



HAL
open science

Numerical Analysis of Layered Materials Models

Mégane Bati, Romain Pacanowski, Pascal Barla

► **To cite this version:**

Mégane Bati, Romain Pacanowski, Pascal Barla. Numerical Analysis of Layered Materials Models. [Research Report] Institut d'Optique Graduate School; Université de Bordeaux; CNRS; INRIA. 2019. hal-02157966v2

HAL Id: hal-02157966

<https://inria.hal.science/hal-02157966v2>

Submitted on 17 Jul 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.

Numerical Analysis of Layered Material Models

Mégane Bati¹, Romain Pacanowski², Pascal Barla³
¹IOGS ²Université de Bordeaux - CNRS ³Inria

July 17, 2019

Abstract

Most real-world materials are composed of multiple layers, whose physical properties impact the appearance of objects. The accurate reproduction of layered material properties is thus an important part of physically-based rendering applications. Since no exact analytical model exists for arbitrary configurations of layer stacks, available models make a number of approximations.

In this technical report, we propose to evaluate these approximations with a numerical approach: we simulate BRDFs and BTDFs for layered materials in order to compare existing models against a common reference. More specifically, we consider 60 layered material configurations organized in three categories: plastics, metals and transparent slabs. Our results show that: (1) no single model systematically outperforms the others on all categories; and (2) significant discrepancies remain between simulated and modeled materials. We analyse the reasons for these discrepancies and introduce immediate corrections that improve models accuracy with little effort. Finally, we provide a few challenging cases for future layered material models.

1 Introduction

The modeling of surface scattering is an essential part of physically-based rendering. For applications ranging from visual effects and video games to architecture and design, the choice of material model depends on a number of factors: physical accuracy of course, but also ease of use and efficiency. Material models thus rely on a number of approximations in order to strike the right balance between these competing factors. The chief approximation is geometric optics, in which case material models are usually based on the microfacet theory [TS67]. This ignores effects due to wave interferences, even though in practice they may still be re-incorporated in the microfacet framework (e.g., [HP17, BB17]).

Most materials are composed of multiple layers, which makes the modeling of surface scattering even more challenging. Practical models used in production (e.g., [GPA⁺19]) rely on simple (linear) combinations of layers, which potentially introduce severe physical approximations. As detailed in Section 2, only a few analyti-

cal layered material models exist in the literature, with the *a priori* goal of reaching the highest possible level of physical accuracy while remaining in the context of geometric optics. However, each of these models necessarily makes approximations to light transport in layered material structures, and the impact of these approximations remains, to date, unclear.

The main objective of this paper is to investigate *a posteriori* which layered material models provide results closest to the ground truth in which material configurations. We draw our inspiration from the study of Ngan et al. [NDM05], where they compared various material models to measured materials coming from the MERL database [MPBM03]. Our approach differs in that the ground truth we compare to is not given by measured materials, but by simulated materials obtained by path tracing (i.e., with a geometric optics assumption). Relying on simulation confers several advantages: the same set of material parameters is used in both simulations and models without requiring fitting; the light is perfectly collimated and lateral transport can be ignored; the scattering profiles may be analyzed on a per-layer basis, providing for a better understanding of the impact of model approximations.

In practice, we restrict our attention to three representative classes of layered materials, differing only by the type of their base layer: either diffuse, conductor or dielectric. We chose a total of 60 layered material configurations (12 for diffuse bases, 36 for conductors, 12 for dielectrics), covering a range of material effects that we have found to be of interest: these mostly include BRDFs, but also BTDFs in the case of dielectric bases. Not all existing analytical models are able to accommodate these different types of layered materials: we focus on two models, each declined in two variants (Section 3) in our study. We not only provide a quantitative evaluation (Section 4) but also analyze the design choices that lead to model inaccuracies and propose immediate improvements as well as challenges for future models (Section 5).

2 Previous work

As previously mentioned, we consider material models in a geometric optics context. More specifically, we assume that a rough interface between two media is prop-




Figure 1: We study how the approximations made by layered material models impact their accuracy, and ultimately material appearance. Here we compare four models side by side with our reference simulation on a frosted metal – one of the 60 material configurations we have considered in our study. This specific choice is particularly problematic for the model of Weidlich and Wilkie [WW09], which creates oddly-colored reflections away from normal incidence. This occurs to be a bug in the implementation, as mentioned in Section 3.1. The variant of Elek [Ele10] follows the original model [WW07] in this respect, and is thus devoid of these artefacts; nevertheless, it clearly overestimates the intensity of the metallic base. Belcour’s models [Bel18] (forward and symmetric) produce more accurate results, even though the intensity of the metallic base remains slightly higher. They still deviate from the reference simulation, especially at grazing angles as seen for instance at the bottom of the spheres. Our analysis in BRDF (and BTDF) space provides explanations for such departures from the reference.

erly described by a distribution of microfacets. Another assumption is that for layered materials, the distance between interfaces is sufficiently large (at least a few microns) to prevent interference effects. Despite these limitations, physically-based layered BSDF models are still expected to conserve energy and obey reciprocity¹.

Microfacet theory In many material models commonly used in Computer Graphics, a rough surface is assumed to be composed of a collection of tiny flat Fresnel reflectors (i.e., microfacets) at the microscopic scale. The orientations of these microfacets is defined statistically through a distribution function, which should also control the statistical visibility of microfacets for physical consistency [Hei14]. Historically, microfacet-based models have first used the Beckmann distribution [TS67, CT82]. Several alternative distributions have been proposed in the literature. In modern renderers, the GGX distribution [TR75, WMLT07] remains the solution of choice as it strikes a good balance between physical plausibility (long tails) and artistic control. More control over distribution tails may be achieved by a complex distribution (e.g., [LKYU12]) or a combination of simple distributions (e.g., [BPV18]), at the expense of an increased number of parameters.

Layered models Previous work on layered models usually consider material layers to consist of either interfaces or media, stacked in a plane-parallel configuration. Except for simple configurations such as stackings of perfectly smooth interfaces [Yeh05], no exact solution is known. Hence existing models necessarily make a number of assumptions and simplifications, as summarized in Table 1.

¹In a unidirectional path tracer, the reciprocity constraint may be lifted.

The model of Weidlich and Wilkie [WW07] recursively combines arbitrary BRDFs that represent successive interfaces, while handling absorption in media through the Beer-Lambert law. The base layer is considered opaque, and may be defined with any BRDF (e.g., Lambertian or conductor). Variants of this model have been proposed by the same authors [WW09, WW11] and by Elek [Ele10], where the combination of layers follows different approximations, as detailed in Section 3.1. These models are both simple and versatile, but are limited to the modeling of BRDFs (not BTDFs) and do not consider the contribution of multiple scattering (inside or among layers).

In contrast, the method of Guo et al. [GQGP17] is able to handle BTDFs, and in principle it also accounts for the main inter-reflections (though no actual renderings are provided in the paper). However, the model is limited to transparent media bounded by rough interfaces, themselves defined by von Mises-Fisher distributions. This means that the model cannot handle Lambertian base layers, or distributions with longer tails like GGX. Moreover, the resulting BSDF does not obey reciprocity.

Belcour [Bel18] takes a similar approach but proposes a statistical representation that accommodates GGX distributions for interface layers and forward scattering inside medium layers. The method works as if a bundle of rays were traced through the layer stack, with the bundle statistics (energy, mean and variance) updated at each layer event (reflection, refraction, absorption and forward scattering). A compelling feature of the model is its ability to handle all inter-reflections among layers through an extended adding-doubling procedure. Two variants are proposed as detailed in Section 3.2: the forward model, which is not reciprocal; and the symmetric model, which obeys reciprocity but is not readily adapted to handle BTDFs. None of the two variants can handle Lambertian base layers.

Finally, some models focus on transmission through rough surfaces. This is the case of the slab model of Dai et al. [DWL⁺09], where analytical formula are provided for two special cases (smooth-on-rough and rough-on-smooth interfaces); for general transparent slabs these two extreme configurations are interpolated in a heuristic way. The authors do not provide an analytic solution for the BRDF part of these same slabs, only the BTDF. For transparent objects of more complex shapes, de Rousiers et al. [dRBS⁺11] introduce a spherical Gaussian approximation that is best suited for real-time rendering.

