Int J Comput Vis (2010) 90: 183–197
DOI 10.1007/s11263-008-0151-7

A Basis Illumination Approach to BRDF Measurement

Abhijeet Ghosh · Wolfgang Heidrich · Shruthi Achutha · Matthew O'Toole

Received: 6 February 2008 / Accepted: 19 June 2008 / Published online: 19 July 2008

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 high speed approach for the acquisition of bidirectional reflectance distribution functions (BRDFs). We develop a new 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. For specular materials, we experiment with alternative basis acquisition approaches such as compressive sensing with a random subset of the higher order orthonormal zonal basis functions, as well as measuring the response to basis defined by an analytical model as a way of optically fitting the BRDF to such a representation. 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.

A. Ghosh (✉) · W. Heidrich · S. Achutha · M. O'Toole
The University of British Columbia, Vancouver, Canada V6T 1Z4
e-mail: ghosh@cs.ubc.ca

W. Heidrich
e-mail: heidrich@cs.ubc.ca

Present address:
A. Ghosh
USC Institute for Creative Technologies, Marina Del Rey, CA 90292, USA

Keywords Reflectance · Computational illumination · Object scanning and acquisition · Optics · Compressive sensing

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 (discounting any sub-surface scattering) is formalized by the notion of the Bidirectional Reflectance Distribution Function (BRDF) (Nicodemus et al. 1977), 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 (Cook and Torrance 1982; He et al. 1991, 1992; Ashikhmin et al. 2000; Ashikhmin and Shirley 2000) 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. This task has typically involved measuring the response of various samples using some version of a gonioreflectometer (Cornell 2005; CUReT 1999; NIST 2003; Ward 1992). More recently, several researchers have employed image based techniques in order to make acquisition

of BRDFs more efficient (Marschner et al. 2000; Lensch et al. 2001; Dana 2001; Matusik et al. 2003; Ngan et al. 2005; Mukaigawa et al. 2007).

The data acquired with 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 (Ward 1992; Lafortune et al. 1997; Ngan et al. 2005) or projected into a suitable basis (Westin et al. 1992; Ramamoorthi and Hanrahan 2002; Basri and Jacobs 2003; Ng et al. 2003). 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 extend 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 (Ghosh et al. 2007). 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.
- Introduction of a compresive sensing approach to BRDF measurement for efficient acquisition of specular materials.
- Constructing data-dependent illumination basis based on an analytical BRDF model and optically projecting BRDFs into such a representation.
- A novel design for a curved reflector catadioptric imaging setup resulting in an efficient image based BRDF acquisition without involving any moving parts.

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 (Cornell 2005), STARR (NIST 2003), and CUReT (CUReT 1999) 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 (Marschner et al. 2000; Matusik et al. 2003), or spatially varying BRDFs (Lensch et al. 2001). 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 (Ngan et al. 2005).

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 (Ward 1992) was the first to use this approach. Malzbender et al. (2001) use a dome with attached, individually controlled light sources to photograph a surface under varying lighting conditions. Han and Perlin (2003) developed a device to capture bidirectional texture functions (BTFs) based on a kaleidoscope. Dana (2001) 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. Kuthirummal and Nayar (2006) 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. Recently, Mukaigawa et al. (2007) have presented a catadioptric imaging system based on an ellipsoidal mirror and no moving parts for high speed BRDF measurement. They use a projector as the illumination source to project circular blobs of light to discretely sample the incident illumination.

Our work is most closely related to the last three papers. Like Kuthirummal and Nayar (2006) and Mukaigawa et al. (2007), we use a rotationally symmetric optical design, although ours is not a cylinder or ellipsoid, but a freeform surface. Our design lets us measure BRDFs over a continuous region of directions, much like the work of Dana (2001) and Mukaigawa et al. (2007). 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. Also, point-sampling of the lighting directions with a projector like in the work of Mukaigawa suffers from reduced available dynamic range, and hence reduced signal-to-noise ratio (SNR) for the measurements. This SNR problem is avoided for measurements with basis illumination.

Measurements with Basis Functions. Basis function approaches have been used in the past for measuring other visual effects, including light fields (Goesele et al. 2003), reflectance fields (Wenger et al. 2005) and environment mattes (Peers and Dutré 2005). While similar is 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. (2003) 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 (Marschner et al. 2000). 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 (Ward 1992; Lensch et al. 2001; Gardner et al. 2003), or generic function bases such as polynomials (Koenderink et al. 1996; Malzbender et al. 2001), spherical harmonics (Westin et al. 1992; Ramamoorthi and Hanrahan 2002; Basri and Jacobs 2003), Zernike polynomials (Koenderink et al. 1996), hemispherical basis (Gautron et al. 2004) or wavelets (Schröder and Sweldens 1995; Lalonde and Fournier 1997; Ng et al. 2003).

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 approach 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 manufacturing process that we used for our prototype system allowed us to produce a section of that range corresponding to 51% of the hemisphere. Even then, the optical setup achieves greater coverage than a lot of previous measurement systems that have been successfully used for BRDF measurement. For example, Ngan et al. (2005) sampled only 20% of the hemisphere for their measurements of anisotropic BRDFs, and Dana (2001) measured a much smaller zone with a parabolic reflector than our optical setup. Interestingly, the optical setup of Mukaigawa et al. (2007) also covers a similar (albeit larger) spherical zone of directions. Our basis acquisition approach can also be readily applied with their physical setup. Hence, hemispherical coverage is not a fundamental limitation of the basis measurement approach.

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 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 (Sect. 4), and then describe the physical setup (Sect. 6). Finally, we present results in Sect. 7 and conclude with a discussion in Sect. 8.

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 6 then deals with the physical realization of these concepts.

