Czech Technical University in Prague
Faculty of Information Technology
Department of Theoretical Computer Science

Adaptive Measurement of Material Appearance

by

Ing. Radomír Vávra

A dissertation thesis submitted to
the Faculty of Information Technology, Czech Technical University in Prague,
in partial fulfilment of the requirements for the degree of Doctor.

Dissertation degree study programme: Informatics

Prague, April 2017

Supervisor:

Prof. Ing. Michal Haindl, DrSc.
Department of Pattern Recognition
Institute of Information Theory and Automation
Czech Academy of Sciences
Pod Vodárenskou věží 4
182 08 Prague 8
Czech Republic

Abstract

One of the ultimate challenges of computer graphics is the realistic visualization of appearance of real-world materials. The appearance can be captured by various approaches, but they are often only approximative or usually require an excessively long measurement time. Therefore, this thesis deals with the precise measurement of material appearance utilizing time-reducing adaptive methods. To better understand the behavior of material appearance, we propose an affordable setup for its instantaneous analysis that is based on an ellipsoidal reflector. Also, we study a human's ability to distinguish structure of a material in a virtual environment in dependence on an observation distance. Although, the first proposed method of adaptive measurement does not require a database of already measured materials, it is precise and very flexible. In this approach, the measured space is filled by one-dimensional continuous signals, which are sampled adaptively. We study the optimal deployment of the signals and propose an interpolation method that enables a quick reconstruction of an arbitrary value. Next, we introduce adaptive approaches that rely on the database. Our template-based methods use precomputed sampling patterns for the measurement of a new material and they achieve better results than conventional methods for more than several hundred samples. On the other hand, our minimal-sampling method achieves outstanding results for less than one hundred samples. It is based on the acquisition of a few samples for each rotation of a material around its normal. Therefore, a measurement setup can be very simple or even industrial multi-angle reflectometers can be used. Among adaptive methods, we introduce a non-adaptive image-based approach for acquisition of a huge number of material samples from a large homogeneous specimen. Also, we use gathered knowledge on material appearance to build an inexpensive setup for the rapid acquisition of approximative datasets and propose a novel method for the correct registration of multi-view images. To sum up, our approaches to analysis and measurement have great potential to improve the efficiency of current material appearance acquisition methods.

Keywords:
material appearance, BRDF, apparent BRDF, BTF, anisotropy, adaptive measurement, sparse sampling, portable setup, ellipsoidal reflector.

As a collaborator of Ing. Radomír Vávra and a co-author of his papers, I agree with Ing. Radomír Vávra's authorship of the research results, as stated in this dissertation thesis.

. . . .
Ing. Jiří Filip, Ph.D.

Acknowledgements

First of all, I would like to express my gratitude to my dissertation thesis supervisor, Prof. Ing. Michal Haindl, DrSc. He has been my mentor and helped me with numerous problems and professional advancements. Above all, I would like to thank my colleague and friend, Ing. Jiří Filip, Ph.D., who provided me the opportunity to work on his projects and has been a constant source of new interesting ideas. My research has been then supported by the Czech Science Foundation grants 103/11/0335, 14-02652S, 17-18407S.

Special thanks go to the staff of the Department of Pattern Recognition, Institute of Information Theory and Automation, Czech Academy of Sciences, who maintained a pleasant and flexible environment for my research. I would like to express special thanks to Ing. Jiří Grim, CSc., who encouraged me to wade through various mathematical expressions, which helped me to derive equations in the thesis. Finally, my greatest thanks go to my family members, including my fiancée and her family members, for their infinite patience and care.

Dedication

For Kačenka, my fiancée,
who encouraged me to apply
for a doctoral study program.

