An Inverse Method for the Exploration of Layered Material Appearance

Mégane Bati, Pascal Barla, Romain Pacanowski

► To cite this version:

Mégane Bati, Pascal Barla, Romain Pacanowski. An Inverse Method for the Exploration of Layered Material Appearance. ACM Transactions on Graphics, 2021, 40 (4), pp.1-15. 10.1145/3450626.3459857 . hal-03255469

HAL Id: hal-03255469

<https://hal.inria.fr/hal-03255469>

Submitted on 9 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.

An Inverse Method for the Exploration of Layered Material Appearance

MÉGANE BATI, IOGS, Université de Bordeaux, France

PASCAL BARLA, Inria, France

ROMAIN PACANOWSKI, Inria, France

Fig. 1. Our method permits to control layered material appearance by automatically retrieving physical parameters. (a) Modification of the color and sharpness of reflections produced by a 2-layer material. (b) Isolated control over grazing-angle colors obtained via a 2-layer material with a conductor base. (c) Coating an existing 2-layer material without modifying its underlying appearance. (d) Control over haze extent and color on a 2-layer material with a conductor base.

Layered materials exhibit a wide range of appearance, due to the combined effects of absorption and scattering at and between interfaces. Yet most existing approaches let users set the physical parameters of all layers by hand, a process of trial and error. We introduce an inverse method that provides control over BRDF lobe properties of layered materials, while automatically retrieving compatible physical parameters. Our method permits to explore the space of layered material appearance: it lets users find configurations with nearly indistinguishable appearance, isolate grazing angle effects, and give control over properties such as the color, blur or haze of reflections.

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

Additional Key Words and Phrases: layered materials, inverse method

ACM Reference Format:

Mégane Bati, Pascal Barla, and Romain Pacanowski. 2021. An Inverse Method for the Exploration of Layered Material Appearance. *ACM Trans. Graph.* 40, 4, Article 176 (August 2021), 15 pages. <https://doi.org/10.1145/3450626.3459857>

1 INTRODUCTION

Many materials are layered, with appearance varying depending on the number of layers and their nature (e.g., frosted glasses, coated metals, etc). The modeling of physically-realistic layered BRDFs (Bidirectional Reflectance Distribution Functions) has made significant progress in recent years. Yet, the *design* of layered BRDFs still

relies on trial and error, which becomes tedious with an increasing number of layers: the relationships between the physical parameters of a layer stack and the BRDF lobes it produces remain unknown.

The general process of relating visual properties to physical parameters is called inverse design. As detailed in Section 2, many such methods exist to retrieve physical BRDF parameters from various inputs (images, measurements, etc). However, they almost never account for layered materials; a notable exception being the work of Ershov [2004] specifically tailored to a two-layer car paint. The inverse design of layered materials is a difficult problem since BRDF lobes are interdependent. Lobes that are due to light paths propagating deep into the layer stack depend on all the physical parameters of the layers they interact with, in an intricate manner. For instance, a change of roughness of the top layer affects all BRDF lobes. As a result, many BRDF lobe combinations are not physically achievable (e.g., a perfectly specular lobe underneath a rough one).

Our main contribution is to introduce a method to explore the space of layered material appearance by means of inverse procedures, restricting ourselves to the context of geometric optics and non-scattering media. As outlined in Section 3, we focus on the properties of BRDF lobes *at normal incidence*, and explicitly retrieve the physical parameters of layer stacks that can produce the desired properties (i.e., energy and variance). Interestingly, several different layer configurations may be valid – they form equivalence classes – even though their resulting BRDFs may differ at grazing angles. This permits to identify classes of layered materials with markedly different physical parameters but similar visual appearance; an idea similar to the concept of *similarity relations* (e.g., see [Zhao et al. 2014]) at normal incidence. An additional advantage of explicitly characterizing such equivalent classes is that it permits to restrict input BRDF lobe properties to *physically-achievable* configurations.

Authors' addresses: Mégane Bati, IOGS, Université de Bordeaux, Bordeaux, France, megane.bati@institutoptique.fr; Pascal Barla, Inria, Bordeaux, France, pascal.barla@inria.fr; Romain Pacanowski, Inria, Bordeaux, France, romain.pacanowski@inria.fr.

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only.

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

0730-0301/2021/8-ART176 \$15.00

<https://doi.org/10.1145/3450626.3459857>

2 PREVIOUS WORK

2.1 Modeling

The most straightforward method for obtaining the BRDF model of a layer stack is to rely on the simulation of light transport at microscopic scales. This has the advantage of making no restriction on the type of layer: absorption, scattering and interference may be accounted for. The main drawback is of course computation time. In order to speed up computations, recent work [Gamboa et al. 2020; Guo et al. 2018; Xia et al. 2020] has relied on optimizations tailored to plane-parallel configurations. In our context, the main limitation of simulation approaches is not speed: they do not help relate BRDF lobe properties to physical stack parameters.

In the specific case of *smooth interfaces*, several models exist, such as the well-known adding equations [Stokes 1862]. The reflectance of thick layers is modeled through the combination of reflectance and transmittance computed using Fresnel equations. They can be generalized to diluted media with attenuation, relying on the Beer-Lambert-Bouguer law; and they may be combined with a Lambertian base layer as well [Simonot et al. 2006]. When dealing with thinner layers, the adding equations must be generalized to handle interferences, as done by Abeles [1948] and Yeh [2005]. In Computer Graphics, Hirayama [2001] and Icart & Arquès [2000] have introduced models that deal with an arbitrary number of such thin layers, with the restriction that roughness at each interface should be small compared to visible wavelengths.

The case of a stack of *rough interfaces* and thick media is more complex, even though it remains restricted to the geometric optics regime. Practical implementations in rendering engines (e.g., Arnold renderer [Georgiev et al. 2019]) usually take a non-physical approach: each interface is modeled through Microfacet theory and the resulting BRDFs are linearly combined. The most common microfacet distributions are those of Beckmann [1963] and GGX [Walter et al. 2007], each relying on a roughness parameter and a ratio of refractive indices. Dai et al. [2009] also proposed a simple combination strategy for a pair of rough interfaces, using log-interpolation between smooth-on-rough and rough-on-smooth models. None these interpolation-based solutions are physically-realistic.

Another class of layered material models targets physical plausibility. In their pioneering work, Weidlich and Wilkie [2007] proposed an analytical model that is particularly simple to implement. Unfortunately, it does not properly handle roughness of the upper interfaces and multiple bounces inside a medium are not accounted for. The model of Guo et al. [2017] offers a way to better handle roughness: they focus on a pair of interfaces with Gaussian-like microfacet distributions and avoid multiple bounces, even though their model could be extended to more complex configurations. The model of Belcour [2018] instead relies on GGX microfacet distributions and explicitly handles multiples bounces between layers. It is the first approach to model scattering inside media, even though it is restricted to forward propagation. Scattering in itself is studied in models like Gu et al.'s [2007], for dust layers where single scattering is the dominant effect; but it does not account for roughness.

A solution that sits inbetween simulation and modeling is the LayerLab system of Jakob et al. [2014]. It consists of a framework that handles both roughness at interfaces and scattering in media,

relying on a Fourier basis. Scattering is achieved through adding-doubling equations, also employed by Ergun et al. [2016] to extend LayerLab to handle translucent and iridescent flakes in a medium. The main drawback of LayerLab is that it requires an increasing amount of coefficients to handle specular lobes in a BRDF. It is thus unlikely that the relationship between BRDF lobe properties and physical parameters could be unveiled with this approach, even though some appearance exploration techniques have been proposed in the anisotropic version of the system [Zeltner and Jakob 2018], where the impact of a layer removal is characterized.

2.2 Inverse Methods

Inverse techniques may be used to retrieve many aspects of a 3D scene, from light positions and shapes, to object geometries or materials (e.g., see the survey of Patow et al. [2003]). Nowadays, differentiable rendering is a powerful general numerical tool to solve this kind of problem. However, in the context of going from a BRDF to physical parameters, several specific solutions exist.

Image-based methods rely on one or more photographs to infer material properties such as diffuse albedo, normal, specular albedo or roughness maps (e.g., [Riviere et al. 2017], [Descaintre et al. 2018]). Such methods are popular in Computer Graphics applications, as they are well adapted to artistic practice; but they fail to retrieve the more complex appearance afforded by layered materials.

Inverse methods are common in Optics, in which case physical parameters are retrieved from optical measurements. For instance, the complex refractive index of a conductor may be retrieved by measuring polarized reflectance at grazing angle configurations (e.g., [Querry 1969]). The measurements could be replaced by user inputs, but such an approach would be obviously far from intuitive, for instance regarding the difficulty in providing valid inputs (see [Armaly et al. 1972] for details on the conductor example).

An interesting middle-ground between image- and measurement-based approaches consists in providing navigation methods to explore the appearance space of BRDFs. Existing solutions rely on a set of pre-established BRDFs (either coming from measured data [Matusik et al. 2003] or from models [Ngan et al. 2006], eventually including micro-geometry [Wu et al. 2013]); each BRDF becomes a point in a high-dimensional space, and the space of valid BRDFs on which to navigate corresponds to a manifold. This approach has been recently combined with perceptual traits obtained via experiments [Serrano et al. 2016] that help further reduce the dimensionality of the appearance space. Navigation methods are user-friendly as they permit to quickly converge to a BRDF from simple user interaction. However, the BRDF itself is not described by physical parameters, but using pre-established BRDFs; some methods may even output a BRDF that is the interpolation of single- and multi-layered BRDFs.

When working with a known analytical BRDF model, a more direct approach consists in finding a mapping from visual to physical parameters. In the case of a single smooth conductor interface, Gulbrandsen's mapping [2014] permits to control the complex refractive index through a pair of parameters: reflectance at normal incidence, and edge tint at grazing angles. For a single rough interface described through Ward's BRDF model [1992], Pellacini [2000] introduced artist-friendly parameters based on a user study. The

control of several rough lobes has been studied in the context of BRDF shop [Colbert et al. 2006], making use of a painting interface. The particular relationship between a pair of BRDF lobes leading to hazy gloss effects has been studied by Barla et al. [2018], along with a visual-to-physical parameter mapping. All these methods have in common that they only deal with a single more complex interface.

