



HAL
open science

Transfer matrix based layered materials rendering

Joël Randrianandrasana, Patrick Callet, Laurent Lucas

► **To cite this version:**

Joël Randrianandrasana, Patrick Callet, Laurent Lucas. Transfer matrix based layered materials rendering. ACM Transactions on Graphics, 2021, Proceedings of ACM SIGGRAPH 2021, 40 (4), Article No. 177, 16 p. 10.1145/3450626.3459859 . hal-03262831

HAL Id: hal-03262831

<https://hal.univ-reims.fr/hal-03262831v1>

Submitted on 21 Jun 2021

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.

Transfer matrix based layered materials rendering

JOËL RANDRIANANDRASANA, University of Reims Champagne-Ardenne, France and United Visual Researchers, France

PATRICK CALLET, MINES ParisTech, France

LAURENT LUCAS, University of Reims Champagne-Ardenne, France



Fig. 1. The BSDF model introduced in this paper addresses the efficient rendering of layered materials including rough interfaces and arbitrary scattering volumes based upon a transfer matrix approach. This scene depicts the ability of our model to reproduce the variety of appearances resulting from these material configurations.

A statistical multi-lobe approach was recently introduced in order to efficiently handle layered materials rendering as an alternative to expensive general-purpose approaches. However, this approach poorly supports scattering volumes as the method does not account for back-scattering and resorts to single scattering approximations. In this paper, we address these limitations with an efficient solution based upon a transfer matrix approach which leverages the properties of the Henyey-Greenstein phase function. Under this formalism, each scattering component of the stack is described through a lightweight matrix, layering operations are reduced to simple matrix products and the statistics of each BSDF lobe accounting for multiple scattering effects are obtained through matrix operators. Based on this representation, we leverage the versatility of the transfer matrix approach to efficiently handle forward and backward scattering which occurs in arbitrary layered materials. The resulting model enables the reproduction of a wide range of layered structures embedding scattering volumes of arbitrary depth, in constant computation time and with low variance.

Authors' addresses: Joël Randrianandrasana, rasolonjatovo-joel.randrianandrasana@etudiant.univ-reims.fr, University of Reims Champagne-Ardenne, Reims, France, 51100, United Visual Researchers, Paris, France; Patrick Callet, president@centrefrancaisdela couleur.fr, MINES ParisTech, Paris, France, 75006; Laurent Lucas, laurent.lucas@univ-reims.fr, University of Reims Champagne-Ardenne, Reims, France, 51100.

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

© 2021 Association for Computing Machinery.

0730-0301/2021/8-ART177 \$15.00

<https://doi.org/10.1145/3450626.3459859>

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

Additional Key Words and Phrases: BRDF, Layered materials, Transfer matrix

ACM Reference Format:

Joël Randrianandrasana, Patrick Callet, and Laurent Lucas. 2021. Transfer matrix based layered materials rendering. *ACM Trans. Graph.* 40, 4, Article 177 (August 2021), 16 pages. <https://doi.org/10.1145/3450626.3459859>

1 INTRODUCTION

Realistic simulation of light-matter scattering plays a crucial role in many rendering applications. Whether targeting entertainment sources such as video games and animation movies, or addressing industrial needs such as architectural and product prototyping, high-quality appearance models are core ingredients of successful digital experiences.

In the realm of surface materials, a wide range of everyday life objects fall in the category of *layered materials*. This class of materials consist of a superposition of one or several layers of varying composition on top of a substrate. Typical examples include all various kinds of surface paints and finishes applied to human-made objects (e.g., clear and colored coats, metallic paints, varnished woods, glazed ceramics). An important subset of layered materials consists of materials undergoing natural or anthropogenic phenomena referred to as *weathering* or *aging*. For instance, the exposition of metals to their surrounding environment over a long period of time usually implies the formation of a patina layer due to corrosion. Additionally, surface appearances may significantly change within

short periods of time with the accumulation of dust or atmospheric pollutants.

In general, such layer combinations give rise to visually rich and complex appearance characteristics not exhibited by the individual components when observed separately. Unfortunately, this visual diversity dramatically increases the difficulty of designing scattering models that are simultaneously general, accurate, and efficient.

Researchers generally rely on multiple hypotheses in order to simplify the derivation of the BSDFs of layered materials. For instance, a common framework is the *plane-parallel approximation* which assumes that the material can be locally described as a stack of layers whose optical properties are independent of the lateral position. Moreover, the thickness of these layers is usually assumed to be much larger than the wavelength of light, making it possible to ignore interference effects. In practice, these simplifications are reasonable for a wide range of layered materials and enable a convincing reproduction of their appearance.

Based on these assumptions, highly accurate solutions have been proposed in recent years. However, these solutions remain impractical for low-time budget rendering due to expensive per-material precomputations [Jakob et al. 2014; Zeltner and Jakob 2018] or because of their stochastic nature [Guo et al. 2018]. As an alternative, more efficient methods usually make a compromise between accuracy and performance to meet these requirements [Elek 2010]. To this date, the most advanced method in this category is the recent multi-lobe approach of Belcour [2018]. This model relies on a statistical framework to estimate low-order moments, namely *energy*, *mean*, and *variance* of multiple scattering BSDF lobes, and extends the well-known *adding equations* for this framework. His approach is not only efficient, but also supports arbitrarily layered interfaces as well as spatially varying appearances. Unfortunately, scattering volumes are poorly supported, as the method resorts to a *single scattering approximation*. In practice, this approximation incurs severe energy loss, resulting in an inconsistent dark appearance with increasing volume scattering (Fig. 2). This represents a serious impediment, as many of the aforementioned natural phenomena involve multiple scattering volume effects. Moreover, many weathering effects, such as dust accumulation, typically entail significant *back-scattering* not accounted for by the methodology.

The work presented in this paper addresses the aforementioned issues by introducing an efficient solution based upon a *transfer matrix* model. While our approach is inspired by the work of Belcour [2018], the main difference lies in the underlying mathematical framework: We use the Henyey-Greenstein phase function [1941] as an intermediate representation to compute the shapes of the outgoing lobes. This crucial difference allows us to compute the energies and the shapes of multiple scattering lobes within a unified transfer matrix treatment.

The transfer matrix is widely used in many fields of physics involving layered systems [Sánchez-Soto et al. 2012]. Under this formalism, the optical properties of the layers are expressed as lightweight matrices, typically of size 2×2 or 4×4 , and layering operations are reduced to simple matrix products. The key benefits of the transfer matrix are twofold: First, it provides a compact yet versatile formalism to express and resolve multiple scattering problems as already demonstrated in the optics literature. Second, resorting to the



Fig. 2. The model of Belcour [2018] suffers from severe energy loss with scattering volumes due to its single-scattering approximation and its lack of support for back-scattering. Our transfer matrix approach overcomes these issues and provides results close to the ground-truth computed with the stochastic approach of Guo et al. [2018]. In these examples, a unit depth non-absorbing medium with scattering coefficient $\sigma_s = 0.755$ lies on top of a smooth gold substrate. Henyey-Greenstein asymmetry parameters are respectively set to $g = 0.9$ and $g = 0$ in the top and bottom row.

linear algebra of the transfer matrix allows getting rid of expensive *doubling* operations in favor of efficient closed-form expressions. We show how to jointly leverage these beneficial properties to efficiently approximate the forward and backward light transport of layered materials through this work. Our approach supports homogeneous participating media of arbitrary depth, while also ensuring energy preservation.

We summarize our contributions as follows:

- We introduce the classical two-flux transfer matrix as a convenient alternative to the adding equations in the context of smooth parallel interfaces (Section 3). Then, we generalize the matrix formalism to rough interfaces by revisiting Belcour’s [2018] statistical approach from a transfer matrix modeling perspective. To this end, we review some fundamental properties of the Henyey-Greenstein phase function (Section 4) and show how to leverage these properties to compute the energies and the shapes of the BSDF lobes within a unified transfer matrix formalism (Section 5).
- Then, we leverage the versatility of the transfer matrix approach in order to efficiently handle the forward and backward light transport occurring in layered materials with scattering volume layers. To this end, we introduce the six-flux transfer matrix and derive its specialized forms for interface components and homogeneous participating media of arbitrary depth (Section 6).
- Finally, we propose an improvement to the BSDF model introduced in the previous work to ensure energy preservation (Section 6.6).

2 RELATED WORK

2.1 Specialized layered models

The last few decades have seen an important number of scattering models targeting specific layered materials. Examples include the

model for metallic car paints proposed by Ergun et al. [2016], the approach of Meneveaux et al. [2017] addressing rough Lambertian substrates interfaced with a dielectric coat, but also extend to many organic materials whose internal structures consist of complex stratifications [Baranoski and Rokne 2001; Stam 2001]. Weathering effects are a topic of broad interest in the computer graphics community. For instance, Dorsey and Hanrahan [1996] address the natural effect of corrosion by modeling the constitutive layers of metallic patinas and resort to the *Kubelka-Munk* theory [1948] to reconstitute the appearance of the layers. The natural effect of dust accumulation has also been addressed in the work of Hsu and Wong [1995], as well as Guo and Pan [2014], and Gu et al. [2007] proposed an approach to reproduce the appearance of contaminated transparent surfaces. Despite their practicality, all these models suffer from an inherent lack of extensibility since they were tailored for specific layer configurations.