Contents

List of Figures

List of Tables

List of Algorithms

Notation

$\mathbb{Z}$ Integer numbers set
ω_i Illumination direction, $\omega_i = \{\theta_i, \varphi_i\}$
θ_i Illumination elevation angle, $\theta_i \in [0, 90°]$
φ_i Illumination azimuthal angle, $\varphi_i \in [0, 360°)$, $\varphi_i = \varphi_i + k \cdot 360°$, $\forall k \in \mathbb{Z}$
ω_v View direction, $\omega_v = \{\theta_v, \varphi_v\}$
θ_v View elevation angle, $\theta_v \in [0, 90°]$
φ_v View azimuthal angle, $\varphi_v \in [0, 360°)$, $\varphi_v = \varphi_v + k \cdot 360°$, $\forall k \in \mathbb{Z}$
$\Delta\varphi$ Azimuthal difference angle, $\Delta\varphi = \varphi_v - \varphi_i$, $\Delta\varphi \in [0, 360°)$, $\Delta\varphi = \Delta\varphi + k \cdot 360°$, $\forall k \in \mathbb{Z}$
$\vec{x}$ Coordinate axis tangential to surface or spatial coordinate
$\vec{y}$ Coordinate axis tangential to surface or spatial coordinate
N Normal vector
H Half vector, $H = \frac{\omega_i + \omega_v}{\|\omega_i + \omega_v\|}$, $H = \{\theta_h, \varphi_h\}$
θ_h Half vector elevation angle, $\theta_h \in [0, 90°]$
φ_h Half vector azimuthal angle, $\varphi_h \in [0, 360°)$, $\varphi_h = \varphi_h + k \cdot 360°$, $\forall k \in \mathbb{Z}$
D Difference vector, $D = \{\theta_d, \varphi_d\}$, $D = rot_{\vec{y}, -\theta_h} rot_{N, -\varphi_h} \omega_i$ [Rus98]
θ_d Difference vector elevation angle, $\theta_d \in [0, 90°]$
φ_d Difference vector azimuthal angle, $\varphi_d \in [0, 360°)$, $\varphi_d = \varphi_d + k \cdot 360°$, $\forall k \in \mathbb{Z}$
α Axial slice placement parameter, $\alpha \in [0, 360°)$, $\alpha = \alpha + k \cdot 360°$, $\forall k \in \mathbb{Z}$
β Diagonal slice placement parameter, $\beta \in [0, 360°)$, $\beta = \beta + k \cdot 360°$, $\forall k \in \mathbb{Z}$
λ Color spectrum or light wavelength

$\boldsymbol{b}$	Vector $\boldsymbol{b}$
b_i	the i^{th} element of vector $\boldsymbol{b}$
$\lVert\boldsymbol{b}\rVert$	Norm of vector $\boldsymbol{b}$
$\boldsymbol{A}$	Matrix $\boldsymbol{A}$
$A_{i,j}$	Element of matrix $\boldsymbol{A}$ at the i^{th} row, and the j^{th} column
$\boldsymbol{A}^{-1}$	Inverse matrix to matrix $\boldsymbol{A}$
$\boldsymbol{A}^T$	Transposed matrix to matrix $\boldsymbol{A}$
$\lVert\boldsymbol{A}\rVert_F$	Frobenius norm of matrix $\boldsymbol{A}$
$\mathrm{Tr}(\boldsymbol{A})$	Trace of matrix $\boldsymbol{A}$
$\boldsymbol{I}$	Identity matrix
$\boldsymbol{J}^{a,b}$	Matrix of ones of the size of $a \times b$
$\kappa(\boldsymbol{A})$	Condition number of matrix $\boldsymbol{A}$

Abbreviations

AFC	Alternative Forced Choice
ASI	Adaptive Sampling Iteration
BRDF	Bidirectional Reflectance Distribution Function
BTF	Bidirectional Texture Function
CCD	Charge-Coupled Device
CPU	Central Processing Unit
CSF	Contrast Sensitivity Function
DI	Displacement Interpolation
DOF	Degree of Freedom
DPI	Dots Per Inch
EV	Exposure Value
EXIF	Exchangeable Image File Format
FPS	Frames Per Second
GPU	Graphics Processing Unit
HD	High-Definition
HDR	High Dynamic Range
HVS	Human Visual System
HW	Hardware
IS	Instability of Performance
LAN	Local Area Network
LED	Light-Emitting Diode
MERL	Mitsubishi Electric Research Laboratories
nD	n-Dimensional
PCA	Principal Component Analysis
PCHIP	Piecewise Cubic Hermite Interpolating Polynomial
RBF	Radial Basis Function
STAF	Space-Time Factorization
SVBRDF	Spatially-Varying Bidirectional Reflectance Distribution Function
SVD	Singular Value Decomposition

TPS	Thin Plate Spline
TVBTF	Time-Varying Bidirectional Texture Function
UTIA	Institute of Information Theory and Automation of the CAS, v. v. i.

