

A Unified Manifold Framework for Efficient BRDF Sampling based on Parametric Mixture Models

Sebastian Herholz¹ Oskar Elek² Jens Schindel¹ Jaroslav Křivánek² Hendrik P. A. Lensch¹

¹Tübingen University ²Charles University, Prague




Figure 1: Sampling quality of our unified sampling framework for an analytic isotropic BRDF model with spatially varying roughness values (left) and a measured BRDF from the MERL database (right). By using a PMM representation for the BRDF sampling function, our method outperforms the current state-of-the-art sampling methods (as seen in the insets at the top and the false-color variance plots at the bottom).

Abstract

Virtually all existing analytic BRDF models are built from multiple functional components (e.g., Fresnel term, normal distribution function, etc.). This makes accurate importance sampling of the full model challenging, and so current solutions only cover a subset of the model’s components. This leads to sub-optimal or even invalid proposed directional samples, which can negatively impact the efficiency of light transport solvers based on Monte Carlo integration. To overcome this problem, we propose a unified BRDF sampling strategy based on parametric mixture models (PMMs). We show that for a given BRDF, the parameters of the associated PMM can be defined in smooth manifold spaces, which can be compactly represented using multivariate B-Splines. These manifolds are defined in the parameter space of the BRDF and allow for arbitrary, continuous queries of the PMM representation for varying BRDF parameters, which further enables importance sampling for spatially varying BRDFs. Our representation is not limited to analytic BRDF models, but can also be used for sampling measured BRDF data. The resulting manifold framework enables accurate and efficient BRDF importance sampling with very small approximation errors.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Keywords: BSDF, Importance Sampling, Product Importance Sampling, Global Illumination, Monte-Carlo Sampling

1. Introduction

The recent widespread adoption of physically-based rendering in the visual effects industry [Jar13, FWKH17] has had, among others, the following two effects. First, complex analytical or measured *bi-directional reflectance distribution functions* (BRDFs) are the prevalent way of modeling material appearance nowadays [Bur12]. Second, the need to accurately sample the BRDFs in the probabilistic context of Monte Carlo rendering methods has become essential. This is because the efficiency of Monte Carlo solvers largely depends on their ability to sample transport paths proportional to the illumination integrand of the rendering equation [Kaj86, Vea97].

Since the BRDF is one of the two key factors in the integrand, any mismatch between its shape and the sampling *probability density function* (PDF) will negatively impact the variance of the resulting light transport solution.

Our work addresses exactly this issue: to provide a comprehensive sampling methodology for complex BRDFs. Since analytic physically-based BRDFs typically consist of multiple nontrivial factors [PJH16], deriving a tight PDF proportional to their shape is a research challenge – especially because such PDFs need to be model-specific [WMLT07, Hd14]. For measured BRDFs, there ex-

ists no inherently correct sampling strategy, hence one has to resort to expensive tabulated sampling.

Considering the above, the core idea of our approach is to use a generic parametric model to agnostically represent the shape of each BRDF. Our approach builds on the pioneering work of Lafor-tune et al. [LFTG97] to represent BRDFs using a unified, non-linear, *parametric mixture model* (PMM, Sec. 3). While we use standard numerical optimization techniques for fitting the PMM to the BRDF’s shape, there are two key improvements (Sec. 4). First, our objective function is designed to optimize the fit for sampling, not direct visualization. And second, we perform the fitting in the *meta-space of the PMM parameters*. We show that with proper regularization, in this space the parameters of the PMM form a smooth *manifold*, and can thus be represented by an interpolating basis such as generalized B-splines. We can therefore represent the sampling PDFs in a very compact yet complete manner, i. e., covering the full directional variation as well as the variation of the BRDF’s shape parameter (in case of analytic ones). The main advantages of this approach are:

- **Generality:** no reliance on any specific assumptions about the form of the represented BRDFs, with the exception that they are smooth functions.
- **Versatility:** the PMM model can be configured to suit the needs of different target applications, by using different component kernels and their number in the mixture. We evaluate this by fitting two mixture models with different degrees of expressiveness: Gaussians and skewed Gaussians.
- **Uniformity:** the same sampling interface be used for a range of different BRDFs by supplying different meta-parameters. This simplifies implementation, requiring only a single sampling routine to be optimized and maintained.

Compared to the current state of the art (Sec. 2), our sampling framework yields convincingly better BRDF sampling, both in terms of overall sampling accuracy and the number of valid samples (Sec. 6). In addition, our approach:

- supports the sampling of spatially varying BRDFs thanks to the smooth manifold representation in the meta-parametric space;
- allows easy inclusion of additional terms (most prominently the cosine of the incident angle);
- provides an approximation of the material’s directional hemispherical reflectance (i. e., directional albedo);
- allows sampling according to the full illumination integral in PMM-based product sampling frameworks such as Herholz et al. [HEV*16].

Currently, the main shortcoming of our method is its limitation to isotropic BRDFs. However, this is only the property of our current fitting methodology (Sec. 4), and we firmly believe an extension to anisotropic BRDFs is only a quantitative issue. Even in its current form, our framework provides superior sampling for a range of both analytic BRDFs (Phong, Ward, Torrance-Sparrow using GGX and Beckmann microfacet distributions [PJH16]) and measured ones (from the MERL database [MPBM03]), as we show in Fig. 1 and demonstrate in Sec. 6 both numerically and in rendered images.

2. Related Work

Material modelling. For image synthesis purposes, material appearance is predominantly modeled by BRDFs, potentially enriched by spatially varying components (SVBRDFs) or further generalized to high-dimensional texturing functions (BTFs). The number of works in this area is particularly extensive, hence we refer the interested reader to the excellent overview by Dorsey et al. [DRS08]. For physically based analytical models (which are our main target), surface interaction is most commonly captured by the Torrance-Sparrow model [TS67, CT82] with the Beckmann or more recently the GGX function [WMLT07] applied to model the micro-facet distribution. A well-known application from the industrial context is the Disney “principled” BRDF [Bur12], among others. Finally, most recent progress in this direction is well represented by the works of Heitz et al. [Hd14, HHdD16, HDHN16], as well as Ribardière et al. [MRS17] who considered a bi-variate micro-facet distribution for the first time.

For BRDF importance sampling, Walter et al. [WMLT07] provide a great overview and analysis of available methods. They also proposed an improved approach by sampling the normal distribution function (NDF) weighted by the cosine between the half vector and normal itself. Heitz and d’Eon [Hd14] further increase the sampling efficiency by only sampling the distribution of *visible* normals, decreasing the incidence of invalid samples on the bottom hemisphere. These methods usually achieve very good results, but also demonstrate the difficulty of sampling the entire BRDF model by jointly considering each term, which only gets worse with increasing complexity of the model. Our approach circumvents this issue by treating the BRDF *holistically*, therefore not caring about its underlying complexity but only fitting to its resulting shape. Lawrence et al. [LRR04] achieve comparable results based on a factored BRDF representation, but akin to Herholz et al. [HEV*16] they tabulate their fits over some dimensions of the BRDF, which makes the methods impractical for SVBRDFs (due to storage requirements and lack of interpolation in the discretized dimensions).

Model fitting. The main limitation in representing realistic materials in terms of analytic models (be it for direct visualization or sampling) is the frequent lack of expressiveness of the latter to capture the complexity of the former. As Ngan et al. [NDM05] have demonstrated by fitting a selection of analytic BRDFs to a large database of measured materials, it is very challenging to find a single analytical model that is both sufficiently compact and expressive.

