

BRDF Acquisition with Basis Illumination

Abhijeet Ghosh, Shruthi Achutha, Wolfgang Heidrich, Matthew O’Toole
The University of British Columbia*

Abstract

Realistic descriptions of surface reflectance have long been a topic of interest in both computer vision and computer graphics research. In this paper, we describe a novel and fast approach for the acquisition of bidirectional reflectance distribution functions (BRDFs). We develop a novel theory for directly measuring BRDFs in a basis representation by projecting incident light as a sequence of basis functions from a spherical zone of directions. We derive an orthonormal basis over spherical zones that is ideally suited for this task. BRDF values outside the zonal directions are extrapolated by re-projecting the zonal measurements into a spherical harmonics basis, or by fitting analytical reflection models to the data. We verify this approach with a compact optical setup that requires no moving parts and only a small number of image measurements. Using this approach, a BRDF can be measured in just a few minutes.

1. Introduction

Accurate descriptions of how light reflects off a surface have long been a topic of research in both computer vision and computer graphics. Real world materials exhibit characteristic surface reflectance, such as glossy or specular highlights, anisotropy, or retro-reflection. Descriptions of such effects find their applications for example in shape from shading algorithms and realistic rendering. The surface reflectance of a material is formalized by the notion of the Bidirectional Reflectance Distribution Function (BRDF) [25], which is a 4 dimensional function describing the response of a surface in a certain exitant direction to illumination from a certain incident direction over a hemisphere of directions.

Numerous analytical models of BRDFs exist in the literature [5, 14, 13, 2, 3] that observe the laws of energy conservation and reciprocity, and hence are physically plausible. However, these models generally do not capture the reflectance properties of all kinds of materials. Furthermore, selecting appropriate model parameters for representing different kinds of real-world materials can be a non-intuitive and time-consuming process. Therefore, acquisition of real

world BRDF data has been a very active area of research over the last few years.

The data acquired in such a process is generally not used directly due to its large size, the noise present in the measurement process, and missing data for certain incident and exitant directions. Instead, the data is usually either fitted to an analytical model [33, 17, 24] or projected into a suitable basis [35, 31, 4, 23]. This fitting process results in the loss of some of the captured high frequency details in the original data, possibly making the high sampling density of acquisition an overkill. At the same time, reducing the sampling density during acquisition would result in aliasing artifacts for sharp features that would then fall below the Nyquist limit.

In this paper, we propose an alternative approach to the acquisition of reflectance data where we *optically* project the data into a suitable basis function directly *during* the capture process. This approach results in optical low-pass filtering of the data at capture time, and thus addresses aliasing issues and minimizes high-frequency noise. An added benefit is that this prevents any redundancy in data capture as we can use all of the data we acquire. We focus on a very fast capture of object appearance for vision and graphics purposes, and do not seek to replace high-precision measurement devices such as gonio-reflectometers. Our approach speeds up acquisition time to one or two *minutes* compared to a few hours required by previous acquisition approaches. The main contributions of this paper are:

- The theory behind, and a practical implementation of the concept of measuring the response of a surface to a basis function as a way of optically filtering and encoding the BRDF data.
- Development of a set of orthogonal basis functions defined over the measurement space, as well as basis transformation as a way of data extrapolation.
- A novel design for a curved reflector catadioptric imaging setup resulting in an efficient image based BRDF acquisition without involving any moving parts.

*email:{ghosh, heidrich}@cs.ubc.ca

2. Related Work

As an alternative to analytical BRDF models, one can use measurements of BRDFs in a rendering system. Such data is available from many sources, including the Cornell [6], STARR [26], and CURET [7] databases. For example, the BRDF data in the CURET database represents 256 reflectance measurements uniformly distributed over the hemisphere of 60 different materials. However, comprehensive data is still not readily available for a large number of day-to-day materials, and thus the development of rapid new measurement methods has been a focus of research activities.

Measurement Setup. The wide availability and decreasing cost of digital cameras has led researchers to explore various image based BRDF acquisition approaches. One way of reducing the number of images that need to be taken is by using curved surfaces for recovering homogeneous [21, 22], or spatially varying BRDFs [19]. Generally, these methods require knowledge of the geometric shape, and are not well-suited for capturing fabric or sheet materials. Such materials can be measured by wrapping them around a cylinder at various orientations [24].