To the best of our knowledge, only the work of Erchov et al. [2004] provides an inversion method for a physically-realistic layered model [Ershov et al. 2001]. However, it is specific to a mirror-flaked binder above a Lambertian base: the inversion is based on visual sliders (gloss, glitter and shade) to edit an input BRDF, where the physical parameters are retrieved using a two-step optimization. The drawback of such an optimization is that it might be stuck in a local minimum. In contrast, our objective is to span the full appearance space of layered materials, in order to identify multiple physical configurations that lead to a similar material appearance.

3 OUR APPROACH

We propose to explore the space of layered material appearance with a *inverse mapping* approach, which takes as input BRDF lobe properties, and outputs physical parameters that achieve the desired appearance. Such an explicit mapping approach raises two difficulties: lobes are interdependent as they correspond to sets of light paths propagating through several interfaces; and multiple physical configurations may lead to visually similar BRDF lobes.

We thus need to characterize *explicitly* the impact of the propagation of light paths on BRDF lobe properties. This motivates us to rely on Belcour’s model, where the energy and variance of each lobe is built from such a propagation, as illustrated in Figure 2. We choose this model over Guo’s because it handles inter-reflections among layers and the stacking of several interfaces more explicitly.

We can now restate our objective as that of inverting Belcour’s model, to retrieve physical parameters from BRDF lobe properties, as shown in Figure 2. To this end, we have decided to focus on BRDF lobes at normal incidence for two reasons: they characterize most of the material appearance; and this makes technical derivations tractable. Note that for transparent materials, we only control the reflectance properties. These choices are discussed in Section 8.

Fig. 2. Our approach inverts Belcour’s model, retrieving physical parameters from BRDF lobe properties. Two interfaces may yield a single lobe (in red).

Hypothesis. Following Belcour [2018], we consider thick dielectric layers modeled by GGX interfaces, each controlled by an *achromatic* index of refraction η and a roughness α . The base layer might also be a colored conductor, modelled with a complex index of refraction $\eta + i\kappa$; we use bold symbols to denote chromatic parameters (i.e., RGB vectors). The media between two interfaces may be absorbing, with absorption controlled by the optical depth at normal incidence

τ . We chose not to address scattering since it is not handled in full in Belcour’s model (as discussed in the work of Bati et al. [2019]).

We consider that a *physically-achievable* configuration must respect the following constraints: $\alpha \in [0, 1]$ and $\tau \in \mathcal{R}^+$; $\eta \in [1, \eta_{max}]$ for dielectrics; $\kappa > 0$ and $\kappa \geq \eta$ for conductors (see Section 5.1.2 for details). Even though dielectric media with $\eta < 1$ are achievable, we disregard them as they are very rarely encountered in practice.

A central hypothesis of our approach is that the space of layered materials forms *equivalence classes*: sets of physical configurations that yield identical BRDF lobe properties at normal incidence. For instance, several configurations of smooth interfaces can yield a single dirac lobe with the same reflectance at normal incidence; differences will only show up at grazing angles. In practice, we restrict our approach to at most two interfaces per lobe, which we found to be a good compromise between the number of parameters and visual effects. For each BRDF lobe, we thus have to retrieve up to six physical parameters for dielectrics (two chromatic absorptions $\tau_{0,1}$, achromatic indices of refraction $\eta_{1,2}$ and roughnesses $\alpha_{0,1}$), and seven for conductors (since $\eta + i\kappa$ is complex). Multiple lobes may be stacked to achieve materials with an arbitrary number of layers.

Fig. 3. We consider (a) each BRDF lobe to be composed of either (b) 1 layer or (c) 2 layers, each corresponding to a specific physical configuration controlled by roughnesses α , optical depths τ , and refractive indices (η for dielectrics, $\eta + i\kappa$ for conductors).

Figure 3 recapitulates the physical configurations we consider: 1-layer and 2-layer lobes, with either an achromatic dielectric or colored conductor base. Notice that η_0 is excluded because it is already fixed (either by the layer on top, or to $\eta_0 = 1$ for the top layer). In contrast, τ_0 is always included as it impacts light paths reflecting off interfaces attached to the current lobe; whereas for dielectric bases, τ_1 and τ_2 are excluded in Figures 3b and 3c respectively, as they only impact light paths propagating downward. Our approach handles materials made of an arbitrary number of layers through the stacking of multiple 1-layer or 2-layer lobes.

Outline. After recalling some essential technical background in Section 4, we present the three main contributions of our work:

- We present in Section 5 an in-depth theoretical analysis of equivalence classes yielding a *single lobe* of desired energy and variance for the configurations of Figure 3.
- We address the practical problem of controlling independently *multiple distinct lobes* while preserving their properties at normal incidence in Section 6.
- We explore the space of layered material appearance in Section 7. We demonstrate that our approach provides unprecedented control over layered material appearance as shown in Figure 1: the color, blur and haze of reflections are directly manipulated, while edge-tint and coating effects are handled.

4 TECHNICAL BACKGROUND

We first recall optics formula used for rendering layered materials (forward arrow in Figure 2). Each color channel is considered independently in this context. Notations are summarized in Table 1 and illustrated in Figures 3a and 13–right.

Table 1. Notations for physical parameters and lobe properties.

Physical parameters	
α_i	roughness (GGX distribution) of i -th interface
η_i	real part of refractive index below i -th interface
κ_i	imaginary part of refractive index below i -th interface
τ_i	optical depth at normal incidence below i -th interface
θ_i	angle below i -th interface (where θ_0 is the incident angle)
Lobe properties	
R_k, σ_k	k -th lobe energy and variance, $k \in [1..K]$
R_{ij}, T_{ij}	reflectance/transmittance from media i to j
R'_{ij}	reflectance from media i to j attenuated by τ_i
$\sigma_{ij}^R, \sigma_{ij}^T$	reflected/transmitted variance from media i to j
$\tilde{\sigma}_{ij}$	unnormalized variance (e.g., $\tilde{\sigma}_{ij}^R = R_{ij}\sigma_{ij}^R$)

Smooth interfaces. The reflectance of a smooth interface between two media of refractive indices η_0 and $\eta_1 + i\kappa_1$ is characterized by the Fresnel equations (assuming $\kappa_1 = 0$ for dielectrics). At normal incidence, the reflectance at the interface between media 0 and 1 is:

$$R_{01} = \frac{(\eta_0 - \eta_1)^2 + \kappa_1^2}{(\eta_0 + \eta_1)^2 + \kappa_1^2}. \quad (1)$$

The Beer-Lambert-Bouguer law characterizes the attenuation of light through a diluted absorbing medium of depth d ; in our case between two interfaces. Considering incident light of angle θ and intensity I_0 , the transmitted intensity is given by $I = I_0 e^{-\tau/\cos\theta}$ where $\tau = \mu_a d$ is the optical depth at normal incidence. In theory, the absorption coefficient μ_a and the imaginary part of the index of refraction κ should be related. However, in practice, this relationship can safely be ignored as explained in supplemental material.

The reflectance of a stack of smooth interfaces is obtained by iteratively computing the reflectance between two interfaces using the reflectance adding equation [Simonot et al. 2006; Stokes 1862]:

$$R_{02} = R_{01} + \frac{T_{01}R'_{12}T_{10}}{1 - R'_{12}R_{10}}. \quad (2)$$

Note that the $R_{0\times}$ terms depend on the incident angle θ_0 , while the $R_{1\times}$ terms depend on the refracted angle θ_1 , computed using Snell's law: $\eta_0 \sin \theta_0 = \eta_1 \sin \theta_1$. Here R_{02} denotes the reflectance of the stack (between media 0 and 2), while all the other reflectance and transmittance terms are computed using the Fresnel equations, where $R'_{12} = e^{-2\tau_1/\cos\theta_1} R_{12}$ is the reflectance attenuated by Beer-Lambert-Bouguer absorption. Similar adding equations exist for transmission (T_{02}), and for the reversed stack (R_{20}, T_{20}).

Rough interfaces. The stacking of rough interfaces yields a BRDF that is not purely specular. In Belcour's model, such a BRDF is given, for a pair (ω_i, ω_o) of light and view directions, by [Belcour 2018]:

$$\rho(\omega_i, \omega_o) = \sum_{k=1}^K R_k(\theta_0) \cdot \rho_k(f^{-1}(\sigma_k(\theta_0)), \omega_i, \omega_o), \quad (3)$$

where $\theta_0 = \arccos(\omega_i \cdot \mathbf{n})$ and $\mathbf{n}$ is the surface normal.

It is defined as a sum of K GGX lobes (excluding the Fresnel term) denoted by ρ_k , each gathering all the light paths reaching at most the interface k . The energy of the k -th lobe is written R_k , which matches a Fresnel reflectance in the case of a Dirac. For simplicity, we keep the same notation for the energy of a lobe irrespective of its variance σ_k . The equivalent roughness of each GGX lobe is obtained from this variance using $f^{-1}(\sigma_k)$, where the roughness-to-variance mapping function $f(\alpha) = \frac{\alpha^{1.1}}{1-\alpha^{1.1}}$. Note that R_k and σ_k depend on the light direction ω_i , following Belcour's *forward model*.

The computation of lobe energies $R_k(\theta_0)$ requires to take into account interface roughnesses. To this end, Belcour uses the directional albedo of each interface, denoted FGD , instead of Fresnel reflectance in Equation 2. A total internal reflection (TIR) term may optionally be taken into account when computing energies.

