

# Position-free Multiple-bounce Computations for Smith Microfacet BSDFs

BEIBEI WANG, Nankai University, China

WENHUA JIN, JIAHUI FAN, Nanjing University of Science and Technology, China

JIAN YANG, Nankai University, China

NICOLAS HOLZSCHUCH, University Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, France

LING-QI YAN, University of California, Santa Barbara, USA

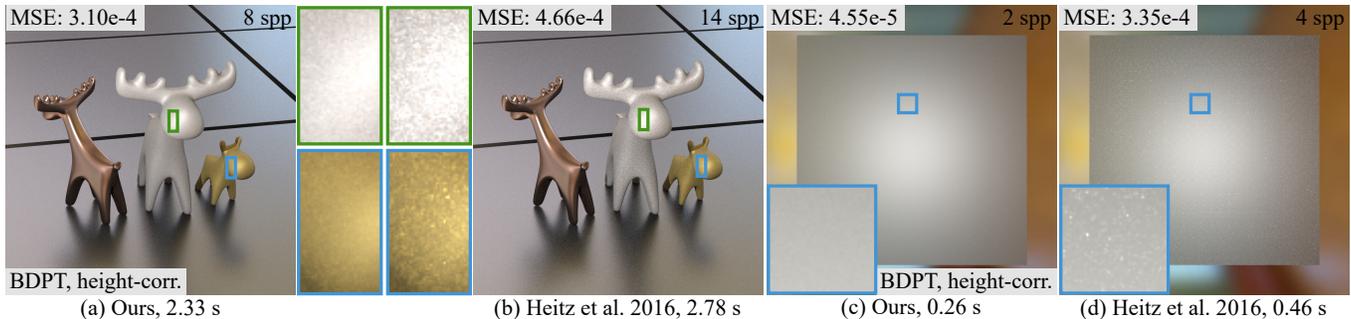


Fig. 1. We present a new method to compute multiple bounces for microfacet BSDFs. Our method uses a position-free approach, eliminating the need to trace the heights of previous approaches, significantly reducing noise (see insets). It applies to both rough conductors and rough dielectrics. (a) and (b): equal-time comparison with Heitz et al. [2016] for rough conductors (from left to right: copper (GGX model,  $\alpha = 0.1$ ), aluminum (GGX model,  $\alpha = 0.6$ ) and gold (GGX model,  $\alpha = 0.5$ )). (c) and (d): comparison with Heitz et al. [2016] for a rough dielectric (GGX model,  $\alpha = 1.0$ ). We support both height-correlated and height-uncorrelated shadowing-masking functions. This figure uses the height-correlated model.

Bidirectional Scattering Distribution Functions (BSDFs) encode how a material reflects or transmits the incoming light. The most commonly used model is the microfacet BSDF. It computes the material response from the microgeometry of the surface assuming a single bounce on specular microfacets. The original model ignores multiple bounces on the microgeometry, resulting in an energy loss, especially for rough materials. In this paper, we present a new method to compute the multiple bounces inside the microgeometry, eliminating this energy loss. Our method relies on a position-free formulation of multiple bounces inside the microgeometry. We use an explicit mathematical definition of the path space that describes single and multiple bounces in a uniform way. We then study the behavior of light on the different vertices and segments in the path space, leading to a reciprocal multiple-bounce description of BSDFs. Furthermore, we present practical, unbiased Monte Carlo estimators to compute multiple scattering. Our method is less noisy than

existing algorithms for computing multiple scattering. It is almost noise-free with a very-low sampling rate, from 2 to 4 samples per pixel (spp).

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

Additional Key Words and Phrases: microfacet, position-free, multiple-bounce, full-spherical

## ACM Reference Format:

Beibei Wang, Wenhua Jin, Jiahui Fan, Jian Yang, Nicolas Holzschuch, and Ling-Qi Yan. 2022. Position-free Multiple-bounce Computations for Smith Microfacet BSDFs. *ACM Trans. Graph.* 41, 4, Article 1 (July 2022), 14 pages. <https://doi.org/10.1145/3528223.3530112>

## 1 INTRODUCTION

Material properties or reflectance encodes how materials interact with the incoming light. Having good material appearance properties is essential for photorealistic rendering. Microfacet models are widely used both in real-time applications and in the high-quality offline rendering. They predict the appearance of the material from the statistical properties of its surface's microgeometry. The most common model, Cook-Torrance [Cook and Torrance 1982; Walter et al. 2007], assumes that the surface is made of planar specular microfacets, and computes the material response by integrating a single bounce over this microgeometry. The BSDF is connected to the distribution of microfacet normals. By nature, the Cook-Torrance

Authors' addresses: Beibei Wang, Nankai University, College of Computer Science, China, [beibei.wang@nankai.edu.cn](mailto:beibei.wang@nankai.edu.cn); Wenhua Jin, Jiahui Fan, Nanjing University of Science and Technology, China, [jwh,fjh@njust.edu.cn](mailto:jwh,fjh@njust.edu.cn); Jian Yang, Nankai University, College of Computer Science, China, [csjyang@nankai.edu.cn](mailto:csjyang@nankai.edu.cn); Nicolas Holzschuch, University Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, France, [nicolas.holzschuch@inria.fr](mailto:nicolas.holzschuch@inria.fr); Ling-Qi Yan, University of California, Santa Barbara, USA, [lingqi@cs.ucsb.edu](mailto:lingqi@cs.ucsb.edu).

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](mailto:permissions@acm.org).

© 2022 Association for Computing Machinery.

0730-0301/2022/7-ART1 \$15.00

<https://doi.org/10.1145/3528223.3530112>

model ignores the light that has bounced several times on the microgeometry, resulting in an energy loss. The effect is especially visible when the surface has a high roughness.

Multiple algorithms have tried to enhance the microfacet models by computing multiple bounces. The common idea is that the shadowing-masking term of the Cook-Torrance model actually encodes the proportion of light that was not reflected in the first bounce. This light that was blocked by the microgeometry is used as input to compute the multiple-bounce term.

There are several models for the shadowing-masking term of the Cook-Torrance BRDF model: the V-groove model assumes that for each microfacet, there is another one next to it forming a V-shaped groove with it; the Smith model assumes the independent distribution of microfacets heights and normals, treating the microfacets as randomly distributed microflakes. The former is easier for numerical analysis and provides explicit analytical solutions, the latter requires a double indefinite integral to compute the shadowing-masking term; some microfacet distributions still have an analytical term for the shadowing-masking.

Two algorithms use the V-groove model to compute multiple bounces in the microgeometry, and provide an analytical formula for the missing energy [Lee et al. 2018; Xie and Hanrahan 2018]. However, the V-groove model has several issues that can make it undesirable: discontinuities or singularities in the shadowing-masking term, an overall shiny appearance even for very rough surfaces, and the inability to model transparent materials.

The Smith model results in a better overall appearance for the material, but there is no explicit formula to compute multiple bounces of light. Heitz et al. [2016] showed that it was nevertheless possible to compute multiple scattering effects inside the microgeometry using random walks. Their simulation takes into account the heights inside the microgeometry, leading to a very accurate result, at the expense of computation time.

In this paper, we present a new understanding of multiple-bounce microfacet BSDFs. Inspired by the position-free approach that Guo et al. [2018] applied to layered materials, we analyze and formulate the *path space* as the light undergoes an arbitrary number of bounces inside general microfacet BSDFs. In this path space, we study the behavior of light at the vertices and segments along different paths, introducing the *vertex term* and the *segment term*, respectively.

At the core of our method is a deeper understanding of the *position-free* property of all BSDFs, extended from macro-scale [Guo et al. 2018] to micro-scale. In Sec. 3, we demonstrate that with our position-free formulation, all points at micro scale share the same statistics, especially the normal distribution. Therefore, any light path bouncing away from a shading point and inflicting the next bounce will be *incident to the same shading point again* from its last outgoing direction. This immediately leads to a nice height-free property — we do not need the micro-scale height distribution, and we do not have to trace the height, keeping its height as micro-scale.

To quantify the energy of the light that undergoes further bounces, we propose two kinds of different shadowing-masking functions in our multiple-bounce BSDF: the height-uncorrelated and the height-correlated function. The height-uncorrelated shadowing-masking function studies the visibilities completely irrelevant to the height,

Table 1. Notations.

Mathematical notation	
$\Omega^\pm$	full spherical domain
$\Omega^+$	upper spherical domain
$\omega_i \cdot \omega_o$	dot product
$ \omega_i \cdot \omega_o $	absolute value of the dot product
$\langle \omega_i, \omega_o \rangle$	dot product clamped to 0
$\chi^+(a)$	Heaviside function: 1 if $a > 0$ and 0 if $a \leq 0$
Physical quantities used in microfacet models	
$\omega_g = (0, 0, 1)$	geometric normal
$\omega_m$	microfacet normal
$\omega_i$	incident direction ( $\omega_i \cdot \omega_g$ could be $< 0$ )
$\omega_o$	outgoing direction
$\Lambda(\omega)$	the Smith Lambda function
$D(\omega_m)$	normal distribution function
$F(\omega_i, \omega_m)$	Fresnel factor
$G_1(\omega, \omega_m)$	masking function
$G_1^{\text{local}}(\omega, \omega_m)$	masking (local)
$G_1^{\text{dist}}(\omega)$	masking (distant)
$\rho(\omega_i, \omega_o)$	multiple-bounce BSDF

while the height correlated one averages/integrates different visibilities at different heights, thus is also independent of specific heights, and provides a better match with the Heitz et al. [2016]. However, note specifically that both of them are reasonable approximations to the real visibility, both of them preserve energy conservation, and neither of them requires tracking the height.

With our explicit position-free formulation, acquiring the multiple scattered energy from BSDFs becomes much more efficient, because better Monte Carlo estimators can be immediately applied to solve the integration. We propose practical Monte Carlo estimators, path tracing (PT) and bidirectional path tracing (BDPT). Our method provides result similar to Heitz et al. [2016], with a significantly decreased noise level for evaluation-only tasks, even with very low sample counts, as low as 2 – 4 spp. It passes the white furnace test. It works with dielectrics, anisotropic materials and commonly-used normal distributions such as Beckmann and GGX.

An open source implementation of our methods is available at <https://github.com/wangningbei/sourceCodeMBSDF>.

## 2 RELATED WORK

*Microfacet models.* Torrance and Sparrow [1967] introduced the microfacet model for reflection on rough surfaces. They extract the overall material reflectance from a statistical description of the surface microgeometry, made of specular microfacets. The model focuses on a single bounce over this microgeometry, and gives a full BRDF model from the surface characteristics. The most important parameter is the probability distribution of the microfacet normals (normal distribution function, or NDF). The model depends on two other terms: the Fresnel term, connected to the composition of the material, and the shadowing-masking term, encoding how much of the incoming light goes into this first bounce. The model was introduced to the graphics community by Cook and Torrance [1982] and extended to rough dielectrics by Stam [2001]. Later, Walter

et al. [2007] introduced more normal distribution functions to the microfacet model, like GGX.

