


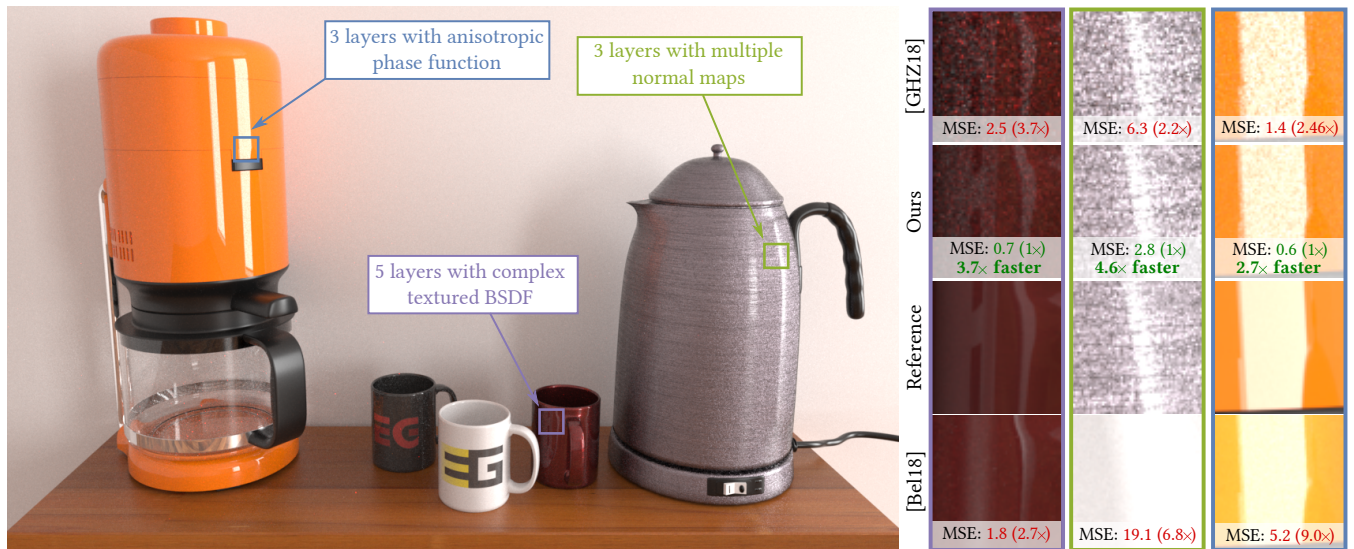


# An Efficient Transport Estimator for Complex Layered Materials

Luis E. Gamboa<sup>1</sup>  Adrien Gruson<sup>2</sup>  Derek Nowrouzezahrai<sup>2</sup> 

<sup>1</sup>Université de Montréal, Canada    <sup>2</sup>McGill University, Canada



**Figure 1:** Left: complex layered materials with high-frequency, spatially-varying reflectance (e.g., albedo, roughness), micro-geometry (i.e., normal maps) and anisotropic participating media. Right: 2x zoom-in insets of equal-time comparisons – Belcour’s model [Bel18] is fast but approximate, whereas our method reduces MSE by about 2–4x compared to the state-of-the-art in offline layered models [GHZ18].

## Abstract

Layered materials capture subtle, realistic reflection behaviors that traditional single-layer models lack. Much of this is due to the complex subsurface light transport at the interfaces of – and in the media between – layers. Rendering with these materials can be costly, since we must simulate these transport effects at every evaluation of the underlying reflectance model. Rendering an image requires thousands of such evaluations, per pixel. Recent work treats this complexity by introducing significant approximations, requiring large precomputed datasets per material, or simplifying the light transport simulations within the materials. Even the most effective of these methods struggle with the complexity induced by high-frequency variation in reflectance parameters and micro-surface normal variation, as well as anisotropic volumetric scattering between the layer interfaces. We present a more efficient, unbiased estimator for light transport in such general, complex layered appearance models. By conducting an analysis of the types of transport paths that contribute most to the aggregate reflectance dynamics, we propose an effective and unbiased path sampling method that reduces variance in the reflectance evaluations. Our method additionally supports reflectance importance sampling, does not rely on any precomputation, and so integrates readily into existing renderers. We consistently outperform the state-of-the-art by  $\sim 2 - 6\times$  in equal-quality (i.e., equal error) comparisons.

## CCS Concepts


• Computing methodologies → Rendering;

## 1. Introduction

Physically-based rendering (PBR) models have led to methods capable of producing images indistinguishable from photographs, given sufficiently accurate geometry, material and lighting data. Recent developments in surface appearance modeling have been fundamental to the growing ubiquity of PBR: spatially-varying albedo,

roughness and normal distributions can produce visually rich and convincing results when combined with faithful simulations of the underlying light transport.

The diversity of real-world appearance necessitates flexible, expressive appearance models. We are interested in the *layered material* appearance, where simpler “single-layer” models fail to prop-



**Figure 2:** Complex layered materials can include many layers, each with spatially-varying surface and volume scattering. While flexible, computing light transport inside these materials is challenging.

erly capture the aggregate appearance that results from subsurface transport within-and-between the layers. Car paint, for instance, comprises both a glossy and matte layer, each with their own (potentially high-frequency) surface reflectance variations. Applying a simple, e.g., affine, combination of the individual profiles fails to capture the complex light transport occurring beneath the surface.

