

The Layer Laboratory: A Calculus for Additive and Subtractive Composition of Anisotropic Surface Reflectance

TIZIAN ZELTNER and WENZEL JAKOB, École polytechnique fédérale de Lausanne (EPFL)



Fig. 1. All materials in this rendering of an interior scene were generated using the techniques proposed in this article. Insets on the left side reveal the corresponding layer structures, which make use of dielectric and conductive microfacet models, scattering and absorbing volumes, and measured materials acquired using dense goniophotometric measurements.

We present a versatile computational framework for modeling the reflective and transmissive properties of arbitrarily layered anisotropic material structures. Given a set of input layers, our model synthesizes an effective BSDF of the entire structure, which accounts for all orders of internal scattering and is efficient to sample and evaluate in modern rendering systems.

Our technique builds on the insight that reflectance data is sparse when expanded into a suitable frequency-space representation, and that this property extends to the class of anisotropic materials. This sparsity enables an efficient matrix calculus that admits the entire space of BSDFs and considerably expands the scope of prior work on layered material modeling. We show how both measured data and the popular class of microfacet models can be expressed in our representation, and how the presence of anisotropy leads to a weak coupling between Fourier orders in frequency space.

In addition to additive composition, our models supports subtractive composition, a fascinating new operation that reconstructs the BSDF of a material that can only be observed indirectly through another layer with known reflectance properties. The operation produces a new BSDF of the desired layer as if measured in isolation. Subtractive composition can be

interpreted as a type of deconvolution that removes both internal scattering and blurring due to transmission through the known layer.

We experimentally demonstrate the accuracy and scope of our model and validate both additive and subtractive composition using measurements of real-world layered materials. Both implementation and data will be released to ensure full reproducibility of all of our results.¹

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

Additional Key Words and Phrases: Layered materials, adding equations, subtracting equations, BRDF, BSDF

ACM Reference Format:

Tizian Zeltner and Wenzel Jakob. 2018. The Layer Laboratory: A Calculus for Additive and Subtractive Composition of Anisotropic Surface Reflectance. *ACM Trans. Graph.* 37, 4, Article 74 (August 2018), 14 pages. <https://doi.org/10.1145/3197517.3201321>

© 2018 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, <https://doi.org/10.1145/3197517.3201321>.

¹ Available at <http://rgl.epfl.ch/publications/Zeltner2018Layer>

1 INTRODUCTION

Photorealistic images created using computer simulations have become an integral part of our society: humans are fundamentally visual creatures, hence almost any creative process requires our ability to reason about the visual nature of objects or materials that are yet to be created. Given a detailed description of materials, shapes, and their environment, physically-based rendering algorithms can produce images that are indistinguishable from photographs of manufactured objects subject to the same specifications, and they have therefore become the tool of choice for applications that require such predictive capabilities.

High-quality material models are a crucial ingredient in this pursuit of realism: they encode the precise spatial and directional displacement of light following a scattering interaction, allowing an object to be visualized in an arbitrary simulated environment. Modern rendering systems typically ship with a large catalog of standard parametric models that reproduce specific phenomena required by a given scene. Although tremendous advances over the last decades have produced models that are in excellent agreement with measured data, they are almost always limited to basic physical phenomena observed in isolation, e.g. reflection by surfaces with microscopic imperfections, refraction due to index of refraction changes, diffraction by thin films, etc. This is a serious impediment, as most real-world materials exhibit behaviors that are significantly more complex than these idealized models, involving multiple different types of interactions along with internal scattering and a nonlinear dependence on the internal material composition.

In this article, we focus on the class of *layered materials*, which are composed of a stack of layers that can each be clear or subject to internal scattering and/or absorption. Each pair of adjacent layers is separated by a smooth or rough interface representing an optional index of refraction change. Classic examples of layered materials include any type of painted or glazed surface (ceramic tiles, car paint, finished wood, paint on primer, etc.), dielectric slabs with two interfaces (e.g. frosted glass panels), potentially with significant internal scattering (e.g. most types of plastics), and organic materials such as leaves or skin. Layered material models are appealing because each new type of layer or interface fuels a combinatorial explosion that greatly expands the space of material models that are available for use in a physically-based renderer. Thinking in terms of layers also creates a helpful connection between material appearance and composition that is absent in models that merely attempt to fit or interpolate reflectance data.

Previous works on layered materials have proposed heuristic approximations [Weidlich and Wilkie 2007] and numerical methods for special cases, such as human skin [Stam 2001] and the class of isotropic materials [Jakob et al. 2014]. The latter work is the most accurate and general approach to date, but the restriction to isotropy constitutes a significant limitation. Virtually all types of industrial manufacturing processes—sawing, milling, turning, drawing, honing, grinding, etc.—involve some type of tool or mask that interacts with the material in a directional sense, producing a rich spectrum of subtle to extreme anisotropic surface appearance. Anisotropy is thus the norm and not the exception, which motivates

our technique: we want to extend the layered material framework so that it can represent and transform the entire space of BSDFs.

Our approach heavily relies on a directional Fourier basis, whose convenient numerical properties enable efficient layering computations for material parameters ranging from mild to pronounced anisotropy. A key aspect of our representation is that anisotropy gives rise to a sparse coefficient structure, which allows for efficient matrix computations involving banded block matrices, typically containing 0.01 – 0.5% nonzero entries.

A crucial aspect of prior work [Jakob et al. 2014] was that the Fourier treatment caused the azimuthal and latitudinal dimensions to separate into independent lower-dimensional sub-problems, which made the layering computations tractable. In the anisotropic case, we find that this separation no longer takes place—instead, the sub-problems are coupled and must be solved jointly. Fortunately, this coupling is *weak* in the low frequency sense, which we exploit to create a stable and efficient numerical solver for arbitrary layered structures.

Our matrix representation admits additive composition via the well-known *adding equations*, a highly nonlinear process that computes the aggregate properties of a layer (or layered structure) observed through another layer (or layered structure), while accounting for arbitrary orders of inter-reflection.

However, our representation is not just useful for combining layers: it can also be used to solve a variety of interesting inverse problems using algebraic transformations. We introduce the *subtracting equations*, which remove either the top or bottom of an arbitrarily layered structure, revealing the layers underneath. Being the inverse of the adding equations, subtractive layer composition must undo the transmissive effects (e.g. directional blur) of the removed layer, while also removing complex nonlinearities due to multiple scattering. Naturally, this is only feasible when enough information is available: a specular material below an opaque or diffuse transmissive layer cannot be reconstructed. Like any type of deconvolution, subtractive composition is a challenging inverse problem that requires regularization to ensure robustness to noise. We believe that subtractive composition could be an asset in situations where an individual layer is hard or impossible to observe directly, e.g. when attempting to measure a single interface of a refractive material that also has a back surface. Other interesting application areas are reverse engineering of materials and cultural heritage preservation, where the desired removal of a layer (e.g. lacquer) may be destructive to an artifact.

Our contributions are as follows:

- We propose a directional representation that enables layered composition of anisotropic materials.
- We derive two types of subtracting equations, which constitute an inverse of the adding equations.
- We show how to project the popular class of microfacet models into this representation and discuss practical considerations to improve efficiency and robustness.
- We validate additive and subtractive composition using Monte Carlo simulations and dense ground-truth measurements of real-world layered materials performed using a gonio-photometer.

2 PRIOR WORK

Blinn [1982] and Hanrahan and Krueger [1993] introduced the first layered material models to computer graphics; their BRDFs account for a single scattering event in a homogeneous medium, the latter potentially involving several layers with different scattering properties and indices of refractions.

A number of BRDF models combine diffuse and specular terms to approximate the optics of a single dielectric layer with multiple internal scattering, an important special case that is representative of extremely common material types including plastics and ceramics. Examples include the classic Phong [1975], Blinn-Phong [1977], and Ward [1992] BRDFs as well as later works by Shirley [1997] and Kelemen [2001] that consider the effects of the dielectric interface to distribute energy between the specular and diffuse portions.

Beckmann and Spizzichino [1963] introduced the idea of modeling scattering due to microscopic surface imperfections by representing the surface as a random process with Gaussian autocorrelation. Building on this work, Torrance, Sparrow, and Cook [1982; 1967] introduced microfacet BRDFs, which model the surface as a complex specular reflector with shadowing and masking. An extension by Walter et al. [2007] made this approach compatible with the refractive case. Due to their simplicity and good agreement with real materials, microfacet models have become the predominant way of representing roughness at layer boundaries. One limitation is their builtin assumption that light will scatter at most once at the interface, which ceases to be the case as the surface roughness increases. Heitz et al. [2016] recently proposed a volumetric generalization of microfacet theory that addresses this flaw. Dai et al. [2009] proposed a specialized layered model for a two-sided dielectric with rough microfacet boundaries. Burley's Disney BRDF [2012] has become a popular microfacet variant due to its artist-friendly controls and ability to smoothly interpolate between diffuse and specular as well as reflective and transmissive conditions. The model supports an optional coating layer but does not account for interreflection between the coating and substrate.

