Accurate fitting of measured reflectances using a Shifted Gamma micro-facet distribution

Mahdi M. Bagher, Cyril Soler, Nicolas Holzschuch

Accurate fitting of measured reflectances using a Shifted Gamma micro-facet distribution

M. M. Bagher, C. Soler and N. Holzschuch

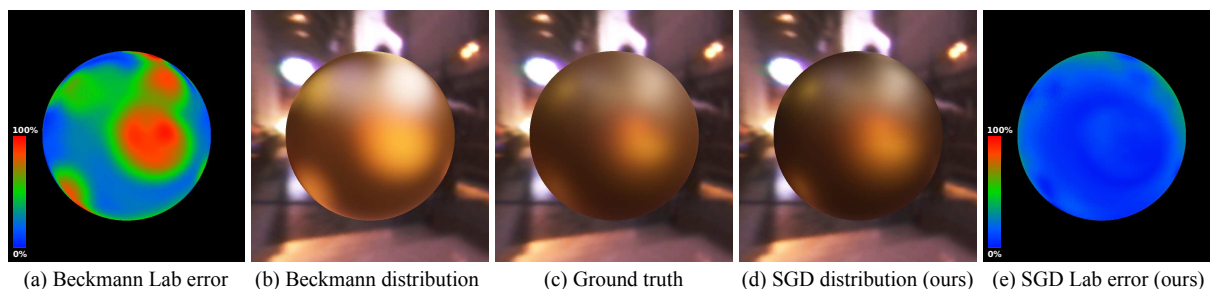Maverick, INRIA Grenoble-Rhône-Alpes and LJK (University of Grenoble and CNRS)

| (a) Beckmann Lab error | (b) Beckmann distribution | (c) Ground truth | (d) SGD distribution (ours) | (e) SGD Lab error (ours) |

Figure 1: gold-metallic-paint *material from the MERL database. Comparison between measured data (c) and Cook-Torrance BRDF using (b) Beckmann distribution or (d) SGD distribution. (a) and (e): difference measured in Lab color space between ground truth and (b) and (d) respectively.*

Abstract

Material models are essential to the production of photo-realistic images. Measured BRDFs provide accurate representation with complex visual appearance, but have larger storage cost. Analytical BRDFs such as Cook-Torrance provide a compact representation but fail to represent the effects we observe with measured appearance. Accurately fitting an analytical BRDF to measured data remains a challenging problem. In this paper we introduce the SGD micro-facet distribution for Cook-Torrance BRDF. This distribution accurately models the behavior of most materials. As a consequence, we accurately represent all measured BRDFs using a single lobe. Our fitting procedure is stable and robust, and does not require manual tweaking of the parameters.

Categories and Subject Descriptors (according to ACM CCS): http://www.acm.org/class/1998/ I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Color, shading, shadowing, and texture

1. Introduction

Photorealism in image synthesis depends strongly on the material representations we use. Measured BRDFs [MPBM03] provide a high degree of realism, but have storage issue (33 MB for a single isotropic material), making them harder to fit in a rendering pipeline. It is also difficult to use them with importance sampling, as they have no analytical definition.

BRDF models such as Phong, Lafortune or Cook-Torrance are much more compact: 8-10 parameters are enough to model an isotropic material, and are easy to combine with importance sampling in Monte-Carlo integration. We would like to compute the parameters of a BRDF model so that it photo-realistically resembles a measured material. Ngan *et al.* [NDM05] proved that diffuse and glossy materi-

als are well approximated using single lobe Cook-Torrance or Lafortune BRDFs; the models are less accurate for specular materials such as metals, metallic paints or shiny plastics.

Adding more lobes to the BRDF model makes it more accurate, at the expense of the stability of the fitting process. 3 lobes are usually considered to be the limit for automatic fitting. Highly accurate fitting for all BRDFs, ranging from diffuse to highly specular, is still impractical.

This paper presents accurate and robust fitting for any BRDF (including metals) with a single lobe (in addition to the diffuse component). We use the Cook-Torrance BRDF model but with a different micro-facet distribution. The Cook-Torrance BRDF [TS67] simulates the micro geometry of a material as specular micro-facets and derives the BRDF mathematically from the probability distribution of

micro-facets. Although it is typically used with the Beckmann distribution, the model can be used with many other distributions. Our study of measured BRDFs [MPBM03] shows that shiny materials do not follow the Beckmann distribution: micro-facets are more likely to be aligned with the surface normal, resulting in a sharper peak, consistent with the smooth appearance.

We introduce the SGD probability distribution function for micro-facets. This distribution provides a more accurate fitting for all measured materials, comparing to existing BRDF models, from highly shiny to diffuse. We use the Cook-Torrance BRDF with SGD distribution to approximate measured materials resulting in a compact and accurate representation. Since the SGD distribution approximates measured data very well, we were able to design a simpler fitting method. This method converges quickly (2.5 minutes on average) and it only uses two slices of the measured data to fit the unknown parameters.

This paper is organized as follows: in the next section, we review previous work on BRDF models and measurements. Then we present the Cook-Torrance model and micro-facet distributions. In section 4, we describe the SGD probability distribution function, and in section 5 our fitting method. Section 6 presents our results and compares them with previous work. In section 7, we conclude and discuss potential avenues for future work.

2. Previous work

Analytical BRDF models provide a reasonably realistic reflectance behavior at a very low computational and memory cost [TS67, CT82, War92, LFTG97]. Among these, micro-facet BRDF models assume that the surface geometry is made of specular micro-facets and the BRDF model depends on the facets' slope distribution.

In the Cook-Torrance BRDF model [TS67, CT82], the reflectance is decomposed into a distribution term, a shadowing term, accounting for self-masking of the micro-geometry, and a Fresnel term. For isotropic materials, the shadowing term is derived from the micro-facet distribution, through a double indefinite integration [Smi67, Bro80, BBS02]. Oren and Nayar [ON94] accounted for inter-reflections in the micro-geometry in addition to masking and shadowing. Ashikhmin *et al.* [APS00] proposed a formulation of micro-facet models that includes reciprocity and energy conservation. Weidlich and Wilkie proposed an extension of micro-facet models to multi-layer materials [WW07]. The Trowbridge-Reitz distribution [TR75] has stronger peak and larger tail than the Beckmann distribution. It is identical to the GGX distribution [WMLT07]. Low *et al.* [LKYU12] introduced a new reflectance model for glossy surfaces, with a different distribution, based on the Rayleigh-Rice theory, but do not include masking effects.