Faithfully simulating these effects is challenging, and can often bottleneck a rendering system: a single point-wise evaluation of a physically-accurate layered material model (i.e., for fixed incident and outgoing directions) should account for the full light transport simulation inside the material. Several accurate methods are limited to non-spatially-varying, isotropic [JdJM14] and anisotropic [ZJ18] interface layers. Recently, a simplified formulation more accurately simulates multi-layer transport which, in turn, allows for arbitrarily complex per-layer models [GHZ18].

We present a light transport model for layered materials that simplifies the throughput computation inside, and across, layers. We found that while full bidirectional path sampling of layered material subpaths [GHZ18] is a sample-efficient strategy, the overhead associated with path construction and multiple importance sampling (MIS) does not justify the per-sample variance reduction. We demonstrate that a carefully constructed unidirectional estimator with a specialized next-event estimation to support an arbitrary number of interfaces both significantly increases performance and simplifies implementation. We also identify and provide justification for estimator choices that significantly increase the performance of our approach without introducing any bias.

Specifically, we present the following technical contributions:

- an analysis of which path sampling strategies contribute the most to microscale light transport in complex layered materials,
- an unbiased, efficient path construction method to sample and evaluate high-throughput, low-variance paths through an arbitrary number of interfaces and media layers, and
- a benchmark across many scenes, consistently demonstrating  $\sim 2 - 6\times$  performance gains over the state-of-the-art.

Our method supports *complex layered materials* with arbitrary *per-interface* properties (Figure 2), including spatially-varying roughness, albedo, high-frequency normal map textures, as well as isotropic and anisotropic scattering media with spatially-varying scattering properties *between* interfaces.

## 2. Related Work

Our approach builds upon the *position-free* framework of Guo et al. [GHZ18], where microscale light transport depends only on depth and direction. We briefly review and relate prior work on layered materials and specialized sampling schemes.

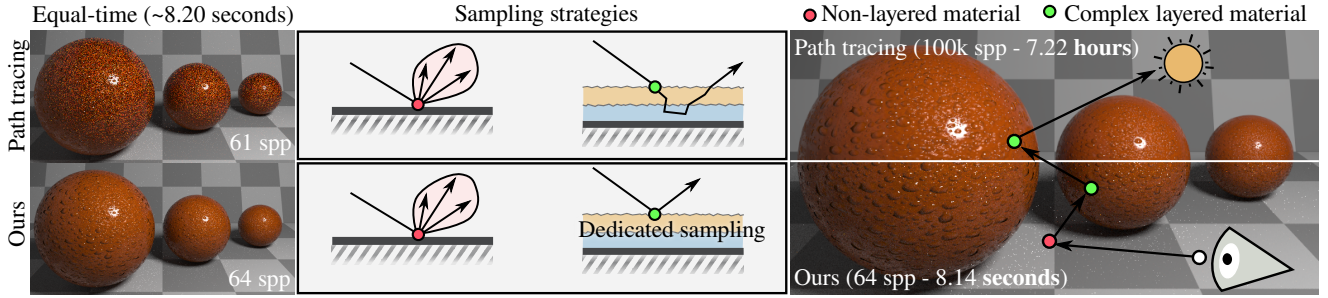
**Microfacet Models.** Microfacet theory [BS63, CT81, BSN16] is a powerful tool that allows us to model the scattering of light on a surface in a well founded way. Recently, these models have been extended to support energy conserving methods that account for multiple scattering in the material [HHdD16, LJJ\*18]. While these improvements have pushed further the realism offered by the model, the multiple scattering operations occur in the microfacets of a single material and they cannot capture the resulting appearance of a layered configuration in which light interactions depend on the properties of each layer present in the material. While some methods extend microfacet models to account for layer-like structures, e.g., thin transparent layers [QGQP17], these are restricted to such fixed configurations.

**Approximate or Precomputed Layered Materials.** Hanrahan and Krueger introduced the seminal work [HK93] that supports multiple material layers by means of a first-order approximation to single scattering events in the media and multiple scattering using a Monte Carlo simulation on discretized directions. Weidlich and Wilkie provided a flexible and fast method [WW07] that used a combination of the BSDFs at each interface with a simplified transport assumption and no support for multiple scattering events. While the results are plausible, the complex light paths that form between and across the layers mandate the incorporation of complex light transport algorithms instead of phenomenological methods to accurately model the material.

Jakob and colleagues provided a robust framework [JdJM14] that allowed to compose aggregate scattering behaviors for layered materials with different layer components using the adding-doubling method [vdH80] in a frequency-space representation. Follow-up work [ZJ18] extends the framework to anisotropic interfaces. While these methods offer high performance and are accurate, their main drawback is that they rely on potentially large precomputed data for a fixed set of parameters, making them virtually unsuitable for modeling arbitrary spatially-varying properties present in *complex layered materials*. Our method does not require precomputation, does not rely on basis space expansions and seamlessly supports spatial variation of all parameters.

Belcour presents a statistical model of layered light transport [Bel18] that operates on first- and second-order moments using a novel adding-doubling scheme. Their resulting approximation uses a mixture of GGX [WMLT07] lobes, yielding a high-performance solution for isotropic GGX BSDFs with moderate roughness. Our method does not approximate any transport, is not restricted to distribution-specific isotropic interfaces and supports a wider range of (potentially spatially-varying) parameters.

**Explicit Simulation.** A naïve way to treat layered materials could explicitly model the layers' geometry, solving transport without any disambiguation of micro- and macro-level scattering. Such an



**Figure 3:** Rendering complex layered materials with explicit geometric models for the layers is impractical (top left). We instead perform a separate microscale light transport simulation inside the material, correctly simulating the multi-layer surface and volume scattering effects – orders of magnitudes faster than an explicit model – including multi-scale subpixel antialiasing due to high-frequency spatial variations in the micromorphologies, interface roughnesses and volumetric scattering/absorption parameters.

approach would be prohibitively costly while not exploiting any of the structure inherent to such layered materials (Figure 3).