In many cases, planar samples are, however, more convenient. Other researchers have therefore focused on special optics to cover a large range of incident or exitant light directions for a planar sample in a single photograph. Ward’s reflective dome design [33] was the first to use this approach. Malzbender et al. [20] use a dome with attached, individually controlled light sources to photograph a surface under varying lighting conditions. Han and Perlin [12] developed a device to capture bidirectional texture functions (BTFs) based on a kaleidoscope. Dana [8] designed an acquisition device using a parabolic mirror that densely covers a relatively small solid angle. The system also involves planer translations of the light source to cover various incident directions and translations of the sample in order to scan the surface for spatial variations in reflectance. Finally, Kuthirummal and Nayar [16] have developed a class of radial imaging systems for image-based acquisition of geometry, texture, and BRDFs. Their BRDF measurement setup can image 4 radial lines of reflectance of a given material for a fixed light source direction.

Our work is most closely related to the last two papers. Like Kuthirummal and Nayar [16], we use a rotationally symmetric optical design, although ours is not a cylinder or cone, but a freeform surface. Our design lets us measure BRDFs over a continuous region of directions, much like the work of Dana [8]. However, unlike any other previous work, we use basis function illumination, rather than individual point or directional lights, which allows us to very rapidly acquire BRDFs, including anisotropic ones.

Measurements with Basis Functions. Basis function approaches have been used in the past for measuring

other visual effects, including light fields [11], reflectance fields [34] and environment mattes [27]. While similar in spirit, these approaches measure different physical properties, and are therefore both mathematically and optically very different from ours. Our proposed zonal basis illumination is very similar in principle to *harmonic lights* as proposed by Sato et al. [32] for encoding the appearance function of a scene. Hence our optical setup can be seen as a step towards the realization of such harmonic lights.

BRDF Representations. After the BRDF has been measured, one can use it directly in tabular form in a rendering system [21]. However, the data often contains holes and is noisy, so that some form of post-processing and interpolation is almost always required. Furthermore, the inherent dimensionality of the BRDF data, and the need to sample it at a high resolution leads to unwieldy storage problems. Most researchers therefore represent their BRDFs as either analytical reflectance models [33, 19, 10], or generic function bases such as polynomials [15, 20], spherical harmonics [35, 31, 4], or wavelets [18, 23].

The key difference of our work from these approaches is that we *directly* acquire the BRDF in a basis representation, rather than measuring a tabulated representation that is later fit with the basis functions. The concept of basis function acquisition, in combination with the optical setup that we devised, allows for extremely rapid and easy measurement of BRDFs.

3. Overview

The distinguishing characteristic of our BRDF measurement system is that it captures the response of the surface to illumination in the form of smooth basis functions, while existing methods measure impulse response using thin pencils of light that approximate Dirac peaks. For this concept to be practical, we require an optical setup that allows us to simultaneously project light onto the sample from a large range of directions, and likewise to measure the reflected light distribution over a similarly large range of directions. Developing such optics also has the advantage that no moving parts are required, which is one reason for the speed of our acquisition.

In this paper, we choose a spherical zone of directions as the acquisition region for both incident and exitant light directions. Spherical zones have several advantages over regions of other shape. First, they allow us to develop basis functions that align nicely with the symmetries present in many BRDFs, thus minimizing the number of basis functions required to represent a given BRDF. Alignment also simplifies extrapolation of data into missing regions. Second, a zonal setup allows us to design optics that could, in principle, cover over 98% of the hemisphere, with only a small hole near the zenith, where BRDF values are usually smoother compared to more tangential directions. The man-




Figure 1. Left: Physical setup of our reflectance acquisition device. A camera focused on the mirrored components views a zone of reflected directions. A projector illuminates the corresponding zone of incident directions using a beam splitter. Right: A prototype demonstrating the concept in 2-D. Here, we focus illumination on the mirrored components using a laser pointer and observe that the beam bounces back to its origin.

ufacturing process that we used for our prototype system allowed us to produce a section of that range corresponding to 51% of the hemisphere.

Figure 1 shows a diagram and a 2D mockup of such an optical setup. A camera focused on the mirrored components can capture the full zone of reflected directions in our setup. Simultaneously, a projector focused on the the mirrored components can cover the corresponding zone of incident directions.

In the following, we will first discuss the theoretical underpinnings for basis function BRDF acquisition (Section 4), and then describe the physical setup (Section 5). Finally, we present results in Section 6 and conclude with a discussion in Section 7.