Fitting measured or simulated BRDFs (and other functions) with linear models has a long tradition in computer graphics, using piecewise-continuous representations [Kaj85, GMN94], spherical harmonics [CMS87, WAT92], spherical wavelets [SS95] or Zernike polynomials [KDS96], to name a few examples. Among the first to recognize that capturing the qualitative features of BRDFs might be preferable to a numerically optimal fit were Lafor-tune et al. [LFTG97]. In this seminal paper they proposed to use a non-linear fitting model based on normalized cosine lobes, with the underlying argument that linear basis functions are generally agnostic to features of the target BRDFs, such as glossy and off-specular reflections. Albeit with different conditions and requirements, our work applies this approach to obtaining sampling PDFs by extend-

ing the fitting to the higher-dimensional manifold space of model meta-parameters.

For directly representing glossy SVBRDFs in static scenes, Wang et al. [WRG*09] efficiently fit sums of *spherical Gaussians* to the dominant parts of the model; Xu et al. [XSD*13] generalize the method to anisotropic Gaussians and distributions thereof. On a related note, Heitz et al. [HHdD16] employ non-linear optimization for fitting a generalized, transformed cosine lobe to model glossy reflections of area light sources. This representation is unimodal, however, and so it remains an open question how to efficiently fit a mixture model of such transformed cosine lobes. In the context of volumetric light transport, Jakob et al. [JRJ11] use expectation-maximization to fit photon distributions for approximate rendering.

We point out that fitting a BRDF for *direct representation* versus *sampling* conforms to very distinct criteria. For the former, the representation needs to preserve salient perceptual features of the BRDF, while for sampling, the focus is on global numerical accuracy. This is because even small deviations from the actual BRDF can cause significant variance spikes, especially in those parts of its domain where the BRDF values are small. Our work therefore focuses on establishing a robust framework for the sampling case.

Light transport sampling. Several works [LW95, Jen95, HP02], and more recently Vorba et al. [VKŠ*14] and Dahm and Keller [DK17], have shown that including an approximation of the (indirect) incident illumination helps to guide the transport paths to important regions in the scene. This significantly decreases the variance of Monte Carlo solvers, especially for difficult scenes dominated by complex indirect transport. Herholz et al. [HEV*16] followed up on Vorba et al.’s concept, by using a Gaussian mixture model (GMM) representation of the BRDF. Through the calculation of the product of the GMMs for the BRDF and the incident illumination, they were able to sample an approximate representation of the product integrand. However, similar to Lawrence et al. [LRR04] they rely on an impractical, densely tabulated representation of the GMM parameters, which is one of the shortcomings we address.

3. Background

This section introduces the two key concepts underlying our method: we first generally elaborate the optimal ‘local’ BRDF sampling PDF in the Monte Carlo context (Sec. 3.1), followed by a summary of parametric mixture models which we use to represent the actual BRDF-specific PDFs (Sec. 3.2).

3.1. Rendering Equation and BRDF Sampling

The canonical problem of rendering—synthesizing physically-plausible and realistic images—is governed by the rendering equation [Kaj86]. The key challenge here is to evaluate the surface radiance for a given point \mathbf{x} reflected into the direction ω_o from all directions on the upper hemisphere Ω^+ :

$$L_r(\mathbf{x}, \omega_o) = \int_{\Omega^+} L_i(\mathbf{x}, \omega_i) f_r(\mathbf{x}, \omega_o, \omega_i) \cos \theta_i d\omega_i. \quad (1)$$

In the Monte Carlo context, sampling exactly in proportion to the integrand would yield a perfect zero-variance solution [Hoo08, KW08].

Since the integrand itself is the unknown solution, one possible approach is to pre-compute an approximation of the incident radiance L_i , and calculate and sample the product with the cosine-weighted BRDF f_r in run-time [VKŠ*14, HEV*16].

But regardless whether the approximation of L_i is available, obtaining a tightly fitting PDF for sampling the BRDF is still essential for minimizing the estimator variance for Eq. 1. The optimal *local* sampling strategy therefore encapsulates all factors of the integrand, beside the unknown radiance L_i , into a single composite PDF:

$$p(\omega_i | \mathbf{x}, \omega_o) \propto f_r(\mathbf{x}, \omega_i, \omega_o) \cos \theta_i. \quad (2)$$

Deriving this optimal PDF is still far from trivial, as virtually all physically-based BRDFs consist of multiple terms (e. g., Fresnel, normal distribution function (NDF), and shadowing term). Common sampling approaches then typically try to identify the dominant term in Eq. 2 (such as the NDF [LRR04, Hd14] or even the cosine term in case of near-diffuse BRDFs) and sample according to it. This in turn introduces additional variance into the solution in proportion to the neglected terms, and can even yield invalid samples (as we elaborate in Sec. 6.1).

Our solution to this problem is to obtain an analytic fit to the entire right side of Eq. 2, in particular using parametric mixture models. Beside the fact that this solution is decoupled from the complexity of the underlying BRDF model, it can also be directly incorporated into the product sampling approaches discussed above.

3.2. Parametric Mixture Models

Parametric mixture models (PMMs) are used in numerous statistical contexts, for instance classification or model representation in machine learning. A PMM consists of a sum of K weighted parametric ‘kernel’ functions and their corresponding parameter sets Θ_k .

$$\mathcal{P}(\mathbf{y} | \Theta) = \sum_{k=1}^K \pi_k \cdot \mathcal{K}(\mathbf{y} | \Theta_k) \quad (3)$$

$$\Theta = \{\Theta_1, \dots, \Theta_k\}. \quad (4)$$

The weights π_k define the contribution of each component to the mixture. Based on the underlying kernel function the resulting mixture model can have a variety of properties (sampling, integration, product calculation, etc.) suitable for different applications. For example, if the used kernel function is a probability density function and the weights of the mixture fulfill the constraints that $\pi_k \geq 0$ and $\sum \pi_k = 1$, the resulting mixture also represents a valid PDF.

Gaussian mixture model (GMM). One of the most common PMMs is the Gaussian mixture model (GMM), which is based on multivariate normal distributions [Bis06].

$$\mathcal{K}_{\mathcal{G}}(\mathbf{y} | \Theta_k) = \mathcal{N}(\mathbf{y} | \mu_k, \Sigma_k). \quad (5)$$

The parameters for each component of a GMM are the mean μ_k and the covariance matrix Σ_k of each normal distribution.

Skewed Gaussian mixture model (SMM). An extension to the GMM is the skewed Gaussian mixture model (SMM). It replaces the multivariate normal distribution by its skewed counterpart

(SND) [AC99], allowing the kernels to adapt to non-symmetric features of the data or function represented through the SMM:

$$\mathcal{K}_S(\mathbf{y}|\Theta_k) = 2 \cdot \mathcal{N}(\mathbf{y}|\mu_k, \Sigma_k) \cdot \tilde{\mathcal{N}}((\mathbf{y} - \mu_k)^T s_k). \quad (6)$$

The SND is a product of the multivariate normal distribution \mathcal{N} and the cumulative distribution function $\tilde{\mathcal{N}}$ of the standard *univariate* normal distribution. The additional vector parameter $s_k \in (-\infty, \infty)^d$ defines the skewness of the distribution, where d is the dimension of the SND – if null, the SND reduces to the regular multivariate ND. Notably, the skewness of the SND is a desirable when it comes to representing the cosine term and the hemispherical boundary for BRDF lobes at grazing angles (Sec. 6.1).