Numerical techniques When physical accuracy is of primary interest, one may turn to the simulation of light transport in layered materials, with the downside of greatly increased rendering times. Such simulations have been presented for specific families of materials in the literature, such as wet surfaces [JLD99], skin [HK93, Sta01], and car paints [EKM01, EOO16] for instance. Recently, a solution targeting general layered materials has been proposed by Guo et al. [GHZ18], where they rely on a position-free light transport simulation in the layer stack to reduce variance, as well as to introduce spatial geometric variations through normal maps. However, the method remains costly to evaluate.

An efficient alternative is the Layer Lab system [JdJM14, ZJ18], which can accommodate arbitrary combinations of interfaces (even with anisotropic distributions) and media. It combine layers with adding-doubling formula, each layer being expressed in a sparse Fourier-based representation. The method is particularly efficient for modeling rough reflection or transmission through multiple layers, but requires longer precomputation times and storage requirements when smoother interfaces are involved; ringing artifacts may appear otherwise, especially at grazing angles. Perfectly smooth interfaces (i.e., Dirac distributions) are thus not handled. Even though the Fourier-based representation can be directly instantiated from Beckmann interfaces, it requires to go through the MERL format to deal with other types of interfaces (e.g., GGX). To the best of our knowledge, this only works for reflection; hence such MERL-based interfaces may only be used as opaque base layers. We provide more details of our experiments with LayerLab in Appendix B.

In this paper, we have chosen to rely on our own numerical simulation using path tracing in order to gain maximal flexibility for comparison purposes. As opposed to previous work where comparisons are most often made on rendered images, we evaluate layered material models directly on BRDF or BTDF slices, which permits to better understand the impact of model approximations.

3 Models

We have chosen to compare layered material models that can accommodate GGX distributions at interfaces, due

Properties	[WW09] [Ele10]	[GQGP17]	[Bel18] fwd sym	[JdJM14] [ZJ18]
Dirac	✓	✓	✓	✗
Rough interfaces	✓	vMF	GGX	Beckmann MERL*
Slab	✗	✓	✓ ✗	✓
Medium	absrp.	absrp.	absrp. fwd scatt.	✓
Lambert	✓	✗	✗	✓
Inter-reflections	✗	✓ [†]	✓	✓
Reciprocity	✓	✗	✗ ✓	✓

Table 1: Properties of analytical layered material models and of the LayerLab system. *Interfaces provided in MERL format are only defined for reflection. [†]The authors provide formula for incorporating inter-reflections in their model, but do not produce results.

to their wide adoption in Computer Graphics applications. Hence we will not consider the model of Guo et al. [GQGP17] since it relies on the von Mises Fisher (vMF) distribution, which lacks the heavier tails of the GGX distribution. This is discussed in greater details in Appendix D where we provide a simple ad-hoc formula to convert between GGX and vMF distributions, and show a comparison (Figure 15).

Looking at Table 1, this leaves us with two variants of two models. They all rely on microfacet theory; a BRDF is then given by:

$$f_r(\mathbf{l}, \mathbf{v}) = \frac{D(\mathbf{h})G(\mathbf{l}, \mathbf{v})F(\mathbf{l} \cdot \mathbf{h})}{4|\mathbf{l} \cdot \mathbf{n}||\mathbf{v} \cdot \mathbf{n}|}, \quad (1)$$

where D is the microfacet distribution, G is the geometric attenuation factor, F is the Fresnel reflectivity, \mathbf{l} is the incoming (or light) direction, \mathbf{v} is the outgoing (or view) direction, \mathbf{n} is the geometric normal and $\mathbf{h} = \frac{\mathbf{l} + \mathbf{v}}{\|\mathbf{l} + \mathbf{v}\|}$ is the half-way vector.

3.1 Weidlich & Wilkie’s model and variants

The original model of Weidlich and Wilkie [WW07] accommodates any type and number of interfaces through a recursive combination approach. Focusing on a pair of interfaces, it is defined as:

$$\begin{aligned} f_r^{ww}(\mathbf{l}_0, \mathbf{v}_0) &= f_{r_0}(\mathbf{l}_0, \mathbf{v}_0) + T_{01} f_{r_1}(\mathbf{l}_1, \mathbf{v}_1) a t; \quad (2) \\ a &= e^{-\sigma_a d \left(\frac{1}{|\mathbf{n} \cdot \mathbf{l}_1|} + \frac{1}{|\mathbf{n} \cdot \mathbf{v}_1|} \right)}, \\ t &= (1 - G) + T_{10} G. \end{aligned}$$

The layered BRDF model f_r^{ww} is obtained as a combination of the BRDFs f_{r_0} and f_{r_1} characterizing the top and bottom interfaces respectively. While the former is evaluated in the directions \mathbf{l}_0 and \mathbf{v}_0 like f_r^{ww} , the latter is evaluated in the transmitted directions \mathbf{l}_1 and \mathbf{v}_1 using Snell’s law, and attenuated by: the Fresnel transmissivity T_{01} ; Beer-Lambert attenuation a based on the absorption

Variant	T_{01}	$(\mathbf{l}_1, \mathbf{v}_1)$	f_{r_1}	t
[WW09]	\mathbf{h}_0	\mathbf{h}_0	$(\mathbf{l}_0, \mathbf{v}_0)$	$G(\mathbf{l}_1, \mathbf{v}_1)$
[Ele10]	\mathbf{h}_0	\mathbf{n}	$(\mathbf{l}_1, \mathbf{v}_1)$	$G(\mathbf{l}_0, \mathbf{v}_0)$

Table 2: Comparison of two variants of Weidlich and Wilkie’s model [WW07]. The first two columns indicate the reference interface used in Fresnel and Snell formula. The last two columns describe when the refracted rays pair $(\mathbf{l}_1, \mathbf{v}_1)$ is effectively used.

coefficient σ_a and layer depth d ; an approximation t of the effect of total internal reflection (TIR) based on the geometric attenuation factor G from microfacet theory. The recursion is applied to f_{r_1} through Equation 2 when more than two layers are considered.

This model makes a number of obvious approximations, as apparent in Table 1: it ignores inter-reflections among layers, and does not consider transmission through the layer stack (i.e., BTDFs).

A more subtle approximation lies in the choice of the interface normal used for computing Fresnel transmissivity as well as \mathbf{l}_1 and \mathbf{v}_1 through Snell’s law: one may either use the geometric normal \mathbf{n} or the half-way vector \mathbf{h}_0 (i.e., the top microfacet normal). Different solutions have been chosen in two variants of the model [WW09, Ele10], as summarized in Table 2. While both variants rely on \mathbf{h}_0 for computing transmissivity, they differ in the way they compute the refracted vectors. In the limit of a perfectly smooth top interface, the refraction of rays based on \mathbf{h}_0 does not make physical sense (unless $\mathbf{h}_0 = \mathbf{n}$), which is problematic for the variant of Weidlich and Wilkie [WW09].

Also shown in the table are the cases where the refracted ray pair $(\mathbf{l}_1, \mathbf{v}_1)$ is effectively used. Perhaps surprisingly, the variant of the authors [WW09] does not make use of the refracted directions when evaluating the BRDF f_{r_1} of the bottom interface²³. The computation of G in the TIR term t also differs.

Additionally, the variant of Elek uses a post-hoc modification of the roughness of the microfacet distribution on the bottom interface: $\alpha_1 \leftarrow \max(\alpha_0, \alpha_1)$. In a recent variant [WW11], a similar modification is proposed, albeit specific to the shininess parameter of the Blinn distribution; hence we do not consider it here. The goal of these heuristics is to compensate for the under-estimation of the top interface roughness in the original model [WW07]. Indeed, since only a pair of rays is transmitted through the top interface, whether it is smooth or rough does not have much impact on the BRDF, unless it is somehow re-incorporated in the roughness of the bottom interface as in the discussed variants.

²See slide 46 on page 8 of [WW09], or page 39 of [WW11].

³According to the authors, this is actually a bug in the implementation. However, Elek’s variant [Ele10] correctly uses refracted rays, so there is no need to implement another version. Note that this has no impact on plastic configurations, defined in Section 4.1.

3.2 Belcour’s models

The approach of Belcour [Bel18] expresses a layered BRDF as the sum of K GGX lobes (in practice, K is the number of layers):

$$f_r^B(\mathbf{l}_0, \mathbf{v}_0) = \sum_{k=0}^{K-1} e_k(\mathbf{l}_0) \rho_k(\mathbf{l}_0, \mathbf{v}_0; \alpha_k(\mathbf{l}_0)), \quad (3)$$

where e_k stands for the directionally-dependent energy of the k -th lobe, and ρ_k is of the form of Equation 1, except that the roughness α_k of the GGX distribution D is allowed to vary with the incoming direction as well. Each lobe is implicitly pointing in the direction of the reflected incoming direction; hence only e_k and α_k must be computed from the physical parameters of the layer stack. This is done by first tracking ray bundles and updating their directional statistics with atomic operators, then combining layers by an extended version of the adding-doubling procedure, which handles inter-reflections among layers (see the paper for details).