The normal distribution function has a strong impact on the overall aspect of the BRDF. Initial works used the Beckmann distribution [Beckmann and Spizzichino 1963; Cook and Torrance 1982; Torrance and Sparrow 1967]. Trowbridge and Reitz [1975] introduced a different distribution, corresponding to microfacets distributed on half-ellipsoids. It was rediscovered by Walter et al. [2007] as the GGX distribution. Other statistical distributions have been introduced, see e.g., [Bagher et al. 2012; Ribardi re et al. 2017]. The work by Dong et al. [2016] shows that microfacet theory is accurate up to a certain scale, comparing it directly with scalar diffraction theory.

Ashikhmin and Premo ze [2007] extracted the NDF from measured BRDFs. The technique was extended by Dupuy et al. [2015] and Ribardi re et al. [2019]. None of them can solve the multiple-bounce issue of the microfacet model.

*Shadowing-masking.* The shadowing-masking term is important for energy conservation in the microfacet model. It encodes how much of the incoming light was blocked by the microgeometry (shadowing) as well as how much of the reflected light was blocked (masking). To compute it, we need a model of the surface microgeometry. Initial work [Cook and Torrance 1982; Torrance and Sparrow 1967] relied on the V-groove model: for each microfacet, there is another microfacet facing it with the same slope. The V-groove model results in simple computations, as occlusion only depends on the current microfacet slope. The resulting shadowing-masking term has discontinuous derivatives.

Smith [1967] computes the shadowing-masking term from the NDF, assuming that the orientations and positions of the microfacets are independent. The shadowing-masking term is computed from the NDF through a double integration. The resulting term is smooth, and varies more consistently with the roughness of the NDF. Walter et al. [2007] and Heitz [2014] explain and expand the Smith shadowing-masking term for more distributions and take into account the correlation between incoming and outgoing direction.

*Multiple-bounce in microfacet models.* By nature, the microfacet models only express the light reflected after a single bounce on the surface microgeometry. Light that bounces several times is not represented, resulting in an energy loss. The effect is particularly visible for rough surfaces. Kelemen and Szirmay-Kalos [2001] introduce multiple-scattering to the microfacet BSDF by computing the portion of light blocked by the shadowing-masking term and reintroducing it as a diffuse component. The method was extended by Jakob et al. [2014] on dielectric and conductor in layered materials.

Heitz et al. [2016] proposed a multiple-bounce method treating the microfacets randomly distributed microflakes, resulting a random walk solution, which reaches an agreement with the simulated data from surfaces [Heitz and Dupuy 2015]. Dupuy et al. [2016] introduced a unified model between multiple bounce in microsurfaces and microflakes. Sch ussler et al. [2017] extended the approach to normal-mapped surfaces. Westin et al. [1992] encoded multiple scattering by using random walks in microgeometry, and Falster et al. [2020] combined Westin’s approach with wave optics. These methods match the simulated data very well, but do not have an

explicit solution. The random walk simulation results in large variance in the rendered results. Compared to theirs, our method has an explicit formula, although our method still relies on Monte Carlo methods to solve this formula. However, without the need of tracing the height during random walks, our method produces less noise. This explicit formulation enables the use of more advanced light transport methods such as bidirectional path tracing, further reducing the variance.

All these methods use the Smith shadowing model. By contrast, using the V-groove model allows for analytic solutions for multiple-bounce [Lee et al. 2018; Xie and Hanrahan 2018]. The drawbacks are those of the V-groove model: too shiny for rough surfaces, discontinuous derivatives, and not compatible with transparent materials. Lee et al. [2018] redistribute energy to mask the discontinuities, but at the cost of reintroducing randomness.

Kulla and Conty [2017] approximate multiple bounces in microfacets by mixing the single scattering an azimuthally invariant lobe. Turquin [2019] proposed an even simpler multiple bounce computation approach, by scaling the single bounce results. The scaling factor is precomputed based on the surface roughness, the outgoing angle and the index of refraction for dielectrics. These methods are fast, but the multiple bounce term does not have the properties observed in simulations.

Meneveaux et al. [2018] proposed an analytical model for the multiple reflections of light between the interface and the substrate for interfaced Lambertian materials, but ignores the multiple reflections between microfacets. Xie et al. [2019] proposed to represent the multiple scattering with Gaussians or the Real NVP neural network, and used these models for rendering at run-time. Both of two models produce close to energy-conserving results, but with no performance reported.

*Position-free formulation for layered materials.* The position-free path integral formula was proposed by Guo et al. [2018] for the evaluation and sampling of layered materials, and is recently improved by Xia et al. [2020] and Gamboa et al. [2020] with a more advanced sampling method or a more efficient estimator.

The biggest advantage of the position-free formulation is that, it allows explicit representation of light transport in the local subspace of the shaded point. Then, advanced methods and estimators can be exploited to reduce variance. Inspired by this line of work, we formulate the multiple bounce of light transport within BSDFs using the position-free path integral.

### 3 POSITION-FREE MULTIPLE-BOUNCE BSDF FORMULATION

In this section, we describe our path formulation of general light transport for any bounces of microfacet BSDFs. We first introduce our position-free formulation with the definition of a path, then dive into its components on vertices (vertex term) and segments (segment term). We show that the vertex term controls the local light transport that reflects / refracts according to the Fresnel and NDF, while the segment term is responsible for global light transport that accounts for occlusions and multiple scattering. After that, we present detailed derivations of both terms and analyze their properties.

### 3.1 Motivation

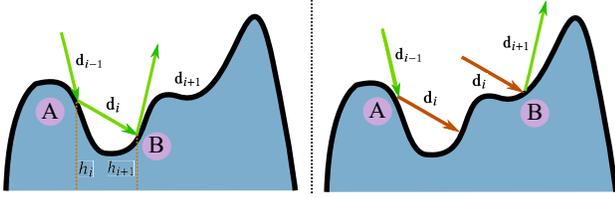


Fig. 2. From the micro scale, Heitz et al. [2016] trace the height for each bounce (left), while our method is based on our position-free observation (right). With our position-free formulation, all points (e.g., A and B) at micro scale share the same statistics, especially the normal distribution. Therefore, any light path bouncing away from a shading point and inflicting the next bounce will be incident to the same shading point again from its last outgoing direction. The independence also explains why both directions require the shadowing-masking function — shadowing from  $\mathbf{d}_i$  to determine at A the proportion of energy able to proceed to the next bounce, and masking from  $-\mathbf{d}_i$  to determine the visible normals at B.

Our key idea is to use a position-free formulation for the integration of multiple-bounces inside the microgeometry. It is based on two observations: first, at the macro scale, BSDFs use a position-free formulation: the point at which the incident light arrives and the point from which the outgoing light leaves are considered to be identical, regardless of how many bounces the light undergoes in the microgeometry. Second, at the micro scale, the Smith shadowing theory [1967] assumes that the positions and normals of the microfacets are uncorrelated. This leads to an obvious but important observation: all points in the micro scale can be considered the same, statistically. BSDFs are also position-free in the micro scale (Fig. 2).

Note specifically that the position-free formulation is irrelevant to specific height distributions. Different height distributions (often referred to as  $P_1$  in related literature, and typically assumed to be either Gaussian or uniform), are only used to derive the shadowing-masking functions. Interestingly, as discussed in Heitz et al. [2016] and implicitly suggested in the Appendix by Walter et al. [2007], the choice of height distribution functions does not even affect the final result of the shadowing-masking functions, since the height distribution is canceled out in the derivation. Our method, following the position-free formulation, is independent of any specific type (e.g., uniform) of height distributions, similar to previous work.

A direct consequence of the position-free formulation is that the outgoing direction for the current bounce is identical to the incoming direction for the next bounce. This leads to two considerations: First, the occlusion comes in pairs between consecutive bounces, so we consider vertices and segments separately. Second, it is possible that an incident ray reaches a microfacet while coming from lower hemisphere of the macrosurface. Thus, we will need a full-spherical formulation of the shadowing-masking functions, instead of the usual hemispherical formulation. Heitz et al. [2016] implicitly suggested the full-spherical definition for their visible normal distribution function (VNDF), resulting in a mixed shadowing-masking function computation (independent of the height for the incoming direction and dependent of the height for the outgoing direction).

However, the full-spherical formulation of the height-correlated shadowing-masking function is still unknown.

### 3.2 Position-free path integral

We define the light transport at any shading point  $\mathbf{s}$ , potentially undergoing multiple bounces, as a path integral for a given pair of query directions  $\omega_i$  and  $\omega_o$ .

The light might bounce several times before exiting the microsurface, and we define each bounce as a **vertex**  $b_i$ . We treat the position of the vertices as identical and focus on the two adjacent directions, rather than positions or depth. This makes our position-free path formulation completely independent of positions, even simpler than for layered materials which requires a depth to be recorded and is only position-free “horizontally”.

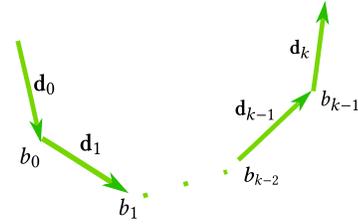


Fig. 3. A light path is defined by  $\bar{x} = (\mathbf{d}_0, b_0, \mathbf{d}_1, b_1, \dots, b_{k-1}, \mathbf{d}_k)$ .

A **direction** is mostly the same as that in macro scale. We use  $\mathbf{d}_i$  as a unit vector on  $\mathcal{S}^2$  to denote the light bouncing among the microfacets. The only difference is that we use the natural flow of light, i.e., assuming the incident pointing inwards instead of outwards any vertex  $b_i$ .

Now we define a **light path**  $\bar{x}$  as a sequence of vertices and directions:  $\bar{x} = (\mathbf{d}_0, b_0, \mathbf{d}_1, b_1, \dots, b_{k-1}, \mathbf{d}_k)$ , as shown in Figure 3. The first and last directions are aligned with the macro incident and outgoing directions of a BSDF query, i.e.,  $\mathbf{d}_0 = -\omega_i$  and  $\mathbf{d}_k = \omega_o$ .

