



HAL
open science

Glint Rendering based on a Multiple-Scattering Patch BRDF

Xavier Chermain, Frédéric Claux, Stephane Merillou

► **To cite this version:**

Xavier Chermain, Frédéric Claux, Stephane Merillou. Glint Rendering based on a Multiple-Scattering Patch BRDF. Computer Graphics Forum, 2019, 38 (4), pp.27-37. 10.1111/cgf.13767 . hal-02364885

HAL Id: hal-02364885

<https://hal.science/hal-02364885v1>

Submitted on 15 Nov 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Glint Rendering based on a Multiple-Scattering Patch BRDF

Xavier Chermain, Frédéric Claux and Stéphane Mérillou
Univ. Limoges, CNRS, XLIM, UMR 7252, F-87000 Limoges, France.

2019-10-31

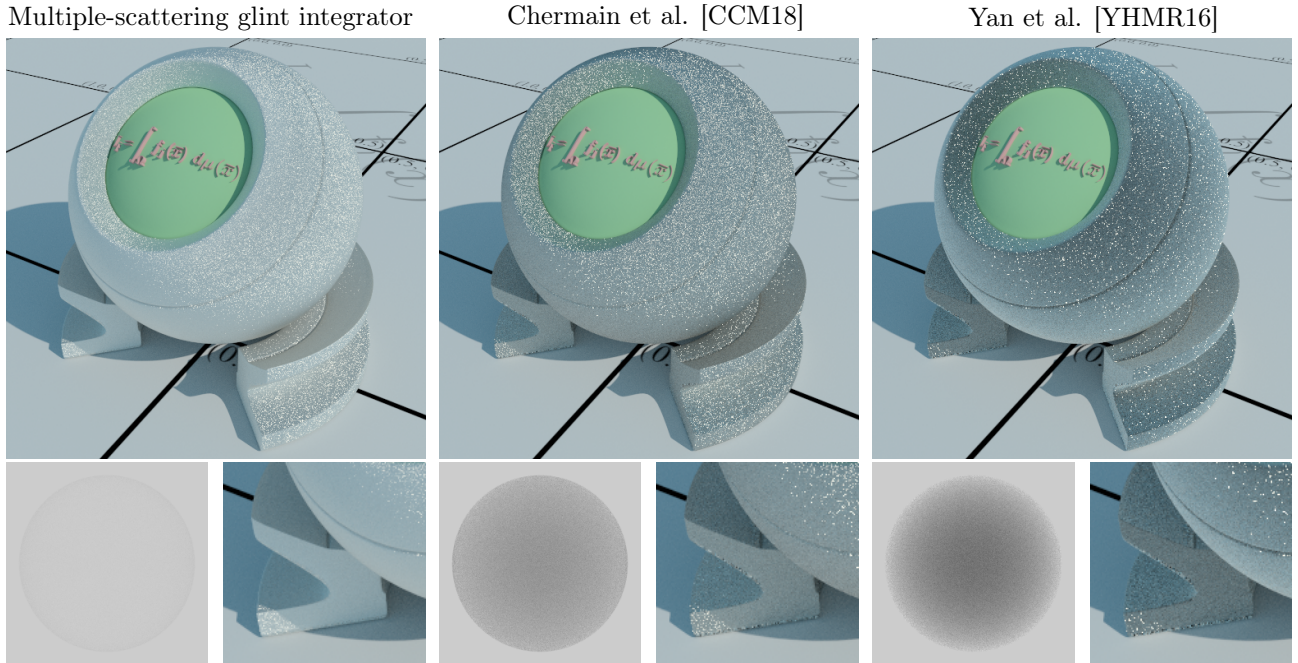


Figure 1: Glittery orb illuminated by an environment map using 1,024 spp. The micro-surface is modeled by a specular normal map with high RMS roughness ($\sigma = 1$). Previous methods based on classic normal mapping and single-scattering BRDF evaluations darken the appearance and create black holes (insets and white furnace tests). Importance sampling is suboptimal and causes fireflies. Our new multiple-scattering glint integrator tackles these issues. It produces renderings with no artifacts, and almost passes the white furnace test, for an extra 36% rendering time for this scene.

Abstract

Rendering materials such as metallic paints, scratched metals and rough plastics requires glint integrators that can capture all micro-specular highlights falling into a pixel footprint, faithfully replicating surface appearance. Specular normal maps can be used to represent a wide range of arbitrary micro-structures. The use of normal maps comes with important drawbacks though: the appearance is dark overall due to back-facing normals and importance sampling is suboptimal, especially when the micro-surface is very rough. We propose a new glint integrator relying on a multiple-scattering patch-based BRDF addressing these issues. To do so, our method uses a modified version of microfacet-based normal mapping [SHHD17] designed for glint rendering, leveraging symmetric microfacets. To model multiple-scattering, we re-introduce the lost energy caused by a perfectly specular, single-scattering formulation instead of using expensive random walks. This reflectance model is the basis of our patch-based BRDF, enabling robust sampling and artifact-free rendering with a natural appearance. Additional calculation costs amount to about 40% in the worst cases compared to previous methods [YHMR16, CCM18].

1 Introduction

Glittery materials, metallic paints, rough, brushed and scratched plastics and metals are surface types that are challenging to render in computer graphics. These materials have tiny mirror-like details visible only under

sharp and powerful lighting. A major feature of these materials is the spatial coherence of their hundreds of high frequency specular lobes, where small light or camera translations reveal the micro-surface. Using smooth normal distribution functions (NDF) to model material appearance is common for rendering specular surfaces, because these functions integrate well with common, field-proven shading models and frameworks. They are however ineffective at representing real life surfaces having complex micro-structures. Traditional surface filtering in computer graphics uses low-pass filters to reduce calculation costs, restricting NDFs to simple definitions, poorly replicating material appearance.

To capture the glints featured by specular surfaces, the reflectance model cannot use traditional Bidirectional Reflectance Distribution Functions (BRDF) based on smooth NDFs. More sophisticated models have to be used, such as *glint integrators*. This family of methods never simplify the micro-surface, and guarantees the rendering of all the micro-details falling into a screen pixel. The complexity of the material is fully taken into account and its related appearance is meant to be faithfully reproduced.

A Recent Interest. Rendering of specular micro-structures has received a growing interest in the last five years. The contributions of Jakob et al. [JHY⁺14] and Yan et al. [YHJ⁺14] have inspired many recent works exploring procedural materials [AK16, ZK16, WWH18], primitive-based BRDF [RGB16] or normal map-based glint integrators [YHMR16, CCM18, GGN18]. Each kind of method has associated strengths and weaknesses. Some methods specialize into a particular type of micro-structure (scratches, mirror-like flakes), sometimes achieving multiple-scattering effects and proposing good sampling routines. Others are more versatile, handling a large range of materials (glittery surfaces, brushed and scratched metals), often at the expense of memory consumption and rendering quality.

Remaining Issues. Most glint integrators based on normal maps or procedural surfaces simulate one bounce of light on the micro-surface. The rendered images produced with these integrators come out unnaturally dark, especially when using very rough surfaces (Figure 1, middle and right). A lot of surface normals are back-facing when the viewing or lighting direction is grazing. These unfortunate cases prevent light from being reflected. The second issue is an inefficient sampling procedure, yielding high sample weights, introducing isolated but nonetheless disturbing firefly artifacts (Figure 12). Current state-of-the art glint integrators randomly pick a normal of the micro-surface without taking into account the masking and interreflections that are inherent to the geometry of the micro-surface. At very grazing angles, about half of the samples are wasted and high weights compensate for these losses.

Position of our Work. The aim of this paper is to address all these issues. We choose to represent the micro-surface by a *normal map*, for its versatility. Our representation is based on a dual normal model, inspired by the work of Schüssler et al. [SHHD17]. A key difference is that we use a symmetric dual, not a tangent one (vertical), so as to not alter material roughness and related overall appearance. We take into account multiple-scattering by re-introducing the energy lost by a sole one-bounce BRDF formulation, a quantity that we analytically derive. This lost energy quantity is approximated and compensated for by using an energy-compensation BRDF, which is added to the single-scattering expression, eventually forming our final *local* BRDF. For importance sampling, we sample the single-scattering BRDF and keep the sample only if it is not shadowed or transmitted (Figure 5). Otherwise, we sample the energy-compensation BRDF, to account for multi-scattering.

Our local BRDF is then used as part of a footprint or *patch* based glint integrator. Our patch-BRDF is a discrete weighted sum of local BRDF instances, and is energy conserving.

Our contributions in the field of glint rendering are the following:

- We propose the first multiple-scattering glint integrator based on a specular normal map.
- The proposed patch-BRDF does not have the rendering-quality related issues of the previous methods: black holes on the surface, dark overall appearance and fireflies.
- The sampling procedure is optimal: no samples are wasted, the probability density function (PDF) is analytic and sample weights are one when the surface is non-absorptive.

