Fast-MSX: Fast Multiple Scattering Approximation

Enrique Rosales*
Huawei Technologies
Burnaby, Canada
enrique.rosales@huawei.com

Fatemeh Teimury*
Huawei Technologies
Burnaby, Canada
fatemeh.teimury@huawei.com

Joshua Horacsek
Huawei Technologies
Burnaby, Canada
j.horacsek@huawei.com

Aria Salari
Huawei Technologies
Burnaby, Canada
aria.salari@huawei.com

Xuebin Qin
Huawei Technologies
Burnaby, Canada
xuebin.qin@huawei.com

Adi Bar-Lev
Huawei Technologies
Burnaby, Canada
adi.bar.lev@huawei.com

Xiaoqiang Zhe
Huawei Technologies
Shanghai, China
zhexiaoqiang@huawei.com

Ligang Liu
Huawei Technologies
Shanghai, China
liuligang4@huawei.com

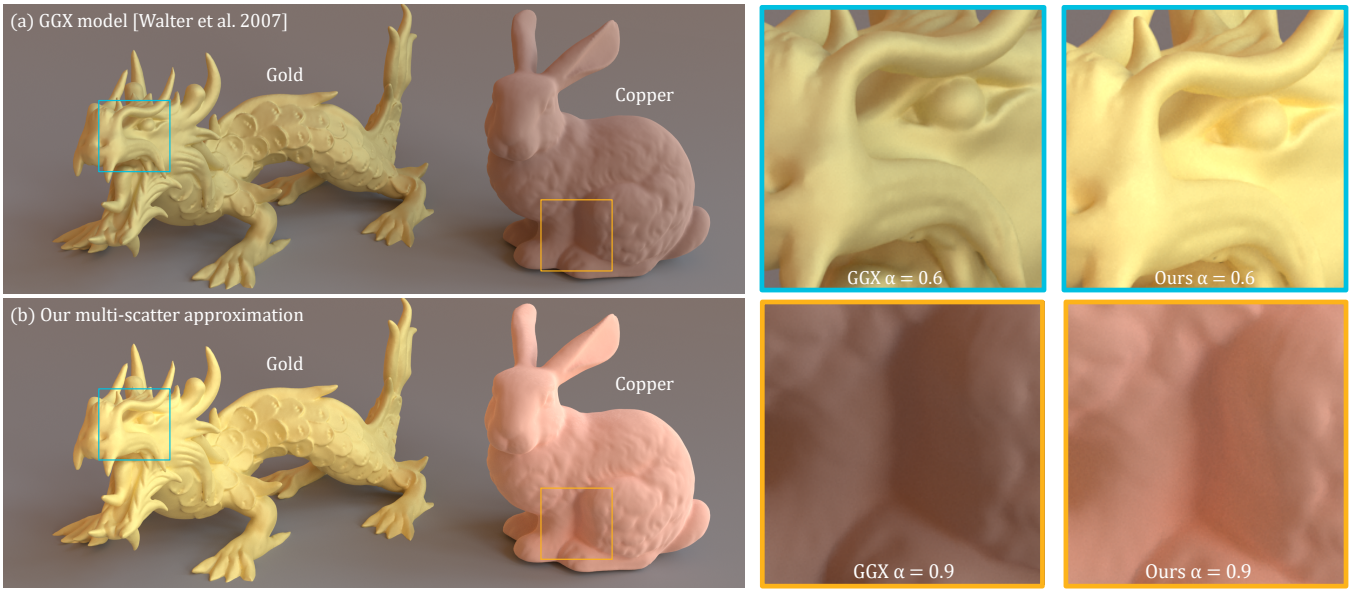

Figure 1: We present a fast multiple scattering extension to the classical GGX model [Walter et al. 2007]. Traditionally, the GGX model only assumes single bounce of light transport; this assumption leads to a loss of outgoing radiant flux (and therefore a darkening of rendered images). Our method approximates the energy from second bounce light transport, which produces more saturated colors and visually pleasing results (b) compared with the darker response of the classical model (a); α is the roughness. Images rendered with Mitsuba 0.6 [Jakob 2010]. Dragon and Bunny: © Stanford Computer Graphics Laboratory

ABSTRACT

Classical microfacet theory suffers from energy loss on materials with high roughness due to the single bounce assumption of most microfacet models. When roughness is high, there is a large chance of multiple scattering occurring among the microfacets of the surface. Without explicitly modelling for this behaviour, rough surfaces appear darker than they should. To address this issue, we present a novel method to estimate the multiple scattering contribution from a second light bounce. Our method is inspired by Zipin's geometric construction approach, which simplifies the calculation of the light transport inside a V-groove cavity. Our experimental

*Joint first authors

SA Conference Papers '23, December 12–15, 2023, Sydney, NSW, Australia

ACM ISBN 979-8-4007-0315-7/23/12…$15.00
https://doi.org/10.1145/3610548.3618231

results demonstrate that our method is visually pleasing, physically plausible, and artifact-free compared to recent multiple scattering works. Additionally, the low computational cost makes our model suitable for real-time rendering.

CCS CONCEPTS

• **Computing methodologies → Reflectance modeling**.

KEYWORDS

BRDF, BSDF, multiple scattering, microfacet theory, microfacet distributions, appearance representation, geometry

ACM Reference Format:

Enrique Rosales, Fatemeh Teimury, Joshua Horacsek, Aria Salari, Xuebin Qin, Adi Bar-Lev, Xiaoqiang Zhe, and Ligang Liu. 2023. Fast-MSX: Fast Multiple Scattering Approximation. In *SIGGRAPH Asia 2023 Conference Papers (SA Conference Papers '23), December 12–15, 2023, Sydney, NSW, Australia.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3610548.3618231

1 INTRODUCTION

Material representation is a fundamental concept in computer graphics and is a critical component of all photo-realistic offline and real-time renderers. In order to achieve realistic renders, contemporary material models are typically physically based [Walter et al. 2007]. These models commonly rely on microfacet theory as a basis for their underlying mathematical formulation. At its core, microfacet theory [Cook and Torrance 1982] assumes that the microscopic surface of a rough material is irregular (Fig. 2) and it is modeled as composed of small planar mirrors, each acting as a perfect specular reflector. Modelling the interaction of light at a microscopic level is challenging, thus, most models resort to some form of simplification of the microfacet geometry. These simplifications lead to computationally tractable shading models.

One key component of the microfacet model is the geometric "masking-shadowing" term, which accounts for localized occlusion among microfacets. This model is a bidirectional function composed of two terms. Loosely speaking, the first term is the probability that light travelling towards a microfacet is *not* blocked by some other microfacet before it hits the surface. The second term represents the probability of the reflected rays *not* being blocked by another microfacet after reflection. This naturally excludes the case in which light travels between multiple microfacets before reaching the eye (multiple scattering), which can have a profound impact on the appearance of the rendered material [Heitz 2014].

