

Microfacet BSDFs Generated from NDFs and Explicit Microgeometry

MICKAËL RIBARDIÈRE, University of Poitiers CNRS, XLIM, UMR 7252, France
BENJAMIN BRINGIER, University of Poitiers CNRS, XLIM, UMR 7252, France
LIONEL SIMONOT, University of Poitiers CNRS UPR 3346 Institut Prime, France
DANIEL MENEVEAUX, University of Poitiers CNRS, XLIM, UMR 7252, France




Fig. 1. Example of microfacet-based BSDFs, with very different appearances created from four configurations of Normal Distribution Functions (illustrated in insets). Microfacet properties are: Lambertian floor, perfect reflector elephant, gold conductor sphere, glass dielectric vase, and copper conductor cup. For each configuration: First row provides a rendering with multiple scattering estimated from our surface generation process; Second row provides a rendering without multiple scattering and a $CIE \Delta E_{00}$ color-difference image illustrating appearance changes between single and multiple scattering images (false color scale is provided in Figure 12).

Authors' addresses: Mickaël Ribardière, University of Poitiers, CNRS, XLIM, UMR 7252, Poitiers, France, mickael.ribardiere@univ-poitiers.fr; Benjamin Bringier, University of Poitiers, CNRS, XLIM, UMR 7252, Poitiers, France, benjamin.bringier@univ-poitiers.fr; Lionel Simonot, University of Poitiers, CNRS UPR 3346 Institut Prime, Poitiers, France, lionel.simonot@univ-poitiers.fr; Daniel Meneveaux, University of Poitiers, CNRS, XLIM, UMR 7252, Poitiers, France, daniel.meneveaux@univ-poitiers.fr.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Association for Computing Machinery.
XXXX-XXXX/2019/5-ART \$15.00
<https://doi.org/10.1145/nmnnnnn.nmnnnnn>

Microfacet distributions are considered nowadays as a reference for physically plausible BSDF representations. Many authors have focused on their physical and mathematical correctness, while introducing means to enlarge the range of possible appearances. This paper is dedicated to Normal Distribution Functions (NDFs), and the influence of their shape on the rendered material aspect. We provide a complete framework for studying the impact of NDFs on the observed Bidirectional Scattering Distribution Functions (BSDFs). In order to explore very general NDFs, manually controlled by the user, and including anisotropic materials, we propose to use a piecewise continuous representation. It is derived with its associated Smith shadowing-masking function and importance sampling formulations for ensuring efficient global illumination computations. A new procedure is also proposed in this paper for generating an explicit geometric micro-surface, used to evaluate the validity of analytic models and multiple scattering effects. The results are produced with a computer-generated process using path tracing.

They show that this generation procedure is suitable with any NDF model, independently from its shape complexity.

CCS Concepts: • **Computing methodologies** → *Reflectance modeling*;

Additional Key Words and Phrases: Material Appearance, Microfacet BSDF, Microsurface generation, Multiple scattering, Smith model

ACM Reference Format:

Mickaël Ribardière, Benjamin Bringier, Lionel Simonot, and Daniel Meneveaux. 2019. Microfacet BSDFs Generated from NDFs and Explicit Microgeometry. 1, 1 (May 2019), 15 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

An accurate observation of real-world surfaces reveals many effects in terms of anisotropy, glossiness, or reflection sharpness for instance. Computer-generated images rely on Bidirectional Scattering Distribution Functions (BSDFs) to mimic appearance effects of real-world surfaces. This area has motivated much research in the last four decades, and impressive results have been achieved with the proposed models. However, when compared to measurements, they are far from covering the very wide range of possible real surface aspects. Among the existing families of models, microfacet-based BSDFs have been favored by many authors for several reasons: They are able to represent many types of materials [Heitz 2014; Meneveaux et al. 2017; Walter et al. 2007; Zeltner and Jakob 2018]; They rely on physical aspects of the surface such as roughness or refractive index; They can be manipulated intuitively and with only few parameters. This representation is built upon a statistically roughened surface using a microfacet Normal Distribution Function (NDF) [Beckmann and Spizzichino 1963; Cook and Torrance 1982; Heitz 2014; Ribardière et al. 2017; Walter et al. 2007]. A BSDF f^{μ} is associated with microfacets, often considered as purely specular reflective [Bagher et al. 2012; Cook and Torrance 1982], though some authors have also studied other local BSDFs [Meneveaux et al. 2017; Oren and Nayar 1994; Walter et al. 2007].

As demonstrated in this paper and illustrated in Figure 1, choosing microfacets BSDF f^{μ} and NDF offers a large degree of freedom for controlling material appearance. However, only few NDFs are proposed in the literature, all corresponding to bell-shaped distributions [Beckmann and Spizzichino 1963; Ribardière et al. 2017; Trowbridge and Reitz 1975; Walter et al. 2007]. They all favor only one specular lobe corresponding to the same general normal orientation, in most cases centered on the normal direction or slightly shifted [Bagher et al. 2012]. Unfortunately, these bell-shaped functions are generally not adapted to the representation of manufactured surfaces [Halley 2012].

Nevertheless, statistical surface roughness representations using NDFs are convenient to define microfacet-based materials with Smith assumptions, considering that microfacets slopes are not correlated (defining a masking-shadowing model that accounts for a possible height correlation on the surface is still an open problem). This paper proposes an in-depth study dedicated to the influence of surface roughness on material appearance. We first review and classify the mostly used NDFs, which are actually sub-configurations of a more general distribution function relying on the *Skewed Generalized T-Distribution* (SGTD). We also present a complete framework

based on fully controllable NDFs, defined by a piecewise-linear representation (as discrete anchor points chosen by the user, or sampled from any analytical function). Compliance on physical laws and rendering constraints are handled in this analysis: We derive the mathematical formulations for Smith's shadowing-masking term (also called Geometric Attenuation Function or GAF) [Smith 1967] and importance sampling, both mandatory for defining and using any microfacet BSDF [Ashikhmin et al. 2000; Bourlier et al. 2002; Heitz 2014]. In addition, microfacet BSDF models account for direct reflections only, and handling light multiple scattering remains a difficult challenge. Yet, they play a key role in the appearance, especially for very rough surfaces [Heitz et al. 2016; Kulla and Conty 2017; Lagarde 2017]. We also study their impact in this paper, thanks to a new method that generates a plausible microsurface geometry from any isotropic or anisotropic NDF. The generated mesh can be used for precisely precomputing multiple scattering, allowing in-depth studies of these effects, as well as their integration in a physically based rendering engine. More precisely, the contributions proposed in this paper are the following.

- An exhaustive review and a classification of NDFs, and their usage in rendering lighting simulation and rendering engine.
- An anisotropic piecewise-linear representation of NDFs, and the associated derivations for defining the generalized Smith GAF [Bourlier et al. 2002] and importance sampling.
- A new method that generates a plausible geometric mesh from any NDF, compliant with Smith's shadowing and masking assumptions; This method is flexible enough to handle geometric constraints for controlling the final height field aspect.
- An analysis of multiple-scattering effects, based on a virtual gonioreflectometer that gathers the resulting reflected light distribution. The resulting data can be used in a Monte Carlo based rendering method.
- An in-depth study on the relationship between surface roughness and material appearance.

As shown in the results and supplemental material, a wider control of NDF shapes extends the range of appearances produced by microfacet-based BSDFs. Our surface generation process offers a useful tool for the evaluation of future insights in the definition of multiple-scattering functions and more generally microfacet-based BSDFs. It also extends the control of microsurface generation and the associated material aspect.

The remainder of this paper is organized as follows: Section 2 introduces the theoretical background and an overview of our framework. We also provide a classification of the conventional NDFs. Section 3 presents the mathematical derivations for shadowing-masking effects and importance sampling with piecewise-linear NDFs. Section 4 describes our surface generation process used to estimate and discuss multiple-scattering effects. Section 5 presents the results obtained with our methodology: We evaluate the validity of our framework and study relationships between NDF shape and material appearance. Finally, Section 6 discusses future insights on microfacet-based material representations.

2 BACKGROUND AND OVERVIEW

Let us consider a surface sample dS of normal \mathbf{n} , lit by a collimated light source from direction \mathbf{i} , as illustrated in Figure 2. The radiance reflected toward an observer direction \mathbf{o} is given by the Bidirectional Reflectance Distribution Function (or BRDF) $f(\mathbf{i}, \mathbf{o}, \mathbf{n})$. It is defined as the ratio between the radiance $L(\mathbf{i}, \mathbf{o}, \mathbf{n})$ reflected by dS in direction \mathbf{o} and the incident irradiance $E(\mathbf{i}, \mathbf{n})$ coming from direction \mathbf{i} .




Fig. 2. Geometry of reflection and notations.

A BRDF model is considered as *physically plausible*, if it satisfies Helmholtz reciprocity $f(\mathbf{i}, \mathbf{o}, \mathbf{n}) = f(\mathbf{o}, \mathbf{i}, \mathbf{n})$, and energy conservation $\int_{\Omega_+} f(\mathbf{i}, \mathbf{o}, \mathbf{n}) |\mathbf{o} \cdot \mathbf{n}| d\omega_o \leq 1, \forall \mathbf{i} \in \Omega_+$.

