Parameterization of Tabulated BRDFs

Agatha Mallett
University of Utah
agatha@geometrian.com

Cem Yuksel
University of Utah
cem@cemyuksel.com

ABSTRACT

Measured or simulated data for bidirectional reflectance distribution functions (BRDFs) of various materials is typically stored in a tabulated form by discretizing incoming and outgoing directions over the hemisphere. Since reflected radiance is usually strongly nonuniform over the hemisphere, a uniform discretization of a BRDF usually leads to poor quality at a given memory cost. Unfortunately, there is no universally optimal nonuniform parameterization. In this paper, we provide a mathematical framework for analyzing them with respect to the BRDFs they represent. We use our framework to motivate a parameterization based on Bézier curves: a novel approach that adapts to any given BRDF. Our test results using measured and analytical BRDF models show that it provides superior computational and spatial efficiency for equal quality, as compared to simpler alternatives.

1. INTRODUCTION

A bidirectional reflectance distribution function (BRDF) defines the reflectance properties of a material for all combinations of input and output directions. Therefore, modeling this function accurately is crucial for realistically reproducing material appearance. One way to achieve this is using measured data.

Measured BRDFs are typically stored in a tabulated form. However, because of BRDFs' non-uniformity, not all table axes have the same importance. For example, direction combinations near specular- or retro-reflections usually comprise high-frequency variation. If the distribution of tabulated BRDF samples over the hemisphere does not match these frequencies, parts of the hemisphere get over-sampled while other parts get under-sampled.

The distribution of the BRDF samples can be adjusted using a parameterization. However, since BRDFs can differ substantially, a single parameterization cannot be optimal for all BRDFs. While the existence of non-uniform parameterizations is mentioned in previous work (e.g. in [Matusik et al. 2003]), efficient parameterization methods for tabulated BRDFs have not been explored in computer graphics to our knowledge.

In this paper, we discuss parameterizations of BRDFs and introduce a new mathematical framework for developing and evaluating them. Using this framework, we introduce *Bézier parameterization*, a computationally efficient model that can adapt to a variety of BRDFs. We demonstrate its higher fidelity for real-world examples–in particular, for measured data–for validating our framework.

CGI'16 June 28-July 01, 2016, Heraklion, Crete, Greece

ISBN 978-1-4503-4123-3/16/06…$15.00

DOI: http://dx.doi.org/10.1145/2949035.2949047

2. BACKGROUND

The BRDF f_r is defined as the differential ratio of outgoing radiance dL_o in the direction of ω_o to the incoming irradiance in the direction of ω_i at a surface point with surface normal $\mathbf{n}$. In general, the BRDF is a 4D or higher function, but for isotropic materials the BRDF can be reduced to 3D.

There are many empirical BRDF models in computer graphics [Blinn 1977; Ward 1992; Lafortune et al. 1997; Neumann et al. 1999; Ashikhmin and Shirley 2000] as well as physically inspired models based on analysis of small-scale surface structure [Cook and Torrance 1982; He et al. 1991; Oren and Nayar 1994; Löw et al. 2012]. While these analytic BRDF formulations provide a compact representation, they often fail to represent the subtle variations of a realistic material appearance [Ngan et al. 2005]. Realism can be achieved by measuring BRDFs of real materials [Matusik et al. 2003]. For approximating measured data, researchers investigated fitting analytical BRDFs [Ngan et al. 2005; Bagher et al. 2012; Brady et al. 2014], multi-lobe models [Lafortune et al. 1997; Rump et al. 2008; Yu et al. 2011], spherical Gaussian lobes [Wang et al. 2009], and dimensionality-reduction approaches [Ashikhmin 2007; Pacanowski et al. 2012].

The generic coordinate frame for representing the BRDF is the *normal-vector basis* (Figure 1). The incoming direction ω_i is represented using two angles θ_i and ϕ_i; and the outgoing direction ω_o is represented similarly with θ_o and ϕ_o, such that $0 \leq \theta_{i,o} < \pi/2$ and $0 \leq \phi_{i,o} < 2\pi$. Isotropic BRDFs can be represented by replacing ϕ_i and ϕ_o with a single coordinate $\phi_\Delta = \phi_o - \phi_i$, such that $0 \leq \phi_\Delta \leq \pi$.

