Successive Height Preintegration for a Height-Interdependent Path Formulation of Multiple Bounces in Smith Microfacet BRDFs

SIYUAN ZHANG, Chiba University, Japan
TAKUYA FUNATOMI, Kyoto University, Japan
YUKI FUJIMURA, Nara Institute of Science and Technology, Japan
YASUHIRO MUKAIGAWA, Nara Institute of Science and Technology, Japan
HIROYUKI KUBO, Chiba University, Japan

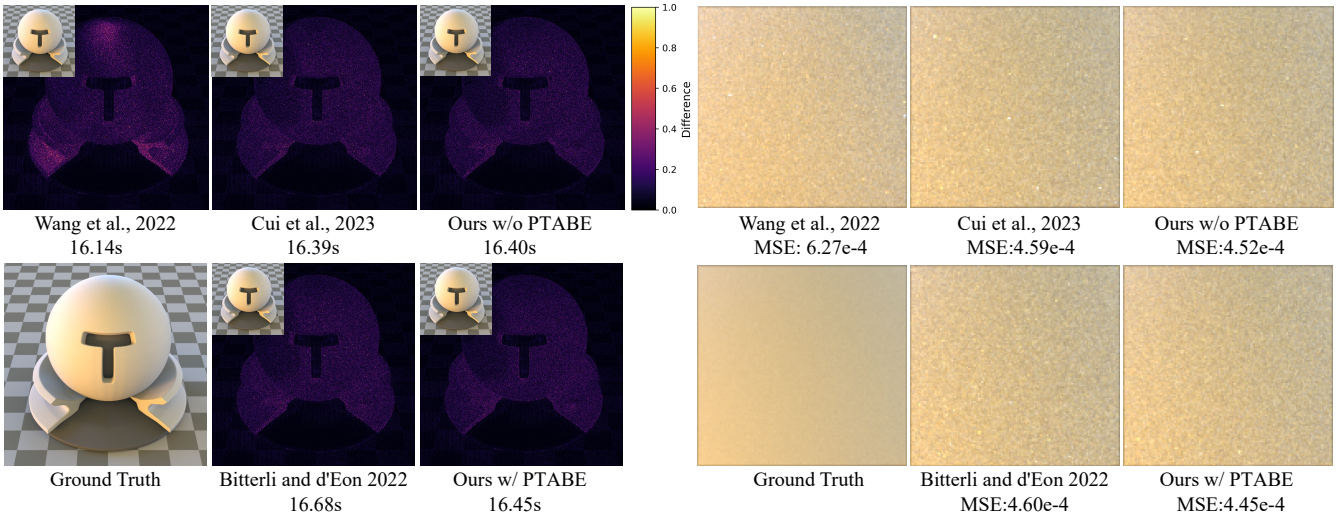

Fig. 1. A silver Matpreview scene with GGX roughness ($\alpha = 1.0$) rendered with 6 spp. We propose a novel multiple-bounce BRDF model along with a new BRDF evaluation method (PTABE), which reduces variance. Our model (without PTABE) produces results with lower variance compared to existing models [Bitterli and d'Eon 2022; Cui et al. 2023]. Furthermore, our BRDF evaluation method (PTABE) can further reduce the variance (Ours with PTABE).

Smith microfacet models are widely used to describe the interaction of light with surfaces. However, conventional microfacet models account for only a single bounce on microsurfaces, which leads to energy loss. To address the problem of missing energy, several unbiased methods have been proposed to compute multiple bounces on Smith microsurfaces. In addition, a position-free multiple-bounce microfacet model has been introduced by assuming that each bounce of a light path is height-independent, leading to lower variance but biased results. In this paper, we propose a novel derivation of multiple-bounce Smith microfacet bidirectional reflectance distribution functions (BRDFs) based on a partial adoption of the independent-bounce assumption. Our model treats most bounces as height-interdependent, leading to a height-interdependent path formulation. We further propose a successive height preintegration to make this formulation position-free. Although our model still exhibits an extremely small bias, it achieves lower variance than prior unbiased methods. Furthermore, we introduce a path type-aware BRDF evaluation that further reduces variance.

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

ACM Reference Format:
Siyuan Zhang, Takuya Funatomi, Yuki Fujimura, Yasuhiro Mukaigawa, and Hiroyuki Kubo. 2026. Successive Height Preintegration for a Height-Interdependent Path Formulation of Multiple Bounces in Smith Microfacet BRDFs. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Papers (SIGGRAPH Conference Papers '26), July 19–23, 2026, Los Angeles, CA, USA.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3799902.3811194

Authors' Contact Information: Siyuan Zhang, Chiba University, Chiba, Japan, 24rd1006@student.gs.chiba-u.jp; Takuya Funatomi, Kyoto University, Kyoto, Japan, funatomi.takuya.2c@kyoto-u.ac.jp; Yuki Fujimura, Nara Institute of Science and Technology, Nara, Japan, fujimura.yuki@is.naist.jp; Yasuhiro Mukaigawa, Nara Institute of Science and Technology, Nara, Japan, mukaigawa@is.naist.jp; Hiroyuki Kubo, Chiba University, Chiba, Japan, hkubo@chiba-u.jp.

This work is licensed under a Creative Commons Attribution 4.0 International License.
SIGGRAPH Conference Papers '26, Los Angeles, CA, USA
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2554-8/26/07
https://doi.org/10.1145/3799902.3811194

1 INTRODUCTION

Material models describe how light interacts with surfaces and are essential for reproducing realistic appearance. The microfacet model [Cook and Torrance 1982; Walter et al. 2007] represents surfaces as collections of specular microfacets whose normals follow a given distribution, with occlusion modeled by a shadowing-masking

term. However, classic microfacet models account only for single-bounce light on the microsurface, discarding multiple-bounce contributions and thus causing energy loss, especially for rough surfaces. To address this issue, multiple-bounce computation requires modeling microfacet blocking. We focus on the Smith model [Smith 1967], which assumes independent microfacet heights and normals and implies an equivalence between Smith microfacets and a volumetric medium. Many works exploit this equivalence to derive unbiased multiple-bounce models [Bitterli and d'Eon 2022; Cui et al. 2023; Dupuy et al. 2016; Heitz et al. 2016]. In contrast, Wang et al. [2022] proposed an approximate position-free formulation based on a height-averaged Smith shadowing function by treating light paths as sequences of height-independent bounces. While this approach achieves lower variance, the independent-bounce assumption introduces significant bias compared to the unbiased models [Bitterli and d'Eon 2022; Cui et al. 2023; Heitz et al. 2016].

In this paper, we propose a novel derivation for multiple-bounce computation of Smith microfacet BRDFs. Instead of directly using the independent-bounce assumption [Wang et al. 2022] to make the path formulation position-free, we first define a height-interdependent path formulation for light paths (Sec. 4.2), where "height" refers to the heights of microsurface intersections. In the definition of this height-interdependent path formulation, we partially adopt the independent-bounce assumption [Wang et al. 2022]. We demonstrate that this formulation can be successively preintegrated over the height dimension (Sec. 4.3), yielding a more accurate position-free path formulation. Compared with Wang et al. [2022], our model significantly reduces bias. Although an extremely small bias remains, the independent-bounce assumption leads to lower variance than prior unbiased methods [Bitterli and d'Eon 2022; Cui et al. 2023].

Furthermore, we propose a path type-aware BRDF evaluation method (Sec. 4.4). The multiple-bounce BRDF model is evaluated using Monte Carlo sampling over paths on the microsurface. Paths on the microsurface can be categorized into different types (Sec. 4.2). We empirically observe that allocating more samples to certain path types can lead to a noticeable reduction in estimator variance when Russian roulette [Arvo and Kirk 1990] is applied. Our method is purely empirical in nature, and experimental results demonstrate its effectiveness in reducing variance.

2 RELATED WORK

In this section, we briefly review previous works related to multiple-bounce microfacet models.