The microfacet representation of a BRDF $f(\mathbf{i}, \mathbf{o}, \mathbf{n})$ corresponds to a statistical description of a microfacet distribution. Given the BRDF $f^\mu(\mathbf{i}, \mathbf{o}, \mathbf{m})$ of an individual microfacet associated with a normal \mathbf{m} , its contribution is weighted by the distribution $D(\mathbf{m})$ and a geometric attenuation factor $G(\mathbf{i}, \mathbf{o}, \mathbf{m})$. $D(\mathbf{m})$ also defines the surface roughness, indicating the proportion of microfacets corresponding to a normal \mathbf{m} ; The attenuation factor $G(\mathbf{i}, \mathbf{o}, \mathbf{m})$ determines the portion of a microfacet of normal \mathbf{m} visible from both the light source and the observer. It has a major influence at grazing angles. Many authors have studied the use of various distributions and geometric attenuation factors [Ashikhmin et al. 2000; Bagher et al. 2012; Bourlier et al. 2002; Cook and Torrance 1982; Dupuy et al. 2013; Kelemen and Szirmay-Kalos 2001; Oren and Nayar 1994; Torrance and Sparrow 1967; Walter et al. 2007], which have to be carefully chosen together [Ashikhmin et al. 2000; Heitz 2014]. The equation for the general case of microfacet-based BRDFs is:

$$f(\mathbf{i}, \mathbf{o}, \mathbf{n}) = \int_{\Omega_+} \frac{|\mathbf{i} \cdot \mathbf{m}|}{|\mathbf{i} \cdot \mathbf{n}|} f^\mu(\mathbf{i}, \mathbf{o}, \mathbf{m}) \frac{|\mathbf{o} \cdot \mathbf{m}|}{|\mathbf{o} \cdot \mathbf{n}|} D(\mathbf{m}) G(\mathbf{i}, \mathbf{o}, \mathbf{m}) d\omega_m. \quad (1)$$

Microfacets are supposed to be oriented toward the upper hemisphere ($D(\mathbf{m}) = 0$ if $\mathbf{m} \cdot \mathbf{n} \leq 0$). The projected areas of all microfacets have to be equal to the macroscopic surface:

$$\int_{\Omega_+} D(\mathbf{m}) |\mathbf{m} \cdot \mathbf{n}| d\omega_m = 1.$$

In the case of purely specular microfacets [Cook and Torrance 1982; Torrance and Sparrow 1967; Walter et al. 2007], Equation 1 simplifies to:

$$f(\mathbf{i}, \mathbf{o}, \mathbf{n}) = \frac{F(\mathbf{i}, \mathbf{h}) D(\mathbf{h}) G(\mathbf{i}, \mathbf{o}, \mathbf{h})}{4 |\mathbf{i} \cdot \mathbf{n}| |\mathbf{o} \cdot \mathbf{n}|}, \quad (2)$$

where $\mathbf{h} = \frac{\mathbf{i} + \mathbf{o}}{|\mathbf{i} + \mathbf{o}|}$ is the half-angle vector between \mathbf{i} and \mathbf{o} , and $F(\mathbf{i}, \mathbf{h})$ corresponds to Fresnel's reflectance, depending on n_i , the relative

refractive index between the material and the exterior medium refractive index (generally the air with $n_i = 1$). In other cases, finding a simplified formulation of Equation 1 is difficult due to scattering effects [Meneveaux et al. 2017; Oren and Nayar 1994].

Shadowing and Masking

The geometrical attenuation function $G(\mathbf{i}, \mathbf{o}, \mathbf{m})$ accounts for self-masking and self-shadowing. The widely used function proposed by Torrance and Sparrow [Torrance and Sparrow 1967] makes the assumption that the microsurface corresponds to a set of two dimensional V-cavity profiles. This model is mathematically consistent but physically unrealistic [Heitz 2014]. Nowadays, it is commonly admitted that Smith's shadowing-masking term [Bourlier et al. 2002; Smith 1967] is closest to the physical behavior of rough surfaces. Shadowing and masking are considered as independent, and $G(\mathbf{i}, \mathbf{o}, \mathbf{m})$ is thus approximated using the product of the same two functions G_1 :

$$G(\mathbf{i}, \mathbf{o}, \mathbf{m}) = G_1(\mathbf{i}, \mathbf{m}) G_1(\mathbf{o}, \mathbf{m}). \quad (3)$$

The most important assumption in Smith's term is that microfacet normals are not correlated, even in close proximity. Mathematically, this assumption can be written as follows:

$$G_1(\mathbf{v}, \mathbf{m}) = \begin{cases} G_1(\mathbf{v}) & \mathbf{v} \cdot \mathbf{m} \geq 0 \\ 0 & \mathbf{v} \cdot \mathbf{m} < 0 \end{cases} \quad (4)$$

It has been used by Ashikhmin et al. [2000] to derive the following expression:

$$G_1(\mathbf{v}) = \frac{(\mathbf{v} \cdot \mathbf{n})}{\int_{\Omega_+(\mathbf{v})} (\mathbf{v} \cdot \mathbf{m}) D(\mathbf{m}) d\omega_m}, \quad (5)$$

but the denominator integral remains difficult to estimate. One solution is to precompute and tabulate G_1 for a fixed set of values $\mathbf{v} = (\theta_v, \varphi_v)$.

More convenient expressions of this function can be obtained starting from the work of Bourlier et al. [2002]. Walter et al. [2007] and more recently Heitz [2014] have expressed the normal distribution in slopes space $P_{22}(p, q)$:

$$P_{22}(p, q) = \cos^4 \theta_m D(\mathbf{m}), \quad (6)$$

where p and q correspond to the microsurface slopes in the macro-surface local coordinate system, defining the normal \mathbf{m} , such that $p^2 + q^2 = \tan^2 \theta_m$. The one-dimensional distribution of slopes is given by:

$$P_2(q) = \int_{-\infty}^{+\infty} P_{22}(p, q) dp. \quad (7)$$

Finally, $G_1(\mathbf{v})$ is obtained by the integration of P_2 :

$$G_1(\mathbf{v}) = \frac{1}{1 + \Lambda(\mathbf{v})}, \quad (8)$$

where $\Lambda(\mathbf{v}) = \frac{1}{\mu} \int_{-\infty}^{+\infty} (q - \mu) P_2(q) dq$ and $\mu = \cot \theta_v$. A formal proof of these derivations is given by Heitz [2014].

Note that this shadowing and masking term is used by most rendering systems, though Ross et al. [2005] or Heitz et al. [2013] express correlated versions between masking and shadowing, considered as physically more plausible. The following height-correlated

formulation is considered as more accurate for masking and shadowing, with an equivalent computational cost:

$$G(\mathbf{i}, \mathbf{o}, \mathbf{m}) = \frac{1}{1 + \Lambda(\mathbf{i}) + \Lambda(\mathbf{o})}. \quad (9)$$

Our study is based on Smith's assumptions, where the shadowing-masking term directly depends on the chosen distribution. Thus, the distribution function has to be carefully chosen, so that an analytical expression can be derived for G_1 (Equation 5) or Λ (Equations 8 and 9).

2.1 Normal Distribution Function

This paper focuses on how the normal distribution impacts the appearance of materials described by a microfacet-based BSDF. Controlling the NDF shape is mandatory for managing more complex appearances, as pointed out by several authors [Bagher et al. 2012; Burley 2012; Butler and Marciniak 2014; Dong et al. 2015; Hoffman 2016; Holzschuch and Pacanowski 2017; Ribardière et al. 2017].

Moreover, to be efficiently included in rendering systems, an NDF has to come with an analytical derivation for Smith's shadowing-masking term and importance sampling schemes. The latter should be ideally implemented with visible normal importance sampling as proposed by Heitz and d'Eon [2014]. To our knowledge, only Beckmann's [Beckmann and Spizzichino 1963] and GGX [Walter et al. 2007] distributions offer all these features. This is why they are currently considered as standards in computer graphics.

They are controlled by one roughness parameter σ that drives the lobe width. Unfortunately, real material normal distributions exhibit more complex shapes, as mentioned by Holzschuch and Pacanowski [2017], who also propose a Generalized Beckmann's Distribution (GND). GND includes an additional parameter γ that characterizes the lobe's flattening (called here *tailedness*), defining the drop at the origin. Ribardière et al. [Ribardière et al. 2017] have used an even more general NDF called *Student's T-Distribution* (*Student's TD*), also controlled using an additional *tailedness* parameter γ . *Student's TD* encompasses GGX when $\gamma = 2$ and tends to Beckmann's when $\gamma \rightarrow \infty$. They also derive a more general formulation that allows to include the *Hyper-Cauchy Distribution* (HCD) as a specific configuration of *Student's TD*. HCD was first introduced by Wellems et al. [Wellems et al. 2006] and used to fit measured BRDFs by Butler and Marciniak [2014]. All these configurations control the shape through two parameters: σ for the lobe width and γ for its *tailedness*. Contrary to GND, *Student's TD* comes with the analytic Λ function to derive Smith's shadowing-masking expression. An open question concerns the degree of control required on the NDF shape for fitting parameters, designing, and rendering materials, and if such a control can be reached by an existing distribution. All the previous distributions are in fact sub-families of a more general distribution, the *Skewed Generalized T-Distribution* (SGTD), introduced in financial statistics by Theodossiou [Theodossiou 1998]. Figure 3 illustrates a complete classification from Beckmann and GGX to SGTD, including *Student's TD* or GND, and *shape invariant* sub-configurations. *Shape invariance*¹ property [Heitz 2014; Hoffman 2016] defines the invariance of both the NDF shape and the

masking function when the surface configuration is stretched. This helps to derive anisotropic formulations of the distribution. More details and discussions are provided in the supplemental material files about SGTD and its sub-configurations.

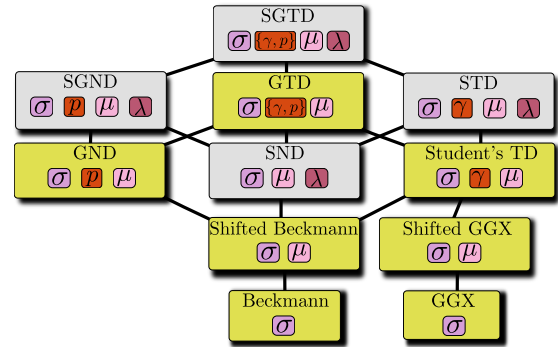


Fig. 3. *Skewed Generalized T-Distribution* [Wikipedia 2018] (SGTD) tree: Yellow blocks represent shape invariant distributions coming with anisotropy.