A popular alternative is the *half-vector basis* [Rusinkiewicz 1998], based on the half vector $\mathbf{h}$ between ω_i and ω_o, defined as $\mathbf{h} = (\omega_i + \omega_o)/\left|\omega_i + \omega_o\right|$. $\mathbf{h}$ is represented using θ_h and ϕ_h (Figure 1). θ_d and ϕ_d are evaluated on the rotated hemisphere aligned with $\mathbf{h}$. In this formulation, $0 \leq \theta_{h,d} < \pi/2$, $0 \leq \phi_h < 2\pi$, and $0 \leq \phi_d < \pi$. Isotropic BRDFs are represented without ϕ_h. More azimuthal samples are automatically provided near the specular lobe $\mathbf{h} \approx \mathbf{n}$, the primary reason it has been favored by previous work [Matusik et al. 2003; Edwards et al. 2006]. On the other hand, the half-vector basis does not automatically provide denser longitudinal sampling. Also, there are values of θ_h, θ_d, and ϕ_d that correspond to invalid directions for ω_i and ω_o. Therefore, more than 38% of the table's data is thus wasted.

Other bases are possible. For example, several related barycen-

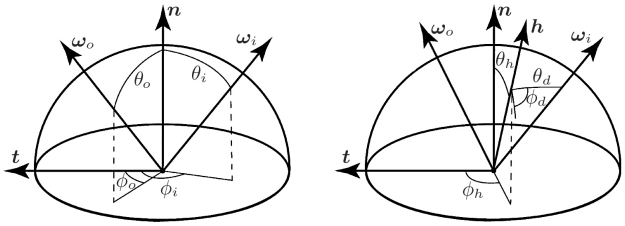

Normal-Vector Basis Half-Vector Basis

Figure 1. *Two major coordinate bases for representing reflectance function parameterizations, where ω_i and ω_o are incoming and outgoing directions, $\mathbf{n}$ is the surface normal and $\mathbf{t}$ is a tangent vector.*

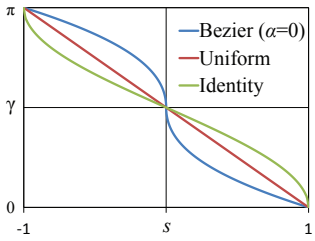

Figure 2. *Plots of three parameterizations. Our framework allows us to reason about and derive good models, such as the Bézier parameterization pictured here.*

tric bases were proposed by [Stark et al. 2005], and a variety of bases were tested by [Marschner 1998]. For clarity of exposition, in this paper we focus our analysis and provide example results for the normal-vector basis only. However, our analysis and framework apply to any basis based on angles.

3. PARAMETERIZATION FRAMEWORK

Given vectors ω_i and ω_o, the respective θ and ϕ angles in normal-vector basis are typically computed using dot-products of these vectors with $\mathbf{n}$ or $\mathbf{t}$ (Figure 1). Let γ denote one of the angles we would like to parameterize (one of $\theta_{i,o}$, $\phi_{i,o}$ or one of $\theta_{h,d}$, $\phi_{h,d}$). For $\gamma \in [0, \pi]$ we can write $\gamma = \arccos(\cos\gamma)$ where the $\cos\gamma$ term can be computed inexpensively using a dot-product of two vectors. We can introduce $s \in [-1, 1]$ and a bijection $F : [-1, 1] \to [-1, 1]$ for representing the value of the dot-product, such that $\cos\gamma = F^{-1}(s)$. Thus:

$$s = F(\cos\gamma) \quad \text{and} \quad \gamma = \arccos\left(F^{-1}(s)\right) . \tag{1}$$

This formulation allows us to relate, via a parameterization F, the angle γ and the parameter s. In other words, the separation between samples along γ is determined by $d\gamma/ds$ and ultimately by F. This allows existing parameterizations to be understood and reasoned about in terms of angular density.

Using this formulation, a simple parameterization model is the *identity parameterization* $F_{id}(x) = x$, which produces a sample separation of $d\gamma/ds = -1/\sqrt{1-s^2}$. Since $d\gamma/ds$ approaches negative infinity near $s = \pm 1$, fewer samples are placed as γ approaches 0 or π. Thus, the identity parameterization usually provides poor sampling, since these regions are direct specular lobes in most BRDFs with high-frequency changes.

