

# An Adaptive Parameterization for Efficient Material Acquisition and Rendering

JONATHAN DUPUY, Unity Technologies

WENZEL JAKOB, École Polytechnique Fédérale de Lausanne (EPFL)



Fig. 1. Spectral rendering of isotropic and anisotropic materials acquired from real-world samples using our method; insets show corresponding reflectance spectra. We measured these BRDFs using a motorized goni-photometer, leveraging our novel adaptive parameterization to simultaneously handle BRDF acquisition, storage, and efficient Monte Carlo sample generation during rendering. Our representation requires 16 KiB of storage per spectral sample for isotropic materials and 544 KiB per spectral sample for anisotropic specimens.

One of the key ingredients of any physically based rendering system is a detailed specification characterizing the interaction of light and matter of all materials present in a scene, typically via the Bidirectional Reflectance Distribution Function (BRDF). Despite their utility, access to real-world BRDF datasets remains limited: this is because measurements involve scanning a four-dimensional domain at sufficient resolution, a tedious and often infeasibly time-consuming process.

We propose a new parameterization that automatically adapts to the behavior of a material, warping the underlying 4D domain so that most of the volume maps to regions where the BRDF takes on non-negligible values, while irrelevant regions are strongly compressed. This adaptation only requires a brief 1D or 2D measurement of the material's retro-reflective properties. Our parameterization is unified in the sense that it combines several steps that previously required intermediate data conversions: the same mapping can simultaneously be used for BRDF acquisition, storage, and it supports efficient Monte Carlo sample generation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2018 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.  
0730-0301/2018/11-ART1 \$15.00  
<https://doi.org/10.1145/3272127.3275059>

We observe that the above desiderata are satisfied by a core operation present in modern rendering systems, which maps uniform variates to direction samples that are proportional to an analytic BRDF. Based on this insight, we define our adaptive parameterization as an invertible, retro-reflectively driven mapping between the parametric and directional domains. We are able to create noise-free renderings of existing BRDF datasets after conversion into our representation with the added benefit that the warped data is significantly more compact, requiring 16KiB and 544KiB per spectral channel for isotropic and anisotropic specimens, respectively.

Finally, we show how to modify an existing goni-photometer to provide the needed retro-reflection measurements. Acquisition then proceeds within a 4D space that is warped by our parameterization. We demonstrate the efficacy of this scheme by acquiring the first set of spectral BRDFs of surfaces exhibiting arbitrary roughness, including anisotropy.

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

Additional Key Words and Phrases: microfacet theory, spectral BRDF, photorealistic rendering

## ACM Reference Format:

Jonathan Dupuy and Wenzel Jakob. 2018. An Adaptive Parameterization for Efficient Material Acquisition and Rendering. *ACM Trans. Graph.* 37, 6, Article 1 (November 2018), 14 pages. <https://doi.org/10.1145/3272127.3275059>

## 1 INTRODUCTION

Physically based rendering algorithms simulate real-world appearance by means of an intricate simulation of the interaction of light and matter. Scattering by surfaces denotes the most important type

of interaction, and the physics of this process are typically encoded in a quantity known as the *bidirectional reflectance distribution function* (BRDF). High-fidelity BRDF models thus constitute a crucial ingredient to any type of realistic rendering. In this work, we propose a practical BRDF representation that simultaneously serves as a mechanism for its acquisition, storage and manipulation within rendering software.

*BRDF Acquisition.* Given a unit amount of incident illumination arriving from a specific direction, the BRDF specifies the directional profile of light scattered by a surface. This definition implies a conceptually straightforward mechanism for BRDF acquisition from real-world measurements that entails illuminating the surface with parallel illumination and measuring the scattered radiance traveling into a dense set of outgoing directions. The availability of such measurements is important not only because they improve realism, but also because they serve as validation against theoretical models that help advance our understanding of light-material interactions. For instance, the MERL database [Matusik et al. 2003a] has served as inspiration and validation of numerous ray- and wave-optics BRDF models over the last fifteen years [Ashikhmin and Premoze 2007; Bagher et al. 2012; Dupuy et al. 2015; Holzschuch and Pacanowski 2017; Löw et al. 2012; Ngan et al. 2005].

*Challenges in Acquisition.* Unfortunately, obtaining rich sets of BRDF measurements is challenging for several reasons:

- *High dimensionality.* The BRDF is typically a function of two directions and wavelength, i.e., a five-dimensional quantity. Any type of measurement or computation involving such high-dimensional data is invariably costly by nature due to the curse of dimensionality.
- *Arbitrary frequency.* A BRDF can contain abrupt changes in both the directional and spectral domains, and it is difficult to infer the (5D) locations of where such changes occur without performing dense sets of measurements. Due to the five-dimensional domain, this is generally not feasible.
- *Physical constraints.* Performing measurements with a physical instrument carries its own set of practical challenges: beam occlusion and stretching at grazing angles, mechanically impossible angular configurations, high dynamic range, exposure times, etc. Any practical acquisition scheme must confront these realities and develop mechanisms to efficiently explore the 5D domain and fill in missing data at reasonable time budgets.

Because of these difficulties, existing measurements have been limited to coarse resolutions in either or both the directional and spectral domains: To our knowledge, high resolution spectral BRDF measurements do not exist at all, while RGB measurements of anisotropically rough surfaces are limited to low-frequency specimens.

*Contributions and Outline.* In this work, we provide a new set of all-frequency spectral BRDF measurements from surfaces of arbitrary roughness, i.e., possibly anisotropic. Our BRDF measurements rely on a novel directionally adaptive BRDF parameterization, which allows us to infer the location of the most representative samples for arbitrary BRDF measurements and sample them at low resolution. More specifically, our contributions are

- In Section 3, we show to build an efficient BRDF parameterization for acquisition based on microfacet theory.
- In Section 4, we describe our pipeline for material acquisition that leverages this parameterization.
- In Section 5, we compare our method against related BRDF representations. Following this, we introduce a new spectral BRDF database alongside validation, analysis and renderings.

## 2 RELATED WORK

*BRDF Parameterizations.* As for any other type of signal, the acquisition of a BRDF relies on a parameterization that determines where the measurements should be taken. Since the Whittaker-Nyquist-Kotelnikov-Shannon sampling theorem bounds the lowest measurement resolution to twice a signal’s maximum frequency, any practical BRDF parameterization should ensure that the BRDF’s maximum frequency is the lowest possible in parametric space. To this end, various parameterizations have been proposed and/or compared [Barla et al. 2015; Guarnera et al. 2016; Löw et al. 2012; Matusik et al. 2003b; Nielsen et al. 2015; Rusinkiewicz 1998; Simonot and Obein 2007; Stark et al. 2005]. The fundamental limitation of these parameterizations is that they are independent of the material under measurement. As such, they can only be expected to work for a certain category of materials. For instance, parameterizations built around the half-angle [Barla et al. 2015; Löw et al. 2012; Rusinkiewicz 1998; Simonot and Obein 2007; Stark et al. 2005] can only capture high-frequency materials accurately at high sampling rates, while those obtained from optimizations through datasets [Matusik et al. 2003b; Nielsen et al. 2015] are restricted to materials similar to those in the dataset. In this work, we avoid such limitations by procedurally generating a parameterization for each material we acquire using a small 2D subset of BRDF measurements, which we perform at the beginning of our acquisition process.

*Microfacet Theory.* Our parameterization builds upon microfacet theory. Microfacet theory is a geometric optics model, which represents a material as an arrangement of microscopic specular facets that deviate incident light. Note that this idea is quite old: it dates back to the work of Bouguer [1760] according to Trowbridge and Reitz [1975]. In the computer graphics literature, microfacet theory was first introduced by Blinn [1977] and has undergone considerable development ever since [Belcour and Barla 2017; Dupuy et al. 2016, 2013, 2015; Heitz 2014; Heitz and d’Eon 2014; Heitz et al. 2016; Jakob et al. 2014; Walter et al. 2007]. What makes this theory especially relevant here is its efficiency at reproducing the behavior of a wide range of real-world materials [Ashikhmin and Premoze 2007; Bagher et al. 2012; Dupuy et al. 2015; Holzschuch and Pacanowski 2017; Löw et al. 2012; Ngan et al. 2005; Torrance and Sparrow 1967]. Based on this observation, we build our parameterization with the expectation that it inherits this versatility in two steps, which effectively define our parameterization. First, we adapt the inversion mechanism of Dupuy et al. [2015] to compute a compatible microfacet configuration from retroreflective measurements. Second, we use this configuration in an importance sampling scheme based on the approach of Heitz and d’Eon [2014] to automatically and adaptively measure the material on the full (4D) directional domain.