Our representation is Jacobian free, i.e. it does not take into account surface curvature. It is asymmetric, therefore it limits the use of our method to forward path tracing. Our BRDF compensation scheme is restricted to specular surfaces. The dual normal approach may change the look of some materials in unwanted ways, even though this has not been significant for the examples in the paper.

After a detailed review of the related works (Section 2), we present our local multiple-scattering BRDF in Section 3. We use this local model to build the patch-BRDF in Section 4, which filters the surface and guarantees the spatial coherence of the glints. We evaluate our method and discuss the results in Section 5.

2 Related Work

Our work is related to three different areas of research in computer graphics: glint integrators, microfacet-based normal mapping and multiple-scattering BRDFs. We review here the most relevant works in these areas.

2.1 Glint Integrators

Glint integrators can be seen as special types of BRDFs where specular micro-structures are rendered using either geometrical primitives randomly placed on the surfaces or high definition normal maps. Instead of relying on stochastic Monte Carlo sampling to estimate shading, glint integrators calculate all the radiance leaving a ray footprint (pixel-bound for primary rays). The BRDF associated with the footprint is usually referred to as a patch BRDF.

Stochastic surfaces. The first method making use of stochastic specular micro-structures is the method of Jakob et al. [JHY⁺14] and focuses on glittery materials. The micro-surface is composed of mirror flakes, randomly placed on the surface. These are used to compute a discrete, patch-related NDF used as part of a microfacet-based BRDF. An acceleration hierarchy speeds up contributing reflector queries. Their method only models single-scattering and micro-flakes can be back-facing and consequently dismissed. The BRDF normalization is global, not footprint (patch) based. Importance sampling is also not optimal, as non-contributing reflectors may inadvertently be sampled. The extension proposed by Atanasov and Koylazov [AK16] alleviates this issue, at the price of going through all the flakes falling into a footprint, which is inefficient. Other works concentrate on improving performance [WWH18]. Selected methods offer real time performance [ZK16, WB16] but cannot really be used for photorealistic rendering.

The work of Raymond et al. [RGB16] models specular scratches randomly distributed over a surface. Users have control over the profile, the orientation and the density of the scratches. They simulate multiple-scattering by pre-computing the reflectance distributions of scratch profiles with a 2D ray tracer. The full 3D-BRDF is reconstructed by interpolating these 2D-BRDFs. Mipmapped filtering is used to calculate the relative area of the scratches in a footprint. Their method gives realistic results, its importance sampling procedure is efficient, multiple-scattering is taken care of, but they only handle scratched surfaces.

Glint Integrators based on a Normal Map. As opposed to methods using procedural surfaces, glint integrators based on normal maps provide more flexibility over material appearance. Scratched, brushed, rough metals and glittery objects can be modeled by changing the normal map contents. The work of Yan et al. 2014 [YHJ⁺14] was the first to introduce this micro-surface representation, obtaining realistic renderings of a wide range of materials. Several papers improve the performance of their method by either using 4D NDF mixtures [YHMR16] or spherical harmonics [GGN18]. The work of Chermain et al. [CCM18] focuses on rendering quality and proposes footprint-relative normalization factors. The importance sampling procedure is efficient since it uses footprint-based NDFs, however visibility information is not used, causing fireflies and zero-radiance paths. The method of Gamboa al. [GGN18] does not use importance sampling. Instead, they efficiently integrate all the incident radiance, at the cost of a large bunch of pre-computed data. The common shortcoming of these methods is their single-scattering formulation, causing black overall appearance for rough surfaces. They also rely on classic normal mapping, where back-facing normals lead to black holes on the surface. Our work tackles these issues.

2.2 Microfacet-Based Normal Mapping

Classic normal mapping causes black fringes on surfaces and requires complex re-normalization factors, as demonstrated in the work of Schüssler et al. [SHHD17]. To solve these issues, they propose a microfacet-based BRDF faking normal perturbations, using tangent facets to seal the micro-surface. The tangent facets make really continuous an otherwise discontinuous normal map composed of unrelated, individual discrete facets. They rely on random walks to model infinite multiple-scattering on the micro-surface, and derive an interesting closed-form expression for the two-scattering order. They obtain much more realistic results than classic normal mapping. Our model is inspired by their work but uses a symmetric dual normal, more appropriate for glints rendering that would otherwise be negatively impacted by vertical facets (Figure 2). Another major difference is that Schüssler et al. use expensive random walks to account for multiple-scattering, whereas we estimate the energy going into indirect bounces and compensate for it by using an alternate (or energy-compensation) BRDF to increase performance. Figure 3 gives a quick hint at the differences between classic normal mapping, the model of Schüssler et al. and our model.




Figure 2: Comparison between a smooth Beckmann surface (left), a microfacet-based normal mapped surface using tangent normals (middle, [SHHD17]) and using symmetric normals (right, our micro-surface). The normal map being very dense in this Figure, tangent facets introduce unwanted roughness compared to the smooth model while our micro-surface does not. Even with 262,144 spp in this scene, there still are fireflies as random walks are used to model multiple-scattering.

2.3 Multiple-Scattering BRDFs

Recently, the multiple-scattering of light over micro-surfaces has gained a lot of interest in the graphics community. This is thanks to the work of Heitz et al. [HHdD16] on microfacet BSDFs using the Smith model. They provide ground truth results by matching stochastic reflectance functions with simulated data. The bottleneck of their method is the random walk approach that requires several evaluations to provide converged results. Recent works propose more practical multiple-scattering reflectance models, using distributions of specular V-grooves [LJJ⁺18, XH18, SPSH19]. These V-cavity-based methods are too expensive for our purpose. Finally, to efficiently model multiple-scattering, several works directly re-integrate the energy lost by a single-scattering formulation, focusing on either reciprocity [CK17, Hil18, FA19] or efficient sampling [Tur19]. We also use energy compensation, but do it in the context of glint rendering using microfacet-based normal mapping. Our method offers robust importance sampling, and a customizable energy-compensation BRDF.

3 Local Multiple-Scattering BRDF

Our model is built around the definition of a local and a patch-wide BRDFs. We start with the formulation of the local BRDF.

3.1 Micro-Surface Model Overview

Our micro-surface model is based on an interpolated, continuous normal map typically built from a discrete texture map. Each position u over the map has a perturbed normal ω_p as well as an associated dual normal ω_s . This dual normal is symmetric to ω_p with respect to the geometric normal ω_g , forming a V-shaped cavity




Figure 3: Left: classic normal mapping. Middle: normal mapping of Schüssler et al. [SHHD17]. Right: our model. The symmetric facets ω_s do not change overall material roughness the way tangent facets ω_t do. For glint rendering, where the normal map is highly repeated, tangent facets would alter too much the appearance.

(Figure 3). Locally, at any position u over the surface, the orientation of the micro-normals ω_p and ω_s are used to derive the probability that an incident or outgoing ray direction hits either normals, taking into account the shading induced by its symmetric counterpart. All V-cavities have the same micro-BRDF. We detect situations incurring a loss of energy caused by local single-scattering within the V-cavity and account for this loss by also evaluating an additional BRDF. Stochastic Monte-Carlo rendering leverages this model to importance sample the surface locally.

3.2 Normal Distribution Function and Projected Areas

Normal distribution functions (NDF) are widely used to model micro-surfaces. They give the density of micro-facets ω_m at a surface position u . In our case, this function is composed of two Dirac distributions centered around the perturbed normal ω_p and the symmetric normal ω_s (we drop the dependence on u for the rest of the section):

$$D(\omega_m) = \frac{1}{2} \frac{\delta_{\omega_p}(\omega_m)}{\omega_p \cdot \omega_g} + \frac{1}{2} \frac{\delta_{\omega_s}(\omega_m)}{\omega_s \cdot \omega_g}. \quad (1)$$

This NDF is a Cook-Torrance V-cavity NDF. Its projected area onto the geometric normal is exactly one [CT81, Hei14]. The projected areas of a facet and its dual towards the observation direction ω_o are

$$a_p(\omega_o) = \frac{\langle \omega_o, \omega_p \rangle}{\langle \omega_p, \omega_g \rangle}, \quad a_s(\omega_o) = \frac{\langle \omega_o, \omega_s \rangle}{\langle \omega_s, \omega_g \rangle} \quad (2)$$

where $\langle \omega_1, \omega_2 \rangle$ is the clamped dot product. From these definitions, the probabilities $\lambda_p(\omega_o)$ and $\lambda_s(\omega_o)$ that an observation ray intersects ω_p or ω_s can be derived (Figure 4):

$$\lambda_p(\omega_o) = \frac{a_p(\omega_o)}{a_p(\omega_o) + a_s(\omega_o)}, \quad \lambda_s(\omega_o) = \frac{a_s(\omega_o)}{a_p(\omega_o) + a_s(\omega_o)}. \quad (3)$$

The terms λ_p and λ_s will be used to evaluate (Section 3.4) and sample (Section 3.6) the BRDF.