In this work, we develop an empirically tuned shading model that compensates for the missing energy in the Cook-Torrance [1982] single bounce approximation. We start by relaxing a V-cavity model [Torrance and Sparrow 1967; Zipin 1966], and extending the cavities to allow for more light transfer. Our main contribution is a novel geometric model for two bounce scattering that is

(1) easy to implement; our approximation is a simple additive term that relies on simple operations to compute
(2) real-time capable; most other multiple scattering methods require some form of iteration (i.e. Monte Carlo simulation)
(3) noise free and artifact free; by virtue of simplicity, since we rely on no stochastic numerical processes, we avoid unpleasant artifacts (i.e. fireflies and noise).

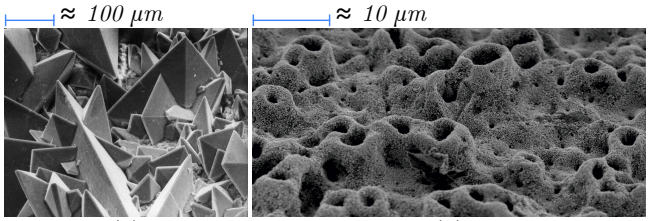

Figure 2: Scanning electron micrographs of tetragonal crystals of Weddellite (a) [EK 2012], and an orthopaedic titanium screw with "volcano" pores coated with zinc oxide nanoparticles (b) [Woźniak 2014] reveals the complex geometry of real materials with high roughness.

We also provide a qualitative and quantitative assessment of our methodology.

2 BACKGROUND

The Cook and Torrance [1982] model describes the interaction between light and microfacet geometry with:

$$f(v,l,\alpha) = \frac{D(h,\alpha)G(v,l,\alpha)F(v,h,f_0)}{4(n\cdot v)(n\cdot l)}, \tag{1}$$

where n, v and l represent the *normal*, *view* and *light* direction vectors; and D, G and F represent the normal distribution, the geometric attenuation and the Fresnel reflectance terms, respectively. The D term is a micronormal distribution function that models the concentration of microfacets whose normal is aligned with the half vector $h = l + v/||l + v||$ (Fig. 3a). A widely used version of the D term is the Trowbridge-Reitz model [1975] introduced by Blinn [1977], and defined in its modern form called GGX by Walter et al. [2007].

The F term models the amount of light that is reflected (instead of absorbed) from the microfacets.

The G term models the concentration of microfacets that are visible from both v and l directions. As seen in Fig. 3a, only microfacet m_4 is lit and visible from the camera thus, the only one able to transmit energy. There are two popular versions of the G term, the *separable* also called *uncorrelated* Smith in which each direction is computed separately and then multiplied together. And the *height-correlated* Smith that was demonstrated by Heitz [2014] to produce better results. In both cases, the major limitation of the geometry term G is that ignores cases in which light may bounce between multiple facets of the surface before ultimately leaving the surface (i.e. multiple scattering). As seen in Fig. 3b, a microfacet m_3 that is visible from the camera can indirectly observe the light reflected from another microfacet m_1, even if m_3 is directly occluded from the light. The G formulation does not account for this occlusion, as shadowed microfacets are directly discarded. Thus, materials with higher roughness may appear darker than expected due to the energy lost from this phenomenon [Heitz 2014].

3 RELATED WORK

Several attempts have been made to introduce an extra energy correction term. For instance, Burley [2012] adds a simple sheen effect

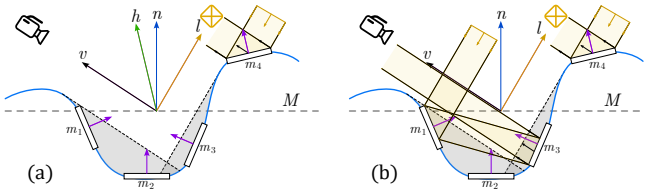

Figure 3: Microfacets m_1 and m_2 are lit but not visible from the camera; Microfacet m_3 is visible from the camera but occluded from the light (shadowed). In a single-scattering approach (a), only the microfacet m_4 is aligned with the half vector, lit and visible from the camera. Accounting for a second bounce multiple scattering (b) allows to get the light reflected from microfacet m_1 to microfacet m_3 and then to the view vector v.

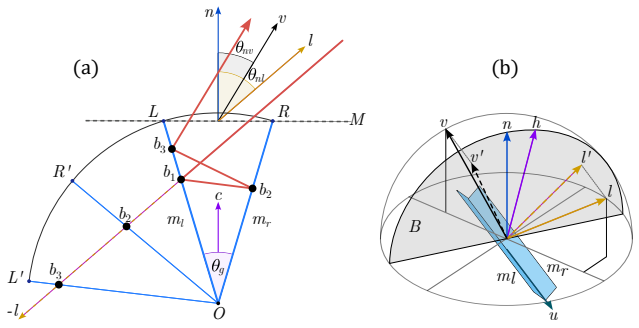

Figure 4: (a) Zipin's [1966] geometric construction works by creating a series of reflected virtual V-grooves inside a unitary circle. By extending the light vector l in $-l$ direction, the places in which that ray crosses the V-grooves' facets (points b_1, b_2, and b_3) correspond to the light bounces inside the original V-groove (OLR). The light transport path is then represented as the thick red line that goes from l to v in three bounces. To compute this method, we need the angles from l and v to the normal n (θ_{nl}, θ_{nv}) and the angle of the V-groove opening θ_g. (b) In 3D, Lee et al. [2018] and Xie and Hanrahan [2018] project the vectors l and v into the plane B (grey semicircle). Which is the bisector plane for l and v formed by the normal n and the half-vector h.

based on a Fresnel Schlick approximation [Schlick 1994], but the final result is not physically plausible for all roughness values. Kulla and Conty [2017] inspired by Kelemen and Szirmay-Kalos [2001] and Jakob et al. [2014] proposed a diffuse term to model the missing energy. Similarly, Hill et al. [2020] compensates the energy loss by introducing a Lambertian term that models multiple scattering within the microsurface. Later, Turquin [2018] based on the work of Kulla et al. [2017] scales the single scattering lobe to approximate multiple scattering. This work can properly compensate the single bounce energy loss and color saturation which in some cases could be an optional feature for artistic purposes. This approximation is suitable for real-time rendering, but is not reciprocal, a limitation that our proposed method also has.