Due to the ability to efficiently generate samples, fit and represent distributions via expectation maximization or calculating the product of two GMMs, they have been proven to be useful in computer graphics and rendering [WRG*09, JRJ11, XSD*13, VKŠ*14, HEV*16]. Since the skewed normal distribution is based on the normal distribution, it inherits some of its features, such as efficient sample generation or product calculation [AC99].

On the flip side, simple operations like calculating distances [LH00, MS05] or the interpolation between multiple PMMs are not trivial. This is due to the fact that the correspondence between the components of two mixtures is generally unknown, unless given a-priori. Our solution is to define the PMM parameters in a smooth manifold space over the dimensions of the BRDF parameters. This way, we are not only able to achieve a compact BRDF-specific representation of the PMMs, but also overcome the hurdle of interpolating two PMM representations between different (continuously varying) sets of BRDF parameters.

4. Unified BRDF Sampling Framework

The primary motivation for this work has been to provide a unified sampling framework specifically designed for BRDFs (according to Eq. 2), but without requiring any knowledge about the complexity of each individual BRDF. We especially aim to make the design of new material appearance models (that is, BRDFs) easier by removing the need to manually design sampling functions for them. The secondary objective was for the resulting representation to be useful in a wider array of sampling tasks, e. g., product importance sampling [HEV*16], sampling mixtures of BRDFs, or sampling spatially varying BRDFs.

Our solution lies in the realization that probabilistically sampling a BRDF conforms to different optimality criteria than their direct visualization. Most prominently, sampling error is *relative*, since even a small BRDF value can cause visible variance spikes if sampled with a still significantly smaller PDF. On top of a generalized parametric mixture model (PMM) as a representation for the PDF, we utilize a loss function that reflects the relative nature of the sampling error. Then, to overcome the issues with interpolating between instances of PMM (cf. Sec. 3.2), we perform the fitting *jointly* over all dimensions of the BRDF, and define all its meta-parameters on a smooth manifold in the respective domains. This *manifold space* can then be compactly represented by smooth interpolating functions (we use generalized B-splines), which serves as an implicit regularizer during the fitting process, and at the same time provides the benefit of continuous queries in the meta-parametric domain.

Later in Sec. 6, we provide and evaluate PDF fits for a range of common analytic BRDFs (Phong, Ward, Torrance-Sparrow using GGX and Beckmann microfacet distributions), as well as measured ones (from the MERL database). We also compare two different PMM models: Gaussians and skewed Gaussians (which are particularly suitable for representing asymmetric distributions). While our framework can potentially handle anisotropic BRDFs (thanks to the generalized spline representation of the manifold surfaces), due to technical limitations of our fitting we focus on isotropic BRDFs (details in Sec. 4.3). We demonstrate that our PDFs yield numerically and visually better sampling than existing BRDF-specific PDFs (Secs. 6.1 and 6.2), and are useful for other sampling tasks.

The following exposition has three stages with increasing levels of abstraction:

- We start in Sec. 4.1 by detailing how to represent and sample PDFs using generalized PMMs.
- Then in Sec. 4.2 we describe the manifold space for representing the PMM's meta-parameters, and explain how this is modeled via generalized multi-dimensional B-splines.
- Finally in Sec. 4.3 we detail the joint fitting procedure, including the optimal objective function which minimizes the PDF's sampling error.

4.1. BRDF Representation

Common analytic BRDFs usually consist of a diffuse and a glossy (directionally dependent) component: we focus on sampling the latter, since the former is trivial. For measured BRDFs we make no such distinction. In either case, to decouple our representation from the complexity of the underlying BRDF, we model its sampling PDF using general PMMs (see illustration in Fig. 2):

$$\mathcal{P}(\omega_i|\Theta_q) \approx f_r(q, \omega_i) \cdot \cos\theta_i. \quad (7)$$

Here, for any parameter set q of a given BRDF, our framework proposes a PMM *meta-parametrization* Θ_q that then represents the corresponding sampling PDF over the domain of incident directions ω_i (cf. Eq. 1). Given our assumption of isotropic BRDFs, we will hereinafter work with $q = \{\omega_o, \alpha\}$, where α stands for the shape parameter of typical BRDF models (such as the microfacet roughness). For measured BRDFs the shape parameter is not necessary.

Following Herholz et al. [HEV*16], we use the disk mapping of the hemi-spherical domain of the BRDF to the $[-1, 1]^2$ domain. This converts the space, in which the PMM from Eq. 7 is defined,




Figure 2: An optimal sampling PDF (cosine-weighted BRDF) and its corresponding PMM representation with meta-parameters Θ_q .

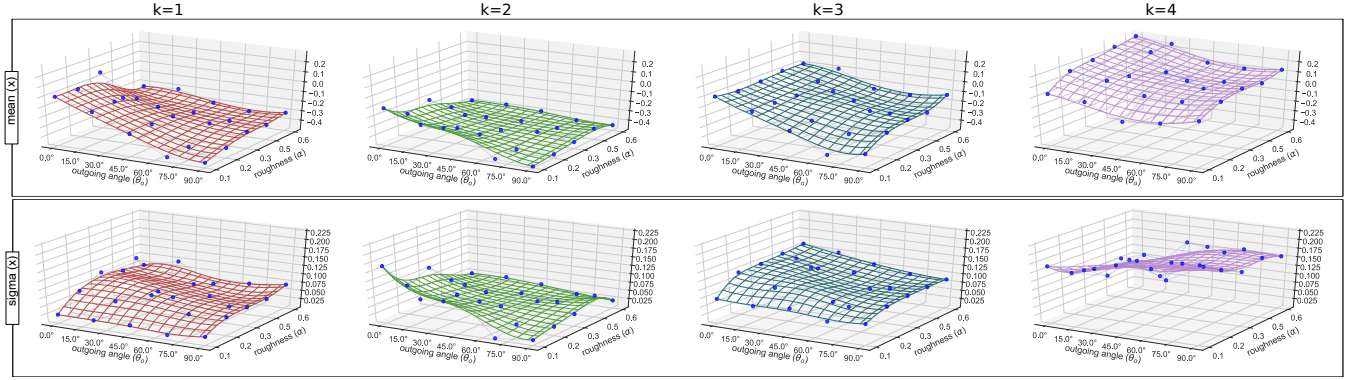


Figure 3: Visualization of the B-spline manifolds for a subset of the GMM meta-parameters (μ_x and σ_x) for a rough conductor BRDF using the GGX NDF, with $K = 4$ components. The blue points visualize the control points of the B-splines coefficient vectors for the PMM parameter.

from the (hemi-)sphere to a two-dimensional XY-space. Similarly, since we are only considering isotropic BRDFs (which are invariant to azimuthal rotations) we can reduce the dimension of the outgoing direction ω_o to its zenith angle θ_o and adjust the PMM representation later, by rotating it to the desired ϕ_o orientation. The same fact allows us to reduce the PMM meta-parameterization to $\Theta_k = \{\mu_x, \sigma_x, \sigma_y\}_k$ for a GMMs and $\Theta_k = \{\mu_x, \sigma_x, \sigma_y, s_x\}_k$ for SMMs (cf. Sec. 3.2).

