Acquisition and Modeling of Material Appearance

by

Wai Kit Addy Ngan

B.S.E., Princeton University (2001)
S.M., Massachusetts Institute of Technology (2003)

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

Massachusetts Institute of Technology

August 2006

Author . . . . . .
Department of Electrical Engineering and Computer Science
August 30, 2006

Certified by . . . . . .
Frédo Durand
Associate Professor
Thesis Supervisor

Accepted by . . . . . .
Arthur C. Smith
Chairman, Department Committee on Graduate Students

Acquisition and Modeling of Material Appearance

by

Wai Kit Addy Ngan

Submitted to the Department of Electrical Engineering and Computer Science
on August 30, 2006, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

In computer graphics, the realistic rendering of synthetic scenes requires a precise description of surface geometry, lighting, and material appearance. While 3D geometry scanning and modeling have advanced significantly in recent years, measurement and modeling of accurate material appearance have remained critical challenges. Analytical models are the main tools to describe material appearance in most current applications. They provide compact and smooth approximations to real materials but lack the expressiveness to represent complex materials. Data-driven approaches based on exhaustive measurements are fully general but the measurement process is difficult and the storage requirement is very high. In this thesis, we propose the use of hybrid representations that are more compact and easier to acquire than exhaustive measurement, while preserving much generality of a data-driven approach.

To represent complex bidirectional reflectance distribution functions (BRDFs), we present a new method to estimate a general microfacet distribution from measured data. We show that this representation is able to reproduce complex materials that are impossible to model with purely analytical models. We also propose a new method that significantly reduces measurement cost and time of the bidirectional texture function (BTF) through a statistical characterization of texture appearance. Our reconstruction method combines naturally aligned images and alignment-insensitive statistics to produce visually plausible results. We demonstrate our acquisition system which is able to capture intricate materials like fabrics in less than ten minutes with commodity equipments.

In addition, we present a method to facilitate effective user design in the space of material appearance. We introduce a metric in the space of reflectance which corresponds roughly to perceptual measures. The main idea of our approach is to evaluate reflectance differences in terms of their induced rendered images, instead of the reflectance function itself defined in the angular domains. With rendered images, we show that even a simple computational metric can provide good perceptual spacing and enable intuitive navigation of the reflectance space.

Thesis Supervisor: Frédo Durand
Title: Associate Professor

Acknowledgments

First and foremost, I would like to thank my advisor Professor Frédo Durand for his support and guidance for the past four years. His ever-flowing stream of insights and ideas in broad areas of computer graphics have never ceased to amaze me. Working with him has been a fruitful and enjoyable learning experience. I am also very grateful that he has made himself available to me most of the time despite his immense workload. In particular, my thesis would not have been completed in time if Frédo had not presented the two papers at EGSR 2006 on my behalf. Our discussions also go beyond research - I have enjoyed chatting about photography and travel with him. Frédo has played an important role in making my time in graduate school a pleasant experience.

I would like to thank Professor Leonard McMillan for advising me during my first year of graduate studies. I would also like to thank Wojciech Matusik for his collaborations on more than a few projects, including most work presented in this thesis as well as other ones. His devotion and relentless energy are contagious and have propelled me to try to think and work harder. I also would like to express my gratitude to Hanspeter Pfister, who co-advised me on several projects at MERL.

I would like to thank my thesis committee members Professors Jovan Popovic and Edward Adelson for their time and input on the thesis, especially given the tight schedules. I want to thank Eric Chan for his insightful inputs on both graphics research and photography, and being a great officemate. I would also like to thank Sara Su who has been a supportive officemate and practically proofread every paper I have written, including this thesis. Soonmin Bae, Jiawen Chen, Barb Cutler, Paul Green, Tilke Judd, Jan Kautz, Jaakko Lehtinen, Tom Mertens, Sylvain Paris, Marco da Silva, Daniel Vlasic, Robert Wang, Matthias Zwicker have all provided valuable feedbacks to my research. Tom Buehler has helped me to produce several paper videos, and Bryt Bradley has been the magical force behind the scenes of everything the graphics group has accomplished. I would like to thank all members of the MIT Computer Graphics Group for their support.

Finally, I would like to thank my family for their love and support over many years. My girlfriend Joyce has provided more than spiritual support - she has provided tremendous help in the construction of various hardware setups and measurement targets. Of course, I am even more grateful for her love, patience and support. This thesis is dedicated to her.

Contents

List of Figures

Chapter 1

Introduction

Recent advances in computer graphics have enabled the production of highly realistic rendered images when a precise description of a scene is provided. A complete scene description includes surface geometry, lighting, and material appearance. While 3D geometry scanning and modeling have advanced significantly in recent years, measurement and modeling of accurate material appearance have remained critical challenges. Analytical reflectance models are the main tools used to describe appearance in most current applications. They provide compact and smooth approximations to real materials but lack the expressiveness to represent complex materials. In the production industry, the realism of materials in graphics renderings has improved tangibly, but the improvement is not entirely the result of advancement in appearance techniques. To some extent, increase in raw computation power has provided improved image quality with existing models. More importantly, graphics artists have become more experienced in exploiting existing techniques creatively. However, several classes of materials remain difficult to reproduce, most notably human skin and fabrics. Direct exhaustive measurement of real materials is a viable alternative for materials that are difficult to model (e.g. [10]), but its use is far from widespread due to the high cost of measurement time and equipment setup. In this thesis, we propose the use of hybrid representations that are more compact and easier to acquire than exhaustive measurement, while preserving much generality of data-driven approaches.

1.1 Context

In the early days of computer graphics, simple empirical models were developed to approximate local reflection. Models like Lambert's and Phong's provide only crude qualitative approximations to material appearance. Over the past few decades, numerous models have been proposed for improved fidelity, flexibility and computation efficiency. However, there has been a lack of studies to compare these models quantitatively. The first contribution of this thesis is to provide such a study. We quantitatively validate a number of reflectance models based on high quality measured data. In the process, we gain important insights into the merits of different model formulations. In addition, our new measurements of anisotropic materials show that existing analytical models lack expressiveness needed to handle complex appearance. A semi-parametric representation, based on general microfacet distributions [1], is able to provide much better approximations to the measurement while still achieving good compactness compared to full-blown data. We demonstrate an efficient method to fit measured data to this representation.

Direct acquisition of material appearance has become very popular over the past decade for several reasons. As light transport algorithms mature, they have become more sensitive to the quality of the input data. Measured data are either used to fit analytical models or are directly used for rendering. If the measurement resolution is high enough, direct rendering can provide the highest quality since it is not bounded by the expressiveness of any particular model. Moreover, the increase in computation power and storage has made measurement and processing of reflectance data, often in size of gigabytes, more tractable. Also, the image-based approaches introduced recently (e.g. [65]) allow for more efficient measurement by collecting many reflectance samples in a single image. However, quantities such as spatially-varying reflectance are still unwieldy to acquire due to the high dimensionality. State-of-the-art measurement devices require expensive robotics setup, and the measurement process is very time-consuming. In this thesis, we propose a method to reduce acquisition cost and time significantly. The main idea of our approach is to measure visually important statistics which are insensitive to precise registration. By removing the requirement of precise alignment, the measurement setup can be much simplified: no

robotics setup is required, and multiple views can be acquired in a single image. Such dramatic reduction in measurement precision and sampling does come at a cost: materials with strong parallax or sharp specularities are not reproduced faithfully. Nonetheless, we will show that our technique is general enough for a large class of materials and degrades gracefully for the failure cases.

Another important element of the appearance reproduction pipeline is user design. In most current graphics applications (e.g. movies, games), material appearance is typically specified with analytical models rather than measured data. Choosing the right parameters for a particular model, however, is often difficult because these parameters can be non-intuitive for users and can have non-uniform effects on the rendered image. Ideally, a given step in the control parameter space should produce a predictable and perceptually-uniform change of the rendered image. To alleviate this issue, software developers often include ad-hoc remapping of the parameters. However, a more systematic solution is desirable in order to make material modeling more intuitive. Systems that employ psychophysics have produced important advances in this direction [85], but the main drawback of these systems is the requirement of extensive user studies, which limits their scalability. In this thesis, we introduce a metric in the space of material appearance which corresponds roughly to perceptual measures. The main idea of our approach is to evaluate appearance differences in terms of the rendered images, instead of the radiometric quantity itself defined in the angular domains. With rendered images, we show that even a simple computational metric can provide good perceptual spacing and enable more intuitive navigation of the appearance space.

To sum up, the work presented in this thesis tackles a number of issues in different stages of the material appearance reproduction pipeline. For acquisition, our work on statistical reconstruction of texture appearance allows efficient measurement of high-dimensional reflectance data. For modeling, we validate existing models with measured data and provide guidance for their use. In addition, we demonstrate the limited expressiveness of pure analytical models and show that a semi-parametric model is needed for approximating complex materials. Finally, for user design, we introduce a image-based metric that corresponds to perceptual measures and facilitates an effective navigation in-

terface. In the following sections, we will introduce each component of the thesis in more details.

1.2 Modeling of Homogeneous Reflectance

The Bidirectional Reflectance Distribution Function (BRDF) describes how light interacts at a surface point. It is a function of the incoming and outgoing directions that describes how much of the incoming light is reflected. The BRDF can fully describe the appearance of homogeneous materials, e.g. most plastic and smooth metals. Early models such as Phong [87] and Lambertian [54] were not introduced formally as BRDF models, but they govern the same relationship. BRDF models can be divided roughly into two classes: empirical models seek to emulate the characteristics of material reflection qualitatively [6, 118, 52], while physically-based models derive the quantities from first principles of physics[13, 42]. While the latter models are physically plausible, they should not be immediately considered as *true* models as they all make significant simplifying assumptions. While new BRDF models are often introduced together with a few measured materials to illustrate their capabilities, no quantitative comparisons between different models have been published. We believe the main reason for the lack of validation is the difficulty of obtaining high-resolution measurement. In 2003, Matusik et al.[68] published a data set of over a hundred BRDFs measured at high dynamic range and high angular resolution. With the presence of such a data set, we have the opportunity to provide a detailed quantitative comparison between a number of popular models. To the best of our knowledge, this is first time multiple BRDF models are validated and compared using a large measured data set. The detailed validation results will be presented in Chapter 3.

An interesting alternative to pure analytical BRDF models was introduced by Ashikhmin et al.[1]. In their microfacet-based BRDF generator, a physically plausible BRDF can be created from an arbitrary distribution of microfacets. We derive a simple method to invert the process and estimate a discretely sampled microfacet distribution from BRDF measurements. Given the much higher degrees of freedom in this BRDF representation, it is not surprising that this representation provides higher fidelity than pure analytical models.

Since Matusik et al.'s data set is limited to isotropic materials, we have extended their measurement setup and measured a few anisotropic materials. We show that for some complex anisotropic materials, the sampled microfacet distribution is the only method that can qualitatively resembles the true BRDF. Our inversion procedure and the associated results will be presented in Chapter 4.

1.3 Measurement and Modeling of Spatially-Varying Reflectance

The most common solution to model spatially-varying appearance is the use of texture mapping. Texture mapping typically maps each surface location to a color value and is suitable for cases where the spatial variation is solely due to color differences. While a BRDF can be coupled with the texture by multiplication, intricate effects due to meso-scale geometry is not reproduced. A number of techniques have been proposed to augment texture mapping. Bump mapping simulates the effect of surface roughness by perturbing surface normals [7]. Horizon mapping enhances bump mapping by handling self shadowing effects [72, 105]. In both cases, an auxiliary map has to be provided by the user through estimation. A more general representation for spatially-varying reflectance is the Bidirectional Texture Function (BTF) [23], a 6D function that depends on the incoming and outgoing directions, as well as spatial locations. Current BTF measurements all require precisely calibrated equipments, as well as robotics setup for positioning the light, camera and the texture patch itself. Due to the large size of the BTF representation, most recent research has focused on data compression. In this thesis, we instead focus on an alternative representation that allows significant reduction in terms of measurement cost and time. Our representation is a hybrid representation that combines images with perfect correspondence and aggregate statistics that are alignment-insensitive. By removing part of the required correspondence in a traditional BTF measurement, we are able to measure BTF appearance through a much simplified setup. Complex materials like fabrics can be captured in less than ten minutes without elaborate calibration, and their appearance are reproduced in a visually plausible

way. In Chapter 5, we will describe the acquisition setup and our reconstruction procedure, followed by results and discussions.

1.4 Material Appearance Model Navigation

Although measurement of real materials has become popular in recent years, in most graphics applications, material appearance is still dominantly specified with analytical models. Despite the fact that many models have been introduced since Phong, in many commercial applications, only the simplest models are employed. The reason is twofold. First, many models are complex mathematically, and it is difficult for users to have an intuitive understanding. Second, the parameters of the models are often quantities that do not necessarily map to appearance changes uniformly. Using data from human experiments, Pellacini et al.[85] redefine and rescale the parameter axes of the Ward model [118] in a perceptually uniform way. However, the requirement of human experiments on such an approach limits its scalability to more complex models. In this thesis, we introduce a computation metric that roughly corresponds to perceptual spacing in the space of reflectance. The main idea of our approach is to evaluate reflectance differences in image space: instead of comparing reflectance in its native angular domain, we compare them in terms of their induced rendered images. We show that even a simple metric in the image space can produce visually plausible spacing. However, while we believe that an image-driven metric is a step in the right direction, we do not claim it to be a perceptual metric as it has not been validated by psychophysical experiments. In Chapter 6, we will explain the metric in details. We will also demonstrate a user interface based on this metric and show that it offers intuitive and uniform navigation in the space of analytical reflectance models.

1.5 Thesis Overview

In Chapter 2, we introduce the fundamental concepts and terminology in the field of material appearance and provide a broad overview of previous research in measurement and modeling. In Chapter 3, we present our quantitative validation of homogeneous reflectance.

In Chapter 4, we describe our method to estimate a generalized microfacet distribution from measured data and demonstrate its effectiveness on complex anisotropic materials. In Chapter 5, we present our work in statistical acquisition and reconstruction of spatially-varying reflectance. In Chapter 6, we describe our image-driven reflectance metric, and a visual interface based on the metric that allows intuitive navigation of appearance models. Finally, we conclude the thesis and provide directions for future work in Chapter 7.

Chapter 2

Background

The study of material appearance is an important subject in computer graphics, as well as physics, computer vision and perceptual science. In computer graphics, accurate reproduction of material appearance is a crucial element in realistic rendering. The study of material appearance can be roughly categorized into measurement, representation, modeling and rendering. To make our discussion more concrete, we will first introduce the common notations and definitions in radiometry, as well as important reflectance representations such as the Bidirectional Reflectance Distribution Function (BRDF) and the Bidirectional Texture Function (BTF). We will then provide a comparison of these different representations and discuss their limitations. Next, we will review the different measurement techniques in previous research. We will then discuss previous work on statistical studies of reflectance and texture. In the final section, we will briefly describe previous work on human perception of materials and their implications on graphics techniques.

2.1 Definitions and Notations

2.1.1 Radiometry

Radiometry is the field that studies the measurement of electromagnetic radiation. In computer graphics, light transport is mostly simulated at the geometric optics level. Wave phenomena like diffraction and interference are typically ignored (Exceptions include [76,

3, 106]). In this thesis, we also ignore effects of polarization, fluorescence, and phosphorescence.

Radiant Power or Flux

Flux, often denoted as Φ, expresses the amount of energy per unit time. Total emission from light sources is often described in terms of flux. The unit of flux is Watt $[W]$.

Irradiance

Irradiance (E) is defined as flux per unit surface area and is often used to describe the amount of energy striking the surface from a certain direction. The unit of irradiance is Watt per meter squared. $[Wm^{-2}]$

$$E = \frac{d\Phi}{dA} \tag{2.1}$$

Radiance

Radiance is the most fundamental quantity in radiometry. Radiance is measured in flux per unit projected area per unit solid angle. Solid angle is measured in steradians and represent the subtended surface area when the region is projected to the unit sphere. One favorable property of radiance is that it remains constant along a ray through empty space. This makes it a convenient quantity to propagate in raytracing algorithms. The unit of radiance is Watt per steradian per meter squared. $[Wsr^{-1}m^{-2}]$

$$L = \frac{d^2\Phi}{d\omega dA^{\perp}} = \frac{d^2\Phi}{d\omega dA\cos\theta} \tag{2.2}$$

2.1.2 Reflectance

The following descriptions of the reflectance quantities mostly follow the convention of Nicodemus et al. [81]. Similar descriptions can be found in [33, 86].

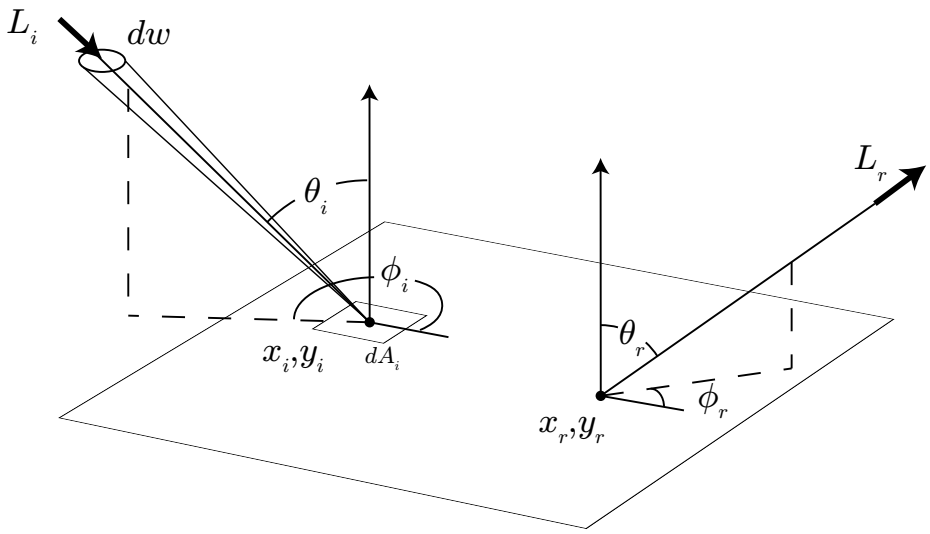

Figure 2-1: Geometry of surface reflection

Bidirectional Scattering-Surface Reflectance-Distribution Function

The *bidirectional scattering-surface reflectance-distribution function* (BSSRDF) describes the relationship between incoming irradiance and outgoing radiance on a general surface. In Figure 2-1 we show a surface patch approximated by a plane. Let the incident radiance from direction (θ_i, ϕ_i), within the element of solid angle $d\omega_i$, be L_i. The portion of incident flux, which strikes an element of area dA_i centered at point (x_i, y_i), is denoted by $d\Phi_i$, which equals $L_i \cdot \cos\theta_i \cdot d\omega_i \cdot dA$. This incident flux is then reflected and scattered before leaving the surface. Due to multiple (subsurface) scattering, the reflected radiance can possibly leave the surface at any location. The scattered radiance in the direction (θ_r, ϕ_r) at a certain location (x_r, y_r) which comes from $d\Phi_i$ is termed dL_r. Even though the exact form of light transport is unspecified, dL_r and $d\Phi_r$ should be linearly related due to the linear nature of reflections. We have:

$$dL_r = S \cdot d\Phi_i = S \cdot L_i \cdot \cos\theta_i \cdot d\omega_i \cdot dA \tag{2.3}$$

The factor S depends on the location of incoming and outgoing locations, as well as the directions of the incoming and outgoing rays. Thus, S is a 8-dimensional quantity:

$$S = S(\theta_i, \phi_i, x_i, y_i; \theta_o, \phi_o, x_o, y_o) \tag{2.4}$$

This quantity S is called the *bidirectional scattering-surface reflectance-distribution function* (BSSRDF). Its unit is per steradian per meter squared [$sr^{-1}m^{-2}$]. Given S and a complete description of incoming radiance from all directions, we can compute the outgoing radiance on every point of the surface by integrating Equation 2.3 over the incident locations and incident directions.

The BSSRDF is a very general property of a surface and captures all kinds of light transport, including subsurface scattering, in a black-box manner. However, the high dimensionality of the BSSRDF makes it very difficult to measure and use. As a result, it is often reduced to simpler representations by imposing certain assumptions or restrictions. We will now discuss two of these cases.

Bidirectional Reflectance Distribution Function

If we assume that the material is uniform, the BSSRDF S and the differential outgoing radiance dL_r should both be independent of (x_r, y_r). Without loss of generality, we can set (x_r, y_r) equal to $(0,0)$. Now, we assume the entire surface is irradiated by radiance $L_i(\theta_i, \phi_i)$, from direction (θ_i, ϕ_i) over the solid angle element $d\omega_i$. We can integrate Equation 2.3 over the incident locations:

$$\begin{aligned} dL_r(\theta_i, \phi_i; \theta_r, \phi_r, 0, 0) &= \int_{A_i} S(\theta_i, \phi_i, x_i, y_i; \theta_r, \phi_r, 0, 0) \cdot L_i \cdot \cos\theta_i \cdot d\omega_i \cdot dA_i \\ &= L_i \cdot \cos\theta_i \cdot d\omega_i \cdot \int_{A_i} S(\theta_i, \phi_i, x_i, y_i; \theta_r, \phi_r, 0, 0) dA_i \end{aligned} \tag{2.5}$$

Nicodemus et al. [81] define the bidirectional reflectance distribution function (BRDF) as:

$$\rho_1(\theta_i, \phi_i; \theta_r, \phi_r) = \int_{A_i} S(\theta_i, \phi_i, x_i, y_i; \theta_r, \phi_r, 0, 0) dA_i \tag{2.6}$$

In this formulation, the BRDF sums up all the scattering contribution over the entire area. Substituting into Equation 2.5, we have:

$$\rho_1(\theta_i, \phi_i; \theta_r, \phi_r) = \frac{dL_r(\theta_i, \phi_i; \theta_r, \phi_r, 0, 0)}{L_i \cdot \cos\theta_i \cdot d\omega_i} \tag{2.7}$$

Intuitively, the BRDF relates the outgoing radiance at a particular location to the incoming irradiance on a nearby flat surface patch. Given the incoming irradiance over the full hemisphere, the BRDF fully specifies the outgoing radiance in all directions.

However, the BRDF is commonly defined in a different way. Since the intensity of scattered rays falls off very quickly for many materials, one way to simplify the BSSRDF is to completely ignore contributions from the neighborhood. In this case, the BRDF can be seen as a slice of the BSSRDF:

$$\rho_2(\theta_i, \phi_i; \theta_r, \phi_r) = S(\theta_i, \phi_i, 0, 0; \theta_r, \phi_r; 0, 0) \tag{2.8}$$

While the second definition is more commonly assumed, we should note that with most measurement devices, the measured BRDFs correspond more closely to the first definition. This is because, in most measurement devices, the entire sample is lit by the light source. To accurately measure a BSSRDF slice, we would need the ability to illuminate a narrow area of the sample (e.g. using a laser beam).

Bidirectional Texture Function

Dana et al. [23, 24] introduced the term *bidirectional texture function* (BTF) to represent spatially-varying reflectance. It is a 6D quantity $R(\theta_i, \phi_i; \theta_o, \phi_o, x, y)$. Again it can be defined either as the BSSRDF integrated over the incident locations, or simply a slice of the BSSRDF:

$$R_1(\theta_i, \phi_i; \theta_r, \phi_r, x, y) = \int_{A_i} S(\theta_i, \phi_i, x_i, y_i; \theta_r, \phi_r, x, y) dA_i \tag{2.9}$$

$$R_2(\theta_i, \phi_i; \theta_r, \phi_r, x, y) = S(\theta_i, \phi_i, x_i = x, y_i = y; \theta_o, \phi_o, x_o = x, y_o = y) \tag{2.10}$$

With the BTF, nonlocal subsurface scattering effects are ignored or pre-integrated. It encodes all other effects such as shadowing, masking and multiple scattering.

2.2 Comparison of Reflectance Representations

Among the different representations of reflectance described in the previous section, it is natural to ask which one to use when we want to model the appearance of a particular surface. The representations can be ordered in an ascending order of complexity: BRDF, BTF, BSSRDF. In addition, at a even higher complexity, we can model the surface with the actual micro-geometry at a fine resolution. Assuming we have an ideal light transport simulator, the micro-geometry representation can be readily converted to the BSSRDF or its derived quantities. For example, Westin et al. [119] use Monte Carlo techniques to approximate BRDFs from micro-geometry. Becker and Max [5] describe a method to convert between displacement map, bump map, and BRDF such that transition between the three representations is smooth when they are used at the same time.

If we assume that the explicit micro-geometry of a surface is available through modeling or measurement, it may seem reasonable to always use this representation directly. However, there are a number of reasons that this is often not a good solution. First, if the scale of geometric variations is small enough with respect to our rendering resolution, using the full micro-geometry is simply overkill. In addition, we would need to pay special attention to avoid aliasing due to undersampling. Second, the perfect conversion from micro-geometry to BSSRDF is only possible with a ideal light transport simulator, which would need to simulate all kinds of physical phenomena present. In practice, the conversion is only an approximation. On the contrary, by measuring the less complex representations such as BRDF or BTF directly, all the physical interactions can be treated in a black-box manner. In a setting where rendering time is limited (e.g. 3D games), it is typically impractical to simulate global illumination effects. In these cases, BRDF and BTF can provide realism that is otherwise not achievable, since the material level scattering is pre-integrated. Finally, it is a strong assumption that we can obtain the micro-geometry of real materials as the three-dimensional measurement at that scale is generally difficult.

2.2.1 Limitations

When employing the various reflectance quantities, it is important to understand their limitations. First of all, the BSSRDF and its derived quantities are all defined on a flat geometry and do not provide accurate silhouettes. This is particularly visible when the BTF is applied due to the abrupt transition at the silhouettes.

The second limitation is concerned with surface curvature. Since the BRDF is defined at a confined location, it does not depend on the curvature of the surface. However, BSSRDF and BTF are typically defined and measured on flat surfaces. When these quantities are applied to a parameterized surface with non-zero curvature, the reflectance can change in a non-trivial manner. However, the curvature assumption is seldom mentioned in the literature, and BTFs are often applied to surfaces without such considerations. Notable exceptions are the works by Wang et al. [116] and Heidrich et al. [44]. In the former work, the view-dependent displacement map is pre-computed for various surface curvatures. In the latter case, global illumination effects are precomputed for height fields on a flat geometry and changes due to surface curvature are approximated when the height field is applied to curved surfaces.

2.3 Reflectance Measurement

2.3.1 BRDF Measurement

Traditionally, the BRDF is measured using gonio-reflectometers [82, 14]. A gonio-reflectometer consists of a light source and a detector, both of which can be positioned at any specified directions by motors or robotic arms. A flat sample of the measured material is placed in the center. Individual BRDF samples can be measured by positioning the light source and the detector from the desired directions. Since the detector is only responsible for measuring one sample at a time, sophisticated spectro-radiometer with high dynamic range and high spectral resolution can be employed. However, since only one value is recorded at a time, it is very time-consuming to capture a 3D or 4D BRDF. As a result, it is impractical to use a gonio-reflectometer to measure BRDFs at high angular resolution.

Ward [118] accelerates the BRDF measurement process by measuring multiple samples at the same time. He uses a hemispherical half-silvered mirror with a flat sample in the center. The outgoing radiance from the flat sample to all directions can be observed at the same time, using a CCD camera with a fish-eye lens. With this setup, the measurement time is much reduced since only the light direction needs to be varied. However, quality from this setup is relatively low due to the multiple optical elements, and angular resolution near grazing angle is limited. More recently, Dana and Wang [19] have proposed a similar setup to measure spatially varying reflectance, using parabolic mirrors and translational stages.

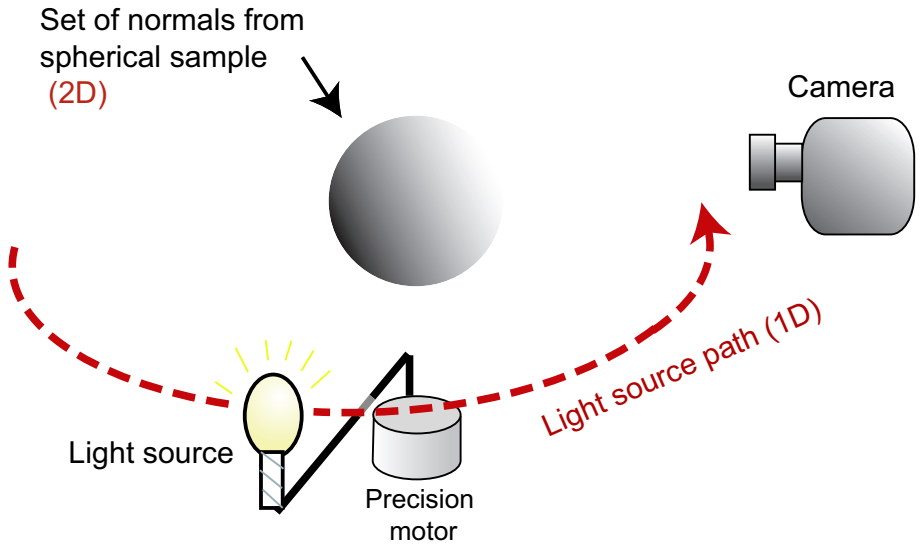

Figure 2-2: Image-based measurement of isotropic BRDFs [66, 68]. The surface normals of the spherical sample provide two degrees of freedom, the rotating light provides the third for the 3D isotropic BRDF.

Instead of using mirrors, Marschner et al. [66, 65] use a curved material sample (e.g. sphere) such that a set of varying surface normals are visible in a single image (Figure 2-2). Since the material is homogeneous, each surface point with a different normal corresponds to a BRDF sample with a different viewing direction. With a fixed camera and an orbiting light source, this setup can measure BRDFs at a very high resolution efficiently. High dynamic range imaging technique using multiple exposures are often necessary to capture the high dynamic range of specular BRDFs. Although the setup is limited to isotropic BRDFs, the technique is efficient and robust. Matusik et al. [68] have used a similar setup to measure over a hundred BRDFs. This data set is now publicly available and is used for

our quantitative validation of isotropic BRDFs.

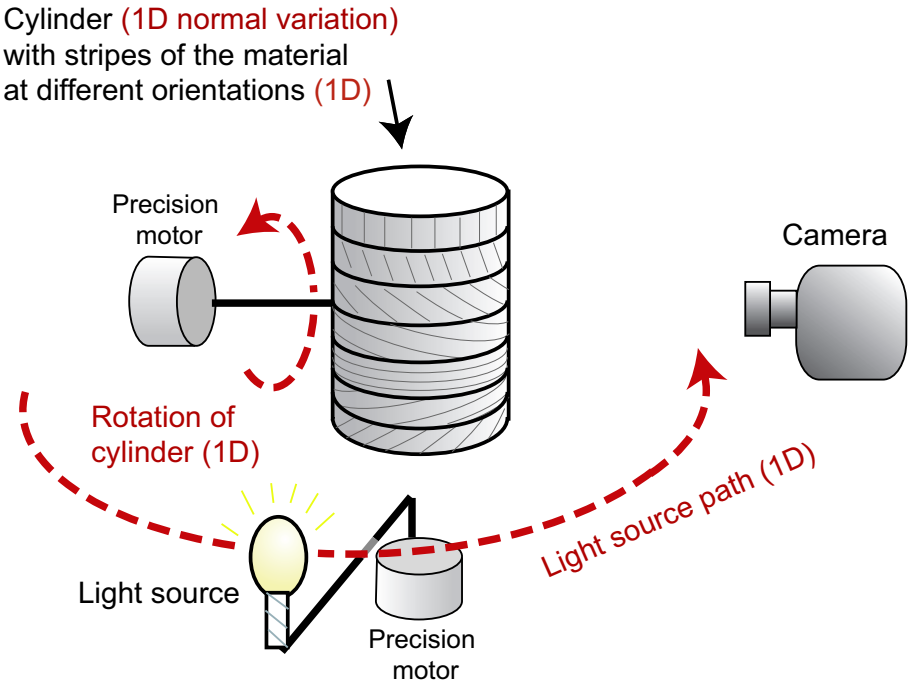

Figure 2-3: Our acquisition setup for anisotropic materials. The cylinder tilt, surface normal variation, light position and strip orientation each contributes to one degree of freedom for acquiring the 4D BRDF.

We have extended Matusik's image-based setup to handle anisotropic BRDFs (Figure 2-3). Our extension is similar to the method proposed by Lu et al. [61, 62], who measure the BRDFs of velvet and shot fabric with a manual setup. Our setup is automated with precision motors. A cylinder is used instead of a sphere as the measurement target so that we can paste material samples on top without distortion. We compensate for the lost degree of freedom in normal variations by mounting the cylinder on a precision motor and performing measurements at different tilt angles. In order to account for anisotropy we cut multiple strips from a planar sample of the material at different orientations. The strips are then pasted onto the cylinder target. Together with the degree of freedom for the light position, we are able to acquire the full 4D BRDF using a large set of two dimensional images. Figure 4-1 shows an example input image of a acquired velvet sample. More details about our measurement will be described in Section 4.3.1.

2.3.2 BTF Measurement

The first BTF measurement device was built by Dana et al. [24]. In this device, a robot arm is used to orient the texture sample at arbitrary orientations, and the camera and light orbit around the sample. 205 combinations of light and view directions are sampled for each material, and more than 60 materials have been measured and published [17]. Due to the sparse sampling, it is not practical to use the measured data for rendering directly. More recently, researchers have built similar setups and provided measurements at higher angular resolutions [97, 51, 38]. Analogous to the gonio-reflectometers for BRDFs, only one *sample* is measured at a time for a particular lighting and viewing directions, even though each sample is a texture.