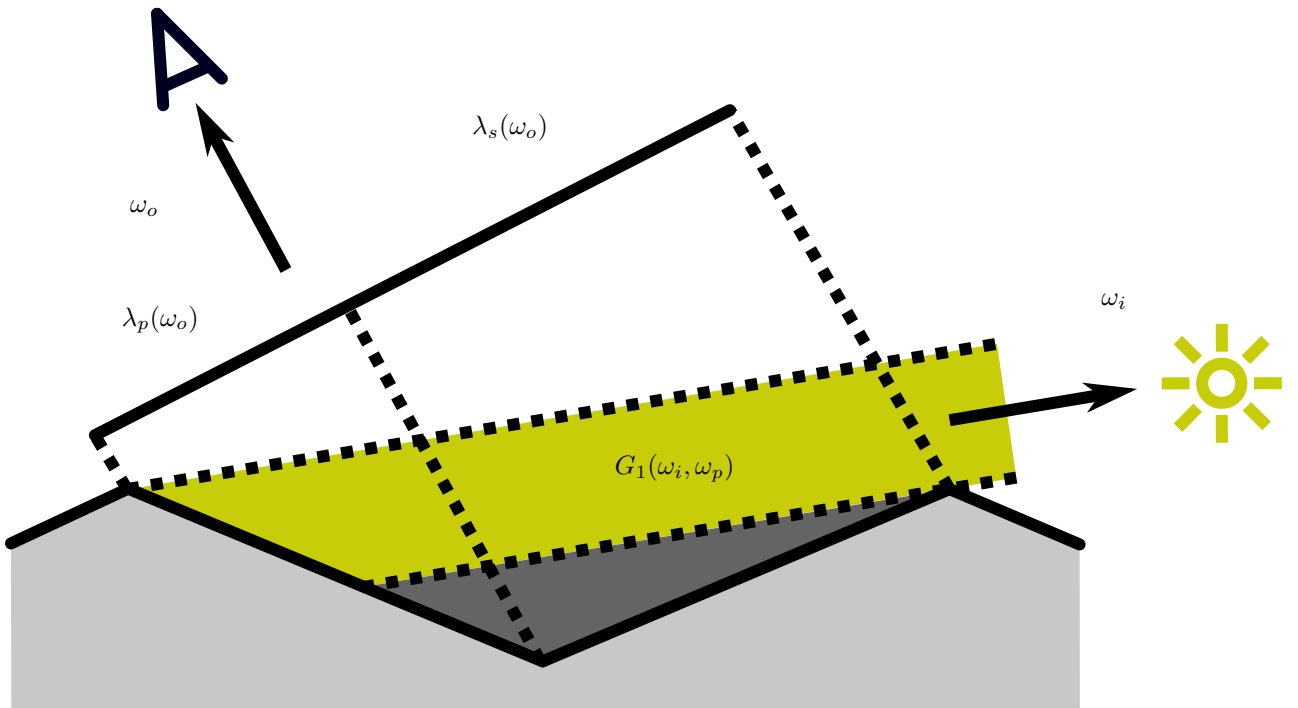


Figure 4: Illustration of the intersection probabilities and the shadowing function.

3.3 Masking Function

The masking function gives the proportion of microfacets ω_p or ω_s that are visible from the observation direction ω_o . It is called the shadowing function when using the incident direction ω_i as parameter (Figure 4). The V-

Cavity masking function is [CT81, Hei14]:

$$\begin{aligned} G_1(\omega_o, \omega_p) &= H(\langle \omega_o, \omega_p \rangle) \min \left(1, \frac{\langle \omega_o, \omega_g \rangle}{\frac{1}{2}a_p(\omega_o) + \frac{1}{2}a_s(\omega_o)} \right) \\ &= H(\langle \omega_o, \omega_p \rangle) \min \left(1, \frac{2\langle \omega_o, \omega_g \rangle \langle \omega_p, \omega_g \rangle}{\langle \omega_o, \omega_p \rangle + \langle \omega_o, \omega_s \rangle} \right) \end{aligned} \quad (4)$$

for the facet with normal ω_p . For the micro-normal ω_s , it is

$$G_1(\omega_o, \omega_s) = H(\langle \omega_o, \omega_s \rangle) \min \left(1, \frac{2\langle \omega_o, \omega_g \rangle \langle \omega_s, \omega_g \rangle}{\langle \omega_o, \omega_p \rangle + \langle \omega_o, \omega_s \rangle} \right) \quad (5)$$

where $H(-)$ is the Heaviside function.

As in Schüssler et al. [SHHD17], we assume no correlation of masking and shadowing, giving a simple separable masking-shadowing function (here for the micro-normal ω_p):

$$G_2(\omega_o, \omega_i, \omega_p) = G_1(\omega_o, \omega_p)G_1(\omega_i, \omega_p) \quad (6)$$

where ω_i is the incident direction. This choice has an important implication for the importance sampling procedure (Section 3.6). We have to determine when a ray is shadowed and with this assumption, we can directly use the shadowing function $G_1(\omega_i, \omega_m)$ instead of using explicit ray-facet intersections [SHHD17].

3.4 Single-Scattering BRDF

In the context of glint rendering, the material of the two facets is very specular. Assuming that the two facets of our micro-surface have the same micro-BRDF f_m , the cosine-weighted single-scattering BRDF f_1 for our V-cavity is

$$\begin{aligned} f_1(\omega_o, \omega_i, \omega_p) \langle \omega_i, \omega_g \rangle &= \lambda_p(\omega_o) f_m(\omega_o, \omega_i) \langle \omega_i, \omega_p \rangle G_1(\omega_i, \omega_p) \\ &\quad + \lambda_s(\omega_o) f_m(\omega_o, \omega_i) \langle \omega_i, \omega_s \rangle G_1(\omega_i, \omega_s). \end{aligned} \quad (7)$$

The BRDF expression f_1 is the result of the integral of the contribution of each micro-normal ω_m over the hemisphere Ω [SHHD17]. The shadowing function and intersection probability of each facet (Figure 4) weight the corresponding facet's cosine-weighted micro-BRDF f_m , rotated on either ω_p or ω_s . Note that f_1 does not create any energy and is symmetric as long as the same goes for f_m .

3.5 Multiple-Scattering BRDF

In this section, we demonstrate how we estimate f_∞ , the multiple-scattering BRDF, for a single cavity. Its definition is based on f_1 and an energy-compensation BRDF f_{2+} meant to represent second-or-more order scattering. We estimate the energy $1 - E_1$ going through this latter BRDF, and weight its evaluation using this value.

Energy Term Assuming that the micro-BRDF f_m is perfectly specular and non-absorptive, the directional albedo E_1 of f_1 , also called the energy term, is

$$\begin{aligned} E_1(\omega_o, \omega_p) &= \int_{\Omega} f_1(\omega_o, \omega_i, \omega_p) \langle \omega_i, \omega_g \rangle d\omega_i \\ &= \lambda_p(\omega_o) G_1(\omega_r(\omega_o, \omega_p), \omega_p) \\ &\quad + \lambda_s(\omega_o) G_1(\omega_r(\omega_o, \omega_s), \omega_s) \end{aligned} \quad (8)$$

where $\omega_r(\omega_1, \omega_2)$ is the reflected direction of ω_1 around the normal ω_2 . The term E_1 gives the radiance that escapes the micro-surface after the first bounce of light in the outgoing direction ω_o . For the single-scattering BRDF f_1 , this value ranges from 0 to 1 for any different couple of (ω_o, ω_p) .

In order to achieve 100% energy conservation, we define our local multiple scattering BRDF f_∞ as

$$\begin{aligned} f_\infty(\omega_o, \omega_i, \omega_p) &= \\ &= f_1(\omega_o, \omega_i, \omega_p) + (1 - E_1(\omega_o, \omega_p)) f_{2+}(\omega_o, \omega_i). \end{aligned} \quad (9)$$

Weighted by the cosine, the BRDF integrates to one over incident directions

$$\begin{aligned} &\int_{\Omega} f_\infty(\omega_o, \omega_i, \omega_p) \langle \omega_i, \omega_g \rangle d\omega_i \\ &= E_1(\omega_o, \omega_p) + (1 - E_1(\omega_o, \omega_p)) \int_{\Omega} f_{2+}(\omega_o, \omega_i) \langle \omega_i, \omega_g \rangle d\omega_i \\ &= 1 + E_1(\omega_o, \omega_p) - E_1(\omega_o, \omega_p) = 1 \end{aligned} \quad (10)$$

only if the energy-compensation BRDF $f_{2+}(\omega_o, \omega_i)\langle\omega_i, \omega_g\rangle$ does. The choice of f_{2+} is arbitrary and is discussed in Section 5.2. Our formulation for f_∞ has a simple closed-form expression and can be evaluated efficiently. It is non-symmetric as can be seen in Equation 9, swapping ω_o and ω_i yields a different E_1 term ($E_1(\omega_o, \omega_p) \neq E_1(\omega_i, \omega_p)$). This is only a deal breaker when using bidirectional path tracing, and does not negatively impact visual appearance otherwise.