The **path contribution**  $f(\bar{x})$  of a light path is the product of vertex terms  $v_i$  (on each vertex) and segment terms  $s_i$  (on each direction):

$$f(\bar{x}) = s_0 v_0 s_1 v_1 \cdots v_{k-1} s_k. \quad (1)$$

Based on our earlier analysis, we define the **vertex term**  $v_i$  to represent local interactions between the light and the microfacets. It consists of everything except the shadowing-masking term, i.e., the normal distribution function  $D$ , the Fresnel term  $F$  and the Jacobian term together. This vertex term for reflection is defined as:

$$v_i = \frac{F(-\mathbf{d}_i, \omega_h^i) D(\omega_h^i)}{4|\omega_g \cdot (-\mathbf{d}_i)|}, \quad (2)$$

where  $\omega_g$  is the macrosurface normal.

We also define a vertex term for transmission. It has the same formulation as the bidirectional transmission distribution function defined in Walter et al. [2007], without the shadowing-masking

term:

$$v_i = \frac{|\mathbf{-d}_i \cdot \boldsymbol{\omega}_{ht}^i| |\mathbf{d}_{i+1} \cdot \boldsymbol{\omega}_{ht}^i|}{|\boldsymbol{\omega}_g \cdot (-\mathbf{d}_i)|} \frac{\eta_o^2 (1 - F(-\mathbf{d}_i, \boldsymbol{\omega}_{ht}^i)) D(\boldsymbol{\omega}_{ht}^i)}{(\eta_i (-\mathbf{d}_i \cdot \boldsymbol{\omega}_{ht}^i) + \eta_o (\mathbf{d}_{i+1} \cdot \boldsymbol{\omega}_{ht}^i))^2}, \quad (3)$$

where  $\eta_i$  and  $\eta_o$  are indices of refraction for the media on the incident and transmitted sides of the surface, respectively, and

$$\boldsymbol{\omega}_{ht}^i = \frac{\eta_i (-\mathbf{d}_i \cdot \boldsymbol{\omega}_{ht}^i) + \eta_o (\mathbf{d}_{i+1} \cdot \boldsymbol{\omega}_{ht}^i)}{\|\eta_i (-\mathbf{d}_i \cdot \boldsymbol{\omega}_{ht}^i) + \eta_o (\mathbf{d}_{i+1} \cdot \boldsymbol{\omega}_{ht}^i)\|}$$

is the half vector for transmission.

The **segment term**  $s_i$  describes the amount of energy leaving the previous vertex (if any) and arriving at the next (if any). We formulate it in detail in the next subsection.

The path space  $\Omega(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$  is the set of all possible paths with their first directions equal to  $-\boldsymbol{\omega}_i$  and the last directions equal to  $\boldsymbol{\omega}_o$ . Denoting the length of a path  $\bar{x}$  as the number of directions in this path, a subspace  $\Omega_k$  is then the set of all possible paths with the same length  $k$ , and we immediately have  $\Omega = \cup_{k \geq 2} \Omega_k$ .

The path space measure  $\mu(\bar{x})$  is a product of solid angle measures  $\sigma$  at all vertices along the path toward their outgoing directions. That is, for a path with length  $k$ ,

$$\mu(\bar{x}) = \prod_{i=0}^{k-1} \sigma(d_i). \quad (4)$$

Finally, we define **the multiple-bounce BSDF**  $\rho(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$  as an integral over the set of paths  $\Omega(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o)$ :

$$\rho(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o) = \int_{\Omega(\boldsymbol{\omega}_i, \boldsymbol{\omega}_o)} f(\bar{x}) d\mu(\bar{x}). \quad (5)$$

### 3.3 Full-spherical segment term and height-uncorrelated shadowing-masking term

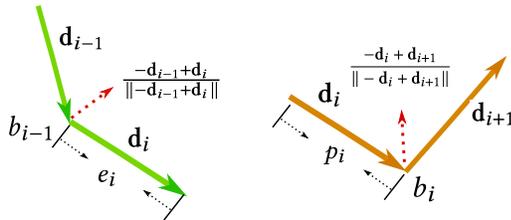


Fig. 4. The segment term  $s_i$  includes two different parts:  $e_i$  and  $p_i$ , where  $e_i$  represents the exit probability for  $\mathbf{d}_i$  at bounce  $i - 1$  and  $p_i$  accounts for the effect of occluding microfacets preventing the direction  $\mathbf{d}_i$  from hitting the next vertex  $b_{i+1}$  at bounce  $i$ .

The meaning of the segment term  $s_i$  is intuitive. It tells the outgoing directional energy distribution at vertex  $b_i$  along direction  $\mathbf{d}_{i+1}$ . But before receiving energy at the next bounce, we need to ensure that the reflected / refracted light will participate in the next bounce. We formulate the segment term into two different parts (Fig. 4) as

$$s_i = e_i \cdot p_i. \quad (6)$$

The first part  $e_i$  is the exit probability and the second part  $p_i$  is the incident probability for the bounce  $i + 1$ . When the light bounces

on vertex  $b_{i-1}$ , suppose there is no shadowing and masking, 100% of the reflected / refracted energy towards  $\mathbf{d}_i$  (given by the vertex term  $v_{i-1}$ ) will remain not occluded as the light leaves the vertex  $b_{i-1}$ . But with potential shadowing and masking, part of the energy will be occluded as the light exits, inflicting the next bounce, while the other part of the energy will never touch the microstructure again, thus stopping further bounces of the light.

**Reflection.** We mathematically define the exit probability as

$$e_i^{(0 < i < k)} = \begin{cases} 1 - G_1(\mathbf{d}_i, \frac{-\mathbf{d}_{i-1} + \mathbf{d}_i}{\|-\mathbf{d}_{i-1} + \mathbf{d}_i\|}), & \text{if } \mathbf{d}_i \cdot \boldsymbol{\omega}_g > 0 \\ 1, & \text{if } \mathbf{d}_i \cdot \boldsymbol{\omega}_g \leq 0, \end{cases} \quad (7)$$

where  $G_1$  is the usual single-sided shadowing-masking term [Smith 1967], the proportion of microfacets that are not occluded from a given direction. Also, when the light bounces downwards the macrosurface, the next bounce is always guaranteed to happen, as the macrosurface is watertight.

There are two special cases — the first and the last segments. For the first segment, since it does not have to exit any previous vertex, its value is always  $e_0 = 1$ . The other special case on the last bounce is easily understood, as we would like to continue the next bounce for all other vertices except the last one, where we actually need the path to stop bouncing to keep its total length to  $k$ . Therefore, the exit probability becomes the “inverse” of others as

$$e_i^{(i=k)} = \begin{cases} G_1(\mathbf{d}_i, \frac{-\mathbf{d}_{i-1} + \mathbf{d}_i}{\|-\mathbf{d}_{i-1} + \mathbf{d}_i\|}), & \text{if } \mathbf{d}_i \cdot \boldsymbol{\omega}_g > 0 \\ 0, & \text{if } \mathbf{d}_i \cdot \boldsymbol{\omega}_g \leq 0, \end{cases} \quad (8)$$

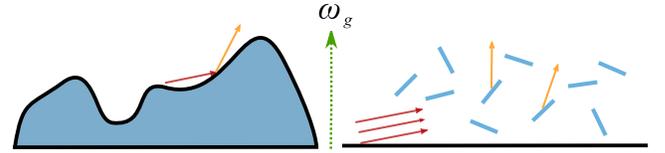


Fig. 5. Under the microfacet theory, it is possible that an incident ray coming from the lower hemisphere (the other side of the macro normal) be reflected, either by the actual microsurfaces (left), or by the microflakes in the Smith theory (right).

Now we take a look at the incident probability  $p_i$ , which accounts for the effect of occluding microfacets preventing the direction  $\mathbf{d}_i$  from hitting the next vertex  $b_{i+1}$  (if any). From the definition, we immediately know that this is the single-sided shadowing-masking term  $G_1$ . However, in this case, we must explicitly deal with the possible incident direction  $\mathbf{d}_i$  from below the macrosurface to the vertex  $b_{i+1}$ , as seen in Fig. 5.

The full-spherical shadowing-masking term  $G_1$  was implicitly indicated by Heitz et al. [2016] as the absolute value of its original hemispherical version. But to our knowledge, there is no explicit explanation or derivation available. Therefore, we provide a detailed derivation in the supplemental materials, and provide our conclusion:

$$G_1(\boldsymbol{\omega}, \boldsymbol{\omega}_m) = G_1^{\text{local}}(\boldsymbol{\omega}, \boldsymbol{\omega}_m) G_1^{\text{dist}}(\boldsymbol{\omega}), \quad (9)$$

where  $G_1^{\text{local}}(\boldsymbol{\omega}, \boldsymbol{\omega}_m) = \chi^+(\boldsymbol{\omega} \cdot \boldsymbol{\omega}_m)$ , which is 1 when  $\boldsymbol{\omega}$  and  $\boldsymbol{\omega}_m$  are facing in the same direction, i.e.,  $\boldsymbol{\omega} \cdot \boldsymbol{\omega}_m > 0$ , and is 0 otherwise.

$G_1^{\text{dist}}$  is the distant shadowing / masking term, in our full-spherical case:

$$G_1^{\text{dist}}(\omega) = \left| \frac{1}{1 + \Lambda(\omega)} \right| = \begin{cases} 1/(1 + \Lambda(\omega)), & \text{if } \omega \cdot \omega_g > 0, \\ -1/(1 + \Lambda(\omega)), & \text{if } \omega \cdot \omega_g \leq 0, \end{cases} \quad (10)$$

$\Lambda(\omega)$  is the Smith Lambda function, which is analytical for both Beckmann and GGX models. More details could be found in Heitz et al. [2016].

With the explicit full-spherical shadowing-masking function, we are able to define  $p_i$  as

$$p_i^{(i < k)} = G_1 \left( -\mathbf{d}_i, \frac{-\mathbf{d}_i + \mathbf{d}_{i+1}}{\|-\mathbf{d}_i + \mathbf{d}_{i+1}\|} \right), \quad (11)$$

and  $p_k = 1$  since the last segment exiting the surface will not hit any more vertices.

