$\mathbf{f} = \frac{1}{\lambda} \text{proj}(\text{refl}(\mathbf{i}) - \mathbf{o}) \quad (2)$$

For isotropic materials, diffraction effects only depend on f , the norm of \mathbf{f} , which can be expressed using the half-angle parametrization [Rus98]:

$$f = \frac{2}{\lambda} \sin \theta_h \cos \theta_d \quad (3)$$

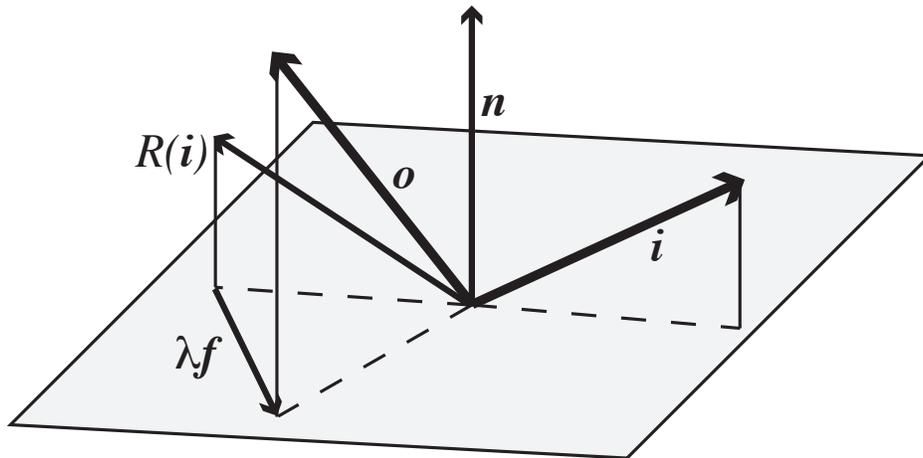


Figure 2: Diffraction effects depend on f , projected vector between reflected incoming direction and outgoing direction, divided by the wavelength.

3.2 Micro-Facet Model

The micro-facet model [CT82] is widely used in Computer Graphics. It is based on the hypothesis that surface details are much larger than light wavelength, so we can apply the rules of optical geometry. The model has been described in great details in the literature [CT82, WMLT07, Hei14] so we focus here on the key points.

The model assumes that the surface is made of micro-facets, with a normal distribution function D . Each micro-facet acts as a specular mirror; its reflectance is a Dirac delta function, multiplied by the Fresnel term for reflection:

$$\rho_{\mu F}(\mathbf{i}, \mathbf{o}) = F(\mathbf{i}, \mathbf{o})\delta(\text{refl}(\mathbf{i}), \mathbf{o}) \quad (4)$$

The Fresnel term F is responsible for the color of the facet. It depends on the index of refraction η of the material; the index of refraction depends on the wavelength, and this dependency is responsible for the color of the material. For dielectrics (transparent materials), η is a real number, related to the speed of light inside the material. For conductors, such as copper and nickel, η is a complex number, with real and imaginary parts: $\eta = n + ik$.

Summing the contributions of all micro-facets results in the Cook-Torrance reflectance model:

$$\rho_{CT}(\mathbf{i}, \mathbf{o}) = \frac{F(\theta_d)D(\theta_h)G(\mathbf{i}, \mathbf{o})}{4 \cos \theta_i \cos \theta_o} \quad (5)$$

The shadowing-masking term G expresses the probability for light to be blocked before or after the specular reflection on the micro-facet. It is usually computed using Smith's method, through a double integration of D [Smi67, WMLT07, Hei14]. First, we approximate G as the product of separate, one-dimensional shadowing and masking functions, then compute these using Λ function:

$$G(\mathbf{i}, \mathbf{o}) \approx G_1(\theta_i)G_1(\theta_o) \quad (6)$$

$$G_1(\theta) = \frac{1}{1 + \Lambda(\theta)} \quad (7)$$

Λ is computed through a double integration of D , through auxiliary functions P_{22} and P_2 :

$$P_{22}(\tan^2 \theta) = D(\theta)\chi_{[0, \frac{\pi}{2}]}(\theta) \cos^4 \theta \quad (8)$$

$$P_2(r) = \int_{-\infty}^{\infty} P_{22}(r^2 + q^2) dq \quad (9)$$

$$\Lambda(\theta) = \int_{1/\tan \theta}^{\infty} (r \tan \theta - 1)P_2(r) dr \quad (10)$$

Note that D and G only depend on the *geometry* of the micro-facets, not on the wavelength of the incoming light. The only wavelength dependency comes from the Fresnel term F . For this model, the reflection lobes for the different wavelengths should be scaled version of each other, with the same relative width.

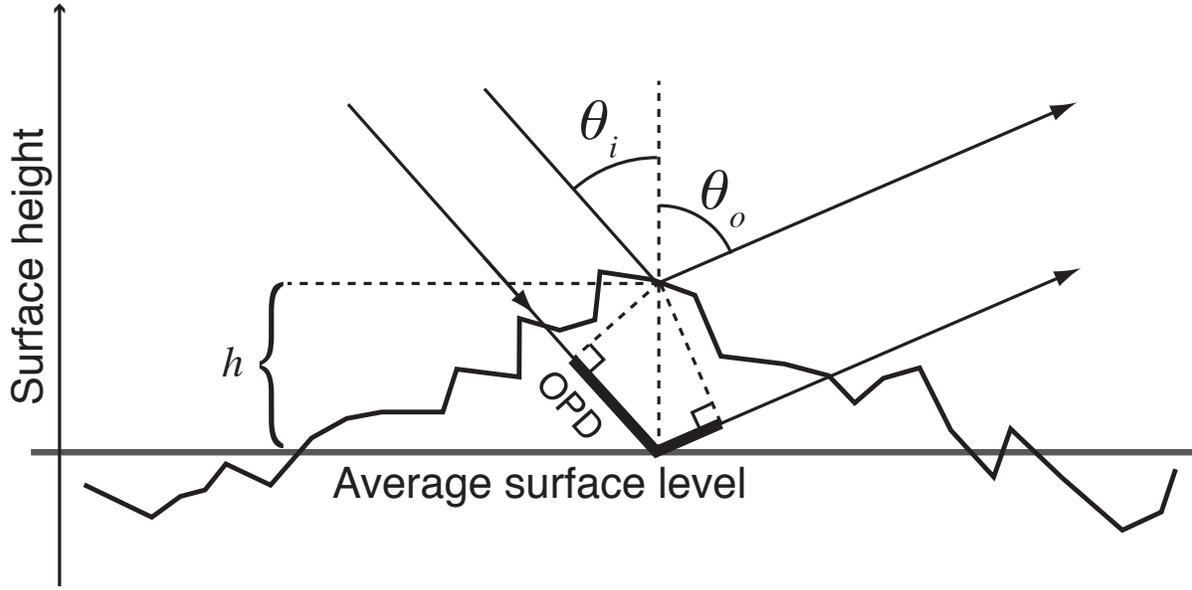


Figure 3: Light reflected off the surface at a height h has travelled a shorter distance. This Optical Path Difference (OPD) causes interference between parallel rays and is responsible for diffraction effects. Here, $OPD = (\cos \theta_i + \cos \theta_o)h(x, y)$ [Kry06]. Heights h are relative to the average surface level.

3.2.1 Color and Fresnel term for reflection

The material color comes from the Fresnel term. For unpolarized light, the Fresnel term is computed using these equations:

$$F_s(\theta) = \left| \frac{\cos \theta - (\eta^2 - \sin^2 \theta)^{1/2}}{\cos \theta + (\eta^2 - \sin^2 \theta)^{1/2}} \right|^2 \quad (11)$$

$$F_p(\theta) = \left| \frac{\eta^2 \cos \theta - (\eta^2 - \sin^2 \theta)^{1/2}}{\eta^2 \cos \theta + (\eta^2 - \sin^2 \theta)^{1/2}} \right|^2 \quad (12)$$

$$F = \frac{F_s + F_p}{2} \quad (13)$$

These equations remain valid even if η is a complex number, as happens for conductors. F_s , F_p and F are always real numbers. In the presence of polarized light, F_s gives the material reflectance for perpendicular polarized light, and F_p the reflectance for parallel-polarized light.

3.3 Diffraction: Generalized Harvey-Shack Theory

In this section, we review the basis of the diffraction theory for predicting material reflectance. We focus on the Harvey-Shack theory, because it is applicable to a wide range of configurations and has been found to provide a good approximation to measured reflectance. In the diffraction theory, we consider the distance travelled by light as it is reflected on the surface. Because the surface is not flat, light being reflected at a point of height $h(x, y)$ has travelled a shorter distance than if it was reflected at the reference height (see Figure 3). This difference in optical path length (OPD) is:

$$OPD = (\cos \theta_i + \cos \theta_o)h(x, y) \quad (14)$$

This results in phase variations at the reflected wavefront of $(2\pi/\lambda)(\cos \theta_i + \cos \theta_o)h(x, y)$. Averaging these phase variations gives the reflectance function.

According to the Harvey-Shack theory, reflectance is separated in two different lobes: a peak in the direction of specular reflection, corresponding to a Dirac delta function, surrounded by a halo of scattered light [Har75, Kry06]. The relative

intensity of these two components depend on the surface roughness σ_s^2 , defined as the variance of height distribution h :

$$\rho_{\text{diff}}(\mathbf{i}, \mathbf{o}) = A\delta(\text{refl}(\mathbf{i}), \mathbf{o}) + (1 - A)S_{GHS}(f) \quad (15)$$

$$\text{with } A = e^{-(2\pi\frac{\sigma_s}{\lambda}(\cos\theta_i + \cos\theta_o))^2} \quad (16)$$

where S_{GHS} is an angle spread function.

