Joint MAM MANER APPMAT Workshop on Material Appearance Modeling (2026)
D. Muselet and R. Pacanowski and H. Rushmeier (Editors)

Approximating Sparse BRDF Measurements

R. Yu[1,2], F. Margall[2], P. Mézières[2], R. Pacanowski[2], C. Soler[2]

[1]Owens Corning
[2]Inria, France

Abstract

Sparse BRDF measurements are quite common in the optical community, which uses spectro-goniophotometers. Although physics-based analytical models can approximate this type of data, they have the disadvantage of making assumptions about the underlying materials and, for the most part, requiring optimization algorithms with no guaranteed convergence. In this short paper, we compare (in terms of time and quality) various numerical methods (Spherical Harmonics, Wavelets, Radial Basis Functions, kernel-based methods, etc.) that allow us to approximate or even interpolate this type of sparse measurement without any prior assumptions about the material's optical behavior. To precisely assess the reconstruction quality of each method, we apply them on decimated dense measurements along different types of measurement layouts coming from known BRDF acquisition devices.

1. Motivation

Although digital cameras have become commonplace, BRDF databases do not necessarily provide dense measurements (e.g., Matusik et al. [MPBM03]) from an angular perspective. This is particularly the case when these measurements come from more traditional optical instruments, grouped under the name spectro-goniophotometer, or when the BRDF measurements are 4D (e.g., Ngan [NDM05]) and remains sparse due to the curse of dimensionality, measuring time and cost. In the case of sparse measurements, one can either rely on fitting analytical BRDF models or on interpolating using numerical methods. Analytical (e.g., microfacet-based [WMLT07]) or neural [SRRW21] models provide a closed-form expression that can approximate the measurements thanks to different numerical optimization methods (e.g., Levenberg-Marquardt, Adam). Since the optimization problem is often non-linear, the result depends heavily on the starting point (i.e., only local convergence can be guaranteed), and the result is constrained by the inherent BRDF model or by the learning space of the neural reconstruction method.

In this paper, we focus on comparing methods that approximate the measurements in a vector space of spherical functions, which encompasses a lot of traditional techniques. Within this class, the different methods vary in the function types (density of the function space, bandwidth localization, etc) and in the fitting algorithm itself (stability, computational complexity, memory cost). Since BRDF measurements are generally not uniform in the space of directions and follow geometric patterns that are characteristic of the measurement apparatus, different basis functions and fitting algorithms may behave very differently on each of them.

Since the goal of this communication is to provide practical insights into which method should be used preferably on various measurement setups (and ultimately publish the code for each), we limit ourselves to a finite spectrum of relevant options:

- A kernel-based interpolation methods
- Spherical Harmonics [WAT92, SBN15]
- Radial Basis Functions [ZREB06, FVH14]
- Wavelet Compressive Sensing [SS95, CBP07]
- Spherical subdivision.

We will compare these methods over various sparse measurement layouts: the MERL anisotropic layout [NDM05], the Coupole acquisition system [?], and the PAB gonio measurements [AB10]. The different reconstruction methods will be compared over dense materials measurements provided by the UTIA material database [VF18] (serving as a reference), which we subsample along these sparse layouts.

Furthermore, we separate for each technique the computational complexity of computing the representation itself (*e.g.* time to compute coefficients of Spherical Harmonics), from that of evaluating this representation for any light-view direction pair. Finally, we compare and evaluate their fitting precision both in BRDF space as L_2 and rendered images space using a perceptual image distance.

2. Methodology

Dimensionality. BRDFs are inherently 4-dimensional functions. However when measurement devices provide separable data (several measurements for a given ω_i or ω_h, it becomes possible to reconstruct each "outgoing lobe" in 2D and then interpolate theses w.r.t. to the incident direction.

In this short paper, we focus mainly on 2D slices (i.e., a lobe) for separable BRDFs aligned along ω_i. The extension of each method to non-separable 3D or 4D measurements is discussed in Section **??**.

Measurement layouts. BRDF acquisition setups have different layouts (also called geometrical configurations) of measurements. As shown in Figure 1, which exhibits 3 examples of such a layout, some are more regular/uniform than others, and some have holes due to mechanical constraints.

We evaluate, for each tested method, its reconstruction error w.r.t. to each layout, for a material *densely* measured that we use as a reference. We chose for this the dense UTIA database [FVH14].