We first review multiple-bounce microfacet models based on the Smith model [Smith 1967]. After the equivalence between Smith microfacets and a volumetric medium was established, many multiple-scattering models exploiting this equivalence were proposed. Heitz et al. [2016] formulated Smith microfacet scattering as volumetric scattering by modeling the microsurface as an inhomogeneous microflake volume, leading to a random-walk solution for multiple-bounce computation. While their model closely matches simulated reference data, explicitly tracing height results in high variance. Dupuy et al. [2016] simplified this approach by assuming a semi-infinite homogeneous medium; however, as their method also relies on random walks, it similarly suffers from high variance. These approaches have been extended to normal-mapped surfaces [Schüssler et al. 2017] and combined with wave optics [Falster et al. 2020]. Wang et al. [2022] proposed a biased position-free formulation for multiple bounces using products of the Smith shadowing function, treating light paths as sequences of spatially independent bounces. Although this formulation achieves lower variance than that of Heitz et al. [2016], the independent-bounce assumption introduces significant bias. Following the assumption of Dupuy et al. [2016], Bitterli and d'Eon [2022] and Cui et al. [2023] independently proposed unbiased multiple-bounce microfacet models. Bitterli and d'Eon [2022] derived the first unbiased, low-variance multiple-bounce BRDF by analytically preintegrating collision distances, yielding an unbiased position-free formulation with reduced noise, albeit with potential numerical singularities. Cui et al. [2023] proposed a simpler unbiased derivation based on the invariance principle [Ambartsumian 1943], achieving the same variance while avoiding numerical singularities and reducing noise at equal time cost.

Beyond Smith microfacet models, multiple-bounce microfacet models have also been developed based on the V-groove model, yielding analytic solutions for multiple-bounce computation [Lee et al. 2018; Rosales et al. 2023; Xie and Hanrahan 2018]. While these models are typically more efficient, they are less physically accurate and incompatible with transparent materials. In addition, several approximate models have been proposed to account for multiple bounces on microgeometry, including adding an extra lobe to compensate for energy loss [Kulla and Conty 2017], scaling the single-bounce result [Turquin 2019], and using neural networks to represent multiple-bounce effects [Xie et al. 2019]. Although efficient, these approaches are generally less accurate, and we therefore do not discuss them further.

3 BACKGROUND

In this section, we review the Smith shadowing-masking term in the single-bounce Smith BRDF. Then, we introduce the multiple-bounce BRDF framework proposed by Wang et al. [2022].

3.1 The Smith shadowing-masking term

The Smith Masking Term. A random microsurface can be statistically modeled by its distributions of heights $P^1(h)$ and normals $D(\omega_m)$, where $h \in (-\infty, \infty)$ represents the displacement from the average plane of the surface at $h = 0$. $P^1(h)$ is a probability density function (PDF) and satisfies the normalization condition $\int_{-\infty}^{+\infty} P^1(h)\, dh = 1$. We denote the cumulative distribution function (CDF) of heights as $C^1(h)$. According to the Smith assumption, the visibility of a non-backfacing microfacet on the microsurface depends only on its height h and is independent of its normal ω_m. Thus, the visibility probability of a microfacet from direction ω_i is given by the product of the local masking term $G_1^{\text{local}}(\omega_i, \omega_m) = \chi^+(\omega_i \cdot \omega_m)$ and the distant masking term:

$$G_1^{\text{dist}}(\omega_i, h) = \left(C^1(h)\right)^{\Lambda(\omega_i)}, \tag{1}$$

which is the probability that the ray is not blocked by a distant microfacet. Heitz [2014] derived the above distant masking term based on projected area. $\Lambda(\omega_i)$ is the Smith Lambda function [Heitz 2014].

The visibility of a microfacet at a certain height on the microsurface is:

$$G_1(\omega_i, \omega_m, h) = G_1^{\text{local}}(\omega_i, \omega_m)\, G_1^{\text{dist}}(\omega_i, h). \tag{2}$$

Besides, the distant masking term can be derived by tracing a ray on the microsurface [Walter et al. 2007]. The method gives a probability that a ray, starting from h_r along direction ω_r, travels a distance τ without being occluded:

$$G_1^{\text{dist}}(\omega_r, h_r, \tau) = \left(\frac{C^1(h_r)}{C^1(h_r + \tau\cot\theta_r)}\right)^{\Lambda(\omega_r)}, \tag{3}$$

where ω_g is the macrosurface normal. θ_r denotes the angle between ω_r and ω_g. The distant masking term from Eq. (1) is derived by taking the distance τ in Eq. (3) to the limit at ∞:

$$G_1^{\text{dist}}(\omega_r, h_r) = \lim_{\tau\to\infty} G_1^{\text{dist}}(\omega_r, h_r, \tau). \tag{4}$$

In this paper, we re-parameterize $G_1^{\text{dist}}(\omega_r, h_r, \tau)$ as:

$$G_1^{\text{dist}}(\omega_r, h_r, h_{r+1}) = \left(\frac{C^1(h_r)}{C^1(h_{r+1})}\right)^{\Lambda(\omega_r)}, \tag{5}$$

which represents the probability that a ray, starting from h_r along direction ω_r, hits the microfacet located at height h_{r+1} without being blocked, where $h_{r+1} = h_r + \tau\cot\theta_r$.

The distant masking term (Eq. (1)) averaged over all heights is:

$$G_1^{\text{dist}}(\omega_i) = \int_{-\infty}^{+\infty} G_1^{\text{dist}}(\omega_i, h)\, P^1(h)\, dh = \frac{1}{1+\Lambda(\omega_i)}. \tag{6}$$

The height-averaged form of the masking term is:

$$G_1(\omega_i, \omega_m) = G_1^{\text{local}}(\omega_i, \omega_m)\, G_1^{\text{dist}}(\omega_i). \tag{7}$$

Shadowing-masking term. The shadowing-masking term accounts for the visibility of microfacets in both the incident and outgoing directions. It has a separable form using the product of height-averaged masking and shadowing terms [Walter et al. 2007]:

$$G_2(\omega_i, \omega_o, \omega_m) \;=\; G_1(\omega_i, \omega_m)\, G_1(\omega_o, \omega_m), \tag{8}$$

which neglects the correlation between masking and shadowing arising from the heights. Therefore, it is called the height-uncorrelated shadowing-masking term. However, the visibility of a microfacet in the incident and outgoing directions increases or decreases together as the height changes. The height-correlated shadowing-masking term considers the correlation of heights [Heitz 2014; Ross et al. 2005]:

$$G_2(\omega_i, \omega_o, \omega_m) = G_1^{\text{local}}(\omega_i, \omega_m)\, G_1^{\text{local}}(\omega_o, \omega_m)\, G_2^{\text{dist}}(\omega_i, \omega_o), \tag{9}$$

which is derived by integrating the shadowing and masking term over the heights:

$$G_2^{\text{dist}}(\omega_i, \omega_o) = \int_{-\infty}^{+\infty} G_2^{\text{dist}}(\omega_i, \omega_o, h)\, P^1(h)\, dh, \tag{10}$$

where $G_2^{\text{dist}}(\omega_i, \omega_o, h) = G_1^{\text{dist}}(\omega_i, h) G_1^{\text{dist}}(\omega_o, h)$. For the reflection case, it has the closed form:

$$G_2^{\text{dist}}(\omega_i, \omega_o) = \frac{1}{1+\Lambda(\omega_i)+\Lambda(\omega_o)}. \tag{11}$$

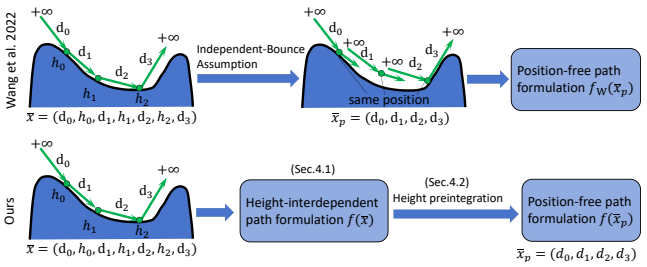

Fig. 2. Wang et al. [2022] assume that a light path consists of height-independent bounces and thus obtain a position-free light path $\bar{x}_p$, which in turn results in a position-free path formulation. In contrast, we first define a height-interdependent path formulation for a light path $\bar{x}$ and then preintegrate it over height to make it position-free.