Computing lobe variances $\sigma_k(\theta_0)$ requires specific adding equations for the *propagation* of variances through the layer stack, in both reflection and transmission. For the case of a pair of interfaces, the reflected variance is given by $\sigma_{02}^R = \frac{\tilde{\sigma}_{02}^R}{R_{02}^R}$, where:

$$\tilde{\sigma}_{02}^R = \underbrace{R_{01}\sigma_{01}^R}_{1^{st} \text{ lobe}} + \underbrace{\frac{T_{01}R'_{12}T_{10}}{1 - R'_{12}R_{10}}\sigma_{02}^R}_{2^{nd} \text{ lobe}}. \quad (4)$$

However, for the final BRDF *evaluation* (Equation 3), the two terms of Equation 4 are output separately since they correspond to two different lobes. They are trivially normalized to obtain $\sigma_1 = \sigma_{01}^R$ and $\sigma_2 = \sigma_{02}$. In Equation 4, σ_{02} corresponds to the variance of the lobe reaching the bottom interface and “seen through” the rest of the stack on top of it. It is written as (see supplemental material):

$$\sigma_{02} = \sigma_{10}^T + \frac{\eta_1}{\eta_0} \left(\sigma_{01}^T + \frac{\sigma_{12}^R + R'_{12}R_{10}\sigma_{10}^R}{1 - R'_{12}R_{10}} \right), \quad (5)$$

where the unitary operators between media i and $i+1$ are given by:

$$\begin{cases} \sigma_{i,i+1}^R = \sigma_{i+1,i}^R = f(\alpha_{i+1}), \\ \sigma_{i,i+1}^T = f\left(\frac{1}{2}\left|1 - \frac{\eta_i}{\eta_{i+1}}\right|\alpha_{i+1}\right), \\ \sigma_{i+1,i}^T = f\left(\frac{1}{2}\left|1 - \frac{\eta_{i+1}}{\eta_i}\right|\alpha_{i+1}\right). \end{cases} \quad (6)$$

5 THEORETICAL ANALYSIS OF A SINGLE LOBE

We first consider a BRDF made of a single lobe (i.e., $K = 1$), and study the mapping from its properties $\mathbf{R}_1$ and σ_1 to physical parameters for the configurations in Figure 3. The 1-layer configurations have been studied in the literature for the case of $\eta_0 = 1$; in Section 5.1, we show the impact of varying η_0 (which routinely happens when layers are stacked) on the retrieval of other physical parameters. The more complex 2-layer configurations are addressed in Section 5.2.

Regardless of the number of layers, we disregard the τ_0 parameter in the following, since an equivalent configuration can be achieved by replacing $\mathbf{R}_1$ by $e^{2\tau_0}\mathbf{R}_1$ and then assuming $\tau_0 = 0$.

5.1 1-layer Configurations

We first consider one interface per lobe (Figure 3a); hence we have $\mathbf{R}_1 := \mathbf{R}_{01}$ and $\sigma_1 := \sigma_{01}^R$. Using Belcour's model, roughness is directly obtained from variance: $\alpha_1 = f^{-1}(\sigma_1)$. We are thus solely concerned by the retrieval of refractive indices from lobe energy in this section.

Fig. 4. Equivalent classes for 1-layer lobes at a fixed normal reflectance R_{01} in an ambient medium of index η_0 . The retrieved physical parameters are (a) η_1 for dielectrics and (b,c) (η_1, κ_1) for conductors. Dots indicate dielectric configurations in air. Crosses indicate configurations retained in Figure 5(a-c).

Fig. 5. Renderings showing different elements of an equivalent class (left/right halves) for a given normal reflectance: (a) $R_{01} = 0.04$ for dielectrics, (b,c) $R_{01} = 0.6$ for conductors. In (a), we add a clear coat, which shows that two solutions become physically-achievable for dielectrics when η_0 is high enough. In (b), we show a conductor in air with a colored edge tint. We add a clear coat in (c) while preserving reflectance at normal incidence for a better comparison (this is done manually here; see Section 6 for an automatic solution). As shown on the right half, the edge tint becomes much more subtle. As illustrated in (d), this is due to Snell's law, which restricts the range of angles accessible in Fresnel reflectance (dashed and dotted vertical bars, same legend as in Figure 4c)

5.1.1 Dielectrics. Considering a single dielectric interface, the only physical parameter left to retrieve is η_1 . For a smooth interface (i.e., $\alpha_1 = \sigma_1 = 0$), given an achromatic reflectance $R_{01} \in [0, 1)$ and the refractive index η_0 of the ambient medium, η_1 is obtained by inversion of Fresnel's equation at normal incidence (i.e., Equation 1 with $\kappa_1 = 0$), which yields two roots:

$$\eta_1 = \eta_0 \left(\frac{1 \pm \sqrt{R_{01}}}{1 \mp \sqrt{R_{01}}} \right). \quad (7)$$

When $\eta_0 = 1$, only one of the two roots is valid (i.e., $\eta_1 > 1$). However, when increasing η_0 , the smaller root becomes valid when $\eta_0 > \frac{1+\sqrt{R_{01}}}{1-\sqrt{R_{01}}}$. This is illustrated in Figure 4a where the smaller root is shown with a dashed line; and in Figure 5a, where we compare the appearance of the two roots on a rendering. We further impose that refractive indices remain below η_{max} (we use $\eta_{max} = 4$ throughout the paper for illustration purposes). As a result, R_{01} should also be limited to a maximum reflectance, simply computed through Equation 1 with $\kappa_1 = 0$ and $\eta_1 = \eta_{max}$. We have selected values of R_{01} that have solutions within these bounds in Figure 4a. Note that invalid configurations are always visualized with hatched regions.

For a rough interface, the reflectance at normal incidence also depends on roughness, as shown in Figure 6 which plots the FGD term at $\theta_0 = 0$ as a function of $\eta = \frac{\eta_1}{\eta_0}$ for various values of α_1 .

Note that in air ($\eta_0 = 1$), we have $\eta > 1$ and Equation 7 remains valid as seen in Figure 6. However, in the general case, we must instead numerically invert the FGD term, which again yields one or two valid roots. The impact of roughness on reflectance has further repercussions in the 2-layer configurations, detailed in Section 5.2.

A 1-layer dielectric lobe may also have a colored reflectance $\mathbf{R}_{01}$ when put in an ambient medium different than air, since then we may have a colored $\tau_0 \neq 0$. In this case, $\mathbf{R}_{01}$ is replaced by $e^{2\tau_0} \mathbf{R}_{01}$. However, this requires special care since η_1 must be achromatic. A simple solution consists in using the color channel of maximal reflectance in $\mathbf{R}_{01}$ in place of R_{01} in Equation 7. The colored τ_0 is retrieved by inverting the Beer-Lambert-Bouguer law: $\tau_0 = \frac{1}{2} \log \frac{R_{01}}{\mathbf{R}_{01}}$.

Fig. 6. Visualization of FGD at $\theta_0 = 0$. Fresnel (Eq. 1 in dotted black) is a good approximation for $\eta > 1$ and low α , but not for $\eta < 1$ when $\alpha > 0.1$.

5.1.2 Conductors. For a conductor layer, the real and imaginary parts of the complex index of refraction must be retrieved, denoted η_1 and κ_1 . Although a conductor may be colored, each color channel can be treated independently; hence we write R_{01} to denote the reflectance of a single arbitrary color channel. Contrary to dielectrics, we have found no significant impact of roughness on energy at normal incidence; hence we directly work from Fresnel equations.

Gulbrandsen [2014] has characterized the corresponding equivalence class in the air, which is shown in Figure 4b. Its extension to an arbitrary $\eta_0 \neq 1$ is directly obtained by inverting Equation 1:

$$\kappa_1 = \eta_0 \sqrt{\frac{R_{01} \left(\frac{\eta_1}{\eta_0} + 1 \right)^2 - \left(\frac{\eta_1}{\eta_0} - 1 \right)^2}{1 - R_{01}}}. \quad (8)$$

The effect of varying η_0 is visualized in Figure 4c. Since we explicitly distinguish dielectrics from conductors in our approach, we need to avoid configurations where κ_1 vanishes, which appear as dots on the abscissa in Figure 4b. This is done by restricting η_1 to:

$$\eta_{\max}^c = \eta_0 \frac{1 + R_{01} + \sqrt{6R_{01} - 1 - R_{01}^2}}{2(1 - R_{01})}, \quad (9)$$

$$\eta_{\min}^c = \eta_0 \frac{1 - R_{01}}{1 + R_{01}}. \quad (10)$$

The upper bound is obtained by imposing $\kappa_1 \geq \eta_1$ (see supplemental material), which is a defining condition for conductors. Since $\eta_{\max}^c$ may be undefined in Equation 9, reflectance should at least be constrained to $R_{01} > 3 - 2\sqrt{2} \approx 0.17$. In contrast to dielectrics, conductors hence cannot produce very low reflectance. For the lower bound, we have chosen to rely on the configuration where the edge tint is maximal, as derived in Gulbrandsen's work [2014].

We show the differences between these two bounds on pairs of renderings in Figures 5b in air ($\eta_0 = 1$) and in 5c under a clear coat ($\eta_0 = 1.5$); in the latter case, the colored edge tint is hardly noticeable. This cannot be explained by changes in Fresnel equations since they solely depend on the ratio $\frac{\eta_0}{\eta_1 + i\kappa_1}$, and hence yield identical results irrespective of η_0 (Equations 8 through 10 are proportional to η_0). This is due to Snell's law: a smaller portion of the Fresnel reflectance curve is spanned when increasing η_0 : as shown in Figure 5d, the difference between the two bounds decreases as η_0 increases. As a result, *edge tint nearly vanishes under a coating*.

5.2 2-layer Configurations

We now study how to produce a single lobe from a pair of interfaces (Figure 3b); hence $\mathbf{R}_1 := \mathbf{R}_{02}$ and $\sigma_1 := \sigma_{02}$. However, contrary to the previous section, we are faced with the problem that both $\mathbf{R}_1$ and σ_1 are affected by the same set of physical parameters, namely $\eta_1, \alpha_1, \tau_1, \eta_2 + i\kappa_2$ and α_2 (recall that we assume $\tau_0 = 0$), as seen from Equations 2 and 5. One solution to disentangle energy and variance would be to assume that roughness has no impact on energy at normal incidence. Unfortunately, this is not a valid hypothesis, as can be seen in Figure 6: even if $\eta_0 = 1$, we may have $\eta_1 > \eta_2$, requiring to take the roughness α_2 into account in the computation of R_{12} (i.e., using the FGD term instead of Fresnel Equations). In contrast, we do not make use of Belcour's TIR term as we have found it to have a negligible impact at normal incidence.

In the following, we progressively build up complexity by first considering smooth dielectric interfaces without and with achromatic absorption, then incorporating interface roughness. The general procedure to recover the constrained equivalence class in the *achromatic* case for a 2-layer dielectric is summarized in Algorithm 1. The cases of colored absorption in media and of a colored conductor base layer are derived from it at the end of the section.