In practice, we mollify the Dirac distribution of f_m to a NDF with tiny roughness α_m [YHJ⁺14, YHMR16, CCM18] in order to evaluate f_1 (Equation 7). Note that the energy term of Equation 8 still guarantees energy conservation with a mollified Dirac. Section 5.1 shows that f_∞ passes the white furnace test when the f_m NDF has a tiny roughness, validating this approximation.

Fresnel Term. The above model does not incorporate Fresnel effects, which give the object its color and is, in the case of multiple scattering, responsible for the saturation effect that is visible on rough surfaces. We therefore need to choose the absorption coefficient for the energy-compensation BRDF f_{2+} , but it is not trivial. We choose to approximate it by the Fresnel term F_1 of the micro-BRDF f_m . The local multiple-scattering BRDF f_∞ for an absorptive normal map is now

$$f_\infty(\omega_o, \omega_i, \omega_p) = f_1(\omega_o, \omega_i, \omega_p) + (1 - E_1(\omega_o, \omega_p))F_1(\omega_o, \omega_g)f_{2+}(\omega_o, \omega_i), \quad (11)$$

where $F_1(\omega_o, \omega_g)$ colors the non-absorptive f_{2+} BRDF the same way it does for f_1 . Note that Turquin [Tur19] performs a similar approximation for its compensated microfacet BRDF. His model exhibits similar results as the ground truth [HHdD16].

3.6 Importance Sampling the Local BRDF

To importance sample f_∞ , we first sample the cosine-weighted single-scattering BRDF f_1 by choosing a facet using the intersection probabilities $\lambda_p(\omega_o)$ and $\lambda_s(\omega_o)$. A direction ω_i is then picked by importance sampling the micro-BRDF f_m of the sampled micro-normal. If the sampled direction ω_i is shadowed (Figure 5, middle) or goes inside the surface (Figure 5, left), multiple-scattering occurs and f_{2+} is sampled. The probability that the sampling of the single-scattering BRDF f_1 fails, i.e. ω_i is shadowed or transmitted, is exactly equal to $1 - E_1(\omega_o, \omega_p)$. Our importance sampling procedure is summarized in Algorithm 1 (\mathcal{U} is a uniform random number in the $[0, 1[$ range) and is illustrated in Figure 5. It is guaranteed to give weights w between 0 and 1 as long as f_m and f_{2+} also do. The weight will always be 1 if there is no Fresnel effect, considering f_m has a near-zero roughness. Otherwise, w is multiplied by the Fresnel term. In all cases, arbitrary large weights cannot happen and firefly artifacts are avoided.

Algorithm 1: Sample $f_\infty(\omega_o, \omega_i, \omega_p)\langle\omega_i, \omega_g\rangle$

```

 $\omega_m \leftarrow [\mathcal{U} < \lambda_p(\omega_o) ? \omega_p : \omega_s]$  // Choose a facet
 $(\omega_i, w) \leftarrow \text{sample } f_m(\omega_o, \omega_i, \omega_m)\langle\omega_i, \omega_m\rangle$  // Sample  $f_m$ 
if  $\omega_i \cdot \omega_g \leq 0$  or  $\mathcal{U} \geq G_1(\omega_i, \omega_m)$  then
    /* Ray goes inside the surface or is shadowed */
     $(\omega_i, w) \leftarrow \text{sample } f_{2+}(\omega_o, \omega_i)\langle\omega_i, \omega_g\rangle$  // Sample  $f_{2+}$ 
end

```




Figure 5: Different cases of local multiple-scattering BRDF sampling.

4 Multiple-Scattering Patch BRDF

In this section, we introduce the theoretical background to define, evaluate and sample patch BRDFs (\mathcal{P} -BRDFs) based on a normal map. We then explain how to efficiently compute our \mathcal{P} -BRDF based on the local multiple-scattering BRDF f_∞ introduced in the previous section. We close this section with our importance sampling procedure.

4.1 Patch BRDF Definition

Our \mathcal{P} -BRDF computes the density of outgoing radiance for a flat geometric surface \mathcal{P} instead of a surface point u . The patch \mathcal{P} matches the ray footprint. It is represented by a normalized low-pass filter $k_{\mathcal{P}}$, i.e. $\int_{\mathbb{R}^2} k_{\mathcal{P}}(u) du = 1$ and antialiases the radiance of all the glints in the footprint. Several techniques for surface filtering [Ige99, SW01, BYRN17] give methods to construct \mathcal{P} during path tracing.

Our \mathcal{P} -BRDF $f_{\mathcal{P}}$ is defined as an integral of local BRDFs weighted by the ray footprint $k_{\mathcal{P}}(u)$ at each position u :

$$f_{\mathcal{P}}(\omega_o, \omega_i) = \int_{\mathbb{R}^2} f(\omega_o, \omega_i, \omega_p(u)) k_{\mathcal{P}}(u) du \quad (12)$$

where $f(\omega_o, \omega_i, \omega_p(u))$ is the surface's BRDF perturbed by $\omega_p(u)$. Note that this sharply contrasts previous works [JHY⁺14, YHJ⁺14, YHMR16, CCM18, GGN18] for which the \mathcal{P} -BRDF formulations use a \mathcal{P} -NDF.

4.2 Patch-BRDF Evaluation

By replacing f with our multiple-scattering BRDF f_∞ perturbed by $\omega_p(u)$, we get

$$\begin{aligned} f_{\mathcal{P}}(\omega_o, \omega_i) &= \int_{\mathbb{R}^2} f_\infty(\omega_o, \omega_i, \omega_p(u)) k_{\mathcal{P}}(u) du \\ &= f_1(\omega_o, \omega_i, \mathcal{P}) + E_{2+}(\omega_o, \mathcal{P}) F_1(\omega_o, \omega_g) f_{2+}(\omega_o, \omega_i). \end{aligned} \quad (13)$$

The Fresnel term $F_1(\omega_o, \omega_g)$ and the energy-compensation BRDF $f_{2+}(\omega_o, \omega_i)$ can be taken out of the integral because these terms do not depend on u . The term $f_1(\omega_o, \omega_i, \mathcal{P})$ is the single scattering patch-BRDF, defined as

$$f_1(\omega_o, \omega_i, \mathcal{P}) = \int_{\mathbb{R}^2} f_1(\omega_o, \omega_i, \omega_p(u)) k_{\mathcal{P}}(u) du. \quad (14)$$

The term $E_{2+}(\omega_o, \mathcal{P})$ is the energy lost by the single-scattering BRDF f_1 in the footprint, and its expression is

$$\begin{aligned} E_{2+}(\omega_o, \mathcal{P}) &= \int_{\mathbb{R}^2} 1 - E_1(\omega_o, \omega_p(u)) k_{\mathcal{P}}(u) du \\ &= 1 - \int_{\mathbb{R}^2} E_1(\omega_o, \omega_p(u)) k_{\mathcal{P}}(u) du. \end{aligned} \quad (15)$$

We now develop how we calculate $f_1(\omega_o, \omega_i, \mathcal{P})$ and $E_{2+}(\omega_o, \mathcal{P})$.

4.2.1 Single-Scattering Patch BRDF Evaluation

Until now our model, based on Equation 14, implies an infinite set of local BRDFs. In practice, we use discrete local BRDFs and assume a finite set of m perturbed normals ω_p :

$$f_1(\omega_o, \omega_i, u) = \sum_j^m f_1(\omega_o, \omega_i, \omega_p(u_j)) k_j(u). \quad (16)$$

Seed points u_j separated by a distance h are uniformly distributed across the normal map. Each BRDF at position u_j is weighted by its distance to u to ensure smooth transitions between two adjacent local BRDFs. As our footprint $k_{\mathcal{P}}$ is Gaussian, we use a Gaussian weight k_j for the BRDF in order to have an analytic weighting of f_1 . The standard deviation of k_j is set to $\sigma_h = h/\sqrt{8 \log 2}$ as in the methods of Yan et al. and Chermain et al. [YHMR16, CCM18] to achieve an equal contribution of two adjacent BRDFs when the evaluation point u is halfway between the two (Figure 6).




Figure 6: Discrete set of Gaussian weighted BRDFs, and associated discrete set of normals $\omega_p(u_j)$.