*BRDF Measurement Devices.* Once the measurement schedule has been set, the acquisition process requires an opto-mechanical device to automate the large number of measurements that must be taken. Various designs with different accuracy and cost trade-offs exist—we refer the reader to a survey by Guarnera et al. [2016] and limit our discussion to the two most common types:

- Gonio-photometers/reflectometers [Nicomemus et al. 1977; White et al. 1998] use up to four motorized axes to illuminate and observe the sample from arbitrary directions on the sphere. Alternatively, the light source or sensor can be fixed, while the sample rotates. These devices achieve excellent measurement characteristics, but the need to advance the motorized axes for each sample can lead to impractically long acquisition times if dense 4D measurements are desired.
- Image-based devices [Ghosh et al. 2007; Marschner et al. 2000, 1999; Mattison et al. 1998; Ngan et al. 2005; Ren and Zhao 2010; Ward 1992] use CCD or CMOS cameras to acquire many BRDF samples at once, which greatly accelerates the acquisition process. However, such devices have more stringent limitations on their precision on the observable angle combinations [Ward 1992], or they impose a specific sample geometry (typically nonplanar), which impedes measurement of anisotropic samples [Marschner et al. 2000, 1999]. Optical aberrations present additional challenges in image-based devices: for instance, Karamata and Andersen [2013] observe that parallel catadioptric gonio-photometers are fundamentally flawed due to severe aberrations inherent in this design.

Our acquisition pipeline relies on a modified PAB pgII [2018] gonio-photometer, which acquires spectral BRDF samples at a rate of approximately 0.5-1Hz. In conjunction with our adaptive parameterization, this leads to measurement times of approximately 2.5 hours for isotropic samples and 2-3 days for anisotropic samples.

*BRDF Databases.* Several reflectance measurement databases have been created over the last two decades [Apian-Bennewitz 2013; Dana et al. 1999; Filip and Vávra 2014; Filip et al. 2014; Löw et al. 2012; Marschner et al. 1999; Matusik et al. 2003a; Ngan et al. 2005]. Of these, only the MERL [Matusik et al. 2003a] and UTIA databases [Filip and Vávra 2014; Filip et al. 2014] provide sufficient directional resolution to be used straightforwardly for either material modeling and/or rendering. Unfortunately, the former is limited to isotropic materials, while the latter is limited to low-frequency specimens. In addition, both specify RGB values, which require an awkward heuristic conversion from RGB to spectra to be usable in spectral renderers. Burley [2012] points out a number of measurement artifacts present in the MERL database, including optical aberrations, discontinuities, and extrapolation artifacts for angles below 75 degrees; Section 5 and the supplemental material show slices of the MERL database warped through our parameterization, which clearly reveals these problems. A higher-level issue of this dataset is the lack of information on what parts of the data are “real”, and what parts are extrapolated or post-processed. In this work, we provide all-frequency spectral BRDF measurements of isotropic and anisotropic materials. Due to the use of large (>20 × 20cm) flat samples, grazing angle configurations are feasible up to an angle of 85 degrees. Thus far, we have measured 32 isotropic and 4 anisotropic materials using

the techniques described in the paper and will release this material database to the community. Whenever data was post-processed (e.g., to fill in unreachable angle configurations), we also retained the raw measurements and will provide them along with the processed data. We plan to extend our dataset over time so that the database becomes a comprehensive repository of all major material classes.

*BRDF fitting and compression.* A large body of work has investigated analytic and semi-analytic models [Bagher et al. 2016, 2012; Holzschuch and Pacanowski 2017; Löw et al. 2012; Ngan et al. 2005] that can be fit to measured data using numerical optimization. These techniques produce succinct approximations of reflectance data that are potentially orders of magnitude more compact than what is proposed in this article. A major limitation common to all of these methods in our context of BRDF acquisition is that the underlying optimization procedures assume that sufficiently dense BRDF measurements are already available. In contrast, our method directly stores the measurement without any further optimization but does so at carefully chosen locations. In order to validate the accuracy of our approach, we devise a virtual acquisition experiment based on the MERL database, where we effectively mimic the process of a fitting algorithm; Section 5 discusses the results of this experiment.

### 3 OUR BRDF PARAMETERIZATION

In this section, we introduce our adaptive parameterization. We begin by observing a fundamental connection between the concepts of BRDF importance sampling and BRDF parameterization (Section 3.1). Motivated by this insight, we build our parameterization from a parametric importance sampling method driven by a measurement of the material’s retro-reflective properties (Section 3.2). This initial parameterization is only adaptive in the outgoing direction argument, and so we finally introduce further adaptation on the incident domain (Section 3.3).

#### 3.1 Motivation

*The BRDF from a Rendering Perspective.* The main objective of the BRDF in a rendering context is the evaluation of spherical illumination integrals of the form

$$I(\omega_i, \lambda) = \int_{S^2} f_r^\perp(\omega_i, \omega_o, \lambda) L_i(\omega_o, \lambda) d\omega_o, \quad (1)$$

where  $L_i$  is the incident radiance,  $\lambda$  is the wavelength of light, and

$$f_r^\perp(\omega_i, \omega_o, \lambda) = f_r(\omega_i, \omega_o, \lambda) \cos \theta_o \quad (2)$$

is the cosine-weighted BRDF; note that this weighted quantity is the one that is typically measured by a gonio-photometer. In the following derivations, we will omit the wavelength parameter  $\lambda$  for readability. To efficiently solve Equation (1), state-of-the-art rendering systems importance sample the spherical domain. Intuitively, importance sampling provides directions  $\omega_o$  that characterize the most important features of the BRDF, which is illustrated in the left part of Figure 2. As we will show, importance sampling is also a desirable tool in the context of BRDF acquisition.

*Importance Sampling a BRDF.* An importance sampling strategy is a mapping  $g$ —typically a diffeomorphism  $g: \mathcal{U}^2 \rightarrow S^2$ , where




Fig. 2. Fundamental link between parameterization and importance sampling: If a mapping  $g$  that takes parametric space to a space of interest (here, the directional domain) constitutes an efficient importance sampling mechanism, then its inverse  $g^{-1}$  constitutes an efficient parameterization of the space of interest, and vice-versa.

$\mathcal{U}^2$  is the unit square—that is used to perform a change of variables  $\omega_o = g(\mathbf{u})$  in Equation (1) so that it takes on the form

$$I(\omega_i) = \int_{\mathcal{U}^2} f_r^\perp(\omega_i, g(\mathbf{u})) L_i(g(\mathbf{u})) \|J_g(\mathbf{u})\| d\mathbf{u}, \quad (3)$$

where  $J_g$  is the Jacobian of  $g$ :

$$J_g(\mathbf{u}) = \frac{\partial g(\mathbf{u})}{\partial \mathbf{u}}. \quad (4)$$

A “good” strategy makes the integrand nearly constant so that low-variance estimates can be obtained at low sample counts. In a general rendering context, a perfect strategy typically does not exist, because  $L_i$  is not known a priori. High-quality sampling strategies therefore focus on the remaining terms and attempt to make  $\|J_g(\mathbf{u})\|$  inversely proportional to the cosine-weighted BRDF  $f_r^\perp$ . A direct mathematical consequence of this choice is that the Jacobian of the inverse of  $g$ , denoted  $g^{-1}$ , satisfies

$$\|J_{g^{-1}}(\omega_o)\| \propto f_r^\perp(\omega_i, \omega_o). \quad (5)$$

Figure 2 illustrates some implications of Equation (5): By construction, the inverse mapping  $g^{-1}$  warps the BRDF to  $\mathcal{U}^2$  in such a way that regions with negligibly small values are strongly compressed, while most space maps to regions where  $f_r^\perp$  is large; further weighting the BRDF in this  $\mathcal{U}^2$  space turns it into a near-constant signal.

*Summary.* We remark on the appealing characteristics of this construction: if used for acquisition, a suitably chosen forward map  $g$  can be used to focus the measurement on salient parts of the domain, dramatically shortening measurement time. On the other hand, re-parameterizing existing BRDF data via the inverse map  $g^{-1}$  and weighting them by the inverse Jacobian yields low-frequency functions that can be stored at coarse resolutions.

To summarize, we have shown that there is a close relationship between importance sampling and parameterization of BRDFs that can be leveraged for acquisition and storage of measured data.

### 3.2 Microfacet-Based Parameterization

*Microfacet-Based Parameterization.* In an acquisition context, we lack a-priori knowledge about the BRDF  $f_r^\perp$  to be measured, hence we seek a parametric strategy that is sufficiently flexible to handle a vast range of behaviors through its set of parameters. To this end, we

assume that the material’s scattering behavior is in reasonably good agreement with microfacet theory, which models reflection from an arrangement of infinitesimally small specular facets; Figures 3 (1, 2) illustrate this intuition. Note that materials in violation of this property can still be measured, although they will be less smooth in our parameterization and hence require a higher resolution.

