

Metalog Materials

HIROYUKI SAKAI*, Technische Universität Wien (TU Wien), Austria

SAIP-CAN HASBAY*, Independent Researcher, Austria

CHRISTIAN FREUDE, Technische Universität Wien (TU Wien), Austria

MICHAEL WIMMER, Technische Universität Wien (TU Wien), Austria

DAVID HAHN, Technische Universität Wien (TU Wien), Austria




Fig. 1. We showcase metalog materials exhibiting a diverse range of reflectance behaviors, including an **analytic** Trowbridge–Reitz (TR) material, **measured** materials from the RGL material database [Dupuy and Jakob 2018], and materials acquired via **differentiable rendering**. Our representation naturally supports **interpolation** between materials, enabling applications such as spatially varying bidirectional reflectance distribution functions (SVBRDFs). The plots on the left and right show bivariate metalogs for selected incident directions of the indicated materials.

A central aspect of modern physically based rendering is the practical yet realistic modeling of material appearance, which requires efficient and accurate representations of bidirectional reflectance distribution functions (BRDFs). In this paper, we present a unified material reflectance representation based on metalog distributions. Introducing quantile-based metalog interpolation and metalog mixture models, our representation captures a wide range of complex reflectance behaviors of both analytic models and real-world materials, while remaining interpretable and efficient. Our approach achieves competitive runtime performance and a small memory footprint, while enabling continuous interpolation between fundamentally different reflectance behaviors. Furthermore, our model is fully differentiable with respect to the coefficients defining the distributions, thus easily integrates into existing differentiable rendering pipelines. In summary, metalog materials can effectively represent real-world reflectance data, interpolate between multiple materials within a unified representation, and be acquired from real-world images via differentiable rendering.

*Joint first authors

Authors' Contact Information: Hiroyuki Sakai, sakai@cg.tuwien.ac.at, Rendering and Modeling Group, Technische Universität Wien (TU Wien), Vienna, Austria; Saip-Can Hasbay, saipcanhasbay@gmail.com, Independent Researcher, Vienna, Austria; Christian Freude, freude@cg.tuwien.ac.at, Rendering and Modeling Group, Technische Universität Wien (TU Wien), Vienna, Austria; Michael Wimmer, wimmer@cg.tuwien.ac.at, Rendering and Modeling Group, Technische Universität Wien (TU Wien), Vienna, Austria; David Hahn, david.hahn@tuwien.ac.at, Rendering and Modeling Group, Technische Universität Wien (TU Wien), Vienna, Austria.



This work is licensed under a Creative Commons Attribution 4.0 International License.
© 2026 Copyright held by the owner/author(s).
ACM 1557-7368/2026/12-ART182
<https://doi.org/10.1145/3842563>

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

Additional Key Words and Phrases: material modeling, BRDF representation, importance sampling, Monte Carlo rendering, physically based rendering, metalog distributions

ACM Reference Format:

Hiroyuki Sakai, Saip-Can Hasbay, Christian Freude, Michael Wimmer, and David Hahn. 2026. Metalog Materials. *ACM Trans. Graph.* 45, 6, Article 182 (December 2026), 13 pages. <https://doi.org/10.1145/3842563>

1 Introduction

Physically based rendering is crucial in modern computer graphics for accurately simulating real-world light–material interactions. A key challenge in this field is efficiently and realistically representing materials' optical properties through bidirectional reflectance distribution functions (BRDFs). Despite extensive research, no unified framework exists; current methods are often domain-specific, such as analytic models for parametric materials or data-driven techniques for measured ones. These disparate approaches require distinct implementations and thus lack interoperability across applications. To overcome these limitations, we introduce a novel material representation building on metalog distributions, initially proposed by Keelin [2016]. Metalogs are continuous probability distributions known for their flexibility but, to our knowledge, have previously been unexplored in the rendering community.

Building on these distributions constitutes the core innovation of our approach. By leveraging their unique properties, we introduce smooth metalog interpolation via optimal transport and formulate metalog mixture models. Integrating linear least-squares (LLSQ)

fitting into an expectation–maximization (EM) algorithm enables joint estimation of the mixture coefficients and component distributions. With these novel contributions, we significantly expand the expressivity of multivariate metalog models, facilitating their application to multidimensional scenarios involving highly multimodal distributions with diverse shapes. In particular, we create a unified, flexible, and compact representation of material reflectance, capable of modeling a wide range of reflectance behaviors.

In summary, we present the following contributions:

- A material reflectance model built on metalog distributions, therefore consisting of interpretable coefficients that enable efficient storage, sampling, and editing.
- Continuous interpolation between materials with fundamentally different reflectance behaviors.
- Effective compression of measured materials, leveraging the compactness and shape flexibility of metalogs, while achieving competitive runtime performance.
- Our material model is differentiable with respect to the distribution coefficients and therefore seamlessly integrates into differentiable rendering pipelines, enabling optimization of material parameters to match real-world images.

Taken together, these contributions make our metalog materials a powerful foundation for material modeling, appearance representation, and manipulation. We provide supplementary materials, such as the source code, at <https://www.cg.tuwien.ac.at/MM>.

2 Related Work

In this section, we provide an overview of material models in rendering and review prior work on unifying these models. Following the taxonomy introduced by Guarnera et al. [2016], we categorize existing bidirectional reflectance distribution functions (BRDFs) into three major groups: phenomenological models, physically based models, and data-driven models. We then review material representation and sampling methods before introducing metalogs.

2.1 BRDF Models

Phenomenological models describe reflectance using *analytic functions* that are either artist-designed or loosely fitted to observed materials. Early glossy BRDFs include the Blinn–Phong model [Blinn 1977; Phong 1975] and the (anisotropic) Ward [1992] model, offering efficient Monte Carlo (MC) sampling. These nonphysical models often lose energy, especially at grazing angles [Guarnera et al. 2016]. More recent work includes simplified Blinn–Phong variants [Neumann et al. 1999] and extensions to the Ward model [Dür 2006; Geisler-Moroder and Dür 2010].

Physically based models aim to describe material reflectance according to physical principles, primarily conservation of energy. Early work includes self-shadowing [Smith 1967], the Oren–Nayar [1994] model (and its energy-preserving variant [Portsmouth et al. 2025]), which remain widely used for their physical realism, as well as microfacet BRDFs [Blinn 1977; Cook and Torrance 1982; Torrance and Sparrow 1967] for metals and dielectrics, supporting a broad range of roughness and view-dependent color changes. Walter et al. [2007] used what is now known as the Trowbridge–Reitz (TR) distribution [1975] as “GGX” in microfacet theory, including

transmission and Smith’s shadowing–masking function. Extending microfacet models for specific materials or applications is still an active field of research [Bagher et al. 2012; Hill et al. 2025; Jakob et al. 2014; Kurt et al. 2010; Lee et al. 2018; Löw et al. 2012; Rosales et al. 2023; Turquin 2019; Xie and Hanrahan 2018]. The Disney BRDF [Burley 2012] builds on the generalized Trowbridge–Reitz (GTR) distribution, further popularizing TR-based microfacet models through artist-friendly controls. Heitz et al. unify masking–shadowing with visible–normal sampling and extend microfacet theory to multiple scattering and volumetric regimes [Heitz 2014; Heitz and d’Eon 2014; Heitz et al. 2015, 2016]. Dupuy et al. [2015] reconstruct microfacet BRDFs from measured reflectance data via eigenanalysis.

Data-driven models often rely on separable approximations, expressing the 4D BRDF as a sum of products of lower-dimensional terms. Early approaches used singular value decomposition (SVD) [Press et al. 2007], including Fournier’s [1995] approximation of the Blinn–Phong model from Ward’s [1992] measurements. While accurate, SVD is expensive and may yield negative factors [Guarnera et al. 2016]; normalized decomposition provides a cheaper, non-optimal alternative [Kautz and McCool 1999]. Other decompositions have also been investigated [McCool et al. 2001; Rusinkiewicz 1998]. A major milestone was Matusik et al.’s [2003] BRDF database. Their complementary wavelet analysis also identifies necessary sampling frequencies. Further research has focused on compact and editable representations, often employing wavelets or PCA linearization [Bilgili et al. 2011; Claustres et al. 2004; Lawrence et al. 2006; Sun et al. 2007]. While convenient for editing [Serrano et al. 2016], PCA imposes a global linear subspace, which may limit expressivity for data-driven reflectance models. Dupuy and Jakob [2018] generalize microfacet fitting and present a comprehensive database of real-world material measurements. Fitting our metalog models to their data achieves comparable visual fidelity and runtime performance at much lower memory consumption.