Substituting the sum expression of f_1 (Equation 16) into the single-scattering patch-BRDF (Equation 14) gives

$$\begin{aligned}
 f_1(\omega_o, \omega_i, \mathcal{P}) &= \int_{\mathbb{R}^2} \sum_j^m f_1(\omega_o, \omega_i, \omega_p(u_j)) k_j(u) k_{\mathcal{P}}(u) du \\
 &= \sum_j^m f_1(\omega_o, \omega_i, \omega_p(u_j)) \int_{\mathbb{R}^2} k_j(u) k_{\mathcal{P}}(u) du \\
 &= \sum_j^m f_1(\omega_o, \omega_i, \omega_p(u_j)) W_j
 \end{aligned} \tag{17}$$

where the weight W_i is the analytic solution of the integral of the product of two 2D Gaussians, here the local weight k_j and the ray footprint $k_{\mathcal{P}}$.

Acceleration Hierarchy. As explained in Section 3, the local BRDF f_1 has two symmetric facets with a near-perfectly specular material f_m . This micro-BRDF is a delta Dirac function, but is modeled for practical reasons by a microfacet-based BRDF with very low roughness α_m . This implies that f_1 is different than zero only when the half vector $\omega_h = \frac{\omega_o + \omega_i}{\|\omega_o + \omega_i\|}$ approximately matches the perturbed normal ω_p or its symmetric ω_s , and is used to accelerate the sum evaluation of Equation 17 by pruning out non-contributing discrete BRDFs.

A 4D Bounding Volume Hierarchy (BVH) is built using the discrete set of m perturbed normals ω_p . Each normal has a position u_j , a slope \tilde{p}_j and a 4D bounding box defining the (u, \tilde{p}) domain where the perturbed facet of f_1 reflects a non-negligible amount of light, for a given couple of position u and half vector slope \tilde{h} . The 4D bounding box is $(u_j \pm 3\sigma_h, \tilde{p}_j \pm 3\sigma_m)$, where $\sigma_m = \alpha_m/\sqrt{2}$, assuming that the contribution is zero beyond 3σ . This is mostly similar to the hierarchies used in Yan et al. [YHMR16] and Chermain et al. [CCM18]. Because of the use of a symmetric facet, we have to intersect the 4D BVH with (u, \tilde{h}) as in the previous methods, but also with $(u, -\tilde{h})$ to account for the contribution of the ω_s micro-normals.

4.2.2 Energy Lost in the Patch

Evaluating the energy lost in the patch requires us to process all the perturbed normals ω_p falling into the ray footprint. For performance reasons, we assume that normals ω_p follow a Gaussian distribution D of normals ω_p having an anisotropic footprint roughness $(\alpha_x(\mathcal{P}), \alpha_y(\mathcal{P}))$. In other words, we change the integral domain from spatial positions u to perturbed normals ω_p , similarly to Heitz et al. [HNPN13]

$$\begin{aligned}
 E_{2+}(\omega_o, \mathcal{P}) &= 1 - \int_{\mathbb{R}^2} E_1(\omega_o, \omega_p(u)) k_{\mathcal{P}}(u) du \\
 &\approx 1 - \int_{\Omega} E_1(\omega_o, \omega_p) D(\omega_p, \alpha_x(\mathcal{P}), \alpha_y(\mathcal{P})) d\omega_p \\
 &\approx 1 - \left(\cos^2(\phi_o) \int_{\Omega} E_1(\theta_o, \omega_p) D(\omega_p, \alpha_x(\mathcal{P})) d\omega_p \right. \\
 &\quad \left. + \sin^2(\phi_o) \int_{\Omega} E_1(\theta_o, \omega_p) D(\omega_p, \alpha_y(\mathcal{P})) d\omega_p \right).
 \end{aligned} \tag{18}$$

The terms θ_o and ϕ_o are respectively the polar and azimuthal angles of observation direction ω_o . We use ϕ_o to weight the spherical integrals implying an isotropic distribution $D(\omega_p, \alpha)$ of normals ω_p . By doing so, we can

precompute integrals of the form $\int_{\Omega} E_1(\theta_o, \omega_p) D(\omega_p, \alpha) d\omega$, and store the values into a small 2D array (128×128 floats, Figure 7) for different couples of roughness α and polar angles θ_o . To calculate $\alpha_x(\mathcal{P})$ and $\alpha_y(\mathcal{P})$ at render time, we use the method described in LEADR mapping [DHI⁺13], based on the first and second moments of a slope map.




Figure 7: Let $D(\omega_p, \alpha)$ be a Gaussian normal distribution with isotropic roughness α over the normal map space ω_p . D approximates the accurate normal distribution within a ray footprint. We use it to precompute footprint-dependent first-bounce energies $\int_{\Omega} E_1(\theta_o, \omega_p) D(\omega_p, \alpha) d\omega$. The values are stored in a 2D array, mapping energy quantities to (α, θ_o) value pairs.

4.3 Importance Sampling the Patch-BRDF

Importance sampling is used to efficiently solve the patch scattering equation:

$$\begin{aligned} L(\omega_o) &= \int_{\Omega} L(\omega_i) f_{\mathcal{P}}(\omega_o, \omega_i) \langle \omega_i, \omega_g \rangle d\omega_i \\ &= \int_{\Omega} L(\omega_i) \int_{\mathbb{R}^2} f(\omega_o, \omega_i, \omega_p(u)) \langle \omega_i, \omega_g \rangle k_{\mathcal{P}}(u) du d\omega_i \end{aligned} \quad (19)$$

where $L(\omega_i)$ is the incident radiance from direction ω_i . The sampling of a patch-BRDF $f_{\mathcal{P}}(\omega_o, \omega_i)$ is done by sampling (u, ω_i) with the following probability density function (PDF)

$$\text{PDF}(\omega_i, u) = f(\omega_o, \omega_i, \omega_p(u)) \langle \omega_i, \omega_g \rangle k_{\mathcal{P}}(u). \quad (20)$$

The marginalized random variable u is first chosen, followed by the conditional random variable $\omega_i|u$. For the first step, the marginalized PDF is

$$\begin{aligned} \text{PDF}(-, u) &= \int_{\Omega} f(\omega_o, \omega_i, \omega_p(u)) \langle \omega_i, \omega_g \rangle d\omega_i k_{\mathcal{P}}(u) \\ &= E(\omega_o, \omega_p) k_{\mathcal{P}}(u) \end{aligned} \quad (21)$$

where E is the energy term. This simplifies to $k_{\mathcal{P}}(u)$ only if the local BRDF f is 100% energy conserving, that is, if $\int_{\Omega} f(\omega_o, \omega_i, \omega_p(u)) \langle \omega_i, \omega_g \rangle d\omega_i = 1$ for all $u \in \mathbb{R}^2$. Because our local BRDF honors this condition, the sampling first consists in sampling the ray footprint $k_{\mathcal{P}}(u)$. During the second step, the conditional PDF is $f(\omega_o, \omega_i, \omega_p(u)) \langle \omega_i, \omega_g \rangle$ at the sampled position u .

To our knowledge, no other method has such a simple and efficient sampling procedure. Other methods either disregard the observation direction when sampling the patch [YHJ⁺14, JHY⁺14, YHMR16, CCM18], or need to walk through all normals if they want to do it [AK16] (Figure 8), which is inefficient.

5 Results and Discussion

In this section, we evaluate our method and compare it with state-of-the-art methods. All results have been rendered with the PBRT-v3 [PJH16] renderer on a dual Intel Xeon Gold 6138 CPU computer. Figures 2 and 9 do not use pixel footprints, while all other figures in the paper do. After the first light bounce, we never compute any footprint and rely on a smooth BRDF using the overall roughness and Fresnel term of the normal map, as in Jakob et al. [JHY⁺14]. We only use forward path tracing since our BRDF is asymmetric. All figures use $h = 1$ (Section 4.2.1) except Figures 14 and 15 (scratches part only).

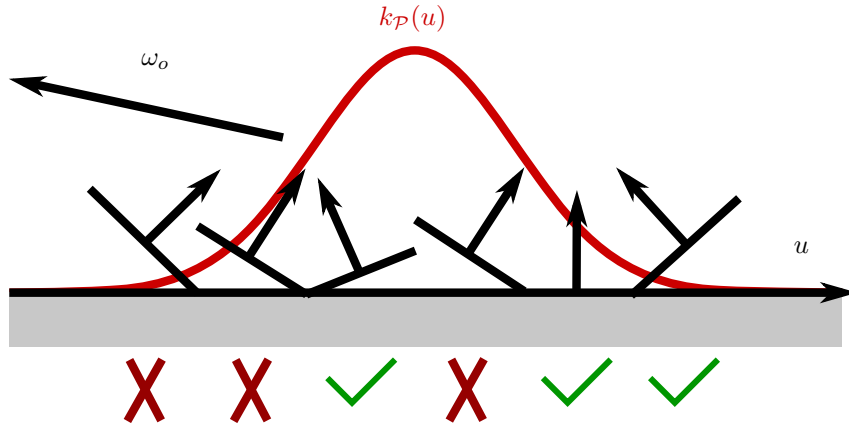


Figure 8: Classic normal-mapped surface. Importance sampling visible normals can only be done if all normals in the footprint are examined one by one. Our method does not have to go through this.