By construction, this equation guarantees conservation of energy: if the surface is rough, σ_s/λ is large; the energy of the diffraction lobe increases, and the energy of the specular peak decreases. The relative importance of the specular peak and diffraction lobe depends on the directions: at grazing angles, $\cos\theta_i \approx 0$ and A is almost equal to 1; most of the energy goes into the specular reflection. The influence of diffraction is stronger at normal incidence, where $\cos\theta_i \approx 1$.

Krywonos [Kry06] showed that the Generalized Harvey-Shack theory provides a good fit with measured reflectance, outside of the specular peak. It provides a better match than competing diffraction models, such as Rayleigh-Rice and the original Harvey-Shack theory. It provides equivalent results to the Modified Beckmann-Kirchoff theory.

3.3.1 Smooth Surface Linearization

If the reflecting surface is smooth, we can assume that $\sigma_s/\lambda \ll 1$ and linearize the terms in Equation 16:

$$A \approx 1 - \left(2\pi\frac{\sigma_s}{\lambda}(\cos\theta_i + \cos\theta_o)\right)^2 \quad (17)$$

The diffraction lobe becomes:

$$\rho_{\text{diff}}(\mathbf{i}, \mathbf{o}) = 4\pi^2\frac{\sigma_s^2}{\lambda^2}(\cos\theta_i + \cos\theta_o)^2 S_{GHS}(f) \quad (18)$$

Under the smooth surface linearization, we compute the diffraction function from the Fourier transform of the autocovariance function of the surface height, which is the Power Spectral Density (PSD) of the surface:

$$\rho_{\text{diff}}(\mathbf{i}, \mathbf{o}) = \frac{4\pi^2}{\lambda^4}(\cos\theta_i + \cos\theta_o)^2 \text{PSD}(f) \quad (19)$$

We recognize here the usual formula for Harvey-Shack diffraction for smooth surfaces. This expression depends on the wavelength in two places: the *intensity* of the diffraction lobe decreases with $1/\lambda^4$; the *width* of the lobe increases with λ , since f contains a $1/\lambda$ term. The diffraction lobe for blue colors is sharper, but with higher intensity, while the diffraction lobe for red colors spreads more, but with lower maximum intensity.

The Generalized Harvey-Shack diffraction theory is not limited to this smooth surface approximation. In the generic case, the angle spread function S_{GHS} changes with incoming direction θ_i and is computed by a Fourier transform of surface height.

3.3.2 K-correlation Model

If the surface height distribution follows a Gaussian, then the PSD should also be Gaussian, making computations simpler. Most optical surfaces however do not have a Gaussian distribution of heights. They tend to have an inverse power law falloff at large spatial frequencies. For these surfaces, we use the K-correlation model, or abc model, to express the power spectral density:

$$\text{PSD}(f) = \frac{1}{2\sqrt{\pi}} \frac{\Gamma((c+1)/2)}{\Gamma(c/2)} \frac{ab}{(1+b^2f^2)^{\frac{c+1}{2}}} \quad (20)$$

Integrating this distribution gives the surface roughness σ_s^2 . It is finite only for $c > 1$ and is equal to:

$$\sigma_s^2 = \frac{\sqrt{\pi}}{c-1} \frac{\Gamma((c+1)/2)}{\Gamma(c/2)} \frac{a}{b} \quad (21)$$

We use this surface roughness to express the Power Spectral Density in a simpler expression:

$$\text{PSD}(f) = \frac{c-1}{2\pi} \frac{\sigma_s^2 b^2}{(1+b^2f^2)^{\frac{c+1}{2}}} \quad (22)$$

3.3.3 Renormalization

The angle spread function S_{GHS} used in Equation 16 is defined over the entire plane for its parameter, f . But f is limited in its variation, as the projections of incoming and outgoing directions must fall over the unit disc in the plane. For energy conservation, S_{GHS} must be renormalized by dividing by the integral over the set of possible values for vector f :

$$\hat{S}_{GHS} = \frac{\sigma_0^2}{\sigma_{rel}^2} S_{GHS} \quad (23)$$

$$\sigma_0^2 = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} S_{GHS}(f) df_x df_y \quad (24)$$

$$\sigma_{rel}^2(\theta_i) = \int_{\|\theta\| \leq 1} S_{GHS}(f) df_x df_y \quad (25)$$

This renormalization ensures that the energy of the diffraction lobe is constant. Usually, angle spread functions are normalized so that $\sigma_0^2 = 1$. In that case, Equation 23 is equivalent to:

$$\hat{S}_{GHS} = \frac{1}{\sigma_{rel}^2} S_{GHS} \quad (26)$$

3.3.4 Color and Polarization

As for specular reflection and micro-facet models, the color of the diffraction lobe comes from an index of refraction dependent term. This term is noted Q ; it is different from the Fresnel term for reflection, F , defined in Section 3.2.1 [Kry06, BNM15a]. For unpolarized light, it is defined as:

$$\begin{aligned} \phi &= \phi_o - \phi_i - \pi \\ Q_{ss} &= \frac{(\eta^2 - 1) \cos \phi}{(\cos \theta_i + \sqrt{\eta^2 - \sin^2 \theta_i})(\cos \theta_o + \sqrt{\eta^2 - \sin^2 \theta_o})} \\ Q_{sp} &= \frac{(\eta^2 - 1) \sqrt{\eta^2 - \sin^2 \theta_o} \sin \phi}{(\cos \theta_i + \sqrt{\eta^2 - \sin^2 \theta_i})(\eta^2 \cos \theta_o + \sqrt{\eta^2 - \sin^2 \theta_o})} \\ Q_{ps} &= \frac{(\eta^2 - 1) \sqrt{\eta^2 - \sin^2 \theta_i} \sin \phi}{(\eta^2 \cos \theta_i + \sqrt{\eta^2 - \sin^2 \theta_i})(\cos \theta_o + \sqrt{\eta^2 - \sin^2 \theta_o})} \\ Q_{pp} &= (\eta^2 - 1) \times \\ &\quad \frac{\sqrt{\eta^2 - \sin^2 \theta_i} \sqrt{\eta^2 - \sin^2 \theta_o} \cos \phi - \eta^2 \sin \theta_i \sin \theta_o}{(\eta^2 \cos \theta_i + \sqrt{\eta^2 - \sin^2 \theta_i})(\eta^2 \cos \theta_o + \sqrt{\eta^2 - \sin^2 \theta_o})} \\ Q &= |Q_{ss}|^2 + |Q_{sp}|^2 + |Q_{ps}|^2 + |Q_{pp}|^2 \end{aligned} \quad (27)$$

Q depends on the azimuthal angle between reflected incoming direction and outgoing direction. With polarized light, part of the parallel-polarized light is reflected at perpendicular-polarized, and reciprocally. The Q_{sp} and Q_{ps} terms express this.

Along the specular reflected direction, where $\theta_i = \theta_o$ and $\phi_o = \phi_i + \pi$, we have $Q_{sp} = 0$ and $Q_{ps} = 0$, and $Q = 2F$. Outside of this specific case, the two color terms are substantially different.

4 Our Model

Our reflectance model is based on the following observation: reflectance and diffraction effects are *not* mutually exclusive. The Generalized Harvey-Shack theory explicitly allocates energy to the reflection and diffraction lobes, combining a specular reflection with a diffraction lobe (see Equation 16).

Most surfaces in Computer Graphics are not smooth, at least not in the sense used by the Optics community. They have variations in surface normal. These variations cause a larger distribution of specular reflections, instead of just a Dirac. The Cook-Torrance model is designed precisely to connect the variations of surface normals with the general reflectance properties of the surface.

4.1 Derivation from micro-surface properties

More precisely, we assume that the diffraction model is valid at the micro-geometry level. Each micro-facet reflects light according to Equation 15. We assume that we have a micro-facet normal distribution $D(\mathbf{m})$, describing the statistical distribution of micro-facet normals \mathbf{m} . This micro-facet normal distribution has a larger scale than the micro-geometry responsible for diffraction.

We compute the surface reflectance at the macroscopic level by integrating the contributions over all visible microfacets:

$$\rho(\mathbf{i}, \mathbf{o}, \mathbf{n}) = \int \left| \frac{\mathbf{i} \cdot \mathbf{m}}{\mathbf{i} \cdot \mathbf{n}} \right| \rho_{\text{diff}}(\mathbf{i}, \mathbf{o}, \mathbf{m}) \left| \frac{\mathbf{o} \cdot \mathbf{m}}{\mathbf{o} \cdot \mathbf{n}} \right| G(\mathbf{i}, \mathbf{o}, \mathbf{m}) D(\mathbf{m}) d\omega_m \quad (28)$$

The integral gives a reflectance model with two contributions:

$$\rho(\mathbf{i}, \mathbf{o}, \mathbf{n}) = \rho_{\text{refl.}}(\mathbf{i}, \mathbf{o}) + \rho_{\text{diff.}}(\mathbf{i}, \mathbf{o}) \quad (29)$$

The reflection part is simpler; the dirac inside the integral transforms the integral into a function evaluation, after taking into account the proper Jacobian for a change of variable (see, e.g. [WMLT07]). Angles between micro-facet normal and incoming or outgoing direction become angles between half-vector and incoming or outgoing directions, that is θ_d :

$$\rho_{\text{refl.}}(\mathbf{i}, \mathbf{o}) = e^{-(4\pi\sigma_s \cos \theta_d)^2} \frac{F(\theta_d) D(\theta_h) G(\mathbf{i}, \mathbf{o})}{4 \cos \theta_i \cos \theta_o} \quad (30)$$