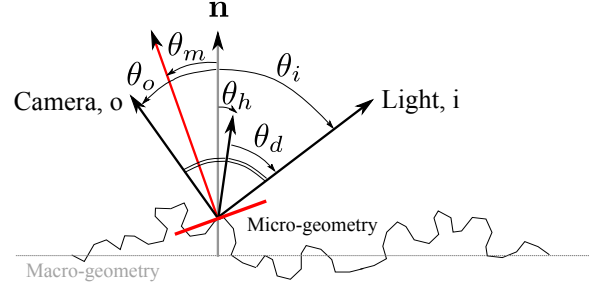The Fresnel term depends on the wavelength, although

Figure 2: *For Cook-Torrance BRDF, the material micro-geometry is a random distribution of specular micro-facets. Angles are relative to the normal* **n** *of the macro-geometry.*

this is often ignored. Schlick [Sch94] approximates the components of BRDF models with rational fractions for faster computations. His approximation of the Fresnel term is widely used today. Lazanyi *et al.* [LSK05] introduced a more accurate approximation of the Fresnel term.

Matusik *et al.* [MPBM03] measured reflectance properties for a large range of materials. We use their data-base for all our tests. Ngan *et al.* [NDM05] examined how classical BRDF models such as Lafortune [LFTG97] and Ashikhmin-Shirley [APS00] can be used to approximate measured reflectances. They showed that Cook-Torrance [CT82] provide the best fits, and specular materials are poorly approximated with a single lobe; the quality of the fitting improves with multiple lobes, but the fitting process becomes unstable. A better fitting is recommended by global optimization with branching and bounding for multi-lobe cases [YSL11].

Ashikhmin and Premoze [AP07] computed the BRDF model directly from the data using back-scattering: if the input and output directions are equal, the entire BRDF model collapses into a 1D function, representing the micro-facet distribution. This function is read from the data and stored in a compact form. Wang *et al.* [WRG*09] used spherical Gaussians to approximate existing reflectance models, including Cook-Torrance and Lafortune. The convolution of two spherical Gaussians is another spherical Gaussian, so BRDF expressed this way are easier to convolve with environment maps or normal distributions in the same model.

Romeiro *et al.* [RVZ08] and Pacanowski *et al.* [PSCS*12] project measured BRDFs on a two-dimensional space (θ_d, θ_h). Pacanowski approximates this projection with rational fractions, compressing the BRDF data down to 200 KB.

3. The Cook-Torrance reflectance model

The Cook-Torrance BRDF model [CT82] assumes that surfaces are made of random specular micro-facets (see Figure 2). The surface normal *m* follows a probability distribution function *D*. The full BRDF is:

$$\rho(\boldsymbol{i}, \boldsymbol{o}) = \frac{\rho_d}{\pi} + \frac{\rho_s}{\pi} \frac{F(\boldsymbol{i} \cdot \boldsymbol{h}) D(\theta_h) G(\boldsymbol{i}, \boldsymbol{o})}{\cos(\theta_i) \cos(\theta_o)} \qquad (1)$$

where i is the incoming direction, o is the outgoing direction, h is the half-vector ($h = i + o$, normalized). θ_i, θ_o and θ_h are the respective angles between these vectors and the surface normal n. ρ_d and ρ_s are the diffuse and specular coefficients, respectively (see notations on Figure 2).

D is the probability distribution function for micro-facet normals. G is the shadowing and masking coefficient, expressing the amount of light that is blocked by other micro-facets before reaching the point (shadowing), or after reflection (masking). For isotropic materials, G is approximated by a product of one-dimensional functions:

$$G(i,o) \approx G_1(i)G_1(o)$$

and G_1 is computed from D through a double integration [Smi67, Bro80, BBS02, WMLT07].

F is the Fresnel term, depending on the refraction index η of the material. It is usually close to constant at normal incidence, and increases to unity at grazing angles.

3.1. The normal probability density, D

Given a direction $\omega = (\theta, \phi)$ and an infinitesimal solid angle $d\omega$ centered around this direction, the probability for the normal of a micro-facet to be inside the cone $(\omega, d\omega)$ is $D(\omega)(\omega.\mathbf{n})\,d\omega$.

In order to be physically plausible, D should follow several requirements: it should be positive, the projected area of the micro-facets in a direction should be equal to the projected area of the macro-surface and in particular, $\int D(\omega)(\omega.\mathbf{n})\,d\omega = 1$. D is usually expressed as:

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

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. With this expression,

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

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

We use this expression for importance sampling [WMLT07]: for a given incoming direction i, we first pick a random micro-facet normal m with probability $D(\theta_m)\cos\theta_m$, then compute the reflected direction o so that m is the half-vector. Finally, we compute the sampling weight as the BRDF divided by the probability, multiplied by the Jacobian of o:

$$w_o = \frac{\rho(i,o)\cos\theta_o}{D(\theta_m)\cos\theta_m}\left\|\frac{\partial\omega_o}{\partial\omega_m}\right\|$$

To pick m with the probability $D(\theta_m)\cos\theta_m$, we need the associated Cumulative Distribution Function F:

$$F(\theta_m, \phi_m) = \int_0^{\phi_m}\int_0^{\theta_m} D(\theta)\cos\theta\,d\omega = \frac{\phi_m}{2}\int_0^{\tan^2\theta_m} P_{22}(u)\,du$$

	Beckmann	Exponential	TR/GGX
$P_{22}(x)$	$\frac{1}{\pi\alpha^2}e^{-x/\alpha^2}$	$\frac{1}{2\pi\alpha^2}e^{-\sqrt{x/\alpha^2}}$	$\frac{\alpha^2}{\pi(\alpha^2+x)^2}$

Table 1: *P_{22} functions corresponding to classical micro-facet distributions, with $x = \tan^2\theta$.*

F is separable: $F = F_\phi F_\theta$, with F_ϕ and F_θ from 0 to 1:

$$F_\phi(\phi_m) = \frac{\phi_m}{2\pi} \quad (4)$$

$$F_\theta(\theta_m) = \pi \int_0^{\tan^2\theta_m} P_{22}(u)\,du \quad (5)$$

Assuming we have two uniform random variables u_1 and u_2 in $[0, 1)$, we get the sampling equations:

$$\phi_m = F_\phi^{-1}(u_1) = 2\pi u_1$$

$$\theta_m = F_\theta^{-1}(u_2)$$

For any micro-facet distribution D following Equation 2, importance sampling reduces to computing the integral of P_{22} and inverting it.