Since the approach relies on GGX lobes, it cannot handle Lambertian layers. One might argue that a similar diffusive effect could be obtained via scattering in a medium layer. However, even though Belcour’s model does handle scattering, it is limited to forward propagation. This means that a bundle of rays entering a base scattering medium layer will never exit the surface on reflection. We demonstrate this limitation in Appendix E.

The model necessarily makes a number of approximations. The energy e_k must be updated at an interface layer according to both Fresnel reflectance (i.e., indices of refractions) and roughness. This is made possible by precomputing the directional albedo of the BRDF of an interface – called FGD – and making the approximation that it is decoupled from incident radiance. In practice, FGD is stored in a 4D table, parametrized by elevation, roughness and a complex index of refraction. To avoid further numerical approximations due to interpolation in this table, we precompute it for the specific material parameters used in our evaluation. Note that Belcour uses the model of Heitz et al. [HHdD16] for precomputing a table FGD^∞ that accounts for inter-reflections at the interface; however, due to numerical issues potentially violating energy conservation, we do not consider that solution. The roughness α_k is not directly updated in the propagation of a ray bundle: its variance is. A mapping between the GGX roughness parameter and the variance is thus employed, which might include additional approximations, especially at high roughness values.

Another limitation of this so-called *forward* model is that it is not reciprocal. Belcour suggests an alternative *symmetric* model, whereby the half-way vector at the top interface \mathbf{h}_0 is used in place of the geometric normal \mathbf{n} when updating directional statistics. Both e_k and α_k in Equation 3 then become symmetric across a flip of \mathbf{l}_0 with \mathbf{v}_0 , which makes the BRDF reciprocal. From a physical

point of view, this post-hoc symmetrization not only assumes that \mathbf{l}_0 and \mathbf{v}_0 share the same microfacet as in the model of Weidlich and Wilkie [WW07], but also that the same interface normal (i.e., \mathbf{h}_0) will be used at all subsequent events occurring at interfaces below the top layer. In spite of this approximation, the symmetric model seems to produce satisfying results in Belcour’s paper; we thus also include it in our evaluation.

4 Evaluation

We have implemented a virtual gonio-spectrophotometer in order to evaluate the impact of the model approximations described in the previous section. It simulates light transport in the layer stack and performs intensity estimation on an hemisphere of outgoing directions. In its current version, it is restricted to isotropic materials since we only compare to isotropic layered BRDFs.

We have also considered using the python-based toolbox version of the LayerLab system [Jak15], but found it limiting for two reasons, already noted in Table 1. First, it is not possible to *directly* express a GGX interface in the Fourier-based representation; one must instead store the corresponding BRDF in the MERL format, which can then be projected onto the internal representation. Unfortunately, this only works for BRDFs, not BTDFs; as a result, it does not seem possible to use a GGX interface except for the base layer of an opaque material. Second, the system is prone to an increasing number of ringing artifacts when decreasing interface roughness, which can only be alleviated by increasing the number of basis coefficients. As detailed in Appendix B, the LayerLab toolbox is impractical for such material configurations, and thus cannot be used in our evaluation.

4.1 Methodology

In this section we describe the choices we made in the design of our virtual gonio-spectrophotometer. Additional implementation details are provided in Appendix B.

A material is defined by a stack of layers of three types: Lambertian diffusers, rough interfaces and media. A Lambertian layer is parametrized by its albedo ρ . Rough interfaces are based on the model of Walter et al. [WMLT07], parametrized by a roughness α and a complex index of refraction $\eta + i\kappa$. We do not account for inter-reflections at interfaces, since none of the models we consider properly account for them; hence we restrict roughness to intermediate values (see Section 4.2). A medium layer is parametrized by its absorption coefficient σ_a and depth d . We do not consider scattering in a medium layer, since only Belcour’s models implement this feature for the restricted case of forward propagation. We thus rely on Beer-Lambert law, which is in practice parametrized by the optical depth $\tau = \sigma_a d$. Nevertheless, we do provide a specific comparison of scattering in medium layers in




Figure 2: The two parametrizations used to store BRDF or BTDF slices in our simulations. They are both shown for a same incident elevation of 60° , and visualized with a reduced resolution of 10×10 cells. The **classic-slice** parametrization (a) is invariant to \mathbf{l}^m , whereas the **half-slice** parametrization (b) concentrates more cells in the specular direction with respect to \mathbf{l}^m .

Appendix E, this time using the LayerLab system as it is particularly well adapted to such configurations. Apart from α , all material parameters are defined as RGB vectors. We could have instead used spectral material parameters, but preferred to retain RGB channels to facilitate comparisons with the models of Section 3.

Our light transport simulation itself relies on a Monte-Carlo path-tracer (e.g., [PJH16]): N rays are sent onto the layer stack from one of M incoming direction \mathbf{l}^m ($m \in [0..M-1]$); they propagate by unidirectional path tracing until they exit the stack either by reflection or transmission; they are finally accumulated in a hemisphere of directions and their density is used to either compute a BRDF slice f_r^m or a BTDF slice f_t^m . The set of M slices constitutes a raw simulated BRDF f_r^* or BTDF f_t^* . The process is parallelized for each of the M incoming lighting directions.

The accumulation of rays on the hemisphere of outgoing directions for a fixed incident direction \mathbf{l}^m may be done in two different parametrizations. In the **classic-slice** parametrization (Figure 2(a)), the throughput of a light path exiting in a direction \mathbf{v} is accumulated at angular coordinates (θ_v, ϕ_v) , irrespective of \mathbf{l}^m . In the **half-slice** parametrization (Figure 2(b)), \mathbf{v} is first mapped to a halfway vector $\mathbf{h}(\mathbf{v}) = \frac{\mathbf{v} + \mathbf{l}^m}{\|\mathbf{v} + \mathbf{l}^m\|}$, and the light path throughput is accumulated at angular coordinates $(\sqrt{\theta_h \frac{\pi}{2}}, \phi_h)$. Both slice parametrizations are discretized into 90×360 cells. The choice of a non-linear mapping for θ_h in the **half-slice** parametrization is borrowed from the way BRDF data are stored in the MERL database [MPBM03]. As detailed in Section 4.2, we use **classic-slice** for predominantly rough materials, while **half-slice** is used for more specular materials since it better represents thin outgoing lobes in the specular direction.

In order to perform a fair comparison with material models, we need to integrate their response at the specific set of incoming directions \mathbf{l}^m , on the chosen slice parametrization, using the same resolution. This is done by discretizing each slice cell in 5×5 sub-cells, evaluating




Figure 3: Summary of layered material configurations considered in our evaluation. The top layer acts as a coating onto one of three base layers: Lambertian, conductor, and dielectric. The absorption coefficient τ and the diffuse albedo ρ may be chromatic or achromatic. The imaginary part κ of the complex index of refraction is always chromatic. Colors in the diagram indicate the tint for the chromatic configurations.

the BRDF or BTDF at sub-cell centers, and accumulating their contribution weighted by the solid angle of each sub-cell. In the following, we directly use these raw integrated data for comparison. For rendering, we simulate or integrate materials for a large number M (typically 30) of BRDF or BTDF slices. Given a (\mathbf{l}, \mathbf{v}) pair, the two slices corresponding to the incoming directions closest to \mathbf{l} are interpolated at render time. When using the `half-slice` parametrization, we simply use the half-way vector of \mathbf{l} and \mathbf{v} to lookup the slices in the interpolation process.

4.2 Results

We establish a representative set of layered material configurations for our evaluation, as summarized in Figure 3.

We consider three categories of materials: plastics, metals, and transparent slabs. They all share the same type of top interface: a coating layer, controlled by an optical depth τ (colored or achromatic), a constant index of refraction η_0 and a variable roughness α_0 . Plastics use a Lambertian base layer, which is controlled by its albedo ρ (colored or achromatic). Metals are obtained by using a conductor base layer, controlled by its complex index of refraction $\eta_1 + i\kappa_1$ (we only vary η_1) and a variable roughness α_1 . Transparent slabs are achieved with a dielectric base layer of index 1.0, controlled by its variable roughness α_1 .