The diffraction part is expressed as:

$$\rho_{\text{diff.}}(\mathbf{i}, \mathbf{o}) = \int \left| \frac{\mathbf{i} \cdot \mathbf{m}}{\mathbf{i} \cdot \mathbf{n}} \right| (1 - A) S_{GHS}(\mathbf{f}) \left| \frac{\mathbf{o} \cdot \mathbf{m}}{\mathbf{o} \cdot \mathbf{n}} \right| G(\mathbf{i}, \mathbf{o}, \mathbf{m}) D(\mathbf{m}) d\omega_m \quad (31)$$

We cannot compute this integral; we approximate it using a different angle spread function, $S'_{GHS}(\mathbf{f})$:

$$\rho_{\text{diff.}}(\mathbf{i}, \mathbf{o}) \approx (1 - A) S'_{GHS}(\mathbf{f}) G(\mathbf{i}, \mathbf{o}) \quad (32)$$

$$(33)$$

This gives the complete model:

$$\rho(\mathbf{i}, \mathbf{o}) = \rho_{\text{refl.}}(\mathbf{i}, \mathbf{o}) + \rho_{\text{diff.}}(\mathbf{i}, \mathbf{o}) \quad (34)$$

$$\rho_{\text{refl.}}(\mathbf{i}, \mathbf{o}) = e^{-(4\pi\sigma_s \cos \theta_d)^2} \frac{F(\theta_d) D(\theta_h) G(\mathbf{i}, \mathbf{o})}{4 \cos \theta_i \cos \theta_o} \quad (35)$$

$$\rho_{\text{diff.}}(\mathbf{i}, \mathbf{o}) = (1 - A) S'_{GHS}(\mathbf{f}) G(\mathbf{i}, \mathbf{o}) \quad (36)$$

$$\text{where } A = e^{-(2\pi \frac{\sigma_s}{\lambda} (\cos \theta_i + \cos \theta_o))^2} \quad (37)$$

Compared to the Generalized-Harvey-Shack diffraction model, we make the following assumptions: we do not use the smooth surface approximation, A is expressed as an exponential instead of the linearized version; we represent the combined effect of surface normal and diffraction using an angle spread function S_{GHS} and assume it remains the same for all θ_i . However, it is not necessarily the Power Spectral Distribution of the surface height.

The color for both micro-facet and diffraction terms comes from the index of refraction of the material, through the Fresnel term F for the micro-facet component, and through the reflectance Q for diffraction. Figure 1 shows examples of the contributions of diffraction and reflection for several materials. You can see slight differences in color between the components.

In the following sections, we review the components of the model: the micro-facet distribution for the Cook-Torrance model and the angle spread function for the diffraction component. We also combine our surface reflectance model with two multi-layer models: sub-surface scattering and plastics with specular varnish.

4.2 Micro-Facet distribution

For our micro-facet distribution, we use a Generalized Beckmann distribution; it can model the sharp peak we observe on measured reflectance for many surfaces (see Figures 11 and 12):

$$D(\theta) = \frac{\chi_{[0, \frac{\pi}{2}]}(\theta)}{\cos^4 \theta} P_{22}(\tan^2 \theta) \quad (38)$$

$$P_{22}(x) = \frac{P}{\pi \beta^2 \Gamma(1/p)} e^{-\left(\frac{x}{\beta^2}\right)^p} \quad (39)$$

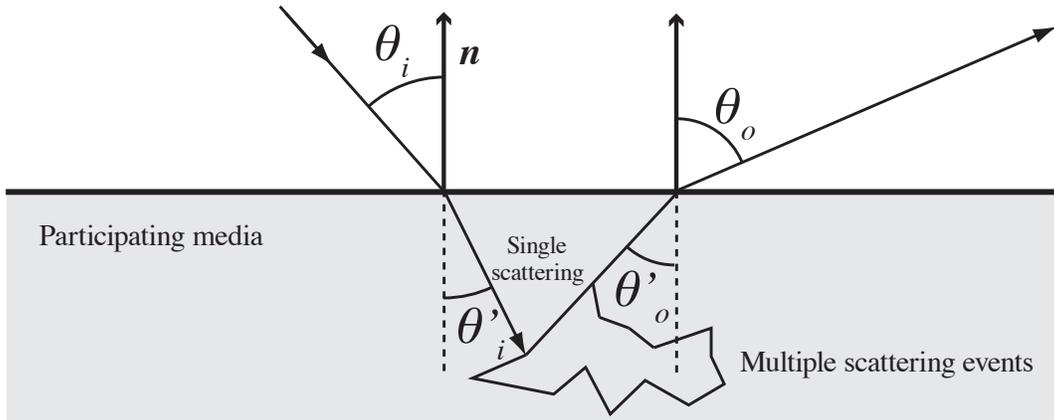


Figure 4: We also consider multiple layer materials. In the *subsurface* model, light is reflected on the surface as well as refracted inside the underlying participating media, where it is scattered before leaving again through the refractive interface.

It has two parameters: β controls the width of the peak, while p controls its kurtosis, that is whether it drops sharply at the origin or not. This distribution is similar to the “BRDF model A” obtained by Brady et al. [BLPW14] by analysing measured BRDFs with genetic programming; we added the $1/\cos^4 \theta$ term, which is required for analytical integration. This allows us to provide normalization, Smith shadowing and importance sampling.

We compute the shadowing term G for this distribution using Smith’s method [Smi67, WMLT07, Hei14], as described in Section 3.2. For this distribution:

$$P_2(r) = \frac{p}{\pi\Gamma(1/p)} \int_0^\infty e^{-(r^2+q^2)^p} dq \quad (40)$$

$$\Lambda(\beta \tan \theta) = \int_{1/\beta \tan \theta}^\infty (r\beta \tan \theta - 1)P_2(r) dr \quad (41)$$

$$G_1(\theta) = \frac{1}{1 + \Lambda(\beta \tan \theta)} \quad (42)$$

We observe that Λ , and therefore G_1 , can be expressed as a function of $\beta \tan \theta$ and p , reducing the number of parameters to two. We precompute G_1 and store it in a 2D array, to be accessed at render time (a single 2D array for all values of the three parameters, with a memory cost of 10 kB) Finally, we compute G as: $G(\mathbf{i}, \boldsymbol{\theta}) = G_1(\theta_i)G_1(\theta_o)$.

4.3 Diffraction component

We did not use the smooth surface approximation of the Generalized Harvey-Shack model; as a consequence, S'_{GHS} is not necessarily the Power Spectral Distribution of the surface height. We nevertheless chose to express it using the K-correlation model described in Section 3.3.2, as it has good qualities. We used the simplified form defined in Equation 22, and renormalize it using Equation 26:

$$S'_{GHS}(f) = \frac{\sigma_s^2}{\sigma_{rel}^2} \frac{c-1}{2\pi\lambda^2} \frac{b^2}{(1+b^2f^2)^{\frac{c+1}{2}}} \quad (43)$$

This component has three parameters: b , corresponding to the *width* of the diffraction lobe, c , corresponding to the falloff and σ_s^2 , the surface roughness term.

4.4 Multi-layer Models

Some materials have multiple layers: for example diffuse plastic with a layer of transparent varnish, or a semi-precious stone with participating media. These layers contribute to the overall response of the material. We combine our surface model with two different multi-layer materials: sub-surface scattering BRDF and diffuse plastics. The first is the BRDF approximation of subsurface scattering, introduced by Jensen et al. [JMLH01] (see Figure 4). The second is the multi-layered model introduced by Weidlich and Wilkie [WW07] with a transparent layer of varnish over a diffuse layer (see Figure 5).

In both models, we use our reflectance model only for direct reflection by the top interface. For transmission through this interface, we approximate it as a flat specular surface. The transmission term for light passing through this interface at \mathbf{i} and

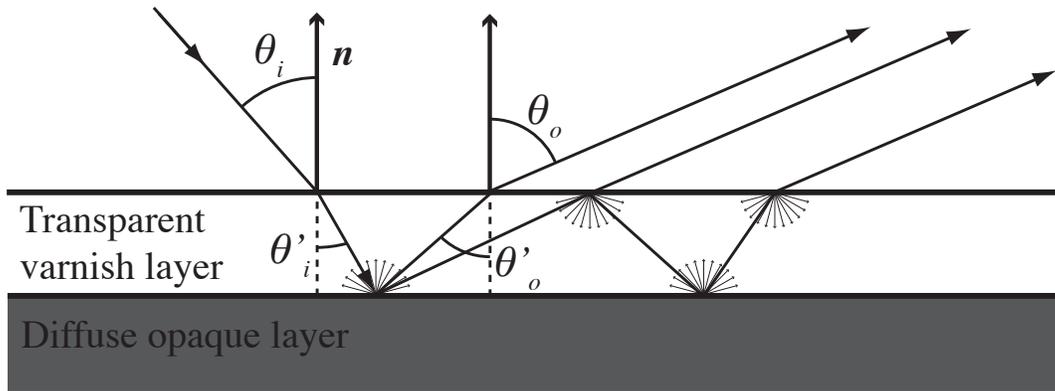


Figure 5: In the *plastic* two-layer model, we assume we have a transparent layer of refractive varnish over a diffuse material. We model light bouncing several times between the two interfaces before leaving the object.