The height-correlated shadowing-masking term has its integral form:

$$G_2(\omega_i, \omega_o, \omega_m) = \int_{-\infty}^{+\infty} G_2(\omega_i, \omega_o, \omega_m, h)\, P^1(h)\, dh, \tag{12}$$

where $G_2(\omega_i, \omega_o, \omega_m, h)$ is the height-correlated shadowing-masking term at a certain height:

$$G_2(\omega_i, \omega_o, \omega_m, h) = G_1^{\text{local}}(\omega_i, \omega_m)\, G_1^{\text{local}}(\omega_o, \omega_m)\, G_2^{\text{dist}}(\omega_i, \omega_o, h). \tag{13}$$

3.2 Multiple-bounce BRDF framework

Wang et al. [2022] assumed that all points on the microsurface can be considered statistically identical, resulting in the position-free light path $\bar{x}_p = (\mathbf{d}_0, \mathbf{d}_1, \ldots, \mathbf{d}_k)$ that remains in the angular domain only (Fig. 2). The first and last directions correspond to the macro incident and outgoing directions of a BRDF query, i.e., $\mathbf{d}_0 = -\omega_i$ and $\mathbf{d}_k = \omega_o$. Note that the direction of ω_i points upwards. The path space $\Omega(\omega_i, \omega_o)$ is the set of all possible paths whose first direction is $-\omega_i$ and whose last direction is ω_o. Their multiple-bounce BRDF $\rho(\omega_i, \omega_o)$ is defined as an integral over the set of paths $\Omega(\omega_i, \omega_o)$:

$$\rho(\omega_i, \omega_o) = \int_{\Omega(\omega_i,\omega_o)} f_W(\bar{x}_p)\, \mathrm{d}\mu(\bar{x}_p), \tag{14}$$

where $f_W(\bar{x}_p)$ is the path formulation of light transport and indicates the contribution of the path $\bar{x}_p$ with bounce k. The measure of the path space $\mu(\bar{x}_p) = \prod_{i=0}^{k-1} \sigma(d_i)$ is defined as the product of the solid angle measures σ.

Wang et al. [2022] treat a light path as a set of height-independent bounces (Fig. 2), and each bounce involves a vertex term V and a segment term S. This leads to the definition of the path formulation using the product of the vertex term V and the segment term S:

$$f_W(\bar{x}_p) = \prod_{i=0}^{k-1} V_i(\mathbf{d}_i, \mathbf{d}_{i+1}) S_i(\mathbf{d}_i, \mathbf{d}_{i+1}), \tag{15}$$

where the vertex term V_i allows for local interactions between the light and the microfacets, composing of the normal distribution D,

the Fresnel term F and the Jacobian term:

$$V_i(\mathbf{d}_i, \mathbf{d}_{i+1}) = \frac{F\left(-\mathbf{d}_i, \mathbf{d}_h^i\right) D\left(\mathbf{d}_h^i\right)}{4\left|\omega_g \cdot (-\mathbf{d}_i)\right|}, \tag{16}$$

where $\mathbf{d}_h^i = \frac{-\mathbf{d}_i + \mathbf{d}_{i+1}}{\|-\mathbf{d}_i + \mathbf{d}_{i+1}\|}$. ω_g is the macrosurface normal.

The role of the segment term S_i is similar to that of the shadowing-masking term. At each bounce, it encodes how much incoming light is received by the microsurface without being blocked, as well as how much of the reflected light is blocked.

4 OUR MODEL

In this section, we first give an overview of our model. Then, we provide a detailed definition of the height-interdependent path formulation and propose a successive height preintegration method to make it position-free.

4.1 Overview

Wang et al. [2022] defined the path formulation as the product of height-independent bounces (Eq. (15)), thereby ignoring the spatial correlation between bounces arising from the heights. This leads to highly biased results. Their definition is analogous to the height-uncorrelated shadowing-masking term (Eq. (8)), which neglects the correlation between masking and shadowing due to the heights.

In this paper, we define a height-interdependent path formulation for the light path $\bar{x} = (\mathbf{d}_0, h_0, \ldots, h_{k-1}, \mathbf{d}_k)$ (Fig. 2):

$$f(\bar{x}) = \left(\prod_{i=0}^{k-1} V_i(\mathbf{d}_i, \mathbf{d}_{i+1})\right) \underbrace{S_k(\mathbf{d}_0, h_0, \ldots, h_{k-1}, \mathbf{d}_k)}_{S_{\mathrm{p}_k}(\mathbf{d}_0, \ldots, \mathbf{d}_k)}, \tag{17}$$

where S_k is a height-interdependent segment term for an entire path that depends on both the spatial domain and the angular domain. This definition is analogous to the height-correlated shadowing-masking term at a certain height (Eq. (13)), which considers the correlation of heights. We will now show how S_k can be defined in detail and successively preintegrated over height to make it position-free. The resulting position-free segment term is denoted as S_{p_k}.

4.2 Height-interdependent path formulation

The vertex term V in the path formulation (Eq. (17)) is the same as Eq. (16). We focus on S_k in Eq. (17).

Since no downward-facing microfacets exist in a microsurface, once a light ray is reflected upward, it can no longer be reflected downward. Therefore, the light path $\bar{x} = (\mathbf{d}_0, h_0, \ldots, h_{k-1}, \mathbf{d}_k)$ with k bounces on the microsurface can be classified into three types depending on when the light is first reflected upward (Fig. 3):

(1) *Path type 1*: the light path begins to reflect upward from the k-th bounce ($\mathbf{d}_{k-1} \cdot \mathbf{d}_k < 0$),
(2) *Path type 2*: the light path begins to reflect upward from the first bounce ($\mathbf{d}_0 \cdot \mathbf{d}_1 < 0$),
(3) *Path type 3*: the light path begins to reflect upward from any bounce other than the first or the last ($\mathbf{d}_{k-1} \cdot \mathbf{d}_k > 0$ and $\mathbf{d}_0 \cdot \mathbf{d}_1 > 0$).

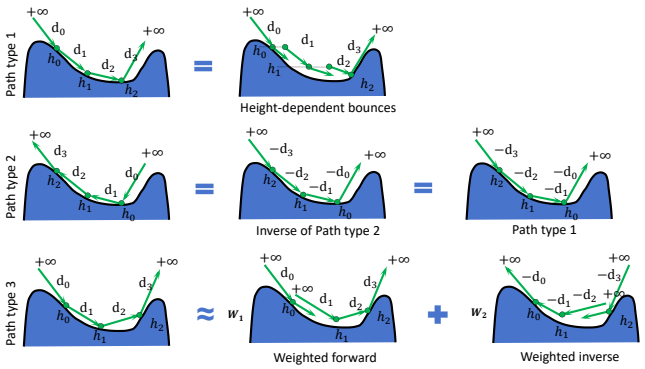

Fig. 3. Basic idea of defining segment terms for the three path types on microsurfaces. *Path type 1:* We view *path type 1* as a sequence of height-dependent bounces, leading to a height-cascaded product of segment terms (Eq. 18). *Path type 2:* BRDF reciprocity requires *path type 2* to be equivalent to its inverse. Since the inverse of *path type 2* falls into *path type 1*, we define *path type 2* based on the formulation of *path type 1*. *Path type 3:* We observe that a three-bounce *path type 3* can be decomposed into an independent bounce followed by a *path type 2* by assuming that its first bounce is height-independent of the subsequent bounces. To ensure reciprocity, we define *path type 3* as a blend of its forward and inverse results.