*The Microfacet NDF and VNDF.* The statistical orientation of the facets is described by a 2D directional distribution  $D$  known as the *microfacet normal distribution function* (NDF). The NDF is defined such that  $D(\omega_m) d\omega_m$  gives the surface area of facets oriented towards the directions associated with the solid angle element  $d\omega_m$ . In addition, the NDF tells how much facet area is visible from any observation direction  $\omega_i$  via the *distribution of visible normals* (VNDF) [Heitz and d’Eon 2014]

$$D_{\text{vis}}(\omega_m, \omega_i) = \frac{\langle \omega_m \cdot \omega_i \rangle D(\omega_m)}{\sigma(\omega_i)}, \quad (6)$$

where the notation  $\langle \mathbf{a} \cdot \mathbf{b} \rangle = \max\{0, \mathbf{a} \cdot \mathbf{b}\}$  denotes a clamped dot product, and  $\sigma$  denotes the projected area of the facets, which also acts as a normalization constant for the VNDF

$$\sigma(\omega) = \int_{S^2} \langle \omega_m \cdot \omega \rangle D(\omega_m) d\omega_m. \quad (7)$$

Figure 3 (3) illustrates the concept of microfacet projected area. The VNDF serves as predictor of the BRDF to be acquired, thus we rewrite Equation (5) as

$$\begin{aligned} \|J_{g^{-1}}(\omega_o)\| &= D_{\text{vis}}(\omega_m, \omega_i) \frac{d\omega_m}{d\omega_o} \\ &= \frac{D(\omega_m)}{4\sigma(\omega_i)}, \quad \omega_m = \frac{\omega_i + \omega_o}{\|\omega_i + \omega_o\|}, \end{aligned} \quad (8)$$

which corresponds to Heitz and D’Eon’s [2014] importance sampling scheme (see Figure 4). In order to instantiate such an importance sampling scheme, we need to find a microfacet NDF that faithfully describes the behavior of the material under measurement; the rest of the subsection is dedicated to this problem.

*Retrieving the Microfacet NDF.* In the retro-reflective domain, we have  $\omega_i = \omega_o =: \omega$ , so that only microfacets oriented towards the incident direction contribute to the retro-reflection; Figure 3 (4)

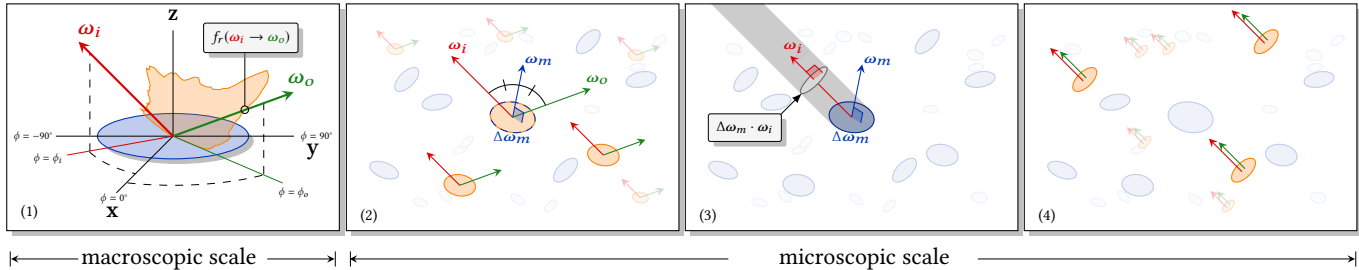


Fig. 3. Geometric configuration and illustration of (1) the BRDF. In microfacet theory, the BRDF emanates from (2) the interactions of light rays and specular facets, whose directional distribution is given by the NDF. The BRDF is derived from the NDF and (3) the projected area of the facets. In this work, we invert the BRDF from (4) backscattering configurations.

illustrates this property. From Equations (5) and (8), we have

$$f_r^\perp(\omega, \omega) \propto \frac{D(\omega)}{4\sigma(\omega)}, \quad (9)$$

which, using Equation (7), can be rewritten as a Fredholm equation of the second kind [Polyanin and Manzhirov 2012]

$$D(\omega) \propto \int_{S^2} K(\omega, \omega_m) D(\omega_m) d\omega_m, \quad (10)$$

where  $K$  denotes the kernel of the equation:

$$K(\omega, \omega_m) = f_r^\perp(\omega, \omega) \langle \omega \cdot \omega_m \rangle. \quad (11)$$

Equation (10) provides a mechanism for retrieving a material's NDF from its retro-reflective response. This is especially convenient in the context of acquisition since retro-reflection is a low-dimensional material property that can be acquired relatively cheaply—it is two-dimensional for anisotropic materials and only one-dimensional for isotropic ones. Since we lack a priori knowledge of the shape of the kernel, we solve Equation (10) numerically. Let  $w_k$  denote the  $k$ -th weight of an arbitrary quadrature rule (we use uniform weights). We obtain the new relation

$$D(\omega_j) = \sum_{k=1}^N w_k K(\omega_j, \omega_k) D(\omega_k), \quad (j = 1, \dots, N) \quad (12)$$

where  $\omega_1, \dots, \omega_N$  are the quadrature nodes. Now, let  $\mathbf{v}$  denote the vector  $\mathbf{v} = (D(\omega_1), \dots, D(\omega_N))^T$ , and  $\mathbf{K}$  the matrix

$$\mathbf{K} = \begin{bmatrix} w_1 K(\omega_1, \omega_1) & \cdots & w_N K(\omega_1, \omega_N) \\ \vdots & \ddots & \vdots \\ w_1 K(\omega_N, \omega_1) & \cdots & w_N K(\omega_N, \omega_N) \end{bmatrix}, \quad (13)$$

Equation (12) becomes an eigenvalue problem

$$\mathbf{v} \propto \mathbf{K} \mathbf{v}. \quad (14)$$

To retrieve the microfacet NDF from retro-reflective measurements, we must thus find a non-negative eigenvector of the kernel matrix  $\mathbf{K}$ . We use the method of Dupuy et al. [2015] to robustly obtain the NDF from the largest eigenvalue's eigenvector via power iterations; Algorithm 1 provides pseudocode for this step.




Fig. 4. Our parameterization builds upon a state-of-the-art sampling scheme by Heitz and D'Eon [2014]. Originally proposed to reduce variance in renderings, this approach also significantly improves the quality of interpolants produced by our method. **(top)** Discretization of an ideal Lambertian BRDF using classic NDF sampling. The leftmost column corresponds to perpendicular incidence ( $\theta_i = 0$ ). Note the large red regions, which are parts of the parameter domain that produce directions below the horizon. **(bottom)** Visible normal sampling considerably reduces this “wasted” space and produces overall smoother interpolants.

#### Algorithm 1 NDF Extraction from retro-reflective responses

```

1: function EXTRACTNDF( $f_r^\perp, N$ )
2:   for each  $j, k \in [1, N]$  do           ▶ Build kernel matrix
3:      $\mathbf{K}[j, k] \leftarrow w_j f_r^\perp(\omega_j, \omega_j) \langle \omega_k \cdot \omega_j \rangle$ 
4:   end for
5:    $\mathbf{v} \leftarrow \mathbf{1}$                        ▶ Create a vector of size  $N$ 
6:   for  $0 \leq i < M$  do                 ▶ Power iterations (we use  $M = 4$ )
7:      $\mathbf{v} \leftarrow \mathbf{K} \mathbf{v}$ 
8:   end for
9:   return  $\mathbf{v}$ 
10: end function
    
```

### 3.3 Parameterization of the Outgoing Domain

Once we have computed the microfacet NDF, we turn to the problem of importance sampling its associated VNDF. We achieve this via a composition of several invertible mappings, which are illustrated in Figure 5. In practice, we evaluate the mappings in a forward sense during acquisition and sampling, and in an inverse sense during BRDF evaluation.

