



A physically accurate reflectance model combining reflection and diffraction

Nicolas Holzschuch, Romain Pacanowski

► To cite this version:

Nicolas Holzschuch, Romain Pacanowski. A physically accurate reflectance model combining reflection and diffraction. [Research Report] RR-8807, INRIA. 2015, pp.24. hal-01224702

HAL Id: hal-01224702

<https://inria.hal.science/hal-01224702>

Submitted on 5 Nov 2015

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 accurate reflectance model combining reflection and diffraction

Nicolas Holzschuch , Romain Pacanowski

**RESEARCH
REPORT**

N° 8807

November 2015

Project-Teams Maverick and
Manao



A physically accurate reflectance model combining reflection and diffraction

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

Project-Teams Maverick and Manao

Research Report n° 8807 — November 2015 — 24 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, it has a metallic or plastic appearance. The reflectance model (BRDF) is essential for photorealistic pictures. Measured reflectance provide high realism, at the expense of memory cost. Parametric models are compact, but it is difficult to find the right parameters from measured reflectance. In this research report, we show that two different physical phenomena are present in measured reflectance: reflection and diffraction. Taking both into account, we present a reflectance model that is compact and a very good approximation of measured reflectance. Designers can act on model parameters, related to surface properties, to create new materials.

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

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

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

**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

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

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 est 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

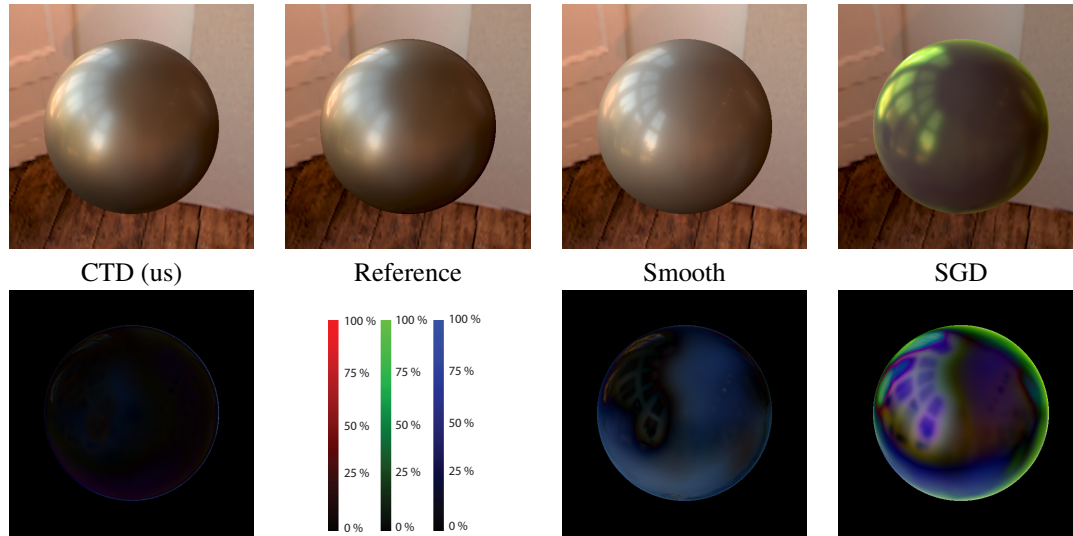


Figure 1: **gold-metallic-paint2**: Comparison between our model, reference solution using measured data and state-of-the-art solutions: the Smooth BRDF from [LKYU12] and the SGD distribution for Cook-Torrance BRDF from [BSH12]. Bottom row: symmetric mean absolute percentage error (sMAPE) between each solution and the reference. Our model is the only one to capture all reflectance effects.

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 well the appearance of the material, but have storage issues (33 MB for a single isotropic material). It is also more difficult to edit measured materials or combine them.

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 aspect 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 aspect as a given measured material. This would provide a compact representation and allow importance sampling. It will also make editing and combining materials easier.

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] assumes that light follows the principles of optical geometry: it is reflected by the surface micro-geometry but also potentially occluded. Another reflectance model [CT95] assumes that the micro-geometry diffracts the incoming light. This reflectance model has an intrinsic wavelength dependency.

Previous experiments have shown that fitting measured materials to parametric models is hard. Ngan *et al.* [NDM05] proved that diffuse and glossy materials are well approximated using the Cook-Torrance or Lafortune models. The models are less accurate for specular materials such as metals, metallic paints and shiny plastics.

Bagher *et al.* [BSH12] introduced a new micro-facet distribution for the Cook-Torrance model. It provides a very good fit with measured data, if you use different parameters for each color channel, implying that the geometry of micro-facets is different for each channel. Löw *et al.* [LKYU12] showed that

the diffraction model provides a good fit for measured data, but they removed the wavelength dependency from the model.

In summary, it seems that both the diffraction and reflection models provide a good fit for measured reflectance, if you remove one of the fundamental physical assumptions of the model. In this paper, we explain this apparent contradiction: both physical effects are present in most materials. Physically-based BRDF models have to take both effects into account.

We first present a simple test to separate between parametric domains where diffraction effects dominate, and domains where both diffraction and reflection are present (Section 4). We study each effect separately. We use this to propose a new model for reflectance, combining reflection and diffraction effects in Section 5. This model provides a very strong fit with all measured materials (Section 6). It relies on a small number of parameters related to the micro-geometry of the surface. These parameters are easy to edit for graphists and designers.

2 Previous Work

Several physically-based BRDF models connect the surface micro-geometry with the way the surface interacts with light.

The *Cook-Torrance model* [TS67, CT82] assumes that the micro-geometry is made of specular micro-facets, at a large enough scale that light follows the principles of optic geometry. The overall reflectance depend mainly on the micro-facet slope probability distribution. Cook and Torrance [CT82] used a Gaussian distribution; Trowbridge and Reitz [TR75] and Walter *et al.* [WMLT07] introduced the TR/GGX probability distribution, with a narrower peak and larger tail; it provides a better fit for some measured materials. Bagher *et al.* [BSH12] showed that a shifted Gamma distribution (SGD) provides a very good fit for all measured materials, if the parameters are allowed to vary across color channels.

The micro-facet distribution explains most of the visual aspect for the Cook-Torrance BRDF. The shadowing/masking term is essential for energy conservation, but its effects are less visible. 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 taking into account correlation between input and output directions.

The *diffraction model* is widely used in the Optical Engineering community [CT95]. It assumes that surface micro-geometry is roughly the same size as the wavelength of the incoming light, and diffracts the incoming light instead of reflecting it. It predicts the distribution of outgoing light, as a function of the Power Spectral Distribution (PSD) of surface height. The effect of the PSD depend on the wavelength: the diffraction lobe is larger for the red channel than for the blue channel, but usually with lower intensity. Löw *et al.* [LKYU12] showed that this model is consistent with the lobe shape in measured BRDFs. They also showed that the diffraction model provides a good fit with measured reflectance. To conduct the fitting, they removed the wavelength dependency from the diffraction model. Their model also does not account for energy conservation.