We define S_k separately for different types of light paths. This classification is motivated by the observation that *path type 2* can be directly derived from *path type 1* via BRDF reciprocity, while *path type 3* can be approximately derived from *path type 2* using the independent-bounce assumption [Wang et al. 2022]. Consequently, only S_k for *path type 1* needs to be explicitly defined, which constitutes the core idea of our S_k formulation.

Path type 1. A light path can be viewed as a sequence of height-dependent bounces (Fig. 3). Therefore, we define S_k in Eq. (17) for *path type 1* as a height-cascaded product of segment terms, where each segment term inherits its height from the preceding bounce:

$$S_k^{t_1}(\mathbf{d}_0, h_0, \ldots, h_{k-1}, \mathbf{d}_k) = \prod_{i=0}^{k-1} S_{\mathrm{b}_i}(h_{i-1}, \mathbf{d}_i, h_i, \mathbf{d}_{i+1}), \tag{18}$$

where $S_k^{t_1}$ represents the segment term for *path type 1*. S_b denotes the height-dependent segment term for a bounce. Each bounce is associated with a corresponding S_b. We formulate it into two different parts as follows:

$$S_{\mathrm{b}_i} = p_{\mathrm{b}_i}(h_{i-1}, \mathbf{d}_i, h_i, \mathbf{d}_{i+1}) \cdot o_{\mathrm{b}_i}(\mathbf{d}_i, h_i, \mathbf{d}_{i+1}), \tag{19}$$

where p_b denotes the incident probability that a ray, originating from the incident height h_{i-1} and traveling along the incident direction $\mathbf{d}_i$, reaches the microfacet located at the intersection height h_i without being occluded. o_b is defined as the probability that a ray reflected from the intersection height h_i along the outgoing direction $\mathbf{d}_{i+1}$ is occluded. A subsequent bounce occurs only if the reflected ray is blocked.

We mathematically define the incident probability p_b as:

$$p_{\mathrm{b}_i} = \begin{cases} G_1(-\mathbf{d}_i, \mathbf{d}_h^i, h_i), & \text{if } h_{i-1} = +\infty, \\ G_1^{\mathrm{local}}(-\mathbf{d}_i, \mathbf{d}_h) G_1^{\mathrm{dist}}(-\mathbf{d}_i, h_i, h_{i-1}), & \text{otherwise.} \end{cases} \tag{20}$$

If the incident height h_{i-1} of a bounce is $+\infty$, its incident probability can be computed using the height-dependent Smith masking term (Eq. (2)). This case occurs at the first bounce ($i = 0$) of *path type 1* (Fig. 2). For a bounce with a finite incident height h_{i-1}, the incident probability depends on both the incident height h_{i-1} and the intersection height h_i, which can be accounted for using Eq. (5) over a specific height interval. In addition, a local masking term is still required to eliminate backfacing microfacets. For *path type 1*, finite incident heights occur in all bounces except the first one.

The occlusion probability o_{b} is defined as:

$$o_{\mathrm{b}_i} = \begin{cases} 1, & \text{if } \mathbf{d}_{i+1} \cdot \omega_g < 0, \\ G_1(\mathbf{d}_{i+1}, \mathbf{d}_h^i, h_i), & \text{otherwise.} \end{cases} \tag{21}$$

If the reflected ray points downward, it is always occluded by the microsurface, and the occlusion probability equals 1. For *path type 1*, the only remaining case is the last bounce ($i = k-1$), whose outgoing direction $\mathbf{d}_{i+1}$ is upward. This last bounce is treated as a special case: since the path must terminate to maintain a total length of k, its occlusion probability is defined as the probability of *not* being occluded, which can be computed using the height-dependent Smith shadowing term (Eq. (2)).

Path type 2. The reciprocity of the BRDF requires that the segment term also exhibits reciprocity:

$$S_k^{t_2}(\mathbf{d}_0, h_0, \ldots, h_{k-1}, \mathbf{d}_k) = S_k^{t_{2\mathrm{inv}}}(-\mathbf{d}_k, h_{k-1}, \ldots, h_0, -\mathbf{d}_0), \tag{22}$$

where $S_k^{t_2}$ and $S_k^{t_{2\mathrm{inv}}}$ denote the segment terms of *path type 2* and its inverse, respectively. It is evident that the inverse of *path type 2* is *path type 1* (Fig. 3). Therefore, the segment terms of the two path types satisfy the following equivalence relation:

$$S_k^{t_{2\mathrm{inv}}}(-\mathbf{d}_k, h_{k-1}, \ldots, h_0, -\mathbf{d}_0) = S_k^{t_1}(-\mathbf{d}_k, h_{k-1}, \ldots, h_0, -\mathbf{d}_0). \tag{23}$$

This leads to an obvious but important observation: the segment term of *path type 2* can be directly defined using *path type 1*. This observation underlies our definition of $S_k^{t_2}$ for *path type 2*:

$$\begin{aligned} S_k^{t_2}(\mathbf{d}_0, h_0, \ldots, h_{k-1}, \mathbf{d}_k) &= S_k^{t_1}(-\mathbf{d}_k, h_{k-1}, \ldots, h_0, -\mathbf{d}_0) \\ &= \prod_{i=k-1}^{0} S_{\mathrm{b}_i}(h_{i+1}, -\mathbf{d}_{i+1}, h_i, -\mathbf{d}_i), \end{aligned} \tag{24}$$

where since the last bounce ($i = k-1$) of *path type 2* corresponds to the first bounce of *path type 1*, the product is taken in decreasing order of i, from $k-1$ down to 0. S_{b} is the height-dependent segment term for a bounce, which is given by Eq. (19).

Path type 3. We illustrate our idea using the simplest case of *path type 3* with three bounces (Fig. 3). By partially adopting the independent-bounce assumption [Wang et al. 2022], we treat the first bounce of *path type 3* as height-independent of the subsequent bounces, allowing the three-bounce *path type 3* to be viewed as a single bounce followed by a *path type 2*. To ensure reciprocity, we define it as a weighted blend of its forward and inverse results. This idea naturally extends to $S_k^{t_3}$ with k bounces:

$$S_k^{t_3}(\mathbf{d}_0, h_0, \ldots, h_{k-1}, \mathbf{d}_k) = w_1 S_k^{t_{3\mathrm{fwd}}} + w_2 S_k^{t_{3\mathrm{inv}}}, \tag{25}$$

where $w_1 = G_1(\mathbf{d}_k, \mathbf{d}_h^{k-1}, h_{k-1})$ and $w_2 = G_1(-\mathbf{d}_0, \mathbf{d}_h^0, h_0)$. Note that the choice of these two weights is motivated by reciprocity and does not form a convex combination. The definitions of $S_k^{t_{3\mathrm{fwd}}}$ and $S_k^{t_{3\mathrm{inv}}}$ are obtained by regarding the first bounce as height-independent of the subsequent bounces:

$$\begin{aligned} S_k^{t_{3\mathrm{fwd}}}(\mathbf{d}_0, h_0, \ldots, h_{k-1}, \mathbf{d}_k) &= S_{\mathrm{b}_0}(+\infty, \mathbf{d}_0, h_0, \mathbf{d}_1) \cdot \\ &\qquad S_{k-1}^{\mathrm{sub}_{\mathrm{fwd}}}(\mathbf{d}_1, h_1, \ldots, h_{k-1}, \mathbf{d}_k), \\ S_k^{t_{3\mathrm{inv}}}(-\mathbf{d}_k, h_{k-1}, \ldots, h_0, -\mathbf{d}_0) &= S_{\mathrm{b}_0}(+\infty, -\mathbf{d}_k, h_{k-1}, -\mathbf{d}_{k-1}) \cdot \\ &\qquad S_{k-1}^{\mathrm{sub}_{\mathrm{inv}}}(-\mathbf{d}_{k-1}, h_{k-2}, \ldots, h_0, -\mathbf{d}_0), \end{aligned} \tag{26}$$