5.1 Validation of the Local BRDF

In this section, we validate our multiple-scattering local BRDF f_∞ . To do so, we compare renderings and white furnace tests of classic and microfacet-based normal mapped surfaces. In Figure 9, the normal map is a Beckmann surface ($\alpha = 0.5$). We compare f_1 and f_∞ to classic normal mapping. Black fringes are visible on the surface with this latter method. The single-scattering BRDF f_1 partially eliminates these artifacts while the multiple-scattering BRDF f_∞ makes them completely disappear. Furthermore, the BRDF f_∞ passes the white furnace test. We use the following parameters: f_m is a classic microfacet-based BRDF with a Beckmann NDF ($\alpha_m = 0.02$). Its Fresnel term F_1 models a gold metallic material. The energy-compensation BRDF f_{2+} is lambertian (diffuse), with an albedo equal to F_1 . Notice how the diffuse compensation and coloring incurred by F_1 is natural and plausibly simulates multiple-scattering on the micro-surface. For all other results, we use $\alpha_m = 0.02$ and a diffuse f_{2+} , unless otherwise stated. In Figure 9, bottom rows, the normal map is repeated 20 times more than in the top rows, with other parameters unchanged. The back-facing normals associated with classic normal mapping aggregate, leading to an overall dark appearance (Figure 9, bottom left). Our local multiple-scattering BRDF f_∞ does not have this issue (Figure 9, bottom right).

5.2 Patch-BRDF Results

Figure 10 shows different materials that our method can model. An animation is provided in the supplemental materials to evaluate the temporal coherence offered by our method. Brushed metals with dents can be obtained using a Beckmann surface instance with very high anisotropy ($\alpha_x = 0.5$, $\alpha_y = 0.01$, see the bigger spoon). The knife uses a texture with scratches in only one direction, while the dessert spoon uses random directions, reproducing a completely different appearance. The forks have a specular f_{2+} matching a plastic appearance. The green fork uses a Beckmann normal map with high isotropic roughness $\alpha = 1$, and the purple one uses a Voronoi texture with $\alpha = 1.41$. Finally, the silver spoon has a Beckmann surface with $\alpha = 0.1$.

Memory footprint. For the normal map representation and the roughness estimation in the footprint, we need two mipmapped textures, one for the normal map slopes and another one for the second moments. All our results use 2048×2048 mipmapped textures, each requiring 96 MiB. The lost energy approximation of Section 4.2.2 takes up 128×128 floats (64 KiB). The acceleration hierarchy of Section 4.2 uses a similar size than the one used by the methods it is based on [YHMR16, CCM18] for the same value of h (512 MiB for $h = 1$, 2048 MiB for $h = 0.5$).

Impact of the Patch-Energy Approximation Our \mathcal{P} -BRDF relies on an approximation of the energy lost in the patch (Section 4.2.2) to speed up evaluations. Figure 11 compares this approximation with the accurate quantity, directly calculated with a discrete weighted sum of f_∞ across the whole ray footprint and fully passing the white furnace test. Our approximation method provides a $2\times$ speed up, except for the scratched surface which includes a lot of planar normals. While the white furnace tests don't come out perfect, they are significantly better than those of previous methods, and we offer a plausible appearance, without any black holes on the surface, unsaturated colors or overly dark general appearance.

Importance Sampling. The weights of our importance sampling procedure are never greater than one, guaranteeing renderings without unnaturally bright pixels (Figure 12, right). On the contrary, previous methods




Figure 9: Comparison between classic normal mapping (left) and either our single (middle) or multiple (right) scattering microfacet-based normal mapped surface. Our multiple-scattering BRDF f_∞ is 100% energy conserving and removes the black holes (top rows) and the overall dark appearance (bottom rows) on the surface. In the bottom rows, the normal map repetition is twentyfold, leading to a coarse Beckman specular lobe. We use 131,072 spp to render the majority of glints and achieve noise-free renderings.

can give arbitrary large sampling weights to compensate for zero-radiance samples (Figure 12, left). When the surface is partially shadowed as in Figure 12, no tiny reflections should be visible. This is the case for our model and not for the previous method.

Choosing the Energy-Compensation BRDF. Most figures in the paper use a diffuse energy-compensation BRDF f_{2+} to model multiple-scattering. This choice gives results comparable to the previous methods, where their overall, dark appearance can be seen as a diffuse material with a low albedo. Figure 13 shows the difference between the use of a diffuse f_{2+} and a specular f_{2+} ($\alpha = 0.02$). Metallic paint can be modeled with a near-perfectly specular f_{2+} , giving the impression that metallic flakes are coated with a transparent varnish. Using a lambertian f_{2+} suits materials with rough appearance.

Comparison with Methods using Gaussian Mixtures. We compare the renderings of glittery surfaces using two different methods in Figure 1 (and supplemental video). The methods of Chermain et al. and Yan et al. suffer from backfacing normals, related black fringes and dark overall appearance. The outcome of the white furnace test is significantly better with our method. Figure 1 uses the flat variant of the Gaussian mixtures [YHMR16], that is best suited to represent glittery surfaces (constant micro-BRDF roughness $\alpha_m = 0.02$).

Figure 14 (and supplemental video) shows an object with a structured micro-surface rendered using different methods. Gaussian mixtures-based methods exhibit unnatural dark or bright scratches. They can be observed in the right area of the orb, due to back-facing normals or incorrect \mathcal{P} -BRDF normalization. Structured micro-surface renderings use $h = 0.5$ (Figures 14 and 15) for our method for a reason that we now explain.

The models of Chermain et al. and Yan et al. reconstruct an interpolated normal map with NDF mixtures, where each NDF is a Gaussian with a variance that best fit locally the curvature of the input surface. Per-Gaussian reconstruction kernels k_j further interpolate BRDF values between one another at render time, resulting in a surface model fundamentally continuous, even when the surface has relatively high frequency contents locally. Our model resembles the Jacobian-free, flat element variant (with fixed Gaussian roughness) of Yan et al., where normals are not interpolated continuously. It defines two normals for a given surface position instead of one. Equation 8 of our model forbids the use of high local roughness for the micro-BRDF. Our method therefore solely relies on the k_j reconstruction filter to interpolate neighboring BRDF values at render time. Consequently, local high frequency contents are approximated less accurately than when using mixtures of Gaussians with varying roughness. Scratches appear sharper with our model, less smoothed down. This issue can be alleviated by reducing the h sampling distance (increasing the density of V-cavities). On the upside, the intrinsic shadowing of V-cavities naturally weights the intensities of the glints, favoring glints continuity when animating. Note that high roughness severely impacts performance of previous methods, causing BVH node extents to be extremely large. For example, the orb of Figure 14 rendered with a fixed micro-BRDF roughness $\alpha = 0.02$ (flat elements variant) requires about 90 s with the method of Yan et al. and Chermain et al. Decreasing the value of h does of course affect performance, but to a lesser extent.

The quality of the lost energy approximation (Section 4.2.2) depends on the size of the ray footprint. Large ray footprints lead to favorable Gaussian approximations of normals distributions (Equation 18). For close-up views, smaller footprints cause lost energies to be underestimated, leading to poor white furnace tests (Figure 15). At low scales, roughness is not a reliable indicator for the lost energy estimation. Other methods exhibit even poorer results. Yan et al. normalize their footprint-dependant BRDF in a footprint-independent way. Chermain et al. suffer from the lack of multiple-scattering.

6 Conclusion and Future Works

In this work, we have presented the first practical multiple-scattering glint integrator based on normal maps. Our method leverages a local energy preserving microfacet-based BRDF that fakes normal perturbations. Our patch-BRDF can be used in Monte Carlo forward path tracing using multiple importance sampling. A large collection of material appearances can be reproduced through the use of different texture maps. Compared to the previous methods based on classic normal mapping, we obtain renderings that are free of artifacts such as black holes or a too dark overall appearance.

Currently available glint integrators only model opaque surfaces. To our knowledge, no method addresses the problem of rendering rough or scratched glass. Transparent materials are more sensitive to energy loss. We hope that our new approach will make solving this problem more accessible.