For the purpose of evaluation, we trace $N = 160$ million rays for each of $M = 7$ incident lighting directions. The resulting BRDF slices are accumulated in the `classic-slice` parametrization whenever $\alpha_0 = 0.2$ and in the `half-slice` parametrization otherwise. We compute the RMSE (Root Mean Squared Error) on a single color channel of, say Weidlich and Wilkie’s model [WW09]:

$$E_r^{ww} = \sqrt{\frac{\sum_m \int_{\Omega^+} \|f_r^*(\mathbf{l}^m, \mathbf{v}) - f_r^{ww}(\mathbf{l}^m, \mathbf{v})\|^2 \cos \theta_v d\mathbf{v}}{M}}, \quad (4)$$

where $f_r^*(\mathbf{l}^m, \mathbf{v}) := f_r^m(\mathbf{v})$ and the integral over the upward hemisphere Ω^+ is discretized according to the choice

of slice parametrization. We use similar formula for Elek’s variant [Ele10], and Belcour’s models [Bel18]. For concision, we average these errors over color channels as shown in Figure 4. We have also found pertinent to discard direct reflections off the first interface in error computation for two reasons: first, such reflections can be computed exactly and should thus be error-free; and second, it avoids the common pitfall of over-emphasized highlights compared to more diffuse reflections in the RMSE.

In the following, we evaluate the performance of layered material models for each category independently. The analysis of the impact of different model approximations is delayed to Section 5. Beside error plots, we also report detailed results by selecting a few interesting materials and showing a pertinent subset of incident elevations. The full set of comparisons is provided in supplemental material, using a HTML interface that facilitates browsing through the space of material configurations.

Plastics There are 12 different plastic material configurations: (2 for τ) \times (3 for α_0) \times (2 for ρ). Only the variants of Weidlich and Wilkie’s model accommodate this material category. Since they are exact for reflection off the first interface (Equation 2), we compute errors for *internal paths only*.

As shown in Figure 4(a), Weidlich and Wilkie’s model is superior to Elek’s in terms of RMSE for this category of materials. The errors of both models increase when either ρ or τ increase, which is likely due to the occurrence of inter-reflections.

However, when looking at actual BRDF slices, the evaluation is not as clear-cut, as shown in Figure 5 on one particular plastic material configuration. As expected, the specular peaks of the two variants match, but their diffuse lobes differ in intensity, color and shape, both from the reference and from each others. Somehow surprisingly, both models overestimate the diffuse intensities, even though they ignore inter-reflections in the layer stack. Elek’s model overestimation is even more pronounced than Weidlich and Wilkie’s, which explains the difference in errors. Nevertheless, Elek’s model reasonably preserves the greenish color resulting from the choice of a colored diffuse base and a slightly absorptive coating layer; whereas the model of Weidlich and Wilkie produces a more yellowish-green color. The shapes of the diffuse lobes also differ from the reference: they are not radially-symmetric around the geometric normal \mathbf{n} (most noticeable at grazing angles) and show only slight variations with α_0 (see the supplemental material for comparisons on other values of α_0). More specifically, the lobe shape in Weidlich and Wilkie’s model is closer to the $\alpha_0 = 0.001$ configuration, whereas that of Elek’s is similar to the $\alpha_0 = 0.2$ configuration (i.e., Lambertian-like).

Metals There are 36 different metallic material configurations: (2 for τ) \times (3 for α_0) \times (2 for α_1) \times (3 for η_1).




Figure 4: All plots show RMSE for the three material categories we consider, for all three values of α_0 and various models (see legend at top). Reflection off the first interface is *discarded* in plastic and metal categories. Errors are only provided for transmission in the case of slabs. The errors are averaged over color channels in all plots unless mentioned otherwise, and are shown on a logarithmic scale. The naming for material configurations (on the horizontal axis) is described in Appendix A. For plastics (a) Weidlich and Wilkie’s model is consistently more accurate than Elek’s. Metals (b) are separated into two groups. For $\alpha_1 = 0.2$, Belcour’s models provide the best results (the forward and symmetric variants are similar), followed by Elek’s model, and Weidlich and Wilkie’s model. For $\alpha_1 = 0.001$, Elek’s model achieves errors close to Belcour’s for $\alpha_0 = 0.05$, or even undistinguishable for $\alpha_0 = 0.2$. Our variant of Weidlich and Wilkie’s model (Equation 5) outperforms the other variants on plastics, as well as on metals when $\alpha_1 = 0.001$, but not when $\alpha_1 = 0.001$. In the case of slabs (c), the errors of Belcour’s forward model significantly increase with decreasing α_0 and/or α_1 .




Figure 5: Evaluation results for a plastic material ($\alpha_0 = 0.05$, τ and ρ both colored). We remap values using $\log(1 + \sqrt{f_r})$ for visualization purposes. The reference simulation is plotted with a dashed black curve in BRDF slice diagrams, whereas the models are drawn with solid colored curves (red [WW09] and orange [Ele10]). The last three columns show the BRDF slice with a top orthographic view, with the yellow arrow indicating the incoming direction I^i . The two models over-estimate diffuse lobe intensity, and differ in terms of color and shape with the reference.

As before, the variants of Weidlich and Wilkie’s model are exact for reflection off the first interface, and are thus not compared to the ground truth. However, this is not the case of Belcour’s models, as seen in Figures 6 through 8: none of his variants is exact, even though the symmetric variant is slightly more accurate.

If we instead focus on the BRDF corresponding to *internal paths*, then Belcour’s variants clearly produce the smallest RMSE as shown in Figure 4(b). This time Elek’s model produces more accurate results than Weidlich and Wilkie’s model. Its errors are comparable to Belcour’s errors when $\alpha_1 = 0.05$ (even identical with $\alpha_0 = 0.2$), which is likely due to its post-hoc roughness modification. In order to make sense of these patterns of errors, we next

study three material configurations of interest, which are similar to those considered in the original paper of Weidlich and Wilkie [WW07].

In Figure 6, we show a metallic paint, characterized by a smooth coating on top of a rough base. The model of Weidlich and Wilkie [WW09] exhibits an extra achromatic lobe at grazing angles, which is due to the implementation bug of using non-refracted vectors (recall Table 2). Elek’s model does not make this approximation, and as a result is devoid of such artefacts. The models of Belcour produce results similar to Elek’s in this case: they all tend to slightly overestimate the angular spread of the wide reddish lobe due to reflection off the conductor base layer.

In Figure 7, we show a frosted metal, characterized by a rough coating on top of a smooth base. This is the most problematic configuration for Weidlich and Wilkie’s model, which ignores the roughness of the coating layer, resulting in an overly sharp lobe shape (i.e., a reddish BRDF peak). The post-hoc modification of α_1 in Elek’s variant tries to correct this issue. However, this is not sufficient: the shape of the lobe is much better captured, even though it slightly departs from the reference toward normal incidence; but its color is off as the whole BRDF slice takes on a reddish tint. The models of Belcour capture the lobe shape very accurately at normal incidence, even though their shape starts to depart from the reference at moderately grazing incidences.

In Figure 8, we show a metallic patina, characterized by a rough coating on top of a rough base. The limitation of Weidlich and Wilkie’s model exhibited for frosted metals is here less pronounced at normal incidence, even though the artefacts at grazing angles are still present. The BRDF obtained with Elek’s model is identical to the frosted metal configuration by construction, due to the post-hoc modification of α_1 . As before, the lobe shapes of Belcour’s model differ from the reference at grazing angles; whereas lobe shapes in Elek’s model accurately reproduce




Figure 6: Evaluation results for a metallic paint material ($\alpha_0 = 0.001$, achromatic τ , $\alpha_1 = 0.2$, $\eta_1 = 1.45$). Four models are drawn in BRDF slice diagrams with solid colored curves (red [WW09], orange [Ele10], cyan [Bel18](fwd) and blue [Bel18](sym)). The model of Weidlich and Wilkie exhibits artifacts at grazing angles; all other models only slightly differ in the extent of the wide redish reflection lobe.




Figure 7: Evaluation results for a frosted metal material ($\alpha_0 = 0.2$, achromatic τ , $\alpha_1 = 0.001$, $\eta_1 = 1.45$). The model of Weidlich and Wilkie cannot reproduce the redish rough reflection lobe seen in the reference. The variant of Elek captures the shape of that lobe, but not its color. Belcour’s models accurately reproduce the expected lobe shape at normal incidence, but it departs from the reference toward grazing angles.