Our approach is most similar to Guo et al.’s method [GHZ18], consisting of a path-space simplification tailored to the layered material context. Here, transport is modeled in a *position-free* manner, where tangential translation is ignored. This method employs a bidirectional estimator for multi-layer transport, however, we will show that a full multiple importance sampled bidirectional method is not necessary to effectively simulate layered light transport. Instead, we will tailor our subsurface path construction to only form those subpaths contributing to low-variance transport estimates – simplifying implementation and improving performance.

**Advanced Sampling.** Bidirectional path tracing (BDPT) [LW93, VG95a] improves robustness by combining sensor and light subpaths using Multiple Importance Sampling (MIS) [VG95b]. BDPT is particularly effective when variance reduction is due to asymmetries in the effectiveness of sensor and light subpath sampling; in the position-free formulation, however, these asymmetries are not present (i.e., importance and radiance are symmetric). VCM/UPS [HPJ12, GKDS12] techniques combined density estimators and BDPT to further improve robustness when non-separable paths are present. However, this type of path is non-existent in the position-free formulation as it is always possible to “move” path vertices horizontally to discard delta interactions. The same reasoning applies for Manifold Next-Event Estimation [HDF15] where the use of Manifold Exploration [JM12] to solve the Dirac function analytically becomes unnecessary.

### 3. Position-free Layered Formulation

Guo et al. [GHZ18] introduced the position-free formulation that simplifies the light transport inside the material by using depth-direction –instead of position-direction– dependent paths. Its flexibility allows it to accurately model a large variety of layered materials as a BSDF. This approach admits Monte Carlo solutions for simulating light transport interactions within the material. We will review the fundamentals of this theory and highlight important aspects that motivate our more effective path sampling strategies as opposed to sophisticated, e.g., bidirectional, estimators.

### 3.1. Light Transport and Assumptions

The path integral formulation [Vea97] models light’s interactions in an environment in the form of an integral over light paths,

$$I = \int_{\mathcal{P}} f(\bar{\mathbf{x}}) d\mu(\bar{\mathbf{x}}) \quad (1)$$


where  $\mathcal{P}_k = \{\bar{\mathbf{x}} | \mathbf{x}_0 \mathbf{x}_1 \dots \mathbf{x}_k\}$  is the set of length- $k$  paths with scattering vertices  $\mathbf{x}_i$ ,  $\mathcal{P} = \bigcup_{k=1}^{\infty} \mathcal{P}_k$  is the set of transport paths of all lengths,  $\mu$  is a measure on this space of paths, and  $f$  is called the measurement contribution function.

We could apply this formulation to layered materials by explicitly modeling the geometry of each layer as separate objects in the scene before applying, e.g., Monte Carlo path tracing (Figure 3, middle). This is inefficient since surface scattering interactions may now require additional ray-tracing acceleration structures, and no macro- vs. microscale disambiguation is leveraged. Moreover, this integration process remains agnostic to the context of a material’s layered structure, and so it cannot explicitly exploit this structure to more intelligently construct light transport paths. For example, direct illumination-like evaluations will not treat underlying layers that are occluded by overlying layers, increasing the variance of such paths (Figure 3, left).

**Complex Layered Materials.** A layered material consists of a set of stacked parallel slabs, each with an associated BSDF (interface layer) and/or participating medium, as in Figure 2. We treat interfaces with high-frequency, spatially-varying: anisotropic roughness, indexes of refraction, specular coefficients and micro-normal variation. Medium layers can have isotropic or anisotropic phase functions and spatially-varying scattering/absorption coefficients.

Similar to previous work, we assume a far-field approximation in which incoming and outgoing directions are considered to be located at the same point. Despite the small displacement between these points, the relative distances from the camera and light source with respect to the layer thickness make this a reasonable assumption. This significantly simplifies the light transport problem as the spatially-varying properties of the layers are assumed to remain constant at the evaluation point. Moreover, we do not rely on any acceleration structure for computing intersections as all the interfaces are parallel and it suffices to track depth. The effect of this assumption can be assessed by comparing images with those obtained by performing a full simulation (Figure 3).





**Figure 4:** Macroscale path tracing through a model of the material as, e.g., a set of independent layered objects (left), compared to the same paths in the position-free formulation – i.e., where all events are dependent on depth  $z$  for the given directions, being located at the same horizontal  $(x, y)$  position. Eliminating tangential spatial-variation in reflectance simplifies the layered BSSRDF to a BSDF.

**Motivation.** By using multiple analytic BSDF/phase function models while accurately simulating light interactions, our method offers a large panel of expressiveness. However, it is necessary to use numerical estimates of light interactions inside the layered material as there is no general analytical solution to model the interactions of the different layers. Under the far-field assumption, the formulation does not account for color bleeding or caustics produced from neighboring positions and their respective shading normals. Additionally, importance from the sensor or light sources is symmetric inside this model, see Figure 5.

These observations indicate that a bidirectional method is unnecessary. We are going to exploit these properties to efficiently compute the transport by intelligently sampling the paths using simpler and less heavy weight methods.

