Eurographics Symposium on Rendering - Experimental Ideas & Implementations (2015)
J. Lehtinen and D. Nowrouzezahrai (Editors)

Practical Rendering of Thin Layered Materials with Extended Microfacet Normal Distributions

Jie Guo[1], Jinghui Qian[1,2] and Jingui Pan[1]

[1]State Key Lab for Novel Software Technology, Nanjing University, China
[2] College of Computer Science and Technology, Nanjing Tech University, China

(a) Chess scene (16K triangles, 21.1 min.) (b) Teapot scene (19K triangles, 33.6 min.) (c) Vase scene (24K triangles, 38.9 min.)

Figure 1: *Several complex scenes rendered with our BSDF model using 1,600 samples per pixel (spp). Note that every object in these images is covered by a thin transparent layer with different roughness.*

Abstract

We propose a practical reflectance model for rendering thin transparent layers with different sides varying in roughness and levels of gloss. To capture the effect of subsurface reflection, previous methods rely on importance sampling for each light-surface interaction. This soon becomes a computationally demanding task since a recursive sampling scheme is required to handle multiple internal reflections. In this paper, we first provide a comprehensive analysis of the relationship between the directional distribution of scattered light and the roughness of each layer boundary using joint spherical warping. Based on the analysis, we generalize the traditional microfacet theory for layered materials and introduce the extended normal distribution function (ENDF) to accurately model the behavior of subsurface reflection. With the ENDF, the number of sampling processes can be reduced to only once for each bounce of subsurface reflection. We demonstrate that our BSDF model based on the ENDF is easy to be implemented on top of Monte Carlo sampling based offline renderers and it incurs little computational overhead. Moreover, it can be also efficiently used in real-time applications with the help of GPU acceleration.

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

1. Introduction

In the natural world, surfaces of many objects have a laminar structure incorporating a thin transparent or semi-transparent layer. Examples include glazed ceramics, waxed floor, metallic car paint, and some biological structures like skin or leaves. Simulating these materials in a physically cor-

DOI: 10.2312/sre.20151162

rect way requires accurately describing both surface reflection and subsurface reflection. For a thin transparent layer, there is no significant absorption and scattering inside the layer such that light scattering happens only at a boundary with a discontinuity in the refractive index. In this situation, subsurface reflection can be viewed as the net effect of multiple bounces of light interactions with the layer's boundaries. To deal with full subsurface reflection, infinitely many bounces of reflections and refractions at the layer's boundaries should be taken into consideration, leading to a daunting task. Some BSDF models adopt a simple ideal diffuse term to approximate subsurface reflection. However, such approximation is not physically plausible for many materials, and adjusting parameters to account for subtle variations in subsurface reflection from different materials is forbidden [HK93]. In general, most practical BSDF models fail to capture the physically correct effects of subsurface reflection [SSHL97, MHM*13], except for solving a very complex 1D radiative transfer equation with slow convergence rate [HK93, PH00] or with additional precomputation time and storage [Sta01, JdJM14].

In this paper, we propose an inexpensive and flexible BSDF model accounting for both surface reflection and subsurface reflection in the context of geometric optics. We use microfacet theory as the foundation to derive our model since this theory is considered as a fundamental part in physically-based, energy-conserving shading [MHM*13]. In microfacet-based BSDF models, estimating the statistical distribution of the microfacet's orientation which is defined as the normal distribution function (NDF), is a task of pivotal importance. Currently, microfacet theory is mainly used to predict the appearance under single bounce surface reflection. In order to handle subsurface reflection induced by multiple bounces of surface reflections and refractions, we improve microfacet theory by introducing a new representation, the extended normal distribution function (ENDF), to model visually perceived overall roughness of multiple bounces of reflections. With the ENDF, subsurface reflection can be treated within the same working framework as the single bounce surface reflection using microfacet theory.

One major contribution of this paper is a thorough study of how each bounce of internal reflection influences the observed surface reflectance, especially the perceived surface roughness. Based on the insights gained from the study, we describe an efficient method called joint spherical warping to faithfully estimate the ENDFs for multiple bounces of subsurface reflections. Furthermore, the corresponding modifications of shadowing-masking term and Fresnel reflectance are also analyzed. Finally, a practical BSDF model for layered materials is designed, which can be directly used in existing rendering pipeline without lengthy precomputation. We also show that this BSDF model can be efficiently utilized in real-time applications with pre-filtered environment maps.

2. Related Work

Microfacet-Based BSDF Models. Microfacet theory was introduced to computer graphics by Cook and Torrance [CT82] to quantify light reflection at rough surfaces. Experimental validations against real-world measurements have proven that BSDF models based on microfacet theory compare favorably against other families of parametric BSDF models [NDM05]. In microfacet-based BSDF models, the microfacet orientations are statistically described by the normal distribution function [Fou92, APS00]. The original NDF is a 2D function parameterized by the orientation of microgeometry normal, and it is further extended to a 4D function by including spatial variation (SVNDF) [WRG*09]. Bidirectional visible normal distribution function (BVNDF) [WDR11, IDN12] takes into account important shadowing and masking effects over appearance by employing densely sampled bidirectional visibility.

Besides surface reflection, microfacet theory has been extended to simulate transmitting effects through rough surfaces. Stam [Sta01] derived a microfacet-based BTDF model for transmissive materials with rough surfaces as part of his layered model for the reflectance of skin. This model was further extended by Walter et al. [WMLT07] with proper normalization and importance sampling strategies. Our work bears some similarity to that of Dai et al. [DWL*09], which proposed a dual-microfacet model to treat a special case of thin slabs with spatially-varying roughness, omitting light transport inside the objects. They estimate the overall NDF with an empirical solution, while we provide an analytical model with clear physical concepts. Furthermore, multiple internal reflections are also correctly captured in our work.

For microfacet-based BSDF models, correct shadowing-masking functions are essential for energy conservation [MHM*13, Hei14]. Several different analytical expressions appear in the literature, such as the Smith functions [Smi67] and the V-cavity masking function [CT82]. Heitz [Hei14] provided a good survey on this topic and also discussed how to derive shadowing-masking function according to a given microfacet distribution. In our work, we also provide an analytical expression for the shadowing-masking function based on the vMF distribution. Recently, a novel importance sampling strategy for microfacet-based BSDFs is presented, considering the influence of shadowing-masking term [Hd14].

BSDF Models For Layered Materials. Layered materials are widely adopted in computer graphics to describe the complete surface and subsurface scattering. Although rendering them can be computationally intensive, they offer a great potential for producing physically convincing results. Blinn [Bli82] was the first to generate the subsurface scattering effect in computer graphics in the context of dust-covered surfaces, based on a single-scattering assumption. Hanrahan and Krueger [HK93] extended Blinn's model and developed a more accurate scattering model for layered surfaces in terms of 1D linear transport theory. The efficiency of

Symbol	Description
$\mathbf{o}$	outgoing light direction
$\mathbf{i}$	incident light direction
$\mathbf{n}$	macro-scale geometric normal
$\mathbf{m}$	microfacet normal
$\mathbf{r}$ / $\mathbf{t}$	scattered light direction
$d\omega_{\mathbf{m}}$	an infinitesimal solid angle centered around $\mathbf{m}$
F	Fresnel reflectance
G	shadowing-masking term
D	normal distribution function
$\hat{D}$	extended normal distribution function
χ^+	Heaviside function
$\langle , \rangle$	clamped dot product between two directions
$*$	spherical convolution operator
η	relative index of refraction ($=\eta_{medium}/\eta_{air}$)

Table 1: *Notation used in this paper.*

this model was improved in [PH00] by deriving the scattering equations in an integral form. Apart from their accuracy, these solutions are prohibitively slow to converge.

To lower the computational complexity, a wide variety of approximate methods are available in the CG literature. Some inexpensive methods are designed for specific types of layer materials with moderate accuracy. For highly scattering and optically thick materials, a dipole or multipole diffusion approximation can be used to capture multiple scattered reflections [JMLH01, DJ05, PdMJ14]. For simulating absorption and scattering of tiny pigments inside a planar sheet, the Kubelka-Munk model [DH96, Pre02, GP14] works quite well if all incident and scattered light is assumed to be perfectly diffuse. To generate fast and accurate estimates of light distributions in multiple slabs, the adding-doubling method [EKM01, JdJM14] is highly recommended. Our work is inspired by that of Weidlich and Wilkie [WW07], which proposed a flexible family of layered BRDFs combining several microfacet-based surface layers. Since their work is based on the assumption that all secondary rays scattered from one layer interface are supposed to be collimated and meet at a single point on the next layer interface, the effect of one-bounce subsurface reflection seems to be glossier than it should be. Moreover, multiple bounces of internal reflections are not correctly handled. Our work circumvents these limitations simultaneously, and also offers additional benefits which will be discussed in the following sections.