Figure 8: Evaluation results for a metallic patina material ($\alpha_0 = 0.2$, achromatic τ , $\alpha_1 = 0.2$, $\eta_1 = 1.45$). The model of Weidlich and Wilkie provides a better approximation compared to Figure 7, but still exhibits the grazing-angle artifacts of Figure 6. Elek’s model is identical (by construction) on metallic patina and frosted metal (see the supplemental material for a comparison on all θ_l). The two variants of Belcour’s model here depart from each other at grazing angles, both in terms of color and extent of the reflection lobe due to the conductive base.

the reference. Nevertheless, all models differ in terms of color from the reference. The forward and symmetric variants of Belcour also differ among themselves at grazing angles, the former exhibiting a wide achromatic lobe of

over-estimated intensity, the latter being off in terms of saturation.

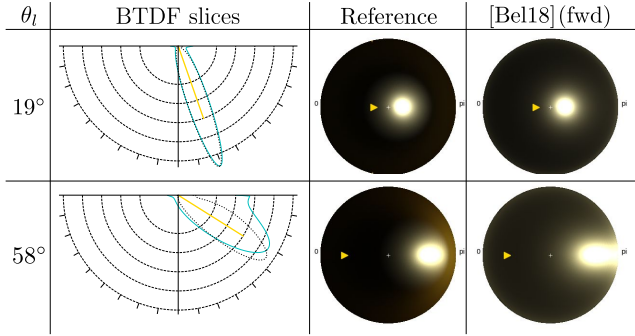


Figure 9: Evaluation results for a transparent slab material ($\alpha_0 = 0.001$, $\alpha_1 = 0.2$, colored τ). The BTDF model of Belcour [Bel18] (cyan solid curve) captures the shape of the transmitted lobe rather accurately around normal incidence (the lobe tails are less pronounced in the reference). The lobe shape departs significantly from the reference toward grazing angles: it is wider, points in a different direction, and does not reproduce the orange lobe.

Transparent slabs There are 12 different slab material configurations: (2 for τ) \times (3 for α_0) \times (2 for α_1). Only the forward model of Belcour accommodates this category of materials. We follow Belcour’s implementation⁴ and do not use any of the precomputed tables (*FGD* or *TIR*).

We focus here on BTDFs, and show errors with respect to the simulated reference in Figure 4(c). The RMSE markedly increases with decreasing α_0 and/or α_1 . Moreover, when τ is yellow-tinted (configurations *S1* and *S3*), the errors based on the blue channel are smaller than those based on luminance. This is likely due to the reduced impact of inter-reflections in an absorptive slab.

In Figure 9, we show comparisons for a smooth-on-rough yellow-tinted slab. The transmitted lobes in the simulation have moderately complex shapes: the main lobe exhibits a relatively sharp falloff, likely due to total internal reflection at critical angles inside the slab; while a secondary, more colored lobe appears towards grazing angles, due to inter-reflections in the slab. The model of Belcour accurately reproduces the reference lobe shape and color around normal incidence, except for the lobe tails which are less pronounced in the simulated results. However, at grazing angles, its shape significantly departs from the reference, as well as its direction; the secondary colored lobe is not present.

5 Analysis

In the previous section, we have shown that existing layered models exhibit notable departures from simulated references. Here we first explain the reasons for such departures, leveraging these insights to suggest simple immediate improvements (Sections 5.1 and 5.2). We then present challenges for future layered material models, identifying effects that should be reproduced in order to reach a higher physical accuracy (Section 5.3).

⁴We have used the `layered.dielectric` Mitsuba plugin.




Figure 10: Compared to Elek’s variant, our variant of Weidlich and Wilkie’s model (Equation 5, in pink) provides more accurate results for plastics (first row) and metallic paints (second row). For frosted metals (third row), both variants still depart significantly from the reference simulation. Our simple modification of Belcour’s model (Equation 6, in green) systematically improves its accuracy, especially for frosted metals (last row).

5.1 Variants of Weidlich and Wilkie’s model

The two variants of the original model of Weidlich and Wilkie [WW09] exhibit a common departure from the reference: they over-estimate the intensity of the lobe reflected off the base layer, be it a Lambertian or a conductor. This is unexpected since none of the variants handle inter-reflections; these intensities should instead be underestimated. We suggest this is due to a missing term in Equation 2: not only should f_{r_1} be evaluated in the refracted incoming and outgoing directions \mathbf{l}_1 and \mathbf{v}_1 , but the differential solid angle should also be modified. As detailed in Appendix C, this leads to an additional Jacobian term that has an analytical form; in the case of Equation 2, it is simply given by $\frac{1}{\eta_1^2}$. Note that such a Jacobian term had already been described in the literature (e.g., Hanrahan and Krueger [HK93] cite Ishimaru [Ish78][pp. 154-155]).

Contrary to the original paper [WW07], the variant of Weidlich and Wilkie [WW09] does not use refracted directions, which leads to odd BRDF results. For plastic materials, it exhibits angular variations in terms of colors and intensity, which are especially pronounced at grazing angles (see Figure 5). These are likely due to the choice of reference normal (i.e., \mathbf{h}_0) in the computation of the attenuation term a . Our choice of a yellow-tinted coating for this figure explains the color variations. For metallic materials, the model exhibits an additional lobe toward grazing angles (see Figures 6 and 8). This is due to the

base layer BRDF f_{r_1} being evaluated at $(\mathbf{l}_0, \mathbf{v}_0)$, reaching more grazing angles compared to $(\mathbf{l}_1, \mathbf{v}_1)$. The variant of Elek [Ele10] is devoid of these limitations as it properly uses the refracted directions $(\mathbf{l}_1, \mathbf{v}_1)$ to evaluate f_{r_1} .

Both variants choose \mathbf{h}_0 to evaluate the transmissivity T_{01} , which does not make physical sense when the top layer is smooth. We suggest another variant of Weidlich and Wilkie’s model that systematically uses the geometric normal \mathbf{n} in the computation of both T_{01} and $(\mathbf{l}_1, \mathbf{v}_1)$, and corrects intensity issues thanks to the Jacobian:

$$f_r^{ww'}(\mathbf{l}_0, \mathbf{v}_0) = f_{r_0}(\mathbf{l}_0, \mathbf{v}_0) + T_{01} f_{r_1}(\mathbf{l}_1, \mathbf{v}_1) \frac{a}{\eta_1^2} T_{10}, \quad (5)$$

where a is the same as in Equation 2 and we have removed the TIR term since we have not found sufficient physical justification for it. We also use Elek’s post-hoc roughness modification.

Our variant yields smaller RMSE on most material configurations, as seen in Figure 4. It yields a consistently smaller error on plastics, which is confirmed by the top row of Figure 10 and the first column of Figure 11. This is due to the inclusion of the Jacobian term J , which also makes the result slightly darker than the reference, as expected. Equation 5 also yields smaller errors on metallic paints and patinas, corresponding to $\alpha_1 = 0.2$ in Figure 4. As seen in Figures 10 and 11, the result for metallic paint is indistinguishable from the reference. For frosted metals (i.e., with $\alpha_1 = 0.001$ in Figure 4), Equation 5 yields smaller errors than Elek’s only when $\alpha_0 = 0.001$. For other values of α_0 , the two variants mainly differ in terms of color with the reference, and both are inferior to Belcour’s models as seen in Figures 10 and 11.

In summary, our variant remains limited in a number of respects: (1) inter-reflections are not considered, which produces slightly darker results; (2) the post-hoc roughness modification is not sufficient to account for lobe color and shape in frosted metals; (3) transmission is not handled.

5.2 Belcour’s models

As shown in Section 4.2, the models of Belcour [Bel18] depart from the reference simulation mainly toward grazing angles.