Müller et al. [77] present a setup to capture BTFs more efficiently by simultaneously capturing multiple views using multiple cameras. They built a hemispherical gantry consisting of 151 consumer digital cameras. The on-camera flashes serve as light sources. By synchronizing the cameras, $151 \times 151 = 22801$ images can be captured in 151 time steps, and the authors report a measurement time of about 40 minutes. While this is a big improvement in terms of measurement time, the setup is large and costly.

Han and Perlin [41] introduced a measurement setup based on a kaleidoscope, which allows viewing a sample from multiple angles simultaneously through multiple reflections. Illumination is provided by a projector pointing into the kaleidoscope. By selectively illuminating a small group of pixels, the light direction can be controlled. Since there is no moving part in this setup, measurement is very fast. However, the equipment is difficult to build and calibrate. In addition, due to multiple reflections in the optical path, the result quality tends to be rather low. Dana and Wang [19] proposed a setup based on a parabolic mirror. While their setup can provide higher quality measurements than the kaleidoscope setup, they can only capture a single spatial location at a time. As a result, it does not offer any acceleration compared to the gonio-reflectometer-like approaches.

Due to the nature of spatial variations for BTFs, multiplexing approaches based on curved geometry cannot be directly applied to BTF measurement. Suppose we have a curved sample of the material and the spatial variation is periodic. To map each measured

pixel to the domain of the BTF, we will need to know the texture coordinate the surface point corresponds to, in addition to its normal. Indeed, real textured objects are seldom perfectly periodic, so the correct correspondence can be ill-defined. In Chapter 5, we will present our BTF representation, which avoids the need of precise correspondence. As a result, we are able to build a new measurement setup that allows capturing of multiple texture samples at the same time, even if they are not perfectly identical.

2.3.3 Reflectance Fields

The BSSRDF is defined on a planar geometry (Section 2.1.2). However, its definition can be easily generalized to a general parameterized surface. In this case the BSSRDF is commonly known as a *reflectance field*. The full reflectance field is a 8D quantity like the BSSRDF, and it is often reduced to lower dimensions by fixing some of the dimensions.

Single View

Many works in reflectance field measurement assume distant illumination: the entire object is lit by parallel rays from a certain direction. This assumption removes the dependency on incident locations. By further restricting the measurement to a single view, the reflectance field can be reduced to 4D, depending only on the surface location and direction of incident light. Debevec et al. [26] built a dome-shaped setup which records the reflectance field of a human face under varying light directions and a fixed view. Similarly, Malzbender et al. [63] have built a hemispherical gantry with 50 lights to capture the single-view reflectance field of small objects. Masselus et al. [67] present a method to lift the distant illumination assumption and acquire a fixed-view 6D reflectance field using a local lighting basis provided by a projector.

Multiple Views

The following works all assume distant lighting. Lensch et al. [56] capture a small number of images of an object under varying light and viewing directions. By assuming that the object can be represented with a small number of distinct BRDFs, they cluster similar

reflectance from different surface locations to improve angular resolution. The clustered BRDFs are then fit to parametric models. Matusik et al. [69, 70] present a system which acquires the reflectance field from multiple viewpoints, using turntables to rotate the object and lights. A relatively sparse sampling of the 6D reflectance field can be captured, but the acquisition process take many hours. Weyrich et al. [121] built a dome with 15 cameras and more than a hundred lights for measuring face reflectance, similar to Debevec's Light Stage. The measurement time is about 15 seconds, making it possible to record live subjects. With this setup, 15 views are captured, providing sparse slices of the 6D reflectance field.

Subsurface Scattering

Jensen et al. [45] propose a model for rendering the BSSRDF by using a dipole approximation to a diffusion model. Although the method is limited to homogeneous materials, it has become very popular in the field due to its efficient evaluation and the tremendous improvement in rendering quality. They also introduce a simple method to estimate the diffusion model parameter from a single image. More recently, Goesele et al. [39] present a measurement setup for translucent inhomogeneous objects by using a laser beam as the illuminant. To limit the dimensionality of measurement, their work is limited to diffuse surface reflection. Tong et al. [109] measure subsurface scattering in quasi-homogeneous material by assuming only short distance subsurface scattering is inhomogeneous.

2.4 BRDF Representations

2.4.1 Analytical Models

Many analytical BRDF models have been proposed in the past few decades. They can be roughly divided into two classes: empirical models seek to emulate the characteristics of material reflection qualitatively [6, 118, 52], while physically-based models derive the quantities from first principles of physics [13, 42]. A number of the more well-known models will be listed here. We change the notations for some of the models from their original form for more consistent notation. For further information on the models, the

$\rho(L,V)$	BRDF evaluated at (L,V)
N	normal
V, (v_x, v_y, v_z)	outgoing vector (view)
L, (l_x, l_y, l_z)	incoming vector (light)
R	mirror reflection of L
H	half-way vector between L and V
(θ_h, ϕ_h)	spherical coordinates of the half vector H
p_s	specular lobe scaling factor
p_d	diffuse lobe scaling factor

Table 2.1: Common notations for BRDF models.

reader is referred to the original papers.

In Table 2.4.1 we list the common notations we use for all the BRDF models listed. BRDF is always defined in a *local* coordinate system - defined by the normal, the tangent and the binormal. A BRDF model is typically composed of two components, the *diffuse* term and the *specular* term. The diffuse term generally describes the aggregate effects of multiple reflections. As a result, it is mostly independent of the light and view direction. In the ideal case it is Lambertian, i.e. constant for all direction pairs. The specular term generally describes the directional component of reflection. For common materials which are not entirely diffuse, the overall reflection for a particular incident ray has a maximum near the mirror direction. The specular term specifies the location of this peak, together with the falloff surrounding it. For the BRDF models listed here, many of them are focused on the specular term, and either assume a Lambertian diffuse term or do not specify it explicitly. For the Lambertian case, the diffuse term is simply specified by the diffuse color, a separate p_d for each channel. For the specular term, it is generally a nonlinear function of several parameters, scaled by p_s for each channel. The nonlinear parameters are typically the same for all channels. This is not a requirement for physical plausibility, but different parameters per channel will lead to unnatural color shift in most cases.

There are two properties that a BRDF model need to satisfy in order to be physically plausible. First, a BRDF has to be reciprocal: the value of the BRDF should be the same when the incoming and outgoing direction is exchanged [114]. Second, a BRDF has to be

energy-conserving, meaning the outgoing energy must be less than or equal to the incoming energy. With the presence of the scalar factors p_d and p_s, this can always be achieved for any model. Under a more restrictive definition, a BRDF model is considered energy-conserving if given $p_d + p_s = 1$, the BRDF is energy-conserving for all possible parameters. Finally, in cases where the incoming and outgoing energy are strictly equal, we call the model strongly energy-conserving.

Phong Model

The Phong model [87] is the most well-known shading model in computer graphics. The core idea is the empirical observation that for specular/glossy materials, the reflection is strongest along the mirror direction R with respect to the light L, with a smooth falloff surrounding R. To describe this falloff, Phong suggests to use the cosine between V and R, raised to a certain power n. The original formulation is:

$$\rho(L,V) = p_d + p_s \cdot \frac{(V \cdot R)^n}{N \cdot L} \tag{2.11}$$

The original Phong model is not physically plausible - it does not satisfy energy conservation or reciprocity. A simple modification that ensure these properties is as follows:

$$\rho(L,V) = \frac{p_d}{\pi} + p_s \cdot \frac{n+2}{2\pi}(V \cdot R)^n \tag{2.12}$$

There is a variant of the Phong model proposed by Blinn [6], based on the halfway vector H and surface normal N:

$$\rho(L,V) = \frac{p_d}{\pi} + p_s \cdot \frac{n+2}{2\pi}(H \cdot N)^n \tag{2.13}$$

This variant is also the default shading model implemented in OpenGL. In Chapter 3 we will demonstrate the subtle but important difference between the Blinn version and the original Phong model.

Ward Model

Ward [118] proposes an empirical BRDF model that accounts for anisotropy. Instead of the cosine lobe used by Phong, the Ward model is based on the elliptical Gaussian. In addition, the model is defined with respect to the half-vector H. It has two parameters, α and β, that define the spread of the lobe in the two directions. For isotropic materials $\alpha = \beta$ and the Ward model has a single parameter. The formula is as follows:

$$\rho(L,V) = \frac{p_d}{\pi} + p_s \cdot \frac{1}{\sqrt{(N \cdot L)(N \cdot V)}} \cdot \frac{\exp[-\tan^2 \theta_h (\cos^2 \phi_h / \alpha^2 + \sin^2 \phi_h / \beta^2)]}{4\pi\alpha\beta} \tag{2.14}$$

The model is reciprocal, a fact immediate from the equation shown. However, recently Dür [30] shows that the model is not strongly energy-conserving. He proposes a modification to the model to make it strongly energy-conserving:

$$\rho(L,V) = \frac{p_d}{\pi} + p_s \cdot \frac{1}{(N \cdot L)(N \cdot V)} \cdot \frac{\exp[-\tan^2 \theta_h (\cos^2 \phi_h / \alpha^2 + \sin^2 \phi_h / \beta^2)]}{4\pi\alpha\beta} \tag{2.15}$$

Cook-Torrance Model

The Cook-Torrance model [13] is a physically-based model which adopts the microfacet formulation from Torrance and Sparrow [110]. The basic assumption is that the surface is composed of many tiny smooth mirrors. As a result, the reflectance at a particular light/view combination is a function of the number of aligned microfacets, the amount of shadowing and masking between the microfacets, and the reflectance of the ideal mirrors. The formula is as follows:

$$\rho(L,V) = \frac{p_d}{\pi} + \frac{p_s}{\pi} \frac{DGF}{(N \cdot L)(N \cdot V)} \tag{2.16}$$

where

$$G = min\{1, \frac{2(N\cdot H)(N\cdot V)}{(V\cdot H)}, \frac{2(N\cdot H)(N\cdot L)}{(V\cdot H)}\} \tag{2.17}$$

$$D = \frac{1}{m^2\cos^4\theta_h} e^{-[(tan\theta_h)/m]^2} \tag{2.18}$$

G is the geometric attenuation term which approximates the shadowing and masking effects, while D is the distribution of the microfacets. F is the Fresnel term which governs the reflection for smooth mirrors.

Lafortune

While the work by Lafortune et al. [52] is presented as a non-linear basis for reflectance representation, it is commonly considered as a phenomenological BRDF model as well. The non-linear basis can be seen as a generalized cosine lobe:

$$B = [C_x l_x v_x + C_y l_y v_y + C_z l_z v_z]^n \tag{2.19}$$

This expression is equivalent (within a constant factor) to the cosine lobe used by Phong when $C_x = C_y = -1, C_z = 1$. The authors suggest the use of multiple lobes to fit measurement or a physically-based BRDF model. Since its introduction, the Lafortune model has become one of the most popular models because of its generality and computational efficiency. However, as we will show in Chapter 3, the model produces unnatural highlights which limits its plausibility for a significant subclass of materials.

Oren-Nayar

Oren and Nayar [83] developed a model that improve upon the common Lambertian diffuse model. Similar to Cook-Torrance, it is also based on a collection of microfacets, but these microfacets are modeled as perfectly diffuse surfaces. However, due to complex shadowing effects, the overall BRDF is not Lambertian. They show that their model matches diffuse objects in the real world much closer than Lambertian. The reader is referred to their paper for the full formulation.

He-Torrance-Sillion-Greenberg Model

The model proposed by He et al. [42] is considered the most comprehensive analytical BRDF model. Unlike most other works, the He model takes polarization and other wave phenomenon into account. The formula for the model is very complex and the reader is referred to the paper for further information. In this thesis, we use the implementation of the unpolarized version of the He model by Rusinkiewicz [95].

2.4.2 Generalized microfacet model

While physically-based models like Cook-Torrance and He are derived based on a microfacet distribution, they typically put strong assumptions on the shape of this distribution, e.g. Gaussian. These distributions are typically controlled by one or two parameters only. Ashikhmin et al. [1] introduced a method to create a BRDF from an arbitrary distribution of microfacets.

The expression for the Ashikhmin et al. model is as follows:

$$\rho(L,V) = \frac{p(H)\langle H \cdot N \rangle F(L \cdot H)}{4g(L)g(V)} \tag{2.20}$$

where $p(H)$ is the microfacet distribution; $g(\cdot)$ is the shadowing/masking term; $F(L \cdot H)$ is the Fresnel term for the mirrors aligned with H; and $\langle \cdot \rangle$ denotes the average over the distribution p. The model will be explained in more detail in Chapter 4, where we also present a method to estimate the microfacet distribution from measured data.

2.4.3 Data-driven Representations

To overcome the limitations of analytical models, a data-driven approach considers the BRDF as a fully general 4D function. By exhaustive sampling of the 4D domain, it is theoretically possible to represent any BRDF. However, in practice, this representation is limited by the measurement resolution and noise. In addition, it requires significant storage and is difficult to edit.

The full 4D function can also be compressed into more compact representations. For

example, Westin et al. [119] propose the use of spherical harmonics to store anisotropic BRDF data. Theoretically, spherical harmonics can encode any function on the hemisphere, but many coefficients will be necessary for functions with high angular frequencies. As a result, representations based on spherical harmonics are only efficient for relatively diffuse BRDFs. Schroeder and Sweldens [100] propose spherical wavelets to represent a 2D slice of the BRDF by fixing one of the direction. Lalonde and Fournier [53] use a wavelet coefficient tree to represent full BRDFs. Other representations based on compression include Zernicke polynomials [50], singular value decomposition[48] and positive matrix factorization [74].

2.5 Spatially-Varying Reflectance

2.5.1 Approximations

A number of techniques have been proposed in the past to approximate spatially-varying reflectance with different levels of simplification. The most primitive example is standard texture mapping, where different surface locations map to different colors. Bump mapping [7] simulates the effect of surface roughness by perturbing surface normals. Horizon mapping [72, 105] enhances bump mapping by handling self shadowing effects. Heidrich et al.[44] further improves microgeometry rendering by efficient interreflection estimation using precomputed visibility. The displacement map [12] is typically not considered a reflectance technique since it changes the underlying geometry. However, recent extensions [116, 117] to the displacement map are applied at shading time. Occlusion and shadowing are precomputed to allow for interactive rendering. While the above techniques offer different trade-offs between computation cost and quality, they all depend on accurate descriptions of the microgeometry which are unavailable and often difficult to acquire from real materials in most cases. As a result, these techniques are often only applied to render synthetic materials.

2.5.2 BTF

The BTF is first introduced by Dana et al. [24] and has gained popularity recently due to advancement in related techniques. It serves the same goal with the approximations described in the previous section, but it is more general due to its black-box nature. It describes the relationship between incoming and outgoing radiance, without prescribing the means of light transport in between. As a result, it is directly measurable and can be applied directly to synthetic scenes. The reader is referred to a recent survey by Müller et al.[79] which provides a comprehensive study of recent work on BTFs. In the remainder of this section, we will summarize some of the important aspects of BTF research.

BTF Compression

Because of the large size of BTF data even at a low sampling (over one gigabyte for 81 views and 81 lights, 256×256 texture patch), there has been a great deal of previous work on compression methods. The methods can be roughly divided into two groups. The first group treats the BTF as a spatially-varying BRDF, and fits analytical BRDF models or their variants per texel. The second group considers BTF as a general 6D function and applies linear basis decomposition for compression.

McAllister et al. [73] fit Lafortune lobes to each texel separately. The compressed data is very compact but the method is limited as a BRDF model is not suitable for the complex shadowing and masking effects typical in a BTF. To better handle the meso-structure, Daubert et al. [25] add an extra multiplicative view-dependent term to the Lafortune lobes. Meseth et al. [75] further improves compression quality at the expense of space by fitting separate Lafortune lobes to the BTF per pixel *per view*.

The other group of methods compresses BTF by basis decomposition or factorization. Matusik et al. [69] compress 6D reflectance field by applying principal component analysis (PCA) on image blocks. Kautz and McCool [48] factorize BTF into product of 2D textures. Koudelka et al. [51] apply principal component analysis (PCA) to the full 6D matrix. Vasilescu and Terzopoulos [113] arrange the BTF into a 3-mode tensor and applied 3-mode SVD. This allows for more flexible compression by reducing view and light dimensions in-

dependently. Compared to PCA, this method leads to a higher root-mean-squared error, but the authors claim that their method provides perceptually more satisfactory results.

BTF Rendering

BTFs are straightforward to be incorporated into both offline and online rendering system. For offline systems, the BTF represented as a collection of textures can be used directly provided there is enough memory to hold the textures. In a real-time system, however, it is typically impossible to store even a single BTF in current graphics hardware. All the compression techniques mentioned in the previous subsection can be used for interactive rendering. The reduced data, either in the form of parameters for spatially-varying BRDFs [73, 25, 75], or coefficients of linear bases [108, 97], are stored in the texture units of the graphics hardware. BTF shading can then be implemented in the programmable pixel shader.

2.6 Statistical Analysis of Reflectance and Texture

In computer vision, the study of material appearance is concerned with recognition and identification of material and geometry rather than image generation. As a result, precise representation of material reflectance is not required. On the other hand, due to the presence of uncertainties and calibration errors, aggregate statistics that summarize the property of material robustly play a more important role. Statistical approaches are employed in computer graphics as well, most notably in the field of texture synthesis, where a large texture is generated from a small sample patch. In this section we will highlight some of the statistical studies of reflectance and texture which form the basis of our hybrid BTF representation (Chapter 5).

2.6.1 Texture Modeling and Synthesis

Statistical characterization of texture appearance is motivated by research on human texture perception, which states that two textures are difficult to discriminate when their statistics

across different scales and orientations are similar. Heeger and Bergen make use of image pyramids of a source texture sample for synthesizing a similar texture [43]. Their method is based on matching marginal statistics, i.e. histogram matching. To create a new texture, a pyramid is first initialized with white noise. A source pyramid is built by decomposing the source texture. The statistics of each subband in the noise pyramid is then matched to the corresponding subband in the source pyramid, through histogram matching of the coefficients. The modified pyramid is collapsed to produce a new intermediate image. By repeating the entire process for several iterations the modified image will have very similar statistics to the source texture in every scale and orientation. While their method is very successful in reproducing textures with a stochastic nature, it fails for highly coherent textures, e.g. a wall of bricks.

DeBonet [9] also uses multi-scale pyramid analysis to perform texture synthesis. In a two-phase process, he first computes joint probability of texture features across multiple resolutions of the input texture. In the second phase, the texture is generated from coarse to fine, obeying the conditional probabilities between different levels.

Portilla and Simoncelli [89] propose a statistical texture model based on first and second order properties of joint wavelet coefficients. This statistical model is used to both analyze and synthesize new texture. It can capture both stochastic and repeated textures very well, but again fails to reproduce highly structured patterns.

2.6.2 Light- and View-Dependent Statistics

Leung and Malik [57] study the occlusion properties of a special class of 3D textures by assuming a surface with perpendicular elements. With some simple analysis, they derive the probability of view-dependent masking and show that the color distribution in the texture changes significantly when the view is changed. A canonical example they cite is the sunflower field, where one can often see much denser yellow at grazing angles.

Pont and Koenderink [88] use a simple micro-facet Lambertian model to predict texture contrast at different lighting and viewing configurations. Ginneken et al. [112] investigate the pixel histograms of a wider class of materials in the CUReT data set [17] as a function

of view and light direction. In particular, in the case of surfaces with uniform albedo, they show that a single texture modulated with histogram matching offers improved material appearance compared to plain texture mapping. Cula and Dana [16, 15] present a method for material recognition using histograms of BTF features described by 3D textons [58, 59].

Dana and Nayar's work [20, 22] is closely related to our reconstruction technique described in Chapter 5. They study BTF under varying light/view configurations and propose analytical histogram and correlation models in the special case of random isotropic surfaces with Lambertian reflectance. They also propose BTF synthesis based on histogram transfers [21]. We will further compare their work to our reconstruction in Chapter 5.

The BTF synthesis work by Liu et al. [60] is based on geometry statistics. They use a sparse set of images to estimate an approximate height field by shape-from-shading, and synthesize new geometry that is statistically similar to the acquired sample. Pixel samples from the input are then copied to the synthesized image based on feature matching to reconstruct the bidirectional appearance.

2.7 Perceptual Study of Material Appearance

Since a human is involved in both designing the virtual scene and judging the final rendering, it is important to consider perceptual aspects of material appearance. In this section, we will highlight the previous work in the field of human perception related to material appearance, and their implications on graphics applications.

2.7.1 Reflectance Perception under Unknown Illumination

Fleming et al. [36, 37] investigate the way humans determine material reflectance. They observe that humans can identify reflectance from a single image reliably even the illumination is unknown, a task which is theoretically under-constrained. Through experiments with human subjects, they show that without contextual cues to illumination, humans can match surface reflectance properties with good precision. The subjects are more accurate in reflectance matching under real world illuminations. The authors suggest that this is due to some statistical regularities of real-world illuminations. Based on this observation, Dror

et al. [28, 27] describe a machine vision system that identifies surface reflectance from a single image, without knowledge of the illumination except that it shares the same statistics of typical real-world illumination.

Our work on reflectance metric is inspired by the results from these works, since our approach relies on the assumption that a rendered object under natural illumination is a *good* representation of the underlying reflectance. The reader is referred to Chapter 6 for the description of our image-driven reflectance metric.

2.7.2 Perceptual Distance in Reflectance Space

In order to leverage the industry's expertise in material specification, Westlund and Meyer [120] apply appearance standards to establish correspondences between measurement scales and parameters of analytic BRDF models. They *measure* instances of the analytical models by rendering BRDFs at several angles specified by the industry appearance standards. This produces a one-to-one mapping between the industry gloss values and the gloss parameters of a few analytical models.

Pellacini et al. [85], inspired by the work on perceptually uniform color spaces (e.g. CIELAB), proposed a perceptual adaptation of Ward's BRDF model [118]. They reparameterize the model based on psychophysical experiments, in which subjects are asked to assign numbers to describe the apparent differences between rendered images of different BRDFs. Next, they apply multidimensional scaling (MDS) techniques to recover the perceptual axes and scaling of the gloss space according to the reported differences. They demonstrate that the reparameterized model is easier to use because it is perceptually uniform. However, it is difficult to extend the same method to more general BRDF models since the higher dimensionality and increase in the number of samples quickly render human experiments impractical.

In Chapter 6 we will describe our computational approach which serves as a replacement for the psychophysical experiments. While our approach is not technically perceptual, we show that it produces results that are consistent with the experiments by Pellacini et al.

Chapter 3

Analysis of Isotropic BRDF Models

The Bidirectional Reflectance Distribution Function (BRDF) fully describes the appearance of an opaque material by specifying the ratio of reflection for all possible view and light directions. In most graphics applications, BRDFs are defined using analytical models. While the use of measured BRDFs can potentially offer better quality, measurement is costly and time-consuming. In addition, raw measured data is often limited in resolution, noisy, and difficult to store or edit. Analytical models, on the other hand, provide a smooth and compact representation that can be controlled with a few parameters. Many analytical models have been introduced and they offer different balance of efficiency, physical validity and complexity. While new models are occasionally validated against physical measurements, the number of materials used in the validation is typically small and often confined to the category in which the new model is primarily focused. In this chapter, we present a quantitative evaluation of analytical BRDF models based on their ability to represent real materials. In our study we compare the performance of seven analytical BRDF models on a data set of 100 isotropic[1] BRDFs acquired at high resolution and covering a wide range of materials. In the evaluation process we have gained important insights about various formulations in BRDF modeling. We have also acquired a few anisotropic materials with our new measurement setup, these materials will be studied in Chapter 4.

The work described in this chapter has been published previously in [80]. In the fol-

[1] A BRDF is *isotropic* if it is invariant with respect to rotation of the material patch around the surface normal.

lowing sections, we start by describing the data set used in our study. Next we describe the pre-processing of the raw BRDF data. We then explain our method of fitting BRDF models to the measured data, followed by the fitting result and detailed analysis. In particular we study two common formulations of the BRDF in depth and illustrate the important differences between them. We conclude the chapter with a summary of our findings.

3.1 Data Set

We use a data set of 100 BRDFs measured with high precision acquired by Matusik et al. [68]. The data set includes metals, plastics, painted surfaces, and fabrics. We acknowledge that our data set does not cover the full range of materials, and effects such as retroreflection might be underrepresented. In fact, our BRDFs have no valid samples within 3 degrees from the retroreflection direction since the light source occludes the detector. The BRDFs are acquired by a setup similar to Marschner et al. [65], described in Section 2.3. This image-based method allows for high angular resolution measurement since many radiance samples can be recorded in one single image. However, this setup is limited to isotropic BRDFs and requires spherical samples.

3.2 Data Processing

The isotropic data from Matusik et al. was originally represented using the Rusinkiewicz parameterization [96] based on the half-vector and tabulated non-uniformly to optimize for rendering quality. To avoid bias due to the choice of parameterizations and non-uniform sampling, we reprocess the isotropic data directly from the raw acquired images. Each pixel in the input images with non-zero irradiance is considered a BRDF sample. In addition, we notice that for some highly specular materials, strong flare is observed due to internal reflection in the lens system. The flare generally takes the shape of the aperture, and the size of the flare is correlated with the strength of the reflection. To remove the flare, we take a semi-automatic approach: for each material we manually specify the radius of the strongest flare (typically an image with grazing angle illumination), and the software scales

the flare size for the remaining images based on the intensity of the brightest pixel. The pixels covered by the flare are discarded. We group all the valid samples into bins uniformly spaced in the 3D domain of an isotropic BRDF defined by $\theta_{in}, \theta_{out}, \phi_{diff}$. We set the BRDF values to the median value in each bin. The domain is divided into intervals of 1 degrees, and generally about 85% of the bins are non-empty. The processed data can be used directly for fitting by ignoring the empty bins.

3.3 Methodology

3.3.1 BRDF models

We choose seven analytical isotropic models for our analysis: Ward [118], Blinn-Phong[6], Cook-Torrance[13], Lafortune et al.[52], Ashikhmin-Shirley [2], He et al. [42], and a variant of the Ward model proposed by Duer [30], which fixes the normalization of the specular lobe. The formulas and brief explanations of the models are listed in Section 2.4, and the reader is referred to the original papers for further details. These models differ in their degrees of freedom and goals, ranging from complex physically-based models to ad-hoc empirical models. Comparing them is in a sense "unfair", as we only focus on the numerical ability to fit measured data. Other important qualities such as computation speed and ease of sampling are ignored.

To evaluate the performance of the analytical models, we first restrict them to have only one specular lobe and the diffuse contribution to be Lambertian. The shape of the specular lobe is enforced to be identical across the color channels, and hence all the BRDF models can be expressed with six parameters for the diffuse and specular color, and a variable number of parameters for the specular lobe:

$$M = (\rho_{dr}, \rho_{dg}, \rho_{db}) diffuse + (\rho_{sr}, \rho_{sg}, \rho_{sb}) specular(p_0, p_1, ..., p_n) \tag{3.1}$$

3.3.2 Fitting BRDFs

A core part of our analysis is to fit measured BRDFs to analytical models. We apply a constrained nonlinear optimization technique based on Sequential Quadratic Programming(SQP)[2] over the specular lobe parameters $p_0, p_1, ..., p_n$ to minimize the error metric described in the next paragraph. The diffuse and specular color (6 parameters) are computed analytically as a sub-procedure based on linear least squares. Like any nonlinear optimization, the quality of the fit is dependent on a good initial guess. To ensure the optimization converges to the global minimum, we visually inspect the fitting quality of the result and if necessary, restart the optimization from a different set of initial guesses.

Our fitting procedure does not take into account knowledge about the meaning of physically-based parameters for models such as Cook-Torrance or He et al. While this might result in physically-inconsistent values, it provides the best approximation of the data.

3.3.3 Error Metric

The objective function of the optimization is defined as the mean squared error between the measured BRDF R, and the target model M given parameter vector $\mathbf{p}$:

$$E(\mathbf{p}) = \sqrt{\frac{\sum w[R(\omega_i, \omega_o)cos\theta_i - M(\omega_i, \omega_o; \mathbf{p})cos\theta_i]^2}{\sum w}} \tag{3.2}$$

The sum is over the non-empty bins of R, ω_i and ω_o are the incident and outgoing directions respectively, and θ_i is the elevation angle of the incident direction. The weight w is the solid angle correction term. We ignore data with incident or outgoing angle larger than 80 degrees as measurements close to extreme grazing angles are in general unreliable. The metric is essentially a L^2 metric over the hemispheres, with part of the domain removed where there is a lack of reliable measurements. The form of this metric allows us to separate the linear parameters (diffuse/specular color) from the main optimization, which greatly improves the efficiency and stability of the fitting routine.

[2]The routine fmincon in MATLAB is employed.

The $\cos\theta_i$ term in the metric weighs the BRDF under the assumption of uniform incoming radiance, and in practice we observe that the optimal fit with the cosine factor produces renderings which are visually superior to the fit without it. Lafortune et al.[52] define the objective function for fitting as the BRDF difference multiplied by the cosines of both the incident and outgoing elevation angles. We find that the additional cosine factor gives results which are visually similar to ours.

We have experimented with other fitting metrics based on the logarithm or cubic root of the BRDF. As the logarithm behaves badly near zero, it imposes an arbitrary choice of scaling the BRDF, and in practice it is difficult to make a single choice that consistently performs well for all materials. For the metric based on the cubic-root, the best-fit BRDF often result in renderings with highlights that are too blurry compared to the measured data.

It is clear that choosing an appropriate metric for fitting BRDF is difficult, and our choice of the direct L^2 metric is based on two reasons: i) it provides plausible fits in the rendered images for most materials in our data set, ii) the linear nature of the metric makes the minimization lower dimensional and more stable. We should emphasize that the best fit according to our metric might not always correspond to the best visual match, which is highly dependent on scene geometry and illumination. Developing a perceptually-based metric that is simple and stable for optimization would be interesting direction for future work.

3.4 Analysis

3.4.1 Fitting quality of the models

We fit the isotropic data set of 100 materials to the seven models listed in the previous section. The fitting errors of 5 of the models are plotted in Figure 3-1. The errors are normalized by the maximum albedo of each BRDF and plotted in logarithmic scale. The albedo is computed by integrating the BRDF over the outgoing hemisphere, given an incident direction. We compute the albedos of the BRDF over the range of 0 to 80 degrees incidence, and we use the maximum value to normalize the errors plotted in Figure 3-1.

This normalization allows us to compare the residual errors for different materials while discounting the factor of relative scale. In certain cases, the absolute error can be more meaningful (e.g. darker materials are easier to approximate). The readers should refer to our website[3] for complete fitting results and instructions for raw data access.

We sort the materials according to the errors of the Lafortune fits. The errors for the Ashikhmin-Shirley and Ward-Duer models are not shown in the figure: the Ashikhmin-Shirley errors are very close to Cook-Torrance in most cases, and the Ward-Duer errors are very close to the original Ward, though consistently lower.

Figure 3-1 shows that He, Cook-Torrance and Ashikhmin-Shirley typically have the lowest errors. One factor of their good performance is probably their explicit modeling of the Fresnel effect. In the case of the He model, the number of degrees of freedom is also bigger. In contrast Blinn-Phong and Ward consistently yield higher errors. The errors for the Lafortune model lie in between for a majority of the materials. To the left end of the plot are mostly diffuse or mildly glossy materials like fabrics, rubber and paints, while materials close to the right end are mostly smooth metals and highly specular plastics. The normalized errors span several order of magnitudes as the shape of the BRDF is vastly different for the wide spectrum of materials.