He *et al.* [HTSG91] derived the most comprehensive reflectance model for isotropic surfaces; it takes into account all scales in the micro-geometry in a single model.

Stam [Sta99] showed that both the diffraction model and the Cook-Torrance model can be derived from Maxwell equations, depending on the size of the surface microgeometry. His model takes into account anisotropic surface effects, and reverts to He's model for isotropic surfaces. Cuyper *et al.* [CHB*12] extended the rendering of diffraction effects by keeping trace of the phase of the reflected light, allowing for destructive interference.

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




Figure 2: Notations for BRDFs: i is the incoming direction, o is the outgoing direction, h is the half-vector. θ_h is the angle between h and the normal to the macro-surface, n .

measured data. They found the best fits for the Cook-Torrance [CT82] and Lafortune [LFTG97] models. They also reported that 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. Romeiro *et al.* [RVZ08] and Pacanowski *et al.* [PSCS*12] extend this idea by projecting measured BRDFs on a two-dimensional space (θ_h, θ_d) . They show that fitting in this space does not degrade the quality of the fitting. Pacanowski *et al.* approximate the projected measured BRDF with rational fractions, compressing BRDF data to roughly 300 Bytes.

Brady *et al.* [BLPW14] applied genetic programming to discover new analytical BRDF models from measured BRDFs. The micro-facet distribution term we propose in section 5 has some similarities with the “BRDF Model A” they introduced.

Burley [Bur12] described the reflectance model used at Disney: a two-layer model, with both layers using the Cook-Torrance model. The first layer (the base material) uses the TR/GGX distribution. The second layer represents a clearcoat layer over the base material, with a Generalized Trowbridge-Reitz distribution (GTR), with an exponent of 1. He also describes the mapping between BRDF parameters and artist-friendly parameters. In this paper, we describe a model that is close to Burley’s in terms of functionality, while having a more physical basis.

3 Scientific Background

3.1 Parameters

In the following, we consider a surface of normal n . We consider the BRDF on this surface, with incoming direction i and outgoing direction o .

We use the half-angle parametrization [Rus98] (cf. Figure 2). It depends on four parameters, $(\theta_h, \phi_h, \theta_d, \phi_d)$, where:

- h is the half-vector, between the incoming and outgoing direction: $h = (i + o) / \|i + o\|$,
- (θ_h, ϕ_h) are the spherical coordinates of the half-vector h in the frame centered on the macro-surface normal, n ,
- (θ_d, ϕ_d) are the spherical coordinates of the incoming direction i in the frame centered on h .

Table 1: Probability functions for classical micro-facet distributions, with $x = \tan^2 \theta$, with normalization constants.

	Beckmann	TR/GGX	GTR ($\gamma \neq 1$)	GTR ($\gamma = 1$)	SGD
$P_{22}(x)$	$\frac{1}{\pi\alpha^2} e^{-x/\alpha^2}$	$\frac{\alpha^2}{\pi(\alpha^2+x)^2}$	$\left(\frac{(\gamma-1)(1-\alpha^2)}{1-\alpha^{2(\gamma-1)}}\right) \frac{\alpha^{2(\gamma-1)}(1+x)^{\gamma-2}}{\pi(\alpha^2+x)^\gamma}$	$\frac{1-\alpha^2}{2\pi \log(1/\alpha)} \frac{1}{(1+x)(\alpha^2+x)}$	$\frac{\alpha^{p-1}}{\Gamma(1-p, \alpha)} \frac{e^{-\frac{\alpha^2+x}{\alpha}}}{(\alpha^2+x)^p}$

For isotropic reflectance, we can drop the dependency on ϕ_h . The entire BRDF can be expressed as a function of three parameters, $(\theta_h, \theta_d, \phi_d)$.

3.2 Cook-Torrance model

The Cook-Torrance model [CT82] assumes that the surface of the object is made of specular micro-facets, large enough that we can assume geometrical optics. The resulting BRDF model then depends mainly on the probability distribution of the micro-facets normals, $D(\theta_h)$:

$$\rho_s(\mathbf{i}, \mathbf{o}) = \frac{k_d}{\pi} + k_s F(\mathbf{i}, \mathbf{h}) \frac{G(\mathbf{i}, \mathbf{o}) D(\theta_h)}{4(\mathbf{i} \cdot \mathbf{n})(\mathbf{o} \cdot \mathbf{n})} \quad (1)$$

k_d and k_s are the diffuse and specular reflection coefficients respectively. F is the Fresnel reflection coefficient for each micro-facet, G is the masking and shadowing term, expressing how much of the incoming and outgoing light is masked by local geometry (see Figure 3). It is computed from D by a double indefinite integration [Smi67, WMLT07, Hei14].

The Cook-Torrance BRDF model is generic, and can be used with any probability distributions D :

- Cook and Torrance [CT82] used a Beckmann distribution.
- Trowbridge and Reitz [TR75] and Walter *et al.* [WMLT07] used the TR/GGX distribution. It has a sharper peak and larger tail. It provides a better fit to measured data [WMLT07].
- Burley [Bur12] used a generalization of the TR/GGX distribution, the GTR distribution.
- Bagher *et al.* [BSH12] used a Shifted Gamma Distribution.

D is usually expressed as:

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




Figure 3: Shadowing/masking in micro-facet models: part of the incoming light is blocked by other micro-facets (shadowing). Part of the reflected light is also blocked (masking).

where $P_{22}(x)$ is a positive function of the variable $x \in [0, \infty)$ and $\chi_{[0, \frac{\pi}{2}]}(\theta)$ ensures sidedness: it is equal to 1 if $\theta < \frac{\pi}{2}$, and 0 otherwise. This expression simplifies computations and normalization, by allowing a change of variables $x = \tan^2 \theta$. With this expression,

$$D(\theta) \cos \theta d\omega = \frac{1}{2} P_{22}(\tan^2 \theta) d(\tan^2 \theta) d\phi \quad (3)$$

Since $\int D(\theta) (\omega \cdot \mathbf{n}) d\omega = 1$, we have $\int_0^\infty P_{22}(x) dx = 1/\pi$.

Table 1 provides the P_{22} functions corresponding to these distributions. We provide normalization factors for the GTR distributions; these are not currently available in the literature.

To compute the Smith shadowing term G , we need to compute two successive indefinite integrals of D :

$$P_2(r) = \int_{-\infty}^{\infty} P_{22}(r^2 + q^2) dq \quad (4)$$

$$\Lambda(\theta) = \int_{1/\tan \theta}^{\infty} (r \tan \theta - 1) P_2(r) dr \quad (5)$$

We then express the shadowing term G from Λ ; usually [Smi67]:

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

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

More accurate expressions for G taking into account correlation between \mathbf{i} and \mathbf{o} can be found in [Hei14].

The Fresnel term F expresses how specular reflection becomes more important at grazing angles. For dielectrics and conductors, it can be expressed as a function of the indices of refraction. In Computer Graphics, it is usually expressed using Schlick approximation [Sch94]:

$$F(\theta) = F_0 + (1 - F_0)(1 - \cos \theta)^5 \quad (8)$$

The Cook-Torrance model is widely used in the Computer Graphics community. It has few parameters that are related to surface properties, and is easy to use with importance-sampling in with Monte-Carlo integration.