3.2. Shadowing and masking

In early versions of the Cook-Torrance BRDF [TS67] the shadowing function G was independent from the micro-facet distribution. Further research [Smi67, Bro80, BBS02] proved that we get a more accurate shadowing term by integration from the micro-facet distribution D. G is approximated as the product of two one-dimensional functions: $G(i,o) \approx G_1(\theta_i)G_1(\theta_o)$. G_1 is computed from D through a double integration:

$$P_2(r) = \int_{-\infty}^{\infty} P_{22}\left(r^2 + q^2\right)dq \quad (6)$$

$$\Lambda(\theta) = \int_{\frac{1}{\tan\theta}}^{\infty} (r\tan\theta_v - 1)P_2(r)\,dr$$

$$G_1(\theta) = \frac{\chi_{[0,\frac{\pi}{2}]}(\theta)}{1 + \Lambda(\theta)}$$

For further details, please refer to the original papers or the short summary in [WMLT07].

3.3. Existing distributions

Table 1 lists the P_{22} functions corresponding to existing micro-facet distributions. Beckmann is frequently used [CT82]; it is a Gaussian in $\tan\theta$, and is also very close to the Phong distribution [WMLT07]. The exponential distribution is mostly used for the propagation of radio waves [BBS02, Bro80], and has a sharper peak and larger tails than Beckmann. The Trowbridge-Reitz/GGX distribution [TR75, WMLT07] has a sharper peak and larger tails than Beckmann and Exponential distributions.
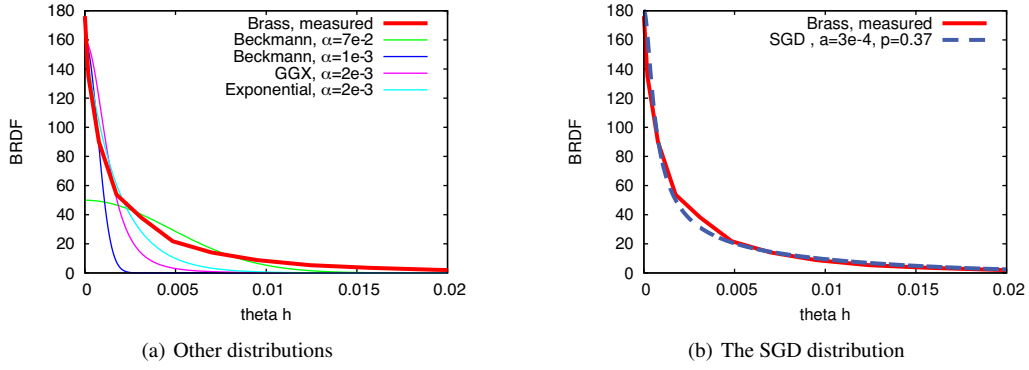
(a) Other distributions

(b) The SGD distribution

Figure 3: *The micro-facet probability distribution for brass extracted from the measured data. The SGD distribution (right) fits the data accurately both at the peak and the tail, unlike existing distributions (left)*

4. Our reflectance model

4.1. Data observation

It is hard to compare measured data with BRDF models, since models depend on position and several functions. For simple comparisons, we take $i = o$. As a consequence, $h = i$: we get an equation that only depends on θ_h, with two unknown functions D and G_1 [AP07]:

$$\rho(\theta_h) = \frac{\rho_d}{\pi} + \frac{\rho_s}{\pi} \frac{F(0)D(\theta_h)G_1(\theta_h)^2}{\cos^2(\theta_h)} \quad (7)$$

For small values of θ_h and specular materials, G_1 is almost constant, equal to 1 (see Figure 5). The measured data varies only with D, and we can observe directly the micro-facet distribution.

Figure 3 compares the measured distribution D for *brass* from the MERL database [MPBM03] with approximations using different distributions. The peak of the measured distribution is much sharper than the predictions, while the tail is larger. If we fit the tail accurately (for example with the Beckmann distribution, $\alpha = 7 \times 10^{-2}$), then we miss the peak. On the other hand, if we fit the peak well, we miss the tail. To fit this data with existing distributions we would need several lobes. This observation holds for many materials in the MERL database: metals, metallic paints and shiny plastics.

At large angles, the measured distribution seem to decrease exponentially. But for small angles, the decrease is sharper than an exponential. Based on this observation, we suggest a slope in x^{-p}, where $x = \tan^2 \theta_h$ and p is a parameter that depends on the model, multiplied by an exponential factor:

$$P_{22}(x, \alpha, p) = \frac{K_{\alpha,p} e^{-\frac{x}{\alpha}}}{(x + \alpha^2)^p} \quad (8)$$

where $K_{\alpha,p}$ is a normalization factor. With this distribution, we fit the measured data accurately (see Figure 3(b)), both at the peak and the tail.

4.2. The SGD distribution

We compute the normalization factor $K_{\alpha,p}$ so that $\int_0^\infty P_{22}(x)\,dx = 1/\pi$. The result is a shifted gamma distribution (SGD):

$$P_{22}(x) = \frac{\alpha^{p-1}}{\Gamma(1-p,\alpha)} \frac{e^{-\frac{\alpha^2+x}{\alpha}}}{(\alpha^2+x)^p}$$

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

Γ is the incomplete Gamma function: $\Gamma(s,x) = \int_x^\infty t^{s-1} e^{-t}\,dt$
The SGD fits accurately a large range of materials, from diffuse to specular (see Figures 3(b), 4 and the supplemental material).

4.3. Shadowing and masking

Once we have D, we can compute the shadowing and masking function, G_1, using Equation 6:

$$P_2(r) = \frac{1}{\sqrt{\pi}\alpha\Gamma(1-p,\alpha)} U\left(p, p+\frac{1}{2}, \alpha+\frac{r^2}{\alpha}\right) e^{-\alpha-\frac{r^2}{\alpha}}$$

$$\Lambda(\theta) = \int_{\frac{1}{\tan\theta}}^\infty (r\tan\theta - 1) P_2(r)\,dr$$

$$G_1(\theta) = \frac{\chi_{[0,\frac{\pi}{2}]}(\theta)}{1 + \Lambda(\theta)}$$

where U is the confluent hyper-geometric function of the second kind. We do not have a closed form expression for G_1. We either pre-compute G_1 using Gauss integration and store its values for varying θ or use the following approximation:

$$G_1(\theta) \approx \begin{cases} 1 + \lambda\left(1 - e^{c(\theta-\theta_0)^k}\right) & \text{if} \quad \theta > \theta_0 \\ 1 & \text{if} \quad \theta \le \theta_0 \end{cases} \quad (9)$$