For importance-sampling a BRDF according to Eq. 2, we normalize the weights of the PMM and use the corresponding sampling methods of the used mixture (which are defined analytically for both GMMs and SMMs, see Apx. A). Conceptually, any PMM capable of representing hemispherical functions and having an efficient sampling procedure can be used instead of Gaussian-based ones.

Compared to existing works, our representation:

- models the optimal sampling PDF including the cosine factor (not only the BRDF or even only some of its terms), and
- the PMM represents the *un-normalized* version of Eq. 2.

This has the advantage that it allows accurately estimating the hemispherical directional reflectance (or ‘directional albedo’ hereinafter) of a BRDF parametrization by simply summing the PMM weights:

$$\bar{f}(q) = \int f_i(q, \omega_i) \cdot \cos \theta_i d\omega_i \approx \sum_{k=1}^K \pi_k \quad (8)$$

We take advantage of this property in Sec. 4.3, and later in Sec. 6.3 we show its benefit for sampling BRDF compositions.

4.2. Manifold Parametric Mixture Model Representation

An overwhelming majority of BRDFs exhibits a good degree of smoothness – for instance analytic BRDFs are typically defined by products of smooth functions (e. g., NDF, Fresnel, geometry term). Consequently, they change smoothly in respect to their non-directional parameters too (such as the shape parameter α). This implies that the meta-parameters of the representing model should likewise follow this behavior, even in the case of a mixture model.

Building up on this assumption the valid PMM meta-parameter sets used to represent a complete BRDF model would lie on a

set of smooth manifolds (one for each actual PMM parameter $\{\pi^k, \mu_x^k, \sigma_x^k, \dots\}$) – in the PMM’s meta-parameter domain with the dimensions of $R^{|q|}$ (i. e., same as the BRDF parameters q). With a K -component PMM, and a kernel function with N parameters, we need $K(N + 1)$ manifolds to represent a complete BRDF model.

Thanks to the assumed smoothness we are able to represent these manifolds by an interpolation basis, such as generalized B-splines. For the example of a GMM:

$$\{\pi_k, \Theta_k\} = \{B_{C_\pi^k}(q), B_{C_\mu^k}(q), B_{C_\Sigma^k}(q)\}. \quad (9)$$

The coefficient vectors C_π^k , C_μ^k and C_Σ^k define the B-spline manifold space of the k -th component of the GMM. Each vector contains $M^{|q|}$ control points regularly distributed in the BRDF parameter subset q . In this work we use $M = 5$, which in the case of an analytic isotropic BRDF model with $q = \{\theta_o, \alpha\}$ results in a coefficient vector with the modest size of 5×5 for each PMM meta-parameter. The order of the B-splines used to represent the manifolds is 3 for each dimension, which has proven to be flexible enough to cover the PMM meta-parameter variation. For completeness we provide the formulas for evaluating a generalized B-spline in Appendix B.

To visualize the smoothness of the meta-parameter manifolds, Fig. 3 shows an example of such a B-Spline representation for an isotropic rough conductor BRDF using the GGX normal distribution. We now follow by describing how the coefficients for the B-Spline manifold representation can be obtained using non-linear optimization. Readers interested more in the practical application of this framework might skip to Sec. 5.

4.3. PMM Representation Fitting

Fitting the manifold representation of the PMM meta-parameters faces two major challenges:

- the PMM for a given BRDF parameter set needs to be optimal for importance-sampling the BRDF according to Eq. 2, and
- the PMM meta-parameters for a given BRDF model should span a smooth manifold (i. e., should not change abruptly), in order to be well represented by generalized B-splines.

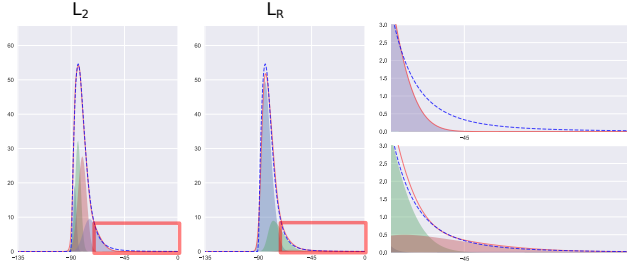


Figure 4: *Left:* the effect of different residual functions on the resulting parametric mixture, which is used for importance-sampling the BRDF. While the \mathcal{L}_2 residual ensures a good fit in high-valued regions, it handles areas with lower values with less importance (top), compared to our relative loss \mathcal{L}_r (bottom).

Relative loss function. The task of the loss function is to quantify the quality (induced error) of a proposed PMM meta-parameter set Θ_q for a given BRDF parameter q . During the optimization the loss is calculated for N different directional samples ω_n of Eq. 2 to compare them to the evaluation of the proposed PMM. For simplicity we label the BRDF samples as $d_n = f_r(q, \omega_n) \cos \theta_n$ and the model evaluations as $m_n = \mathcal{P}(\omega_n | \Theta_q)$.

Since the resulting PMM is used by a Monte Carlo estimator to sample proportionally to Eq. 2, we design a relative loss function \mathcal{L}_r which minimizes the variance of this estimator:

$$\mathcal{L}_r(d_n, m_n) = \left[\frac{d_n - m_n}{m_n + \epsilon} \right]^2. \quad (10)$$

In contrast to the commonly-used squared error loss function (\mathcal{L}_2) our loss considers the *relative* error between the BRDF and the PMM representation of its PDF. This mainly ensures that the loss function strongly penalizes significant under-sampling, in regions of the BRDF where a small absolute difference in its PDF could cause hyperbolic increase of variance. A small epsilon $\epsilon = 0.001$ prevents numerical instabilities.

Fig. 4 shows the difference between the \mathcal{L}_2 loss and our relative loss \mathcal{L}_r . The \mathcal{L}_2 loss concentrates on fitting the high-valued peak of the BRDF but neglects the tail(s), which are equally critical for a well-behaved estimator.

Manifold fitting. Instead of starting with individual fits of the PMM meta-parameters over the whole parameter sub-space of a given BRDF model, and then trying to fit the B-spline manifolds to these individual fits, we decided to fit the B-spline coefficients directly. The advantage is that the smoothness of the B-splines acts as an implicit regularizer for the PMM parameters, which ensures their smooth variation. To fit the B-spline coefficients $C_{\{\tau_1, \dots, s\}}$ for a given PMM model we use the following objective function:

$$\arg \min_{C_{\pi, \dots, C_s}} \sum_{m=0}^M \left[\sum_{n=0}^N \mathcal{L}_r(f_r(q_m, \omega_n) \cos \theta_n, \mathcal{P}(\omega_n | \Theta_{q_m})) \right] + R(\bar{f}_r(q_m), \Theta_{q_m}) \quad (11)$$

where the spline coefficients $C_{\{\tau_1, \dots, s\}}$ are encapsulated in the evaluation of Θ_{q_m} . To solve this optimization problem the BRDF parameter space is discretized into M parameter sets $\{q_1, \dots, q_M\}$. For each




Figure 5: *Left:* visualization of the progressive fitting process using 4 iterations $\{I_1, \dots, I_4\}$. *Right:* example for the three sample sets (cf. Sec. 4.3) generated for a microfacet BRDF model for an outgoing angle $\theta_0 = \pi/4$ and roughness $\alpha = 0.3$.

set of N samples we then calculate the loss, and add an additional regularization term ensuring that the fitted PMM can be used to approximate the directional albedo \bar{f}_r for each BRDF parametrization:

$$R(\bar{f}_r(q), \Theta_q) = \left(N \cdot \frac{\sum \pi \tau_k}{\bar{f}_r(q)} \right)^2. \quad (12)$$

For uni-variate BRDF models such the MERL database [MPBM03], the M parameter sets are distributed over a 1D space of U different outgoing angles ($M = U$). For an analytical isotropic BRDF where the BRDF parameter sub-space is 2D ($q = \{\theta_0, \alpha\}$) the M parameter sets cover this space with U angles and V roughness values ($M = U \cdot V$). For all our fits we used the same number of discretization steps $U = V = 63$ (see Fig. 5, left).