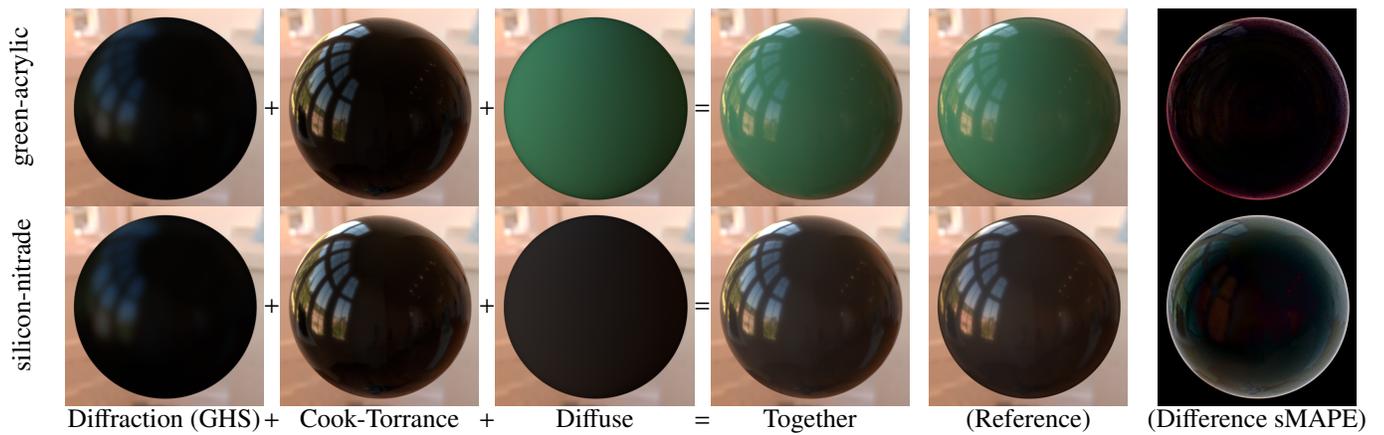


Figure 6: Examples of multi-layer materials: a plastic (green-acrylic) and a material with sub-surface scattering (silicon-nitride).

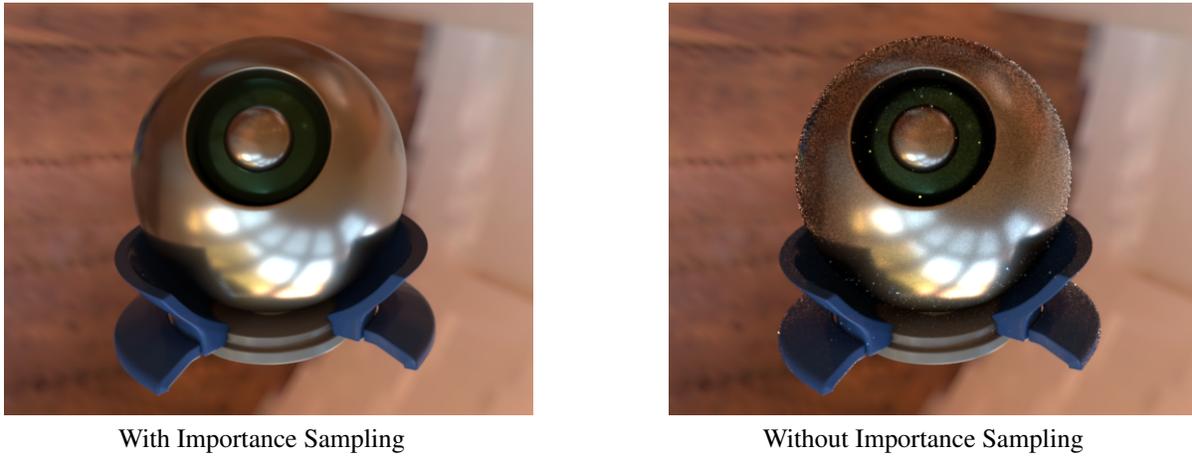


Figure 7: Importance sampling reduces the noise in global illumination computations (1000 samples per pixel in both pictures).

leaving in direction \mathbf{o} is:

$$T = (1 - F(\theta_i))(1 - F(\theta_o)) \quad (44)$$

where F is the Fresnel term for reflection, defined in Equation 13. In both models, we also use the Fresnel diffuse reflectance, F_{dr} , the average of Fresnel reflectance times the cosine of the angle with the normal over all directions. We use the rational approximation given by d'Eon and Irving [DI11].

The **sub-surface** BRDF model combines single and multiple scattering. Assuming a material with albedo α and an isotropic phase function, the resulting lobe is:

$$\text{single} = \frac{\alpha T}{4\pi} \frac{1}{\cos \theta'_i + \cos \theta'_o} \quad (45)$$

$$A_F = \frac{1 - F_{dr}}{1 + F_{dr}} \quad (46)$$

$$\text{multi} = \frac{\alpha T}{2\pi} e^{-\sqrt{3(1-\alpha)}} \left(1 + e^{-\frac{4}{3}A_F \sqrt{3(1-\alpha)}}\right) \quad (47)$$

$$\rho_{\text{SSS}} = \text{single} + \text{multi} \quad (48)$$

where θ'_i and θ'_o are the angles of the refracted rays with the surface normal (see Figure 4).

In the **plastic** BRDF model, we assume that we have a diffuse layer of reflectance ρ_d under a transparent layer. We account for light entering through the interface, being reflected on the diffuse layer, and leaving through the interface. We also account for possible multiple bounces between the diffuse layer and the refractive interface. The resulting component is:

$$\rho_{\text{plastic}} = T \frac{\rho_d}{\pi} \left(1 - \frac{\rho_d}{\pi} F_{dr} \left(\frac{1}{\eta}\right)\right)^{-1} \quad (49)$$

4.5 Importance sampling

The A coefficient express the energy repartition between the reflection and diffraction lobes. It depends on the incoming and outgoing directions and on the σ_s/λ ratio (see Figure 8). For $\sigma_s/\lambda > 0.3$, it is close to 0 on most of the parameters range, and close to 1 for grazing angles. We need to take it into account for importance sampling.

We perform importance sampling in three steps: first, we select whether to importance sample the diffuse lobe (if there is one) or the glossy lobe. For the glossy lobe, we select between the micro-facet and the diffraction lobe depending on the incoming direction. Once we have picked a lobe, we sample it using the appropriate strategy. Figure 7 displays the impact of importance sampling.

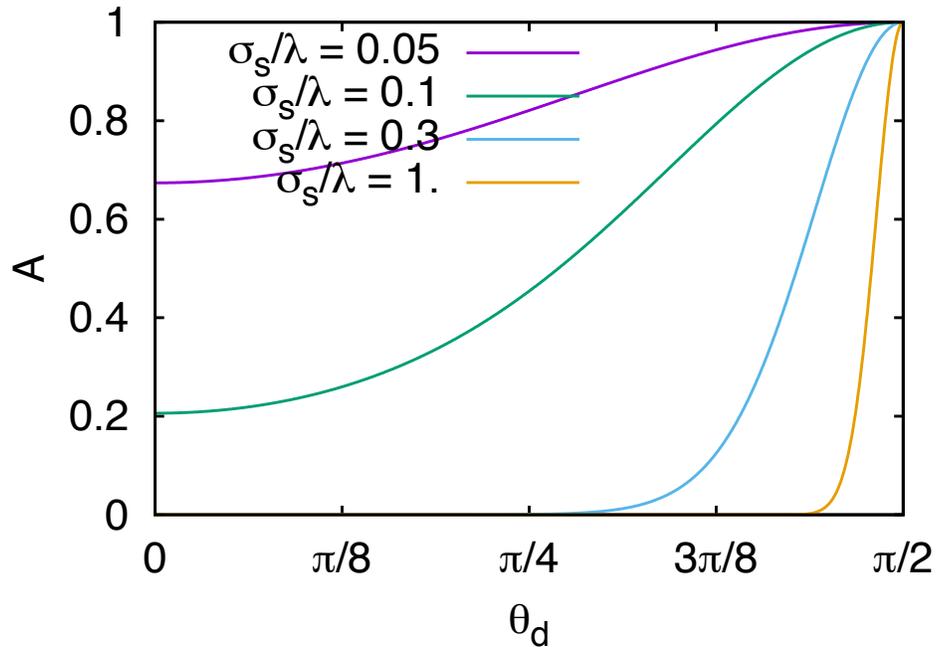


Figure 8: The A coefficient express the ratio of energy going into the reflection and diffraction lobes. It depends on the incoming angle and the σ_s/λ ratio.

4.5.1 Approximated Energy for each Lobe

To compute the approximate energy for each lobe, we begin by computing the values of the Fresnel reflection coefficient at normal incidence:

$$F_0 = \frac{n^2 + k^2 - 2n + 1}{n^2 + k^2 + 2n + 1} \quad (50)$$

$$(51)$$

We use this as the approximate energy of the combined micro-facet and diffraction lobes: $E_C = F_0$.

For multi-layers material, we also compute the energy of the diffuse lobe: For the plastic model:

$$E_{D \text{ Plastic}} = (1 - F_0)^2 \rho_d \left(1 - \frac{\rho_d}{\pi} F_{dr} \left(\frac{1}{\eta}\right)\right)^{-1} \quad (52)$$

For the sub-surface scattering model:

$$E_{D \text{ SSS}} = \alpha(1 - F_0)^2 \frac{1 + R_d}{2} \quad (53)$$

$$\text{where } R_d = e^{-\sqrt{3(1-\alpha)}} \left(1 + e^{-\frac{4}{3}A_F \sqrt{3(1-\alpha)}}\right) \quad (54)$$

4.5.2 Lobe Selection

To importance sample the BRDF, we begin by selecting which lobe to sample. We take a random variable u_1 in $[0, 1)$ to select between the diffuse lobe and the combined micro-facet and diffraction lobes:

$$p_C = E_C / (E_C + E_D) \quad (55)$$

$$p_D = E_D / (E_C + E_D) \quad (56)$$

We sample the combined micro-facet and diffraction lobes if $u_1 < p_C$ and the diffuse lobe otherwise. We use cosine-weighted sampling for the plastic or sub-surface scattering diffuse lobes.