Recently, data-driven material models have also employed neural networks [Kuznetsov et al. 2021; Sztrajman et al. 2021; Zeltner et al. 2024; Zhou et al. 2026] providing compression for offline rendering and fast inference for real-time applications. A key advantage of our representation is that reflectance and sampling behavior are intrinsically linked through metalog distributions, allowing both BRDF evaluation and importance sampling without intermediate representations. Compared to the neural BRDF (NBRDF) approach by Sztrajman et al. [2021], which maps back to parameters of an analytic BRDF model for importance sampling, our metalog approach achieves lower variance as demonstrated in §5.4.

While analytic models often provide editable and interpretable parameters, data-driven or neural models cover a much wider range of materials, but are often tedious to manipulate or fine-tune. Generally, the main objectives for BRDF models are to represent a wide range of materials (flexibility), enable artistic expression (editability), and reduce memory consumption and runtime cost (storage and sampling efficiency) [Hill et al. 2025; Weidlich 2025].

2.2 General Representations and Sampling

The manifold sampling framework of Herholz et al. [2018] employs parametric mixture models (PMMs) to construct efficient importance samplers for analytic and measured BRDFs. We pursue similar goals of unified representation and sampling, but rely on metalog distributions instead. For material interpolation, Herholz et al. [2018] provide smoothly varying PMMs within the parameter space of an underlying material model, enabling efficient sampling while interpolating material parameters. Our representation, on the other hand, operates directly on the resulting reflectance distributions, providing a lightweight interpolation scheme using metalogs.

Recent efforts, such as OpenPBR Surface [Andersson et al. 2025], also seek to provide a standardized material representation and parameterization across rendering systems. Rather than relying on a shared parameter space, our approach operates directly on the resulting reflectance behavior and thus allows consistent interpretation across material models regardless of their underlying parameterization. Our model-agnostic approach increases the generality of the framework and, unlike the parameter-space manifolds of Herholz et al. [2018], naturally enables interpolation *across all materials* represented within our framework.

Concurrently, Colom et al. [2026] developed a model-agnostic representation of BRDFs based on Gaussian-process regression to obtain compact reflectance models without assuming a particular material model. The resulting representation accurately reconstructs reflectance behavior. For importance sampling, they convert their model to spherical harmonics and sample using the technique of Jarosz et al. [2009]. Our work is also related to learned importance sampling approaches such as neural importance sampling by Müller et al. [2019], which employs normalizing-flow-based neural networks to learn sampling distributions for MC rendering. Our non-neural approach instead provides a representation with small memory footprint, exposes interpretable coefficients, and supports both evaluation and importance sampling.

2.3 Metalog Distributions

The metalog distribution family, introduced by Keelin [2016; 2023], emerged from efforts to create a flexible, analytically tractable alternative to traditional parametric distributions in statistical modeling [Keelin and Powley 2011]. Via smooth quantile functions, metalogs capture diverse distribution shapes, including skewness, tail behavior, and multimodality, and extend naturally to multivariate marginal-conditional distributions. To our knowledge, these distributions have not been used in the rendering literature. Our metalog materials closely approximate most microfacet distributions, including TR, and are capable of capturing a wide variety of physical reflectance behaviors. They thus provide an alternative basis for BRDFs, combining efficient storage and sampling with interpolation and differentiability.

3 Background and Notation

Light transport is typically formalized by the rendering equation [Kajiya 1986]:

$$L_o(\mathbf{x}, \omega_o) = L_e(\mathbf{x}, \omega_o) + \int_{\Omega} f(\mathbf{x}, \omega_i, \omega_o) L(\mathbf{x}, \omega_i) |\mathbf{n} \cdot \omega_i| d\omega_i. \quad (1)$$

Here, we focus on the case of surface-based reflection, so Ω is the hemisphere at point $\mathbf{x}$ with surface normal $\mathbf{n}$. Combining emission L_e and reflection of incident radiance L from ω_i via the bidirectional reflectance distribution function (BRDF) f produces the outgoing radiance L_o in direction ω_o . Due to its recursive structure and high dimensionality, rendering commonly relies on Monte Carlo (MC) integration. A MC estimator of Eq. (1) is

$$\hat{L}_o(\mathbf{x}, \omega_o) = L_e(\mathbf{x}, \omega_o) + \frac{1}{n} \sum_{i=1}^n \frac{f(\mathbf{x}, \omega_i, \omega_o) \hat{L}(\mathbf{x}, \omega_i) |\mathbf{n} \cdot \omega_i|}{p(\omega_i | \mathbf{x}, \omega_o)}, \quad (2)$$

where $\hat{\cdot}$ denotes an estimator, ω_i refers here to importance-sampled directions drawn from the probability density function (PDF) p , and n is the number of samples. In practice, p is chosen using a priori information, such as the BRDF or the distribution of light sources. For BRDF sampling, the main challenge is to find a distribution p that is proportional to f and can be efficiently sampled and evaluated.

To this end, we build on Keelin’s [2016] *metalog distributions* (“metalogs”) to construct BRDF representations. In the following, we summarize their definition and most relevant properties, deferring details on our material models to §4. Metalogs are defined via their quantile function as a linear combination of t basis functions:

$$M_t^u(y | \mathbf{a}) = \mathbf{a} \cdot \boldsymbol{\phi}(y), \quad (3)$$

where $\mathbf{a}$ is a vector of coefficients and $\boldsymbol{\phi}(y)$ is a vector of basis functions of the same size defined as $\boldsymbol{\phi}(y) = (1, \ln(y/(1-y)), (y-0.5)\ln(y/(1-y)), (y-0.5), (y-0.5)^{(l-1)/2} \text{ for odd } l \geq 5, (y-0.5)^{l/2-1}\ln(y/(1-y)) \text{ for even } l \geq 6)^T$ with $l \leq t$ denoting the basis-function index. Here, $y \in (0, 1)$ denotes the cumulative probability and the quantile function $x = M_t^u(y)$ returns the corresponding quantile. When $\mathbf{a}$ is fixed or clear from context, we omit it for brevity. Note that the corresponding cumulative distribution function (CDF) is the inverse of the quantile function, $P = (M_t^u)^{-1}$.

The *metalog probability density function* (PDF) $m_t^u(y)$ as a function of the cumulative probability y can be derived in closed-form using $m_t^u(y) = (dM_t^u/dy)^{-1}$; we provide Keelin’s result in the supplementary document for completeness. This result makes metalog-based BRDFs particularly convenient for importance sampling: a sample is generated via inverse transform sampling using the quantile function $M_t^u(y)$, where y is uniformly distributed, while the corresponding PDF value is obtained directly by evaluating $m_t^u(y)$. Conversely, evaluating the PDF at a specific value x requires first numerically solving $M_t^u(y) = x$ for y , before evaluating $\hat{m}_t^u(x) = m_t^u((M_t^u)^{-1}(x))$. This root-finding problem is well-suited for a 1D Newton–Raphson method due to the monotonicity of the CDF. In the context of metalog-based BRDFs, this case arises for example when using multiple importance sampling (MIS).

Metalogs provide high shape flexibility, allowing the representation of varying degrees of skewness, tailedness, and (to some extent) multimodality. Coefficients can also be intuitively interpreted: a_1 corresponds to the median, while the remaining coefficients control the distribution’s shape; in particular, a_2 , a_3 , and a_4 are measures of scale, skewness, and kurtosis [Keelin 2016]. A larger number of coefficients t yields a more flexible distribution. Clearly, the coefficient vector can be zero-padded to a larger size when needed, without affecting the distribution. Figure 2 demonstrates the expressivity of a metalog fitted to heavy-tailed bimodal data.




Fig. 2. A single metalog distribution represents highly non-Gaussian data with multiple modes: we fit a 10-term metalog to a Trowbridge–Reitz (TR) distribution with anisotropic roughness parameters $\alpha_u = 0.2$ and $\alpha_v = 0.3$.

Metalog Feasibility. Not all coefficient vectors $\mathbf{a}$ yield a valid, i.e., *feasible*, metalog distribution. Feasibility requires the metalog PDF $m_t^u(y)$ to be strictly positive on $(0, 1)$, or equivalently, that the quantile function $M_t^u(y)$ is strictly increasing. Formally, we require $dM_t^u(y)/dy > 0$ in $(0, 1)$, leading to the *feasibility condition* on $\mathbf{a}$:

$$\frac{a_2}{y(1-y)} + a_3 \left(\frac{y-0.5}{y(1-y)} + \ln \left(\frac{y}{1-y} \right) \right) + a_4 + \dots > 0, \quad \forall y \in (0, 1). \quad (4)$$

We use this condition to detect and prevent infeasible metalog parameterizations in our framework.