3.3 Diffraction-based model

For the diffraction reflectance model [CT95, LKYU12], the dimensions of the surface micro-geometry are close the incoming light wavelength. Diffraction effects are dominant. They depend on the Power Spectral Distribution (PSD) S_z of the height distribution on the surface:

$$\rho_w(\mathbf{i}, \mathbf{o}, \mathbf{n}) = \frac{k_d}{\pi} + \frac{k_w}{\lambda^4} F(\mathbf{i}, \mathbf{h}) S_z(f) \quad (9)$$

$$f = \frac{1}{\lambda} (\|(\mathbf{i} + \mathbf{o}) \times \mathbf{n}\|) \quad (10)$$

f encodes the wavelength dependency; it is the norm of the projection of the deviation vector \mathbf{d} : the difference between outgoing direction \mathbf{o} and reflected incoming direction $R(\mathbf{i})$, divided by the wavelength λ (see Figure 4). F is the Fresnel term (see Equation 8).

With the half-angle parameterization, f can also be expressed as:

$$f = \frac{2}{\lambda} \sin \theta_h \cos \theta_d \quad (11)$$




Figure 4: Diffraction-based models are parameterized by f , proportional to the norm of projection \mathbf{d}_p of the deviation vector \mathbf{d} , difference between outgoing direction \mathbf{o} and reflected incoming direction $R(\mathbf{i})$.

S_z is usually approximated using the ABC model:

$$S_z(f) = \frac{a}{(1 + b^2 f^2)^{\frac{c+1}{2}}} \quad (12)$$

where the parameters $(a, b, c) \in \mathbb{R}^3$.

The diffraction model is widely used in the optical engineering community. It provides a convenient way to predict surface properties from measured reflectance.

The diffraction model contains intrinsic dependency on the wavelength, both in the width and in the amplitude of the lobe.

4 Analyzing Measured Reflectance

4.1 Reducing the number of parameters

Reflectance models are a combination of several effects: distribution, shadowing, masking, Fresnel term, which makes more difficult to separate and analyse their relative contribution. As in previous studies, we observe that both the shadowing and Fresnel terms are constant on a large area of the parametrization domain. The shadowing term $G_1(\theta)$ is close to 1 except for grazing angles or very rough materials. The Fresnel term $F(\theta_d)$ (see Equation 8) is approximately equal to F_0 for non-grazing angles ($\theta_d < \pi/4$).

Thus, on a large subset of parameter space, both BRDF models depend mainly on two parameters: (θ_h, θ_d) . In this paper, we use this property for exploration and local classification of measured BRDFs.

4.2 Identifying diffraction effects using derivatives

We have two competing physical models to explain material reflectance: geometrical optics (Cook-Torrance) and diffraction. We also have access to a large database of measured reflectance [MPBM03]. We want to identify whether one model is predominant for certain areas in the parametric domain, for these measured reflectance.

If the Fresnel term is locally approximated by a constant, the diffraction reflectance model is mainly a function of f , plus a diffuse term:

$$\rho(\mathbf{i}, \mathbf{o}) \approx \frac{k_d}{\pi} + \frac{k_w}{\lambda^4} S_z(f) \quad (13)$$




Figure 5: Displaying the values of $\det(\nabla\rho, \nabla f)$ for representative BRDFs, in (θ_h, θ_d) space. z scale: BRDF intensity (in log space). Color scale: values of $\det(\nabla\rho, \nabla f)$. Diffraction effects are dominant for areas colored in blue and purple. For specular BRDFs (e.g. hematite, ss440, green-plastic) it is easy to separate this effect from other components. For more diffuse materials (e.g. colonial-maple-223, pearl-paint, red-fabric), the separation is less obvious.

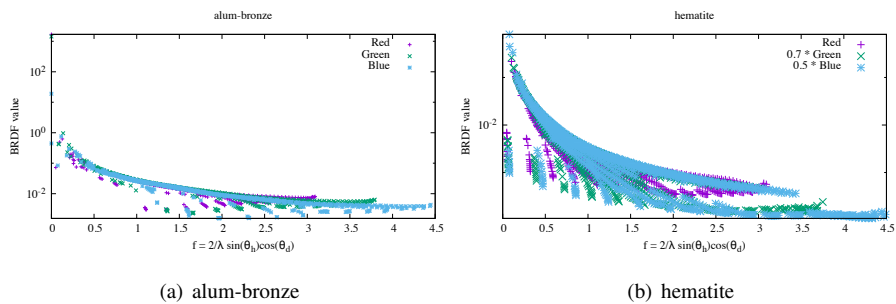


Figure 6: BRDF values for sample points where diffraction effects are likely to be dominant ($\det(\nabla\rho, \nabla f) < 0.01$). Observe how f is a good parameter across all channels: sample points superimpose over each other.




Figure 7: BRDF values for sample points where diffraction effects are unlikely to be dominant ($\det(\nabla\rho, \nabla f) > 0.01$), with $\theta_d < 0.3$ to limit the influence of the Fresnel term. For these points, θ_h is a good parameter for all channels.

If this is true, then the derivatives of ρ and f should be colinear:

$$\nabla\rho \approx \frac{k_w}{\lambda^4} \frac{dS_z}{df} \nabla f \quad (14)$$

We test vector colinearity by computing the determinant of these 2D vectors:

$$\det(\nabla\rho, \nabla f) = 0 \quad (15)$$

Equation 15 provides a necessary condition that must be verified if the reflectance properties depend on diffraction effects. We use it for an analysis of measured reflectance data:

- first, we average measured data over ϕ_d to get a 2D function over (θ_h, θ_d) . We also compute and store the variance over ϕ_d and the number of sample points.
- We then compute the partial derivatives of the reflectance over θ_h and θ_d .
- The partial derivatives of f are computed analytically, using Equation 11.
- Having these two derivatives, we compute $\det(\nabla\rho, \nabla f)$ and visualize it using a color ramp in Figure 5. In this figure, the height of the surface corresponds to the BRDF intensity, while its color depicts the value of $\det(\nabla\rho, \nabla f)$. Areas in blue and purple correspond to $\det(\nabla\rho, \nabla f) < 10^{-2}$, where Equation 15 is approximately true, and diffraction effects can be considered to be dominant.

Looking at the representative data in Figure 5, several things appear clearly:

- The separation between diffraction effects and other effects is clearly visible for specular-type BRDFs, such as metals and shiny plastics. For more diffuse BRDFs, such as woods, paints and fabrics, the separation is less marked.
- Diffraction effects, when present, correspond to wide-angle scattering, and to lower values of the BRDF. This wide-angle scattering plays an important role in the visual aspect of the BRDF. It provides the color of the material outside of the specular peak.
- The specular peak, around $\theta_h = 0$, does not appear to correspond to diffraction effects.