3. Microfacet Theory

To motivate our approach, we start with a brief introduction of microfacet theory. In computer graphics, microfacet theory assumes that any rough surface can be modeled as a collection of randomly oriented microfacets whose statistical distribution can be described by a normal distribution function $D(\mathbf{m})$, and each microfacet (with surface normal $\mathbf{m}$) acts as a perfect specular reflector or refractor. A general BSDF model based on microfacet theory is derived in [WMLT07], which consists a BRDF model (see notation in Table 1):

$$f_r(\mathbf{o},\mathbf{i}) = M_r(\mathbf{o},\mathbf{i})D(\mathbf{m}) \tag{1}$$

$$M_r(\mathbf{o},\mathbf{i}) = \frac{F(\mathbf{i},\mathbf{m})G(\mathbf{o},\mathbf{i})}{4\,|\mathbf{o}\cdot\mathbf{n}|\,|\mathbf{i}\cdot\mathbf{n}|} \tag{2}$$

and a BTDF model:

$$f_t(\mathbf{o},\mathbf{i}) = M_t(\mathbf{o},\mathbf{i})D(\mathbf{m}) \tag{3}$$

$$M_t(\mathbf{o},\mathbf{i};\eta) = \frac{|\mathbf{o}\cdot\mathbf{m}|\,|\mathbf{i}\cdot\mathbf{m}|\,\eta^2(1-F(\mathbf{i},\mathbf{m}))G(\mathbf{o},\mathbf{i})}{|\mathbf{o}\cdot\mathbf{n}|\,|\mathbf{i}\cdot\mathbf{n}|\,(\mathbf{o}\cdot\mathbf{m}+\eta\mathbf{i}\cdot\mathbf{m})^2} \tag{4}$$

Here the remaining terms $M_r(\mathbf{o},\mathbf{i})$ and $M_t(\mathbf{o},\mathbf{i};\eta)$ contain the shadowing-masking term $G(\mathbf{o},\mathbf{i})$ and the Fresnel term $F(\mathbf{i},\mathbf{m})$, as well as proper normalization. It has been verified that both $M_r(\mathbf{o},\mathbf{i})$ and $M_t(\mathbf{o},\mathbf{i};\eta)$ are smooth, and have relatively little effect on the shape of BSDF [NDM05, WMLT07, WRG*09].

When adopting a microfacet-based BSDF model in physically-based rendering, there are several choices to be made.

Choosing D. We first decide on a distribution function that is used to fit the NDF $D(\mathbf{m})$. Several probability distributions exist for modeling the normal data on the unit sphere such as von Mises-Fisher (vMF), Fisher-Bingham, Kent distributions [MJ00]. We opt to use the vMF distribution for its simplicity and analytical tractability in deriving several formulas related to the NDF, especially the joint spherical warping approximation described in the following sections. Actually, a vMF can be regarded as a normalized spherical Gaussian on S^2. A trivariate vMF representation in 3D Euclidean space is parameterized by mean $\boldsymbol{\mu}$ and concentration κ and has density function:

$$\mathcal{M}(\mathbf{x};\boldsymbol{\mu},\kappa) = C_3(\kappa)e^{\kappa\boldsymbol{\mu}\cdot\mathbf{x}} \tag{5}$$

in which $C_3(\kappa) = \frac{\kappa}{4\pi\sinh\kappa}$. Obviously, this distribution is unimodal, and a multi-modal extension is the mixture of vMF distributions (movMF).

To make the vMF distribution valid for representing the NDF, a normalization factor is required to ensure that $\int_\Omega (\mathbf{n}\cdot\mathbf{m})D(\mathbf{m})d\omega_{\mathbf{m}} = 1$, in which $\mathbf{n}$ is the macro-scale geometric normal. This constraint equation yields:

$$D(\mathbf{m}) = \frac{\mathcal{M}(\mathbf{m};\boldsymbol{\mu}_{\mathbf{m}},\kappa_{\mathbf{m}})}{A_3(\kappa_{\mathbf{m}})(\mathbf{n}\cdot\boldsymbol{\mu}_{\mathbf{m}})} \tag{6}$$

where function $A_3(\kappa_{\mathbf{m}}) = \coth\kappa_{\mathbf{m}} - \frac{1}{\kappa_{\mathbf{m}}}$ returns the mean resultant length of a vMF with concentration parameter $\kappa_{\mathbf{m}}$. If bump mapping is not used, we can assume $\boldsymbol{\mu}_{\mathbf{m}} = \mathbf{n} = (0,0,1)^T$ in the tangent space of $\mathbf{n}$, such that $\mathbf{n}\cdot\boldsymbol{\mu}_{\mathbf{m}} = 1$. Moreover, $A_3(\kappa_{\mathbf{m}})$ approaches 1 very fast as $\kappa_{\mathbf{m}}$ increases. For instance, $A_3(100) \approx 0.99$. Therefore, it turns out that in many cases $D(\mathbf{m}) \approx \mathcal{M}(\mathbf{m};\boldsymbol{\mu}_{\mathbf{m}},\kappa_{\mathbf{m}})$. In the following, for the sake of simplicity, we will use the original form

of the vMF distribution to represent the NDFs. However, it is preferable to apply a normalization factor $1/A_3(\kappa_\mathbf{m})$ for small $\kappa_\mathbf{m}$ (< 100).

Choosing G. The shadowing-masking function $G(\mathbf{o},\mathbf{i})$ ($\in [0,1]$) describes the fraction of the microfacet which is visible either in the incident light direction $\mathbf{i}$ (shadowing term, $G_1(\mathbf{i},\mathbf{m})$) or in the viewing direction $\mathbf{o}$ (masking term, $G_1(\mathbf{o},\mathbf{m})$). Correct masking term $G_1(\mathbf{o},\mathbf{m})$ (or shadowing term $G_1(\mathbf{i},\mathbf{m})$) should be derived mathematically from the NDF $D(\mathbf{m})$ since they are not independent [MHM*13]. We follow the derivation in [Hei14] to compute $G_1(\mathbf{o},\mathbf{m})$ from the NDF $D(\mathbf{m})$ via

$$\begin{aligned} G_1(\mathbf{o},\mathbf{m}) &= \chi^+(\mathbf{o}\cdot\mathbf{m})\frac{\mathbf{o}\cdot\mathbf{n}}{\int_\Omega \langle \mathbf{o},\mathbf{m}\rangle D(\mathbf{m})\mathrm{d}\omega_\mathbf{m}} \\ &= \chi^+(\mathbf{o}\cdot\mathbf{m})\frac{A_3(\kappa_\mathbf{m})(\mathbf{n}\cdot\boldsymbol{\mu}_\mathbf{m})(\mathbf{o}\cdot\mathbf{n})}{\int_\Omega \langle \mathbf{o},\mathbf{m}\rangle \mathcal{M}(\mathbf{m};\boldsymbol{\mu}_\mathbf{m},\kappa_\mathbf{m})\mathrm{d}\omega_\mathbf{m}} \end{aligned} \tag{7}$$

in which a heaviside function is used to discard back-facing microfacets. Unfortunately, the integral at the denominator has no closed-form expression because of the nonlinearity of the clamped dot product. Based on the observation that the masking term is affected by both the concentration parameter $\kappa_\mathbf{m}$ and the cosine of viewing angle $\mathbf{o}\cdot\boldsymbol{\mu}_\mathbf{m}$, we find the following approximation to this integral:

$$\int_\Omega \langle \mathbf{o},\mathbf{m}\rangle \mathcal{M}(\mathbf{m};\boldsymbol{\mu}_\mathbf{m},\kappa_\mathbf{m})\mathrm{d}\omega_\mathbf{m} \approx a\cos(b\arccos(\mathbf{o}\cdot\boldsymbol{\mu}_\mathbf{m})) \tag{8}$$

in which $a = \frac{1}{4}(A_3(\kappa_\mathbf{m})+1)^2$ and $b = A_3(\kappa_\mathbf{m})^{\frac{1}{3}}$. Experimental validation in Fig. 2 reveals that our approximation matches the ground truth quite well for any roughness. In addition, from the plots we can see that $G_1(\mathbf{o},\mathbf{m})$ is almost constant except for very rough surface (e.g., $\kappa_\mathbf{m} = 1$). The final shadowing-masking function $G(\mathbf{o},\mathbf{i})$ can be obtained via $G(\mathbf{o},\mathbf{i}) = G_1(\mathbf{o},\mathbf{m})G_1(\mathbf{i},\mathbf{m})$, if the shadowing term and masking term are assumed to be uncorrelated [Hei14].

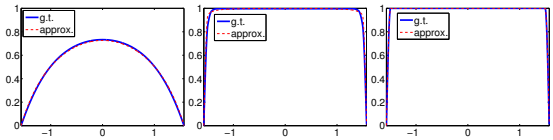

(a) $\kappa_\mathbf{m} = 1$ (b) $\kappa_\mathbf{m} = 10^2$ (c) $\kappa_\mathbf{m} = 10^4$

Figure 2: *The validation of our approximation to the masking term $G_1(\mathbf{o},\mathbf{m})$ based on Equation 8. Note that the approximated results (dashed red curves) are very close to the ground truth results (blue curves) for any rough surface.*

Choosing F. The Fresnel term $F(\mathbf{i},\mathbf{m})$ is relatively easy to choose. Following [CT82], it is given by

$$F(\mathbf{i},\mathbf{m}) = \frac{(g-c)^2}{2(g+c)^2}\left(1+\frac{(c(g+c)-1)^2}{(c(g-c)+1)^2}\right) \tag{9}$$

where $g^2 = \eta^2 + c^2 - 1$ and $c = |\mathbf{i}\cdot\mathbf{m}|$.

4. Extended Normal Distribution Function

It is important to note that the original microfacet BSDF model discussed in the previous section is based on the assumption that for a given screen pixel the source of incident light (here we regard the viewing direction as the incident direction based on the principle of reversibility) is at a great distance from the surface point to be shaded. This means all the incident light rays inside a pixel footprint are nearly parallel to one another, which is sketched in Fig. 3. For light scattering from a single surface, this assumption is valid, and it is the foundation for a well-known fact in microfacet theory which states that only microfacets with normals lying within solid angle $\mathrm{d}\omega_\mathbf{m}$ can reflect light into the solid angle $\mathrm{d}\omega_\mathbf{r} = 4|\mathbf{o}\cdot\mathbf{m}|\mathrm{d}\omega_\mathbf{m}$ as shown in Fig. 4(a).