Bounded Metalogs. In their original form, Eq. (3), metalogs are unbounded in their value range. In this paper, we aim to represent directional distributions, for instance parameterized by elevation and azimuth angles $\theta \in [0, \pi/2]$, $\phi \in [-\pi, \pi]$. For such settings, Keelin introduced *bounded metalogs*, whose range is restricted to given bounds by transforming the original quantile function to

$$M_t(y) = \frac{b_l + b_u e^{M_t^u(y)}}{1 + e^{M_t^u(y)}}, \quad (5)$$

where b_l and b_u denote the lower and upper bounds. Note that this transformation preserves monotonicity, and therefore feasibility, so Eq. (4) still holds. The PDF of a bounded metalog similarly follows from the unbounded PDF $m_t^u(y)$ (see the supplementary document):

$$m_t(y) = m_t^u(y) \frac{(1 + e^{M_t^u(y)})^2}{(b_u - b_l) e^{M_t^u(y)}}. \quad (6)$$

Multivariate Metalogs. Keelin [2023] provides a marginal–conditional construction for extending metalogs to higher dimensions, where the CDF of a multivariate distribution P is expressed as the product of a marginal distribution CDF on the first variable P_1 and subsequent CDFs $P_2, \dots, P_k$ conditioned on prior outcomes:

$$P(x_1, x_2, \dots, x_k) = P_1(x_1) P_2(x_2 | x_1) \dots P_k(x_k | x_1, \dots, x_{k-1}),$$

where the 1D CDF components are inverse metalog quantile functions. This construction ensures that the result is a proper probability distribution and preserves the sampling efficiency of 1D metalogs. The multivariate distribution can be sampled by sequentially sampling each conditional component. We also employ this construction for our BRDFs, but crucially add interpolation between multiple conditional metalog “slices” to extend the expressive capabilities of our material models.

Parameter Estimation. Because the (unbounded) metalog quantile function is linear in its coefficients, estimating these parameters from input CDF data can be done simply by linear least-squares (LLSQ) fitting. This parameter estimation method can be interpreted as the maximum likelihood estimator assuming a Gaussian noise model [Keelin 2016]. For the bounded case, one first inverse-transforms the observed data and fits the unbounded distribution. Applying the transformation, Eq. (5), then yields the bounded distribution that fits the original data. We refer to Keelin [2023] for details and provide the explicit formula and derivations in the supplementary document.

Mixture Models. Beyond fitting a specific distribution to observed data, mixture models provide a general framework for representing more complicated, multimodal probability densities as weighted sums of simpler component distributions [Bishop and Bishop 2024]. Existing methods have primarily focused on mixtures of Gaussians. When fitting a Gaussian mixture model to actual data, the mixture coefficients, component means, and covariance matrices are typically estimated by maximizing the data likelihood using the expectation–maximization (EM) algorithm. It alternates between computing responsibilities, which quantify the probability that each mixture component generated a given data point (E-step), and reestimating the mixture coefficients, means, and covariances using these responsibilities (M-step), yielding a locally optimal maximum likelihood solution. In the following sections, we adopt the same approach, but replace the Gaussian components with metalogs. This results in metalog mixture models, which retain the flexibility of classical mixtures and further improve expressivity through the increased shape flexibility of metalog distributions.

4 Our Approach

Having established the necessary background, we now present a framework for representing material models in rendering using metalogs. To make metalog-based materials practical, we first address the issues of metalog interpolation and constructing metalog mixtures in general, before focusing on material modeling.

The main building block of our metalog BRDF is the probability distribution p such that

$$f(\omega_i, \omega_o) \cos(\theta_o) \propto p(\omega_o | \omega_i), \quad \omega_o = (\theta_o, \phi_o). \quad (7)$$

Note that in the following, we adopt the convention of Dupuy and Jakob [2018], considering the incident light direction ω_i as given, and the outgoing direction ω_o as sampled, which corresponds to real-world light transport and therefore measured data acquisition. Of course, in a path-tracing context, sampling actually constructs camera-to-light paths in the reverse direction.

We define this sampling PDF $p(\omega_o | \omega_i)$ implicitly as the probability density associated with the CDF $P(\omega_o | \omega_i)$:

$$P(\omega_o | \omega_i) = P(\theta_o, \phi_o) = P^m(\theta_o) P^c(\phi_o | \theta_o), \quad (8)$$

following a marginal–conditional construction. Here, θ_o may be cosine-transformed to compensate for the non-uniform area distortion across the hemisphere and θ_o may be interchanged with ϕ_o in practice. In the most basic case, we could define the marginal and conditional distributions via metalog quantile functions. In order

to expand the expressive range of our distribution, however, we instead use 1D metalog mixtures of the general form:

$$P_M(x | \{\pi_k\}, \{\mathbf{a}_k\}) = \sum_k^K \pi_k M_t^{-1}(x | \mathbf{a}_k), \quad (9)$$

where $\{\pi_k \geq 0\}$ is a sequence of K mixture coefficients ($\sum_k^K \pi_k = 1$), and $\{\mathbf{a}_k\}$ is a sequence of metalog coefficient vectors (each of length t) for the individual components. Furthermore, instead of a single conditional distribution, we use multiple parallel “slices” at fixed locations along the marginal axis, see Fig. 7, and interpolate between them based on their relative distance $\lambda_i(\theta_o)$. Finally, we define the marginal and conditional distributions:

$$\begin{aligned} P^m(\theta_o) &= P_M(\theta_o | \{\pi_k^m\}, \{\mathbf{a}_k^m\}), \\ P^c(\phi_o | \theta_o) &= \text{interp}(\lambda_i(\theta_o), \\ &\quad P_M(\phi_o | \{\pi_{k,i}^c\}, \{\mathbf{a}_{k,i}^c\}), P_M(\phi_o | \{\pi_{k,i+1}^c\}, \{\mathbf{a}_{k,i+1}^c\})). \end{aligned} \quad (10)$$

Our results, Fig. 13, show that extending the construction to multiple slices significantly enhances the expressive range of multivariate metalogs. Furthermore, for fitting the metalog BRDF to given reflectance data, we need to determine the coefficients of all metalog distributions $\{\mathbf{a}_k^m, \mathbf{a}_{k,i}^c\}$ together with the mixture coefficients $\{\pi_k^m, \pi_{k,i}^c\}$. As both interpolation and mixture construction are problems that apply to metalog distributions independently from material modeling, we first address these two issues in a general framework (§4.1) before turning to our BRDF formulation in §4.2.

4.1 Generalized Metalog Framework

Here, we first define *metalog interpolation*, which enables smooth transitions between conditional slices and also between different materials. We then address fitting our *multivariate metalog mixtures* to represent continuous multidimensional data. In the supplementary document, we provide details on the fitting process and on how to evaluate the metalog CDF M_t^{-1} and associated PDF $\hat{m}_t$.

4.1.1 Metalog Interpolation. While linearly blending PDFs, i.e., averaging values at fixed locations, is a common way to interpolate distributions, doing so typically introduces artificial multimodality when interpolating between unimodal distributions. Instead, we use *quantile interpolation*, which moves probability mass continuously rather than blending densities. In one dimension, this corresponds to interpolating quantile functions (i.e., inverse CDFs), so that each quantile shifts smoothly from its initial to its final location. In the unbounded case, this construction coincides with Wasserstein optimal transport, meaning the transformation minimizes total displacement of probability mass [Peyré and Cuturi 2019]. As a result, salient structural features, such as unimodality, are preserved, and intermediate distributions show intuitive, structure-preserving deformations.

Recall that the unbounded metalog quantile function is linear in its coefficients: $M_t^u(y | \mathbf{a}) = \mathbf{a} \cdot \boldsymbol{\phi}(y)$. Therefore, interpolating two metalogs with coefficient vectors $\mathbf{b}$ and $\mathbf{c}$

$$M_t^u(y | \mathbf{a}) = \lambda M_t^u(y | \mathbf{b}) + (1 - \lambda) M_t^u(y | \mathbf{c}), \quad (11)$$

results in the quantile function defined by the coefficient vector

$$\mathbf{a} = \lambda \mathbf{b} + (1 - \lambda) \mathbf{c}, \quad \lambda \in [0, 1], \quad (12)$$

i.e., quantile interpolation follows directly from linear interpolation of the metalog coefficients. Instead of directly interpolating *bounded*




Fig. 3. Linearly interpolating the metalog coefficients produces smooth intermediate metalog distributions, illustrated here by the resulting PDFs for unbounded (left) and bounded (right) metalogs.

metalogs as in Eq. (11), which would be complicated by the non-linear transformation in Eq. (5), we also define their interpolation via the coefficients $\mathbf{a}$, Eq. (12), as

$$\text{interp}(\lambda, M_t(y, \mathbf{b}), M_t(y, \mathbf{c})) = M_t(y, \mathbf{a}). \quad (13)$$

We thus interpolate their underlying unbounded quantile functions before transforming to bounds, performing optimal transport in the distance metric of the unbounded space while preserving smooth probability-mass transport. Finally, for metalog mixture interpolation, we sort the mixture components by their median and interpolate the corresponding mixture coefficients. The interpolated conditional in Eq. (10) therefore becomes

$$P^c(\phi_o | \theta_o) = P_M(\phi_o | \{\hat{\pi}_k^c\}, \{\mathbf{a}_k^c\}),$$

where all $\mathbf{a}_k^c$ are interpolated from adjacent slices $\mathbf{a}_{k,i}^c, \mathbf{a}_{k,i+1}^c$ based on $\lambda_i(\theta_o)$ following Eq. (12) and analogously $\hat{\pi}_k^c = \lambda_i \pi_{k,i} + (1 - \lambda_i) \pi_{k,i+1}$.