where $S_{k-1}^{\mathrm{sub}_{\mathrm{fwd}}}$ and $S_{k-1}^{\mathrm{sub}_{\mathrm{inv}}}$ denote the segment terms of the subpath of *path type 3* and its inverse, respectively, each consisting of all bounces except the first one:

$$S_{k-1}^{\mathrm{sub}_{\mathrm{fwd}}} = \begin{cases} S_{k-1}^{t_2}(\mathbf{d}_1, h_1, \ldots, h_{k-1}, \mathbf{d}_k), & \text{if } \mathbf{d}_1 \cdot \mathbf{d}_2 < 0, \\ S_{k-1}^{t_3}(\mathbf{d}_1, h_1, \ldots, h_{k-1}, \mathbf{d}_k), & \text{otherwise,} \end{cases}$$

$$S_{k-1}^{\mathrm{sub}_{\mathrm{inv}}} = \begin{cases} S_{k-1}^{t_2}(-\mathbf{d}_{k-1}, h_{k-2}, \ldots, h_0, -\mathbf{d}_0), & \text{if } \mathbf{d}_{k-1} \cdot \mathbf{d}_{k-2} < 0, \\ S_{k-1}^{t_3}(-\mathbf{d}_{k-1}, h_{k-2}, \ldots, h_0, -\mathbf{d}_0), & \text{otherwise,} \end{cases} \tag{27}$$

where $S_{k-1}^{\mathrm{sub}_{\mathrm{fwd}}}$ has two cases. If the reflection direction points upward starting from $\mathbf{d}_2$, the subpath belongs to *path type 2*, and its segment term can be computed using Eq. (24). Otherwise, it belongs to *path type 3*, which is recursively defined by Eq. (25). Similarly, $S_{k-1}^{\mathrm{sub}_{\mathrm{inv}}}$ also has two cases.

Given the definitions of $S_k^{t_1}$ (Eq. (18)), $S_k^{t_2}$ (Eq. (24)), and $S_k^{t_3}$ (Eq. (25)), we can express the height-interdependent S_k. We provide the explicit formulation of S_k in the supplementary.

4.3 Successive height preintegration

With the height-interdependent segment term S_k defined, we now turn to successively preintegrating it over the height dimension, thereby making it position-free. Note that since *path type 2* is defined from *path type 1*, and *path type 3* is defined from *path type 2*, it suffices to discuss how to preintegrate $S_k^{t_1}$ (Eq. (18)) over height.

Our preintegration method is inspired by the height-dependent shadowing–masking term used in the single-bounce BRDF (Eq. (10)), which depends only on the intersection height and eliminates this dimension by averaging over it. The segment term $S_k^{t_1}$ (Eq. (18)) is a height-cascaded product of multiple bounces with two key properties: the first bounce depends only on the intersection height, while each subsequent bounce depends on both the incident and intersection heights, with the incident height inherited from the previous bounce. If a bounce can eliminate its intersection-height dependence in a manner analogous to Eq. (10), it becomes dependent only on its incident height. The previous bounce can then further eliminate its own intersection-height dependence, which includes the incident height of the current bounce. This observation leads to a successive height-reduction process applied in reverse order to $S_k^{t_1}$, which we refer to as *successive preintegration*. Based on this idea, we eliminate all height dimensions of $S_k^{t_1}$ by starting from the

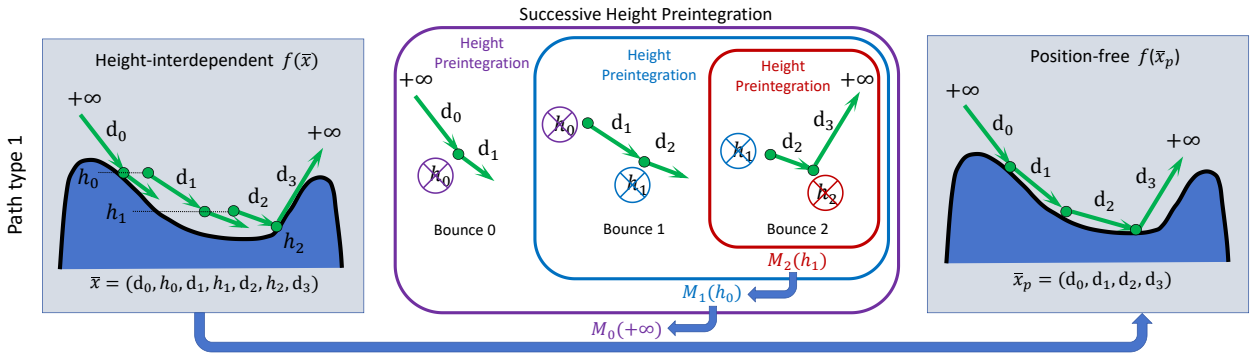

Fig. 4. Our successive preintegration method refers to a successive height-reduction process applied in reverse order. We illustrate this process using *path type 1* with three bounces as an example. As shown in Fig. 3, *path type 1* is viewed as a sequence of height-dependent bounces. Starting from the last bounce, we successively preintegrate each bounce over height to eliminate its intersection-height dimension, yielding intermediate results M (Eq. 28). Specifically, Bounce 2 is first preintegrated over h_2, resulting in $M_2(h_1)$, which depends only on h_1. Next, Bounce 1 and the preintegrated Bounce 2 are preintegrated over h_1, producing $M_1(h_0)$, which depends only on h_0. Finally, Bounce 0 together with the preintegrated Bounce 1 and Bounce 2 are preintegrated over h_0, yielding $M_0(+\infty)$, which is position-free.

last bounce and proceeding backward, as described below (Fig. 4):

$$M_i(h_{i-1}) = \begin{cases} \int_{-\infty}^{h_{i-1}} S_{\mathrm{b}_i}(h_{i-1}, h_i) P^1_{h_{i-1}}(h_i)\, dh_i, & \text{if } i = k-1, \\ \int_{-\infty}^{h_{i-1}} S_{\mathrm{b}_i}(h_{i-1}, h_i) M_{i+1}(h_i) P^1_{h_{i-1}}(h_i)\, dh_i, & \text{otherwise,} \end{cases} \tag{28}$$

where S_{b} (Eq. (19)) denotes the height-dependent segment term of a bounce in $S^{t_1}_k$ (Eq. (18)); its direction dependence is omitted for simplicity. The intersection height of the i-th bounce is denoted by $h_i \in (-\infty, h_{i-1})$, with the upper bound given by the incident height h_{i-1}. Preintegration starts from the last bounce ($i = k-1$). We denote the intermediate result obtained through successive preintegration as M, which depends only on the incident height of the current bounce. Since h_i is conditioned on h_{i-1}, we introduce the conditional probability density function $P^1_{h_{i-1}}(h_i)$:

$$P^1_{h_{i-1}}(h_i) = \frac{1}{C^1(h_{i-1})} P^1(h_i), \tag{29}$$

which satisfies the normalization condition $\int_{-\infty}^{h_{i-1}} P^1_{h_{i-1}}(h_i)\, dh_i = 1$. When $h_{i-1} = +\infty$, the conditional density reduces to the unconditional one: $P^1_{h_{i-1}}(h_i) = P^1(h_i)$.

The position-free segment term of *path type 1* can be derived as follows:

$$S^{t_1}_{\mathrm{p}_k}(\mathbf{d}_0, \ldots, \mathbf{d}_k) = M_0(+\infty), \tag{30}$$

where M_0 denotes the intermediate result obtained by successively preintegrating all bounces over their intersection heights in reverse order.

The position-free segment terms of *path type 2* and *path type 3* can also be derived following the same procedure as *path type 1*, thereby completing the derivation of the position-free segment term S_{p_k} (Eq. (17)). The details of S_{p_k} are provided in the supplementary.

4.4 Path type-aware BRDF evaluation

With the explicit position-free path formulation, any Monte Carlo method can be used to compute the integral (Eq. 14) for BRDF evaluation. In this section, we introduce an empirical BRDF evaluation method called path type-aware BRDF evaluation (PTABE) to further reduce variance.