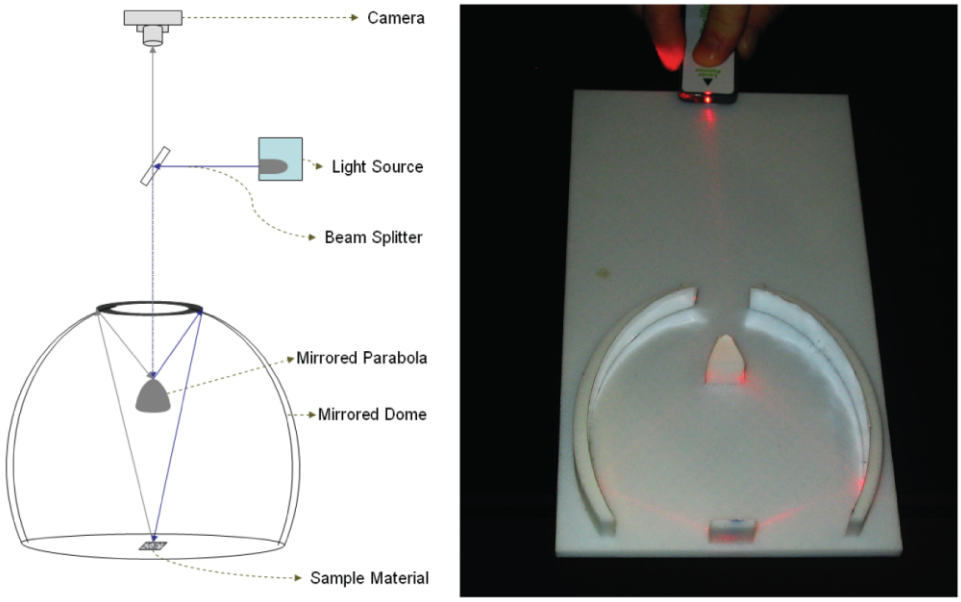

Fig. 1 *Left*: Physical setup of our reflectance acquisition device. A camera focused on the mirrored components views a zone of reflected directions. A projector illuminates the corresponding zone of incident directions using a beam splitter. *Right*: A prototype demonstrating the concept in 2D. Here,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

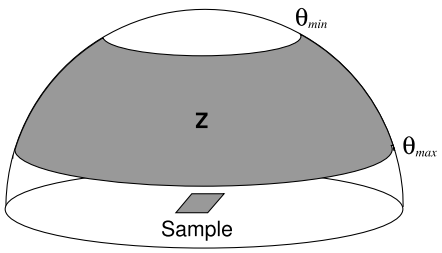

Fig. 2 The measurement zone **Z**

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 **Z** centered around the surface normal. **Z** corresponds to longitudinal angles $\phi \in [0 \dots 2\pi]$ and latitudinal angles $\theta \in [\theta_{min} \dots \theta_{max}]$, as shown in Fig. 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.

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

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

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

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

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* **Z**, then the coefficients are given as

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

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. (2003).

There are several ways in which one can define a suitable orthonormal basis over **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:

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

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

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

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 (4) does *not* hold for spherical harmonics. Instead, we have

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

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

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

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

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

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

$$C_{l,p}^{m,q} = \int_{\mathbf{Z}} Z_p^q \hat{Y}_l^m d\omega. \tag{10}$$

Interestingly, the construction of orthonormal basis functions restricted to a zone of a sphere from a linear combination of the SH basis (similar to our construction of the dual functions $\hat{Y}_l^m$) has also been investigated in the astrophysics literature by Gorski (1994) for studying cosmic microwave background temperature distribution. However, we go beyond this in our work and provide an anlytical formulation for orthonormal zonal basis functions that are a generalization of the SH basis over an arbitrary zone of a sphere.

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, as well as for compressive sensing approaches as discussed in Sect. 4.2.1.

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 (Fig. 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 (Fig. 3, right). In our experiments, we worked with the distribution based BRDF model by Ashikhmin (2006) due to the data-driven nature of the model and the simplicity of the fitting procedure. The D-BRDF model is also a generalization of the Ashikhmin-Shirley-Phong model (Ashikhmin and Shirley 2000), which was recently found to be particularly well-suited for fitting to measured data (Ngan et al. 2005). However, the measured zonal data can be fitted to any other suitable analytic model using a numerical procedure such as Press et al. (1992).

Note that such fitting of BRDF measurements to analytic models is in general sub-optimal as one is limited by the expressiveness of the particular analytic model employed for the fitting process. We try and preserve as much of the measurements as possible in this context by employing a data-driven distribution BRDF model for the fitting process, where the micro-facet distribution is extracted from the measured data instead of assuming a Gaussian or cosine lobe distribution in the fitting process.

4.2.1 Compressive Sensing for Specular BRDFs

A very interesting related approach for efficiently encoding high frequency signals can be found in the rapidly evolving compressive sensing literature (Candés and Romberg 2005; Candés et al. 2006). The compressive sensing theory states that if a signal $f \in \Re^N$ has a sparse representation in a given basis, then it is possible to accurately recover f with a small

Fig. 3 An illustration of the suppression of ringing through fitting analytical BRDF models. *Left*: original acrylic blue paint BRDF acquired by Matusik et al. (2003). Center: 10th order zonal reconstruction, rendered after transformation into SH, exhibiting severe ringing artifacts. *Right*: Corresponding D-BRDF fit to the zonal reconstruction

number of projections onto *randomly chosen* subspaces via a tractable optimization program.

A finite signal $f \in \Re^N$ can be recovered from a set of K linear measurements

$$y_k = \langle f, z_k \rangle, \quad k = 0, \ldots, K-1 \text{ or } y = \mathbf{M}f \tag{11}$$

taken against vector $z_k \in \Re^N$ of the *measurement ensemble* $\mathbf{M}$ (z_k is the kth row of $\mathbf{M}$) with $K \ll N$ as the solution to the following convex optimization problem:

$$\min_{g \in \Re^N} ||g||_{l_1} := \sum_t g(t) \quad \text{s.t.} \quad y = \mathbf{M}g. \tag{12}$$

In our case, such a measurement ensemble can be formulated by randomly choosing K bases from the basis space spanned by higher order orthonormal zonal basis functions:

$$\mathbf{M} = \mathbf{Z}_K.$$

This results in a significant reduction in the number of projected zonal basis functions, and hence acquisition time, for encoding specular materials. Note that the BRDF function may not necessarily be very sparse in the zonal basis representation. The compressive sensing theory still holds as long as the BRDF function is *compressible* in this basis, i.e., the coefficients of the basis decay with a power law (Candés and Romberg 2005). This is true for most band-limited signals and hence is a valid approach for very glossy and moderately specular BRDFs. However, the need to extrapolate the BRDF data outside the measurement zone still exists with a compressive sensing approach.

Hence, we propose to further fit the output of the optimization g to the D-BRDF model as a way of extrapolation of the BRDF, and for overcoming the Gibbs phenomenon when projecting very specular BRDFs into the zonal basis. To simplify the optimization in (12), we solve only for a back-scattering BRDF function $f_{back\text{-}scatter}$, thereby reducing the dimensionality of f from 4D to 2D. Finally, we numerically fit the 2D function $f_{back\text{-}scatter}$ to a back-scattering D-BRDF ($\hat{k}_1 = \hat{k}_2 = \hat{h}$) to obtain the appropriate coefficients of the D-BRDF model.

5 Analytical Model as Measurement Basis

While the orthonormal zonal basis function introduced in Sect. 4 is well suited for encoding low-frequency BRDFs, it suffers from ringing artifacts in the reconstruction for specular BRDFs. One way of overcoming this is by fitting the zonal reconstruction to an analytical BRDF model as discussed in Sect. 4.2. We observe that analytic models are often used for representing the reflectance property of specular materials in computer vision and graphics applications. Hence for such applications, it would be ideal to directly measure the response of such materials in a basis defined by an analytic BRDF model.