References

- [AK16] Asen Atanasov and Vladimir Koylazov. A practical stochastic algorithm for rendering mirror-like flakes. In *ACM SIGGRAPH 2016 Talks*, SIGGRAPH '16, pages 67:1–67:2, New York, NY, USA, 2016. ACM.
- [BYRN17] Laurent Belcour, Ling-Qi Yan, Ravi Ramamoorthi, and Derek Nowrouzezahrai. Antialiasing complex global illumination effects in path-space. *ACM Trans. Graph.*, 36(1):9:1–9:13, January 2017.
- [CCM18] Xavier Chermain, Frédéric Claux, and Stéphane Mérillou. A microfacet-based brdf for the accurate and efficient rendering of high-definition specular normal maps. *The Visual Computer*, Oct 2018.
- [CK17] Alejandro Conty and Christopher Kulla. Revisiting physically based shading at Imageworks. In *ACM SIGGRAPH 2017 Courses*, SIGGRAPH '17, pages 7:1–7:8, New York, NY, USA, 2017. ACM.
- [CT81] Robert L. Cook and Kenneth E. Torrance. A reflectance model for computer graphics. *SIGGRAPH Comput. Graph.*, 15(3):307–316, August 1981.
- [DHI⁺13] Jonathan Dupuy, Eric Heitz, Jean-Claude Iehl, Pierre Poulin, Fabrice Neyret, and Victor Ostromoukhov. Linear efficient antialiased displacement and reflectance mapping. *ACM Trans. Graph.*, 32(6):211:1–211:11, November 2013.
- [FA19] Carmelo J. Fdez-Agüera. A multiple-scattering microfacet model for real-time image based lighting. *Journal of Computer Graphics Techniques (JCGT)*, 8(1):45–55, January 2019.
- [GGN18] Luis E. Gamboa, Jean-Philippe Guertin, and Derek Nowrouzezahrai. Scalable appearance filtering for complex lighting effects. *ACM Trans. Graph.*, 37(6):277:1–277:13, December 2018.
- [Hei14] Eric Heitz. Understanding the masking-shadowing function in microfacet-based brdfs. *Journal of Computer Graphics Techniques (JCGT)*, 3(2):48–107, June 2014.
- [HHdD16] Eric Heitz, Johannes Hanika, Eugene d’Eon, and Carsten Dachsbacher. Multiple-scattering microfacet bsdfs with the smith model. *ACM Trans. Graph.*, 35(4):58:1–58:14, July 2016.
- [Hil18] Stephen Hill. A multi-faceted exploration (part 2), jun 2018.
- [HNPN13] Eric Heitz, Derek Nowrouzezahrai, Pierre Poulin, and Fabrice Neyret. Filtering color mapped textures and surfaces. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, I3D '13, pages 129–136, New York, NY, USA, 2013. ACM.
- [Ige99] Homan Igehy. Tracing ray differentials. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '99, pages 179–186, New York, NY, USA, 1999. ACM Press/Addison-Wesley Publishing Co.
- [JHY⁺14] Wenzel Jakob, Miloš Hašan, Ling-Qi Yan, Jason Lawrence, Ravi Ramamoorthi, and Steve Marschner. Discrete stochastic microfacet models. *ACM Trans. Graph.*, 33(4):115:1–115:10, July 2014.
- [LJJ⁺18] Joo Ho Lee, Adrian Jarabo, Daniel S. Jeon, Diego Gutierrez, and Min H. Kim. Practical multiple scattering for rough surfaces. *ACM Trans. Graph.*, 37(6):275:1–275:12, December 2018.
- [PJH16] Matt Pharr, Wenzel Jakob, and Greg Humphreys. *Physically Based Rendering: From Theory to Implementation*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 3rd edition, 2016.
- [RGB16] Boris Raymond, Gaël Guennebaud, and Pascal Barla. Multi-scale rendering of scratched materials using a structured sv-brdf model. *ACM Trans. Graph.*, 35(4):57:1–57:11, July 2016.
- [SHHD17] Vincent Schüssler, Eric Heitz, Johannes Hanika, and Carsten Dachsbacher. Microfacet-based normal mapping for robust monte carlo path tracing. *ACM Trans. Graph.*, 36(6):205:1–205:12, November 2017.
- [SPSH19] Dorian Saint-Pierre, Lionel Simonot, and Mathieu Hébert. Reflectance computation for a specular only v-cavity. In Shoji Tominaga, Raimondo Schettini, Alain Trémeau, and Takahiko Horiuchi, editors, *Computational Color Imaging*, pages 289–303, Cham, 2019. Springer International Publishing.
- [SW01] Frank Suykens and Yves D. Willems. Path differentials and applications. In Steven J. Gortler and Karol Myszkowski, editors, *Rendering Techniques 2001*, pages 257–268, Vienna, 2001. Springer Vienna.

- [Tur19] Emmanuel Turquin. Practical multiple scattering compensation for microfacet models, 2019.
- [WB16] Beibei Wang and How Bowles. A robust and flexible real-time sparkle effect. In *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations*, EGSR '16, pages 49–54, Goslar Germany, Germany, 2016. Eurographics Association.
- [WWH18] Beibei Wang, Lu Wang, and Nicolas Holzschuch. Fast global illumination with discrete stochastic microfacets using a filterable model. *Computer Graphics Forum*, 37(7):55–64, 2018.
- [XH18] Feng Xie and Pat Hanrahan. Multiple scattering from distributions of specular v-grooves. *ACM Trans. Graph.*, 37(6):276:1–276:14, December 2018.
- [YHJ⁺14] Ling-Qi Yan, Miloš Hašan, Wenzel Jakob, Jason Lawrence, Steve Marschner, and Ravi Ramamoorthi. Rendering glints on high-resolution normal-mapped specular surfaces. *ACM Trans. Graph.*, 33(4):116:1–116:9, July 2014.
- [YHMR16] Ling-Qi Yan, Miloš Hašan, Steve Marschner, and Ravi Ramamoorthi. Position-normal distributions for efficient rendering of specular microstructure. *ACM Trans. Graph.*, 35(4):56:1–56:9, July 2016.
- [ZK16] Tobias Zirr and Anton S. Kaplanyan. Real-time rendering of procedural multiscale materials. In *Proceedings of the 20th ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, I3D '16, pages 139–148, New York, NY, USA, 2016. ACM.




Figure 10: Cutlery scene rendered at 1024 spp and exhibiting a wide range of materials that our method can handle. From left to the right: brushed and scratched metals, green coated rough plastic, machined metal, rough metal and finally purple coated glittery plastic.

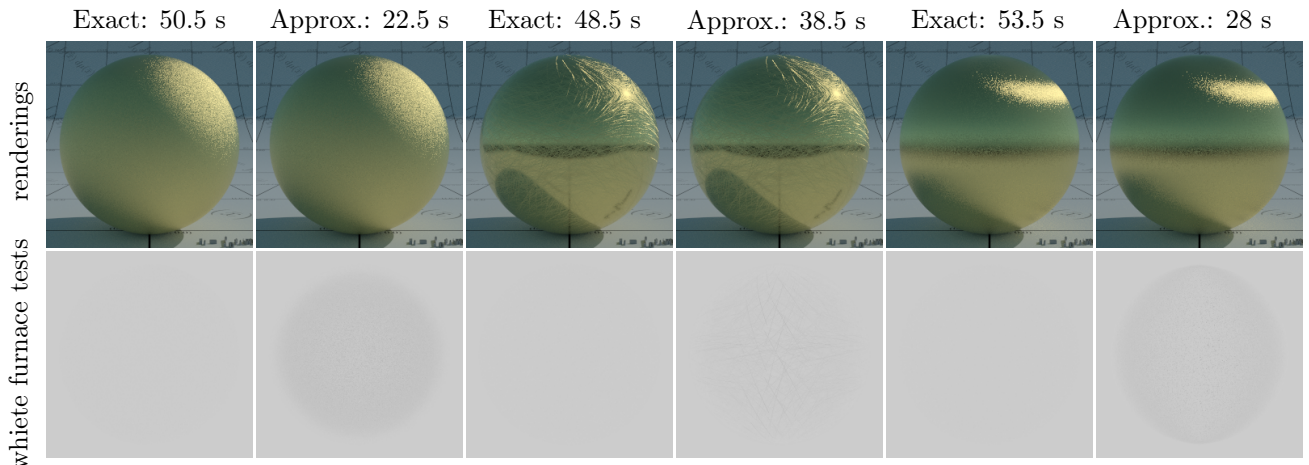


Figure 11: Renderings and white furnace tests of our \mathcal{P} -BRDF using the accurate and approximated energy term, for different normal maps. From left to right: isotropic Beckman surface with $\alpha = 0.5$, scratched surface and anisotropic Beckmann surface with $\alpha_x = 0.5$ and $\alpha_y = 0.1$. Our exact model passes the white furnace test. Our approximated \mathcal{P} -BRDF is not perfect, but has good energy properties, better performance and gives results similar to the reference.