However, this assumption is violated when dealing with subsurface reflection for layered materials, as illustrated in Fig. 3. Although the incident light beam can be assumed to be collimated for the first surface, the reflected or refracted light beams incident upon the second surface become smoothly distributed due to the surface bumpiness. As a consequence, the redirected light from the second surface seems to have a broader range compared with its intrinsic light scattering behavior (see Fig. 4(b)). In other words, the directional distribution of incident light beam will affect the perception of surface roughness, i.e., $D(\mathbf{m})$. Inspired by this observation, we seek to find an extended normal distribution function (ENDF) to represent this altered surface roughness such that subsurface reflection can be treated in the same way as the original microfacet BSDF.

Figure 3: *Because of rough reflection (left) or rough refraction (right) from the first surface, the light beam incident upon the second surface is no longer collimated even if the original light beam is.*

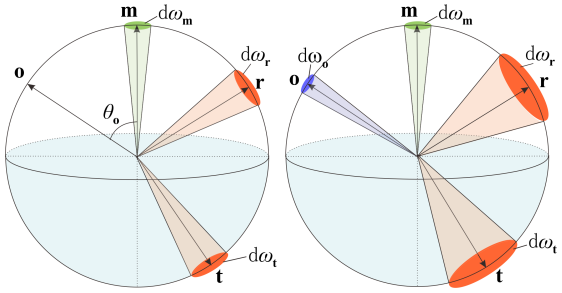

(a) collimated light beam (b) generalized light beam

Figure 4: *The vectors and infinitesimal solid angles involved in our derivation of the ENDF in this section.*

To tackle this problem, we present a study which investigates the relationships between directional distribution of incident light ($D(\mathbf{o})$), microfacet normal distribution ($D(\mathbf{m})$) and directional distribution of re-scattered light ($D(\mathbf{r})$ or $D(\mathbf{t})$) using joint spherical warping strategy, and includes an examination of how incident light beams with various distributions affect the shape of the ENDFs.

4.1. Spherical Warping

Warping Functions. First of all, we regard perfect specular reflection and refraction as two warping functions, which can be mathematically summarised as:

$$\Phi(\mathbf{x};\mathbf{m}) = 2(\mathbf{m}\cdot\mathbf{x})\mathbf{m} - \mathbf{x} \tag{10}$$

$$\Psi(\mathbf{x};\mathbf{m},\eta) = \left(\frac{\mathbf{m}\cdot\mathbf{x}}{\eta} - \sqrt{1-\frac{1-(\mathbf{m}\cdot\mathbf{x})^2}{\eta^2}}\right)\mathbf{m} - \frac{\mathbf{x}}{\eta} \tag{11}$$

Given a normalized direction, say $\mathbf{o}$, these two warping functions return the reflected direction $\mathbf{r} = \Phi(\mathbf{o};\mathbf{m})$, and refracted direction $\mathbf{t} = \Psi(\mathbf{o};\mathbf{m},\eta)$, respectively.

With these two functions, spherical warping strategy can be used to obtain the expression for the NDF in terms of scattered lighting vector $\mathbf{r}$ or $\mathbf{t}$, i.e., $D(\mathbf{r})$ or $D(\mathbf{t})$.

The Case of Reflection. For collimated incident beam with a normalized direction $\mathbf{o}$ (Fig. 4(a)), determining $D(\mathbf{r})$ is quite straightforward if the Jacobian determinant of the transform between $\mathbf{m}$ and $\mathbf{r}$ with respect to Φ (i.e., $\left\|\frac{\partial\mathbf{m}}{\partial\mathbf{r}}\right\|_\Phi$) is available. Since $D(\mathbf{m})\mathrm{d}\omega_\mathbf{m} = D(\mathbf{r})\mathrm{d}\omega_\mathbf{r}$, it is easy to verify that $D(\mathbf{r}) = \left\|\frac{\partial\mathbf{m}}{\partial\mathbf{r}}\right\|_\Phi D(\mathbf{m})$, in which $\left\|\frac{\partial\mathbf{m}}{\partial\mathbf{r}}\right\|_\Phi = \lim_{\mathrm{d}\omega_\mathbf{r}\to 0}\frac{\mathrm{d}\omega_\mathbf{m}}{\mathrm{d}\omega_\mathbf{r}} = \frac{1}{4|\mathbf{o}\cdot\mathbf{m}|}$.

Recall that our NDF ($D(\mathbf{m})$) is fitted with the vMF distribution, but the directional distribution of $D(\mathbf{r})$ does not necessarily agree with a vMF distribution. Inspired by the work of [WRG*09], we provide a vMF approximation to $D(\mathbf{r})$ that well matches the exact expression. Mathematically, the mean direction $\boldsymbol{\mu}_\mathbf{r}$ and the concentration parameter $\kappa_\mathbf{r}$ for $D(\mathbf{r})$ are approximated as:

$$\begin{aligned}\boldsymbol{\mu}_\mathbf{r} &= \Phi(\mathbf{o};\boldsymbol{\mu}_\mathbf{m})\\ \kappa_\mathbf{r} &= \left\|\frac{\partial\mathbf{m}}{\partial\mathbf{r}}\right\|_\Phi \kappa_\mathbf{m} = \frac{\kappa_\mathbf{m}}{4|\mathbf{o}\cdot\boldsymbol{\mu}_\mathbf{m}|}\end{aligned} \tag{12}$$

The Case of Refraction. The same strategy can be applied to the case of refraction by replacing Φ with Ψ and replacing $\left\|\frac{\partial\mathbf{m}}{\partial\mathbf{r}}\right\|_\Phi$ with $\left\|\frac{\partial\mathbf{m}}{\partial\mathbf{t}}\right\|_\Psi = \lim_{\mathrm{d}\omega_\mathbf{t}\to 0}\frac{\mathrm{d}\omega_\mathbf{m}}{\mathrm{d}\omega_\mathbf{t}} = \frac{\eta^2|\mathbf{t}\cdot\mathbf{m}|}{(\mathbf{o}\cdot\mathbf{m}+\eta\mathbf{t}\cdot\mathbf{m})^2}$ [WMLT07]. Then the mean direction $\boldsymbol{\mu}_\mathbf{t}$ and the concentration parameter $\kappa_\mathbf{t}$ for $D(\mathbf{t})$ are approximated as:

$$\begin{aligned}\boldsymbol{\mu}_\mathbf{t} &= \Psi(\mathbf{o};\boldsymbol{\mu}_\mathbf{m},\eta)\\ \kappa_\mathbf{t} &= \left\|\frac{\partial\mathbf{m}}{\partial\mathbf{t}}\right\|_\Psi \kappa_\mathbf{m} = \frac{\eta^2|\boldsymbol{\mu}_\mathbf{t}\cdot\boldsymbol{\mu}_\mathbf{m}|}{(\mathbf{o}\cdot\boldsymbol{\mu}_\mathbf{m}+\eta\boldsymbol{\mu}_\mathbf{t}\cdot\boldsymbol{\mu}_\mathbf{m})^2}\kappa_\mathbf{m}\end{aligned} \tag{13}$$

Validation. The accuracy of our approximation to both $D(\mathbf{r})$ and $D(\mathbf{t})$ using vMF distributions is provided in Fig. 5. Here the blue curves show the exact results of $D(\mathbf{r})$ (Fig. 5(a)) and $D(\mathbf{t})$ (Fig. 5(b)) with respect to four incidence angles ($\theta_\mathbf{o} = 0, \frac{\pi}{6}, \frac{\pi}{4}, \frac{\pi}{3}$), while the dashed red curves show corresponding single lobe vMF approximation. It is clear that in many cases the approximated results are very close to the ground truth results, while error increases as $\theta_\mathbf{o}$ approaches $\frac{\pi}{2}$. However, the error is subtle for the case of refraction, as is evident in Fig. 5(b). For the case of reflection, the shadowing-masking term $G(\mathbf{o},\mathbf{i})$ will partly reduce the error at grazing angles since $G(\mathbf{o},\mathbf{i})$ becomes unignorable as $\kappa_\mathbf{m}$ decreases. Other powerful models, such as mixture of vMF distributions or anisotropic distributions, can be employed to further reduce this error.

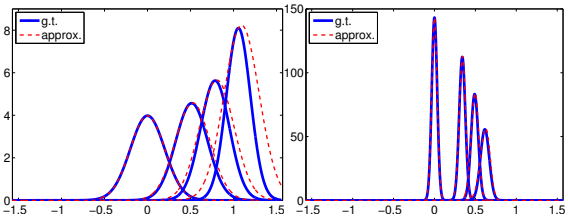

(a) $D(\mathbf{r})$ (b) $D(\mathbf{t})$

Figure 5: *2D plots of the ground truth results of $D(\mathbf{r})$ and $D(\mathbf{t})$ compared to the single lobe vMF approximated results. Here $\kappa_\mathbf{m} = 100$ and $\eta = 1.5$.*

4.2. Joint Spherical Warping

For the case of generalized light beam (Fig. 4(b)), simple spherical warping fails to correctly estimate $D(\mathbf{r})$ and $D(\mathbf{t})$, since both the directional distributions of $\mathbf{o}$ and $\mathbf{m}$ will affect their shapes. To address this issue, we propose joint spherical warping to obtain the expression of $D(\mathbf{r})$ or $D(\mathbf{t})$ given $D(\mathbf{o})$ and $D(\mathbf{m})$.