Structures with layer thicknesses on the order of the wavelength of light are subject to wave-optical interference, which causes a strong coupling between the wavelength, intensity, and direction of scattered light. This mechanism is exploited in anti-reflection coatings for optical elements, dichroic filters for microscopy and effect coatings for architectural applications. BRDF models for multi-layer iridescence include the works of Hirayama [2001], Icart [2000] and Ershov [2001]. The tools developed in this article target macroscopic layers, whose behavior can be described using geometric optics. Uncorrelated iridescence effects that are local to each layer could be incorporated using an extension to microfacet models proposed by Belcour and Barla [2017].

BRDF models for macroscopic layered structures were proposed by Stam [2001] for the special case of human skin and Weidlich and Wilkie [2007] for arbitrarily layered surfaces. Their model can in principle support anisotropy but does not account for multiple internal scattering. The layering framework by Jakob et al. [2014] is closest to the methods described in this article but is restricted to isotropy. We considerably extend their framework to robustly

support anisotropy at moderate additional cost and perform comprehensive validations against real-world measurements of layered structures.

We are not aware of any works that have attempted to reconstruct the full BRDF of an indirectly observed material. Most closely related are techniques that use special sensing modalities such as time-of-flight cameras to classify BRDFs [Su et al. 2016] or fit simple parametric models [Naik et al. 2011]. Tanaka et al. [2015] extract the spatially varying albedo of a layered structure using high-frequency projector illumination.

3 ANISOTROPIC LAYER CALCULUS

This section introduces our layering calculus for anisotropic materials, starting with the frequency-space representation that lies at its foundation. We also discuss relevant background material and point out differences to prior work where applicable.

3.1 Representation of radiance

We denote the radiance in a layered structure as $\Phi(\mu, \phi)$, where $\mu = \cos \theta \in [-1, 1]$ is the cosine of the elevation angle and ϕ is the azimuth measured in the layer's tangent frame. Although regular spherical coordinates could also be used, we prefer this formulation as it simplifies some of the subsequent expressions and numerical computations.

We discretize radiance using a representation originally proposed by Chandrasekhar [1960], which is particularly well-suited for layering computations that must evaluate large quantities of numerical integrals. It samples Φ at a fixed number of elevation cosines μ_1, \dots, μ_n that are chosen based on the resolution requirements of the material under consideration. Instead of placing the elevation samples uniformly, Chandrasekhar uses a numerical quadrature scheme, i.e. integration nodes $\{\mu_1, \dots, \mu_n\} \subset [-1, 1] \setminus \{0\}$ and weights $\{w_1, \dots, w_n\}$ so that integrals over elevations angles are easily approximated as

$$\int_0^\pi f(\theta) \sin \theta \, d\theta = \int_{-1}^1 f(\mu) \, d\mu \approx \sum_{k=1}^n w_k f(\mu_k). \quad (1)$$

Similar to prior work [Jakob et al. 2014], we use a Gauss-Lobatto quadrature rule, which maximizes the order of exactly integrable polynomials subject to the constraint of including the endpoints (i.e. $\mu_1 = -1$ and $\mu_n = 1$).


Rather than discretizing azimuths at a similar set of discrete positions, they must be projected into frequency space: this is key to the efficiency of the method, as will be seen later. Both Chandrasekhar and Jakob et al. expand the azimuthal dependence of the radiance function into a Fourier series containing only the even cosine terms:

$$\Phi(\mu_i, \phi) = \sum_{l=1}^{\infty} \Phi_l(\mu_i) \cos(l\phi).$$

This is possible in their case due to the restriction to isotropic materials: scattered light is symmetric around the incident azimuth causing odd Fourier terms to vanish. However, this property no longer holds when working with anisotropic layers.

Table 1. Overview of used notation

Term	Meaning
$\Phi(\mu, \phi)$	Continuous radiance function
θ_i, θ_o	Incident and outgoing elevation angle
μ_i, μ_o	Incident and outgoing elevation angle cosine
ϕ_i, ϕ_o	Incident and outgoing azimuth angle
$\phi_s = \phi_o + \phi_i$	Azimuthal sum
$\phi_d = \phi_o - \phi_i$	Azimuthal difference
n	Number of discretizations in μ
m_s	Number of Fourier coefficients in ϕ_s
m_d	Number of Fourier coefficients in ϕ_d
$\delta_{i,j}$	Kronecker delta
α_u, α_v	Anisotropic Beckmann roughness in u and v
f	Surface BSDF of a layer
$\Phi_j(\mu)$	Fourier expansion of Φ
Φ_j	μ -discretization of $\Phi_j(\mathbb{R}^n)$
$f_{s,d}(\mu_i, \mu_o)$	2D Fourier expansion of f
$F_{s,d}$	μ -discretization of $f_{s,d}(\mathbb{R}^{n \times n})$



In contrast, our method represents this non-symmetry using a complex exponential series of the form

$$\Phi(\mu_i, \phi) = \sum_{l \in \mathbb{Z}} \Phi_l(\mu_i) \exp(il\phi), \quad (2)$$

which is somewhat redundant given that Φ is always real-valued (in particular, $\Phi_l = \Phi_{-l}$). This symmetry can be exploited by the implementation to reduce storage and computation overheads roughly by half. By convention, light traveling along a ray with $\mu > 0$ is considered to propagate deeper into the layered structure, while directions with $\mu < 0$ correspond to light traveling towards the top surface.

3.2 Representation of Layers

We partition the incident radiance Φ^i at a given layer into directions that are incident at the top ($\mu > 0$) and bottom ($\mu < 0$) surfaces. Scattered radiance Φ^o uses the same convention, i.e. light leaving the bottom surface satisfies $\mu > 0$. Due to linearity, the effect of a single layer on light arriving on both sides can then be expressed via the following integral equation

$$\Phi^o(\mu_o, \phi_o) = \int_0^{2\pi} \int_{-1}^1 \Phi^i(\mu_i, \phi_i) f(\mu_i, \phi_i, \mu_o, \phi_o) |\mu_i| d\mu_i d\phi_i, \quad (3)$$

where f characterizes the flow of energy. What constitutes a layer is intentionally broad: it could be a volume with internal scattering or a single smooth or rough surface separating two regions of space with different indices or refraction. In the case of a surface layer, f is the well-known *bidirectional scattering distribution function* (BSDF).

The four-dimensional nature of f leads to a number of practical challenges: dense discretizations in $(\mu_i, \phi_i, \mu_o, \phi_o)$ are typically needed to represent arbitrary layers accurately, but this leads to impracticably high storage requirements. The 4D space can also be re-parameterized using alternative coordinates with superior properties, allowing for coarser discretizations, e.g. using the half-vector mapping of Rusinkiewicz [1998]. The downside of such warped representations is that they impede efficient evaluation of integrals over (μ_i, ϕ_i) , a crucial operation in layered material models.

In the isotropic case, the optical behavior of the layer is invariant under rotation around the surface normal, and f therefore simplifies to a three-dimensional function $f(\mu_i, \mu_o, \phi_d)$ involving the azimuth difference $\phi_d := \phi_i - \phi_o$. Prior work takes advantage of this property by expanding f into a Fourier basis in ϕ_d , i.e.

$$f(\mu_i, \mu_o, \phi_d) = \sum_{l=0}^{\infty} f_l(\mu_i, \mu_o) \cos(l\phi_d), \quad (4)$$

and inserting the expansion into (3). This step leads to a remarkable simplification: the two-dimensional illumination integral decouples into an infinite series of simpler one-dimensional integrals that only make reference to elevation angles.

$$\Phi_l^o(\mu_o) = \pi(1 + \delta_{0l}) \int_{-1}^1 \Phi_l^i(\mu_i) f_l(\mu_i, \mu_o) |\mu_i| d\mu_i. \quad (l = 0, \dots) \quad (5)$$

Each equation characterizes the material's response to a particular oscillation mode in azimuth. We refer to Jakob et al. [2014] for details on this derivation.