Another simple parameterization is the *uniform parameterization*, placing samples with constant angular separation. The constant derivative $d\gamma/ds = -\pi/2$, yields $F_{eq}(x) = 1 - (2/\pi)\arccos(x)$. Uniform parameterization is an improvement over Identity parameterization, but still tends to oversample low-frequencies while under-sampling high-frequencies. Additionally, it introduces an expensive $\arccos$ function, which is computed every time the BRDF is accessed.

Motivated by these two examples, we observe that by carefully picking a transformation function F, we can in-principle produce any desired sample rate $d\gamma/ds$. Of course, one should balance sampling performance and computational complexity–computationally complex functions can be produced by, for example, integrating some forms of $d\gamma/ds$. We introduce the *Bézier parameterization* as a simple formulation that can produce a desirable sample distribution. Using a specific quadratic Bézier curve

$$\mathbf{B}(t) = (1-t)^2\mathbf{P}_0 + 2t(1-t)\mathbf{P}_1 + t^2\mathbf{P}_2 \tag{2}$$

with three control points

$$\mathbf{P}_0 = (0, 0), \quad \mathbf{P}_1 = (1, \alpha), \quad \text{and} \quad \mathbf{P}_2 = (1, 1), \tag{3}$$

where α is a parameter that can be adjusted according to the properties of the BRDF, the transformation is defined as

$$F_{bz}(B_x(t)) = \quad B_y(t) \qquad \text{for} \quad t \in [0, 1] \quad \text{and} \tag{4}$$

$$F_{bz}(-B_x(-t)) = -B_y(-t) \qquad \text{for} \quad t \in [-1, 0) \ , \tag{5}$$

yielding

$$F_{bz}(x) = \mathrm{sgn}(x)(2 - 2\alpha)\left(1 - \sqrt{1 - |x|}\right) - (1 - 2\alpha)x.$$

Note that $F_{bz}(x)$ is monotonically increasing for $\alpha \in [0, 1]$ within the range $x \in [-1, 1]$. When $\alpha = 1$ Bézier parameterization becomes identical to identity parameterization and when $\alpha \approx 0.5841$ it approximates uniform parameterization. However, unlike either, it supports different angular densities by further variation of α, allowing to non-linearly and flexibly adjust sample density (Figure 2). In fact, Bézier parameterization can represent a series of mappings between the blue and the green curves in Figure 2.

4. IMPLEMENTATION AND RESULTS

We implemented the above parameterizations using GLSL in a deferred renderer, with BRDF samples stored in four-byte RGB+exponent. We tested using all 100 materials in the MERL BRDF database [Matusik et al. 2003] as well as analytical BRDF models [Phong 1975; Ward 1992]. We use 3D textures for isotropic BRDFs with texture dimensions corresponding to θ_i, θ_o, and ϕ_Δ using different parameterizations with resolutions N, N, and $2N$ respectively. For anisotropic BRDFs, we emulate 4D textures by packing slices into 2D textures. The additional texture dimension of ϕ_i is always sampled using the uniform parameterization with resolution $2N$. For Bézier parameterization, we use the same α value α_θ for both θ_i and θ_o, and another value α_ϕ for ϕ_Δ. We construct the BRDF using different combinations of α_θ and α_ϕ to determine the right α_θ and α_ϕ values for a given BRDF.

For simplicity, our coarse evaluation of each combination is based on the total intensity of the generated BRDF texture, measured in angular space. This total intensity metric is directly indicative of how well high-energy portions of the BRDF are sampled. Roughly speaking, the larger the intensity, the more resolution is allocated to the key (high-energy) features of the BRDF. We pick the α_θ and α_ϕ combination that produces the largest total intensity without unduly affecting quality elsewhere.