### 3.2. Position-free Formulation

In the path integral formulation, path vertices implicitly include all the directional information of a path. A position-free (i.e., depth-direction) path parameterization, however, requires a change of variables from path vertices to in/out-directions and depths. For example, a path  $\bar{\mathbf{x}} = \omega_0 z_1 \omega_1 z_2 \omega_2 \dots z_n \omega_n$  that enters the material along direction  $\omega_i$  and exits along  $\omega_o$  has  $\omega_0 = \omega_i$  and  $\omega_n = \omega_o$ . To represent vertex  $\mathbf{x}_j$ ,  $\omega_{j-1}$  and  $\omega_j$  correspond to the incoming and outgoing directions at depth  $z_j$  inside the layered material. As such, this reparameterization collapses all tangentially-translated paths onto a single depth-direction configuration (Figure 4).

With this simplification, spatially-varying parameters at the interfaces or participating media remain constant for any single/fixed path sample, i.e., during Monte Carlo integration, with their value(s) corresponding to those located at the incident surface point.

Consider first a scattering event at depth  $z \in [0, d]$  inside one of the participating media layers (with thickness  $d$ ) of our material. We define the radiance coming from direction  $\omega$  as:

$$L_o(z, \omega) = \sigma_s \int_0^{r(z, \omega)} L_s(z^*, \omega) \text{Tr}(\sigma_t, z, z^*, \omega) dz^* \quad (2)$$

where the integration bounds depend on whether we are in the process of tracing a path into ( $z$  to  $d$ ) or out of ( $0$  to  $z$ ) the material. We encapsulate this subtlety by always integrating from  $0$  to  $r$ , and determining the correct depth-length in the following way:

$$r(z, \omega) = \begin{cases} z & \text{if } \cos \theta > 0 \\ d - z & \text{otherwise} \end{cases}$$

where  $\cos \theta$  is the third Cartesian component of direction  $\omega$ .

In depth-space, the transmittance term is defined for segments between depths  $z$  and  $z^*$  in the layer instead of the distance between two points in three-dimensional space:

$$\text{Tr}(\sigma_t, z, z^*, \omega) = \exp(-\sigma_t |z - z^*| / |\cos \theta|) \quad (3)$$

and multiple scattering  $L_s$  with a phase function  $\rho$  is:

$$L_s(z^*, \omega) = \int_{S^2} L_i(z^*, \omega_i) \rho(\omega, \omega_i) d\omega_i \quad (4)$$

Finally, when the depth  $z$  is at a BSDF interface  $f_s^z$ , we have:

$$L_o(z, \omega) = \int_{S^2} f_s^z(\omega, \omega_i) |\cos \theta_i| L_i(z, \omega_i) d\omega_i \quad (5)$$

We can resolve Equation 1 using Monte Carlo integration, starting at the first surface interface and sampling surface and volume interactions in-and-out of the layered material until we form a valid transport path. Here, sampling a path is straightforward, but constraining the Monte Carlo process to start and finish along fixed incident and outgoing directions (i.e., the ones required when evaluating the BSDF and PDF) is far more challenging. The manner in which we build these paths determines both performance and variance. Unlike the state-of-the-art position-free technique [GHZ18] we do not exhaustively sample all bidirectional path construction techniques, instead, we use forward path tracing and make explicit light connections across an arbitrary number of interface layers.

### 4. Efficient Estimator for Layered Materials

At the macroscopic scale, MIS-based path sampling algorithms estimate surface light transport as:


$$L(\omega_i) \approx \overbrace{w_1 L_e(\omega_o) \frac{f_s(\omega_i, \omega_o)}{p(\omega_o | \omega_i)}}^{\text{PDF Computation (Sec. 4.3)}} + \overbrace{w_2 \frac{L_e(\omega_l)}{p(\omega_l)} f_s(\omega_i, \omega_l)}^{\text{Evaluation (Sec. 4.1)}} \quad (6)$$

Sampling (Sec. 4.2)

where  $\omega_o$  and  $\omega_l$  are outgoing directions sampled according to BSDF and light distributions,  $L_e$  is the emitted luminance,  $w_1$  and  $w_2$  are the MIS weights to blend the BSDF and light sampling strategies. The strategies themselves involve computing the evaluation and sample operations of the layered BSDF. Complementarily, the probability to sample an outgoing direction given an incoming direction is required to compute MIS weights. Note that we group function and PDF evaluations as a quotient – i.e.,  $f_s/p$  and  $L_e/p$  – to disambiguate the distribution from which samples are drawn. Moreover, our sampler computes this ratio directly, without independently evaluating  $f_s$  and  $p$ , ensuring unbiasedness.

We describe the design of an estimator that enables fast and practical evaluation for an arbitrary number of layers, illustrating the main differences compared to Guo et al.'s bidirectional integrator.





**Figure 5:** Visualizing the impact that different bidirectional sampling strategies have, across a 2D parameterization of many 3-interface-2-media layered material configurations:  $(s,t)$  sampling strategies (yellow) are visualized – without MIS – across path lengths. For each path length, more elaborate bidirectional strategies do not significantly reduce variance (red), relative to the simpler  $(s,1)$  strategies. The variance’s uniformity across strategies indicates that bidirectional construction is not necessary for transport estimation in layered materials.

#### 4.1. Evaluation

Any single evaluation of the layered BSDF  $f_s$  (e.g., in a Monte Carlo estimator for Equation 6) – for a fixed pair of incident and outgoing directions – requires solving a nested integration problem over the microscale and subsurface path subspace inside the layered material. We denote this subspace by  $\mathcal{P}^\delta$ , and it is defined as the set of all paths that go through the layered material, starting along  $\omega_i$  and exiting along  $\omega_o$ .