Progressive fitting. To fit the B-spline coefficients we start with the parameter set of size M in the center of the BRDF parameter space (e. g., $\alpha = 0.3$ and $\theta_0 = \pi/4$), then progressively extend the parameter range considered in the evaluated objective function.

Fig. 5, left illustrates this procedure. In the first iteration I_1 only the central set is considered in the B-spline fitting. In the following iterations (I_2 to I_4) the range increases until the full BRDF parameter space is covered. We found that small increments of this range (about 20 iterations) allows the formation of smooth manifolds.

Sample generation. To evaluate the overall error of a proposed PMM meta-parametrization Θ_{q_m} the objective function (Eq. 11) evaluates the loss for N different samples of Eq. 2. There are several criteria for these samples to ensure good fit. They need to:

- represent the shape of the main BRDF lobe,
- cover the area outside the main lobe on the upper hemisphere, and finally
- ensure that the fitted PMM is minimized across the lower hemisphere.

To achieve this, three different batches of samples are generated. The first batch samples 1024 directions using the standard importance-sampling method of the BRDF, the second are 512 samples uniformly distributed over the upper hemisphere, and finally the third batch consists of a belt of 512 samples around the lower hemisphere

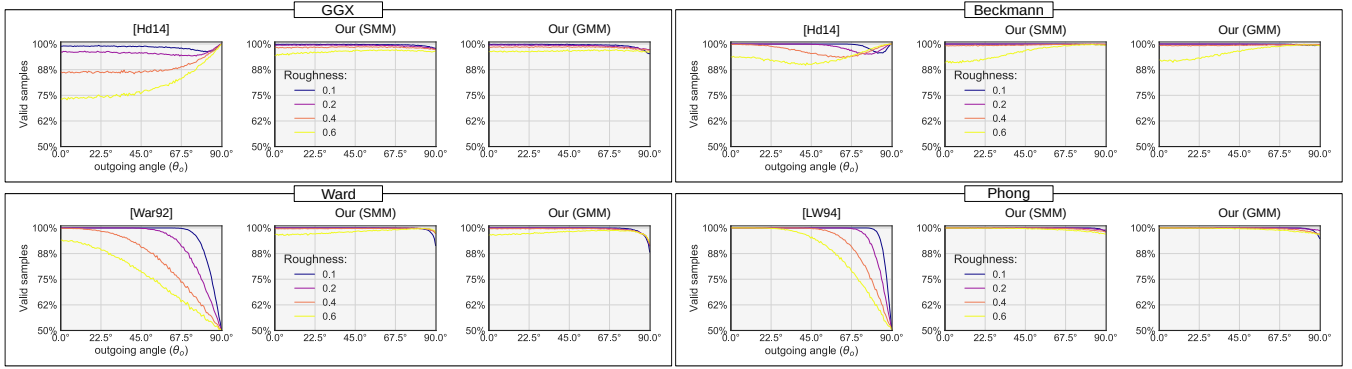


Figure 6: Comparing the percentages of valid directional samples for different sample generation methods: respectively, **left:** the analytic state-of-the-art sampling, **center and right:** our SMM and GMM based sampling.

boundary. Fig. 5, **right** gives an example of such a sample set generated for a microfacet BRDF using the Beckmann NDF ($\alpha = 0.3$).

5. Implementation

We integrated our unified PMM-based sampling framework in the scientific Mitsuba renderer [Jak10]. Since our algorithm only replaces the sampling and PDF methods of a BRDF, it can be implemented using the standard interfaces, while leaving the actual integrator code (e. g., the path tracer we use) untouched. The PMM-based BRDF sampling code, as well as the resulting fits for our experiments (Sec. 6) will be made public on the project website.

The fitting of the B-spline surfaces, which represent the PMM meta-parameter manifolds, is done in a pre-processing step the Ceres [AMO16] optimization package. Since these manifolds contain all the information needed to generate a PDF for each fitted BRDF, we only need to store one set of manifolds per model. Using $K = 4$ GMM/SMM components, these occupy only about 2kB of storage for the analytic BRDFs, and even less for the MERL [MPBM03] BRDFs (as they have one dimension less). To translate the exponent of the Phong model to a roughness value $\alpha \in [0, 1]$ we use the conversion method described by Walter et al. [WMLT07]. From the MERL database we fitted and tested four different materials: *blue metallic*, *nickel*, *nylon* and *purple paint*.

The overhead of our sampling compared to the state-of-the-art methods varied from 0% for the Beckmann model (i. e., same cost as [Hd14]) to about 100% for Ward [War92]. However, since the directional sampling as such is only a small component in the path generation process, the average *total overhead* was only 4% for a path segment in our renderer evaluation.

6. Results

We evaluate our BRDF sampling framework in the perspective of its use in stochastic Monte Carlo light transport algorithms. In Sec. 6.1 we start by numerically analyzing the quality of produced samples against state-of-the-art sampling methods. Then in Sec. 6.2 we demonstrate the positive impact of our sampling on the rendering performance, when used in a complete rendering pipeline. Finally, Sec. 6.3 showcases additional useful applications of our framework.

6.1. Numerical Sampling Evaluation

We evaluate the sampling quality using two metrics: the proportion of valid samples produced, and the overall error of the sampling. The comparison is done against Phong, Ward, and GGX- and Beckmann-based microfacet BRDFs using their established sampling routines [War92, LW94, WMLT07, Hd14].

Sample validity. We define a sample as invalid, if generated by the sampling PDF in a null region of the associated BRDF. This most prominently happens near the hemispherical boundary, i. e., at grazing angles. Such samples need to be avoided, since they can dramatically increase the solution variance in renderers with throughput-based path termination strategies.

Fig. 6 visualizes the percentage of valid samples generated from the standard sampling methods and our PMM-based methods (using both GMM and SMM for fitting) at different roughness values $\alpha \in \{0.1, 0.2, 0.4, 0.6\}$ and for outgoing angles $\theta_o \in [0, \pi/2]$. Thanks to our dedicated fitting strategy which accounts for the hemispherical boundary (Sec. 4.3) we obtain superior rates across the board. Also observe that the skewness of the SMM adapts to the hemispherical boundary better than GMM. While the microfacet model sampling of Heitz et al. [Hd14] achieves slightly higher validity rates at near-boundary angles (since they prevent the invalid samples for these configurations explicitly) we still achieve better results overall.