*Transmission.* The formula for transmission is more complex than the reflection case. In the microfacet model, the incident angle between a ray and the macrosurface normal cannot decide whether the ray originates from the outside or the inside of an object, especially with the shadowing-masking function extended to the full-spherical domain. Therefore, it is common to keep track of and switch the status (inside or outside) along a path [Heitz et al. 2016] whenever a refraction event happens. We follow this approach, and we also analyze the exit probability based on the incident angle, resulting in four cases. The exit probability for the middle bounce is:

$$e_i^{(0 < i < k)} = \begin{cases} 1 - G_1 \left( \mathbf{d}_i, \omega_{ht}^i \right), & \text{if } \mathbf{d}_i \cdot \omega_g > 0 \text{ and } \mathbf{d}_{i-1} \text{ is inside} \\ 1 - G_1 \left( -\mathbf{d}_i, \omega_{ht}^i \right), & \text{if } \mathbf{d}_i \cdot \omega_g < 0 \text{ and } \mathbf{d}_{i-1} \text{ is outside} \\ 1, & \text{if } \mathbf{d}_i \cdot \omega_g \leq 0 \text{ and } \mathbf{d}_{i-1} \text{ is inside} \\ 1, & \text{if } \mathbf{d}_i \cdot \omega_g \geq 0 \text{ and } \mathbf{d}_{i-1} \text{ is outside.} \end{cases} \quad (12)$$

These four cases correspond to a combination of positions and directions for  $\mathbf{d}_i$  and  $\mathbf{d}_{i-1}$ . The first two cases are similar to the first case of Eqn. 7, the last two cases to the second case. In the first case, the ray goes from the inside to the outside of the surface, and the outgoing direction  $\mathbf{d}_i$  points upwards. The outgoing ray intersects the surface with probability  $1 - G_1 \left( \mathbf{d}_i, \omega_{ht}^i \right)$ . The second case is the reverse: the ray goes from the outside to the inside, and the outgoing direction points downwards. In the third case, the ray goes from the inside of the surface towards the outside, and the outgoing direction  $\mathbf{d}_i$  points downwards. As the macrosurface is continuous, the outgoing ray will always intersect the surface. The fourth case is similar, with reversed incoming and outgoing directions.

For the last bounce, the exit probability is:

$$e_i^{(i=k)} = \begin{cases} G_1 \left( \mathbf{d}_i, \omega_{ht}^i \right), & \text{if } \mathbf{d}_i \cdot \omega_g > 0 \text{ and } \mathbf{d}_{i-1} \text{ is inside} \\ G_1 \left( -\mathbf{d}_i, \omega_{ht}^i \right), & \text{if } \mathbf{d}_i \cdot \omega_g < 0 \text{ and } \mathbf{d}_{i-1} \text{ is outside} \\ 0, & \text{otherwise.} \end{cases} \quad (13)$$

The first and the second cases are the same as Eqn. 12, except that the ray  $\mathbf{d}_i$  should leave the surface with probability  $G_1 \left( \mathbf{d}_i, \omega_{ht}^i \right)$ , similar to Eqn. 8. In the other cases, the ray will be blocked, resulting in zero probability to exit.

Similar to the reflection case, the  $p_i$  for transmission is defined as

$$p_i^{(i < k)} = \begin{cases} G_1 \left( -\mathbf{d}_i, \omega_{ht}^i \right), & \text{if } -\mathbf{d}_i \text{ is outside,} \\ G_1 \left( \mathbf{d}_i, \omega_{ht}^i \right), & \text{if } -\mathbf{d}_i \text{ is inside,} \end{cases} \quad (14)$$

and  $p_k = 1$  since the last segment exiting the surface will not hit any more vertices.

So far, we have the complete segment term  $s_i$  derived. Note that there is no double counting of occlusion along the same direction  $\mathbf{d}_{i+1}$ , as  $e_i$  is for exiting vertex  $b_i$  and  $p_{i+1}$  is for entering vertex  $b_{i+1}$ . They are different vertices, even with our position-free formulation.

### 3.4 Height-correlated shadowing-masking term

In the previous section, we derive our multiple-bounce BSDF model using a separable fashion: the incident and outgoing contributions are completely separated into two different segment terms. We used the height-uncorrelated version of the shadowing-masking function at each bounce, a product of two  $G_1$  functions, each independent of specific heights. This first model provides energy conservation and passes the white furnace test (Sec. 5). The height-correlated shadowing-masking function is more accurate than the height-uncorrelated version, since it models the correlation between masking and shadowing due to the height of the microsurface [Ross et al. 2005]. Therefore, we want to use it to compute the multiple-bounce part of the BSDF.

In this section, we derive a height-correlated multiple-bounce BSDF model, which is full-spherical and fits our position-free formulation. We present below our formulations for reflection and transmission, then organize them in the segment term for the multiple-bounce part of the BSDF. The height-correlated shadowing-masking function analyzes the visibilities at all heights and then averages/integrates them at a specific bounce, thus is conceptually different to “keeping specific heights at different bounces” as in previous work [Heitz et al. 2016].

*Reflection.* The height-correlated distant shadowing-masking function for reflection is [Ross et al. 2005]:

$$G_2^{\text{dist}}(\omega_i, \omega_o) = \frac{1}{1 + \Lambda(\omega_i) + \Lambda(\omega_o)}, \quad (15)$$

where both  $\omega_i$  and  $\omega_o$  are from the upper hemisphere:  $\omega_i \cdot \omega_g > 0$  and  $\omega_o \cdot \omega_g > 0$ .

To extend Eqn. 15 to the full-spherical domain, the derivation must start from the definitions of  $G_2^{\text{dist}}$  and  $\Lambda$ . We leave the details in the supplementary material, and provide our re-organized formulation for Eqn. 15:

$$\frac{1}{1 + \Lambda(\omega_i) + \Lambda(\omega_o)} = \frac{\cos\theta_i \cos\theta_o}{\int_{\Omega^+} (\langle \omega_i, \omega_m \rangle \cos\theta_o + \langle -\omega_o, \omega_m \rangle \cos\theta_i) D(\omega_m) d\omega_m}, \quad (16)$$

which is currently still limited to  $\cos\theta_i = \omega_i \cdot \omega_g > 0$  and  $\cos\theta_o = \omega_o \cdot \omega_g > 0$ .

We denote the height-correlated distant shadowing-masking function as  $G_2^{(i=k)}(\omega_i, \omega_o)$  for the last bounce and  $G_2^{(i < k)}(\omega_i, \omega_o)$  for the middle bounce for clarity.

From Eqn. 16, we first propose our full-spherical formulation for the last bounce. This is similar to the traditional shadowing-masking function which gives the probability that neither  $\omega_i$  nor  $\omega_o$  will be blocked:

$$G_2^{(i=k)}(\omega_i, \omega_o) = \begin{cases} \frac{1}{-\Lambda(-\omega_i) + \Lambda(\omega_o)}, & \text{if } \omega_i \cdot \omega_g > 0, \\ \frac{1}{\Lambda(-\omega_i) + \Lambda(\omega_o)}, & \text{if } \omega_i \cdot \omega_g \leq 0. \end{cases} \quad (17)$$

Since only one of  $\omega_i$  and  $\omega_o$  can come from the lower hemisphere, we assume that it's  $\omega_i$ , without loss of generality.

The shadowing-masking function for the middle bounces were never computed before. To compute them, we need to extend our understanding of shadowing-masking functions: we define them as the probability that *a specific bounce can happen*. This corresponds with our earlier analysis regarding the segment terms. This definition reduces automatically to the traditional shadowing-masking function for single-bounce BSDFs, and it is suitable for our last bounce as well. At the middle bounces, our extended definition immediately tells us that we need the incident direction unblocked and the outgoing direction blocked, so the middle bounce can happen (otherwise, it must be the last bounce since it will exit without being blocked):

$$G_2^{(i<k)}(\omega_i, \omega_o) = \begin{cases} \frac{1}{|\Lambda(-\omega_i)|} - \frac{1}{|\Lambda(-\omega_i)| + \Lambda(\omega_o)}, & \text{if } \omega_o \cdot \omega_g > 0, \\ \frac{1}{|\Lambda(-\omega_i)|}, & \text{if } \omega_o \cdot \omega_g \leq 0, \end{cases} \quad (18)$$

Again, we leave the detailed derivation in the supplementary.

*Transmission.* Similar to the reflection case, we also derive the height-correlated shadowing-masking functions for transmission. The original height-correlated shadowing-masking function [Pinel et al. 2005] is

$$G_2^{\text{dist}}(\omega_i, \omega_o) = B(1 + \Lambda(\omega_i), 1 + \Lambda(\omega_o)), \quad (19)$$

where  $B$  is the Beta function and  $\omega_i$  and  $\omega_o$  are on opposite sides of the microsurface.

This formulation is valid on a hemisphere; it does not hold in the full-spherical case: If the ray  $\omega_i$  locates at the front-side of the microsurface, Eqn. 19 requires that it should also at the front side of the macrosurface. However,  $\omega_i$  could come from the lower hemisphere of the macrosurface, and still at the front-side of the microsurface.

Similar to the reflection case, we first extend Eqn. 19 to the full-spherical domain for the last bounce as follows:

$$G_2^{(i=k)}(\omega_i, \omega_o) = B(|\Lambda(-\omega_i)|, 1 + \Lambda(\omega_o)), \quad (20)$$

then, for the middle bounce, our new formulation is:

$$G_2^{(i<k)}(\omega_i, \omega_o) = \begin{cases} \frac{1}{|\Lambda(-\omega_i)|}, & \text{if } (\omega_o \cdot \omega_m)(\omega_o \cdot \omega_g) \leq 0, \\ \frac{1}{|\Lambda(-\omega_i)|} - B(|\Lambda(-\omega_i)|, 1 + \Lambda(\omega_o)), & \text{otherwise.} \end{cases} \quad (21)$$

The derivation is similar to the reflection case.

*Height-correlated segment term.* We can now express the amount of energy that undergoes further bounces, using our new definition for the height-correlated shadowing-masking term. We rewrite the exit probability  $e_i$  at each bounce using our height-correlated

shadowing-masking functions:

$$e_i^{(i=k)} = \begin{cases} \frac{G_2^{(i=k)}(-\mathbf{d}_{i-1}, \mathbf{d}_i)}{G_1^{\text{dist}}(-\mathbf{d}_{i-1})}, & \text{if } \mathbf{d}_i \cdot \omega_g > 0 \\ 0, & \text{if } \mathbf{d}_i \cdot \omega_g \leq 0, \end{cases} \quad (22)$$

and

$$e_i^{(i<k)} = \frac{G_2^{(i<k)}(-\mathbf{d}_{i-1}, \mathbf{d}_i)}{G_1^{\text{dist}}(-\mathbf{d}_{i-1})}. \quad (23)$$

The  $p_i$  term is unchanged (Eqn. 11). When multiplying the segment terms for all bounces together, the denominator  $G_1^{\text{dist}}$  in  $e_i$  will be canceled out by  $p_{i-1}$ , leaving only the height-correlated term for each bounce.

### 3.5 Properties and analysis

Now that we have a complete multiple-bounce BSDF formulation, we briefly analyze it to check that it has the right properties.

*Position-free.* Our BSDF formulation is completely independent of the positions of individual vertices along a light path. This immediately demonstrates that both in the macro scale and in the micro scale, our method is position-free. As a result, there is no need to keep track of the height of a vertex as in [Heitz et al. 2016]. Our far-field assumption comes from the commonly-used derivation of the height-correlated or uncorrelated shadowing masking terms [Heitz 2014; Walter et al. 2007]. In both height-correlated and uncorrelated cases, heights are implicitly integrated.