Computational Error Measures

We use the following measures for comparison of rendered images with ground-truth images:

CIE ΔE Difference Empfindung (sensation),
$\Delta E = \sqrt{(L_1^* - L_2^*)^2 + (a_1^* - a_2^*)^2 + (b_1^* - b_2^*)^2}$,
where (L_1^*, a_1^*, b_1^*), (L_2^*, a_2^*, b_2^*) are two images in $L^*a^*b^*$ color space.
The lower value is better. $\Delta E \in [0, \infty)$.

MRE Mean Relative Error, $MRE = \frac{1}{N} \cdot \sum\limits_{\lambda,\theta_i,\theta_v,\varphi_i,\varphi_v} \frac{\left|f_{r_\lambda} - f'_{r_\lambda}\right|}{f_{r_\lambda}} \cdot 100[\%]$, where $f_r(\theta_i, \theta_v, \varphi_i, \varphi_v)$ is the reference BRDF, $f'_r(\theta_i, \theta_v, \varphi_i, \varphi_v)$ is its approximation, $\lambda \in \{R, G, B\}$ is color channel and N is the count of data points. Individual color channels are treated separately and the results are then summed up. The lower value is better. $MRE \in [0, \infty)$.

MSE Mean Squared Error, $MSE = \frac{\sum\limits_{i=1}^{M} \sum\limits_{\lambda=\{R,G,B\}} (I_1(i,\lambda) - I_2(i,\lambda))^2}{3M}$,
where I_1, I_2 are compared images and M is the number of their pixels.
The lower value is better. $MSE \in [0, \infty)$.

RMSE Root Mean Squared Error, $RMSE = \sqrt{MSE}$. The lower value is better. $RMSE \in [0, \infty)$.

PSNR Peak Signal to Noise Ratio, $PSNR = 10 \cdot \log_{10} \frac{255^2}{MSE}$. The higher value is better. $PSNR \in (-\infty, \infty)$.

Log PSNR Logarithmic PSNR [AMS08]. The higher value is better.
$Log\ PSNR \in (-\infty, \infty)$.

SSIM Structural Similarity Index [WBSS04]. The higher value is better.
$SSIM \in [-1, 1]$.

VDP2 Visual Difference Predictor 2 (HDR VDP 2) [MKRH11]. The higher value is better. $VDP2 \in [0, 100]$.

Moreover, we often depict difference images, whose each pixel (x, y) is computed as $255 - s \times |I_1(x, y) - I_2(x, y)|$, where s is a scale, which is always denoted under each figure that contains difference images. The difference images are computed between two compared images I_1, I_2.

Chapter 1

Introduction

Reproduction of appearance of the real-world materials in virtual environments has been one of the ultimate challenges of computer graphics. Therefore, methods for representation, acquisition, and rendering of material appearance have already received a lot of attention. Despite that, methods for acquisition of material appearance presented so far are often not very elaborated, are focused on a certain small class of materials, provide only oversimplified representation of materials or demand excessively long measurement time. Hence, in this thesis, we focus ourselves on acquisition methods that enable us to measure a wide range of materials in appropriate representations and simultaneously keep measurement time short.

There are several ways how to reduce measurement time of directionally-dependent datasets, i.e., the datasets that enable us to reproduce material appearance that changes when we observe a material from one direction while it is illuminated from another direction. Correct reproduction of the directionally-dependent behavior is very important for realistic visualization of materials; therefore, we focus on it in this thesis.