In the anisotropic setting, the difference angle ϕ_d no longer fully characterizes the azimuthal behavior of a layer, hence a different set of parameters is needed. Instead of reverting to the original incident and outgoing azimuth values, we found it preferable to retain the difference angle and complete it to an orthogonal basis, i.e.:

$$\phi_d := \phi_o - \phi_i, \quad \phi_s := \phi_o + \phi_i, \quad (6)$$

where ϕ_s is the sum of azimuths. Holding ϕ_d constant and changing ϕ_s by δ radians is equivalent to rotating the layer by $\delta/2$ radians around the surface normal. For each incident and outgoing elevation, the anisotropic frequency-space representation of f now turns into a two-dimensional array of Fourier coefficients

$$f(\mu_i, \phi_i, \mu_o, \phi_o) = \sum_{d,s \in \mathbb{Z}} f_{d,s}(\mu_i, \mu_o) e^{i d (\phi_o - \phi_i)} e^{i s (\phi_o + \phi_i)}. \quad (7)$$

It is instructive to consider behavior of a microfacet model in this parameterization. In the case of an isotropic material with roughness parameter α , the difference coordinate ϕ_d suffices to fully characterize the dependence on azimuth. Generalizing to an anisotropic material ($\alpha_v = c \cdot \alpha_u$) for some constant $c > 1$ introduces blur, which cannot increase the frequency content in ϕ_d . The sum coordinate ϕ_s is now needed to account for anisotropy: it describes how the (formerly isotropic) Fourier coefficients must update as the material rotates about the surface normal. Figure 2 visualizes

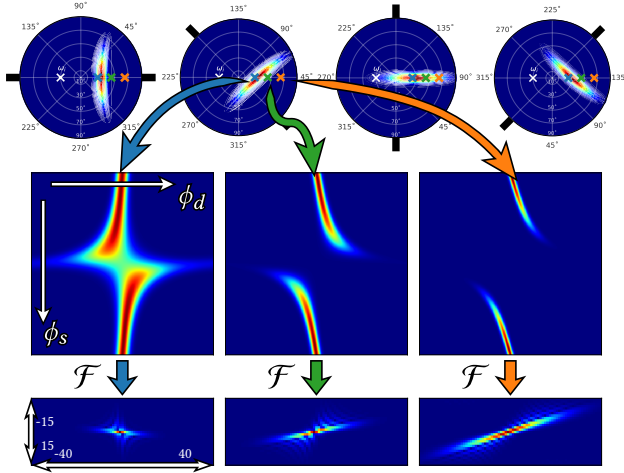


Fig. 2. Frequency-space response of a conductive layer with anisotropic Beckmann roughness $\alpha = (0.05, 0.3)$. **Top row.** Four BRDF slices for $\theta_i = 30^\circ$ visualized on the projected unit disk. Three outgoing directions are highlighted, and the bold black line denotes rotation around the surface normal. **Middle row.** Matching BRDF plots in the (ϕ_s, ϕ_d) parameterization. **Bottom row.** Fourier transform of the above row. The Fourier coefficients are sparse and quickly decay when moving away from a central band. For most configurations, the high-frequency behavior is predominantly constrained to the ϕ_d parameter.

the low-frequency behaviour of ϕ_s —it causes a slight widening and shift of the lobe, requiring comparatively few coefficients.

Another important observation is that the Fourier coefficients rapidly decay when moving away from a central band that is centered around zero—they are *sparse* in the (d, s) -domain. Our implementation exploits this property to accelerate layering computations.

Surface illumination integral. Having decided on a basis, we must now project the surface illumination integral onto this representation to derive a linear operator that encapsulates the optics of a single layer. Let

$$\mathcal{F}_l[g(\cdot)] := \frac{1}{2\pi} \int_{-\pi}^{\pi} g(\phi) e^{-il\phi} d\phi$$

denote the l -th complex Fourier coefficient of a 2π -periodic function g . Transforming both sides of Equation 3 and reversing the order of integration yields

$$\Phi_l^o(\mu_o) = \int_{-1}^1 \mathcal{F}_l \left[\int_0^{2\pi} \Phi^i(\mu_i, \phi_i) f(\mu_i, \phi_i, \mu_o, \cdot) d\phi_i \right] |\mu_i| d\mu_i.$$

Neglecting the μ -dependence for readability, the Fourier integral in the above expression then simplifies to

$$\begin{aligned} & \frac{1}{2\pi} \sum_{k,d,s \in \mathbb{Z}} \Phi_k^i f_{d,s} \int_{-\pi}^{\pi} e^{i\phi_o(s+d-l)} \int_0^{2\pi} e^{i\phi_i(k+s-d)} d\phi_i d\phi_o \\ &= \sum_{k,s \in \mathbb{Z}} \Phi_k^i f_{k+s,s} \int_{-\pi}^{\pi} e^{i\phi_o(2s+k-l)} d\phi_o = 2\pi \sum_{s \in \mathbb{Z}} \Phi_{l-2s}^i f_{l-s,s}. \end{aligned} \quad (8)$$

and the final set of equations thus reads

$$\Phi_l^o(\mu_o) = 2\pi \int_{-1}^1 \sum_{s \in \mathbb{Z}} \Phi_{l-2s}^i f_{l-s,s}(\mu_i) |\mu_i| d\mu_i. \quad (l \in \mathbb{Z}) \quad (9)$$

These equations reveal a surprising structure that was absent in the isotropic case: the l -th Fourier mode of the scattered radiance is now coupled to modes $l \pm 0, 2, 4, 6, \dots$ of the incident radiance. At first, this appears to be a fundamental problem: the independence of each Fourier mode and resulting dimensionality reduction were both crucial factors in making the isotropic layering model of Jakob et al. computationally feasible.

Fortunately, only a small number of adjacent modes must be considered in practice—the frequency-space coupling is *weak*. This is a consequence of the rapid decay of Fourier coefficients in the azimuth sum parameter ϕ_s observed in Figure 2. Figure 3 shows the visual effect of using an increasing numbers of coupling coefficients $s \in [-m_s, \dots, m_s]$ in a rendering of an anisotropic conductive surface. A similar parameter $d \in [-m_d, \dots, m_d]$ determines the number of coefficients used to model the dependence on ϕ_d .

Finally, we apply the quadrature rule (1) to Equation 9 and rewrite the resulting expression using matrix-vector products:

$$\Phi_l^o = 2\pi \sum_{s=-m_s}^{m_s} \mathbf{F}_{l-s,s} \mathbf{M} \Phi_{l-2s}^i. \quad (l = -m_d, \dots, m_d) \quad (10)$$

Here, Φ_l^o and Φ_l^i are n -dimensional vectors tabulating the incident and outgoing radiance over elevation cosines for a given Fourier mode, $\mathbf{F}_{d,s}$ is an $n \times n$ matrix containing the (d,s) Fourier coefficients of f , and $\mathbf{M} = \delta_{ij} w_i |\mu_i|$ is a diagonal matrix containing quadrature weights and cosine foreshortening factors.

The equations in (10) specify the scattered radiance for a single frequency due to incident radiance along neighboring frequencies. We now group all of these equations into a single matrix \mathbf{S} that




Fig. 3. **Left to right.** The effect of increasing the number of coefficients m_s used to represent the coupling between Fourier modes due to anisotropy, visualized using a conductive layer. **Top row.** Moderate anisotropy, $\alpha = (0.1, 0.2)$, **Bottom row.** Pronounced anisotropy, $\alpha = (0.05, 0.4)$. The last column matches the corresponding analytic references.

left and right, which yields

$$\left(\mathbf{T}_1^{bt}\right)^{-1}\left(\tilde{\mathbf{R}}^t-\mathbf{R}_1^t\right)\left(\mathbf{T}_1^{tb}\right)^{-1}=\left(\mathbf{I}-\mathbf{R}_2^t\mathbf{R}_1^b\right)^{-1}\mathbf{R}_2^t \quad (13)$$

Let \mathbf{X} denote the left hand side of Equation 13, i.e.

$$\mathbf{X} := \left(\mathbf{T}_1^{bt}\right)^{-1} \left(\tilde{\mathbf{R}}^t - \mathbf{R}_1^t\right) \left(\mathbf{T}_1^{tb}\right)^{-1}.$$

Multiplying (13) by $(\mathbf{I} - \mathbf{R}_2^t \mathbf{R}_1^b)$ and expanding produces

$$\mathbf{X} - \mathbf{R}_2^t \mathbf{R}_1^b \mathbf{X} = \mathbf{R}_2^t,$$

which has the solution

$$\mathbf{R}_2^t = \mathbf{X} \left(\mathbf{I} + \mathbf{R}_1^b \mathbf{X} \right)^{-1}. \quad (14)$$

The remaining removing equations read

$$\begin{aligned} \mathbf{T}_2^{tb} &= \tilde{\mathbf{T}}^{tb} \left(\mathbf{T}_1^{tb} \right)^{-1} \left(\mathbf{I} - \mathbf{R}_1^b \mathbf{R}_2^t \right) \\ \mathbf{T}_2^{bt} &= \left(\mathbf{I} - \mathbf{R}_2^t \mathbf{R}_1^b \right) \left(\mathbf{T}_1^{bt} \right)^{-1} \tilde{\mathbf{T}}^{bt} \\ \mathbf{R}_2^b &= \tilde{\mathbf{R}}^b - \mathbf{T}_2^{tb} \left(\mathbf{I} - \mathbf{R}_1^b \mathbf{R}_2^t \right)^{-1} \mathbf{R}_1^b \mathbf{T}_2^{bt} \end{aligned} \quad (15)$$

An alternative set of equations removes the bottom layer instead:

$$\begin{aligned} \mathbf{Y} &:= \left(\mathbf{T}_2^{tb}\right)^{-1} \left(\tilde{\mathbf{R}}^b - \mathbf{R}_2^b\right) \left(\mathbf{T}_2^{bt}\right)^{-1} \\ \mathbf{R}_1^b &= \mathbf{Y} \left(\mathbf{I} + \mathbf{R}_2^t \mathbf{Y}\right)^{-1} \\ \mathbf{T}_1^{tb} &= \left(\mathbf{I} - \mathbf{R}_1^b \mathbf{R}_2^t\right) \left(\mathbf{T}_2^{tb}\right)^{-1} \tilde{\mathbf{T}}^{tb} \\ \mathbf{T}_1^{bt} &= \tilde{\mathbf{T}}^{bt} \left(\mathbf{T}_2^{bt}\right)^{-1} \left(\mathbf{I} - \mathbf{R}_2^t \mathbf{R}_1^b\right) \\ \mathbf{R}_1^t &= \tilde{\mathbf{R}}^t - \mathbf{T}_1^{bt} \left(\mathbf{I} - \mathbf{R}_2^t \mathbf{R}_1^b\right)^{-1} \mathbf{R}_2^t \mathbf{T}_1^{tb} \end{aligned} \quad (16)$$

Similar to other problems that resemble deconvolution, the inverse terms have high condition numbers and exhibit oscillatory behavior that requires an appropriate regularization strategy to avoid magnifying small measurement errors. Figure 4 visualizes an inverse transmission matrix $(\mathbf{T}_1^{tb})^{-1}$ computed for a dielectric microfacet model. Note the large-valued positive and negative bands centered around the specular peak.