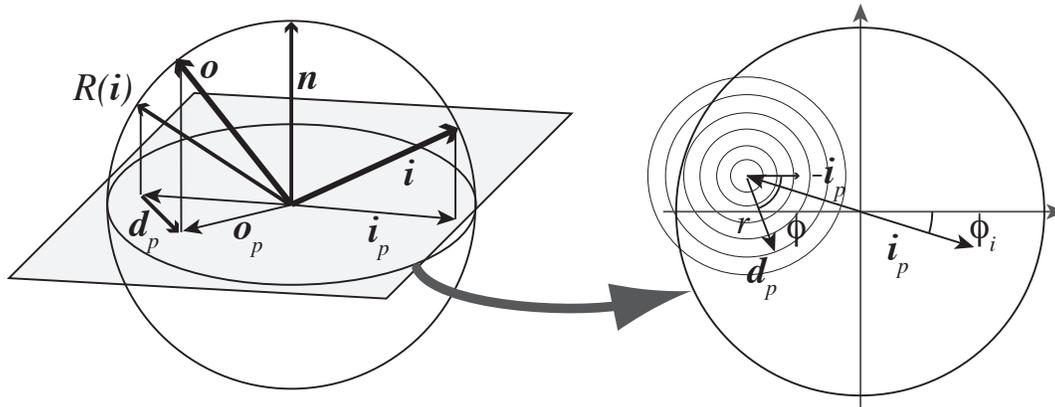


Figure 9: To importance sample the diffraction lobe, we start from the projection of the reflection of incoming direction \mathbf{i} , $-\mathbf{i}_p$. We compute the difference between projected incoming direction and outgoing directions, \mathbf{d}_p and use it to compute the outgoing direction \mathbf{o} .

We then compute A_{imp} , using the known incoming direction:

$$A_{\text{imp}} = e^{-(2\pi \frac{\sigma_s}{\lambda} \cos \theta_i)^2} \quad (57)$$

We take a second random variable u_2 in $[0, 1)$, and sample the micro-facet lobe if $u_2 < A$, and the diffraction lobe otherwise. To avoid using too many random variables, we renormalise each random variable after use.

4.5.3 Importance sampling for the micro-facet lobe

To importance sample the micro-facet lobe, we take two uniform random variables (u_1, u_2) in $[0, 1)$ and use them to build the half-vector \mathbf{h} through (θ_h, ϕ_h) :

$$\phi_h = 2\pi u_1 \quad (58)$$

$$\theta_h = \arctan \left(\beta \left(\gamma_u^{-1}(1/p, u_2) \right)^{\frac{c}{2}} \right) \quad (59)$$

where γ_u^{-1} is the inverse of the normalized upper incomplete gamma function. We then build the outgoing direction \mathbf{o} from \mathbf{h} .

4.5.4 Diffraction Lobe Importance Sampling

To importance sample the diffraction lobe, we work in the unit disk corresponding to the projections of all directions (see Figure 9). We begin with the projection of the incoming direction, $\mathbf{i}_p = (\sin \theta_i \cos \phi_i, \sin \theta_i \sin \phi_i)$. The projection of the *reflected* incoming direction is $-\mathbf{i}_p$. We build a 2D vector \mathbf{d}_p , the difference between $-\mathbf{i}_p$ and the projection of the outgoing direction \mathbf{o}_p . From this, we extract the projection $\mathbf{o}_p = -\mathbf{i}_p + \mathbf{d}_p$ and then the outgoing direction $\mathbf{o} = \mathbf{o}_p + \sqrt{1 - \|\mathbf{o}_p\|^2} \mathbf{n}$. Our importance sampling method does not give equal probability for all points in the disc. We compensate for this point by dividing the sampling weight for each point by its probability.

We take two uniform random variables (u_1, u_2) in $[0, 1)$ and use them to build \mathbf{d}_p in polar coordinates (r, ϕ) :

$$r = \frac{\lambda}{b} \sqrt{(1 - M u_1)^{-\frac{2}{c-1}} - 1} \quad (60)$$

$$\text{where } M = 1 - \left(1 + \frac{b^2}{\lambda^2} (1 + \sin \theta_i)^2 \right)^{-\frac{c-1}{2}} \quad (61)$$

This gives $r \in [0, 1 + \sin \theta_i]$. We then select ϕ so that \mathbf{o}_p remains inside the unit disc:

$$\phi_{\max} = \arccos \left(\max \left(-1, \frac{r^2 + \sin^2 \theta_i - 1}{2r \sin \theta_i} \right) \right) \quad (62)$$

$$\phi = \phi_i + \pi + (2u_2 - 1)\phi_{\max} \quad (63)$$

Our importance sampling of the diffraction lobe is simpler than the method provided in [LKYU12], and works for all valid values of c . It does not provide an equal coverage of the unit disk, so each sample must be multiplied by ϕ_{\max}/π .

5 Validation with measured materials

To test the validity of our model, we have used it to fit all 100 materials in the MERL database [MPBM03]. We searched for optimal parameters using Levenberg-Marquadt optimisation, as implemented by Lourakis [Lou04]. We conduct our fitting over the entire parameter domain $(\theta_i, \theta_o, \phi_o - \phi_i)$, using an L^2 norm:

$$E = \int_{\theta_i} \int_{\Omega} (\text{measured}(\mathbf{i}, \mathbf{o}) - \text{predicted}(\mathbf{i}, \mathbf{o}))^2 d\omega_o \cos \theta_i d\theta_i \quad (64)$$

We restrict the fitting to $\theta_i > 1.4$, $\theta_o > 1.4$, to avoid grazing angles, where measured data is unreliable. We fit over all three channels simultaneously. We enforce physical consistency by restricting these parameters: $c > 1$, $|\eta| > 1$.

5.1 Parameters

All our models have 5 parameters related to the surface micro-geometry: the surface roughness σ_s , 2 parameters for the normal distribution β and p and 2 parameters for the angle spread function b and c . Two of these parameters correspond to distances in the diffraction micro-geometry: b is its characteristic length, σ_s is the variance of its height. As these parameters only appear in fractions with the wavelength λ : b/λ , σ_s/λ , they are expressed in the same unit as λ (micro-meters in our implementation, to avoid numerical issues).

Our single layer model has 6 additional parameters, corresponding to the index of refraction $\eta = n + ik$ of the material for each color channel, with real and imaginary parts. For our multi-layer models, there are only 4 extra parameters: the albedo or diffuse color of the underlying material, and the index of refraction η . For the multi-layer models, we made the assumption that the index of refraction would be a constant, independent of the wavelength. Index of refraction for transparent materials in the visible domain is well approximated using Cauchy's formula: $\eta = B + C/\lambda^2$. In practice, the variations with wavelength appear to be negligible compared with other sources of imprecision.

Our single layer model thus has 11 total parameters, the multi-layer models have 9. For a more compact representation, it is possible to remove the imaginary part k of the index of refraction, as its influence over material behaviour is relatively small. This would bring the total number of parameters to 8 for the single layer model, at the expense of physical consistency.

Our model is extremely constrained: the micro-facet and diffraction lobes are connected by the surface roughness and the index of refraction. The number of degrees of freedom is close to a minimum. By enforcing the connection between the different components, we sometimes make it harder for the fitting algorithm to converge.

The supplemental material includes a detailed comparison between measured data, our model and two state-of-the-art material fits: Shifted Gamma Distribution [BSH12] and Smooth reflectance model [LKYU12]. We provide lobe shape visualization, both in the incident plane and outside, rendering comparisons and error curves. Both models we are comparing against are not physically consistent: the normal distribution in [BSH12] changes with the color channel, while [LKYU12] models diffraction effects without wavelength dependency. We get similar results while being physically consistent.

Please note that Bagher et al. [BSH12] use a total of 18 parameters, significantly more than our model; Low et al. [LKYU12] use 9 parameters, but do not use the imaginary part of the index of refraction. They also have a diffuse lobe, with different color. Differences in fitting quality are not related to our model having more parameters. Holzschuch and Pacanowski [HP15a] provide a much better fit with measured materials, but use 15 parameters and separate the materials for the specular lobe and diffraction lobe.

5.2 Implementation Details

5.2.1 Two-Step Fitting Process

The time it takes to run an optimisation process is correlated with the number of function evaluations: fitting over the entire domain $(\theta_i, \theta_o, \phi_o - \phi_i)$ requires $O(n^3)$ function evaluations. To speed-up the fitting, we designed a two-steps fitting process: in a first step, we fit only over the backscattering data, with $\mathbf{i} = \mathbf{o}$. As pointed out by Ashikmin and Premože [AP07], this slice of the data contains most of the information about BRDF behaviour. This first step is relatively quick as we only fit over a 1-dimension slice of data.

In a second step, we fit over the entire data, using as starting points the values we found in the first step. This second step converges in a small number of iterations (less than 20), as we are already close to a minimum.

We also tried a single step fitting process. It converges for some materials but not all, and tends to leave some parameters unaffected.

5.2.2 Color Channels Wavelengths

The most common wavelengths for the red, green and blue channels in color displays are 645 nm, 526 nm and 444 nm. We used these values in all our computations, with positive results.

The BRDFs in the MERL database were captured using a QImaging Retiga 1300 camera [MPBM03]. This camera has continuous color sensitivity for each channel, with peaks at roughly 610 nm, 540 nm and 470 nm. The raw numerical data stored in the MERL database is scaled before use, with coefficients of 1 for the red channel, 1.15 for the green channel and 1.66 for the blue channel. We found that we also obtain good fits with measured data if we use the peak sensitivity wavelengths of the Retiga 1300 camera and the unscaled data. We used the standard wavelength and the scaled data for compatibility with other systems.