5.2.1 Smooth Dielectric. The reflectance of a smooth 2-layer dielectric lobe is modeled using the adding equation (Equation 2). In the rest of this section, we rely on its inversion, which is given by:

$$R'_{12} = R_{12} e^{-2\tau_1} = \frac{R_{02} - R_{01}}{1 + (R_{02} - 2)R_{01}}, \quad (11)$$

where we have used $R_{01} = R_{10} = 1 - T_{01} = 1 - T_{10}$ in the derivation.

No absorption. We first consider $\tau_1 = 0$ for simplicity. Given η_0 , for each possible value of η_1 we compute R_{01} , then use Equation 11 to retrieve R_{12} and finally compute roots for η_2 using Equation 7. The resulting *unconstrained* equivalent class $\mathcal{E}$ is displayed in Figure 7a for $\eta_0 = 1$: it is a 1D manifold in the 2D space of physical parameters.

Here the intersections with the vertical dashed line (where $\eta_1 = 1$) and the diagonal dashed line (where $\eta_1 = \eta_2$) correspond to a single interface, whereas the intersections with the horizontal dotted line (where $\eta_2 = \eta_0 = 1$) correspond to a slab. The points A and D hence correspond to the roots of Equation 7.

As shown in Figure 7b, decreasing R_{02} makes $\mathcal{E}$ shrink toward the singular case (η_0, η_0) when $R_{02} = 0$, which effectively makes interfaces vanish. When increasing η_0 , $\mathcal{E}$ translates on the diagonal line, but also grows in size, as shown in Figure 7c.

Not all points on $\mathcal{E}$ are physically-achievable: for instance, A in Figure 7a corresponds to a configuration where both refractive indices are smaller than 1. We thus restrict (η_1, η_2) to the $[1, \eta_{\max}]^2$ domain, which has the effect of clipping $\mathcal{E}$ in both dimensions. We write the resulting *constrained* equivalence class $\hat{\mathcal{E}} = \mathcal{E} \cap [1, \eta_{\max}]^2$. To navigate over $\hat{\mathcal{E}}$, we must find a subset $\mathcal{N}_1$ of η_1 values that guarantee that at least one point of $\hat{\mathcal{E}}$ is reached. Formally, we write:

$$\mathcal{N}_1 = \{\eta_1 / \exists \eta_2, (\eta_1, \tau_1 = 0, \eta_2) \in \hat{\mathcal{E}}\}. \quad (12)$$

For some values of R_{02} , we might have $\mathcal{N}_1 = \emptyset$. These impossible configurations may be avoided by bounding R_{02} by simply computing the adding equation (Equation 2) using the boundaries of $\mathcal{N}_1$. The practical computation of $\mathcal{N}_1$ is given in supplemental material.

Compared to a 1-layer lobe, a 2-layer dielectric lobe has the ability to reach higher reflectance while respecting physical constraints. For instance, with $\eta_{\max} = 4$, the highest achievable reflectance at normal incidence is 0.36 and 0.53 for 1- and 2-layer lobes respectively. Moreover, for two points of the same equivalence class $\hat{\mathcal{E}}$, reflectance may differ at grazing angles. This is visualized in Figure 8, which shows the difference in reflectance between different points on $\hat{\mathcal{E}}$ and the 1-layer configuration. The slab configuration (i.e., point E in Figure 7a) is the one that maximizes grazing angle reflectance because in this case, rays at a grazing incidence on the top interface are refracted to a direction that *reaches the critical angle* on the second interface. Figure 10a compares the 1-layer and slab configurations on a rendering. We use the Glacier environment as it brings out differences at grazing angles.

Fig. 7. Equivalence classes for 2-layer dielectric lobes at a fixed normal reflectance R_{02} in a medium of refractive index η_0 . (a) The retrieved physical parameters are (η_1, η_2) pairs, bounded on η_1 by the 1-layer configurations (A and E). (b) Decreasing R_{02} makes the classes shrink toward the (η_0, η_0) configuration. (c) Increasing η_0 shifts the classes along the diagonal and makes them grow, potentially increasing the number of invalid configurations (hatched regions).

Fig. 8. We compare reflectance at grazing angles for several points of an equivalence class $\mathcal{E}$ shown in the inset ($R_{02} = 0.15$, $\eta_0 = 1$). The curves display the difference between each point and point D (the 1-layer configuration), revealing that grazing angle effects vary non-monotonically while walking along $\mathcal{E}$.

Achromatic absorption. We now consider $\tau_1 > 0$, here in the achromatic case. The procedure to recover the *unconstrained* equivalence class $\mathcal{E}$ is the same as before, except this time we retrieve R'_{12} instead of R_{12} in Equation 11, with τ_1 another degree of freedom. $\mathcal{E}$ then becomes a 2D manifold in the 3D space of physical parameters, as shown in Figure 9a. Bounds on η_1 (i.e., points A and D) are preserved since they correspond to 1-layer configurations. This is best seen in Figure 9b, which shows that all τ_1 -isolines have the same projection on the η_1 axis. Any point that lies above or below the 0-isoline (in green) and outside of the hatched regions thus corresponds to a physically-achievable configuration. Figure 9c provides yet another view of $\mathcal{E}$ with different values of R_{02} . It shows that τ_1 remains limited in magnitude by constraints on refractive indices.

Consequently, we have not found the need to bound explicitly τ_1 and we navigate the corresponding *constrained* equivalence class $\hat{\mathcal{E}}$ by picking a valid (η_1, η_2) pair, then retrieving τ_1 using Equation 11:

$$\tau_1 = \frac{1}{2} \log \left(\frac{R_{12}}{R'_{12}} \right). \quad (13)$$

In order to navigate over $\hat{\mathcal{E}}$, we must now also find a subset $\mathcal{N}_2$ of η_2 values that guarantee that at least one point of $\hat{\mathcal{E}}$ is reached:

$$\mathcal{N}_2 = \{\eta_2 / \exists \tau_1 \geq 0, (\eta_1, \tau_1, \eta_2) \in \hat{\mathcal{E}}\}. \quad (14)$$

The main effect of reaching configurations with higher τ_1 is to darken reflectance at grazing angles, since light paths then *travel longer distances in the absorbing medium*. This is shown in Figure 10b. The effect is subtle in the achromatic case, but much more noticeable with colored absorption, as described in Section 5.2.3.

5.2.2 Rough Dielectric. The main challenge with a rough 2-layer dielectric configuration is to guarantee that a *single lobe* of desired variance $\sigma_{02} > 0$ occurs at normal incidence, requiring a proper choice of interface roughnesses α_1 and α_2 , which also affects the way the lobe energy R_{02} is inverted to yield physical parameters.

In order to *merge the two lobes into a single one* at normal incidence, we impose that $\sigma_{01}^R = \sigma_{02}$ in Equation 4. As a result, the top interface roughness is directly given by $\alpha_1 = f^{-1}(\sigma_{02})$. Since roughness affects reflectance in the general case as shown in Figure 6, the first modification to make to Equation 11 is to compute R_{01} using α_1 . For each possible value of η_1 , R_{01} is thus computed using the *FGD* term. We then apply Equation 11 to obtain the attenuated reflectance R'_{12} . The bottom interface roughness $\alpha_2 = f^{-1}(\sigma_{12}^R)$ may now be retrieved by inverting Equation 5:

$$\sigma_{12}^R = (1 - R_{10}R'_{12}) \left(\frac{\eta_0}{\eta_1} (\sigma_{02} - \sigma_{10}^T) - \sigma_{01}^T \right) - R_{10}R'_{12}\sigma_{10}^R, \quad (15)$$

where the variance terms are defined in Equation 6. To complete the retrieval of physical parameters, for each possible value of η_2 , we need to compute R_{12} to evaluate Equation 13. Once again, we rely on the *FGD* term, except that we replace α_2 with an effective roughness $\alpha'_2 = f^{-1}(\sigma_{12}^R + \sigma_{01}^T)$ as in Belcour's approach. Indeed, the first interface impacts the transmitted light lobe's variance.

This procedure yields the unconstrained equivalence class $\mathcal{E}$, which is shown in Figure 11 and compared to the smooth 2-layer configuration. In order to derive the constrained equivalence class $\hat{\mathcal{E}}$, we must restrict $\mathcal{E}$ to physically-achievable parameter ranges. This requires to ensure that $\sigma_{12}^R \geq 0$ in addition to constraints on refractive indices. For a given α_1 , we may define a subset $\mathcal{V}$ so that $(\eta_1, \eta_2) \in \mathcal{V}$ yields a positive σ_{12}^R through Equation 15. The constrained equivalence class is then given by $\hat{\mathcal{E}} = \mathcal{E} \cap [1, \eta_{\max}]^2 \cap \mathcal{V}$.

Fig. 9. Equivalence classes for 2-layer dielectric lobes with absorption at a given normal reflectance R_{02} in air ($\eta_0=1$). The retrieved physical parameters are (η_1, τ_1, η_2) triplets. In (a), we show the equivalence class as a 3D surface, with several τ_1 slices also visualized in (b) with colors, and one η_1 slice shown in black and visualized in (c). In (c), we also show η_1 slices for different values of R_{02} . Crosses correspond to configurations selected for rendering in Figure 10b.

Fig. 10. Renderings comparing two elements of an equivalent class for a given reflectance and variance at normal incidence: $R_{02} = 0.15$; $\sigma_{02} = 0$ for (a,b) and $\sigma_{02} = 0.02$ for (c). In (a), we assume $\tau_1 = 0$ and compare two refractive index pairs: the right half (slab configuration) shows a brighter reflectance at grazing angles. In (b), we incorporate achromatic absorption and compare two different triplets of parameters: the right half is this time darker due to absorption. In (c), we consider roughness and compare two different quadruplets of parameters: the two halves are nearly indistinguishable. As shown in (d), this is due to a decrease of the energy of the bottom lobe (in orange) at grazing angles which tends to mask its higher variance compared to the top lobe (in blue).

Unfortunately, we have not found an exact solution for $\mathcal{V}$. However, we provide an approximation in supplemental material, which places an additional constraint $\eta_1 \leq \eta_{\max}^\sigma$, visualized by the vertical dashed line in Figure 11. In practice, we use this constraint to further restrict $\mathcal{N}_1$. Note that our approximation is not conservative, hence we still need to explicitly check that $\sigma_{12}^R \geq 0$; but many of the unachievable configurations can now be avoided.