4 PROJECTING ANALYTIC AND MEASURED BSDFS

We now discuss the conversion of existing BRDF models into our frequency-space representation, which entails finding 2D Fourier expansion coefficients $f_{d,s}$ of the form

$$f_{d,s}(\mu_i, \mu_o) = \mathcal{F}_{d,s}[f(\mu_i, \phi_i(\phi_d, \phi_s), \mu_o, \phi_o(\phi_d, \phi_s))] \quad (17)$$

$$= \frac{1}{4\pi^2} \int_{-\pi}^{\pi} \int_{-\pi}^{\pi} f(\mu_i, \frac{1}{2}(\phi_s - \phi_d), \mu_o, \frac{1}{2}(\phi_s + \phi_d)) e^{-id\phi_d - is\phi_s} d\phi_s d\phi_d.$$

Our system currently supports reflective and transmissive BSDF measurements as well as dielectric and conductive microfacet BSDFs using the anisotropic Beckmann normal distribution function.²




Fig. 4. Visualization of an inverse transmission scattering matrix $(T_1^{tb})^{-1}$ that occurs as part of the subtracting equation (14). These terms are highly oscillatory and require an appropriate regularization strategy to avoid magnifying small measurement errors. (Rough dielectric with isotropic Beckmann roughness $\alpha = 0.2$.)

4.1 Anisotropic Microfacet BSDFs

Consider the reflection component of a standard microfacet model

$$f(\mu_i, \phi_i, \mu_o, \phi_o) = \frac{F(\mu_h) G_1(\mu_i, \phi_i) G_1(\mu_o, \phi_o) D(\mu_h, \phi_h)}{4|\mu_i \mu_o|}, \quad (18)$$

where F is the Fresnel term, D is normal distribution function, G_1 represents Smith's separable shadowing and masking terms, and μ_h, ϕ_h denote elevation cosine and azimuth of the half-direction vector between incident and outgoing direction. Using the (μ_h, ϕ_h) parameterization, the anisotropic Beckmann distribution D is defined as

$$D(\mu_h, \phi_h) = \frac{1}{\pi \alpha_u \alpha_v \mu_h^4} e^{-\frac{1-\mu_h^2}{\mu_h^2} \left(\frac{\cos^2 \phi_h}{\alpha_u^2} + \frac{\sin^2 \phi_h}{\alpha_v^2} \right)}. \quad (19)$$

We use a combination of analytic and numeric integration techniques to compute the needed Fourier expansion coefficients $f_{d,s}$. Specifically, we split the BSDF into the product

$$f_D(\mu_i, \mu_o, \phi_d, \phi_s) \cdot f_F(\mu_i, \mu_o, \phi_d) \cdot f_{G_1}(\mu_i, \phi_i) \cdot f_{G_1}(\mu_o, \phi_o), \quad (20)$$

which separates the NDF, Fresnel, and separable shadowing-masking terms. As the main source of high-frequency variation, the micro-facet NDF f_D requires special precautions and is projected using a semi-analytic approach. Note that is also the only factor in need of a full 2D Fourier expansion, as the Fresnel and shadowing-masking factors only vary with respect to a single azimuth parameter. Similar to Jakob et al. [2014], we then exploit the convolution theorem to convert the multiplications in Equation 20 into convolutions of complex Fourier coefficients defined as

$$(g * h)_i = \sum_{k=-\infty}^{\infty} g_k \cdot h_{i-k}. \quad (21)$$

In our case, the convolutions have the following new structure on the two-dimensional domain (ϕ_s, ϕ_d) :

$$\mathcal{F}[f] \otimes \mathcal{F}[D] = \mathcal{F}[F] \otimes \mathcal{F}[G_1(\omega_i)] \otimes \mathcal{F}[G_1(\omega_o)]$$

²In principle, other microfacet distributions (such as GGX) are also supported by treating them like measured data.

Assuming that the individual Fourier series are available, these convolutions have negligible cost: all terms except for the first are short and only involve at most dozens of coefficients. Secondly, only one-dimensional convolutions are required if the evaluation is carried out in a suitable order. We now focus on each term in turn:

4.1.1 Exponential Component. Ignoring constant normalization factors, the microfacet distribution term f_D can be written as

$$f_D(\phi_s, \phi_d) = e^{c_0 + c_1 \cos \phi_d + c_2 \cos \phi_s + c_3 \cos(\phi_d + \phi_s) + c_4 \cos(\phi_d - \phi_s)} \quad (22)$$

for constants c_0, c_1, c_2, c_3, c_4 that only depend on variables not related to azimuths, specifically $\mu_i, \mu_o, \alpha_u, \alpha_v$ and the relative index of refraction η :

$$\begin{aligned} c_0 &= \frac{1}{2} \gamma \left(\alpha_u^2 + \alpha_v^2 \right) \left(\mu_i^2 - 1 + \eta^2 (\mu_o^2 - 1) \right) \\ c_1 &= \gamma \left(\alpha_u^2 + \alpha_v^2 \right) \eta \sqrt{1 - \mu_i^2} \sqrt{1 - \mu_o^2} \\ c_2 &= \gamma \left(\alpha_v^2 - \alpha_u^2 \right) \eta \sqrt{1 - \mu_i^2} \sqrt{1 - \mu_o^2} \\ c_3 &= \frac{1}{2} \gamma \left(\alpha_v^2 - \alpha_u^2 \right) \eta^2 \left(\mu_o^2 - 1 \right) \\ c_4 &= \frac{1}{2} \gamma \left(\alpha_v^2 - \alpha_u^2 \right) \left(\mu_i^2 - 1 \right) \end{aligned}$$


where $\gamma = \frac{1}{\alpha_u^2 \alpha_v^2 (\mu_i - \eta \mu_o)^2}$ and $\eta = \frac{\eta_o}{\eta_i}$.

They are valid for both refraction and reflection, in which case $\eta = 1$. For isotropic materials ($\alpha_u = \alpha_v$), only c_0 and c_1 take on non-zero values and are equivalent to the constants given by Jakob et al. [2014].

Unfortunately, we are not aware of an analytic mechanism to directly compute the Fourier series of Equation 22. However, note that if the exponential argument only depended on a single azimuth argument, a variant of the Jacobi-Anger expansion (Abramowitz and Stegun [1964], 9.6.34) yields

$$e^{A + B \cos(\phi + C)} = e^A \left[\sum_{k \in \mathbb{Z}} I_k(B) e^{i k (\phi + C)} \right] \quad (23)$$

where $I_k(z)$ are modified Bessel functions of the first kind. Our approach, then, is as follows: assuming that the azimuth sum parameter ϕ_s is a constant, we can use phasor addition



to bring Equation 22 into the above form:

$$f_D(\phi_d) = e^{A + B \cos(\phi_d + C)}, \text{ where} \quad (24)$$

$$A = c_0 + c_2 \cos \phi_s,$$

$$B = |z|,$$

$$C = \arg(z),$$

$$\text{and } z = c_1 + (c_3 + c_4) \cos \phi_s + i(c_3 - c_4) \sin \phi_s.$$

In other words, for fixed ϕ_s , it is straightforward to compute the Fourier series over ϕ_d . The angle sum parameter ϕ_s , in turn, is easy to handle as it has significantly lower frequency content. We therefore use Filon [1928] quadrature to determine the Fourier coefficients along this second dimension, using 30 evenly spaced samples on the interval $\phi_s \in [-\pi, \pi]$. The main differences compared to Jakob et al. [2014] are the higher dimension of the transform and addition of the phase shift C , which captures the bilateral asymmetry due to anisotropy. The remaining three projections of the Fresnel and geometric terms rely on pure quadrature (i.e. without nested analytic integration).