2.2 General layered models

Jakob et al. [2014] introduced the first general solution for arbitrarily layered materials, supporting interfaces of arbitrary roughness and scattering volumes. Their work was extended by Zeltner and Jakob [2018] to handle anisotropic interfaces. The framework of their work relies on zenithal discretizations of the radiance field combined with azimuthal frequency mode expansions. Based on this representation, the approach computes the light transport due to multiple scattering via the adding-doubling method [Van de Hulst 1980]. While this approach supports a rich diversity of layered material and provides high levels of accuracy, it suffers from expensive per-material pre-computation. Typically, hundreds to thousands of frequency modes might be required to capture all appearance features. Consequently, this model remains impractical for spatially varying structures such as weathered surfaces. Recently, Guo et al. [2018] proposed a general solution based on a Monte Carlo approach, free from per-material pre-computation, and supporting rough anisotropic interfaces. To achieve this, the authors use a path integral formulation explicitly designed for the plane-parallel approximation along with efficient estimators. Their work was recently extended by Xia et al. [2019] and Gamboa et al. [2020], who propose new variance reduction techniques. Unfortunately, the application of these methods in the context of low-budget rendering remains difficult on account of their stochastic nature.

2.3 Efficient layered models

When the rendering time is a strong constraint, approximate models making a compromise between accuracy and computing times are usually preferred. For instance, Hanrahan and Krueger [1993] proposed to approximate the effects of internal volume layers based on single scattering approximations. Weidlich and Wilkie [2007] introduced a solution for layered rough interfaces based on multiple assumptions about the layer structure. Elek [2010] subsequently adapted their work for real-time needs. Guo et al. [2016] proposed to use the von Mises-Fisher (vMF) distribution as a microfacet normal distribution function to approximate light-matter interactions with vMF convolution products. These approaches are computationally efficient but, unfortunately, lack support for volume scattering.

The closest work to our method is due to Belcour [2018], who introduced an efficient multi-lobe approach accounting for multiple scattering between rough interfaces. Detailed analysis of the models proposed by Belcour [2018] and Weidlich and Wilkie [2007] have been recently carried by Bati et al. [2019]. Extensions of Belcour's framework have been recently proposed by Yamaguchi et al. [2019] and Weier and Belcour [2020] for anisotropic interfaces.

Our work extends the multi-lobe approach of Belcour [2018] in order to provide an efficient support for volume layers. We introduce an efficient transfer matrix-based approach, handling backward and forward scattering in homogeneous volume layers of arbitrary depth, at constant cost. The framework introduced in this paper can be seen as an approximate version of Jakob et al. [2014] designed to track arbitrary appearance features, which we summarize as lightweight Henyey-Greenstein lobes, in contrast to their expensive angular / Fourier mode matrix representation.

2.4 The transfer matrix

The transfer matrix finds its roots in the pioneering work of Abelès [1948]. It has since been successfully used in various fields including quantum mechanics, acoustics, geometrical optics [Sánchez-Soto et al. 2012] and color reproduction applications [Mazauric et al. 2014]. In their most popular form, transfer matrices are 2×2 matrices in so-called *two-flux transfer* models and rely on strong assumptions about the geometry of the radiation traveling through the system. For instance, Abelès [1948] assumed the collimated propagation of electromagnetic waves between smooth parallel interfaces to obtain reflection and transmission coefficients accounting for interference effects. The matrix formalism introduced by Abelès has been used in computer graphics to reproduce the visual effects of interferences by Icart and Arques [2000] and more recently by Benamira and Pattanaik [2020] to simulate anti-reflective coatings.

As another remarkable result, a transfer matrix calculus directly leads to the Kubelka-Munk equations [1948]. In this case, the system is modeled as a homogeneous scattering volume of arbitrary depth, and the light distribution is assumed to be perfectly diffuse everywhere in the layer. This model is especially efficient with highly scattering media, which explains its wide adoption in the pigment and printing industries, computer graphics [Callet 1996; Dorsey and Hanrahan 1996], and material synthesis applications [Hašan et al. 2010]. Alternatively, *three-flux* and *four-flux* models decompose the total flux into collimated and diffuse components. For instance, four-flux models have been proposed to address layered systems mixing smooth interfaces and scattering volumes [Gali et al. 2017; Slovick et al. 2017], and stacks mixing smooth and rough interfaces on top of a substrate [Simonot et al. 2016]. However, considering only diffuse and/or collimated flux represents a serious limitation, as virtually all materials entail light distributions falling somewhere between these two limiting cases. As a consequence, these models fail to capture important features such as blurring effects due to subsurface scattering.

In contrast, our transfer matrix approach avoids these rigid separations and enables arbitrary changes of the angular flux distributions. As a result, our approach subsumes the previous ideal models and enables a wide range of layer configurations mixing smooth and




Fig. 3. Light transport of a collimated beam between two smooth parallel interfaces.

rough interfaces, and volumes with arbitrary scattering properties, in a unified treatment.

3 BACKGROUND

This section reviews the light transport of a ray of light illuminating a glass plate, as introduced in the pioneering work of Stokes [1860]. We introduce the transfer matrix formalism, then we move to the more complicated case of rough stacks. In this context, we review the statistical approach of Belcour [2018]. We discuss the applicability limitation of this framework to scattering volumes and we explain our motivation for reformulating it in a transfer matrix compliant form.

3.1 The glass plate problem

Each medium is labeled with index $m = 0, 1, \dots$ according to its depth in the stack and is assigned a real refractive index η_m . A ray of light arriving from medium 0 impinges on the upper interface and gives rise to a reflected ray and a transmitted ray inside the dielectric medium. Due to the same mechanism, this refracted ray may in turn trigger a sequence of internal bounces inside the plate, letting each time a fraction of energy escape from either side of the medium (Fig. 3). In this simplest case, all the rays traveling in the structure are governed by Snell's laws, and the fraction of reflected and transmitted energies between media i and j directly relate to Fresnel equations. We respectively denote these quantities r_{ij} and t_{ij} and refer to them as *transfer factors*. Given a unit amount of incident illumination, the transfer factors r_{02} and t_{02} from the slab are obtained by summing the contributions from either side of the medium:

$$r_{02} = r_{01} + \sum_{k=0}^{\infty} t_{01} r_{12} (r_{10} r_{12})^k t_{10} = r_{01} + \frac{t_{01} r_{12} t_{10}}{1 - r_{10} r_{12}} \quad (1)$$

$$t_{02} = \sum_{k=0}^{\infty} t_{01} (r_{12} r_{10})^k t_{12} = \frac{t_{01} t_{12}}{1 - r_{10} r_{12}}. \quad (2)$$

A similar calculus leads to the transfer factors r_{20} and t_{20} of the slab in the opposite case of a ray arriving from medium 2. These equations have been known for a while in the optics community and form the building blocks of the adding equations in the radiative transfer field [Van de Hulst 1980].

In essence, the previous approach can be seen as a chronological treatment of the events that light undergoes in the structure. From the transfer matrix perspective, each scattering component of the




Fig. 4. The transfer matrix of a component relates the total upwelling and downwelling fluxes sitting on each side of the component.

system is regarded as a linear system relating the *total* incoming and outgoing flux located on a given side of the component to the total incoming and outgoing flux sitting at the opposite side. In our study case, the scattering components reduce to the interfaces between the different media and we denote i_k^+ (resp. i_k^-) the total downward (resp. upward) flux traveling in the medium k (Fig. 4). The flux balance at the boundaries of the first interface is given as:

$$i_1^+ = t_{01} i_0^+ + r_{10} i_1^-, \quad i_0^- = r_{01} i_0^+ + t_{10} i_1^- \quad (3)$$

which can equivalently be written under the matrix notation

$$\begin{bmatrix} i_0^+ \\ i_0^- \end{bmatrix} = \frac{1}{t_{01}} \begin{bmatrix} 1 & -r_{10} \\ r_{01} & t_{01} t_{10} - r_{01} r_{10} \end{bmatrix} \begin{bmatrix} i_1^+ \\ i_1^- \end{bmatrix} \quad (4)$$

with M_{01} denoting the transfer matrix of the interface separating media 0 and 1. By transitivity and assuming no volume effects in medium 1, it directly follows that $I_0 = M_{01} M_{12} I_2 = M_{02} I_2$ with M_{02} denoting the transfer matrix resulting from the ordered superposition of the two interfaces. Obviously, this relation also holds for any ordered sequence of scattering components, i.e. $I_i = M_{ij} I_j$ with $M_{ij} = M_{i+1} M_{i+2} \dots M_{j-1}$. Thus, complex layered systems can be easily modeled through straight matrix products from the knowledge of the transfer matrices of the individual components. Note that any transfer matrix has a structure similar to the matrix of Equation (4).

Given a transfer matrix M_{ij} , the downward r_{ij} and t_{ij} transfer factors of the corresponding component (or stack of components) can be obtained via the matrix operators

$$r_{ij} = R(M_{ij}) = \frac{(M_{ij})_{21}}{(M_{ij})_{11}} \quad (5)$$

and

$$t_{ij} = T(M_{ij}) = \frac{1}{(M_{ij})_{11}} \quad (6)$$

where $(M_{ij})_{xy}$ denotes the entry located at the x^{th} row and y^{th} column of the transfer matrix of the interface separating media i and j . Analogous operators exist for the upward r_{ji} and t_{ji} transfer factors. Also, one can easily verify that applying operators (5) and (6) to the transfer matrix M_{02} yields relations similar to the summation treatment (Equations (1) and (2)).




Fig. 5. In the case of rough interfaces, each light-matter interaction gives rise to complex angular light distributions in the structure (a). With Belcour’s [2018] statistical approach, the BSDF of the stack is approximated as a mixture of GGX lobes approximating the effects of each interface (b).

3.2 Rough stacks treatment with Belcour’s statistical framework

The main difficulty with rough interfaces is that each light-matter interaction gives rise to complex angular light distributions in the layered structure (Fig. 5(a)). The core idea of Belcour [2018] is to approximate complex light distributions due to interactions between rough interfaces with simpler ones with equivalent low-order statistical moments. To this end, the author introduces a set of *atomic operators* simulating the statistical effects of individual layer interactions, namely, reflection and refraction, on incoming light distributions. In practice, the approach focuses on the GGX microfacet based model [Walter et al. 2007] and only accounts for the three lowest order moments corresponding to the energy e of the distribution, its mean direction μ , and its variance σ^2 (denoted as σ for convenience). Moreover, the mean directions of the light distributions are assumed to obey Snell’s laws to simplify the analysis further.