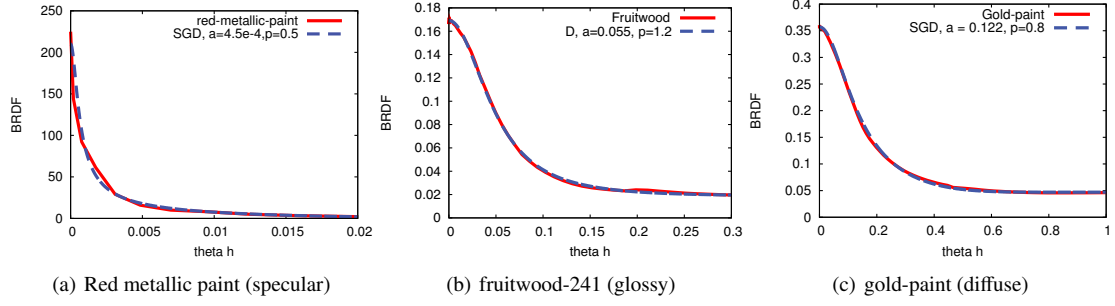$$\theta_0 = \frac{\pi}{2} - \left(\frac{\log\left(1 + \frac{1}{\lambda}\right)}{c}\right)^{\frac{1}{k}}$$

(a) Red metallic paint (specular) (b) fruitwood-241 (glossy) (c) gold-paint (diffuse)

Figure 4: *The SGD distribution is an accurate fit for a wide range of materials.*

(a) Red metallic paint (specular) (b) fruitwood-241 (glossy) (c) gold-paint (diffuse)
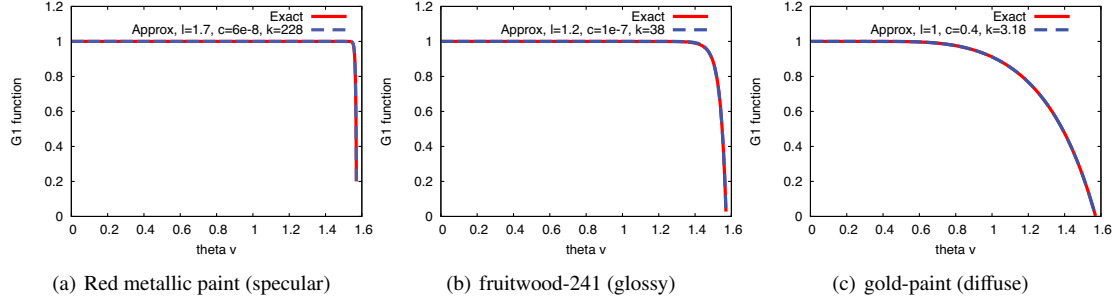
Figure 5: *The shadowing term G_1 is expensive for SGD. Our approximation is very accurate; the L_2 error is below 10^{-5}.*

θ_0 ensures that $G_1(\pi/2) = 0$. We get the (λ, c, k) parameters by fitting the approximation over the precomputed values, using Levenberg-Marquardt [Lou04]. This approximation is a very good fit for G_1 for all materials (the L_2 error is negligible, mostly below 10^{-5}, see Figure 5).

The choice between using stored values for G_1 and the approximation depends on the processing power, memory bandwidth and storage capacities. For offline rendering, where storage is not an issue, we use precomputed values. For GPU rendering, we use the approximation.

4.4. The Fresnel term

Since we have an expression for D and G_1, we can express the Fresnel term using the measured data:

$$F(\mathbf{i} \cdot \mathbf{h}) = \left(m(\mathbf{i}, \mathbf{o}) - \frac{\rho_d}{\pi} \right) \frac{\pi \cos\theta_i \cos\theta_o}{\rho_s D(\theta_h) G_1(\theta_i) G_1(\theta_o)} \quad (10)$$

where $m(\mathbf{i}, \mathbf{o})$ is the measured data. For simplicity, we will note $\theta_d = \arccos(\mathbf{i} \cdot \mathbf{h})$. Looking at the data for the Fresnel term, we make two important observations:

- The measured data behaves strangely at grazing angles. In theory, the Fresnel term should be increasing with θ_d up to 1 when $\theta_d = \pi/2$. The Fresnel term extracted from the measured data starts decreasing after a certain angle (from 70° to 80°), and is equal to 0 for $\theta_d = \pi/2$ (see Figure 6(a)). This is probably caused by the acquisition process: at grazing angles, the sensor is saturated by direct

illumination from the light source[†]. We discard data for $\theta_d > 70°$.

- For some materials such as *gold-paint* or *nickel*, the Fresnel term does not behave as predicted by the Schlick approximation (see Figure 6(b)). Instead of a plateau followed by a sharp increase, we observe linear slope, followed by an increase. To model these materials, we suggest the following generalization of Schlick approximation:

$$F(\theta_d) = F_0 + (1 - F_0)(1 - \cos\theta_d)^5 - F_1 \cos\theta_d \quad (11)$$

For $F_1 = 0$, we get the original approximation.
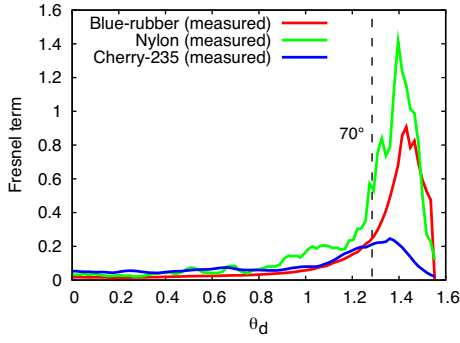
4.5. Importance sampling

Efficient Monte-Carlo integration requires importance sampling. For this, we need the inverse of the integral of P_{22} (see Equation 5). We easily have F_θ:

$$F_\theta = 1 - \frac{\Gamma\left(1 - p, \alpha + \frac{\tan^2\theta}{\alpha}\right)}{\Gamma(1 - p, \alpha)}$$
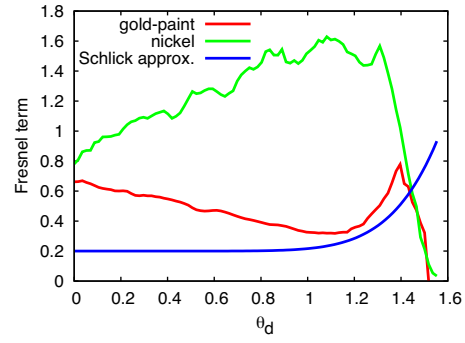
but there is no closed form expression for the inverse for F_θ, so we can't use it for exact importance sampling. As an approximation, we importance sample the GGX distribution [WMLT07], and compute weights, w, to fit our model:

$$\theta_m = \arctan\left(\frac{\alpha \sqrt{u_2}}{\sqrt{1 - u_2}}\right)$$

[†] F. Durand, personal communication.

(a) At grazing angles, it decreases to 0 instead of increasing to 1.