In metallic materials, we have found that the reflection off the top interface is approximate at grazing angles, especially with a rough coating. As seen in the BRDF slices of Figures 7 and 8, the symmetric variant (in blue) is slightly more accurate than the forward variant (in cyan). This departure from the reference is likely due to the approximation made by the FGD term, which is based on \mathbf{h}_0 instead of \mathbf{n} for the symmetric variant. Interestingly, when the two variants are compared to the reference for *internal paths* in Figure 4(b), they are similar in terms of RMSE. The remaining errors in this case are due to relatively small differences in the shape and color of the BRDF lobe at grazing angles. Regarding lobe colors, the forward variant produces desaturated colors at grazing angles, whereas

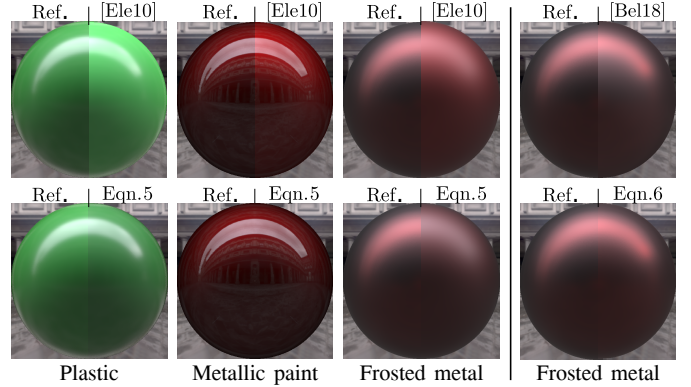


Figure 11: The differences in terms of BRDF observed in Figure 10 are visible in renderings. Equation 5 is closer to (albeit darker than) the reference for plastics. It is indistinguishable from the simulation in the case of metallic paints, but does not improve on Elek’s variant for frosted metals. Belcour’s forward model provides better results on frosted metals, and the simple fix of Equation 6 brings it closer to the reference.

the symmetric variant slightly over-estimates saturation (see Figure 8); this is most likely due to the approximation brought by the FGD term, this time on transmission. Regarding lobe shapes, both variants produce reflections off the base layer that are wider compared to the reference, which we attribute to the projection of BRDF lobes on GGX-based functions.

In slab materials, the difference in lobe shapes is more pronounced (see Figure 9). At normal incidence, refracted lobes have stronger tails in Belcour’s model compared to the reference, again due to projection on GGX lobes. Toward grazing angles, the shape of the refracted lobe departs even more from the reference. This is likely due to the multiple approximations that accumulate with successive transmission events in Belcour’s model: the conversion between roughness and lobe variance, the “fake” rough refraction, etc. The direction of refracted lobes also differ, which is due to the lack of updating mechanism for mean statistics in Belcour’s approach; as a result, reflected and refracted lobes always point in the specular direction. Finally, the orange secondary lobe observed in the reference is absent from the model results. This issue could be addressed by letting Belcour’s model output two lobes in this case; however, they would still be based on GGX functions and hence would depart from the reference lobe shape in this case.

We propose a trivial fix for issues found in reflection off the top interface, which simply consists in using f_{r_0} for the 0-th lobe:

$$f_r^{B'}(\mathbf{l}_0, \mathbf{v}_0) = f_{r_0}(\mathbf{l}_0, \mathbf{v}_0) + \underbrace{\sum_{k=1}^{K-1} e_k(\mathbf{l}_0) \rho_k(\mathbf{l}_0, \mathbf{v}_0; \alpha_k(\mathbf{l}_0))}_{f_{r_1}^B(\mathbf{l}_0, \mathbf{v}_0)}, \quad (6)$$

where the forward variant is used to compute $f_{r_1}^B$, since it provides slightly more accurate results, while being more adapted to real-time rendering applications [Bel18].

Errors for the model of Equation 6 are identical to Bel-

cour’s forward model in Figure 4 since internal paths are not modified in our variant. As shown in Figures 10 and 11, reflections off the first interface in the case of a frosted metal are properly corrected with our modified model. The effect of this correction on metallic patinas is very similar, but it is much more subtle on metallic paints.

5.3 Challenges for future work

A practical contribution of this technical report lies in the database of simulated layered BRDFs, of which we show a subset in the renderings of Figure 12. It could be used for validation in future models, and we plan to extend it with additional material configurations and properties (e.g., scattering in medium layers). A direct benefit of this database is that it permits to identify material configurations that lead to interesting BRDF or BTDF properties, as discussed in the following. These represent challenging results to reproduce for future layered material models.

Lobe shapes Different families of light paths in simulations may be studied as separate BRDF or BTDF components. In particular, for the material configurations studied in this paper, we propose a decomposition in terms of the number of events (reflection or transmission) on the top layer, starting from 1 for light paths that are directly reflected. As shown in Figure 13, this clearly exposes how the lobe shapes and colors of the various material components differ among each other. They represent a challenge since an accurate reproduction would potentially require different models for each component – or a complex-enough model.

Grazing-angle colors As shown in Figure 13(left), we have observed surprisingly-colored lobes at grazing angles for metallic materials, most notably when $\alpha_1 = 0.2$. The tint is significantly different from the red conductor base, even though the color remains weakly saturated and requires a specific light source placement to be clearly observed in renderings, as shown in Figure 14. This effect is due to the specific choices of refractive indices: $\eta_0 = 1.5$ and $\eta_1 = 1.2 + i(1.0, 0.01, 0.01)$. The base layer acts as a conductor for the red channel, but as a dielectric for the green and blue channels. As a result, Fresnel equations show markedly different behaviors among color channels, as seen at the top right of Figure 14; in particular, the two curves cross at $\theta \approx 51^\circ$. When the top interface is smooth, the most grazing angle that can be achieved for an incoming ray on the base layer corresponds to the critical angle of refraction $\theta_c \approx 42^\circ$ (vertical dashed line in Figure 14); hence few color effects are observed. When the top interface is rough, rays are more likely to reach incident angles past 51° , especially after multiple bounces; this yields the observed color effects.

Such effects cannot be reproduced by the models of Elek [Ele10] or Belcour [Bel18], since the directions they

consider never go passed the critical angle by construction. Interestingly, the model of Weidlich and Wilkie [WW09] does reach these angles since it uses non-refracted rays, but this strongly over-estimates the effect. The odd yellowish appearance of Weidlich and Wilkie’s result in Figure 1 is due to this issue: here we use a frosted metal with $\alpha_0 = 0.05$, $\alpha_1 = 0.2$, a yellow-tinted τ and, most importantly, $\eta_1 = 1.2$. We conjecture that these color effects are also responsible for the rise in RMSE in Figure 4(b).

6 Conclusions

We have conducted a numerical evaluation of two families of analytical BSDF models on 60 layered material configurations organized in three categories: plastics, metals and transparent slabs.

The two variants of Belcour [Bel18] have yielded the most accurate results for metals. The forward and symmetric variants yield very similar errors when the evaluation is restricted to internal paths. Dealing explicitly with reflection off the top layer corrects some of the issues observed at grazing angles. Nevertheless, the model still exhibits departures with the reference in terms of lobe colors and shapes at grazing angles (especially for BTDFs) and does not handle Lambertian base layers.

Even though the two variants of Weidlich and Wilkie’s model [WW09, Ele10] are less accurate, we have shown that with a few modifications, this approach yields quite accurate results, comparable or even better in some cases than the models of Belcour. It would thus be interesting to investigate whether the two approaches could be combined in a single model that could accommodate all material categories, and ideally ensure reciprocity.

We have made a number of simplifying assumptions in our simulation, which we would like to lift in future work. We have restricted simulations to three color channels in order to simplify comparisons with analytical models. However, physically-realistic materials should be defined spectrally; in particular, indices of refraction vary with wavelengths, sometimes significantly as in the case of conductors. It would also be interesting to consider more complex scattering at interfaces, including wavelength-dependent effects such as diffraction, and multiple scattering. Finally, we have not considered scattering in medium layers due to the limitations of existing models; in future work, we would like to extend our analysis to diffusion phenomena, which will likely require to simulate lateral transport in layer stacks (e.g., layered BSSRDF models).

References

[BB17] Laurent Belcour and Pascal Barla. A Practical Extension to Microfacet Theory for the Modeling of Varying Iridescence. *ACM Transactions on Graphics*, 36(4):65, July 2017.

[Bel18] Laurent Belcour. Efficient rendering of layered materials using an atomic decomposition with statistical op-




Figure 12: A subset of 5 out of the 60 materials from our database. Naming conventions are detailed in Appendix A.




Figure 13: BRDF and BTDF slices for separate bounce orders on a patina (same as Figure 8 with $\eta_1 = 1.2$) and a slab (same as Figure 9) respectively. The intensity for higher-order bounces is doubled for visualization. We also show isovalues in false colors to reveal the varying lobe shapes at different orders.




Figure 14: Comparison of two metal patina ($\alpha_0 = \alpha_1 = 0.2$, yellow-tinted τ , different η_1). Top left: when using $\eta_1 = 1.2$, a yellowish reflection occurs at grazing lighting angles, whereas it is absent when $\eta_1 = 1.9$. Bottom left: the effect is subtle but noticeable when rendering in complex lighting. Right column: these color effects are due to Fresnel reflectivities that differ among color channels (the cyan curves correspond to the blue and green channels).