5.2.3 Renormalization

The renormalization step is costly, but important for accuracy outside of normal incidence. We have precomputed the value of σ_{rel}^2 for all values of b , c and $\sin \theta_i$ and stored them in a large look-up table, which we access at runtime.

5.3 Analysis of the Fitting

The fitting results bring the following observations:

- Materials named after semi-precious stones (aventurine, alumina-oxide, white-marble. . .) are well approximated with the subsurface model.
- Materials named after specular plastics (specular-blue-phenolic, blue-acrylic. . .) are well approximated with the plastic model.
- Diffraction effects are visually more important for smooth surfaces, such as chrome, nickel, specular-blue-phenolic. . .
- Several materials previously regarded as multi-lobes are well approximated by our model: red-metallic-paint, alum-bronze, two-layer-silver and two-layer-gold.
- Our model does not perform so well for very rough and diffuse surfaces, including diffuse paints and fabrics. This may be due to two different cause: other effects are present including backscattering and multiple scattering inside the micro-geometry, and the fitting process has difficulties separating between the different phenomena.
- An interesting effect appears for rough surfaces (where $\sigma > 0.5\lambda$). For these materials, A is negligible for almost all directions (see Figure 8). The impact of the micro-facet lobe is only felt at grazing angles. As we restricted fitting to angles smaller than 1.4 radians, we do not get any input on the micro-facet distribution from the fitting. This results in surprising values for the micro-facet distribution, including highly specular lobes on some fabrics. The only solution here is to fit over the entire interval. The fitting performs better on materials where the effects of both lobes are being felt over the entire domain, such as alum-bronze.

5.4 Editing Material Parameters

Our reflectance model introduces a separation between two shape components: diffraction and reflection. Figure 10 shows the effects of editing each parameter independently:

- Reducing σ_s reduces the importance of the diffraction lobe. With a low value, we get only the specular reflection. The color shift between the diffraction and micro-facet lobes appears clearly on Figure 10(b). This color shift is also present in Figure 10(a): the material is bluer at normal incidence, more orange at grazing angles, but as the transition is smooth, the effect is less perceptible.
- Acting on b changes the width of the diffraction lobe. Smaller values of b correspond to sharper lobes, as visible on Figure 10(c). Larger values of b result in larger lobes, more diffuse.
- Acting on β has the same effect for the micro-facet lobe. If we keep the original value for σ_s , the effects are not visible, as $A \approx 0$ for most of the scene. We reduce σ_s in Figure 10(d) to increase the importance of the micro-facet lobe. Comparing Figure 10(c) and (d), you see the color difference between diffraction and micro-facet lobes.

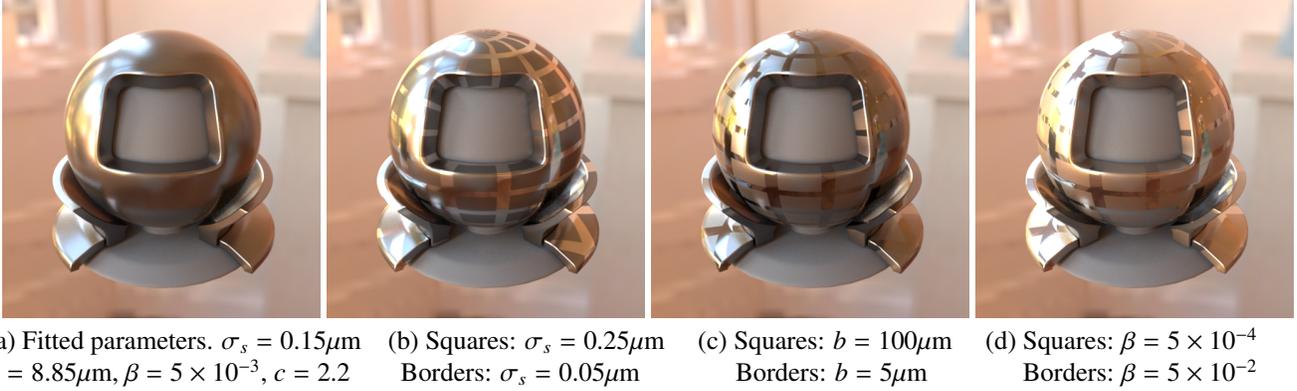


Figure 10: Effects of parameter edition for micro-facet geometry for our reflectance model. (a) nickel, with parameters from fitting. (b), (c) and (d): effect of modifying each of the parameters, in a grid pattern. We used different parameters values for the squares and the borders of the grid. For (d), $\sigma_s = 0.05\mu\text{m}$ to increase the impact of the micro-facet lobe.

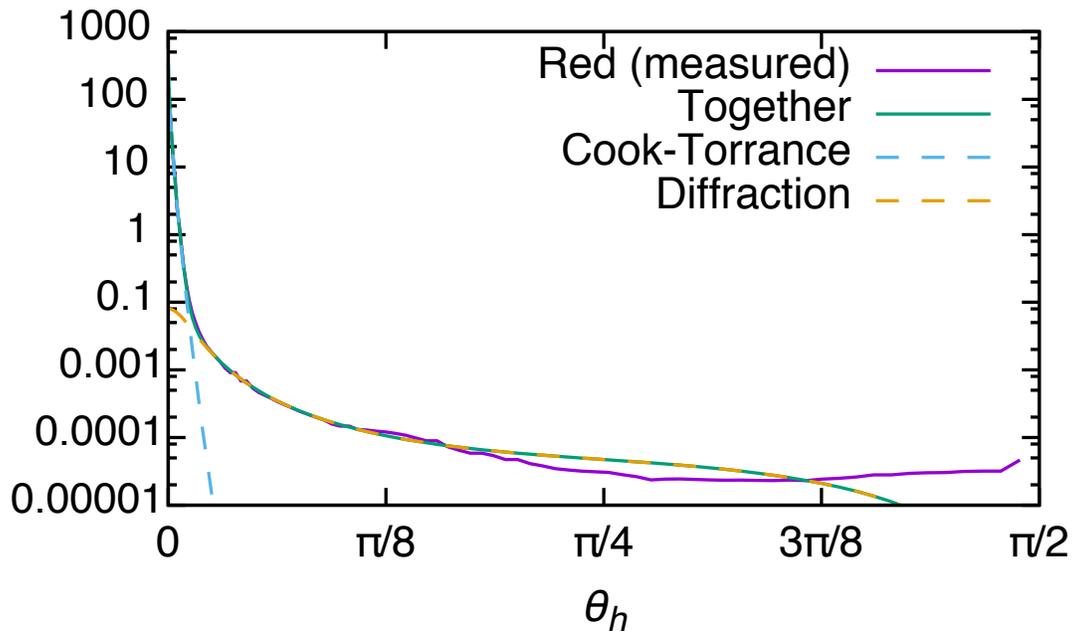


Figure 11: Measured reflectance data for tungsten-carbide as a function of θ_h and our fitting. This reflectance has two separate behaviours: a sharp peak and a softer tail. Our model provides a simple explanation: these correspond to different physical phenomena, micro-facet reflection and diffraction.

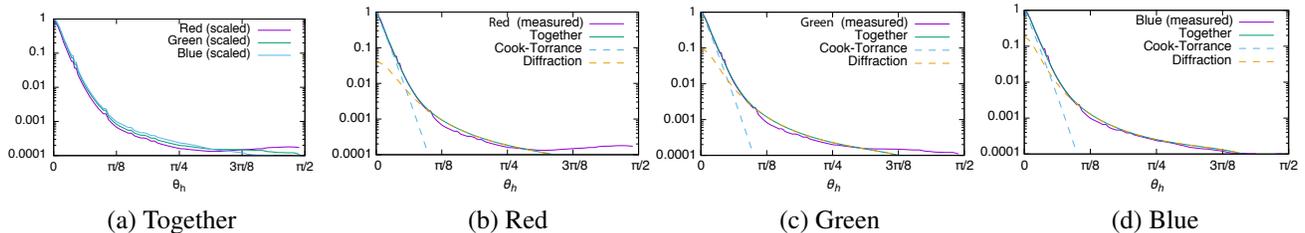


Figure 12: (a) Measured reflectance data for nickel as a function of θ_h , renormalized by the value at the origin. The width of the BRDF lobe changes with the wavelength, in contradiction with the micro-facet theory. (b) (c) and (d): The presence of a diffraction lobe, combined with the reflection, explains this variation of lobe width. The diffraction lobe intensity increases as the wavelength decreases (from red to blue).

5.5 Importance of the Diffraction Component

5.5.1 Material Features

We visualise material response by plotting BRDF values over backscattering data (where $\mathbf{i} = \mathbf{o}$), as a function of θ_h . This provides a good visual summary of each BRDF. We visualize this data using a logarithmic scale for the y axis, as BRDF values typically span several orders of magnitude (from 10^3 to 10^{-4} for tungsten-carbide).

For smooth materials, such as tungsten-carbide (see Figure 11), we see two distinct behaviours: a sharp peak near the origin, a smoother tail after the peak. Our model explains these two behaviours as two different physical phenomena: reflection for the peak, diffraction for the tail, and computes the relevant parameters.

The visual impact of the diffraction lobe depends on the surface properties. For smooth surfaces, including metals, it explains most of the behaviour. The reflection lobe only becomes active at grazing angles. This is the observed behaviour for nickel (see Figure 1, top row), but also chrome and most smooth plastics. For rougher surfaces, the two lobes contribute together, sometimes with very different effects. That is the behaviour for alum-bronze (see Figure 1, middle row). Finally, for other materials, the impact is less visible, and most of the behaviour is well approximated with a Cook-Torrance lobe. That is the case for green-metallic-paint2 (see Figure 1, bottom row).