Denoting $\phi_i = (e_i, \mu_i, \sigma_i)$ the statistics of an incoming light distribution, the statistics ϕ^R of the reflected distribution are approximated as

$$e^R = e_i \times \text{FGD}, \quad \mu^R = \text{reflect}(\mu_i), \quad \sigma^R = \sigma_i + f(\alpha) \quad (7)$$

where α denotes the roughness of the interface and $f(\alpha)$ maps GGX roughnesses to a space where light variance behaves additively between each light bounce. The total amount of reflected energy is approximated with the directional albedo of the surface evaluated at incident mean direction and is encoded in the FGD term:

$$\text{FGD}(\omega_i, \alpha, \eta, \kappa) = \int_{\Omega^+} \frac{F(\omega_i, \mathbf{h}, \eta, \kappa) G(\omega_i, \omega_o, \alpha) D(\mathbf{h}, \alpha)}{4(\omega_i \cdot \mathbf{n})} d\omega_o \quad (8)$$

with \mathbf{n} the geometric normal, \mathbf{h} the halfway vector of ω_i and ω_o , $\tilde{n} = \eta + i\kappa$ the complex refractive index of the interface, F the Fresnel term, G the shadowing-masking term, α the roughness parameter, and D the GGX normal distribution function. As no closed-form solution exists for this integral, the FGD term is precomputed in a 4D table over a predefined range of incidence elevations, roughnesses, and refractive indices.

For transmission, Belcour [2018] proposes to mimic a rough refraction with a mirrored rough reflection occurring at the opposite side of the interface. Following this methodology, the statistics ϕ^T

of the transmitted distribution are obtained as

$$e^T = e_i \times (1 - \text{FGD}), \quad \mu^T = \text{refract}(\mu_i, \eta), \quad \sigma^T = \frac{\sigma_i}{\eta} + f(s \times \alpha) \quad (9)$$

$$\text{with } s = \frac{1}{2} \left[1 + \eta \frac{\cos \theta_i}{\cos \theta_t} \right]$$

where s is a roughness scaling factor used to match the shape of the real transmitted lobe with the fake mirror reflection.

Outgoing light statistics resulting from arbitrary sequences of interactions in the structure are afterward approximated by successively applying these operators. For a ray of light impinging on a material consisting of two interfaces, the total amount of reflected energy can be obtained via the geometric series

$$e_{02}^R = r_{01} + \sum_{k=0}^{\infty} t_{01} r_{12} (r_{10} r_{12})^k t_{10} = r_{01} + \frac{t_{01} r_{12} t_{10}}{1 - r_{10} r_{12}} \quad (10)$$

with $r_{ij} = \text{FGD}$ and $t_{ij} = (1 - \text{FGD})$. In this form, the r_{ij} and t_{ij} terms do not account for total internal reflections between the media. The author corrects this behavior by introducing a TIR term depending on the incident elevation, the relative index of the two media, and the adjacent interface’s roughness. As this factor does not admit a closed-form solution, the latter is pre-computed and stored in a 3D LUT interpolated during the rendering to correct the reflection and transmission terms as $r_{ij} \rightarrow r_{ij} + (1 - \text{TIR}) \times t_{ij}$ and $t_{ij} \rightarrow t_{ij} \times \text{TIR}$.

The variance of the distribution is obtained by computing the weighted sum of the individual series terms variances. The unnormalized variance $\tilde{\sigma}_{02}^R$ of the reflected light can then be expressed as an arithmetico-geometric series whose analytic solution is given as

$$\tilde{\sigma}_{02}^R = r_{01} \sigma_{01}^R + \left[\frac{t_{01} r_{12} t_{10}}{1 - r_{10} r_{12}} \right] \times \left[\sigma_{10}^T + J_{10} \left(\sigma_{01}^T + \sigma_{12}^R + \left(\sigma_{12}^R + \sigma_{10}^R \right) \left(\frac{r_{10} r_{12}}{1 - r_{10} r_{12}} \right) \right) \right] \quad (11)$$

where J_{ij} denotes a variance scaling factor due to transmission. An analogous relation is obtained for the unnormalized variance $\tilde{\sigma}_{02}^T$ of the light transmitted through the two interfaces.

Based on these results, the BSDF of a stack composed of an arbitrary number of interfaces is modeled as a mixture of GGX lobes approximating the reflection due to each interface (the TR^+T paths) and a single lobe for transmission (Fig. 5(b)). The resulting BSDF is formally given as

$$\rho(\omega_i, \omega_o) = \sum_k e_k \rho(\omega_k, \omega_o, \sigma_k) \quad (12)$$

with

$$\omega_k = \text{reflect}(\mu_k)$$

$$\rho(\omega_k, \omega_o, \sigma_k) = \frac{G(\omega_k, \omega_o, \alpha_k) D(\mathbf{h}, \alpha_k)}{4(\omega_k \cdot \mathbf{n})(\omega_o \cdot \mathbf{n})}$$

$$\alpha_k = f^{-1}(\sigma_k).$$

We refer the reader to the original paper of Belcour [2018] for further details.

3.3 Motivation

While the statistical framework proposed by Belcour [2018] remains simple yet powerful, extending this approach to scattering volume layers remains very challenging. One of the main difficulties is that volume scattering generally involves both forward and backward flux and continuous interactions between them in the overall structure. In Belcour's terminology, this entails tracking at least two sets of statistics, formulating the associated *multi-flux* atomic operators, and deriving the corresponding energy and variance multi-flux adding equations. Another critical difficulty lies in the treatment of volumes of arbitrary depth. Scattering in volumes of arbitrary depth is typically achieved by repeating the adding equations starting from a very thin layer. This process is referred to as doubling and may incur expensive and non-constant computation costs as volume layer thicknesses generally vary on surfaces. Unfortunately, the statistical framework of Belcour [2018] is neither well suited to multi-flux treatment nor derivation of efficient methods for volumes of arbitrary depth.

On the other hand, the transfer matrix offers an ideal framework for multi-flux treatment as its compact formalism allows straightforward extensions [Hébert and Emmel 2015] and multiple scattering operator derivations. Furthermore, the underlying linear algebra is particularly well suited for volume layers and allows the derivation of efficient closed-form transfer matrices [Kubelka 1948]. Inspection of Belcour's [2018] statistical framework suggests partial compliance with the transfer matrix formalism as energy statistics obey the classical adding equations. Unfortunately, as suggested by the form of Equation (11), variance statistics can not be computed using transfer matrices. Therefore, we propose to overcome this limitation by switching to a representation better suited to a transfer matrix modeling.

4 HENY-STEIN REPRESENTATION

The Henyey-Greenstein (HG) phase function [1941] is widely used in computer graphics to describe the scattering properties of participating media. This function is controlled by a single parameter $g \in [-1, 1]$ called the *asymmetry parameter* and is defined as

$$p(\theta; g) = \frac{1}{4\pi} \frac{1 - g^2}{(1 + g^2 - 2g \cos \theta)^{\frac{3}{2}}} \quad (13)$$

with θ denoting the deviation angle from the forward direction. As it depends only on a deviation angle, this function is axially symmetric around the forward direction. When $g = 0$, the resulting function is uniform, while it degenerates to a forward (resp. backward) delta function when $g = 1$ (resp. $g = -1$).

An important property of this phase function is its *stability under convolution* [Davis 2006]: Convoluting two arbitrary HG phase functions $p(g_1)$ and $p(g_2)$ yields another HG phase function such that

$$(p(g_1) * p(g_2))(\theta) = p(\theta; g_1 \times g_2). \quad (14)$$

where $*$ denotes the spherical convolution product. Note that the vMF distribution also known as *Spherical Gaussian* used in many previous works does not, strictly speaking, possess this stability [Mardia and Jupp 2009].

We build upon this property to approximate light-matter interactions in participating media and scattering due to GGX based rough surfaces. To this end, we propose to approximate hemispherical light distributions with HG lobes, respectively defined with an energy, a mean direction, and an asymmetry parameter.

In practice, our approach is similar to the work of Belcour [2018] except that we summarize the statistics of each lobe with sets of the form $\phi = (e, \mu, g)$ where asymmetry parameters g are used in place of variance descriptors σ . Similarly to the prior work, we also always assume that the mean directions of the HG lobes obey Snell's laws which considerably simplifies the derivations. Based on this representation, we leverage HG convolution property to compute the statistics of the outgoing HG lobes. Finally, we map back the outgoing HG lobes to GGX lobes with equivalent perceived roughnesses since HG lobes cannot reproduce the light distributions resulting from microfacet surface scattering.

In contrast to the prior work, tracking HG asymmetry parameters presents two core benefits. First, this representation provides native support for participating media as they are commonly described with this phase function. Most importantly, its multiplicative convolution property allows the resolution of layer multiple scattering with the transfer matrix formalism, which enables efficient support for scattering volume layers.

5 A TWO-FLUX MODEL FOR ROUGH INTERFACES

This section consists of a transfer matrix compliant reformulation of Belcour's [2018] statistical framework. We start by approximating rough reflections and transmissions by HG convolution products and show how to compute multiple-layer light transport statistics with transfer matrices.

