

# A Non-Parametric Factor Microfacet Model for Isotropic BRDFs

Mahdi M. Bagher<sup>1,2</sup> John Snyder<sup>2</sup> Derek Nowrouzezahrai<sup>1</sup>  
<sup>1</sup>University of Montreal <sup>2</sup>Microsoft Corporation

We investigate the expressiveness of the microfacet model for isotropic BRDFs measured from real materials by introducing a *non-parametric factor model* that represents the model’s functional structure but abandons restricted parametric formulations of its factors. We propose a new objective based on *compressive weighting* that controls rendering error in high dynamic range BRDF fits better than previous factorization approaches. We develop a simple numerical procedure to minimize this objective and handle dependencies that arise between microfacet factors. Our method faithfully captures a more comprehensive set of materials than previous state-of-the-art parametric approaches, yet remains compact (3.2KB per BRDF). We experimentally validate the benefit of the microfacet model over a naïve orthogonal factorization, and show that fidelity for diffuse materials is modestly improved by fitting an unrestricted shadowing/masking factor. We also compare against a recent data-driven factorization approach [Bilgili et al. 2011] and show that our microfacet-based representation improves rendering accuracy for most materials while reducing storage by more than 10 $\times$ .

CCS Concepts: •Computing methodologies  $\rightarrow$  Reflectance modeling;

Additional Key Words and Phrases: BRDF compression and factorization, microfacet theory, non-parametric models, robust statistics

## ACM Reference Format:

Mahdi M. Bagher, John Snyder, and Derek Nowrouzezahrai. 2016. A Non-Parametric Factor Microfacet Model for Isotropic BRDFs. *ACM Trans. Graph.* V, N, Article XXX (Month 2016), 15 pages.

## 1. INTRODUCTION

In addition to the accurate simulation of global light transport, synthesizing realistic images requires true-to-life models of local reflectance. The Cook-Torrance microfacet model [Cook and Torrance 1981] is one such model ubiquitous in computer graphics, that represents the local reflection from a surface as the product of lower-dimensional factors derived from the statistical distribution of micro-scale geometry. We investigate the validity of this model from real-world measurements, using the MERL database [Matusik et al. 2003] of *isotropic* reflection. This neglects *anisotropic* materials with directionally aligned surface micro-geometry like scratches (on brushed metal) or fibers (in hair or fabric), but still comprises a wide and interesting class of appearance.

We test whether the microfacet model spans the space of isotropic reflectance or at least provides benefit over other models,

and whether assumptions commonly made in applying this model reduce its ability to capture real materials. In particular, we focus on the question of how existing parametric models used for its factors [Cook and Torrance 1981; Walter et al. 2007; Bagher et al. 2012; Löw et al. 2012] may limit its representational power.

To do this, we devise a new representation and a model fitting strategy based on the underlying structure of the microfacet model. By separating the model’s functional structure from specific, constrained forms of its factors, we show the resulting generalized representation is still compact and improves rendering fidelity for a more comprehensive set of real-world materials. We call it the *non-parametric factor microfacet model*. The factors we obtain are complex and not realizable by previous parametric models (Figure 6).

Data-driven factorization [McCool et al. 2001; Lawrence et al. 2004; Lawrence et al. 2006; Bilgili et al. 2011] is another alternative to parametric microfacet models for compressing unwieldy BRDF measurements (35MB in double precision for each material in the MERL database). Surprisingly, these methods often fall short of simple parametric fits though they command enormously more fitting variables. There are at least two possible reasons for this. Existing approaches apply generic factorizations not based on microfacet theory. They also apply simple optimization metrics whose fits produce visibly high error when rendered in our experiments (see Figure 7, bottom row). Fitting with many variables is computationally difficult and motivates simple, easy-to-solve optimization metrics, such as a weighted sum of squared errors.

We solve the metric problem by formulating a new objective which automatically specializes itself to each material’s high dynamic range (HDR) measurements. Like previous work in data-driven factorization, it takes a simple weighted sum of squared errors to make minimization tractable, but it combines a new *compressive weight* (Equation 11) with standard material-independent weights based on the BRDF’s parametric volume form (Equation 9) and its importance in the local rendering integral (Equation 10). In essence, we extend techniques in robust statistics to increase fidelity in the BRDF’s darker but broader areas relative to its bright peaks, dramatically reducing error when the fits are rendered. Another way to understand this idea is that rendering fidelity is improved by minimizing relative rather than absolute BRDF fit error.

We also develop a new *alternating weighted least squares* (AWLS) solver to perform the minimization. While similar to previous approaches in tensor approximation [Lawrence et al. 2006; Kolda and Bader 2009; Bilgili et al. 2011], our approach handles nonlinear dependencies between factor variables arising in the microfacet model. It is simple to implement, rapidly converging, robust to random initialization, and completely automatic with no per-material parameter tweaking needed. It yields single-term fits that are consistently more accurate than previous models (see Figure 1), both quantitatively and visually, and requires only 3.2KB per BRDF, a compression factor of more than 5000.

For each color channel, our model comprises three tabulated 1D vectors, one for the shadowing-masking  $G$ , microfacet distribution  $D$ , and Fresnel  $F$  factors, as well as the scalar diffuse and specular coefficients  $\rho_d$  and  $\rho_s$  (Equation 3). We investigate two formula-

---

We acknowledge funding from MITACS, Microsoft Research, and NSERC (Discovery Grant 2012-418519). We thank Maurice Ko for his input, Iliyan Georgiev for help with diagrams, and Stephen Hill for his WebGL base.

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 ACM 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).

© Year ACM. \$15.00

DOI: 10.1145/XXXXXXX.YYYYYY




Fig. 1. Comparing (left to right) [Bagher et al. 2012], [Löv et al. 2012], [Bilgili et al. 2011], our method, to BRDF measurements. The inset visualizes squared error. The Buddha uses GREEN-METALLIC-PAINT, the bumpy sphere is CHROME-STEEL, and the floor is POLYETHYLENE [Matusik et al. 2003]. Compared to parametric approaches and data-driven non-microfacet factorization, our fit better matches ground-truth. Note the lack of brightness in the floor’s background, the incorrect hue and bright underside of the bumpy sphere, the washed out appearance of Löv’s GREEN-METALLIC-PAINT fit, and the failure of [Bilgili et al. 2011] to improve upon a simple parametric model (visually and in PSNR) despite its vastly increased degrees of fitting freedom.

tions for  $G$ : one that determines it from  $D$  using a standard shadowing model [Smith 1967] ( $G$ -from- $D$ ), and another that is completely data-driven (*independent-G*). We also compare a *naïve model* not based on microfacet theory: a simple orthogonal factorization in the measurement coordinates (Equation 4). Though it fits with only half as many optimization variables, G-from-D provides rendering results comparable to the naïve model while the independent-G method provides a superior rendering match using 3/4 as many variables. The independent-G model also yields a moderate increase in rendering fidelity compared to the G-from-D model for many of the more diffuse materials. These results demonstrate the benefit of the microfacet model’s non-orthogonal but physically-motivated factorization, fit using a robust metric that compensates for measurement magnitude.

Our findings (e.g. Figure 1) confirm that existing parametric models are too weak to capture a broad range of realistic reflectance, and that generalizing to non-parameterized factors improves rendering fidelity and captures a larger set of materials. Furthermore, the combination of our representation and fitting metric improves rendering results even compared to state-of-the-art data-driven factorizations with over  $10\times$  as many parameters. Only a few MERL materials are not visually captured by our non-parametric model; we show that this is because their specular behavior can not be factored into the independent product of the normal distribution ( $D$ ) and Fresnel ( $F$ ) functions as microfacet theory predicts. This limitation is fundamental and can not be overcome by substituting more sophisticated parametric models for its factors.

## 2. PREVIOUS WORK

Appearance modeling spans four decades of research and we discuss only the most relevant previous work. We forward readers to Dorsey et al.’s comprehensive review [2008].

**Data-driven models.** BRDF data can be directly captured with a gonireflectometer [Marschner et al. 1999]. Matusik et al.’s [2003] MERL dataset comprises a large set of (100) isotropic BRDFs captured at high angular resolution (1.5 million samples), although sampling noise and the well-known difficulty measuring at grazing angles require corrective post-processing. Burley has further analyzed error in this dataset and conjectured additional sources (e.g., cloth wrinkles when laying over a sphere, some materials like GREASE-COVERED-STEEL are almost cer-

tainly anisotropic, etc.) [Burley 2012; McAuley et al. 2012]. Recognizing both the limitations and benefits of physical measurement, we test the expressive power of the microfacet model using this dataset, assuming it as ground truth. We show that a carefully designed weighting scheme (Section 4.1) is necessary to make BRDF fitting robust to measurement errors in real-world HDR data.

**Parametric and microfacet models.** Analytic BRDF models are lightweight but approximate alternatives to data-driven representations. Blinn [1977] introduced a now commonly used specular reflection model, Schlick [1994] incorporated Fresnel effects, and Lafortune et al. [1997] proposed a more powerful and realistic phenomenological model. Pacanowski et al. [2012] use rational functions to obtain low error fits to some measured examples, but their model is costly to fit and unstable across materials.

Ngan and colleagues [2005] were the first to experimentally validate parametric models using least-squares fits to the MERL dataset, and provided guidelines for using parametric versus data-driven models. Since then, more sophisticated parametric models have been introduced [Ashikhmin and Premože 2007; McAuley et al. 2012; Bagher et al. 2012; Löv et al. 2012] but none consistently obtain good fits across real-world BRDFs.