A standard method for measurement of this behavior utilizes a gonioreflectometer that consists of a sensor (a camera, a spectrometer, etc.) and a light source, which are positioned on a hemisphere above a material. The measurement time can be reduced by simultaneous usage of many sensors (as material appearance is captured in parallel) or even by usage of many light sources (as light sources do not need to be positioned above a material or they can even be used simultaneously to increase intensity of reflected light and reduce integration time of sensors). A disadvantage of this solution is increased cost of the setup.

Another solution that reduces the measurement time captures many viewing directions in a single photograph. To achieve that, we can utilize mirrors or exploit knowledge of shape (a sphere, a cylinder, a large plane, etc.) of a material. However, results might be influenced by any imperfection in the shape of a mirror or a material, which must be spatially-uniform when we exploit knowledge of their shape. Also, a range of captured directions or a dynamic range of captured photographs is often limited.

Alternatively, we can reduce the measurement time by adaptive selection of important directions, which are measured by a gonioreflectometer. Then, we can invest more money into a sensor and a light source, because there is only single piece of each. Also, we usually need only a small specimen of a material, which can be easily prepared. Hence, we can reach uncompromising precision of individual measured samples of a material and final quality is mainly influenced only by the number of samples and a strategy of their acquisition.

While a non-adaptive regular sampling provides stable results for a wide range of materials, adaptive acquisition adapts a measurement process on-the-fly based on already collected information on the measured material. Additionally, information from a database of already measured materials, or even a prior knowledge of typical behavior of materials can be used. Then, only important features of the material are captured. As a result, the adaptive measurement often outperforms quality of the regular sampling for many materials or it can reach similar quality with the lower number of samples reducing the measurement time significantly.

1.1 Theoretical Background

A material representation required for appropriate reproduction of material appearance depends on complexity of the appearance. Spatially homogeneous materials can be usually represented by means of a *bidirectional reflectance distribution function* (BRDF) as introduced by Nicodemus et al. [NRH*77]. The BRDF describes a distribution of energy reflected from a material in the viewing direction when illuminated from a specific direction. It is the ratio of reflected radiance exiting along the viewing direction to incoming irradiance from the illumination direction.

As the BRDF cannot capture a spatial structure or a texture of a material, it can be extended to a *spatially-varying BRDF* (SVBRDF). This representation allows an already quite efficient approximation of material appearance, mostly based on a wide range of analytical BRDF models. However, constraints of the BRDF (mainly a reciprocity of illumination and viewing directions) limit applicability of the SVBRDF to nearly flat opaque materials.

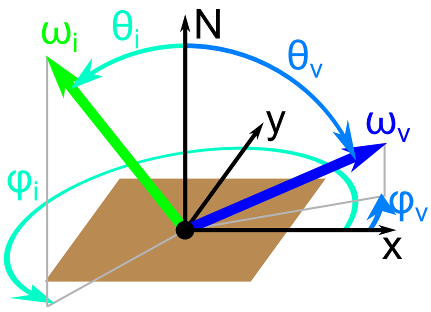

Figure 1.1: Parameterization of the BRDF.

On the contrary, these restrictions of the BRDF and the SVBRDF do not apply to a *bidirectional texture function* (BTF) [DvGNK99]. Therefore, the BTF can capture non-local effects in rough material structures, such as occlusions, masking, shadowing, subsurface scattering, or inter-reflections as shown in [FH09].

As the BRDF, SVBRDF and BTF datasets enable us to represent photo-realistic appearance of materials, they have high application potential mainly in the areas that require physically correct visualizations ranging from computer-aided interior design, visual safety simulations and medical visualizations in dermatology, to digitization of objects of cultural heritage as we show in [A.7].

1.1.1 Parameterizations and Properties of the BRDF

The BRDF is a four-dimensional *vector-valued* function $f_r(\theta_i, \theta_v, \varphi_i, \varphi_v)$ of the illumination direction $\omega_i = \{\theta_i, \varphi_i\}$ and the viewing direction $\omega_v = \{\theta_v, \varphi_v\}$, where $\theta \in [0, 90°]$ is the elevation angle and $\varphi \in [0, 360°)$ is the azimuthal angle of a spherical coordinate system as shown in Figure 1.1. The vector-valued function returns all color values (e.g., $\lambda = \{R, G, B\}$) at once.