This evaluation can thus be written as an integral over paths as:

$$f_s(\omega_i, \omega_o) = \int_{\mathcal{P}^\delta} f(\bar{\mathbf{x}}) d\mu(\bar{\mathbf{x}}) \quad (7)$$

for which we can estimate a solution using Monte Carlo and the position-free parameterization. Here, Guo et al.’s method [GHZ18] employs a bidirectional estimator that exhaustively samples all  $(s,t)$  strategies, combining them with MIS. Their method also eliminates the geometry factor by sampling directions at the subpath endpoints proportional to their local scattering model (Figure 6).

Employing a bidirectional estimator in this setting is computationally wasteful for two key reasons: first, the different strategies within a same path length do not exhibit variance reduction benefits and indeed, some strategies occasionally increase it; and second, the BDPT pyramid formed by all  $(s,t)$  strategies is symmetric.

We illustrate this behavior in Figure 5, across a wide variety of surface and media scattering settings in a layered material. As is evident here, the majority of the variance reduction in the Monte Carlo estimate of Equation 7 is due only to a consistent subset of these bidirectional strategies. As such, the gains of exhaustively considering every bidirectional strategy does not justify the overhead associated to their construction, let alone their combination using MIS. We propose a more selective and efficient strategy.

**Our Approach.** Motivated by these empirical observations, we design a path-based sampler that employs specialized *multi-layer* next-event estimates. More precisely:

1. we trace a full path through (i.e., in and out of) the layered material, starting along  $\omega_i$ , keeping track of every scattering vertex,
2. we sample directions, at each interface, along  $-\omega_o$  (i.e., a path in the opposite direction), and

3. we use these directions to connect every vertex – explicitly with next-event estimation – generated in #1.

In doing so, we are able to more efficiently support an arbitrary number of layers (i.e., interfaces and media), each with potentially high-frequency spatial variation in their scattering properties.


Performing connections in this way potentially leads to double counting path contributions. To avoid this, we do not perform an explicit connection if:

- the vertex we wish to connect from is located on an interface,
- that interface connects to the path arriving from  $-\omega_o$ , and
- the connection incurs a transmission event across the interface.

We illustrate this procedure in Figure 6.3, where we connect all vertices except for the last two. We can interpret these last two vertices as residing on the emitter, and the next event estimations as performing explicit path tracing. Adding the emission of such vertices yields an incorrect result, since emission has already been accounted for in previous vertices (i.e., as direct illumination).

**Discussion.** Note that Guo et al. [GHZ18] also suggested the use of a unidirectional method with next-event estimation for integrating light transport in layered materials. However, their next-event estimation is only able to connect vertices that are directly visible from the top interface (i.e., vertices in the adjacent medium or on the upper side of the second interface). For deeper layers, they used a nested integrator to compute the outgoing direction and the BSDF evaluation. The main motivation for this approach is to avoid double counting light paths. Moreover, as they were performing next-event estimation, they need to compute the MIS weights for each sampling technique. More specifically, if a vertex lies on an interface represented by a nested model, making a connection involves performing a sample, evaluation and two PDF computations of their model to make an MIS connection. This approach is correct but introduces significant overhead which in turn makes it uncompetitive against their own bidirectional estimator.

On the contrary, our approach generates one path for all layers and performs next-event estimation explicitly without a nested approach. Specifically, we need only add path segments when forming a connection, as opposed to grouping deep layers through a combination of sampling, function evaluation and (two) PDF eval-



**Figure 6:** Various light transport sampling strategies in layered materials. Green checkmarks denote an explicit next-event connection that exits the material along direction  $\omega_o$ . Guo et al.’s method traces two subpaths (left), then samples a direction and “virtually” translates an endpoint to connect them (left middle). Guo et al.’s unidirectional strategy (right middle) can only make explicit connections to vertices that are above the second interface, otherwise it requires an expensive recursive and nested sampling process along directions  $(\tilde{\omega}_i, \tilde{\omega}_o)$ . Our method estimates transport across all layers with a single path, improving efficiency and robustness (right).

uations over the *nested* material. When compared to Guo et al.’s unidirectional method, our approach is very efficient and competes favourably against more complicated estimators. As our approach only considers a subset of bidirectional sampling strategies, this simple but powerful algorithm allows our method to excel at computing BSDF evaluations, as can be observed in the evaluation-only comparison in Figure 7.

## 4.2. Sampling

Sampling an outgoing direction and computing the resulting sampling weight (i.e.,  $eval/PDF$ ) is fairly straightforward. For this operation, only the incident direction is known and in this regard, our approach and that in the state-of-the-art are equivalent.

More precisely, we generate a path through the material, with initial direction  $\omega_i$ , using importance sampling at each scattering event. Once the path exits the material, the sample is its outgoing direction and its throughput is the sampling weight. One fundamental advantage of this approach is that the exact PDF of the sampled direction is already included in the path throughput. This sampling operation, coupled with our evaluation method (Section 4.1), leads to two unbiased BSDF and light sampling estimators. We can combine these estimators with MIS (Equation 6).

## 4.3. Probability Density Function

Modern Monte Carlo rendering techniques leverage MIS to robustly combine the BSDF and light sampling strategies in order to reduce variance. At the macroscale level integrator, the one sample value given by Equation 6 requires not only the evaluation of the BSDF (Section 4.1) but also the evaluation of its probability density function, which appears two times within the strategy weights  $w_1$  and  $w_2$ . Just like evaluation, the PDF is also a nested integration problem with the same domain  $\mathcal{P}^\delta$ .

The probability density function of the layered BSDF for an outgoing direction  $\omega_o$  given an incoming direction  $\omega_i$  is defined as:


$$p(\omega_o|\omega_i) = \int_{\mathcal{P}^\delta} P(\bar{x}) d\mu(\bar{x}) \quad (8)$$

where  $P(\bar{x})$  is the product of the sampling PDFs at each vertex and along the segments in path  $\bar{x}$ .