Cook and Torrance [1981] proposed the microfacet model for computer graphics and analytically derived the shadowing-masking factor from a statistical distribution of orientations on micro-scale, perfectly specular facets. Their model’s factors and their angular arguments will be discussed further in Section 3. Schlick [1994] approximated the Cook-Torrance model factors with more efficient rational functions. Different parametric forms for the individual factors in this model have been investigated, including more efficient Fresnel approximations [Schlick 1994; Lazányi and Szirmay-Kalos 2005] and different shadowing-masking formulations [Smith 1967; Brown 1980; Ashikhmin et al. 2000; Bourlier et al. 2002]

Recent work proposes two-parameter models for the microgeometry distribution factor  $D$  via the *shifted Gamma* [Bagher et al. 2012] or *ABC* [Löv et al. 2012] functions. We compare in part to these two fits as the most accurate parametric models to date. Bagher et al. build on the work of Ashikhmin and Premože [2007], which assumes  $D(\theta_h)$  to be the dominant factor and fits it from the 1D retro-reflective ( $\theta_d = 0$ ) BRDF slice. For the Fresnel factor, Ashikhmin et al. fit the Schlick model using rendering comparisons, while Bagher et al. fit a more general Fresnel model using the  $\theta_h = 0$  BRDF slice. Bagher et al. thus fit BRDFs using just two

2D data slices within  $\theta_h, \theta_d \in [0, 70^\circ]$ . Löw et al. fit over the 3D set of angular samples restricted to  $\theta_h, \theta_d \in [0, 80^\circ]$ . We fit over the entire angular range, ensuring fidelity even for glancing reflections.

Löw et al. notably apply a type of compressive metric by transforming scaled BRDF samples  $\hat{\rho} = \cos \theta_i \rho$  via  $\log(1 + \hat{\rho})$  before summing squared errors. Nielsen et al. [2015] also transform samples by applying a log to the ratio of BRDF measurements and the median BRDF value over the entire MERL dataset.<sup>1</sup> Log-transformed values are better suited to principal component analysis, which they use to accelerate material capture. These nonlinear transforms preclude our simple AWLS minimization strategy and are challenging to apply to tabulated models with more than a few parameters; no previous data-driven factorization method has done so. We compare Löw’s microfacet model, fit using both their metric and ours (Löw et al. also propose another model based on Rayleigh-Rice theory that they report is generally a little less accurate than their microfacet-based one.) We find that our metric is comparable to Löw’s, with the advantage that ours can be efficiently minimized for tabulated models, so that our gain comes from the more general non-parametric model we propose.

Microfacet models with a tabulated  $D$  (microfacet distribution) factor have also been proposed [Ngan et al. 2005; Ashikhmin and Premoze 2007; Wang et al. 2008; Ghosh et al. 2008]. These models are only partially “non-parametric”, retaining a parametric model for the Fresnel factor  $F$  based on a single index of refraction [Cook and Torrance 1981] and a shadowing factor  $G$  derived from  $D$  using a simple model. We study the effectiveness of general tabulated factors for all three microfacet factors.

Dupuy et al. [2015] efficiently fit microfacet slope distributions, which are analytically related to the surface’s underlying *normal distribution function*, from the retro-reflective slice ( $\theta_d = 0$ ) of real-world BRDF data. They then compute the remaining microfacet factors based on a parametric model. Dupuy et al. identify the harder problem of independently fitting *all* microfacet factors (which they refer to as “inversion”), which is our contribution.

*Factored representations.* Several BRDF factorizing strategies have been proposed; each varies in the form and parameterization of its factors, its fitting objective and method, and the tradeoff it yields between accuracy and storage. Steigleder et al. [2002] tabulate an anisotropic Phong model [Ashikhmin and Shirley 2000] to make it more suitable for interactive rendering. Kautz and McCool [1999] apply decomposition to a BRDF matrix (optionally retabulated according to the halfway parameterization and with 2D directions unfolded into 1D column-vectors, resulting in a one-term product of these 2D factor column-vectors. McCool et al. [2001] fit a one-term model of three 2D factors in log space. Their fitting implicitly enforces non-negativity and de-emphasizes high BRDF values. Instead of fitting in log space, we re-weight a simple sum-of-squares metric. This avoids problems with near-zero measurements and magnification of small differences with the inverse (exponentiation) transform. To accelerate importance sampling, Lawrence et al. [2004] propose a sum of 2D products for view variation and a nested sum of 1D products parameterized by halfway elevation and azimuthal angles for light variation (139-332KB per material). They adopt non-negative matrix factorization (NMF) [Lee and Seung 2000] to fit their factors. Our model uses two orders of magnitude less memory. We observe that BRDF fitting of single-term models naturally avoids negative components with only trivial constraint enforcement in an otherwise unconstrained minimization.

<sup>1</sup>The measurements and median are cosine-weighted, similarly to Löw et al., and slightly biased to avoid numerical issues due to the log transform.

Lawrence et al.’s *inverse shade trees* [2006] enforce an overall functional structure to infer purely 1D factors. They consider the broader class of spatially-varying reflectance while we focus on point-wise BRDFs. Their model is not based on microfacet theory but is instead similar to the naive model we propose. We note that each microfacet factor is not a direct function of a separate independent variable, so that our approach can be interpreted as a “non-orthogonal” tensor factorization. Specifically, non-linear dependencies arise from the shadowing-masking ( $G$ ) factor’s dependence on  $\theta_i$  or  $\theta_o$  (see Equations 1 and 2). The  $G$  factor also appears twice, each time as a function of a different derived angle. Finally, the G-from-D version of our model determines  $G$  using a double integration of the microfacet distribution factor  $D$ .

Ben-Artzi et al. [2006] consider editing/rendering of a factored representation fit using existing metrics and techniques. As with other methods, their factorization is not based on microfacet theory, but rather a two-term, two-factor expansion without a  $G$  factor in the specular term. We believe our microfacet factors support similar intuitive user editing but have not yet investigated this.

Bilgili et al. [2011] apply Tucker tensor decomposition to measured BRDF data in order to obtain the product of four tabulated 1D factors, parameterized via the elevation and azimuthal angles of the halfway and outgoing vectors. Their model involves a (multi-term) sum of 20 factor products, for a total of 80 1D factors, requiring 77KB (double precision) per material and making it difficult to edit the resulting representation. We directly compare our model and show that ours provides better rendering fidelity for most materials in the MERL dataset. Our fits also require less than a tenth their storage. Note that Bilgili already includes rendering comparisons showing improvement over [Lawrence et al. 2004].

Without compressive weighting, fitting factors to measured BRDFs with a simple sum of squares metric, as proposed by Lawrence et al. or Pacanowski et al., leads to high rendering error in single-term product models (Figure 7). This conclusion broadly aligns with previous work using compressive BRDF transforms [McCool et al. 2001; Löw et al. 2012; Nielsen et al. 2015].

*Weighted linear regression.* While robust techniques for regression analysis are ubiquitous in statistics and experimental analysis [Graybill and Iyer 1994; Carroll and Ruppert 1998], previous work in CG reflectance modeling ignores the well-known problem of *heteroscedasticity* [White 1980]: the fact that statistical variables (i.e. BRDF samples) differ in terms of variances and expected measurement errors. A solution to this problem is *weighted* linear regression, where weights are inversely proportional to sample variance. Often weights are simply made a direct function (e.g. reciprocal squared) of the dependent variable or measured response, called least-squares percentage regression [Tofallis 2008]; see [NIS 2012] and [EPA 2003, Section 11.5.2] for specific examples. We extend these ideas to define a *compressive weighting* in Section 4.2; it applies a novel compressive numerator to ensure the weighting does not grow arbitrarily as a BRDF measurement approaches 0.

### 3. MICROFACET MODEL

We use a canonical local coordinate frame with surface normal  $\mathbf{n} = z = (0, 0, 1)$ . We denote the light direction by  $\mathbf{i}$ , view direction by  $\mathbf{o}$ , and *halfway vector* as the bisector of these two vectors:  $\mathbf{h} = (\mathbf{i} + \mathbf{o}) / \|\mathbf{i} + \mathbf{o}\|$  (see Figure 2). All directions are 3D unit vectors.

An isotropic BRDF is invariant to rotation of the halfway vector around the normal, and can be parameterized by three angles  $(\theta_h, \theta_d, \phi_d)$  derived from  $\mathbf{n}$ ,  $\mathbf{i}$ , and  $\mathbf{o}$ . The angle between the halfway vector and normal is denoted  $\theta_h \in [0, \pi/2]$  where




Fig. 2. Isotropic BRDFs are functions of  $\theta_h$ ,  $\theta_d$ , and  $\phi_d$ . The angle  $\phi_d$  (not shown) represents rotation of the view vector  $\mathbf{o}$  (or light vector  $\mathbf{i}$ ) around the halfway vector  $\mathbf{h}$ , out of the plane formed by  $\mathbf{h}$  and  $\mathbf{n}$ .

$\cos \theta_h = \mathbf{n} \cdot \mathbf{h}$ . The angle  $\theta_d \in [0, \pi/2]$  is formed between the view or light direction and the halfway vector:  $\cos \theta_d = \mathbf{i} \cdot \mathbf{h} = \mathbf{o} \cdot \mathbf{h}$ . Finally, the angle  $\phi_d \in [0, \pi]$  denotes the rotation of the view direction about  $\mathbf{h}$ . It is defined by  $\cos \phi_d = \mathbf{o}_h^\perp \cdot \mathbf{n}_h^\perp$ , where  $\mathbf{u}_v^\perp = \frac{\mathbf{u} - (\mathbf{u} \cdot \mathbf{v})\mathbf{v}}{\|\mathbf{u} - (\mathbf{u} \cdot \mathbf{v})\mathbf{v}\|}$  denotes perpendicular projection of  $\mathbf{u}$  onto  $\mathbf{v}$ .

The angle between the light direction and normal,  $\theta_i = \cos^{-1}(\mathbf{i} \cdot \mathbf{n})$ , and between the view direction and normal,  $\theta_o = \cos^{-1}(\mathbf{o} \cdot \mathbf{n})$ , can be determined from these three angular coordinates via:

$$\cos \theta_i = \cos \theta_h \cos \theta_d + \sin \theta_h \sin \theta_d \cos \phi_d, \quad (1)$$

$$\cos \theta_o = \cos \theta_h \cos \theta_d - \sin \theta_h \sin \theta_d \cos \phi_d. \quad (2)$$

BRDF samples must lie in the positive hemisphere of  $\mathbf{n}$ ; in other words,  $\theta_i, \theta_o \in [0, \pi/2]$ . The above formulas are easily derived using the parameterization introduced in the appendix. BRDFs satisfy *reciprocity*: reflectance is unchanged if  $\mathbf{i}$  and  $\mathbf{o}$  are exchanged. This implies that  $\theta_i$  and  $\theta_o$  can also be exchanged; the above formulas follow the conventions of the MERL dataset access code.