Feasibility of Interpolated Metalogs. For interpolation to be robust, it is essential that interpolating between two feasible metalog distributions yields another feasible metalog. Recall that the feasibility condition, Eq. (4), applies to both unbounded and bounded metalog distributions and is linear in the coefficients. Therefore, (linearly) interpolating the coefficients of two feasible metalogs yields another feasible metalog. Figure 3 illustrates interpolation results for both unbounded and bounded metalogs, demonstrating smooth and feasible intermediate distributions across λ .

4.1.2 Metalog Mixture Models. While metalogs exhibit considerable shape flexibility, highly complex distributions with multiple modes or pronounced skewness and kurtosis may lie beyond the expressive power of the standard framework. In practice, we observed a limit of approximately three modes, largely independent of the number of coefficients. To overcome this limitation and achieve effectively unbounded expressiveness, we introduce *metalog mixtures* in this section, building upon the well-established concept of parametric mixture models (PMMs). Our metalog mixture models allow us to fit highly multimodal data, as shown in Fig. 4. The PDF of a metalog mixture, Eq. (9), is $p(x | \{\pi_k\}, \{\mathbf{a}_k\}) = \sum_k \pi_k \hat{m}_t(x | \mathbf{a}_k)$. We estimate the mixture parameters $\{\pi_k\}, \{\mathbf{a}_k\}$ by maximizing the data log-likelihood using the common EM algorithm. Given observed data $X = \{x_n\}$, the log-likelihood is

$$\ln p(X | \{\pi_k\}, \{\mathbf{a}_k\}) = \sum_n \ln p(x_n | \{\pi_k\}, \{\mathbf{a}_k\}). \quad (14)$$




Fig. 4. Fitting luminance data over elevation and azimuth (in the data’s directional parameter space) for two materials from the RGL database [Dupuy and Jakob 2018]: insets report the Wasserstein-1 distance to the empirical distribution. For an equal number of components, metalog mixtures are noticeably more expressive than Gaussian and von Mises mixtures, even when using just three metalog coefficients (the other component distributions use two parameters).

In the E-step, we compute the responsibilities γ_k for all components k and samples x_n :

$$\gamma_k(x_n) = \frac{\pi_k \hat{m}_t(x_n | \mathbf{a}_k)}{\sum_l \pi_l \hat{m}_t(x_n | \mathbf{a}_l)}. \quad (15)$$

In the M-step, we then update the mixture coefficients π_k by averaging the responsibilities for each component $\pi_k = \sum_n \gamma_k(x_n) / |X|$, where $|X|$ is the number of data points, and the metalog coefficients $\mathbf{a}_k$ with a LLSQ fit, solving $(\mathbf{Y}^T \mathbf{Y} + \Lambda) \mathbf{a}_k = \mathbf{Y}^T \hat{\mathbf{q}}_k$. Here, Λ is a diagonal regularizer, and the design matrix $\mathbf{Y}$ contains the metalog basis function values $\phi(y_j)$ for a preselected sequence of representative quantile levels $\{y_j\}$. Using normalized weights $w_{nk} = \gamma_k(x_n) / (\sum_i \gamma_k(x_i))$, we compute the empirical CDF value for each data point $P_k(x_n) = \sum_{i=1}^n w_{ik}$. From the sorted input data $\{x_n\}$ and their corresponding empirical probabilities $P_k(x_n)$, we then build the right-hand side vector $\mathbf{q}_k$ by interpolating the data values at the preselected quantiles, i.e., considering piecewise linear $P_k(x)$, we set $\mathbf{q}_k = (x_j : P_k(x_j) = y_j)^T$. For bounded metalogs (with bounds b_l, b_u), we transform the quantiles component-wise to $\hat{\mathbf{q}}_k = g(\mathbf{q}_k)$, where $g(x) = \ln((x - b_l)/(b_u - x))$, see also Keelin [2016]. (In the unbounded case $\hat{\mathbf{q}}_k = \mathbf{q}_k$.) Because this LLSQ fitting corresponds to the maximum likelihood estimate of the metalog parameters [Keelin 2016], the EM approach maximizes Eq. (14) (up to regularization). For multivariate mixtures, the LLSQ fitting applies independently to the marginal and conditional slices, updating their respective coefficient vectors sequentially. Please also refer to the supplementary document (§S2) for details on multivariate metalog mixture fitting.




Fig. 5. Overview of our metalog material fitting and rendering pipeline.

Sampling Metalog Mixtures. In the 1D case, the mixture CDF is $P(x) = \sum_k \pi_k M_t^{-1}(x | \mathbf{a}_k)$. We can sample from this mixture efficiently using the standard two-stage mixture sampling strategy: first, choosing a random component according to $\{\pi_k\}$, and then drawing a sample from the corresponding metalog via its quantile function. In the multivariate case, we represent a joint distribution by a marginal distribution for the first variable and a sequence of nested conditional distributions for the remaining variables. In particular, our material models rely on bivariate metalog mixtures, Eq. (10). This hierarchical structure naturally induces a sequential sampling procedure: drawing the first variable from the marginal metalog distribution and then successively evaluating the corresponding conditional metalog quantiles. Each step is driven by an independent uniform random variable, enabling direct and efficient sampling from the full multivariate distribution. Additionally, with this random variable, we can directly evaluate $m_t(y)$ in closed form (no further root finding required in this case).

4.2 Metalog Materials

Having described metalog interpolation and mixtures, we continue with building our reflectance model, which is capable of representing a wide range of surface scattering behaviors in a unified BRDF framework, as demonstrated in Fig. 1. Specifically, we first address *model construction* by fitting reflectance data from either analytic BRDFs or measured real-world materials. This fitting is carried out offline as a one-time precomputation, resulting in reusable material models that can be employed across multiple renderings. We then address *rendering* of our metalog materials, where we need to perform two essential operations: efficient *sampling* and *evaluation* of the material response; Fig. 5 summarizes this process.

4.2.1 Model Construction. The BRDF $f(\omega_i, \omega_o)$ (for a fixed surface point) is a four-dimensional function describing directional scattering density for outgoing directions ω_o given an incident direction ω_i (each direction represented by elevation θ and azimuth ϕ). By Helmholtz reciprocity, the roles of ω_i and ω_o are interchangeable.

Model Structure. When querying the BRDF during rendering, the incident direction is known and can therefore be treated as a fixed parameter rather than a random variable subject to sampling. Consequently, representing the full 4D BRDF with a single multivariate metalog would be impractical, as it would require recovering the corresponding distribution for a given incident direction using the Newton–Raphson solver described in the supplementary document. Instead, we discretize the incident-direction domain into a regular grid parameterized by $\cos \theta_i$ and ϕ_i , and store a bivariate metalog, Eq. (10), describing the distribution of outgoing directions at each grid node, Fig. 7. This choice is purely conventional, as incident and outgoing directions are interchangeable without loss of generality.




Fig. 6. Overview of our metalog material sampling process: Given an incident direction, we interpolate the metalog distributions stored at adjacent grid nodes, producing a bivariate distribution from which we sample the outgoing direction. We first sample the marginal (red), then interpolate the conditional slices adjacent to the marginal sample, and finally sample the resulting conditional distribution.




Fig. 7. Illustration of our sampling data structure: we store a bivariate metalog distribution per incident direction on a hemispherical grid, modeling outgoing directions via a marginal–conditional construction of one-dimensional metalogs. This representation enables efficient path sampling and weighting during MC rendering.

The overall resolution of our model is determined by the number of (1) points in the incident direction grid, (2) conditional slices per bivariate metalog, (3) mixture components, and (4) metalog coefficients. Following prior work [Dupuy and Jakob 2018], we choose a grid resolution of 8 bins in $\cos \theta_i$ and 17 bins in ϕ_i , and use 6 conditional slices as a reasonable default, which we found to provide a good overall trade-off. This grid setup is therefore directly compatible with the data of Dupuy and Jakob, allowing us to fit their distributions using the same incident resolution. We leave the investigation of alternative parameterizations such as HEALPix [Górski et al. 2005] for future work.

Regarding the number of mixture components and metalog coefficients, multiple combinations can yield an adequate fit; both higher-order metalogs and more mixture components enhance the expressive range of the model. As the optimal parameterization depends on the characteristics of the underlying distribution and varies across materials, we determine these parameters in a data-driven manner during fitting rather than fixing them a priori.