4. Measurement with Basis Functions

In this section, we discuss the mathematical concepts behind a basis function approach for BRDF measurement, and derive the specific basis that we use in our work. Section 5 then deals with the physical realization of these concepts.

Assume that we want to measure a BRDF $f_r(\omega_i, \omega_o)$ for combinations of incident light direction ω_i and exitant light direction ω_o restricted to a spherical zone \mathbf{Z} centered around the surface normal. \mathbf{Z} corresponds to longitudinal angles $\phi \in [0 \dots 2\pi]$ and latitudinal angles $\theta \in [\theta_{min} \dots \theta_{max}]$, as shown in Figure 2.

We would like to approximate the BRDF over this zone with a linear combination of basis functions $\{Z_k(\omega_i)\}$ over the incident light directions. We will include the $\cos \theta_i$ term in this basis representation for convenience and numerical stability, i.e.

$$\hat{f}_r(\omega_i, \omega_o) = f_r(\omega_i, \omega_o) \cos \theta_i \approx \sum_k Z_k(\omega_i) z_k(\omega_o), \quad (1)$$




Figure 2. The measurement zone \mathbf{Z} .

so that we can write the reflected radiance for any outgoing direction ω_o as

$$L_o(\omega_o) = \int_{\mathbf{Z}} f_r(\omega_i, \omega_o) L_i(\omega_i) \cos \theta_i d\omega_i \quad (2)$$

$$\approx \sum_k z_k(\omega_o) \int_{\mathbf{Z}} Z_k(\omega_i) L_i(\omega_i) d\omega_i. \quad (3)$$

In this framework, BRDF measurement can be seen as the process of determining the coefficients $z_k(\omega_o)$ for each basis Z_k and each exitant light direction ω_o . If we have chosen the Z_k such that they form an orthonormal basis over the zone \mathbf{Z} , then the coefficients are given as

$$z_k(\omega_o) = \int_{\mathbf{Z}} Z_k(\omega_i) f_r(\omega_i, \omega_o) \cos \theta_i d\omega_i. \quad (4)$$

In other words, we can measure $z_k(\omega_o)$ by recording the reflected light along each direction $\omega_o \in \mathbf{Z}$ for different incident illumination patterns $Z_k(\omega_i)$. In practice, we separately project the positive Z_k^+ and the negative Z_k^- parts of the basis function Z_k , and subtract the resulting coefficients in software, similar to the work by Goesele et al. [11].

There are several ways in which one can define a suitable orthonormal basis over \mathbf{Z} . In Appendix A we derive the set of orthonormal Zonal Basis (ZB) functions $Z_l^m(\phi, \theta) \in [0, 2\pi] \times [\theta_{min}, \theta_{max}]$ that we use for our purposes:

$$Z_l^m(\theta, \phi) = \begin{cases} \sqrt{2} \hat{K}_l^m \cos(m\phi) \hat{P}_l^m(\cos \theta) & \text{if } m > 0 \\ \sqrt{2} \hat{K}_l^m \sin(-m\phi) \hat{P}_l^{-m}(\cos \theta) & \text{if } m < 0 \\ \hat{K}_l^0 \hat{P}_l^0(\cos \theta) & \text{if } m = 0 \end{cases}, \quad (5)$$

where the zonal normalization constant \hat{K}_l^m is

$$\hat{K}_l^m = \sqrt{\frac{(2l+1)(l-|m|)!}{2\pi \cdot (\cos \theta_{min} - \cos \theta_{max}) \cdot (l+|m|)!}}. \quad (6)$$

For practical applications, we of course need to extrapolate from the data measured over the zone to incident and exitant directions that have not been measured. This task is simplified by the global support of our basis functions, and would be much more difficult for a basis with local support, such as a Wavelet basis. In general, we would also like to transform the data into a different representation that

is more convenient for rendering purposes, such as a tensor-product Spherical Harmonics (SH) basis, or coefficients of an analytical reflection model. Interestingly, format conversion and extrapolation can be achieved in a single, inexpensive step, as described in the following.