One example fit, to the measurement "green-metallic-paint" is shown in Figure 3-2 with spheres rendered under an environment map. The polar plots in the incidence plane for three models are also shown in Figure 3-3. The relative quality of the fits from the different models for this material is typical of the data set. The Ward, Ward-Duer and Blinn-Phong images are all noticeably different from the original data. The highlights near the center of the sphere are significantly brighter than the original, and near grazing angle the highlights are significantly less pronounced. Of the three models, Ward-Duer has the strongest highlight near grazing angle, but it is still noticeably dimmer than the measured data. The images from the Ashikhmin-Shirley, Cook-Torrance and He fits are very similar to each other and are mostly faithful to the direct rendering. The Lafortune fit shows an abrupt increase in intensity near grazing angle, however values closer to normal incidence are too low. In addition, the highlight is excessively blurred along the circumference of the

[3] http://groups.csail.mit.edu/graphics/brdf/

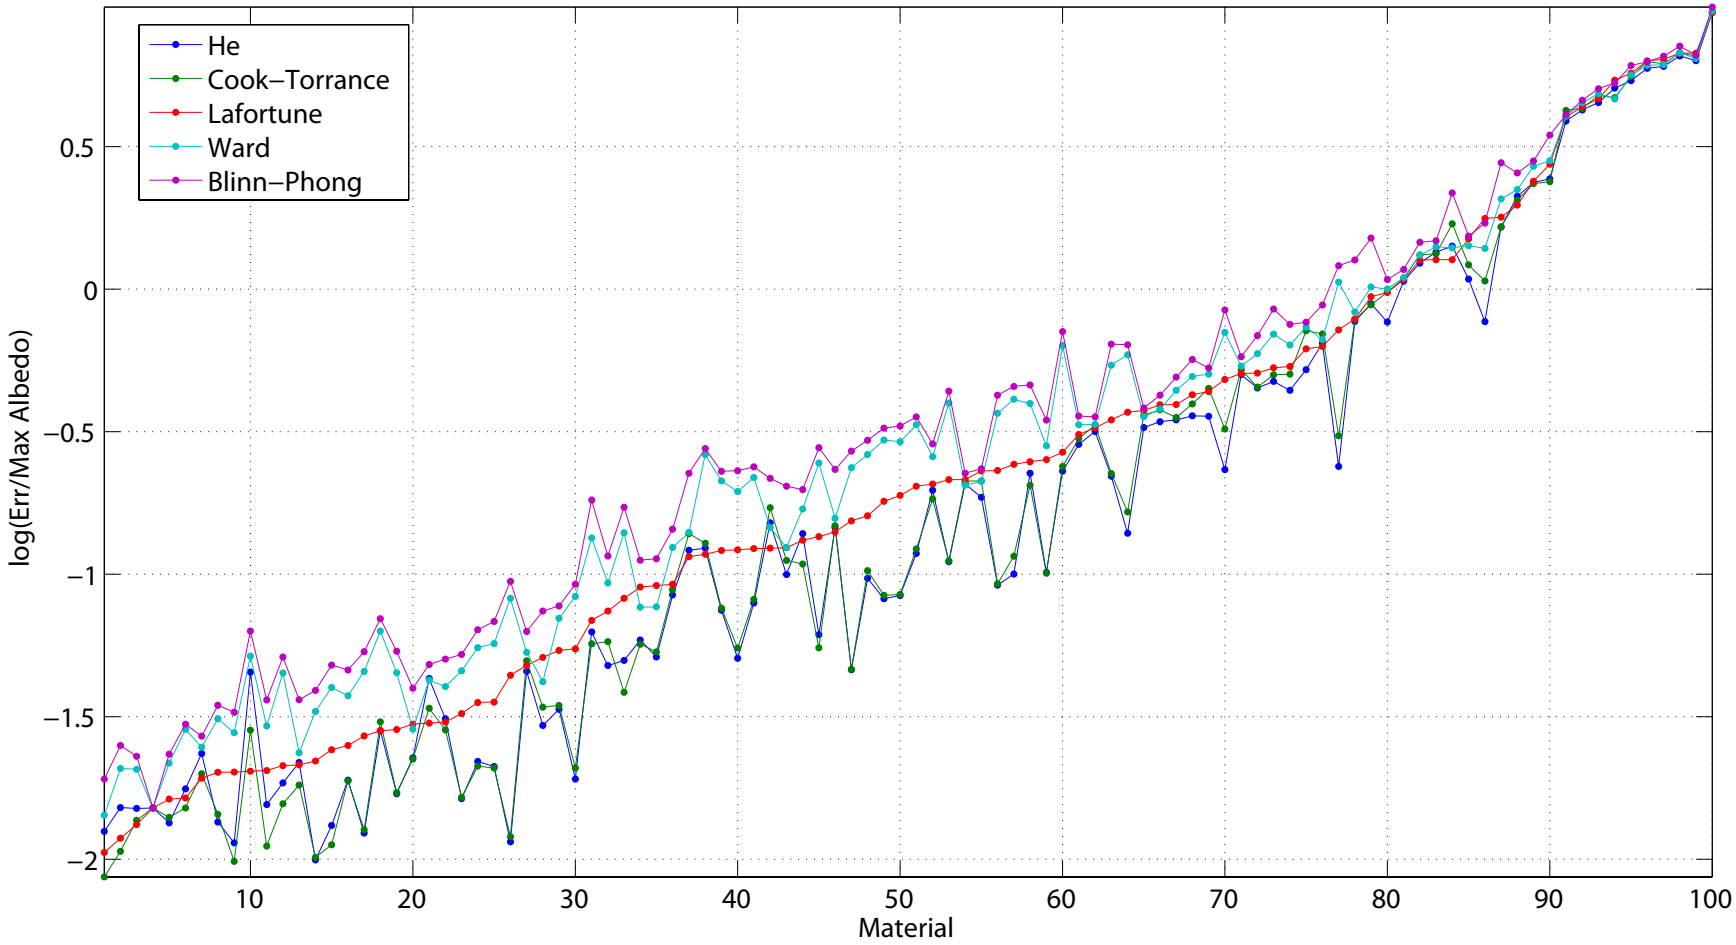

Figure 3-1: The normalized fitting errors (logarithmic scale) of five analytic models to our isotropic data set of 100 BRDFs. The BRDFs are sorted by the errors in the Lafortune model (Red) for the purpose of visualization.

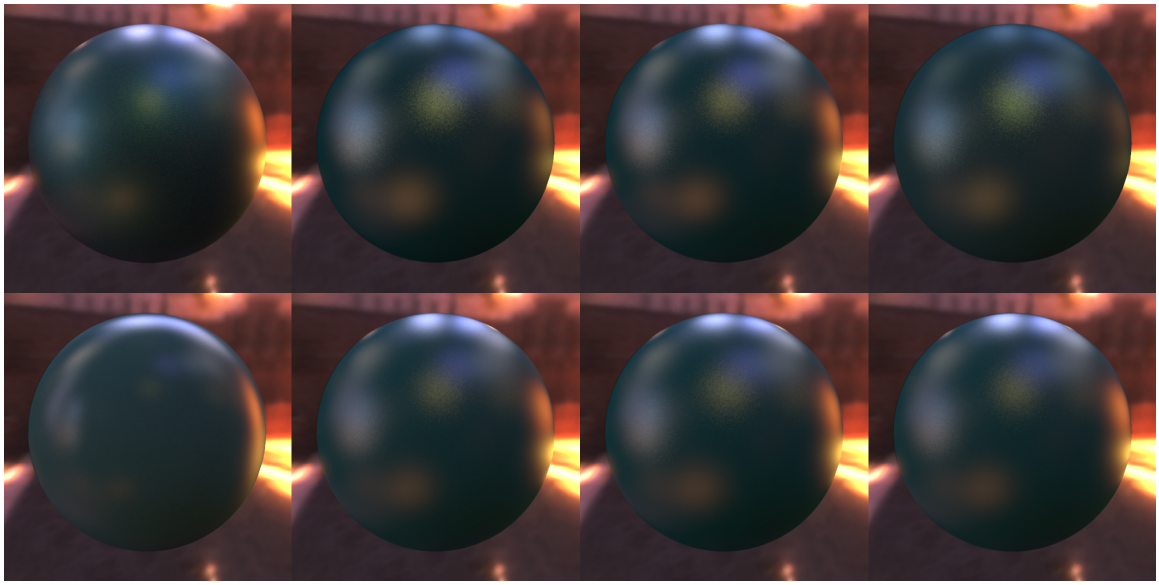

Figure 3-2: Fitting the measured "green-metallic-paint" BRDF with the seven models. Clockwise from upper left: Measured, Ward, Ward-Duer, Blinn-Phong, He, Ashikhmin-Shirley, Cook-Torrance, Lafortune.

sphere, an effect that we will look into in Section 3.5.

With a broader view over the entire data set, we find that physically-based models like Cook-Torrance and He often produce the best fits. The He model, which is known to be the most comprehensive BRDF model, does not produce noticeably superior visual results to the Cook-Torrance model. However, it should be noted that our measurement and analysis ignore polarization and spectral dependence, two important aspects modeled by He. The Lafortune model, without an explicit Fresnel factor, yields reasonable fits for materials that do not show significant intensity increase near grazing angle. In addition, glossy materials are not good candidates for the Lafortune model due to the anisotropic blurriness near grazing angle. The Blinn-Phong and Ward model both have near constant reflection power independent of incident angle, and thus they are only suitable to a subclass of the materials, including some metals and plastics. The Duer version of the Ward model with increased strength near grazing improves the fitting performance, however the improvement is limited as the grazing angle increase is not controllable.

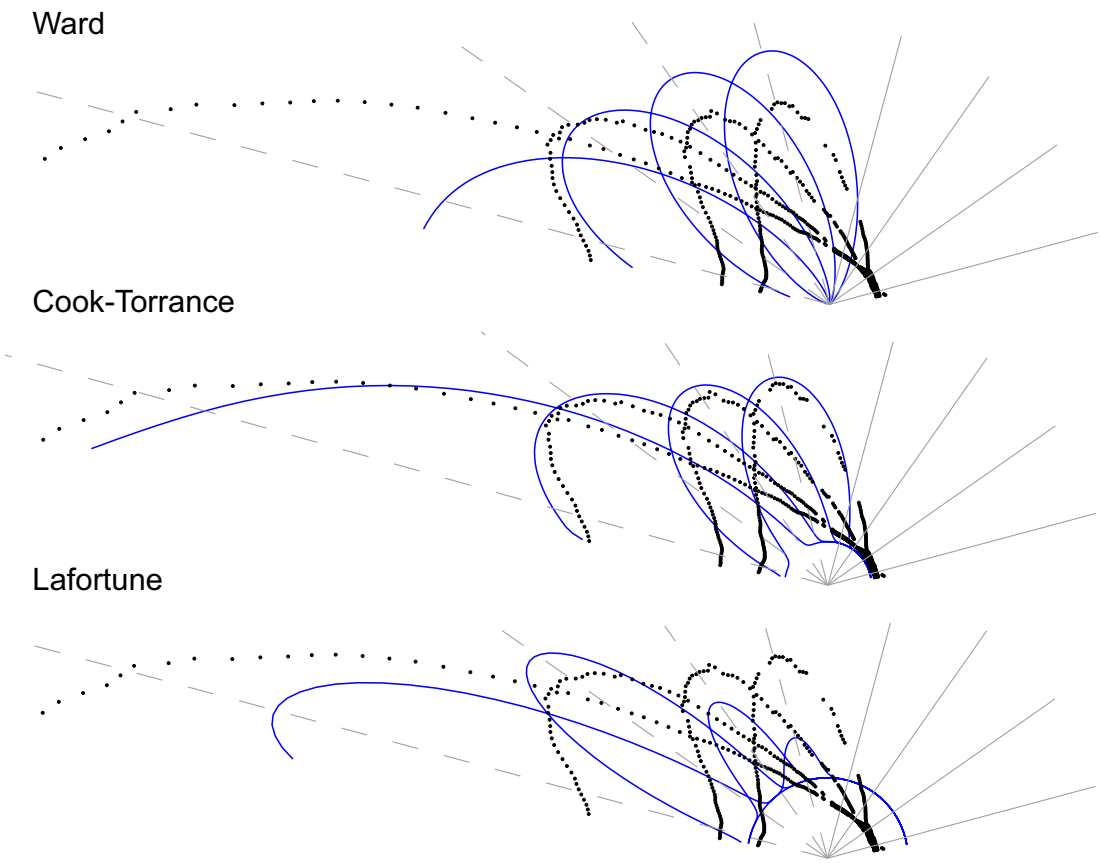

Figure 3-3: Polar plots of the measured "green-metallic-paint" in the incidence plane compared to the fits from the Ward, Cook-Torrance and Lafortune models. Cubic root applied.

3.4.2 Multiple lobes

Cook and Torrance [13] suggest that two or more specular lobes are necessary to model surfaces where there are multiple scales of roughness. Lafortune et al's generalized cosine lobes [52] serve as basis functions for representing BRDFs and thus multiple lobes are expected in general. We fit the data set with the two lobes version of both the Cook-Torrance and Lafortune models. In 26 cases for Cook-Torrance and 31 cases for Lafortune, the extra lobe reduces the error by more than 25%. These typically include materials that have multiple layers of finishing (e.g. metallic paints) and also some brushed metals. The fitting errors with the Cook-Torrance and the Lafortune models with one and two lobes for the 30 most improved materials are shown in Figure 3-4. Figure 3-5 compares the one-lobe and two-lobes fits of the nickel sample.

Lafortune et al. [52] employed three lobes for their example materials. However, we

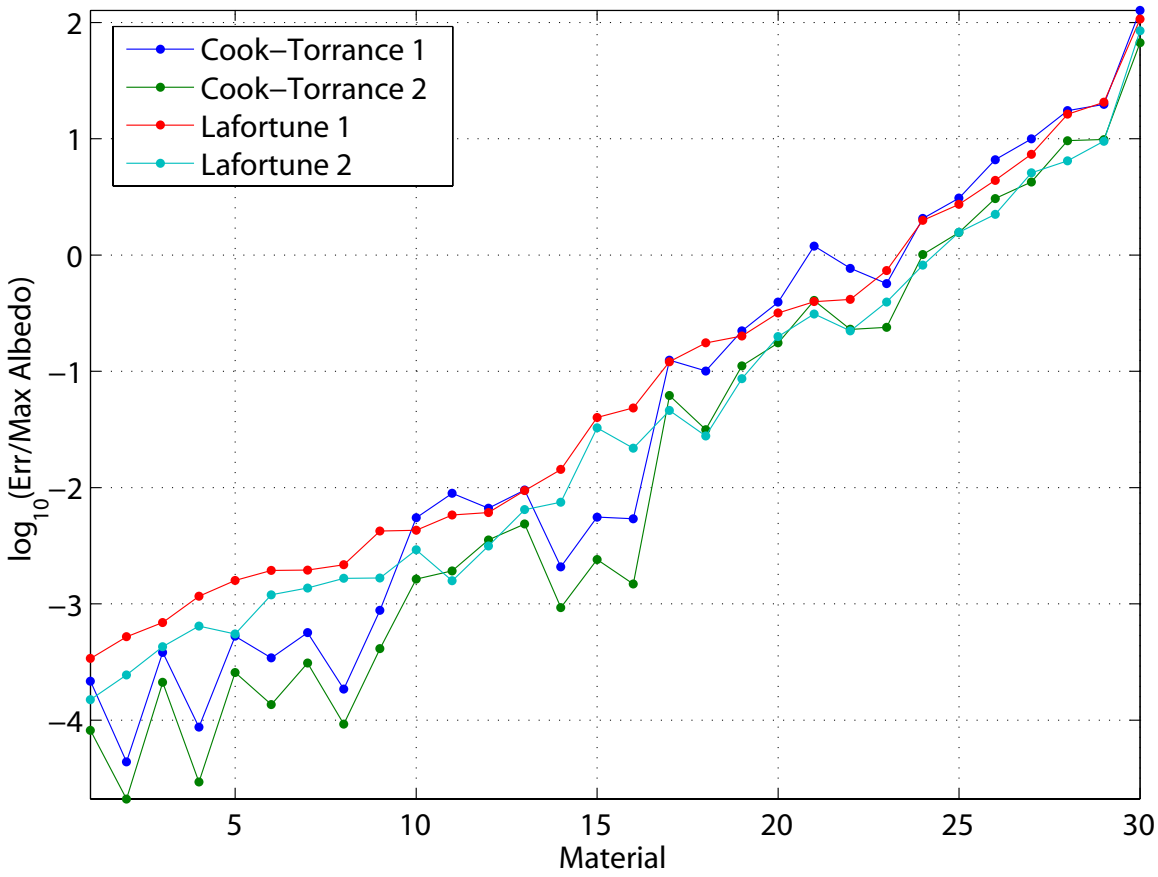

Figure 3-4: The fitting errors (logarithmic scale) with one/two lobes of the Cook-Torrance and the Lafortune models.

found that in practice, fitting with three lobes is very unstable, and hence in the current version we omit results from the three-lobe Lafortune model. Qualitatively, we believe the addition of a third lobe would only offer marginal improvements to the shape mismatch of the main lobe.

3.4.3 Evaluation of different microfacet distributions

There are three common choices of the shape of the specular lobe: Gaussian, cosine power, and the Beckmann distribution function. Our study shows that the fitting quality of the Cook-Torrance model with the three different formulations are nearly identical both numerically and visually.

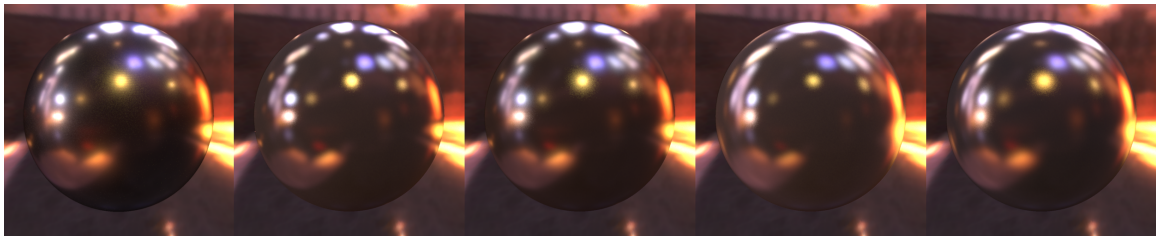

Figure 3-5: Fitting the measured BRDF "nickel" with the multi-lobe Cook-Torrance and Lafortune models. Left to Right: Input, Cook-Torrance 1 lobe, Cook-Torrance 2 lobes, Lafortune 1 lobe, Lafortune 2 lobes.

3.4.4 Fresnel effect

To study the importance of the Fresnel effect, we compare the fit quality of the Cook-Torrance model with and without the standard Fresnel term, and also with the Schlick approximation [98]. While the errors without the Fresnel term are significantly higher (20% on average), the Schlick approximation does not introduce any additional errors compared to the original formulation.

3.5 Formulation of the specular lobe

Our fitting results highlight the profound difference between two popular formulations of the specular lobe and their effects at near grazing angles. Most previous discussions of the behavior of reflectance when the incident angle moves toward grazing have focused on the increase in intensity of the reflection. As such, BRDF models are often only compared to measurements in the incidence plane. We find that the shape of the lobe also has important implications. The mirror angle construction employed by models like Phong and Lafortune, results in excessively blurry reflections near grazing angle. While the shape difference between these two lobe constructions is known to BRDF specialists, we believe it is important to restate it in the context of experimental evaluation, so that practitioners understand the tradeoff in models such as Lafortune that sacrifice lobe shape for computation speed.

In the original Phong model [87], the specular lobe is defined around the mirror direction symmetrically ($V \cdot R$). In the Blinn variation of the Phong model and most microfacet-

based models, the specular lobe is expressed with respect to the half vector H and the normal N, where H is the average vector between the incident and outgoing vector. The Lafortune model, on the other hand, employs a generalized cosine lobe similar to the original Phong lobe (Section 2.4). While the generalized dot product gives flexibility in specifying the principal reflection direction, the lobe remains symmetric around the main reflection direction. Hence, the Lafortune lobe is similar to the $V \cdot R$ lobe where R can be seen as a generalized *mirror* direction.

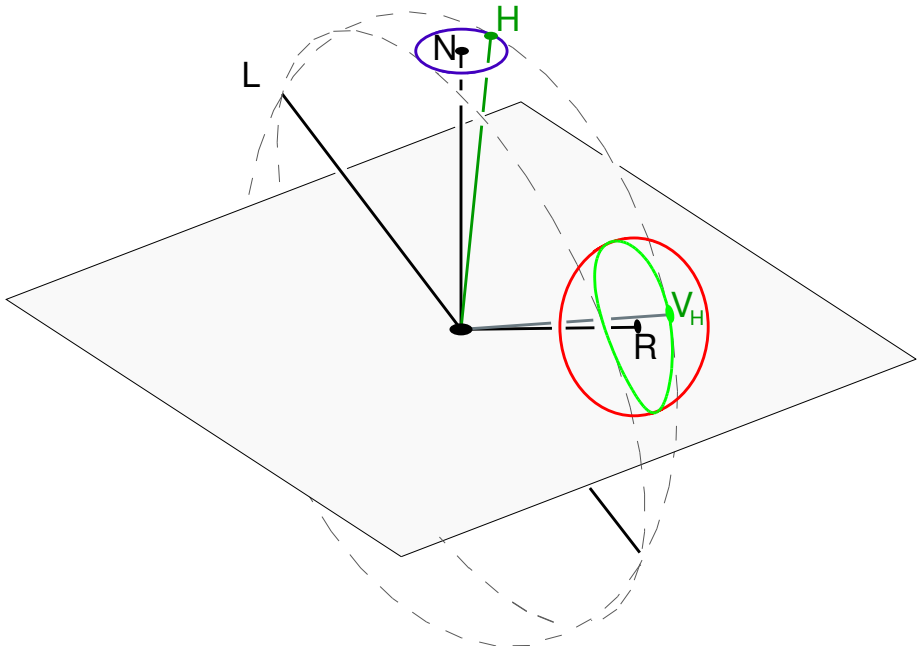

Figure 3-6: The $V \cdot R$ lobe compared with the $H \cdot N$ lobe remapped to the outgoing directions.

Since all analytical models we have considered are functions of either $H \cdot N$ or $V \cdot R$, we focus the rest of our discussion on isocontours of these dot products in the space of directions. We illustrate the difference between the two representations in Figure 3-6. Consider a given incident direction L. The blue circle around the normal N denotes the cone of directions where the dot product with N is constant. One of the half-vectors H in this set is shown. We can map H to its corresponding outgoing direction V_H, following a great circle on the sphere of directions. This can be done for the entire set of half-vectors in the blue circle, and we show the corresponding set of outgoing directions in green around the mirror vector R. Comparing it to the red circle, representing the cone of directions V

with constant $V \cdot R$, the remapped contour is asymmetric and is narrower in the direction perpendicular to the incident plane. By similar construction, we can visualize that this contour is more circular when L is near N, and becomes narrower when L moves towards grazing angle. The $V \cdot R$ lobe is, however, always circular around the mirror direction R.

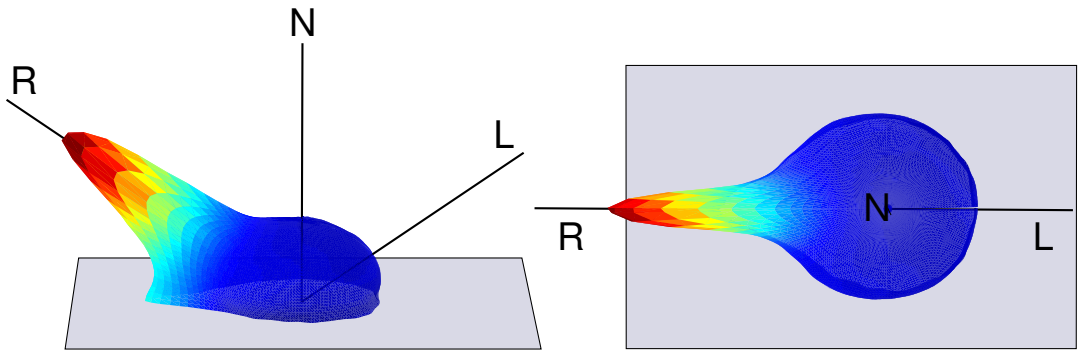

Figure 3-7: Side and top-down view of the measured "PVC" BRDF, at 55° incidence. (Cubic root is applied to the BRDF for visualization.) Its specular lobe exhibits a similar asymmetry to the $H \cdot N$ lobe.

A plot of the measured BRDF of PVC plastic at 55° incident angle is shown in Figure 3-7. It displays a similar asymmetry to the remapped contour associated with the $H \cdot N$ lobe. This asymmetry is observed in the data set for most materials with an apparent specular lobe. Such behavior can be explained by the microfacet theory, e.g. [13]. The theory models the surface as a collection of microfacets, with normals oriented according to some distribution function. Assuming that the microfacets are perfect mirrors, the reflectance from L to V is then proportional to the fraction of microfacets with normals aligned with the half-vector H, with some adjustments for shadowing and masking. Thus, the reflectance function for a microfacet model is naturally defined with respect to H. For isotropic models, the distribution of normals is radially symmetric; therefore, a cone of half-vector H around N should have constant reflectance value if the shadowing and masking terms are not considered. Hence, it is reasonable to define reflectance as a function of $H \cdot N$ under this theory. On the contrary, the $V \cdot R$ contour is not consistent with any distribution of normals, as a circular cone of vectors around different Rs with the same spread would map

to a different set of half-vectors around N.

The difference between the two formulations explains the anisotropic blurring for the Lafortune fits in Figure 3-5 and Figure 3-2. Compared to a half-vector based lobe, the symmetric mirror lobe does not get narrow at grazing angles, resulting in blurring in the direction orthogonal to the incident plane, which is along the circumference for the case of a sphere.

We have shown that the $H \cdot N$ representation is a more accurate way to model specular reflections. This is in strong opposition to some beliefs that the half-vector representation is just a more convenient way of interpreting the specular lobe. We believe this has been overlooked in the past mainly due to the fact that the two representations are trivially related in the incidence plane.

Note that this observation is different from the popular re-parameterization proposed by Rusinkiewicz [96] that also emphasizes the importance of the half vector. In his work, he shows that the parameterization based on half-angle is a more compact representation in the context of compression, compared to the standard BRDF parameterization based on the incident and outgoing directions. In our work, we compare the *shape* of the reflection lobe defined by the half vector to the one defined by the *mirror direction*.

Stark et al.[107] show that Ward and Lafortune can both be mathematically reduced to bivariate functions, and thus they are inherently unable to represent general isotropic BRDFs of three variables. They went on to introduce a new set of 2D coordinates which can be considered as a hybrid of the two families of parameterization. They showed that their parameterization produces a superior approximation to two measured materials when the BRDFs are projected onto the 2D domain. Further investigating the performance of Stark et al.'s proposed model and variants based on the hybrid parameterization would be an interesting avenue for future work.

3.6 Summary

We have analyzed the performance of several analytical BRDF models by their ability to fit real BRDFs. Our experimental results suggest that using a single specular lobe, the Cook-

Torrance, Ashikhmin-Shirley and the He models perform well for most of the 100 isotropic BRDFs in the data set. Moreover, our results show that for a number of materials, the fit quality is much improved when an additional lobe is employed. This is mostly attributed to the multiple layer finish on these materials.

Our study also illustrates that the difference between the two ways of modeling the specular lobes is profound. The $H \cdot N$ formulation naturally corresponds to the microfacet-based theory, and its viability is confirmed by the data set. Meanwhile, models based on the $V \cdot R$ lobes are visually inaccurate in a rendered scene. The fitting results of these models are systematically inferior to their counterparts. The Lafortune model, though highly expressive with its basis of generalized cosine lobes, cannot accurately model real BRDFs with a single lobe. While the computation for each generalized cosine lobe is cheaper relative to a Cook-Torrance lobe, fitting more than a few of these lobes to measured data is unstable.

3.7 Discussion

There are several directions for future work. First, we expect that a new set of basis functions, similar to the generalized cosine lobes, but expressed with the half-vector, would yield improved fitting results with fewer lobes. Stark et al.[107]'s barycentric coordinates are also possible candidates. These basis functions should be reasonably inexpensive to compute and be flexible enough to represent complex reflectance phenomena, e.g. retro-reflection. Second, while our metric for BRDF fitting gives reasonable results in most cases, it does not directly correspond to perceptually important differences. A perceptually-based metric would improve the visual quality of the fits [85, 94], especially when the fit is relatively poor and different choices of approximation tradeoffs are present. However, perceptual metrics are often computationally expensive and may not be practical as an objective function for optimization. A possible alternative is our image-driven metric (Chapter 6), which can be accelerated by precomputing and embedding the models in a linear space.

Chapter 4

Representing BRDFs with a Sampled Microfacet Distribution

There are a variety of methods to represent BRDFs. Raw measured data can be stored in a tabular format that partitions the 4D BRDF domain into uniform bins. Assuming the resolution is sufficiently high and every table entry is filled, the BRDF can be evaluated through table lookups. This is the most general representation as it does not pose any assumption on the BRDF. However, this representation requires dense measurements and consumes a large amount of storage. Compressed representations decompose the BRDF into a linear basis such as spherical harmonics and discard the low-energy components. They are more compact than the 4D table, and they can often smoothly fill in the missing bins. In current graphics applications, the most popular way of representing BRDFs is the use of analytical models. Analytical models are typically defined with a single equation of a few parameters. They are compact, smooth, adjustable and do not generally require measured data as input. In the evaluation described in the previous chapter, we find that good models such as Cook-Torrance and He provide reasonable approximations to many isotropic materials. Unfortunately, as we will demonstrate in this chapter, the situation is different for more complex anisotropic materials such as velvet and satin. These materials exhibit behaviors that are outside the modeling power of any pure analytical model.

In this chapter, we show that a compact 2D representation based on microfacet theory, proposed by Ashikhmin et al. [1], provides the required expressiveness for complex mate-

rials. Ashikhmin's work is focused on the generation of a physically plausible BRDF from an arbitrary microfacet distribution. Due to the high degrees of freedom in this representation, fitting measured data with the optimization approach described in the last chapter is impractical. Instead, by exploiting the smoothness property of the model, we have developed an efficient iterative method to estimate the microfacet distribution from measured data.

In contrast to our study on isotropic BRDFs, the validation based on measured data in this chapter is far less exhaustive. This is in large part due to the difficulty in measuring anisotropic materials. Moreover, we observe that purely analytical models are fundamentally limited by simple geometric assumptions that fail on complex materials. Consequently, comparison of these models would not be very meaningful.

In the following sections, we will first review the basic microfacet theory and Ashikhmin et al.'s generalized model. Next we describe our method to estimate a microfacet distribution from BRDF measurements. We then study a few measured materials and show that in some cases pure analytical models are completely incapable of approximating their appearance. We conclude by showing results of our microfacet estimation.

The estimation technique and the evaluation of anisotropic materials has been published previously in [80]. In this chapter, we show additional results by applying the technique to isotropic materials.

4.1 Background

4.1.1 Microfacet Theory

The microfacet theory is first developed by Torrance and Sparrow [110] in optics and later introduced to computer graphics by Blinn [6] and Cook and Torrance [13]. The main idea of the microfacet theory is to assume that the surface consists of a large number of small flat perfect mirrors. Intuitively, the evaluation of the BRDF at a particular pair of directions involves counting the number of aligned mirrors that are visible and unshadowed, and computing the Fresnel coefficient for the perfect mirrors. After some geometrical

derivation (see [1]), we can express a general microfacet-based BRDF model with Equation 4.1, following the notation convention defined in Section 2.4.

$$\rho(L,V) = \frac{p(H)P(L,V,H)F(L\cdot H)}{(L\cdot N)(V\cdot N)} \tag{4.1}$$

where $p(H)$ is the probability density of the microfacets aligned in the half-vector direction H, and $F(L\cdot H)$ is the Fresnel reflectance for these aligned microfacets. $P(L,V,H)$ is the shadowing term, which corresponds to the probability for a microfacet *not* to be shadowed in either light or view directions.

In the work by Torrance and Sparrow [110], a simple Gaussian or Gaussian-like function specified by a single parameter is used to describe $p(H)$. Since the shadowing effect depends on the precise microgeometry, a distribution of microfacet orientations does not uniquely define the shadowing term. To estimate a reasonable shadowing term, Torrance and Sparrow make a strong assumption that the surface can be modeled by long V-shaped cavities. The shadowing term is shown in Equation 4.2. Note that this term has no dependency on the distribution $p(H)$.

$$P(L,V,H) = G = min\{1, \frac{2(N\cdot H)(N\cdot V)}{(V\cdot H)}, \frac{2(N\cdot H)(N\cdot L)}{(V\cdot H)}\} \tag{4.2}$$

4.1.2 Ashikhmin's Microfacet-based BRDF Generator

To generalize the microfacet model, Ashikhmin et al. [1] propose a method to generate a BRDF based on an arbitrary distribution $p(H)$. Similar to the traditional microfacet methods, the goal is to estimate a plausible BRDF given a statistical distribution of the microfacet orientations, without precise description of the actual microgeometry. But unlike the distribution-independent shadowing term proposed by Torrance and Sparrow, a more general shadowing term that is consistent with an arbitrary $p(H)$ is required. We will show a simplified version of the derivation here which only provides the shadowing term up to a scalar factor. The reader is referred to the original paper for the full derivation.

To make the development of such shadowing term tractable, two assumptions are made. First, shadowing with respect to the light and view directions are assumed to be independent

events:

$$P(L,V,H) = P(L,H)P(V,H) \tag{4.3}$$

This assumption allows the independent estimation of the visibility from the light and view direction. However, it can lead to some underestimation of the BRDF if L and V are close together, since in such case the visibility from L and V should be correlated.

The second assumption is that the probability of a microfacet to be visible from a direction K does not depend on the orientation of the microfacet, except for the fact that the microfacet do not face away from K:

$$P(K,H) = \begin{cases} P(K) & \text{if } K \cdot H > 0 \\ 0 & \text{if } K \cdot H < 0 \end{cases} \tag{4.4}$$

With these two assumptions the shadowing term can be estimated based on an argument of visible surface area. First we assume the surface has a planar total area of A, and each microfacet has the same area A_{mf}. The total area projected in the direction K is $(K \cdot N)A$. At the same time, this projected area is also equal to the sum of the projected area over all microfacets visible from the direction K:

$$(K \cdot N)A = \sum_{i \in all\, facets} A_{mf}(K \cdot H_i)_+ P(K) \tag{4.5}$$

Here the subscript '+' denotes the clamping of the dot product to non-negative values. If we convert the summation to averaging over the distribution:

$$(K \cdot N)A \propto A_{mf} P(K) \int (K \cdot H)_+ p(H) d\omega_H \tag{4.6}$$

As a result, we have

$$P(K) \propto \frac{(K \cdot N)}{g(K)} \tag{4.7}$$

where $g(K)$ is defined as:

$$g(K) = \int (K \cdot H)_+ p(H) d\omega_H \tag{4.8}$$

The final formula for the BRDF can be derived by substituting $P(K)$ for $K = V$ and $K = L$ into Equation 4.1, together with the correct constant of proportionality (not derived here) [1]:

$$\rho(L,V) = \frac{p(H)\langle H \cdot N \rangle F(L \cdot H)}{4g(L)g(V)} \tag{4.9}$$

where $\langle \cdot \rangle$ denotes the average over the distribution $p(H)$.