Given this parameterization, our microfacet-based model represents BRDFs in terms of three 1D factor functions via

$$\rho_M(\theta_h, \theta_d, \phi_d) = \rho_d + \rho_s \left( \frac{D(\theta_h) F(\theta_d) G(\theta_i) G(\theta_o)}{\cos \theta_i \cos \theta_o} \right). \quad (3)$$

The  $D$  factor is called the *normal distribution function* (NDF) and represents the probability that a microfacet normal makes an angle of  $\theta_h$  with the macro-scale normal  $\mathbf{n}$ . The  $F$  factor is called the *Fresnel factor* and models reflectance variation in the  $\theta_d$  parameter. The *geometric factor*  $G$  models shadowing, masking, and inter-reflection effects, in terms of the obliquity of the view or light direction,  $\theta_i$  or  $\theta_o$ . Reciprocity implies that the same  $G$  factor must be shared for both the  $\theta_i$  and  $\theta_o$  arguments.<sup>2</sup> For conciseness, we absorb the division by  $\pi$  usually present into  $\rho_d$  and  $\rho_s$ .

Equation 3 represents a *single-term* product model. In particular, the structure of the microfacet model presumes that the BRDF's dependence on  $\theta_h, \theta_d$  can be decomposed into a single product of independent functions  $D(\theta_h)$  and  $F(\theta_d)$ . Other data-driven (and non-microfacet-based) factorization methods such as [Lawrence et al. 2004; Lawrence et al. 2006; Bilgili et al. 2011] apply a *multi-term* model: sums of multiple products over many such factors, with potentially more fitting power. Multi-term models are difficult to fit, and do not necessarily realize this potential, as we will demonstrate in Section 5. Editing and manipulating the parameters of such multi-term models is also more difficult.

To test the microfacet model's validity, we define a *naïve model*

$$\rho_N(\theta_h, \theta_d, \phi_d) = \rho_d + \rho_s D(\theta_h) F(\theta_d) H(\phi_d), \quad (4)$$

<sup>2</sup>Previous work defines a general 2D geometric factor which is approximated by the product of 1D factors:  $G(\theta_i, \theta_o) \approx G_1(\theta_i) G_1(\theta_o)$ . To simplify notation, we use  $G$  for the 1D geometric factor rather than  $G_1$ .

that factorizes directly in terms of the angular coordinates. We additionally compare against the *bivariate* model [Romeiro et al. 2008]

$$\rho_B(\theta_h, \theta_d, \phi_d) = B(\theta_h, \theta_d) \quad (5)$$

that tabulates over  $\theta_h$  and  $\theta_d$ . This model is a 2D factorization and much less compact than our 1D microfacet models, despite completely ignoring dependence on  $\phi_d$ . It will help us to analyze the limitations of our principal model for specular materials.

We fit to the MERL database [Matusik et al. 2003] containing measurements for 100 different isotropic materials. Each material measurement is stored in a uniformly-sampled 3D block of  $(\theta_h, \theta_d, \phi_d)$  coordinate space with an angular sampling of  $90 \times 90 \times 180$ , for each *rgb* color channel. In single precision, this amounts to 17.5MB of data (originally 35MB in double precision). The  $\theta_h$  dimension is further transformed via

$$\theta'_h = \sqrt{\theta_h}. \quad (6)$$

This provides denser sampling close to 0, in the region of the specular highlight;  $\theta'_h$  is then uniformly sampled in  $[0, \sqrt{\pi/2}]$ .

We fit each color channel of the model separately in (3). For each color channel, we must solve for three vectors each containing 90 components (representing sampled factor functions  $D$ ,  $F$ , and  $G$ ), and two scalars ( $\rho_d$  and  $\rho_s$ ). Our choice for the number of factor samples (90) is made so as to match the MERL database's angular sampling resolution, and assumes the database already samples each material's reflectance adequately. Quantizing each vector component and scalar using single precision floating point (four bytes) yields a size of  $3 \cdot (3 \cdot 90 + 2) \cdot 4 = 3264$  bytes per material. In fact, dynamic range clamping (discussed in Section 4.3) requires only 20, rather than 32, bits per value. The naive model's  $H$  factor tabulates 180 samples to match the MERL sampling resolution, yielding 4344 bytes per BRDF. The bivariate model is sampled at a resolution of  $90 \times 90 \times 3$ , yielding 97200 bytes per BRDF.

## 4. MODEL FITTING

For each *rgb* color channel, each of the  $D$ ,  $F$ ,  $G$ , or  $H$  factors becomes a vector representing a 1D sampled function.  $D$ ,  $F$ , and  $G$  consist of 90 vector components while  $H$  comprises 180 components. Each factor component is a single optimization variable denoted by subscript indexing, e.g.  $D_k$ . The diffuse and specular coefficients ( $\rho_d$  and  $\rho_s$ ) yield two additional optimization variables per color channel. We explain how to solve for these variables, detailing both the objective function as well as how we minimize it.

### 4.1 Fitting objective

We minimize the following sum over all BRDF measurements:

$$E = \sum_j w_j (\rho(\theta_h, \theta_d, \phi_d)_j - \rho_j^*)^2, \quad (7)$$

where  $\rho_j^*$  is a BRDF measurement,  $(\theta_h, \theta_d, \phi_d)_j$  is its parametric coordinates, and  $\rho$  is evaluated from (3), (4), or (5) by accessing the factor component vectors at the specified parametric coordinates.

The weight  $w_j$  is given by the product of three subweights as

$$w_j = w_V(\theta_h, \theta_d, \phi_d)_j w_I(\theta_h, \theta_d, \phi_d)_j w_C(\rho_j^*). \quad (8)$$

We detail the three subweights below; global constants in these definitions do not affect the minimization and can be ignored. We visualize the subweight functions across BRDF samples in Figure 3.

We include all measured samples in the objective except those for which  $\theta_i, \theta_o > 89^\circ$ , where cosines in the denominator of (3) cause numerical instabilities.




Fig. 3. Objective weighting. Weighting magnitudes are visualized in false color (red is highest and blue lowest), across the BRDF’s 3D angular coordinates  $(\theta'_h, \theta_d, \phi_d)$ . Compressive weights in the third column are material-dependent and are shown for the example of GOLD-METALLIC-PAINT2.




Fig. 4. The compressive function  $f(x; \alpha)$ .

**Volume form weighting.** The first subweight arises from the three-angle parameterization of isotropic BRDFs. It is given by

$$w_V = \sqrt{8 \sin^2 \theta_d (\cos^2 \theta_d + \sin^2 \theta_d \cos^2 \phi_d)} d\theta_h d\theta_d d\phi_d. \quad (9)$$

Though prior work mentions weighting by solid angle (see e.g. [Ngan et al. 2005]), the proper weighting is determined by the *volume form* of the mapping between the  $(\theta_h, \theta_d, \phi_d)$  parameter space and the  $(\mathbf{i}, \mathbf{o})$  output space of pairs of unit directions. This formula is derived in the appendix; we believe it has not appeared before in the CG literature. Since the BRDF parameterization is in terms of  $\theta'_h$  rather than  $\theta_h$ , (6) implies that  $d\theta_h = 2\theta'_h d\theta'_h = 2\sqrt{\theta'_h} d\theta'_h$ .

**BRDF importance weighting.** The second subweight is

$$w_I = \cos \theta_i \cos \theta_o. \quad (10)$$

The  $\cos \theta_i$  factor arises because it always multiplies the BRDF in the local shading integral over incident lighting directions. The  $\cos \theta_o$  factor is included for reciprocal symmetry and since surface points are more likely to appear as their normal aligns with the view direction. This and similar weighting schemes have been used in earlier work [Lafortune et al. 1997; Ngan et al. 2005].

**Compressive weighting.** There are many reasons to expect more measurement error near a higher-magnitude BRDF sample, including blurring, diffraction, distortion, transmission inefficiency, registration/calibration errors, and veiling glare [Talvala et al. 2007] in the acquisition. This is why errors in optical systems are estimated using signal to noise ratio (SNR), a logarithmic measure relative to signal strength.

The compressive subweight specifies that we should not try as hard to fit higher (and thus likely higher error) values  $\rho_j^*$  as we do lower ones, and is given by

$$w_C = \left( \frac{f(\rho_j^*/\bar{\rho}; \log 2)}{\rho_j^*/\bar{\rho}} \right)^p, \quad (11)$$

where  $f$  is a *compressive function* of the form

$$f(x; \alpha) = 1/\alpha (1 - e^{-\alpha x}), \quad (12)$$

and  $\bar{\rho}$  is the weighted median of the BRDF with weight given by the product  $w_V w_I$ . The power  $p$  is typically chosen in the range  $[1, 2]$ : the larger  $p$ , the higher the weight given to darker regions of the BRDF relative to brighter regions. We always use  $p = 1.4$ .

The compressive function in (12) is shown in Figure 4. It maps all non-negative real numbers  $x$  to the finite interval  $[0, 1/\alpha]$ . The larger  $\alpha$ , the more compressive the transformation; i.e., the more higher values are squeezed closer to the function’s asymptote of  $1/\alpha$ . We fix  $\alpha = \log 2$  in order to map the BRDF weighted median to the center of the output range (i.e., to  $1/2\alpha$ ).

Note that  $f(x) \leq x$  and  $\lim_{x \rightarrow 0} f(x)/x = 1$ . The maximum value of  $w_C$  is therefore 1. From this maximum, at  $x = \rho_j^*/\bar{\rho} = 0$ , it decreases monotonically to 0 as the BRDF value  $\rho_j^* \rightarrow \infty$ . There is no difficulty with weights getting arbitrarily high as  $\rho_j^* \rightarrow 0$ , as is the case with the simpler alternative of weighting by the measurement reciprocal  $w_C = (\rho_j^*)^{-p}$  used in previous work [Tofallis 2008; EPA 2003]. Away from 0, (11) is essentially identical to this simpler alternative scheme.