4.1. Basis Conversion to Spherical Harmonics

One way of extrapolating the acquired zonal data in the zone of missing measurements is by transformation into an alternative basis such as spherical harmonics. The SH basis has been used extensively in the past for representing BRDF data. Unlike the ZB basis functions, however, the restrictions of the SH basis functions Y_l^m to our measurement zone \mathbf{Z} are *not* orthonormal, and therefore, the equivalent of Equation 4 does *not* hold for spherical harmonics. Instead, we have

$$y_l^m(\omega_o) = \int_{\mathbf{Z}} \hat{Y}_l^m(\omega_i) \hat{f}_r(\omega_i, \omega_o) d\omega_i, \quad (7)$$

where $\{\hat{Y}_l^m(\omega_i)\}$ is the *dual basis* to the spherical harmonics *over the zone \mathbf{Z}* , i.e. the basis that fulfills the conditions

$$\int_{\mathbf{Z}} Y_l^m(\omega) \hat{Y}_p^q(\omega) d\omega = \begin{cases} 1 & \text{if } l = p \text{ and } m = q \\ 0 & \text{otherwise} \end{cases}. \quad (8)$$

Since $\{\hat{Y}_l^m(\omega_i)\}$ is a basis for the same function space as the SH basis, we also have

$$\hat{Y}_p^q = \sum_{l,m} c_{l,p}^{m,q} Y_l^m. \quad (9)$$

Equations 8 and 9 together describe a sparse linear system that can be solved to obtain the linear weights that define the duals \hat{Y} . Conversion from ZB to SH is then a simple linear transformation of the zonal coefficients z_p^q of a function \hat{f}_r by a sparse basis change matrix C into corresponding SH coefficients y_l^m . Each element of this matrix is defined by

$$C_{l,p}^{m,q} = \int_{\mathbf{Z}} z_p^q \hat{Y}_l^m d\omega. \quad (10)$$

Note that we could use the dual functions \hat{Y}_l^m for measurements and directly project the BRDF into a spherical harmonic basis. However, having a single orthonormal basis is more convenient for projection into arbitrary function spaces, including analytical BRDF models, as described next.

4.2. Fitting Analytical Reflection Models

For relatively low frequency BRDFs, the spherical harmonic representation produces very good results. For specular materials, it is well known that basis functions such

as spherical harmonics or our zonal basis suffer from oscillations in the proximity of discontinuities or strong gradients. These oscillations are visible in the reconstruction as undesirable ringing artifacts also known as the Gibbs phenomenon (Figure 3, center). Hence, for specular materials, we cannot directly use the acquired coefficients or transform them into SH for final use.

Instead, we propose to fit the higher order zonal representation of specular BRDFs to an analytical model, thereby computing a least-squares fit over the spurious oscillations. Since the Gibbs phenomenon is an oscillation around the true function value, such a least-squares fit produces a very good reconstruction (Figure 3, right). In our experiments, we worked with the distribution based BRDF model by Ashikhmin [1] due to the simplicity of the fitting procedure. The D-BRDF model is also a generalization of the Ashikhmin-Shirley-Phong model [3], which was recently found to be particularly well-suited for fitting to measured data [24]. However, the measured zonal data can be fitted to any other suitable analytic model using a numerical procedure such as Levenberg-Marquardt [29].




Figure 3. An illustration of the suppression of ringing through fitting analytical BRDF models. Left: original acrylic blue paint BRDF acquired by Matusik et al. [22]. Center: 10th order zonal reconstruction, rendered after transformation into SH, exhibiting severe ringing artifacts. Right: Corresponding D-BRDF fit to the zonal reconstruction.

5. Measurement Setup and Calibration




Figure 4. Photograph of the proposed BRDF acquisition setup including a camera, a projector, a beam-splitter, and two curved reflectors mounted on a 40 cm × 40 cm optical bench.




Figure 5. Iterative process for designing the profile of the reflective dome for a fixed convex parabolic reflector.

The primary components of our image-based acquisition setup are a convex parabolic mirror suspended inside a mirrored dome. This optical setup can cover a zone of incident as well as exitant directions of measurement. In addition to the mirrored components, the acquisition system consists of a FireWire machine vision camera (Prosilica EC 1350C), an LED RGB PocketProjector (Mitsubishi PK1), and a beam splitter. The camera has a resolution of 1360×1024 and an acquisition rate of 15 frames per second at 12-bits per color channel. The projector has a resolution of 800×600 with peak illumination intensity specified at 200 Lux. An external 350 mm lens was used to focus the projector at the required focal distance. All reflectance measurements are performed with multiple exposures [9] for high dynamic range (HDR) acquisition.