Most BRDF models are unfortunately not easily separated into a structured illumination and a reflected light basis, which presents a challenge for deriving the appropriate basis illumination. However, the D-BRDF model is an example of a model where such an illumination basis can constructed for measurements in the back-scattering direction. For the back-scattering direction, $\hat{k}_1 = \hat{k}_2 = \hat{k} = \hat{h}$ and the model simplifies to

$$\rho(\hat{k}, \hat{k}) = \frac{c \cdot r_0 \cdot p(\hat{h})}{2(\hat{k} \cdot \hat{n}) - (\hat{k} \cdot \hat{n})^2}, \tag{13}$$

providing a function that is proportional to the distribution $p(\hat{h})$. Here c is a normalization constant and r_0 is the reflectance at normal incidence.

We propose to model specular materials with the usual specular and diffuse separation:

$$\rho(\hat{k}_1, \hat{k}_2) = k_s \cdot \rho_{spec}(\hat{k}_1, \hat{k}_2) + \frac{k_d}{\pi}.$$

Then the measurement process just involves obtaining estimates of the diffuse and specular reflectance coefficients k_d

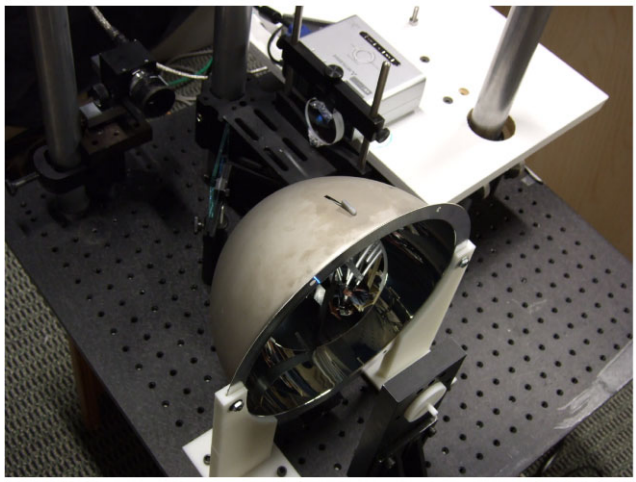

Fig. 4 Photograph of the proposed BRDF acquisition setup including a camera, a projector, a beam-splitter, and two curved reflectors mounted on a 40 cm × 40 cm optical bench

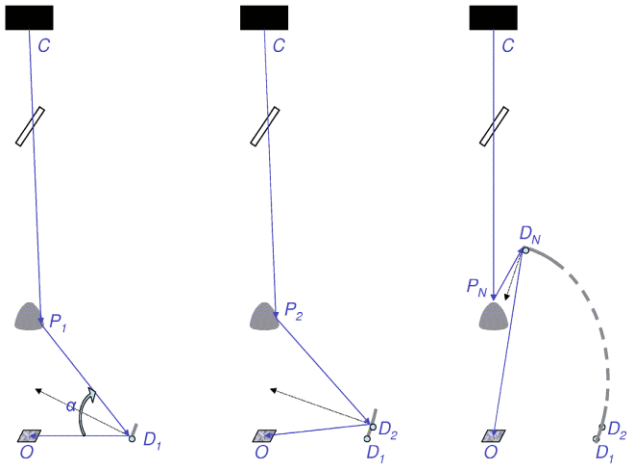

Fig. 5 Iterative process for designing the profile of the reflective dome for a fixed convex parabolic reflector

and k_s respectively with a basis of chosen distribution $p(\hat{h})$. An advantage of such an approach over using a linear basis function such as the zonal basis is that the choice of basis is data-driven, thus making the encoding efficient.

We first project a constant basis $\frac{1}{\pi}$ over the measurement zone. The camera observes $k_d \cdot (\cos^2\theta_{min} - \cos^2\theta_{max})$ in response to this basis for a number of directions over the zone that we then compute an average of to obtain an estimate of k_d.

In order to obtain an estimate of k_s, we sample the measurement zone with the following basis centered at a number of directions $\hat{k}$:

$$\rho_{spec}(\hat{k},\hat{k}) = \frac{c \cdot p(\hat{h})}{2(\hat{k}\cdot\hat{n}) - (\hat{k}\cdot\hat{n})^2}. \tag{14}$$

The camera records $k_s \cdot \cos\theta \cdot r_0$ in response to this basis at each sampling direction $\hat{k}$. We can estimate r_0 by observing the gain in reflectance at closer to grazing angles compared to normal incidence as suggested by Ashikhmin (2006). Thereafter k_s can be estimated similarly from the multiple back-scattering direction measurements by computing averages of the measurements around each $\hat{k}$ weighted by the specular lobe in that direction.

Appropriate distributions $p(\hat{h})$ for many commonly found materials have been well documented in the literature. For example, Ngan et al. (2005) have published cosine lobe and microfacet distribution fits to the 100 measured isotropic materials in the MERL database (Matusik et al. 2003), while distributions for fabrics such as velvet and satin and general anisotropic Gaussian distributions for materials like brushed metal can be found in (Ashikhmin et al. 2000).

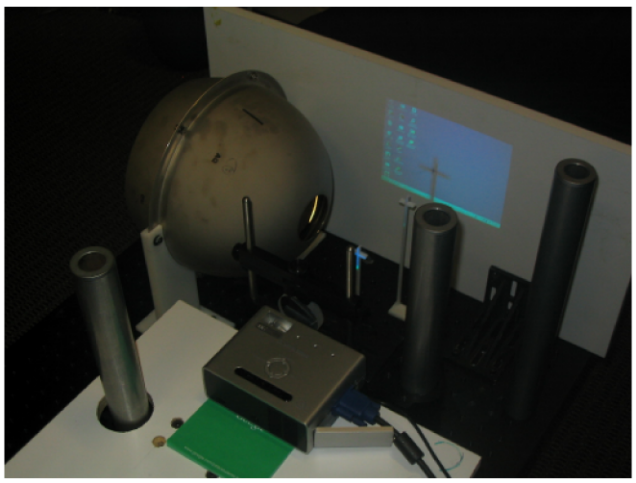

Fig. 6 Projector alignment using two crosses mounted cross along the optical path and observing the shadows on a projection plane

6 Measurement Setup and Calibration

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 (Debevec and Malik 1997) for high dynamic range (HDR) acquisition.