Relative estimator variance. To evaluate our overall sampling accuracy we computed the estimate of the directional albedo \bar{f}_r (defined in Sec. 4.1) as it is the canonical integrated quantity in our problem. We use the *relative estimator variance* \bar{V} as a metric, which itself is conceptually equivalent to our relative loss function (Sec. 4.3):

$$\bar{V}(X) = \mathbb{E} \left[\left(\frac{X - \mathbb{E}[X]}{\mathbb{E}[X]} \right)^2 \right]. \quad (13)$$

Where X is a set of random samples estimating \bar{f}_r , which are proposed by the different sampling strategies, and $\mathbb{E}[X]$ is equal to the ground-truth \bar{f}_r . We use 4k samples to estimate Eq. 13, and 16k samples generated with the standard sampling method to obtain the ground-truth \bar{f}_r .




Figure 7: Comparing the normalized MC variance of estimating the directional albedo for different BRDFs, using our SMM-based sampling versus the standard methods.




Figure 8: Similar comparison as in Fig. 7, focusing on the difference between SMM and GMM fits respectively. Please note the different error scales in the two comparisons.

Fig. 7 and Fig. 8 evaluate the relative variances for the complete 2D space of BRDF parameters q (outgoing angles and roughness values). In Fig. 7 we see that our sampling yields dramatically lower variances than the standard sampling methods, again with the singular exception of extreme grazing angles in comparison with Heitz et al. [Hd14]. In Fig. 8 the SMM/GMM comparison leads to the same conclusion as before, i. e., that the SMM handles the asymmetric configurations in grazing angles better.

6.2. Renderer Integration

Next, we evaluate our sampling framework when incorporated into Mitsuba’s path tracer with next event estimation (however, being a black-box model it can be used by any path-based integrator).

Fig. 9 shows an evaluation of a selection of BRDFs (microfacet GGX and Ward at two different roughness values, and two materials from the MERL database [MPBM03]) when sampled by our PMM-based framework, compared to the respective standard sampling methods. The sampling for the measured BRDFs was done according to Dupuy et al. [DHI*15] (by first fitting a microfacet BRDF to each measured one and then using its sampling routines). From the false-color variance plots, as well as insets, it is clear that our PMM-based sampling performs on par or better, depending on the configuration. Table 1 presents a comprehensive evaluation using the same scene, where we measure the standard RMS error, as well as proportions of valid samples generated for each configuration.

When analyzing these results, it should be noted that the noise in the renders results from the variance of the *entire estimator*, meaning that even a perfect BRDF sampling (according to Eq. 2) would still retain the variance of the illumination sampling. That then explains the discrepancy between the performance gains measured in the numerical analysis (e. g., in Fig. 7) and the rendered results.

6.3. Additional Benefits

Spatially varying BRDFs. Since our manifold PMM representation covers a continuous range of the BRDF parameters, we can seamlessly apply it to SVBRDFs (e. g., with a spatial variance of the roughness parameter α). In Fig. 1 we show a comparison for a GGX microfacet BRDF, using a texture to define the roughness variation.

Table 1: RMS errors and proportions of valid samples for rendered scenes in the same set of configurations as analyzed in Sec. 6.1.

	$\alpha = 0.1$	$\alpha = 0.2$	$\alpha = 0.4$	$\alpha = 0.6$
GGX	rMSE VS	rMSE VS	rMSE VS	rMSE VS
SMM	.0448 99.7	.0454 99.3	.0382 98.4	.0313 96.4
GMM	.0449 99.6	.0456 99.2	.0382 98.4	.0311 96.6
[Hd14]	.0445 98.5	.0484 95.5	.0465 87.3	.0414 78.0
Beck.				
SMM	.0335 99.9	.0398 99.9	.0422 99.4	.0402 95.8
GMM	.0336 99.9	.0401 99.8	.0423 99.3	.0397 95.9
[Hd14]	.0337 99.7	.0404 99.1	.0439 96.5	.0423 92.2
Ward				
SMM	.0365 99.9	.0405 99.9	.0383 99.7	.0332 97.9
GMM	.0365 99.9	.0407 99.8	.0385 99.5	.0331 97.9
[War92]	.0379 99.0	.0444 96.6	.0464 88.5	.0429 78.5
Phong				
SMM	.0243 99.9	.0296 99.9	.0352 99.8	.0381 99.5
GMM	.0241 99.9	.0296 99.9	.0351 99.8	.0384 99.4
[LW94]	.0244 99.7	.0304 99.0	.0389 96.3	.0462 91.7

Given that our PMM representation is very compact (cf. Sec. 5), it could be used in conjunction with illumination-based guiding methods [VKŠ*14, HEV*16] for a practical product sampling with SVBRDFs (instead of resorting to expensive tabulated representations as Herholz et al. [HEV*16] had to).

Directional albedo. In Fig. 10 we demonstrate the accuracy of our estimate of the hemispherical directional reflectance (cf. Sec. 4.1), which we can trivially evaluate from the weights of the PMM (Eq. 8). We compare to an MC-estimated ground truth for two different materials under different outgoing angles.

An example where an accurate estimate of the directional albedo is useful is the importance sampling of mixtures of different BRDF models. Fig. 11 demonstrates the sampling of such a mixture model (diffuse plus glossy microfacet BRDF) with a mixture weight of $w = 0.5$. The default approach is using this weight to randomly decide on the sampled component, neglecting its actual contribution. Our accurate estimate of the directional albedo, on the other hand, allows to make this sampling decision optimally.




Figure 9: Comparing our PMM-based framework to standard sampling routines for two analytic BRDFs (microfacet GGX and Ward at two different roughness values), and two materials from the MERL database. The false-color maps visualize the per-pixel relative estimator variance (Sec. 6.1) of each respective render.




Figure 10: Comparing our approximation of the directional albedo \bar{f}_T to the ground truth (gt) for two different BRDF models: Phong and Torrance-Sparrow (GGX) with a roughness of $\alpha = 0.4$.

7. Discussion and Limitations

We believe that the proposed approach has the potential to subsume and unify the existing approaches to BRDF sampling, for two reasons. First, the formulation is general and agnostic to the complexity of the fitted (SV)BRDF, instead only caring about the extrinsic features (which can be captured by different numbers of PMM components). For the same reason, the BRDF can be defined




Figure 11: Using our approximation of the directional albedo for choosing the optimal sampling weight of a composite BRDF (25% diffuse and 75% specular component using a dielectric specular Torrance-Sparrow lobe) yields an improved result compared to static sampling according to the composition weight alone.

both analytically and by measured data. Second, our framework provides a unified access interface, which is potentially very convenient for industrial rendering software where modularity is important for robust and optimized implementation.

Anisotropic BRDFs. Since the used B-spline surfaces generalize to higher dimensions, a natural extension of this work would be fitting and sampling anisotropic BRDF models. Due to the high memory requirements of our optimization package [AMO16], we have been limited to using the two-dimensional manifolds (representing the outgoing angle and the BRDF shape parameter). Examining or developing optimization methods that would scale better with regard to the dimensionality of the joint fitting problem would therefore be a meaningful question to investigate.

Product importance sampling. As already discussed, our sampling PDFs could be used in a product sampling framework for a *globally optimal* sampling of the full illumination integral. We plan to investigate this possibility, for instance by integrating our method into the GMM-based product sampling of Herholz et al. [HEV*16].

BRDF representation and compression. The PMM framework is robust enough to be considered as a replacement for BRDFs themselves, leading to an out-of-the-box perfect BRDF sampling. An additional benefit would be in compressing measured BRDFs and SVBRDFs, instead of fitting specialized microfacet models akin to Dupuy et al. [DHI*15] and hence having to adopt additional assumptions about the measured materials. Since our fitting is designed for sampling optimality (especially the loss function), a step towards that direction would be using perceptually oriented metrics for such direct representation.