Our approach is based on an empirical observation: allocating more samples to *path type 3* by converting *path type 1* into *path type 3* can reduce estimator variance. To this end, we introduce an additional resampling step into the general BRDF evaluation [Cui et al. 2023; Wang et al. 2022]. Specifically, when the current path belongs to *path type 1*, we resample it to encourage conversion into *path type 3*.

We provide the algorithm of our PTABE method in the supplementary. Note that PTABE alters the path distribution in a non-probabilistic manner, deviating from the probability density function (Eq. 16 in the supplementary) of the sample path. This modification introduces bias, causing the estimated results to deviate from their expected values. To mitigate this bias, we adopt two simple heuristics.

First, when the currently sampled path is of type 1, we resample the next direction with probability $1 - P_{\mathrm{d}}$. Here, P_{d} denotes the downward probability, which we approximate as:

$$P_{\mathrm{d}}(\omega_i) = T(\alpha_x, \alpha_y)\, G_1(\omega_i), \tag{31}$$

where $G_1(\omega_i)$ denotes the masking term (Eq. (7)), and $T(\alpha_x, \alpha_y)$ is a precomputed function that depends on surface roughness. We tabulate T over a range of roughness values (Fig. 12 in the supplementary) and evaluate it using bilinear interpolation at runtime.

This heuristic adaptively controls the resampling probability: a path of type 1 that has a lower probability under the original sampling distribution is more likely to be resampled, thereby increasing the occurrence of *path type 3*.

Second, due to the difference in contributions between *path type 1* and *path type 3* (Fig. 11 in the supplementary), the resampling step changes their relative proportions. This change can lead to

an underestimation of the multiple-bounce BRDF (Fig. 9 in the supplementary). To compensate for this effect, we introduce an empirical weighting term $1/(1 - P_{\mathrm{d}})$ for the resampled *path type 3*.

We provide detailed experimental validation of the above two heuristics in the supplementary.

Moreover, for BRDF evaluation, it is beneficial to early-terminate low-energy paths on the microsurface using Russian roulette (RR) [Bitterli and d'Eon 2022; Cui et al. 2023; Wang et al. 2022]. We validate in the supplementary that the variance reduction achieved by PTABE is always coupled with RR. Without applying RR, allocating more samples to *path type 3* has little effect on reducing the estimator variance. Furthermore, we present a discussion (Sec. 6), supported by experimental validation, explaining why PTABE can reduce estimator variance when RR is applied.

5 RESULTS

We have implemented our algorithm inside the Mitsuba renderer [Jakob 2010]. The implementations of prior work [Bitterli and d'Eon 2022; Cui et al. 2023; Heitz et al. 2016; Wang et al. 2022] are provided by the authors. For Wang et al. [2022], we use the height-correlated version. We also reimplemented Bitterli and d'Eon [2022]. We use the converged results of Heitz et al. [2016] as the ground truth images. Note that all models perform BRDF evaluation using a unidirectional estimator based on path tracing. All models are evaluated with Russian roulette enabled. For all results, the maximum number of bounces is set to 10. Heitz et al. [2016] serve as the ground truth.

We use MSE to measure the difference between each method and the ground truth. All timings in this section are measured on a 2.2 GHz Intel Core i9 (24 cores) with 64 GB of main memory.

In Fig. 5, we compare the mean BRDF reflectance, bias, variance, and inverse efficiency of our method with previous work [Cui et al. 2023; Heitz et al. 2016; Wang et al. 2022]. We compute the mean squared error (MSE) of the BRDF reflectance to quantitatively evaluate the bias of each model. Inverse efficiency is defined as the product of variance and time cost, where a lower value indicates higher performance. We report results for one roughness setting ($\alpha = 1.0$) and two incident directions. The quantitative bias evaluation shows that our model without PTABE exhibits only a very small bias, while achieving lower variance and higher performance than Cui et al. [2023] in all cases (across different incident and outgoing directions). Our model with PTABE further reduces variance significantly and improves performance.

Sphere scene. In Fig. 6, we show two sphere scenes with different materials: an isotropic copper (GGX, $\alpha = 1.0$) and an anisotropic aluminum (GGX, $\alpha_x = 1.0, \alpha_y = 0.7$), under direct lighting. Both scenes are illuminated by a directional light. We compare our methods with previous work [Bitterli and d'Eon 2022; Cui et al. 2023; Wang et al. 2022] at an equal sampling rate of 8 samples per pixel (spp). Although our models (with and without PTABE) introduce a small bias, they achieve lower variance and outperform the other models.

In Fig. 9, we show the MSE curves for the sphere scene. The ground truth is given by the converged results of Heitz et al. [2016]. Our methods nearly converge to the reference, indicating that the bias of our methods is extremely small and has a negligible impact on the rendered results. Moreover, due to higher performance, our models converge faster than the others.

Matpreview scene. We also provide a silver Matpreview scene (GGX, $\alpha = 1.0$) in Fig. 1. These scenes are illuminated by an environment map, which serves to evaluate the performance of our model under more complex lighting conditions. We compare our methods with previous work [Bitterli and d'Eon 2022; Cui et al. 2023; Wang et al. 2022] at an equal sampling rate of 6 spp. For each pixel sample, we employ 16 samples for emitter sampling and 6 samples for BRDF sampling. Under environment map illumination, our model, both with and without PTABE, achieves better results than the other models.

6 Discussion

In this section, we analyze the main properties of our model and discuss the differences from prior work.

Bias and energy conservation. Since our model partially adopts the independent-bounce assumption proposed by Wang et al. [2022], it introduces a slight bias. In addition, our PTABE method further introduces bias. We provide a quantitative evaluation of this bias in Fig. 5. The results show that, although the bias of our model still exists, it remains at a low level. Furthermore, we present the results of the white furnace test in Fig. 8, which indicate that our model conserves energy well.

Analysis of the variance reduction mechanism of PTABE.. It is known that RR generally increases the variance of the estimator. In Fig. 7, we compare the estimator variance at a maximum bounce of 10 for our model without RR, with RR applied to all path types, and with RR applied to each path type individually. We observe that applying RR to *path type 3* leads to a larger increase in estimator variance than applying RR to *path type 1* or *path type 2* across all cases. We provide additional experimental analysis in the supplementary to explain this behavior. These observations suggest that the variance reduction achieved by PTABE is due to allocating more samples to *path type 3*, thereby compensating for the larger variance increase introduced by RR. Please refer to the supplementary for further validation of PTABE.

Variance reduction. Our variance reduction comes from two factors. First, for *path type 3*, our segment term partially adopts the independent-bounce assumption to simplify its definition, thereby reducing the influence of the intersection height of the first bounce. This leads to a lower-variance BRDF model (Fig. 5). A similar variance reduction can also be observed in the model of Wang et al. [2022]. Since their entire formulation is based on the independent-bounce assumption, their model achieves even lower variance (Fig. 5). Second, our PTABE method further reduces the variance during BRDF evaluation.

Reciprocity. Our BRDF model is reciprocal, as both the vertex terms and the segment term satisfy reciprocity. We further prove the reciprocity of the segment term in the supplementary.

Limitations. Our multiple-bounce BRDF model partially adopts the independent-bounce assumption, which leads to biased estimates. Moreover, our current model, including the proposed PTABE,

considers reflection only. Preintegration for the height-interdependent path formulation of transmission remains challenging. In addition, PTABE does not strictly sample paths according to the underlying probability distribution; instead, it modifies the path distribution through resampling, which introduces bias into the estimator. Furthermore, PTABE is designed for unidirectional estimators based on path tracing. Extending it to more efficient bidirectional estimators is challenging, as the final path type cannot be determined solely from a single sampled direction in the bidirectional setting.