The inversion procedure to recover $\hat{\mathcal{E}}$ in the achromatic dielectric case is summarized in Algorithm 1. Note that due to existing bounds on R_{02} , τ_0 must also be constrained to a subset $\mathcal{T}_0$ given by:

$$\mathcal{T}_0 = \frac{1}{2} \log \frac{[R_{\min}, R_{\max}]}{R_{02}}, \quad (16)$$

where $R_{\min}$ and $R_{\max}$ denote the bounds on R_{02} .

As before, one should remember that $\hat{\mathcal{E}}$ is an equivalence class for BRDF lobe properties only at normal incidence. At grazing angles, the two lobes due to each interface may not fuse anymore into a single one; yet, as shown in Figure 10c, the difference between 1-layer and 2-layer lobes next to the sphere silhouette is extremely subtle. Figure 10d provides an explanation: for the bottom lobe, the energy increases when the variance decreases toward grazing

angles. Our approach thus successfully permits to produce a single BRDF lobe out of two physical layers; the main advantage over a 1-layer lobe is that it affords not only more intense, but also colored dielectric lobes as shown in Figure 18 and explained next.

Fig. 11. Equivalence classes for a 2-layered dielectric lobe of variance $\sigma = f(\alpha_1)$ and normal reflectance $R_{02} = 0.2$ in air ($\eta_0 = 1$). The retrieved physical parameters are $(\eta_1, \eta_2, \alpha_2)$. When increasing roughness, the smooth class (dotted black line) is modified and might get interrupted when $\alpha_2 < 0$. We approximate this limit by a condition on η_1 , shown as vertical dashed lines.

5.2.3 Colored media. When considering color, the input lobe reflectance R_{02} can be specified for each color channel. In the case of a 2-layer dielectric, color only comes from absorption in media. The first modification to Algorithm 1 is to use a colored optical depth $\tau_0 \in \mathcal{T}_0$, which merely changes the input R_{02} . Assuming a colored optical depth τ_1 is more problematic: it yields a colored R'_{12} in Equation 11, but this cannot be directly used when inverting variance in Equation 15 since α_2 and thus σ_{12}^R are expected to be scalars. In practice, we pick the color channel of maximum intensity in R'_{12} . Thus, note that the desired lobe variance σ_{02} is only achieved for this color channel, while variances may differ for the other channels.

We are then left with a final constraint to resolve: even though τ_1 might differ per color channel, η_1 and η_2 must remain achromatic. The equivalent class that satisfies this final constraint is obtained by intersection of the per-channel *constrained* equivalent classes: $\hat{\mathcal{E}} = \hat{\mathcal{E}}^R \cap \hat{\mathcal{E}}^G \cap \hat{\mathcal{E}}^B$. An example is shown in Figure 12, where each per-channel constrained equivalence class is depicted with a fill color, while their intersection $\hat{\mathcal{E}}$ appears in light gray surrounded by a black outline. In practice, navigating this constrained equivalence class is done as in the achromatic case, except that we now use refractive index ranges that are intersections of their per-channel equivalent: $N_1 = N_1^R \cap N_1^G \cap N_1^B$ and $N_2 = N_2^R \cap N_2^G \cap N_2^B$.

5.2.4 Conductor base. A 2-layer lobe where the bottom medium is a conductor can be made equivalent to the dielectric case at normal incidence. To this end, the refractive index η_2 is replaced by an equivalent *colored* refractive index $\tilde{\eta}_2$, no longer bounded by $\eta_{\max}$.

Each color channel of $\tilde{\eta}_2$ is first computed independently with Algorithm 1. The complex refractive index $\eta_2 + i\kappa_2$ of the base layer is then simply obtained by navigating the equivalence class of a conductor (see Figure 4b) given by a reflectance of $\frac{(\tilde{\eta}_2 - \eta_1)^2}{(\tilde{\eta}_2 + \eta_1)^2}$.

Algorithm 1: Achromatic 2-layer dielectric inversion

Input: $\eta_0, R_{02}, \sigma_{02}$
Output: $\hat{\mathcal{E}}$
 $\alpha_1 = f^{-1}(\sigma_{02})$ ▷ w/ $\sigma_{01}^R = \sigma_{02}$
 $\hat{\mathcal{E}} = \emptyset$
 Compute $\mathcal{T}_0$ ▷ eq. 16
forall $\tau_0 \in \mathcal{T}_0$ **do**
 $R_{02} \leftarrow e^{2\tau_0} R_{02}$
 Compute N_1 ▷ eq. 12
 forall $\eta_1 \in N_1$ **do**
 $R_{01} = FGD(\eta_1/\eta_0, \alpha_1)$
 Compute R'_{12} using (R_{01}, R_{02}) ▷ eq. 11
 Compute σ_{12}^R using $(\eta_0, \eta_1, R'_{12})$ ▷ eq. 15
 if $\sigma_{12}^R \geq 0$ **then**
 $\alpha_2 = f^{-1}(\sigma_{12}^R)$
 $\alpha'_2 = f^{-1}(\sigma_{12}^R + \sigma_{01}^T)$
 Compute N_2 ▷ eq. 14
 forall $\eta_2 \in N_2$ **do**
 $R_{12} = FGD(\eta_2/\eta_1, \alpha'_2)$
 Compute τ_1 using (R_{12}, R'_{12}) ▷ eq. 13
 $\hat{\mathcal{E}} = \hat{\mathcal{E}} \cup (\tau_0, \alpha_1, \eta_1, \tau_1, \alpha_2, \eta_2)$

Fig. 12. The equivalence class $\hat{\mathcal{E}}$ in the case of colored absorption is the intersection (in light grey with dark outlines) of per-channel constrained equivalent classes (filled with their respective color). The solid curves show per-channel $\tau_1 = 0$ isolines, here with $\eta_0 = 1.5$ and $\sigma_{02} = 0$.

6 PRACTICAL STACKING OF MULTIPLE LOBES

We now consider a stack of several lobes ($K > 1$) and assume they follow a specific order: the first lobe corresponds to the top interface(s), the second lobe to the one(s) below, etc. Unfortunately, the lobes cannot be set independently and then combined together since each lobe is affected by the interfaces on top of it (e.g., changing the roughness of the top interface will affect *all* BRDF lobes).

A logical solution is to set lobes from top to bottom: starting with the top lobe, which can be treated as in the previous section; then iteratively proceeding to lobes below. However, in the latter case, we must take into account the effect of the previously chosen layers, as illustrated in Figure 13. This process is outlined in Algorithm 2, whereas the specifics of stacked 1-layer and 2-layer lobes are detailed in Sections 6.1 and 6.2 respectively.

Fig. 13. **Left:** the inversion of a 1-layer ($l = 1$) or 2-layer ($l = 2$) lobe at the bottom of a n -layer stack requires to retrieve its energy $R_{n,n+l}$ and variance $\sigma_{n,n+l}^R$ after having travelled through the n layers. **Right:** this requires to summarize the n -layer stack with upward/downward lobe statistics.

In both Figure 13 and Algorithm 2, we use $l[k] \in 1, 2$ to characterize whether the k -th lobe is made of 1 or 2 layers, and we simply write l in subscripts to simplify notations. Observe that contrary to Algorithm 1 that outputs a constrained equivalence class, the output of Algorithm 2 is an explicit layer stack $\mathcal{S}$ (i.e., one instance of the multi-lobe equivalence class). Even though we will only consider the achromatic case throughout the section, a chromatic inversion follows the same process. In addition, Algorithm 2 assumes that the bottom (i.e., K -th) lobe is dielectric; the case of a base conductor lobe is obtained as in Section 5.2.4, through an equivalent refractive index $\tilde{\eta}_K$ to retrieve the complex refractive index of the last layer.

Algorithm 2: Achromatic multi-lobe inverse design

Input: $[R_k], [\sigma_k], k \in [1..K]$
Output: stack S
 $S = []; n = 0$
forall lobe k **do**
 if $k = 1$ **then**
 $R_{0,l}, \sigma_{0,l}^R, \sigma_{0n}^T = R_0, \sigma_0, 0$ ▷ 1st lobe ($n = 0$)
 else
 $R'_{n,n+l}, \sigma_{n,n+l}^R, \sigma_{0n}^T = \text{BaseStats}(R_k, \sigma_k, S)$ ▷ Alg. 3
 if $l[k] = 1$ **then**
 $S = \text{1-layerParams}(R'_{n,n+1}, \sigma_{n,n+1}^R, \sigma_{0n}^T, S)$ ▷ Alg. 4
 $n += 1$
 else
 $S = \text{2-layerParams}(R'_{n,n+2}, \sigma_{n,n+2}^R, \sigma_k, S)$ ▷ Alg. 5
 $n += 2$

6.1 Stacking onto a 1-layer lobe

Let's first consider the case of a 1-layer lobe, under a stack of n layers for which we already know physical parameters. The case $n = 0$ is straightforward as it boils down to the process of Section 5.1, with $\tau_0 = 0$ and $\eta_0 = 1$ since we assume air for the ambient medium.

In the case where $n > 0$, the properties used to control the k -th 1-layer lobe correspond to the energy $R_k := R_{0,n+1} - R_{0n}$ and the variance $\sigma_k := \sigma_{0,n+1}^R$. The term R_{0n} must be subtracted as it is already incorporated in the $(k - 1)$ -th lobe, and should not be counted twice. As illustrated at the left of Figure 13, these properties characterize the lobe in the ambient medium outside of the stack; whereas the inverse process of Section 5 expects energy and variance in the medium just below the n -th interface. In the case of a 1-layer lobe ($l = 1$), these are given by $R_{n,n+1}$ and $\sigma_{n,n+1}^R$ respectively.

Starting with the energy, we must compute $R_{n,n+1}$, knowing R_k and the physical parameters of the n -layer stack. To do so, we use Belcour's model to compute the upward/downward reflected and transmitted variances and energies. As illustrated in Figure 13-right, they are denoted by R_{0n}, σ_{0n}^R (resp. R_{n0}, σ_{n0}^R) for the upward (resp. downward) reflection, and T_{0n}, σ_{0n}^T (resp. T_{n0}, σ_{n0}^T) for the downward (resp. upward) transmission¹. We write R_k using adding equations:

$$R_k = R_{0,n+1} - R_{0n} = \frac{T_{0n}R'_{n,n+1}T_{n0}}{1 - R'_{n,n+1}R_{n0}}, \quad (17)$$

where $R'_{n,n+1} = R_{n,n+1}e^{-2\tau_n}$ denotes the attenuated energy.