Alternative lines of works generalize microfacet models by incorporating multiple scattering [Bitterli and d'Eon 2022; Dupuy et al. 2016; Heitz et al. 2016; Lee et al. 2018; Wang et al. 2022; Xie and Hanrahan 2018]. Heitz et al. [2016] offers a volumetric solution based on the Smith self-shadowing model, that accurately models multiple scattering and color saturation by computing an approximation of multiple-scattering with stochastic path tracing. Although Heitz et al. [2016] is reciprocal, energy-conserving, and supports the anisotropic GGX and Beckmann distributions, it is computationally expensive and suffers from noise due to its stochastic evaluation. Thus, not suitable for real-time rendering. Dupuy et al. [2016], further explores the connection between microflake theory and the standard microfacet model. Later, Wang et al. [2022] proposed an approximate position-free approach which eliminates the need to compute the position of each surface sample, allowing for faster computation of multiple bounces. The method is based on Heitz et al. [2016] and Dupuy et al. [2016], but produces almost noise-free results with low sample rates. Concurrent to Wang et al. [2022], Bitterli and d'Eon [2022] propose a position-free path integral for evaluating the scatter distribution in a homogeneous slab. Their solution preintegrates collision distances to simplify the Monte Carlo simulation allowing less variance than Heitz et al. [2016], Dupuy et al. [2016] and Wang et al. [2022].

On the other hand, Lee et al. [2018] and Xie and Hanrahan [2018] illustrate an efficient closed-form solution for multiple scattering based on the Cook-Torrance model [1982] and Zipin's [1966] geometric construction (Fig. 4). The models of Lee et al. [2018] and Xie and Hanrahan [2018] are reciprocal, energy conserving, and support anisotropy. However, they have a singularity in the direction of mirror reflection, which leads to unrealistic artifacts for high roughness materials (Fig. 5). Our work is inspired in Lee et al. [2018] and Xie and Hanrahan [2018]. Beyond close-form solutions, learning-based methods for approximating multiple scattering [Xie et al. 2019] have recently gained momentum and are a promising direction for future research.

4 OUR METHOD

Our goal is to approximate the missing multiple scattering energy and add it to a standard single-scattering Cook-Torrance specular lobe. For the single-scattering lobe we use: GGX distribution [Walter et al. 2007] for the D term, uncorrelated Smith for the G term [Heitz 2014], and Schlick [1994] approximation for the F term. We do this by first modelling the two bounce light transport for V-cavities, we then relax our model to introduce additional energy, which ultimately compensates for additional missing bounces within a localized region on a surface $\mathcal{M}$. Fig. 13 shows that, in the range from $\alpha = 0.3$ to $\alpha = 0.7$ [1], most of the transmitted energy is carried by the first and second bounces. Thus, we focus on the second bounce, adding a small amount of additional energy to compensate for higher order bounces. Empirically, our second bounce accounts for most of the missing energy, avoiding additional calculations for further bounces, and drastically simplifying our implementation. To model second bounce energy, we appeal to V-cavity theory, but discard two assumptions.

[1]The most realistic roughness values for most materials [Jakob 2010]

First, we do not assume that all V-cavities are aligned with the surface $\mathcal{M}$. Second, we do not assume that V-cavities are symmetric. Discarding these two assumptions leads to a more realistic model of the behaviour of the microfacets on the surface $\mathcal{M}$). Other works that make these two assumptions [Lee et al. 2018; Xie and Hanrahan 2018], have the ability to model an arbitrary number of bounces within a V-cavity, but have a higher specular response when the surface roughness α is high. This is inherently the incorrect behaviour, since a surface with high roughness should scatter light over a wider angle, this is demonstrated in Fig. 5.

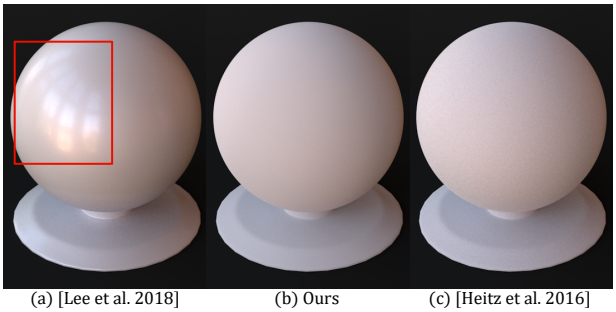

(a) [Lee et al. 2018] (b) Ours (c) [Heitz et al. 2016]

Figure 5: (a) Lee et al. [2018] method, shows a high portion of the microfacets reflecting radiance towards the mirror reflection in high roughness values ($\alpha = 0.9$). In contrast, our method (b) correctly approximates the light scattering, and it is visually similar to Heitz et al. [2016] (c).

We seek an indirect multiple scattering specular correction term that is artifact-free, physically plausible, suitable for real-time rendering, and easy to implement (Fig. 1). We proceed by describing how our assumptions prescribe the behaviour of the surface at the microfacet level. We then use the intuition gained to model new masking and distribution terms that, when combined with an appropriate Fresnel term, produce an estimate of the missing energy.

4.1 Computing the V-groove Coordinate Frame

We start our discussion by finding the V-cavities that are able to transmit light from l to v in exactly two bounces. With our current assumptions there are infinitely many V-cavities with this property. Rather than using Zipin's[1966] construction (Fig. 4) to construct second bounce paths within infinitely many V-grooves, we strictly assume the second bounce direction b is aligned with the horizontal axis (i.e. $\langle b \cdot n \rangle = 0$, see Fig. 6). This assumption drastically reduces the complexity of our formulation; with this assumption there are exactly two such V-cavities that transmit light from l to v (one that has an acute V-grove angle and another that is obtuse). If a V-groove has an obtuse cavity angle θ_g, then we discard this case, since it is likely covered by the original Cook-Torrance term.

Without loss of generality, we initially choose a coordinate frame defined by n and s (with $s := ((l \times v) \times n)$) as in Fig. 6. If l intersects with m_l and reflects into ray b, then m_l's normal direction is uniquely constrained by $\theta_l + 2\theta_b = \pi$, where $\cos(2\theta_b) = l \cdot b$ is the inner reflection angle, it follows from geometric construction that $\theta_{l_r} = \theta_b$. A similar derivation gives θ_{v_r} for m_r

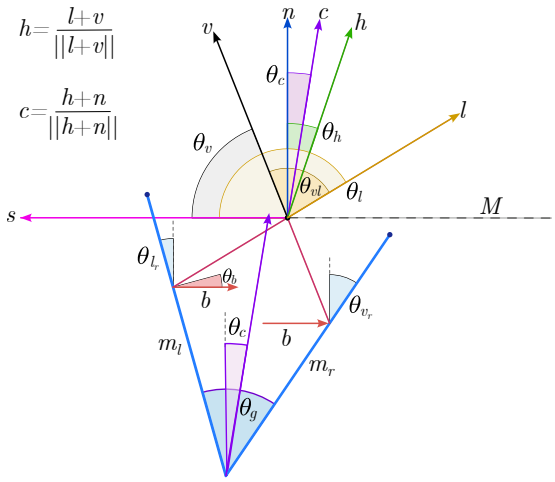

Figure 6: The first step of our method consists of defining a microfacet m_l that is able to reflect energy from the l vector to a horizontal bouncing path b, and a microfacet m_r that is able to reflect energy from the same horizontal bouncing path b to the vector v. These microfacets are defined by their angles with respect to the normal n (θ_{l_r} and θ_{v_r}). The angle of the V-groove opening θ_g is then the sum of these angles. The bisecting axis of our V-groove cavity (denoted by c), defines the coordinate frame in which we compute our variant of the Zipin's method.

$$\theta_{l_r} = \frac{\pi - \theta_l}{2}, \ \theta_{v_r} = \frac{\theta_v}{2}. \tag{2}$$

From this, it naturally follows that

$$\theta_g = \theta_{l_r} + \theta_{v_r} \tag{3}$$

simplfying this futher yeilds

$$\theta_g = \frac{1}{2}(\pi + \theta_v - \theta_l). \tag{4}$$

$$\theta_g = \frac{1}{2}(\pi - \theta_{vl}). \tag{5}$$

Where θ_{vl} is the angle between v and l so we no longer need the vector s. An interesting note here is that this identity holds regardless of the orientation of the cavity. That is to say, in general, the only acute V-grooves that reflect light from l to v are those with cavity angles given by Eq. 4.