*Data-driven warps.* We heavily rely on warping schemes that transform uniform variates on the unit square  $\mathcal{U}^2$  into samples that are distributed according to an arbitrary bilinear interpolant (also

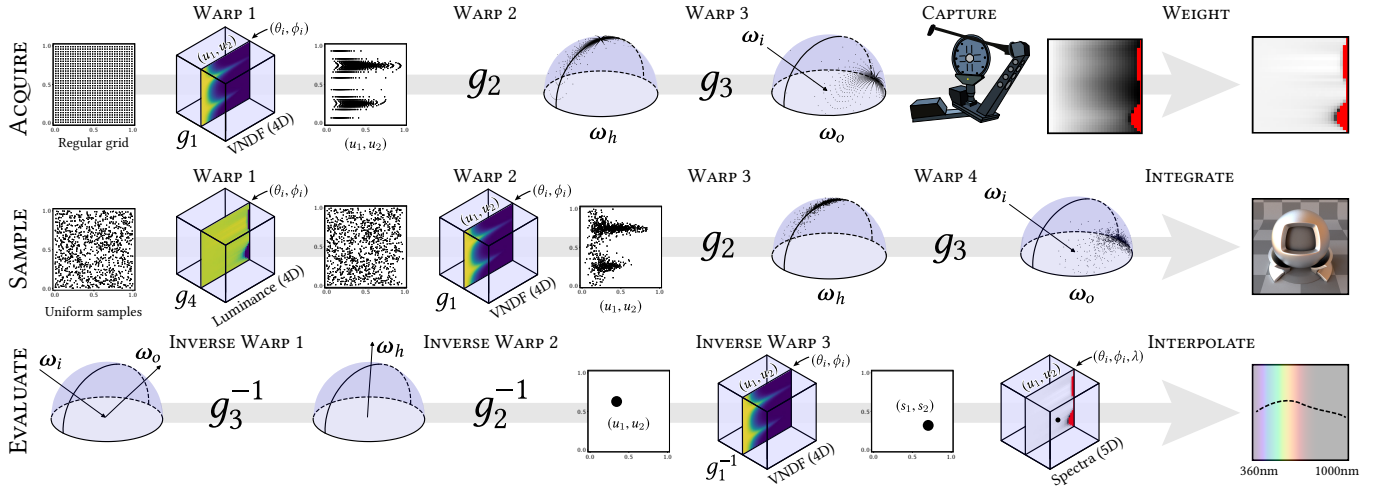


Fig. 5. Illustration of three central operations—acquisition, sampling, and evaluation—performed using the proposed parameterization. **(top)** For a fixed incident direction  $(\theta_i, \phi_i)$ , we warp a regular grid so that its local density matches the material’s visible normal distribution (VNUF) in local coordinates. We then transform these coordinates into normal directions and compute corresponding outgoing directions. The goniometer measures each configuration, resulting in a BRDF slice that is finally weighted by the Jacobian of the parameterization. Red pixels indicate invalid directions that lie below the horizon. **(middle)** To importance sample the BRDF, we first warp uniform variates on  $[0, 1]^2$  so that their density is proportional to the luminance of the associated BRDF slice. The resulting samples pass through the same sequence of transformations used during acquisition. **(bottom)** To evaluate the BRDF or its sampling density, we apply the inverse of these transformations in reverse order: for a given pair of angles, we compute the half-direction vector and map it to local coordinates. An inverse warp by the NDF yields the “random numbers”  $(s_1, s_2)$  that would have generated the given outgoing direction. Finally, we use the five-tuple  $(\theta_i, \phi_i, s_1, s_2, \lambda)$  to perform a lookup into a spectral 5-tensor storing the measurement and weight the linearly interpolated value by the inverse Jacobian of the mapping.

defined on  $\mathcal{U}^2$ ). There are a number of ways in which such a scheme can be constructed, but we have found that not all of them work equally well for our application. In Figure 6, we show several options we experimented with: (a) hierarchical sample warping [Clarberg et al. 2005] introduces discontinuities that are problematic for interpolation. Another common approach involves sampling from the marginal followed by the conditional distribution, in both cases using the inversion method. (b) A naive implementation that samples the row and column followed by the fractional position within the chosen bilinear patch produces the correct density, but the resulting mapping is not smooth enough. (c) A “rigorous” implementation that directly samples continuous versions of the marginal and conditional distributions results in a smooth warping scheme and visibly improved results. Note that this approach easily generalizes to density functions  $f(u_1, u_2 | p_1, \dots, p_n)$  that are conditional on additional parameters  $p_1, \dots, p_n$  as all steps are linear.




Fig. 6. BRDF slices of an anisotropic GGX microfacet conductor ( $\alpha_u = 0.05$ ,  $\alpha_v = 0.3$ ,  $\theta_i = 0$ ,  $\phi_i = 0$ ) generated using multiple warping schemes. **(a)** Hierarchical. **(b)** Naive marginal-conditional. **(c)** Rigorous marginal-conditional. **(d)** Slice from (c) with inverse Jacobian weighting term applied.

**Warp 1: VNUF.** The first invertible mapping  $g_1 : \mathcal{U}^2 \rightarrow \mathcal{U}^2$  applies a marginal-conditional warp to transform uniformly distributed samples  $(s_1, s_2) \in \mathcal{U}^2$  into ones that are distributed according to the VNUF  $D_{\text{vis}}(\omega_m, \omega_i)$ . The normal direction  $\omega_m$  is represented in local coordinates discussed next.

**Warp 2: normal parameterization.** We use the following function to map from  $\mathcal{U}^2$  to the space of microfacet normals in spherical coordinates.

$$g_2(u_1, u_2) = \left[ \frac{\pi}{2} u_1^2 \right], \quad g_2^{-1}(\theta_m, \phi_m) = \left[ \frac{\sqrt{\frac{2}{\pi} \theta_m}}{\frac{1}{2\pi}(\phi_m + \pi)} \right]. \quad (15)$$

This function was designed to concentrate samples around the normal direction, which facilitates storage of highly specular NDFs with few samples. Note that the sample placement is not uniform on the sphere, hence a Jacobian correction factor of  $\sqrt{8\pi^3 \theta_m} \sin \theta_m$  must be applied to interpolants used for sample generation.

We also rely on the mapping  $g_2$  during the retro-reflection measurement. Here, we simply map a regular grid through  $g_2$  and measure the amount of retro-reflection for each direction.

**Warp 3: Specular reflection.** The third mapping  $g_3$  determines the outgoing direction  $\omega_o$  by specular reflection of the incident direction  $\omega_i$  over the microfacet normal  $\omega_m$

$$g_3(\omega_m) = 2(\omega_m \cdot \omega_i) \omega_m - \omega_i, \quad g_3^{-1}(\omega_o) = \frac{\omega_i + \omega_o}{\|\omega_i + \omega_o\|}. \quad (16)$$

For completeness, conversion from spherical to Cartesian coordinates involves the mapping

$$S(\theta_m, \phi_m) = \begin{bmatrix} \sin \theta_m \cos \phi_m \\ \sin \theta_m \sin \phi_m \\ \cos \theta_m \end{bmatrix}, \quad S^{-1}(\omega_m) = \begin{bmatrix} \cos^{-1} z_m \\ \text{atan2}(y_m, x_m) \end{bmatrix}.$$

*Acquisition.* Figure 3.3 (top) illustrates the mappings we employ during acquisition: starting from positions  $\mathbf{u}_j \in \mathcal{U}^2$  arranged on a regular grid, we retrieve outgoing directions as:

$$\omega_o^{(i)} = g_3(g_2(g_1(\mathbf{u}_j; \omega_i))).$$

We subsequently scale the measured values  $f_r^\perp$  by the inverse Jacobian of the mapping and commit to disk. By construction, Monte Carlo sampling of our model is possible by simply replacing  $\mathbf{u}_i$  by a uniform variate  $\xi \in \mathcal{U}^2$  in the above expression, which produces the density (8).

*Warp 4: Luminance warp.* In the context of rendering, the sampling scheme with density (8) is problematic when the discretized BRDF is very non-uniform in our parameterization; this happens when the behavior of the measured BRDF is markedly different from a microfacet BRDF. Therefore, we introduce an optional fourth mapping, which we use to reduce variance in such a case. We precompute a 4D tensor storing the luminance of captured spectral data and use it to perform another marginal-conditional warp  $g_4$ ; Figure 5 (second row) illustrates this strategy. The improved Monte Carlo sampler then uses the sequence

$$\omega_o = g_3(g_2(g_1(g_4(\xi); \omega_i))).$$

to importance sample the BRDF proportional to its luminance.

### 3.4 Parameterization of the Incident Domain

So far, we introduced the mapping from  $\mathcal{U}^2$  to outgoing directions, given an incident direction. We also adapt to the material behavior in the incident direction domain, leveraging additional insights from microfacet theory.

*Incident elevation parameterization.* For the incident elevation angle, we apply the inversion method to the projected area term  $\sigma$  from Equation (7). We thus effectively warp the elevation between a uniform and cosine-weighted angular distribution (specular materials tend toward cosine-like distributions, while diffuse materials are more uniform). This has the effect of concentrating more samples at low elevation angles to capture Fresnel-related grazing angle variation found in specular materials.