5.5.2 Lobe Width Variation with Wavelength

Rougher materials present a more interesting case. Measured nickel, for example, appears to have a single behaviour, with monotonic decrease (see Figure 12(a)). Renormalizing each lobe with the value at origin, we see that the lobes for the red, green and blue channels have different widths, in contradiction with the micro-facet model. Butler et al. [BNM15c] observed a similar variation in lobe width in other measured reflectance data.

Our model explains this variation of the lobe width with wavelength by the presence of a large diffraction lobe, combined with the micro-facet reflection lobe (see Figure 12(b), (c) and (d)). Summing these two lobes gives the appearance of a single lobe with varying width.

5.5.3 Efficient Fitting

The presence of the diffraction component also helps when fitting measured materials to the model. When we fit material parameters without a diffraction component, we get two possible behaviours: either the fitting process ignores wide angle scattering effects and focuses on the specular peak, or it tries to fit the wide angle scattering effects with the micro-facet model, with unsatisfying results (see Figure 13).

We used here our fitting algorithm, working over the entire data ($\theta_i, \theta_o, \phi_o - \phi_i$). Other fitting approaches [AP07, BSH12, DH1*15] avoided the issue by explicitly fitting only over θ_h with the back-scattering data. By construction, this gives a micro-facet model that approximates the BRDF well, but it does not take into account data outside of the incident plane.

5.5.4 Behavior at Grazing Angles

Several papers [BSH12, NDM05, Bur12] have noted the discrepancy between the behaviour of measured materials and the behaviour predicted by the Cook-Torrance model at grazing angles. Two components of the Cook-Torrance model rise at grazing angles: the Fresnel term $F(\theta_d)$ and the $1/(\cos \theta_i \cos \theta_o)$ term. For rough surfaces, the latter is partially counterbalanced by the shadowing and masking term G . For smooth surfaces, we have $G \approx 1$ for almost all directions. Figure 14 shows the

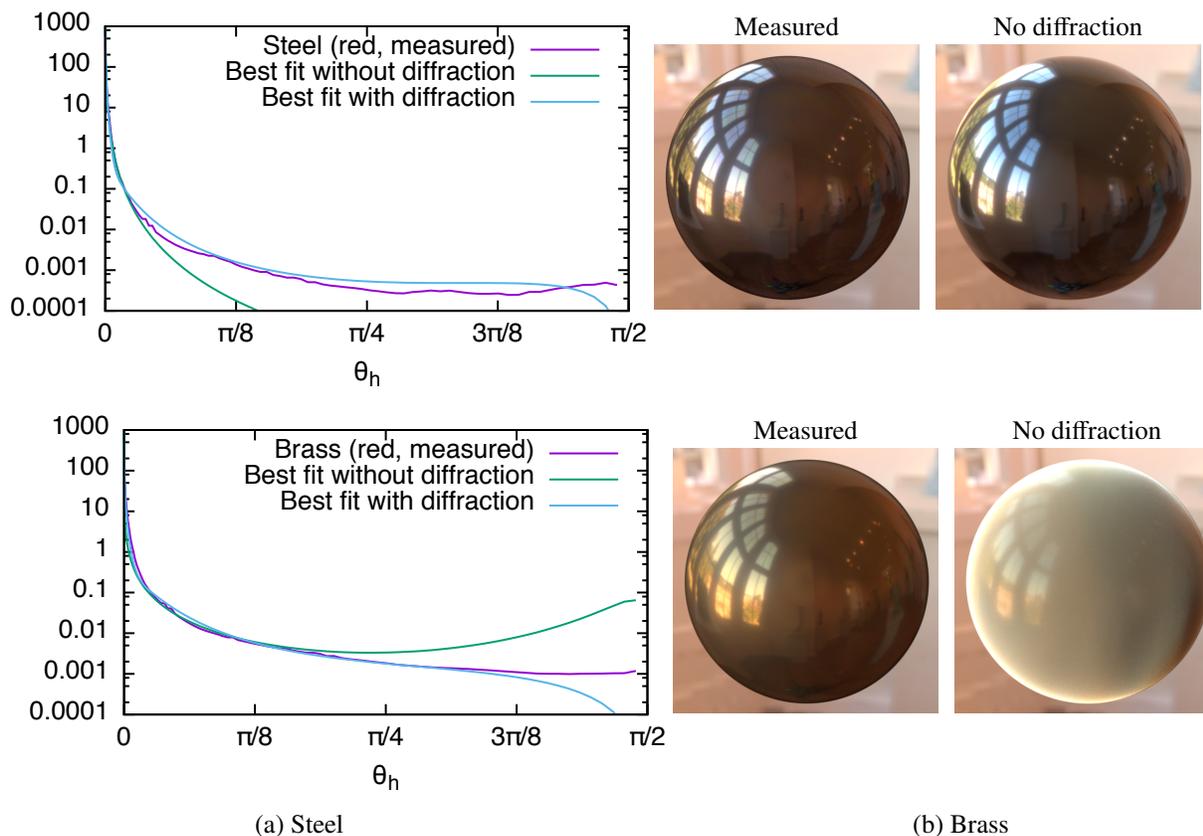


Figure 13: We tried fitting measured materials without a diffraction lobe. For steel, left, the fitting process ignores wide-angle scattering, but provides satisfying results. For brass, right, the fitting process tries approximating wide-angle scattering with a larger micro-facet distribution, resulting in unsatisfying results.

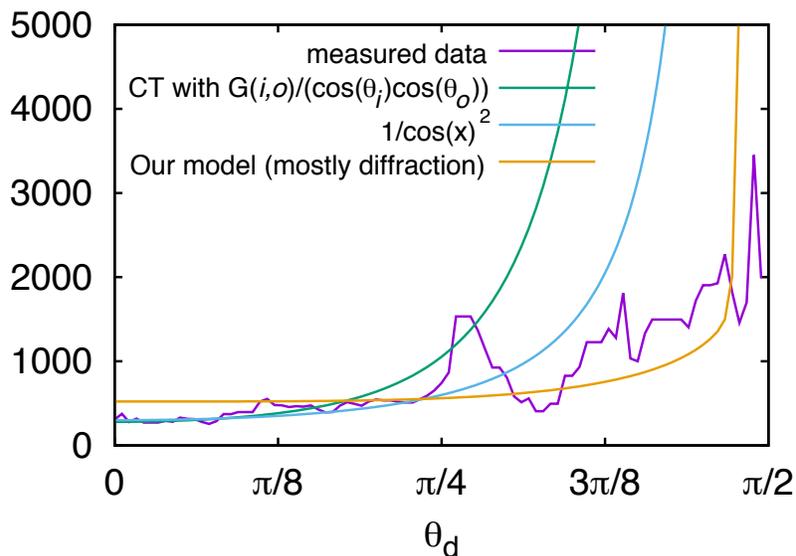


Figure 14: BRDF behaviour as a function of θ_d . The Cook-Torrance approximation rises quickly for $\theta_d > \pi/4$, mainly because of the $1/(\cos \theta_i \cos \theta_o)$ term. Our model, using diffraction, rises much later.

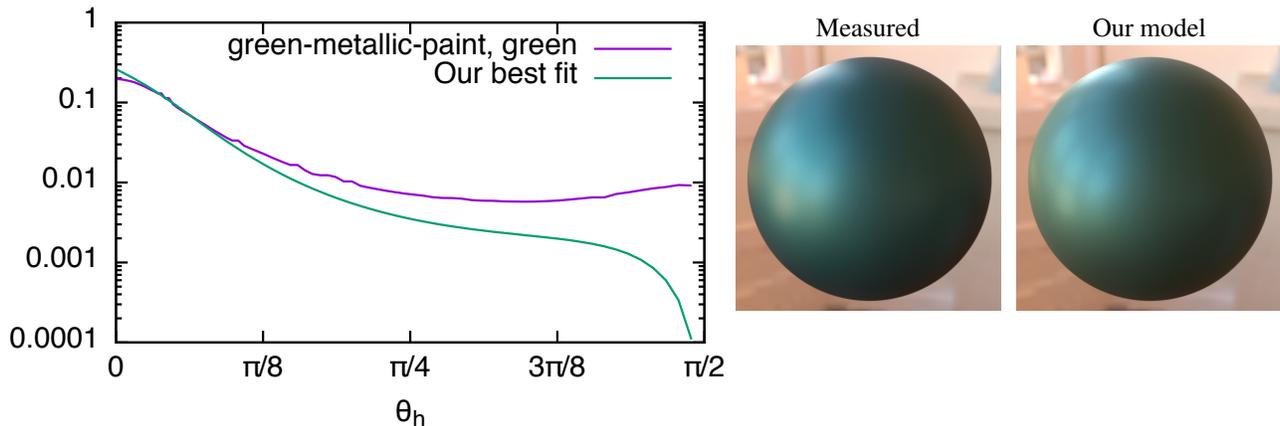


Figure 15: Our model is limited to a single interaction between the light and the material. It has problems approximating more glossy materials, such as `green-metallic-paint`, probably because it does not model multiple light bounces inside the micro-geometry.

measured data for chrome, as a function of θ_d , along with a fit using only Cook-Torrance model and a fit using our model, a combination of diffraction and Cook-Torrance. For this material, the fitting algorithm predicted that most of the energy was going into the diffraction lobe.

The Cook-Torrance approximation begins rising quickly for $\theta_d > \pi/4$. For comparison, we plot $1/\cos(\theta_d)^2$. The two curves are parallel, showing that most of the rise comes from the $1/(\cos \theta_i \cos \theta_o)$ term. The Fresnel term explains the difference between these two curves. By comparison, our model rises later, around $\theta_d = 1.4$. For most of the interval, $A \approx 0$, and the diffraction lobe explains most of the behaviour. It is only for grazing angles that $A \approx 1$ and we revert to the Cook-Torrance model.