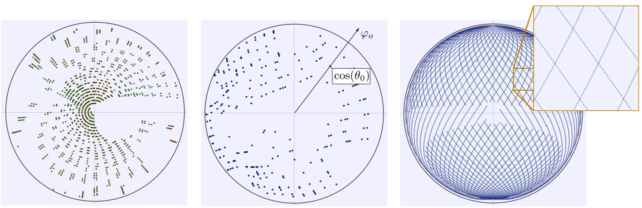

MIT ("Red Velvet") Coupole PAB

Figure 1: Three examples of layout for measurement points coming from actual measurement benches. *Left:* anisotropic "red velvet" (MIT) [NDM05]; *Middle:* La Coupole data [**?**]; *Right:* PAB measurement [AB10]. Even measurements with high point density (right) remain highly anisotropic and challenge some interpolation methods in their own way.

2.1. 2D Reconstruction/interpolation methods

We compare multiple methods to reconstruct an unknown BRDF lobe $\omega \mapsto f(\omega)$ from measurement samples (ω_i, y_i), for $i \in [1..N]$. The computational complexity and parameters for each method are listed in table 1.

Nearest neighbor averaging. Given a kernel function $\kappa(\omega, \omega')$ with local support, the reconstructed function is

$$f(\omega) = \sum_{i \in n(\omega)} k(\omega, \omega_i) y_i / \sum_{i \in n(\omega)} k(\omega, \omega_i),$$

where $n(\omega)$ is the set of n closest measurements for direction ω. Our implementation scales up to millions of measurements, by associating two layers: samples are first stored into directional bins, which resolution ensures an average of 50 to 200 samples per bin, then neighbor samples of any given direction are collected using a max-heap stack that feeds on a priority queue of bins.

Local Gaussian kriging. Gaussian kriging aims at interpolating the measurements by a mixture of Gaussian kernels:

$$f(\omega) = \sum_{i=1}^{N} \alpha_i g(\omega . \omega_i)$$

Since the number of kernels is exactly the number of samples (kernels are centered on samples) the solution is simply

$$\alpha = \mathbf{M}^{-1} y,$$

where $\mathbf{M} = \|g(\omega_i, \omega_j)\|_{ij}$. To avoid a prohibitive cost of solving a large linear system, we only collect the n closest samples for each output direction ω and solve one $n \times n$ linear system to determine $f(\omega)$ for each direction ω. The method being sensitive to noise (at close sample pairs) we perturbate $\mathbf{M}$ with matrix $\delta \mathbf{I}$ with $\delta \ll 1$. Since δ is chosen to be very small (typically $\delta = 10^{-3}$), the method can still be considered to be interpolating.

Wavelet compressive sensing. Given a space of wavelets on the sphere, represented by its synthesis operator $\mathbf{Q}$ (transforming wavelet coefficients into a spherical function) we aim at optimizing

$$\alpha^* = argmin_{\alpha} \|\mathbf{Q}\alpha - \mathbf{y}\|_2^2 + \lambda \|\alpha\|_1.$$

Wavelets having local support in both space and frequency, this allows to fit the input samples using a function of minimal local variation. We discarded the faster SpaRSA and FISTA algorithms [BT09] for a slower but much more stable iterative thresholding algorithm which steps are:

1. enforcing values: $f(\omega_i) = y_i$
2. wavelet analysis: $\alpha = \mathbf{Q}^{-1} f$
3. soft-thresholding α
4. wavelet synthesis: $f = \mathbf{Q}\alpha$

This algorithm requires wavelets to be orthogonal for better stability [Mey93], which on the sphere trades for rotational invariance, localization or smoothness [MP11]. Equivalently we chose to first convert f to $\theta - \phi$ space and use 2D Daubechie-12 wavelets for the fitting before turning back the result to the spherical domain.