(b) For some materials, we observe a linear slope, different from Schlick's approximation.

Figure 6: *The Fresnel term extracted from the measured data does not behave as predicted by the theory*
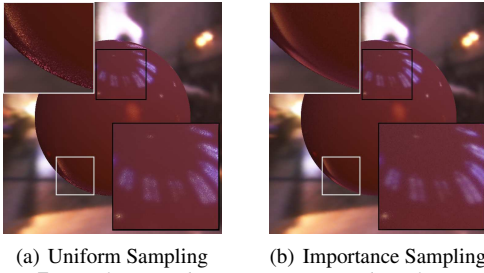

(a) Uniform Sampling   (b) Importance Sampling

Figure 7: *Uniform and importance sampling for* specular-violet-phenolic *with the same rendering time (6 mn).*

$$\phi_m = 2\pi u_1$$
$$\boldsymbol{m} = (\cos\phi_m \sin\theta_m, \sin\phi_m \sin\theta_m, \cos\theta_m)$$
$$\boldsymbol{o} = 2|\boldsymbol{i} \cdot \boldsymbol{m}|\boldsymbol{m} - \boldsymbol{i}$$

where u_1 and u_2 are realizations of a uniform random variable in $[0, 1)$. Our BRDF model contains a diffuse and a specular component. We alternatively sample either component according to the following probabilities:

$$p_{specular} = \frac{\rho_s}{\rho_d + \rho_s}$$
$$p_{diffuse} = 1 - p_{specular}$$

The importance sampling weights for our distribution is:

$$w_{specular} = \frac{4\rho_s}{p_{specular}} F(\boldsymbol{i} \cdot \boldsymbol{h}) D(\theta_h) G(\boldsymbol{i}, \boldsymbol{o}) \cos^3\theta_m \frac{(\alpha^2 + \tan^2\theta)^2}{\alpha^2}$$
$$w_{diffuse} = \frac{\rho_d}{p_{diffuse}}$$

This expression accounts for the jacobian of sampling $\boldsymbol{m}$ instead of $\boldsymbol{o}$ [WMLT07]. Figure 7 shows the same material rendered with and without importance sampling.

5. Fitting algorithm

5.1. Fitting RGB channels separately

Fitting BRDF models with measured materials is usually done in a two-step approach [NDM05]: first compute the dif-

fuse and specular colors using a linear least square optimization, then compute the BRDF parameters (α, F_0) through a second optimization.

This approach places constraints on the second optimization process: the parameters must fit all RGB channels. We used a different approach: we compute the BRDF parameters separately for each channel, red, green and blue. Although we compute and store more data for each material, the fitting process is faster and more efficient, as we only fit real-valued functions, with less constraints for each fitting.

We also get wavelength dependent effects. The Fresnel term should change with the wavelength [LSK05] and we do find different values for F_0 and F_1 for the red, green and blue channels, for all materials in the MERL database (see supplemental material). Often, we get very similar values for the α and p parameters for the micro-facet distribution.

We do not take any explicit step to conserve the color of the material, since we treat each channel independently. We still get BRDF models that match very well the color of the measured materials (see Figures 1, 8, and 10 and the supplemental). We attribute this to the very good match between our distribution function and the data behavior.

5.2. Non-linear optimization

Given the measured data for a single channel, we find the optimal value for the parameters ($\rho_d, \rho_s, \alpha, p, F_0, F_1$) so that the BRDF model is as close as possible to the measured data.

We could compute the L_2 norm of the difference between the BRDF model and the measured BRDF over the whole hemisphere, as in [NDM05]. But regular sampling in θ_i and θ_o results in under-sampling near the specular peak, resulting in poor quality fits. Instead we fit in two steps:

1. We fit the probability distribution parameters (α, p) to the measured data for varying θ_h, with $\theta_d = 0$.
2. We fit the Fresnel term parameters (F_0, F_1) to the measured data for varying θ_d, with $\theta_h = 0$.

Each step is the fitting of a real-valued function of several parameters, which is simple and stable. Although we fit the parameters using only two slices of the measured reflectance, we still capture the full behavior of the BRDF.

We define the following quantities:

- $m(\boldsymbol{i}, \boldsymbol{o})$: the full measured reflectance, as extracted from the database,
- $m_h(\theta_h) = m(\boldsymbol{h}, \boldsymbol{h})$: the slice of the BRDF corresponding to $\theta_d = 0$,
- $m_d(\theta_d) = m(\boldsymbol{i}, \text{reflection}(\boldsymbol{i}))$: the slice of the BRDF corresponding to $\theta_h = 0$.

In a first step, we fit $m_h(\theta_h)$ with:

$$\rho_h(\theta_h) = \frac{\rho_d}{\pi} + \frac{F_r}{\pi} \frac{D(\theta_h)G_1^2(\theta_h)}{\cos^2(\theta_h)}$$

The function ρ_h depends on 4 parameters: ρ_d, F_r, α and p. Using Levenberg-Marquadt optimization [Lou04], we find the value of these parameters that minimizes the error function:

$$E_1(\rho_d, F_r, \alpha, p) = \sum_i w_i (m_h(\theta_h) - \rho_h(\theta_h))^2 \qquad (12)$$

At the end of this step, we have the value for 3 parameters: ρ_d, α and p. The last coefficient, F_r is the product of ρ_s and the Fresnel term for $\theta_d = 0$.

In the second step, we first extract the specular coefficient times the Fresnel term from $m_d(\theta_d)$ using the values ρ_d, α and p we just computed:

$$\rho_s F_D(\theta_d) = \left(m_d(\theta_d) - \frac{\rho_d}{\pi}\right) \frac{\pi \cos^2(\theta_d)}{D(0)G_1^2(\theta_d)}$$

We fit this function with our Fresnel term approximation (Equation 11), using Levenberg-Marquadt optimization. The data gives us $\rho_s F_0$ and $\rho_s F_1$. Using $F_D(\pi/2) = 1$, we get ρ_s, F_0 and F_1.

6. Results and Comparison

6.1. Visual comparison

Figures 8 and 10 shows a side-by-side comparison for *aventurine* and *colonial-maple-223* between images rendered using the measured reflectance from the MERL database and several BRDF models: Lafortune and Cook-Torrance with the Beckmann, TR/GGX and SGD distributions. For each model, we show the difference with the reference image, measured in Lab space. We chose the Lab color space because it is perceptually uniform. Qualitatively, our BRDF model provides the best approximation to the measured data.