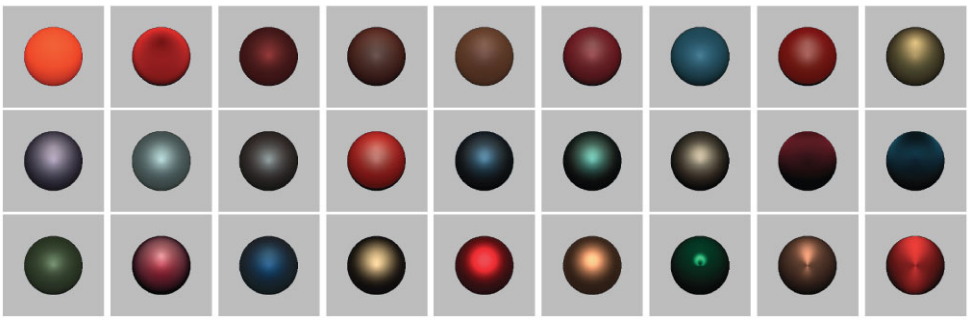

Fig. 7 Various BRDFs acquired with our prototype setup using zonal basis functions. *Top row*: from *left* to *right*—bright orange paper, red velvet, maroon synthetic fabric, brown leather, coated brown envelope, red printer toner, blue rubber band, glossy red paper, glossy blue-gray paper, Lindt™ chocolate box paper. *Center row*: from *left* to *right*—magenta plastic with grain finish, retro-reflective plastic, dark brown plastic coffee lid, Krylon™ banner red paint, Krylon™ 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 KitKat™ wrapping paper, copper colored Lindt™ chocolate wrapping paper, anisotropic plastic guitar pick, anisotropic copper coin, anisotropic red satin

Our optical setup consists of two mirrored components, a convex parabola and a concave reflective dome as shown in Figs. 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 (Fig. 5, left). The surface normal at the rim defines a tangent plane in D_1. For the next camera ray reflecting of P_2, we compute the intersection D_2 of the reflected ray with the tangent plane of D_1 (Fig. 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 (Fig. 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. The parabolic profile for the small convex mirror element was chosen after simulation with various alternative shapes, including a hemispherical mirror. The parabolic design provided the most uniform sampling density across the measurement zone in our experiments.

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 Fig. 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. Zonal basis functions defined over this range are presented in Fig. 16.

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 (Fig. 4).

6.1 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 (Fig. 6).

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 28 cm. 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 mea-

Fig. 8 The Audi-TT model rendered with acquired BRDFs of 2 different paint samples. The BRDFs were acquired using 25 4th order basis functions as defined in this paper, and then rendered with a basis transformation into spherical harmonics. *Top*: Metallic teal automotive paint. *Bottom*: Krylon™ 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

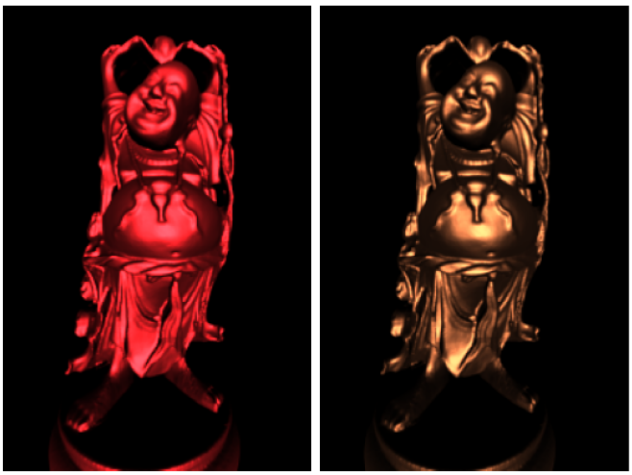

Fig. 9 Specular chocolate wrapping papers acquired using higher order zonal basis functions, and then fit to an analytical model for rendering. *Left*: Red KitKat™ wrapping paper. *Right*: Copper colored Lindt™ chocolate wrapping paper

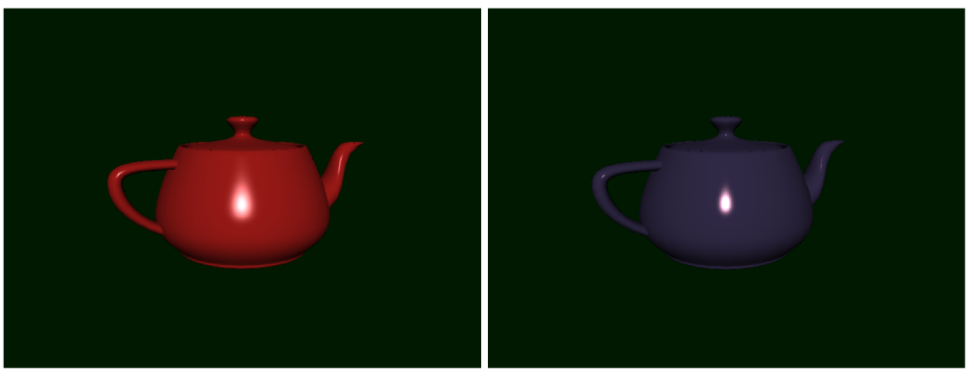

Fig. 10 Specular materials acquired using compressive sensing with higher order zonal basis functions, and then fit to an analytical model for rendering. *Left*: Red electrical tape (specular exponent $s = 270$). *Right*: Cadbury™ chocolate wrapping paper (specular exponent $s = 928$)

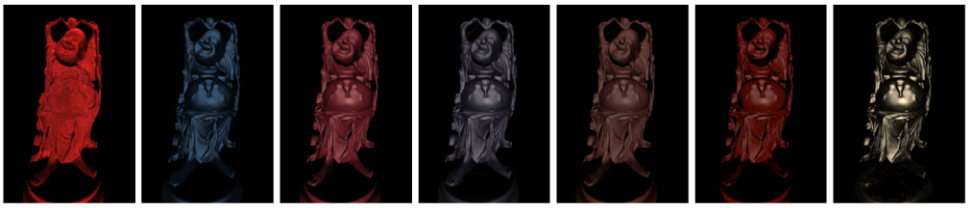

Fig. 11 Representative set of BRDFs acquired with lower order zonal basis functions rendered under directional lighting. From *left* to *right*: red velvet, synthetic blue fabric, red printer toner, magenta plastic sheet, brown leather,glossy red paper, chrome gold dust automotive paint

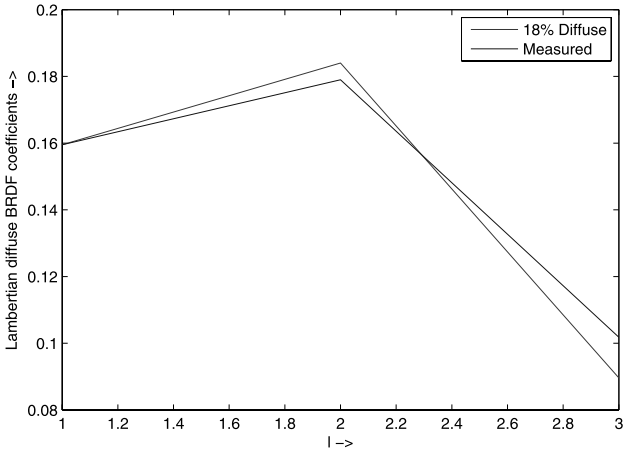

Fig. 12 Plot of measured coefficients of gray card vs. that of pure Lambertian diffuse reflectance. Note that the measured coefficients have been scaled so that the DC term ($l = 0, m = 0$) matches the DC term of 18% Lambertian diffuse reflectance

surements 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 (Ramamoorthi and Hanrahan 2001). The relative scaling factors for each color channel are obtained by white-balancing the results of the gray card measurements.