Spherical harmonic fitting. The spherical lobe is approached by a linear combination of spherical harmonics as

$$f(\omega) = \sum_{l=0}^{L} \sum_{-l \le m \le l} \alpha_l^m \mathcal{Y}_l^m(\omega),$$

were $\mathcal{Y}_l^m$ is the harmonic of degree l and order m. The optimization is performed as

$$\alpha = argmin_{\alpha} \|\mathbf{M}\alpha - \mathbf{y}\|^2 + \|\Gamma\alpha\|^2,$$

where $\mathbf{M}_{ij} = \mathcal{Y}_i(\omega_j)$ and Γ acts as a Tikhonov regularization matrix to limit Gibbs effects.This is a very classical optimization problem which solution is known to be

$$\alpha = (\mathbf{M}^\intercal \mathbf{M} + \mathbf{\Gamma}^\intercal \mathbf{\Gamma})^{-1} (\mathbf{M} + \mathbf{\Gamma}) \mathbf{y}.$$

For the regularization, we use a diagonal matrix that multiplies each α_l^m by e^{-l^2/σ^2}.

Spherical subdivision. This methods directly interpolates the samples in the spherical domain with a function with good derivability. It follows these steps:

1. build a spherical Delaunay triangulation of the measurement samples
2. apply several steps of triangular subdivision to this mesh
3. export the function using barycentric interpolation at the finest scale

The subdivision scheme we're using is the *modified butterfly* scheme. Equivalent results are obtained with *loop* [WW01]. This method naturally fills holes.

3. Results

2D slice reconstruction. The convergence of all methods for an increasing number of uniformly selected directions on the hemisphere is presented on Figure 3, while we show in table 2 the error maps in (θ, ϕ) space for sparse uniform sampling of the hemispheres for all five methods. Table 3 compares the L_2 reconstruction error for the different layouts of Figure 1 as well as the SSIM distance for rendered spheres with the reconstructed data.

Among the five methods, Subdivision achieves the best overall reconstruction accuracy, with the lowest relative L2 error and SSIM dissimilarity scores on most layouts, while nearest neighbor averaging is the fastest method across all tested configurations. A clear trade-off between accuracy and speed is observed: although Subdivision provides the best reconstruction quality, its computation time increases most steeply as the number of measurements grows. In Table 2 and Table 3, where the number of grid points for Wavelet Compressive Sensing is manually reduced, the method shows smaller differences in accuracy compared to the other methods than what is observed in Figure 3. Finally, reconstruction quality depends heavily on the spatial distribution of input measurements, as evidenced by the large differences in error between the MIT and PAB layouts across all methods. As shown in Figure 1, the Pab layout provides a dense and uniform coverage of hemisphere, whereas the MIT layout leaves large angular regions unsampled, resulting a more irregular distribution after downsampling.

3D reconstruction. We tested 3D Gaussian krigging for Isotroptic Measurements. Contrary to the method of Zickler et al. [ZREB06], which modifies the Rusinkiewicz parametrization to establish a 3D Euclidean distance function that is used for the RBF kernel, we define our distance function as the geodesic distance between two bidirections in $S \times S$. This distance function guarantees the absence of deformation on the bi-hemispherical domain and also preserves Helmholtz reciprocity and optionally bilateral symmetry properties. Furthermore, we adapt the bandwidth (σ_h and σ_d) of Gaussian kernels between the two half-vector and difference vector subdomains. In order to simulate sparse measurements on MERL materials, we construct the 3D RBF representation using a decimated version of the original MERL data. The decimation is done by first removing a subset of θ_d slices, with 10% kept (i.e., this is roughly equivalent to take less photograph during the measurement process) and then removing data points uniformly, on each of the remaining slices (2% kept on each slice). This is more or less equivalent to keep 0.2% of all measurements. Figure 2 illustrates interpolation results using this method on two different MERL materials.

4. Discussion

Automatic parameter selection is applied before each method (except subdivision) to adapt to different input measurements. For both local averaging and Gaussian krigging, the kernel variance σ is derived from the angular distance to a fixed number of $n = 20$ nearest neighbours. For the later the σ is influenced by four descriptors: median far-neighbour distance, coefficient of variation of inter-sample spacing, local roughness, and noise ratio. The degree of spherical harmonics depends on the number of measurements, balancing over-fitting and computation time at higher degrees against over-smoothing at lower degrees. And for wavelet compressive sensing, the smoothness parameter λ and the convergence tolerance ϵ are set as piece-wise constant functions of the number of measurements; the size of the 2D grid is not treated as a free parameter to avoid imposing assumptions on the frequency of the data. Although some methods could benefit from additional parameters (*e.g.,* number of neighbors, (θ, ϕ) parameter space tabulation, etc), finding a formula that jointly accounts for the various materials remains an open problem, and the parameters described here represent compromises rather than optimal solutions.