Our optical setup consists of two mirrored components, a convex parabola and a concave reflective dome as shown in Figures 1 and 5. The dome has a rotationally symmetric shape with a freeform profile, as detailed in the following.

Dome Shape: For a fixed configuration of parabola, sample, camera, and projector, the freeform profile of the dome is determined as follows. First, the location of the dome's rim D_1 is found by intersecting a camera ray reflecting off the bottom edge P_1 of the paraboloid with the tangent plane of the sample (Figure 5, left). The surface normal at the rim defines a tangent plane in D_1 . For the next camera ray reflecting off P_2 , we compute the intersection D_2 of the reflected ray with the tangent plane of D_1 (Figure 5, center). The normal in D_2 defines a new tangent plane that we can use in the same way to obtain the next point on the dome. Proceeding iteratively with this approach, we can determine the full shape of the dome (Figure 5, right) in what amounts to an Euler integration procedure. Note that these simulations are run at orders of magnitude higher resolution than actual camera or projector pixel resolution.

Design Simulations: The design parameters, i.e. the spatial location of parabola, sample, camera, and projector, were optimized using detailed simulations with a ray-tracer. We modeled the camera and projector as thin lens devices.

Our simulations took into account various parameters such as focal distances, finite apertures and pixel resolutions of cameras and projectors, and stability under minor misalignments of the various optical components to the optical axis.

Final Design: After extensive simulations, we decided on a design that lets us project over 100 pixels between the vertex and the tangent of the parabolic mirror in order to provide at least 1 measurement per degree along the latitudinal directions. For this setup, the distance between the center of projection of the camera and the vertex of the parabolic mirror is 27 cm, and the distance between the parabola vertex and the sample at the bottom is 13.5 cm. The dimensions of the full dome are $11'' \times 11'' \times 10''$ for this setup. Our design provides us > 1 pixel/degree measurements over the full measurement zone. The full dome as simulated in Figure 5 would cover the zone from 9° to 90° off the normal to the sample. This range corresponds to over 98% of the full hemisphere.

Physical Implementation: For the manufacturing of the dome and parabola, we chose electroforming process, in which a mandrel of the dome is first machined and polished, and then the actual dome is deposited on this mandrel in an electrolyte bath. This process allows the production of optical quality free-form surfaces at moderate cost. However, a downside of this approach is that it only allows for convex holes, since the mandrel has to be removed after the electroforming process. For this reason, we were only able to build a dome covering the zone from 9° to 57° off normal, corresponding to about 51% of the hemisphere (Figure 4).

5.1. Calibration

Geometric calibration is necessary in order to align the camera and the projector to the optical axis of the acquisition setup. We also need to perform photometric calibration in order to recover the absolute scaling factors for our measurements with respect to some known reflectance standard.

Optical Axis Calibration: The optical axis of the camera and projector need to be aligned with that of the parabolic mirror and dome. We mount the dome on an optical table, and mark its optical axis with crosses that are attached to the dome with precision mounts. The camera is moved with a manual translation stage until all crosses line up. Likewise, the projector is moved manually until the shadows of all crosses line up.

Sample Mounting: Due to the large aperture of our optical system, the depth-of-field is very shallow, about 2 mm. As a result, the material samples have to be mounted with fairly high precision, which is easily achieved with a mechanical stop.

Projector Flat-Fielding: We account for any spatial variation of the projector illumination by acquiring an HDR photograph of a full screen image set to medium gray, projected on to a diffuse white screen at the required focal distance of



Figure 6. The Audi-TT model rendered with acquired BRDFs of 2 different paint samples. The BRDFs were acquired using 25 4^{th} order basis functions as defined in this paper, and then rendered with a basis transformation into spherical harmonics. Left: Metallic teal automotive paint. Right: KrylonTM true blue paint. In each case, the time taken for the entire BRDF measurement process including data capture and re-projection into the spherical harmonic basis was about one minute.

28cm. All the basis images are then modulated by this image.

Reflectance Calibration: An important aspect of the calibration is to recover the relative scaling factors for our measurements with respect to some known reflectance standard. For this, we take advantage of an 18 % diffuse gray card commonly used in photography. We measure the diffuse reflectance of the gray card with our setup using low order zonal basis functions [30]. The relative scaling factors for each color channel are obtained by white-balancing the results of the gray card measurements.