Fig. 6 shows the bounce vector b from two different points on the surface. Fig. 6 also illustrates that the position of the V-groove in the 2D plane is irrelevant; at the microfacet scale, light transport theory assumes l and v are constant vector fields. Thus, only vector orientation is relevant.

We assume that the second bounce light transport is horizontal, this necessarily constrains the orientation of the V-groove cavity with respect to the macrosurface $\mathcal{M}$. We can explicitly characterize the cavity orientation as the bisector of the two microfacets with:

$$\theta_c = \frac{1}{2}|\theta_{l_r} - \theta_{v_r}|. \tag{6}$$

Substituting Eq. 2 into this gives us the following

$$\theta_c = \frac{1}{2}\left|\frac{\pi-\theta_l}{2}-\frac{\theta_v}{2}\right| \quad (7)$$

$$= \frac{1}{2}\theta_h. \quad (8)$$

This has a simple consequence on c; that is, light travels horizontally within an acute V-groove cavity if and only if the cavity is oriented along the vector

$$c = \frac{h+n}{||h+n||}. \quad (9)$$

This specific choice of c lies between two special cases of the unconstrained V-cavity model. If $c = n$, then we degenerate to the case in which all V-grooves are aligned with the surface (which is a model we are avoiding). If $c = h$, then the geometric masking is minimised, which means that we would account for too much energy on average. Our choice of c is somewhat of a compromise between these two cases; empirically it produces results that are convincing, while also being inexpensive to compute.

4.2 Indirect multiple scattering Geometric Term

Now that microfacet structure has been characterized, we construct an indirect geometric attenuation factor G_I using the cavity's opening angle θ_g and the cavity's orientation c.

Similar to the formulation in Lee et al. [2018], we next compute offset distances (d, d', k, k') from the unitary circle center O with respect of $l_\perp$ (Fig. 7a). We further simplify Lee's equations and only take two bounces into consideration, where these distances are computed as:

$$d = \sin(\theta_{l_c} + \theta_g/2), \quad (10)$$

$$d' = \sin(\theta_{l_c} - \theta_g/2), \quad (11)$$

$$k = \sin(\theta_{l_c} + 3\theta_g/2), \quad (12)$$

$$k' = \sin(\theta_{l_c} + 5\theta_g/2), \quad (13)$$

where the segment $\overline{KK'}$ is the portion of the left microfacet m_l that if lit, would be able to transmit energy from l to v in z bounces (where $z = 2$). In the segment $\overline{LK}$, the light would bounce $z-1$ times. Moreover, in the segment $\overline{K'O}$ the light would bounce $z+n$ times where $n \in \mathbb{N}$. Furthermore, the segment $\overline{DD'}$ is the portion of the left microfacet m_l that is able to receive energy from l. Thus, in the traditional V-cavity theory used by [Lee et al. 2018], and [Xie and Hanrahan 2018], the segment $\overline{KD'}$ is the only portion that is able to transmit energy from l to v in exactly two bounces, because the segment $\overline{D'K'}$ is under the shadow of the right microfacet m_r.

Therefore, the geometric attenuation factor from [Lee et al. 2018] is equivalent to:

$$G = (k-k')/d. \quad (14)$$

4.2.1 Asymmetrical Light Transport. We consider the microfacets m_l and m_r as asymmetrical regarding their light transport. This means that each microfacet is reflective from one side (the one inside the V-groove cavity) and transmissive from the other one (the one outside of the V-groove cavity). Strictly speaking, asymmetrical light transport is not physically-based; there is no existing material that exhibits such behaviour. However, this assumption enables us to indirectly compute a symmetrical light transport that is indeed physically plausible.

As seen in Fig. 7b, neglecting shadowing by allowing light to pass through the V-groove's outside, is based on the assumption that the microfacets are not connected to each other in the center of a unitary circle. Thus, reflection could be a distant phenomenon similar to the traditional Smith [1967] model that assumes masking and shadowing as uncorrelated distant properties [Heitz 2014; Heitz et al. 2016]. Therefore, for the right microfacet m_r it could exist a corresponding distant microfacet $m_r{}'$ with the same orientation lying in a position along the bouncing path b that casts no shadows on m_l (Fig. 7c). Furthermore, this implies that our geometric term is unidirectional because it only accounts for masking from v. Thus, our second bounce approximation is brighter because assumes that the microfacet m_l is always lit. The resulting overestimation help us to compensate for the energy transmitted by further bounces (rows b2 to b5 in Fig. 13) effectively getting our result closer to [Heitz et al. 2016] while simplifying our computation. This also means that our method is non-reciprocal, a limitation that we later discuss.

Therefore, we compute our G_I term with respect to the right asymmetrical microfacet m_r because the point K projects along the path vector b to the point R. Similarly, the point K' projects along the path vector b to the point P (Fig. 7c) (note that this simplification only occurs at the second bounce). So P is inherently the masking point from the view vector v over m_r. Thus, our final indirect geometry term G_I can be formulated as $1 - max(0, \overline{OP})$ (Fig. 8) where $\overline{OP}$ is obtained from

$$\overline{OP} = \frac{\overline{OP'}}{cos(\gamma)}, \quad (15)$$

where $\overline{OP'}$ is the shadowing of v and is derived by $\overline{OP'} = sin(\theta_{v_c} - \theta_m)$, γ is the angle between $v_\perp$ and m_r defined as $\gamma = \left(\frac{\pi}{2} - \theta_{v_c}\right) - \theta_m$ where θ_m is defined as

$$\theta_m = \frac{\theta_g}{2}. \quad (16)$$

Thus, G_I simplifies to

$$G_I = 1 - max\left(0, \frac{sin(\theta_{v_c} - \theta_m)}{sin(\theta_{v_c} + \theta_m)}\right). \quad (17)$$