4.3 Comparison between channels

Our condition on the derivatives (Equation 15) is only a necessary condition. We combine it with a second test, by plotting BRDF values as a function of f . According to Equation 13, if diffraction effects dominate, BRDF sampled values should be mainly functions of f , up to multiplicative and additive constants. Figure 6 shows this assumption is valid, for points where $\det(\nabla\rho, \nabla f) < 10^{-2}$: observe how data points for all three channels are overlaid on top of each other.

Figure 7 displays, on the contrary, data points where Equation 15 does *not* hold: $\det(\nabla\rho, \nabla f) > 10^{-2}$. To avoid parasite variations due to the Fresnel term, we also limit $\theta_d < 0.3$. This time, θ_h appears to be the main parameter. Data points as a function of θ_h are overlaid on top of each other.

4.4 Combining both models

Our study shows that measured reflectance behave as if they were a combination of two separate phenomena:

- Diffraction for wide-angle scattering,
- Geometric Optic reflection for the specular peak.

The former is well described by the diffraction model (Section 3.3), the latter by the Cook-Torrance model (Section 3.2). For fitting and analysis, we can identify parameter domains where diffraction effects are dominant by testing the partial derivatives.

5 Our Model: Cook-Torrance and Diffraction

Based on our observation of measured data, we propose that material reflectance are a sum of three lobes: a diffuse lobe, a diffraction lobe and a reflection lobe. Their effects are combined in the full BRDF.

Most basis functions used so far in analytical BRDF models, such as ABC and SGD, are heuristic functions, designed to fit the data. This explains the number of parameters, since they have to fit both the behaviour near the specular peak and the behaviour at wide-angle scattering.

Knowing that reflectance is a combination of two different physical phenomena allows us to design separate basis functions for each of them. In this section, we study measured data for diffraction and reflection, and present new simpler basis functions. For each lobe, we provide energy conservation and importance sampling.

5.1 Diffraction lobe

The measured reflectance where diffraction effects dominate are slowly varying, with a single large lobe (see Figures 6). We found that in the general case, the diffraction lobe is well approximated by a simple function:

$$S_z(f) = \frac{1}{\left(1 + \frac{f^2}{a^2}\right)^2} \quad (16)$$

This function provides a good approximation for measured data for all materials, where diffraction effects dominate. (see Figure 8 for an example).

In terms of energy conservation, the maximum energy of this BRDF lobe is $E_{a,\lambda} = \pi \left(1 + \frac{1}{a^2\lambda^2}\right)$. Division by $E_{a,\lambda}$ ensures energy conservation for the diffraction lobe. Our full diffraction lobe is then:

$$\rho_w(\mathbf{i}, \mathbf{o}, \lambda) = \frac{k_w(\lambda)}{\lambda^4} \frac{1}{E_{a,\lambda}} S_z(f) \quad (17)$$




Figure 8: Focusing on areas where diffraction is the dominant phenomenon, we can approximate it with a simpler function. Example with the brass material.

Importance sampling this diffraction lobe is difficult because it is a function of f , so importance sampling will give values for f . But some values of f are not valid, depending on incoming direction \mathbf{i} . Instead, we observe that the diffraction lobe is mainly low frequency (see Figure 11), and sample it as a diffuse lobe.

5.2 Reflection lobe (Cook-Torrance)

The reflection, or Cook-Torrance, lobe is the remainder of the reflectance data: measured values minus what is predicted by the diffraction and diffuse lobes. It is a function of θ_h , and its value drops sharply as θ_h increases (see Figure 7). We found that the following function provides a good approximation:

$$P_{22}(x) = \beta e^{-\left(\frac{x}{\sigma}\right)^p} \quad (18)$$

(See Figure 9 for an example). This distribution has two parameters: the first, σ , is a measure of the width of the peak; the second, p governs the kurtosis of the distribution, whether the peak at the origin drops sharply or not. Our distribution is similar to the “BRDF model A” obtained from measured data by [BLPW14], but we keep the shadowing and Fresnel term from the Cook-Torrance model.

β is the normalization constant. To ensure energy conservation, it must be such that $\int_0^\infty P_{22}(x) dx = 1/\pi$. Thus:

$$\beta = \frac{p}{\pi \sigma^2 \Gamma(1/p)} \quad (19)$$

where Γ is the gamma function.

The shadowing term associated to this probability distribution must be computed numerically:

$$P_2(r) = \frac{p}{\pi \Gamma(1/p)} \int_0^\infty e^{-(r^2+q^2)^p} dq \quad (20)$$

$$\Lambda(\sigma \tan \theta) = \int_{1/\sigma \tan \theta}^\infty (r \sigma \tan \theta - 1) P_2(r) dr \quad (21)$$

We observe that Λ , and therefore G_1 , can be expressed as a function of $(\sigma \tan \theta)$ and p . We precompute G_1 and store it in a 2D array, to be accessed at render time. We then compute the shadowing term G using Equation 7.

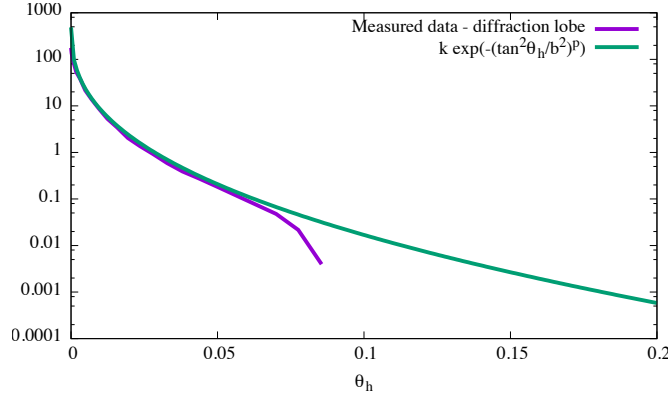


Figure 9: Focusing on areas where reflection is the dominant phenomenon, removing the diffraction and diffuse components, we can approximate reflectance with a simple function. Example with the brass material.

Importance sampling this new distribution is done using the classical method: 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) :

$$\phi_h = 2\pi u_1 \quad (22)$$

$$\theta_h = \arctan\left(\sigma\left(\gamma_u^{-1}(1/p, u_2)\right)^{\frac{p}{2}}\right) \quad (23)$$

where γ_u^{-1} is the inverse of the normalized upper incomplete gamma function.

5.3 Fresnel term and shadowing

The main parameter of the diffraction lobe, f , is proportional to $\sin \theta_h \cos \theta_d$. The diffraction lobe should behave symmetrically if we exchange θ_h with $\pi/2 - \theta_d$. Specifically, the maximum we observe for $\theta_h = 0$ should correspond to a similar maximum for $\theta_d = \pi/2$. This is not what we observe on measured data (see Figure 5). On the contrary, BRDF values go towards 0 as θ_d approaches $\pi/2$.

The easiest way to reconcile theory with observed data is by using the shadowing and masking term from the reflection lobe: micro-facets blocking light for reflection are likely to also block it for diffraction.