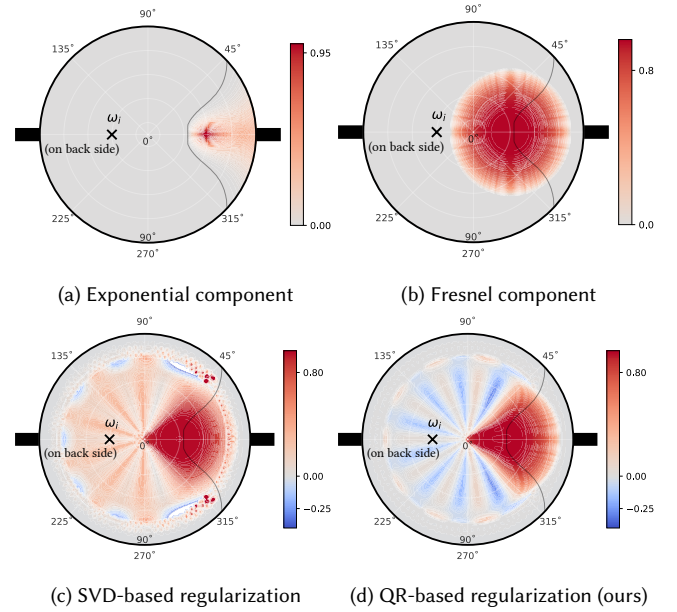


Fig. 5. Transmission plots (T^{bt}) of an isotropic dielectric ($\eta = 1.5, \alpha = 0.2$). **(a)** Exponential component D . The thin outline indicates the region of interest, where the BSDF takes on non-negligible values. **(b)** Analytic reference of the Fresnel component F . **(c-d)**: Localized Fourier expansions using the two different regularization strategies.

4.1.2 Fresnel Component. The Fresnel term F is not affected by anisotropy and only depends on the azimuth difference ϕ_d . It is normally very well-behaved, but a problematic case arises when passing into a material of lower density—here, F contains high frequencies near the critical angle, where the reflectance rapidly jumps to 1. Although the peaked exponential term of the microfacet model typically masks this behavior, any high frequencies in F will cause global ringing in the Fourier series that reduce accuracy

elsewhere. Jakob et al. [2014] address this problem via a custom Fourier projection that optimizes L_2 error on a subinterval $\phi_d \in [\phi_{\min}, \phi_{\max}]$, where the BRDF actually reflects appreciable amounts of light. This involves a restricted inner product

$$\langle g, h \rangle_\Omega = \int_\Omega g(\phi) h(\phi) d\phi$$

in conjunction with a change of basis of the projections $\langle q_i, F \rangle$

$$\begin{bmatrix} F_0 \\ \vdots \\ F_d \end{bmatrix} = \underbrace{\begin{bmatrix} \langle q_0, q_0 \rangle_\Omega & \dots & \langle q_0, q_d \rangle_\Omega \\ \vdots & & \vdots \\ \langle q_d, q_0 \rangle_\Omega & \dots & \langle q_d, q_d \rangle_\Omega \end{bmatrix}^{-1}}_{=: \mathbf{M}(\Omega)} \begin{bmatrix} \langle q_0, F \rangle_\Omega \\ \vdots \\ \langle q_d, F \rangle_\Omega \end{bmatrix}$$

where $q_l(\phi) = \cos(l\phi)$ and $\Omega = [\phi_{\min}, \phi_{\max}]$. The authors solve this linear system using a truncated pseudo-inverse, as it is severely ill-conditioned when $\phi_{\min} \approx \phi_{\max}$.

We found that the necessary SVD factorizations are a significant bottleneck, requiring up to 90% of the total time needed to project a microfacet model. We therefore use the following scheme with improved behavior: we extend the linear system with two soft constraints $\mathbf{M}([-\pi, \phi_{\min}])\mathbf{F} = 0$ and $\mathbf{M}([\phi_{\max}, \pi])\mathbf{F} = 0$ weighted by $\varepsilon = 0.00001$ —effectively a type of Tikhonov regularization. Their objective is to smoothly force the Fourier series to zero outside of the region of interest. This turns the linear system into an over-constrained system that can be solved using a much simpler QR factorization, while yielding functions with overall smoother behavior (Figure 5).

4.1.3 Shadowing-Masking component. We use Smith’s separable shadowing-masking function which allows us to build two 1D Fourier series along ϕ_i and ϕ_o (again using Filon quadrature). We then transform them into the common (ϕ_s, ϕ_d) space and perform a diagonal convolution to combine them with the 2D coefficient array of the remaining Microfacet BSDF.

4.2 Choice of discretization parameters

When converting anisotropic Microfacet BSDFs into the Fourier representation we must choose suitable discretization parameters n , m_s , m_d . We reuse the approach of Jakob et al. [2014] to set n and m_d based on the smaller roughness value $\min\{\alpha_u, \alpha_v\}$. To select a suitable anisotropy-related parameter m_s , we instantiated many layers of different sizes, in each case searching for the minimal value of m_s that achieves a relative L_2 error below 1%. We found that the required number of coefficients is related to the *roughness ratio* $\bar{\alpha} = \max(\frac{\alpha_u}{\alpha_v}, \frac{\alpha_v}{\alpha_u})$, i.e. the amount of anisotropy. The dependence of m_s on $\bar{\alpha}$ is closely approximated by the following linear fit:

$$m_s(\bar{\alpha}) = \lceil ax + b \rceil, \text{ where } a = 2.58 \text{ and } b = -1.82. \quad (25)$$

4.3 Measured BSDFs

Our system also supports measured BSDF data, which entails numerical evaluation of the Fourier projections (17) using an interpolant of the discrete set of BSDF measurements. We assume that the measurement is dense in the (μ_o, ϕ_o) parameters for a smaller set of incident light directions (μ_i, ϕ_i) . To enable evaluation of the BSDF for arbitrary incident light directions, we transform each (μ_i, ϕ_i) -slice into

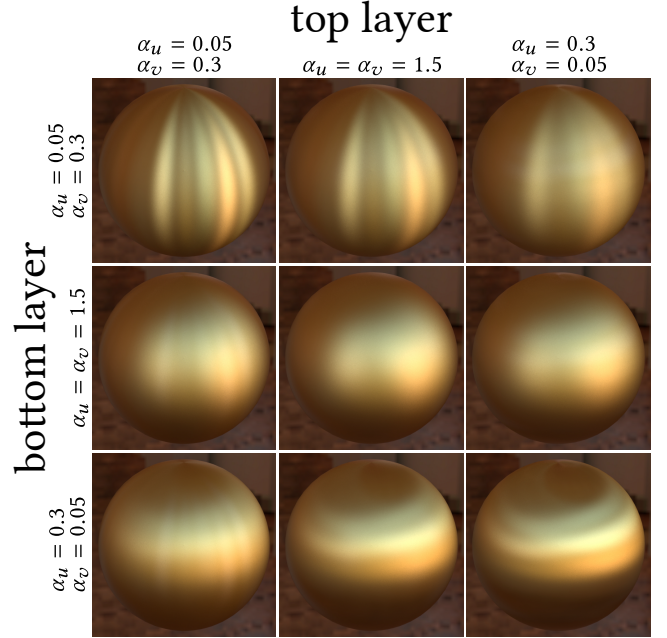


Fig. 6. Different combinations of a conductor base layer (gold) with dielectric coating ($\eta = 1.5$) using varying values of anisotropy along two axes.

the half-vector domain, where we construct a linear interpolant. Finally, for each pair of elevation angles (μ_i, μ_o) , we densely evaluate the interpolant on the (ϕ_s, ϕ_d) domain and use FFTW [Frigo and Johnson 2005] to transform the values into the Fourier domain.

5 RESULTS

We implemented our method in a modified version of the Mitsuba renderer [2010], using the Eigen [2010] library for most sparse matrix manipulation and UMFPACK [Davis 2004] for sparse LU factorizations. The generated reflectance models admit a perfect importance sampling scheme that involves mapping uniform variates to the integrated Fourier series. This portion closely resembles the sampling scheme of Jakob et al. [2014] and just needs to be adapted to an exponential Fourier series—we provide additional detail in the supplemental material. For the special case of isotropic materials, all our results are in perfect agreement with prior work.

The remainder of this section illustrates the visual impact of anisotropic layer interfaces and shows validations of both additive and subtractive layer composition using ground-truth Monte Carlo simulations and measured reflectance data.

5.1 Kitchen scene

Figure 1 contains a wide variety of anisotropic surfaces including coated anisotropic metals, combinations of measured transmissive and analytic conductive layers, and wall tiles made of anisotropic dielectrics with internal scattering that simulate the anisotropy of machined plastic. All surface except for the base of the table are anisotropic, and layering computations for all 20 materials took a total of 20min 3s on a 4-core i7-6700K CPU. The scene involves

Table 2. Layer adding statistics for typical configurations of conductive bottom layer with dielectric coating ($\eta = 1.5$).