2.2 Discussion and Methodology Overview

The NDF is a key element in a microfacet-based material appearance model. It is currently controlled by one parameter (the standard deviation for Beckmann's distribution or GGX) or two (the tailedness for *Student's TD* or GND). These parameters drive the NDF shape and consequently the microsurface roughness. The most widely used NDFs, and by extension all the SGTD sub-configurations, correspond to bell-shaped distributions and they all favor only one specular lobe, corresponding to the same general normal orientation (in most cases centered on the normal direction). Unfortunately, the measured NDFs associated with manufactured surfaces do generally not comply with this type of representation [Halley 2012].

This is why we have chosen to introduce in this paper a new framework for designing multi-lobes NDF and analyze their effect on material appearance (summarized in Figure 4). We first propose to employ a piecewise-linear NDF formulation associated to its importance sampling scheme and Smith's shadowing-masking term (Section 3). It allows to design multi-lobes NDFs as well as an approximate representation of any analytic distribution as SGTD for instance (step 1 in Figure 4).

With this methodology, very rough surfaces can be associated with microfacets with steep slopes. In this case, multiple-scattering of light (between microfacets) strongly impacts the visual appearance of a material. As noted by several authors, it should be handled otherwise the rendering process will lead to abnormally dark surfaces, due to incorrect loss of energy. A recent contribution of Heitz et al. [Heitz et al. 2016] handles these effects for Beckmann's and GGX NDFs, considering shape-invariant surfaces. Unfortunately, with a more general NDF, this assumption is not true anymore. Kulla and Conty [Kulla and Conty 2017] propose to precompute and tabulate a non physically based approximation of the lost energy, according to the NDF roughness parameter. Note that multiple scattering has been recently addressed in the context of (non physical)

¹A shape invariant NDF can be written as follows: $f\left(\frac{\tan\theta}{\sigma}\right) / (\sigma^2 \cos^4 \theta)$.

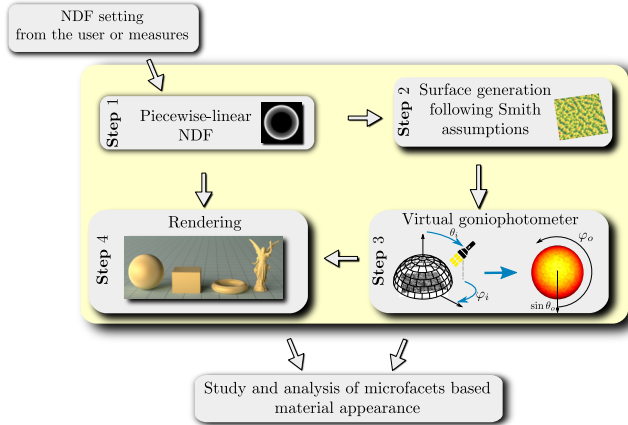


Fig. 4. Complete process proposed in our study. The bottom right drawing illustrates how the BSDF is presented in this paper (see Section 5).

V-cavity profile assumptions [Lee et al. 2018; Xie and Hanrahan 2018].

The framework proposed in this paper makes a step forward in the study of the effects modeled with such representations (in the context of Smith’s assumptions). Our method consists in explicitly generating a surface mesh (step 2 in Figure 4), consistent with any piecewise-linear NDF (Section 4), and applying lighting simulation techniques to estimate multiple light reflections. This strategy is physically based and can thus be employed for validating many types of models.

Let us consider a virtual gonioreflectometer with equal-area cells (i.e., iso-solid angle and with aspect ratio for all cells to mitigate acquisition bias). A hemispherical sensor [Beckers and Beckers 2012] is used. Similarly to a real gonioreflectometer system, a collimated light beam illuminates the microsurface from a fixed direction, defined by spherical angles θ_i and φ_i (step 3 in Figure 4). All the reflected light leaving the microsurface after one bounce (L_1), two bounces (L_2), or more (L_∞) is captured by the hemispherical sensor cells. The resulting histogram (i.e., values gathered in each sensor’s cell) is proportional to $f(\mathbf{i}, \mathbf{o}) \cos \theta_o$, with \mathbf{o} the center direction of one sensor’s cell. This system is straightforwardly extended for transmitted energy with a second hemispherical sensor placed below the surface. Virtual gonioreflectometers have been used for long in computer graphics [Cabral et al. 1987; Westin et al. 1992]. However, our system differs in (i) accounting for all light bounces over the surface, and (ii) using a fine hemispherical sensor instead of basis functions (spherical harmonics in [Cabral et al. 1987; Westin et al. 1992]), prone to ringing artifacts and/or require a large number of coefficients.

In practice, multiple scattering (in the following, L_{2+} corresponds to indirect lighting only, i.e., L_∞ without L_1) is challenging to manage, but the actual impact on appearance cannot be neglected. The explicitly reconstructed surface could be used directly with a Monte Carlo based engine, but the time required for final path integration would be drastically increased. Instead, we propose to tabulate the data captured by the virtual gonioreflectometer, considering all the

contributions $L_\infty = L_1 + L_{2+}$. Such a method has proven to be a good compromise and an easy to implement solution. During the rendering process, L_∞ is evaluated with a bilinear interpolation through a lookup performed in the L_∞ -table (step 4 in Figure 4). A possible step further consists in separating L_1 and L_{2+} , where L_1 is computed using the analytical microfacet BSDF, and L_{2+} is approximated by a projection of the captured data in a basis of functions (e.g., wavelets [Claustres et al. 2007; Lewis 1998], spherical/hemispherical harmonics [Gautron et al. 2004; Mahajan et al. 2008; Westin et al. 1992], or spherical Gaussians [Soler et al. 2018]) or fit to the parameters of a second BSDF model as suggested by several authors [Heitz 2014; Lagarde 2017].

Our work is related to that of Yan et al. [Yan et al. 2014] and its following extensions [Yan et al. 2016, 2018], that introduce the concept of *pixel NDF* (or \mathcal{P} -NDF). These contributions concern the management of complex surface geometry, observed through a pixel, for highly specular materials. These papers address the representation of visible surface structures as glints or scratches but they did not investigate the impact of shadowing and masking nor multiple scattering. Our paper focuses on another level of detail, where such effects are not distinguishable and smoothed in the complete BSDF. The surface is considered from an observer placed at a distance where these structures are very small compared to the size of a pixel footprint on the object surface. The management of levels of detail is not addressed in this article.

3 A PIECEWISE-LINEAR NDF

With microfacet-based BSDFs, anisotropy is directly controlled by NDF and for common models as GGX or Beckmann’s, obtained through a trigonometric interpolation of the roughness parameter. This interpolation method restricts the use to only specific distributions. Let us define an alternative representation, that can be used in the general case. An anisotropic distribution $D(\mathbf{m})$, with φ and θ polar coordinates of a given normal orientation \mathbf{m} , can be defined as a trigonometric interpolation of two normalized isotropic NDF curves: The first one (D_x) is defined for $\varphi = 0$, while the second one (D_y) is defined for $\varphi = \frac{\pi}{2}$ (Figure 5). The anisotropic NDF representation proposed in this paper is defined by:

$$D(\mathbf{m}) = \frac{\cos^2 \varphi D_x(\theta) D_y(0) + \sin^2 \varphi D_y(\theta) D_x(0)}{c}, \quad (10)$$

where $c = \frac{1}{2}(D_y(0) + D_x(0))$ is the normalization factor (please see the supplemental material for mathematical details). This formulation is different from the conventional anisotropic NDF (e.g. GGX [Heitz 2014] or Beckmann’s [Ashikhmin et al. 2000; Kurt et al. 2010]). For instance, when using Beckmann’s functions for both D_x and D_y , Equation 10 defines an NDF shape slightly different from the anisotropic Beckmann’s model (though it is exactly the same with an equal roughness for both D_x and D_y). It offers a mean to blend two different NDF models (such as GGX and Beckmann’s, or two STDs with different tailedness and roughness parameters). However, in that case, the normalization factor c is not trivial to compute and neither is the GAF.

We propose to make these computations possible with piecewise linear functions for two isotropic distributions D_x and D_y . Let us

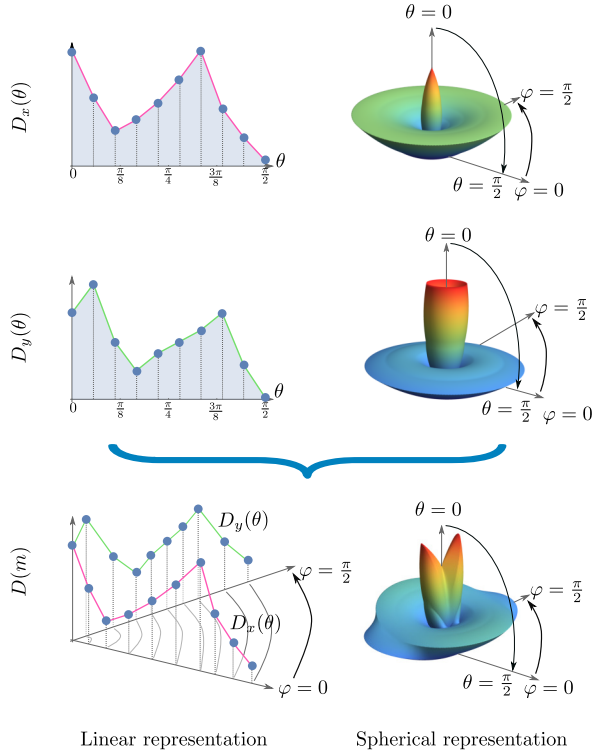


Fig. 5. Examples of piecewise-linear NDFs, with $n=10$ control points. First and second rows: Isotropic distributions with two lobes. Last row: Anisotropic NDF with multiple lobes, resulting from the two previous distributions, respectively defined for $\varphi=0$ and $\varphi=\frac{\pi}{2}$.