7 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 acquired with the orthonormal zonal basis functions (Sect. 4), 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 8 presents the BRDFs of 2 different paint samples that we acquired using 4th 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 (Pharr and Humphreys 2004).

A representative set of the BRDFs acquired using lower order ($l \leq 6$) zonal basis functions is shown in Fig. 11. For this class of materials, the entire process of acquisition followed by a basis transformation into the SH basis took under three minutes.

Figure 9 demonstrates the specular materials, in this case shiny metal foil chocolate wrapping papers, that we acquired using higher order zonal basis functions and 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 the satin samples (Fig. 7, bottom row).

Figure 10 presents results of our preliminary experiments on acquiring specular materials with compressive sensing with the zonal basis. Here, we projected a random sub-set of $K = 25$ zonal basis functions from a space spanned by 15 orders of the zonal basis (corresponding to $N = 256$) and then fit the result g of the optimization (12) specifically to the specular component of a D-BRDF with a cosine lobe distribution $p(\hat{h}) = (\hat{n} \cdot \hat{h})^s$. We separately approximated the diffuse component of the D-BRDF in these examples using the first 9 (2nd order) zonal basis coefficients. Thus, we employ compressive sensing in our work as a way of acquiring

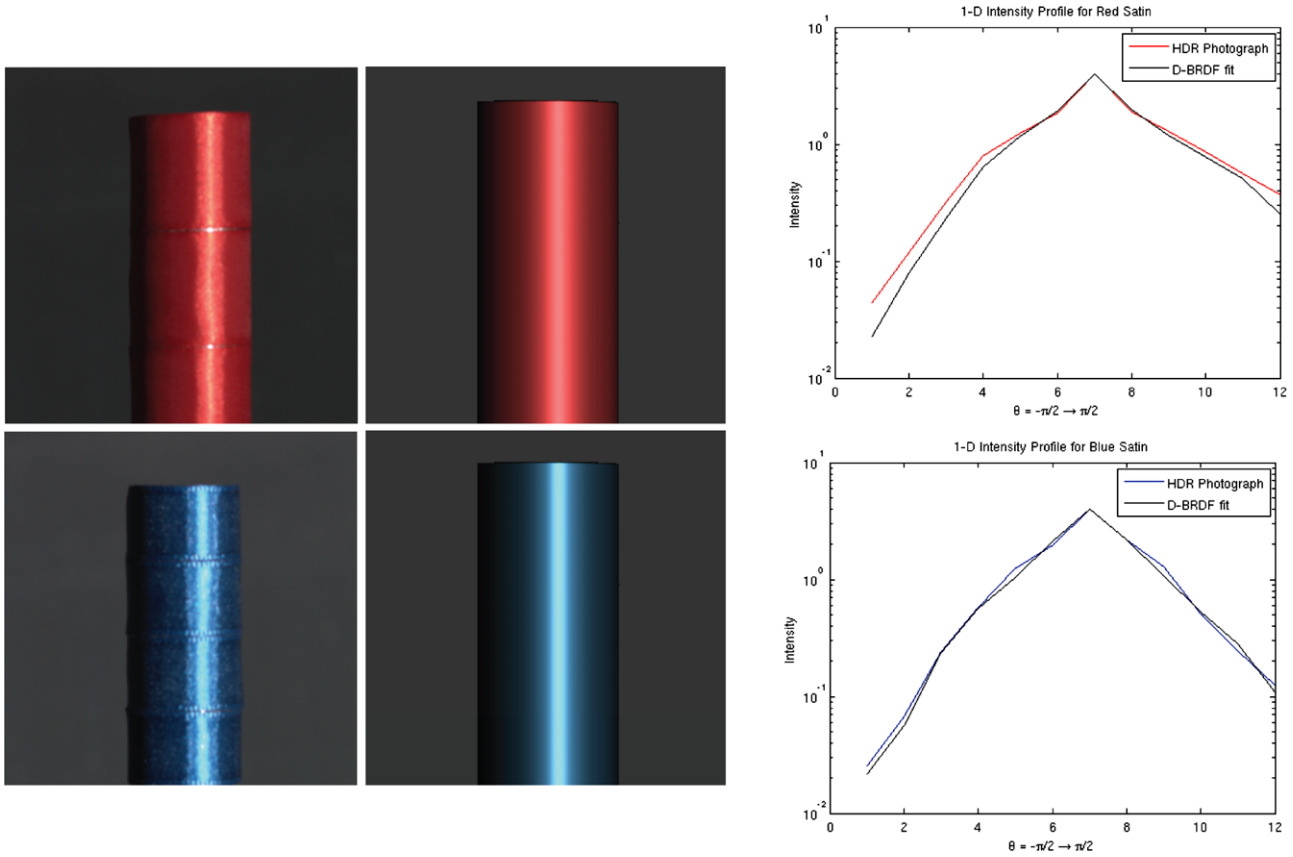

Fig. 13 Comparison of two kinds of acquired satin samples wrapped around a cylinder as lit by a point source against real photographs. Here, the response to 8th order zonal basis functions were fit to the D-BRDF analytic model in a post-process. *Left column*: Photographs of the red and blue satin samples. *Center column*: Rendering of the D-BRDF fits. *Right column*: Plots of 1-D intensity profiles for a cross-section of the cylinder in the photographs and the corresponding renderings. Plots represent data for a single color channel (red and blue respectively for the two satin samples)