The Fresnel term is responsible for a sharp climb in BRDF values when θ_d approaches $\pi/2$. Its effects are visible on the specular lobe in Figure 5. We do not observe the same effects on the diffraction lobe. In our model, we apply the Fresnel term only to the reflection lobe.

5.4 Full reflectance model

Our full reflectance model is:

$$\rho(i, o) = \frac{k_d}{\pi} + G(i, o) \left(k_s \frac{D(\theta_h) F(\theta_d)}{4 \cos \theta_i \cos \theta_o} + \frac{k_w}{\lambda^4} \frac{S_z(f)}{E_{a,\lambda}} \right) \quad (24)$$

$$S_z(f) = \frac{1}{\left(1 + \frac{f^2}{a^2}\right)^2} \quad (25)$$



(a) Car model (Path tracing with adaptive sampling, (b) Kitchen model (BDPT, 4096 samples per pixel), with 128 to 16384 samples per pixel), with smoother colonial-maple-223 cupboards, chrome tap, sink and red-metallic-paint body, chrome wheels, handles, nickel kitchen wall, alumina-oxide oven door, white-marble floor. Inside the car: pickled-oak-260, white-marble and black-obsidian tiles, brass bowls, specular-white-phenolic aluminium glasses.

Figure 10: Example scenes using our model inside a global illumination renderer.

$$D(\theta) = \frac{\chi_{[0, \frac{\pi}{2})}(\theta)}{\cos^4 \theta} \frac{p}{\pi \sigma^2 \Gamma(1/p)} e^{-\left(\frac{\tan^2 \theta}{\sigma^2}\right)^p} \quad (26)$$

$$F(\theta) = F_0 + (1 - F_0)(1 - \cos \theta)^5 \quad (27)$$

It has three parameters related to the micro-geometry of the surface: (σ, p) corresponding to the distribution of micro-facet normals and a corresponding to the spatial frequency of height perturbations. It also has three color parameters: the diffuse color k_d , the specular color k_s and the diffraction color k_w . Finally, the Fresnel term is also wavelength-dependent. Using RGB color representation, the total number of parameters is 15: 3 for surface geometry, 9 for material colors and 3 for the Fresnel term.

5.5 Implementation details

We implemented our model inside two different renderers: the Mitsuba renderer [Jak10] and the Mallia Rendering Framework. Source code for our model in the Mitsuba renderer is available as supplemental material for this paper.

Figure 10 shows two examples of complex scenes, where all opaque materials are defined with our model.

Wavelength In all our computations we used the following wavelengths for the red, green and blue channels: $\{0.645, 0.526, 0.444\} \mu\text{m}$. We expressed the wavelengths in μm for better numerical stability.

Shadowing and masking term We precomputed the G_1 shadowing term and stored it in a table. Data in this table is indexed by:

$$\left(\frac{5}{p}, e^{-\left(\frac{1}{\sigma \tan \theta}\right)^{\frac{1}{20}}}\right) \quad (28)$$

We found that this parameterization encodes most of the variations in the shadowing term. The resulting data can be linearly interpolated with great accuracy. At runtime, we compute index values from $(p, \sigma \tan \theta)$ and interpolate bi-linearly between precomputed samples.




Figure 11: Contribution from each component of our model for the gold-metallic-paint2 material: diffuse, diffraction and Cook-Torrance reflection. Although the diffraction component is mostly low-frequency, its contribution to the overall aspect of the material is important.

Multi-lobe importance sampling To importance sample the combined reflectance model, we follow the algorithm for multiple lobe BRDFs described by Pharr and Humphreys [PH04]:

- compute overall energy for the BRDF, $k_T = k_s + k_w + k_d$ (for energy conservation, we must have $k_T < 1$)
- pick a random number u_3 in $[0, 1]$,
- select which reflectance lobe to sample:
 - if $u_3 k_T < k_s$: sample Cook-Torrance lobe,
 - if $k_s < u_3 k_T < k_s + k_w$: sample diffraction lobe,
 - if $u_3 k_T > k_s + k_w$: sample diffuse lobe,
- importance sample selected lobe using its own method, resulting in outgoing direction,
- evaluate BRDF at computed outgoing direction,
- divide result by sum of PDFs for all lobes at outgoing direction.

5.6 Component contribution

Figure 11 shows, for the gold-metallic-paint2 material in the MERL database, the contribution from each component of our reflectance model: diffuse, diffraction lobe, and Cook-Torrance reflectance lobe, along with their combined effect.

All high-frequency effects come from the Cook-Torrance lobe; the diffraction lobe contribution is low-frequency, but contains more variations than the diffuse lobe. It contributes significantly to the overall aspect of the material.

5.7 Editing material parameters

Our reflectance model introduces a separation between two shape components: diffraction and reflection. Figure 12 shows the effects of editing each parameter independently:

- Acting on σ has a direct influence on the *width* of the specular lobe. Reducing b produces a smoother material. With $\sigma = 0.001$, we get an almost specular metal.



(a) Nickel with fitted parameters: $a = 0.91$, $\sigma = 0.027$, $p = 0.51$. (b) Squares with $\sigma = 0.001$, borders with $\sigma = 0.05$. (c) Squares with $p = 1.4$, borders with $p = 0.3$. (d) Squares with $a = 0.3$, borders with $a = 1.4$.

Figure 12: 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.

- Acting on p changes the *shape* of the specular lobe. Decreasing p gives a lobe with a wider peak, but shorter tail. Increasing p provides a longer tail with a sharper peak. Overall effects on lobe shapes are subtle.
- Reducing a makes diffraction effects more high-frequency while pulling the overall tint of the material towards blue. Increasing a makes the overall material redder, while simultaneously making diffraction effects smoother.

5.8 Fitting our model to measured data

We designed two different methods to fit the parameters of our model to measured data: the first begins by fitting the parameters of the diffraction and diffuse lobes for areas of the parameter domain where they are predominant, then fits parameters for the Cook-Torrance lobe over the difference. The second simply fits all parameters over measured data, in a single step.

The first method is highly stable, and almost completely independent from the initial values of the parameters. The second has a stronger dependency over initial values, but with enough attempts converges to a better solution.

For both fitting methods:

- We operate over the full parameter domain $(\theta_i, \theta_o, \phi_o)$ (assuming $\phi_i = 0$ since we're dealing with isotropic BRDF).
- We sample θ_i regularly, then sample the outgoing directions (θ_o, ϕ_o) non-linearly, to place more samples near the mirror reflected direction. We weight our sample points accordingly.
- We define the error function, to be minimized in the fitting process as the Mean Square Error between measured values and predicted values:

$$E = \int (\text{meas.}(\theta_i, \theta_o, \phi_o) - \text{pred.}(\theta_i, \theta_o, \phi_o))^2 \cos^2 \theta_i d\omega_o d\theta_i \quad (29)$$

- We use Levenberg-Marquardt fitting method as implemented by Lourakis [Lou04] to find the set of parameters that minimize the overall error.