The Case of Reflection. Recall that $\mathbf{r} = 2(\mathbf{m}\cdot\mathbf{o})\mathbf{m} - \mathbf{o}$, and this equation can be regarded as a binary function that maps two random directions $\mathbf{o}$ and $\mathbf{m}$ into a random direction $\mathbf{r}$. We will use the method of transformation to find the directional distribution of $\mathbf{r}$. To do so, we define another function $\mathbf{s} = \mathbf{o}$, and the Jacobian of the transformation is

$$J(\mathbf{r},\mathbf{s}) = \begin{vmatrix}\frac{\partial\mathbf{m}}{\partial\mathbf{r}} & \frac{\partial\mathbf{m}}{\partial\mathbf{s}}\\ \frac{\partial\mathbf{o}}{\partial\mathbf{r}} & \frac{\partial\mathbf{o}}{\partial\mathbf{s}}\end{vmatrix} = \left\|\frac{\partial\mathbf{m}}{\partial\mathbf{r}}\right\|_\Phi \tag{14}$$

since $\mathbf{o}$ and $\mathbf{m}$ are independent ($\frac{\partial\mathbf{m}}{\partial\mathbf{s}} = 0$). Then the joint distribution of $\mathbf{r}$ and $\mathbf{s}$ is given by

$$D(\mathbf{r},\mathbf{s}) = D(\mathbf{m},\mathbf{o})J(\mathbf{r},\mathbf{s}) = \left\|\frac{\partial\mathbf{m}}{\partial\mathbf{r}}\right\|_\Phi D(\mathbf{m})D(\mathbf{o}) \tag{15}$$

We can then find the marginal distribution of $\mathbf{r}$ by integrating over $\mathbf{s}$ ($= \mathbf{o}$) as follows:

$$D(\mathbf{r}) = \int_\Omega D(\mathbf{r},\mathbf{s})\mathrm{d}\omega_\mathbf{s} = \int_\Omega \left\|\frac{\partial\mathbf{m}}{\partial\mathbf{r}}\right\|_\Phi D(\mathbf{m})D(\mathbf{o})\mathrm{d}\omega_\mathbf{o} \tag{16}$$

This equation is the exact form of joint spherical warping, and it is easy to verify that spherical warping is actually a special case of joint spherical warping in which $D(\mathbf{o})$ collapses to a delta function.

If both $D(\mathbf{m})$ and $D(\mathbf{o})$ follow vMF distributions, we are able to find another vMF distribution that fits $D(\mathbf{r})$. Here we resort to the moment preserving method, and we begin with the derivation of the first-order moment of $D(\mathbf{r})$:

$$\begin{aligned}
&\int \mathbf{r} D(\mathbf{r}) \mathrm{d}\omega_{\mathbf{r}} = \int \mathbf{r} \left(\int \left\| \frac{\partial \mathbf{m}}{\partial \mathbf{r}} \right\|_{\Phi} D(\mathbf{m}) D(\mathbf{o}) \mathrm{d}\omega_{\mathbf{o}} \right) \mathrm{d}\omega_{\mathbf{r}} \\
&\approx \int \left(\int \mathbf{r} \mathcal{M}\left(\mathbf{r}; \Phi(\mathbf{o}; \boldsymbol{\mu}_{\mathbf{m}}), \left\| \frac{\partial \mathbf{m}}{\partial \mathbf{r}} \right\|_{\Phi} \kappa_{\mathbf{m}} \right) \mathrm{d}\omega_{\mathbf{r}} \right) D(\mathbf{o}) \mathrm{d}\omega_{\mathbf{o}} \\
&= \int A_3 \left(\left\| \frac{\partial \mathbf{m}}{\partial \mathbf{r}} \right\|_{\Phi} \kappa_{\mathbf{m}} \right) \Phi(\mathbf{o}; \boldsymbol{\mu}_{\mathbf{m}}) \left\| \frac{\partial \mathbf{o}}{\partial \mathbf{r}} \right\|_{\Phi} D(\mathbf{o}) \mathrm{d}\omega_{\mathbf{r}} \\
&\approx A_3 \left(\left\| \frac{\partial \mathbf{m}}{\partial \mathbf{r}} \right\|_{\Phi} \kappa_{\mathbf{m}} \right) \int \mathbf{r} \mathcal{M}\left(\mathbf{r}; \Phi(\boldsymbol{\mu}_{\mathbf{o}}; \boldsymbol{\mu}_{\mathbf{m}}), \left\| \frac{\partial \mathbf{o}}{\partial \mathbf{r}} \right\|_{\Phi} \kappa_{\mathbf{o}} \right) \mathrm{d}\omega_{\mathbf{r}} \\
&= A_3 \left(\left\| \frac{\partial \mathbf{m}}{\partial \mathbf{r}} \right\|_{\Phi} \kappa_{\mathbf{m}} \right) A_3 \left(\left\| \frac{\partial \mathbf{o}}{\partial \mathbf{r}} \right\|_{\Phi} \kappa_{\mathbf{o}} \right) \Phi(\boldsymbol{\mu}_{\mathbf{o}}; \boldsymbol{\mu}_{\mathbf{m}})
\end{aligned} \tag{17}$$

Notice that two spherical warping approximations are involved in this derivation. The first one (shown in red) transforms $D(\mathbf{m})$ into the space of $\mathbf{r}$, and the other one (shown in blue) transforms $D(\mathbf{o})$ into the space of $\mathbf{r}$. To approximate $D(\mathbf{r})$ as a vMF distribution, say $D(\mathbf{r}) = \mathcal{M}(\mathbf{r}, \hat{\boldsymbol{\mu}}_{\mathbf{r}}, \hat{\kappa}_{\mathbf{r}})$, a strong candidate is the one whose first-order moment (i.e., $A_3(\hat{\kappa}_{\mathbf{r}})\hat{\boldsymbol{\mu}}_{\mathbf{r}}$) matches $A_3\left(\left\|\frac{\partial \mathbf{m}}{\partial \mathbf{r}}\right\|_{\Phi} \kappa_{\mathbf{m}}\right) A_3\left(\left\|\frac{\partial \mathbf{o}}{\partial \mathbf{r}}\right\|_{\Phi} \kappa_{\mathbf{o}}\right) \Phi(\boldsymbol{\mu}_{\mathbf{o}}; \boldsymbol{\mu}_{\mathbf{m}})$. It turns out that the mean direction and the concentration parameter should be

$$\begin{aligned}
\hat{\boldsymbol{\mu}}_{\mathbf{r}} &= \Phi(\boldsymbol{\mu}_{\mathbf{o}}; \boldsymbol{\mu}_{\mathbf{m}}) \\
\hat{\kappa}_{\mathbf{r}} &= A_3^{-1}\left(A_3 \left(\left\| \frac{\partial \mathbf{m}}{\partial \mathbf{r}} \right\|_{\Phi} \kappa_{\mathbf{m}} \right) A_3 \left(\left\| \frac{\partial \mathbf{o}}{\partial \mathbf{r}} \right\|_{\Phi} \kappa_{\mathbf{o}} \right) \right)
\end{aligned} \tag{18}$$

respectively. Here the Jacobian determinant $\left\|\frac{\partial \mathbf{o}}{\partial \mathbf{r}}\right\|_{\Phi}$ is easy to calculate, and the value happens to be 1. It is worth noting that the results in Equation 18 are quite similar to the spherical convolution results of two vMF distributions in [MJ00]. Hereafter we will use the notation of spherical convolution, i.e., $D(\mathbf{r}) = \left(\left\|\frac{\partial \mathbf{m}}{\partial \mathbf{r}}\right\|_{\Phi} D(\mathbf{m})\right) * \left(\left\|\frac{\partial \mathbf{o}}{\partial \mathbf{r}}\right\|_{\Phi} D(\mathbf{o})\right)$, to represent our calculation of $D(\mathbf{r})$ according to Equation 17.

It is fairly straightforward to sample $D(\mathbf{r})$ directly to estimate the final radiance for each screen pixel, or we can step further to find an ENDF $\hat{D}(\mathbf{m})$ such that generalized incident light beam can be treated in the same way as traditional collimated light beam that importance samples $\hat{D}(\mathbf{m})$ and then applies the adjustment for a change of variables $\left\|\frac{\partial \mathbf{m}}{\partial \mathbf{r}}\right\|_{\Phi}$. In this paper, we will use the latter sampling scheme by default. Since $\hat{D}(\mathbf{m}) = \left\|\frac{\partial \mathbf{r}}{\partial \mathbf{m}}\right\|_{\Phi} D(\mathbf{r})$, we find the following vMF approximation to the ENDF:

$$\hat{D}(\mathbf{m}) \approx \mathcal{M}\left(\mathbf{m}; \boldsymbol{\mu}_{\mathbf{m}}, \left\| \frac{\partial \mathbf{r}}{\partial \mathbf{m}} \right\|_{\Phi} \hat{\kappa}_{\mathbf{r}}\right) \tag{19}$$

With such expression of the ENDF, the incident light beam can be viewed as collimated (i.e., $D(\mathbf{o}) = \delta(\boldsymbol{\mu}_{\mathbf{o}})$), while keeping the final shape of $D(\mathbf{r})$ unchanged.