In the supplemental material, we provide a similar comparison for all 100 materials in the MERL database: reference image and Cook-Torrance BRDF with Beckmann and SGD distributions, along with differences in Lab space.
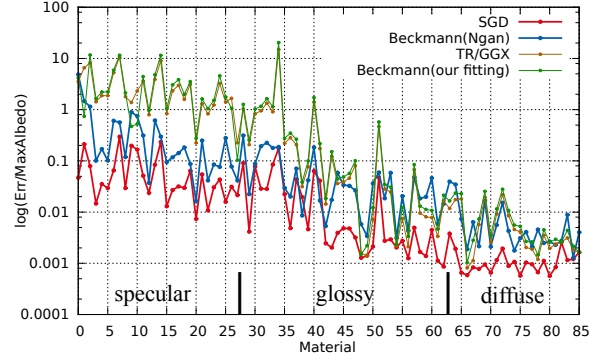
Figure 11: *Normalized fitting errors (logarithmic scale) of the Cook-Torrance model with Beckmann and SGD distributions. In most cases, SGD provides a much better fit.*

For 13 materials in the database, we couldn't find the corresponding parameters in the supplemental of [NDM05], so we only provide the comparison between the measured material and our distribution. Visually, the Cook-Torrance BRDF with the SGD distribution looks very close to the measured reflectance.

6.2. BRDF Lobes comparison

Figures 8(j), 10(j) and 9, plot the BRDF lobes for *aventurine*, *colonial-maple-223*, *hematite*, *black-obsidian*, *nickel*, and *black-oxidized-steel*, with the Beckmann, GGX and SGD distributions compared to the measurements, for four incoming directions ($10°$, $30°$, $50°$, and $70°$).

The SGD distribution closely fits the measured data, while the Beckmann distribution either misses the tail by trying to fit the peak of the lobe, or vice-versa. This is particularly evident for *black-obsidian* where the Beckmann distribution underestimates the peak. In the *hematite*, *aventurine*, *nickel*, *black-oxidized-steel*, and *colonial-maple-223* cases, which are ranging from highly-specular to glossy-specular and glossy-diffuse, the Beckmann distribution significantly overestimates the lobes. The GGX distribution performs better than Beckmann and SGD is very close to the actual shape of the lobes.

The supplemental material provides the same comparison for all 100 materials in the MERL database: lobes from the measured data and lobes from the Cook-Torrance BRDF, using Beckmann and SGD distributions. For most materials, the SGD distribution is very close to the measured data.

6.3. Quantitative Error Measure

For each material, we define the fitting error as the L_2 norm of the difference between measured data and the Cook-Torrance BRDF:

$$E = \left(\int (m(\boldsymbol{i}, \boldsymbol{o}) - \rho(\boldsymbol{i}, \boldsymbol{o}))^2 \cos\theta_i \cos\theta_o \, d\omega_i\right)^{\frac{1}{2}}$$

(a) Lafortune   (b) Beckmann   (c) TR/GGX   (d) SGD (ours)   (e) Ground truth

(f) Lafortune Lab error   (g) Beckmann Lab error   (h) TR/GGX Lab error   (i) SGD Lab error (ours)   (j) BRDF Lobes
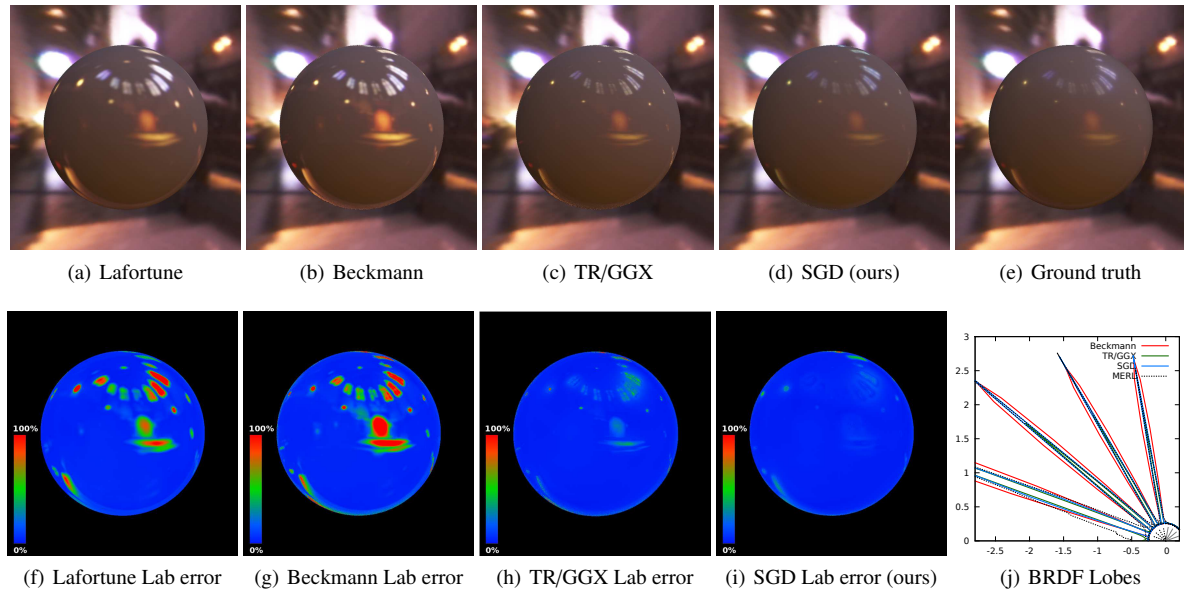
Figure 8: *Comparison between the Cook-Torrance model with Beckmann, TR/GGX and SGD distributions, Lafortune and ground truth for* aventurine *from the MERL database (spelled* aventurnine *in [NDM05]).*
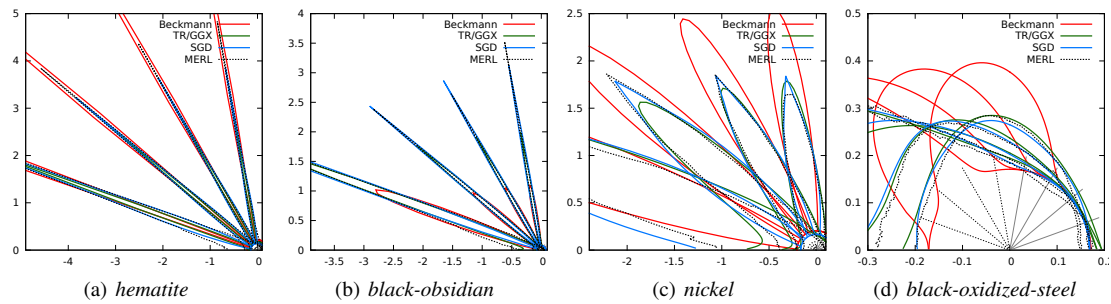