Estimating this integral is similar to the layered BSDF evaluation and so we could readily replicate the procedure detailed in Section 4.1; however, we can leverage the fact that these estimates will be used exclusively for MIS weight computation to further accelerate our approach. To do so, we will rely on an approximation that resembles the true PDF but is much faster to estimate. Note that the final estimator remains unbiased if the MIS weights:

1. are estimated using an independent sampling process [VG95b] – namely,  $E[w_0 f] + E[w_1 f] = E[w_0]E[f] + E[w_1]E[f]$ , and
2. have a normalized expected value of their sum, yielding  $E[f](E[w_0] + E[w_1]) = E[f]$ .

In practice, stochastically estimating the PDF violates condition #2 due to Jensen’s inequality [Cha87], however the error introduced is constrained to MIS weights where its effect will become negligible after combining the two unbiased estimators. To enforce condition #1, we do not reuse paths from the evaluation or sampling processes and, instead, resample paths with a new strategy described below.



**Figure 7:** Equal-time direct illumination with light (left) and multiple importance (right) sampling, comparing the effective cost of material evaluations. Our method renders  $\sim 11\times$  more samples using only light sampling. With MIS, we improved the sample-per-pixel cost by a factor of  $\sim 7\times$ .



**Figure 8:** Log-scale, 1-sample Monte Carlo estimates of the PDF: our method (middle) and [GHZ18] (left). As expected, both methods yield the same PDF profile and similar variance, however ours is  $1.7\times$  faster. The ground truth PDF (right) is very costly to compute and has the same dominant features.

Guo et al. [GHZ18] optimize their implementation by:

- omitting scattering in participating media, since its impact on the PDF profile is minimal, and
- restricting the maximum path length to the number of layers.

These simplifications can alter the PDF's profile, yielding incorrect results. To mitigate these issues, their method adds a constant/diffuse term to compensate for any missing transport. Doing so also ensures a non-zero PDF, and their PDF estimator relies on a bidirectional method similar to their BSDF evaluation.

**Our Approach.** We use a similar technique to the BSDF evaluation (Section 4.1) following the approximations stated above. Aside from being unidirectional, we benefit from using importance sampling connections as opposed to *local* MIS connections as we found that the benefits are minimal.

In Figure 8 we show a comparison between Guo et al.'s approximate PDF estimation using their bidirectional path tracer and our efficient approximate PDF. Both techniques converge to the same PDF shape. For lower and equal sampling count, our approach is significantly faster with marginally noisier results. Having a fast PDF estimation method is essential for high-performance light transport algorithms (Figure 7, right).

## 5. Results

We render our scenes in Mitsuba [Jak10] on a dual Intel Xeon Gold 6148 with 40 cores @ 2.4 GHz, reporting equal-time and equal-quality render times (excluding scene loading) in Table 1. Reference images are rendered using Guo's bidirectional method and a high sampling rate.

**Implementation Details.** For rendering all our scenes, we further tuned the evaluation method so that explicit connections between two interfaces are performed using MIS. When connecting medium scattering events, we importance sample the connecting interface and perform an evaluation of the phase function. Indeed, in this case, the probability of sampling the correct direction with the phase function is low, which results in low MIS weights in general. The application of MIS for these vertices is trivial, but comes at a computational cost particularly when dealing with dense scattering media and long paths.

The simplicity of our implementation affords a  $2 - 3\times$  performance gain over Guo's implementation when sampling directions proportional to the complex layered material. The reason of this speedup is Guo et al.'s sampling code is shared with their bidirectional and unidirectional estimator. This has the consequence that unnecessary extra information is computed and stored for the unidirectional sampling strategy. In contrast, while our implementation also shares the tracing procedure, we use the unidirectional estimator for all layered material procedures. Our implementation is easier to maintain and allows low-level optimization.

**Performance Analysis.** We modified the RED MUG so that the top layer had a higher roughness of  $\alpha = 0.05$  (Figure 7). This modification is motivated by the poor performance of the light sampling on smooth BSDFs. Without this modification, most of the variance will come from the top layer, making the performance analysis of this layered material impractical. The light sampling case only uses the evaluation routine over the complex layered material. The BSDF-Light MIS case (Equation 6) shows the performance of all components of our method. We observe that our method outperforms the state-of-the-art methods by a large margin.

**Equal-time Comparisons.** The COFFEE TABLE scene is a closed room environment with all lighting due to two area light windows on the right. The orange layered material on the coffee machine consists (from the bottom, up) of a diffuse interface followed by a colored absorbing and scattering medium and a rough dielectric interface ( $\alpha = 0.001$ ). The table's layered material comprises a normal-mapped diffuse interface with spatially-varying albedo, followed by a thin absorbing and scattering medium, and a rough dielectric surface interface ( $\alpha = 0.15$ ). The left and center mugs have layered materials with a base, unit-albedo diffuse interface underneath three (left) and two (center) layers of participating media, one of which has spatial variation and an additional rough dielectric interface ( $\alpha = 0.001$ ). The red mug has a high-frequency brushed aluminum normal map base interface ( $\alpha = 0.05$ ) followed by two dense colored absorbing and scattering media, and a rough dielectric top interface ( $\alpha = 0.001$ ). Finally, the teapot's layers comprise a high-resolution normal-mapped metallic surface layer ( $\alpha = 0.025$ ), followed by a uniform absorbing medium and a high-resolution scratched normal map at the top-most surface interface.