As discussed in Ashikhmin et al.'s paper, while the shadowing term is designed to be consistent with any microfacet distribution, it is a simple approximation and does not apply to situations when intricate shadowing effects dominate the appearance of the BRDF.

By manually designing the distributions for a few materials including velvets and satins, Ashikhmin demonstrates that the generalized model can qualitatively simulate these complex BRDFs. However, the core missing component of the model is a way to use it to approximate measured data. Fitting via an optimization approach similar to the one described in Chapter 3 is impractical due to the high degrees of freedom in the specification of $p(H)$. In the next section we describe a simple and efficient method to estimate the best $p(H)$ for measured BRDF.

4.2 Estimation of the Microfacet Distribution from Measured Data

We can rewrite Equation 4.9 in the following way:

$$p(H) = \rho(L,V) \frac{4g(L)g(V)}{\langle H \cdot N \rangle F(L \cdot H)} \tag{4.10}$$

First, note that the Fresnel term, $F(L \cdot H)$, is a function fully specified by a single parameter[1]. For the moment, we assume the Fresnel parameter and the shadowing function

[1]Either the index of refraction, or the reflectance at normal incidence can fully define the Fresnel re-

$g(\cdot)$ are given. Each measured BRDF sample $\rho(L,V)$ then directly provides an estimate to the probability density $p(H)$. In practice, we divide the hemispherical domain of $p(H)$ into uniform bins. We obtain a least-squares estimate of the density for each bin by taking the average.

The main difficulty, however, is the function g which depends on the distribution $p(H)$ itself. Fortunately, we observe that g is extremely smooth and we have found that an iterative approach works very well. We start by assuming a constant g, and we alternate the estimation of $p(H)$ using the current g and the computation of g from the updated $p(H)$. In practice, the algorithm converges in less than six iterations. The converged distribution can be then used to generate the full BRDF with Equation 4.9 and the difference from the measured data is computed as the fitting error. To find the Fresnel parameter, we use a simple line-search algorithm in a loop that optimizes for the parameter with the lowest fitting error.

4.3 Data Analysis

To illustrate the performance of the generalized microfacet model, we have measured a few anisotropic materials using our new setup (See Section 2.3.1). We will first describe the materials, and then show fitting results to analytical models as well as the generalized microfacet model.

4.3.1 Measurement of anisotropic materials

To perform measurement of anisotropic materials, we have built a new image-based setup described in Section 2.3.1. Our anisotropic data set currently includes 4 materials: velvet, two satins, and brushed aluminum. One of the acquired images for the velvet is shown in Figure 4-1. The light position is rotated at intervals of 2 degrees and covers most of the full circle. Positions where the light occludes the camera or is occluded by the cylinder are discarded. The motor driving the cylinder is incremented at intervals of 5 degrees, from 0 to 180. The strips on the cylinder are cut from a planar sample of the material using

flectance at all angles.

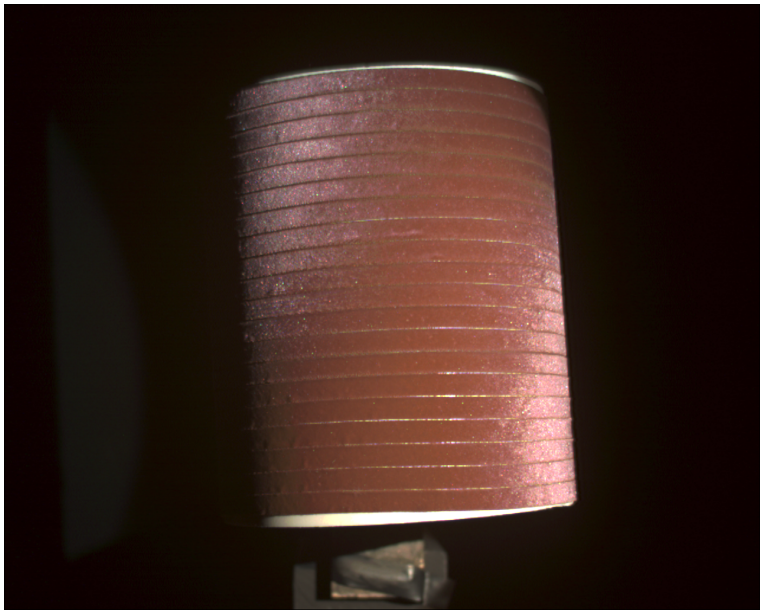

Figure 4-1: Target cylinder covered with velvet.

a laser cutter to obtain precise orientations. The orientations of the strip range from 0 to 180 degrees spaced at 9 degree. Combining this with the cylinder rotation provides a full 360 coverage. For each object/light position combination, we take 8 images with exposure time from 40 microseconds to 1 second, and combine them to form a high dynamic range image. The acquisition time is about 16 hours for each material and the raw uncompressed data size is 30GB.

We discretize the 4D domain of the BRDF into bins at 2 degrees intervals. The measured data are distributed to the appropriate bins and approximately only 25% of the bins are non-empty with setup. The typical sampling pattern of our measurement is shown in Figure 4-2. To the best of our knowledge, our anisotropic measurements have the highest resolution among publicly available BRDF data. However, due to the non-uniform sampling, the raw data is unsuitable for direct rendering.

4.3.2 Analytical Models

In this section we describe our attempt to use purely analytical reflectance models to represent the measured anisotropic materials. We find that two materials in our data set (brushed aluminum and yellow satin) are suitable for analytical modeling. We fit two different

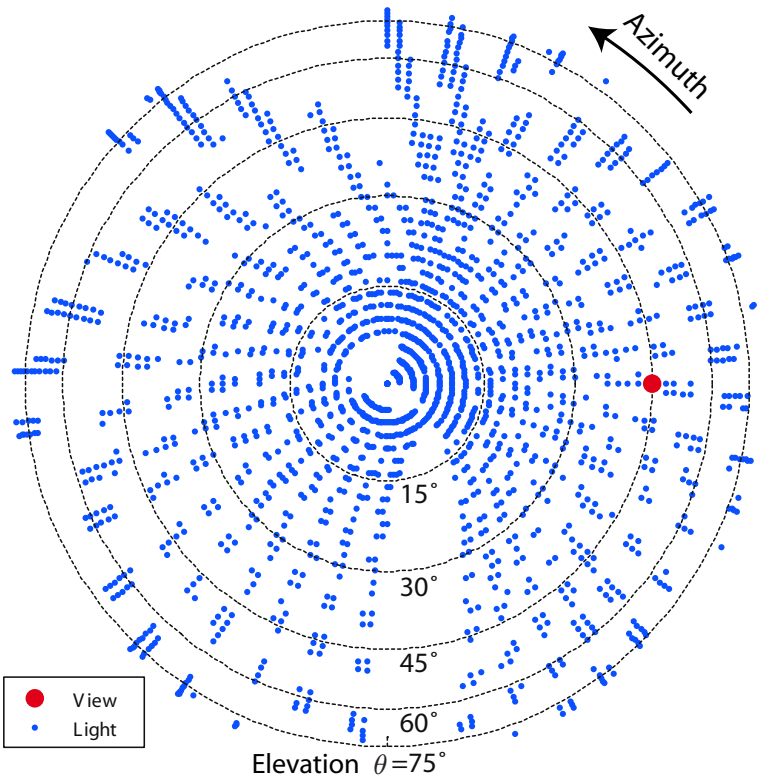

Figure 4-2: Sampling density for a typical measured BRDF at a particular view direction ($\theta = 45, \phi = 0$). Blue: data bins containing valid measurements.

anisotropic analytical models to these measurements – anisotropic Ward [118] and Poulin-Fournier [90] models, using the same fitting procedure as for the isotropic materials. We present the visual comparison of the fit for the brushed aluminum to the original measurements in Figure 4-3. These models can qualitatively represent brushed aluminum relatively faithfully. The yellow satin (not shown) has blurrier highlights, but the anisotropy is qualitatively similar to the brushed aluminum. As a result it can also be modeled reasonably well with the analytical models.

Some fabrics that exhibit anisotropic reflectance properties can have much more complicated micro-geometry since the way threads are woven together can be potentially much more complex than the relatively simple structure of a brushed metal. We observe that the BRDFs of the purple satin and red velvet exceed the expressive power of simple analytical models. Notably, while most materials exhibit the peak intensity near the center of the highlight, the purple satin sample in Figure 4-4 shows two peaks away from the center.

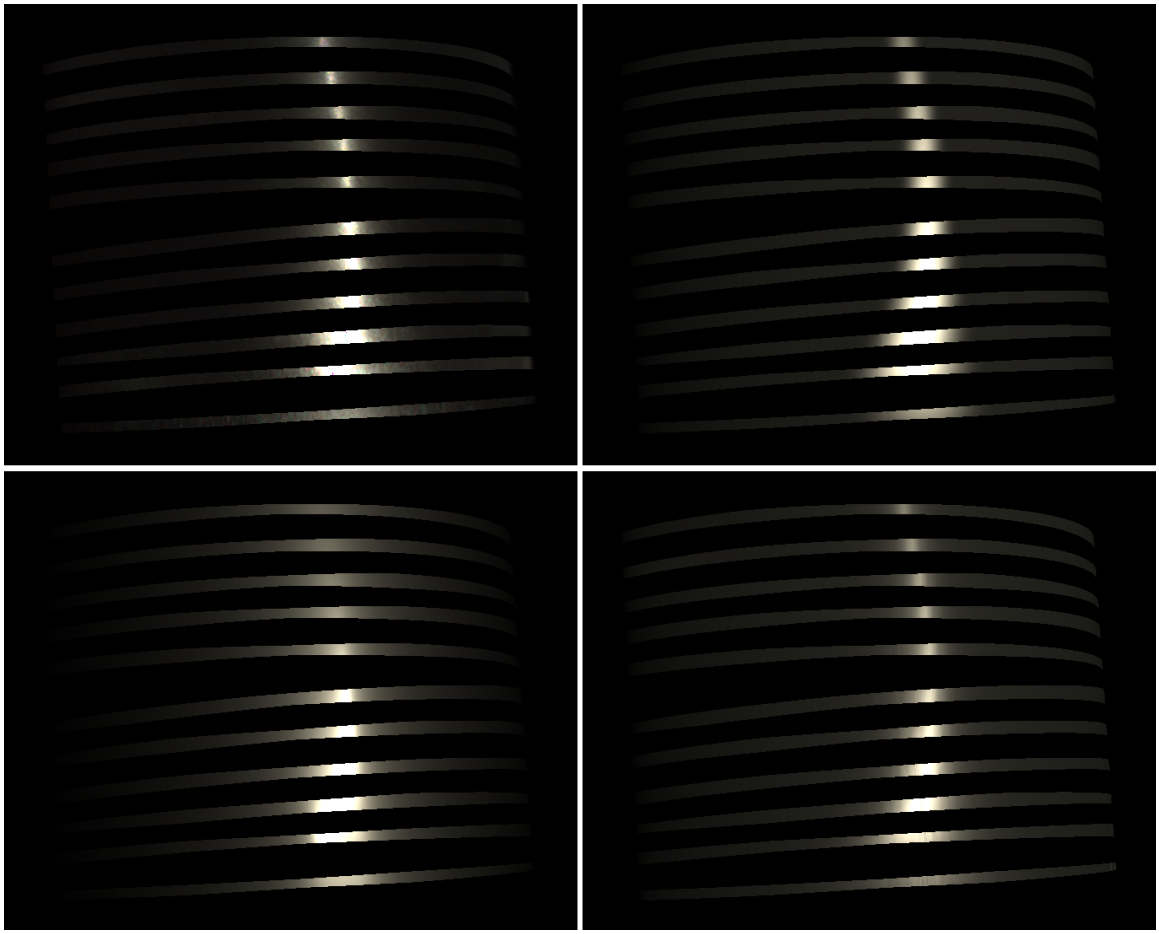

Figure 4-3: Brushed Aluminum: One measurement photograph (upper left), reconstruction using Ward model ($\alpha_x = 0.04, \alpha_y = 0.135$) (upper right), reconstruction using Poulin-Fournier model ($n = 3000, d = 0.18, h = 0.5$)(lower left), reconstruction using Ashikhmin model with sampled microfacet distribution (lower right).

4.3.3 Generalized Microfacet Model

Using the method described in Section 4.2, we compute estimated microfacet distributions for our four measured anisotropic materials. Figure 4-3 and Figure 4-4 show comparisons between photographs of the measurement cylinder and renderings using the estimated microfacet distribution. The renderings show good qualitative agreement with the photographs and the shape of the highlights is well-preserved, although their sharpness is not reproduced exactly. There are possibly multiple reasons for the loss of sharpness. Our microfacet distribution is tabulated at finite intervals and may not have enough resolution to represent the material. Any calibration errors due to hardware and software can also lead

to blurring of the BRDF.

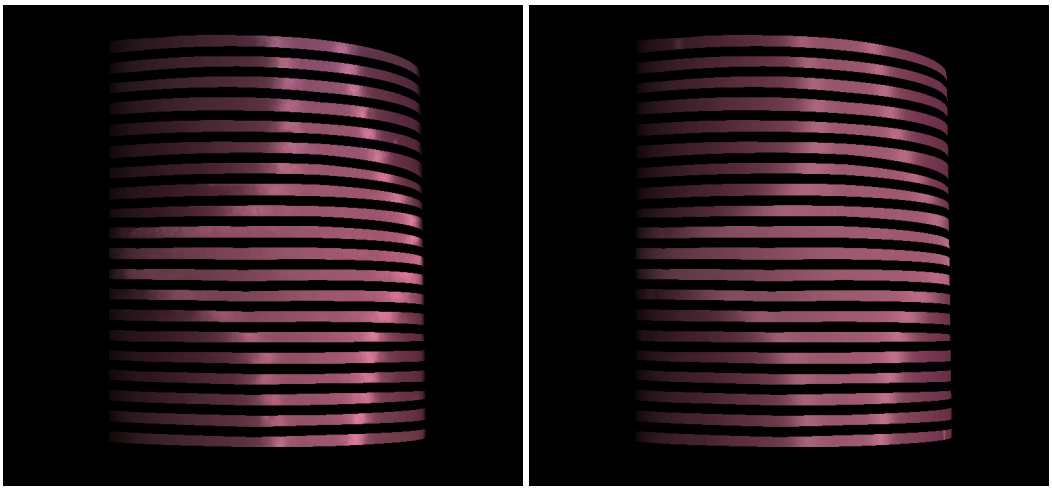

Figure 4-4: Left: Purple satin - one input photograph. Right: - reconstruction using sampled microfacet distribution.

The computed microfacet distributions also match the microstructure of the materials. Compare the spherical plot of the microfacet density in Figure 4-5 with a macro photograph of the corresponding sample of brushed aluminum. The density is high around a vertical great circle, which is the type of normal distribution expected from cylindrical micro-geometry.

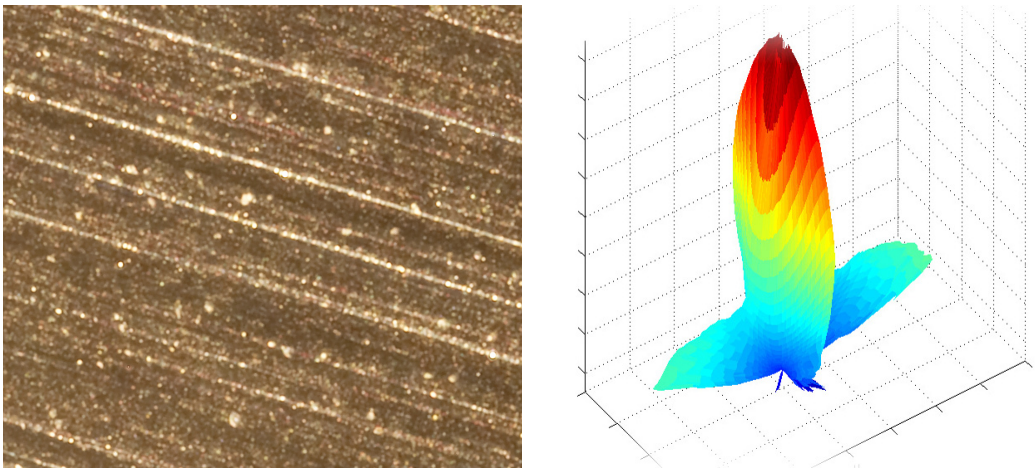

Figure 4-5: Left: Brushed aluminium macro photograph. Right: deduced microfacet distribution (log plot).

The purple satin sample has an intriguing distribution: a large density around two great circles that are symmetrically slightly tilted from the vertical direction (Figure 4-6). This

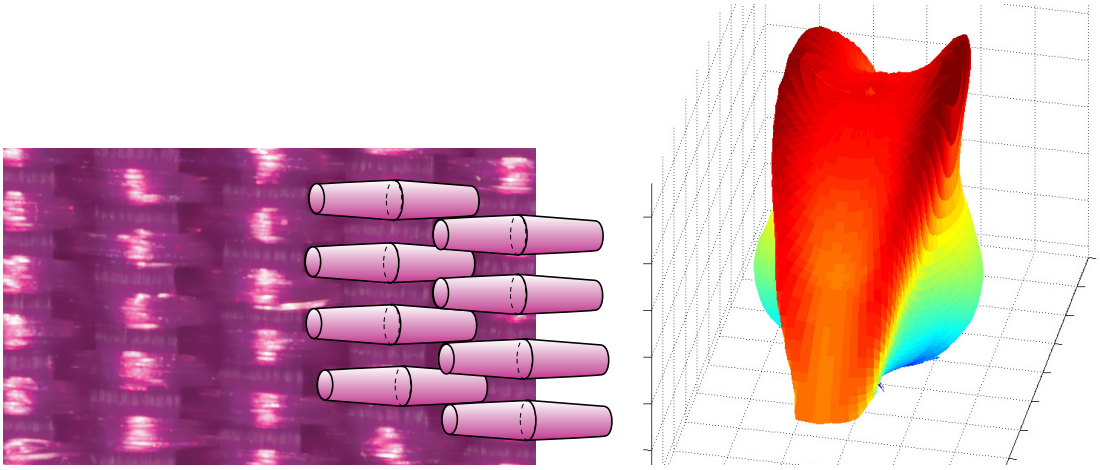

Figure 4-6: Left: Purple satin macro photograph with sketched double cones. Right: deduced microfacet distribution (log spherical plot).

type of distribution can be caused by symmetric cone pairs. When we examine the macro photograph (Figure 4-6) we note that the longest visible segments of thread are high in their middle and have a symmetrical slope that might be similar to a cone, as sketched on the right part of the image.

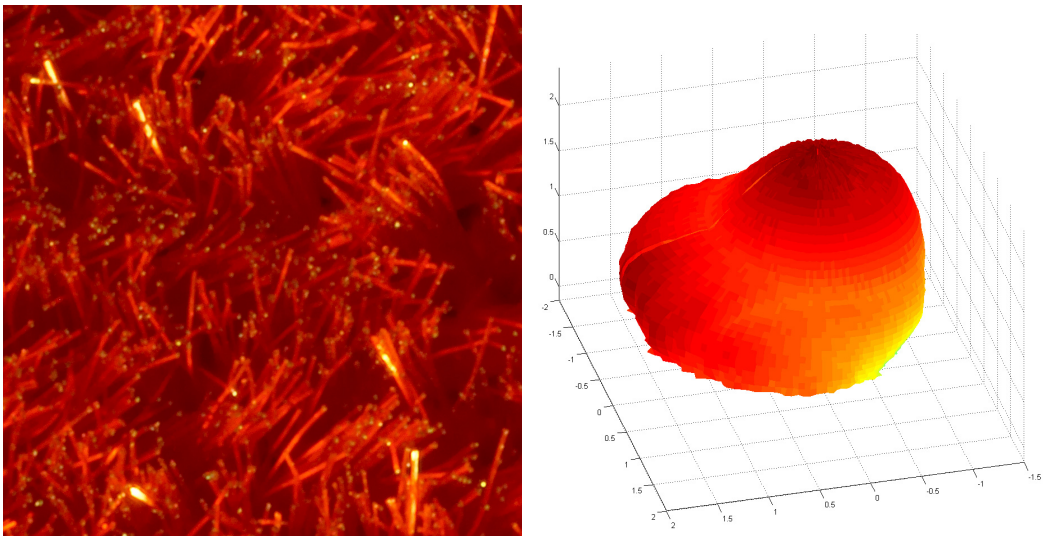

Figure 4-7: Left: Red velvet macro photograph. Right: deduced microfacet distribution (log spherical plot).

The red velvet sample exhibits a secondary lobe in addition to the main lobe around the normal. This is mainly due to the dominant combing direction of the velvet sample. This shows an important type of anisotropy that cannot be modeled by the standard analytical

models. We should also note that we lack measurements near extreme grazing angles, and therefore we may have missed important reflectance properties of the velvet in those directions. In addition, it is difficult to enforce the exact combing direction uniformly on the entire piece of velvet. Small deviations can be tolerated, however, as the average BRDF of velvet is relatively low-frequency.

4.4 Reconstruction

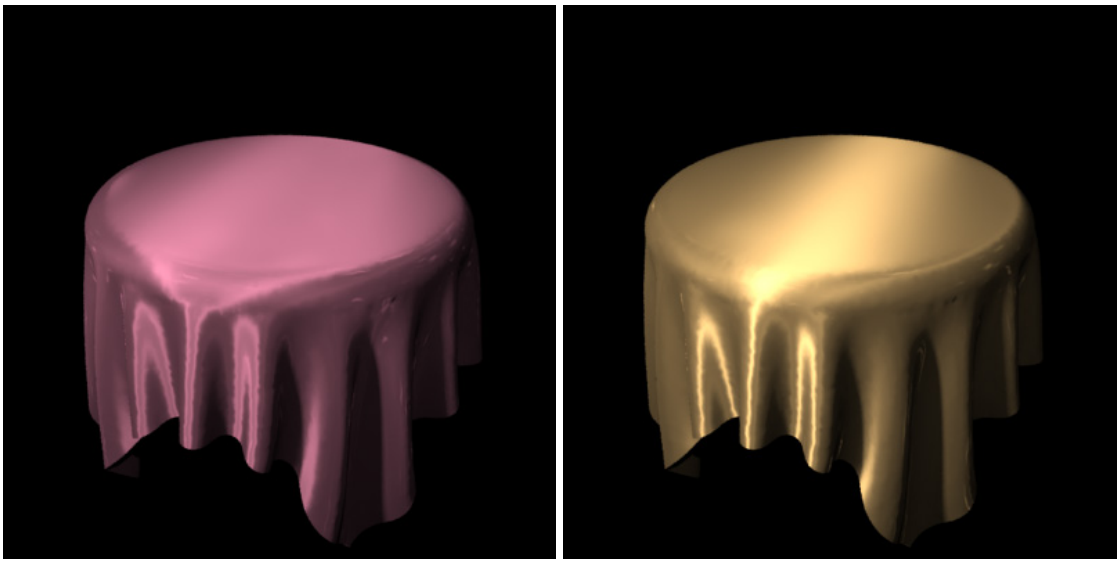

Figure 4-8: Full reconstruction of purple satin (left) and yellow satin (right).

The sparse and non-uniform sampling (See Figure 4-2) of our anisotropic measurements makes them unsuitable for direct rendering. However, since each microfacet normal corresponds to many combinations of incoming and outgoing directions, our measurement generally maps to a dense sampling of the microfacet distribution. In practice, for a two-degree binning interval, our microfacet distributions are completely covered except for extreme grazing angles. Given a complete microfacet distribution, we can reconstruct the full 4D BRDF with Equation 4.9. Two examples of the reconstructed BRDF are shown in Figure 4-8.

4.5 Application to Isotropic BRDFs

Since isotropic BRDFs are simply a subset of general 4D anisotropic BRDFs, our estimation method can be applied to these cases without any modification. Each BRDF sample from an isotropic measurement, defined by three parameters, can be replicated to simulate a set of samples of the full 4D BRDF. However, for efficiency purposes, we can modify Equation 4.9 to a simplified form for isotropic materials:

$$\rho(L,V) = \frac{p(\theta_h)\langle H \cdot N \rangle F(L \cdot H)}{4g(\theta_i)g(\theta_o)} \tag{4.11}$$

In this version, the microfacet distribution $p(\cdot)$ and the shadowing and masking function $g(\cdot)$ are both reduced to 1D functions that only depend on the angle from the normal. The estimation equation can be derived similarly. In terms of implementation, we now only use a one-dimensional array to represent the microfacet distribution $p(\cdot)$.

4.5.1 Results

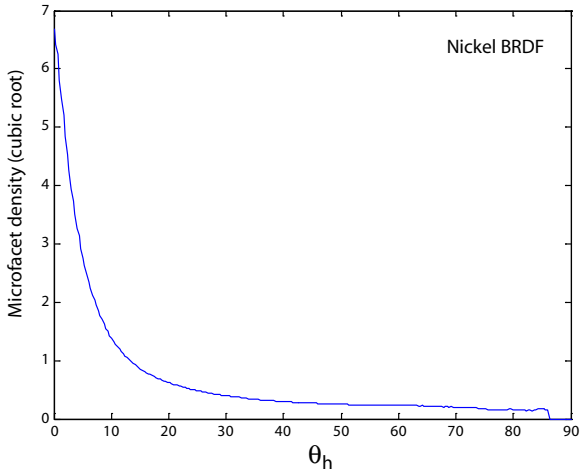

Figure 4-9: Estimated microfacet distribution of the measured BRDF nickel.

Using Equation 4.11 and the iterative technique described in Section 4.2, we have estimated the microfacet distributions for the same data set of isotropic BRDFs we used in the previous chapter, measured by Matusik et al. [68]. One example distribution, of the *nickel* BRDF, is shown in Figure 4-9. The comparison between the microfacet distribution approx-

Figure 4-10: Comparing the microfacet distribution approximation (right) to the measured BRDF nickel (left).

imation and the original data is shown in Figure 4-10. The approximation by the microfacet distribution looks remarkably similar to the original measurement. With analytical models like Cook-Torrance, two lobes are necessary to provide a good visual match to this material (Figure 3-5). However, in a somewhat surprising way, the fitting errors with the microfacet distribution estimation are sometimes higher than the single-lobe Cook-Torrance fit. Indeed, out of the 100 isotropic BRDFs in the data set, approximation with the microfacet distribution technique provides a lower fitting error than Cook-Torrance in only 46 of them. We hypothesize that for some BRDFs, the analytical models already provide a very good approximation to the distribution of microfacet normals. The additional degrees of freedom in our method does not improve the approximation quality significantly. It is also possible that the shadowing and masking terms of the isotropic models are superior to Ashikhmin's microfacet model for these materials and produce a lower error as a result. At the same time, even when the microfacet estimation does not provide better accuracy, we observe that the estimation procedure is much more stable than the non-linear optimization for analytical models. This method consistently produces well-matched results visually, even for materials where more than one lobe is required with the pure analytical models.

4.6 Summary

In this chapter we have reviewed the microfacet theory and Ashikhmin's generalized microfacet-based BRDF model. We then show a simple technique to estimate the microfacet distribution from measured data. We have built a new setup for high resolution measurement of anisotropic BRDFs using material strips on a rotating cylinder. We used it to measure four typical anisotropic materials and showed that only two of them can be modeled by analytical models. Our fabric measurements exhibit a complexity that cannot be accounted for by the simple functions used to introduce anisotropy in analytical models. We showed that the generalized microfacet model provides the necessary expressiveness to model these complex BRDFs.

4.7 Discussion

The technique proposed in this chapter is based on a hybrid representation of analytical and data-driven approaches. The generation of the full BRDF from the microfacet distribution $p(H)$is based on the microfacet theory and described with analytical functions. On the contrary, the microfacet distribution $p(H)$ is estimated through measurement and no assumption is made on the underlying properties of the material. This representation is more general than pure analytical models as it inherits some generality from the microfacet distribution estimated through measurement. At the same time, the analytical formula derived from the microfacet theory allows us to work with a lower-dimensional function which is significantly easier to measure and store.

Since the microfacet distribution is a 2D function, it is theoretically impossible to represent an arbitrary 4D function. To provide reasonable approximation to real BRDFs, our technique relies on the assumption that the microfacet theory is *correct*. There are two main assumptions in the theory: a) each microfacet is a perfect mirror, b) spatial arrangement is unimportant and shadowing and masking effects can be approximated with some simple functions. Any violations of the two assumptions would pose deviation from the microfacet theory, and thus produce variations that cannot be captured with the method.

Ashikhmin et al. [1] suggest that the generalized microfacet model is only suitable for BRDFs that are dominated by the effect of microfacet normal distribution, and not the ones that are dominated by shadowing and masking effects. Determining the kinds of materials in the latter class, both on theoretical and empirical grounds, would be an interesting avenue for future work.

While the generalized microfacet distribution is a powerful tool to represent complex BRDFs, it is not easy for an artist to edit these distributions. However, we observe that our estimated distributions all exhibit enough smoothness to leave hope for expressive and easy-to-manipulate parametric models. Finally, since the microfacet representation is a 2D representation, it is theoretically possible to acquire the full distribution from a single image. Constructing a measurement device to achieve this task would be a logical next step to our work.

Chapter 5

Statistical Acquisition of Texture Appearance

The most common solution to model spatially-varying material is the use of photographic textures. However, plain texture mapping is only a crude approximation for most real world materials as it is only suitable for representing perfectly smooth surfaces with albedo variation. It is unable to simulate the effect of any underlying mesostructure of the material. The Bidirectional Texture Function (BTF) is a more general representation that describes the full light/view dependence of a planar surface. Unfortunately, the measurement of BTFs is a formidable task that requires the imaging of a single sample for all pairs of view and light directions. A robotic setup is typically used, and precise calibration is necessary to align all these images, a task made harder by the various moving parts.

In this chapter, we propose an efficient method to capture complex spatially-varying materials. Key to our approach is a novel reconstruction technique that takes sparse measurements and reconstructs a full BTF. Our reconstruction is based on the interpolation of statistical properties, thereby alleviating the need for precise registration and avoiding cross-fading artifacts that could be caused by simple data interpolation.

We will first present our acquisition setup, followed by our reconstruction technique. Next, we present a real-time implementation to render a approximate reconstruction using graphics hardware. We then demonstrate our measurement and reconstruction results, together with a validation using publicly available data sets. We conclude with a discussion

of limitations and future directions for extensions.

5.1 Acquisition

Acquiring a full 6D BTF is a difficult problem. For a material patch, two degrees of freedom are required for both the light and view directions. To cover the light/view hemisphere, robot arms are often used to either move or rotate the sample, the light source or the camera [38, 97, 51]. As a result, acquiring a single material often takes several hours or more. Recent setups use multiple cameras/lights [77] or sophisticated optical system [41] to speed up the acquisition, but these devices are difficult and expensive to build.

We observe that the view and light dimensions are not completely symmetrical in terms of BTF measurement. While calibration errors of the view and light positions both translate to errors in the angular domain, view miscalibration also leads to spatial misalignment of the texture. When the BTF data is used directly during rendering, spatial misalignment will lead to undesirable blurring of the texture. In many cases, spatial blurring leads to more objectionable artifacts than reflectance blurring in the angular domain.

Our reconstruction technique allows us to avoid the view alignment problem. We only require the images of the frontal view of the material sample to be aligned, which is easy to achieve with fixed camera and material samples, and a moving light source. The images with different view directions do not need to be registered with the frontal view. Thus, we can combine multiple views per images by using multiple samples of the material placed at different angles (see Fig. 5-1), which greatly accelerates acquisition. This multiplexing technique exploits the high resolution of current digital camera, and is previously only possible for BRDF measurements (Section 2.3.1).