7 Conclusion

In this paper, we present a multiple-bounce BRDF model. Theoretically, we introduce a novel derivation of multiple-bounce BRDFs; practically, we propose a new definition of the segment term for a light path, which leads to a height-interdependent path formulation, and further introduce a successive height preintegration method to make it position-free. Although our model exhibits an extremely small bias due to the partial adoption of the independent-bounce assumption of Wang et al. [2022], it benefits from lower variance than prior unbiased methods [Bitterli and d'Eon 2022; Cui et al. 2023], resulting in higher performance. Furthermore, we propose a novel BRDF evaluation method, termed path type-aware BRDF evaluation, which enables additional variance reduction.

Acknowledgments

We would like to thank all anonymous reviewers for their constructive feedback. This work was partly supported by JST FOREST PMJFR206I, JST SPRING JPMJSP2109, JSPS Kakenhi JP24K02953, and JSPS JPJS00420230002.

References

V. A. Ambartsumian. 1943. The problem of diffuse reflection of light by a turbid medium. *Dokl. Akad. Nauk SSSR* 19 (1943), 30–41.

J. R. Arvo and D. S. Kirk. 1990. Particle transport and image synthesis. *ACM SIGGRAPH Computer Graphics* (1990).

Benedikt Bitterli and Eugene d'Eon. 2022. A Position-Free Path Integral for Homogeneous Slabs and Multiple Scattering on Smith Microfacets. *Computer Graphics Forum* 41, 4 (2022), 93–104. doi:10.1111/cgf.14589

Robert L. Cook and Kenneth E. Torrance. 1982. A Reflectance Model for Computer Graphics. *ACM Trans. Graph.* 1, 1 (Jan. 1982), 7–24.

Yuang Cui, Gaole Pan, Jian Yang, Lei Zhang, Ling-Qi Yan, and Beibei Wang. 2023. Multiple-bounce Smith Microfacet BRDFs using the Invariance Principle. In *Proceedings of SIGGRAPH Asia 2023.*

Jonathan Dupuy, Eric Heitz, and Eugene d'Eon. 2016. Additional Progress Towards the Unification of Microfacet and Microflake Theories. In *EGSR - Experimental Ideas & Implementations.* The Eurographics Association, 55–63.

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.

Eric Heitz. 2014. Understanding the Masking-Shadowing Function in Microfacet-Based BRDFs. *Journal of Computer Graphics Techniques (JCGT)* 3, 2 (30 June 2014), 48–107. http://jcgt.org/published/0003/02/03/

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.

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.

Enrique Rosales, Fatemeh Teimury, Joshua Horacsek, Aria Salari, Xuebin Qin, Adi Bar-Lev, Xiaoqiang Zhe, and Ligang Liu. 2023. Fast-MSχ: Fast Multiple Scattering Approximation. In *SIGGRAPH Asia 2023 Conference Papers.* Association for Computing Machinery, New York, NY, USA, 9 pages. doi:10.1145/3610548.3618231

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.

Emmanuel Turquin. 2019. Practical multiple scattering compensation for microfacet models. 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.

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 134 (jul 2022), 14 pages.

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* (Los Angeles, California). Article 70, 2 pages.

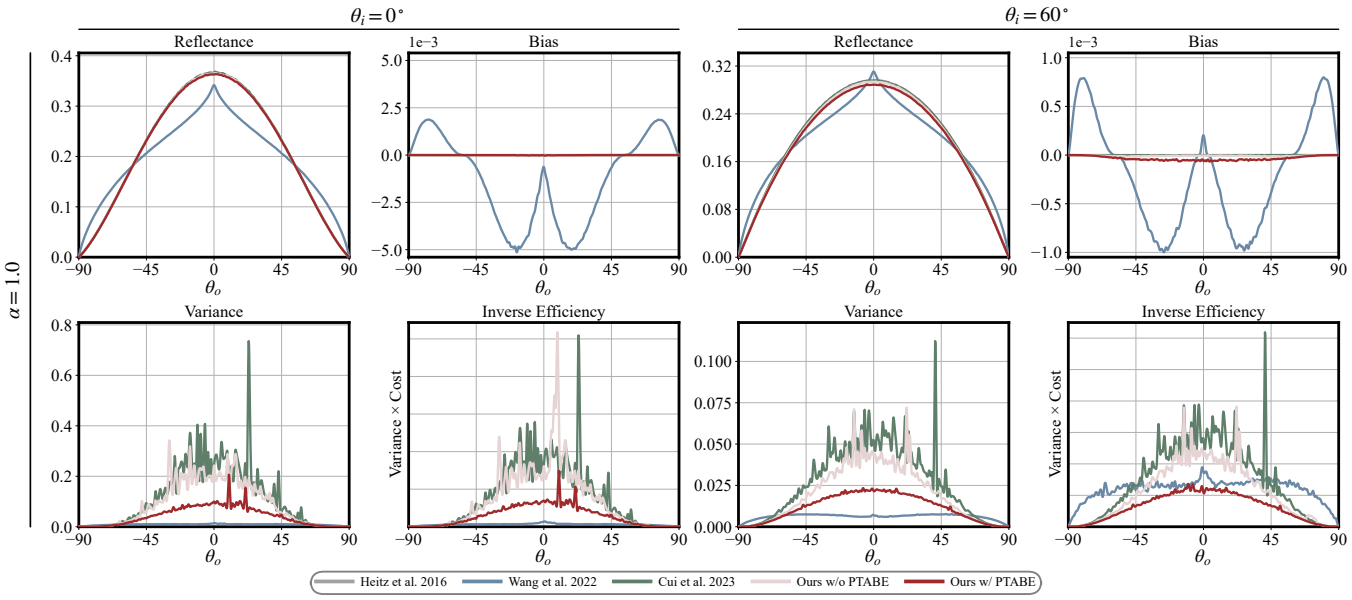

Fig. 5. Comparison between our models (with and without PTABE), Cui et al. [2023], Wang et al. [2022], and Heitz et al. [2016] in terms of BRDF reflectance, bias, variance, and inverse efficiency (variance × time cost) for a rough conductor with roughness $\alpha = 1.0$. Heitz et al. [2016] serve as the ground truth. We compute the mean squared error (MSE) of the BRDF reflectance to quantitatively evaluate the bias of each model. The bias evaluation shows that our models (with and without PTABE) exhibit only a very small bias. The variance results indicate that our model without PTABE achieves lower variance than Cui et al. [2023], while the model with PTABE further significantly reduces variance. Moreover, our models (with and without PTABE) achieve lower inverse efficiency, indicating higher overall performance. The model of Wang et al. [2022], which fully adopts the independent-bounce assumption, introduces large bias and therefore cannot match the ground truth. However, due to this assumption, it achieves the lowest variance among all methods.