## 4.2 Fitting method

Alternating weighted least squares (AWLS) is an iterative procedure that repeatedly updates each factor in sequence until convergence. To update a single factor, we hold everything else constant and solve for its optimal weighted least-squares value. This guarantees that the iteration proceeds downhill in the objective.

To understand how each factor component is computed, consider the simplified objective

$$E(x) = \sum_j w_j (z_j - x y_j)^2, \quad (13)$$

where  $x$  represents an individual factor component to be solved for. The minimizing solution is

$$x = \frac{\sum_j w_j y_j z_j}{\sum_j w_j y_j^2}. \quad (14)$$

For each factor component, we therefore accumulate the weighted sums representing the above numerator and denominator over all BRDF samples mapping to that component’s parametric location. The optimal value of the component is then given by the division in (14).

Specifically, assume we want to update the factor  $D$ . The other factors, as well as  $\rho_d$  and  $\rho_s$ , are held constant. A component of  $D$ , labeled  $D_k$ , corresponds to a particular value of its argument  $(\theta_h)_k$ .

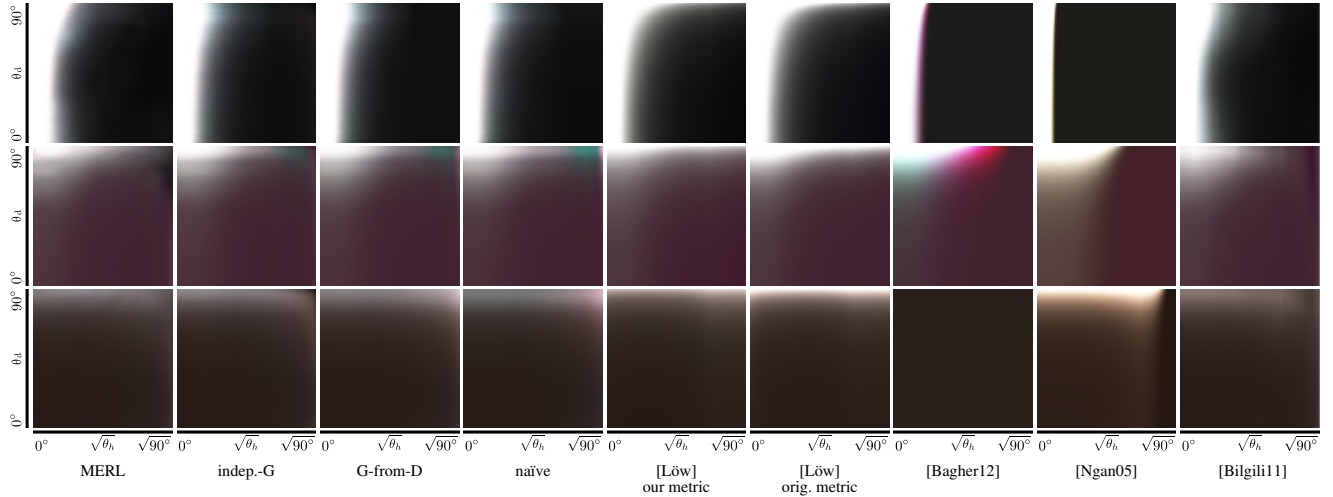


Fig. 5. BRDF fit visualizations. The three rows represent materials (top to bottom): ALUMINIUM, VIOLET-RUBBER and POLYURETHANE-FOAM. Only the single slice  $\phi_d = 90^\circ$  is shown. The visualization uses tone-mapping via the compressive function in (12).

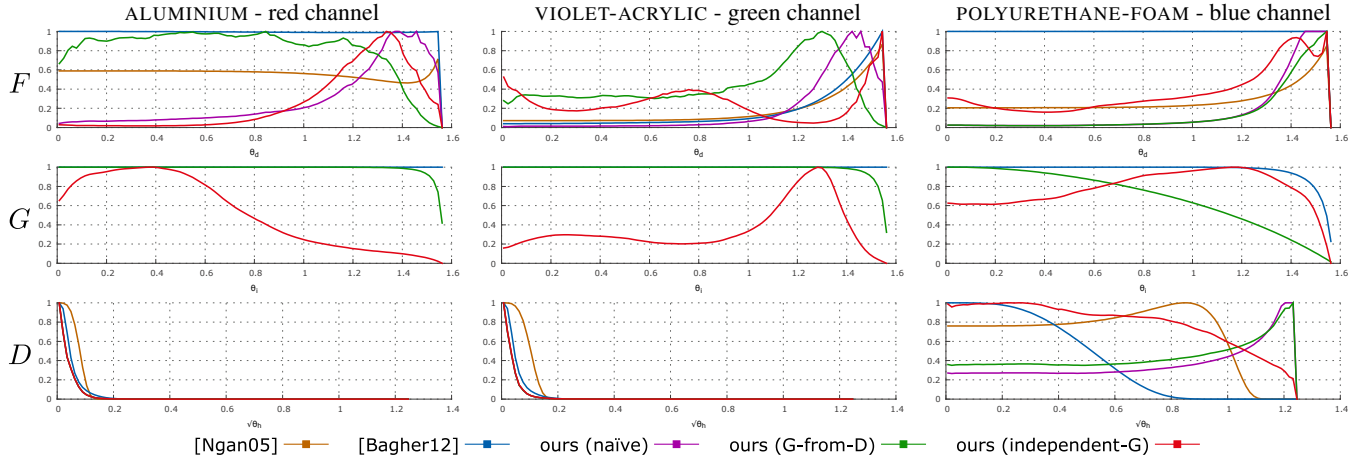


Fig. 6. Microfacet factors for three example materials. (left to right): ALUMINIUM, VIOLET-ACRYLIC, and POLYURETHANE-FOAM. The factors plotted are (top to bottom):  $D$ ,  $G$ , and  $F$ . The corresponding BRDFs are visualized in Figure 5.

Let the parametric locations of all BRDF samples whose  $\theta_h$  coordinate maps to  $(\theta_h)_k$  be indexed by  $j$ , yielding  $(\theta_d, \phi_d, \theta_i, \theta_o)_{kj}$ . To update  $D_k$  we seek the least-squares best solution to the weighted system of equations indexed by this  $j$ :

$$\rho_{kj}^* = \rho_d + \rho_s \frac{D_k [F(\theta_d)]_{kj} [G(\theta_i)]_{kj} [G(\theta_o)]_{kj}}{[\cos\theta_i \cos\theta_o]_{kj}},$$

where, as before, equation  $j$  is weighted by  $w_j$ . The optimal solution is given by (14) with

$$\{x, y_j, z_j\} = \left\{ D_k, \frac{\rho_s [F(\theta_d)]_{kj} [G(\theta_i)]_{kj} [G(\theta_o)]_{kj}}{[\cos\theta_i \cos\theta_o]_{kj}}, \rho_{kj}^* - \rho_d \right\}.$$

The  $F$  factor, and  $H$  in the naïve model, are updated similarly.

To solve for  $\rho_d$  and  $\rho_s$ , we observe that the optimal solution to the simplified objective

$$E(a, b) = \sum_j w_j (a + b x_j - y_j)^2 \quad (15)$$

can be computed by Cramer's rule as:

$$a = \left\| \frac{\sum_j w_j y_j \quad \sum_j w_j x_j}{\sum_j w_j x_j y_j \quad \sum_j w_j x_j^2} \right\| / \left\| \frac{\sum_j w_j \quad \sum_j w_j x_j}{\sum_j w_j x_j \quad \sum_j w_j x_j^2} \right\|, \quad (16)$$

$$b = \left\| \frac{\sum_j w_j \quad \sum_j w_j y_j}{\sum_j w_j x_j \quad \sum_j w_j x_j y_j} \right\| / \left\| \frac{\sum_j w_j \quad \sum_j w_j x_j}{\sum_j w_j x_j \quad \sum_j w_j x_j^2} \right\|. \quad (17)$$

If the optimal  $a < 0$ , we clamp it to 0 and apply the following simpler formula to determine an optimal  $b$  via

$$b = \frac{\sum_j w_j x_j y_j}{\sum_j w_j x_j^2}. \quad (18)$$

The optimization proceeds by updating each factor component (e.g.  $D$  or  $F$ ) using (14), and then updates  $\rho_d$  and  $\rho_s$  which become  $a$  and  $b$  in (16) and (17). We show results of this fitting procedure for a few representative MERL materials in Figure 5, with their corresponding derived factors in Figure 6.

Fitting the bivariate model is also similar, but can be computed in a single step without iterative relaxation. Each of the components




Fig. 7. Rendering results for BRDF fits with and without compressive weighting. We compare four materials, from left to right NICKEL, COLOR-CHANGING-PAINT1, MAROON-PLASTIC, and WHITE-FABRIC2. For specular materials (leftmost three columns), compressive weighting (middle row) is necessary to obtain a fit that matches the rendered ground truth; the reflectance match is very poor without it (bottom row). This is true for all three of our non-parametric models: (a) independent-G, (b) G-from-D, and (c) naive. For more diffuse materials (rightmost column), results with and without compressive weighting are similar.

in the 2D  $B$  matrix in (5) is computed by a simple weighted sum over all BRDF samples mapping there via the formula in (14). More specifically, each  $\theta_h, \theta_d$  component is computed over all samples in  $\phi_d$  with weights determined by (8), as before. This fitting method differs from [Romeiro et al. 2008], which takes a simple average over all  $\phi_d$  samples.

### 4.3 Solving for $G$

$G$  can either be treated as an independent tabulated factor or derived from  $D$  using a shadowing model [Smith 1967; Brown 1980; Bourlier et al. 2002; Walter et al. 2007; Bagher et al. 2012]. When evaluating  $G$  or the cosine factors in (3), we apply linear interpolation in the  $\theta_i$  or  $\theta_o$  spaces. This is not necessary for  $D$  and  $F$  since all BRDF samples  $\rho_j^*$  are uniformly parameterized in  $\theta'_h$  and  $\theta_d$ , and map exactly onto the corresponding vector component.