(a) *hematite*   (b) *black-obsidian*   (c) *nickel*   (d) *black-oxidized-steel*

Figure 9: *BRDF lobes (cubic root applied) with Beckmann, GGX and SGD distributions for* hematite, black-obsidian, nickel, *and* black-oxidized-steel. *The SGD distribution provides a much closer fit to measured reflectance than the Beckmann and GGX distribution.*

To ensure a fair comparison between bright and dark materials, we normalize this error by dividing it with the maximum albedo of the BRDF, as in [NDM05] (the supplemental has the un-normalized error value). Figure 11 compares the fitting errors for the Cook-Torrance BRDF using the SGD and Beckmann distributions (parameters for Beckmann from [NDM05]). We sorted the materials by increasing α roughness, going from highly specular to diffuse. For all but 3 of the materials (*black-fabric*, *yellow-matte-plastic*, and *white-paint*), the SGD distribution provides the best approximation. For many materials, we improve the accuracy by one order of magnitude, and for some, such as *hematite*, by two orders of magnitude.

For the materials where we do not improve over Beckmann distribution, we think that the problem comes from the fitting algorithm rather than the SGD distribution. Man-

ually setting the parameters for these materials improves the accuracy (all the parameters we use in this paper were computed in a single run of our optimization process, without any tweaking for specific materials).

To evaluate our fitting method, we used it to compute parameters for the Beckmann and TR/GGX distributions, and reported the error for each material in Figure 11. For almost all materials, the error is higher than with Beckmann distribution and parameters from [NDM05]. Our fitting method works better if the shape of the distribution matches the measured data.

Figure 12 provides a pixel-by-pixel comparison for a horizontal scanline of rendered image between: measured reflectance, Cook-Torrance with Beckmann distribution and SGD distribution, with and without importance sampling. The Beckmann distribution overestimates specular lobes;
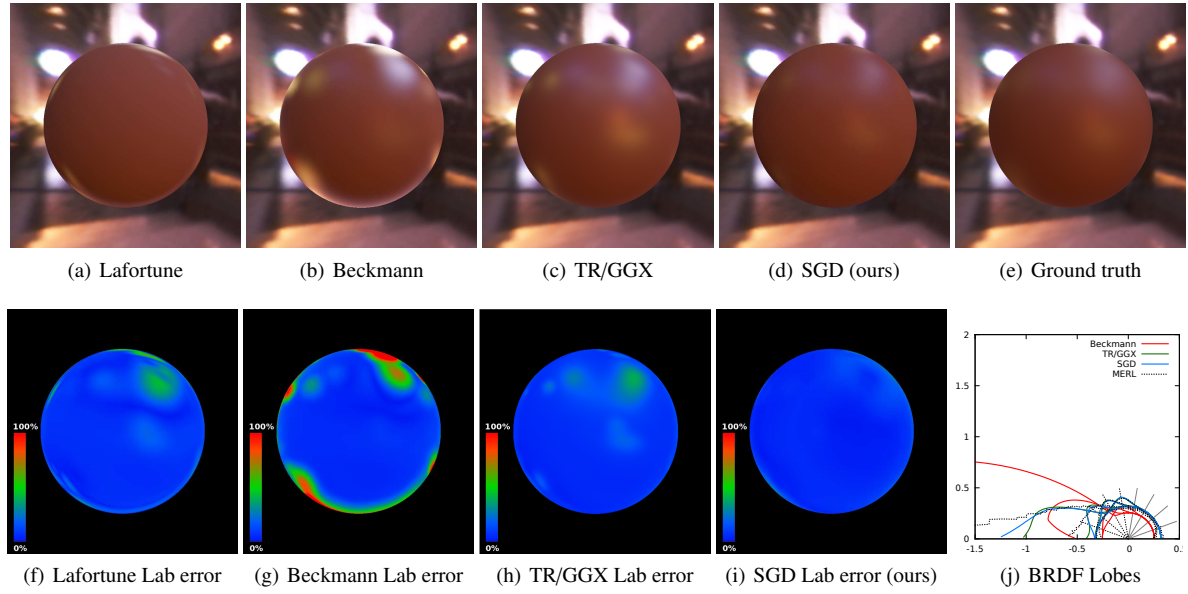
(a) Lafortune	(b) Beckmann	(c) TR/GGX	(d) SGD (ours)	(e) Ground truth

(f) Lafortune Lab error	(g) Beckmann Lab error	(h) TR/GGX Lab error	(i) SGD Lab error (ours)	(j) BRDF Lobes

Figure 10: *Comparison between the Cook-Torrance model with Beckmann, TR/GGX and SGD distributions, Lafortune and ground truth for* colonial-maple-223 *from the MERL database.*
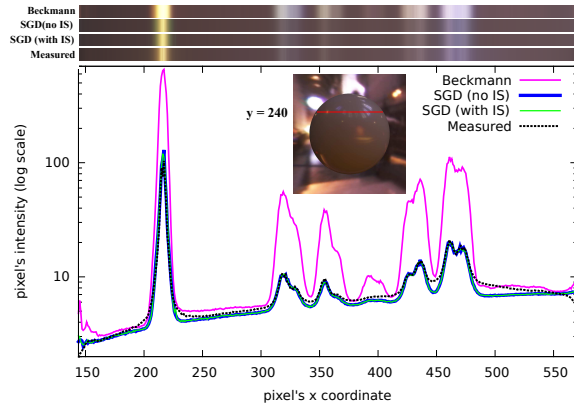

Figure 12: *Pixel by pixel comparison between the reference, Beckmann and SGD (with and without importance sampling) for* aventurine *from MERL database.*

SGD remains very close to the reference. The figure also validates that using GGX importance sampling for SGD distribution gives the same results, but reduces the noise.

6.4. Timings

Since our algorithm only uses two slices of the BRDF measurement for fitting the parameters, it is extremely fast: 2.5 minutes on average to fit all the parameters for a single material on a single core 2.57 GHz Intel CPU. For 65 % of the materials, the fitting process took less than 1 minute.
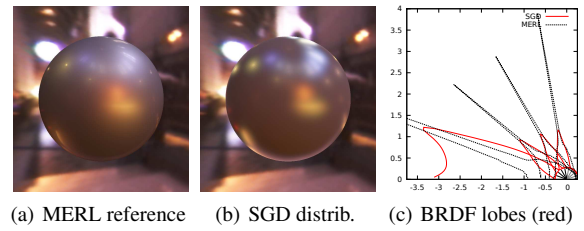

(a) MERL reference	(b) SGD distrib.	(c) BRDF lobes (red)

Figure 13: *An example of a two-layer material,* two-layer-silver, *that requires two lobes for accurate fitting.*

We also measured the cost of evaluating our reflectance model inside a ray-tracer, using valgrind profiling tool. A single evaluation takes 3044 CPU cycles, compared to 1416 with Beckmann and 551 with Lafortune. Our model is more expensive because we estimate the BRDF independently for each color channel.