5.1 Rough reflection and transmission HG representations

We obtained the equivalent asymmetry parameter of a GGX based microfacet BRDF with roughness $\alpha \in]0, 1]$ by adapting the L^3 optimization method provided by Heitz et al. [2016]. Figure 6 outlines some fitting results for varying incidence and GGX roughness. In general, the equivalent asymmetry parameter does not remain constant with respect to incident elevation due to the deformation that microfacet-based lobes undergo toward grazing angles. In practice, we restrict our fit to normal incidences for which we were able to find the analytic approximation

$$g(\alpha) = -0.085 + \frac{1.085}{1 + \left(\frac{\alpha}{0.5}\right)^{1.3}}. \quad (15)$$

The accuracy level of this approximation is reported in the supplementary document. Thus, given an incoming light distribution with asymmetry parameter g_i , we approximate the asymmetry parameter of the reflected distribution following HG convolution property (14) as

$$g^R = g_i \times g(\alpha). \quad (16)$$

When crossing a dielectric interface, light distributions undergo expansion or contraction due to the index change. In the case of Henyey-Greenstein, this additionally translates as a loss of the axial symmetry due to Snell's refraction law. However, while these deformations increase with the incident elevation, we observed that




Fig. 6. We approximate a GGX BRDF model with an Henyey-Greenstein lobe to compute subsurface transport. GGX roughness is $\alpha = 0.1$ for columns (a) and (b), and $\alpha = 0.2$ for columns (c) and (d). Incident elevation angle is set to $\theta_i = 0^\circ$ for columns (a) and (c), and $\theta_i = 60^\circ$ for columns (b) and (d).

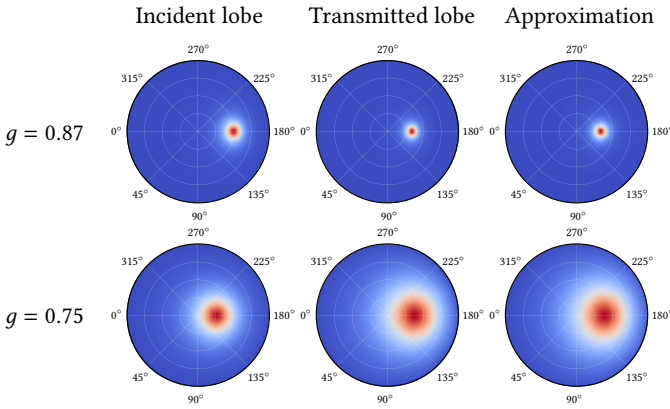


Fig. 7. At moderate incidences, the distribution transmitted through a smooth interface (middle column) due to an incident HG lobe (left column) can be approximated by another HG lobe. The right column shows the results obtained with our approximation. In the first row, the incident and transmitted media's indices are respectively set to $\eta_i = 1$ and $\eta_t = 1.5$, and the incident elevation to $\theta_i = 30^\circ$. The second row depicts a transmission with $\eta_i = 1.5$ and $\eta_t = 1$ and an incident elevation defined to $\theta_i = 15^\circ$. The incident asymmetry parameters used in the two rows respectively correspond to microfacet-based reflected lobes with GGX roughnesses $\alpha = 0.1$ and $\alpha = 0.2$.

the transmitted distribution of an incident HG lobe could be approximated by another HG lobe for moderate elevations. Based on this observation, we approximate the transmitted lobe's asymmetry parameter due to a smooth interface with the operator

$$h(g) = \sqrt{1 - \max(0, \min(t, 1))} \quad \text{with} \quad t = \left(1 - g^2\right) \left(\frac{\eta_i}{\eta_t}\right)^{\frac{3}{4}} \quad (17)$$

where η_i and η_t respectively denote the refractive indices of the incident and transmitted media. Figure 7 depicts some results obtained thanks to this approximation. Additional validation plots are available in the supplementary document. In the general case of a rough interface, the transmitted distribution additionally undergoes blurring effects due to the interface's micro-geometry. Leveraging

Henyey-Greenstein's convolution property, we compute the asymmetry parameter of a transmitted distribution as

$$g^T = h(g_i) \times g(s \times \alpha) \quad (18)$$

where s is Belcour's [2018] fake rough transmission scaling factor defined in Equation (9).

5.2 Transfer matrix formulation

In the case of two rough interfaces, an incident ray of light gives rise to an infinite sequence of reflected lobes which we summarize as a single HG lobe with energy e_{02}^R and asymmetry parameter g_{02}^R .

The core idea of this paper is to compute these quantities based on transfer matrices of the form

$$M_{ij} = \frac{1}{\tau_{ij}} \begin{bmatrix} 1 & -\rho_{ji} \\ \rho_{ij} & \tau_{ij}\tau_{ji} - \rho_{ij}\rho_{ji} \end{bmatrix}. \quad (19)$$

As studied in Section 3.2, the total amount of reflected energy is given as the geometric series

$$e_{02}^R = \mathbf{r}_{01} + \mathbf{t}_{01}\mathbf{r}_{12}\mathbf{t}_{10} + \dots \quad (20)$$

Similarly, the asymmetry parameter of the reflected lobe can be expressed as the weighted sum

$$\frac{\tilde{g}_{02}^R}{e_{02}^R} \quad \text{with} \quad \tilde{g}_{02}^R = \mathbf{r}_{01} g_{01}^R + \mathbf{t}_{01} g_{01}^T \mathbf{r}_{12} g_{12}^R \mathbf{t}_{10} g_{10}^T + \dots \quad (21)$$

where g_{ij}^R and g_{ij}^T , respectively denote constant asymmetry parameters approximating the effects of rough reflections and transmissions between media i and j . While the constant asymmetry parameter terms for reflections are directly given as $g_{ij}^R = g(\alpha_{ij})$, special care must be given to the g_{ij}^T transmission terms due to the non-linearity of operator (17). We address this limitation by linearizing the effects of rough transmissions on incident asymmetry parameters as

$$g_{ij}^T = \frac{g_j^{(1)}}{g_i^{(1)}} \quad \text{with} \quad g_j^{(1)} = h_{ij}(g_i^{(1)}) \times g(s_{ij} \times \alpha) \quad (22)$$

where $g_i^{(1)}$ is a first-order asymmetry parameter not accounting for multiple bounces between interfaces which we iteratively track in each medium thanks to the operators studied in Section 5.1. We compute the opposite transmission's transfer factor g_{ji}^T following the same methodology except that we account for the next interface's roughness when computing the first-order asymmetry parameters terms.

The main result of this construction is that the statistics of the reflected lobe can now be computed using transfer matrices since energy and unnormalized asymmetry parameter take the form of geometric series. More precisely, we compute these quantities with specialized forms of the matrix (19) which we refer to as *energy* and *asymmetry* transfer matrices. We respectively denote these matrices as E and G .

For two interfaces, the total amount of reflected energy is given as $e_{02}^R = R(E_{01}E_{12})$, where E_{ij} denote instances of the transfer matrix (19) with transfer factors $\rho_{ij} = \mathbf{r}_{ij}$ and $\tau_{ij} = \mathbf{t}_{ij}$, and R is the reflection matrix operator (5). Similarly, we compute the unnormalized asymmetry parameter \tilde{g}_{02}^R of the reflected lobe as $\tilde{g}_{02}^R = R(G_{01}G_{12})$

where G_{ij} denote instances of the transfer matrix (19) with transfer factors $\rho_{ij} = \mathbf{r}_{ij} g_{ij}^R$ and $\tau_{ij} = \mathbf{t}_{ij} g_{ij}^T$.

5.3 Multiple-layers statistics

We iteratively compute the statistics of the light transport due to an arbitrary number of interfaces based on the matrix formulation introduced in Section 5.2. The iterative approach is described in Algorithm 1.

Algorithm 1: Multi-lobe BSDF transfer matrix algorithm

Input: Incident direction ω_i and properties of the n stack components ordered by increasing depth in the stack

Output: Outgoing lobes *energies, means and roughnesses*

begin

```

     $\mu_i \leftarrow -\omega_i$  // Mean direction
     $g_i^{(1)} \leftarrow 1$  // First-order asymmetry parameter term
     $E_{0i} \leftarrow I$  // Accumulated energy transfer matrix
     $G_{0i} \leftarrow I$  // Accumulated asymmetry transfer matrix

```

// Stack components traversal

for j in $1..n$ **do**

- 1) Compute transmission's mean direction μ_j and first-order asymmetry parameter term $g_j^{(1)}$
- 2) Instantiate component's E_{ij} and G_{ij} transfer matrices
- 3) Compute E_{0j} and G_{0j} transfer matrices including top components
- 4) Compute the *energy, the mean and the roughness* of the component's reflection lobe

// Update iteration variables

```

     $\mu_i \leftarrow \mu_j$ ,  $g_i^{(1)} \leftarrow g_j^{(1)}$ ,  $E_{0i} \leftarrow E_{0j}$ ,  $G_{0i} \leftarrow G_{0j}$ 

```

- 5) Compute the *energy, the mean and the roughness* of the lobe transmitted through the stack
-

In the case of a stack only composed of interfaces, the properties of the scattering components reduce to the refractive index and the roughness of each interface. For each interface (Step 1), we compute the mean transmission direction based on Snell's refraction law and the first-order transmission asymmetry parameter according to Equation (22). We use the latter to instantiate the energy and asymmetry transfer matrices E_{ij} and G_{ij} of the interface (Step 2) according to the methodology studied in Section 5.2, and we compute the transfer matrices accounting for top components as $E_{0j} = E_{0i} E_{ij}$ and $G_{0j} = G_{0i} G_{ij}$ (Step 3). Then, we extract the TR^+T paths statistics due to the interface (Step 4) as

$$\begin{aligned}
 e_{0j}^{TRT} &= R(E_{0j}) - R(E_{0i}) \\
 \mu_{0j}^{TRT} &= \text{reflect}(\omega_i) \\
 \tilde{g}_{0j}^{TRT} &= R(G_{0j}) - R(G_{0i}).
 \end{aligned} \tag{23}$$




Fig. 8. When traveling in a medium, a photon may cross the volume without suffering any interaction (a), it may be absorbed (b), or it may be scattered (c). To make the problem more tractable without sacrificing important appearance features, we isolate unscattered light from forward- and back-scattered light. We respectively denote the corresponding fluxes as i , j^f , and j^b (d).