*Incident azimuth parameterization.* When dealing with analytic microfacet distributions such as GGX or Beckmann, the VNDF scheme involves a warping that is applied to the microfacet normals—and its dual—to the incident direction. We found a way to mimic this behavior on a tabulated NDF by applying the following transformation:

$$\phi_i(\mathbf{u}) = 2\pi g_1(\mathbf{u}, \omega_z) - \pi, \quad \phi_i^{-1}(\phi) = g_1^{-1}\left(\frac{1}{2\pi}(\phi - \pi), \omega_z\right), \quad (17)$$

where  $\omega_z = (0, 0, 1)$ ; Figure 7 illustrates the distribution of azimuthal angles we obtain for a highly anisotropic material. Notice how the distribution concentrates in a dual form with respect to that of the




Fig. 7. Our method relies on an adaptive parameterization to discretize the BRDF over incident directions. The black dots in this figure represent sample locations for an anisotropic GGX conductor with roughness parameter  $\alpha = (0.05, .2)$ .

anisotropy. We found azimuthal adaptivity to be essential to acquire anisotropic materials at low resolution.

*Optimization for isotropic materials.* Due to their rotational symmetry, isotropic materials do not require a full measurement of the 4D domain. Note that this symmetry also manifests in the NDF—we therefore reduce the dimension of both acquisition steps for isotropic samples.

## 4 EXPERIMENTAL SETUP

In this section, we show how to leverage the theory from Section 3 to adaptively measure BRDFs using a motorized gonio-photometer. We discuss our experimental setup and the modifications we made to the gonio-photometer, which were required by our method.

*Gonio-photometer.* Our acquisition pipeline relies on a modified PAB pgII [2018] gonio-photometer; Figure 8 shows an overview of the setup including modifications for retro-reflective measurements. This device illuminates a flat vertically mounted sample from a light source that is mounted one of two stationary rails (we use both for the two phases of our pipeline). Changing the incident light direction entails rotating the motorized sample mount around the vertical axis and the normal direction of the sample. A sensor head mounted on a motorized two-axis arm measures the illumination that is scattered from the sample. The machine has two “dead spots” that cannot be measured with the combination of sensor arm and light source: first, the sensor arm cannot view the sample from straight below, as this space is occupied by the pedestal holding the sample mount. This is usually not a problem, since the BRDF scatters little light in these directions, which can easily be reconstructed from nearby samples. Second, shadowing becomes unavoidable when the sensor is within approximately 3 degrees of the beam. This is problematic since it makes any kind of retro-reflection measurement impossible. Instead of using the sensor arm, we therefore rotate it out of the way and measure the retro-reflection using a sensor that is directly mounted on the light rail (this is feasible since the retro-reflection measurement is only two dimensional, hence the two motorized axes of the sensor arm are not needed). We then fill in all missing regions by solving a Poisson problem on the four-dimensional domain.




Fig. 8. Our measurement setup relies on a pgl gonio-photometer with a customized lighting setup for retro-reflective measurements. During the first measurement phase, a laser (mounted on the left light rail) illuminates a vertically mounted sample in a separate room partition. On the way to the sample, the laser beam passes through a beamsplitter, which redirects retro-reflectively scattered light onto a silicon photodiode. A broadband light source illuminates the sample from the right rail during the second measurement phase, and a spectrometer on the pgl’s main sensor head records scattered illumination in the 360-1000nm range. **(left)**: schematic view. **(middle)**, **(right)**: photographs of the illumination path and motorized parts of the device.

#### 4.1 Retro-Reflection Measurement

To acquire a materials’ retro-reflective properties, we illuminate the sample with a 532nm 40mW diode-pumped solid-state laser. This facilitates calibration, measurement of high-frequency retro-reflection, and ensures that the intense unscattered beam is in perfect alignment when it reaches the subsequent optical elements.

The beam is first reflected by a pair of kinematic mirror mounts, which enable precise calibration of both offset and direction. Next, it passes through a Faraday isolator, which shields the laser cavity from destabilizing feedback caused by retro-reflected light. The beam is then widened from 0.3 to 1.5cm using Galilean beam expander ( $5\times$ ) to reduce laser speckle, which would add unwanted noise to the measured NDFs.

The laser passes through a polarization-insensitive 50:50 plate beamsplitter, which reflects half of the energy into a beam dump (i.e. an absorptive cavity) and transmits the other portion towards the sample. Light that is retro-reflected from the sample travels back into the beam splitter, which reflects  $\sim 50\%$  of the light onto a lens that focuses it onto the receptive area of a reverse-biased silicon photodiode. The photodiode is used in photoconductive mode, i.e. it produces a faint current on the order of a few nano-Amperes that is linearly proportional to the received illumination. This current is then amplified and digitized by a Picoammeter (Keithley 6485) that is read out by a connected workstation. This process is very sensitive to stray light, hence the retro-reflective beam path must be shielded.

We noted that the choice of beamsplitter has a significant effect on the feasibility of this setup. We initially used a cube beamsplitter, whose optical cement leaked a large portion of the unscattered laser light towards the photodiode. We had also initially mounted the beam expander after the beamsplitter, which caused significant retro-reflection from the expander’s two lens elements despite their optical coating. Intrinsic retro-reflection from the measurement setup is undesirable because it requires a large constant term to be subtracted from the measurement, causing cancellation that reduces measurement accuracy. The use of a plate beamsplitter with no further optical elements downstream resolved this issue.

Both the laser source and the white light source used in the main measurement process are controlled via electronic shutters that open and close in the respective measurement phases. We found this preferable to controlling their AC power, as the light sources (including the laser) are not stable enough to be used for measurements before reaching their operating temperature.

For mirror-like materials (e.g. `chm_orange` in Figure 10), the tails of the acquired NDF are many orders of magnitude below the peak of the distribution and tend to be noisy. We apply a Gaussian filter to the tails in a few such cases.

#### 4.2 Measurements Over the 4D Domain

The main acquisition phase illuminates the sample with broadband illumination produced by a xenon arc light source (Osram XBO 150W). A condenser lens concentrates the divergent beam onto a pinhole, whose output is then focused onto the sample. We record the illumination scattered from sample using a fiber-coupled Zeiss MMS1 spectrometer installed on the gonio-photometer’s sensor arm. This spectrometer uses a reflective grating to separate wavelengths onto a co-located NMOS photodiode array yielding a spectral resolution of approximately 3.3 nanometers digitized at 16 bit precision. The spectrometer can nominally record wavelengths between 310-1100nm, although the product of the emission spectrum and spectrometer sensitivity restrict the usable range to roughly 360-1000nm (Figure 9).




Fig. 9. Relative spectral sensitivity of our measurement setup. Sharp peaks are caused by the XBO light source emission spectrum and are removed after calibration.