As proper measurement of such a four-dimensional function is very complicated and time demanding, several simplifications of the function were introduced. The three-dimensional simplification of the BRDF is called the isotropic BRDF $f_r(\theta_i, \theta_v, \Delta\varphi = \varphi_v - \varphi_i)$. A material is considered isotropic when its reflectance is invariant to its rotation around its normal. In contrast, materials whose reflectance is not constant are considered anisotropic. Anisotropic materials are, due to their atypical rich appearance, often used to achieve eye-catching look of many man-made products as shown in Figure 1.2. Fabrics for fancy apparel are probably the most typical example. Anisotropic finishing of plastics or metals is often used to create resonant appearance of personal items or design electronics. Genuine wooden materials are also very popular due to their specific anisotropic behavior.

The standard parameterization of the BRDF by the spherical coordinate system is not

Figure 1.2: Real-world examples of anisotropic appearance of fabric, hair, brushed metal and polished wood caused by orientation of cylindric threads, hair fibers, parallel grooves in metal (all causing highlights perpendicular to their direction) and wooden fibers (exhibiting both specular (left) and anisotropic (right) highlights), respectively.

the only possibility how to parameterize the BRDF. In many situations, it can be better to use the half-difference parameterization introduced by Rusinkiewicz [Rus98]. The BRDF is then represented as $f_r(\theta_h, \theta_d, \varphi_h, \varphi_d)$, where $H = \{\theta_h, \varphi_h\}$ is the half vector between the illumination direction and the viewing direction, $H = \frac{\omega_i+\omega_v}{\|\omega_i+\omega_v\|}$ and $D = \{\theta_d, \varphi_d\}$ is the difference vector. The difference vector $D = rot_{\vec{y},-\theta_h} rot_{N,-\varphi_h} \omega_i$ is the illumination direction parameterized by spherical angles with respect to the half vector H. When a material is isotropic, its appearance is invariant to changes of φ_h. Another parameterization of the BRDF is, e.g., the onion-like parameterization introduce by Havran et al. [HFM10], where the BRDF is represented along meridian slices of the hemisphere of incoming directions.

Romeiro et al. [RVZ08] proposed an even simpler bivariate representation of the BRDF $f_r(\theta_h, \theta_d)$. They use half-difference parameterization and their simplification is based on the assumption that a material is isotropic and bilaterally symmetric; therefore, its reflectance is almost invariant to changes of φ_d. Different parameterizations suggested for several analytical BRDF models are studied in [SAS05].

The physically-based BRDF has three important properties: Helmholtz reciprocity ($f_r(\omega_i, \omega_v) = f_r(\omega_v, \omega_i)$), positivity ($f_r \geq 0$) and energy conservation ($\forall \omega_v \int_\Omega f_r \cos\theta_i d\omega_i \leq 1$). Spatially-varying extension of the BRDF only brings two extra dimensions (x, y) into parameterization of all its variants but shares all its properties. Therefore, it is limited to nearly flat opaque surfaces. On the other hand, the BTF is a six-dimensional *vector-valued* function $f_r(x, y, \theta_i, \theta_v, \varphi_i, \varphi_v)$, where directionally-dependent behavior of individual pixels does not comply with properties of the BRDF. Therefore, one directionally-dependent pixel $f_r(\theta_i, \theta_v, \varphi_i, \varphi_v)$ of the BTF cannot be called the BRDF even if its behavior is very similar to behavior of the BRDF. Hence, it is called the Apparent BRDF (ABRDF). A recent survey of other parameterizations of material appearance can be found in [HF13].

1.1.2 Visualization of the BRDF

As values of a 4D function cannot be directly visualized, the easiest way to visualize at least a part of a tabulated BRDF is to fix two dimensions and show the resulting 2D subspace as a color image. Iterating through fixed dimensions and compositing resulting images next to each other brings together a comprehensive preview of the entire BRDF.