Note that for a conducting base, the transfer matrices are undefined due to the zero transmission factors. In this case, we compute the reflected lobe's statistics using the explicit reflectance operator

$$R(M, \rho) = \frac{M_{21} + M_{22}\rho}{M_{11} + M_{12}\rho} \tag{24}$$

where M and ρ respectively denote the transfer matrices of the upper layers and the reflectance transfer factors of the conducting base. Derivation of this operator is given in the supplementary document.

For a transmissive substrate, we compute the statistics of the light transmitted through the stack (Step 5) as

$$e^T = T(E_{0n}), \quad \mu^T = \text{refract}(\omega_i, \eta_{0n}), \quad \tilde{g}^T = T(G_{0n}) \tag{25}$$

where T is the transmission matrix operator (6).

Finally, we compute the equivalent GGX roughness of each outgoing lobe by applying the inverse HG fit to the normalized lobes asymmetry parameters.

6 A SIX-FLUX MODEL FOR SCATTERING VOLUMES LAYERS

In this section, we introduce the six-flux approach which extends the two-flux model studied in Section 5 to account for scattering volumes effects. We start by introducing the general approach in Section 6.1, and we present the general form of the six-flux transfer matrix and its associated multiple-scattering operators in Section 6.2. Then, we derive specialized forms of the six-flux transfer matrix for homogeneous participating media and interface components in Sections 6.3 and 6.4. Finally, we show how to compute the statistics of the outgoing lobes with the six-flux approach in Section 6.5, and we introduce an improvement to the resulting BSDF in Section 6.6.

6.1 The six-flux approach

Three outcomes are possible for a photon, i.e., a light particle carrying a finite amount of energy, when traveling in a scattering volume: It may cross the entire volume without undergoing any interaction (Fig. 8(a)), it may be absorbed (Fig. 8(b)), or its direction of flight may change due to scattering by particles in the medium (Fig. 8(c)).




Fig. 9. With the six-flux approach, we isolate light having undergone one or more scattering events in a medium (secondary j fluxes depicted as orange arrows) from light not scattered by any medium (primary i fluxes depicted in black). We additionally split the scattered light part into forward and backward contributions denoted as j^f and j^b . The left inset depicts the primary-to-primary and primary-to-secondary flux transfers occurring between depths $z = 0$ and $z = 1$ of a generic scattering component (or stack of components). The right inset focuses on secondary-to-secondary flux transfers.

Determining the escape probability of a photon in an arbitrary direction given its incoming direction falls into the radiative transfer realm and, in general, does not admit closed-form solutions.

To simplify the problem without sacrificing important appearance features, we propose to isolate the unscattered light part from light having undergone one or more scattering events in a medium. We respectively denote the corresponding fluxes i and j and refer to them as *primary* and *secondary fluxes*. Moreover, a significant amount of light might be scattered back with participating media layers. Thus, we propose to split the scattered light part into forward and backward contributions, which we respectively denote as j^f and j^b (Fig. 8(d)).

In terms of our HG representation, we express these light distributions as hemispherical HG lobes. Also, we recall that according to the assumptions made in Section 4, we always assume that the mean directions of the HG lobes obey Snell's laws. We only make a distinction for secondary backward lobes whose mean directions are oriented in the opposite direction of light travel.

Accounting for the vertical direction of propagation in the structure, we thus describe the light distribution at any depth z of the layered structure as a six-flux vector $I_z = [i_z^+, i_z^-, j_z^{f+}, j_z^{f-}, j_z^{b+}, j_z^{b-}]^T$.

6.2 The six-flux transfer matrix

Under the six flux approach, each component of the stack may give rise to primary-to-primary, primary-to-secondary, and secondary-to-secondary flux transfers. We respectively denote the corresponding reflection and transmission transfer factors as $\rho|\tau$, $\rho|\tau^x$, and $\rho|\tau^{xy}$, where the superscripts specify the lateral directions of the interacting secondary fluxes. For instance, we denote primary-to-secondary back reflection as ρ^b , secondary-to-secondary forward transmission as τ^{ff} , and so on. Figure 9 depicts the possible flux transfers occurring between depths $z = 0$ and $z = 1$ of a generic stack component, may it be an interface, a scattering volume, or any combination of them.

As in the two-flux case, the flux balance between the boundaries of the component can be compactly written as

$$I_0 = M_{01}I_1 \quad \text{with} \quad M_{01} = \begin{bmatrix} M_{pp} & 0 & 0 \\ M_{pf} & M_{ff} & M_{bf} \\ M_{pb} & M_{fb} & M_{bb} \end{bmatrix} \quad (26)$$

where each M_{01} sub-block is a 2x2 matrix. In this form, primary-to-primary, primary-to-secondary, and secondary-to-secondary flux transfers are respectively expressed through the M_{pp} , M_{px} , and M_{xy} matrix blocks.

Given a transfer matrix M_{ij} , the primary reflectance transfer factor ρ_{ij} of the corresponding component (or stack of components) can be obtained as $\rho_{ij} = R(M_{ij})$ where R is the matrix operator (5). The primary-to-secondary forward and backward reflectances ρ_{ij}^f and ρ_{ij}^b accounting for internal multiple-scattering can respectively be obtained with the matrix operators R^f and R^b , defined as

$$R^f(M) = \frac{M_{41}}{M_{11}} + \frac{M_{45}|M|_{31,53} - M_{43}|M|_{31,55}}{M_{11}|M|_{33,55}} \quad (27)$$

and

$$R^b(M) = \frac{M_{61}}{M_{11}} + \frac{M_{65}|M|_{31,53} - M_{63}|M|_{31,55}}{M_{11}|M|_{33,55}} \quad (28)$$

where $|M|_{kl,mn} = M_{kl}M_{mn} - M_{ml}M_{kn}$ denotes the 2x2 matrix determinant operator. Similarly, the primary transmittance can be evaluated as $\tau_{ij} = T(M_{ij})$ where T is the matrix operator (6), and the primary-to-secondary forward transmission is obtained with the matrix operator T^f defined as

$$T^f(M) = \frac{-|M|_{31,55}}{M_{11}|M|_{33,55}}. \quad (29)$$

Derivation of these operators is given in the supplementary document.

6.3 Homogeneous participating media

In this section, we derive the transfer matrix of a scattering volume of arbitrary thickness. To this end, we start by modeling the flux transfers occurring in a thin volume slab, and we derive the transfer matrix for an arbitrary depth thanks to a matrix exponentiation. We restrict our study to homogeneous participating media and respectively denote σ_t and σ_s , the extinction and scattering cross-sections *per unit length* traveled in the medium. We also assume Henyey-Greenstein with asymmetry parameter g as the phase function of the medium.

When propagating in the medium, each flux decreases due to volume extinction, while each secondary flux also undergoes an increase due to primary and secondary flux in-scattering. Assuming that the mean direction of each flux obeys Snell's laws, the flux transfers are moreover independent of the vertical and lateral directions of propagation. Consequently, under these simplifications, five transfer factors τ , τ^f , τ^{ff} , ρ^b , and ρ^{fb} are required to describe a homogeneous participating medium. Figure 10 depicts the flux transfers occurring in a thin volume slab of height Δz .




Fig. 10. Six-flux transfers occurring in a thin homogeneous volume slab of height Δz resulting from our simplifications.

We can express the back-scattering probability per unit length traveled of a photon propagating in the medium, with respect to its flight direction, as the back-scattering cross-section σ_{sb} defined with respect to the phase function's lower hemisphere integral as

$$\sigma_{sb} = \sigma_s \times p_-(g) \quad \text{with} \quad p_-(g) = 2\pi \int_{-1}^0 p(\cos \theta; g) d \cos \theta. \quad (30)$$

This integral can be accurately approximated as

$$p_-(g) = \begin{cases} Q(|g|) & \text{if } g \geq 0 \\ 1 - Q(|g|) & \text{otherwise} \end{cases} \quad (31)$$

with

$$Q(u) = 0.5039 - 0.8254 u + 0.3226 u^2.$$

The accuracy level of this approximation is reported in the supplementary document. We analogously define the forward scattering σ_{sf} cross-section as $\sigma_{sf} = \sigma_s - \sigma_{sb}$, and we approximate the extinction and the forward- and back-scattering of each flux with cross-sections σ_t , σ_{sf} , and σ_{sb} evaluated at mean directions. In practice, we find that these simplifications turn out to be acceptable compromises, as shown in Section 7.2.

Passing down to the limit $\Delta z \rightarrow 0$, we can express the energy change rate of each flux as the differential system

$$\frac{dI_z}{dz} = \frac{A}{\cos \theta_i} I_z \quad (32)$$

with

$$A = \begin{bmatrix} -\sigma_t & 0 & 0 & 0 & 0 & 0 \\ 0 & \sigma_t & 0 & 0 & 0 & 0 \\ \tau^f & 0 & -\sigma_t + \tau^{ff} & 0 & 0 & \rho^{fb} \\ 0 & -\tau^f & 0 & \sigma_t - \tau^{ff} & -\rho^{fb} & 0 \\ 0 & \rho^b & 0 & \rho^{fb} & -\sigma_t + \tau^{ff} & 0 \\ -\rho^b & 0 & -\rho^{fb} & 0 & 0 & \sigma_t - \tau^{ff} \end{bmatrix}$$

where

$$\rho^b = \rho^{fb} = \sigma_{sb} \quad \text{and} \quad \tau^f = \tau^{ff} = \sigma_{sf}. \quad (33)$$

Similarly, we can express the change rates of the unnormalized asymmetry parameters of the light distributions as an instance of the differential system (32) with transfer factors defined as

$$\rho^b = \rho^{fb} = -\sigma_{sb} g \quad \text{and} \quad \tau^f = \tau^{ff} = \sigma_{sf} g \quad (34)$$

where the negated asymmetry parameter accounts for the reversal of back-scattered light propagation directions.