*Generality.* One can easily verify the generality of our path formulation in Eqn. 5. It reduces to classic single-bounce BSDFs, if we limit the length of paths to 2. Therefore, our formulation is a general definition of BSDFs.

*Normal mapping support.* In Schussler et al. [2017], it is pointed out that regular normal mapping on hemispherical BSDFs will inevitably confuse the sides of incident and outgoing directions, leading to black regions when the specified normals deviate much from the original. However, since our BSDF formulation is full-spherical, directly applying normal mapping will never cause similar issues. Therefore, no additional effort needs to be done to support correct normal mapping. We demonstrate this in Fig. 7.

*Variance reduction.* Our variance reduction comes from two factors: firstly, our position-free formulation does not explicitly trace heights, leading to much less variation in the integrand. Secondly, we can deploy more advanced Monte Carlo estimators such as a bidirectional path tracers, leading to further variance reduction. We consider the aggregated contribution from all heights at each bounce in our model, while Heitz et al. [2016] explicitly trace the height, leading to more variance. As mentioned before, our explicit path integration enables any Monte Carlo solutions to it. This property allows us to introduce much more efficient estimators, such as Bidirectional Path Tracing (BDPT) than previous random walk methods, which reduces the variance significantly, as we show next in Sec. 4. Note specifically that Heitz et al. [2016] use multiple importance sampling (MIS) to combine the contribution from two random walk paths, one pure forward and the other backward. This approach is far from a complete BDPT, since it ignores all the connections

between internal bounces. Our explicit formulation is different from theirs and allows full bidirectional approaches, enabling connection of half paths from both directions, reusing much more samples and resulting in less variance.

*Reciprocity.* From the structure of the path integral, we can see that its overall reciprocity lies in the individual vertex terms and segment terms. The reciprocity of the vertex terms is trivial to verify, since they are essentially traditional microfacet BSDFs without the shadowing-masking terms. In the supplemental materials, we prove that the height-uncorrelated segment terms are also reciprocal, and provide an experiment to demonstrate this. Therefore, our entire height-uncorrelated BSDF formulation is reciprocal. However, the reciprocity of height-correlated BSDF formulation does not hold. A counter example is also provided in the supplementary. The different reciprocity properties are because we rely on different height correlation, which will be discussed next.

*Height correlation.* We have presented our position-free BSDF model using both the height-correlated and the height-uncorrelated shadowing-masking functions. We first emphasize again that both are reasonable approximations to the real visibility/occlusion problem at micro scale, both satisfy energy conservation and pass the white furnace test, but neither requires a specific height at each bounce. The height-correlated shadowing-masking function ignores the height, while the height-uncorrelated averages/integrates the height.

Quality-wise, we demonstrate that our height-correlated model has a better match with the simulated ground truth (Fig. 6), and is also closer to Heitz et al. [2016]. The main difference between our height-correlated multiple-bounce BSDF and Heitz et al. [2016] is that our model assumes the position-free property among bounces, which means the incident rays for each bounce always come from the distance. Therefore, we can see that the random walk method by Heitz et al. [2016] is neither height-uncorrelated nor height-correlated. Precisely speaking, it should be labeled *height-specific* in light transport, and *height-mixed* in its shadowing-masking functions. The former is straightforward to understand, because it does record specific heights at different bounces. The second term is slightly out of the scope of our focus, and we simply point out that Heitz et al. [2016] use height-correlated shadowing-masking term for the first bounce and use separable shadowing-masking function for other bounces: considering the height for the masking function, while ignoring the height for the shadowing function.

A brief summary on different height correlation models: our height-uncorrelated model is simple to compute, reciprocal, but less accurate (only in terms of the closeness to simulated multiple-bounce energy distribution); our height-correlated model can still be explicitly written, is closer to the ground truth, but is not reciprocal; Heitz et al. [2016] is height-specific, slow to converge, similarly close to the ground truth as compared to our height-correlated model, but is reciprocal. We can conclude from the comparison that different height correlation models indeed play an important role in the final appearance model.

## 4 MONTE CARLO PATH INTEGRAL ESTIMATORS

With our explicit and position-free path formulation, any Monte Carlo method can be used to compute the integral. In this section, we propose two estimators for BSDF evaluation: unidirectional path tracing (PT) and bidirectional path tracing (BDPT) to evaluate the multiple scattering path integral, inspired by the position-free integral that solves the BSDFs of layered materials [Guo et al. 2018]. We show how to efficiently sample our multiple-bounce BSDFs, and the computation of the corresponding probability density functions. Throughout this section, we use conductors as examples, without loss of generality.

### 4.1 Path Tracing

We first propose a unidirectional estimator using path tracing for BSDF evaluation:

$$\rho(\omega_i, \omega_o) \approx \frac{1}{N} \sum_{j=0}^{j=N} \frac{f(\bar{x})}{\text{pdf}(\bar{x})}, \quad (24)$$

where  $N$  is the sample count,  $\bar{x}$  is a sampled path starting from  $\mathbf{d}_0$  and  $\text{pdf}(\bar{x})$  is the probability density function (PDF) of the sample path. We set  $N$  as 1 for each BSDF evaluation. Since the path has to be ended with  $\omega_o$ , it's impossible to reach such a direction with directional sampling only, thus we perform the next event estimation (NEE) from the final outgoing direction  $\omega_o$  for each bounce, resulting in:

$$\rho(\omega_i, \omega_o) \approx \frac{1}{N} \sum_{j=0}^{j=N} \sum_{i=2}^{i=k+1} \frac{f(\bar{x}_i)}{\text{pdf}(\bar{x}_i)}, \quad (25)$$

where  $\bar{x}_i = (\mathbf{d}_0, b_0, \mathbf{d}_1, b_1, \dots, b_{i-2}, \mathbf{d}_{i-1})$  represents a path with length  $i$  from the path space  $\Omega_i$  and  $f(\bar{x}_i)$  is computed with Eqn. 1. The PDF of a path  $\text{pdf}(\bar{x}_i)$  is computed as the product of all the PDF to sample the internal directions, from  $\mathbf{d}_1$  to  $\mathbf{d}_{i-2}$ :

$$\text{pdf}(\bar{x}_i) = \prod_{j=0}^{j=i-3} \text{pdf}(-\mathbf{d}_j, \mathbf{d}_{j+1}), \quad (26)$$

where  $\text{pdf}(-\mathbf{d}_j, \mathbf{d}_{j+1})$  represents the PDF of sampling  $\mathbf{d}_{j+1}$  from  $-\mathbf{d}_j$ . Note that both  $f(\bar{x}_i)$  and  $\text{pdf}(\bar{x}_i)$  are evaluated recursively.

For the specific sampling method and corresponding PDF values, we simply refer to the distribution of visible normals (VNDF) sampling technique [Heitz and d'Eon 2014].

Finally, note that our path tracing method provides an unbiased estimation of the multiple-bounce BSDF with Russian roulette [Arvo and Kirk 1990]. This is in essence different from the multiple-bounce BSDFs under the V-groove assumption [Lee et al. 2018], in which a maximum bounce must be specified, balancing between potential energy cutoff and computational overhead.

### 4.2 Bidirectional path tracing

We now present an even more efficient bidirectional estimator, following the classical bidirectional approach in light transport. We first trace rays from both  $\mathbf{d}_0$  and  $-\mathbf{d}_k$  with maximum length  $k + 1$  and generate a camera path and a light path. Then we combine them, choosing  $s$  directions from camera path, and  $t$  directions from the light path, where  $0 < s < k + 1$  and  $0 < t < k + 2 - s$ . If  $s = 1$ , then only  $\mathbf{d}_0$  is chosen from the camera path. Similarly, if  $t = 1$ , then

only  $\mathbf{d}_k$  is taken from the light path. For each generated path  $\bar{x}_i$ , we compute its contribution  $f(\bar{x}_i)$ ,  $\text{pdf}(\bar{x}_i)$ , and the MIS weight  $w(\bar{x}_i)$ .

The path contribution  $f(\bar{x}_i)$  of path  $\bar{x}_i$  with length  $i$  is computed with Eqn. 1, by accumulating all the vertex terms and the segment terms along the path. Since our height-correlated BRDF is not reciprocal, to ensure consistency, we compute the contribution of all paths from the macro-scale incident direction  $\omega_i$ , to the macro-scale outgoing direction  $\omega_o$ , regardless of whether the paths are constructed by backward path tracing (in PT) or by connecting the vertices from the two subpaths (in BDPT).

The PDF is computed by accumulating the  $s$  PDFs from the camera path and  $t$  PDFs from the light path:

$$\text{pdf}(\bar{x}_i) = \prod_{j=1}^{j=s} \text{pdf}(-\mathbf{d}_{j-1}, \mathbf{d}_j) \prod_{j=k-s}^{j=k} \text{pdf}(\mathbf{d}_j, -\mathbf{d}_{j-1}), \quad (27)$$

Note that, we start from 1 rather than 0, as the both  $\mathbf{d}_0$  and  $\mathbf{d}_k$  are not sampled.

Regarding the MIS weight, for a given path  $\bar{x}_i$  with length  $i$ , there are  $i - 1$  possible ways to generate this length, by taking different numbers of directions from the camera path and the light path. We sum up all the PDF for each possible way as  $\sum \bar{x}_i^j$ , and then compute the MIS weight with the balance heuristic, as:

$$w(\bar{x}_i) = \frac{\text{pdf}(\bar{x}_i)}{\sum \text{pdf}(\bar{x}_i^j)}. \quad (28)$$

Finally, we get the bidirectional estimator for multiple-bounce BSDF as:

$$\rho(\omega_i, \omega_o) = \sum_{i=2}^{i=k+1} \frac{f(\bar{x}_i)w(\bar{x}_i)}{\text{pdf}(\bar{x}_i)}. \quad (29)$$

The bidirectional estimator produces results with less variance, since there are implicitly more paths used for estimation. The paths are also weighted in a proper way, which allows to further reduce variance.

Also, from the definition of paths in our formulation, we can see that they are completely consisted with directions, thus is completely position-free. This is different from Guo et al. [2018], where the position-independence is only in the horizontal direction, while they keep trace of the depths into the different layers for their path formulations.

### 4.3 Importance sampling and PDF of multiple bounces

Importance sampling is required to fit our BSDF in a path tracing framework. It's straight forward to do sampling in our BSDF. For a given incident direction  $\omega_i$ , the sample function should answer the final  $\omega_o$  and compute the sampling weights.