5.1.1 Implementation

Our current setup (Figure 5-1) captures 13 views of a material at a time using different samples at various orientations. The camera is fixed and the light direction is sampled by moving a light source. We paste a number of planar patches of the material onto square backing boards with known dimensions, which are then positioned to form a pyramid-

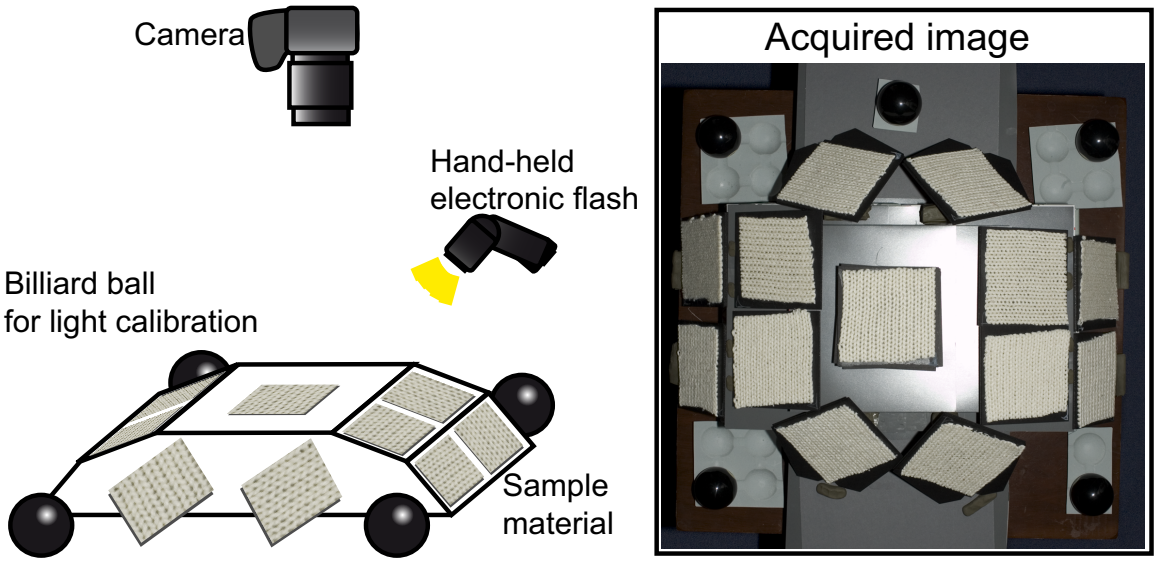

Figure 5-1: Acquisition setup - both the camera and the measured target are fixed, and a handheld wireless flash is used as the light source. During measurement, the user moves the flash source around to roughly cover all possible directions, and remotely triggers the camera shutter to take pictures.

like target. The arrangement provides 4 views at about 30° incidence angle, 8 views at about 60° and the top view. We use an 8 megapixel digital SLR camera with a hand-held electronic flash. The camera is set up on a tripod at a fixed position about 1 meter above the measurement target, and the size of the top-view patch in the image is roughly 500×500 pixels. We put specular spheres (billiard balls) around the target for light position estimation. The user holds the flash directed at the target from various directions, and a remote control is used to trigger the camera. We take about 100 pictures for each material in about $10 - 15$ minutes.

The camera is calibrated automatically using several images of a checkerboard [122]. For each sequence, the user manually marks the position of the corners of each square backing boards. With the known dimensions of the boards, we compute the position and orientation of each patch using a least-square optimization. The flash source is small, and we approximate it as a point source. For each image, the mirror peak on each specular sphere is located automatically, if available, and the light source position is estimated with a least-square fit.

For each measured image, we resample each lit patch into a texture image at a resolution

of 512×512. As the extent of a single patch is relatively small, we assume the light and view directions (ω_i, ω_o) are constant across each patch. The resampled texture is then normalized by the estimated irradiance to form a 2D spatial slice of the BTF. Even though the flash power is manually set to constant power, in practice the flash output can deviate from the specified power significantly. We put diffuse gray cards next to the pyramid to get a reliable estimate of the flash power. For cases where the gray cards are shadowed, the mean flash power is assumed. To further compensate for the flash power estimation error, we compute the average BRDF of the measured data by averaging over each normalized texture. We then smooth this BRDF and use it to re-scale our measured data so that it is consistent with the smoothed BRDF. In practice, the BRDF blurring does not degrade the quality of measurement as our ability to measure specular materials is limited by the sparse sampling of the view.

5.1.2 Alternative Setup

While the acquisition setup we use is relatively simple to build, different setups could be potentially employed with our reconstruction technique as long as they provide aligned images for at least one view direction. Our setup is based on the assumption that multiple samples of the material is easy to gather. If we are limited to a single material sample, we would need to measure the sample multiple times at different orientations to gather different views. However, we can modify the setup to maintain efficiency. Since only one material sample is measured at a time in this case, the requirement on camera resolution is decreased and a video camera can be used instead. Combined with a continuous light source (instead of a flash), we can measure the material at a much higher frame rate than the original setup, compensating for the higher number of captured frames required.

5.2 Reconstruction

The central idea of our reconstruction is to use the actual images only from the top view and to characterize the visual properties of other views using alignment-insensitive statistics. We reconstruct a BTF one texture at a time, for each pair of view-light directions. A texture

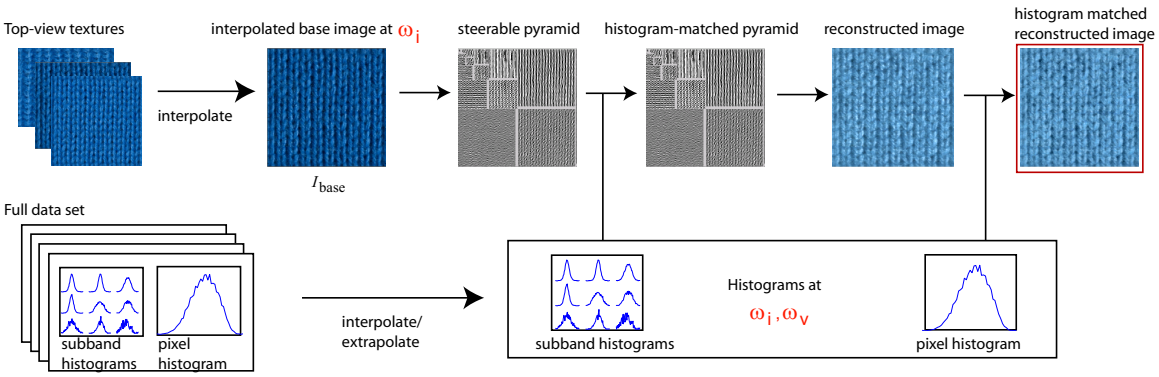

Figure 5-2: Reconstruction pipeline for producing a texture at light/view direction ω_i, ω_v. From the set of top-view textures, we interpolate the neighbors of ω_i to produce a base texture I_{base}. From the full data set comprising all light/view combinations, we locate the closest neighbors of (ω_i, ω_v). We interpolate (and potentially extrapolate) the subband histograms and pixel histogram to form the desired histograms of our target texture. The base texture I_{base} is decomposed into a multi-scale oriented pyramid, and each subband is matched to the desired histogram. We then collapse the pyramid to reconstruct the image, which is then matched to the desired pixel histogram. This produces our final texture.

is reconstructed in two steps (Fig. 5-2). We first use our set of top-view images, which are assumed to be aligned, and interpolate them linearly based on the light direction to obtain a *base texture*. This texture contains appropriate shadowing effects but might exhibit cross-fading artifacts and does not include view-dependent effects such as masking, BRDF, and asperity scattering. Our second step improves interpolation quality and reproduces these effects by enforcing statistics corresponding to the appropriate view-light directions.

5.2.1 Histogram Statistics

Before presenting our full reconstruction pipeline, we discuss our choice of statistics to characterize material appearance. We observe that the pixel histogram encodes variation in the color distribution due to effects such as shadowing and masking. Fig. 5-3(a) shows a BTF slice of the measured knitwear at $[\omega_v = (60,0), \omega_i = (60,180)]$.[1] Fig. 5-3(b) shows the *base texture* from the top-view with approximately the same light direction ω_i. We observe prominent shadowing due to the low elevation of the light in the base texture, but the effect is significantly reduced when viewed from the side in Fig. 5-3(a), both due to masking and

[1] We represent direction vectors ω in spherical coordinates (θ, ϕ), where θ and ϕ are the incidence and azimuth angles respectively.

(a) Target texture (b) Top-view base texture (c) Pixel histogram matched (d) Full reconstruction (e) Texture synthesis from (a)

Figure 5-3: Multi-scale statistics transfer. We seek to reproduce the statistics of a target texture (a) starting from a top-view base texture (b). Note how our full technique (d) improves the reconstruction quality beyond pixel histogram matching (c). Without the spatial structure provided by the base texture, texture synthesis technique by Heeger and Bergen [43] (e) is unable to reproduce the target texture.

the increased scattering path length. By matching the pixel histogram of the base image to the target image, we can recover the overall distribution of intensity, and as a result shadows are mostly eliminated (Fig. 5-3(c)). This pixel histogram matching technique has been proposed by Dana and Nayar [21]. However, pixel histogram matching does little to the structure of the image as it is insensitive to multi-scale effects such as blurring.

To further improve the appearance transfer, we employ the steerable pyramid in an approach similar to Heeger and Bergen[43]. The base and target image are both decomposed into an image pyramid with multiple scaled and oriented subbands, and the coefficient histogram of each subband is matched independently. Both pixel histograms and pyramid coefficient histograms can be computed without image registration and are well suited for our goal. The subband histogram matching captures effects such as the fuzziness that some materials exhibit at grazing angle. In Fig. 5-3(d), we first transfer the subband histograms from the target image to the base image, followed by the pixel histogram matching.

Heeger and Bergen[43] generate new textures starting from noise and iteratively enforcing the histogram of pixel color and pyramid coefficients. As a result, their method is generally unable to reproduce extended spatial structure. Fig. 5-3(e) shows the result when we directly apply their texture synthesis technique with five iterations using statistics from Fig. 5-3(a). In contrast, we use the multi-scale statistics to *modify* the base texture obtained for a given light direction, which already contains a very similar spatial structure to the target.

Dana and Nayar's work [20, 22] is closely related to our reconstruction technique. They study BTF under varying light/view configurations, and propose analytical histogram and correlation models in the special case of random isotropic surfaces with Lambertian reflectance. They also propose BTF synthesis based on histogram transfers [21]. The top-view image with correct lighting is first synthesized and then the image is transformed to an arbitrary view through pixel histograms transfer. They showed that the technique works well for a sample Lambertian material with gaussian height distributions. Our reconstruction technique extends this idea to a wider class of materials by using multi-scale statistics. In particular, we show that the idea of pyramid matching traditionally used to synthesize textures from white noise [43] can be adapted to modify a base texture and enforce the most salient visual variation due to a material's appearance. In addition, our reconstruction scheme is specifically designed for sparsely sampled data and allows for meaningful histogram interpolations and extrapolations.

5.2.2 Data Preprocess

To prepare our data for reconstruction, we pre-process each captured texture patch to gather the required statistics. We apply the steerable pyramid transform [102, 103], using Simoncelli's MATLAB implementation [101]. The transform decomposes the image into spatial frequency subbands at multiple orientations. We decompose the image into 4 orientations and all levels of power-of-two scale. To prepare the data for histogram matching during texture generation, we compute the cumulative density function (CDF) of the coefficients for each subband. The CDF maps the range of coefficients to the interval $[0, 1]$. We compute the CDF of the image pixels before decomposition as well. We then compute the inverse of each CDF and store them. Since a CDF is a monotonic function, the inverse can be computed simply by a table lookup and linear interpolation. At this stage, all original patches not from the top view can be discarded.

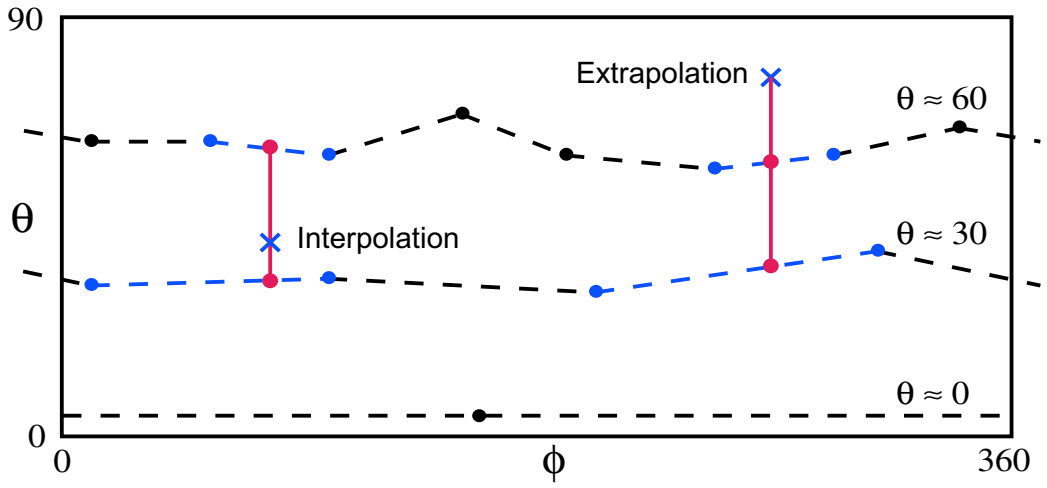

Figure 5-4: View interpolation/extrapolation in spherical coordinates. The view directions are grouped into classes with similar incidence angle θ and form rings on the hemisphere. With this semi-uniform structure, the view direction can be interpolated or extrapolated in a bilinear fashion.

5.2.3 Texture Generation

To reconstruct the full BTF, we reconstruct 2D image slices at discretely sampled view/light directions. The pipeline for reconstruction is shown in Fig. 5-2. To generate the texture at a particular light/view configuration (ω_i, ω_v), we first find neighbors of ω_i in the set of top-view textures. We then gather all the textures from the top view, and we project the light direction of each texture in the unit hemisphere onto the unit disk. We perform a Delaunay triangulation of the set, and find the three neighboring textures of ω_i by searching the containing triangle. If ω_i is outside the convex hull, we project ω_i radially onto the convex hull. We blend the corresponding textures using the barycentric weights to form I_{base}, which roughly corresponds to the texture lit from the desired direction ω_i but viewed from the top.

To obtain the desired histogram statistics, we need to find the close neighbors and determine the appropriate interpolation weights. As the set of view and light directions are separable ($2D \times 2D$), we can perform a two-level interpolation: 1) first we find the weights for view interpolation, 2) for each neighboring view, we find the interpolation weights for the neighboring light directions. As the view direction is assumed to be sparsely sampled

and roughly structured, we use a different interpolation scheme to avoid skewed triangles from triangulations, and to allow for meaningful extrapolation. We represent the view directions in spherical coordinates (θ, ϕ) and group them into classes based on the incidence angle θ (Fig. 5-4). For our setup, we have three classes $\theta \approx 0, \theta \approx 30$ and $\theta \approx 60$. We join the view directions within each class with a polyline, and we assume the polylines do not cross each other. These polylines divide the spherical domain into rings, with which a bilinear interpolation/extrapolation can be well-defined. The second-level interpolation of the light direction for each view is performed using the same strategy as the base texture.

We interpolate the pixel and subband inverse CDFs according to the computed blending weights. As proposed by Matusik et al. [71], interpolation of inverse CDFs gives much more natural transitions than interpolating the histograms themselves. Also, inverse CDFs can be readily extrapolated with negative weights that sum to one, while directly extrapolating histograms can lead to invalid negative density.

We decompose I_{base} into the same steerable pyramid described in Section 5.2.2. Then we match the histogram of each subband using the corresponding blended inverse CDFs, with the routine `match_histogram`:

```
match_histogram( subband, subband_blended_icdf )
    cdf = compute_cdf ( subband )
    for i = 1 to  { number of coefficients in subband }
        pct = lookup( cdf, subband[i] )
        matched[i] = lookup( subband_blended_icdf, pct )
    end
```

Next, we collapse the pyramid to recover an image I'_{base}. Finally, we match the pixel histogram of I'_{base} to the interpolated pixel histogram to produce our final reconstructed texture.

Note that in the texture synthesis technique proposed by Heeger and Bergen [43], multiple iterations of the subband and pixel histogram matching are required to synthesize the final texture. In our reconstruction, since our starting point is a base texture instead of random noise, we find that a single iteration is enough for all our measurements.

5.2.4 Texture Tiling

An optional final step to the reconstruction pipeline is to make the reconstructed BTF tileable. In other words, there should be no visible seams when multiple copies of the texture is placed adjacent to each other. While this is not a requirement for the reproduction of texture appearance, it is a nice property to ensure easy mapping on general surfaces.

To make the textures tileable, we adopt a semi-automatic approach. First we assume a rectangular sub-patch of the original texture is tileable. Such rectangle can be specified by its location and size, as well as a rotation angle θ : $R_\theta(x, y, width, height)$. To find the best rectangle in terms of tiling ability, we first define a border size (b_x, b_y) for overlaps. The objective function is defined as the sum of errors in overlap regions - for the horizontal overlap, we define the overlap error as the sum of squared difference between the strips $R_\theta(x, y, b_x, height)$ and $R_\theta(x + width, y, b_x, height)$. The vertical overlap error is computed similarly. To compute the total error for a particular candidate rectangle, we sum the overlap errors over a sparse subset of all textures in the BTF.

To find the best rectangular sub-patch, the user supplies initial guess to the ideal rectangle by inspection. Next, we do a brute-force search in the neighborhood of the specified parameters to find the optimal sub-patch with the lowest overlap error. Once the optimal parameters are found, we go through each texture of the BTF and resample the rectangular sub-patch to produce the tileable texture. We also replace the border strips by an average strip: e.g. for horizontal overlap, we replace $R_\theta(x, y, b_x, height)$ with the average of $R_\theta(x, y, b_x, height)$ and $R_\theta(x + width, y, b_x, height)$ to ensure smooth transition when the texture is not perfectly tileable.

Our texture tiling technique can be seen as a simplified version of the image quilting work by Efros and Freeman [34]. In their work an optimal boundary is computed in the overlap region to avoid cross-blending artifacts. While our approach has worked well for most materials we measured, it is limited to cases when the original texture is periodic. Minor discrepancies can be smoothed out by the border strips averaging. However, tiling of general aperiodic texture is difficult and beyond the scope of our work.

5.3 Hardware Implementation

By reconstructing the full BTF as a set of images, we can employ any of the existing BTF compression/rendering techniques (see Section 2.5.2). Meanwhile, our statistical characterization of the BTF can also be viewed as a compact representation of a BTF. In this section, we present a method to render from this representation directly using graphics hardware, without generating and compressing the full BTF.

To adopt our reconstruction pipeline to graphics hardware, a number of changes is necessary. First of all, instead of reconstructing the entire texture at a time, we need to be able to generate single texel from different view and light independently. As a result, subband histogram transfer and pyramid reconstruction is not feasible. To compensate for the loss of multi-scale statistics transfer during render time, we precompute and store reconstructed textures for several views in addition to the top view. Pixel histogram matching also needs to operate on a full texture generally since it depends on the CDF of the full texture. We can, however, remove that requirement by precomputing the CDFs. Since the CDFs are very compact (1D functions), we can afford to precompute them for different light and view directions at a high sampling rate.

5.3.1 Data Preparation

Three sets of data are required for our hardware implementation: a sparse set of precomputed reconstructed textures, a dense set of pixel CDFs from interpolated textures, and a dense set of inverse CDFs from target statistics. We use the method described in Section 5.2.3 to reconstruct a sparse set of texture images (in practice 5 views $\times$ 13 lights). To precompute the CDFs, we sample the view and light directions at a denser rate (81 views $\times$ 81 lights), and blend our sparse texture set to form an interpolated image. The histogram of the interpolated image is computed and converted to the form of CDF. Finally, to precompute the target inverse CDFs, we employ the interpolation/extrapolation scheme described in Section 5.2.3. The sets of texture images, CDFs and inverse CDFs are each packed into 2D textures, which can then be sent to the graphics hardware.

5.3.2 Real-time Rendering

The rendering routine is implemented as a pixel shader on a programmable graphics hardware. First of all, a color value is interpolated from the precomputed set of textures, at a particular texture location. Next, we lookup the CDF to find the percentile of the interpolated value. The precomputed CDFs allow us to compute the texel percentile without having the full interpolated texture. Finally, we use this percentile to lookup the inverse CDFs to complete the histogram matching step. The pseudo code for the pixel shader is as follows:

```
function btf_shader( light, view, uv,  textures, cdfs, icdfs )
    value = interpolate_sparse( textures, light, view, uv )
    pct = interpolate_dense( cdfs, light, view, value )
    final = interpolate_dense( icdfs, light, view, pct )
    return final
```

5.4 Results

We evaluate our approach in two different ways. First, we perform new measurements using the simple setup described in Section 5.1, and show that we are able to capture important visual characteristics of realistic textures. Second, we use publicly-available BTF databases measured using robotic setups as ground truth and perform our reconstruction using a subset of the data.

5.4.1 Acquisition and reconstruction

We have measured 16 materials including different kinds of carpets and fabrics and reconstructed the full BTF from the measurements; rendered images are shown in Fig. 5-5. The cloth geometry used in the rendered images is generated through standard cloth simulation without considering physical properties of the rendered materials. As a result, the amount of folding and stretching may not be consistent with the real materials. In addition, our texture representation inherits the limitation of BTFs: silhouettes are not captured, and effects

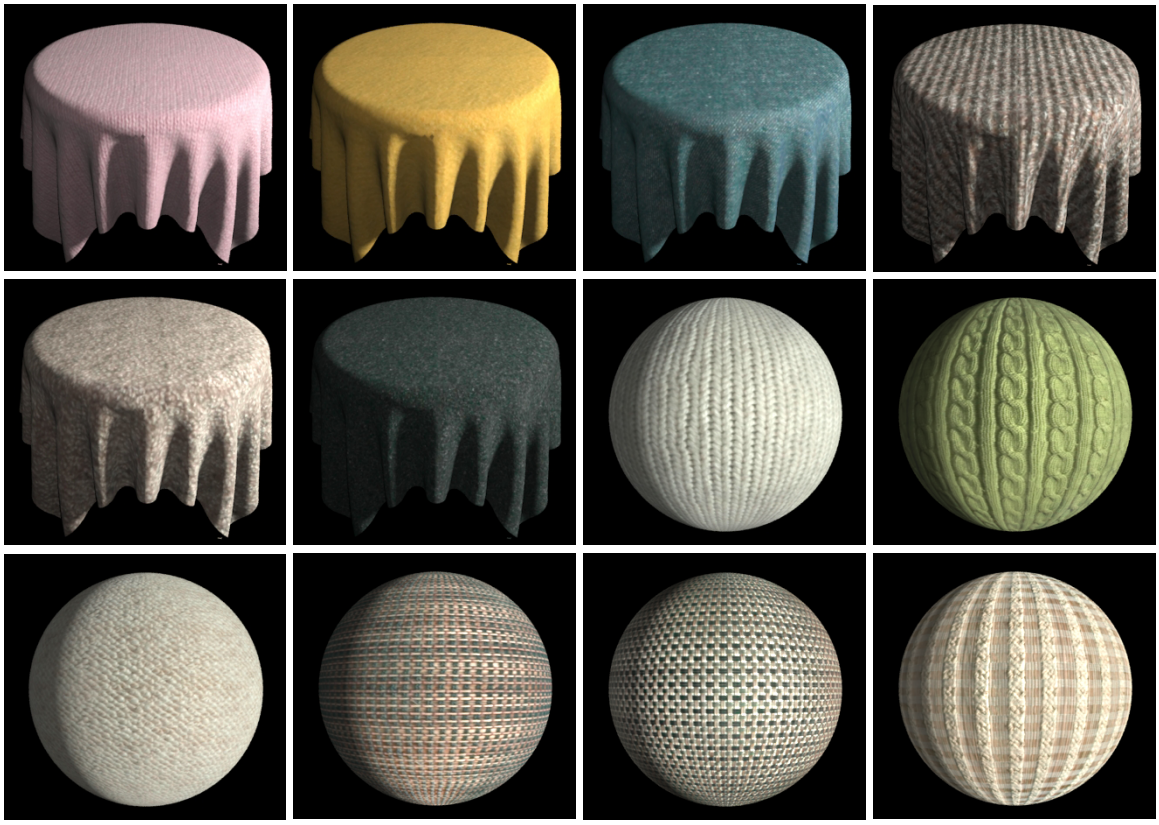

Figure 5-5: Rendered images of the measured materials. First row: pink-knitwear, fleece, denim, carpet-2, second row: carpet-3, velcro, knitwear-1, green-knitwear and third row: carpet-1,pattern-3,pattern-2 and pattern-1.

due to surface curvature are ignored. Nonetheless, in most cases our reconstruction is able to reproduce the visual quality of real materials faithfully.

Acquisition time is dependent on the sampling density of the light directions; in practice, we capture about 100 images for a material in about $10-15$ minutes. Next the images are processed and individual patches (512×512 pixels) are resampled. This processing takes $20-30$ minutes on a P4 3.0GHz PC and about 800 slices of the BTF is typically captured.

We decompose each texture slice of the captured BTF into a steerable pyramid as described in Section 5.2, and we compute and store a histogram for each subband, in addition to the pixel histogram of the image. This step takes about $2-3$ hours on our PC. Texture reconstruction takes about 20 seconds per image, and it takes 36 hours to reconstruct a full BTF with 81×81 views and lights. Our reconstruction is implemented using MATLAB

and could be greatly optimized.

5.4.2 Validation

To further validate our reconstruction method, we compare our results against a number of materials in the Bonn BTF database [97]. Each material in the Bonn database is measured at 81 light directions $\times 81$ view directions, for a total of 6561 texture images. To test the capability of our reconstruction, we pick 13 views roughly corresponding to our acquisition setup, including 8 views at 60° incidence, 4 at 30° and the top view. We use all 81 light directions for each of the 13 views. The input to our reconstruction thus includes the 81 light-dependent textures from the top view, and 13×81 sets of pixel and subband histograms. This is roughly equal to the number of samples we collect in our acquisition, despite the fact that our sampling directions are not as uniform.

We reconstruct the BTF at each of the original 6561 direction pairs with our technique. The first two columns of Fig. 5-6 shows the comparison between the reconstruction and original BTF for the materials corduroy, wool and proposte.

We also compare our reconstruction with other approximation techniques. The third column of Fig. 5-6 shows renderings when only the top-view textures are used, without any statistics transfer. Notice that this can be seen as an upper bound for the approximation quality of view-independent methods, e.g. [63]. The fourth column uses a single texture (light and view both from the top), modulated by the average BRDF computed from the original Bonn data (i.e. 6561 samples). We do the same comparison for our measured materials *knitwear-1* and *green-knitwear* in Fig. 5-7. Note the complex visual features not reproduced in the texture-mapped and top-view versions, notably the reduction of the shadowed texels due to masking, and how the texture becomes more blurry at grazing angles, revealing the fluffiness of the materials. Interestingly, we have observed that, in the absence of context, casual observers often prefer the top-view version because it exhibits more contrast and sharpness. This is a well-known bias in image quality evaluation, e.g. [46]. We however emphasize that this excessive sharpness results in hyper-realistic images that look artificial in complex scenes. Contrast/sharpness reduction is an important effect that many

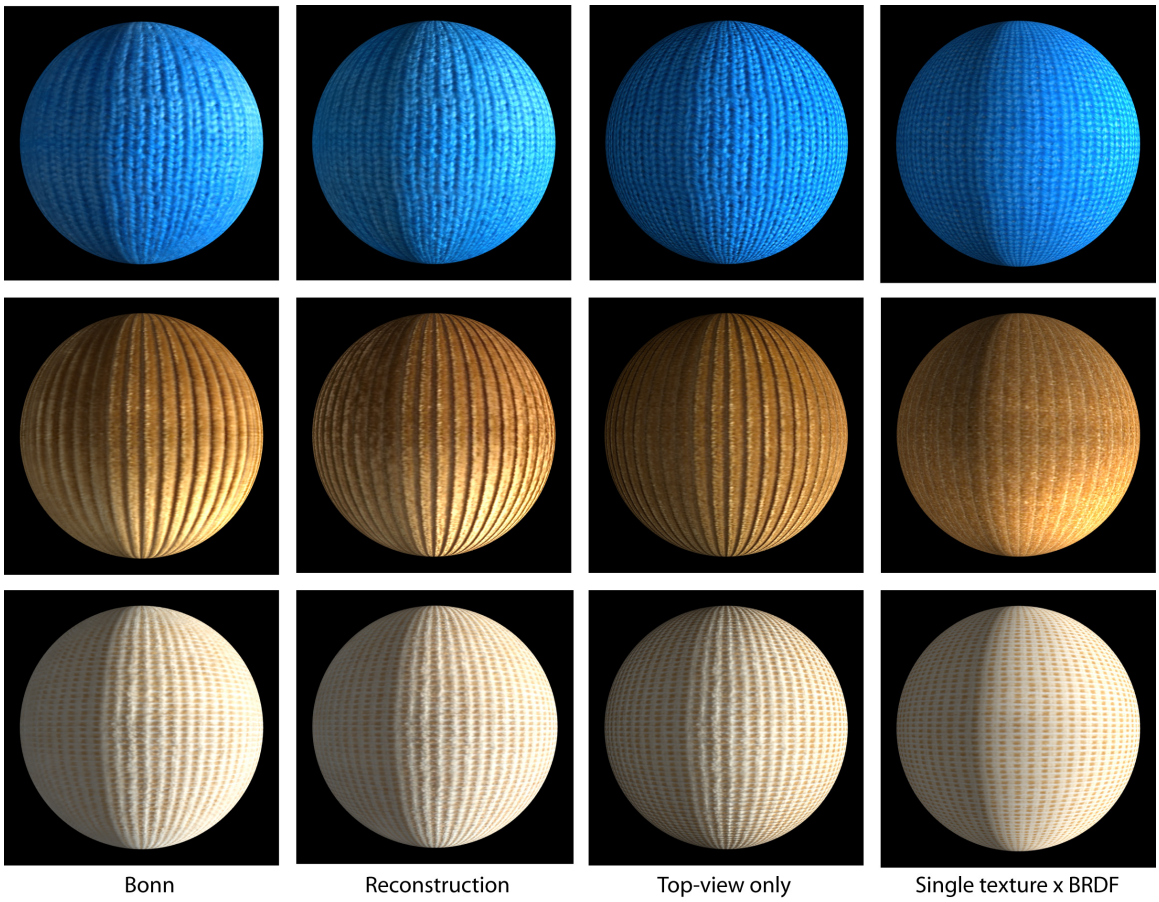

Figure 5-6: With 3 materials from the Bonn BTF database (Wool, Corduroy and Proposte), we compare our reconstruction against the original data. The scene is lit with a main directional light from the right, and a second directional light from the front. The third column shows the results when only the top-view textures are used, while the fourth column shows a single texture (top view, top light), modulated by the per-slice BRDF averaged over the texture.

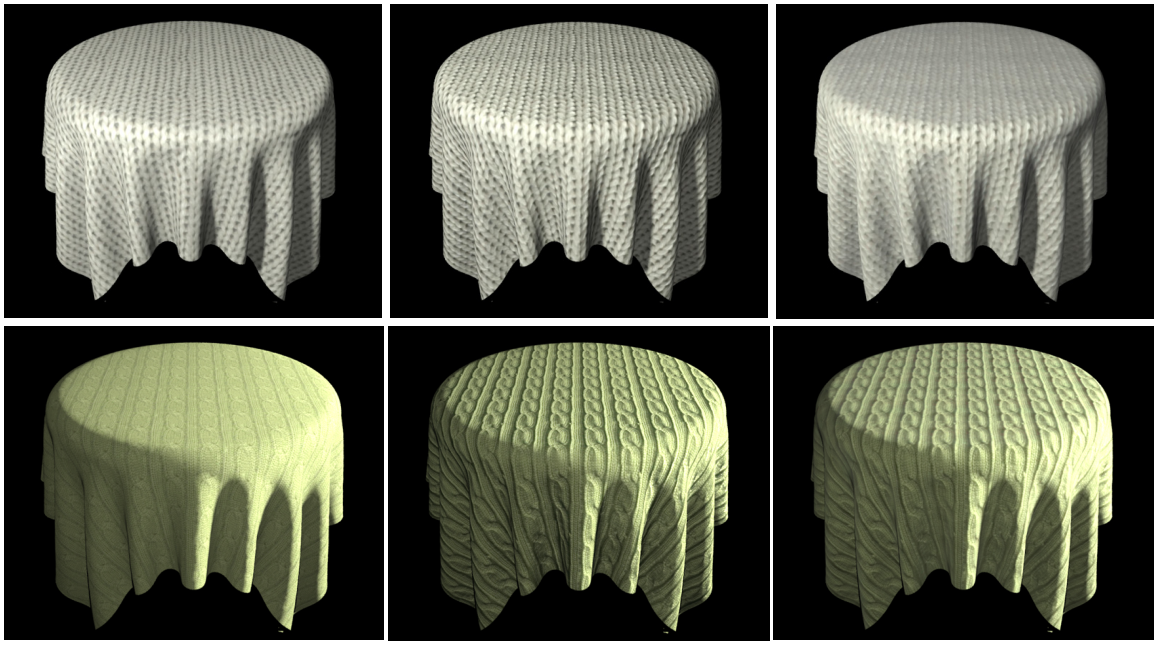

Figure 5-7: Comparing approximations to the measured materials knitwear-1 and green-knitwear. First column: single texture modulated by acquired BRDF, second column: light-varying textures from top view, and third column: our reconstruction.

real materials exhibit and that it is particularly critical to reproduce fuzzy materials like fabric. In addition reduced contrast due to masking is also important in more solid materials such as tree bark and plaster [88].

5.4.3 Statistical characterization

In Figs. 5-3 and 5-8, we show results of our reconstruction for a few example slices viewed from 60° incidence. Our reconstruction provides good visual matches to the target images in terms of brightness, color, contrast, blurriness and other subtle multi-scale effects. Notice that these side-view images are rectified to the normal view for display: for $\omega_v = (60, 0)$, the width of the texture would be halved when used for actual rendering. As a result, any discrepancy along the horizontal dimension would be diminished.

Our method captures effects such as the Fresnel term to a first-order approximation. As described in Section 5.2.3, we extrapolate the histograms for slices beyond 60 degrees, and in practice, it is able to provide plausible increase of the brightness at grazing angle (Figure 5-9). It is in particular important that we extrapolate the *inverse cumulative*

Figure 5-8: Texture reconstruction at $\omega_v = (60, 0)$. Row 1: Carpet from Koudelka et al. [51], Row 2: Measured material *carpet-1*.