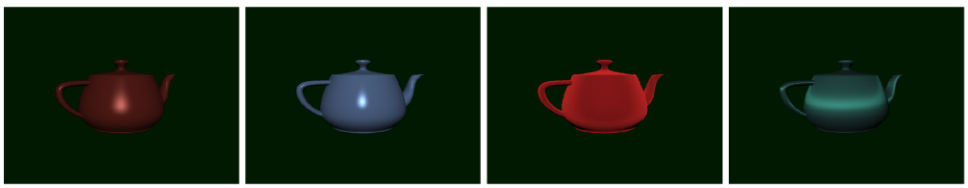

Fig. 14 Specular materials acquired using data-dependent basis functions based on the D-BRDF analytic model. From *left* to *right*: red wax ($p(\hat{h}) = (\hat{n} \cdot \hat{h})^{100}$), blue electrical tape ($p(\hat{h}) = (\hat{n} \cdot \hat{h})^{300}$), red velvet ($p(\hat{h}) = 1 + 4\exp(-\cot^2\theta)$), green satin ($p(\hat{h}) = 0.7p_{flats} + 0.3p_{ends}(\hat{h})$, where $p = \exp(-\tan^2\theta\cos^2\phi/\sigma_x^2)$ with $\sigma_x = 0.1$ for p_{flats}, and $\sigma_x = 0.3$ for p_{ends})

data with fewer measurements than higher order linear basis functions. This enables us to reduce the number of measurements by a factor of 5 compared to the linear measurements which is consistent with compressive sensing literature.

As a step towards quantitative validation of our approach of measuring low frequency BRDFs with the zonal basis functions and transforming into the SH basis for extrapolation, we used an 18% gray card as a diffuse reflectance standard, and compared the recovered coefficients to the ones expected for a diffuse target (Ramamoorthi and Hanrahan 2001). For coefficients of up to 2nd order basis functions, the error is within a few percent of the expected value (Fig. 12).

In order to validate our measurement and fitting approach for high frequency BRDFs, we photographed (in HDR) two satin samples wrapped around a cylinder in a dark room and lit by a collimated point light source. Figure 13 presents the comparisons of these photographs with the corresponding renderings of the D-BRDF fits to these samples. The high-

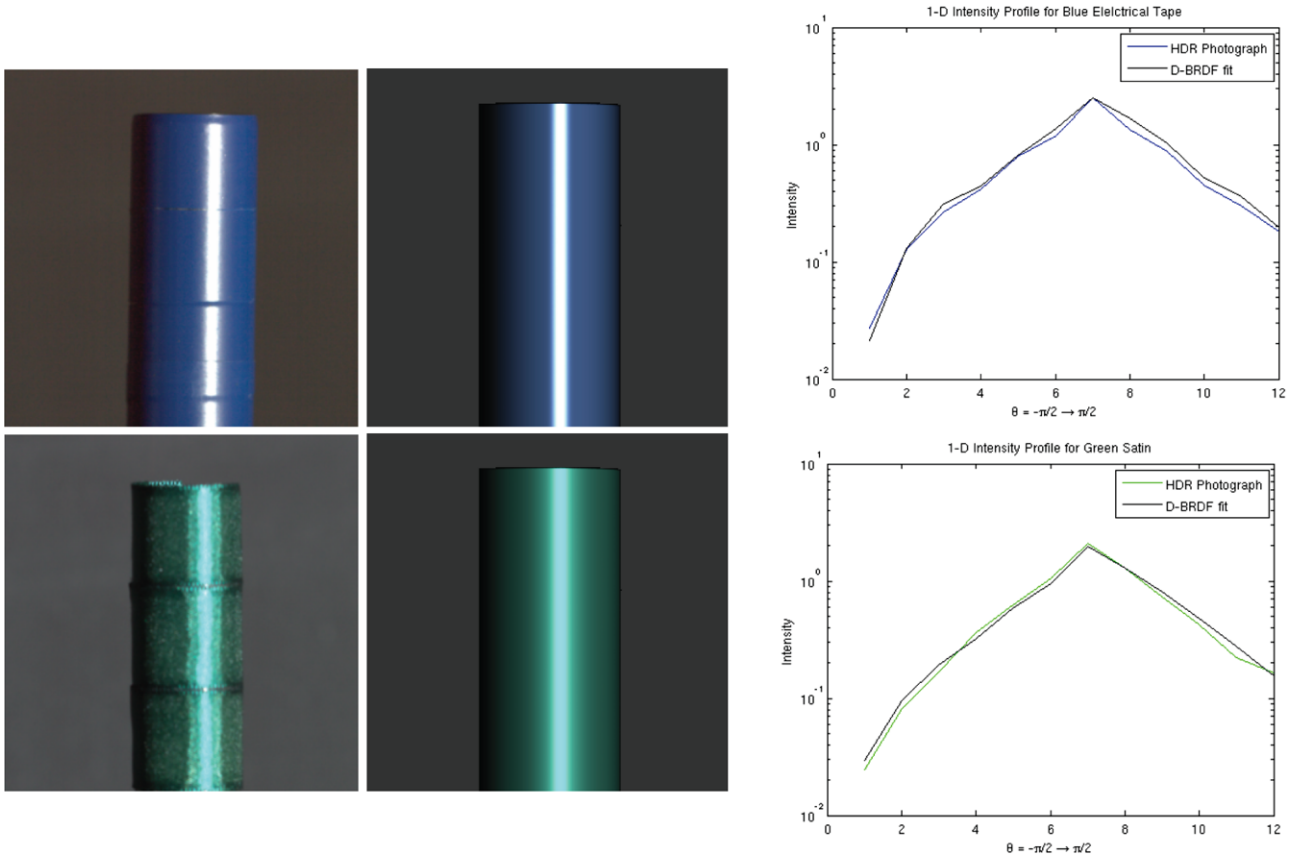

Fig. 15 Comparison of D-BRDF fits of two kinds of acquired samples wrapped around a cylinder as lit by a point source against real photographs. In both the cases the D-BRDF fits were obtained by projecting just 25 basis images defined by the analytic model. *Left column*: Photographs of the samples. *Center column*: Rendering of the D-BRDF fits. *Right column*: Plots of 1-D intensity profiles for a cross-section of the cylinder in the photographs and the corresponding renderings. Plots represent data for a single color channel (blue and green respectively for the tape and satin samples). *Top row*: Blue electrical tape. *Bottom row*: Green satin

lights in the rendered images (center column) are a close match to the real photographs (left column). As an additional step towards quantitative validation, we plot the 1D intensity profiles for a cross-section of the cylinders in the HDR photographs versus the same profile in the renderings (right column). Here, we scaled the D-BRDF fits to match the peak intensities of the specular highlights of the corresponding HDR photographs. As seen in the plots, the two profiles match very closely validating the measurement and fitting process.