As a first-order ordinary differential system, Equation (32) admits general solutions of the form $I_z = \exp(A(z - z_0)/\cos \theta_i) I_{z_0}$. In particular, with $z = 0$ and $z_0 = h$, we have $I_0 = M I_h$ where M denotes the transfer matrix of the medium of depth h defined as

$$M = e^{-h'A} = \frac{1}{\gamma} \begin{bmatrix} M_{pp} & 0 & 0 \\ M_{pf} & M_{ff} & M_{fb} \\ M_{fb} & M_{fb} & M_{ff} \end{bmatrix} \quad (35)$$

with

$$M_{pp} = \begin{bmatrix} \gamma e^{\sigma_t h'} & 0 \\ 0 & \gamma e^{-\sigma_t h'} \end{bmatrix} \quad M_{ff} = \begin{bmatrix} C\gamma + S\alpha & 0 \\ 0 & C\gamma - S\alpha \end{bmatrix}$$

$$M_{fb} = \begin{bmatrix} 0 & -S\beta \\ S\beta & 0 \end{bmatrix} \quad M_{pf} = \begin{bmatrix} C\gamma + S\alpha - \gamma e^{\sigma_t h'} & 0 \\ 0 & C\gamma - S\alpha - \gamma e^{-\sigma_t h'} \end{bmatrix}$$

where $h' = h/\cos \theta_i$, $\alpha = \sigma_t - \tau^f$, $\beta = \rho^b$, $\gamma = \sqrt{\alpha^2 - \beta^2}$, $S = \sinh(\gamma h')$, and $C = \cosh(\gamma h')$. While the M_{pp} block directly follows from Beer-Lambert's law, it is worth noting that the M_{ff} and M_{fb} blocks take a form analogous to the Kubelka-Munk [1948] solutions for the case of two-flux diffuse transfer.

Based on this result, we respectively obtain the energy and asymmetry transfer matrices E and G of an arbitrary depth homogeneous volume as instances of matrix (35) with respective transfer factors (33) and (34).

6.4 Interface components

In the case of an interface component, no transfer occurs between primary and secondary fluxes. Assuming that most light is forwardly reflected and transmitted at the interface, we express the corresponding transfer matrix as

$$M = \begin{bmatrix} M_{pp} & 0 & 0 \\ 0 & M_{ff} & 0 \\ 0 & 0 & M_{bb} \end{bmatrix} \quad (36)$$

where each M sub-block takes the form of a two-flux transfer matrix (19). As a result, we compute the energy and asymmetry transfer




Fig. 11. With the six-flux approach, we approximate the BSDF of the stack as a mixture of forward and backward microfacet lobes sharing the same mean directions.

matrices E and G of an interface by instantiating their respective 2×2 sub-blocks following the two-flux methodology studied in Section 5.2.

6.5 Multiple-layers statistics

We compute the light transport statistics due to each layer following the iterative approach described in Algorithm 1, except that scattering components now consist of interfaces and participating media. With the six-flux approach, we summarize the light transport due to each component with two additional sets of statistics corresponding to back-scattered and forward-scattered secondary fluxes. As a result, we approximate the BSDF of the stack as a mixture of forward and backward microfacet lobes sharing the same mean directions (Fig. 11).

When iterating through the stack, we instantiate the energy and asymmetry transfer matrices E_{ij} and G_{ij} of each component (Step 2) based on the definitions given in Sections 6.3 and 6.4. Furthermore, we specialize the matrix product operations depending on the transfer matrix of the scattering component being multiplied when computing E_{0j} and G_{0j} transfer matrices accounting for top components (Step 3). This significantly decreases the rendering time as the specialized transfer matrices contain large numbers of zero entries. In addition to the primary reflected lobes statistics (Step 4), we compute the energies of the secondary lobes and their roughnesses, by sequentially evaluating the intermediate transfer matrices (23) with the R^f and R^b operators defined in Section 6.2. For backward reflection statistics, we define the mean directions of the lobes as the incident direction.

Note that for an interface, computing the secondary lobes is necessary only if the upper components contain a scattering volume layer. Additionally, we compute primary and secondary forward reflected lobes only if the components include an interface.

As for the two-flux case, explicit reflectance operators must be used for conducting bases as their transfer matrices can not be instantiated. In this case, we evaluate the primary reflectance with the $R(M, \rho)$ operator defined in Section 5.3, and we provide the explicit $R^f(M, \rho)$ and $R^b(M, \rho)$ operators for conductors secondary reflectances in the supplementary document.




Fig. 12. In its original form [Belcour 2018], the BSDF is not energy preserving as it relies on single-scattering microfacet lobes (left inset). We restore this property by normalizing each lobe with a pre-integrated GD term (right inset). Here, a non-absorbing scattering volume with parameters $\sigma_s=0.755$, HG $g=0$ and depth $d=4$ lies on top of an ideally reflecting conductor with roughness $\alpha=0.1$.

6.6 BSDF normalization

It is important to note that in its original form (Equation (12)), the resulting BSDF is not energy preserving as it relies on single-scattering microfacet-based lobes. While this does not significantly affect the appearance of stacks of moderately rough interfaces, this may have a considerable impact on participating media as their equivalent roughnesses can be arbitrarily high. As a consequence, a substantial amount of energy may be lost in the presence of scattering media, leading to inconsistent dark appearances. We correct this misbehavior by normalizing each lobe of the BSDF (Fig. 12): Letting e_k , μ_k , and $\alpha_k = g^{-1}(g_k)$ denote the k^{th} lobe's energy, mean, and perceived roughness, we instantiate the corresponding BSDF lobe as

$$e_k \frac{\rho(\omega_k, \omega_o, \alpha_k)}{\text{GD}(\omega_k, \alpha_k)} \quad \text{with} \quad \omega_k = \text{reflect}(\mu_k), \quad (37)$$

and

$$\text{GD}(\omega_k, \alpha_k) = \int_{\Omega_+} \rho(\omega_k, \omega_o, \alpha_k) \cos \theta_o d\omega_o. \quad (38)$$

As no closed-form exists for this integral, we precompute it in a 2D LUT indexed over a predefined range of incidence elevations and roughnesses. Finally, we sample the BSDF according to the balance MIS scheme introduced in the previous work (see [Belcour 2018] for details).

7 RESULTS

The results shown in this section are computed using the path integrator of the Mitsuba rendering engine [Jakob 2010] for which we provide implementations of the two- and six-flux approaches. We use a pre-computed GGX TIR LUT of resolution 64^3 occupying 1MB of RAM and a GGX FGD LUT of resolution 64^4 occupying 64MB RAM. For the sake of reproducibility, we do not account for multiple bounces in the micro-geometry when computing the FGD LUT. All the results were generated on an Intel[®] i7-8750H Core[™] processor with 16 GB of RAM, and the images used for reference were computed using the stochastic approach of Guo et al. [2018]. The quality of the results is evaluated using a ΔE_{00} perceptual difference with respect to the reference method. We respectively label TM2 and TM6 the results obtained with the two- and six-flux approaches in the figures. Sections 7.1 and 7.2 respectively present the results of the two- and six-flux approaches, and limitations as well as future work are discussed in Section 7.3.

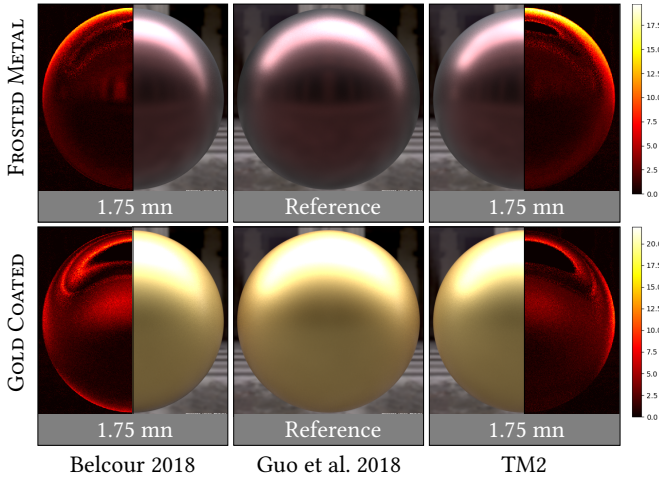


Fig. 13. We reproduced the FROSTED METAL and GOLD COATED examples from Belcour [2018] involving a clear coat with $\eta_1=1.5$ and $\alpha_1 = 0.1$ on top of a conducting base. The roughnesses of the conducting bases are respectively set to $\alpha_2=0.01$ and $\alpha_1=0.1$. The results were obtained using 128 samples per pixel.

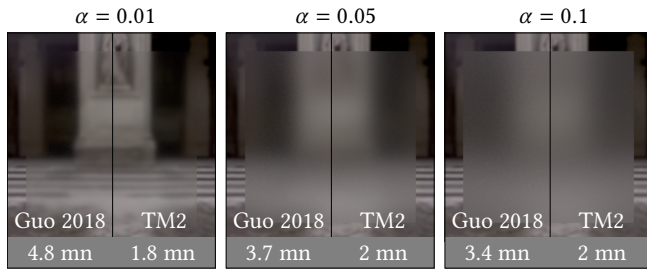


Fig. 14. In the above scenes, the glass plate is modeled as a single quad surface. The slab has a refractive index $\eta=1.5$ and the two sides of the plate share the same roughness α . The results were obtained using 128 samples per pixel.