Let's consider bounce  $i$ . Starting from  $\mathbf{d}_i$ , we use the VNDF importance sampling to generate an outgoing direction  $\mathbf{d}_{i+1}$ . With this sampled direction, we compute the Fresnel factor and use it as a probability to decide reflecting or refracting the ray. Then we check whether  $\mathbf{d}_{i+1}$  points towards the macrosurface. If so, we continue sampling. Otherwise, we compute the masking function  $G_1(\mathbf{d}_{k+1})$  and use it as a probability to choose to leave the surface or performing more bounces. If choosing to leave the surface,  $\omega_o$  is obtained.

If the maximum bounces are reached, but the ray has not left the surface, the sampling fails, which is usual in the light transport.

For conductor material, the sampling weights include all the Fresnel factor along the path, since all the other terms are canceled out by the VNDF sampling and the exit surface sampling. For a rough dielectric BSDF, the weight is 1, since all the terms are canceled out.

The PDF function of a BSDF is used in MIS. Given an  $\omega_i$  and  $\omega_o$ , it should answer the PDF for this setting. Since we are using a random walk in our BSDF, it's impossible to get the exact PDF for an  $\omega_i$  and  $\omega_o$  pair. Hence, we use the same method as Heitz et al. [2016], combining the PDF of single-scattering and a diffuse term PDF.

## 5 RESULTS AND COMPARISON

We have implemented our algorithm inside the Mitsuba renderer [2010] for both rough conductor and rough dielectric BSDFs. The implementation of Heitz et al. [2016] is from the author's website, with bidirectional random walk. All timings in this section are measured on a 2.20GHz Intel i7 (48 cores) with 32 GB of main memory. For the reference images, we simply refer to the converged result using the same method. This is because in theory there is no ground truth, and Heitz et al. [2016] and our method converge to different results. Therefore, to have a fair estimation of noise level, we compare different methods with the converged results of these methods as references.

*Comparison of lobes for individual bounces.* In Fig. 6 and the supplemental materials, we compare the visualized lobes for individual bounce between our method (both height-uncorrelated and height-correlated), Heitz et al. [2016] and the simulated data which is obtained by ray tracing on a generated surface with Beckmann distribution [Heitz and Dupuy 2015]. We perform the comparison on rough diffuse (albedo set as 1), rough conductor (Fresnel set as 1) and rough dielectric BSDFs, considering both isotropic ( $\alpha = 0.25$ ,  $\alpha = 0.5$  and  $\alpha = 1$ ) and anisotropic ( $\alpha = (1.0, 0.1)$ ) cases. We visualize the lobes with  $\omega_i$  elevation angles of 0.0, 0.5 1.0 and 1.5 radians.  $E_r$  and  $E_t$  denote the total amount of reflected and transmitted energies, respectively. We only show the images for materials with roughness  $\alpha = 1$  with  $\omega_i$  elevation angles of 0.0 and 1.5 radians in Fig. 6, and more images are shown in the supplemental materials.

For all the bounces with all the incident angles, our height-correlated model produces very similar results as Heitz et al. [2016], while our height-uncorrelated model has larger difference from Heitz et al. [2016] mostly at grazing angles. Note that all these three methods pass the white furnace test.

*Evaluation-only comparison.* Our method is especially suitable to render under sharp lighting and in evaluation-heavy situations, thus, we first show some results with direct lighting only. In Fig. 8, we show a yellow rough diffuse material (GGX model,  $\alpha = 1.0$ ) lit by a directional light source. Even our model (PT) shows less variance than Heitz et al. [2016]. In Fig. 9, we show a copper sphere (GGX model,  $\alpha = 1.0$ ) lit by a directional light. Both of our PT and BDPT methods produce much less noise than Heitz et al. [2016], with a slight performance overhead, and our BDPT approach produces the best result. Fig. 1 (c) and (d) show a single slab with a rough dielectric BSDF (GGX model,  $\alpha = 1.0$ ). Both our path tracing and our BDPT

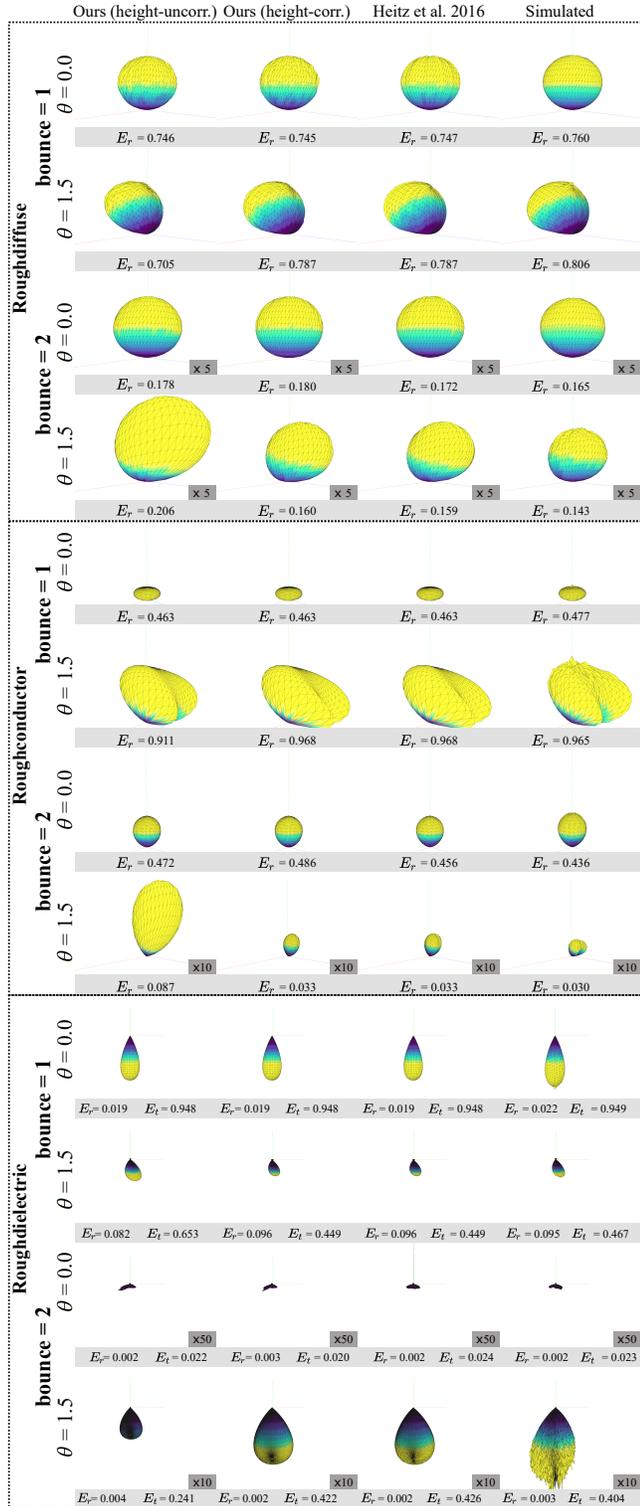


Fig. 6. Comparison between our multiple-bounce BSGF models (both height-uncorrelated and height-correlated), Heitz et al. [2016] model and simulated data, for rough diffuse, conductor and dielectric materials with roughness 1.0.  $\theta$  is the angle between the incident direction and the normal to the macrosurface.

ACM Trans. Graph., Vol. 41, No. 4, Article 1. Publication date: July 2022.

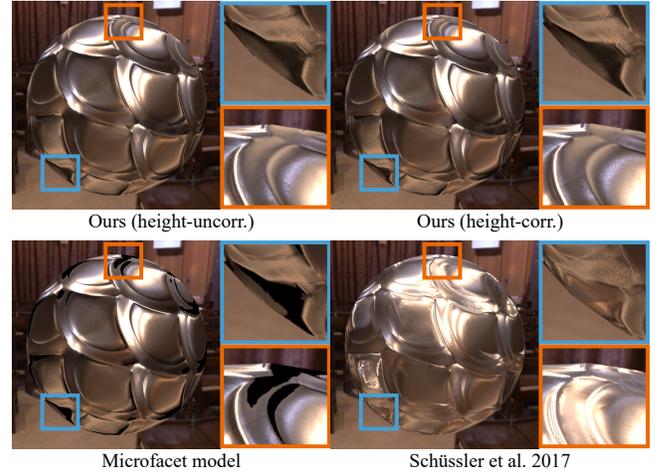


Fig. 7. Comparison between our methods, Schüssler et al. [2017] and the microfacet model on a metal sphere with a normal map. Our method does not have the black artifacts in microfacet models. Since Schüssler et al. [2017] use a different model, it produces a different result from ours, as expected.

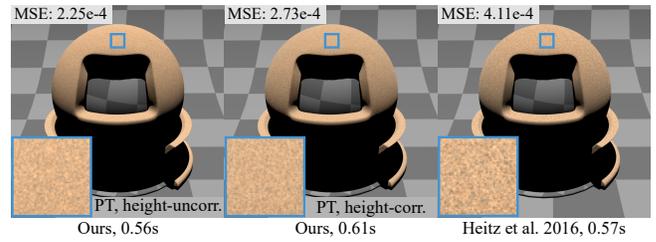


Fig. 8. Comparison between our methods (PT) and Heitz et al. [2016] on a rough diffuse material (GGX model,  $\alpha = 1.0$ ) under a directional lighting with 64 spp.

produce better results than Heitz et al. [2016], while our path tracing method is faster than Heitz et al. [2016] and our BDPT reduces the noise significantly with acceptable extra time cost. With only two samples per pixel, our method (BDPT) produces results close to noise-free. In Fig. 10, we show the Mean Square Error (MSE) as a function of varying sampling rate for our method (BDPT and PT) and Heitz et al. [2016] in the Sphere scene and the Single Slab scene, considering directional lighting only. With only two samples per pixel, our method is able to produce very close result to the ground truth, while Heitz et al. [2016] produces result with a lot of noise. Increasing the sampling rate improves the quality for both methods, but our method remains consistently better. In the supplemental materials, we show more convergence comparisons between our method and Heitz et al. [2016] over varying roughness. For all these configurations, our method shows better quality.