- We eliminate directions that are too close to the horizon (i.e., $\theta_i \geq 1.4$, $\theta_o \geq 1.4$), where measured data is unreliable.
- We pick initial values randomly: all parameters start with random values in $[0, 1)$. Each color starts with the same value over the three channels, the other parameters are chosen independently.
- We placed a small number of constraints over the parameters:
 - All parameters must remain positive.
 - The maximum value for p is 5 during fitting, because we only have precomputed the shadowing function for $p < 5$. In practice, the largest value observed for p after fitting is 0.95.
 - The Fresnel term is constrained to remain inside the interval $[0.02, 1.2]$. The Fresnel term is harder to fit, both because of measurement imprecisions, and because the function with which we are fitting is mostly flat. Our constraints ensure that the fitting algorithm produces values that are physically plausible. Removing them reduces the error in fitting, but produces less plausible values for the parameters.

Two-step fitting In the two-step fitting method,

- first, we restrict the fitting to areas where diffraction effects are predominant (where $\det(\nabla\rho, \nabla f) < 10^{-2}$), and compute parameters for the diffuse and diffraction lobes.
- then we compute parameters for the Cook-Torrance lobe and Fresnel parameters, while keeping the parameters for the diffuse and diffraction lobes to the values computed in the first step.

This method is highly stable, and mostly independent of initial values: it converges towards the same solution for 90 % of randomly chosen initial values.

Single-step fitting Our other fitting method is simply fitting all 15 model parameters over the whole range of measured data, using a single call to Levenberg-Marquadt optimisation.

This method has a stronger dependency to initial values of the parameters. With completely random initial values, it finds the best solution, on average, 30 % of the time. With 10 runs, it converges to a solution that is better or equal to the two-step fitting method.

6 Comparison with Previous Work

6.1 Visual comparison

We have run our fitting algorithm for our model on all 100 materials in the MERL database. We compare our model with the state-of-the-art: the Cook-Torrance model using a Shifted-Gamma Distribution [BSH12] (SGD) and the diffraction model with ABC distribution (Smooth BRDF) [LKJU12].

Figures 1, 13 and 14 show side-by-side comparison between reference solution computed using measured data and these three BRDF models. For each reflectance model, we also provide the symmetric mean absolute percentage error (sMAPE) at each pixel:

$$E_{\text{sMAPE}} = \frac{|V_{\text{measured}} - V_{\text{predicted}}|}{V_{\text{measured}} + V_{\text{predicted}}} \quad (30)$$

This value is computed at each pixel for each channel (red, green and blue). The difference images provide two informations about the reflectance models:




Figure 13: **alum-bronze**: Comparison between our model, reference solution using measured data and state-of-the-art solutions: the Smooth BRDF from [LKYU12] and the SGD distribution for Cook-Torrance BRDF from [BSH12]. Bottom row: symmetric mean absolute percentage error (sMAPE) between each solution and the reference.

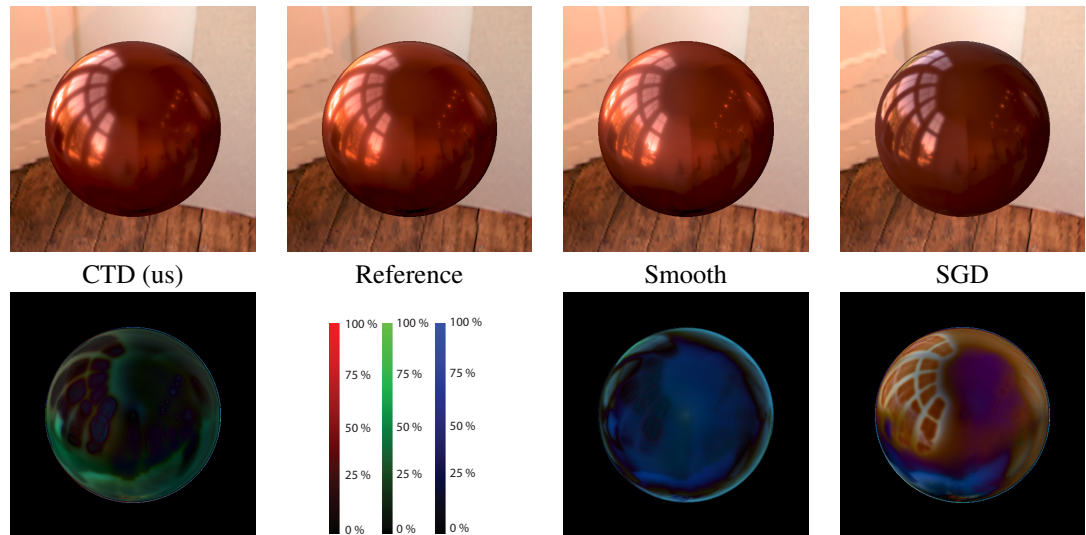


Figure 14: **red-metallic-paint**: Comparison between our model, reference solution using measured data and state-of-the-art solutions: the Smooth BRDF from [LKYU12] and the SGD distribution for Cook-Torrance BRDF from [BSH12]. Bottom row: symmetric mean absolute percentage error (sMAPE) between each solution and the reference.

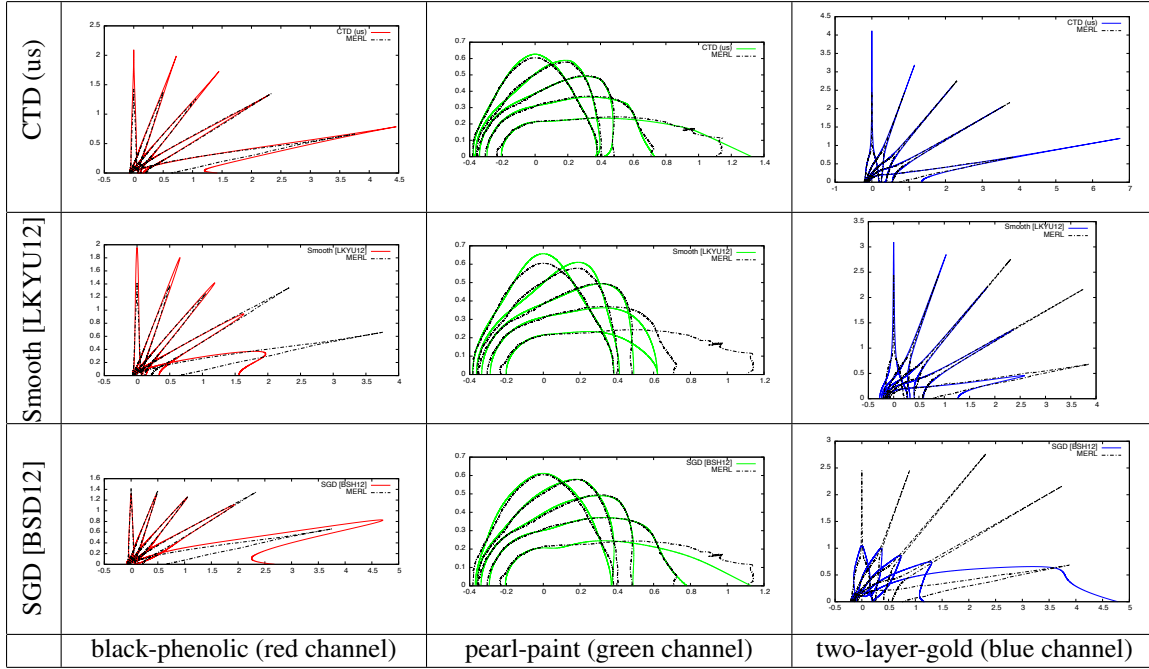


Figure 15: Reflected lobes in the incident plane for representative materials, for $\theta_i = 0, 20, 40, 60, 80^\circ$ (cubic root applied).