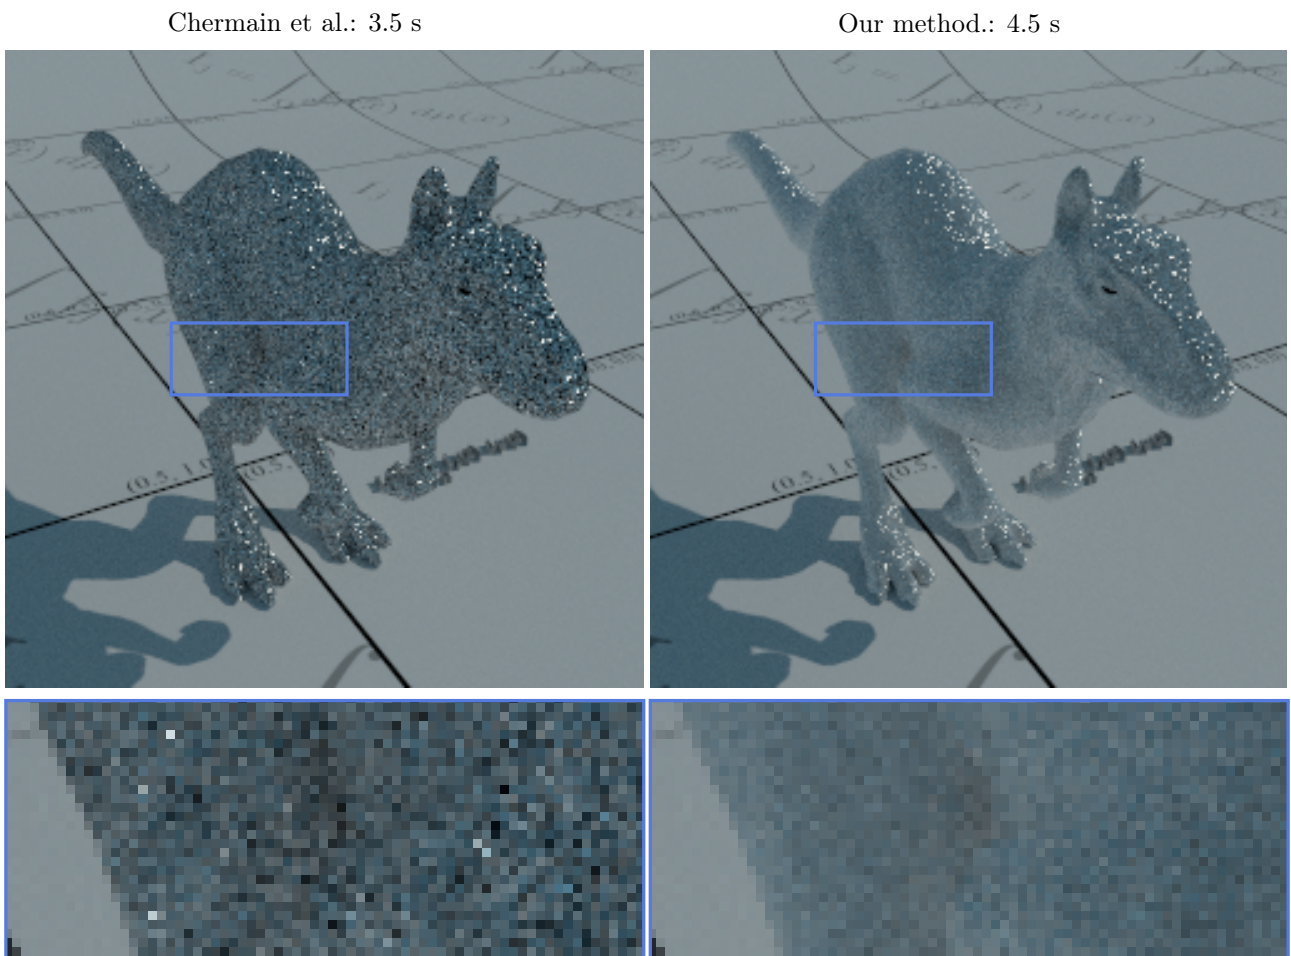


Figure 12: A killeroo rendered at 256 samples per pixel. Fireflies are visible when the \mathcal{P} -BRDF importance sampling does not use the observation direction (left). Our method does and generates no firefly.




Figure 13: A killeroo rendered at 256 samples per pixel with our \mathcal{P} -BRDF using either a diffuse (left) or a near-perfectly specular (right) energy-compensation BRDF f_{2+} , giving different material appearance.

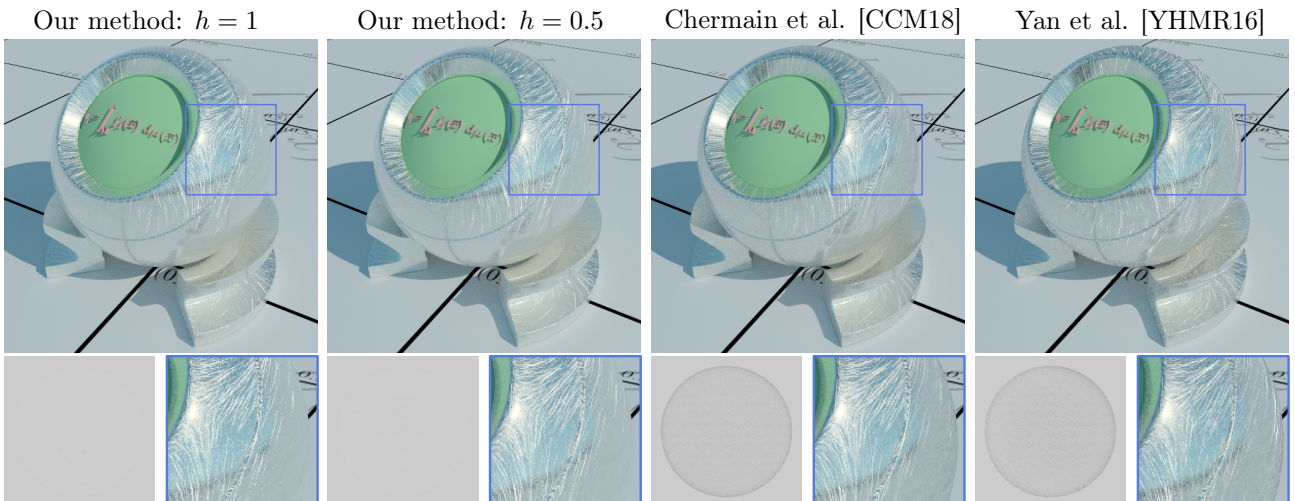


Figure 14: Scratched metallic orb renderings and white furnace tests using our method and state-of-the art works. Note that our multiple-scattering glint integrator does not darken the silhouette of the object or the scratches the way other methods do. Rendering times from left to right: 129.5 s, 202.5 s, 296.5 s and 258.5 s.

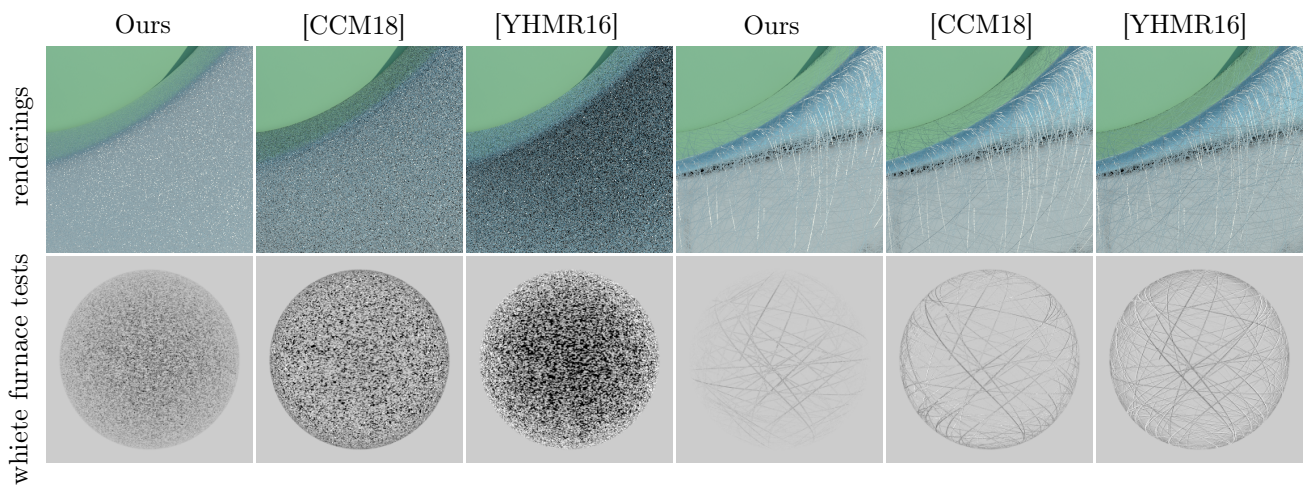


Figure 15: For close-up views, our patch energy approximation gives underestimated values, leading to poor white furnace tests, though better than other methods. Furthermore, our renderings are more realistic. Rendering times from left to right: 84.5 s, 65.5 s, 52.5 s, 194.5 s, 407.6 s and 398.6 s.