Validation of Reflection. Experimental results, which are provided in Fig. 6, reveal that our approximation method using spherical convolution provides a nice fit to the ground truth results (Equation 16). In this figure, we take three $D(\mathbf{o})$ and $D(\mathbf{m})$ pairs as examples (see the first row of Fig. 6). It is apparent from the 3D plots of $D(\mathbf{r})$ in the second row of Fig. 6 that the approximated results ($D(\mathbf{r})$ conv.) are visually identical to the ground truth results ($D(\mathbf{r})$ g.t.). To further explain the accuracy of our approximation, we provide scan line plot of the values on the equator plane of Example 2 in Fig. 8(a). Clearly, the difference between $D(\mathbf{r})$ conv. (dashed red curve) and $D(\mathbf{r})$ g.t. (blue curve) is subtle. Here we also show the 2D plot of the ENDF $\hat{D}(\mathbf{m})$ (dashed black curve). As we have seen, the ENDF has greater dispersion around the mean direction compared with the original NDF (green curve), which is consistent with our intuition.

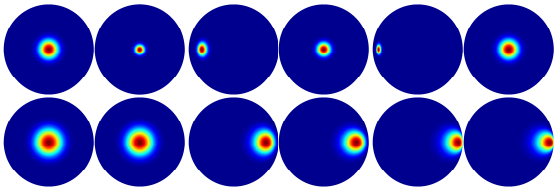

(a) Example 1 (b) Example 2 (c) Example 3

Figure 6: *Three reflection examples of joint spherical warping. In the first row of each example, left is $D(\mathbf{o})$ and right is $D(\mathbf{m})$; in the second row of each example, left is the ground truth result ($D(\mathbf{r})$ g.t.) and right is our approximated result using spherical convolution ($D(\mathbf{r})$ conv.). Example 1: $\boldsymbol{\mu}_{\mathbf{o}} \cdot \boldsymbol{\mu}_{\mathbf{m}} = 1$, $\kappa_{\mathbf{o}} = 50$, $\kappa_{\mathbf{m}} = 200$. Example 2: $\boldsymbol{\mu}_{\mathbf{o}} \cdot \boldsymbol{\mu}_{\mathbf{m}} = 0.707$, $\kappa_{\mathbf{o}} = 100$, $\kappa_{\mathbf{m}} = 100$. Example 3: $\boldsymbol{\mu}_{\mathbf{o}} \cdot \boldsymbol{\mu}_{\mathbf{m}} = 0.5$, $\kappa_{\mathbf{o}} = 200$, $\kappa_{\mathbf{m}} = 50$.*

The Case of Refraction. Similar reasoning allows the directional distribution of refracted light to be determined as:

$$D(\mathbf{t}) = \int_{\Omega} \left\| \frac{\partial \mathbf{m}}{\partial \mathbf{t}} \right\|_{\Psi} D(\mathbf{m}) D(\mathbf{o}) \mathrm{d}\omega_{\mathbf{o}} \tag{20}$$

Likewise, we can approximate $D(\mathbf{t})$ with a single lobe vMF distribution with parameters:

$$\begin{aligned}
\hat{\boldsymbol{\mu}}_{\mathbf{t}} &= \Psi(\boldsymbol{\mu}_{\mathbf{o}}; \boldsymbol{\mu}_{\mathbf{m}}, \eta) \\
\hat{\kappa}_{\mathbf{t}} &= A_3^{-1}\left(A_3 \left(\left\| \frac{\partial \mathbf{m}}{\partial \mathbf{t}} \right\|_{\Psi} \kappa_{\mathbf{m}} \right) A_3 \left(\left\| \frac{\partial \mathbf{o}}{\partial \mathbf{t}} \right\|_{\Psi} \kappa_{\mathbf{o}} \right) \right)
\end{aligned} \tag{21}$$

in which $\left\|\frac{\partial \mathbf{m}}{\partial \mathbf{t}}\right\|_{\Psi} = \frac{\eta^2 |\boldsymbol{\mu}_{\mathbf{t}} \cdot \boldsymbol{\mu}_{\mathbf{m}}|}{(\boldsymbol{\mu}_{\mathbf{o}} \cdot \boldsymbol{\mu}_{\mathbf{m}} + \eta \boldsymbol{\mu}_{\mathbf{t}} \cdot \boldsymbol{\mu}_{\mathbf{m}})^2}$ and $\left\|\frac{\partial \mathbf{o}}{\partial \mathbf{t}}\right\|_{\Psi} = \frac{\eta^2 |\boldsymbol{\mu}_{\mathbf{t}} \cdot \boldsymbol{\mu}_{\mathbf{m}}|}{|\boldsymbol{\mu}_{\mathbf{o}} \cdot \boldsymbol{\mu}_{\mathbf{m}}|}$.

Then it follows that the ENDF for refraction can be approximated as

$$\hat{D}(\mathbf{m}) \approx \mathcal{M}\left(\mathbf{m};\boldsymbol{\mu}_{\mathbf{m}}, \left\|\frac{\partial \mathbf{t}}{\partial \mathbf{m}}\right\|_{\Psi} \hat{\kappa}_{\mathbf{t}}\right) \tag{22}$$

Validation of Refraction. Similar validation of such approximation is illustrated in Fig. 7, and we use the same vMF parameters for $D(\mathbf{o})$ and $D(\mathbf{m})$ as shown in Fig. 6. Again, approximated results using spherical convolution ($D(\mathbf{t})$ conv.) are visually indistinguishable from the accurate ones ($D(\mathbf{t})$ g.t.). 2D plots of Example 2 are shown in Fig. 8(b), as well as the ENDF for refraction (dashed black curve).

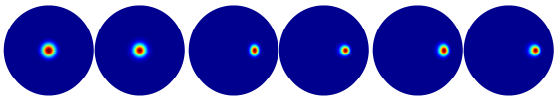

(a) Example 1 (b) Example 2 (c) Example 3

Figure 7: *Three refraction examples of joint spherical warping. The parameters of $D(\mathbf{o})$ and $D(\mathbf{m})$ are the same as in Fig. 6 while $\eta = 1.5$. In each example, left is $D(\mathbf{t})$ g.t. and right is $D(\mathbf{t})$ conv..*

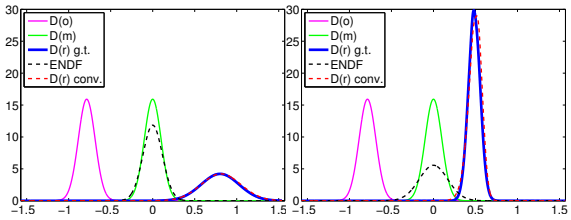

(a) (b)

Figure 8: *Scan line plots of Example 2 for reflection (a) and refraction (b), respectively. Here we also show the ENDFs with dashed black curves.*

4.3. A Case Study: Rough Refraction From Thin Slab

To prove the effectiveness of our joint spherical warping strategy proposed in the previous subsection, we show its application in rendering a thin transparent slab bounded by two rough surfaces, as presented in [DWL*09]. In that paper, Dai et al. introduced a dual-microfacet model for capturing refraction through flat slabs, and the overall NDF for the BTDF is obtained using an iterative model fitting method. Here we show that the overall NDF for a thin slab can be derived analytically using joint spherical warping.

The geometry of this application is shown in Fig. 9. In this figure, the slab has two rough surface boundaries, whose NDFs are characterized by $D(\mathbf{m}_1)$ and $D(\mathbf{m}_2)$, respectively. According to our approximated joint spherical warping strategy, if both $D(\mathbf{m}_1)$ and $D(\mathbf{m}_2)$ follow vMF distributions, the

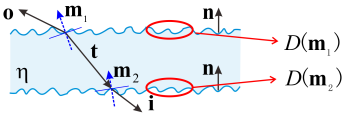

Figure 9: *Geometry for rough refraction from a thin transparent slab.*

directional distribution of $\mathbf{i}$ can be well approximated by a vMF distribution with the following parameters:

$$\begin{aligned}\hat{\boldsymbol{\mu}}_{\mathbf{i}} &= \Psi\left(-\Psi(\mathbf{o};\boldsymbol{\mu}_{\mathbf{m}_1},\eta);\boldsymbol{\mu}_{\mathbf{m}_2},\frac{1}{\eta}\right) = -\mathbf{o}\\ \hat{\kappa}_{\mathbf{i}} &= A_3^{-1}\left(A_3\left(\left\|\frac{\partial \mathbf{m}_1}{\partial \mathbf{i}}\right\|_{\Psi}\kappa_{\mathbf{m}_1}\right)A_3\left(\left\|\frac{\partial \mathbf{m}_2}{\partial \mathbf{i}}\right\|_{\Psi}\kappa_{\mathbf{m}_2}\right)\right)\end{aligned} \tag{23}$$

As before, we assume that $\boldsymbol{\mu}_{\mathbf{m}_1} = \boldsymbol{\mu}_{\mathbf{m}_2} = \mathbf{n}$. Obviously, these parameters are very similar to those in Equation 21, except that the directional distribution of the incident light beam to the bottom surface $D(\mathbf{t})$ is further determined by the NDF of the top surface $D(\mathbf{m}_1)$. Here the Jacobian determinant $\left\|\frac{\partial \mathbf{m}_1}{\partial \mathbf{i}}\right\|_{\Psi}$ can be obtained via the following chain rule:

$$\left\|\frac{\partial \mathbf{m}_1}{\partial \mathbf{i}}\right\|_{\Psi} = \left\|\frac{\partial \mathbf{m}_1}{\partial \mathbf{t}}\right\|_{\Psi}\left\|\frac{\partial \mathbf{t}}{\partial \mathbf{i}}\right\|_{\Psi} = J(\mathbf{m}_1 \xrightarrow{\Psi} \mathbf{t} \xrightarrow{\Psi} \mathbf{i}) \tag{24}$$

For brevity, we use the notation $J(\mathbf{m}_1 \xrightarrow{\Psi} \mathbf{t} \xrightarrow{\Psi} \mathbf{i})$ to denote the Jacobian determinant chain which gives the overall Jacobian determinant for the transformation along the path $\mathbf{m}_1 \xrightarrow{\Psi} \mathbf{t} \xrightarrow{\Psi} \mathbf{i}$.