consider a NDF defined by a discrete curve, over $\theta \in [0; \frac{\pi}{2}]$ with n control points (Figure 5). The isotropic NDF $D_{iso}(\theta)$ is a piecewise continuous function defined by:

$$D_{iso}(\theta) = \begin{cases} a_1\theta + b_1 & \theta_0 = 0 \leq \theta \leq \theta_1 \\ a_2\theta + b_2 & \theta_1 \leq \theta \leq \theta_2 \\ \dots & \\ a_n\theta + b_n & \theta_{n-1} \leq \theta \leq \theta_n = \frac{\pi}{2}, \end{cases} \quad (11)$$

where a_i and b_i are the linear coefficients associated with the θ_i values given by the corresponding control points. This distribution has to be normalized: $\frac{1}{c_{iso}} \int_{\Omega^+} D_{iso}(\mathbf{m})(\mathbf{m} \cdot \mathbf{n}) d\mathbf{m} = 1$. Therefore:

$$\begin{aligned} c_{iso} &= \int_{\varphi=0}^{2\pi} \int_{\theta=0}^{\frac{\pi}{2}} D_{iso}(\theta) \cos \theta \sin \theta \, d\theta \, d\varphi \\ &= 2\pi \sum_{i=1}^n \int_{\theta=\theta_{i-1}}^{\theta_i} (a_i\theta + b_i) \cos \theta \sin \theta \, d\theta. \end{aligned} \quad (12)$$

For the sake of clarity, we consider D_{iso} as already normalized ($c_{iso} = 1$) in the remainder of this paper.

Using two D_{iso} functions in Equation 10 provides a fully controllable anisotropic NDF model. It can also be used to provide combinations of sampled conventional distributions (such as GGX

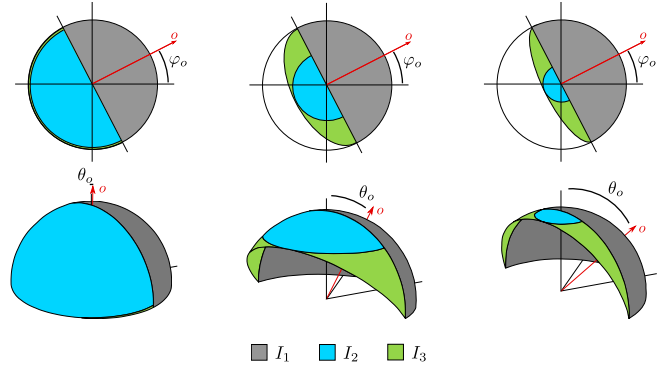


Fig. 6. GAF computation $G_1(\mathbf{o}, \mathbf{m})$: The integration domain Ω is split in three sub-domains depending on direction \mathbf{o} , illustrated for three polar angles: left $\theta_o = 0$, middle $\theta_o = \pi/6$ and right $\theta_o = \pi/4$.

for $\varphi = 0$, and Beckmann's for $\varphi = \pi/2$ for instance) or more complex analytic ones as SGTD (see Section 2 and the supplemental file).

The representation of BSDFs based on 1D curves is not new. Kautz and McCool [1999] replace the evaluation of the BSDF by a sum of products of lower dimensional functions, leading to interesting data compression. More recently, inverse shade trees [Lawrence et al. 2006] account for different levels of separation, organized in a tree where leaves are editable 1D curves. The work presented in this paper rather focuses on the representation of NDFs with 1D curves (instead of the complete BSDF), providing compliance with the microfacet theory while extending material microsurface representation.

3.1 Shadowing and Masking

The microfacet theory requires to handle shadowing and masking, mandatory for energy conservation. As mentioned in Section 2, Heitz proposes to express the distribution in slope space [Heitz 2014], convenient for deriving the Smith model corresponding to conventional NDFs (Beckmann's, GGX, or STD). Unfortunately, this process leads to intractable computations in the general case, and especially with a manually defined NDF. This is why we have used the expression proposed by Ashikhmin et al. [Ashikhmin et al. 2000; Heitz 2014] (Equation 5) where the denominator integral is denoted as $I = \int_{\Omega} \langle \mathbf{o} \cdot \mathbf{m} \rangle D(\mathbf{m}) \, d\omega_m$, such that:

$$I = \int_{\varphi_m=0}^{2\pi} \int_{\theta_m=0}^{\frac{\pi}{2}} \chi^+(\mathbf{o} \cdot \mathbf{m})(\mathbf{o} \cdot \mathbf{m}) D(\mathbf{m}) \sin \theta_m \, d\theta_m \, d\varphi_m, \quad (13)$$

with $\chi^+(a)$ the heaviside function (1 if $a > 0$ and 0 if $a \leq 0$). With a careful subdivision of the integration domain (Figure 6), $\chi^+(a)$ can

be dropped and the integral becomes:

$$\begin{aligned}
 I &= \int_{\theta_m=0}^{\frac{\pi}{2}} \int_{\varphi_m=\varphi_o-\frac{\pi}{2}}^{\varphi_o+\frac{\pi}{2}} (\mathbf{o} \cdot \mathbf{m}) D(\mathbf{m}) \sin \theta_m d\varphi_m d\theta_m \\
 &+ \int_{\theta_m=0}^{\frac{\pi}{2}-\theta_o} \int_{\varphi_m=\varphi_o+\frac{\pi}{2}}^{\varphi_o+\frac{3\pi}{2}} (\mathbf{o} \cdot \mathbf{m}) D(\mathbf{m}) \sin \theta_m d\varphi_m d\theta_m \\
 &+ \int_{\theta_m=\frac{\pi}{2}-\theta_o}^{\frac{\pi}{2}} \left(\int_{\varphi_m=\varphi_o+\frac{\pi}{2}}^{\varphi_o+C} (\mathbf{o} \cdot \mathbf{m}) D(\mathbf{m}) \sin \theta_m d\varphi_m + \right. \\
 &\quad \left. \int_{\varphi_m=\varphi_o-C+2\pi}^{\varphi_o+\frac{3\pi}{2}} (\mathbf{o} \cdot \mathbf{m}) D(\mathbf{m}) \sin \theta_m d\varphi_m \right) d\theta_m \\
 &= I_1 + I_2 + I_3, \tag{14}
 \end{aligned}$$

where $C = \arccos\left(-\frac{1}{\tan \theta_o \tan \theta_m}\right)$. The integration domains for I_1 , I_2 , and I_3 are represented in Figure 6 (see supplemental material for mathematical details). The result is a quasi-analytical solution where some factors can be precomputed since they do not depend on \mathbf{o} .

Note that Equations 5 and 8 are equivalent, and the expression of function Λ can be found straightforwardly from previous equations:

$$\begin{aligned}
 \frac{\cos \theta_o}{\int_{\Omega} \langle \mathbf{o} \cdot \mathbf{m} \rangle D(\mathbf{m}) d\omega_m} &= \frac{1}{1 + \Lambda(\mathbf{o})}, \\
 \text{leading to } \Lambda(\mathbf{o}) &= \frac{\int_{\Omega} \langle \mathbf{o} \cdot \mathbf{m} \rangle D(\mathbf{m}) d\omega_m}{\cos \theta_o} - 1.
 \end{aligned}$$

with $\langle \cdot \rangle$ the clamped dot product. When an analytical expression can be found for Λ , it becomes possible to compute the height-correlated masking and shadowing expression of Smith's GAF (Equation 9) considered as physically more plausible.

3.2 Importance Sampling

Importance sampling is mandatory during the rendering process. Microfacet normals \mathbf{m} are sampled according to the probability density function $p(\mathbf{m}) = D(\mathbf{m})|\mathbf{m} \cdot \mathbf{n}|$, and the associated cumulative distribution function (CDF) is:

$$\begin{aligned}
 cdf(\mathbf{m}) &= \frac{1}{c} \int_{\varphi=0}^{\varphi_m} \int_{\theta=0}^{\theta_m} \left(\cos^2 \varphi D_x(\theta) D_y(0) + \right. \\
 &\quad \left. \sin^2 \varphi D_y(\theta) D_x(0) \right) \cos \theta \sin \theta d\theta d\varphi. \tag{15}
 \end{aligned}$$

Let $(\xi_1, \xi_2) \in [0, 1)^2$, two uniformly sampled numbers. For isotropic materials (i.e., $D_x(\theta) = D_y(\theta)$ and $\varphi_m = 2\pi\xi_1$), θ_m is sampled independently. In the anisotropic case, two CDFs are employed: $cdf(\varphi_m)$ and $cdf(\theta_m|\varphi_m)$.

$$\begin{aligned}
 cdf(\varphi_m) &= \frac{1}{c} \int_{\varphi=0}^{\varphi_m} \int_{\theta=0}^{\frac{\pi}{2}} \left(\cos^2 \varphi D_x(\theta) D_y(0) + \right. \\
 &\quad \left. \sin^2 \varphi D_y(\theta) D_x(0) \right) \cos \theta \sin \theta d\theta d\varphi \\
 &= \frac{\varphi_m}{2\pi} + \frac{D_y(0) - D_x(0)}{4\pi(D_y(0) + D_x(0))} \sin(2\varphi_m) = \xi_1. \tag{16}
 \end{aligned}$$

Unfortunately, this expression is not directly invertible, but it wobbles around $\frac{\varphi}{2\pi}$, with equal values for $\varphi = \{0, \frac{\pi}{2}, \pi, \frac{3\pi}{2}, 2\pi\}$ (corresponding to $\xi_1 = \{0, \frac{1}{4}, \frac{1}{2}, \frac{3}{4}, 1\}$). For $\varphi = \left[0, \frac{\pi}{2}\right]$, the shape of the

CDF is smooth and exhibits a close-to-linear behavior (exactly linear when D is isotropic). According to the process suggested by Jakob [2014], we propose to bracket intervals and use Newton's method to invert the CDF, given ξ_1 .