5.6 Multiple Scattering

Our model is limited to a single light interaction with the surface: both reflection and diffraction are modelled by a single bounce. This approximation is valid for smooth materials: metals, shiny plastics, polished stones, but breaks for rougher materials, such as glossy paints. Reflectance values for a glossy paint such as `green-metallic-paint` show a large lobe, followed by a plateau for $\theta_h > \pi/4$ (see Figure 15). Our model is unable to capture this behaviour, although our approximation is visually acceptable.

6 Conclusion and Future Work

We have presented a physical explanation of the different phenomena involved in material reflection. Based on the Generalized-Harvey-Shack theory, our model combines together diffraction and reflection caused by surface micro-geometry: the reflection lobe is responsible for the specular peak of the reflectance, while diffraction is responsible for wide-angle scattering. The energy repartition between the different lobes depends on the surface roughness, which is defined as the variance of surface height variations. We validated our model by comparing it with measured materials.

The relative importance of the diffraction lobe depends on surface roughness: it can be an important part of the surface aspect for semi-rough materials. Our analysis explains many observed behaviours of measured reflectance, such as the difference between the peak and the tail, and the size of the lobe changing with the wavelength.

This new explanation of the influence of surface micro-geometry on material reflectance opens the way for new research on material acquisition and representation. In future work, we want to exploit the knowledge provided by this new model to speed-up material acquisition, possibly using polarized light.

Our model is only based on physical parameters: the index of refraction of the materials, the albedo of participating media. These parameters are useful for physical explanation of surface properties, but not practical for material editing by artists. In the future, we would like to connect these parameters to more convenient parameters for artists, such as color or roughness.

Finally, We have focused on a single light interaction with the surface. This is consistent for relatively smooth surfaces, but

rougher materials break this assumption. For these materials, we would like to model multiple interactions with the surface, as in [JDJM14, HHdD16] but taking into account the color change caused by diffraction.

References

- [AP07] ASHIKHMINE M., PREMOŽE S.: Distribution-based BRDFs. University of Utah., 2007. <http://www.cs.utah.edu/~premoze/dbrdf/>.
- [BLPW14] BRADY A., LAWRENCE J., PEERS P., WEIMER W.: genBRDF: Discovering new analytic BRDFs with genetic programming. *ACM Trans. Graph.* 33, 4 (July 2014), 114:1–114:11.
- [BNM15a] BUTLER S. D., NAUYOKS S. E., MARCINIAK M. A.: Comparison of microfacet BRDF model elements to diffraction BRDF model elements. *Proc. SPIE 9472, Algorithms and Technologies for Multispectral, Hyperspectral, and Ultraspectral Imagery XXI* (2015), 94720C–94720C–11.
- [BNM15b] BUTLER S. D., NAUYOKS S. E., MARCINIAK M. A.: Experimental analysis of bidirectional reflectance distribution function cross section conversion term in direction cosine space. *Opt. Lett.* 40, 11 (Jun 2015), 2445–2448.
- [BNM15c] BUTLER S. D., NAUYOKS S. E., MARCINIAK M. A.: Experimental measurement and analysis of wavelength-dependent properties of the BRDF. *Proc. SPIE 9611, Imaging Spectrometry XX* (2015), 96110G–96110G–15.
- [BS87] BECKMANN P., SPIZZICHINO A.: *The scattering of electromagnetic waves from rough surfaces*. Artech House, 1987.
- [BSH12] BAGHER MAHDI M., SOLER C., HOLZSCHUCH N.: Accurate fitting of measured reflectances using a Shifted Gamma micro-facet distribution. *Computer Graphics Forum* 31, 4 (June 2012).
- [Bur12] BURLEY B.: Physically-based shading at Disney. In *Siggraph course: Practical Physically Based Shading in Film and Game Production*, Hill S., McAuley S., (Eds.). ACM, 2012.
- [CT82] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. *ACM Trans. Graph.* 1, 1 (1982), 7–24.
- [DHI*15] DUPUY J., HEITZ E., IEHL J.-C., POULIN P., OSTROMOUKHOV V.: Extracting Microfacet-based BRDF Parameters from Arbitrary Materials with Power Iterations. *Computer Graphics Forum* (2015), 10.
- [DI11] D’EON E., IRVING G.: A quantized-diffusion model for rendering translucent materials. *ACM Trans. Graph. (proc. Siggraph)* 30, 4 (July 2011), 56:1–56:14.
- [DWMG15] DONG Z., WALTER B., MARSCHNER S., GREENBERG D. P.: Predicting appearance from measured microgeometry of metal surfaces. *ACM Trans. Graph.* 35, 1 (2015), 9:1–9:13.
- [Har75] HARVEY J. E.: *Light-Scattering Characteristics of Optical Surface*. PhD thesis, University of Arizona, 1975. Adviser: R. V. SHACK. <http://www.dtic.mil/dtic/tr/fulltext/u2/a095132.pdf>.
- [Hei14] HEITZ E.: Understanding the Masking-Shading Function in Microfacet-Based BRDFs. *Journal of Computer Graphics Techniques* 3, 2 (June 2014), 32–91.
- [HHdD16] HEITZ E., HANIKA J., D’EON E., DACHSBACHER C.: Multiple-scattering microfacet BSDFs with the Smith model. *ACM Trans. Graph. (Proc. SIGGRAPH 2016)* 35, 4 (2016).
- [HP15a] HOLZSCHUCH N., PACANOWSKI R.: *A physically accurate reflectance model combining reflection and diffraction*. Research Report RR-8807, INRIA, Nov. 2015.
- [HP15b] HOLZSCHUCH N., PACANOWSKI R.: Identifying diffraction effects in measured reflectances. In *Eurographics Workshop on Material Appearance Modeling* (Darmstadt, Germany, June 2015).
- [HTSG91] HE X. D., TORRANCE K. E., SILLION F. X., GREENBERG D. P.: A comprehensive physical model for light reflection. In *Computer Graphics (ACM SIGGRAPH ’91 Proceedings)* (July 1991), vol. 25, pp. 175–186.
- [JDJM14] JAKOB W., D’EON E., JAKOB O., MARSCHNER S.: A comprehensive framework for rendering layered materials. *ACM Trans. Graph. (Proc. SIGGRAPH 2014)* 33, 4 (2014).

- [JMLH01] JENSEN H. W., MARSCHNER S., LEVOY M., HANRAHAN P.: A practical model for subsurface light transport. In *Computer Graphics Proceedings, Annual Conference Series (SIGGRAPH 2001)* (August 2001).
- [Kry06] KRYWONOS A.: *Predicting Surface Scatter using a Linear Systems Formulation of Non-Paraxial Scalar Diffraction*. PhD thesis, University of Central Florida, 2006. Adviser: J. E. HARVEY. http://etd.fcla.edu/CF/CFE0001446/Krywonos_Andrey_200612_PhD.pdf.
- [LFTG97] LAFORTUNE E. P. F., FOO S.-C., TORRANCE K. E., GREENBERG D. P.: Non-linear approximation of reflectance functions. In *SIGGRAPH '97* (1997), pp. 117–126.
- [LKYU12] LÖW J., KRONANDER J., YNNERMAN A., UNGER J.: Brdf models for accurate and efficient rendering of glossy surfaces. *ACM Trans. Graph.* 31, 1 (Feb. 2012), 9:1–9:14.
- [Lou04] LOURAKIS M. I. A.: levmar: Levenberg-marquardt nonlinear least squares algorithms in C/C++. <http://www.ics.forth.gr/~lourakis/levmar/>, July 2004.
- [MPBM03] MATUSIK W., PFISTER H., BRAND M., McMILLAN L.: A data-driven reflectance model. *ACM Trans. Graph.* 22, 3 (July 2003).
- [NDM05] NGAN A., DURAND F., MATUSIK W.: Experimental analysis of BRDF models. In *Eurographics Symposium on Rendering* (2005).
- [Rus98] RUSINKIEWICZ S.: A new change of variables for efficient brdf representation. In *Rendering Techniques '98 (Proceedings of Eurographics Rendering Workshop '98)* (1998), Drettakis G., Max N., (Eds.), Springer Wien, pp. 11–22.
- [Smi67] SMITH B.: Geometrical shadowing of a random rough surface. *IEEE Transactions on Antennas and Propagation* 15, 5 (Sept. 1967).
- [Sta99] STAM J.: Diffraction shaders. In *SIGGRAPH '99* (1999), ACM.
- [TR75] TROWBRIDGE T. S., REITZ K. P.: Average irregularity representation of a rough surface for ray reflection. *J. Opt. Soc. Am.* 65, 5 (1975), 531–536.
- [TS67] TORRANCE K. E., SPARROW E. M.: Theory for off-specular reflection from roughened surfaces. *J. Opt. Soc. Am.* 57, 9 (Sept. 1967), 1105–1112.
- [VH98] VERNOLD C. L., HARVEY J. E.: A modified beckmann-kirchhoff scattering theory. *Proc. SPIE 3426, Scattering and Surface Roughness II* (1998), 51–56.
- [WMLT07] WALTER B., MARSCHNER S., LI H., TORRANCE K. E.: Microfacet models for refraction through rough surfaces. In *Eurographics Symposium on Rendering* (2007), pp. 195–206.
- [WW07] WEIDLICH A., WILKIE A.: Arbitrarily layered micro-facet surfaces. In *GRAPHITE '07 (International conference on Computer graphics and interactive techniques in Australia and Southeast Asia)* (2007), pp. 171–178.



**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

Inovallée
655 avenue de l'Europe Montbonnot
38334 Saint Ismier Cedex

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399