- erators. *ACM Trans. Graph.*, 37(4):73:1–73:15, July 2018.
- [BPV18] Pascal Barla, Romain Pacanowski, and Peter Vangorp. A Composite BRDF Model for Hazy Gloss. *Computer Graphics Forum*, 37, 2018.
- [CT82] R. L. Cook and K. E. Torrance. A reflectance model for computer graphics. *ACM Trans. Graph.*, 1(1):7–24, January 1982.
- [dRBS⁺11] Charles de Rousiers, Adrien Bousseau, Kartic Subr, Nicolas Holzschuch, and Ravi Ramamoorthi. Real-time rough refraction. In *Symposium on Interactive 3D Graphics and Games, I3D '11*, pages 111–118 PAGE@7, New York, NY, USA, 2011. ACM.
- [DWL⁺09] Qiang Dai, Jiaping Wang, Yiming Liu, John Snyder, Enhua Wu, and Baining Guo. The Dual-microfacet Model for Capturing Thin Transparent Slabs. *Computer Graphics Forum*, 2009.
- [EKM01] Sergey Ershov, Konstantin Kolchin, and Karol Myszkowski. Rendering Pearlescent Appearance Based On Paint-Composition Modelling. *Computer Graphics Forum*, 2001.
- [Ele10] Oscar Elek. Layered materials in real-time rendering. In *Central European Seminar on Computer Graphics, CESC10*, 2010.
- [EOO16] S. Ergun, S. Önel, and A. Ozturk. A general microflake model for predicting the appearance of car paint. In *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations, EGSR '16*, pages 65–71, Goslar Germany, Germany, 2016. Eurographics Association.
- [GHZ18] Yu Guo, Miloš Hašan, and Shaung Zhao. Position-free monte carlo simulation for arbitrary layered bsdfs. *ACM Trans. Graph.*, 37(6), 2018.
- [GPA⁺19] Iliyan Georgiev, Jamie Portsmouth, Zap Andersson, Adrien Herubel, Alan King, Shinji Ogaki, and Frederic Servant. A surface standard. <https://autodesk.github.io/standard-surface/>, 2019.
- [GQGP17] J. Guo, J. Qian, Y. Guo, and J. Pan. Rendering thin transparent layers with extended normal distribution functions. *IEEE Transactions on Visualization & Computer Graphics*, 23(9):2108–2119, Sept. 2017.
- [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.
- [HK93] Pat Hanrahan and Wolfgang Krueger. Reflection from layered surfaces due to subsurface scattering. In *Proceedings of the 20th Annual Conference on Computer*

Graphics and Interactive Techniques, SIGGRAPH '93, pages 165–174, New York, NY, USA, 1993. ACM.

[HP17] Nicolas Holzschuch and Romain Pacanowski. A Two-Scale Microfacet Reflectance Model Combining Reflection and Diffraction. *ACM Transactions on Graphics*, 36(4):12, July 2017. Article 66.

[Ish78] A. Ishimaru. *Wave Propagation and Scattering in Random Media*. Number v. 2 in IEEE/OUP series on electromagnetic wave theory. Academic Press, 1978.

[Jak15] Wenzel Jakob. layerlab: A computational toolbox for layered materials. In *SIGGRAPH 2015 Courses*, SIGGRAPH '15, New York, NY, USA, 2015. ACM.

[JdJM14] Wenzel Jakob, Eugene d'Eon, Otto Jakob, and Steve Marschner. A comprehensive framework for rendering layered materials. *ACM Trans. Graph.*, 33(4):118:1–118:14, July 2014.

[JLD99] Henrik Wann Jensen, Justin Legakis, and Julie Dorsey. Rendering of wet materials. In *Proceedings of the 10th Eurographics Conference on Rendering*, EGWR'99, pages 273–282, Aire-la-Ville, Switzerland, Switzerland, 1999. Eurographics Association.

[LKYU12] Joakim Löw, Joel Kronander, Anders Ynnerman, and Jonas Unger. Brdf models for accurate and efficient rendering of glossy surfaces. *ACM Trans. Graph.*, 31(1):9:1–9:14, February 2012.

[MPBM03] Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. A data-driven reflectance model. *ACM Transactions on Graphics*, 22(3):759–769, July 2003.

[NDM05] Addy Ngan, Frédo Durand, and Wojciech Matusik. Experimental analysis of brdf models. In *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques*, EGSR '05, pages 117–126, Aire-la-Ville, Switzerland, Switzerland, 2005. Eurographics Association.

[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.

[Sta01] Jos Stam. An illumination model for a skin layer bounded by rough surfaces. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, pages 39–52, London, UK, UK, 2001. Springer-Verlag.

[TR75] T. S. Trowbridge and K. P. Reitz. Average irregularity representation of a rough surface for ray reflection. *J. Opt. Soc. Am.*, 65(5):531–536, May 1975.

[TS67] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces*. *J. Opt. Soc. Am.*, 57(9):1105–1114, Sep 1967.

[WMLT07] Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques*, EGSR'07, pages 195–206, Aire-la-Ville, Switzerland, Switzerland, 2007. Eurographics Association.

[WW07] Andrea Weidlich and Alexander Wilkie. Arbitrarily layered micro-facet surfaces. In *Proceedings of the 5th International Conference on Computer Graphics and Interactive Techniques in Australia and Southeast Asia*, GRAPHITE '07, pages 171–178, New York, NY, USA, 2007. ACM.

[WW09] Andrea Weidlich and Alexander Wilkie. Exploring the potential of layered brdf models. In *ACM SIGGRAPH ASIA 2009 Courses*, SIGGRAPH ASIA '09, pages 7:1–7:58, New York, NY, USA, 2009. ACM.

[WW11] Andrea Weidlich and Alexander Wilkie. Thinking in layers: Modeling with layered materials. In *SIGGRAPH Asia 2011 Courses*, SA '11, pages 20:1–20:43, New York, NY, USA, 2011. ACM.

[Yeh05] P. Yeh. *Optical Waves in Layered Media*. Wiley Series in Pure and Applied Optics. Wiley, 2005.

[ZJ18] Tizian Zeltner and Wenzel Jakob. The layer laboratory: A calculus for additive and subtractive composition of anisotropic surface reflectance. *ACM Trans. Graph.*, 37(4):74:1–74:14, July 2018.

A Full set of results

The full set of results is provided in the `html/` directory of the supplemental material, under the name `show_comparisons.html`. It permits to quickly browse through all three material categories, change parameters, observe simulation results and compare to analytical models. We provide two different views of BRDF slices: the RGB view permits to observe differences in colors, while the isovalues view more precisely reveals lobe shapes (for luminance only). We also present renderings for the plastic and metal categories. A separate html interface is provided to compare models to the reference on rendered images, under the name `rendering_compare_ldr.html`. We do not show rendered images for the slab category due to current limitations of our in-house renderer. This will be implemented in the near future and does not have any impact on the conclusions drawn in the paper, since we rely on a different software for the simulation of layered materials.

The naming conventions we use in Figures 4 and 12 of the paper are detailed in the following tables. The C and A symbols refer to the chromatic and achromatic configurations respectively, as defined in Figure 3 in the paper.

Name	τ	ρ
P_1	C	C
P_2		A
P_3	A	C
P_4		A

Table 3: Naming conventions for the plastic material category.

Name	τ	η_1
M_1	C	1.9
M_2		1.45
M_3	A	1.9
M_4		1.2
M_5	C	
M_6	A	

Table 4: Naming conventions for the metallic material category.

Name	τ	α_1
S_1	C	0.2
S_2	A	
S_3	C	0.001
S_4	A	

Table 5: Naming conventions for the slab material category.

B Implementation details

We considered using the LayerLab toolbox [Jak15] but we encountered serious technical problems: out of our 60 material configurations, 20 yielded segmentation faults (those with $\alpha_0 = 0.2$), while 20 others still exhibited significant ringing artifacts (those with $\alpha_0 = 0.001$) when using 1000 nodes in elevation and 4000 Fourier modes in azimuth. The toolbox also provides heuristics to compute nodes and modes that reduce those artifacts, but in our case it yielded 38436 nodes and 78156 mdoes, for which the Gauss-Lobatto method did not even converge. Nevertheless, we found the LayerLab system very well suited to materials having scattering layers, and thus used it as a reference in the additional comparisons (Section E).

We have thus opted for a dedicated path tracing simulator, implemented on the CPU and parallelized with OpenMP in terms of incident directions. Both the path tracing simulation in layer stacks and the integration of layered material models is done per BRDF (or BTDF) slice, for a given incoming direction. Each slice is subdivided in 90×360 cells. Simulation of a single slice using 160 million rays takes on average 50 seconds on a Intel(R) Core(TM) i7-4790K CPU @ 4.00GHz computer.