Validation. To validate the accuracy of the approximated directional distribution of $\mathbf{i}$ parameterized with $\hat{\boldsymbol{\mu}}_{\mathbf{i}}$ and $\hat{\kappa}_{\mathbf{i}}$, we conduct two experiments in Fig. 10 with a flat slab and a curved slab, respectively. In both experiments, for each viewing direction $\mathbf{o}$, we generate samples of $\mathbf{i}$ according to the vMF distribution given by Equation 23, and for each sampled direction, we return the radiance of the intersection point directly without multiplying the remaining terms. The ground truth results are synthesized by explicitly sampling both $D(\mathbf{m}_1)$ and $D(\mathbf{m}_2)$ and applying appropriate adjustment.

As shown, for the flat slab model (2,304 triangles), even a close examination reveals no discriminable differences between our approximated result and the ground truth result. For the curved slab (6,400 triangles), there are large errors at grazing viewing angles due to the anisotropic distortion introduced by spherical warping, but those errors are still not significant. Since our method only sample once compared with the ground truth solution, the rendering time of our method reduces by nearly half. If more rough surfaces are stacked together, the rendering time of the ground truth solution will grow linearly with the number of surface layers,

while our method with appropriate ENDF will have nearly constant time cost.

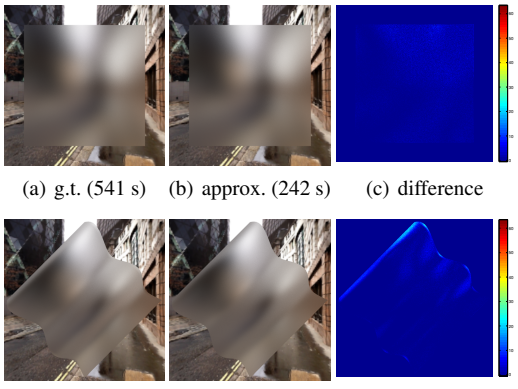

(a) g.t. (541 s) (b) approx. (242 s) (c) difference

(d) g.t. (1334 s) (e) approx. (562 s) (f) difference

Figure 10: *Image quality comparison between the ground truth (g.t.) and our method (approx.). Here we show two types of glass slabs with the same NDF ($\kappa = 10^2$) on all faces (without shadowing-masking and Fresnel effects).*

5. Reflectance Model for Thin Layered Materials

So far, we have derived the ENDFs for both reflection and refraction, and also have verified their accuracy against the ground truth. In this section, we investigate how to obtain physically plausible reflectance model for thin layered materials based on the extended normal distribution functions. In what follows, we restrict our consideration to one thin transparent layer composed of two smooth or rough boundaries. We further assume the lower boundary of the layer is opaque such that light is only reflected at this boundary. When a light beam hits such a layer, it is both reflected and refracted at the outer boundary depending on its surface properties. Light refracted into the layer may undergo multiple internal reflections before it finally exits the medium (see Fig. 11).

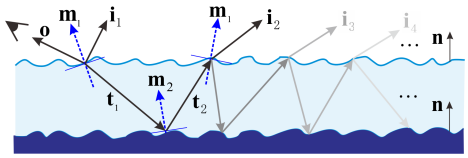

Figure 11: *Light scattering geometry for a thin transparent layer with two rough surface boundaries.*

According to the principle of independently propagating of light, the total outgoing radiance $L(\mathbf{o})$ is the superposition of all bounces of reflections from incident light $L(\mathbf{i}_1), L(\mathbf{i}_2), L(\mathbf{i}_3), \ldots$, respectively, applied with corresponding BSDFs, i.e.,

$$\begin{aligned} L(\mathbf{o}) = & \underbrace{\int_\Omega f_r(\mathbf{o},\mathbf{i}_1) L(\mathbf{i}_1) \langle \mathbf{i}_1, \mathbf{n} \rangle \mathrm{d}\omega_{\mathbf{i}_1}}_{\text{surface reflection}} \\ & + \underbrace{\int_\Omega f_r(\mathbf{o},\mathbf{i}_2) L(\mathbf{i}_2) \langle \mathbf{i}_2, \mathbf{n} \rangle \mathrm{d}\omega_{\mathbf{i}_2}}_{\text{one-bounce subsurface reflection}} \\ & + \underbrace{\int_\Omega f_r(\mathbf{o},\mathbf{i}_3) L(\mathbf{i}_3) \langle \mathbf{i}_3, \mathbf{n} \rangle \mathrm{d}\omega_{\mathbf{i}_3} + \cdots}_{\text{multi-bounce subsurface reflection}} \end{aligned} \tag{25}$$

The first integral on the right-hand side of Equation 25 gives the surface reflection, and it can be resolved by traditional microfacet-based BSDF models with actual NDF $D(\mathbf{m}_1)$. The remaining integrals model the subsurface reflection components induced by multiple internal reflections, and these integrals can be well resolved by the ENDFs.

5.1. One-Bounce Subsurface Reflection

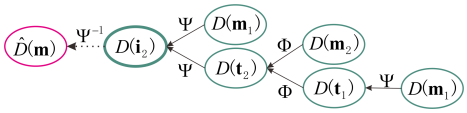

Figure 12: *The NDF tree for the determination of the ENDF for one-bounce subsurface reflection.*

We first demonstrate how to obtain the ENDF for one-bounce subsurface reflection. Recall that $D(\mathbf{i}_2)$ can be approximated by the spherical convolution between $D(\mathbf{m}_1)$ and $D(\mathbf{t}_2)$ with appropriate Jacobian determinant, while $D(\mathbf{t}_2)$ is further determined by $D(\mathbf{m}_2)$ and $D(\mathbf{t}_1)$. This procedure will iterate several times, and finally forms a NDF tree as shown in Fig. 12. Using the notation of the Jacobian determinant chain, we can write $D(\mathbf{i}_2)$ as

$$\begin{aligned} D(\mathbf{i}_2) = & [J(\mathbf{m}_1 \xrightarrow{\Psi} \mathbf{i}_2) D(\mathbf{m}_1)] * [J(\mathbf{m}_2 \xrightarrow{\Phi} \mathbf{t}_2 \xrightarrow{\Psi} \mathbf{i}_2) D(\mathbf{m}_2)] \\ & * [J(\mathbf{m}_1 \xrightarrow{\Psi} \mathbf{t}_1 \xrightarrow{\Phi} \mathbf{t}_2 \xrightarrow{\Psi} \mathbf{i}_2) D(\mathbf{m}_1)] \end{aligned} \tag{26}$$

Now, given the expression of $D(\mathbf{i}_2)$, it is straightforward to get its corresponding ENDF $\hat{D}(\mathbf{m})$ by multiplying an inverse Jacobian determinant as illustrated in Fig. 12. Akin to Equations 18 and 19, $D(\mathbf{i}_2)$ and its corresponding ENDF $\hat{D}(\mathbf{m})$ can be further approximated by vMF distributions according to the spherical convolution rule.

Correctly solving the integral for one-bounce subsurface reflection also requires proper handling of the remaining terms in the microfacet-based BSDF, i.e., M_r and M_t in Equations 1 and 3. As these remaining terms are rather smooth compared with the NDF and the ENDF, they can be pulled out of the integral safely. In our implementation,

these remaining terms are simply estimated using the mean directions of vMF distributions. For instance, $\mathbf{m}$ is replaced with $\boldsymbol{\mu}_{\mathbf{m}}$, $\mathbf{i}$ is replaced with $\boldsymbol{\mu}_{\mathbf{i}}$, etc. Furthermore, the remaining terms for each light-surface interaction should be multiplied together to give final result. To summarize, the reliable remaining term for one-bounce subsurface reflection is $M_t(\mathbf{o},\boldsymbol{\mu}_{\mathbf{t}_1})M_r(-\boldsymbol{\mu}_{\mathbf{t}_1},\boldsymbol{\mu}_{\mathbf{t}_2})M_t(-\boldsymbol{\mu}_{\mathbf{t}_2},\boldsymbol{\mu}_{\mathbf{i}_2})$.

5.2. Multi-Bounce Subsurface Reflection

The analysis of one-bounce subsurface reflection can be extended to multiple bounces. Observing that multi-bounce subsurface reflection contains repeated multiple internal reflection terms before light exits the layer, and this recursive process, together with the rough boundaries of the layer, tends to make the radiation distribution inside the medium roughly diffuse.

In addition, with the increase of iteration number, the exitant radiance degrades rapidly due to the existence of Fresnel effects at the upper boundary. As clearly shown in Fig. 13, the net Fresnel effect in two-bounce subsurface reflection is very low compared with that in once-bounce case, and this effect falls off rapidly as the number of bounce increases. Therefore, it can be concluded that the case of once-bounce contributes most in the final effect of subsurface reflection.