Algorithm 1: Computation of G_I

$c \leftarrow normalize(h+n)$
$\theta_{v_c} \leftarrow acos(c \cdot v)$
$\theta_m \leftarrow (PI - acos(v \cdot l)) * 0.25$
$OP \leftarrow sin(\theta_{v_c} - \theta_m)/sin(\theta_{v_c} + \theta_m)$
$G_I \leftarrow 1.0 - max(0.0, OP)$

4.3 The full second bounce model

Comparatively, our constructions for D_I and F_I terms are less exotic. We choose our indirect normal distribution function $D_I(\theta_m)$ as a standard Trowbridge-Reitz distribution (GGX) [2013; 1975; 2007] using the cosine of θ_m (Eq. 16).

$$D_I = \frac{\alpha^2}{\pi(cos^2\theta_m(\alpha^2-1)+1)^2} \quad (18)$$

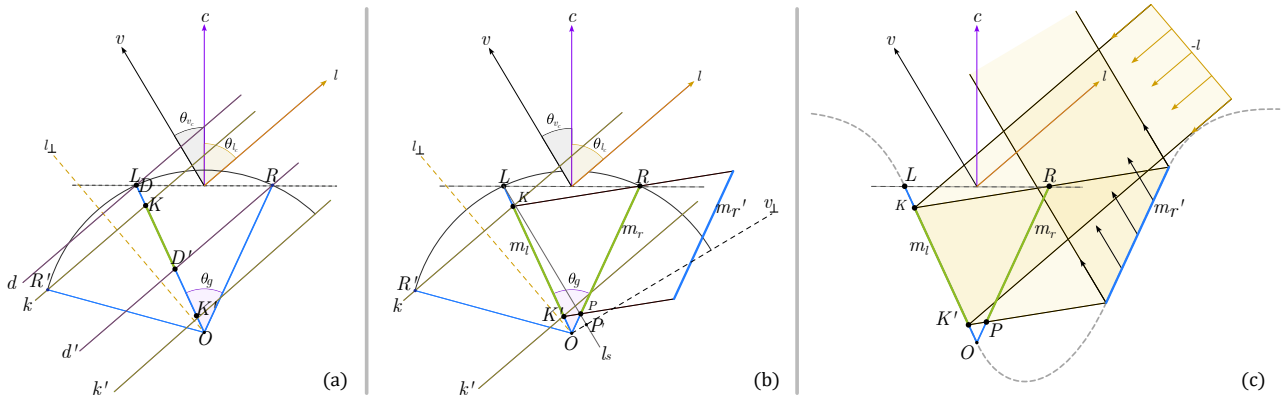

Figure 7: (a) In Zipin's method, the segment $\overline{KD'}$ is the only portion of the microfacet OL that is able to transmit energy from l to v in exactly two bounces. (b) We extend the portion of the microfacet m_l that is able to transmit energy from l to v through the segment $\overline{KK'}$, considering the microfacets as asymmetrical in terms of their light transport. This assumption is not consistent with the physical behaviour of real materials. However, it is able to represent the reflection from a distant microfacet $m_r{}'$ that casts no shadows over the microfacet m_l. (c) Our asymmetrical light transport assumption allows us to estimate a physically-plausible light transport event, with the simplicity of Zipin's geometric construction.

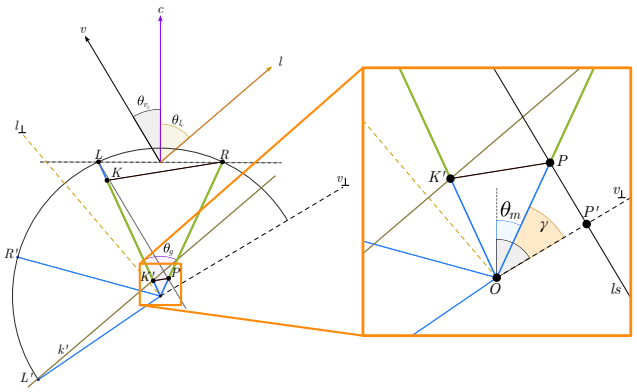

Figure 8: Our indirect multiple scattering geometric term G_I reduces to find the segment $\overline{OP}$ given the angle γ and the segment $\overline{OP'}$ computed with Zipin's approach.

effectively treating each asymmetric V-cavity as a single light reflecting entity oriented in θ_m.

Finally, the indirect Fresnel term F_I requires two separate Fresnel terms for each bounce. When computing the Schlick [1994] approximation, the angles θ_{l_r} and θ_{v_r}, which are equivalent to the $(i \cdot h)$ (i is the incident angle, and their h vector is aligned to their normal) of each microfacet, should be used. To further reduce the computation, we approximate the above process by simply squaring the F term from the direct contribution. Empirically we find the difference caused by this approximation to be visually negligible. Thus we take

$$F_I = F^2. \tag{19}$$

Combining all of these ingredients, we obtain our final indirect multiple scattering specular term:

$$f_{s_I} = \frac{D_I G_I F_I}{2(c \cdot v)}. \tag{20}$$

Note that this is an additive term, so we add Eq. 20 to Eq. 1 to obtain our final multiple scattering specular lobe.

4.4 Importance Sampling

We used the default sampling technique in Mitsuba 0.6 based on Heitz and d'Eon [2014] which restricts the sampling to the visible microfacet normals. Their approach is based on the GGX distribution and a one-way geometric function, our G_I function is also unidirectional, so we choose it as a practical solution for our offline implementation. However, using the same importance sampling could be valid for a direct scaling method such as Turquin [2018] but this is not exactly our case, so deriving a more accurate importance sampling approach for our V-cavity model is an interesting avenue for future work

5 RESULTS AND COMPARISONS

5.1 Experiment Configurations

To evaluate our method, we rendered images with a resolution of 512×512 with 512 samples per pixel in Mitsuba 0.6 [Jakob 2010] (Windows 10, 64 bits) using an Intel(R) Core(TM) i9-11900K at 3.50GHz with 16 Logical Processors and 64 GB of RAM. Renders from the competing methods are based on the original implementations provided by the corresponding authors.

As noted by Xie et al. [2021], the multiple scattering model from Heitz et al. [2016] is less reliable in smooth surfaces, mainly due to the variance caused by the stochastic nature of their method. However, for high roughness surfaces Heitz et al. [2016], it is often treated as an approximation of the ground truth and a reference point of comparison for other methods.

5.2 Comparisons

To validate our model, we perform both *quantitative* and *qualitative* experiments, comparing all of [Heitz et al. 2016; Lee et al. 2018;

Walter et al. 2007; Wang et al. 2022] to our model. We render a sample scene [Pilo 2010] of a shaderball illuminated by an environment map [Vogl 2010] for various roughness values.

Quantitative Comparisons. We use the Mean Square Error (MSE) to quantitatively assess the results from different models. In Fig. 9 Top we show the MSE versus roughness to compare our method, Lee et al. [2018], Walter et al. [2007], and Wang et al. [2022] against Heitz et al. [2016]. The data shows that for lower roughness values, our method and Wang et al. [2022] generate the lowest error, while for higher roughness values our method generates less error and more similar results to Heitz et al. [2016].