We found that the metric favors $\alpha_\phi = 0$ for all materials in the MERL BRDF database, allocating more samples near the specular, Fresnel, and retroreflection peaks of the materials, which correspond to $\phi_\Delta \approx 0$ and $\phi_\Delta \approx \pi$. We also found that the most com-

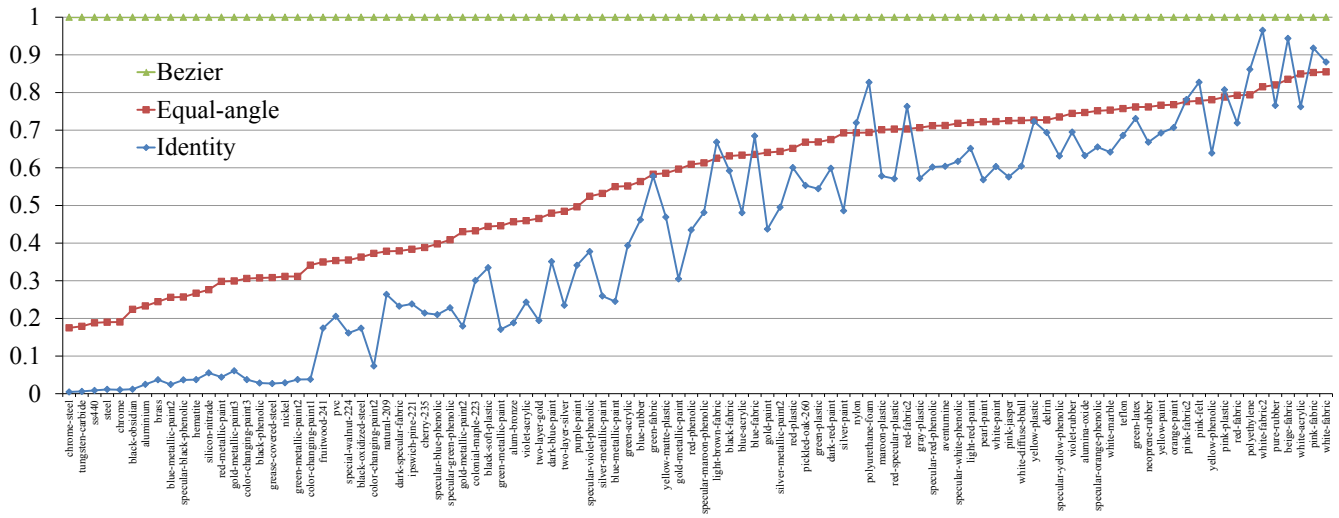

Figure 3. *The total BRDF texture intensity of the Bézier, uniform, and identity parameterizations, normalized to the total intensity of Bézier.*

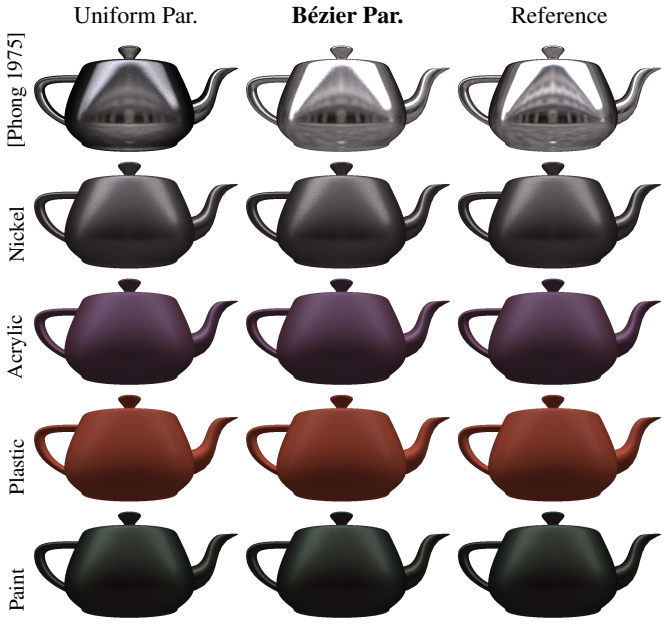

Figure 4. *Utah teapot rendered using several BRDFs and environment mapping with 2000 samples per pixel using* $64 \times 64 \times 128$ *BRDF textures with (left) uniform parameterization and (middle) our Bézier parameterization. (Right) the reference images are computed using* $256 \times 256 \times 512$ *BRDF textures or an analytic model.*

Figure 5. *An anisotropic BRDF rendered using Bézier parameterization under different illuminations.*