¹Contrary to Belcour's optimized implementation, we do not average lobe statistics over color channels, as it would make inversion impractical.

Algorithm 3: Compute base statistics

Input: R_k, σ_k, S
Output: $R_{n,n+l}, \sigma_{n,n+l}^R, \sigma_{0n}^T$
 Compute $R_{0n}, R_{n0}, T_{0n}, T_{n0}$ from S ▷ [Belcour 2018]
 Compute $\sigma_{0n}^R, \sigma_{n0}^R, \sigma_{0n}^T, \sigma_{n0}^T$ from S ▷ [Belcour 2018]
 Compute $R'_{n,n+l}$ from R_k ▷ Eq. 18
 Compute $\sigma_{n,n+l}^R$ from σ_k ▷ Eq. 19

The attenuated energy $R'_{n,n+1}$ is obtained by inverting Equation 17:

$$R'_{n,n+1} = \frac{1}{\frac{T_{0n}T_{n0}}{R_k} + R_{n0}}. \quad (18)$$

To recover the variance $\sigma_{n,n+1}^R$, we generalize Equation 15 to the case of a n -layer stack acting at the top layer, yielding:

$$\sigma_{n,n+1}^R = (1 - R_{n0}R'_{n,n+1}) \left(\frac{\eta_0}{\eta_n} (\sigma_k - \sigma_{n0}^T) - \sigma_{0n}^T \right) - R_{n0}R'_{n,n+1}\sigma_{n0}^R. \quad (19)$$

This is summarized in Algorithm 3, which is written for the general case (1- or 2-layer, arbitrary l) as we will need it again in Section 6.2.

The physical parameters of the $(n + 1)$ -th layer can now be retrieved, as summarized in Algorithm 4 (assuming $m = n$). For a chosen refractive index η_{n+1} , the energy $R_{n,n+1}$ is computed through the *FGD* table, using σ_{0n}^T to account for the transmitted variance onto the $(n + 1)$ -th layer. The optical depth τ_n is recovered from $R_{n,n+1}$ and $R'_{n,n+1}$ as in Equation 13. The roughness α_{n+1} is obtained from $\sigma_{n,n+1}^R$. All three physical parameters are added to the stack S .

Figure 14a illustrates this process with a simple configuration: a 1-layer colored conductor onto which a dielectric coating is applied. When the physical parameters of the base layer are left unchanged, its appearance is modified by the coating; in particular, the lobe variance is significantly increased, causing blurrier reflections. Our inverse procedure identifies the physical configurations that maintain the statistics of the bottom lobe at normal incidence, hence preserving its appearance under the coating. We use the Ennis lighting environment as it helps bringing out differences around normal incidence thanks to its strongly directional light source.

6.2 Stacking onto a 2-layer lobe

One could think that the case of a stacked 2-layer lobe can be made equivalent to the stacked 1-layer lobe problem of Section 6.1, by first turning the 2-layer lobe into a 1-layer lobe using adding equations. As detailed in supplementary material, this is unfortunately not possible in the general case involving rough interfaces since the variance adding equations are not distributive. We thus take another approach where our inversion is performed in a top-down fashion.

The properties used to control the k -th 2-layer lobe correspond to the energy $R_k := R_{0,n+2} - R_{0n}$ and the variance $\sigma_k := \sigma_{0,n+2}^R$. As before, we first recover $R'_{n,n+2}$ and $\sigma_{n,n+2}^R$ with Algorithm 3, this time using $l = 2$. The computation of the physical parameters for a 2-layer lobe is then described in Algorithm 5. The refractive index and

Algorithm 4: Compute 1-layer parameters

Input: $R'_{m,m+1}, \sigma_{m,m+1}^R, \sigma_{0m}^T, S$
Output: stack S
 Choose η_{m+1}
 $R_{m,m+1} = \text{FGD}\left(f^{-1}(\sigma_{m,m+1}^R + \sigma_{0m}^T), \frac{\eta_{m+1}}{\eta_m}\right)$
 $\tau_m = \frac{1}{2} \log\left(\frac{R_{m,m+1}}{R'_{m,m+1}}\right)$
 $\alpha_{m+1} = f^{-1}(\sigma_{m,m+1}^R)$
 $S += \{\eta_{m+1}, \tau_m, \alpha_{m+1}\}$

Fig. 14. We compare a material before (left halves) and after (right halves) adding a coating, either by directly modifying physical parameters (top row) or using our method to preserve base lobe properties (bottom row). The insets show physical parameters, in red when retrieved by our method. In (a), the 1-layer conductor base ($\mathbf{R} = [0.9, 0.7, 0.5]$, $\sigma = 0.03$) appears blurrier with a naive approach, as confirmed by an increase in the base lobe variance (bottom row). In (b), the 2-layer dielectric lobe ($\mathbf{R} = [0.1, 0.2, 0.1]$, $\sigma = 0.03$) looks darker and sharper with a naive approach and is even split in two as shown by the separate variance curves (bottom row).

optical depth for the $(n+1)$ -th layer are directly chosen by the user, while $\alpha_{n,n+1}$ is computed from $\sigma_{n,n+2}^R$, since we want both interfaces to have the same roughness. We are now in a situation similar to that of Section 6.1, this time having identified a $(n+1)$ -layer stack. We may thus once again rely on Algorithm 3 to obtain lobe statistics under the $(n+1)$ -th layer, and then use Algorithm 4 (assuming $m = n+1$) to retrieve the physical parameters of the $(n+2)$ -th layer and update the layer stack $\mathcal{S}$, as detailed in Algorithm 5.

However, there is a remaining subtlety: the desired lobe energy should be adjusted to discard light paths that reach the $(n+1)$ -th layer but do not go further down. We introduce a temporary variable $\bar{R}_k := R_{0,n+2} - R_{0,n+1}$ as our target energy, in practice computed by:

$$\bar{R}_k = \sum_{i=0}^k R_i - R_{0,n+1}, \quad (20)$$

where $R_{0,n+1}$ is obtained with adding equations using the FGD term.

Algorithm 5: Compute 2-layer parameters

Input: $R'_{n,n+2}$, $\sigma_{n,n+2}^R$, σ_k , $\mathcal{S}$
Output: stack $\mathcal{S}$
 Choose η_{n+1} and τ_n
 $\alpha_{n+1} = f^{-1}(\sigma_{n,n+2}^R)$
 $\mathcal{S} += \{\eta_{n+1}, \tau_n, \alpha_{n+1}\}$
 Compute $\bar{R}_k$ ▷ Eq. 20
 $R'_{n+1,n+2}, \sigma_{n+1,n+2}^R, \sigma_{0,n+1}^T = \text{BaseStats}(\bar{R}_k, \sigma_k, \mathcal{S})$ ▷ Alg. 3
 $\mathcal{S} = \text{1layerParams}(R'_{n+1,n+2}, \sigma_{n+1,n+2}^R, \sigma_{0,n+1}^T, \mathcal{S})$ ▷ Alg. 4

As seen in Figure 14b, when the physical parameters of a 2-layer dielectric lobe are left unchanged when applying a coat layer on top, the resulting appearance is not only darker, but also appears sharper. Actually, the 2-layer lobe is split in two separate lobes, as shown by the orange dashed curves at the bottom of the figure. In contrast, the procedure described above finds physical parameters that yield the desired lobe properties, preserving the base lobe appearance.

7 EXPLORATION RESULTS

The inverse procedures presented in the previous two sections constitute the core of our contribution. In practice, they are implemented in a prototype code in Python, which performs the inverse mapping from lobe properties to physical parameters. Since this mapping is not bijective, we provide interactive jupyter notebooks that permit to navigate equivalence classes. The renderings are produced with a slightly modified version of Belcour's Layered_forward Mitsuba plugin, using physical parameters output by our approach in a dedicated XML format. Only Figures 23 and 24 are obtained using respectively Belcour's layered_dielectric and layered_ref plugins. We also employ a different discretization of the FGD term for inversion: our lookup table has a higher resolution and is restricted to $\theta_0 = 0$. All implementation details are presented in supplementary material.

We begin our exploration of the space of layered material appearance by comparing layer stack configurations belonging to a same single- or multi-lobe equivalence class in Section 7.1. We then show how our approach permits to obtain complex appearance by providing direct control over the BRDF lobe properties in Section 7.2.

Figure 1 shows a collection of the materials displayed in this section, applied to a variety of objects and placed in a same scene. Several additional results can be found in supplemental material, including a comparison between Belcour's model and Monte Carlo simulations (using the work of Gamboa et. al [2020]) and the list of retrieved physical parameters for all the figures in this section.

7.1 Exploring equivalence classes

If we pick several materials from a same equivalence class, they will provide identical BRDF lobes at normal incidence. However, they may also differ away from $\theta_0 = 0$, as seen in Figure 10d, or in the bottom row of Figure 14. Navigating through an equivalence class then permits to isolate and explore potential variations in material appearance around grazing angles. In Figures 15 through 17, we compare pairs of layered materials coming from a same equivalence class; a more thorough comparison is afforded by the interactive

Fig. 15. Navigation in equivalence classes permits to explore grazing-angle effects. In all images we compare pairs of 2-layer dielectric materials inside their equivalence class (shown on top), with either a high τ_1 (dark square) or a low τ_1 (bright square). When the material is (a) achromatic smooth ($R_1 = 0.15$), differences (left/right halves) are subtle. When it is (b,c) colored ($R_1 = [0.09, 0.09, 0.15]$), differences occur at grazing angles with the presence/absence of a purple tint, irrespective of roughness ($\sigma_1 = 0.03$ in (c)).

HTML interface provided in supplemental material. We also demonstrate our dielectric 2-layer prototype in the supplemental video. We use the Uffizi and Doge lighting environments as their large rectangular light source helps distinguish effects at grazing angles.