Fitting. Given BRDF sample data, we aim to fit an importance-sampling distribution proportional to $f(\omega_i, \omega_o) \cos \theta_o$, Eq. (7), at each incident direction. Given bounded candidate sets for the discrete parameters, our fitting procedure evaluates all admissible configurations and selects the one minimizing a normalized Wasserstein distance [Bishop and Bishop 2024] between the fitted model and the data. For each configuration and at each grid point, we estimate the metalog and mixture coefficients with the EM procedure described in §4.1.2. Please refer to the supplementary document for details.

Having thus defined the distribution such that $p(\omega_o | \omega_i) \propto f(\omega_i, \omega_o) \cos \theta_o$, we now complete the reflectance model by deriving

the factor required to recover the BRDF f . The normalization factor for the PDF p is the directional-hemispherical reflectance:

$$\rho_{dh}(\omega_i) = \int_{\Omega} f(\omega_i, \omega_o) \cos \theta_o d\omega_o. \quad (16)$$

Therefore, for our metalog BRDF, we use

$$f(\omega_i, \omega_o) = p(\omega_o | \omega_i) \rho_{dh}(\omega_i) / \cos \theta_o. \quad (17)$$

Consequently, when importance sampling according to $p(\omega_o)$, the associated MC weight is $f(\omega_i, \omega_o) \cos \theta_o / p(\omega_o) = \rho_{dh}(\omega_i)$, which only depends on the incident direction. This observation allows us to precompute $\rho_{dh}(\omega_i)$ for each grid cell, storing the resulting values in a texture. Thus, our entire reflectance model is represented by this texture and the corresponding sampling distribution.

4.2.2 Rendering. When rendering metalog materials, we perform two core operations: (1) importance sampling of outgoing directions, also providing the PDF values and MC weights, or (2) evaluation of the PDF and BRDF in arbitrary directions for MIS.

Sampling. Given an incident direction, we first interpolate the metalog mixture coefficients for this direction from the data stored at the four adjacent grid nodes. From the resulting bivariate metalog mixture, we then draw an outgoing direction by sequential sampling of the mixture component, followed by the marginal and conditional quantile functions; see also Fig. 6. (Note that evaluating the exact PDF of the generated direction requires computing the full mixture density.) This approach is closely related to the anisotropic sampling framework of Dupuy and Jakob [2018]; however, instead of sampling luminance or visible-normal distribution slices, we directly sample directions from our fitted metalog material. The associated MC weight follows analogously from bilinearly interpolating the precomputed grid values of $\rho_{dh}(\omega_i)$ for the incident direction.

PDF and BRDF Evaluation for MIS. For multiple importance sampling (MIS), we need to evaluate both the PDF and the BRDF including the cosine term ($f(\omega_i, \omega_o) \cos \theta_o$). For the PDF, we interpolate the metalog mixture coefficients for the incident direction as above and then compute $p(\omega_o)$ via root finding (see also the supplementary document). Finally, interpolating $\rho_{dh}(\omega_i)$ from the incident-direction grid yields $f(\omega_i, \omega_o) \cos \theta_o = p(\omega_o) \rho_{dh}(\omega_i)$.

5 Results

We implement metalog material fitting in PyTorch [Paszke et al. 2019] and employ Mitsuba 3 with the `cuda_ad_rgb` variant, built on the Dr.Jit framework [Jakob et al. 2022], for (differentiable) rendering.




Fig. 8. Metalog fitting to some of the isotropic materials in the RGL material database. We also compare to neural BRDF (NBRDF) fitting using the publicly available implementation of Sztrajman et al. [2021].




Fig. 9. Metalog fitting to some of the anisotropic materials in the RGL material database. Note that “Reference” always refers to the original material model rendered with the same settings as ours.




Fig. 10. Comparison of image variance (MSE) for various materials: we render high-SPP ground-truth images independently for both our metalog materials and the reference materials, then compare how quickly images for each reflectance model approach the true solution with increasing samples.

Please also refer to the supplementary document for additional implementation details. All results are obtained on an AMD EPYC 7413 CPU and a NVIDIA A40 GPU. We evaluate our material model on a range of representative scenarios, including approximating Mitsuba’s analytic TR-based roughconductor BRDF (Fig. 15), fitting measured BRDF data (Figs. 8 and 9), smoothly interpolating between various materials (Fig. 17), manual editing through parametric adjustments (Fig. 19), and appearance capture via optimization with differentiable rendering (Fig. 16). Together, these scenarios illustrate the expressiveness, flexibility, and practical usability of metalog




Fig. 11. Error and timing comparison plots for the results in Figs. 8 and 9.




Fig. 12. Sampling and PDF evaluation times for metalog materials fitting aniso_sari_silk_2color with varying number of marginal mixture components, metalog coefficients, and conditional slices. Each plot shows two of these dimensions, while the third is set to the maximal value within the dataset (top: 11 coefficients per metalog, bottom: 8 mixture components).

materials. We restrict our evaluation to BRDFs (surface reflectance) and report both fitting quality and rendered appearance, with direct comparisons to corresponding analytic and measured reference models. Note that the memory footprint of our sampling distribution is fully determined by the number of coefficients of our metalog mixture representation. In particular, the total memory footprint scales linearly with the number of bivariate metalog slices; please refer to the supplementary document (Tables S1 and S2) for details.

Fitting a metalog material to existing BRDF data, either from analytic models or real-world measurements is the most natural starting point for most of our results. We show metalog materials closely approximating various real-world materials (Figs. 8 and 9) from the RGL measured material database [Dupuy and Jakob 2018]. In Fig. 10, we compare the equal-sample variance of our metalog materials against the corresponding reference materials. Our approach achieves equal or slightly lower variance, which is expected, as the PDF values produced by our method are exactly proportional to the cosine-weighted BRDF values. Render times remain comparable to the reference models, demonstrating that metalog BRDFs integrate efficiently into the rendering pipeline, Fig. 11. In Fig. 12, we additionally report timings separately for sampling and PDF evaluation, quantifying the impact of our material model on MIS performance.




Fig. 13. Ablation study: we fit metalog materials to an isotropic (left) and anisotropic (right) reference material and compare our results to restricted variants that do not use mixtures ($K = 1$), or use only a single conditional distribution instead of interpolating multiple slices. Fitting distributions with single conditionals to the anisotropic material fails to find a feasible solution.




Fig. 14. Extended ablation study: we fit two anisotropic materials from the RGL database with varying metalog mixture configurations and report the overall image error (FLIP) for the standard view, as in Fig. 13. Here, K_m refers to the number of (marginal) mixture components and white areas indicate failures to fit feasible distributions.

Generally, a single component (i.e., not using mixtures) is sufficient for most isotropic materials, while anisotropy often benefits from 3 to 4 mixture components (see Table S2 in the supplementary document). In Fig. 13, we show an ablation study for isotropic and anisotropic materials, where we compare our results (using both metalog mixtures and interpolation across multiple conditional slices) to restricted results without mixtures (only one component) or without slice interpolation (only one conditional distribution). For isotropic materials, the difference between a single and multiple mixture components is typically very small, but becomes significant for anisotropy. Conversely, not using multiple conditional slices leads to noticeable errors for isotropic and sometimes fails to deliver feasible metalogs for anisotropic materials; please also refer to extended results in the supplementary document. We extend this ablation study to fitting challenging anisotropic materials with specific metalog configurations in Fig. 14. When using too few mixture components, we observe an increased failure rate in fitting feasible distributions, which is typically not resolved by simply using more metalog coefficients. Increasing the number of conditional slices, on the other hand, consistently lowers observed errors, enabled by our interpolation scheme. Generally, fitting highly specular materials tends to cause feasibility issues; thus, in practice, we test multiple configurations and select the best one in terms of the Wasserstein-1




Fig. 15. We fit our metalog materials to TR materials with different isotropic roughness values α (blue boxes) and interpolate the results (top row), producing behavior similar to directly varying TR roughness (bottom row).

distance to the reference material. Depending on the material, fitting requires between 1.3 min and 1.2 h; we report all fitting times in the supplementary document (Tables S1 and S2) and note that fitting is a precomputation step with no impact on runtime performance.

5.1 Material Interpolation

The advantages of quantile interpolation for metalog distributions naturally extend to the level of entire material models, enabling interpolation between complete BRDF representations, which yields a continuous transition between appearances. In Fig. 15, we first interpolate between a rough and a specular material. The result closely matches the effect of varying the analytic Trowbridge-Reitz (TR) roughness and exhibits a smooth progression, demonstrating that interpolation in our representation can effectively replace parameter variation in the original model. We also demonstrate interpolation of vastly different materials originating from incompatible input models, Fig. 17; our metalog representation enables a smooth transition grounded in the Wasserstein optimality of quantile interpolation. Interpolating metalogs across a surface yields continuous spatially varying bidirectional reflectance distribution functions (SVBRDFs) without resorting to stochastic material selection (example provided in the supplementary document).




Fig. 16. Materials acquired via differentiable rendering from real-world images (using 8 views per sample in 45° rotation increments).