Figure 14 presents results of our measurements with data-dependent basis functions based on the D-BRDF analytic model. Here, we approximated the diffuse component of the BRDF by averaging the response to the constant basis $\frac{1}{\pi}$. The specular component of the BRDF was estimated from an average of 24 back-scattering measurements uniformly distributed over the measurement zone in response to a basis illumination constructed from an appropriate distribution $p(\hat{h})$ (14). For the red wax sample, we constructed the basis illumination for the specular lobe based on the cosine lobe fit to a similar wax sample in the MERL database obtained by Ngan et al. (2005). Similarly, we chose the specular exponent for the blue electrical tape based on the numerical fit obtained for a similar sample (red electrical tape) in Fig. 10. The distributions for the velvet and satin samples are based on the Ashikhmin model (Ashikhmin et al. 2000).

Figure 15 presents the comparisons of photographs of the blue tape and green satin samples with the corresponding renderings of the D-BRDF fits to these samples. The highlights in the rendered images are visually a close match to the real photographs. Moreover, the 1-D intensity profiles in the photographs also match very closely to the corresponding renderings, validating our approach encoding BRDFs with data-dependent basis illumination based on an analytic model. We were able to obtain very good estimates for the parameters of the BRDF model with *very few* images with this approach compared to what would be required for similar materials with a higher order zonal basis acquisition (Fig. 13).

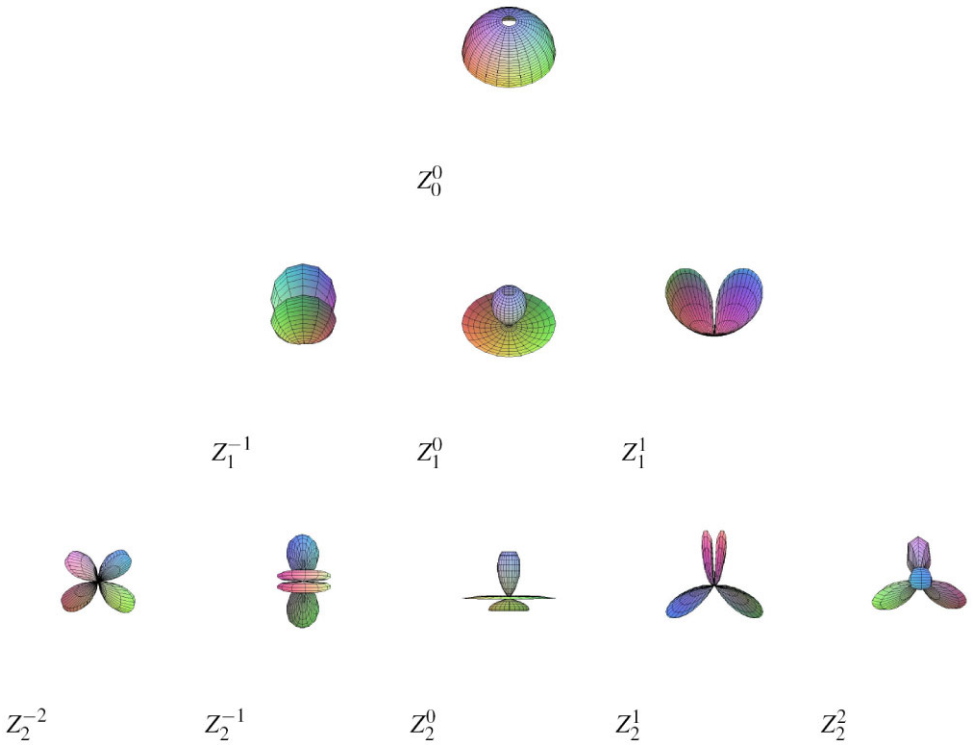

Fig. 16 The plots of zonal basis functions Z_l^m defined over the theoretical measurement space $[\pi/20, \pi/2] \times [0, 2\pi]$, for $l \leq 2$

8 Conclusions

In this paper, we have presented a novel basis function approach to BRDF measurement. Our contributions include a new theory for basis function BRDF acquisition, the development of an orthonormal basis for spherical zones and a data-dependent basis based on an analytic BRDF model, as well as 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. It would be interesting to look into manufacturing techniques that are able to produce such domes. One possible approach for manufacturing our complete dome with electroforming would be to split the dome at its widest point into two parts and then mechanically join the separately electroformed parts. Another option would be to use an off-the-shelf convex profile for the dome such as a hemisphere or an ellipsoid mirror such as the approach of Mukaigawa et al. (2007). The design trade-offs with the choices for the profiles of the curved reflector are the uniformity of angular sampling and coverage of the measurement zone along with cosiderations of camera and projector optics (our approach) versus the ease of manufacturing and availability of a standard profile.

Our preliminary experiments with a compressive sensing approach in conjunction with basis illumination for BRDF acquisition is encouraging. However, our choice of the orthonormal zonal basis may not be optimal for very specular materials. Alternative basis functions such as the wavelet basis or even the Gaussian random ensemble discussed in the compressive sensing literature might be more interesting for such materials. However, for highly specular materials it might be better to employ a diffuse/specular separation and simply point sample the specular highlight instead of trying to encode it with basis functions.

Our experiments with designing data-dependent basis functions based on an analytic BRDF model provided very good results for a range of commonly occurring materials. We were able to obtain good estimates for parameters of the BRDF model with just a few images of the sample in response to the data-driven basis. Further research needs to be done for designing basis illumination for other analytic models in the literature and for materials with other kinds of distributions. Also, the accuracy of the fits to the diffuse and specular components of the BRDF can be further improved with polarization of the incident basis illumination (Ma et al. 2007). Employing such polarization techniques to BRDF acquisition would be a very interesting direction of future work.

Comparing our measurements with basis based on analytic models to experiments with compressive sensing, the number of measurements required for both methods turned out to be very similar for the samples that we captured. However, the advantage of compressive sensing over the basis driven by an analytic model is that it requires no knowledge

of the micro-facet distribution of the material, and hence can be used for measuring more general materials compared to a basis based on specific distributions. We envision that such compressive sensing techniques could also be used for measuring more general light transport in a scene.

In conclusion, due to the basis function illumination and the dispensing of all moving parts, BRDF measurement with our approach 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 for vision and graphics applications.

Appendix 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

$$\int_{-1}^{1} P_l^m(x) P_{l'}^m(x)dx = \frac{2(l+|m|)!}{(2l+1)(l-|m|)!}\delta_{ll'}. \tag{15}$$

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

$$Y_l^m(\theta,\phi) = \begin{cases} \sqrt{2}K_l^m \cos(m\phi) P_l^m(\cos\theta) & \text{if } m>0, \\ \sqrt{2}K_l^m \sin(-m\phi) P_l^{-m}(\cos\theta) & \text{if } m<0, \\ K_l^0 P_l^0(\cos\theta) & \text{if } m=0, \end{cases} \tag{16}$$

where K_l^m is the SH normalization constant:

$$K_l^m = \sqrt{\frac{(2l+1)(l-|m|)!}{4\pi(l+|m|)!}}. \tag{17}$$

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

$$\hat{P}_l^m(x) = P_l^m(n_1 \cdot x - n_2), \tag{18}$$

with

$$n_1 = \frac{2}{\cos\theta_{min} - \cos\theta_{max}},$$

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

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

References

Ashikhmin, M. (2006). Distribution-based BRDFs. http://jesper.kalliope.org/blog/library/dbrdfs.pdf.