Preserving physical constraints. While Helmholtz reciprocity is naturally encoded in the Rusinkiewicz parameterization, interpolating (θ, ϕ) slices doesn't guarantee this property. Contrarywise, albedo (or hemispherical reflectivity) constraints can easily be preserved by interpolating prescaled (θ, ϕ) slices. Non-negativity is not really a problem since it can be guaranteed by interpolating the measured data under a $y \mapsto \ln(1 + y)$ transform.

Interpolating between 2D "outgoing" slices also introduces an additional difficulty: In standard spherical coordinates, the specular peak shifts across slices as the illumination varies; linear interpolation between misaligned peaks then produces spurious intermediate lobes, a *ghosting* artefact common to interpolation problems. This can be addressed either through optimal-transport [BvdPPH11] methods or through a reparameterization that aligns the peaks — most notably the Rusinkiewicz [Rus98] half-vector (θ_h, ϕ_h) - difference (θ_d, ϕ_d) one, which aligns the specular peak at $\theta_h = 0$.

Non-separable data. Most of the methods above can be ported to 3D or 4D with limited modifications: This is the case for kernel interpolation (bins become 3D/4D), Gaussian krigging (3D/4D Gaussian kernels), Wavelet C.Sensing (3D/4D Daubechie wavelets instead of 2D, but probably less stability). Spherical harmonics do not exist in 3D or 4D, but

Method	Interpolation behavior	Fitting complexity	Evaluation complexity	Parameters
Nearest neighbor averaging	Approximating	N/A	$\mathcal{O}\,(N+n)$	n, σ
Gaussian krigging	Interpolating	$\mathcal{O}\,(N+n^3)$	$\mathcal{O}\,(n)$	n, σ
Wavelet Compressive Sensing	Approximating	$\mathcal{O}\,(N)$	$\mathcal{O}\,(1)$	λ
Spherical Harmonics	Approximating	$\mathcal{O}\,(L_{max}^6)$	$\mathcal{O}\,(L_{max}^2)$	δ, L_{max}
Subdivision	Interpolating	$\mathcal{O}\,(N\,\log N)$	$\mathcal{O}\,(\log N)$	N/A

Table 1: Summary table of numerical methods types, computational complexity of fitting and evaluation in terms of the number N of measurements, and method parameters.

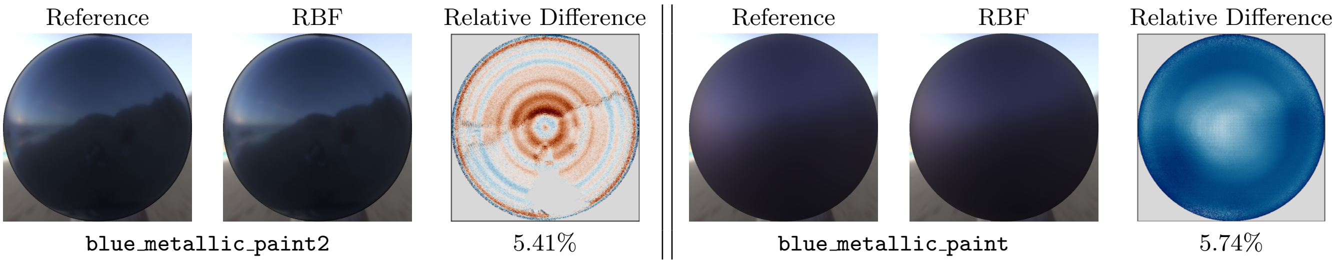

Figure 2: 3D RBF reconstruction using 1755 coefficients per color channel for two MERL Materials, with $\sigma_h = 0.30$ and $\sigma_d = 0.35$. Errors are given in terms of L2 Relative differences, averaged across the 3 color channels.