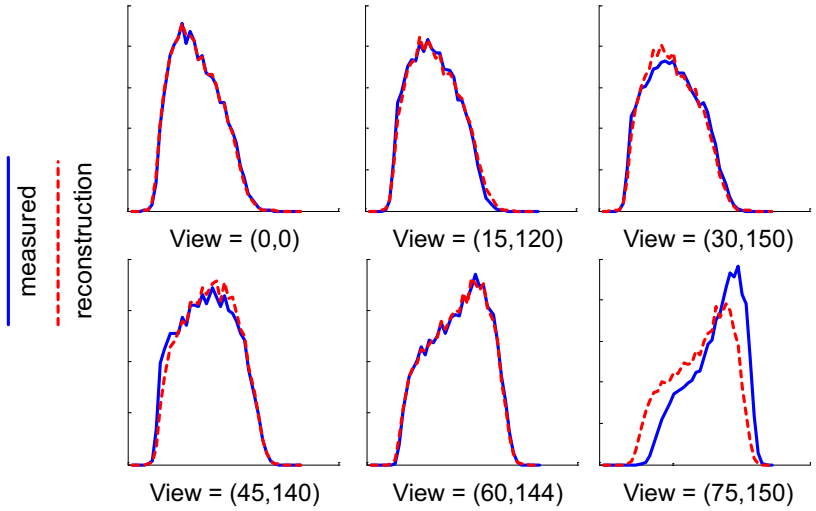

Figure 5-9: Comparing the interpolated pixel histograms to the Bonn measurement of the material proposte. Fixing the light direction at $(75, 0)$, we compare histograms for six viewing directions. All views except (0,0) are not present in the reconstruction data set. The (75,150) view is extrapolated as our most inclined views are at 60 degrees.

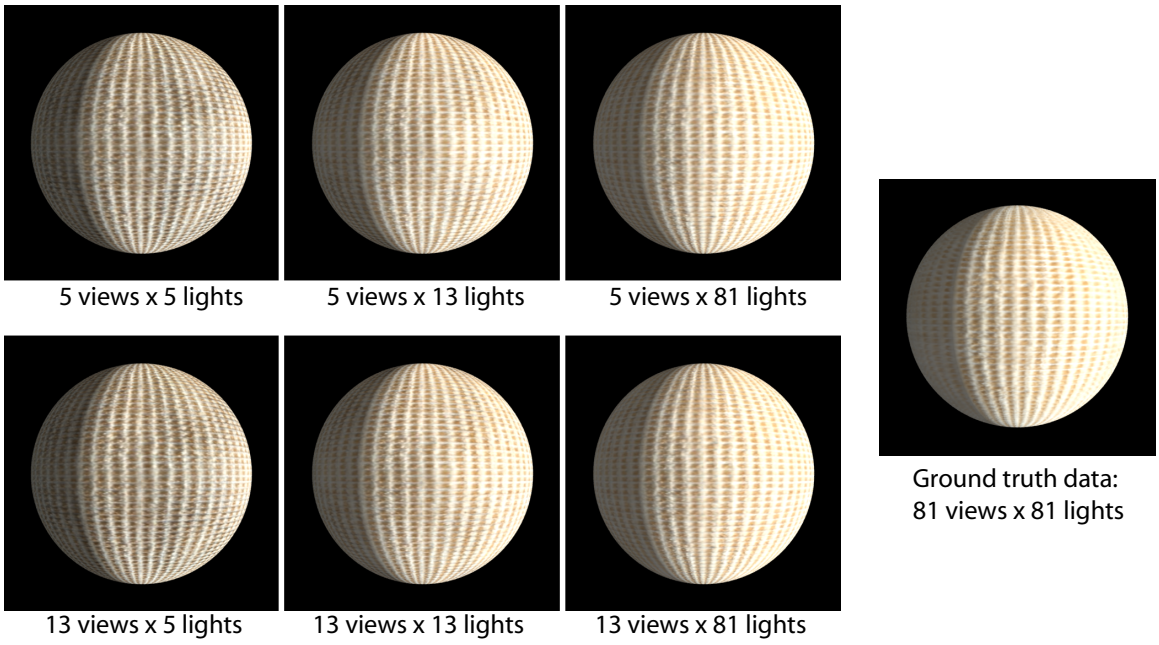

Figure 5-10: Reconstruction of the proposte material at different sampling densities, compared to the full resolution data (81 views x 81 lights) from Bonn BTF Database.

histograms, which faithfully increases brightness when appropriate and avoids numerical problems such as invalid histograms. However, because we perform linear extrapolation, we can underestimate the increase of brightness when the function is concave (Figure 5-9). Unfortunately, we have found it difficult to reliably quantify the extrapolation errors, since available measurements are typically unreliable near grazing angle (e.g. the 75 degree measurements of the Bonn dataset).

5.4.4 Sampling Density

In our measurement setup, the light direction can be sampled more densely by simply taking more pictures with additional flash positions. In contrast, it is more difficult to increase the view sampling as it will involve extra material patches. Our design choice is largely based on the assumption that light-dependent variations are more important than the view-dependent variations.

To verify this assumption, we validate against the measured data from the Bonn BTF Database. We simulate the choice of variable sampling densities for both light and view by

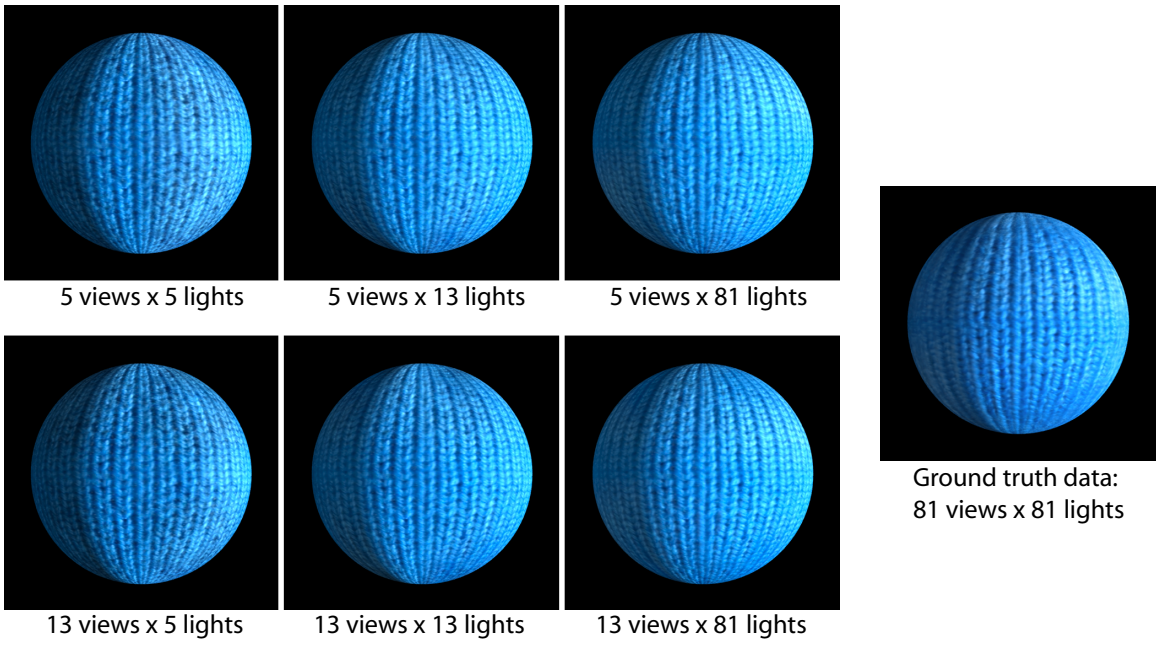

Figure 5-11: Reconstruction of *wool* at different sampling densities, compared to the full resolution data (81 views x 81 lights) from Bonn BTF Database.

selecting different subsets of the measured data and reconstruct the full BTF with our technique. Results for proposte, wool, and impalla are shown in Figures 5-10, 5-11, and 5-12. We observe that our reconstruction quality is much more sensitive to the number of lights than the number of views. In addition, since our views are interpolated through statistics, low sampling rate of the view directions does not introduce cross-blending artifacts. In Figure 5-13, we show the results when the same input images are used for rendering through direct interpolation. At low sampling rates, the BTF is significantly blurred spatially under the direct interpolation method.

5.5 Limitations

Our acquisition is limited to materials without sharp specularities due to the sparse sampling of the view directions. Strong parallax effect is also difficult to reproduce with histogram statistics as spatial structure is not directly encoded. For example, the Lego material measured by Koudelka et al. [51] exhibits strong parallax due to its relatively steep depth, while having an otherwise simple geometry that lacks significant scattering: it is the worst-

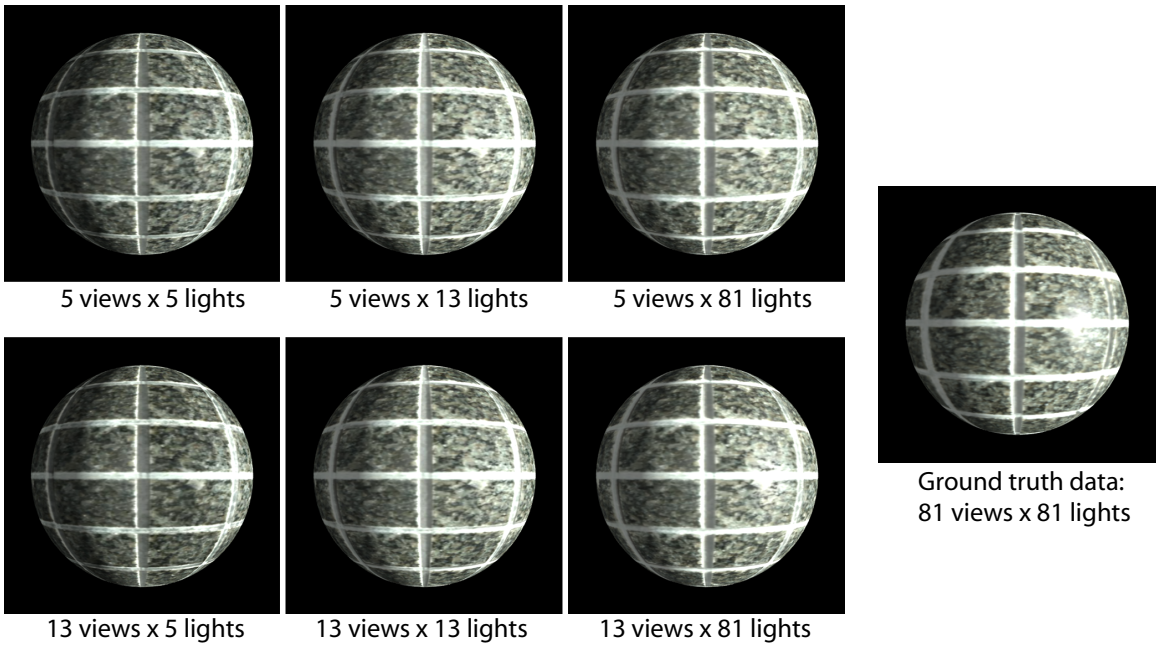

Figure 5-12: Reconstruction of the impalla material at different sampling densities, compared to the full resolution data (81 views x 81 lights) from Bonn BTF Database.

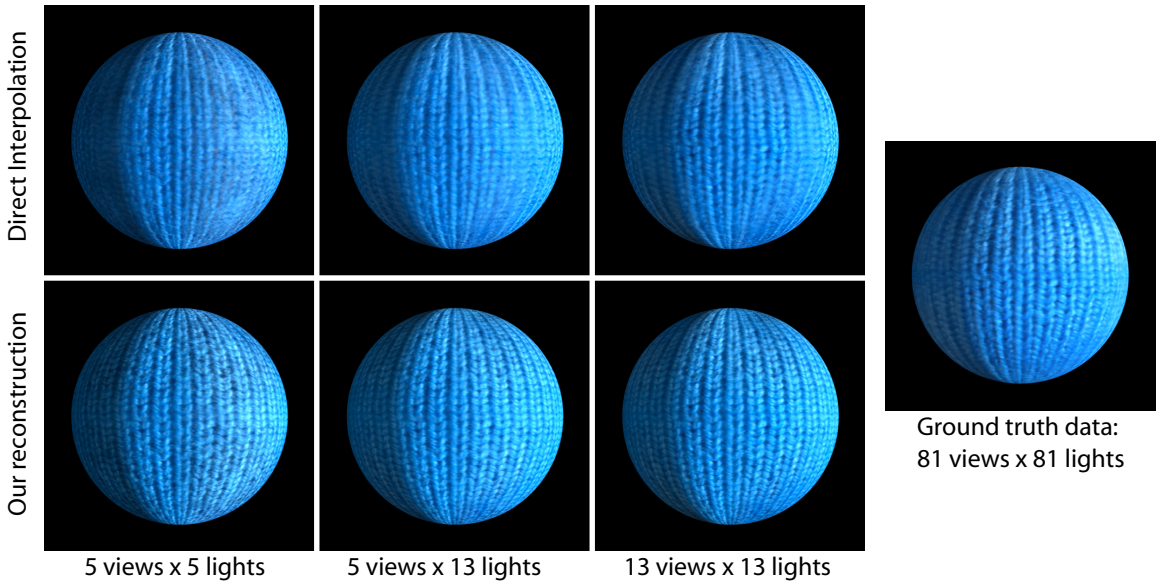

Figure 5-13: Comparing our reconstruction technique against direct interpolation at low sampling densities (Wool BTF).

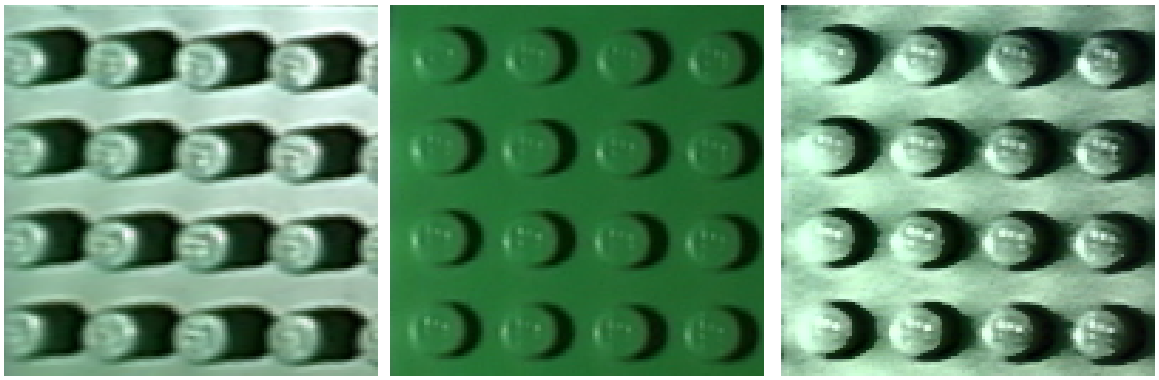

(a) Target texture (b) Base image (c) Reconstruction

Figure 5-14: Failure case for the Lego BTF [$\omega_v = (60, 0)$]. Our reconstruction is unable to reproduce the disparity, but the shadow elongation is partially captured.

case scenario for our approach. In Fig. 5-14(a) we show a slice of the BTF viewed at 60° incidence. Compared to the base image from the top view, the bumps are offset due to disparity, while shadows are stretched due to perspective. Our reconstruction is unable to capture the disparity, but its multi-scale nature allows it to partially reproduce the elongation of shadows. In Fig. 5-15 we compare the original and the reconstructed Lego BTF rendered on a sphere. Note that the light fall-off near the texture boundary in both the Lego BTF and the sponge BTF shown in the video are artifacts from the original measurements. In summary, while we consider the Lego material as a failure case for our approach, the reconstruction still looks surprisingly good and exhibits more 3D effects than a flat texture, in particular in an animated scene.

5.6 Summary

We have presented a simple method to acquire the appearance of materials such as fabric and knitwear that present rich spatial and angular variation. The main contribution is a reconstruction algorithm that generates a full bidirectional texture function (BTF) from a sparse set of measurements. Different views of the material sample do not need to be aligned because we characterize view-dependent effects using alignment-insensitive statistics, namely marginal and multi-scale histograms. Removing the requirement of alignment

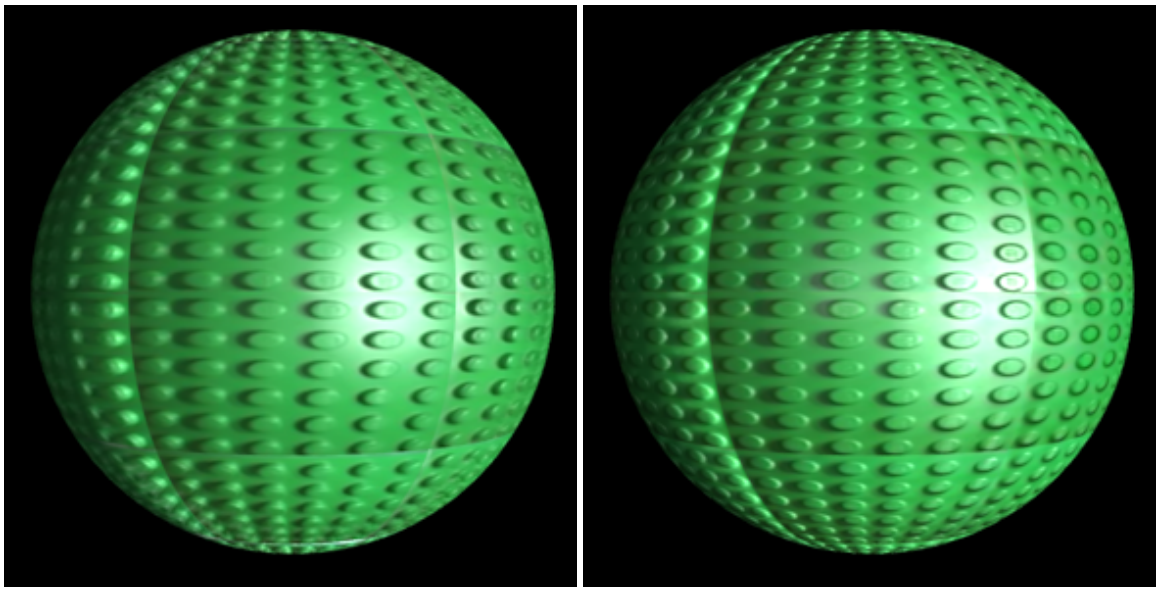

Figure 5-15: Failure case for the Lego BTF: original vs our reconstruction.

allows dramatic reduction in acquisition cost, and we have built one such setup. We also propose an alternative setup where a video camera is used with only one material sample.

As we have limited resolution in the view dimension, it is clear that our technique cannot capture high-frequency effects such as highly specular materials. Our statistical characterization does not handle the geometric effect of parallax but it reproduces some of its effects such as masking. Our statistical reconstruction tends to work best on materials with complex spatial structure (e.g. wool, proposte), as the high-frequency content and the statistical variation dominate the visual appearance. For such materials, it fits an important gap since realistic fuzzy fabric and knitwear appearance has been challenging to measure with simple means. Our statistical characterization of the BTF can also be used as a compact representation for spatially-varying reflectance. We show that our reconstruction can be adapted to graphics hardware, and it can be considered as an alternative method for BTF compression.

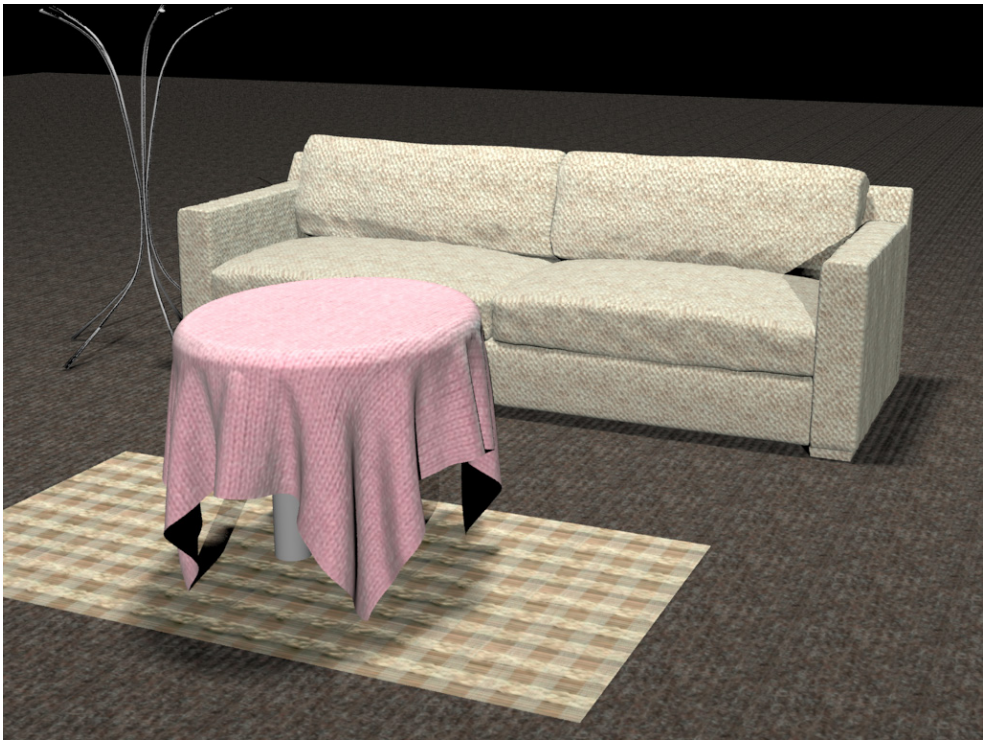

Figure 5-16: Indoor scene rendered using 4 textures acquired and reconstructed with our technique.

5.7 Discussion

Our technique inherits from the limitations of the BTF concept. While BTFs capture statistical effects of multiple scattering, they do not model the spatial component that causes blurred shadows and light bleeding. Moreover, BTFs do not capture the appearance of material silhouettes, a critical visual factor for which appropriate acquisition techniques are needed.

We believe that our statistical reconstruction has potential beyond our simple setup, notably to exploit partial views such as the one obtained with the kaleidoscope acquisition setup. We also believe that statistical material modeling can be applied to the capture of appearance from a single photograph of a general curved surface. Domain-specific knowledge or priors on BTFs might also enable a better use of the higher light sampling rate to improve reconstruction along the view direction and to better handle parallax.

Chapter 6

Image-based Navigation of Analytical BRDF Models

In most current graphics applications (e.g. movies, games), material appearance is typically specified with simple analytical models rather than measured data. There are several reasons. First of all, directly measuring real materials is still a costly and time-consuming process. More importantly, representations from direct measurements often lack control that artists require in specifying the desired appearance. Analytical models, on the other hand, are compact and efficient to compute, and provide more direct manipulation of the appearance through their various parameters. However, even using these models is not always straightforward. The parameters are often non-intuitive, and they do not generally map to appearance changes uniformly. Indeed, much experience and time are required for artists to learn how to use a particular appearance model effectively. As a result, we observe that only the simplest models are employed in most common applications .

Another difficulty often faced by artists is the choice of model. BRDF models differ in the class of materials that they are capable of representing and it takes considerable expertise to pick the right model for a desired material. Various models differ visually in subtle ways that can be important to the user but are hard to guess. In our work, we seek to blur the boundary between models by allowing users to navigate freely across a variety of models and access the nuances afforded by different analytical BRDFs.

In this chapter, we show a new technique to facilitate more uniform and thus effective

navigation of the parameter space of BRDF models. To make a parameter space uniform, first we need to have a meaningful distance metric. By leveraging recent work in material perception, we have designed a simple computational metric that roughly corresponds to perceptual spacing. Based on this metric, we present a visual navigation interface which allows the user directly observe the change in appearance. Our visual navigation interface provides the user with a set of images sampling the reflectance neighborhood of the current material selection. The user simply clicks on the image he likes best, and the reflectance corresponding to this image becomes the new selection. This way of navigation simplifies the user's task by removing the need for a mental model of the parameter mapping. Indeed, the user directly *sees* the potential effect of a parameter change before effecting it. This interface requires that the spacing of the BRDF variations shown to the user be as perceptually uniform as possible.

In addition, our metric enables us to define the notion of a neighborhood in the space of all models. In this context, our interface permits easy conversion between BRDF models, and it reveals local differences between different models. In particular, it emphasizes neighbors that do not have a similar counterpart in the current model.

In the following sections, we will first describe some related works in user design of material appearance. We then describe our computational metric based on rendered images to measure distance between BRDFs. We show that our metric produces visually uniform spacing between BRDFs, and roughly corresponds to previous perceptual studies. Next, we present a method to accelerates the metric computation by embedding precomputed images into a linear subspace. We then present our variations-based interface, and describe an effective way to navigate across different BRDF models. We conclude with a summary and provide directions for future work.

6.1 Related Work

The most popular line of research for intuitive specification in realistic rendering is inverse or goal-based rendering, e.g., [99, 49]. This approach simplifies the task of the user and removes the need for mentally intricate reverse-engineering [31]. However, it requires that

users be able to assign absolute goals and have a precise idea of what they desire to achieve.

The recently-published BRDFShop system [11] also seeks to facilitate the specification of material appearance. It builds on a painting interface and the development of an extended Ward model while we focus on the specification of parameters for existing BRDF models.

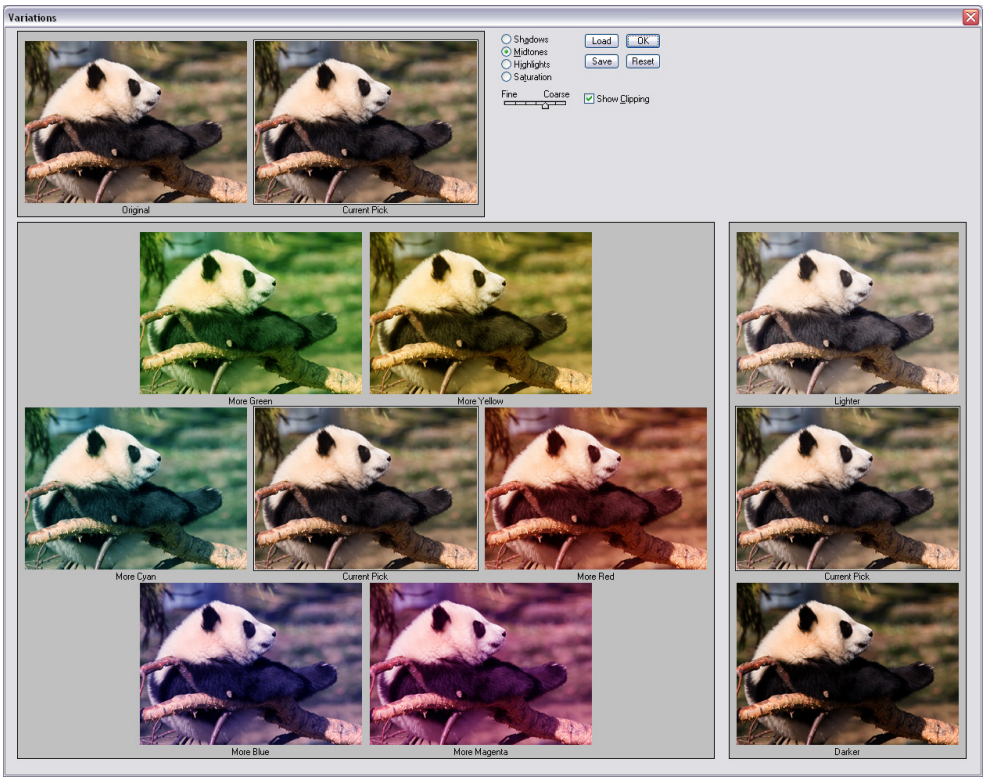

Figure 6-1: Photoshop's Variations interface.

Approaches such as design galleries [64] and the Adobe Photoshop *Variations* interface [4] present the user with various visual options. Design galleries deal with complex multidimensional and discontinuous parameter spaces and use a computational metric to maximize the dispersion of the output images. The Adobe Photoshop "Variations" interface, currently under the image/adjustments menu, facilitates the alteration of chromaticity, brightness and saturation of an image (Fig. 6-1). The interface displays the current modified image as well as several variations along the various chromaticity, saturation, or brightness axes. It makes color correction intuitive because the user does not need to form a mental model of the effect of the various parameters. The options are presented in a visual form that can be directly judged by the user based on relevant subjective criteria [4]. The success

of this interface has strongly influenced our work on BRDF specification.

6.2 BRDF Metric

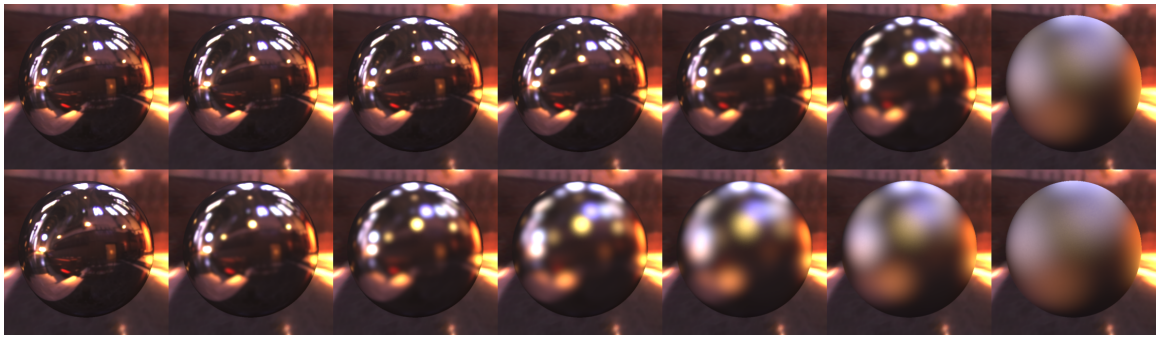

Figure 6-2: The Ward model, varying along the roughness dimension ($\alpha = 0.01$ to 0.37). Row 1: uniformly spaced according to the BRDF space L^2 metric, Row 2: uniformly spaced according to our image metric.

Comparison of BRDFs has typically been used for fitting measured data to an analytic model. One popular choice is the squared difference of the two BRDFs multiplied by the cosine of one or both incident and outgoing angles, integrated over the hemisphere [52]. However, the criteria of a good fitting (near-threshold) metric do not necessarily make it appropriate for our task, which focuses on suprathreshold perceptual uniformity. In particular, BRDF-space metrics like the one mentioned tend to overfit near the mirror direction, as the values are often orders of magnitude higher than the average of the whole BRDF. In the first row of Figure 6-2, we show the Ward model varying along the roughness axis (α) uniformly spaced according to the BRDF-space L^2 metric described above. We observe that the distance is much larger when the BRDF is sharp (mirror-like), and as a result the samples are concentrated near the sharper range. This motivates the development of a metric that better captures the visual effect of BRDFs.

6.2.1 Image-driven Metric

Our metric follows from two decisions. The most important decision is to define the distance between two BRDFs as the difference between the rendered images with the given

BRDFs under a natural environment map. The second decision concerns the precise image difference, and our choice is a compromise between perceptual motivation, simplicity, and computational efficiency.

As BRDFs are used to define object appearance in rendered images in most applications, we choose to define our metric in the image domain. While it is clear that a single image does not have enough information to uniquely define a BRDF, we leverage results from Dror et al. [27] and Fleming et al. [36] that show that a single image can capture a large part of the material characteristics if the illumination is natural (See Section 2.7.1). In this work, we choose to use the image of a sphere rendered with a given natural environment map, in practice the Grace Cathedral, courtesy of Paul Debevec. For the rest of the discussion, we will use the term *BRDF image* to represent this particular scene rendered with the corresponding BRDF. Note that, ideally, the environment map of the intended final rendering could be used at the cost of increased computation.

The image difference that we choose is the L^2 difference between the cubic roots of the RGB channels of two BRDF images. The images are represented in floats with high-dynamic range, without any tone mapping. The choice of the cubic root is inspired by the luminance mapping of the perceptually uniform CIELAB color space [35] and related to early tone mapping operators [111]. In our application, we do not use more perceptually-uniform CIELAB for computational reasons: the cubic root of the RGB channels is a slightly simpler formula that allows us to precompute BRDF images and enable approximation based on principal components analysis (Section 6.3.1). However, we advocate the use of CIELAB if precomputation is not an issue. More comprehensive visual difference predictors could also be used, but would be even more computationally expensive. In practice, we have found that our simple metric yields uniform spacing.