7.1 Two-flux model

Figure 13 reproduces Belcour’s FROSTED METAL and GOLD COATED scenes involving a dielectric coat lying on top of a conducting base. Except for differences at grazing angles, the results obtained with the two approaches remain close to each other. For these incidences, the approach of Belcour [2018] provides a perceived roughness in better agreement with the ground truth. We attribute these differences to the linear space introduced by the author to compute the variance statistics. The latter is explicitly designed to approximate the variance due to multiple reflections between two interfaces yielding better results at high incidences as the number of bounces tends to increase in these configurations. Additional clear coat examples with varying roughness are available in the supplementary document. Figure 14 illustrates the ability of our two-flux model to simulate transmission through dielectric stacks of varying roughness with results close to the stochastic ground truth.

7.2 Six-flux model

All the results introduced in this section have been produced using a pre-computed GGX GD LUT of resolution 64^2 occupying 16KB of RAM, used at runtime to normalize the BSDF (Section 6.6). For a fair comparison, we adapted Belcour’s [2018] Mitsuba plugin to account for this normalization.

Figures 15 and 16 depict a non-absorbing participating medium lying on top of a smooth golden base ($\alpha = 0.02$). Figure 15 analyzes the influence of the thickness of the volume on the appearances produced by the two approaches. In this example, the scattering is strongly forward-oriented ($g = 0.9$), and the thickness of the volume varies between $d = 0.1$ and $d = 4$. Belcour’s approach undergoes severe energy loss increasing with the optical thickness of the medium. The overly dark appearances produced by the approach are due to the underlying single-scattering approximation and its lack of back-scattering support. Note that the GD normalization does not help to recover the missing energy. Moreover, as can be seen on the top row, the approach results in inconsistent perceived roughness even in the case of thin volume layers. In contrast, the six-flux approach produces perceived roughnesses close to the ground truth for all volume depths. Figure 16 analyzes the influence of the scattering profile of the medium. In this example, the thickness of the volume is set to $d = 1$, while the scattering varies from a uniform ($g = 0$) to a forward-oriented profile ($g = 0.7$). Note how the six-flux model well reproduces the complex combinations of specular and diffuse components due to strong back-scattering. Further examples with varying depths and scattering profiles are available in the supplementary document.

Figure 17 illustrates the rendering of materials including a scattering volume in a dielectric layer on top of a conducting base. These configurations are particularly challenging as they lead to complex combinations of specular, directional-diffuse, and diffuse features evolving with the optical thickness of the medium. While some differences with the ground-truth can be observed, the six-flux model yields substantial improvement compared to previous approximations. Additional examples with varying depths and scattering profiles are available in the supplementary document.

The six-flux model entails additional computational costs: Using 6×6 matrices involves higher memory register consumption, additional arithmetic operations, and the evaluation of the secondary lobes incurs additional GD, FGD and TIR LUTs lookups. Therefore, it would be relevant to derive efficient analytical approximations for these integrals in subsequent work to minimize the rendering times. These additional costs are nevertheless amortized against stochastic approaches: Figure 18 illustrates the rendering of a material involving an absorbing scattering volume trapped between two interfaces. While stochastic approaches introduce significant variance, the six-flux approach provides low-variance results even with restricted sample budgets. Also, it is worth noting that the rendering time of the stochastic approach significantly increases with the optical depth of the medium due to the increasing amount of internal samples. In contrast, the computational cost of the six-flux model remains insensitive to the optical properties of the media (Fig. 18, bottom row).



Fig. 15. In the above scene, a non-absorbing scattering medium lies on top of a smooth golden material ($\alpha = 0.02$). The asymmetry parameter and the scattering cross-section of the medium are respectively set to $g = 0.9$ and $\sigma_s = 0.5$, and the thickness of the volume varies from $d = 0.1$ to $d = 4$. Note how the six-flux approach correctly reproduces the apparent roughness which increases with the thickness of the medium. The results were obtained using 1024 samples per pixel.

7.3 Limitations and future work

External absorbing media. Darkening toward grazing angles can be observed with the six-flux model in the case of external absorbing media when the phase function is forward-oriented (Fig. 19). These configurations are poorly supported as our approach assumes that light propagating in the forward direction must reflect off the lower interface before emerging. Consequently, most of the light undergoes extinction as the mean distance traveled through the medium increases toward large values in these configurations.

Back-scattering. As illustrated in Figure 20, departures from the ground truth can be observed with scattering media of increasing optical depth. These divergences are due to our HG back-scattering approximation (Equation 30) which approximates the back-scattering around the mean direction but does not account for multiple scattering events in the volumes. Consequently, our approximation tends to underestimate the back-scattered energy due to thick volumes, which explains the more apparent goldish look of our renderings in Figure 20. As another limitation, our implementation tends to produce brighter appearances near normal incidences in the case of scattering dielectric layers (Fig. 21). This limitation is due to our current implementation, which does not consider total internal reflections for back-scattered light in the volumes. This is a design choice made for simplicity purposes, and we believe that significant improvements could be made in this direction by adequately accounting for HG total internal reflections with a dedicated HG TIR LUT.



Fig. 16. In this example, a non-absorbing scattering medium is placed on top of a smooth golden material ($\alpha = 0.02$). The thickness of the medium is set to $d = 1$, the scattering cross-section is defined to $\sigma_s = 0.755$, and the asymmetry parameter of the phase function varies from $g = 0$ to $g = 0.7$. Note the drastic energy loss endured by the previous approach [Belcour 2018]. In contrast, the six-flux approach provides results close to the ground truth, even in the presence of strongly back-scattering media. The results were obtained using 1024 samples per pixel.




Fig. 17. This example depicts the complex case of a scattering volume trapped in a smooth dielectric layer of index $\eta = 1.5$ sitting on top of a rough conducting substrate ($\alpha = 0.1$). The scattering cross-section and the asymmetry parameter of the medium are respectively set to $\sigma_s = 0.5$, $g = 0.3$, and the thickness d of the medium varies from 0.1 to 2. The results were obtained using 128 samples per pixel.

GGX lobe mixtures. Similarly to previous work, our model is not suited to Lambertian substrates as the resulting BSDFs cannot be

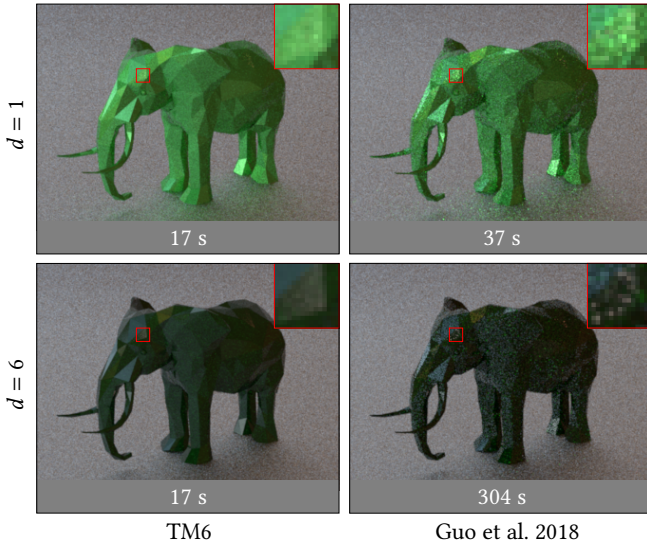


Fig. 18. This example illustrates the low-variance results obtained with the six-flux model with respect to the stochastic approach, using 32 samples per pixel. In this scene, an absorbing scattering dielectric layer lies on top of a conducting base. The absorption and scattering cross-sections of the medium are respectively set to $\sigma_k = (1, 0.2, 1)$ and $\sigma_s = 0.7$, the asymmetry parameter is set to $g = 0.9$, and the thickness d varies between 1 and 6. While the rendering time of the stochastic approach significantly increases with the optical depth of the medium, the computational cost of the six-flux model remains constant (bottom row).

reproduced with GGX lobe mixtures. Also, it is worth emphasizing that this limitation also applies to HG back-scattering which cannot be approximated with GGX lobes.

Multiple bounces between interfaces. As shown in Figures 13 and 22, Belcour’s approach provides perceived roughnesses in better agreement with the ground truth than our approaches at grazing angles, where the number of light bounces tends to increase. This is expected as the method directly accounts for the effects of multiple bounces in its linear variance representation. With this in mind, we believe that accounting for multiple light bounces between interfaces when fitting HG to GGX could be an interesting avenue for future work.

Anisotropic interfaces. Finally, anisotropic interfaces are currently not supported due to the axial symmetry of the Henyey-Greenstein phase function. We believe that an in-depth study of the anisotropic directional distributions existing in the literature, of their properties and associated operators, could inspire future work [Mardia and Jupp 2009; Xu et al. 2013].

8 CONCLUSION

Rendering of layered materials involving arbitrary combinations of rough layers and scattering volumes is a challenging problem. While the community has already proposed reference approaches, the efficient simulation of this class of materials in reasonable and controlled time budgets remains an open problem of broad interest.




Fig. 19. In the case of external absorbing media, darkening can be observed toward high elevations with the six-flux approach when the phase function is forward-oriented. In these examples, an absorbing medium with $\sigma_k = 0.1$, $\sigma_s = 0.5$, $g = 0.9$ and thickness d lies on top of a smooth conducting base.




Fig. 20. Our approach tends to underestimate the back-scattered energy due to thick volumes, as our HG back-scattering cross-section approximation does not account for multiple scattering events in the volumes. In these examples, a non-absorbing scattering medium with asymmetry parameter $g = 0.5$ and scattering cross-section $\sigma_s = 0.755$ lies on top of a golden material.



Fig. 21. In the case of scattering dielectric layers, our implementation tends to produce brighter appearances near normal incidences as we do not consider total internal reflections for back-scattered light in the volumes. In the above scenes, a colored scattering dielectric layer with refractive index $\eta = 1.5$, scattering cross-section $\sigma_s = 0.7$ and asymmetry parameter $g = 0.9$, lies on top of a conducting base.