Choice of PMM. As we demonstrated by using both GMMs and SMMs, the choice of the particular PMM is not instrumental in obtaining high-quality sampling. Still, potential improvements could be gained by considering different mixture kernels (e. g., von-Mises-Fisher or Kent distributions), mainly to satisfy different application-specific criteria.

Manifolds and perception. Having a general BRDF representation defined over a smooth manifold in the meta-parameter space could be conducive to better perceptual characterization of the model [SGM*16]. Examples include linearization of the model and more intuitive navigation in the parametric space of the BRDF.

8. Conclusion

We have presented a unified manifold framework for sampling isotropic BRDFs. The framework is based on parametric mixture models such as Gaussian mixtures and skewed Gaussian mixtures. Our system enables querying a compact PMM representation of a full BRDF model with continuously varying parameters, which can then be used for efficient importance sampling.

In contrast to the current state-of-the-art analytic sampling solutions for physically-based BRDFs, our framework considers all components of the BRDF. Since this full representation leads to sampling PDFs proportional to the full BRDF, we have shown a positive impact on the efficiency of rendering algorithms based on Monte Carlo integration. Our framework can be implemented either as a new PDF model, or as a black-box wrapper around existing models, and therefore be easily integrated in production rendering systems such as Arnold, Manuka, or Corona.

Acknowledgments. This work was supported by the German Research Foundation (DFG): SFB 1233, Robust Vision: Inference Principles and Neural Mechanisms, TP 2. It received further funding from the European Union's Horizon 2020 research and innovation program, under the Marie Skłodowska-Curie grant agreement No 642841 (DISTRO) and was supported by the Czech Science Foundation grant 16-18964S and the Charles University grant SVV-2017-260452.

References

- [AC99] AZZALINI A., CAPITANIO A.: Statistical applications of the multivariate skew normal distribution. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 61, 3 (1999), 579–602. 4, 11
- [AMO16] AGARWAL S., MIERLE K., OTHERS: Ceres solver. <http://ceres-solver.org>, 2016. 7, 10
- [Bis06] BISHOP C. M.: *Pattern recognition and machine learning*. Springer, 2006. 3
- [BM58] BOX G. E. P., MULLER M. E.: A note on the generation of random normal deviates. *The Annals of Mathematical Statistics* 29, 2 (1958), 610–611. 11
- [Bur12] BURLEY B.: Physically-based shading at Disney, 2012. 2nd version. URL: http://disney-animation.s3.amazonaws.com/library/s2012_pbs_disney_brdf_notes_v2.pdf. 1, 2
- [CMS87] CABRAL B., MAX N., SPRINGMEYER R.: Bidirectional reflection functions from surface bump maps. *SIGGRAPH Comput. Graph.* 21, 4 (1987), 273–81. 2
- [CT82] COOK R. L., TORRANCE K. E.: A reflectance model for computer graphics. *ACM Trans. Graph.* 1 (1982), 7–24. 2
- [DHI*15] DUPUY J., HEITZ E., IEHL J.-C., POULIN P., OSTROMOUKHOV V.: Extracting microfacet-based BRDF parameters from arbitrary materials with power iterations. *Computer Graphics Forum* (2015), 10. 8, 10
- [DK17] DAHM K., KELLER A.: Learning light transport the reinforced way. *CoRR abs/1701.07403* (2017). 3
- [DRS08] DORSEY J., RUSHMEIER H., SILLION F.: *Digital modeling of material appearance*. Morgan Kaufmann, 2008. 2
- [FWKH17] FONG J., WRENNINGE M., KULLA C., HABEL R.: Production volume rendering: Siggraph 2017 course. In *ACM SIGGRAPH 2017 Courses* (New York, NY, USA, 2017), SIGGRAPH '17, ACM, pp. 2:1–2:79. 1
- [GMN94] GONDEK J. S., MEYER G. W., NEWMAN J. G.: Wavelength dependent reflectance functions. In *Proc. of SIGGRAPH* (1994), pp. 213–20. 2
- [Hd14] HEITZ E., D'EON E.: Importance sampling microfacet-based BSDFs using the distribution of visible normals. *Computer Graphics Forum* 33, 4 (jul 2014), 103–112. 1, 2, 3, 7, 8
- [HDHN16] HEITZ E., DUPUY J., HILL S., NEUBELT D.: Real-time polygonal-light shading with linearly transformed cosines. *ACM Transactions on Graphics* 35, 4 (jul 2016), 1–8. 2
- [HEV*16] HERHOLZ S., ELEK O., VORBA J., LENSCH H., KŘIVÁNEK J.: Product importance sampling for light transport path guiding. *Computer Graphics Forum* 35 (2016). 2, 3, 4, 8, 10
- [HHdD16] HEITZ E., HANIKA J., D'EON E., DACHSBACHER C.: Multiple-scattering microfacet BSDFs with the Smith model. *ACM Trans. Graph.* 35, 4 (July 2016), 58:1–58:14. 2, 3
- [Hoo08] HOOGENBOOM J. E.: Zero-variance Monte Carlo schemes revisited. *Nuclear Science and Engineering* 160, 1 (sep 2008), 1–22. 3
- [HP02] HEY H., PURGATHOFER W.: Importance sampling with hemispherical particle footprints. In *Proc. of SCCG* (2002), pp. 107–14. 3