*G-from-D.* Determining  $G$  from  $D$  introduces a nonlinear dependence in the relaxation, so that the solved-for  $D$  is no longer optimal. We address this problem by computing  $D$  using the standard AWLS step, deriving  $G$  from it, and then explicitly checking whether the objective function has decreased. If not, we apply a simple 1D minimization based on golden section search (GSS) [Press et al. 2007] that is guaranteed to reduce the objective along the line from  $D$ 's previous state to its new one.

We compute  $G$  from  $D$  by first normalizing  $D$  to be a probability density function (PDF) via  $\bar{D}(\theta_h) = D(\theta_h)/\bar{D}$  where

$$\bar{D} = 2\pi \int_0^{\pi/2} D(\theta) \cos \theta \sin \theta d\theta. \quad (19)$$

Then  $G(\theta) = (1 + \Lambda(\cot \theta))^{-1}$  where

$$\Lambda(\mu) = \frac{1}{\mu} \int_{\tan^{-1} \mu}^{\pi/2} (\tan \omega - \mu)(1 + \tan^2 \omega) p_2(\tan \omega) d\omega,$$

$$p_2(r) = 2 \int_0^{\pi/2} p_{22}(r^2 + \tan^2 \psi)(\tan^2 \psi + 1) d\psi, \text{ and}$$

$$p_{22}(u) = \tilde{D}(\tan^{-1} \sqrt{u}) \cos^4(\tan^{-1} \sqrt{u}).$$

These formulas follow the notation and are easily derived from previous work [Walter et al. 2007; Bagher et al. 2012], using the change of variables  $\psi = \tan^{-1} q$ ,  $\omega = \tan^{-1} r$ , and  $u = \tan^2 \theta_h$ . To numerically evaluate the above integrals, we obtain sufficient accuracy with a 15 point Gaussian quadrature. The tabulated  $D$  function is evaluated using linear interpolation.

*Independent-G.* Solving for  $G$  also introduces a non-trivial factor dependency. Unlike the other two factors, it appears twice in the microfacet model. It is also evaluated at  $\theta_i$  and  $\theta_o$ , which are complicated functions of the BRDF's angular coordinates. We apply a simple heuristic to solve for it. Let the two factors be denoted  $G_i = G(\theta_i)$  and  $G_o = G(\theta_o)$ . Using a standard AWLS update step, we can solve for each separately holding the other (as well as the  $D$  and  $F$  factors, and  $\rho_s$  and  $\rho_d$  scalars) constant. We then average these two results, and apply a simple Gaussian smoothing filter. Gaussian smoothing removes high-frequency oscillations in  $G$  we otherwise obtain for some specular materials and makes the fitting more robust. This procedure is iterated until it converges or increases the objective. In the case where the first iteration increases the objective, we apply a GSS update.

We tried making the 1D factors  $G_i$  and  $G_o$  be independent (i.e. non-reciprocal), but found the constraint that both be identical a useful regularization that sped convergence and eliminated ambiguities between factors. We also experimented with a fully 2D  $G$  factor without obtaining good results; convergence was slow.

For the independent-G solution, note that the cosine factors in (3) are superfluous in the sense that they can be absorbed into the gen-

eral  $G$  factor. We retain them so that the  $G$  factors are comparable between the two methods.

*Dynamic range (DR) clamping.* We keep all factors normalized so that their maximum component is 1 to prevent scaling drift;  $\rho_s$  absorbs the compensatory scaling. To ensure a physically-plausible dynamic range, we clamp each factor component to be no less than  $\epsilon = 1e-6$ . This is necessary in independent-G fitting which involves the product of four factors solved for independently ( $\rho_s, D, F, G$ ). Without DR clamping, we sometimes get large values for  $\rho_s$  and inflated importance of tiny variations in near-zero factor components. For G-from-D fitting, DR clamping is unnecessary and we simply clamp negative factor components to 0.

## 5. RESULTS AND DISCUSSION

Figure 7’s first three columns show that compressive weighting is necessary to obtain good rendered matches to ground truth for many specular materials. For more diffuse materials, results with and without compressive weighting are similar (fourth column). This is true regardless of the form of the non-parametric model we fit: G-from-D (a’), independent-G (b’), or naive (c’). The metrics used in previous work on BRDF data-driven factorization, such as [Lawrence et al. 2004; Ngan et al. 2005; Lawrence et al. 2006; Ben-Artzi et al. 2006; Wang et al. 2008; Pacanowski et al. 2012], lead to poor rendering fidelity, particularly in the case of single-term BRDF expansions and HDR lighting environments. Our single-term non-parametric model would not significantly outperform state-of-the-art parametric and factorized alternatives without magnitude compensation in the weighting of squared differences. This finding parallels that reported in [L ow et al. 2012], where a standard vs. a logarithmically-transformed metric are compared and the latter found superior.

Our test rendering suite combines several HDR environment maps (EMs) and scene geometries: A) a sphere in the Grace EM, B) the Stanford Buddha and sphere in the Ennis EM, C) Buddha and sphere in the Grove EM, and D) a sphere in St. Peter’s EM. EMs are from [Debevec 2001] with the exception of Ennis [Debevec 2008]. EMs acquired from the real world yield realistic imagery when combined with measured BRDFs. They also effectively identify reflectance mismatches across a BRDF’s entire angular extent using a single rendered image.

We compare renderings generated using various parametric and non-parametric models, with quantitative error reported in terms of standard PSNR. PSNR numbers are computed in the linear space of rendered outputs (i.e. without tone-mapping), over all “foreground” pixels, meaning pixels directly covered by an object with MERL reflectance. The PSNR numerator is the maximum intensity in any  $rgb$  color component of any foreground pixel in the ground truth rendering. In all rendered images in this paper and the supplement, we apply the following fixed tone mapping operator which combines exposure compensation with standard gamma correction via:

$$\text{output} = \text{pow}((1 - \exp2(-\text{exposure} * \text{pixel})), 1/\text{gamma}).$$

We use the settings `exposure = 0.02` and `gamma = 2.2` for all renderings except with the darker and lower dynamic range Ennis environment map, which uses `exposure = 1` and `gamma = 2.8`.

We used three different renderers to generate our various results and comparisons. Scene A is rendered with the HQR renderer [Soler and Roche 2014] using BRDF importance sampling and 500K samples per pixel. Scenes B and C are rendered with PBRT [Pharr and Humphreys 2010] using multiple importance sampling (MIS) of the BRDF and lighting [Veach 1998], with 4K

Table I. Fitting statistics over all materials.

technique	fitting time (s)			# of AWLS iterations			GSS calls			# of materials invoking GSS
	min	avg	max	min	avg	max	min	avg	max	
independent-G	59	2838	16421	25	1192	10168	1	47	2296	57
G-from-D	135	765	7222	15	56	1012	2	69	359	46
naive model	349	680	2281	817	2038	7628	0	0	0	0

samples per pixel (including 16x spatial anti-aliasing). Scene D is rendered with a custom GPU renderer with 300k samples per pixel partitioned evenly between distributions proportional to the BRDF, the lighting, and a uniform hemispherical distribution using MIS. For both lighting and BRDF importance sampling, we generate pre-tabulated cumulative distribution functions (CDFs) during initialization and perform sample warping using quasi-Monte Carlo patterns. Further rendering details are included in the Appendix.

We also developed an interactive WebGL demo that renders our independent-G model with HDR environmental lighting. It works directly from our compact factor representation, without precomputing a CDF as in the offline renderers described previously. We first assembled the factors and coefficients for each of the 100 MERL materials into a single  $512 \times 128$  “uber-texture” (top of Fig. 16). Direct rendering using this texture is further described in the Appendix.

We compare to three parametric microfacet models: the Cook-Torrance model of [Ngan et al. 2005], the shifted gamma microfacet model of [Bagher et al. 2012], and the microfacet model of [L ow et al. 2012]), all three with fits across the MERL dataset publicly available from the authors. We investigate L ow’s model fit using both their  $E_2$  (logarithmic) metric in Eqs. 31/32 from their paper and our metric (Eq. 7 in our paper), minimizing our metric using the Ceres solver [Agarwal et al. 2014]. We also compare to the data-driven factorization method of [Bilgili et al. 2011] using fits supplied by the authors.

We provide a comprehensive supplementary document with complete fitting and rendering results across all materials and fitting methods. Each page of the supplement compares rendering results across the various fits. We also include an additional page with the ground truth rendering results repeated in each corresponding image slot, allowing readers to flip back and forth between consecutive pages to obtain a visual impression of the rendering error. Our supplemental video contains similar comparisons between animated rendering results. Among other things, the video confirms that our model retains high rendering fidelity when the light or view rotate continuously. Supplementary material also includes the sampled factor curves we fit for each model, for all the MERL materials. All of our fits were computed on a single core of a 6-core 3.20GHz Intel Xeon E5-1650 CPU with 16 GB of RAM. Our performance statistics are summarized in Table I.

Our non-parametric representation consistently beats parametric factor models for all materials, both visually (Fig. 1) and quantitatively (Fig. 8). Though we fit our model to each color channel independently, we see excellent visual fidelity in rendered reflection color, without the color shifting artifacts in highlights present in alternative methods (see e.g. Figure 1).

Fig. 8 reports bottom-line rendering errors across our experiments. We sort materials by increasing diffuseness (top) and increasing error (bottom), measured from the independent-G method. Diffuseness is computed from  $\bar{D}$  as in Eq. 19, where  $D$  is obtained by fitting to the naive model. Recall that  $D$  is normalized so that its maximum value is 1. It is easy to verify that  $\bar{D} = (\max_{\theta} \bar{D}(\theta))^{-1}$ , where  $\bar{D}$  is the mean-normalized distribution derived from the maximum-normalized  $D$  factor.  $\bar{D}$  has a high peak for specular






Fig. 8. Rendering errors across all materials. Materials are sorted by increasing diffuseness from (19) (top), and decreasing error for independent-G (bottom). All three non-parametric models provide better accuracy than the parametric alternatives. As diffuseness increases, error from the non-parametric models increasingly diverges, with the independent-G method consistently providing the best fit.

materials and a much smaller one for diffuse materials, so the reciprocal of this peak yields a reasonable measure of diffuseness.