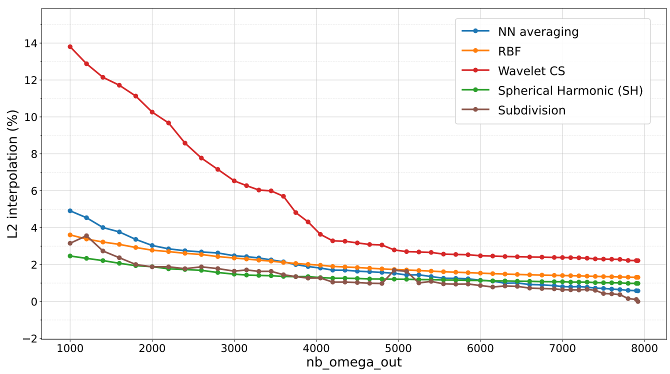

Figure 3: Convergence of interpolation methods w.r.t. data sparsity for the measured UTIA `Fabric112` material. For each incident direction, a fixed number of outgoing measurement directions are randomly subsampled from the dense UTIA directional data. The relative L2 error of the Wavelet Compressive Sensing method exceeds 10% in the low-measurement regime, as the fixed grid resolution leads to a higher proportion of empty cells when fewer measurements are available.

it remains possible to fit within a tensor space of spherical harmonics in 4D, or spherical harmonics on one side plus cosine on the 3$^{\text{rd}}$ dimension. Subdivision stands apart because it would require computing a "Delaunay triangulation" over S^2 or $[0, \pi/2] \times S$, which is very difficult. Of course, the 3D/4D versions of these methods may sort differently in terms of efficiency/accuracy. Additional questions may arise when adapting to other dimensions, such as spectral or spatially varying data.

References

[AB10] Peter Apian-Bennewitz. New scanning goniophotometer for extended BRTF measurements. In Zu-Han Gu and Leonard M. Hanssen, editors, *Reflection, Scattering, and Diffraction from Surfaces II*, volume 7792, page 77920O. International Society for Optics and Photonics, SPIE, 2010.

[BT09] Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, January 2009.

[BvdPPH11] Nicolas Bonneel, Michiel van de Panne, Sylvain Paris, and Wolfgang Heidrich. Displacement interpolation using lagrangian mass transport. *ACM Trans. Graph.*, 30(6):1–12, December 2011.

[CBP07] Luc Claustres, Loïc Barthe, and Mathias Paulin. Wavelet encoding of brdfs for real-time rendering. In *Proceedings of Graphics Interface 2007*, GI '07, page 169–176, New York, NY, USA, 2007. Association for Computing Machinery.

[FVH14] Jirí Filip, Radomír Vávra, and Michal Havlícek. Effective acquisition of dense anisotropic brdf. In *2014 22nd International Conference on Pattern Recognition*, pages 2047–2052, 2014.

[Mey93] Yves Meyer. *Wavelets and Operators*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1993.

[MP11] Domenico Marinucci and Giovanni Peccati. *Random Fields on the Sphere: Representation, Limit Theorems and Cosmological Applications*. London Mathematical Society Lecture Note Series. Cambridge University Press, 2011.

[MPBM03] Wojciech Matusik, Hanspeter Pfister, Matt Brand, and Leonard McMillan. A data-driven reflectance model. *ACM Trans. Graph.*, 22(3):759–769, July 2003.

[NDM05] Addy Ngan, Frédo Durand, and Wojciech Matusik. Experimental analysis of brdf models. In *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques*, EGSR '05, page 117–126, Goslar, DEU, 2005. Eurographics Association.

Interpolation method	M = number of measurements		
	M = 100	M = 300	M = 500
Nearest neighbor averaging	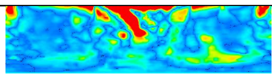 L2 = 11.37%	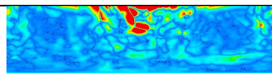 L2 = 8.10%	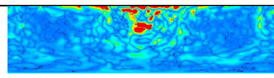 L2 = 6.26%
RBF	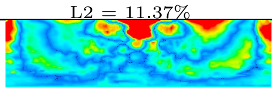 L2 = 12.62%	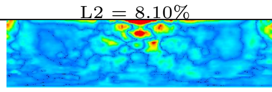 L2 = 6.42%	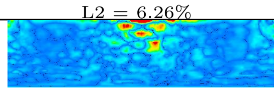 L2 = 4.67%
Wavelet Compressive Sensing	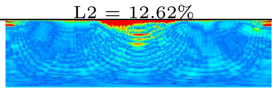 L2 = 15.21%	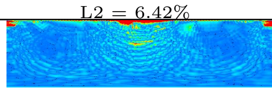 L2 = 12.04%	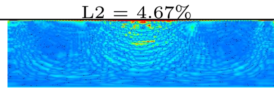 L2 = 7.14%
SH	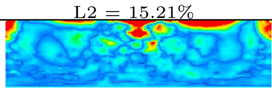 L2 = 9.07%	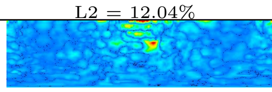 L2 = 3.62%	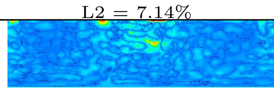 L2 = 2.98%
Subdivision	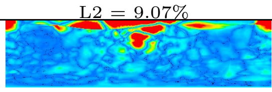 L2 = 10.15%	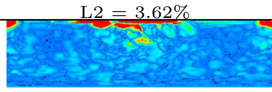 L2 = 11.58%	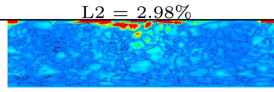 L2 = 5.67%