Roughness (top)		Roughness (bottom)		Time		Sparsity			Storage size
α_u	α_v	α_u	α_v	total	adding only	top	bottom	combined	
0.1	0.1	0.05	0.05	2.6s	0.6s	0.011%	0.019%	0.022%	2.8 MiB
0.1	0.1	0.1	0.1	2.6s	0.6s	0.011%	0.017%	0.02%	2.6 MiB
0.1	0.1	0.3	0.3	2.2s	0.3s	0.011%	0.008%	0.018%	2.3 MiB
0.1	0.1	0.05	0.15	26.8s	16.9s	0.011%	0.489%	0.208%	31.1 MiB
0.1	0.1	0.1	0.3	24.2s	16.9s	0.011%	0.336%	0.144%	27.3 MiB
0.15	0.05	0.05	0.15	15min	13min	0.059%	0.087%	0.01%	214.1 MiB
0.3	0.1	0.1	0.3	67.5s	52.2s	0.29%	0.336%	0.39%	31.3 MiB




Fig. 7. Layered material examples with an anisotropic gold conductor. (a) Base layer only ($\alpha_u = 0.05$, $\alpha_v = 0.3$), (b) covered in a layer of forward scattering HG medium (albedo = 0.95, $g = 0.5$), (c) coated with an anisotropic dielectric ($\eta = 1.5$, $\alpha_u = 0.3$, $\alpha_v = 0.05$), (d) covered with blue scattering HG medium (albedo = [0.6, 0.8, 0.95], $g = 0.2$) and isotropic dielectric coating ($\eta = 1.5$, $\alpha_u = \alpha_v = 0.05$)

complex glossy-to-glossy transport in conjunction with directionally peaked spot light sources, which required a large number of samples per pixel (40K) and a correspondingly long rendering time (80 hours on a 12-core Xeon E5-2680 machine) to reach convergence. The challenges of rendering glossy-to-glossy transport are largely orthogonal to our work and could be alleviated by using more advanced integration techniques (we used simple path tracing with multiple importance sampling).

5.2 Layering anisotropic materials

Figure 6 illustrates the space of anisotropic surface appearance that can be realized by composing a conductive layer with a rough dielectric coating. Both layers are either isotropic or have pronounced anisotropy along one of two tangential axes. Changing the bottom layer generally has a greater effect on the final BSDF, though the top layer also plays a noticeable role due to the addition of highlights and directional blur.

In addition to the anisotropic interfaces, our framework remains compatible with all ingredient layers of prior work [Jakob et al.




Fig. 8. Interactive layer visualization. (a) *Dragon* scene, (b) *Shapes* scene, (c) user interface to change layer roughness parameters.

2014]. This enables the composition of anisotropic materials with layers containing absorbing or scattering media using the Henyey-Greenstein phase function—several examples are shown in Figure 7.

5.3 Interactive viewer

An image viewer provided in the supplemental material enables interactive exploration of a 4D parameter space of anisotropic surface reflectance (Figure 8). It contains two different scenes—the Stanford dragon and a set of cylinders with textured anisotropy—both rendered using a coated gold layer. The roughness parameters (α_u , α_v) of both layers are individually adjustable using an HTML widget, and the layers can also be removed to see the base or coating in isolation.

5.4 Performance and storage cost

While our implementation retains the efficiency of Jakob et al. [2014] when dealing with isotropic layers, the presence of anisotropy and resulting coupling between Fourier modes adds overhead to the system during additive and subtractive composition. Table 2 lists a few key statistics to compare the efficiency of adding typical material configurations of isotropic and anisotropic conductors and dielectrics. Note how the precomputation time is still well under one minute for cases where only one of the layers contains anisotropic scattering.

5.5 Validation

We validate our layering operators using both synthetic and measured BRDF data. For each experiment, we show hemispherical BRDF plots (projected onto the unit disk) and renderings of the Stanford dragon mesh with local tangent vectors provided by a 2-RoS field computed using the Instant Meshes [Jakob et al. 2015]

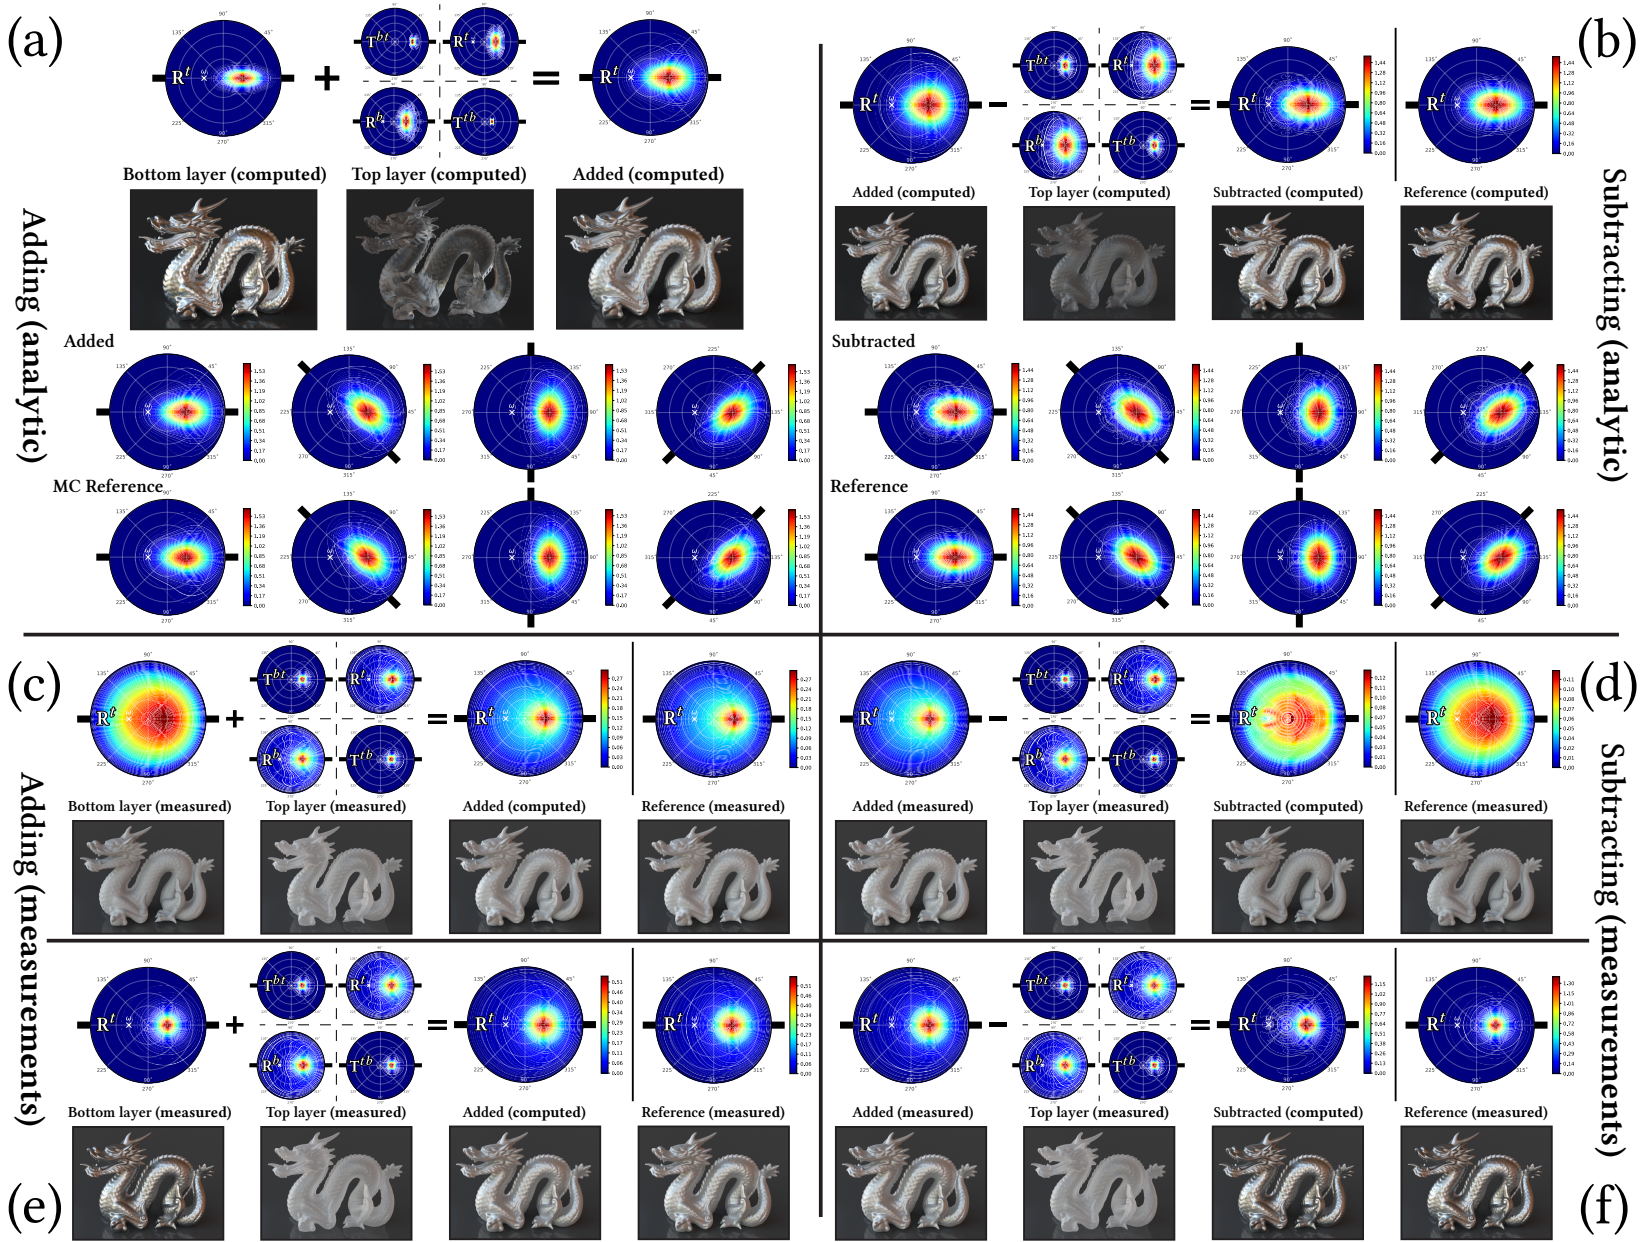


Fig. 9. Additive and subtractive layer composition

algorithm. Due to space constraints, we only show BRDF plots for a single incident light direction ($\theta_i = 30^\circ$). The supplemental material contains expanded versions of these visualizations with many additional configurations.