Based on the sampled azimuthal angle φ_m , θ_m is sampled using the following CDF:

$$\begin{aligned}
 cdf(\theta_m|\varphi_m) &= \frac{1}{c} \int_{\varphi=0}^{2\pi} \int_{\theta=0}^{\theta_m} \left(\cos^2 \varphi_m D_x(\theta) D_y(0) + \right. \\
 &\quad \left. \sin^2 \varphi_m D_y(\theta) D_x(0) \right) \cos \theta \sin \theta d\theta d\varphi \\
 cdf(\theta_m|\varphi_m) &= cdf(\theta_{k+1}|\varphi_m) \\
 &+ \frac{2\pi}{c} \int_{\theta=\theta_{k+1}}^{\theta_m} \left(\cos^2 \varphi_m D_x(\theta) D_y(0) + \right. \\
 &\quad \left. \sin^2 \varphi_m D_y(\theta) D_x(0) \right) \cos \theta \sin \theta d\theta. \tag{17}
 \end{aligned}$$

The polar angle θ_m is sampled from a uniform number $\xi_2 \in [\theta_{k+1}, \theta_{k+2}]$, where θ_{k+1} and θ_{k+2} are two control points. The CDF associated with each control point is precomputed during the distribution instantiation. Knowing the interval containing θ_m , we can write:

$$\int_{\theta=\theta_{k+1}}^{\theta_m} (A^{aniso} \theta + B) \cos \theta \sin \theta d\theta = \frac{c \times (\xi_2' - cdf(\theta_{k+1}|\varphi_m))}{2\pi}, \tag{18}$$

with

$$\begin{aligned}
 A^{aniso} &= \cos^2 \varphi_m D_y(0) a_{k+1}^x + \sin^2 \varphi_m D_x(0) a_{k+1}^y \\
 B^{aniso} &= \cos^2 \varphi_m D_y(0) b_{k+1}^x + \sin^2 \varphi_m D_x(0) b_{k+1}^y \\
 \xi_2' &= \xi_2 * \frac{D_y(0) \cos^2 \varphi_m + D_x(0) \sin^2 \varphi_m}{c}.
 \end{aligned}$$

This function is smooth enough to adopt the same strategy as the sampling of φ_m . The complete importance sampling process is detailed in the supplemental material.

4 SURFACE GENERATION AND MULTIPLE SCATTERING

This section describes a new method for generating a microsurface given any normal distribution function. As mentioned in Section 2, the generated microsurface is placed in a virtual gonioreflectometer that captures the simulated BSDF effects corresponding to light multiple scattering.

4.1 Related Work on Surface Generation

Microfacet BSDFs defined in Equations 1 and 2 do not handle multiple scattering. However, even for glossy microfacets, multiple-scattering effects cannot be ignored, otherwise the surface appears too dark for very rough materials due to loss of energy. In this context, producing a microsurface geometry corresponding to a given distribution is useful for simulating light multiple scattering effects, especially in the case of very general distributions. It is also a powerful tool in applications where simulations are required to understand the effects related to the microsurface geometry. Many approaches have been proposed for generating a height field from distribution functions. In the area of surface tribology, a height probability distribution and autocorrelation behavior are used with two main families of methods: Moving average [Naylor et al. 1966] or Monte Carlo spectral [Thorsos 1988]. They both have been abundantly

investigated or improved for limiting statistical bias [Mack 2013; Uchida et al. 2004; Zou et al. 2007] or for generating anisotropic and non-Gaussian height fields [Manesh et al. 2010; Uchida et al. 2009; Wu 2004]. Unfortunately, the height probability distribution is not necessarily correlated to the slope probability distribution. Some authors propose to generate a C^0 continuous surface from Gaussian slope probability distributions [Cabral et al. 1987; Heitz et al. 2016; Westin et al. 1992], or a continuous C^{-1} surface (i.e., with height discontinuities) from non-Gaussian slope probability distributions [Gondek et al. 1994; Luongo et al. 2017; Pereira et al. 2017]. To the best of our knowledge, the only existing approach suitable for non-Gaussian surfaces and close to C^0 continuity has been introduced by Weyrich et al. [Weyrich et al. 2009] and extended by Papas et al. [Papas et al. 2011]. However as shown by Schwartzburg et al. [Schwartzburg et al. 2014], it does not ensure surface continuity and it is not designed to reproduce smooth distributions. Some other approaches are dedicated to the construction of surfaces with specific collections of microspheres for designing caustics [Hullin et al. 2013; Papas et al. 2011; Schwartzburg et al. 2014]. Such a generated surface must comply with point continuity (incident microfacets have to share edge boundaries with all their neighbors), ensuring that rays are not lost during path tracing simulation. In accordance with the existing methods (that are limited to generate C^0 surfaces only with Gaussian slope probability), we propose a more general solution also compatible with non-Gaussian distributions.

4.2 Surface Generation

Our solution consists in producing a continuous surface mesh from any distribution, based on an iterative shape from gradient method (Algorithm 1 and Figure 7).




Fig. 7. 1D representation of surface generation from a microfacet distribution $D(\theta)$. First row: A set of microfacets is sampled according to $D(\theta)$. Second row: The process re-arranges microfacets to solve difficult configurations (d-e).

The final surface is represented as a height field $h(x, y)$ (or h for short in the following), stored in a 2D regular grid. Given a distribution $D(\theta, \varphi)$, the process is initialized such that each cell of the 2D grid is associated with a normal vector, using the importance sampling process described in Section 3.2. The resulting set of normal vectors defines an initial discrete gradient field ∇h , denoted as \mathcal{Q}_m .

Unfortunately, such a random distribution leads to configurations that cannot be converted directly to a height field without discontinuities. The gradient field should comply with Schwarz's theorem,

ALGORITHM 1: From distribution to surface

input : A distribution $D(\theta, \varphi)$
output : A height field h

Initialize a list of normals \mathcal{Q}_m according to $D(\theta, \varphi)$
Initialize a gradient field ∇h with \mathcal{Q}_m
Initialize a height field h from ∇h

for $it = 0$ **to** $itmax$ **do**

Compute ∇h^+ from h

for each position (x, y) **do**

Compute the Schwarz error $\epsilon(x, y)$ from $\nabla h(x, y)$

if $\epsilon(x, y) > threshold$ **then**

$\nabla h^+(x, y) \leftarrow$ the best matching of $\nabla h^+(x, y)$ from unused values of \mathcal{Q}_m . The best matching value is labelled as already used.

end

else

$\nabla h^+(x, y) \leftarrow \nabla h(x, y)$

end

end

$\nabla h \leftarrow \nabla h^+$
Symmetric padding of ∇h^+
 $h \leftarrow \int \nabla h^+$

end

in its discrete form:

$$\frac{\partial h(x, y)}{\partial x} + \frac{\partial h(x + 1, y)}{\partial y} = \frac{\partial h(x, y)}{\partial y} + \frac{\partial h(x, y + 1)}{\partial x}. \quad (19)$$

Several methods have been proposed to obtain a height field from a noisy gradient field [Agrawal et al. 2006; Horn 1970; Xie et al. 2014]. We have chosen the Frankot-Chellappa algorithm [Frankot and Chellappa 1988], based on a Fourier transform applied to the gradient field. This solution with a fast Fourier transform guarantees Dirichlet boundary condition, and the results are obtained within short computation times, even for large surfaces. It consists in minimizing a least square error function W defined by the following equation:

$$W = \int \int \left(\frac{-\partial h(x, y)}{\partial x} - p \right)^2 + \left(\frac{-\partial h(x, y)}{\partial y} - q \right)^2 dx dy, \quad (20)$$

where $h(x, y)$ is the original surface. Using Parseval's theorem, the expression that links the Fourier transform of the surface $H(u, v)$ and the Fourier transform of the gradients $\left\{ p = \frac{\partial h}{\partial x}, q = \frac{\partial h}{\partial y} \right\}$ is:

$$H(u, v) = \frac{-juP(u, v) - jvQ(u, v)}{u^2 + v^2}, \quad (21)$$

where $(u, v) \neq (0, 0)$. From this relation, the reconstructed surface is directly obtained by the inverse Fourier transform of H . This algorithm is used to obtain an initial height field h with a gradient field ∇h^+ . An iterative process based on the following three steps is used to decrease the bias between ∇h^+ and ∇h .

Step 1: The surface derivative ∇h^+ is approximated by finite central differences.

$$\begin{aligned} \frac{\partial h(x, y)}{\partial x} &= \frac{h(x+1, y) - h(x-1, y)}{2} \\ \frac{\partial h(x, y)}{\partial y} &= \frac{h(x, y+1) - h(x, y-1)}{2}. \end{aligned} \quad (22)$$

Step 2: A matching process between ∇h^+ and \mathcal{L}_m is performed, in order to obtain a new gradient field whose values follow the desired distribution $D(\theta, \varphi)$. The rearrangement process is performed only on values associated with high Schwarz errors. The process is accelerated thanks to a uniform grid, used to find the best matching of $\nabla h^+(x, y)$ from unused values of \mathcal{L}_m . If no value is available in a selected cell of the uniform grid, neighboring cells are used with a spiral path.

Step 3: The Frankot-Chellappa algorithm provides a new height field h . As shown in Figure 8, only few iterations are required for generating a height field for which the slope distribution of microfacet fits with the desired distribution $D(\theta, \varphi)$. In the remainder of this paper, no convergence criterion is used to stop the process (Algorithm 1) and all the surfaces have been obtained after 500 iterations.