*Equal-time comparison.* In Fig. 1(a) and (b), we show three deer statues (copper (GGX model,  $\alpha = 0.1$ ), aluminum (GGX model,  $\alpha = 0.6$ ) and gold (GGX model,  $\alpha = 0.5$ )) on an aluminum floor (GGX model,  $\alpha = 0.1$ ), lit by an environment map and a point light. To better show the effect of the BSGF evaluation, we use 64 spp for

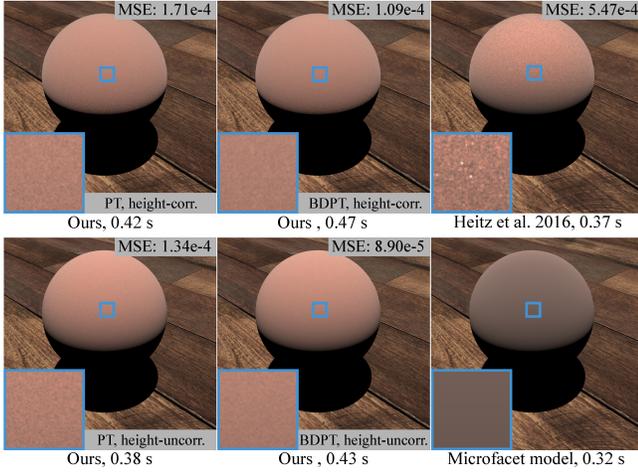


Fig. 9. Comparison between our methods (PT), our methods (BDPT), Heitz et al. [2016] and the classical microfacet model on a rough conductor BSDF (GGX model,  $\alpha = 1.0$ ) with 4 spp. Both our unidirectional and bidirectional methods produce less noise than Heitz et al. [2016], while our bidirectional approach produces even higher-quality results than our path tracing approach, with some extra time cost.

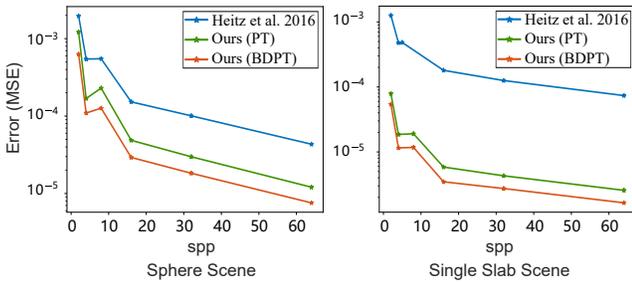


Fig. 10. The error (MSE) with logarithm scale of our method (both BDPT and PT) and Heitz et al. [2016] over varying sampling rate on the Sphere scene (Figure 9) and Single Slab scene (Figure 1(c) and (d)). In both scenes, our method with PT is constantly better than Heitz at all different sampling rates, and our method with BDPT always gives the best results.

the environment map lighting. To achieve equal time, we use 8 spp for our method and 14 spp for Heitz et al. [2016] for the point light source. Our results have much less noise than Heitz et al. [2016]’s result. We also report the MSE of the entire image, which confirms the high quality of our results.

*Comparison with other methods.* In Fig. 11, we compare our method with Lee et al. [2018] (nonsymmetric). The result of Lee et al. [2018] inherits the drawback of V-groove approaches, and the results look too glossy. our method and Heitz et al. [2016] are based on the Smith shadowing method, and produce similar results. In Fig. 12, we compare our method with Kulla and Conty [2017]. Their method is based on an average Fresnel term, instead of accumulating the Fresnel term contributions during multiple scattering. This results in a

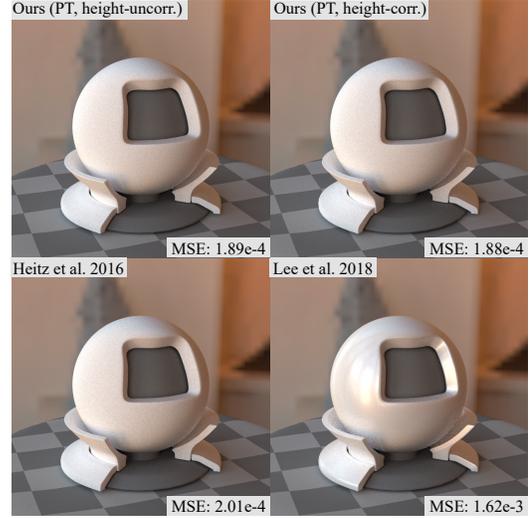


Fig. 11. Comparison between our methods (PT), Heitz et al. [2016] and Lee et al. [2018] for a rough conductor BSDF (Beckmann model,  $\alpha = 1.0$ ). The result by Lee et al. [2018] is too glossy; this issue is a typical drawback of the V-groove shadowing-masking model.

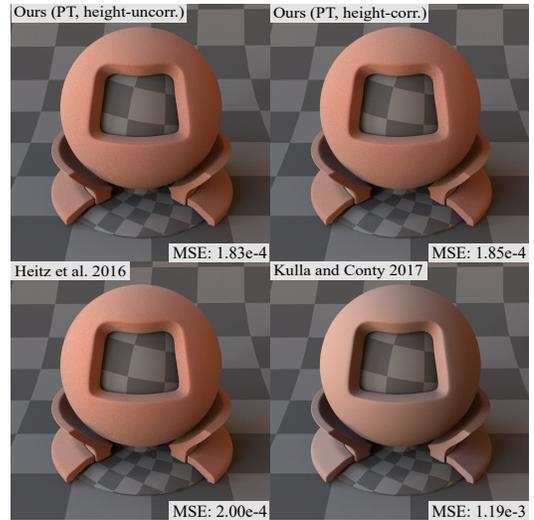


Fig. 12. Comparison between our methods (PT), Kulla and Conty [2017] and Heitz et al. [2016] for a rough conductor BSDF (GGX model,  $\alpha = 1.0$ ). Kulla and Conty [2017] average the Fresnel term over all directions, resulting in a significantly different color.

significant difference in color. Our method and Heitz et al. [2016] use the Fresnel term at each bounce, resulting in the correct color.

*Complex lighting comparison.* In Fig. 13, we compare against Heitz et al. [2016] on more complex lighting, considering both indirect illumination and environment lighting. Our results are almost identical to Heitz et al. [2016]. In this scene, our method does not significantly reduce the noise. There are two reasons: first, in this scene, light

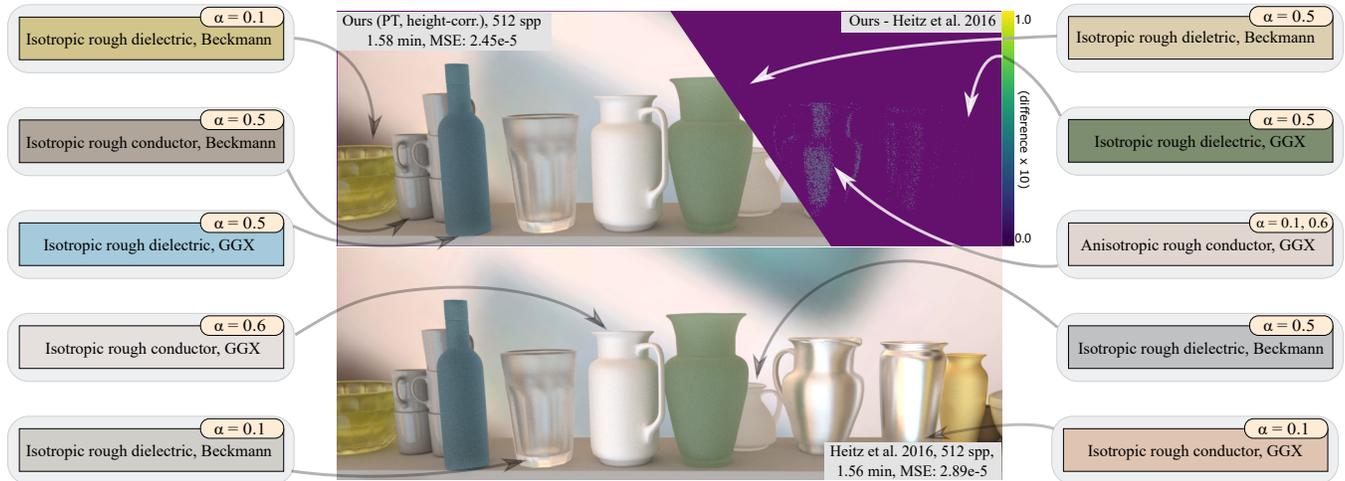


Fig. 13. Our result (rendered with PT, height-correlated) and Heitz et al. [2016] on different materials at roughly equal time. The material settings are shown in the image. We visualize the difference ( $5\times$  scaled) between ours and Heitz et al. [2016].

transport is complex and responsible for most of the noise; second, our method for sampling BSDFs produces similar noise as Heitz et al. [2016]. Hence, our method is especially suitable to render under sharp lighting and in evaluation-heavy situations.

In Fig. 14, we show the results of our methods (PT) with height-correlated and height-uncorrelated segment terms, Heitz et al. [2016] and single-bounce microfacet model for rough dielectric BSDFs with varying roughness. Both of our methods produce results very similar to Heitz et al. [2016] and has almost identical costs. We also show the white furnace test results of the four methods, by rendering these materials lit by a constant white environment map. Both of our methods and Heitz et al. [2016] pass the white furnace test.

*Real-time rendering.* We implemented our method using shaders inside the Unreal Engine 4 (UE4), to show that it can be used for real-time rendering. We implemented both PT and BDPT versions, using a fixed sampling rate of 1 spp. We limit the number of bounces to two, as it already covers most of the energy. All timings in this section are measured on an NVIDIA RTX 2080 Ti graphics card. We exploit the modern rasterization pipeline, readily available in UE4, to generate high quality image sequences with the help of Deep Learning Super Sampling (DLSS).

Fig. 15 and the companion video show that both our PT and BDPT methods correctly produce the appearance from multiple-bounce BSDFs, and that our BDPT method is almost noise-free. The computational overhead of our method over single scattering is only about 0.10 ms (for PT) and 0.19 ms (for BDPT). Our method has potential applications in real-time rendering. We also provide a ShaderToy implementation of our method (showing both PT and BDPT) in the supplemental materials. We could not implement Heitz et al. [2016] in Unreal Engine 4 (UE4), due to its complexity. According to Figure 10 and Figure 5 in the supplemental materials, even after 5 ms, the results from Heitz et al. [2016] are still much noisier than our BDPT method.

## 6 DISCUSSION AND LIMITATIONS

*Height distribution function.* Our method is independent of the choice of a height distribution function, since it is canceled out during the computations on the  $\Lambda$  function. In this respect, it is identical to Heitz et al. [2016]. As shown in both the paper and supplemental materials of Heitz et al. [2016], both Gaussian and uniform height distribution functions provide the same results. We haven't found existing previous work that studied the impact of different height distributions on multiple scattering; this would be an interesting venue for future work.

In our experiments, BDPT always provides better results than PT, since there are much more paths constructed with BDPT. Moreover, unlike traditional BDPT which requires costly visibility query, our microscale BDPT doesn't require ray-scene intersection, making it much more efficient. However, for real-time applications, the BDPT might still be too complex, and may require extra optimization for an efficient GPU implementation.