The comparison matrix (Figure 9) compares our method to Guo et al. [GHZ18] on their GLOBE scene and our CAR scene. The layered car paint mimics real-world compositions: a metallic base and two coats of dense participating media – which contribute to an underlying matte tone – followed by a clear/shiny coat. Error maps illustrate the reduction in variance afforded by our method. Table 1 similarly highlights our  $1.8 - 3.4\times$  improvement in MSE – at equal-time and across every scene – compared to the state-of-the-art. This improvement is due, in large part, to the  $2.1 - 7.1\times$  reduction in per-sample cost. In summary, this leads to a  $1.9 - 5.8\times$  effective performance gain in equal-quality comparisons.

**Comparison to [Bel18].** We modify the COFFEE TABLE and TEAPOT scenes (Figures 1 and 10) for equal-time comparisons – replacing non-translucent interfaces with rough conductors, adjusting roughness and media parameters, and removing features not



SCENE	equal-time						equal-quality				
	target time	Ours		[GHZ18]		per sample speedup	Ours		[GHZ18]		effective speedup
		spp	MSE	spp	MSE		spp	time	spp	time	
COFFEE TABLE	12.33m	2048	1.00×	960	1.80×	<b>2.13×</b>	2048	12.33m	1860	23.70m	<b>1.92×</b>
RED MUG	14.61s	64	1.00×	9	3.42×	<b>7.11×</b>	64	14.61s	52	1.41m	<b>5.79×</b>
TEAPOT	1.49m	512	1.00×	104	3.01×	<b>4.92×</b>	1024	2.98m	720	10.83m	<b>3.63×</b>
GLOBE	10.46s	256	1.00×	59	2.62×	<b>4.33×</b>	1024	2.19m	680	8.77m	<b>4.00×</b>
CAR	15.89s	128	1.00×	37	2.68×	<b>3.46×</b>	1024	1.91m	1010	6.32m	<b>3.31×</b>

Table 1: Equal-time and equal-quality performance breakdown for our scenes.




Figure 9: Equal-time comparison to [GHZ18], visualizing absolute difference w.r.t. converged renders. Zoom-ins highlight regions where layered materials are both direct and indirectly visible.

supported by Belcour’s method, like normal maps. Their approximation is sample efficient ( $2\times$  more than ours), but at a significant loss of detail compared to renderings on the unmodified scenes.


### 6. Conclusion and Future Work

We presented an efficient position-free path-based approach for evaluating the contribution of – and sampling directly from – complex layered materials. Our method supports transport effects due to arbitrary high-frequency spatial variation of reflectance and microgeometry properties at the layer interfaces, as well as anisotropic participating media between these interfaces. By only constructing a subset of bidirectionally-constructed light transport paths that contribute most to the aggregate multi-layer reflectance, we consistently outperform the state-of-the-art across a variety of challenging material configurations. As a result, our implementation – which we will release as an open-source plugin to the MITSUBA renderer [Jak10] – is more compact, relying on fewer data structures and simple path-construction logic.

Our method inherits some limitations, the most important of which is its lack of support for delta reflectance and/or transmittance distribution profiles at the layer interfaces. Due to the manner in which we form our subsurface light transport paths, we cannot support such distributions in general without fundamentally changing our connection strategy. In practice, we use interfaces with very low roughness values to form “near-delta” distributions. In addition to addressing this limitation, two interesting avenues of future work would be to incorporate a multi-scale wave optics model based on thin layers (motivated in part by recent work [YHW\*18]), and to more explicitly treat indirect lighting contributions due to thicker heterogeneous media between the layers [GKH\*13].

### Acknowledgements

We acknowledge Compute Canada for computing resources, the Consejo Nacional de Ciencia y Tecnología (CONACYT) of Mexico and Fonds de Recherche du Québec (FRQ) for funding, and Bitterli and Guo et al. for scenes [Bit16, GHZ18]. The coffee ma-



**Figure 10:** The original TEAPOT scene using high-resolution normal maps at each interface (top). Belcour's method [Bel18] requires a simpler modified material (middle), which we can also render with our method (bottom). Nevertheless, our method generates 512spp in equal-time for both images, and Belcour's 1170spp.

chine model is by cekuhen, and the environment map for car scene is from HDRI-Hub.com (both under CC BY 3.0).