Fig. 8. Several steps of the microfacet reconstruction progress, for a height field h , with Beckmann's distribution $\sigma = 0.75$ and a surface size of 1024^2 . First row: Comparison between the desired distribution (green curve) and the current surface slope distribution (red curve). Second row: Surface generated at the indicated iteration.

5 RESULTS AND APPLICATIONS

This section presents several results obtained with various distributions, including isotropic and anisotropic surfaces.

Firstly, we propose to evaluate and validate our height field generation process. All the microsurfaces have been produced with a 2D grid composed of 1024^2 cells (except for Figure 13 where the surface sizes vary), with a Schwarz's threshold of 0.02, corresponding to 2% of a cumulative ∂_x and ∂_y error (Equation 19). For this configuration, the microsurface is produced in approximately 15 minutes on a recent desktop computer with an i7 processor, and mono-thread in c++ code.

Secondly, we use our process (Figure 4) to study microfacet-based BSDFs and discuss the real impact of the NDF shape on the material appearance and multiple-scattering contributions. We have implemented several Mitsuba [Jakob 2010] plugins for handling microfacet BSDF models, relying on data precomputed with our virtual gonioreflectometer and stored in a table according to $(\theta_i, \varphi_i, \theta_o, \varphi_o)$. All renderings are performed with tabulated multiple scattering ($L_\infty = L_1 + L_{2+}$), except for specific cases that are clearly mentioned. We have also implemented a single scattering microfacet BSDF model relying on our piecewise-linear NDF, using its associated GAF and importance sampling. In this case, the L_1 BSDF values can be in practice estimated using the predictive model while L_{2+} comes from a precomputed data table.

In this study, we propose comparisons based on the *CIE* ΔE_{00} color-difference [Sharma et al. 2005]. The scale used for all color difference images is illustrated in Figure 12. The perceptual difference limit ranges from 1 to 3, depending on the line of business (e.g., textile, printing or movie production).

We also present two practical cases for the representation of real measured surfaces and surface design.

5.1 Validation of the Surface Generation Process

We have tested our system and its capabilities for generating surfaces from NDFs with Smith's assumptions. A comparison between our microsurface reconstruction process and the reconstruction method proposed by Heitz et al. [Heitz 2015; Heitz et al. 2016] shows that in the case of Beckmann's distribution, the results are similar for both L_1 (with height correlated Smith's GAF) and L_{2+} (Figure 9), considering conductors or Lambertian microfacets. The advantage of our method is its generalization to any type of distribution.

The study proposed by Heitz et al. [Heitz 2015; Heitz et al. 2016] is limited to Beckmann's distribution. Our method has been applied to 46 different NDF configurations, offering a wide variety of microsurfaces. We have selected a subset of 10 representative cases (see Figure 10): 6 coming from SGTD (Section 2.1), 2 based on a cardinal sinus function and 2 for anisotropic cases (the first one is a mixture of case 7 for $\varphi = 0$ and case 8 for $\varphi = \frac{\pi}{2}$ illustrated in Figure 10 and the second one is a mixture of GGX with $\varphi = 0$ and Beckmann's with $\varphi = \frac{\pi}{2}$). All the test cases are approximated using our piecewise-linear NDF (Section 3). The additional material presents all 46 test cases. Figure 10 illustrates the geometry obtained with our microsurface generation process, given an input set of normal vectors \mathcal{L} . Note how the obtained gradient field ∇h^+ is similar to the input data and the desired NDF. All the generated surfaces are characterized by a Gaussian-like height histogram, corresponding to a random distribution of heights. Note that correlation between heights is actually low, compliant with Smith's assumptions.

5.2 Material Appearance Controlled by NDF Shape

Surface generation is a powerful tool for analyzing BSDF behaviors and estimating multiple scattering. Figure 11 illustrates the BSDFs of our 10 selected test cases. It compares the predictive models (denoted as *Model* L_1) with the simulated data (denoted as *Simulated* L_1 and L_{2+}), gathered by the sensor of the virtual gonioreflectometer.




Fig. 9. BSRF comparisons with Beckmann’s distribution ($\sigma = 0.75$), between surfaces generated with Heitz’s method [Heitz 2015] and ours. The incident angle is $\theta_i = 1.5$ rad, with E_r the total amount of reflected energy, L_1 one-bounce BSRF, and L_{2+} multiple-bounce BSRF only. False color images correspond to the absolute difference between L_1 and *Model* L_1 .

The BSRF values are estimated thanks to a standard path tracing approach, performed on the reconstructed surface.

Observing the captured L_1 values shows that Smith’s assumptions are observed in practice, with the correlated GAF (also used for producing the data corresponding to *Model* L_1) independently of the chosen NDF. The BSRF predicted by the analytical model is very similar to the estimation based on a reconstructed rough surface.

Slight differences can although be observed between the predictive model and simulated L_1 , especially when strong slopes are generated. This is particularly visible in cases 3 and 8 in Figure 11; Rendered images exhibit the visual difference in Figure 12-bottom. NDF sub-sampling is not responsible for this difference: We have tested our generation process on different surface sizes (Figure 13) for case 2 (without noticeable differences between predictable and simulated L_1) and case 3. Even though sub-sampling can be detrimental (from 64^2 to 256^2 samples), the surface generation process is stable and the rendered images are similar for several reconstructions with the same NDF (at least from 512^2 to 4096^2). Our study confirms that Smith’s GAF overestimates shadowing for some directions, even when the height-correlated shadowing and masking Smith’s term is used [Heitz 2014]. Defining a convenient and more precise shadowing and masking term, for any NDF, is still an open question.

Figure 1 shows four NDF configurations, extracted from our test case set (Figures 10 and 11). These NDFs are associated with different f^μ (diffuse, pure mirror, gold and copper conductors). The choice of an NDF impacts drastically the visual appearance of a material, sometimes leading to strong backward light reflections at the object silhouette. This effect can be observed in Figure 11 for all the cases

where the NDF represents rough surfaces with strong slopes (except for case 1).

In those cases, large normal variations lead to a darker surface appearance if multiple scattering of light is not handled. Figure 1 exhibits the importance of this phenomenon, which drastically impacts the final material aspect for some examples. However, the definition of a formulation dedicated to multiple scattering remains an open problem. Unfortunately, the assumptions proposed by some authors [Lagarde 2017] are not true in general and the multiple scattering lobe L_{2+} is not simply a scaled version of the first bounce lobe (L_1) as shown in Figure 11, especially for very rough microspheres.

Multi-lobe distributions can be defined with our piecewise-linear representation, used in this paper for managing all our test cases. This representation has been used for the validation and stress-tests of our microsphere generation process, proving its robustness even in extreme cases (cases 2 to 10). Managing such complex distributions is also useful for the analysis of light reflections and for the understanding of the range of appearances produced by microfacet BSDF models. Some examples typically exhibit interesting backward lighting reflection effects at grazing angles (e.g., case 6 in Figure 11).

5.3 Practical Cases

5.3.1 Measured NDF. We have also measured a manufactured surface, obtained from a stereo-photometric acquisition system [Shi et al. 2018; Woodham 1980]. The material is an imitation of leather, with anisotropic structures and a non-Gaussian NDF, as illustrated in Figure 14. From the measured anisotropic distribution, a new set of normals \mathcal{Q}_m and a new height field have been produced thanks to our process. The measured surface is different from the generated one, but the resulting NDF is very close, with observed differences due to the constrained interpolation over D_x and D_y (Section 3). Again, comparisons can be performed for L_1 and L_{2+} , based on the obtained NDF and the microsphere. The first row of Figure 14 shows the reconstructed surface shape, based on a stereo-photometric process, as well as the generated surface. The NDF closely corresponds to the original distribution, and L_1 and L_{2+} values have similar shapes between generated and measured surfaces. The BSRFs produced with the measured and the generated surfaces are closely the same. The combination of predictive models and precomputed multiple scattering effects could be employed to render measured surfaces in a multi-scale rendering process.

5.3.2 Constrained surfaces. Our surface generation process can be used to design geometric features while complying with the input distribution. In Algorithm 1, \mathcal{Q}_m is filled according to $D(\theta, \varphi)$ but the gradient field ∇h is initialized with a given geometry. The rest of the algorithm remains the same. In Figure 15, the input NDF corresponds to a conventional Beckmann’s one with $\sigma = 0.5$. Our generation process is performed with two geometric constraints: (i) With a square lattice of domes, and (ii) with a square lattice of holes. The generated gradient fields ∇h^+ are similar to the desired NDF in both cases. Of course, even if the NDF corresponds to Beckmann’s, this constrained method breaks Smith’s assumptions and the obtained surfaces are correlated (see histograms in Figure 15). However, the captured BSRF shows that these two surfaces are similar and the rendered images are also visually similar. We




Fig. 10. NDF configurations: For each one, the first image corresponds to the surface obtained with our height field generation process (Algorithm 1), the second image compares the input NDF (ψ) with the resulting NDF (∇h^+), and the plot on the right provides the height histogram of the surface (insets correspond to the auto-correlation function).

have performed a more extensive test on six different geometric constraints (Figure 16) and the same observations can be made. We also compare computer-generated images with a pure Beckmann surface generated without any constraint. The produced images are similar and differences are visually indistinguishable. These preliminary observations comply with the conclusions of Bourlier *et al.* [Bourlier *et al.* 2000], and their experiments on the effect of correlation with only Smith's shadowing term (G_1). Interesting future questions are opened for multi-level observations of surface details: Geometric features are visible when the surface is observed closely and progressively disappear when the observer moves away.