*Limitations.* We have identified main limitations for our method: First, our height-correlated shadowing-masking function does not have reciprocity. Second, for performance reasons, our model currently focuses on NDFs with an analytical  $\Lambda$  functions, such as Beckmann and GGX. Third, although it has a low variance, our method still does not have a closed-form solution. Finding an analytical solution would be a strong extension to this work. Fourth, we propose both height uncorrelated and height-correlated shadowing-masking functions. Both of them are approximations to the accurate shadowing-masking functions. Using a height-direction correlated shadowing-masking function would improve accuracy, as mentioned in Heitz [2014], but it would require rewriting the path formulation. Furthermore, the lack of closed-form multi-bounce PDF makes sampling less efficient and leads to the similar variance as in Heitz et al. [2016] under complex lighting. Turqin's approximation [2019] can be used for approximating the PDF, however it



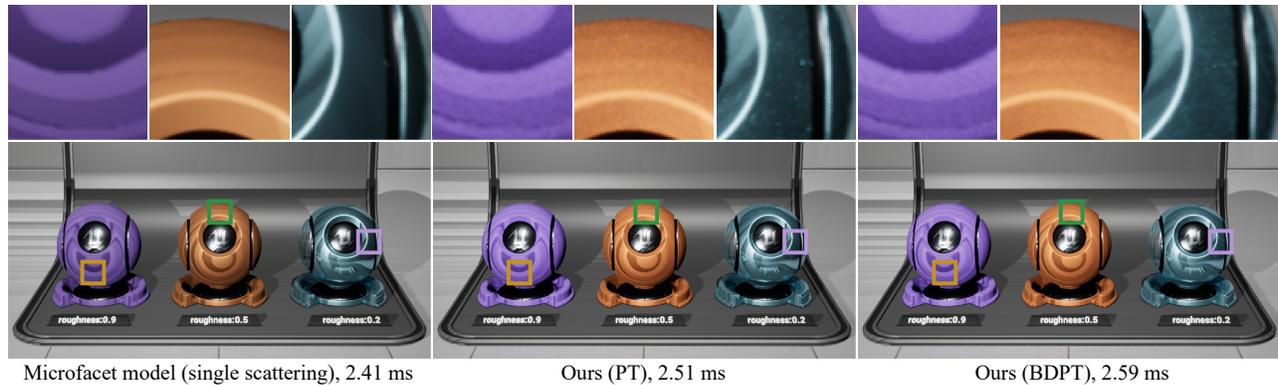


Fig. 15. Comparison between the microfacet model (single scattering only) and our methods (PT and BDPT, both 1 spp) implemented in shaders in UE4 on three material balls with different roughness (0.9, 0.5 and 0.2). Our methods (PT and BDPT) bring obvious color changes, especially when roughness is high. With DLSS that is readily available in UE4, even using only 1 spp, both our PT and BDPT methods achieve low noise level, and our BDPT is almost noise-free for all three materials. We use the height-uncorrelated model in the UE implementation.

## REFERENCES

- James Arvo and David Kirk. 1990. Particle Transport and Image Synthesis. *SIGGRAPH Comput. Graph.* 24, 4 (sep 1990), 63–66.
- Michael Ashikhmin and Simon Premože. 2007. Distribution-Based BRDFs. <https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.214.8243&rep=rep1&type=pdf>.
- Mahdi M. Bagher, Cyril Soler, and Nicolas Holzschuch. 2012. Accurate fitting of measured reflectances using a Shifted Gamma micro-facet distribution. *Computer Graphics Forum* 31, 4 (2012), 1509–1518.
- P. Beckmann and A. Spizzichino. 1963. *The scattering of electromagnetic waves from rough surfaces*. Pergamon Press.
- Robert L. Cook and Kenneth E. Torrance. 1982. A Reflectance Model for Computer Graphics. *ACM Trans. Graph.* 1, 1 (Jan. 1982), 7–24.
- Zhao Dong, Bruce Walter, Steve Marschner, and Donald P. Greenberg. 2016. Predicting Appearance from Measured Microgeometry of Metal Surfaces. *ACM Trans. Graph.* 35, 1, Article 9 (dec 2016), 13 pages.
- Jonathan Dupuy, Eric Heitz, and Eugene d'Eon. 2016. Additional Progress Towards the Unification of Microfacet and Microflake Theories. In *Eurographics Symposium on Rendering - EI & I*. The Eurographics Association, 55–63.
- Jonathan Dupuy, Eric Heitz, Jean-Claude Iehl, Pierre Poulin, and Victor Ostromoukhov. 2015. Extracting Microfacet-based BRDF Parameters from Arbitrary Materials with Power Iterations. *Computer Graphics Forum* 34, 4 (2015), 21–30.
- V. Falster, A. Jarabo, and J. R. Frisvad. 2020. Computing the Bidirectional Scattering of a Microstructure Using Scalar Diffraction Theory and Path Tracing. *Computer Graphics Forum* 39, 7 (2020), 231–242.
- Luis E. Gamboa, Adrien Gruson, and Derek Nowrouzezahrai. 2020. An Efficient Transport Estimator for Complex Layered Materials. *Computer Graphics Forum* 39, 2 (2020), 363–371.
- Yu Guo, Miloš Hašan, and Shuang Zhao. 2018. Position-Free Monte Carlo Simulation for Arbitrary Layered BSDFs. *ACM Trans. Graph.* 37, 6, Article 279 (Dec. 2018), 14 pages.
- Eric Heitz. 2014. Understanding the Masking-Shadowing Function in Microfacet-Based BRDFs. *Journal of Computer Graphics Techniques (JCGT)* 3, 2 (June 2014), 48–107.
- Eric Heitz and Eugene d'Eon. 2014. Importance Sampling Microfacet-Based BSDFs using the Distribution of Visible Normals. *Computer Graphics Forum* 33, 4 (2014), 103–112.
- Eric Heitz and Jonathan Dupuy. 2015. Implementing a Simple Anisotropic Rough Diffuse Material with Stochastic Evaluation. [https://drive.google.com/file/d/0BzvWldpUpRx\\_M3ZmakxHYXZWaUk/view](https://drive.google.com/file/d/0BzvWldpUpRx_M3ZmakxHYXZWaUk/view).
- Eric Heitz, Johannes Hanika, Eugene d'Eon, and Carsten Dachsbacher. 2016. Multiple-Scattering Microfacet BSDFs with the Smith Model. *ACM Trans. Graph.* 35, 4, Article 58 (July 2016), 14 pages.
- Wenzel Jakob. 2010. Mitsuba renderer. <http://www.mitsuba-renderer.org>.
- Wenzel Jakob, Eugene d'Eon, Otto Jakob, and Steve Marschner. 2014. A Comprehensive Framework for Rendering Layered Materials. *ACM Trans. Graph.* 33, 4, Article 118 (July 2014), 14 pages.
- Csaba Kelemen and László Szirmay-Kalos. 2001. A Microfacet Based Coupled Specular-Matte BRDF Model with Importance Sampling. In *Eurographics 2001 - Short Presentations*. The Eurographics Association.
- Christopher Kulla and Alejandro Conty. 2017. Physically Based Shading in Theory and Practice - Revisiting Physically Based Shading at Imageworks. <http://blog.selfshadow.com/publications/s2017-shading-course/>.
- Joo Lee, Adrián Jarabo, Daniel Jeon, Diego Gutiérrez, and Min Kim. 2018. Practical multiple scattering for rough surfaces. *ACM Trans. Graph.* 37, Article 175 (Dec. 2018), 12 pages.
- Daniel Meneveaux, Benjamin Bringier, Emmanuelle Tauzia, Mickaël Ribardière, and Lionel Simonot. 2018. Rendering Rough Opaque Materials with Interfaced Lambertian Microfacets. *IEEE Transactions on Visualization and Computer Graphics* 24, 3 (2018), 1368–1380.
- Nicolas Pinel, Bourlier Christophe, and J. Saillard. 2005. Energy conservation of the scattering from one-dimensional random rough surfaces in the high-frequency limit. *Optics letters* 30 (09 2005), 2007–9. <https://doi.org/10.1364/OL.30.002007>
- Mickaël Ribardière, Benjamin Bringier, Daniel Meneveaux, and Lionel Simonot. 2017. STD: Student's t-Distribution of Slopes for Microfacet Based BSDFs. *Computer Graphics Forum* 36, 2 (2017), 421–429.
- Mickaël Ribardière, Benjamin Bringier, Lionel Simonot, and Daniel Meneveaux. 2019. Microfacet BSDFs Generated from NDFs and Explicit Microgeometry. *ACM Trans. Graph.* 38, 5, Article 143 (June 2019), 15 pages.
- Vincent Ross, Denis Dion, and Guy 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), 2442–2453.
- Vincent Schüssler, Eric Heitz, Johannes Hanika, and Carsten Dachsbacher. 2017. Microfacet-Based Normal Mapping for Robust Monte Carlo Path Tracing. *ACM Trans. Graph.* 36, 6, Article 205 (2017), 12 pages.
- B. Smith. 1967. Geometrical shadowing of a random rough surface. *IEEE Transactions on Antennas and Propagation* 15, 5 (1967), 668–671.
- Jos Stam. 2001. An Illumination Model for a Skin Layer Bounded by Rough Surfaces. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*. Springer-Verlag, Berlin, Heidelberg, 39–52.
- Kenneth Torrance and E. Sparrow. 1967. Theory for Off-Specular Reflection From Roughened Surfaces. *Journal of The Optical Society of America* 57 (Sept. 1967).
- T. S. Trowbridge and K. P. Reitz. 1975. Average irregularity representation of a rough surface for ray reflection. *Journal of The Optical Society of America* 65, 5 (May 1975), 531–536.
- Emmanuel Turquin. 2019. Practical multiple scattering compensation for microfacet models. [https://blog.selfshadow.com/publications/turquin/ms\\_comp\\_final.pdf](https://blog.selfshadow.com/publications/turquin/ms_comp_final.pdf).
- Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. 2007. Microfacet Models for Refraction through Rough Surfaces. In *Rendering Techniques (proc. EGSR 2007)*. The Eurographics Association, 195–206.
- Stephen H. Westin, James R. Arvo, and Kenneth E. Torrance. 1992. Predicting Reflectance Functions from Complex Surfaces. In *Proceedings of SIGGRAPH '92 (SIGGRAPH '92)*. Association for Computing Machinery, 195–206.
- Mengqi (Mandy) Xia, Bruce Walter, Christophe Hery, and Steve Marschner. 2020. Gaussian Product Sampling for Rendering Layered Materials. *Computer Graphics Forum* 39, 1 (2020), 420–435.
- Feng Xie and Pat Hanrahan. 2018. Multiple Scattering from Distributions of Specular V-Grooves. *ACM Trans. Graph.* 37, 6, Article 276 (2018), 14 pages.
- Feng Xie, Anton Kaplanyan, Warren Hunt, and Pat Hanrahan. 2019. Multiple Scattering Using Machine Learning. In *ACM SIGGRAPH 2019 Talks*. Article 70, 2 pages.