## References

- [Bel18] BELCOUR L.: Efficient rendering of layered materials using an atomic decomposition with statistical operators. *ACM Transactions on Graphics* 37, 4 (July 2018), 73:1–73:15. doi:10/gd52pd. 1, 2, 7, 9
- [Bit16] BITTERLI B.: Rendering resources, 2016. URL: <https://benediktbitterli.me/resources/>. 8
- [BS63] BECKMANN P., SPIZZICHINO A.: *The Scattering of Electromagnetic Waves from Rough Surfaces*. Pergamon Press, NY, 1963. 2
- [BSN16] BAGHER M. M., SNYDER J., NOWROUZEZAHRAI D.: A non-parametric factor microfacet model for isotropic BRDFs. *ACM Transactions on Graphics (SIGGRAPH)* 35, 5 (July 2016), 159:1–159:16. doi:10/f85ktv. 2
- [Cha87] CHANDLER D.: *Introduction to Modern Statistical Mechanics*. Oxford University Press, 1987. 6
- [CT81] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. *Computer Graphics (SIGGRAPH)* 15, 3 (Aug. 1981), 307–316. doi:10/br5ps6. 2
- [GHZ18] GUO Y., HAŠAN M., ZHAO S.: Position-free Monte Carlo simulation for arbitrary layered BSDFs. *ACM Transactions on Graphics* 37, 6 (Dec. 2018), 279:1–279:14. doi:10/db3c. 1, 2, 3, 4, 5, 7, 8
- [GKDS12] GEORGIEV I., KRIVÁNEK J., DAVIDOVIČ T., SLUSALLEK P.: Light transport simulation with vertex connection and merging. *ACM Transactions on Graphics (SIGGRAPH Asia)* 31, 6 (Nov. 2012), 192:1–192:10. doi:10/gbb6q7. 3
- [GKH\*13] GEORGIEV I., KRIVÁNEK J., HACHISUKA T., NOWROUZEZAHRAI D., JAROSZ W.: Joint importance sampling of low-order volumetric scattering. *ACM Transactions on Graphics (SIGGRAPH Asia)* 32, 6 (Nov. 2013), 1–14. doi:10/gbd5qs. 8
- [GQGP17] GUO J., QIAN J., GUO Y., PAN J.: Rendering thin transparent layers with extended normal distribution functions. *IEEE Transactions on Visualization and Computer Graphics* 23, 9 (Sept. 2017), 2108–2119. doi:10/gbtbfj. 2
- [HDF15] HANIKA J., DROSKE M., FASCIONE L.: Manifold next event estimation. *Computer Graphics Forum (Proc. of the Eurographics Symposium on Rendering)* 34, 4 (2015), 87–97. doi:10/f7mbxk. 3
- [HHdD16] HEITZ E., HANIKA J., D'EON E., DACHSBACHER C.: Multiple-scattering microfacet BSDFs with the Smith model. *ACM Transactions on Graphics (SIGGRAPH)* 35, 4 (July 2016), 58:1–58:14. doi:10/f89kkm. 2
- [HK93] HANRAHAN P., KRUEGER W.: Reflection from layered surfaces due to subsurface scattering. In *Annual Conference Series (SIGGRAPH)* (New York, NY, USA, 1993), ACM Press, pp. 165–174. doi:10/b4tw3j. 2
- [HPJ12] HACHISUKA T., PANTALEONI J., JENSEN H. W.: A path space extension for robust light transport simulation. *ACM Transactions on Graphics (SIGGRAPH Asia)* 31, 6 (Jan. 2012), 191:1–191:10. doi:10/gbb6n3. 3
- [Jak10] JAKOB W.: Mitsuba renderer, 2010. URL: <http://www.mitsuba-renderer.org>. 7, 8
- [JdJM14] JAKOB W., D'EON E., JAKOB O., MARSCHNER S.: A comprehensive framework for rendering layered materials. *ACM Transactions on Graphics (SIGGRAPH)* 33, 4 (July 2014), 118:1–118:14. doi:10/f6cpsq. 2
- [JM12] JAKOB W., MARSCHNER S.: Manifold exploration: A Markov chain Monte Carlo technique for rendering scenes with difficult specular transport. *ACM Transactions on Graphics (SIGGRAPH)* 31, 4 (July 2012), 58:1–58:13. doi:10/gfzq4p. 3
- [LJJ\*18] LEE J. H., JARABO A., JEON D. S., GUTIERREZ D., KIM M. H.: Practical multiple scattering for rough surfaces. *ACM Transactions on Graphics* 37, 6 (Dec. 2018), 275:1–275:12. doi:10/db3d. 2
- [LW93] LAFORTUNE E. P., WILLEMS Y. D.: Bi-directional path tracing. In *Proc. of the International Conference on Computational Graphics and Visualization Techniques (Compugraphics)* (Alvor, Portugal, Dec. 1993), vol. 93, pp. 145–153. 3
- [vdH80] VAN DE HULST H.: *Multiple light scattering: tables, formulas, and applications*. Multiple Light Scattering: Tables, Formulas, and Applications. Academic Press, 1980. 2
- [Vea97] VEACH E.: *Robust Monte Carlo Methods For Light Transport Simulation*. PhD thesis, Stanford University, 1997. 3
- [VG95a] VEACH E., GUIBAS L. J.: Bidirectional estimators for light transport. In *Photorealistic Rendering Techniques (Eurographics Workshop on Rendering)* (1995), Springer-Verlag, pp. 145–167. doi:10/gfznbh. 3
- [VG95b] VEACH E., GUIBAS L. J.: Optimally combining sampling techniques for Monte Carlo rendering. In *Annual Conference Series (SIGGRAPH)* (Aug. 1995), vol. 29, ACM Press, pp. 419–428. doi:10/d7b6n4. 3, 6
- [WMLT07] WALTER B., MARSCHNER S. R., LI H., TORRANCE K. E.: Microfacet models for refraction through rough surfaces. In *Rendering Techniques (Proc. of the Eurographics Symposium on Rendering)* (June 2007), Eurographics Association, pp. 195–206. doi:10/gfz4kg. 2
- [WW07] WEIDLICH A., WILKIE A.: Arbitrarily layered micro-facet surfaces. In *Proceedings of the 5th International Conference on Computer Graphics and Interactive Techniques in Australia and Southeast Asia* (2007), GRAPHITE '07, pp. 171–178. doi:10/dcnhrf. 2
- [YHW\*18] YAN L.-Q., HAŠAN M., WALTER B., MARSCHNER S., RAMAMOORTHY R.: Rendering specular microgeometry with wave optics. *ACM Transactions on Graphics* 37, 4 (July 2018), 75:1–75:10. doi:10/gd52td. 8
- [ZJ18] ZELTNER T., JAKOB W.: The layer laboratory: A calculus for additive and subtractive composition of anisotropic surface reflectance. *ACM Transactions on Graphics* 37, 4 (July 2018), 74:1–74:14. doi:10/gd52tp. 2