6 CONCLUSION AND FUTURE WORK

This paper proposes a study on the impact of NDFs with microfacet-based BSDF models and their influence on visual appearance. First, we provide a theoretical background for classifying the NDFs mainly

used in computer graphics. Second, we describe a complete framework, based on a new surface generation process combined with a piecewise-linear NDF definition. We show how this latter representation can be employed with Smith's GAF and importance sampling. In practice, our microsurface construction process produces an explicit geometry from any given NDF, making it possible to simulate light single and multiple scattering effects. This process has been used to study the impact of geometric attenuation factors.

Our new definitions extend the representation of anisotropic NDFs, and offer a finer control over the azimuthal angle, in particular for fitting measured NDFs (Figure 14).

Our experiments demonstrate the versatility of Smith's shadowing and masking term, well suited in practice for a large variety of configurations. However, for very rough surfaces, the model does not account for multiple reflexions and energy loss cannot be neglected. Finding a practical and convenient GAF, adaptable to any




Fig. 11. For each test case, the first image shows a rendering with pure specular microfacets (NDFs are illustrated in insets), the middle image compares the predictive BRDF (Smith's assumptions), with the BRDF captured by our virtual goniophotometer. Multiple scattering is illustrated in the right image. $\theta_i = 1$ for all test cases.

NDF, which does not overestimate shadowing, is still challenging. To our knowledge, correlated height fields have not been used yet in microfacet theory. The experiments carried out in this paper tend to show that Smith's formulations can be applied in many cases (Figure 16).

Multiple scattering is another source of interesting future work, notably for defining a practical formulation independently of the chosen NDF, and compliant with Smith's assumptions. Our framework can be useful in this direction for experimentally validating future models. We also plan to derive our system to handle multi-layer




Fig. 12. Comparison between a BSDF predicted using Smith's GAF and simulated data. The results exhibit low energy differences that can be sometimes distinguished (but not systematically) in rendered images. For instance, with case 8, height correlated masking and shadowing from Smith's term overestimates shadowing at grazing angles.

materials, where multiple scattering has to be managed between layers and interface microfacets, to achieve realistic renderings. Recent progress has been made in this direction [Guo et al. 2018], but independent interface roughnesses between layers is not currently supported.

Practical cases mentioned in previous sections are related to multi scale observations and surface design. Our work offers a methodology for defining new improvements on these topics. For instance, a surface could be defined with specific geometric features, visible when observed closely and not distinguishable from a distant observation. Moreover, these features could differ from one another, while producing the same distant appearance (Figures 15 and 16).

Surface design has already been addressed by Wu et al. [Wu et al. 2011], where large-scale material appearance is derived from an already known micro-scale surface. They propose an efficient approximation technique to define a Bidirectional Visible NDF from this surface. However, their system is limited to single scattering components. On the other hand, our framework goes a step further with multiple scattering, mandatory for designing surfaces from given appearances.

REFERENCES

- A. Agrawal, R. Raskar, and R. Chellappa. 2006. What is the Range of Surface Reconstructions from a Gradient Field?. In *Proceedings of the 9th European Conference on Computer Vision - Volume Part I (ECCV'06)*. Springer-Verlag, Berlin, Heidelberg, 578–591.
- M. Ashikhmin, S. Premoze, and P. Shirley. 2000. A microfacet-based BRDF generator. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 2000, New Orleans, LA, USA, July 23–28, 2000*. ACM, 65–74.
- M. M. Bagher, C. Soler, and N. Holzschuch. 2012. Accurate Fitting of Measured Reflectances Using a Shifted Gamma Micro-facet Distribution. *Computer Graphics Forum* 31, 4 (June 2012).
- B. Beckers and P. Beckers. 2012. A general rule for disk and hemisphere partition into equal-area cells. *Computational Geometry* 45, 7 (2012), 275 – 283.
- P. Beckmann and A. Spizzichino. 1963. *The scattering of electromagnetic waves from rough surfaces*. Pergamon Press.
- C. Bourlier, G. Berginc, and J. Saillard. 2002. One- and two-dimensional shadowing functions for any height and slope stationary uncorrelated surface in the monostatic and bistatic configurations. *IEEE Transactions on Antennas and Propagation* 50, 3 (2002).
- C. Bourlier, J. Saillard, and G. Berginc. 2000. Effect of correlation between shadowing and shadowed points on the Wagner and Smith monostatic one-dimensional shadowing functions. *IEEE Transactions on Antennas and Propagation* 48, 3 (2000), 437–446.
- B. Burley. 2012. Physically-Based Shading at Disney. In *ACM SIGGRAPH 2012 Courses*.
- S. D. Butler and M. A. Marciniak. 2014. Robust categorization of microfacet BRDF models to enable flexible application-specific BRDF adaptation. In *SPIE 9205, Reflection, Scattering, and Diffraction from Surfaces IV*.
- B. Cabral, N. Max, and R. Springmeyer. 1987. Bidirectional Reflection Functions from Surface Bump Maps. *SIGGRAPH Comput. Graph.* 21, 4 (Aug. 1987), 273–281.
- L. Claustres, L. Barthe, and M. Paulin. 2007. Wavelet Encoding of BRDFs for Real-time Rendering. In *Proceedings of Graphics Interface 2007 (GI '07)*. 169–176.
- R. L. Cook and K. E. Torrance. 1982. A Reflectance Model for Computer Graphics. In *ACM SIGGRAPH proceedings*.
- Z. Dong, B. Walter, S. Marschner, and D. P. Greenberg. 2015. Predicting Appearance from Measured Microgeometry of Metal Surfaces. *ACM Trans. Graph.* 35, 1 (2015), 9:1–9:13.
- J. Dupuy, E. Heitz, J. Iehl, P. Poulin, F. Neyret, and V. Ostromoukhov. 2013. Linear Efficient Antialiasing Displacement and Reflectance Mapping. *ACM Trans. Graph.* 32, 6 (Nov. 2013), 211:1–211:11.
- R. T. Frankot and R. Chellappa. 1988. A Method for Enforcing Integrability in Shape from Shading Algorithms. *IEEE Trans. Pattern Anal. Mach. Intell.* 10, 4 (July 1988), 439–451.
- P. Gautron, J. Křivánek, S. N. Pattanaik, and K. Bouatouch. 2004. A Novel Hemispherical Basis for Accurate and Efficient Rendering. In *Rendering Techniques 2004, Eurographics Symposium on Rendering*. Eurographics Association, 321–330.
- J. S. Gondek, G. W. Meyer, and J. G. Newman. 1994. Wavelength Dependent Reflectance Functions. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '94)*. ACM, New York, NY, USA, 213–220.
- Y. Guo, M. Hašan, and S. Zhao. 2018. Position-Free Monte Carlo Simulation for Arbitrary Layered BSDFs. *ACM Trans. Graph.* 37, 6 (2018).
- F. Halley. 2012. *Perceptually relevant browsing environments for large texture databases*. Ph.D. Dissertation. Heriot-Watt University, Edinburgh, UK.
- E. Heitz. 2014. Understanding the Masking-Shadowing Function in Microfacet-Based BRDFs. *Journal of Computer Graphics Techniques* 3, 2 (June 2014).
- E. Heitz. 2015. *Generating Procedural Beckmann Surfaces*. Technical Report. <https://eheitresearch.wordpress.com/research/>.
- E. Heitz, C. Bourlier, and N. Pinel. 2013. Correlation effect between transmitter and receiver azimuthal directions on the illumination function from a random rough surface. *Waves in Random and Complex Media* 23, 3 (Aug. 2013).
- E. Heitz and E. D'Eon. 2014. Importance Sampling Microfacet-Based BSDFs using the Distribution of Visible Normals. *Computer Graphics Forum, EGSR proceedings* 33, 4 (July 2014).
- E. Heitz, J. Hanika, E. d'Eon, and C. Dachsbacher. 2016. Multiple-Scattering Microfacet BSDFs with the Smith Model. *ACM Transactions On Graphics, SIGGRAPH Asia proceedings* 35, 4 (August 2016).
- N. Hoffman. 2016. Recent Advances in Physically Based Shading. In *ACM SIGGRAPH 2016 Courses*.
- N. Holzschuch and R. Pacanowski. 2017. A Two-scale Microfacet Reflectance Model Combining Reflection and Diffraction. *ACM Trans. Graph.* 36, 4, Article 66 (July 2017), 66:1–66:12 pages.
- B. K. P. Horn. 1970. *Shape from shading: a method for obtaining the shape of a smooth opaque object from one view*. Technical Report. Cambridge, MA, USA.
- M. B. Hullin, I. Ihrke, W. Heidrich, T. Weyrich, G. Damberg, and M. Fuchs. 2013. Computational Fabrication and Display of Material Appearance. In *Eurographics - State of the Art Reports*.
- W. Jakob. 2010. Mitsuba renderer. (2010). <http://www.mitsuba-renderer.org>.
- W. Jakob. 2014. *An Improved Visible Normal Sampling Routine for the Beckmann Distribution*. Technical Report. ETH Zürich.
- J. Kautz and M. D. McCool. 1999. Interactive Rendering with Arbitrary BRDFs Using Separable Approximations. In *Proceedings of the 10th Eurographics Conference on Rendering (EGWR'99)*. Eurographics Association, 247–260.
- C. Kelemen and L. Szirmay-Kalos. 2001. A Microfacet Based Coupled Specular-Matte BRDF Model with Importance Sampling. In *Eurographics 2001 - Short Presentations*. Eurographics Association.
- C. Kulla and A. Conty. 2017. Physically Based Shading in Theory and Practice - Revisiting Physically Based Shading at Imageworks. In *ACM SIGGRAPH 2017 Courses*.
- M. Kurt, L. Szirmay-Kalos, and J. Křivánek. 2010. An Anisotropic BRDF Model for Fitting and Monte Carlo Rendering. *ACM SIGGRAPH Computer Graphics* 44, 1 (Feb. 2010).
- S. Lagarde. 2017. Open Problems in Real-Time Rendering - Physically-Based Materials: Where Are We?. In *ACM SIGGRAPH 2017 Courses*.
- J. Lawrence, A. Ben-Artzi, C. DeCoro, W. Matusik, H. Pfister, R. Ramamoorthi, and S. Rusinkiewicz. 2006. Inverse Shade Trees for Non-Parametric Material Representation and Editing. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 25, 3 (jul 2006).