6.2.2 Metric Evaluation

To validate our metric with perceptual measures, we compare the reported distance of our metric to the psychophysical experiments by Pellacini et al. [85]. In their work, 11 renderings of the Ward model with different roughness are shown to human subjects, and the

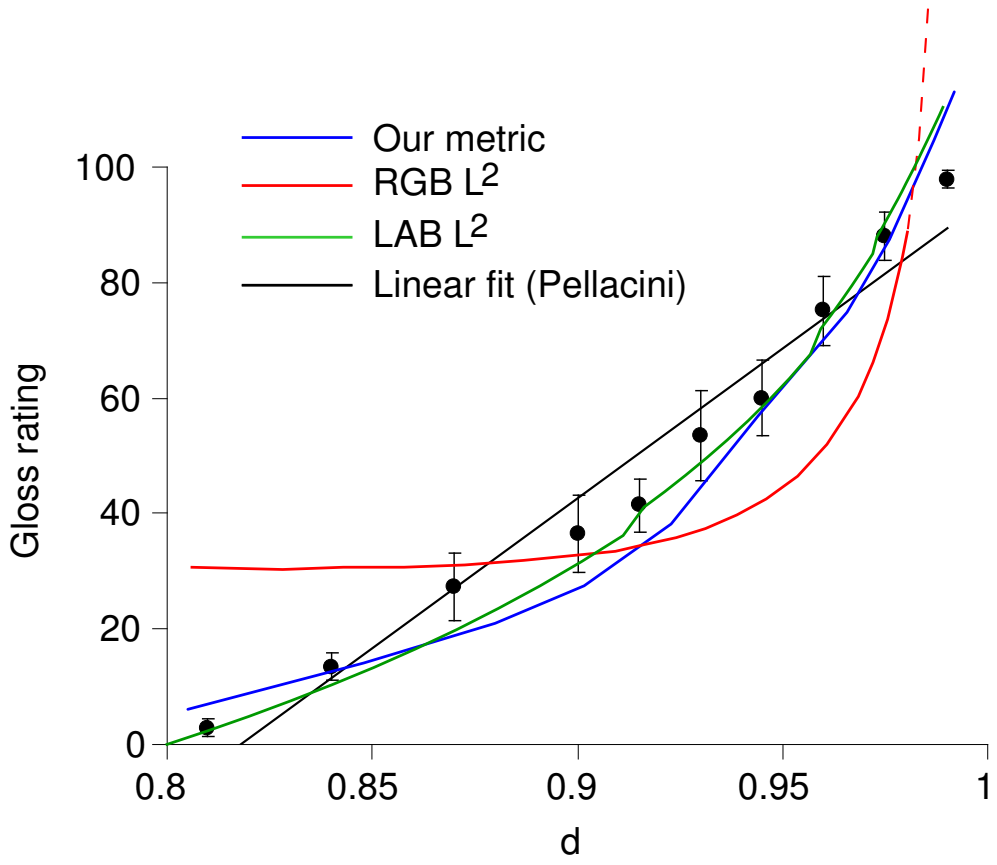

Figure 6-3: Comparing the gloss ratings reported by human subjects and the 10-steps cumulative distances reported by our image-driven metric (Blue), and the direct L^2 metric with LAB (Green) and RGB (Red). The linear fit proposed by Pellacini et al. is also shown (Black). d is the distinctness-of-image parameter defined as $d = 1 - \alpha$, where α is the roughness parameter in the original Ward model. Adapted from [85] with author's permission.

subjective *gloss ratings* are reported. The reported values together with the linear fit proposed are shown in Figure 6-3. In order to provide a comparison with their results, we sample the same parameter range in the Ward model, and the distance of each sample to its next neighbor is computed according to our metric. The cumulative distance starting from the first BRDF ($d = 0.8$) is plotted in blue in the figure. As there is an unknown calibration scale between our metric and the reported ratings, we choose the scale to give a best fit to the data (which explains why the curves do not meet at $d = 0.8$.) The same calibration is applied to two alternative metrics on the BRDF images: RGB (without cubic root) and LAB L^2 difference. Although the shape of the function for our metric is slightly different

from Pellacini et al.'s linear regression, our metric is consistent with the ratings reported. The LAB L^2 result is marginally better than our metric, while the RGB L^2 metric, in contrast, deviates significantly from both estimates and overly emphasize differences close to the mirror-like range. To summarize, the form of this image-difference metric allows us to reduce the computation cost while offering similar performance to the LAB metric.

More importantly, it is the use of rendered images as opposed to BRDF-space differences that makes our metric unique. In contrast, experiments with metric directly computed in the BRDF angular domain are unsuccessful. Even with the cubic root remapping applied directly to the BRDF data, the results are as poor as the RGB L^2 image metric. We hypothesize that the use of rendered images imposes a convolution to the BRDF[91, 32] and emphasizes different features than the original BRDF data, which better captures our perception of material.

Figure 6-4: The Lafortune model, varying along $c_z = 0.54$ to 0.58, exponent $n = 800$. Row 1: linearly spaced along c_z, Row 2: uniformly spaced according to our image-driven metric.

In Figure 6-2 we compare our image-driven metric to the BRDF-space metric defined in the angular domain. In each of the two rows the Ward model is varied over the same parameter range, and samples are chosen uniformly according to the BRDF-space metric (first row) and our image metric (second row). Another example is shown in Figure 6-4 for the Lafortune model, where we compare the uniform parameter spacing (first row) and our image metric (second row), when the parameter c_z is varied with a fixed exponent n. In both cases, the spacing with our metric is much more uniform in terms of appearance, and as such would be easier for the user to control. We have found this observation consistent

across the different parameters of the BRDF models we use.

For the Lafortune model, the c_z parameter is particularly difficult to navigate without our remapping. This is not a critical issue for its initial motivation, data fitting, but the model is more and more used in other applications because it is flexible and efficient to evaluate. When $c_z = -c_x$, the Lafortune lobe is equivalent to the Phong lobe. However, the model is very sensitive with respect to c_z when the exponent is high: even small deviation of c_z can lead to huge bias of the lobe towards normal or grazing angle. In the first row of Figure 6-4 an important region of the parameter space is skipped when c_z is varied linearly, as most of the visually interesting behavior of the model is highly concentrated near the value when $c_z = -c_x$. Our metric is able to offer a much more uniform spacing, and would allow a user to interactively explore the expressiveness of the model more easily.

Scene Dependence Next, we evaluate the sensitivity of our metric to the choice of the particular environment map. We choose a set of 25 Ward BRDFs and render the corresponding images with two additional environment maps. We compute all pairwise distances (25×24) using the image metric on each set of images, and also the BRDF-space L^2 metric for comparison. We find the scale factor between the metrics by least square fitting. We have found that over 85% of the distances for the two additional environment maps are within 20% of our reference distances. In Figure 6-5 we plot the distances of the 25 samples from a particular sample at $\alpha = 0.07$. We conclude that our metric is consistent regardless of the used environment map. We emphasize, however, that it is critical to use a *natural* environment map [29] that exhibits enough complex features. In particular, trivial environment maps such as a constant grey sphere would ignore highlight shapes. Fleming et al. [36] have shown that complex natural illumination greatly facilitates the recognition of materials. Furthermore, Dror et al. [27] have demonstrated that a number of characteristics of rendered BRDFs are invariant to a large class of real-world environment maps.

We have not evaluated our metric's dependence on scene geometry. As each pixel is compared independently, our metric only depends on the distribution of normals instead of the actual shape. For a rendered sphere, normals that are close to grazing (with respect to

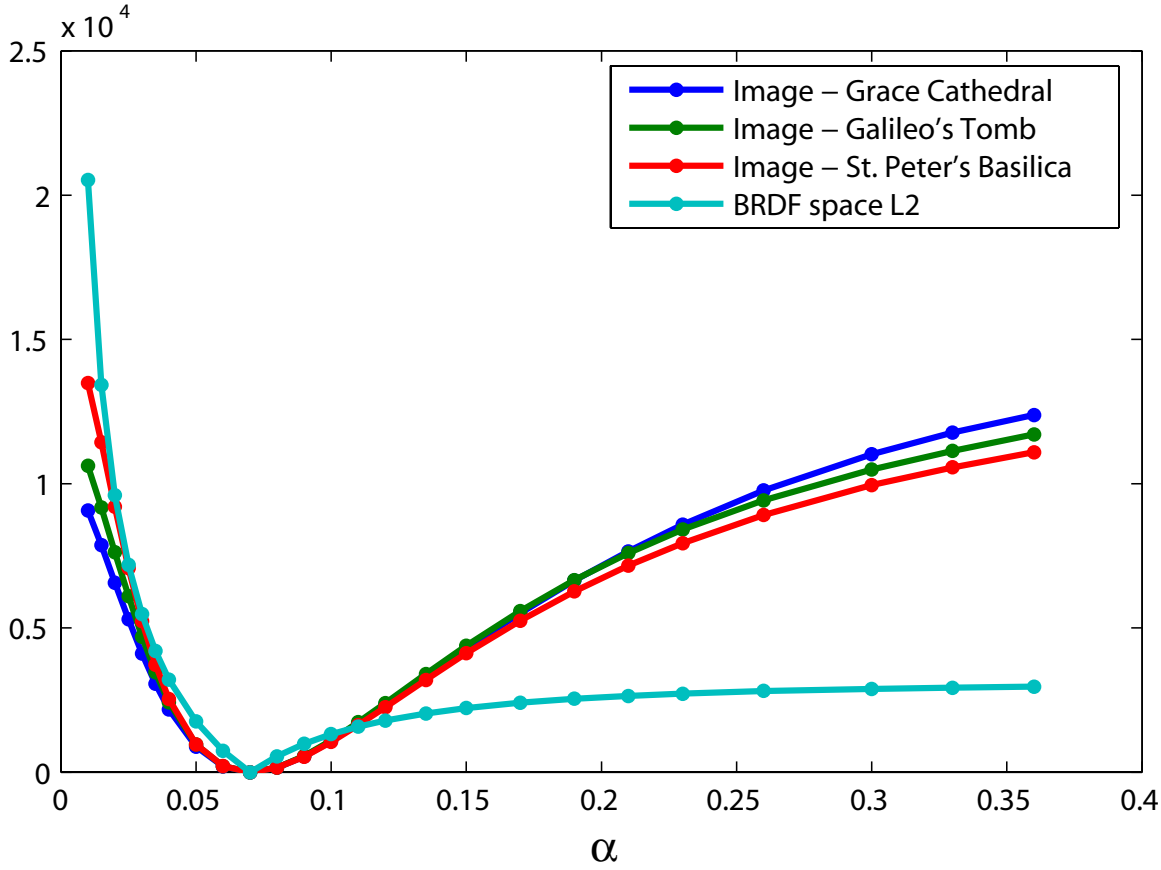

Figure 6-5: Plot of distances from the Ward BRDF at $\alpha = 0.07$ to 25 samples ranging from $\alpha = 0.01$ to 0.36. We compute our image-based metric with renderings using three different environment maps and compare to the BRDF-space L^2 metric. The distances from the different metrics are brought to the same scale by minimizing the least-square errors over all (25×24) pairwise distances.

the view direction) are sampled more sparsely. If a different bias is desired, we can alter our metric to weigh the pixels in a non-uniform way. However, the effects of using non-convex objects would require further study, as self-shadowing and other global illumination effects come into play.

Alternative Metrics A number of computational metrics have been proposed to capture human visual performances. Rushmeier et al. [94] compare different image metrics and suggest that a metric based on the contrast sensitive function (CSF) from Daly [18] performs well according to their criteria. They also suggest that geometric misalignment often hinders the performance of a metric. A number of metrics have had successful applications

in realistic image synthesis, e.g., [8, 78, 93]. Our work deals with a stimulus space that is simpler than most related work because our images are perfectly aligned and the parameter space of analytical BRDF models has a smooth effect on the rendered images. Moreover, the perceptual metrics do not provide the same acceleration that our metric allows by embedding in a linear space (Section 6.3.1).

6.3 Fast Distance Computation

Our navigation interface presents the user with choices that are distributed uniformly according to our metric. This requires on-the-fly image generation and distance computation. We can achieve fast image generation using pre-computation and PCA. However, we still need to compute many image differences and the need for image reconstruction and cubic-root non-linearity makes this computation costly. This is why we introduce a new embedding where our metric corresponds to the Euclidean distance, which affords dramatic speed up. Note that this embedding space is different from the one used for image generation. We then present our interface.

6.3.1 Embedding in a Unified Euclidean Space

In order to facilitate uniform navigation according to the metric, distances in the local neighborhood are required within a model and between models. In theory, one could compute the metric from the images on the fly, but the computational cost is too high when a number of neighbors needs to be considered. This is why we introduce a fast approximation of the metric based on the embedding in a Euclidean space with PCA. This is different from typical usage of PCA for compression purposes: we only want to compute distances and do not need to reconstruct the data points.

Observing that our metric is defined as the Euclidean distance between the cubic root of the two images, we can approximate the metric by embedding the set of post cubic root images in a low-dimensional vector space. However, in this case we can no longer ignore the linear parameters of the BRDF due to the cubic root. We can write the post cubic root image I as

$$
\begin{aligned}
I(\rho_d, \rho_s, \mathbf{p}) &= (\rho_d I_{diffuse} + \rho_s I_{specular}(\mathbf{p}))^{\frac{1}{3}} \\
&= \rho_s^{\frac{1}{3}} (\frac{\rho_d}{\rho_s} I_{diffuse} + I_{specular}(\mathbf{p}))^{\frac{1}{3}}
\end{aligned}
\tag{6.1}
$$

where $\rho_s^{\frac{1}{3}}$ can be seen as a global scale. As a result, in addition to the nonlinear parameters, we also need to sample along the $\frac{\rho_d}{\rho_s}$ parameter. We sample each of the five models with regular grids (Table 6.1). Next we seek to embed all five models into a single embedding space. Conventional PCA, which requires the data to be mean-centered, is unsuitable for our purpose, as linear scaling of the image does not correspond to a simple scaling of the embedding coordinates. We enforce the center to be zero to allow linear scaling and use uncentered PCA [47] to compute the embedding coordinates. While uncentered PCA does not have the optimality properties of standard PCA, the approximation errors still decay very quickly. Distance between the BRDFs according to our image metric can be approximated efficiently in this space using the Euclidean distance. In practice, we conservatively use the first 200 coefficients, which means that our image metric becomes the Euclidean distance for 200 dimensions, as opposed to the difference between cubic roots of 320x320 pixels, thereby achieving dramatic speed-up. This allows us to employ a more sophisticated algorithm for the neighborhood construction when navigating from one model to another (Section 6.5).

6.4 Navigation Interface

Our navigation interface is based on visual variations (Figure 6-6). The user is presented with a number of neighboring BRDF images surrounding the current one. The user can choose the desired image by clicking on it. There is only one real-valued slider in the interface – the radius of the neighborhood. A typical user navigation begins with a large radius, which will then be reduced gradually as the user converges to the desired BRDF.

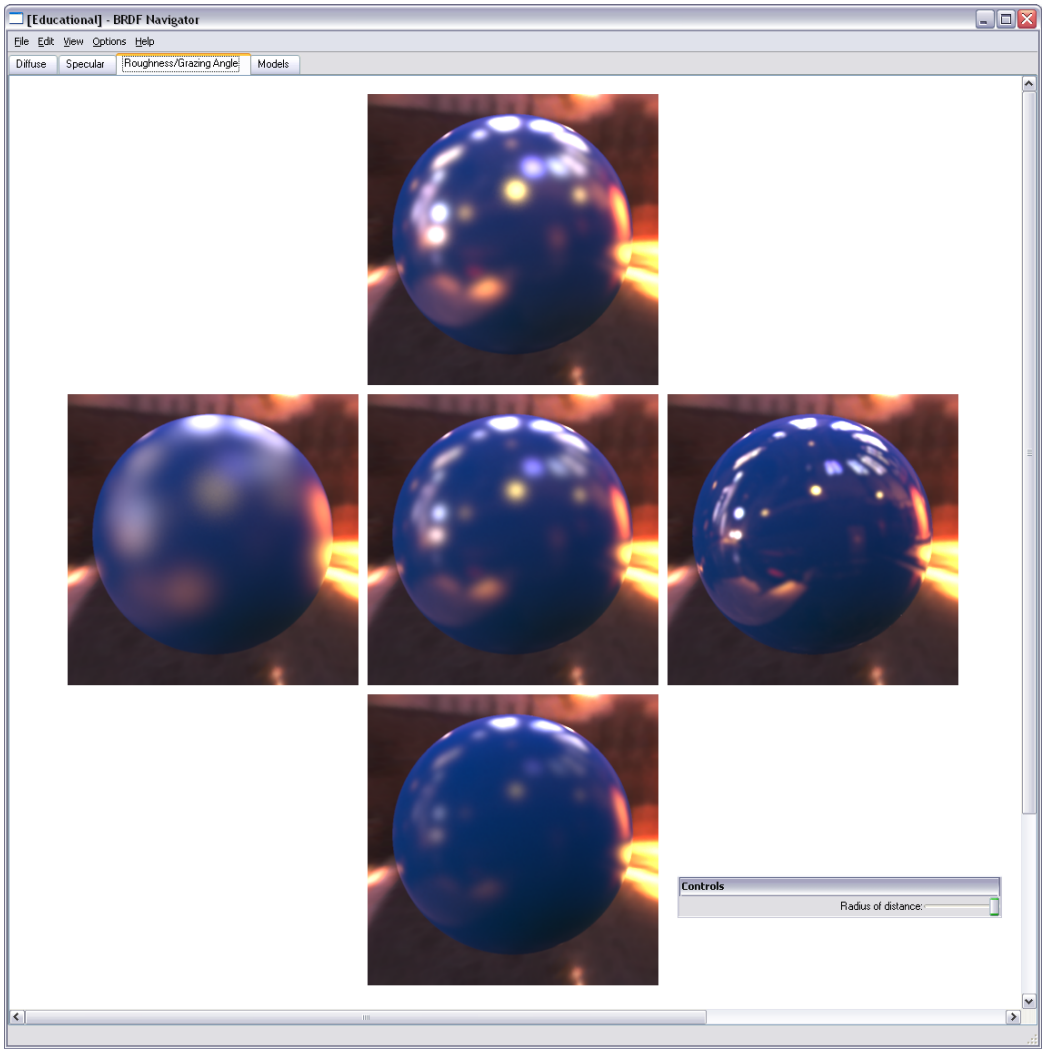

Figure 6-6: Screenshot of the navigation interface. The current model is the Cook-Torrance model, and the user is at the roughness/grazing tab. The center image is the current BRDF, and the surrounding ones are the four equidistant neighbors.

6.4.1 Image Pre-rendering

As images of arbitrary BRDFs under an environment map cannot be rendered interactively, we pre-render a set of images for each model by sampling its non-linear parameters. Linear parameters are applied on demand. The images are rendered at a resolution of 320 by 320 and principal component analysis (PCA) reduces the data size. The sampling density of each model is shown in Table 6.1. For simplicity, the isotropic single-lobe version of the Lafortune model is employed. With this assumption, we can reduce redundant degree of freedom by setting $c_x^2 + c_y^2 + c_z^2 = 1$. As $c_x = c_y$ for an isotropic Lafortune lobe, the

Model	Sampling Grid	Parameters
Ward	$(9) \times 15 \times 15$	α_x, α_y
Blinn-Phong	$(9) \times 15$	n
Cook-Torrance	$(9) \times 11 \times 13$	F_0, m
He et al.	$(9) \times 8 \times 9 \times 5$	σ, τ, n
Lafortune et al.	$(9) \times 14 \times 13$	c_z, n

Table 6.1: Sampling density of various models. The first dimension in parentheses corresponds to the (ρ_d/ρ_s) dimension, which is only applicable to the embedding calculation, but not the linear prerendering.

parameters c_z and n are enough to fully specify the lobe up to a scaling factor. The scaling factor can be ignored as we normalize the lobe approximately such that the BRDF have the same brightness when c_z or n is varied. In the current implementation, we only sample c_x in the negative range which precludes retro-reflection-like behavior. This limitation can be easily lifted by expanding our pre-rendering domain.

6.4.2 Interface Implementation

Most BRDF models have more than a few parameters, which is why we divide the neighbors into three different tabs. In the first two tabs, we show the variations in the color and intensity of the diffuse and the specular scales of the BRDF. In the third tab, we show the variations due to the other parameters depending on the chosen model (roughness, Fresnel factor, etc.), which typically affect the shape of the specular lobe. In this tab, the neighbors are chosen according to our image-driven metric described in Section 6.2.1.

The current implementation of the interface allows the user to navigate within the space of 5 different BRDF models. These include Blinn-Phong [6], Ward [118], Cook-Torrance [13], Lafortune [52] and He [42]. The anisotropic version of the Ward model is employed while the other models are enforced to be isotropic. We use the implementation of the He et al. model by Rusinkiewicz where polarization is ignored [95]. Adding more models to our interface would be straightforward.

Let the user-specified distance be ε. In the diffuse/specular color tabs, we define the neighborhood to be the sphere of radius ε in the LAB space, centered around the current color. We display 2 neighbors in the L (luminance) direction, and 8 neighbors in the a-b

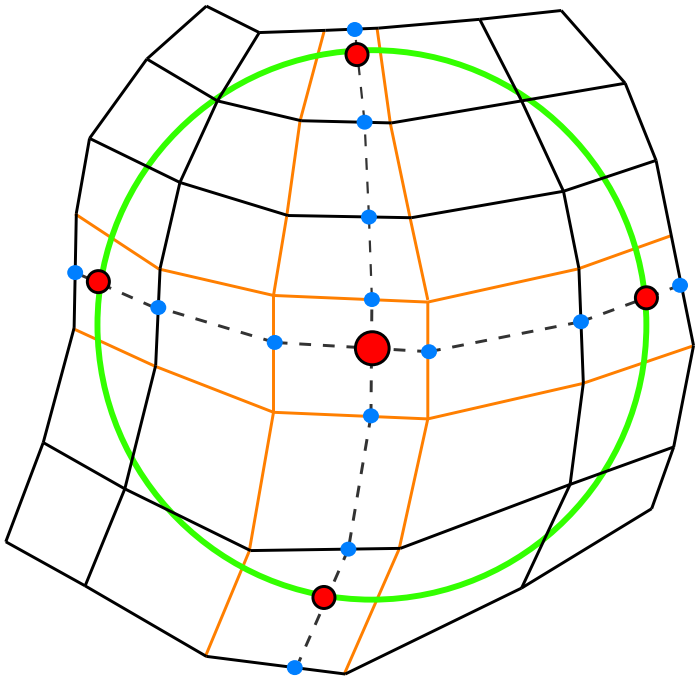

Figure 6-7: Equidistant neighbors (small red circles) are found by walking along the isoparameter lines (dashed lines). The neighbor at the desired distance is found when the segment intersect the circle (sphere). The orange lines highlight the grid samples that are queried during the search. The black lines indicate other grid samples.

chromaticity plane. We could interchangeably have used the cubic root of R, G, and B for these parameters as well, but since we have no computational limitation, we have chosen to use standard CIELAB.

The specular lobe tab (third tab) depends on the parameters used in the chosen analytical model. For these parameters we need to find equidistant neighbors according to our image-driven metric. For illustration, let us assume that the BRDF has only two parameters, and thus the sample grid for the distance precomputation is a 2D lattice (Figure 6-7). First, we find the embedding coordinate of the current BRDF using a multi-linear interpolation. Next we seek neighbors in the 4 different parametric directions. Starting from the current point (red point in the center), we walk along the chosen parameter keeping the other parameters unchanged (dashed line). As we advance along the path, the current segment is intersected with the ε circle. The distance computation for this intersection exploits the post-cubic-root PCA coefficients described above. This algorithm is generalized to higher-dimensions in our application.

6.5 Navigating Across Different Models

An important choice in material reflectance design is that of the analytical model itself. BRDF models differ in their expressiveness and the class of materials that they can represent. Our interface does not require BRDF expertise or trial-and-error of different models. In order to achieve this, we show images of material appearance that are at a certain distance from the current pick but cannot be obtained with the current analytical model. This requires two different features: we need to convert between two different analytical models, and we need to evaluate which material appearance afforded by other BRDF models cannot be achieved using the current analytical model.

We compute conversions between the various analytical models by fitting a discrete set of samples of each model to every other model. In practice, the fit is computed using the L^2 BRDF-space metric because it is more computationally efficient than the image-driven metric. In addition, we are seeking near-threshold matches, which is different from spacing BRDFs uniformly at a suprathreshold distance. Our experiments have indicated that fits obtained using the two metrics are similar, which contrasts the different suprathreshold behavior observed in Section 6.2.1. Fitting a single BRDF to a target model takes about 10 minutes on average on a single PC, but the computation needs to be performed only once. We sample the BRDF models with the same grid as for the embedding space (Table 6.1) and store all the pairwise conversions.

To construct the desired neighborhood, we first want to find neighbors in all models which are at distance ε from the current BRDF. Using the precomputed conversion and multi-linear interpolation, we can find the BRDF A' closest to the current pick A on the manifold spanned by a different analytical model (Fig. 6-8). Starting from point A' in the new model, we can search for neighbors in the same way as described in the previous section. The only difference is that the ε sphere is now centered around A instead of A'. We repeat the same search for each model pair.

For each neighbor (B',C') we look up the best-fit BRDF in the current model and compute the distance $(d_{B'},d_{C'})$. We remove neighbors that are less than a threshold distance δ from the current model. We empirically choose $\delta = 0.3\varepsilon$.

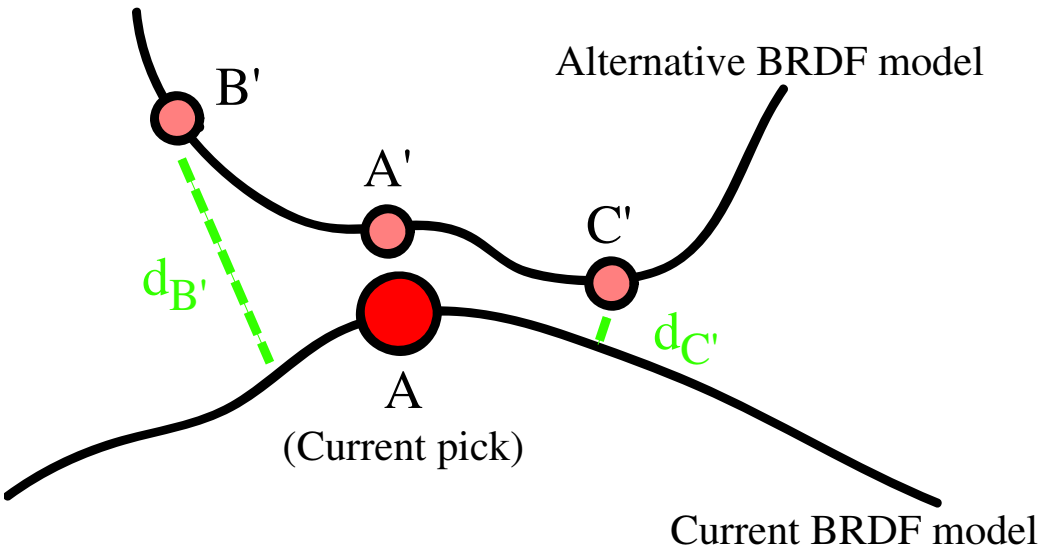

Figure 6-8: Illustration of the manifolds spanned by two analytical BRDF models in an abstract unified BRDF space. Point on the black curves represent instance of two BRDF models. Given a current BRDF, we want to find BRDFs on an alternative model that are close but cannot be represented by the current analytical model. We wish to propose materials such as B' because its distance $d_{B'}$ to the current model is large. See text for detail.

As a final step, we further trim the number of neighbors by clustering. This is necessary for cases when the number of neighbors is too large to display in the interface. In this case, we would like to show neighbors that are most different from each other. A standard solution is to use k-means clustering on the set of neighbors represented by their embedding vectors. A single neighbor is then chosen randomly from each cluster to form the neighborhood. While the k-means algorithm, in general, does not give the global optimum, the method works well in our application. Note that a large number of metric computations happen in this phase and our low-dimensional embedding is crucial for reducing the computation time. In fact, the update time to move to a new selection is less than 0.7 seconds in the worst case.

This solution can be seen as a simplified way of navigating in the manifold spanned by all BRDF models in the spirit of Matusik et al.'s work on data-driven BRDFs [68]. Our manifold-hopping solution allows us to implicitly navigate such a global manifold without explicitly building it. Topological and metric information about this manifold is encoded by the conversion and distance information.

6.6 Summary

In this work we have proposed a new distance metric for BRDFs. Each BRDF is represented by a corresponding rendered image under a natural illumination environment. We show that while this metric is not directly derived from psychophysics measurements, it reasonably reflects the visual differences between materials in a uniform way. Given this metric, we have built a user interface for navigating in the BRDF space. Our interface is intuitive and simple: the only non-visual parameter is the neighborhood radius. Neighbors in different parameter directions at the specified distance from the current BRDF are shown and can be selected by a mouse click. In addition, conversions between models are precomputed and the user can readily see the neighbors in other models and switch to them if desired. With our interface, the user is not required to understand the intricate complexities and differences between the different models. Instead, the user can always see the neighboring BRDFs across different models and freely jump between them.

In our experience, the ability to *previsualize* the result of the next navigation step is key to the effectiveness of the interface. In contrast to slider-based interfaces, our approach allows the user to directly see possible options and does not require creation of a mental model for the effects of various parameters.

6.7 Discussion

The dependence of our metric on an environment map deserves further study. On one hand, we believe that the environment map provides images that better represent real-world usage of BRDFs; and it has been shown that the complexity of real-world environment maps greatly facilitate the recognition of materials by humans. On the other hand, there is something arbitrary in choosing a given environment map, even though our experiments show that the metric is robust to this choice. We believe that the most exciting question is to further characterize what is special about natural environment maps using tools such as Fourier analysis, wavelets, and derivatives [29]. Once natural environment maps are better understood, one can hope to directly define a metric in the BRDF domain. Recent

findings on the signal processing interaction between illumination and BRDFs will likely be important [91, 32].

A second limitation of our current approach is dictated by the original parameterization. The different parameters of a BRDF model are not necessarily "perceptually orthogonal" and while the different choices we present around the current selection are on a perceptual circle, they might not be uniformly distributed on this circle. Pellacini et al. [85] defined new axes for the Ward model based on multi-dimensional scaling of the reported distances. It would be a bigger challenge to define new axes that are meaningful and uniform in the unified space of multiple BRDF models.

There are other avenues for future work. First, we currently only support a single specular lobe in our interface. Additional lobes increase the dimensionality of the space, and organizing them in a meaningful way is challenging. Next, our interface is limited to a predefined scene and environment. To allow for arbitrary scenes/environments, precomputed rendering techniques can be employed [92, 104]. The choice of the initial BRDF at the beginning of navigation is another important issue. An interface similar to design galleries [64] would greatly facilitate this first step and refine our k-means approach. While we have leveraged data from Pellacini's perceptual study to validate our image-driven metric, a logical next step is to perform a formal user study in order to better evaluate the effectiveness of our interface. Finally, we believe that visual interfaces such as design galleries and variations-based interfaces like the one we have presented have important applications in all areas of computer graphics, and can significantly enhance user's experience when dealing with complex parameter spaces.

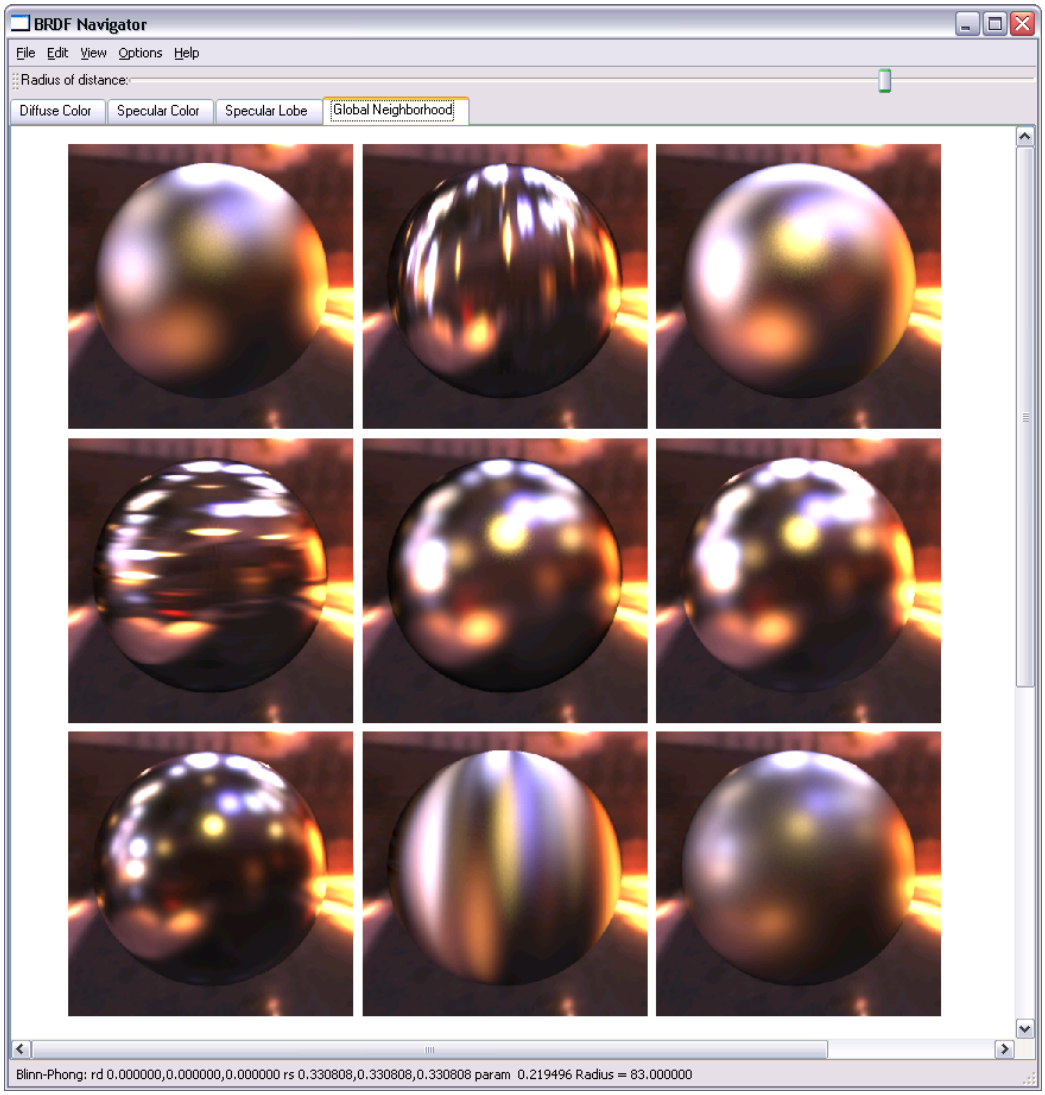

Figure 6-9: The conversion tab showing the neighbors in the union space of all models. Note that all the neighbors show some effects which are not expressible with the current model (Blinn-Phong).