Fig. 17. Interpolating Trowbridge-Reitz (gold, $\alpha = 0.1$; back left), color_connemara_4c (back right), and cc_northern_aurora (front).

5.2 Acquisition via Differentiable Rendering

Metalogs are inherently differentiable: the quantile function $M_t(y)$ is composed of smooth basis functions. Our metalog materials are therefore particularly well-suited for material acquisition, i.e., estimating coefficients π and a , and reflectance ρ_{dh} , via inverse rendering. We show an experiment using a single-camera capture setup, capturing small marbles on a turntable in 45° increments. We then optimize the metalog coefficients to match the captured images in Mitsuba 3 (Dr.Jit); Fig. 18 displays the optimization progress. Our results, Fig. 16, show that our optimized metalog materials, including a spatial albedo texture, closely reproduce the appearance of their real-world counterparts.




Fig. 18. Optimization loss (using Adam with Dr.Jit's default parameters) for the examples in Fig. 16.

5.3 Appearance Editing

Unlike neural models that typically require retraining to edit the appearance, all coefficients of our metalog materials serve identifiable purposes. We can therefore manually edit a material's appearance by hand-tuning these coefficients directly. This editability is highly advantageous for data-driven materials, as edits on raw data are extremely difficult. In Fig. 19, we show manual roughness variation and targeted modification of grazing-angle scattering by adjusting a small number of coefficients.

In our experiments, we focus on simple, direct manipulations of a few metalog coefficients. While such manipulations already enable meaningful appearance editing, it requires some understanding of the underlying model structure and of the visual effect of individual coefficients. Developing production-oriented higher-level editing interfaces that map intuitive user controls to coordinated modifications of multiple coefficients constitutes an interesting direction for future work. Furthermore, we prioritize expressive material control over strict physical validity: we do not enforce reciprocity or energy conservation during editing, although a refitting or optimization pass could in theory restore physically accurate behavior.




Fig. 19. Appearance editing: metallog materials support editing operations by directly modifying metallog coefficients. Top: we first fit a metallog material to a TR material ($\alpha = 0.3$), then edit it to increase specularity. Bottom: starting from a fit to `cc_northern_aurora` from the RGL database, we edit the material for brighter, more diffuse reflectance at grazing angles.

5.4 Comparison to Neural BRDFs

We compare our metallog materials to the neural BRDF (NBRDF) approach of Sztrajman et al. [2021] in Fig. 8 (and in the supplementary document). In our tests, the results are often similar in terms of material appearance, although the neural model sometimes struggles with highly specular materials, which are well resolved by our approach. Since their publicly available implementation is based on Mitsuba 1, whereas our materials are implemented in Mitsuba 3, a direct performance comparison is not meaningful. However, Sztrajman et al. report a modest overall runtime improvement over the approach of Dupuy and Jakob [2018], which we found to exhibit performance comparable to, or slightly slower than, our materials (Fig. 11). We therefore expect both approaches to exhibit similar runtime characteristics overall. We also note that Sztrajman et al. [2021] cannot directly sample the neural reflectance model, but use an approximate analytic model for importance sampling.

5.5 Reciprocity and Energy Conservation

Physically plausible reflectance models should satisfy Helmholtz reciprocity and energy conservation. Helmholtz reciprocity requires that the BRDF remains unchanged when the incident and outgoing directions are exchanged, i.e., $f(\omega_i, \omega_o) = f(\omega_o, \omega_i)$. Energy conservation requires that a material cannot reflect more energy than it receives, i.e., $\forall \omega_i : \int_{\Omega} f(\omega_i, \omega_o) \cos \theta_o d\omega_o \leq 1$. Our metallog materials are constructed from a PDF p and a reflectance factor ρ_{dh} , where p is always a normalized distribution, and it is straightforward to verify whether ρ_{dh} satisfies energy conservation.

Regarding Helmholtz reciprocity, our approach does not currently enforce this constraint. Assuming that the source materials satisfy reciprocity, any violations of this constraint in the fitted materials can be attributed to fitting errors. In Fig. 20, we report relative reciprocity errors for a variety of materials, measured globally as




Fig. 20. Reciprocity errors for isotropic TR materials with varying roughness α : top row shows total error $\mathcal{E}$, while the images show local errors $\epsilon(\omega_o, \omega_i)$ according to Eq. (18) at incident direction $(\theta_i, \phi_i) = (0.75, 0)$.




Fig. 21. Periodicity stress test: bounded metallogs currently force the PDF to zero at the bounds of the azimuth angle, leading to dark stripes when a point light source is placed in exactly this location (middle). These issues quickly fade away as the light source covers a larger solid angle (right). In our supplementary video, we show an animation rotating the light source around the camera as illustrated (left).

$\mathcal{E} = \int_{\Omega} \int_{\Omega} \epsilon(\omega_o, \omega_i) d\omega_o d\omega_i$, where the local error is

$$\epsilon(\omega_o, \omega_i) = \frac{\|f(\omega_i, \omega_o) - f(\omega_o, \omega_i)\|_2 \cos \theta_o \cos \theta_i}{\int_{\Omega} \int_{\Omega} \|f(\omega_o, \omega_i)\|_2 \cos \theta_o \cos \theta_i d\omega_o d\omega_i}. \quad (18)$$

This metric integrates reciprocity violations over all incident and outgoing directions and normalizes the result by the total reflected energy. In our work, we always employ unidirectional path tracing, for which reciprocity violations do not cause significant practical issues. When using algorithms that rely on bidirectional transport, reciprocity could be enforced directly during fitting through additional loss terms in the future.

5.6 Azimuthal Periodicity

Metallogs are aperiodic by construction and therefore do not ensure continuity between $-\pi$ and π in the azimuthal domain, despite these angles representing the same direction. To investigate this issue, we construct a stress test where a point light source rotates around a planar surface, forcing the renderer to evaluate the BRDF at a specific azimuth. We observe a narrow, dark stripe when the light source is close to the azimuth seam, as the bounded metallog PDF must be zero at the bounds, Fig. 21. However, these artifacts disappear once the size of the light source increases and we integrate over a larger range of directions. While these issues are imperceptible when using environment lighting, in principle, we could slightly extend the bounded metallog domain and enforce azimuthal periodicity explicitly during fitting if desired.

6 Limitations and Future Work

In this work, we focus exclusively on local BRDFs. Consequently, phenomena involving transmission and volumetric scattering are beyond its scope. In particular, the characteristic chatoyant appearance of tiger's-eye gemstones (featured in Fig. 16) arises from volumetric

scattering within aligned internal fibers and is strongly influenced by the gemstone's geometry. While certain aspects of this appearance may be captured by a measured BRDF, the underlying volumetric scattering process itself cannot be represented by a local BRDF. Extending our framework beyond local BRDFs to account for such effects therefore constitutes an important direction for future work.

7 Discussion and Conclusion

In this work, we introduce metalog distributions as a general, compact, and versatile representation for complex BRDFs. Our results demonstrate that metalogs capture a wide range of reflectance behaviors while remaining practical for rendering. Their quantile-based formulation enables efficient sampling and supports intuitive quantile-based interpolation, allowing smooth reconstruction across discrete slices and extending naturally to interpolation at the level of complete material models, including SVBRDFs. In terms of memory cost, our metalog model requires around 2 KB for isotropic and 50 KB for anisotropic materials in most cases. One notable limitation of our approach is that evaluating the metalog PDF during MIS requires a small Newton–Raphson solve; however, this typically converges in only a few iterations and represents a reasonable trade-off given the benefits. Our results show that (equal-sample) render times using MIS are comparable to those of the material representation by Dupuy and Jakob [2018], while standard BRDF-based importance sampling is substantially faster. Furthermore, our formulation does not explicitly enforce reciprocity, but approximates reciprocal reference data. In the future, our approach to determining the discontinuous parameters (number of mixture components and metalog coefficients) during fitting could be improved by employing more elaborate mixed-integer optimization methods. Overall, metalog materials offer a flexible and numerically robust framework for data-driven material modeling, providing a promising foundation for future research on unified material representations in graphics and related fields.

Acknowledgments

We thank Thomas Keelin for valuable discussions regarding metalog distributions and our reviewers for their exceptionally thoughtful and constructive feedback. We also thank the creators of the scenes used in this work: Jay-Artist for “The White Room” (Fig. 1), and Wenzel Jakob and Jonas Pilo for “Material Preview” (used in many figures), illuminated by the environment map “At the Window” by Bernhard Vogl. This work has been funded by the Vienna Science and Technology Fund (WWTF) [Grant ID: 1047379/ICT22028]. This research was funded in whole or in part by the Austrian Science Fund (FWF) [10.55776/F77]. For open-access purposes, the author has applied a CC BY public copyright license to any author-accepted manuscript version arising from this submission. The authors acknowledge TU Wien Bibliothek for financial support through its Open Access Funding Programme. Parts of this work were previously presented in the Master's thesis of Saip-Can Hasbay [2026].