6. Results

Using our prototype setup, we have acquired the BRDFs of various types of materials, including velvet, anisotropic synthetic, silk and satin fabrics, leather, various kinds of glossy and shiny papers, paint and plastic samples, printer toners, wax, highly specular metal foil wrapping papers, and anisotropic samples such as a guitar pick and a copper coin. Figure 7 presents a selection of BRDFs as rendered on a sphere under a directional light source. Most of the materials were acquired using lower order ($l \leq 6$) zonal basis functions. The silk and satin fabrics, and the guitar pick were acquired with order $l = 8$ zonal basis function, while the shiny wrapping papers and anisotropic copper coin required acquisition with order $l = 10$ zonal basis function. The total number of images acquired for an order l acquisition is $(l + 1)^2 \times 2 \times 3$, with 2 separate positive and negative parts, and 3 exposures for HDR imaging. The entire acquisition process takes just a few minutes to complete even for higher order zonal basis functions.

Figure 6 presents the BRDFs of 2 different paint samples that we acquired using 4^{th} order zonal basis functions, rendered on the Audi-TT car model, and illuminated by an HDR environment map using the Physically Based Ray Tracing system [28].



Figure 8. Specular chocolate wrapping papers acquired using higher order zonal basis functions, and then fit to an analytical model for rendering. Left: Red KitKatTM wrapping paper. Right: Copper colored LindtTM chocolate wrapping paper.




Figure 9. Visual comparison of two kinds of acquired satin samples wrapped around a cylinder as lit by a point source against real photographs. Left column: Photographs of the red and blue satin samples. Right column: Rendering of the D-BRDF fits.

A representative set of the BRDFs acquired using lower order ($l \leq 6$) zonal basis functions is shown in Figure 10. For this class of materials, the entire process of acquisition followed by a basis transformation into the SH basis took under three minutes.

Figure 8 demonstrates the specular materials, in this case shiny metal foil chocolate wrapping papers, which we then fit to the D-BRDF analytical model. The D-BRDF fitting procedure consists of constructing the distribution of the half-vector ω_h between the incident light direction ω_i and exitant viewing direction ω_o as a function of the back-scattering direction measurements, i.e, the directions where $\omega_i = \omega_o$. In our case, we extract the zonal half-vector distribution from the measured data, and then extrapolate that to cover the full hemisphere of half-vector directions. The entire acquisition and fitting procedure took only a few minutes to complete in all examples. Similarly, we also fit D-BRDFs to the anisotropic guitar pick, the copper coin and




Figure 7. Various BRDFs acquired with our prototype setup using zonal basis functions. Top row: from left to right - bright orange paper, red velvet, maroon synthetic fabric, brown leather, coated brown envelope, red printer toner, blue rubber band, glossy red paper, glossy blue-gray paper, LindtTM chocolate box paper. Center row: from left to right - magenta plastic with grain finish, retro-reflective plastic, dark brown plastic coffee lid, KrylonTM banner red paint, KrylonTM true blue paint, metallic teal automotive paint, chrome gold dust automotive paint, purple anisotropic silk fabric, blue anisotropic silk fabric. Bottom row: from left to right - glossy succulent plant leaf, red wax, shiny blue paper, shiny golden paper, red KitKatTM wrapping paper, copper colored LindtTM chocolate wrapping paper, anisotropic plastic guitar pick, anisotropic copper coin, anisotropic red satin.

the satin samples (Figure 7, bottom row).

Finally, as a way of validating our measurement and fitting approach, we photographed two satin samples wrapped around a cylinder in a dark room and lit by a collimated point light source. Figure 9 presents the comparisons of these photographs with the corresponding renderings of the D-BRDF fits to these samples. The highlights in the rendered images are a close match to the real photographs validating our approach.

As an additional step towards quantitative validation, we used an 18% gray card as a diffuse reflectance standard, and compared the recovered coefficients to the ones expected for a diffuse target. For low order basis functions, the error is within a few percent of the expected value.




Figure 10. Representative set of BRDFs acquired with lower order zonal basis functions rendered under directional lighting. From left to right: red velvet, red printer toner, magenta plastic sheet, chrome gold dust automotive paint.

7. Conclusions

In this paper, we have presented a novel basis function approach to BRDF measurement. Our contributions include a novel theory for basis function BRDF acquisition, the de-

velopment of an orthonormal basis for spherical zones, and the design of an optical setup that allows for basis function illumination of BRDF samples.