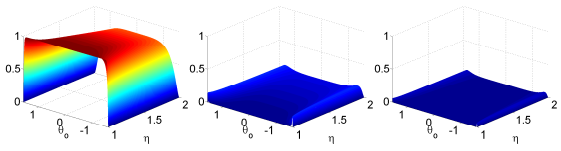

(a) one-bounce (b) two-bounce (c) three-bounce

Figure 13: *Fresnel reflectance as a function of incident angle ($\theta_{\mathbf{o}}$) and index of refraction (η).*

5.3. Light Absorption

Our layered BSDF model, like Weidlich and Wilkie's [WW07], is able to handle light absorption inside the medium. Conceivably, light absorption inside the layer will not alter the ENDF of subsurface reflection, since absorption only reduce the intensity of each light ray. Therefore, for one-bounce subsurface reflection we only have to modify the remaining term $M_t(\mathbf{o},\boldsymbol{\mu}_{\mathbf{t}_1})M_r(-\boldsymbol{\mu}_{\mathbf{t}_1},\boldsymbol{\mu}_{\mathbf{t}_2})M_t(-\boldsymbol{\mu}_{\mathbf{t}_2},\boldsymbol{\mu}_{\mathbf{i}_2})$ by multiplying a new absorption term:

$$A(d) = e^{-\sigma_a d\left(\frac{1}{|\boldsymbol{\mu}_{\mathbf{t}_1}\cdot\mathbf{n}|}+\frac{1}{|\boldsymbol{\mu}_{\mathbf{t}_2}\cdot\mathbf{n}|}\right)} \tag{27}$$

where d is the depth of the medium layer, and σ_a is the absorption coefficient (m^{-1}). Multi-bounce subsurface reflection can be treated in a similar way with repeated multiplying absorption terms along the internal reflection path.

5.4. Importance Sampling

Importance sampling scheme for our BSDF model with ENDF is straightforward. For surface reflection, we can use the conventional sampling scheme designed for microfacet-based BSDF, while for subsurface reflection, we draw samples from the probability density functions derived directly from the ENDFs.

One strength of our BSDF model is that the sampling process can be reduced to only once for each bounce of subsurface reflection if its ENDF is already known, circumventing the need to explicitly sample the corresponding NDF for each light-surface interaction (e.g., computing one-bounce subsurface reflection requires three importance sampling processes).

Furthermore, our ENDF representation has an additional advantage in the challenging lighting condition where the light source is very small (like point light). It has been widely acknowledged that traditional importance sampling is notoriously inefficient for specular-diffuse-specular light interactions, since there is a extremely large chance that a path fails to hit the light source. In dealing with layered materials, such paths occur when the outmost boundary is nearly specular (i.e., $\kappa_{\mathbf{m}_1}$ is very large). In this case, subsurface reflection effects are very difficult to capture using recursive path sampling scheme. However, with our BSDF model we can simply connect the eye sub-path to a point sample chosen from the light source to obtain correct subsurface reflection, avoiding inefficient path sampling. An example of including point light sources is provided in Fig. 1(c).

5.5. Real-Time Rendering

Another potential benefit of our BSDF model is its capability of being used in the real-time rendering applications. One such example can be found in real-time rendering with pre-filtered environmental lighting.

As suggested by Kautz et al. [KVHS00], we first convolve the environment map with vMF distributions of decreasing concentration parameter, and then store the results into a mip-map of 2D texture. During the rendering time, once we respectively obtain the directional distributions of incident light for surface reflection and subsurface reflection inside a screen pixel, we index into the pre-convolved environment map with mean direction served as texture coordinates and concentration parameter served as mip-map level. The returned value is further attenuated by appropriate remaining terms, such as Fresnel coefficients and masking and shadowing effects. It should be noted that the environment map pre-filtering can be performed on the fly as well, for example using summed area table [HSC*05], such that dynamic environmental lighting is also supported.

6. Comparisons and Results

We have implemented the proposed BSDF model in C++ and integrated it into a Monte Carlo path tracing based rendering system, and we have also implemented a GPU version using OpenGL and GLSL to enable real-time rendering. All results are produced on a PC configured with an Intel Core Q8300 CPU, 8G RAM and an NVIDIA GeForce GTS 250 graphics card, using a final output resolution of 512×512. Unless mentioned otherwise, the relative index of refraction η is set to 1.5, since it is a good approximation for medium in practical use. The sampling rate of path tracing is 1,600 spp except where explicitly noted.

Visual Quality Comparisons. We first verify the correctness of our method, and make comparisons with previous solutions. In Fig. 14, we show a rough plane rendered with the traditional microfacet BSDF model of Walter et al. [WMLT07] and the approach presented in this paper, as well as the ground truth solution. The material structures and the boundary roughness are displayed on the top of each image. As already mentioned, traditional microfacet BSDF model only correctly handles surface reflection, while our model also accounts for subsurface reflection. Although Fig. 14(b) is rendered with the same roughness as the bottom surface of the layer, the effect of subsurface reflection in Fig. 14(c) is bit softer in comparison. This is to be expected since the perceived roughness of subsurface is determined by both boundaries, not just the bottom one. In this figure, both our method and the ground truth solution successfully capture surface reflection and one-bounce subsurface reflection, but our method gains a two-fold increase in speed on the same platform (ours: 19.7 min; g.t.: 34.8 min).

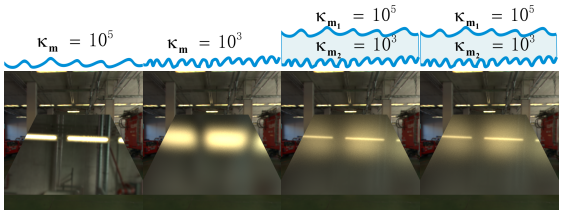

(a) [WMLT07] (b) [WMLT07] (c) ours (d) g.t.

Figure 14: *Comparison with traditional microfacet BSDF model ([WMLT07]) and the ground truth solution (g.t.).*

To further explain the accuracy of the proposed layered BSDF model, we run a step-by-step comparison against the ground truth solution that relies on recursive sampling. From Fig. 15(a) to Fig. 15(c), we show the separate results of one-bounce subsurface reflection and two-bounce subsurface reflection, as well as the total effects including surface reflection. In all cases, our approximated results match the ground truth results quite well. As expected, one-bounce subsurface reflection dominates the final results, and it maintains a high degree of directionality. Apparently, such subsurface reflection effect cannot be captured by a simple diffuse term. The timing performance for different bounces of subsurface reflection is listed in Table. 2, and we can see that the speedup is linear in terms of internal reflection bounces.

Table 2: *The illustrations and comparisons of timing performance in seconds for different bounces of subsurface reflection in Fig. 15.*

bounces	one	two	three	four
g.t.	561	868	1,194	1,519
ours	338	346	349	355
speedup	**1.66**	**2.51**	**3.42**	**4.28**

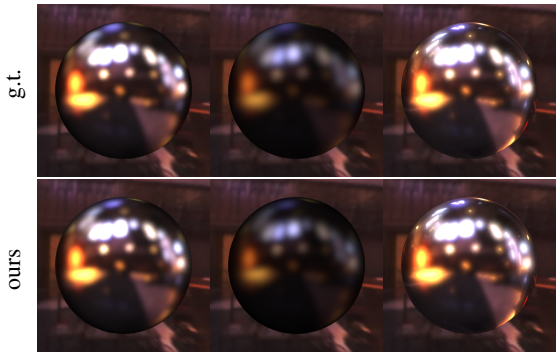

(a) one-bouce (b) two-bounce (c) total

Figure 15: *Detailed comparison against the ground truth solution. Top row: the ground truth results (g.t.); Bottom row: our approximated results. Here* $\kappa_{\mathbf{m}_1} = 10^4$ *and* $\kappa_{\mathbf{m}_2} = 10^3$.

Results of Light Absorption. In Fig. 16, we show examples of rendering results with light absorption, and we also compare our method with the method of Weidlich and Wilkie [WW07]. Still, our method achieves high image quality and the result is very close to the ground truth. Since Weidlich and Wilkie's model is based on the assumption that the refracted light beam from the first layer interface is also collimated, the glossiness for subsurface reflection is not quite correct.

Results of Complex Scenes. In Fig. 1, we show several complex scenes rendered with our method. Every object in these images is coated by a thin transparent layer with different roughness on both sides. Clearly, our model can work with any microfacet distribution. Moreover, our method also supports spatially varying roughness as shown in Fig. 1(b) and small light sources as shown in Fig. 1(c).

Real-Time Rendering with Pre-Filtered Environmental Lighting. Fig. 17 demonstrates the usage of our BSDF model in real-time rendering applications. Here we show a bunny model lit by a distance environmental light. Using

(a) [WW07] (b) ours (c) g.t.

Figure 16: *Rendering results of thin layers with light absorption. For the kitten model (top),* $\kappa_{\mathbf{m}_1} = \kappa_{\mathbf{m}_2} = 10^3$, $\sigma_a = [0.2, 0.8, 0.8]$, $d = 0.5$, *while for the horse model (bottom),* $\kappa_{\mathbf{m}_1} = 10, \kappa_{\mathbf{m}_2} = 10^6$, $\sigma_a = [0.8, 0.8, 0.2]$, $d = 0.5$.

pre-filtered environment map, we can easily achieve real-time performance while still maintaining the phenomenological correctness compared with ground truth. For this real-time rendering application, we obtain a sustained rendering frame rate of around 700 fps, while the path traced ground truth takes 1,085 seconds (with 400 spp). Please refer to the accompanying video for an animated version of this figure.

(a) real-time rendering (b) path tracing

Figure 17: *Image quality comparison between our real-time implementation and the path traced ground truth. Here* $\kappa_{\mathbf{m}_1} = 10^4$, $\kappa_{\mathbf{m}_2} = 10$, *and* $\eta = 2$.