References

Zap Andersson, Paul Edmondson, Julien Guertault, Adrien Herubel, Peter Kutz, Andréa Machizaud, Jamie Portsmouth, Frédéric Servant, and Jonathan Stone. 2025. OpenPBR Surface: An open shading model for physically based materials. In *Proceedings of*

- the Digital Production Symposium, DigiPro 2025, Vancouver, BC, Canada, 9 August 2025*, Lee Kerley, Maryann Simmons, Eva Abramychewa, and Steven Vargas (Eds.). ACM, 1:1–1:12. doi:10.1145/3744199.3744632
- Mahdi M. Bagher, Cyril Soler, and Nicolas Holzschuch. 2012. Accurate fitting of measured reflectances using a Shifted Gamma micro-facet distribution. *Comput. Graph. Forum* 31, 4 (2012), 1509–1518. doi:10.1111/J.1467-8659.2012.03147.X
- Ahmet Bilgili, Aydin Öztürk, and Murat Kurt. 2011. A General BRDF Representation Based on Tensor Decomposition. *Comput. Graph. Forum* 30, 8 (2011), 2427–2439. doi:10.1111/J.1467-8659.2011.02072.X
- Christopher M. Bishop and Hugh Bishop. 2024. *Deep Learning - Foundations and Concepts*. Springer. doi:10.1007/978-3-031-45468-4
- James F. Blinn. 1977. Models of light reflection for computer synthesized pictures. In *Proceedings of the 4th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1977, San Jose, California, USA, July 20-22, 1977*, James George (Ed.). ACM, 192–198. doi:10.1145/563858.563893
- Brent Burley. 2012. Physically-Based Shading at Disney. In *Practical physically-based shading in film and game production (SIGGRAPH Courses 2012)*. ACM, 10:1–10:7. doi:10.1145/2343483.2343493 Available at <https://disneyanimation.com/publications/physically-based-shading-at-disney/>.
- Luc Claustres, Yannick Boucher, and Mathias Paulin. 2004. Wavelet-based modeling of spectral bidirectional reflectance distribution function data. *Optical Engineering* 43, 10 (2004), 2327–2339. doi:10.1117/1.1789138
- Arnau Colom, Christian Bouville, Julien Pettré, Kadi Bouatouch, and Ricardo Marques. 2026. Fast and Accurate Gaussian Process Modelling of Real-World Materials. *ACM Trans. Graph.* 45, 3 (2026), 33:1–33:18. doi:10.1145/3806194
- Robert L. Cook and Kenneth E. Torrance. 1982. A Reflectance Model for Computer Graphics. *ACM Trans. Graph.* 1, 1 (1982), 7–24. doi:10.1145/357290.357293
- Jonathan Dupuy, Eric Heitz, Jean-Claude Iehl, Pierre Poulin, and Victor Ostromoukhov. 2015. Extracting Microfacet-based BRDF Parameters from Arbitrary Materials with Power Iterations. *Comput. Graph. Forum* 34, 4 (2015), 21–30. doi:10.1111/CGF.12675
- Jonathan Dupuy and Wenzel Jakob. 2018. An adaptive parameterization for efficient material acquisition and rendering. *ACM Trans. Graph.* 37, 6 (2018), 274. doi:10.1145/3272127.3275059
- Arne Dür. 2006. An Improved Normalization for the Ward Reflectance Model. *J. Graph. Tools* 11, 1 (2006), 51–59. doi:10.1080/2151237X.2006.10129215
- Alain Fournier. 1995. Separating Reflection Functions for Linear Radiosity. In *Rendering Techniques '95, Proceedings of the Eurographics Workshop in Dublin, Ireland, June 12-14, 1995 (Eurographics)*, Pat Hanrahan and Werner Purgathofer (Eds.). Springer, 296–305. doi:10.1007/978-3-7091-9430-0_28
- David Geisler-Moroder and Arne Dür. 2010. A New Ward BRDF Model with Bounded Albedo. *Comput. Graph. Forum* 29, 4 (2010), 1391–1398. doi:10.1111/J.1467-8659.2010.01735.X
- Dar'ya Guarnera, Giuseppe Claudio Guarnera, Abhijeet Ghosh, Cornelia Denk, and Mashhuda Glencross. 2016. BRDF Representation and Acquisition. *Comput. Graph. Forum* 35, 2 (2016), 625–650. doi:10.1111/CGF.12867
- K. M. Görski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. 2005. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *The Astrophysical Journal* 622, 2 (apr 2005), 759. doi:10.1086/427976
- Saip-Can Hasbay. 2026. *The Fellowship of Metalog BSDFs*. Master's thesis. TU Wien. doi:10.34726/hss.2026.134821
- Eric Heitz. 2014. Understanding the Masking-Shadowing Function in Microfacet-Based BRDFs. *Journal of Computer Graphics Techniques (JCGT)* 3, 2 (30 June 2014), 48–107. <http://jcgt.org/published/0003/02/03/>
- Eric Heitz and Eugene d'Eon. 2014. Importance Sampling Microfacet-Based BSDFs using the Distribution of Visible Normals. *Comput. Graph. Forum* 33, 4 (2014), 103–112. doi:10.1111/CGF.12417
- Eric Heitz, Jonathan Dupuy, Cyril Crassin, and Carsten Dachsbacher. 2015. The SGGX microflake distribution. *ACM Trans. Graph.* 34, 4 (2015), 48:1–48:11. doi:10.1145/2766988
- Eric Heitz, Johannes Hanika, Eugene d'Eon, and Carsten Dachsbacher. 2016. Multiple-scattering microfacet BSDFs with the Smith model. *ACM Trans. Graph.* 35, 4 (2016), 58:1–58:14. doi:10.1145/2897824.2925943
- Sebastian Herholz, Oskar Elek, Jens Schindel, Jaroslav Krivánek, and Hendrik P. A. Lensch. 2018. A Unified Manifold Framework for Efficient BRDF Sampling based on Parametric Mixture Models. In *29th Eurographics Symposium on Rendering, Experimental Ideas & Implementations, EGSR 2018, EI&I Track, Karlsruhe, Germany, 1-4 July 2018*, Wenzel Jakob and Toshiya Hachisuka (Eds.). Eurographics Association, 41–52. doi:10.2312/SRE.20181171
- Stephen Hill, Stephen McAuley, Laurent Belcour, Naty Hoffman, Alain Hostettler, Peter Kutz, Kentaro Suzuki, Hajime Uchimura, Andrea Weidlich, Kenichiro Yasutomi, Pascal Barla, Alban Fichet, and Jamie Portsmouth. 2025. Physically Based Shading in Theory and Practice. In *Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Courses, SIGGRAPH Courses 2025, Vancouver, British Columbia, Canada, August 10-14, 2025*, Ginger Alford and Kalina Borkiewicz (Eds.). ACM, 2:1–2:2. doi:10.1145/3721241.3733991