Additive composition of synthetic data. Figure 9a visualizes additive composition of two anisotropic microfacet BSDF with perpendicular anisotropy (Bottom: conductor $\alpha = (0.2, 0.1)$, top: dielectric $\eta = 1.5$, $\alpha = (0.1, 0.2)$). The addition of the top layer causes directional blur that broadens the reflection lobe along the narrow axis, though the transformation is more complex than a simple spherical convolution. Multiple scattering involving both layers adds a faint diffusive component surrounding the peak.

A sequence of spherical plots in the lower half of the sub-figure compares evaluations of our model to a brute force Monte Carlo simulation that propagates particles through a layer stack, accumulating particles exiting the stack into a spherical histogram. The plots show excellent agreement between the two approaches.

Subtractive composition of synthetic data. Figure 9b demonstrates that composition also works in the opposite direction: we remove a two-sided anisotropic dielectric ($\eta = 1.5$, $\alpha = (0.2, 0.3)$) from a two-layer structure involving an anisotropic conductive base layer with $\alpha = (0.3, 0.2)$.

Deconvolution-type problems require regularization to improve resilience to noise, and the subtracting equations are no exception. We thus add a small amount of uniform Tikhonov regularization ($\epsilon = 5 \cdot 10^{-4}$) when inverting T_1^{bt} and T_1^{tb} in the subtracting equations. When ϵ is set to too small a value, the solution becomes oscillatory and is unusable—we therefore always choose the smallest possible ϵ that still leads to a solution without oscillatory behavior.

Additive composition of measured data. We used a pgII [PAB 2018] gonio-photometer (Figure 10a) to measure several layers—first individually, and then in a stacked configuration. The gonio-photometer illuminates a flat A4-sized sample using a collimated beam generated by an arbitrary static light source. We used a Thorlabs HPLS 345 light emitting plasma lamp with an added infrared filter to generate broadband illumination in the visible range. The samples are mounted in a motorized sample holder, whose rotation axes control the elevation and azimuth of incident illumination. A motorized two-axis sensor head with a monochromatic silicon photodiode (First Sensor PS100-6b) records the scattered flux from arbitrary positions on the sphere. All measurements combine a localized scan of the reflection and transmission peaks along with a broader scanning pattern that captures tail behavior (Figure 10b). Note that we restricted ourselves to isotropic materials in the measurement-based validation experiments due to the significant challenge of acquiring dense 4D reflectance data.

The experiments in the bottom half of Figures 9 all use a polypropylene sheet with two rough interfaces (Figure 10c) as the top layer. This material strongly blurs transmitted light, particularly at low angles of incidence (Figure 10d). The surface finish on the top and bottom surfaces is slightly different, hence it was necessary to measure the material from both sides.

The first two columns of Figures 9c and 9e visualize measurements of the transmissive PP layer and two different types of base




Fig. 10. (a) We use a pgII gonio-photometer to densely measure the reflective and transmissive properties of several layered material structures. (b) Visualization of a measured transmission lobe with lines highlighting locations where the BSDF was sampled. After transforming this data into our frequency-space representation, it can be used for both additive and subtractive composition. (c), (d) For validation, we measured a diffusive polypropylene sheet, whose rough top and bottom surfaces cause a varying amount of directional blur that depends on the angle of incidence. We additionally measure layer combinations, with the PP layer on top of a matte cardboard and a glossy metallic paper sheet. The bottom half of Figure 9 contains several validation experiments using the resulting data.

layers: a matte cardboard with a broad reflection and a paper with a metallic coating that exhibits a strong specular reflection. The third column contains renderings of the combined layer structures computed using the adding equations, and the last column shows a ground-truth reference created by layering the sheets on top of each other and measuring them with the gonio-photometer. We also show a quantitative analysis of our computed BSDF against the ground-truth measurement in Figure 11 including pixel-wise *CIE dE00* error [Sharma et al. 2005] of the rendered results. The images are almost indistinguishable.

Subtractive composition of measured data. This last experiment is the most challenging: we now reverse the steps of the previous computation. Using the measured layer combinations, we subtract the top layer computationally and compare against the ground-truth measurement of the bottom layer in isolation, using Tikhonov regularization constants of $\epsilon = 0.06$ and $\epsilon = 0.09$ for the metallic paper and cardboard, respectively. Figures 9d,f visualize the result of the subtractive composition: the rendering reveals small differences at grazing angles, which are challenging to infer due to the properties of the PP sheet (Figure 10d). This is supported by the analysis in Figure 12: we observe large RMSE when approaching grazing angles especially for the metallic paper. Note however that the rendered results closely match with a perceptual error concentrated only around small regions. Overall, we find that the reconstructions computed using the subtracting equations are in surprisingly good

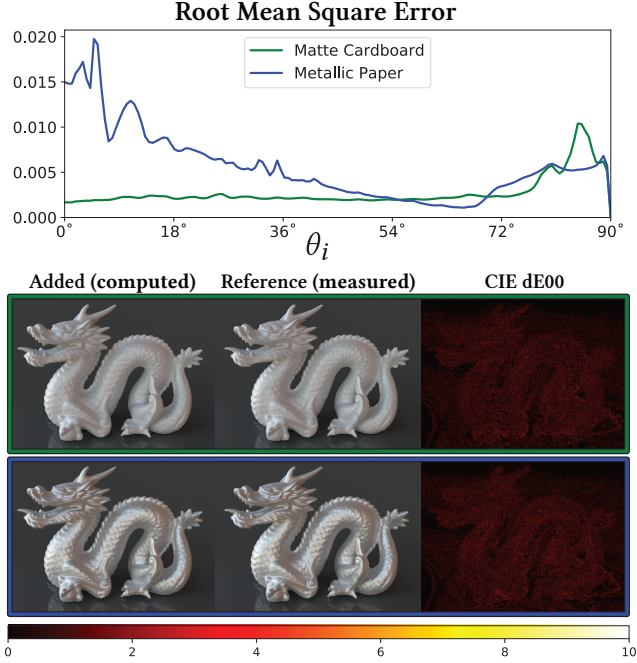


Fig. 11. Error analysis for *additive* composition of measured data. **Top:** Root-mean-square error of the computed BSGF against the measured reference over incident elevation. **Bottom:** Color mapped visualization of the perceptual CIE dE00 error of the rendered results.

agreement to the reference data given the challenging nature of this problem.

6 CONCLUSIONS AND FUTURE WORK

We have presented a new framework for computing the reflective and transmissive properties of anisotropic material structures. Being able to reason about anisotropy in the context of material modeling is crucial, as almost any machined surface is affected by it to some degree. Our technique supports arbitrary composition of real-world measurements and traditional anisotropic microfacet BSDFs and is a major generalization of prior work that focused on the special case of isotropic materials. Using our directional scattering operators, we furthermore introduce subtractive composition, which reconstructs the BSDF of a material that is only indirectly observed through another material with known properties. Modeling materials using layers produces a rich and high-dimensional space of useful appearance models. The direct connection between appearance and internal composition also leads to a useful set of physical parameters that could be used for inverse problems in the context of material design. The large set of possibilities motivates the name of our technique: the flexible operator representation coupled with a number of computational tools serves as a laboratory for a wide variety of future experiments involving layers.