Real-Time Appearance Editing: Since our BSDF model does not require any cumbersome precomputation, and the ENDF can be evaluated on-the-fly, the roughness of each layer boundary can be edited in real-time. This is shown in Fig. 18, where we see the rendering results of a dragon model (100K triangles) with varying surface roughness on both sides. We simply adjust the concentration parameters of two vMF distributions ($\kappa_{\mathbf{m}_1}$ and $\kappa_{\mathbf{m}_2}$, respectively), and our system is able to give the corresponding shading results in around 1.3 ms. More appearance editing results can be found in the accompanying video.

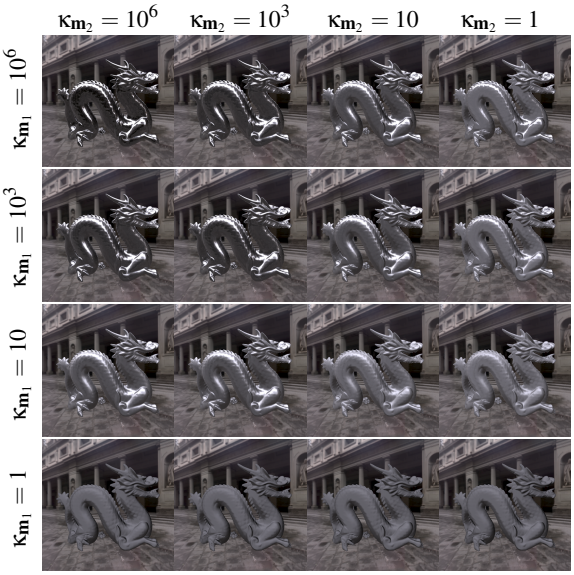

Figure 18: *.Real-time appearance editing examples. From up to down we show the results of changing the first surface's roughness* $\kappa_{\mathbf{m}_1}$, *while from left to right we show the results of changing the second surface's roughness* $\kappa_{\mathbf{m}_2}$.

7. Conclusions and Future Work

To summarize, this paper has presented a framework for the practical rendering of thin transparent layers with both surface reflection and subsurface reflection. The main insight of our work is the extension of the traditional microfacet BSDF model to handle layered rough surfaces. After analyzing the dependence of visually perceived roughness for the subsurface reflection on the material properties of each layer boundary, we employed the ENDF, a new representation designed for depicting the appearance behavior of the subsurface reflection. Unlike previous methods that capture multiple internal reflections using recursive importance sampling, our approach based on the ENDF requires only one sampling process for each bounce of subsurface reflection, and it can also efficiently handle small light sources. Experimental results demonstrate that our BSDF model is applicable to both offline and real-time rendering.

In our current implementation, we use an isotropic vMF lobe to fit both NDF and ENDF. This will cause large anisotropic distortion at grazing angle, as shown in Fig. 10. To reduce such error, we would like to use multi-lobe distributions, such as movMF, or use anisotropic basis functions in our future work. Moreover, in the future we also would like to investigate the influence of light scattering on the shape of the ENDF when dealing with translucent layers.

Acknowledgements

We would like to thank the anonymous reviewers for their valuable feedback. This work is supported by China Postdoctoral Science Foundation funded project (2014M551560) and Postdoctoral Science Foundation of Jiangsu Province (1402029B).

References

[APS00] ASHIKMIN M., PREMOŽE S., SHIRLEY P.: A microfacet-based brdf generator. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 2000), SIGGRAPH '00, ACM Press/Addison-Wesley Publishing Co., pp. 65–74. 2

[Bli82] BLINN J. F.: Light reflection functions for simulation of clouds and dusty surfaces. *SIGGRAPH Comput. Graph. 16*, 3 (July 1982), 21–29. 2

[CT82] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. *ACM Trans. Graph. 1*, 1 (Jan. 1982), 7–24. 2, 4

[DH96] DORSEY J., HANRAHAN P.: Modeling and rendering of metallic patinas. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques* (1996), SIGGRAPH '96, pp. 387–396. 3

[DJ05] DONNER C., JENSEN H. W.: Light diffusion in multi-layered translucent materials. *ACM Trans. Graph. 24*, 3 (July 2005), 1032–1039. 3

[DWL*09] DAI Q., WANG J., LIU Y., SNYDER J., WU E., GUO B.: The dual-microfacet model for capturing thin transparent slabs. *Computer Graphics Forum 28*, 7 (2009), 1917–1925. 2, 7

[EKM01] ERSHOV S., KOLCHIN K., MYSZKOWSKI K.: Rendering pearlescent appearance based on paint-composition modelling. *Computer Graphics Forum 20*, 3 (2001), 227–238. 3

[Fou92] FOURNIER A.: Normal distribution functions and multiple surfaces. In *Proceedings of Graphics Interface Workshop on Local Illumination* (1992), pp. 45–52. 2

[GP14] GUO J., PAN J.-G.: Real-time simulating and rendering of layered dust. *The Visual Computer 30*, 6-8 (2014), 797–807. 3

[Hd14] HEITZ E., D'EON E.: Importance Sampling Microfacet-Based BSDFs using the Distribution of Visible Normals. *Computer Graphics Forum 33*, 4 (July 2014), 103–112. 2

[Hei14] HEITZ E.: Understanding the masking-shadowing function in microfacet-based brdfs. *Journal of Computer Graphics Techniques (JCGT) 3*, 2 (June 2014), 32–91. 2, 4

[HK93] HANRAHAN P., KRUEGER W.: Reflection from layered surfaces due to subsurface scattering. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques* (1993), SIGGRAPH '93, pp. 165–174. 2

[HSC*05] HENSLEY J., SCHEUERMANN T., COOMBE G., SINGH M., LASTRA A.: Fast summed-area table generation and its applications. *Comput. Graph. Forum 24*, 3 (2005), 547–555. 9

[IDN12] IWASAKI K., DOBASHI Y., NISHITA T.: Interactive bi-scale editing of highly glossy materials. *ACM Trans. Graph. 31*, 6 (Nov. 2012), 144:1–144:7. 2

[JdJM14] JAKOB W., D'EON E., JAKOB O., MARSCHNER S.: A comprehensive framework for rendering layered materials. *ACM Trans. Graph. 33*, 4 (July 2014), 118:1–118:14. 2, 3

[JMLH01] JENSEN H. W., MARSCHNER S. R., LEVOY M., HANRAHAN P.: A practical model for subsurface light transport. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 2001), SIGGRAPH '01, ACM, pp. 511–518. 3

[KVHS00] KAUTZ J., VÁZQUEZ P.-P., HEIDRICH W., SEIDEL H.-P.: Unified approach to prefiltered environment maps. In *Proceedings of the Eurographics Workshop on Rendering Techniques 2000* (London, UK, UK, 2000), Springer-Verlag, pp. 185–196. 9

[MHM*13] MCAULEY S., HILL S., MARTINEZ A., VILLEMIN R., PETTINEO M., LAZAROV D., NEUBELT D., KARIS B., HERY C., HOFFMAN N., ZAP ANDERSSON H.: Physically based shading in theory and practice. In *ACM SIGGRAPH 2013 Courses* (New York, NY, USA, 2013), SIGGRAPH '13, ACM. 2, 4

[MJ00] MARDIA K. V., JUPP P. E.: *Directional statistics*. Wiley series in probability and statistics. Wiley, Chichester, 2000. Previous ed. published as: Statistics of directional data. London : Academic Press, 1972. 3, 6

[NDM05] NGAN A., DURAND F., MATUSIK W.: Experimental analysis of brdf models. In *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques* (Aire-la-Ville, Switzerland, Switzerland, 2005), EGSR '05, Eurographics Association, pp. 117–126. 2, 3

[PdMJ14] PAPAS M., DE MESA K., JENSEN H. W.: A physically-based bsdf for modeling the appearance of paper. *Computer Graphics Forum 33*, 4 (2014), 133–142. 3

[PH00] PHARR M., HANRAHAN P.: Monte carlo evaluation of non-linear scattering equations for subsurface reflection. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques* (2000), SIGGRAPH '00, pp. 75–84. 2, 3

[Pre02] PREMOŽE S.: Analytic light transport approximations for volumetric materials. In *Proceedings of the 10th Pacific Conference on Computer Graphics and Applications* (Washington, DC, USA, 2002), PG '02, IEEE Computer Society, pp. 48–. 3

[Smi67] SMITH B.: Geometrical shadowing of a random rough surface. *IEEE Trans. on Antenna and Propagation 15* (1967), 668–671. 2

[SSHL97] SHIRLEY P., SMITS B. E., HU H. H., LAFORTUNE E. P.: A practitioners' assessment of light reflection models. In *5th Pacific Conference on Computer Graphics and Applications (PG '97)* (1997), IEEE Computer Society, p. 40. 2

[Sta01] STAM J.: An illumination model for a skin layer bounded by rough surfaces. In *Proceedings of the 12th Eurographics Conference on Rendering* (Aire-la-Ville, Switzerland, Switzerland, 2001), EGWR'01, Eurographics Association, pp. 39–52. 2

[WDR11] WU H., DORSEY J., RUSHMEIER H.: Physically-based interactive bi-scale material design. In *Proceedings of the 2011 SIGGRAPH Asia Conference* (2011), SA '11, pp. 145:1–145:10. 2

[WMLT07] WALTER B., MARSCHNER S. R., LI H., TORRANCE K. E.: Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques* (2007), EGSR'07, pp. 195–206. 2, 3, 5, 10

[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, 3, 5

[WW07] WEIDLICH A., WILKIE A.: Arbitrarily layered micro-facet surfaces. GRAPHITE '07, pp. 171–178. 3, 9, 10, 11