- the *luminance* tells the overall quality of the prediction (the darker a pixel is, the better the quality of the prediction at this pixel). A white pixel (255, 255, 255) corresponds to a sMAPE of 1, where prediction and measure are as far apart as possible.
- The *hue* tells us for which channels this prediction made an error. White or grey means the same quality for all channels. For example, in Figure 14, the smooth BRDF does a poor estimate of the blue channel at perpendicular angles, and of the red channel at grazing angles.

The supplemental material provides the same visual comparison for all 100 materials in the MERL database. All three models provide a good approximation of measured data. Our model provides a very good approximation even for materials that remained challenging for other methods, such as two-layer-gold and red-metallic-paint. Observe how we can reproduce, for example, the original material blurry specular highlights around the windows in Figures 1 and 13.

6.2 Visualizing BRDF lobes

For better understanding of BRDF behaviour, we plot the BRDF lobes for several incoming directions. We used two different methods to display BRDF lobes:

- in the incident plane ($\phi_o = -\phi_i$), for all outgoing directions, for a set of incoming directions ($\theta_i = 0, 20, 40, 60, 80^\circ$). Figure 15 shows lobe examples for a few representative materials.
- as iso-value lines for the lobes (see Figure 16). This provides a good visualization of reflectance behaviour outside of the incident plane, mainly for wide-angle scattering.

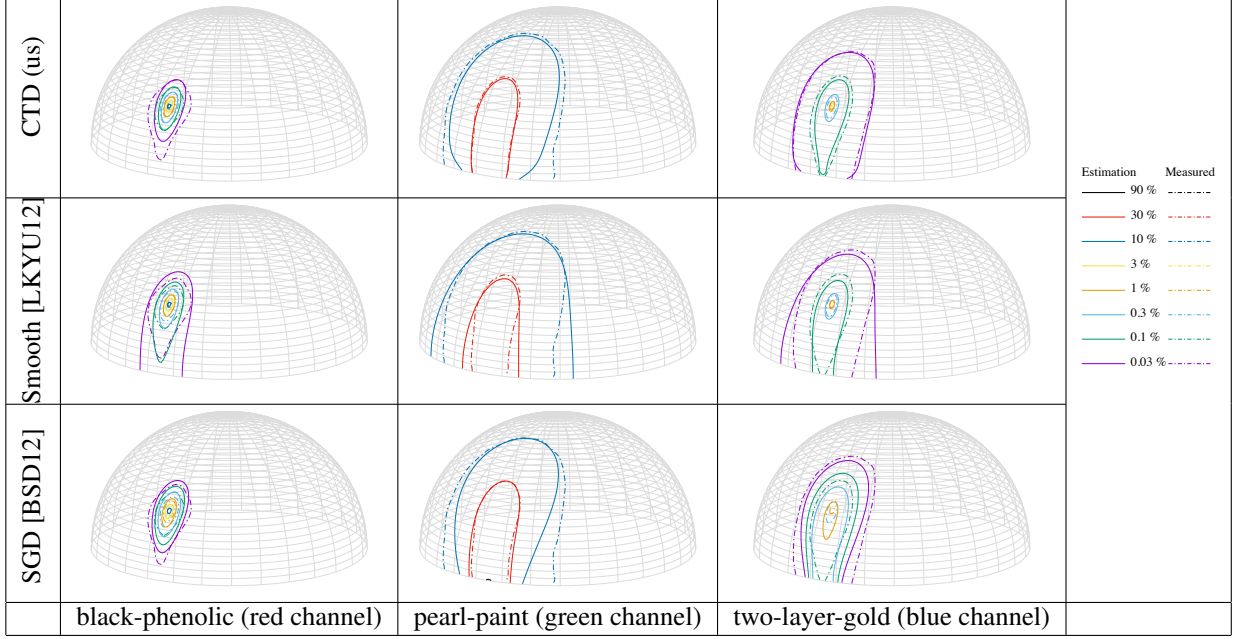


Figure 16: Iso-lines of reflected lobes, for $\theta_i = 60^\circ$, for representative materials. All isovalues are relative to the maximum intensity of the measured BRDF.

On both metrics, our reflectance model provides a very good approximation of the observed behaviour, for all measured materials. The supplemental material provides the same visualizations (incident plane and iso-lines) for all BRDFs in the MERL database.

6.3 Error measures

To add to these qualitative comparisons, we also provide quantitative measurement. For each material approximation, we compute the Root-Mean Square Error (RMSE) between measured reflectance ρ_M and estimated reflectance ρ_E :

$$\text{RMSE} = \left[\int (\rho_M(\theta_i, \theta_o, \phi_o) - \rho_E(\theta_i, \theta_o, \phi_o))^2 \cos^2 \theta_i d\omega_o d\theta_i \right]^{\frac{1}{2}} \quad (31)$$

This error measure is computed by integrating over θ_i , θ_o and ϕ_o , with samples every degree for θ_i , every quarter of a degree for θ_o and ϕ_o . We limit the integration domain to $\theta < 1.4$, to avoid grazing angles. We found that this metric provides reliable information about the quality of the fit.

Figure 17(a) provides a comparison between this error measure for our model and the SGD and Smooth models, for all materials in the MERL database (a larger version is available in the supplemental). We notice several things in this figure:

- All three models provide a good approximation to measured materials, and have similar behaviour: materials that are hard to fit for one model tend to be similarly hard to fit for the other models. For about 20 the materials, the three methods provide approximation with very similar quality (error levels within 10 % of each other).




Figure 17: Error measure (RMSE) for all materials in the MERL database, for the different BRDF models. Both figures use the same order for materials: ascending error according to our model on full data.

- There are a dozen materials where SGD performs much worse than the other two models, with error levels two or three times higher. These correspond to materials where SGD performs visibly worse as well, such as alum-bronze (see Fig. 1).
- For 89 out of 100 materials, our model provides the best quality approximation, in some cases by a significant margin. For 2 materials (white-acrylic and silver-metallic-paint), the Smooth BRDF significantly outperforms our reflectance model. We point out that these do not correspond to failure cases for our model: we provide a good approximation, the Smooth BRDF model provides a better one. For 9 other materials, the Smooth or SGD BRDF outperforms our model, but the difference is smaller than 10 %.
- As observed before [BSH12], the error is roughly proportional to the glossiness of the material, with diffuse materials at the left of the figure, and highly specular materials at the right.