Our method is not without limitations: it currently does not scale to extreme anisotropy, which occurs e.g. in some brushed aluminium samples with RMS roughness ratios exceeding 1:1000. Furthermore, the internal material parameters cannot be controlled by a spatially

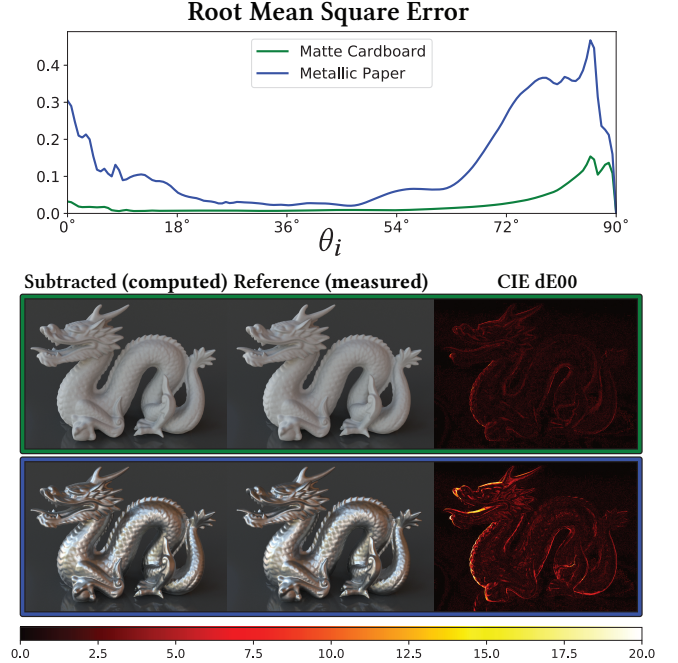


Fig. 12. Error analysis for *subtractive* composition of measured data. **Top:** Root-mean-square error of the computed BSGF against the measured reference over incident elevation. **Bottom:** Color mapped visualization of the perceptual CIE dE00 error of the rendered results.

varying texture, as each set of parameters requires a new layering computation. This means that the model is likely not applicable to visual effects applications in its current form. That said, we believe that our anisotropic layering model could serve as a useful ground-truth in the development of more approximate techniques that do support texturing in the future. The applications of subtractive composition in the context of reverse engineering and cultural heritage preservations are exciting, but the currently used Tikhonov regularization is rather simplistic. We believe that more advanced regularizers that measure distances on the manifold of natural BRDFs [Nielsen et al. 2015; Soler et al. 2018] could considerably boost performance, and we look forward to exploring this topic in the future.

ACKNOWLEDGMENTS

We thank Olesya Jakob for designing the kitchen scene shown in Figure 1, Merlin Nimier-David for building the interactive viewer, and Peter Apian-Bennwitz for support with the gonio-photometric measurements.

This work was supported by EPFL through the use of the facilities of its Scientific IT and Application Support Center.

REFERENCES

- Milton Abramowitz and Irene A. Stegun. 1964. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover, New York.
- Petr Beckmann and Andre Spizzichino. 1963. *The scattering of electromagnetic waves from rough surfaces*. MacMillan.
- Laurent Belcour and Pascal Barla. 2017. A Practical Extension to Microfacet Theory for the Modeling of Varying Iridescence. *ACM Trans. Graph.* 36, 4, Article 65 (July).

- 2017).
- James F. Blinn. 1977. Models of Light Reflection for Computer Synthesized Pictures. *SIGGRAPH Comput. Graph.* 11, 2 (July 1977), 192–198.
- James F. Blinn. 1982. Light Reflection Functions for Simulation of Clouds and Dusty Surfaces. In *Proceedings of the 9th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '82)*. ACM, New York, NY, USA, 21–29.
- Brent Burley. 2012. Physically-based shading at Disney. In *ACM SIGGRAPH 2012 Courses (SIGGRAPH '12)*. ACM, New York, NY, USA.
- Subrahmanyan Chandrasekhar. 1960. *Radiative Transfer*. Dover Publications.
- Robert L. Cook and Kenneth E. Torrance. 1982. A Reflectance Model for Computer Graphics. *ACM Trans. Graph.* 1, 1 (Jan. 1982), 7–24.
- Qiang Dai, Jiaping Wang, Yiming Liu, John Snyder, Enhua Wu, and Baining Guo. 2009. The Dual-microfacet Model for Capturing Thin Transparent Slabs. *Computer Graphics Forum* 28 (10 2009), 1917–1925.
- Timothy A. Davis. 2004. Algorithm 832: UMFPAK V4. 3—an unsymmetric-pattern multifrontal method. *ACM Transactions on Mathematical Software (TOMS)* 30, 2 (2004), 196–199.
- Sergey Ershov, Konstantin Kolchin, and Karol Myszkowski. 2001. Rendering Pearlescent Appearance Based On Paint-Composition Modelling. *Computer Graphics Forum* (2001).
- Louis N.G. Filon. 1928. On a quadrature formula for trigonometric integrals. In *Proc. Roy. Soc. Edinburgh*, Vol. 49, 38–47.
- Matteo Frigo and Steven G. Johnson. 2005. The design and implementation of FFTW3. In *Proceedings of the IEEE*. 216–231.
- Ian P. Grant and Garry. E. Hunt. 1969. Discrete Space Theory of Radiative Transfer. I. Fundamentals. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 313, 1513 (1969).
- Gaël Guennebaud, Benoît Jacob, et al. 2010. Eigen v3. <http://eigen.tuxfamily.org>. (2010).
- Pat Hanrahan and Wolfgang Krueger. 1993. Reflection from Layered Surfaces Due to Subsurface Scattering. In *Proceedings of the 20th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '93)*. ACM, New York, NY, USA, 165–174.
- Eric Heitz, Johannes Hanika, Eugene d'Eon, and Carsten Dachsbacher. 2016. Multiple-scattering Microfacet BSDFs with the Smith Model. *ACM Trans. Graph.* 35, 4, Article 58 (July 2016).
- Hideki Hirayama, Kazufumi Kaneda, Hideo Yamashita, and Yoshimi Monden. 2001. An accurate illumination model for objects coated with multilayer films. *Computers & Graphics* 25, 3 (2001), 391–400.
- Isabelle Icart and Didier Arquès. 2000. A Physically-Based BRDF Model for Multilayer Systems with Uncorrelated Rough Boundaries. In *Proceedings of the Eurographics Workshop on Rendering Techniques 2000*. Springer-Verlag, London, UK, UK, 353–364.
- Wenzel Jakob. 2010. Mitsuba renderer. (2010). <http://www.mitsuba-renderer.org>.
- Wenzel Jakob, Eugene d'Eon, Otto Jakob, and Steve Marschner. 2014. A Comprehensive Framework for Rendering Layered Materials. *ACM Trans. Graph.* 33, 4, Article 118 (July 2014).
- Wenzel Jakob, Marco Tarini, Daniele Panozzo, and Olga Sorkine-Hornung. 2015. Instant Field-aligned Meshes. *ACM Trans. Graph.* 34, 6, Article 189 (Oct. 2015).
- Csaba Kelemen and Laszlo Szirmay-Kalos. 2001. A microfacet based coupled specular-matte BRDF model with importance sampling. In *Eurographics Short Present.*, Vol. 25.
- Nikhil Naik, Shuang Zhao, Andreas Velten, Ramesh Raskar, and Kavita Bala. 2011. Single View Reflectance Capture Using Multiplexed Scattering and Time-of-flight Imaging. *ACM Trans. Graph.* 30, 6, Article 171 (Dec. 2011).
- Jannik Boll Nielsen, Henrik Wann Jensen, and Ravi Ramamoorthi. 2015. On Optimal, Minimal BRDF Sampling for Reflectance Acquisition. *ACM Trans. Graph.* 34, 6, Article 186 (Oct. 2015).
- PAB. 2018. pab advanced technologies Ltd. <http://www.pab.eu>. (2018). Accessed: 2018-01-09.
- Bui Tuong Phong. 1975. Illumination for Computer Generated Pictures. *Commun. ACM* 18, 6 (June 1975), 311–317.
- Szymon Rusinkiewicz. 1998. A New Change of Variables for Efficient BRDF Representation. In *Rendering Techniques '98, Proceedings of the Eurographics Workshop in Vienna, Austria, June 29 - July 1, 1998*. 11–22.
- Gaurav Sharma, Wencheng Wu, and Edul Dalal. 2005. The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations. 30 (02 2005), 21 – 30.
- Peter Shirley, Brian Smits, Helen Hu, and Eric Lafortune. 1997. A practitioners' assessment of light reflection models. In *Proceedings The Fifth Pacific Conference on Computer Graphics and Applications*. 40–49.
- Cyril Soler, Kartic Subr, and Derek Nowrouzezahrai. 2018. A Versatile Parameterization for Measured Material Manifolds. In *Eurographics 2018 - 39th Conference*, Vol. 37. Wiley, Delft, Netherlands, 1–10.
- Jos Stam. 2001. An Illumination Model for a Skin Layer Bounded by Rough Surfaces. In *Rendering Techniques 2001: Proceedings of the Eurographics Workshop in London*. Springer Vienna, Vienna, 39–52.
- Shuochen Su, Felix Heide, Robin Swanson, Jonathan Klein, Clara Callenberg, Matthias Hullin, and Wolfgang Heidrich. 2016. Material Classification Using Raw Time-of-Flight Measurements. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 3503–3511.
- Kenichiro Tanaka, Yasuhiro Mukaigawa, Hiroyuki Kubo, Yasuyuki Matsushita, and Yasushi Yagi. 2015. Recovering inner slices of translucent objects by multi-frequency illumination. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 5464–5472.
- Kenneth E. Torrance and Ephraim M. Sparrow. 1967. Theory for Off-Specular Reflection from Roughened Surfaces. *Journal of the Optical Society of America (JOSA)* 57, 9 (Sept. 1967), 1105–1114.
- Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. 2007. Microfacet Models for Refraction Through Rough Surfaces. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques (EGSR'07)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 195–206.
- Gregory J. Ward. 1992. Measuring and Modeling Anisotropic Reflection. *SIGGRAPH Comput. Graph.* 26, 2 (July 1992), 265–272.
- Andrea Weidlich and Alexander Wilkie. 2007. 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)*. ACM, New York, NY, USA, 171–178.