Qualitative Comparisons. Fig. 14 depicts a qualitative comparison between the proposed approach for an achromatic material, [Walter et al. 2007], [Lee et al. 2018],[Wang et al. 2022] and [Heitz et al. 2016] for different roughness values. The first row shows how the standard method from [Walter et al. 2007] gradually loses energy discarded by its single bounce geometry term. Generally, at low roughness values, all models behave similar to the classic single bounce GGX [Walter et al. 2007] due to the small contribution of multiple scattering (a perfectly smooth surface with $\alpha = 0$ exhibits no multiple scattering effect). The second row shows results from Lee et al. [2018]. Although the results show lower error on average, the error is heavily concentrated in the mirror direction (where h is closer to n). This results in a high-frequency specular response that is not consistent with rough surfaces. This is mainly caused by their assumption of V-cavities being aligned to the surface. The third row demonstrates Wang et al. [2022], which is darker than Heitz et al. [2016]. Fig. 14 shows that overall, the main advantage of our method is that it better approximates light scattering in high roughness, avoiding the mirror-like artifacts from Lee et al. [2018]. We also compared our method vs Turquin [2018] (Fig. 12) in a real-time scenario.The key distinction between both methods is that Turquin [2018] relies on a discretized approximation due to the use of LUTs, while ours is continuous yet fast enough to be used in production. More comparisons are provided in Fig. 11 where copper and gold materials are rendered for Heitz et al. [2016], Lee et al. [2018], Wang et al. [2022] and our method.

5.3 Efficiency

Fig. 9 Bottom shows the computation time for each model in Mitsuba 0.6 [2010]. We rendered images with a resolution of 128×128 with 512 samples per pixel. We repeated each experiment three times and averaged the results in order to reduce the variance in our reported times. Our method is nearly as fast as Walter et al. [2007], indicating its promise for real-time applications. Specifically, our second bounce approximation is an explicit solution and can be implemented in a conventional rasterization pipeline as an additional specular lobe (see Shadertoy and O3DE rendering results in Fig.10 and 12), providing high-fidelity results without sacrificing rendering speed. Unsurprisingly, the method of Heitz et al. [2016] takes longer to render. Unexpectedly, we found Lee et al. [2018] to be the most expensive model.

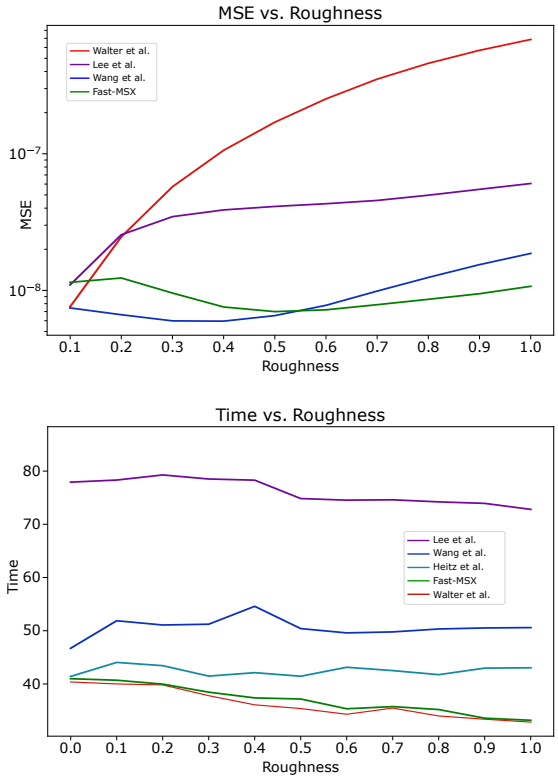

Figure 9: Top: MSE comparison between Walter et al. [2007], Lee et al. [2018], Wang et al. [2022] and our model versus Heitz et al. [2016]. Bottom: Comparison on time costs from the same methods vs our method rendered in Mitsuba 0.6 [2010]. Time in seconds.

6 CONCLUSION

We have proposed a novel methodology for computing multiple scattering reflection, accounting for first and second bounces of light. Our proposed technique efficiently models the second bounce multiple-scattering contribution. This greatly reduces the computational complexity compared to more sophisticated methods that rely on Monte Carlo simulation and thus reduces render times. Our results demonstrate the effectiveness and efficiency of our model in compensating the energy loss. Furthermore, in contrast to existing baselines [Lee et al. 2018; Wang et al. 2022; Xie and Hanrahan 2018], our algorithm shows superiority in avoiding specular artifacts or darkening at grazing angles and noise.

There are some limitations with our current approach. Firstly, by relaxing the assumptions of the model, we overestimate the second light bounce to recover most of the lost energy in later bounces. However, as shown in Fig. 13, there is still a small proportion of energy in the fourth and fifth bounce. Thus, our method is not able to recover all the energy and fails to pass the furnace test (Fig. 10). This is a trade-off between computation time and approximation fidelity.Secondly, similar to Turquin [2018], our method does not exhibit Helmholtz reciprocity hence, if used in a bi-directional path tracer, one would ideally need to compute the corresponding importance sampling procedure for the adjoint [Veach 1997] to avoid variance due to unbalanced contribution from light and camera paths.. Although we do not explore this option, this is an interesting topic for future work. Finally, the phenomenon of transmission is not covered by our model. We plan to investigate this in the future.

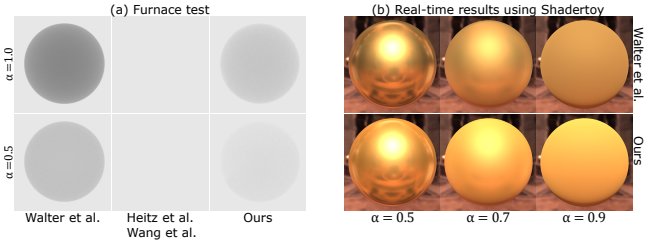

Figure 10: (a) Although [Heitz et al. 2016; Wang et al. 2022] pass the Furnace test, our method relies on the calculation of first and an approximation of the second bounce which leads to fail the Furnace test similar to GGX [Walter et al. 2007]. (b) Real-time results for our method and GGX [Walter et al. 2007] in Shadertoy [Quilez and Jeremias 2017].

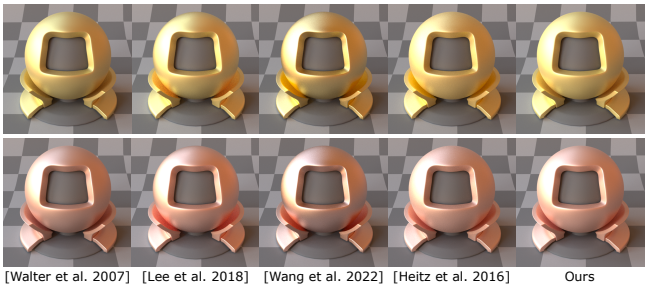

Figure 11: To be consistent with reported results in Heitz et al. [2016], we rendered gold and copper at the most realistic roughness value for metals ($\alpha = 0.3$). Lee et al. [2018] provides saturated results containing artifacts. In contrast, our method is artifact-free and similar to Heitz et al. [2016].