The average PSNR across all 100 materials and all four rendering setups were as follows (in descending order): bivariate 64.07, ours (independent-G) 58.22, ours (G-from-D) 56.42, ours (naive) 56.71, Bilgili 54.02, Löw (orig. metric) 50.96, Löw (our metric) 50.26, Bagher 46.33, and Ngan 39.65.

All three of our non-parametric representations provide a good visual rendering match for nearly all the 100 materials in the MERL database. For highly specular materials, the three perform similarly. This is not surprising since they share the same  $D$  and  $F$  factors and the  $GH$  factors provide little additional expressiveness. As diffuseness increases, independent-G is increasingly accurate compared to the G-from-D and naive models: see the visual results in Figure 9, and the quantitative results in Figure 8 (top). Note the backlighting effect at the silhouette preserved by the independent-G method (i.e., NYLON and POLYETHYLENE in Figure 9), but lost by the other two methods. Even when the images

look similar, independent-G consistently increases quantitative rendering fidelity, up to 14dB compared to the other two methods.

The accuracy of a single-term product of 1D microfacet factors is excellent but not perfect; our goal was to study these limits. As shown in Figure 10, our non-parametric fits yield visually noticeable differences to ground truth for a few materials, like ALUM-BRONZE and VIOLET-ACRYLIC. Figure 11 indicates the reason for these failures: some specular BRDFs are not accurately decomposed as a single product of 1D factors,  $B(\theta_h, \theta_d) \not\approx D(\theta_h)F(\theta_d)$ . In addition to being a heavyweight model (2D rather than 1D and 30 times larger than our independent-G model), we note that the bivariate model does not capture the more diffuse MERL materials well, as shown in the bottom two rows of Figure 11, and in additional examples from the supplement such as the other plastics, GREEN-LATEX, RED-FABRIC, POLYURETHANE-FOAM, and PINK-FELT. We note that differences may be subtle when comparing multiple images on the same page, but become more obvious in



Fig. 9. Rendering comparison for non-parametric fits (close-up in the inset). PSNR over the image is displayed in the lower left corner. The independent-G method better preserves backlighting effects (NYLON and POLYETHYLENE), darkening (SILVER-METALLIC-PAINT), highlights (YELLOW-PLASTIC), and avoids color-shifting (green shift in highlights of naïve rendering for RED-SPECULAR-PLASTIC).

the supplement when flipping (paging back and forth) between the approximate model and the ground truth.

Figure 12 directly compares our independent-G model to Löw’s, the best of the parametric models. Our non-parametric fit improves visual fidelity for many materials. Across all 400 renderings (4 renderings for each of 100 MERL materials), independent-G beats Löw’s result 308 times (77%) and ties it 85 times (21%), where a tie is defined as rendering PSNR within 2db. Many more examples of materials where our model’s rendering fidelity exceeds Löw’s can be examined in the supplement, e.g. BLACK-OBSIDIAN, RED-METALLIC-PAINT, GREEN-ACRYLIC, MAROON-




Fig. 10. Our fits with the greatest visual error.

PLASTIC, AVENTURNINE, HEMATITE, GREEN-PLASTIC, STEEL, PVC and POLYURETHANE-FOAM.

Löw’s model generates slightly better quantitative errors when it is fit with their metric vs. ours: our metric tends to better capture glancing view/light directions, while theirs does better for more normal directions. We conclude that the improvement we observe over Löw is due to our model’s general non-parametric form rather than our fitting metric. As discussed in Section 2, Löw’s logarithmic metric represents a form of magnitude compensation suitable for optimization problems with few degrees of freedom. Our metric’s advantage is that it makes minimization practical for the many optimization variables involved in fitting tabulated factors.

Figure 13 compares our model to the Bilgili data-driven factorization. Despite our model’s relative compactness (a factor of 12 reduction, assuming single precision floating point numbers for both models), it provides a superior quantitative rendering PSNR for 71 out of the 100 MERL materials, and visually better results for many materials, as shown in the figure. Across all 400 renderings, independent-G beats the Bilgili factorization 236 times (61.5%) and ties it 71 times (17.8%). Many more examples of materials where our model’s rendering fidelity exceeds Bilgili’s can be found in the supplement, e.g. BRASS, BLACK-OBSIDIAN, ALUMINIUM, SILICON-NITRADE, GOLD-METALLIC-PAINT2 and NICKEL.

Our fitting strategy considers nearly the entire range of angular measurements in the captured data. Glancing reflection angles may be hard to measure accurately but values there are not arbitrary, and those recorded in the MERL dataset represent a reasonable correction/extrapolation. In contrast, evaluating existing parametric or tabulated models to angles beyond where they were originally fit sometimes yields an unrealistic appearance and a poor rendered match to the reference, as shown in Figure 12 third row and further visualized in Figure 14. Our fits avoid these artifacts. We also tried fitting to a restricted subset of the BRDF’s angular range:  $\theta_h, \theta_d, \theta_i, \theta_o \leq 80^\circ$ , with results shown in Figure 15. Compressive weighting remains critical to rendering accuracy even when the fit is computed over a restricted angular range.

Figure 6 illustrates the form of the factors obtained with our non-parametric fitting.  $D$  typically follows the form of a decreasing bump but sometimes exhibits a slight increase near its peak at  $\theta_h = 0$  (e.g. ALUMINA-OXIDE, YELLOW-PAINT, PVC in the sup-




Fig. 11. Comparing our fits to the bivariate model. For some specular materials (top three rows), the 2D bivariate model provides a better rendering match. Clearly, specular reflectance is not well modeled by the microfacet model’s single-term product separation  $D(\theta_h)F(\theta_d)$ . This accounts for the visual error in our poorest MERL fits (Fig. 10). For more diffuse MERL materials, the bivariate model lacks fidelity (bottom two rows). See the supplement for in-place comparisons to ground truth, and other examples.

plement). For fairly diffuse materials, the  $G$  factors fit with our independent-G model sometimes exhibit complex variation in contrast to the strict monotonic decrease predicted by standard shadowing models (e.g., POLYURETHANE-FOAM). Here,  $G$  may be capturing micro-scale inter-reflection. Our factors also often exhibit large, smooth oscillations not captured by any existing parametric model. We additionally observe a significant bump (rapid increase followed by decrease down to 0) in the Fresnel factor as  $\theta_d \rightarrow 90^\circ$  (ALUMINIUM and VIOLET-ACRYLIC in Fig. 6). This is another behavior not predicted by existing models.

Getting trapped in local minima is a problem in high-dimensional non-linear optimization. A commonly-used strategy is




Fig. 12. Comparison of our fit to the [L w et al. 2012] microfacet model. Note the incorrect color shift in COLOR-CHANGING-PAINT3, washed-out appearance of TWO-LAYER-GOLD, poor fidelity to glancing reflection in SPECULAR-YELLOW-PHENOLIC (with L w’s original metric), and incorrect highlights in RED-FABRIC. HDR error for SPECULAR-YELLOW-PHENOLIC is understated by tone-mapping, but apparent in the numerical PSNR (9.87), and further analyzed in Fig. 14.

to employ multiple starts, each initialized at random locations in the high-dimensional state space. We observe that our G-from-D model is very stable with respect to random initialization, and all of its fits were computed in a single run. Our independent-G model is usually stable, but can get stuck in a local minimum; here, we employ eight fitting runs for each material. Four of these runs were randomly initialized, one with a state-vector of 1’s, one with the results of [Bagher et al. 2012], one with the results of our G-from-D fit, and the last initialized  $D$  with the retro-reflective ( $\theta_h = 0$ ) BRDF slice,  $F$  with the  $\theta_d = 0$  slice, and  $G$  with a constant value of 1 across its domain. Our supplemental video illustrates the factor curve convergence behavior across multiple random starts.

## 6. CONCLUSION

How does the reflectance of real materials behave? What simple models capture this behavior? Does microfacet theory help? Are parametric models proposed for its factors significantly limiting?

To begin to answer these questions, we presented a new BRDF representation comprising the compact product of three generic 1D factors from the non-orthogonal factorization underlying the microfacet model, and a new method (ALWS) to fit this model to isotropic BRDFs. It is simple to implement, quickly converging, automatic (no per-material parameter tweaking), and robust. It is more accurate than previous fits based both on parametric models




Fig. 13. Comparison of our fit to the [Bilgili et al. 2011] factorized model.

and data-driven factorization, and provides visual and quantitative accuracy across nearly all materials in the MERL dataset. We conclude that previous parametric models for the microfacet factors have indeed limited reflectance accuracy.

We also showed that single-term fits that ignore magnitudes of the BRDF samples provide poor rendering fidelity for specular materials. This result was a surprise to us, and we think it will help data-driven factorization for multi-term models as well. Our solution is a simple compressive weighting scheme incorporated into the standard sum-of-squared-error fitting metric. Though we fit independently to each color channel, our metric yields excellent accuracy and color fidelity after rendering.

The Cook-Torrance microfacet model is fundamental in computer graphics. We scientifically validated its underlying structure from real-world measurements by showing that the independent-G model provides significant quantitative and visual error reduction compared to a non-microfacet orthogonal factorization having




Fig. 14. Error visualization for SPECULAR-YELLOW-PHENOLIC from Fig. 12, row three. Red intensity here is linearly mapped to the sum of absolute values of the error over all three color channels at each pixel, divided by the maximum error incurred at any pixel by our independent-G fit, and allowed to saturate above 1. [L w et al. 2012] generates erroneous, extremely bright reflections especially near the sphere’s silhouette. The actual ratio of maximum pixel errors between [L w et al. 2012]’s result and ours is over 1000, so the benefit of our fit is understated in this saturated visualization.