- Wenzel Jakob, Milos Hasan, Ling-Qi Yan, Jason Lawrence, Ravi Ramamoorthi, and Steve Marschner. 2014. Discrete stochastic microfacet models. *ACM Trans. Graph.* 33, 4 (2014), 115:1–115:10. doi:10.1145/2601097.2601186
- Wenzel Jakob, Sébastien Speierer, Nicolas Roussel, and Delio Vicini. 2022. DRJIT: a just-in-time compiler for differentiable rendering. *ACM Trans. Graph.* 41, 4 (2022), 124:1–124:19. doi:10.1145/3528223.3530099
- Wojciech Jarosz, Nathan A. Carr, and Henrik Wann Jensen. 2009. Importance Sampling Spherical Harmonics. *Comput. Graph. Forum* 28, 2 (2009), 577–586. doi:10.1111/j.1467-8659.2009.01398.x
- James T. Kajiya. 1986. The rendering equation. In *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1986, Dallas, Texas, USA, August 18–22, 1986*, David C. Evans and Russell J. Athay (Eds.). ACM, 143–150. doi:10.1145/15922.15902
- Jan Kautz and Michael D. McCool. 1999. Interactive Rendering with Arbitrary BRDFs using Separable Approximations. In *Rendering Techniques '99, Proceedings of the Eurographics Workshop in Granada, Spain, June 21–23, 1999 (Eurographics)*, Dani Lischinski and Gregory Ward Larson (Eds.). Springer, 247–260. doi:10.1007/978-3-7091-6809-7_22
- Thomas W. Keelin. 2016. The Metallog Distributions. *Decis. Anal.* 13, 4 (2016), 243–277. doi:10.1287/DECA.2016.0338
- Thomas W. Keelin. 2023. The Multivariate Metallog Distributions with Application to Strategic Decision-Making in Golf. doi:10.31219/osf.io/cdmrv
- Thomas W. Keelin and Bradford W. Powley. 2011. Quantile-Parameterized Distributions. *Decis. Anal.* 8, 3 (2011), 206–219. doi:10.1287/DECA.1110.0213
- Murat Kurt, László Szirmay-Kalos, and Jaroslav Krivánek. 2010. An anisotropic BRDF model for fitting and Monte Carlo rendering. *SIGGRAPH Comput. Graph.* 44, 1, Article 3 (Feb. 2010), 15 pages. doi:10.1145/1722991.1722996
- Alexandr Kuznetsov, Krishna Mullia, Zexiang Xu, Milos Hasan, and Ravi Ramamoorthi. 2021. NeuMIP: multi-resolution neural materials. *ACM Trans. Graph.* 40, 4 (2021), 175:1–175:13. doi:10.1145/3450626.3459795
- Jason Lawrence, Aner Ben-Artzi, Christopher DeCoro, Wojciech Matusik, Hanspeter Pfister, Ravi Ramamoorthi, and Szymon Rusinkiewicz. 2006. Inverse shade trees for non-parametric material representation and editing. *ACM Trans. Graph.* 25, 3 (2006), 735–745. doi:10.1145/1141911.1141949
- Joo Ho Lee, Adrián Jarabo, Daniel S. Jeon, Diego Gutierrez, and Min H. Kim. 2018. Practical multiple scattering for rough surfaces. *ACM Trans. Graph.* 37, 6 (2018), 275. doi:10.1145/3272127.3275016
- Joachim Löw, Joel Kronander, Anders Ynnerman, and Jonas Unger. 2012. BRDF models for accurate and efficient rendering of glossy surfaces. *ACM Trans. Graph.* 31, 1 (2012), 9:1–9:14. doi:10.1145/2077341.2077350
- Wojciech Matusik, Hanspeter Pfister, Matthew Brand, and Leonard McMillan. 2003. A data-driven reflectance model. *ACM Trans. Graph.* 22, 3 (2003), 759–769. doi:10.1145/882262.882343
- Michael D. McCool, Jason Ang, and Anis Ahmad. 2001. Homomorphic factorization of BRDFs for high-performance rendering. In *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 2001, Los Angeles, California, USA, August 12–17, 2001*, Lynn Pocock (Ed.). ACM, 171–178. doi:10.1145/383259.383276
- Thomas Müller, Brian McWilliams, Fabrice Rousselle, Markus Gross, and Jan Novák. 2019. Neural Importance Sampling. *ACM Trans. Graph.* 38, 5 (2019), 145:1–145:19. doi:10.1145/3341156
- László Neumann, Attila Neumann, and László Szirmay-Kalos. 1999. Compact Metallic Reflectance Models. *Comput. Graph. Forum* 18, 3 (1999), 161–172. doi:10.1111/1467-8659.00337
- Michael Oren and Shree K. Nayar. 1994. Generalization of Lambert's reflectance model. In *Proceedings of the 21th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1994, Orlando, FL, USA, July 24–29, 1994*, Dino Schweitzer, Andrew S. Glassner, and Mike Keeler (Eds.). ACM, 239–246. doi:10.1145/192161.192213
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Z. Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8–14, 2019, Vancouver, BC, Canada*, Hanna M. Wallach, Hugo Larochelle, Alina Beygelzimer, Florence d'Alché-Buc, Emily B. Fox, and Roman Garnett (Eds.). 8024–8035. <https://proceedings.neurips.cc/paper/2019/hash/bdbca288fee792f2bfa9f7012727740-Abstract.html>
- Gabriel Peyré and Marco Cuturi. 2019. Computational Optimal Transport. *Found. Trends Mach. Learn.* 11, 5–6 (2019), 355–607. doi:10.1561/22000000073
- Bui Tuong Phong. 1975. Illumination for Computer Generated Pictures. *Commun. ACM* 18, 6 (1975), 311–317. doi:10.1145/360825.360839
- Jamie Portsmouth, Peter Kutz, and Stephen Hill. 2025. EON: A practical energy-preserving rough diffuse BRDF. *Journal of Computer Graphics Techniques (JCGT)* 14, 1 (31 March 2025), 116–139b. <http://jcgt.org/published/0014/01/06/>
- William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. 2007. *Numerical recipes: the art of scientific computing, 3rd Edition*. Cambridge University Press. <https://www.worldcat.org/oclc/123285342>
- Enrique Rosales, Fatemeh Teimury, Joshua Horacek, Aria Salari, Xuebin Qin, Adi Bar-Lev, Xiaoqiang Zhe, and Ligang Liu. 2023. Fast-MSX: Fast Multiple Scattering Approximation. In *SIGGRAPH Asia 2023 Conference Papers, SA 2023, Sydney, NSW, Australia, December 12–15, 2023*, June Kim, Ming C. Lin, and Bernd Bickel (Eds.). ACM, 11:1–11:9. doi:10.1145/3610548.3618231
- Szymon Rusinkiewicz. 1998. A New Change of Variables for Efficient BRDF Representation. In *Rendering Techniques '98, Proceedings of the Eurographics Workshop in Vienna, Austria, June 29 - July 1, 1998 (Eurographics)*, George Drettakis and Nelson L. Max (Eds.). Springer, 11–22. doi:10.1007/978-3-7091-6453-2_2
- Ana Serrano, Diego Gutierrez, Karol Myszkowski, Hans-Peter Seidel, and Belén Masía. 2016. An intuitive control space for material appearance. *ACM Trans. Graph.* 35, 6 (2016), 186:1–186:12. doi:10.1145/2980179.2980242
- B. Smith. 1967. Geometrical shadowing of a random rough surface. *IEEE Transactions on Antennas and Propagation* 15, 5 (1967), 668–671. doi:10.1109/TAP.1967.1138991
- Xin Sun, Kun Zhou, Yanyun Chen, Stephen Lin, Jiaoying Shi, and Baining Guo. 2007. Interactive relighting with dynamic BRDFs. *ACM Trans. Graph.* 26, 3 (2007), 27. doi:10.1145/1276377.1276411
- Alejandro Sztajman, Gilles Rainer, Tobias Ritschel, and Tim Weyrich. 2021. Neural BRDF Representation and Importance Sampling. *Comput. Graph. Forum* 40, 6 (2021), 332–346. doi:10.1111/CGF.14335
- K. E. Torrance and E. M. Sparrow. 1967. Theory for Off-Specular Reflection From Roughened Surfaces. *J. Opt. Soc. Am.* 57, 9 (Sep 1967), 1105–1114. doi:10.1364/JOSA.57.001105
- T. S. Trowbridge and K. P. Reitz. 1975. Average irregularity representation of a rough surface for ray reflection. *J. Opt. Soc. Am.* 65, 5 (May 1975), 531–536. doi:10.1364/JOSA.65.000531
- Emmanuel Turquin. 2019. *Practical multiple scattering compensation for microfacet models*. Technical Report. https://blog.selfshadow.com/publications/turquin/ms_comp_final.pdf
- Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. 2007. Microfacet Models for Refraction through Rough Surfaces. In *Proceedings of the Eurographics Symposium on Rendering Techniques, Grenoble, France, 2007*, Jan Kautz and Sumanta N. Pattanaik (Eds.). Eurographics Association, 195–206. doi:10.2312/EGWR/EGSR07/195-206
- Gregory J. Ward. 1992. Measuring and modeling anisotropic reflection. In *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1992, Chicago, IL, USA, July 27–31, 1992*, James J. Thomas (Ed.). ACM, 265–272. doi:10.1145/133994.134078
- Andrea Weidlich. 2025. Bridging the Gap Between Offline and Real Time with Neural Materials. In *Physically Based Shading in Theory and Practice (SIGGRAPH Courses 2025)*. ACM, 2:1–2:2. doi:10.1145/3721241.3733991 Available at https://blog.selfshadow.com/publications/s2025-shading-course/weidlich/s2025_pbs_weidlich_notes_v1.1.pdf
- Feng Xie and Pat Hanrahan. 2018. Multiple scattering from distributions of specular v-grooves. *ACM Trans. Graph.* 37, 6 (2018), 276. doi:10.1145/3272127.3275078
- Tizian Zeltner, Fabrice Rousselle, Andrea Weidlich, Petrik Clarberg, Jan Novák, Benedikt Bitterli, Alex Evans, Tomás Davidovic, Simon Kallweit, and Aaron E. Lefohn. 2024. Real-time Neural Appearance Models. *ACM Trans. Graph.* 43, 3 (2024), 33:1–33:17. doi:10.1145/3659577
- C. Zhou, A. Sztajman, G. Rainer, F. Zhong, F. Gokbudak, Z. Guo, W. Xia, R. K. Mantiuk, and C. Oztireli. 2026. Physically Based Neural BRDF: A Framework for Physically Correct Material Reconstruction, Generation and Editing. *Computer Graphics Forum* (2026). doi:10.1111/cgf.70500