To address these efficiency problems, we introduced an original approach based on a flux transfer model leveraging key properties of the Henyey-Greenstein phase function to solve the complex light transport occurring in these materials. We demonstrated that our approach produces results comparable to the state-of-the-art in the case of rough isotropic interfaces. Our approach unveils its strength in the presence of scattering volumes with efficient support of forward- and back-scattering involved in these media and multiple scattering occurring in the overall layered structure. While the previous method fails to reproduce the appearance of these




Fig. 22. The model of Belcour provides perceived roughnesses in better agreement with the ground truth than our approach at grazing angles, where the number of light bounces tends to increase. In this scene, the material is composed of three dielectric layers with refractive indices $\eta_1 = 1.4$, $\eta_2 = 1$, $\eta_3 = 1.4$ and roughnesses $\alpha_1 = 0.1$, $\alpha_2 = 0.01$, $\alpha_3 = 0.1$, on top of a smooth conductor.

materials, the six-flux approach provides better support for these configurations, with a computational cost both lower than the existing reference approaches and independent of the optical properties of the volumes considered. Moreover, the approach should be well suited to volume layers whose optical properties vary as a continuous function of depth in the structure, in which case efficient closed-form transfer matrices might be derived. Also, it is worth mentioning that this framework could find interesting reverse engineering applications as layer removals are simply expressed as lightweight inverse matrix products. It would be interesting to investigate whether the approach is suitable for interactive rendering on the GPU with additional simplifications.

ACKNOWLEDGMENTS

The authors would like to thank Jonathan Dupuy and Laurent Belcour for their comments and feedback at the early stage of this project. Thanks also to Philippe Porral, Thomas Muller, and Loïs Mignard-Debise for their support during the writing of this paper. The authors finally thank the Image Center of the University of Reims-Champagne-Ardenne and Hervé Deleau for his technical support.

REFERENCES

Florin Abeles. 1948. Sur la propagation des ondes électromagnétiques dans les milieux stratifiés. In *Annales de physique*, Vol. 12. EDP Sciences, 504–520.

Gladimir VG Baranoski and Jon G Rokne. 2001. Efficiently simulating scattering of light by leaves. *The Visual Computer* 17, 8 (2001), 491–505.

Mégane Bati, Romain Pacanowski, and Pascal Barla. 2019. Comparative study of layered material models. In *Workshop on Material Appearance Modeling*, H. Rushmeier and R. Klein (Eds.). Strasbourg, France. <https://hal.archives-ouvertes.fr/hal-02184562>

Laurent Belcour. 2018. Efficient rendering of layered materials using an atomic decomposition with statistical operators. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 73.

Alexis Benamira and Sumanta Pattanaik. 2020. Application of the Transfer Matrix Method to Anti-reflective Coating Rendering. In *Computer Graphics International Conference*. Springer, 83–95.

Patrick Callet. 1996. Pertinent data for modelling pigmented materials in realistic rendering. In *Computer Graphics Forum*, Vol. 15. Wiley Online Library, 119–127.

Anthony B Davis. 2006. Effective propagation kernels in structured media with broad spatial correlations, illustration with large-scale transport of solar photons through cloudy atmospheres. In *Computational Methods in Transport*. Springer, 85–140.

Julie Dorsey and Pat Hanrahan. 1996. Modeling and Rendering of Metallic Patinas. In *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '96)*. ACM, New York, NY, USA, 387–396.

Oskar Elek. 2010. Layered Materials in Real-Time Rendering. In *14th Central European Seminar on Computer Graphics*, Vol. 27.

Serkan Ergun, Sermet Önel, and Aydin Ozturk. 2016. A general micro-flake model for predicting the appearance of car paint. In *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations*. Eurographics Association, 65–71.

Marc A Gali, Angus R Gentle, Matthew D Arnold, and Geoffrey B Smith. 2017. Extending the applicability of the four-flux radiative transfer method. *Applied optics* 56, 31 (2017), 8699–8709.

Luis E. Gamboa, Adrien Gruson, and Derek Nowrouzezahrai. 2020. An Efficient Transport Estimator for Complex Layered Materials. *Computer Graphics Forum* 39, 2 (2020), 363–371. <https://doi.org/10.1111/cgf.13936> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13936>

Jinwei Gu, Ravi Ramamoorthi, Peter N Belhumeur, and Shree K Nayar. 2007. Dirty Glass: Rendering Contamination on Transparent Surfaces.. In *Rendering Techniques*. 159–170.

Jie Guo and Jin-Gui Pan. 2014. Real-time simulating and rendering of layered dust. *The Visual Computer* 30, 6-8 (2014), 797–807.

Jie Guo, Jinghui Qian, Yanwen Guo, and Jingui Pan. 2016. Rendering thin transparent layers with extended normal distribution functions. *IEEE transactions on visualization and computer graphics* 23, 9 (2016), 2108–2119.

Yu Guo, Miloš Hašan, and Shuang Zhao. 2018. Position-free monte carlo simulation for arbitrary layered BSDFs. In *SIGGRAPH Asia 2018 Technical Papers*. ACM, 279.

Pat Hanrahan and Wolfgang Krueger. 1993. Reflection from layered surfaces due to subsurface scattering. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*. Citeseer, 165–174.

Miloš Hašan, Martin Fuchs, Wojciech Matusik, Hanspeter Pfister, and Szymon Rusinkiewicz. 2010. Physical reproduction of materials with specified subsurface scattering. In *ACM SIGGRAPH 2010 papers*. 1–10.

Mathieu Hébert and Patrick Emmel. 2015. Two-Flux and Multiflux Matrix Models for Colored Surfaces. *Handbook of Digital Imaging* (2015), 1–45.

Eric Heitz, Jonathan Dupuy, Stephen Hill, and David Neubelt. 2016. Real-time polygonal-light shading with linearly transformed cosines. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 41.

Louis G Henyey and Jesse L Greenstein. 1941. Diffuse radiation in the galaxy. *The Astrophysical Journal* 93 (1941), 70–83.

Siu-chi Hsu and Tien-tsin Wong. 1995. Simulating dust accumulation. *IEEE Computer Graphics and Applications* 15, 1 (1995), 18–22.

Isabelle Icart and Didier Arquès. 2000. A physically-based BRDF model for multilayer systems with uncorrelated rough boundaries. In *Eurographics Workshop on Rendering Techniques*. Springer, 353–364.

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 Transactions on Graphics (ToG)* 33, 4 (2014), 118.

Paul Kubelka. 1948. New contributions to the optics of intensely light-scattering materials. Part I. *Josa* 38, 5 (1948), 448–457.

Kanti V Mardia and Peter E Jupp. 2009. *Directional statistics*. Vol. 494. John Wiley & Sons.

Serge Mazaauric, Mathieu Hébert, Lionel Simonot, and Thierry Fournel. 2014. Two-flux transfer matrix model for predicting the reflectance and transmittance of duplex half-tone prints. *JOSA A* 31, 12 (2014), 2775–2788.

Daniel Meneveau, Benjamin Bringier, Emmanuelle Tauzia, Mickaël Ribardière, and Lionel Simonot. 2017. Rendering rough opaque materials with interfaced lambertian microfacets. *IEEE transactions on visualization and computer graphics* 24, 3 (2017), 1368–1380.

Luis L Sánchez-Soto, Juan J Monzón, Alberto G Barriuso, and José F Cariñena. 2012. The transfer matrix: A geometrical perspective. *Physics Reports* 513, 4 (2012), 191–227.

Lionel Simonot, Roger D Hersch, Mathieu Hébert, and Serge Mazaauric. 2016. Multilayer four-flux matrix model accounting for directional-diffuse light transfers. *Applied optics* 55, 1 (2016), 27–37.

Brian Slovick, Zachary Flom, Lucas Zipp, and Srini Krishnamurthy. 2017. Transfer matrix method for four-flux radiative transfer. *Applied optics* 56, 21 (2017), 5890–5896.

Jos Stam. 2001. An illumination model for a skin layer bounded by rough surfaces. In *Rendering Techniques 2001*. Springer, 39–52.

George Gabriel Stokes. 1860. On the intensity of the light reflected from or transmitted through a pile of plates. *Proceedings of the Royal Society of London* 11 (1860), 545–556.

HC Van de Hulst. 1980. Multiple light scattering.

Bruce Walter, Stephen R Marschner, Hongsong Li, and Kenneth E Torrance. 2007. Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques*. Eurographics Association, 195–206.

Andrea Weidlich and Alexander Wilkie. 2007. Arbitrarily layered micro-facet surfaces.. In *GRAPHITE*, Vol. 7. 171–178.

Philippe Weier and Laurent Belcour. 2020. Rendering Layered Materials with Anisotropic Interfaces. *Journal of Computer Graphics Techniques (JCGT)* 9, 2 (20 June 2020), 37–57. <http://jcgt.org/published/0009/02/03/>

Mengqi Mandy Xia, Bruce Walter, Christophe Hery, and Steve Marschner. 2019. Gaussian Product Sampling for Rendering Layered Materials. In *Computer Graphics Forum*. Wiley Online Library.

Kun Xu, Wei-Lun Sun, Zhao Dong, Dan-Yong Zhao, Run-Dong Wu, and Shi-Min Hu. 2013. Anisotropic Spherical Gaussians. *ACM Transactions on Graphics* 32, 6 (2013), 209:1–209:11.

Tomoya Yamaguchi, Tatsuya Yatagawa, Yusuke Tokuyoshi, and Shigeo Morishima. 2019. Real-time Rendering of Layered Materials with Anisotropic Normal Distributions. In *SIGGRAPH Asia 2019 Technical Briefs*. ACM, 87–90.

Tizian Zeltner and Wenzel Jakob. 2018. The layer laboratory: a calculus for additive and subtractive composition of anisotropic surface reflectance. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 74.