We also computed the RMSE between HDR images computed using all material models and reference images computed using measured reflectance. Figure 17(b) provides these measures. For easier comparison, we kept the materials ordered as in Figure 17(a). From this figure, we observe:

- Computing the error on images provides a different ranking of materials and BRDF models. The two rankings are similar but not identical.
- For diffuse to moderately glossy materials (from 0 to 40), the SGD model performs better on images than it does on the full data. Similarly with the Smooth model for highly specular materials (above 90).
- Our model remains consistently better than the other two models. It performs better than at least one of them on all images. There are 6 glossy materials where the SGD model provides a better approximation, and 8 highly specular materials where the Smooth BRDF provides a better approximation. More importantly, actual differences are hard to spot on the images.

In summary, our combined reflectance model provides a good approximation to *all* materials in the MERL database, from very diffuse materials (fabrics, paints) to highly specular materials (chrome, steel). Larger versions of Figures 17(a) and 17(b) with material names are available in the supplemental material.

7 Conclusion and Future Work

We have presented a study of physical phenomena involved in light reflection by surfaces. By observing the behaviour of measured reflectance, we have shown that it is highly likely that both reflection and diffraction are involved in connecting the surface micro-geometry to the BRDF. Diffraction effects dominate for wide angle scattering, while reflection effects contribute more to the specular lobe.

We used this knowledge to build a reflectance model combining reflection and diffraction. The basis function to represent each phenomena alone is simpler than the functions used in previous models, such as SGD and ABC. Our new reflectance model provides a good fit to all measured materials in the MERL database, including the shape of the lobes.

Our research explains several phenomena reported in previous papers: it explains why Bagher *et al.* [BSH12] had to use channel-dependent parameters, or why Löw *et al.* [LKYU12] had to remove the wavelength-dependency of the model. It also connects to the BRDF model A found using genetic programming by Brady *et al.* [BLPW14].

Our paper opens new directions of research: although we have found that both diffraction and reflection are likely to be present in material reflectance, we have to find *how* they connect to the surface micro-geometry. It could be that there are different scales of micro-facets, with small-scale micro-facets diffracting light and larger scale micro-facets reflecting it, or it could be that the faces of micro-facets are reflecting light while their silhouettes contribute to diffraction. Answering this question will require more research, for example combining acquisition of surface reflectance and surface micro-geometry. Another possibility is to look at anisotropic materials: if both phenomena (reflection and diffraction) are connected to the same surface geometry, then they will likely have the same anisotropy. If we can identify different levels of anisotropy in the reflectance, then it is possible that two different scales of micro-facets are indeed present.

We also want to investigate diffraction effects in more complex materials, such as car paints, with embedded particles; positive interference between the particles could result in different diffraction effects, more complex than predicted by the theory.

Currently, our model describes materials using three different colors: diffuse reflectance, specular reflectance and diffraction reflectance. We have found that these are different on measured materials, but we think they should be connected, as they relate to the same material. Jakob *et al.* [JDJM14] described how diffuse reflectance can be expressed as multiple scattering in micro-facet models, connecting it to specular reflectance. We hope to express this kind of connection between diffuse, specular and diffraction.

8 Acknowledgments

The Mitsuba plugins for rendering MERL and SGD BRDFs were initially developed by Pierre Moreau. Parameters for the Löw *et al.* [LKYU12] model were kindly provided by Joel Kronander.

The Mitsuba renderer was used for Figures 10, 11 and 12 in the paper, the Mallia Rendering Framework for all the other figures as well as for the figures in the supplemental.

Errors Metrics were computed using the ALTA library [BCP*15].

All pictures in the supplemental were generated on the PLAFRIM experimental testbed, being developed under the Inria PlaFRIM development action with support from Bordeaux INP, LABRI and IMB and other entities: Conseil Régional d'Aquitaine, Université de Bordeaux and CNRS (and ANR in accordance to the programme d'investissements d'Avenir).

References

- [AP07] ASHIKHMIN M., PREMOŽE S.: Distribution-based BRDFs. University of Utah, <http://www.cs.utah.edu/~premoze/dbrdf/>, 2007.
- [BCP*15] BELCOUR L., COURTES L., PACANOWSKI R., ET AL.: ALTA: A BRDF Analysis Library. <http://alta.gforge.inria.fr/>, 2013-2015.
- [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.
- [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), 1509–1518.
- [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.
- [CHB*12] CUYPERS T., HABER T., BEKAERT P., OH S. B., RASKAR R.: Reflectance model for diffraction. *ACM Trans. Graph.* 31, 5 (Sept. 2012), 122:1–122:11.
- [CT82] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. *ACM Trans. Graph.* 1, 1 (1982), 7–24.
- [CT95] CHURCH E. L., TAKACS P. Z.: Surface scattering. In *Handbook of optics*, Bass M., (Ed.). McGraw-Hill, 1995.
- [Hei14] HEITZ E.: Understanding the Masking-Shadowing Function in Microfacet-Based BRDFs. *Journal of Computer Graphics Techniques* 3, 2 (June 2014), 32–91.
- [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.
- [Jak10] JAKOB W.: Mitsuba renderer, 2010. <http://www.mitsuba-renderer.org>.
- [JDJM14] JAKOB W., D'EON E., JAKOB O., MARSCHNER S.: A comprehensive framework for rendering layered materials. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2014)* 33, 4 (2014).
- [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/>, Jul. 2004.
- [MPBM03] MATUSIK W., PFISTER H., BRAND M., McMILLAN L.: A data-driven reflectance model. *ACM Trans. Graph.* 22, 3 (July 2003), 759–769.
- [NDM05] NGAN A., DURAND F., MATUSIK W.: Experimental analysis of BRDF models. In *Eurographics Symposium on Rendering* (2005), pp. 117–226.

- [PH04] PHARR M., HUMPHREYS G.: *Physically Based Rendering: from theory to implementation*. Morgan Kaufmann, 2004.
- [PSCS*12] PACANOWSKI R., SALAZAR-CELIS O., SCHLICK C., GRANIER X., PIERRE P., ANNIE C.: Rational BRDF. *IEEE Transactions on Visualization and Computer Graphics* 18, 11 (2012), 1824–1835.
- [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.
- [RVZ08] ROMEIRO F., VASILYEV Y., ZICKLER T.: Passive reflectometry. In *10th European Conference on Computer Vision: Part IV* (2008), pp. 859–872.
- [Sch94] SCHLICK C.: An inexpensive BRDF model for physically-based rendering. *Computer Graphics Forum* 13, 3 (1994), 233–246.
- [Smi67] SMITH B.: Geometrical shadowing of a random rough surface. *IEEE Transactions on Antennas and Propagation* 15, 5 (Sept. 1967), 668 –671.
- [Sta99] STAM J.: Diffraction shaders. In *SIGGRAPH '99* (1999), ACM, pp. 101–110.
- [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.
- [WMLT07] WALTER B., MARSCHNER S., LI H., TORRANCE K. E.: Microfacet models for refraction through rough surfaces. In *Eurographics Symposium on Rendering* (2007).



**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