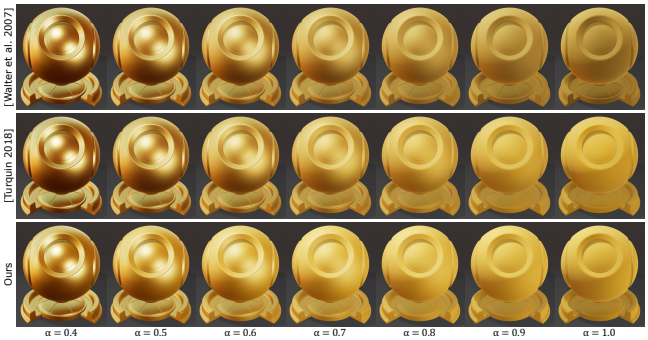

Figure 12: Comparison between Walter et al. [2007] (top) Turquin [2018] (middle) and our method (bottom) at higher roughness values. Both, Turquin [2018] and our method are able to retrieve the missing energy compared with Walter et al. [2007]. Rendered using O3DE [2023], illuminated by three point lights and an environment map [Vogl 2010].

REFERENCES

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

James F. Blinn. 1977. Models of Light Reflection for Computer Synthesized Pictures. *SIGGRAPH Comput. Graph.* (1977). https://doi.org/10.1145/965141.563893

Brent Burley. 2012. Physically-Based Shading at Disney. ACM SIGGRAPH Talks.

R. L. Cook and K. E. Torrance. 1982. A Reflectance Model for Computer Graphics. *ACM Trans. Graph.* (1982), 7–24. https://doi.org/10.1145/357290.357293

Jonathan Dupuy, Eric Heitz, and Eugene d'Eon. 2016. Additional Progress towards the Unification of Microfacet and Microflake Theories. In *Proceedings of the EGSR*. Eurographics Association.

Kempf EK. 2012. Scanning Electron Micrograph of the surface of a kidney stone showing tetragonal crystals of Weddellite.

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

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. https://doi.org/10.1111/cgf.12417 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12417

Eric Heitz, Johannes Hanika, Eugene d'Eon, and Carsten Dachsbacher. 2016. Multiple-Scattering Microfacet BSDFs with the Smith Model. *ACM Trans. Graph.* (2016). https://doi.org/10.1145/2897824.2925943

Stephen Hill, Stephen McAuley, Laurent Belcour, Will Earl, Niklas Harrysson, Sébastien Hillaire, Naty Hoffman, Lee Kerley, Jasmin Patry, Rob Pieké, et al. 2020. Physically based shading in theory and practice. In *ACM SIGGRAPH 2020 Courses*. 1–12.

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.* (2014). https://doi.org/10.1145/2601097.2601139

Csaba Kelemen and László Szirmay-Kalos. 2001. A Microfacet Based Coupled Specular-Matte BRDF Model with Importance Sampling. *In Eurographics Short Presentations* (2001).

Christopher Kulla and Alejandro Conty. 2017. Revisiting physically based shading at imageworks. *SIGGRAPH Course, Physically Based Shading* (2017).

Joo Ho Lee, Adrian Jarabo, Daniel S. Jeon, Diego Gutierrez, and Min H. Kim. 2018. Practical Multiple Scattering for Rough Surfaces. *ACM Trans. Graph.* (2018). https://doi.org/10.1145/3272127.3275016

Stephen McAuley, Stephen Hill, Adam Martinez, Ryusuke Villemin, Matt Pettineo, Dimitar Lazarov, David Neubelt, Brian Karis, Christophe Hery, Naty Hoffman, and Hakan Zap Andersson. 2013. Physically based shading in theory and practice. In *ACM SIGGRAPH 2013 Courses*.

O3DE. 2023. *O3DE - the open source real-time 3D engine*. https://o3de.org/

Jonas Pilo. 2010. . https://www.mitsuba-renderer.org/scenes/matpreview.zip

Iñigo Quilez and Pol Jeremias. 2017. Shadertoy. (2017). https://www.shadertoy.com

Christophe Schlick. 1994. An Inexpensive BRDF Model for Physically-based Rendering. *Computer Graphics Forum* (1994). https://doi.org/10.1111/1467-8659.1330233

Bruce G. Smith. 1967. Geometrical shadowing of a random rough surface. *IEEE Transactions on Antennas and Propagation* 15 (1967), 668–671.

K. E. Torrance and E. M. Sparrow. 1967. Theory for Off-Specular Reflection From Roughened Surfaces. *J. Opt. Soc. Am.* (Sep 1967). https://doi.org/10.1364/JOSA.57.001105

T. S. Trowbridge and K. P. Reitz. 1975. Average irregularity representation of a rough surface for ray reflection. *J. Opt. Soc. Am.* (1975). https://doi.org/10.1364/JOSA.65.000531

Emmanuel Turquin. 2018. Practical multiple scattering compensation for microfacet models. https://api.semanticscholar.org/CorpusID:221737278

Eric Veach. 1997. *Robust Monte Carlo Methods for Light Transport Simulation*. Ph. D. Dissertation. Advisor(s) Guibas, Leonidas J.

Bernhard Vogl. 2010. . http://dativ.at/lightprobes/

Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. 2007. Microfacet Models for Refraction through Rough Surfaces. In *Proceedings of EGSR (EGSR'07)*. Eurographics Association.

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.* 4 (2022). https://doi.org/10.1145/3528223.3530112

Bartosz Woźniak. 2014. Titanium screw "volcano" pores coated with ZnO nanoparticles.

Feng Xie, James Bieron, Pieter Peers, and Pat Hanrahan. 2021. Experimental Analysis of Multiple Scattering BRDF Models. In *SIGGRAPH Asia 2021 Technical Communications*. Association for Computing Machinery. https://doi.org/10.1145/3478512.3488601

Feng Xie and Pat Hanrahan. 2018. Multiple Scattering from Distributions of Specular V-Grooves. *ACM Trans. Graph.* (2018). https://doi.org/10.1145/3272127.3275078

Feng Xie, Anton Kaplanyan, Warren Hunt, and Pat Hanrahan. 2019. Multiple scattering using machine learning. In *ACM SIGGRAPH 2019 Talks*.

Richard B. Zipin. 1966. The apparent thermal radiation properties of an isothermal V-groove with specularly reflecting walls. *Journal of Research of the National Bureau of Standards, Section C: Engineering and Instrumentation* (1966).