Figure 15 shows the case of a 2-layer dielectric BRDF. Differences in appearance remain subtle in the achromatic case, but become more noticeable for either smooth or rough colored configurations: the material acquires a purple tint at grazing angles (purple squares). In Figure 16, we navigate the equivalence class pairs of a 2-layer conductor BRDF. Note that the first equivalence class is vertically unbounded as we use the effective refractive index $\tilde{\eta}_2$; the second one represents the complex refractive index, and we systematically pick the center of the equivalence class since we observed extremely subtle variations (similarly to Figure 5c). We have chosen an achromatic energy at normal incidence on purpose, to let colored grazing angle effects better stand out. Compared to a reference material (black square), variations in η_1 produce extremely subtle differences (grey square), whereas variations in τ_1 let appear a pink tint at grazing angles, more pronounced when η_1 is low (pink squares).

When the layered material is composed of more than one lobe, each has its own equivalence class that may be navigated independently. Another degree of freedom for exploring grazing-angle effects is offered by changing the lobe order. This is shown in Figure 17 for a pair of lobes whose order is swapped, yielding either a reddish or blueish tint at grazing angles. Note that by doing so, our approach entirely recomputes the equivalence classes, shown on the left and right sides of the figure. Using our python implementation on an Intel Core i7-6700HQ CPU, each configuration in Figure 17 takes less than 50ms (see function `invert_multiple_lobes`).

We show a validation of our method in Figure 25 of the Appendix, in which we compare reflectance and variance obtained using our method to the desired lobe properties.

Fig. 16. Within a conductor equivalence class, edge tint is more apparent for low η_1 and high τ_1 . In all images we compare pairs of 2-layer conductor materials from the same equivalence class ($R_1 = 0.45$, $\sigma_1 = 0.03$). Each comparison is made to a same reference material (black square). In (a), we use a smaller value of η_1 (grey square), which leads to a barely noticeable difference (left/right halves). In (b), we instead increase τ_1 for the green channel only, letting appear a pinkish edge tint (pink square). In (c), we apply both variations in η_1 and τ_1 , making the edge tint more prominent.

Fig. 17. Changing the order of lobes also affects grazing-angle colors. The materials in each half are made of a pair of rough 2-layer lobes: one is blue ($R_b = [0.025, 0.045, 0.05]$, $\sigma_b = 0.028$), the other red ($R_r = [0.052, 0.025, 0.025]$, $\sigma_r = 0.014$). Since the sum of the lobes stays the same at normal incidence, the two materials appear identical in the center of the sphere. When the blue lobe is on top of the red one (left half), the edge tint is red due to absorption; whereas it is blue when the order is reversed (right half).

7.2 Exploring complex appearance

Much more dramatic changes in appearance are produced when a material is pushed outside of its initial equivalence class. Thanks to our method, this is easily achieved by editing the input lobe parameters R_k and σ_k , while physical parameters are *automatically retrieved* to build layer stacks that achieve the desired appearance.

Figure 18 shows a first example with a material made of a single 2-layer dielectric lobe, rendered in the Ennis environment. When the overall color saturation or roughness of the material are either increased or decreased, the physical parameters of the two layers are automatically updated to maintain the appearance of a material made of a single lobe. We show in supplemental material the effects of directly manipulating physical parameters on the same scene.

In terms of appearance, the main advantage of having more than one lobe is to produce hazy reflections. We use the Wells and Labo

Fig. 18. *Reflection color and sharpness control.* The Queen chess piece is made of a 2-layer dielectric material ($\mathbf{R}_1 = [0.03, 0.06, 0.09]$, $\sigma_1 = 0.05$). In (a), we change its color saturation, either decreasing it on the Rook piece ($\mathbf{R}_1 = [0.06, 0.06, 0.06]$) or increasing it on the Knight piece ($\mathbf{R}_1 = [0.01, 0.06, 0.11]$, right). In (b) we change its variance, either decreasing ($\sigma_1 = 0$) or increasing it ($\sigma_1 = 0.15$) on the Rook and Knight pieces respectively.

Fig. 19. *Haze effect control.* Starting from (b) a smooth 1-layer dielectric coat ($R_1 = 0.1$, $\sigma_1 = 0$) over a rough 1-layer conductor base ($\mathbf{R}_2 = [0.35, 0.25, 0.3]$, $\sigma_2 = 0.04$) produces a hazy appearance due to the use of different variances. The variance of the base lobe is then easily decreased ($\sigma_2 = 0.01$) to yield (a) a subtle haze effect, or increased ($\sigma_2 = 0.1$) to yield (c) a strong haze effect.

environment lightings as they contain large bright areas with sharp boundaries that help convey haze effects. Figure 19 demonstrates how the extent of the haze effect may be controlled directly through the variance of the base lobe. In this specific case, once an initial hazy conductor has been obtained through our method, varying the haze only affects the roughness of the bottom interface (this is not the case with a 2-layer base lobe, as explained in supplemental material). Compared to a single interface with a mix of microfacet distributions [Barla et al. 2018], a layered material offers the opportunity to choose the base lobe (and hence the haze) color independently. Figure 20 demonstrates this with a pair of materials, each made of two dielectric lobes: a smooth 2-layer coat over a rough 1-layer base. The materials only differ in which of the lobe energies ($\mathbf{R}_1$ or $\mathbf{R}_2$) is colored, yielding either colored highlights or colored halos. This is best seen in inset images that show locations on the object where the materials differ the most. In Figure 21, we consider a pair of materials each made of an achromatic 1-layer dielectric lobe over a 1-layer conductor base. The materials only differ in which of the lobe is smooth, and which is rough (i.e., we swap σ_1 with σ_2), yielding different highlight profiles, as seen in the inset images.

Even though a diverse range of appearance can be achieved with $K = 2$ lobes, we show in Figure 22 an example with $K = 3$ dielectric

Fig. 20. *Haze color control.* Starting from (left) an initial material made of two lobes ($\mathbf{R}_1 = [0.09, 0.14, 0.09]$, $\sigma_1 = 0.005$, $R_2 = 0.09$, $\sigma_2 = 0.1$, dielectric base), we obtain (right) a material where the lobe colors $\mathbf{R}_1$ and $\mathbf{R}_2$ are swapped, as shown in the bottom row where independent lobe contributions are shown.

Fig. 21. *Haze profile control.* Starting from (left) a material made of two lobes ($R_1 = 0.04$, $\sigma_1 = 0.07$, $\mathbf{R}_2 = [0.6, 0.25, 0.25]$, $\sigma_2 = 0.03$, cond. base), we obtain (right) a material where the lobe variances σ_1 and σ_2 are swapped, as shown in the bottom row where we show independent lobe contributions.

lobes. The resulting appearance is similar to an additive blending of the configurations using $K = 2$ lobes, which opens interesting avenues for the mixing of BRDFs using layered materials. Further explanations on Figure 22 are given in supplementary material.

8 DISCUSSION

To recap, our method allows to specify the properties of BRDF lobes produced by layered materials, instead of iterating over high-dimensional physical parameter sets. It is particularly useful for computing configurations where two layers yield a single lobe (Section 5), or for stacking several lobes while preserving their appearance (Section 6). Our method not only provides unprecedented control over the color, blur and haze of reflections, but it also finds physical layer configurations that only differ at grazing angles.

Fig. 22. *Stacking up to 3 lobes.* The top two rows show materials made of two lobes: a smooth 1-layer dielectric coating ($R_1 = 0.03$) onto either of two rough dielectrics: a yellowish 2-layer lobe (top, $\mathbf{R}_Y = [0.1, 0.1, 0.06]$, $\sigma_Y = 0.05$) or a blueish 1-layer lobe (middle, $\mathbf{R}_B = [0.05, 0.05, 0.09]$, $\sigma_B = 0.01$). The last row uses all three lobes, with the yellowish lobe at the base and the blueish lobe in the middle.

Limitations. We have focused on lobe properties at normal incidence, which in retrospect seems reasonable owing to the subtle grazing-angle effects observed in our results overall. In addition, Belcour’s model is most accurate at normal incidence [Bati et al. 2019; Belcour 2018]. Nevertheless, our method is limited in the control it offers over grazing angles, as it relies on the exploration of equivalence classes. We have initially tried to give more direct control at grazing angles, but faced a difficult obstacle: any angle different from $\theta_0 = 0$ will be affected by Snell’s law of refraction, which itself depends on the refractive indices to be retrieved.

A limitation inherent to physically-realistic layered materials is that some lobe properties may not be achievable; then *constrained* equivalence classes become empty. Unfortunately, we have not always found analytical bounds on energy or variance to guarantee that a physically-achievable solution exist. This might complicate the exploration of configurations where only a very small subset of physical parameter space yields a valid solution (e.g., Figure 17).

We have had to make choices on the type of lobes. For instance, the extension to a 3-layer lobe is achievable in principle, but it would expand the dimensionality of equivalence classes, without necessarily offering significant changes in appearance. An important limitation is that we only perform inversion on reflection; hence transmission is obtained as a by-product as shown in Figure 23 and in the supplemental material. Since we rely on Belcour’s model,

Fig. 23. We compare two sets of physical parameters within (a) the 2-layer lobe equivalence class of Figure 15b ($\mathbf{R} = [0.09, 0.09, 0.15]$, $\sigma = 0$). Although reflection is identical and directly editable, transmission is obtained as a by-product with different appearance between (a) low and (b) high absorption.

Fig. 24. Rendering (a) a 2-layer dielectric material ($\mathbf{R} = [0.03, 0.06, 0.09]$, $\sigma = 0.05$) is equivalent to assuming a perfectly absorbing diffuse base layer. It may be replaced by (b) an achromatic Lambertian base ($\rho_d = 0.5$), or (c) a colored Lambertian base ($\rho_d = (0.5, 0.3, 0.1)$), hence affecting the overall color and giving a plastic look to the material. However, our approach neither provide control nor physically-achievable bounds on the diffuse energy, which may be very different from ρ_d , as seen in (c).

we cannot handle either Lambertian base or general scattering media. Of course, we may still render layered materials obtained with our approach and stacked onto a diffuse layer, as shown in Figure 24. Even though the extension of Belcour’s model to diffusion and scattering seems within reach in future work, the handling of interferences among layers would require a different approach.