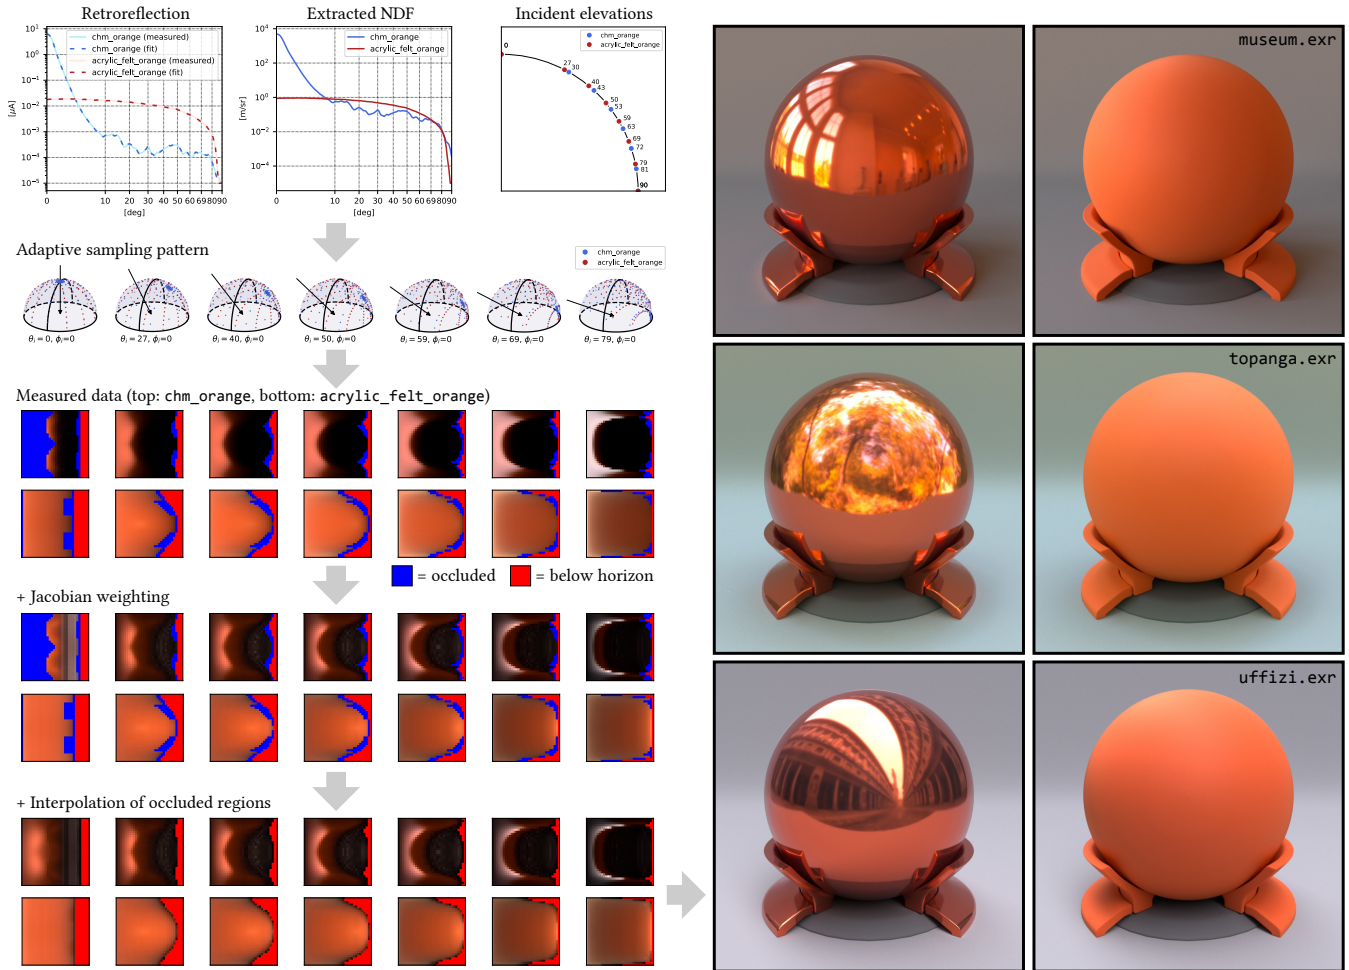


Fig. 10. The stages of our acquisition pipeline for two different materials: starting from a set of retro-reflective measurements, we extract a compatible microfacet NDF; dashed lines in the top left plot show that the fit is in good agreement with the acquired data. This NDF is then used to determine a suitable sampling pattern for both incident (top right) and outgoing directions (second row). The measured data (third row) is weighted by the Jacobian of the transformation (fourth row). The BRDF slices are still incomplete since some configurations are shadowed or cannot be reached by the gonio-photometer (blue regions). We solve a Poisson problem to fill in the missing values. Finally, we are able to render the materials efficiently in a Monte Carlo path-tracer.

Like any light source that is based on stimulated emission, the XBO spectrum contains many sharp peaks that must be accounted for in calibration: we do so by measuring a Spectralon reference target (SRS-99-020) with known reflectance curves before each measurement and dividing all subsequent measurements by the resulting correction factor. Specular materials require extra precautions due to their severe dynamic range: we manually pick an upper exposure time for all samples (typically 512 ms) and repeatedly halve the exposure time of overexposed configurations until successful.

## 5 RESULTS

In this section, we present the measurements performed thus far, along with validations against synthetic measurements.

*Measured Data.* Figure 10 shows each stage of our acquisition pipeline for two materials with different reflectance properties: one

is highly specular (`chm_orange`), while the other is isotropic and mostly diffuse (`acrylic_felt_orange`). These two different behaviors are already noticeable from the retro-reflective measurements, which results in markedly different NDFs. Our adaptive parameterization therefore picks different locations for both the incident elevation angles and outgoing directions. This results in smooth data in parametric space, which is the key feature that allows us to sample our materials at low resolutions.

*Resolution and Memory.* The spectrometer nominally samples the 310-1100nm range using 256 approximately equi-spaced bins, of which 193 covering the 360-1000nm range have sufficient sensitivity to yield usable measurements. We decided to include the near-infrared (NIR) range in our dataset due to renewed interest on NIR imaging [Dümbgen et al. 2018]. Note that there is increased




Fig. 11. Figure 1 rendered using several emission spectra and sensitivities of a RGB camera without IR filter. All images were white-balanced.

noise in portions of the NIR where the sensitivity is comparably low, hence this data will potentially require further processing.

We acquire isotropic materials with a retro-reflective scan of 128 angles at 532nm followed by spectral acquisition of  $32 \times 16$  adaptively sampled outgoing directions per slice. The bottom half of each slice is a mirror copy of the top and is inferred using the bilateral symmetry of isotropic materials. We capture slices for 7 incident elevations, and extrapolate the eighth corresponding to perpendicular incidence ( $\theta_i = 90^\circ$ ).

For anisotropic materials, the preliminary retro-reflection acquisition pass consists of 64 elevations and 64 azimuthal orientations, and we spectrally measure  $32 \times 32$  adaptively sampled outgoing directions over 7 incident elevations and 16 azimuthal orientations. We also tried to exploit symmetries—for instance, the behavior of anisotropic microfacet BRDF can be fully reconstructed from measurements on a reduced interval  $\phi_i \in [0, \frac{\pi}{4}]$ . However, we found that this property did not transfer to real-world materials and introduced considerable error in measurements, hence we do not rely on any symmetries in our representation of anisotropic materials. Acquisition then consists of  $128 + 3,584$  measurements for isotropic materials and  $4096 + 114,688$  for anisotropic ones.

*Material database.* Thus far, we acquired spectral BRDFs of 36 different materials, of which four are anisotropic. We plan to extend this dataset over time so that it covers all major material classes. This database includes many materials that are outside of the realm of classical microfacet theory including: materials with iridescent flakes, opalescent and color-changing paints, and anisotropic cloth (Figure 12). Several material swatches (spray paint on foam core) were provided by Industrial Light & Magic (ILM)—they are interesting because they were made to imitate various materials but behave subtly differently. For instance, `ilm_13_37_metallic`, unlike actual metals, is oscillatory in the spectral domain.

*Spectral rendering.* All renderings were made in a custom full-spectral renderer that integrates over a continuous spectral domain using Monte Carlo sampling. Our BRDF plugin implements BRDF evaluation, a density function, and importance sampling as required by path tracing with multiple importance sampling. Such functions are straightforward to implement using our parameterization and only require a few lines of code. We will release an open source




Fig. 12. Our parameterization also works with materials outside of the scope of microfacet theory. This rendering shows a highly anisotropic silk material with two differently-colored yarns.

reference implementation to enable integration of our captured BRDFs in other rendering systems.

Our materials can be rendered with arbitrary illumination sources including high-frequency fluorescent excitation spectra or even near-infrared illumination (Figure 11).

Figure 1 and Figure 13 show path-traced renderings of the newly acquired material. Thanks to our parameterization efficient importance sampling is available by construction. As a result, we are able to render noise-free images using 1024 samples per pixel, which is something that previous measured BRDF representations could only achieve using additional data-structures and/or pre-computation.

## 5.1 Process Validation

*Methodology.* Before implementing our hardware-based acquisition scheme, we validated our numerical algorithms by performing virtual acquisitions that resampled existing BRDF models according to our parameterization. We experimented with two types of existing BRDF models, which we validated by comparing path-traced renderings of both the input BRDF and our reconstruction.

*Validation Against Microfacet BRDFs.* First, we used microfacet BRDF models. They represent an ideal scenario for our parameterization since the input BRDF exactly matches the assumptions underlying our parameterization. Figure 16 shows a few comparative renderings of our re-sampling experiments against GGX-based microfacet BRDFs.

*Comparison Against the MERL Parameterization.* Figure 14 shows an RMS error plot comparing the (non-adaptive) MERL discretization to our method. The same number of samples used in both cases ( $\theta_h \times \theta_d \times \phi_d = 16 \times 16 \times 28 = 7168$  in the case of MERL.) While both parameterizations perform equally well for rough materials, the error of our parameterization remains relatively constant, while the RMS error of MERL rapidly increases as the BRDF becomes more directionally peaked, demonstrating the benefits of measurements that adapt to the input material.

*Validation Against Measured BRDFs.* Second, we validated our technique against measured BRDFs from both the MERL and UTIA databases, shown in Figure 17. For this, we used the same resolution



Fig. 13. Renderings of the materials we have acquired so far, using 1024 samples per pixel.