33% more degrees of freedom: the naive model. We also showed the benefit of independent-G over a model that determines  $G$  from  $D$  via a standard shadowing model (G-from-D). The gains are visually substantial for some diffuse materials, including an enhanced ability to recover silhouette backlighting effects.

For a few materials, we demonstrated that an important limitation on visual accuracy is the microfacet model’s inherent assumption that specularly can be decomposed into the product of separable univariate functions,  $D(\theta_h)F(\theta_d)$ . We reach this conclusion by testing the bivariate specular model, which retains full flexibility only in the angular coordinates  $\theta_h$  and  $\theta_d$  while totally ignoring  $\theta_i$  and  $\theta_o$ , and showing that this model alone remedies the artifacts we observe in those materials. Evidently, the microfacet model does not comprehensively express real-world specularly. This deficiency can not be corrected by developing more sophisticated parametric models for its 1D factors.

The relevance of these findings relies on our ability to fit non-parametric microfacet factors (which removes their restriction to simple parametric models) and our new metric (which better exploits the fitting power of the model for HDR BRDFs). Both are new contributions. For example, no previous work has investigated a non-parametric representation for the Fresnel factor in the microfacet model, or a geometric factor not derived from  $D$ .

We hope our work may lead to more general factor models to improve fidelity and compactness. We have not yet studied spatial variation in reflectance and are interested in fitting texture-mapped models for the factors perhaps using a low-dimensional factor eigen-basis, addressing the problem of interpolation and spatial coherence, and investigating how factors vary spatially from physical measurement. Our work has explored the limits of a single-term product expansion; a multi-term model is indicated for some materials and in demanding applications. Augmenting the bivariate model by multiplying the 2D  $B(\theta_h, \theta_d)$  factor by a 1D geometric/shadowing  $G$  factor (i.e. replacing  $D(\theta_h)F(\theta_d)$  by  $B(\theta_h, \theta_d)$  in Eq. 3) subsumes our current model and should yield a representation able to capture both specular and diffuse reflectance with high visual accuracy. We are still exploring the difficult problem of how to fit this more general model. We are also studying how to compute a low-dimensional factorization to compress  $B$ , using per-component compressive weighting. As might be expected from our results, singular value decomposition (SVD) and subsequent truncation of small singular values of the *unweighted*  $B$  matrix per-




Fig. 15. Compressive weighting and angular fitting range restriction. We compare renderings of independent-G fits with and without compressive weighting using two different angular ranges in the BRDF fit:  $\theta_i, \theta_o \leq 89^\circ$  (as used for all our other results in this paper and the supplement) and  $\theta_i, \theta_o \leq 80^\circ$  (which eliminates more glancing BRDF measurements from the fitting). Compressive weighting better matches ground truth in terms of both measured PSNR and visual fidelity. This remains true even when we significantly restrict the angular fitting range as is often done in previous methods.

forms poorly in our preliminary tests; weighted low-rank approximation [Nati and Jaakkola 2003] seems necessary.

A final question is whether our factors are unique and physically-motivated, or one among many representational artifacts that do the job. We believe the former is mostly true. For most materials we get consistent results regardless of how the minimization is initialized, especially after invoking the “DR clamping” regularization (Section 4.2), indicating uniqueness and possible physical reality. On the other hand, especially with the independent-G model but even with G-from-D, there is indeed ambiguity between factors for a few materials, meaning that multiple quite different forms of the locally-minimizing factors can be found yielding about the same rendering fidelity. This may be due to single-term insufficiency in the specular microfacet model (i.e., the more powerful bivariate model  $B(\theta_h, \theta_d)$  is needed rather than  $D(\theta_h)F(\theta_d)$ ) along with the presence of nearly equal top singular values for the weighted SVD of  $B$ . Such a factor ambiguity is almost certainly made worse by introducing more degrees of freedom in an independent  $G$  factor. The issue is complicated because there’s no guarantee we find

the global minimum, so we can’t truly distinguish a deficiency in the model from a deficiency in the numerical minimization.

## REFERENCES

- AGARWAL, S., MIERLE, K., AND OTHERS. 2014. Ceres solver. <http://ceres-solver.org>.
- ASHIKHMIN, M. AND PREMOŽE, S. 2007. Distribution-based BRDFs. Tech. rep., Department of Computer Science, University of Utah. March.
- ASHIKHMIN, M. AND SHIRLEY, P. 2000. An anisotropic phong brdf model. *J. Graph. Tools* 5, 2 (Feb.), 25–32.
- ASHIKHMIN, M., PREMOŽE, S., AND SHIRLEY, P. 2000. A microfacet-based BRDF generator. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH ’00. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 65–74.
- BAGHER, M., M., SOLER, C., AND HOLZSCHUCH, N. 2012. Accurate fitting of measured reflectances using a Shifted Gamma micro-facet distribution. *Computer Graphics Forum* 31, 4 (June), 1509–1518.
- BEN-ARTZI, A., OVERBECK, R., AND RAMAMOORTHY, R. 2006. Real-time BRDF editing in complex lighting. *ACM Trans. Graph.* 25, 3 (July).

- BILGILI, A., ÖZTÜRK, A., AND KURT, M. 2011. A general BRDF representation based on tensor decomposition. *Comput. Graph. Forum* 30, 8.
- BLINN, J. F. 1977. Models of light reflection for computer synthesized pictures. *SIGGRAPH Comput. Graph.* 11, 2 (July), 192–198.
- BOURLIER, C., BERGIN, G., AND SAILLARD, J. 2002. One- and two-dimensional shadowing functions for any height and slope stationary uncorrelated surface in the monostatic and bistatic configurations. *Antennas and Propagation, IEEE Transactions on* 50, 3, 312–324.
- BROWN, G. 1980. Shadowing by non-gaussian random surfaces. *Antennas and Propagation, IEEE Transactions on* 28, 6, 788–790.
- BURLEY, B. 2012. Physically-based shading at Disney. Tech. rep., Walt Disney Animation Studios. August.
- CARROLL, R. J. AND RUPPERT, D. 1998. *Translation and Weighting in Regression*. Chapman and Hall, N.Y.
- COOK, R. L. AND TORRANCE, K. E. 1981. A reflectance model for computer graphics. *SIGGRAPH Comput. Graph.* 15, 3 (Aug.), 307–316.
- DEBEVEC, P. 2001. Light probe image gallery. <http://www.pauldebevec.com/Probes/>.
- DEBEVEC, P. 2008. High-resolution light probe image gallery. <http://gl.ict.usc.edu/Data/HighResProbes/>.
- DORSEY, J., RUSHMEIER, H., AND SILLION, F. 2008. *Digital Modeling of Material Appearance*. Morgan Kaufmann Inc., San Francisco, CA, USA.
- DUPUY, J., HEITZ, E., IEHL, J.-C., POULIN, P., AND OSTROMOUKHOV, V. 2015. Extracting Microfacet-based BRDF Parameters from Arbitrary Materials with Power Iterations. *Computer Graphics Forum*, 10.
- EPA 2003. Method 8000C, SW846, Revision 3. Tech. rep., U.S. Environmental Protection Agency. March.
- GHOSH, A., HAWKINS, T., PEERS, P., FREDERIKSEN, S., AND DEBEVEC, P. 2008. Practical modeling and acquisition of layered facial reflectance. *ACM Trans. Graph.* 27, 5 (Dec.), 139:1–139:10.
- GRAYBILL, F. A. AND IYER, H. K. 1994. *Regression Analysis: Concepts and Applications*. Duxbury Press, Belmont, CA.
- KAUTZ, J. AND MCCOOL, M. D. 1999. Interactive rendering with arbitrary BRDFs using separable approximations. In *Rendering Techniques*.
- KOLDA, T. G. AND BADER, B. W. 2009. Tensor decompositions and applications. *SIAM Rev.* 51, 3 (Aug.), 455–500.
- LAFORTUNE, E. P. F., FOO, S.-C., TORRANCE, K. E., AND GREENBERG, D. P. 1997. Non-linear approximation of reflectance functions. In *Proceedings of SIGGRAPH*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 117–126.
- LAWRENCE, J., BEN-ARTZI, A., DECORO, C., MATUSIK, W., PFISTER, H., RAMAMOORTHY, R., AND RUSINKIEWICZ, S. 2006. Inverse shade trees for non-parametric material representation and editing. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 25, 3 (July).
- LAWRENCE, J., RUSINKIEWICZ, S., AND RAMAMOORTHY, R. 2004. Efficient BRDF importance sampling using a factored representation. *ACM Trans. Graph.* 23, 3 (Aug.), 496–505.
- LAZÁNYI, I. AND SZIRMAY-KALOS, L. 2005. Fresnel term approximations for metals. In *WSCG (Short Papers)* (2005-08-01). 77–80.
- LEE, D. D. AND SEUNG, H. S. 2000. Algorithms for non-negative matrix factorization. In *In NIPS*. MIT Press, 556–562.
- LÖW, J., KRONANDER, J., YNNERMAN, A., AND UNGER, J. 2012. BRDF models for accurate and efficient rendering of glossy surfaces. *ACM Transactions on Graphics (TOG)* 31, 1 (January), 9:1–9:14.
- MARSCHNER, S. R., WESTIN, S. H., LAFORTUNE, E. P. F., TORRANCE, K. E., AND GREENBERG, D. P. 1999. Image-based BRDF measurement including human skin. In *Proceedings of the 10th Eurographics Conference on Rendering*. EGWR'99. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 131–144.
- MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. In *ACM SIGGRAPH 2003 Papers*. SIGGRAPH '03. ACM, New York, NY, USA, 759–769.
- MCAULEY, S., HILL, S., HOFFMAN, N., GOTANDA, Y., SMITS, B., BURLEY, B., AND MARTINEZ, A. 2012. Practical physically-based shading in film and game production. In *ACM SIGGRAPH 2012 Courses*. SIGGRAPH '12. ACM, New York, NY, USA, 10:1–10:7.
- MCCOOL, M. D., ANG, J., AND AHMAD, A. 2001. Homomorphic factorization of BRDFs for high-performance rendering. In *SIGGRAPH*.
- NATI, N. S. AND JAAKKOLA, T. 2003. Weighted low-rank approximations. In *In 20th International Conference on Machine Learning*. AAAI Press, 720–727.
- NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental analysis of BRDF models. In *Proceedings of the Eurographics Symposium on Rendering*. Eurographics Association, 117–226.
- NIELSEN, J. B., JENSEN, H. W., AND RAMAMOORTHY, R. 2015. On optimal, minimal BRDF sampling for reflectance acquisition. *ACM Transactions on Graphics (TOG)* 34, 6 (November).
- NIS 2012. e-Handbook of Statistical Methods. Tech. rep., NIST/SEMATECH. April.
- PACANOWSKI, R., SALAZAR-CELIS, O., SCHLICK, C., GRANIER, X., PIERRE, P., AND ANNIE, C. 2012. Rational BRDF. *IEEE Transactions on Visualization and Computer Graphics* 18, 11 (Feb.), 1824–1835.
- PHARR, M. AND HUMPHREYS, G. 2010. *Physically Based Rendering, Second Edition: From Theory To Implementation*, 2nd ed. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.
- PRESS, W. H., TEUKOLSKY, S. A., VETTERLING, W. T., AND FLANNERY, B. P. 2007. *Numerical Recipes 3rd Edition: The Art of Scientific Computing*, 3 ed. Cambridge University Press, New York, NY, USA.
- ROMEIRO, F., VASILYEV, Y., AND ZICKLER, T. 2008. Passive reflectometry. In *Proceedings of the 10th European Conference on Computer Vision: Part IV*. ECCV '08. Springer-Verlag, Berlin, Heidelberg, 859–872.
- SCHLICK, C. 1994. An inexpensive BRDF model for physically-based rendering. *Computer Graphics Forum* 13, 233–246.
- SMITH, B. 1967. Geometrical shadowing of a random rough surface. *Antennas and Propagation, IEEE Transactions on* 15, 5, 668–671.
- SOLER, C. AND ROCHE, J.-C. 2014. High quality renderer (HQR). <http://artis.imag.fr/~Cyril.Soler/HQR/>.
- STEIGLEDER, M. AND MCCOOL, M. D. 2002. Factorization of the Ashikhmin BRDF for real-time rendering. *J. Graphics, GPU, & Game Tools* 7, 4, 61–67.
- TALVALA, E.-V., ADAMS, A., HOROWITZ, M., AND LEVOY, M. 2007. Veiling glare in high dynamic range imaging. *ACM Trans. Graph.* 26, 3.
- TOFALLIS, C. 2008. Least squares percentage regression. *Journal of Modern Applied Statistics Methods* 7, 526–534.
- VEACH, E. 1998. Robust Monte Carlo methods for light transport simulation. Ph.D. thesis, Stanford, CA, USA. AAI9837162.
- WALTER, B., MARSCHNER, S. R., LI, H., AND TORRANCE, K. E. 2007. Microfacet models for refraction through rough surfaces. In *Rendering Techniques*. Eurographics Association, 195–206.
- WANG, J., ZHAO, S., TONG, X., SNYDER, J., AND GUO, B. 2008. Modeling anisotropic surface reflectance with example-based microfacet synthesis. *ACM Trans. Graph.* 27, 3 (Aug.), 41:1–41:9.
- WHITE, H. 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica* 48, 4.