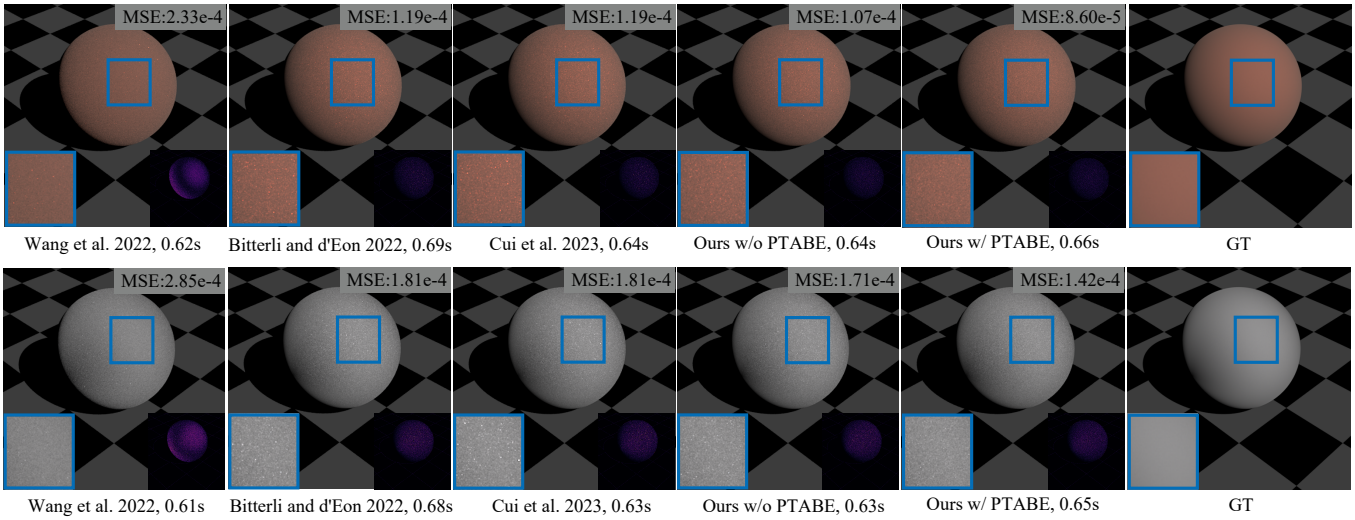

Fig. 6. Comparison between our models (without and with PTABE), Cui et al. [2023], Wang et al. [2022], Bitterli and d'Eon [2022], and Heitz et al. [2016] at an equal sampling rate of 8 spp on the two sphere scenes. The top sphere uses isotropic copper (GGX, $\alpha = 1.0$), while the bottom sphere uses anisotropic aluminum (GGX, $\alpha_x = 1.0, \alpha_y = 0.7$). Our model with PTABE produces the highest quality. Moreover, our model without PTABE also yields lower error than previous work. Cui et al. [2023] and Bitterli and d'Eon [2022] exhibit the same variance, resulting in identical MSE. The results of Wang et al. [2022] exhibit a low noise level; however, their bias results in a larger overall error.

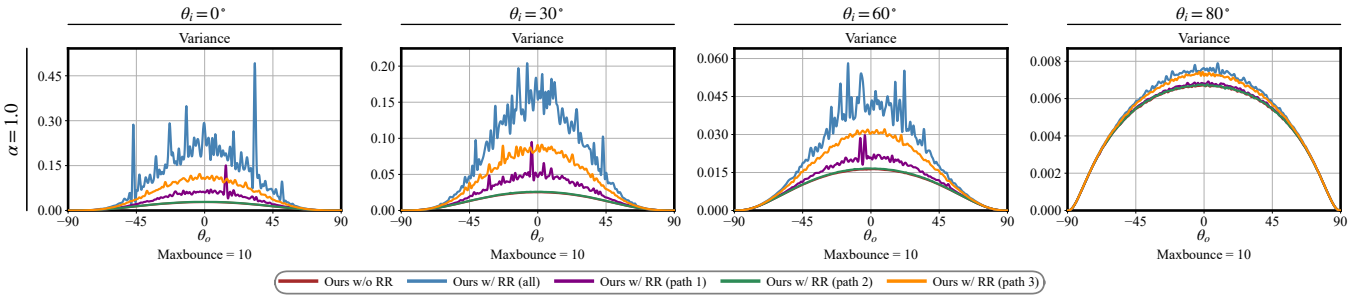

Fig. 7. Comparison of estimator variance for our models (without RR, with RR applied to all path types, and with RR applied to each path type individually) at a maximum bounce of 10. Note that all results are generated using our model without PTABE. The results lead to two observations. First, RR indeed introduces additional variance. Second, applying RR to *path type 3* results in a larger increase in variance than applying RR to *path type 1* or *path type 2* across all cases. These observations suggest that allocating more samples to *path type 3* can compensate for the additional variance introduced by RR, which explains why PTABE is able to reduce variance. We provide further experimental analysis in the supplementary to explain why RR leads to a larger increase in variance for *path type 3*.

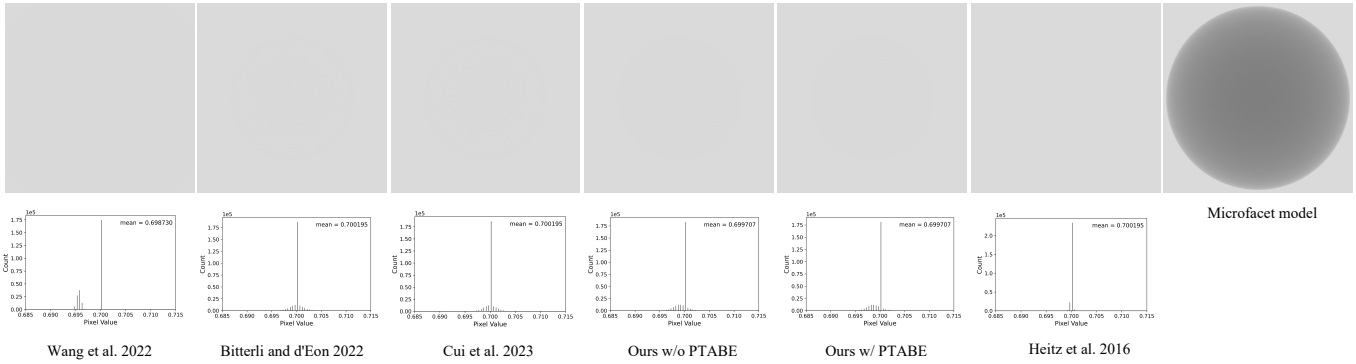

Fig. 8. We show the white furnace test results of BRDF models with roughness 1.0 (top), along with statistics of the pixel values from the rendered output images (512×512) (bottom). The rightmost column shows the result of the standard microfacet BRDF model with roughness 1.0, which serves as a reference. Note that we set the radiance of the constant illumination to 0.7, so the expected result of the white furnace test is 0.7. The statistics indicate that Bitterli and d'Eon [2022]; Cui et al. [2023]; Heitz et al. [2016] preserve energy well. The results of Wang et al. [2022] exhibit a small amount of energy loss, although this loss is difficult to perceive visually from the white furnace test images. Our models (with and without PTABE) also show a very small amount of energy loss, but perform better than Wang et al. [2022] in terms of energy conservation.

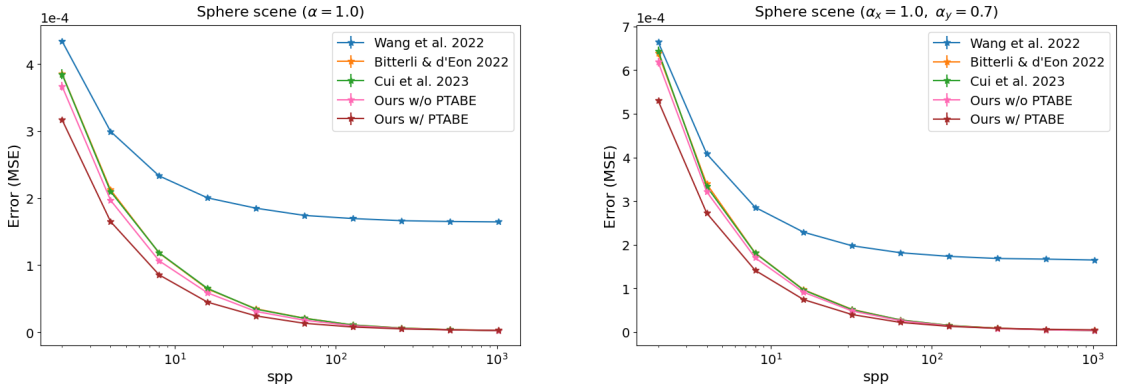

Fig. 9. We show the error (MSE) of our models (with and without PTABE), Wang et al. [2022], Bitterli and d'Eon [2022], and Cui et al. [2023] over varying sampling rates on the sphere scene (Fig. 6). The converged result of Heitz et al. [2016] is treated as the ground truth. Our models can converge to the ground truth, and both converge faster than the other models. Among them, our model with PTABE exhibits the fastest convergence.