6.5. Limitations

For some materials in the MERL database, even with the SGD distribution, we see evidence of a multi-lobe behavior. This is quite obvious for *two-layer-gold* and *two-layer-silver* (see Figure 13), and it appears (but it's less obvious) for *alum-bronze*. For these materials, adding a second lobe would improve the accuracy. However, the number of materials with multi-lobe behavior is 3 out of 100.

Our fitting method is based on only two slices of the material and on the assumption that the material follows the
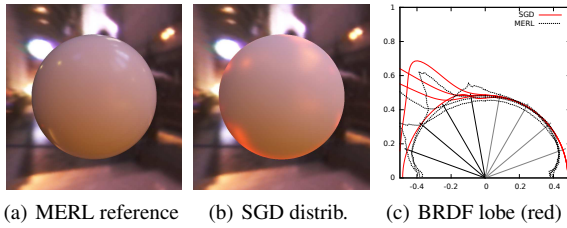
(a) MERL reference   (b) SGD distrib.   (c) BRDF lobe (red)

Figure 14: white-paint, *a failure case for our fitting algorithm.*

Cook-Torrance model. It will miss behavior that are outside of the model, for example *color-changing-paint*.

The *white-paint* material illustrates a failure case of our fitting algorithm (see Figure 14). The error with SGD is larger than with Beckmann. The difference is clearly visible in the rendered pictures (Figure 14(b)). A possible explanation is visible in the lobes (Figure 14(c)): the Fresnel term is almost null at normal incidence, and our fitting algorithm does not pick the right shape for the specular lobe.

7. Conclusion and Future work

We have presented the SGD micro-facet probability distribution function for the Cook-Torrance BRDF. This distribution provides an accurate fit for most measured materials. We have also presented a new fitting algorithm, where we fit each color channel independently (instead of precomputing the diffuse and specular colors), and we fit using only two slices of the BRDF. This fitting algorithm is very efficient when combined with the SGD distribution, taking 2.5 mn to compute the parameters for a measured distribution.

In the supplemental, we provide parameters for the SGD distribution for all materials in the MERL database, as well as GLSL code. We hope that this will help researchers and software developers in using the SGD distribution.

In future work, we want to continue working on micro-facet distributions, finding a distribution that provides the same accurate fit to measured data while being easier to compute and integrating into nicer functions. We also want to target material acquisition: we describe an acquired material accurately with only 18 coefficients. Can we use this property for faster acquisition of new materials?

References

[AP07] Ashikhmin M., Premoze S.: Distribution-based BRDFs. University of Utah, http://www.cs.utah.edu/~premoze/dbrdf/, 2007. 2, 4

[APS00] Ashikmin M., Premože S., Shirley P.: A microfacet-based BRDF generator. In *SIGGRAPH '00* (2000), pp. 65–74. 2

[BBS02] Bourlier C., Berginc G., Saillard J.: One- and two-dimensional shadowing functions for any height and slope stationary uncorrelated surface in the monostatic and bistatic con-figurations. *IEEE Transactions on Antennas and Propagation 50*, 3 (Mar. 2002), 312 –324. 2, 3

[Bro80] Brown G.: Shadowing by non-gaussian random surfaces. *IEEE Transactions on Antennas and Propagation 28*, 6 (Nov. 1980), 788 – 790. 2, 3

[CT82] Cook R. L., Torrance K. E.: A reflectance model for computer graphics. *ACM Trans. Graph. 1*, 1 (1982), 7–24. 2, 3

[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. 2

[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. 2

[Lou04] Lourakis M. I. A.: levmar: Levenberg-marquardt non-linear least squares algorithms in C/C++. http://www.ics.forth.gr/~lourakis/levmar/, Jul. 2004. 5, 7

[LSK05] Lazányi I., Szirmay-Kalos L.: Fresnel term approximations for metals. In *Winter School of Computer Graphics (WSCG), Short papers* (2005). 2, 6

[MPBM03] Matusik W., Pfister H., Brand M., McMillan L.: A data-driven reflectance model. *ACM Trans. Graph. 22*, 3 (July 2003), 759–769. 1, 2, 4

[NDM05] Ngan A., Durand F., Matusik W.: Experimental analysis of BRDF models. In *Eurographics Symposium on Rendering* (2005), pp. 117–226. 1, 2, 6, 7, 8

[ON94] Oren M., Nayar S. K.: Generalization of lambert's reflectance model. In *SIGGRAPH '94* (1994), pp. 239–246. 2

[PSCS*12] Pacanowski R., Salazar-Celis O., Schlick C., Granier X., Pierre P., Annie C.: Rational BRDF. *IEEE Transactions on Visualization and Computer Graphics* (2012). 2

[RVZ08] Romeiro F., Vasilyev Y., Zickler T.: Passive reflectometry. In *10th European Conference on Computer Vision: Part IV* (2008), pp. 859–872. 2

[Sch94] Schlick C.: An inexpensive BRDF model for physically-based rendering. *Computer Graphics Forum 13*, 3 (1994), 233–246. 2

[Smi67] Smith B.: Geometrical shadowing of a random rough surface. *IEEE Transactions on Antennas and Propagation 15*, 5 (Sept. 1967), 668 –671. 2, 3

[TR75] Trowbridge T. S., Reitz K. P.: Average irregularity representation of a rough surface for ray reflection. *J. Opt. Soc. Am. 65*, 5 (May 1975), 531–536. 2, 3

[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. 1, 2, 3

[War92] Ward G. J.: Measuring and modeling anisotropic reflection. *SIGGRAPH Comput. Graph. 26*, 2 (1992), 265–272. 2

[WMLT07] Walter B., Marschner S., Li H., Torrance K. E.: Microfacet models for refraction through rough surfaces. In *Eurographics Symposium on Rendering* (2007). 2, 3, 5, 6

[WRG*09] Wang J., Ren P., Gong M., Snyder J., Guo B.: All-frequency rendering of dynamic, spatially-varying reflectance. *ACM Trans. Graph. 28*, 5 (2009), 133:1–133:10. 2

[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. 2

[YSL11] Yu C., Seo Y., Lee S. W.: Global optimization for estimating a multiple-lobe analytical brdf. *Comput. Vis. Image Understanding 115*, 12 (Dec. 2011), 1679–1688. 2