A first example is shown in Figure 1.3-left where we iterate through the elevation angles θ_i, θ_v and show 2D subspaces of the azimuthal angles φ_i, φ_v. Note that these subspaces are periodical, i.e., $f_r(\theta_i, \theta_v, \varphi_i, \varphi_v) = f_r(\theta_i, \theta_v, \varphi_i + 2\pi k_i, \varphi_v + 2\pi k_v)$, where $k_i, k_v \in \mathbb{Z}$.

Another example is shown in Figure 1.3-right where we show all possible 2D subspaces of the BRDF. In the top-left area are the azimuthal subspaces as in Figure 1.3-left. In the top-right area are the reflectance fields, i.e., the subspaces of the illumination angles for the constant viewing angles. Note that the light fields, i.e., the subspaces of the viewing angles for the constant illumination angles are the same due to Helmholtz reciprocity. In the bottom-left area are the subspaces of the elevation angle of the illumination direction θ_i and the azimuthal angle of the viewing direction φ_v for other angles constant. Again, due to Helmholtz reciprocity, the subspaces of θ_v and φ_i for other angles constant are the same. Finally, the bottom-right area contains the subspaces of the elevation angles for the

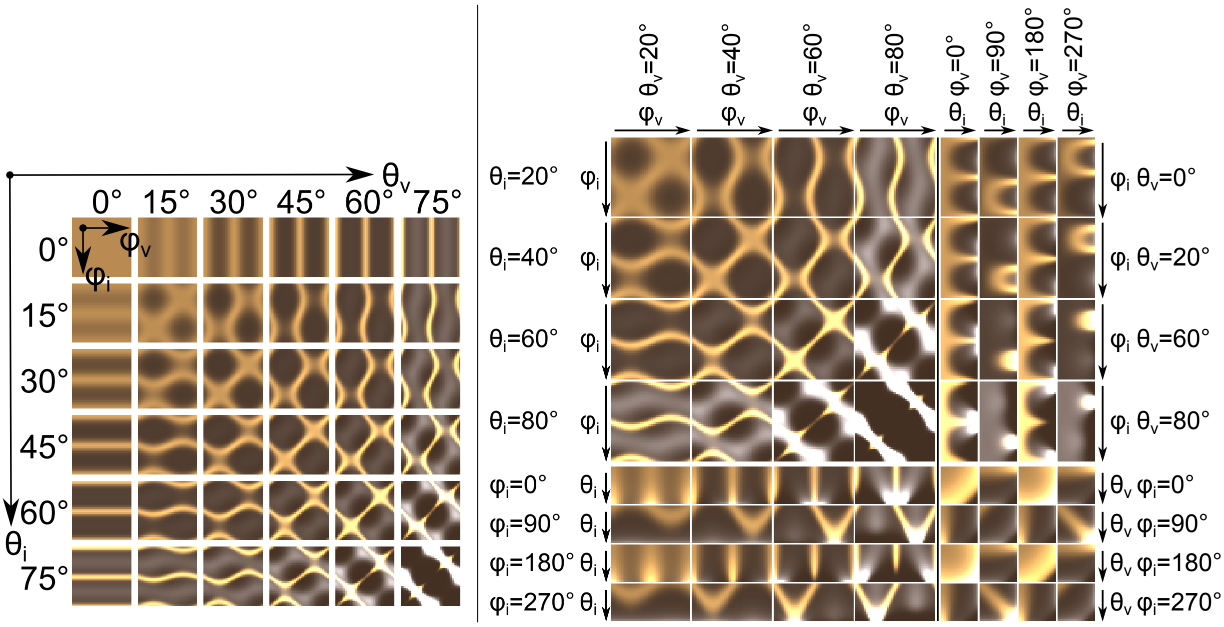

Figure 1.3: Left: The 4D BRDF projected to a 2D color image as a composition of subspaces at given elevations (θ_i, θ_v). Each subspace is two-dimensional, periodical and comprises the entire range of azimuthal angles (φ_i, φ_v). Right: Preview of all possible 2D subspaces of the BRDF.

constant azimuthal angles. Notice that the top-left and bottom-right areas are symmetrical about the main diagonal due to Helmholtz reciprocity.