## Appendix: Volume Form Derivation

In the canonical coordinate system with surface normal  $\mathbf{n} = \mathbf{z}$ , assume the halfway vector  $\mathbf{h}$  is in the  $\mathbf{zx}$ -plane; i.e., perpendicular to  $\mathbf{y}$ . In other words, the azimuthal angle of the halfway vector is

$\phi_h = 0$ . Then the halfway vector can be parameterized by  $\mathbf{h} = \mathbf{z} \cos \theta_h + \mathbf{x} \sin \theta_h$ , and a perpendicular direction (also in the  $\mathbf{zx}$ -plane) by  $\mathbf{h}^\perp = \mathbf{z} \sin \theta_h - \mathbf{x} \cos \theta_h$ . To complete this orthonormal coordinate frame, let  $\mathbf{h}^* = \mathbf{h} \times \mathbf{h}^\perp = -\mathbf{y}$ . Then the light and view directions are given by

$$\begin{aligned} \mathbf{i} &= \mathbf{h} \cos \theta_d - (\mathbf{h}^\perp \cos \phi_d + \mathbf{h}^* \sin \phi_d) \sin \theta_d \\ \mathbf{o} &= \mathbf{h} \cos \theta_d + (\mathbf{h}^\perp \cos \phi_d + \mathbf{h}^* \sin \phi_d) \sin \theta_d \end{aligned}$$

or, in terms of components,

$$\begin{aligned} \mathbf{i}_x &= \sin \theta_h \cos \theta_d + \cos \theta_h \cos \phi_d \sin \theta_d \\ \mathbf{i}_y &= \sin \phi_d \sin \theta_d \\ \mathbf{i}_z &= \cos \theta_h \cos \theta_d - \sin \theta_h \cos \phi_d \sin \theta_d \\ \mathbf{o}_x &= \sin \theta_h \cos \theta_d - \cos \theta_h \cos \phi_d \sin \theta_d \\ \mathbf{o}_y &= -\sin \phi_d \sin \theta_d \\ \mathbf{o}_z &= \cos \theta_h \cos \theta_d + \sin \theta_h \cos \phi_d \sin \theta_d \end{aligned}$$

Let  $T(\theta_h, \theta_d, \phi_d)$  denote the above transformation going from  $(\theta_h, \theta_d, \phi_d)$  parameter space to  $(\mathbf{i}, \mathbf{o})$  output space. Then the *volume form*,  $d\Omega$ , induced by the above transformation is given by

$$d\Omega = \sqrt{\det(\partial T^T \partial T)} d\theta_h d\theta_d d\phi_d,$$

where  $\partial T$  is the  $(6 \times 3)$  Jacobian of the transformation  $T$ . The matrix product  $g = \partial T^T \partial T$  is the  $(3 \times 3)$  *metric tensor* of the transformation; the square root of its determinant scales the volume form after a coordinate transformation.

Taking derivatives and evaluating, we obtain

$$g = \text{diag}[2(\cos^2 \theta_d + \sin^2 \theta_d \cos^2 \phi_d), 2, 2 \sin^2 \theta_d]$$

so that

$$\begin{aligned} d\Omega &= \sqrt{\det g} d\theta_h d\theta_d d\phi_d, \\ &= \sqrt{8 \sin^2 \theta_d (\cos^2 \theta_d + \sin^2 \theta_d \cos^2 \phi_d)} d\theta_h d\theta_d d\phi_d. \end{aligned}$$

Note that this result does not depend on  $\theta_h$ .

## Appendix: Rendering Details

*Offline, CDF-based Rendering.* We sample our microfacet model directly into PBRT's `RegularHalfangleBRDF` representation for a `MeasuredMaterial` [Pharr and Humphreys 2010, Sections 8.6.2 and 9.2.4], which uses the same uniform parameterization internally of  $(\theta_h, \theta_d, \phi_d)$  with resolution  $90 \times 90 \times 180$ . Interpolation is not required when sampling our  $D$  and  $F$  factors; we linearly interpolate between tabulated values when sampling  $G$ .

PBRT constructs the CDF for this representation according to the process in [Pharr and Humphreys 2010, Section 14.6]. First, it re-samples each BRDF view-slice (indexed by  $\cos \theta_o$ ) into a lat-long map with resolution  $(n_\theta, n_\phi) = (90, 360)$ , converting values into luminance. Second, it constructs 2D CDFs for each view-slice with numerical integration of the BRDF luminance. Finally, it inverts the CDF to warp uniform low-discrepancy random numbers proportional to the BRDF. Our HQR and GPU renderers reimplement the same representation and CDF algorithm as in PBRT.

We construct a separate CDF for environment lighting from a lat-long map of resolution  $(n_\theta, n_\phi) = (180, 360)$ . Our Monte Carlo direct illumination estimate  $L_o$  at a shade point is:

$$L_o(\omega_o) \approx \sum_{s=1}^S \sum_{i=1}^{N_s} \frac{w_s(\omega_i) L(\omega_i) \rho(\omega_i, \omega_o) (\mathbf{n} \cdot \omega_i)}{N_s P_s(\omega_i)}, \quad (20)$$




Fig. 16. Direct rendering of our non-parametric microfacet representation from our MERL “uber-texture”. The top image shows our texture map, encoding BRDFs of all 100 MERL materials; the bottom shows the resulting rendering, in the Grace HDR environment map, after accumulating Monte Carlo samples for 60 seconds in our WebGL application.

where  $L$  is the incident radiance (visibility-weighted light emission, for direct illumination) at the shade point, and we draw  $N_s$  samples proportional to the sampling PDFs  $P_s(\omega)$  of the  $S$  sampling strategies. Recall that we use up to three of the following sampling strategies: BRDF sampling, emission sampling, and a uniform hemispherical sampling. Here,  $w_s$  is the MIS weight, and we choose a power heuristic weight with parameter  $\beta = 2$ :

$$w_s(\omega) = (N_s P_s(\omega))^\beta \bigg/ \sum_{s'=1}^S (N_{s'} P_{s'}(\omega))^\beta. \quad (21)$$

*Direct Rendering.* We also implemented an interactive WebGL renderer (Fig. 16). At each pixel, it samples hemispherical light directions using multiple importance sampling and evaluates the BRDF analytically using Eq. 3. Monte Carlo samples of the product of the BRDF and HDR environment map evaluation are progressively accumulated for each light direction. This simple method converges reasonably quickly even for highly specular materials. The  $D$ ,  $F$ , and  $G$  factors, along with  $\rho_d$  and  $\rho_s$ , are queried from a single  $512 \times 128$  *rgb* “uber-texture” which encodes all 100 MERL materials. Factor components and diffuse/specular coefficients are represented as a `Float32Array` in WebGL.