The dome we use in our prototype setup covers a sufficient percentage of the hemisphere to obtain high quality BRDF measurements with our basis function approach. To further increase quality by reducing the amount of extrapolation, a dome with a larger coverage could be used. We are currently looking into manufacturing techniques that are able to produce such domes.

Due to the basis function approach and the dispensing of all moving parts, BRDF measurement with our setup is very fast, reducing the acquisition time to a few minutes even for high-frequency materials. Moreover, the physical dimensions of the setup are quite compact, so that the whole apparatus could be enclosed in a small box for mobile on-site acquisitions.

References

- [1] M. Ashikhmin. Distribution-based BRDFs, 2006. <http://jesper.kalliope.org/blog/library/dbrdfs.pdf>. 4
- [2] M. Ashikhmin, S. Premoše, and P. Shirley. A microfacet-based BRDF generator. In *SIGGRAPH '00: Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 65–74, 2000. 1
- [3] M. Ashikhmin and P. Shirley. An anisotropic phong BRDF model. *J. Graph. Tools*, 5(2):25–32, 2000. 1, 4
- [4] R. Basri and D. W. Jacobs. Lambertian reflectance and linear subspaces. *IEEE PAMI*, 25(2):218–233, Feb. 2003. 1, 2
- [5] R. L. Cook and K. E. Torrance. A reflectance model for computer graphics. *ACM Transactions on Graphics*, 1(1):7–24, 1982. 1
- [6] Cornell. CORNELL light measurement laboratory, 2005. <http://www.graphics.cornell.edu/research/measure/>. 2
- [7] CURET. CURET: Columbia-Utrecht Reflectance and Texture, 1999. <http://www.cs.columbia.edu/CAVE/curet/>. 2

[8] K. Dana. BRDF/BTF measurement device. In *Proc. of ICCV'01*, pages 460–466, 2001. 2

[9] P. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *Proc. of ACM SIGGRAPH '97*, pages 369–378, 1997. 5

[10] A. Gardner, C. Tchou, T. Hawkins, and P. Debevec. Linear light source reflectometry. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 22(3):749–758, 2003. 2

[11] M. Goesele, X. Granier, W. Heidrich, and H.-P. Seidel. Accurate light source acquisition and rendering. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 22(3):621–630, July 2003. 2, 3

[12] J. Y. Han and K. Perlin. Measuring bidirectional texture reflectance with a kaleidoscope. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 22(3):741–748, 2003. 2

[13] X. D. He, P. O. Heynen, R. L. Phillips, K. E. Torrance, D. H. Salesin, and D. P. Greenberg. A fast and accurate light reflection model. In *SIGGRAPH '92: Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, pages 253–254, 1992. 1

[14] X. D. He, K. E. Torrance, F. X. Sillion, and D. P. Greenberg. A comprehensive physical model for light reflection. In *SIGGRAPH '91: Proceedings of the 18th annual conference on Computer graphics and interactive techniques*, pages 175–186, 1991. 1

[15] J. Koenderink, A. van Doorn, and M. Stavridi. Bidirectional reflection distribution function expressed in terms of surface scattering modes. In *ECCV '96. 4th European Conference on Computer Vision*, volume 2, pages 28–39, 1996. 2

[16] S. Kuthirummal and S. K. Nayar. Multiview Radial Catadioptric Imaging for Scene Capture. *ACM Trans. on Graphics (also Proc. of ACM SIGGRAPH)*, 25(3):916–923, Jul 2006. 2

[17] E. Lafortune, S.-C. Foo, K. Torrance, and D. Greenberg. Non-linear approximation of reflectance functions. In *Proc. of ACM SIGGRAPH '97*, pages 117–126, Aug. 1997. 1

[18] P. Lalonde and A. Fournier. A wavelet representation of reflectance functions. *IEEE Transactions on Visualization and Computer Graphics*, 3(4):329–336, 1997. 2

[19] H. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel. Image-based reconstruction of spatially varying materials. In *Eurographics Workshop on Rendering*, pages 104–115, 2001. 2

[20] T. Malzbender, D. Gelb, and H. Wolters. Polynomial texture maps. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 519–528, 2001. 2

[21] S. Marschner, S. Westin, E. Lafortune, and K. Torrance. Image-based measurement of the bidirectional reflection distribution function. *Applied Optics*, 39(16):2592–2600, 2000. 2