Fig. 14. RMS error of a discretized GGX microfacet model, using our parameterization (blue) and the MERL parameterization (red). The MERL sampling pattern is non-adaptive and produces a large error for low roughness values, while adaptation in our method leaves the error approximately constant.




Fig. 15. Mean delta-E difference image on the entire MERL database.

also used in the hardware-based acquisition. Since the original MERL data is stored at  $90 \times 90 \times 360$  resolution, we effectively compress it by two orders of magnitude with little to no loss in rendering quality; our supplemental material provides a detailed analysis of the database.

Burley et al. [2012] observed extrapolation and measurement artifacts in a number of materials that are part of the MERL database. Our experiments appear to confirm these observations: visualizing slice plots using our parameterization reveals strong discontinuities, color oscillation, and non-symmetries that violate isotropy. Figure 18 shows a few such slices taken from our supplemental material.

*Comparison against fitting techniques.* Our representation of the MERL database yields a significant compression of the original data, hence our method could also find application in situations where fitting techniques have traditionally been used. We compare to the state-of-the-art fitting method of Bagher et al [2016], which converts BRDFs into a products of compact lower-dimensional factors.

Figure 15 compares the mean DE00 errors over the MERL database, and the supplemental material contains renderings and error plots for individual materials. These results show that our parameterization introduces less bias, albeit with higher storage requirements (approximately one order of magnitude). Bagher et al.’s method will thus be preferable in cases where a very low storage footprint is desirable. That said, Bagher et al.’s method incurs some errors due to the assumption of factorizability, which will remain even in the limit of an infinitely fine discretization. Our method also makes assumptions on the material behavior but only uses them to warp a regular sample grid, hence the error will tend to zero as the sampling rate increases.

## 6 CONCLUSION AND FUTURE WORK

We have presented the first adaptive BRDF parameterization that simultaneously targets acquisition, storage, and rendering. We demonstrated the efficiency of our method by introducing the first set of spectral BRDF measurements that contain both anisotropy and high frequency behavior.

We find that the availability of spectral BRDF datasets is timely: as of 2018, major industrial systems render in the spectral domain (e.g. Maxwell, Weta’s Manuka renderer), and two major open source renderers, PBRT and Mitsuba, are presently being redesigned for full-spectral rendering [Pharr and Jakob 2017]. A number of recent works [Dong et al. 2015; Löw et al. 2012; Werner et al. 2017] have also proposed diffraction-based reflectance models with a dependence on wavelength. Our dataset contains many materials with spectrally interesting behavior (multiple types of iridescent, opalescent, and color-changing paints). Other materials that initially appear well-described by geometric optics exhibit oscillatory behavior in the spectral domain, indicating wave-optical interference. We believe that access to a comprehensive spectral dataset will enable future advances in the area of material modeling.

Naturally, our parameterization is not free from limitations, which leaves plenty of avenues for future work.

*Layered materials.* While experimenting with the MERL database, we noticed that our parameterization was less efficient when measuring materials with multi-lobed BRDFs. Such behaviors typically arise in layered and/or composite materials such as acrylics and plastics, which usually do not agree with the predictions of microfacet models. While we can still acquire such materials by increasing the resolution, it would be interesting to improve the bijective mapping of our method based on an extension of microfacet theory that supports layered structures.

*Better outgoing directions.* Our parameterization sometimes maps points to invalid outgoing directions that lie below the horizon; such regions are shown in red in our slice plots. Such situations arise because our parameterization only accounts for a single scattering event, and there always exists a set of microfacets oriented so that they reflect incident directions downwards. An interesting avenue for future work would be to derive a novel importance sampling technique that accounts for multiple scattering within the microfacets, which would effectively remove this effect.

*BTDF and participating media acquisition.* Because we restrict our acquisition to that of BRDFs, we can not faithfully reproduce materials that transmit light into the opposite hemisphere; our *maple-leaf* material is an example where this would be desirable.

## ACKNOWLEDGMENTS

We thank Olesya Jakob for designing the counter scene shown in Figure 1, and Peter Apian-Bennwitz for support with the gonio-photometric measurements. We thank Laurent Belcour and Romain Pacanowski for thoughtful discussions regarding backscattering and BRDF acquisition. We thank ILM and particularly Naty Hoffman and André Mazzone for providing us access to material samples from the Star Wars movie: “Solo: A Star Wars Story”.




Fig. 16. Side-by-side comparisons of GGX microfacet BRDFs against our resampled BRDFs.




Fig. 17. Side-by-side comparisons of (top row) MERL and (bottom row) UTIA BRDFs against our resampled BRDFs.

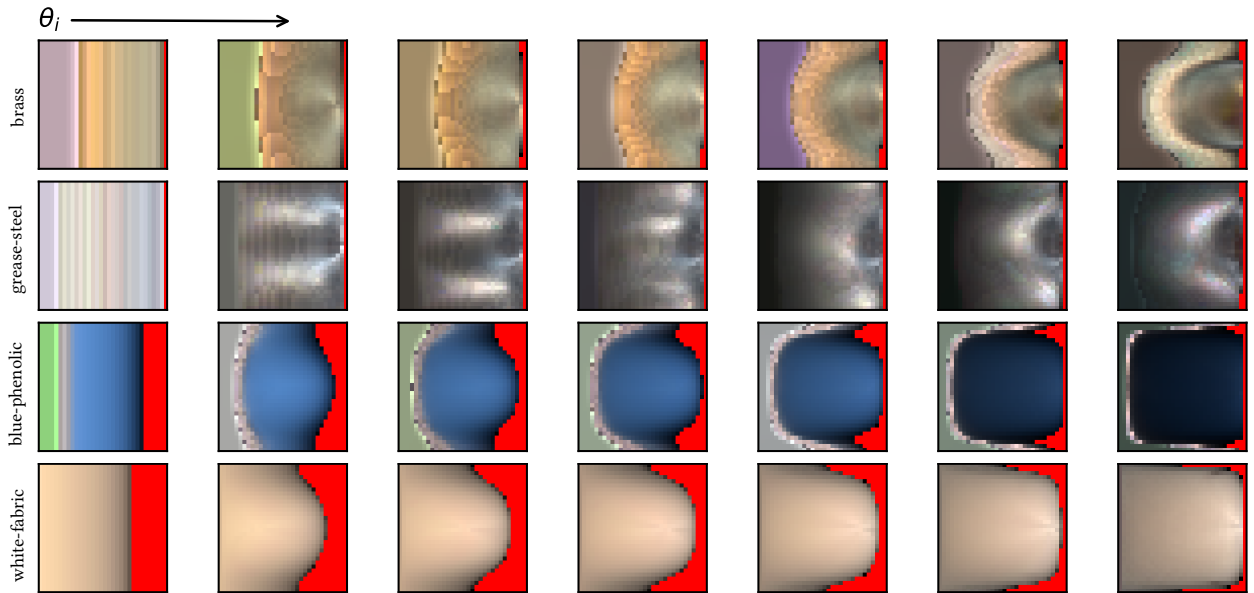


Fig. 18. Visualizing the MERL materials through our parameterization reveals discontinuities, color oscillation, and non-symmetries in violation of isotropy.