Table 2: Comparison of interpolation methods with respect to data sparsity for the dense `Fabric112` material subsampled with 100, 300, and 500 measurements uniformly selected on the hemisphere. Each cell shows the corresponding error map with its L2 relative metric. The relative L2 errors are computed as the root-mean-square deviation between the interpolated and ground-truth intensity values, normalized by the overall magnitude of the ground-truth signal.

Méthode		**MIT layout**	**Coupole layout**	**PAB layout**
Kernel	L2/SSIM	28.18 % / 0.026	16.32 % / 0.067	1.67 % / 1.089e-3
RBF	L2/SSIM	33.53 % / 0.028	13.61 % / 0.081	1.18 % / 6.373e-4
WCS	L2/SSIM	29.86 % / 0.035	23.26 % / 0.081	3.13 % / 6.005e-3
SH	L2/SSIM	34.66 % / 0.053	16.01 % / 0.288	1.45 % / 7.138e-4
Subdivision	L2/SSIM	25.67 % / 0.018	14.91 % / 0.059	0.77 % / 9.394e-5

Table 3: L2 relative reconstruction errors(%) and SSIM dissimilarity scores obtained with five interpolation methods on three test layouts for a fixed incident angle of (20°, 0°) using the UTIA `Fabric112` material. For both metrics, lower values indicate better reconstruction quality. L2 errors are highest on the MIT layout(25–35%) due to sparse distribution and lowest on the PAB (smaller than 3.2%). The Subdivision achieves the lowest L2 relative error, and the lowest SSIM value, on MIT and PAB layouts. On La Coupole layout, RBF yields the lowest L2 error(13.61%) whereas the Subdivision method still attains the best SSIM(0.059).

[Rus98] Szymon M. Rusinkiewicz. A new change of variables for efficient brdf representation. In George Drettakis and Nelson Max, editors, *Rendering Techniques '98*, pages 11–22, Vienna, 1998. Springer Vienna.

[SBN15] Cyril Soler, Mahdi M Bagher, and Derek Nowrouzezahrai. Efficient and accurate spherical kernel integrals using isotropic decomposition. *ACM Transactions on Graphics (TOG)*, 34(5):1–14, 2015.

[SRRW21] Alejandro Sztrajman, Gilles Rainer, Tobias Ritschel, and Tim Weyrich. Neural BRDF Representation and Importance Sampling. *Computer Graphics Forum*, 40(6):332–346, September 2021.

[SS95] Peter Schröder and Wim Sweldens. Spherical wavelets: efficiently representing functions on the sphere. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '95, page 161–172, New York, NY, USA, 1995. Association for Computing Machinery.

[VF18] R. Vávra and J. Filip. Adaptive slices for acquisition of anisotropic brdf. *Computational Visual Media*, (4):55–69, 2018.

[WAT92] Stephen H. Westin, James R. Arvo, and Kenneth E. Torrance. Predicting reflectance functions from complex surfaces. In *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '92, page 255–264, New York, NY, USA, 1992. Association for Computing Machinery.

[WMLT07] Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. Microfacet Models for Refraction through Rough Surfaces. In Jan Kautz and Sumanta Pattanaik, editors, *Rendering Techniques*. The Eurographics Association, 2007.

[WW01] Joe Warren and Henrik Weimer. *Subdivision Methods for Geometric Design: A Constructive Approach*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1st edition, 2001.

[ZREB06] T. Zickler, R. Ramamoorthi, S. Enrique, and P.N. Belhumeur. Reflectance sharing: predicting appearance from a sparse set of images of a known shape. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(8):1287–1302, 2006.