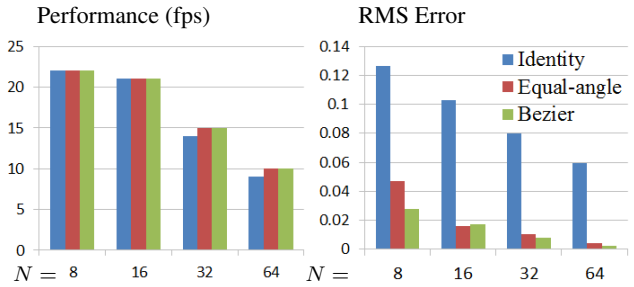

Figure 6. *The performance of different parameterizations measured on an NVIDIA GTX 580M GPU and their RMS errors for* 10 *environment rays per pixel.*

mon α_θ value is either $\alpha_\theta = 0$ (allocating more samples for the Fresnel peak) or $\alpha_\theta = 1$ (allocating more samples for ω_i, o directions near the surface normal). However, this trend is less strong; varying values are due to the conflicting goals of representing both high-angle and glancing reflections well.

Figure 3 shows a comparison of parameterizations based on this metric for all materials in the MERL BRDF database. The ratios of the total intensities show that Bézier parameterization produces significantly higher-intensity textures, suggesting that it can better allocate more samples for the high-energy portions of the BRDF data. The differences between total intensities are especially high for highly glossy materials such as metals, a classic difficult case for tabular BRDFs.

Figure 4 shows rendered images with example BRDFs using Bézier and uniform parameterizations.The difference between the two is particularly prominent for materials with sharp specular lobes, where Bézier parameterization can allocate more samples to it. For other materials, the differences between the two parameterizations are less noticeable. However, Bézier parameterization produces a smaller RMS error in all cases. Figure 5 shows examples of an anisotropic BRDF rendered using Bézier parameterization.

We compare parameterizations' performance and RMS error for different BRDF table resolutions in Figure 6. Although the square root calculation in Bézier parameterization is substantially faster than trigonometric functions, we found that performance is largely bound by working-set size, leading to broadly identical performance characteristics among the three methods presented here. However, Bézier parameterization achieves nearly an order-of-magnitude better accuracy, particularly for higher-resolution data.

In all comparisons, we found that Bézier parameterization out-

Trilinear Filtering Cubic B-spline Filtering

Figure 7. *The effect of higher order filtering. Notice the elliptical contours on the left image due to linear interpolation.*

performs simpler parameterizations. This demonstrates that our framework is capable of generating new parameterizations with superior efficacy at representing BRDFs, particularly those that are highly specular. However, Bézier parameterization does not necessarily provide the optimal sample distribution for any BRDF and no single parameterization can be optimal for all BRDFs.

Particularly prominent for highly specular BRDFs, sampling using trilinear filtering (or quad-linear for anisotropic BRDFs) produces interpolation artifacts (Figure 7). Therefore, it is important to use higher-order reconstruction. In our implementation, we used cubic B-splines [Keys 1981] to filter the sample value, which can also be implemented efficiently on the GPU [Sigg and Hadwiger 2005]. Note that, this filtering problem (and its solution) apply to all tabular representations, and it is orthogonal to our aim of developing a mathematical framework for BRDF parameterization.

5. CONCLUSION AND FUTURE WORK

In this paper we introduce a mathematical framework for generating high-quality parameterizations for representing BRDF data and validate it by using it to develop the well-performing Bézier parameterization.

While we presented results on BRDFs, our parameterization framework can be easily extended to *bidirectional scattering distribution functions* (BSDFs) as well. An interesting future work would be experimenting with other parameterizations that could be generated using our framework and other bases.

6. ACKNOWLEDGMENTS

This work was supported in part by NSF grant #1409129. Lance Williams gave feedback on early versions of this work.

References

ASHIKHMIN, M., AND SHIRLEY, P. 2000. An anisotropic phong BRDF model. *Journal of Graphics Tools 5*, 2 (Feb.), 25–32.

ASHIKHMIN, M. 2007. Distribution-based BRDFs. Tech. rep., University of Utah.

BAGHER, M. M., SOLER, C., AND HOLZSCHUCH, N. 2012. Accurate fitting of measured reflectances using a shifted gamma micro-facet distribution. *CG Forum 31*, 4 (June), 1509–1518.