## REFERENCES

- Peter Apian-Bennet. 2013. Building material example BSDF data. <http://www.pab.eu/gonio-photometer/demodata/bme/>. (2013). Accessed: 2018-01-09.
- Michael Ashikhmin and Simon Premoze. 2007. Distribution-based BRDFs. *Technical Report, University of Utah* (2007).
- Mahdi M. Bagher, John Snyder, and Derek Nowrouzezahrai. 2016. A Non-Parametric Factor Microfacet Model for Isotropic BRDFs. *ACM Trans. Graph.* 35, 5 (2016), 159:1–159:16.
- M. M. Bagher, C. Soler, and N. Holzschuch. 2012. Accurate fitting of measured reflectances using a Shifted Gamma micro-facet distribution. *Comput. Graph. Forum* 31, 4 (2012), 1509–1518.
- Pascal Barla, Laurent Belcour, and Romain Pacanowski. 2015. In praise of an alternative brdf parametrization. In *Workshop on Material Appearance Modeling*.
- Laurent Belcour and Pascal Barla. 2017. A Practical Extension to Microfacet Theory for the Modeling of Varying Iridescence. *ACM Transactions on Graphics* 36, 4 (July 2017), 65.
- James F Blinn. 1977. Models of light reflection for computer synthesized pictures. In *ACM SIGGRAPH Computer Graphics*, Vol. 11. ACM, 192–198.
- Marilyne Andersen Boris Karamata. 2013. Revisiting parallel catadioptric goniophotometers. *Proc.SPIE* 8788 (2013), 8788 – 8788 – 11.
- Pierre Bouguer. 1760. *Traité d'optique sur la gradation de la lumière*. De l'imprimerie de HL Guerin & LF Delatour.
- Brent Burley. 2012. Physically-based shading at Disney. In *ACM SIGGRAPH 2012 Courses (SIGGRAPH '12)*. ACM, New York, NY, USA.
- Petrik Clarberg, Wojciech Jarosz, Tomas Akenine-Möller, and Henrik Wann Jensen. 2005. Wavelet Importance Sampling: Efficiently Evaluating Products of Complex Functions. *ACM Trans. Graph.* 24, 3 (July 2005), 1166–1175.
- Kristin J Dana, Bram Van Ginneken, Shree K Nayar, and Jan J Koenderink. 1999. Reflectance and texture of real-world surfaces. *ACM Transactions On Graphics (TOG)* 18, 1 (1999), 1–34.
- Zhao Dong, Bruce Walter, Steve Marschner, and Donald P. Greenberg. 2015. Predicting Appearance from Measured Microgeometry of Metal Surfaces. *ACM Trans. Graph.* 35, 1, Article 9 (Dec. 2015), 9:1–9:13 pages.
- Frederike Dümbsgen, Majed El Helou, Sabine Susstrunk, and Natalija Gucevska. 2018. Near-Infrared Fusion for Photorealistic Image Dehazing. *IS&T EI Proceedings* (2018).
- Jonathan Dupuy, Eric Heitz, and Eugene d'Eon. 2016. Additional progress towards the unification of microfacet and microflake theories. In *Proceedings of the Eurographics Symposium on Rendering: Experimental Ideas & Implementations*. Eurographics Association, 55–63.
- Jonathan Dupuy, Eric Heitz, Jean-Claude Iehl, Pierre Poulin, Fabrice Neyret, and Victor Ostromoukhov. 2013. Linear Efficient Antialiased Displacement and Reflectance Mapping. *ACM Trans. Graph.* 32, 6, Article 211 (Nov. 2013), 211:1–11 pages.
- Jonathan Dupuy, Eric Heitz, Jean-Claude Iehl, Pierre Poulin, and Victor Ostromoukhov. 2015. Extracting Microfacet-based BRDF Parameters from Arbitrary Materials with Power Iterations. In *Computer Graphics Forum*, Vol. 34. Wiley Online Library, 21–30.
- Jiri Filip and Radomir Vávra. 2014. Template-Based Sampling of Anisotropic BRDFs. In *Computer Graphics Forum*, Vol. 33. Wiley Online Library, 91–99.
- Jiri Filip, Radomir Vávra, and Michal Havlicek. 2014. Effective acquisition of dense anisotropic BRDF. In *Pattern Recognition (ICPR), 2014 22nd International Conference on*. IEEE, 2047–2052.
- Abhijeet Ghosh, Shruithi Achutha, Wolfgang Heidrich, and Matthew O'Toole. 2007. BRDF acquisition with basis illumination. In *Computer Vision, 2007. ICCV 2007. IEEE 11th International Conference on*. IEEE, 1–8.
- Dar'ya Guarnera, Giuseppe Claudio Guarnera, Abhijeet Ghosh, Cornelia Denk, and Mashhuda Glencross. 2016. Brdf representation and acquisition. In *Computer Graphics Forum*, Vol. 35. Wiley Online Library, 625–650.
- 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.
- Eric Heitz and Eugene d'Eon. 2014. Importance Sampling Microfacet-Based BSDFs using the Distribution of Visible Normals. In *Computer Graphics Forum*, Vol. 33. 103–112.
- Eric Heitz, Johannes Hanika, Eugene d'Eon, and Carsten Dachsbacher. 2016. Multiple-scattering microfacet BSDFs with the Smith model. *ACM Transactions on Graphics (TOG)* 35, 4 (2016), 58.
- Nicolas Holzschuch and Romain Pacanowski. 2017. A two-scale microfacet reflectance model combining reflection and diffraction. *ACM Transactions on Graphics* 36, 4 (2017), 12.
- Wenzel Jakob, Miloš Hašan, Ling-Qi Yan, Jason Lawrence, Ravi Ramamoorthi, and Steve Marschner. 2014. Discrete Stochastic Microfacet Models. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2014)* 33, 4 (2014).
- Joakim 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), 9:1–14.
- Stephen R. Marschner, Stephen H. Westin, Eric P. F. LaFortune, and Kenneth E. Torrance. 2000. Image-based bidirectional reflectance distribution function measurement. *Applied Optics* 39, 16 (2000), 2592–2600.
- Stephen R. Marschner, Stephen H. Westin, Eric P. F. LaFortune, Kenneth E. Torrance, and Donald P. Greenberg. 1999. Image-based BRDF Measurement Including Human Skin. In *Proceedings of the 10th Eurographics Conference on Rendering (EGWR '99)*. Eurographics Association, 131–144.
- Phillip R. Mattison, Mark S. Dombrowski, James M. Lorenz, Keith J. Davis, Harley C. Mann, Philip Johnson, and Bryan Foos. 1998. Handheld directional reflectometer: an angular imaging device to measure BRDF and HDR in real time. *Proc.SPIE* 3426 (1998), 3426 – 3426 – 12.
- Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. 2003a. A Data-Driven Reflectance Model. *ACM Trans. Graph.* 22, 3 (July 2003), 759–769.
- Wojciech Matusik, Hanspeter Pfister, Matthew Brand, and Leonard McMillan. 2003b. Efficient Isotropic BRDF Measurement (EGSR). Eurographics Association, 241–247.
- Addy Ngan, Frédo Durand, and Wojciech Matusik. 2005. Experimental Analysis of BRDF Models. In *Proceedings of the Eurographics Symposium on Rendering*. Eurographics Association, 117–226.
- Fred E. Nicodemus, Joseph C. Richmond, J.J. Hsia, W.I. Ginsberg, and T. Limperis. 1977. Geometrical Considerations and Nomenclature for Reflectance. *Applied Optics* 9 (1977), 1474–1475.
- Jannik Boll Nielsen, Henrik Wann Jensen, and Ravi Ramamoorthi. 2015. On Optimal, Minimal BRDF Sampling for Reflectance Acquisition. *ACM Transactions on Graphics (TOG)* 34, 6 (November 2015), 186:1–186:11.
- PAB. 2018. pab advanced technologies Ltd. <http://www.pab.eu>. (2018). Accessed: 2018-01-09.
- Matt Pharr and Wenzel Jakob. 2017. personal communication. (Dec. 2017).
- Andrei D Polyinin and Alexander V Manzhurov. 2012. *Handbook of Integral Equations*. CRC Press.
- Ju Ren and Jianlin Zhao. 2010. Measurement of a bidirectional reflectance distribution and system achievement based on a hemi-parabolic mirror. *Opt. Lett.* 35, 9 (May 2010), 1458–1460.
- Szymon M Rusinkiewicz. 1998. A New Change of Variables for Efficient BRDF Representation. In *Rendering Techniques '98*. Springer, 11–22.
- Lionel Simonot and Gael Obein. 2007. Geometrical considerations in analyzing isotropic or anisotropic surface reflections. *Applied optics* 46, 14 (2007), 2615–2623.
- Michael M Stark, James Arvo, and Brian Smits. 2005. Barycentric parameterizations for isotropic BRDFs. *IEEE transactions on visualization and computer graphics* 11, 2 (2005), 126–138.
- K. E. Torrance and E. M. Sparrow. 1967. Theory for Off-Specular Reflection From Roughened Surfaces. *J. Opt. Soc. Am.* 57, 9 (Sep 1967), 1105–1112.
- 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.
- Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. 2007. Microfacet Models for Refraction Through Rough Surfaces. In *Computer Graphics Forum (Proc. Eurographics Symposium on Rendering) (EGSR)*. 195–206.
- Gregory J Ward. 1992. Measuring and modeling anisotropic reflection. *ACM SIGGRAPH Computer Graphics* 26, 2 (1992), 265–272.
- Sebastian Werner, Zdravko Velinov, Wenzel Jakob, and Matthias Hullin. 2017. Scratch Iridescence: Wave-Optical Rendering of Diffractive Surface Structure. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 36, 6 (Oct. 2017).
- D Rod White, Peter Saunders, Stuart J Bonsey, John van de Ven, and Hamish Edgar. 1998. Reflectometer for measuring the bidirectional reflectance of rough surfaces. *Applied optics* 37, 16 (1998), 3450–3454.