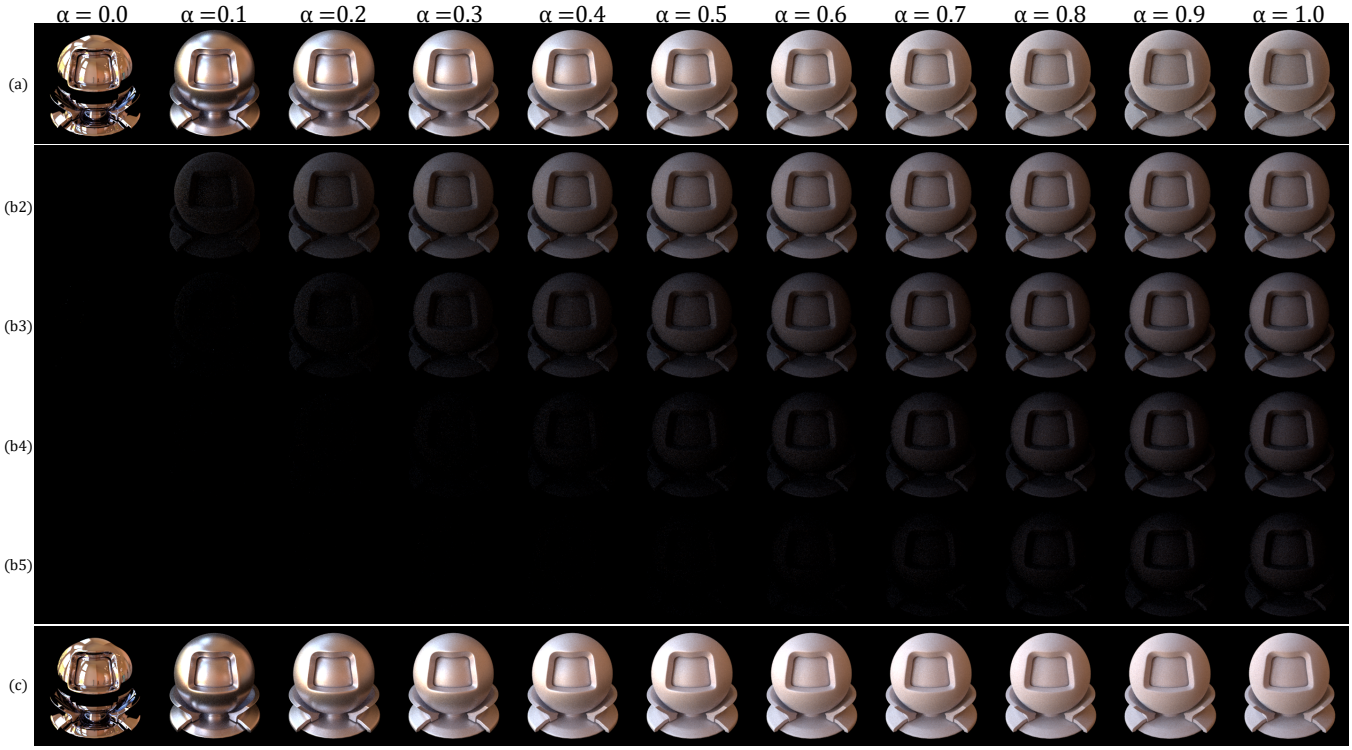

Figure 13: (a) First bounce from [Heitz et al. 2016] equivalent to [Walter et al. 2007]. Rows (b2) to (b5), show the isolated contribution of each bounce using [Heitz et al. 2016] method. (C) final result from [Heitz et al. 2016]. (α) denotes the roughness.

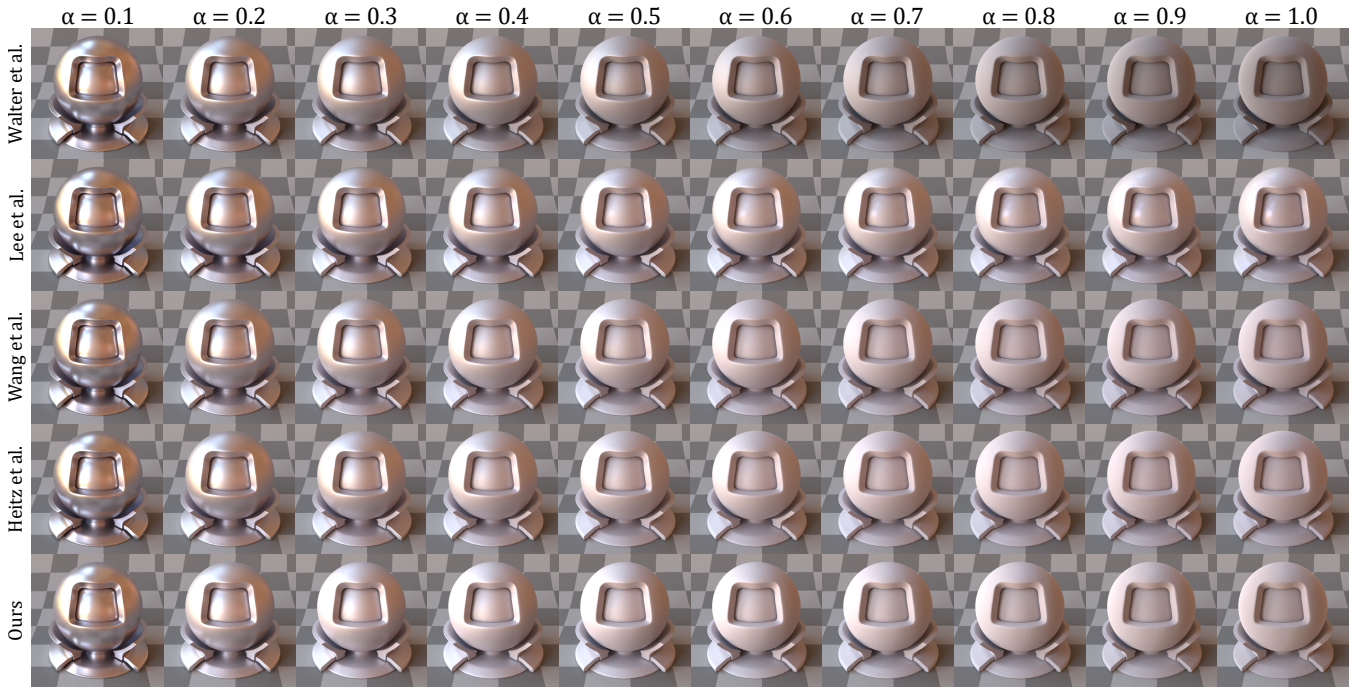

Figure 14: Qualitative comparison between [Walter et al. 2007], [Lee et al. 2018], [Wang et al. 2022], [Heitz et al. 2016], and our method. [Lee et al. 2018] has notable specular artifacts at high roughness values $(0.6 - 1.0)$ whereas our method is artifact free.