- [Jak10] JAKOB W.: Mitsuba renderer, 2010. <http://www.mitsuba-renderer.org>. 7
- [Jar13] JAROSZ W.: The perils of evolutionary rendering research: Beyond the point sample. Eurographics Symposium on Rendering, Invited Talk, 2013. 1
- [Jen95] JENSEN H. W.: Importance driven path tracing using the photon map. In *Rendering Techniques (proc. EGWR)* (1995). 3
- [JRJ11] JAKOB W., REGG C., JAROSZ W.: Progressive expectation-maximization for hierarchical volumetric photon mapping. *Computer Graphics Forum* 30, 4 (2011), 1287–97. 3, 4
- [Kaj85] KAJIYA J. T.: Anisotropic reflection models. *SIGGRAPH Comput. Graph.* 19, 3 (1985), 15–21. 2
- [Kaj86] KAJIYA J. T.: The rendering equation. *SIGGRAPH Comput. Graph.* 20 (1986). 1, 3
- [KDS96] KOENDERINK J. J., DOORN A. J. v., STAVRIDIS M.: Bidirectional reflection distribution function expressed in terms of surface scattering modes. In *Proc. of ECCV* (1996), pp. 28–39. 2
- [KW08] KALOS M. H., WHITLOCK P. A.: *Monte Carlo Methods*. Wiley-VCH, 2008. 3
- [LFTG97] LAFORTUNE E. P. F., FOO S.-C., TORRANCE K. E., GREENBERG D. P.: Non-linear approximation of reflectance functions. In *Proc. of SIGGRAPH* (1997). 2
- [LH00] LIU Z., HUANG Q.: A new distance measure for probability distribution function of mixture type. In *International Conference on Acoustics, Speech, and Signal Processing* (2000), Institute of Electrical and Electronics Engineers (IEEE). 4
- [LRR04] LAWRENCE J., RUSINKIEWICZ S., RAMAMOORTHY R.: Efficient BRDF importance sampling using a factored representation. *ACM Trans. Graph.* 23, 3 (2004), 496–505. 2, 3
- [LW94] LAFORTUNE E. P., WILLEMS Y. D.: *Using the modified Phong reflectance model for physically based rendering*. Report CW 197, Departement Computerwetenschappen, KU Leuven, Celestijnenlaan 200A, 3001 Heverlee, Belgium, November 1994. 7, 8
- [LW95] LAFORTUNE E. P., WILLEMS Y. D.: A 5D tree to reduce the variance of Monte Carlo ray tracing. In *Rendering Techniques* (1995), pp. 11–20. 3
- [MPBM03] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: A data-driven reflectance model. *ACM Transactions on Graphics* 22, 3 (July 2003), 759–769. 2, 6, 7, 8
- [MRS17] M. RIBARDIERE B., BRINGIER D. M., SIMONOT L.: STD: Student's t-distribution of slopes for microfacet based BSDFs. In *Proc. of EG* (2017). 2
- [MS05] MAYBECK P., SMITH B.: Multiple model tracker based on Gaussian mixture reduction for maneuvering targets in clutter. In *2005 7th International Conference on Information Fusion* (2005), Institute of Electrical and Electronics Engineers (IEEE). 4
- [NDM05] NGAN A., DURAND F., MATUSIK W.: Experimental analysis of BRDF models. In *Proc. of EGSR* (2005). 2
- [PJH16] PHARR M., JAKOB W., HUMPHREYS G.: *Physically based rendering: From theory to implementation*, 3rd ed. Morgan Kaufmann, 2016. 1, 2
- [SGM*16] SERRANO A., GUTIERREZ D., MYSZKOWSKI K., SEIDEL H.-P., MASIA B.: An intuitive control space for material appearance. *ACM Transactions on Graphics (SIGGRAPH ASIA 2016)* 35, 6 (2016). 10
- [SS95] SCHRÖDER P., SWELDENS W.: Spherical wavelets: efficiently representing functions on the sphere. In *Proc. of SIGGRAPH* (1995), pp. 161–172. 2
- [TS67] TORRANCE K. E., SPARROW E. M.: Theory for off-specular reflection from roughened surfaces. *J. Opt. Soc. Am.* 57, 9 (1967), 1105–14. 2
- [Vea97] VEACH E.: *Robust Monte Carlo methods for light transport simulation*. PhD thesis, Stanford University, 1997. 1
- [VKŠ*14] VORBA J., KARLÍK O., ŠÍK M., RITSCHER T., KŘIVÁNEK J.: On-line learning of parametric mixture models for light transport simulation. *ACM Trans. Graph. (Proc. of SIGGRAPH)* 33 (2014). 3, 4, 8
- [War92] WARD G. J.: Measuring and modeling anisotropic reflection. *SIGGRAPH Comput. Graph.* 26, 2 (July 1992), 265–272. 7, 8
- [WAT92] WESTIN S. H., ARVO J. R., TORRANCE K. E.: Predicting reflectance functions from complex surfaces. *SIGGRAPH Comput. Graph.* 26, 2 (1992), 255–64. 2
- [WMLT07] WALTER B., MARSCHNER S. R., LI H., TORRANCE K. E.: Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques* (Aire-la-Ville, Switzerland, Switzerland, 2007), EGSR'07, Eurographics Association, pp. 195–206. 1, 2, 7
- [WRG*09] WANG J., REN P., GONG M., SNYDER J., GUO B.: All-frequency rendering of dynamic, spatially-varying reflectance. *TOG* 28, 5 (dec 2009), 1. 3, 4
- [XSD*13] XU K., SUN W.-L., DONG Z., ZHAO D.-Y., WU R.-D., HU S.-M.: Anisotropic spherical Gaussians. *TOG* 32, 6 (nov 2013), 1–11. 3, 4

Appendix A: PMM Importance Sampling

Generating a sample x_i with a PDF, proportional to a PMM model, with normalized weights ($\sum \pi_k = 1$), is a two step process. First, a component k of the PMM is picked, using a random variable ζ_0 and the weights of the PMM, so that:

$$\sum_{i=1}^k \pi_i < \zeta_0 < \sum_{i=1}^{k+1} \pi_i. \quad (14)$$

After a component has been selected, the sample x_i is generated by importance sampling the corresponding ‘kernel’ function. The following paragraph presents the formulas for importance sampling the bi-variate Gaussian ‘kernel’ function used by our GMM to represent the optimal sampling lobe of the BRDF models. A similar analytic sampling method also exists for the bi-variate skew Gaussian distribution used by our SMMs. A detailed explanation of this method and also other interesting features (such as the product calculation) is given by Azzalini and Capitanio [AC99].

Bi-variate Gaussian importance sampling. Importance sampling a bi-variate normal distribution is done, by first sampling the unit bi-variate normal distribution using the Box Muller transform [BM58], followed by transformation to the desired normal distribution. Therefore, two uniform random variables ζ_1 and ζ_2 are need. Alternatively ζ_0 , which was used to select the current PMM components, can be re-normalized and then re-used as ζ_2 :

$$\zeta_2 = \frac{\zeta_0 - \sum_{i=1}^{i-1} \pi_i}{\pi_i}. \quad (15)$$

The Box Muller transformation generates a random sample z_i with a PDF of $p(z_i) = \mathcal{N}(z_i, 0, I)$ by generating the disk coordinates r and the ϕ of z_i from ζ_1 and ζ_2 :

$$r = \sqrt{-2 \ln \zeta_1} \quad (16)$$

$$\phi = \pi \zeta_2 \quad (17)$$

$$z_i = [r \sin(\phi), r \cos(\phi)] \quad (18)$$

To transform z_i into x_i so that $p(x_i) = \mathcal{N}(x_i | \mu_k, \Sigma)$, only the mean μ_k and the lower triangular matrix A_k defined as $\Sigma_k = A_k A_k^T$ are needed: $x_i = A_k z_i + \mu_k$. Methods like the Cholesky decomposition can be used to calculate the lower triangular matrix A_k .

Appendix B: B-Spline Surfaces

B-spline surfaces are an extension of the B-spline curves to the 2D domain. They are defined via a regular set of $(n + 1) \times (m + 1)$ control points, stored inside a coefficient vector C , and two ‘knot’ vectors $U = \{u_0, \dots, u_n + p\}$ and $V = \{v_0, \dots, v_m + q\}$. The orders p and q define the respective degrees of continuity in each dimension. To evaluate a B-spline surface at given coordinates u and v , the support of each control point C_{ij} is defined by the product of the B-spline curves for each dimension. Since the support of a control point depends on the orders p and q , only a subset of control points needs to be evaluated for a given u and v :

$$\begin{aligned}
 B_C(u, v) &= \sum_{i=0}^n \sum_{j=0}^m C_{ij} B_{i,p}(u) B_{j,q}(v) \\
 B_{i,0} &:= \begin{cases} 1, & \text{if } t_i \geq u < t_{i+1} \\ 0, & \text{otherwise} \end{cases} \\
 B_{i,p}(u) &:= \frac{u - u_i}{u_{i+p} - u_i} B_{i,p-1}(u) + \frac{u_{i+k+1} - u}{u_{i+k+1} - u_{i+1}} B_{i+1,p-1}(u).
 \end{aligned} \tag{19}$$