Fig. 13. Robustness of our height field generation process when surface size varies. The two first rows correspond to case 2, and two last rows to case 3 in Figure 11. The middle row illustrates ΔE_{00} between the two models. Between each surface size, we perform a perceptual difference ΔE_{00} . Note that differences start to be imperceptible as soon as a 512^2 resolution is used.

- J. H. Lee, A. Jarabo, D. S. Jeon, D. Gutierrez, and M. H. Kim. 2018. Practical Multiple Scattering for Rough Surfaces. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia 2018)* 37, 6 (2018), 275:1–12.
- Ro, R. Lewis. 1998. *Light-driven Global Illumination with a Wavelet Representation of Light Transport*. Ph.D. Dissertation. Vancouver, BC, Canada, Canada.
- A. Luongo, V. Falster, M. B. Doest, D. Li, F. Regi, Y. Zhang, G. Tosello, J.B. Nielsen, H. Aanæs, and J. R. Frisvad. 2017. Modeling the anisotropic reflectance of a surface with microstructure engineered to obtain visible contrast after rotation. In *ICCV 2017 Workshops*. 159–165.
- C. A. Mack. 2013. Generating random rough edges, surfaces, and volumes. *Applied Optics* 52:7 (2013), 1472–1480.
- D. Mahajan, Yu-T. Tseng, and R. Ramamoorthi. 2008. An Analysis of the In-out BRDF Factorization for View-dependent Relighting. In *Proceedings of the Nineteenth Eurographics Conference on Rendering (EGSR '08)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 1137–1145.
- K. K. Manesh, B. Ramamoorthi, and M. Singaperumal. 2010. Numerical generation of anisotropic 3D non-Gaussian engineering surfaces with specified 3D surface roughness parameters. *Wear* 268 (2010), 1371–1379.
- D. Meneveau, B. Bringier, E. Tauzia, M. Ribardière, and L. Simonot. 2017. Rendering Rough Opaque Materials with Interfaced Lambertian Microfacets. *IEEE Transactions on Visualization and Computer Graphics* (2017).
- T. H. Naylor, J. L. Balintfy, D. S. Burdick, and K. Chu. 1966. *Computer Simulation Techniques*. Wiley. 118–121 pages.
- M. Oren and S. K. Nayar. 1994. Generalization of Lambert's Reflectance Model. In *ACM SIGGRAPH proceedings*.
- M. Papas, W. Jarosz, W. Jakob, S. Rusinkiewicz, W. Matusik, and T. Weyrich. 2011. Goal-Based Caustics. *Computer Graphics Forum (Proc. Eurographics)* 30, 2 (April 2011).
- T. Pereira, C. L. A. P. Leme, S. Marschner, and S. Rusinkiewicz. 2017. Printing Anisotropic Appearance with Magnetic Flakes. *ACM Trans. Graph.* 36, 4, Article 123 (July 2017), 123:1–123:10 pages.
- M. Ribardière, B. Bringier, D. Meneveau, and L. Simonot. 2017. STD: Student's t-Distribution of Slopes for Microfacet Based BSDFs. *Computer Graphics Forum* (2017).
- V. Ross, . Dion, and G. Potvin. 2005. Detailed analytical approach to the Gaussian surface bidirectional reflectance distribution function specular component applied to the sea surface. *J. Opt. Soc. Am. A* 22, 11 (Nov 2005).
- Y. Schwartzburg, R. Testuz, A. Tagliasacchi, and M. Pauly. 2014. High-contrast Computational Caustic Design. *ACM Trans. Graph.* 33, 4, Article 74 (July 2014), 74:1–74:11 pages.
- G. Sharma, W. Wu, and E. N. Dalal. 2005. The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. *Color Research & Application* 30, 1 (2005), 21–30.
- B. Shi, Z. Mo, Z. Wu, D. Duan, S. K. Yeung, and P. Tan. 2018. A Benchmark Dataset and Evaluation for Non-Lambertian and Uncalibrated Photometric Stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2018).
- B. Smith. 1967. Geometrical shadowing of a random rough surface. *IEEE Transactions on Antennas and Propagation* 15, 5 (September 1967), 668–671.
- C. Soler, K. Subr, and D. Nowrouzezahrai. 2018. A Versatile Parameterization for Measured Material Manifolds. *Computer Graphics Forum* 37, 2 (April 2018), 1–10.
- P. Theodossiou. 1998. Financial Data and the Skewed Generalized T Distribution. *Management Science* 44, 12-Part-1 (1998), 1650–1661.
- E. I. Thorsos. 1988. The validity of the Kirchhoff approximation for rough surface scattering using a Gaussian roughness spectrum. *J. Acoust. Soc. Am.* 83 (1988), 78–92.
- K. E. Torrance and E. M. Sparrow. 1967. Theory for Off-Specular Reflection From Roughened Surfaces. *Journal of Optical Society of America* 57, 9 (Sep 1967).
- T. S. Trowbridge and K. P. Reitz. 1975. Average irregularity representation of a rough surface for ray reflection. *Journal of Optical Society of America* 65, 5 (May 1975).
- K. Uchida, J. Honda, and K.-Y. Yoon. 2009. An Algorithm for Rough Surface Generation with Inhomogeneous Parameters. In *IEEE Int. conf. on Parallel Processing Workshops*. 166–173.
- M. Uchida, T. Shimizu, A. Iwabuchi, and K. Yanagi. 2004. Generation of reference data of 3D surface texture using the non-causal 2D AR model. *Wear* 257 (2004), 1288–1295.
- B. Walter, S. R. Marschner, H. Li, and K. E. Torrance. 2007. Microfacet Models for Refraction Through Rough Surfaces. In *Computer Graphics Forum, EGSR proceedings*.
- D. Wellems, S. Ortega, D. Bowers, J. Boger, and M. Fetrow. 2006. Long wave infrared polarimetric model: theory, measurements and parameters. *Journal of Optics A: Pure*




Fig. 14. Real non-Gaussian surface acquired (imitation of leather). Comparison between simulated single and multiple scattering BSDFs (with perfect mirror microfacet BSDF) obtained with measured and reconstructed surfaces.




Fig. 15. Constrained generated surfaces, using a dome (first row) or a hole (second row) pattern from a Beckmann NDF ($\sigma = 0.5$). Smith's assumptions are not reached: resulting surfaces have a skewed and shifted height histogram and correlation between heights is strong (histogram inset).




Fig. 16. Constrained generated surfaces, using different patterns from a Beckmann NDF ($\sigma = 0.5$). The last row compares generated surfaces with one generated from unconstrained Beckmann's ($\sigma = 0.5$).

and Applied Optics 8, 10 (2006).

S. H. Westin, J. R. Arvo, and K. E. Torrance. 1992. Predicting Reflectance Functions from Complex Surfaces. *SIGGRAPH Comput. Graph.* 26, 2 (July 1992), 255–264.

T. Weyrich, P. Peers, W. Matusik, and S. Rusinkiewicz. 2009. Fabricating Microgeometry for Custom Surface Reflectance. *ACM Trans. Graph.* 28, 3, Article 32 (July 2009), 32:1–32:6 pages.

Wikipedia. 2018. Skewed generalized t distribution. (2018). https://en.wikipedia.org/wiki/Skewed_generalized_t_distribution [Online; accessed June 2018].

R. Woodham. 1980. Photometric method for determining surface orientation from multiple images. *Opt. Eng.* 19 (1980), 139–144.

H. Wu, J. Dorsey, and H. Rushmeier. 2011. Physically-Based Interactive Bi-Scale Material Design. *ACM Transactions on Graphics* 30, 145 (dec 2011).

J.-J. Wu. 2004. Simulation of non-Gaussian surfaces with FFT. *Tribology International* 37 (2004), 339–346.

F. Xie and P. Hanrahan. 2018. Multiple Scattering from Distributions of Specular V-grooves. *ACM Trans. Graph.* 37, 6 (Dec. 2018), 276:1–276:14.

W. Xie, Y. Zhang, C. C. L. Wang, and R. C. k. Chung. 2014. Surface-from-Gradients: An Approach Based on Discrete Geometry Processing. In *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23–28, 2014*. 2203–2210.

L.-Q. Yan, M. Hašan, S. Marschner, and R. Ramamoorthi. 2016. Position-Normal Distributions for Efficient Rendering of Specular Microstructure. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2016)* 35, 4 (2016).

L.-Q. Yan, M. Hašan, B. Walter, S. Marschner, and R. Ramamoorthi. 2018. Rendering Specular Microgeometry with Wave Optics. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2018)* 37, 4 (2018).

M. Yan, L.-Q. and Hašan, W. Jakob, J. Lawrence, S. Marschner, and R Ramamoorthi. 2014. Rendering Glints on High-Resolution Normal-Mapped Specular Surfaces. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2014)* 33, 4 (2014).

T. Zeltner and W. Jakob. 2018. The Layer Laboratory: A Calculus for Additive and Subtractive Composition of Anisotropic Surface Reflectance. *Transactions on Graphics (Proceedings of SIGGRAPH)* 37, 4 (2018), 74:1–74:14.

M. Zou, B. Yu, Y. Feng, and P. Xu. 2007. A Monte Carlo method for simulating fractal surfaces. *Physica A* 386 (2007), 176–186.