Chapter 7

Conclusions and Future Work

7.1 Conclusions

Material appearance is an important element for photorealistic rendering of synthetic scenes. In most current applications, analytical models are employed as they are simple and efficient. Other standard tools for material simulation include photographic texture maps, bump maps, displacement maps, as well as custom specialized shaders. While these methods have been used to create many convincing images, they are all inherently limited in their ability to reproduce real materials due to some simplifying assumptions. This limitation is further aggravated by the difficulty of picking the right input parameters for these models. On the other hand, data-driven approaches have recently shown promising results in the research community. The main idea is that by exhaustive measurements, we can perfectly reproduce material appearance without any knowledge or assumptions about the material properties. While data-driven approaches are powerful, they require costly and time-consuming measurement, delicate calibration, as well as large storage size. As a result, data-driven approaches have only been slowly adopted by practitioners.

In this thesis, we have introduced two hybrid representations that preserve much of the flexibility of data-driven approaches while being more compact to measure and store. First, for BRDFs, we employ Ashikhmin et al.'s microfacet-based BRDF generator [1], which generates a physically plausible BRDF from an arbitrary microfacet distribution. By reversing the generation process, we show that complex anisotropic BRDFs can be repre-

sented faithfully. Anisotropic BRDF data in the 4D domain are effectively projected to a 2D function through the use of microfacet theory. The expressiveness of this representation far exceeds that of any pure analytical model, yet its low dimensionality (2D compared to 4D) allows estimation and reconstruction from sparse and non-uniform measured data.

In the second case, we combine raw measured data together with aggregate texture statistics to develop a hybrid BTF representation that facilitates efficient measurement. Existing work on BTFs all require pure brute-force measurement as no analytical model exists. We observe that for many textured materials, light-dependent variations are often more important than view-dependent variations. Based on this observation, our technique employs pure measured data for the light variations for maximal expressiveness, while multi-scale texture statistics are used for view-dependency. Our results show that this combination works remarkably well. An important consequence of such characterization is that it allows a much simplified measurement setup. Multiple physical samples can be used at the same time, allowing multiplexing of several views in a single image, and removing the need for pixel-accurate calibration. Our setup require minimal equipments and high quality BTFs can be measured in about ten minutes.

In this thesis, we have also presented an evaluation of several analytical BRDF models, comparing their ability to fit real isotropic BRDFs. To our knowledge, this is the first attempt to compare BRDF models quantitatively with a wide variety of materials. Our results suggest that using a single specular lobe, the Cook-Torrance, the Ashikhmin-Shirley, and the He models perform well for most of the isotropic BRDFs in the data set. We also illustrate that models based on the reflection direction, particularly the Lafortune model, yield systematically inferior results to the half-vector based models.

Lastly, we have also developed new tools for more effective navigation in the space of BRDF models. Our main contribution is the introduction of an image-driven metric that roughly corresponds to perceptual measures. Each BRDF is represented by a corresponding rendered image under a natural illumination environment, and the difference between two BRDFs is evaluated using a simple image metric. We show that while this metric is not directly derived from psychophysics measurements, it reasonably reflects the visual differences between materials in a uniform way. Given this metric, we have built a

variations-based interface that facilitates intuitive and uniform navigation in the space of BRDF models.

7.2 Analytical and Data-Driven Representations

A core idea in the thesis is the combination of analytical and data-driven techniques to represent material appearance. Analytical models are popular mainly due to their simplicity and compactness. Full BRDFs can be specified with a one-line formula of a few parameters that can be customized by artists. Many of them are also reasonably efficient to evaluate, allowing the use of these models in interactive applications. The main disadvantage of analytical BRDF models is that they are limited in their expressiveness to reproduce real materials. This is unavoidable since all the models make simplifying assumptions about materials so that the BRDF can be formulated with a single equation. Also, the input parameters of these models are often non-intuitive and non-uniform, making parameter specification a difficult task.

Data-driven approaches begin with exhaustive measurement of appearance quantities (BRDFs,BTFs), which are then stored in multi-dimensional tables. This tabulated representation is fully general as it does not rely on any assumption about the material properties. It provides high quality reproduction for real materials and more importantly allows us to better understand materials and evaluate approximate techniques. For example, Matusik's BRDF data set is instrumental to our quantitative evaluation of existing analytical models. Unfortunately, to reliably measure these quantities, complex equipments and delicate calibrations are required. Measurement is also a very time-consuming process, and the measured data requires large storage. Moreover, any real world measurement is in practice limited in spatial and angular resolution and polluted with noise. Finally, due to the large data size and non-parametric nature, the tabulated data is also very difficult to edit.

The microfacet distribution representation of BRDFs and the statistical characterization of BTFs presented in the thesis are both hybrid representations that combine the advantages of data-driven and analytical approaches. Through analytical modeling, high dimensional quantities are reduced into lower dimensional functions which are much easier to measure

and store. Recently, a number of factorization techniques have been introduced to decompose material appearance quantities into similar representations [113, 115, 40, 84, 55]. The common theme of these recent works is to factorize high dimensional measured data (e.g. BTF) into lower dimensional components using numerical techniques, discovering intrinsic structure or dominant components of the original data in the process. Such factorization enables efficient compression as well as more intuitive editing. There is an important difference between these factorization techniques and our work: factorizations require exhaustive measurements of high-dimensional measured data to begin with. In contrast, in our hybrid representations, predetermined structure are enforced as prior knowledge to facilitate efficient measurement - it can be seen as compression at the measurement stage. We believe that further research in this direction, possibly employing new knowledge gained from factorizations of bootstrapping data set, can significantly improve the efficiency and quality of material measurements.

7.3 Future Work

Material appearance acquisition

Our experience show that exhaustive measurement of complex materials at high resolution is difficult and time-consuming. Image-based techniques which measure multiple samples per image yield practical measurement time but our device still run in the order of hours. More importantly, the preparation of the measurement target and the calibration involved is also very tedious. Setups that use optical systems have shown promise in efficient measurement, but they tend to be large, expensive and difficult to build. The development of a portable and easy-to-use reflectance scanner is a major challenge.

Since the generalized microfacet distribution is a 2D representation, it is theoretically possible to acquire the full representation from a single image. Constructing a measurement device to achieve this task would be a logical next step to our work. More generally, it would be desirable to measure complex microfacet distributions from general curved objects that provide many normals from a single view.

Our work on BTF reconstruction shows that the combination of registered light-dependent

images and alignment-insensitive statistics based on pyramid histograms yield plausible results. However, our technique is generally unable to recreate parallax. The use of higher order joint statistics, similar to Portilla and Simoncelli [89], can possibly encode some parallax effects using joint probabilities in the local neighborhood. Our current technique is also limited to relatively diffuse textures. With our measurement setup, it is much easier to sample the light directions at a high resolution compared to the view sampling. Prior knowledge on BTFs might enable a better use of this asymmetry to improve reconstruction along the view direction.

Material appearance modeling

Currently, the most popular way to represent measured BRDF is by fitting the measurement to multiple Lafortune lobes. The main advantages of this representation are its generality and efficient evaluation. However, the generalized cosine lobe proposed by Lafortune et al. is defined with a formulation based on the mirror vector. Since the shape of these basis functions are not consistent with the shape of a typical BRDF, the quality of approximation is limited. We expect that a new set of basis functions, similar to the generalized cosine lobes but expressed with the half-vector, would yield improved fitting results with fewer number of lobes.

While our metric for BRDF fitting gives reasonable results in most cases, it does not directly correspond to perceptually important differences. A perceptually-based metric would improve the visual quality of the fits, especially when the fit is relatively poor and different choices of approximation tradeoffs are present. Applying an image-based metric such as the one described in Chapter 6 is a possible first step.

Moreover, while plenty of analytical models exist for BRDFs, there is exactly none for spatially-varying reflectance. Due to the high degree of variability in spatial variation, we do not expect a pure analytical model to work. A hybrid model similar to ours, partly based on measured data, would be a promising direction.

Perception-driven Navigation

While we have leveraged data from Pellacini's perceptual study to validate our image-driven BRDF metric, a logical next step is to perform a formal user study in order to better evaluate the effectiveness of our interface. We also believe that visual interfaces such as design galleries and variations-based interfaces similar to the one we have presented have important applications in all areas of computer graphics, and can significantly enhance user's experience when dealing with complex parameter spaces.

General Direction

Among the different sub-areas of material appearance research, we believe measurement and reproduction of spatial-variation deserve the most attention. Only very few objects in the real world are spatially homogenous. As a result, even a perfect reproduction of the true BRDF only provides limited realism in a rendered scene. While broad application of high resolution BTFs seems to be a daunting task in terms of both measurement and rendering, our statistical reconstruction suggests that many complex effects do not need to be reproduced exactly. Human observers are seldom sensitive to numerical inconsistency, as long as they are not correlated to form visible artifacts. The high frequency variations on fabrics, the presence of speckles on metals, and the noise-like patterns in car paints all offer important clues to the perception of these materials. The exact locations of these identifying patterns are not as important as their statistical characteristics. We believe further investigation of various appearance statistics can lead to more efficient material measurement and more realistic rendering at a low cost.

Bibliography

[1] M. Ashikhmin, S. Premoze, and P. Shirley. A microfacet-based BRDF generator. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 65–74. ACM Press/Addison-Wesley Publishing Co., 2000.

[2] M. Ashikhmin and P. Shirley. An anisotropic phong BRDF model. *J. Graph. Tools*, 5(2):25–32, 2000.

[3] E. Bahar and S. Chakrabarti. Full wave theory applied to computer-aided graphics for 3D objects. *IEEE Computer Graphics and Applications*, 7(7):46–60, 1987.

[4] A. Baker. Learning from Photoshop's variations tool. http://www.mergers.net/theory/20010308.html, Mar. 2001.

[5] B. G. Becker and N. L. Max. Smooth transitions between bump rendering algorithms. In *SIGGRAPH '93: Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 183–190, New York, NY, USA, 1993. ACM Press.

[6] J. F. Blinn. Models of light reflection for computer synthesized pictures. In *Proceedings of the 4th annual conference on Computer graphics and interactive techniques*, pages 192–198. ACM Press, 1977.

[7] J. F. Blinn. Simulation of wrinkled surfaces. In *Proceedings of SIGGRAPH 1978*, pages 286–292. ACM Press, 1978.

[8] M. R. Bolin and G. W. Meyer. A perceptually based adaptive sampling algorithm. In *Proceedings of SIGGRAPH 98*, Computer Graphics Proceedings, Annual Conference Series, pages 299–310, July 1998.

[9] J. S. D. Bonet. Multiresolution sampling procedure for analysis and synthesis of texture images. In *SIGGRAPH '97: Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 361–368, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co.

[10] G. Borshukov. Measured BRDF in film production-Realistic cloth appearance for "The Matrix Reloaded". *SIGGRAPH 2003 Sketches and Applications Program*, 2003.

[11] M. Colbert, S. Pattanaik, and J. Krivanek. BRDF-shop: Creating physically correct bidirectional reflectance distribution functions. *IEEE Comput. Graph. Appl.*, 26(1):30–36, 2006.

[12] R. L. Cook. Shade trees. In *SIGGRAPH '84: Proceedings of the 11th annual conference on Computer graphics and interactive techniques*, pages 223–231, New York, NY, USA, 1984. ACM Press.

[13] R. L. Cook and K. E. Torrance. A reflectance model for computer graphics. In *Proceedings of the 8th annual conference on Computer graphics and interactive techniques*, pages 307–316. ACM Press, 1981.

[14] Cornell Light Measurement Laboratory: http://www.graphics.cornell.edu/research/measure/.

[15] O. G. Cula and K. J. Dana. Compact Representation of Bidirectional Texture Functions. In *Conference on Computer Vision and Pattern Recognition (CVPR 2001)*, volume 1, pages 1041–1047, Dec. 2001.

[16] O. G. Cula and K. J. Dana. 3d texture recognition using bidirectional feature histograms. *Int. J. Comput. Vision*, 59(1):33–60, 2004.

[17] Curet: Columbia-Utrecht reflectance and texture database. http://www.cs.columbia.edu/CAVE/curet/.

[18] S. Daly. *The Visible Differences Predictor: An algorithm for the assessment of image fidelity*, pages 179–206. Digital Image and Human Vision. MIT Press, Cambridge, MA, 1993.

[19] K. Dana and J. Wang. Device for convenient measurement of spatially varying bidirectional reflectance. *J. Opt. Soc. Am. A*, 21(1):1–12, 2004.

[20] K. J. Dana and S. K. Nayar. Histogram model for 3D textures. In *CVPR '98: Proc. of the 1998 Conference on Computer Vision and Pattern Recognition*, page 618. IEEE Computer Society, 1998.

[21] K. J. Dana and S. K. Nayar. 3D textured surface modeling. In *CVPR workshop on the Integration of Appearance and Geometric Methods in Object Recognition*, 1999.

[22] K. J. Dana and S. K. Nayar. Correlation model for 3d texture. In *ICCV '99: Proc. of the International Conference on Computer Vision-Volume 2*, page 1061. IEEE Computer Society, 1999.

[23] K. J. Dana, S. K. Nayar, B. van Ginneken, and J. J. Koenderink. Reflectance and texture of real-world surfaces. Technical Report CUCS-048-96, 1996.

[24] K. J. Dana, S. K. Nayar, B. van Ginneken, and J. J. Koenderink. Reflectance and texture of real-world surfaces. In *CVPR '97: Proc. of the 1997 Conference on Computer Vision and Pattern Recognition*, page 151. IEEE Computer Society, 1997.

[25] K. Daubert, H. P. A. Lensch, W. Heidrich, and H.-P. Seidel. Efficient cloth modeling and rendering. In *Proceedings of the 12th Eurographics Workshop on Rendering Techniques*, pages 63–70. Springer-Verlag, 2001.

[26] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. In *SIGGRAPH '00: Proceedings of*

the 27th annual conference on Computer graphics and interactive techniques, pages 145–156, New York, NY, USA, 2000. ACM Press/Addison-Wesley Publishing Co.

[27] R. O. Dror, E. H. Adelson, and A. S. Willsky. Recognition of surface reflectance properties from a single image under unknown real-world illumination. In *Proceedings of the Workshop on Identifying Objects Across Variations in Lighting at CVPR*, Hawaii, Dec. 2001.

[28] R. O. Dror, E. H. Adelson, and A. S. Willsky. Surface reflectance estimation and natural illumination statistics. In *Proceedings of the Second International Workshop on Statistical and Computational Theories of Vision*, Vancouver, July 2001.

[29] R. O. Dror, T. K. Leung, E. H. Adelson, and A. S. Willsky. Statistics of real-world illumination. In *Proceedings of CVPR*, Hawaii, December 2001.

[30] A. Duer. On the Ward model for global illumination. Submitted for review, Sept. 2004.

[31] F. Durand. An invitation to discuss computer depiction. In *Proc. NPAR*, 2002.

[32] F. Durand, N. Holzschuch, C. Soler, E. Chan, and F. X. Sillion. A frequency analysis of light transport. *ACM Transactions on Graphics*, 24(3):1115–1126, Aug. 2005.

[33] P. Dutré, P. Bekaert, and K. Bala. *Advanced Global Illumination*. AK Peters, Ltd., 2003.

[34] A. A. Efros and W. T. Freeman. Image quilting for texture synthesis and transfer. *Proceedings of SIGGRAPH 2001*, pages 341–346, August 2001.

[35] M. D. Fairchild. *Color Appearance Models*. Addison-Wesley, 1998.

[36] R. W. Fleming, R. O. Dror, and E. H. Adelson. How do humans determine reflectance properties under unknown illumination? In *Proceedings of the Workshop on Identifying Objects Across Variations in Lighting at CVPR*, Hawaii, Dec. 2001.

[37] R. W. Fleming, R. O. Dror, and E. H. Adelson. Real-world illumination and the perception of surface reflectance properties. *Journal of Vision*, 3(5):347–368, 2003.

[38] R. Furukawa, H. Kawasaki, K. Ikeuchi, and M. Sakauchi. Appearance based object modeling using texture database: acquisition, compression and rendering. In *EGRW '02: Proceedings of the 13th Eurographics workshop on Rendering*, pages 257–266. Eurographics Association, 2002.

[39] M. Goesele, H. P. A. Lensch, J. Lang, C. Fuchs, and H.-P. Seidel. Disco: acquisition of translucent objects. *ACM Trans. Graph.*, 23(3):835–844, 2004.

[40] J. Gu, C.-I. Tu, R. Ramamoorthi, P. Belhumeur, W. Matusik, and S. Nayar. Time-varying surface appearance: Acquisition, modeling and rendering. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 25(3), July 2006.

[41] J. Y. Han and K. Perlin. Measuring bidirectional texture reflectance with a kaleidoscope. *ACM Trans. Graph.*, 22(3):741–748, 2003.

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

[43] D. J. Heeger and J. R. Bergen. Pyramid-based texture analysis/synthesis. In *Proceedings of SIGGRAPH 1995*, pages 229–238. ACM Press, 1995.

[44] W. Heidrich, K. Daubert, J. Kautz, and H.-P. Seidel. Illuminating micro geometry based on precomputed visibility. In *Proceedings of SIGGRAPH 2000*, pages 455–464. ACM Press/Addison-Wesley Publishing Co., 2000.

[45] H. W. Jensen, S. R. Marschner, M. Levoy, and P. Hanrahan. A practical model for subsurface light transport. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 511–518. ACM Press, 2001.

[46] G. Johnson and M. Fairchild. Sharpness rules. In *IST/SID 8th Color Imaging Conference*, pages 24–30, 2000.

[47] I. Jolliffe. *Principal Component Analysis*. Springer Verlag, New York, 2002.

[48] J. Kautz and M. D. McCool. Interactive rendering with arbitrary brdfs using separable approximations. In *Proceedings of the Eurographics Workshop on Rendering*, pages 281–292. Eurographics Association, 1999.

[49] J. K. Kawai, J. S. Painter, and M. F. Cohen. Radioptimization - goal based rendering. In *Proceedings of SIGGRAPH 93*, pages 147–154, Aug. 1993.

[50] J. J. Koenderink, A. J. van Doorn, and M. Stavridi. Bidirectional reflection distribution function expressed in terms of surface scattering modes. In *Proc. of the 4th European Conference on Computer Vision-Volume II*, pages 28–39. Springer-Verlag, 1996.

[51] M. L. Koudelka, S. Magda, P. N. Belhumeur, and D. J. Kriegman. Acquisition, compression, and synthesis of bidirectional texture functions. In *Texture 2003: Third Int. Workshop on Texture Analysis and Synthesis*, pages 59–64, Nice, France, October 2003.

[52] E. P. F. Lafortune, S.-C. Foo, K. E. Torrance, and D. P. Greenberg. Non-linear approximation of reflectance functions. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 117–126. ACM Press/Addison-Wesley Publishing Co., 1997.

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

[54] J. H. Lambert. *Photometria sive de mensure de gratibus lumi-nis, colorum umbrae*. Eberhard Klett, 1760.

[55] J. Lawrence, A. Ben-Artzi, C. DeCoro, W. Matusik, H. Pfister, R. Ramamoorthi, and S. Rusinkiewicz. Inverse shade trees for non-parametric material representation and editing. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 25(3), July 2006.

[56] H. P. A. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel. Image-based reconstruction of spatially varying materials. In *Rendering Techniques 2001: 12th Eurographics Workshop on Rendering*, pages 103–114, June 2001.

[57] T. Leung and J. Malik. On perpendicular texture: Why do we see more flowers in the distance? In *CVPR '97: Proceedings of the 1997 Conference on Computer Vision and Pattern Recognition*, pages 807–813. IEEE Computer Society, 1997.

[58] T. Leung and J. Malik. Recognizing surfaces using three-dimensional textons. In *ICCV '99: Proceedings of the International Conference on Computer Vision-Volume 2*, page 1010, Washington, DC, USA, 1999. IEEE Computer Society.

[59] T. Leung and J. Malik. Representing and Recognizing the Visual Appearance of Materials using Three-dimensional Textons. *International Journal of Computer Vision*, 43(1):29–44, 2001.

[60] X. Liu, Y. Yu, and H.-Y. Shum. Synthesizing bidirectional texture functions for real-world surfaces. In *Proceedings of SIGGRAPH 2001*, pages 97–106. ACM Press, 2001.

[61] R. Lu, A. Kappers, and J. Koenderink. Optical properties (Bidirectional Reflectance Distribution Functions) of velvet. *Applied Optics*, 37(25):5974–5984, 1998.

[62] R. Lu, A. Kappers, and J. Koenderink. Optical properties (Bidirectional Reflectance Distribution Functions) of shot fabric. *Applied Optics*, 39(31):5785–5795, 2000.

[63] T. Malzbender, D. Gelb, and H. Wolters. Polynomial texture maps. In *Proceedings of SIGGRAPH 2001*, pages 519–528. ACM Press, 2001.

[64] J. Marks, B. Andalman, P. A. Beardsley, W. Freeman, S. Gibson, J. K. Hodgins, T. Kang, B. Mirtich, H. Pfister, W. Ruml, K. Ryall, J. Seims, and S. Shieber. Design galleries: A general approach to setting parameters for computer graphics and animation. In *Proceedings of SIGGRAPH 97*, Computer Graphics Proceedings, Annual Conference Series, pages 389–400, Aug. 1997.

[65] S. Marschner, S. Westin, E. Lafortune, K. Torrance, and D. Greenberg. Image-based BRDF measurement including human skin. In *Rendering Techniques '99*, pages 139–152, 1999.

[66] S. R. Marschner, S. H. Westin, E. P. F. Lafortune, and K. E. Torrance. *Applied Optics*, 39(16):2592–2600, 2000.

[67] V. Masselus, P. Peers, P. Dutré, and Y. D. Willems. Relighting with 4d incident light fields. *ACM Trans. Graph.*, 22(3):613–620, 2003.

[68] W. Matusik, H. Pfister, M. Brand, and L. McMillan. A data-driven reflectance model. *ACM Trans. Graph.*, 22(3):759–769, 2003.

[69] W. Matusik, H. Pfister, A. Ngan, P. Beardsley, R. Ziegler, and L. McMillan. Image-based 3D photography using opacity hulls. In *SIGGRAPH '02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques*, pages 427–437, New York, NY, USA, 2002. ACM Press.

[70] W. Matusik, H. Pfister, R. Ziegler, A. Ngan, and L. McMillan. Acquisition and rendering of transparent and refractive objects. In *EGRW '02: Proceedings of the 13th Eurographics workshop on Rendering*, pages 267–278, Aire-la-Ville, Switzerland, Switzerland, 2002. Eurographics Association.

[71] W. Matusik, M. Zwicker, and F. Durand. Texture design using a simplicial complex of morphable textures. *ACM Trans. Graph.*, 24(3):787–794, 2005.

[72] N. Max. Horizon mapping: shadows for bump-mapped surfaces. In *The Visual Computer 4*, volume 2, pages 109–117, Jul 1988.

[73] D. K. McAllister, A. Lastra, and W. Heidrich. Efficient rendering of spatial bi-directional reflectance distribution functions. In *HWWS '02: Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware*, pages 79–88. Eurographics Association, 2002.

[74] M. D. McCool, J. Ang, and A. Ahmad. Homomorphic factorization of brdfs for high-performance rendering. In *SIGGRAPH '01: Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 171–178, New York, NY, USA, 2001. ACM Press.

[75] J. Meseth, G. Müller, and R. Klein. Reflectance field based real-time, high-quality rendering of bidirectional texture functions. *Computers and Graphics*, 28(1):103–112, February 2004.

[76] H. P. Moravec. 3d graphics and the wave theory. In *SIGGRAPH '81: Proceedings of the 8th annual conference on Computer graphics and interactive techniques*, pages 289–296, New York, NY, USA, 1981. ACM Press.

[77] G. Müller, G. Bendels, and R. Klein. Rapid Synchronous Acquisition of Geometry and BTF for Cultural Heritage Artefacts. In *The 6th International Symposium on Virtual Reality, Archaeology and Cultural Heritage (VAST)*, pages 13–20, November 2005.

[78] K. Myszkowski. The visible differences predictor: Applications to global illumination problems. In G. Drettakis and N. Max, editors, *Rendering Techniques '98 (Proceedings of Eurographics Rendering Workshop '98)*, pages 233–236. Springer Wien, 1998.

[79] G. Müller, J. Meseth, M. Sattler, R. Sarlette, and R. Klein. Acquisition, synthesis, and rendering of bidirectional texture functions. *Computer Graphics Forum*, 24(1):83–109, March 2005.

[80] A. Ngan, F. Durand, and W. Matusik. Experimental analysis of BRDF models. In *Proceedings of the Eurographics Symposium on Rendering*, pages 117–126. Eurographics Association, 2005.

[81] F. Nicodemus, J. Richmond, J. Hsia, I. Ginsberg, and T. Limperis. Geometric considerations and nomenclature for reflectance. Monograph 160, National Bureau of Standards (US), October 1977.

[82] NIST reference reflectometer: STARR facility. http://physics.nist.gov/.

[83] M. Oren and S. K. Nayar. Generalization of lambert's reflectance model. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 239–246. ACM Press, 1994.

[84] P. Peers, K. vom Berge, W. Matusik, R. Ramamoorthi, J. Lawrence, S. Rusinkiewicz, and P. Dutre. A compact factored representation of heterogeneous subsurface scattering. *ACM Transactions on Graphics (Proc. SIGGRAPH)*, 25(3), July 2006.

[85] F. Pellacini, J. A. Ferwerda, and D. P. Greenberg. Toward a psychophysically-based light reflection model for image synthesis. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 55–64. ACM Press/Addison-Wesley Publishing Co., 2000.

[86] M. Pharr and G. Humphreys. *Physically Based Rendering : From Theory to Implementation*. Morgan Kaufmann, 2004.

[87] B. T. Phong. Illumination for computer generated pictures. *Communications of ACM*, 18(6):311–317, 1975.

[88] S. C. Pont and J. J. Koenderink. Bidirectional texture contrast function. *International Journal of Computer Vision*, 62(1-2):17–34, 2005.

[89] J. Portilla and E. P. Simoncelli. A parametric texture model based on joint statistics of complex wavelet coefficients. *Int. J. Comput. Vision*, 40(1):49–70, 2000.

[90] P. Poulin and A. Fournier. A model for anisotropic reflection. *Computer Graphics*, 24(4):273–282, Aug. 1990.

[91] R. Ramamoorthi and P. Hanrahan. A signal-processing framework for inverse rendering. In E. Fiume, editor, *SIGGRAPH 2001, Computer Graphics Proceedings*, Annual Conference Series, pages 117–128. ACM Press / ACM SIGGRAPH, 2001.

[92] R. Ramamoorthi and P. Hanrahan. Frequency space environment map rendering. In J. Hughes, editor, *SIGGRAPH 2002 Conference Proceedings*, pages 517–526, 2002.

[93] M. Ramasubramanian, S. N. Pattanaik, and D. P. Greenberg. A perceptually based physical error metric for realistic image synthesis. In *Proceedings of SIGGRAPH 99*, pages 73–82, Aug. 1999.

[94] H. Rushmeier, G. J. Ward, C. Piatko, P. Sanders, and B. Rust. Comparing real and synthetic images: Some ideas about metrics. In P. M. Hanrahan and W. Purgathofer, editors, *Rendering Techniques '95*, Eurographics, pages 82–91. Springer-Verlag Wien New York, 1995.

[95] S. Rusinkiewicz. bv - a BRDF browser. *http://graphics.stanford.edu/~smr/brdf/bv/*.

[96] S. Rusinkiewicz. New change of variables for efficient BRDF representation. In G. Drettakis and N. Max, editors, *Rendering Techniques '98*, Eurographics, pages 11–22. Springer-Verlag Wien New York, 1998.

[97] M. Sattler, R. Sarlette, and R. Klein. Efficient and realistic visualization of cloth. In *EGRW '03: Proceedings of the 14th Eurographics workshop on Rendering*, pages 167–177. Eurographics Association, 2003.

[98] C. Schlick. An inexpensive BRDF model for physically-based rendering. *Computer Graphics Forum*, 13(3):233–246, 1994.

[99] C. Schoeneman, J. Dorsey, B. Smits, J. Arvo, and D. Greenberg. Painting with light. In J. T. Kajiya, editor, *SIGGRAPH 93 Conference Proceedings*, pages 143–146. ACM SIGGRAPH, Aug. 1993.

[100] P. Schröder and W. Sweldens. Spherical wavelets: efficiently representing functions on the sphere. In *SIGGRAPH '95: Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 161–172, New York, NY, USA, 1995. ACM Press.

[101] E. Simoncelli: The Steerable Pyramid - http://www.cns.nyu.edu/~eero/STEERPYR/.

[102] E. Simoncelli, W. Freeman, E. Adelson, and D. Heeger. Shiftable multi-scale transforms. *IEEE Transactions on Information Theory*, 38(2):587–607, 1992.

[103] E. P. Simoncelli and W. T. Freeman. The steerable pyramid: A flexible architecture for multi-scale derivative computation. In *International Conference on Image Processing*, volume 3, pages 444–447, 1995.

[104] P.-P. Sloan, J. Kautz, and J. Snyder. Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. *ACM Transactions on Graphics*, 21(3):527–536, July 2002.

[105] P.-P. J. Sloan and M. F. Cohen. Interactive horizon mapping. In *Proceedings of the Eurographics Workshop on Rendering Techniques 2000*, pages 281–286. Springer-Verlag, 2000.

[106] J. Stam. Diffraction shaders. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, pages 101–110, New York, NY, USA, 1999. ACM Press/Addison-Wesley Publishing Co.

[107] M. M. Stark, J. Arvo, and B. Smits. Barycentric parameterizations for isotropic brdfs. *IEEE Transactions on Visualization and Computer Graphics*, 11(2):126–138, 2005.

[108] F. Suykens, K. Berge, A. Lagae, and P. Dutre. Interactive rendering with bidirectional texture functions. *Comput. Graph. Forum*, 22(3):463–472, 2003.

[109] X. Tong, J. Wang, S. Lin, B. Guo, and H.-Y. Shum. Modeling and rendering of quasi-homogeneous materials. *ACM Trans. Graph.*, 24(3):1054–1061, 2005.

[110] K. Torrance and E. Sparrow. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America*, 57(9):1105–1114, September 1967.

[111] J. Tumblin and H. E. Rushmeier. Tone reproduction for realistic images. *IEEE Computer Graphics & Applications*, 13(6):42–48, Nov. 1993.

[112] B. van Ginneken, J. J. Koenderink, and K. J. Dana. Texture histograms as a function of irradiation and viewing direction. *Int. J. Comput. Vision*, 31(2-3):169–184, 1999.

[113] M. A. O. Vasilescu and D. Terzopoulos. Tensortextures: multilinear image-based rendering. *ACM Trans. Graph.*, 23(3):336–342, 2004.

[114] E. Veach. *Robust Monte Carlo Methods for Light Transport Simulation*. PhD dissertation, Stanford University, 1997.

[115] H. Wang, Q. Wu, L. Shi, Y. Yu, and N. Ahuja. Out-of-core tensor approximation of multi-dimensional matrices of visual data. *ACM Trans. Graph.*, 24(3):527–535, 2005.

[116] L. Wang, X. Wang, X. Tong, S. Lin, S. Hu, B. Guo, and H.-Y. Shum. View-dependent displacement mapping. *ACM Trans. Graph.*, 22(3):334–339, 2003.

[117] X. Wang, X. Tong, S. Lin, S. Hu, B. Guo, and H.-Y. Shum. Generalized displacement maps. In *Proceedings of the Eurographics Symposium on Rendering*, pages 227–234. Eurographics Association, 2004.

[118] G. J. Ward. Measuring and modeling anisotropic reflection. In *Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, pages 265–272. ACM Press, 1992.

[119] S. H. Westin, J. R. Arvo, and K. E. Torrance. Predicting reflectance functions from complex surfaces. In *Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, pages 255–264. ACM Press, 1992.

[120] H. B. Westlund and G. W. Meyer. Applying appearance standards to light reflection models. In ACM, editor, *SIGGRAPH 2001 Conference Proceedings, August 12–17, 2001, Los Angeles, CA*, pages 501–51., New York, NY 10036, USA, 2001. ACM Press.

[121] T. Weyrich, W. Matusik, H. Pfister, B. Bickel, C. Donner, C. Tu, J. McAndless, J. Lee, A. Ngan, H. W. Jensen, and M. Gross. Analysis of human faces using a measurement-based skin reflectance model. *ACM Trans. Graph.*, 25(3):1013–1024, 2006.

[122] Z. Zhang. A flexible new technique for camera calibration. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(11):1330–1334, 2000.