[22] W. Matusik, H. Pfister, M. Brand, and L. McMillan. A data-driven reflectance model. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 22(3):759–769, 2003. 2, 4

[23] R. Ng, R. Ramamoorthi, and P. Hanrahan. All-frequency shadows using non-linear wavelet lighting approximation. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 22(3):376–381, July 2003. 1, 2

[24] A. Ngan, F. Durand, and W. Matusik. Experimental analysis of BRDF models. In *Proceedings of the Eurographics Symposium on Rendering*, pages 117–226, 2005. 1, 2, 4

[25] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. Geometric considerations and nomenclature for reflectance. *NBS Monograph*, 160, 1977. 1

[26] NIST. NIST reference reflectometer: STARR facility, 2003. <http://www.physics.nist.gov/Divisions/Div844/facilities/brdf/starr.html>. 2

[27] P. Peers and P. Dutré. Inferring reflectance functions from wavelet noise. In *Proc. Eurographics Symposium on Rendering*, pages 173–181, 2005. 2

[28] M. Pharr and G. Humphreys. *Physically Based Redering*. Morgan Kaufmann, New York, 2004. 6

[29] W. Press, B. Flannery, S. Teukolsky, and W. Vetterling. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, 1992. 4

[30] R. Ramamoorthi and P. Hanrahan. An efficient representation for irradiance environment maps. In *Proc. of ACM SIGGRAPH '01*, pages 497–500, 2001. 6

[31] R. Ramamoorthi and P. Hanrahan. Frequency space environment map rendering. In *Proc. of ACM SIGGRAPH '02*, pages 517–526, 2002. 1, 2

[32] I. Sato, T. Okabe, Y. Sato, and K. Ikeuchi. Appearance sampling for obtaining a set of basis images for variable illumination. In *Proc. of ICCV'03*, pages 800–807, 2003. 2

[33] G. J. Ward. Measuring and modeling anisotropic reflection. In *SIGGRAPH '92: Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, pages 265–272, 1992. 1, 2

[34] A. Wenger, A. Gardner, C. Tchou, J. Unger, T. Hawkins, and P. Debevec. Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 24(3):756–764, July 2005. 2

[35] S. Westin, J. Arvo, and K. Torrance. Predicting reflectance functions from complex surfaces. In *Computer Graphics 26, Annual Conference Series*, pages 255–264, 1992. 1, 2

A. Orthonormal Zonal Basis

Our Zonal Basis (ZB), like the Spherical Harmonic basis, is derived from the Associated Legendre Polynomials (ALP) $P_l^m(x)$, $m \in \{0, \dots, l\}$, which are orthogonal over $x \in [-1, 1]$ with

$$\int_{-1}^1 P_l^m(x) P_l^m(x) dx = \frac{2(l+|m|)!}{(2l+1)(l-|m|)!} \delta_{ll'}. \quad (11)$$

For defining spherical harmonics Y_l^m , the P_l^m are scaled so that they are orthogonal over $[0, \pi]$, with

$$Y_l^m(\theta, \phi) = \begin{cases} \sqrt{2} K_l^m \cos(m\phi) P_l^m(\cos \theta) & \text{if } m > 0 \\ \sqrt{2} K_l^m \sin(-m\phi) P_l^{-m}(\cos \theta) & \text{if } m < 0 \\ K_l^0 P_l^0(\cos \theta) & \text{if } m = 0 \end{cases}, \quad (12)$$

where K_l^m is the SH normalization constant:

$$K_l^m = \sqrt{\frac{(2l+1)(l-|m|)!}{4\pi(l+|m|)!}}. \quad (13)$$

For our zonal basis, we follow the same principle, and rescale the ALP to the range $[\theta_{\min} \dots \theta_{\max}]$.

$$\hat{P}_l^m(x) = P_l^m(n_1 \cdot x - n_2), \quad (14)$$

with

$$n_1 = \frac{2}{\cos \theta_{\min} - \cos \theta_{\max}}, \quad n_2 = \frac{2 \cos \theta_{\min}}{\cos \theta_{\min} - \cos \theta_{\max}} - 1.$$

The ZB functions $Z_l^m(\phi, \theta) \in [0, 2\pi] \times [\theta_{\min}, \theta_{\max}]$ are then given by Equations 5 and 6 in Section 4.