Ashikhmin, M., & Shirley, P. (2000). An anisotropic phong BRDF model. *Journal of Graphics Tools*, *5*(2), 25–32.

Ashikhmin, M., Premoše, S., & Shirley, P. (2000). A microfacet-based BRDF generator. In *SIGGRAPH '00: Proceedings of the 27th annual conference on computer graphics and interactive techniques* (pp. 65–74).

Basri, R., & Jacobs, D. W. (2003). Lambertian reflectance and linear subspaces. *IEEE PAMI*, *25*(2), 218–233.

Candés, E., & Romberg, J. (2005). Practical signal recovery from random projections. In *Proc. of IS&T/SPIE s 17th Annual Symposium on Electronic Imaging*.

Candés, E., Romberg, J., & Tao, T. (2006). Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, *52*(2), 489–509.

Cook, R. L., & Torrance, K. E. (1982). A reflectance model for computer graphics. *ACM Transactions on Graphics*, *1*(1), 7–24.

Cornell (2005). CORNELL light measurement laboratory. http://www.graphics.cornell.edu/research/measure/.

CUReT (1999). CUReT: Columbia-utrech reflectance and texture. http://www.cs.columbia.edu/CAVE/curet/.

Dana, K. (2001). BRDF/BTF measurement device. In *Proc. of ICCV '01* (pp. 460–466).

Debevec, P., & Malik, J. (1997). Recovering high dynamic range radiance maps from photographs. In *Proc. of ACM SIGGRAPH '97* (pp. 369–378).

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

Gautron, P., Křivánek, J., Pattanaik, S., & Bouatouch, K. (2004). A novel hemispherical basis for accurate and efficient rendering. In *Eurographics symposium on rendering* (pp. 321–330).

Ghosh, A., Achutha, S., Heidrich, W., & O'Toole, M. (2007). BRDF acquisition with basis illumination. In *Proc. of IEEE international conference on computer vision (ICCV)*.

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

Gorski, K. M. (1994). On determining the spectrum of primordial inhomogeneity from the COBE DMR sky maps: method. *The Astrophysical Journal*, *430*, L85–L88.

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

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

He, X. D., Torrance, K. E., Sillion, F. X., & Greenberg, P. (1991). A comprehensive physical model for light reflection. In *SIGGRAPH '91: Proceedings of the 18th annual conference on computer graphics and interactive techniques* (pp. 175–186).

Koenderink, J., van Doorn, A., & Stavridi, M. (1996). Bidirectional reflection distribution function expressed in terms of surface scattering modes. *ECCV '96. 4th European Conference on Computer Vision*, *2*, 28–39.

Kuthirummal, S., & Nayar, S. K. (2006). Multiview radial catadioptric imaging for scene capture. *ACM Transactions on Graphics (also Proc. of ACM SIGGRAPH)*, *25*(3), 916–923.

Lafortune, E., Foo, S.-C., Torrance, K., & Greenberg, D. (1997). Non-linear approximation of reflectance functions. In *Proc. of ACM SIGGRAPH '97* (pp. 117–126).

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

Lensch, H., Kautz, J., Goesele, M., Heidrich, W., & Seidel, H.-P. (2001). Image-based reconstruction of spatially varying materials. In *Eurographics Workshop on Rendering* (pp. 104–115).

Ma, W.-C., Hawkins, T., Peers, P., Chabert, C.-F., Wiess, M., & Debevec, P. (2007). Rapid acquisition of specular and diffuse normal maps from polarized spherical gradient illumination. In *Proc. Eurographics symposium on rendering*.

Malzbender, T., Gelb, D., & Wolters, H. (2001). Polynomial texture maps. In *SIGGRAPH '01: Proceedings of the 28th annual conference on computer graphics and interactive techniques* (pp. 519–528).

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

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

Mukaigawa, Y., Sumino, K., & Yagi, Y. (2007). High-speed measurement of BRDF using an ellipsoidal mirror and a projector. In *Proc. IEEE Computer Vision and Pattern Recognition (CVPR)* (pp. 1–8).

Ng, R., Ramamoorthi, R., & Hanrahan, P. (2003). All-frequency shadows using non-linear wavelet lighting approximation. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, *22*(3), 376–381.

Ngan, A., Durand, F., & Matusik, W. (2005). Experimental analysis of BRDF models. In *Proceedings of the Eurographics Symposium on Rendering* (pp. 117–226).

Nicodemus, F. E., Richmond, J. C., Hsia, J. J., Ginsberg, I. W., & Limperis, T. (1977). Geometric considerations and nomenclature for reflectance. *NBS Monograph*, *160*.

NIST (2003). NIST reference reflectometer: STARR facility. http://www.physics.nist.gov/Divisions/Div844/facilities/brdf/starr.html.

Peers, P., & Dutré, P. (2005). Inferring reflectance functions from wavelet noise. In *Proc. Eurographics symposium on rendering* (pp. 173–181).

Pharr, M., & Humphreys, G. (2004). *Physically based redering*. New York: Morgan Kaufmann.

Press, W., Flannery, B., Teukolsky, S., & Vetterling, W. (1992). *Numerical recipes: the art of scientific computing*. Cambridge: Cambridge University Press.

Ramamoorthi, R., & Hanrahan, P. (2001). An efficient representation for irradiance environment maps. In *Proc. of ACM SIGGRAPH '01* (pp. 497–500).

Ramamoorthi, R., & Hanrahan, P. (2002). Frequency space environment map rendering. In *Proc. of ACM SIGGRAPH '02* (pp. 517–526).

Sato, I., Okabe, T., Sato, Y., & Ikeuchi, K. (2003). Appearance sampling for obtaining a set of basis images for variable illumination. In *Proc. of ICCV'03* (pp. 800–807).

Schröder, P., & Sweldens, W. (1995). Spherical wavelets: efficiently representing functions on the sphere. In *Computer graphics 29, annual conference series* (pp. 161–172).

Ward, G. J. (1992). Measuring and modeling anisotropic reflection. In *SIGGRAPH '92: Proceedings of the 19th annual conference on computer graphics and interactive techniques* (pp. 265–272).

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

Westin, S., Arvo, J., & Torrance, K. (1992). Predicting reflectance functions from complex surfaces. In *Computer graphics 26, annual conference series* (pp. 255–264).