Future Work. Even though our approach is well adapted to appearance exploration, some artists will likely prefer a more direct way to indicate a specific material appearance. One solution would be to restate the problem as an optimization that tries to find the physical configuration that comes as close as possible to the desired appearance. As mentioned by Zsolnai et al. [2020], a good initial solution is needed for optimization, which could be obtained using our method. We believe that the space of equivalence classes uncovered in this paper will likely be useful in this endeavour, as it permits to control the appearance at normal incidence. It could also provide a framework for the fitting of measured layered BRDFs.

Another interesting avenue is material simplification: a BRDF resulting from many interfaces and media stacked together may alternatively be obtained by fewer layers. This is similar in spirit to the use of similarity relations [Zhao et al. 2014] to reduce the amount of multiple scattering in media.

Finally, we have implicitly assumed that material appearance is inherently dependent on BRDF lobe properties. However, it also

depends on shape and environment lighting, as shown in supplemental material. Our approach remains of course valid if one is interested in designing a material irrespective of shape or lighting. Otherwise, material perception should also be taken into account, which goes beyond the scope of this paper. Our method might still be relevant as a tool to perform perceptual experiments, by guaranteeing physical-realism while providing direct control over BRDF lobes instead of indirect control over physical parameters.

ACKNOWLEDGMENTS

We thank Laurent Belcour for his help and advice. We use the Wells and Labo environment maps, properties of Bernhard Vogl; the Material preview scene of Jonas Pilo, the Chess and Maneki-neko models of RescueStudios11 and Laccolith on Blendswap. This work has been supported by the ANR VIDA project (ANR-17-CE23-0017).

REFERENCES

- F. Abelès. 1948. Sur la propagation des ondes électromagnétiques dans les milieux stratifiés. *Ann. Phys.* 12, 3 (1948), 504–520.
- B. F. Armaly, J. G. Ochoa, and D. C. Look. 1972. Restrictions on the Inversion of the Fresnel Reflectance Equations. *Appl. Opt.* 11, 12 (1972), 2907–2910.
- P. Barla, R. Pacanowski, and P. Vangorp. 2018. A Composite BRDF Model for Hazy Gloss. *Computer Graphics Forum* 37, 4 (2018), 55–66.
- M. Bati, R. Pacanowski, and P. Barla. 2019. *Numerical Analysis of Layered Materials Models*. Research Report.
- P. Beckmann and A. Spizzichino. 1963. *The scattering of electromagnetic waves from rough surfaces*. Pergamon Press; [distributed in the Western Hemisphere by Macmillan, New York] goford, New York. viii, 503 p. pages.
- L. Belcour. 2018. Efficient Rendering of Layered Materials Using an Atomic Decomposition with Statistical Operators. *ACM Trans. Graph.* 37, 4, Article 73 (2018), 15 pages.
- M. Colbert, S. Pattanaik, and J. Krivanek. 2006. BRDF-Shop: creating physically correct bidirectional reflectance distribution functions. *IEEE Computer Graphics and Applications* 26, 1 (2006), 30–36.
- Q. Dai, J. Wang, Y. Liu, J. Snyder, E. Wu, and B. Guo. 2009. The Dual-microfacet Model for Capturing Thin Transparent Slabs. *Computer Graphics Forum* 28, 7 (2009), 1917–1925.
- V. Deschaintre, M. Aittala, F. Durand, G. Drettakis, and A. Bousseau. 2018. Single-Image SVBRDF Capture with a Rendering-Aware Deep Network. *ACM Trans. Graph.* 37, 4, Article 128 (2018), 15 pages.
- S. Ergun, S. Onel, and A. Ozturk. 2016. A General Micro-flake Model for Predicting the Appearance of Car Paint. In *Proc. of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations (EGSR '16)*. Eurographics Association, 65–71.
- S. Ershov, R. Durikovic, K. Kolchin, and K. Myszkowski. 2004. Reverse engineering approach to appearance-based design of metallic and pearlescent paints. *The Visual Computer* 20 (2004), 586–600.
- S. Ershov, K. Kolchin, and K. Myszkowski. 2001. Rendering Pearlescent Appearance Based On Paint-Composition Modelling. *Computer Graphics Forum* 20, 3 (2001), 227–238.
- L. E. Gamboa, A. Gruson, and D. Nowrouzezahrai. 2020. An Efficient Transport Estimator for Complex Layered Materials. *Computer Graphics Forum* 39, 2 (2020), 363–371.
- I. Georgiev, J. Portsmouth, Z. Andersson, A. Herubel, A. King, S. Ogaki, and F. Servant. 2019. A Surface Standard. <https://autodesk.github.io/standard-surface/>.
- J. Gu, R. Ramamoorthi, P. Belhumeur, and S. Nayar. 2007. Dirty Glass: Rendering Contamination on Transparent Surfaces. In *Eurographics Symposium on Rendering (EGSR '07)*. Eurographics Association, 159–170.
- O. Gulbrandsen. 2014. Artist Friendly Metallic Fresnel. *Journal of Computer Graphics Techniques (JCGT)* 3, 4 (2014), 64–72.
- J. Guo, J. Qian, Y. Guo, and J. Pan. 2017. Rendering Thin Transparent Layers with Extended Normal Distribution Functions. *IEEE Transactions on Visualization & Computer Graphics* 23, 9 (2017), 2108–2119.
- Y. Guo, M. Hašan, and S. Zhao. 2018. Position-Free Monte Carlo Simulation for Arbitrary Layered BSDFs. *ACM Trans. Graph.* 37, 6, Article 279 (2018), 14 pages.
- H. Hirayama, K. Kaneda, H. Yamashita, and Y. Monden. 2001. An accurate illumination model for objects coated with multilayer films. *Computers & Graphics* 25, 3 (2001), 391–400.
- I. Icart and D. Arques. 2000. A Physically-Based BRDF Model for Multilayer Systems with Uncorrelated Rough Boundaries. In *Proc. of the Eurographics Workshop on Rendering Techniques 2000*. Springer-Verlag, 353–364.
- W. Jakob, E. d'Eon, O. Jakob, and S. Marschner. 2014. A Comprehensive Framework for Rendering Layered Materials. *ACM Trans. Graph.* 33, 4, Article 118 (2014), 14 pages.
- W. Matusik, H. Pfister, M. Brand, and L. McMillan. 2003. A Data-Driven Reflectance Model. *ACM Trans. Graph.* 22, 3 (2003), 759–769.
- A. Ngan, F. Durand, and W. Matusik. 2006. Image-Driven Navigation of Analytical BRDF Models. In *Symposium on Rendering (EGSR '06)*. Eurographics Association, 399–407.
- G. Patow and X. Pueyo. 2003. A Survey of Inverse Rendering Problems. *Computer Graphics Forum* 22, 4 (2003), 663–687.
- F. Pellacini, J. A. Ferwerda, and D. P. Greenberg. 2000. Toward a Psychophysically-Based Light Reflection Model for Image Synthesis. In *Proc. of the 27th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '00)*. ACM Press/Addison-Wesley Publishing Co., 55–64.
- M. R. Querry. 1969. Direct Solution of the Generalized Fresnel Reflectance Equations. *J. Opt. Soc. Am.* 59, 7 (1969), 876–877.
- J. Riviere, I. Reshetouski, L. Filipi, and A. Ghosh. 2017. Polarization Imaging Reflectometry in the Wild. *ACM Trans. Graph.* 36, 6, Article 206 (2017), 14 pages.
- A. Serrano, D. Gutierrez, K. Myszkowski, H. P. Seidel, and B. Masia. 2016. An intuitive control space for material appearance. *ACM Trans. Graph.* 35, 6, Article 186 (2016), 12 pages.
- L. Simonot, M. Hébert, and R. D. Hersch. 2006. Extension of the Williams-Clapper model to stacked nondiffusing colored coatings with different refractive indices. *J. Opt. Soc. Am. A* 23, 6 (2006), 1432–1441.
- G. G. Stokes. 1862. IV. On the intensity of the light reflected from or transmitted through a pile of plates. *Proc. of the Royal Society of London* 11 (1862), 545–556.
- B. Walter, S. Marschner, H. Li, and K. Torrance. 2007. Microfacet Models for Refraction Through Rough Surfaces. In *Eurographics Symposium on Rendering (EGSR '07)*. Eurographics Association, 195–206.
- G. J. Ward. 1992. Measuring and Modeling Anisotropic Reflection. *SIGGRAPH Comput. Graph.* 26, 2 (1992), 265–272.
- A. Weidlich and A. Wilkie. 2007. Arbitrarily Layered Micro-facet Surfaces. In *Proc. of the 5th International Conference on Computer Graphics and Interactive Techniques in Australia and Southeast Asia (GRAPHITE '07)*. ACM, 171–178.
- H. Wu, J. Dorsey, and H. Rushmeier. 2013. Inverse bi-scale material design. *ACM Trans. Graph.* 32, 6, Article 163 (2013), 10 pages.
- M. Xia, B. Walter, C. Hery, and S. Marschner. 2020. Gaussian Product Sampling for Rendering Layered Materials. *Computer Graphics Forum* 39, 1 (2020), 420–435.
- P. Yeh. 2005. *Optical Waves in Layered Media*. Number v. 2 in Wiley Series in Pure and Applied Optics. Wiley.
- T. Zeltner and W. Jakob. 2018. The Layer Laboratory: A Calculus for Additive and Subtractive Composition of Anisotropic Surface Reflectance. *ACM Trans. Graph.* 37, 4, Article 74 (2018), 14 pages.
- S. Zhao, R. Ramamoorthi, and K. Bala. 2014. High-order similarity relations in radiative transfer. *ACM Trans. Graph.* 33, 4 (2014), 104:1–104:12.
- K. Zsolnai-Fehér, P. Wonka, and M. Wimmer. 2020. Photorealistic Material Editing Through Direct Image Manipulation. *Computer Graphics Forum* 39, 4 (2020), 107–120.

A REPRODUCIBILITY AND VALIDATION

We provide supplemental code to reproduce the figures in the paper. We give an exhaustive list of lobe properties and physical parameters in supplemental material. Finally, we validate our approach by comparing expected and obtained lobe properties in Figure 25.

Fig. 25. Comparison between expected and obtained reflectance (top) and variance (bottom) for Figures 15, 16 and 17. Our method is accurate, the slight energy difference observed for conductors (in Figure 16) is due to the low resolution in κ of the FGD table provided in supplemental material.