We make use of two BRDF slice parametrizations, which we have called `classic-slice` and `half-slice`. The latter is best adapted to specular materials as it concentrates the number of cells around the specular direction. It is based on the half-way vector and a non-linear parametrization along the θ_h dimension. With the BRDF slice resolution that we use in the paper, the cells corresponding to $\theta_h = 0$ span an elevation of $\Delta\theta_h = \frac{1}{90}^\circ$. This raises numerical issues as a prohibitive number of rays is required to reduce variance to reasonable values. In practice, we discard these cells in all our comparisons and assign them the average of neighbor cell values. However, the size of discarded cells remains much smaller than the width of the sharpest specular peak used in our analysis, even when $\alpha_0 = 0.001$.

Another difficulty arises with the `half-param` parametrization: since it depends on the incoming direction, a subset of the cells are totally or partially located under the upper hemisphere. We discard both types of cells for both simulated and integrated materials; hence the errors reported in the paper do not consider some of the extremely grazing configurations.

We have also encountered issues with the `layered-dielectric` Mitsuba plugin of Belcour [Bel18].

In our experiments, we have found that its BRDF part does not match the results of the `layered-forward` plugin when setting $\kappa_1 = 0$. We have also observed surprisingly large departures between the simulation and Belcour’s model on transmission at extremely grazing angles, as can be seen using the html interface. For technical reasons, the renderings provided in the html interface only use the `classic-slice` parametrization, with the drawback that configurations where $\alpha_0 = 0.001$ are not precisely reproduced (especially at grazing angles), even with a very high angular resolution. Note that this does not invalidate our analysis in any respect.

C Jacobian term for Weidlich and Wilkie’s model [WW07]

We express the BRDF f_{r_1} for the second interface in Equation 2 (without absorption a) as a function of $(\mathbf{l}_0, \mathbf{v}_0)$:

$$\begin{aligned} f_{r_1}(\mathbf{l}_0, \mathbf{v}_0) &= \frac{dL(\mathbf{v}_0)}{dE(\mathbf{l}_0)} = T_{01} \frac{dL(\mathbf{v}_0)}{dE(\mathbf{l}_1)} \\ &= T_{01} J f_{r_1}(\mathbf{l}_1, \mathbf{v}_1) T_{10}, \end{aligned}$$

where L and E denote radiance and irradiance respectively. The first line is obtained by using $dE(\mathbf{l}_1) = T_{01}dE(\mathbf{l}_0)$. The second line is obtained by first writing:

$$\begin{aligned} dE(\mathbf{v}_0) &= T_{10} dE(\mathbf{v}_1) \\ dL(\mathbf{v}_0) \cos \theta_0 d\mathbf{v}_0 &= T_{10} dL(\mathbf{v}_1) \cos \theta_1 d\mathbf{v}_1; \\ dL(\mathbf{v}_0) &= T_{10} J dL(\mathbf{v}_1), \end{aligned}$$

where we have used $dE(\mathbf{v}) = dL(\mathbf{v}) \cos \theta d\mathbf{v}$ in the second line, and introduced the Jacobian term $J = \frac{\cos \theta_1 d\mathbf{v}_1}{\cos \theta_0 d\mathbf{v}_0}$. If we now write explicitly the differential solid angle as $d\mathbf{v} = \sin \theta d\theta d\phi$, we have:

$$J = \frac{\cos \theta_1 d\theta_1 d\phi_1}{\cos \theta_0 d\theta_0 d\phi_0} \cdot \frac{\sin \theta_1}{\sin \theta_0}.$$

Using Snell’s law of refraction $\eta_0 \sin \theta_0 = \eta_1 \sin \theta_1$, we have that the second term is equal to $\frac{\eta_0}{\eta_1}$. We also obtain the following:

$$\begin{aligned} \eta_0 \frac{d(\sin \theta_0)}{d\mathbf{v}_0} d\mathbf{v}_0 &= \eta_1 \frac{d(\sin \theta_1)}{d\mathbf{v}_1} d\mathbf{v}_1 \\ \eta_0 \frac{d(\sin \theta_0)}{d\theta_0 d\phi_0} d\theta_0 d\phi_0 &= \eta_1 \frac{d(\sin \theta_1)}{d\theta_1 d\phi_1} d\theta_1 d\phi_1 \\ \eta_0 \cos \theta_0 d\theta_0 d\phi_0 &= \eta_1 \cos \theta_1 d\theta_1 d\phi_1. \end{aligned}$$

Hence the first term of the Jacobian is also equal to $\frac{\eta_0}{\eta_1}$ and $J = \frac{\eta_0^2}{\eta_1^2}$.

D Distribution tails in the model of Guo et al. [GQGP17]

We provide additional comparisons with the approach of Guo et al. [GQGP17]. Their BRDF model relies on the

von Mises-Fisher (vMF) distribution:

$$D_{vMF}(\mathbf{h}) = \frac{K}{4\pi \sinh K} \frac{\exp K \cos \theta_h}{A(K)} \quad (7)$$

with $A(K) = \coth K - \frac{1}{K}$. We have empirically found a simple relationship between the concentration parameter K of the vMF distribution and the roughness α of the GGX distribution:

$$K(\alpha) = \frac{6.28733102798}{\pi\alpha^2} - 1.05268783401. \quad (8)$$

The formula gives a reasonable approximation for common values of $\alpha \in (0, 0.4]$, which we have found sufficient for comparing the two distributions. As shown in Figure 15, the vMF distribution D_{vMF} (in blue) lacks the heavy tails of the GGX distribution D_{GGX} [WMLT07] (in orange). The visual impact is significant even with a single-layered material of intermediate roughness, as shown in the rendering insets ($\alpha = 0.2, \eta = 1.5$). This is why we have decided not to use this model in our analysis.




Figure 15: Comparisons between the vMF distribution (in blue) and the GGX distribution (in orange). From left to right, GGX distributions for $\alpha \in \{0.08, 0.2, 0.4\}$, with $K \in \{311.79, 49, 11.46\}$ according to Equation 8

E Layer scattering in Belcour’s models [Bel18]

We compare the (mainly forward) scattering capabilities of Belcour’s model [Bel18] to a reference scattering layer. For the sake of simplicity, we used the Layer Lab toolbox [Jak15] for this reference as it provides faster results compared to a full simulation of the Radiance Transfer Equation (RTE). Belcour’s model works with GGX distributions [WMLT07] for interfaces; unfortunately, LayerLab does not natively supports such interfaces. A workaround that is possible when the interface is only evaluated in reflection is to first store a GGX-based BRDF in the MERL format, then use the result from inside Layer Lab.

We take this approach, and choose a base GGX interface with a roughness $\alpha = 0.2$ and a complex index $\eta = (1.5, 1.5, 1.5)^T + i(\mathbf{3}, 0, 0)^T$ for the base material layer. Without any additional layer, the material is thus a redish rough conductor. We then add a medium layer of varying depth d and anisotropy g , but keeping constant coefficients

$\sigma_s = (0.1, 0.1, \mathbf{1.0})^T$ and $\sigma_a = (0.1, 0.1, 0.1)^T$. The additional scattering layer should thus be blueish and cover the red highlight from the base conductor to different extents depending on d and g .

As shown in Figures 16, the expected results are obtained with the simulation (left halves), but only partially with Belcour’s model. The approximation in Belcour’s model is only valid for $g > 0.7$, hence it is not surprising that the best results are obtained for high values of g . A more fundamental limitation of Belcour’s approach is that ray bundles are only propagated in the forward direction. Hence they cannot “turn around” inside the medium: they first have to be reflected at the base interface layer to eventually exit on reflection. However, the blueish tint is precisely due to light paths scattering inside the medium and exiting on reflection without ever having reached the base redish interface. This limitation is even more problematic toward grazing angles as shown in Figure 17: in these configurations, light paths traverse increased distances in the scattering medium, which results in increased scattering and a higher blueish coverage.

Note that we have not added any additional clear coat layer on top of this configuration, since it would yield artifacts with Layer Lab.




Figure 16: A diffusive material composed of a redish rough conductive base and a blueish scattering layer of varying depth D and anisotropy g . The material is applied to a sphere lit by a single light source making an angle of 30 degrees with the viewing direction. Left halves show simulation results, while right halves show the results of Belcour's forward model.




Figure 17: The same material as in Figure 16 with the light source making an angle of 65 degrees with the viewing direction.