BLINN, J. F. 1977. Models of light reflection for computer synthesized pictures. *SIGGRAPH C. Grap. 11*, 2 (July), 192–198.

BRADY, A., LAWRENCE, J., PEERS, P., AND WEIMER, W. 2014. genBRDF: Discovering new analytic BRDFs with genetic programming. *ACM T. on Graph. 33*, 4 (July), 114:1–114:11.

COOK, R. L., AND TORRANCE, K. E. 1982. A reflectance model for computer graphics. *ACM Trans. Graph. 1*, 1 (Jan.), 7–24.

EDWARDS, D., BOULOS, S., JOHNSON, J., SHIRLEY, P., ASHIKHMIN, M., STARK, M., AND WYMAN, C. 2006. The halfway vector disk for BRDF modeling. *ACM Transactions on Graphics 25*, 1, 1–18.

HE, X. D., TORRANCE, K. E., SILLION, F. X., AND GREENBERG, D. P. 1991. A comprehensive physical model for light reflection. *SIGGRAPH Comput. Graph. 25*, 4 (July), 175–186.

KEYS, R. G. 1981. Cubic convolution interpolation for digital image processing. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 1153–1160.

LAFORTUNE, E. P. F., FOO, S.-C., TORRANCE, K. E., AND GREENBERG, D. P. 1997. Non-linear approximation of reflectance functions. In *Proc. of SIGGRAPH '97*, 117–126.

LÖW, J., KRONANDER, J., YNNERMAN, A., AND UNGER, J. 2012. BRDF models for accurate and efficient rendering of glossy surfaces. *ACM Trans. Graph. 31*, 1 (Feb.), 9:1–9:14.

MARSCHNER, S. R. 1998. *Inverse rendering for computer graphics*. PhD thesis, Stanford University.

MATUSIK, W., PFISTER, H., BRAND, M., AND MCMILLAN, L. 2003. A data-driven reflectance model. *ACM Transactions on Graphics 22*, 3 (July), 759–769.

NEUMANN, L., NEUMANN, A., AND SZIRMAY-KALOS, L. 1999. Reflectance models with fast importance sampling. *Computer Graphics Forum 18*, 4, 249–265.

NGAN, A., DURAND, F., AND MATUSIK, W. 2005. Experimental analysis of brdf models. In *Proc. EGSR*, no. 16th, 117–226.

OREN, M., AND NAYAR, S. K. 1994. Generalization of lambert's reflectance model. In *Proc. of SIGGRAPH '94*, 239–246.

PACANOWSKI, R., CELIS, O. S., SCHLICK, C., GRANIER, X., POULIN, P., AND CUYT, A. 2012. Rational BRDF. *IEEE Trans. on Vis. and Comp. Graphics 18*, 11 (Nov.), 1824–1835.

PHONG, B. T. 1975. Illumination for computer generated pictures. *Communications of the ACM 18*, 6, 311–317.

RUMP, M., MÜLLER, G., SARLETTE, R., KOCH, D., AND KLEIN, R. 2008. Photo-realistic rendering of metallic car paint from image-based measurements. *Computer Graphics Forum 27*, 2 (Apr.), 527–536.

RUSINKIEWICZ, S. M. 1998. A new change of variables for efficient brdf representation. In *Rendering Tech. '98*. 11–22.

SIGG, C., AND HADWIGER, M. 2005. Fast third-order texture filtering. In *GPU Gems 2*. 313–329.

STARK, M. M., ARVO, J., AND SMITS, B. 2005. Barycentric parameterizations for isotropic BRDFs. *IEEE Transactions on Visualization and Computer Graphics 11*, 2, 126–138.

WANG, J., REN, P., GONG, M., SNYDER, J., AND GUO, B. 2009. All-frequency rendering of dynamic, spatially-varying reflectance. *ACM Trans. Graph. 28*, 5 (Dec.), 133:1–133:10.

WARD, G. J. 1992. Measuring and modeling anisotropic reflection. *SIGGRAPH Comput. Graph. 26*, 2 (July), 265–272.

YU, C., SEO, Y., AND LEE, S. W. 2011. Global optimization for estimating a multiple-lobe analytical BRDF. *Computer Vision and Image Understanding 115*, 12 (Dec.), 1679–1688.