In a two-dimensional subspace of constant elevation angles, we can observe azimuthally-dependent behavior of a material. Figure 1.4 depicts the BRDF of a woven fabric material that consists of two threads. In the subspaces, we can see the specular highlights that are usually the most prominent near directions corresponding to the ideal mirror reflection, i.e., $\theta_i = \theta_v$, $\varphi_i = \varphi_v + \pi$. On the other hand, the anisotropic highlights are caused by reflection of light from directional elements in a material. Therefore, a fabric material depicted in Figure 1.4 has two types of the anisotropic highlights. Each of them is caused by one of the two threads and its color corresponds to a color of a thread. Notice especially position and shape of the highlights. Each material has different behavior (see Figures 3.7 and 6.7). While isotropic materials do not have the anisotropic highlights, anisotropic materials can have one, two, or even more anisotropic highlights.

1.2 Motivation and Problem Statement

View- and illumination-dependent reflectance properties of materials are often measured uniformly over a hemisphere. As an example can serve a sampling that consists of $81 \times 81 = 6,561$ directions [SSK03] or a sampling that consists of $151 \times 151 = 22,801$ directions

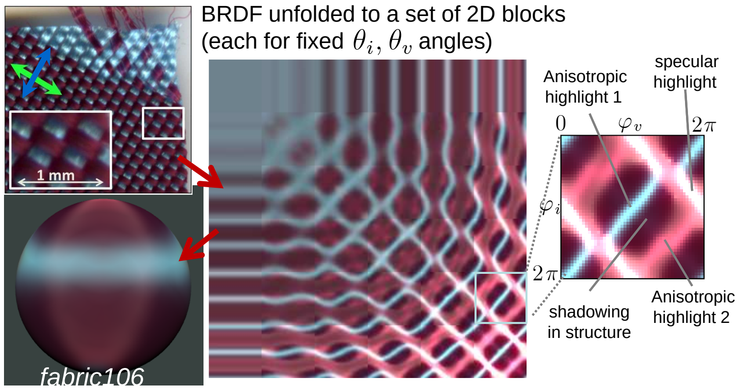

Figure 1.4: Behavior of the anisotropic BRDF of a fabric material. Major visual features are shown in the BRDF subspace of constant elevation angles (right).

[Mül09]. Unfortunately, even such relatively dense samplings are insufficient to accurately capture reflectance behavior of many materials. Figure 1.5 compares renderings of five materials when the ground-truth data and the uniform sampling by 151×151 directions are used. To obtain the ground-truth data, we rendered a sphere in the resolution of 128×128 pixels using a simple ray tracer implemented in Matlab and we generated a list of 10,975 directions, one for each pixel of the sphere. Then all these directions were measured using the UTIA gonioreflectometer (see Sec. 4.1) and they were used for rendering.

As expected, even with the high number of uniformly-distributed samples, the differences are significant, especially for specular materials (see the second column of Fig. 1.5). Therefore, even a more dense uniform sampling is crucial to achieve better precision of the reconstruction or a material-dependent adaptive sampling can be employed.

The main aspect that supports a need for the adaptive measurement is overall time of the measurement with respect to demanded precision. For precise measurement of highlights of even a moderately specular material, we have to use the sampling with the step of 1° in all four angular dimensions $(\theta_i, \theta_v, \varphi_i, \varphi_v)$, which demands $90 \times 90 \times 360 \times 360 \approx 10^9$ samples. Acquisition time of even an extremely fast gonioreflectometer would exceed reasonable limits. Therefore, we need to accelerate the measurement. In this thesis, we focus ourselves especially to adaptive measurement that accelerates the process by sampling in only important directions.

A precise and yet sparse acquisition of material appearance is a tricky task. It is difficult to prepare a list of measured directions for a new material in advance because while positions of the specular highlights can be expected near the mirror reflections, the locations of the anisotropic highlights are unknown. They depend on the local macro-structure as well as on the micro-structure of the measured material.

In our work, we try to develop techniques that enable us to measure material appearance by a small number of samples and yet they should allow a precise reconstruction of material appearance. Some of the techniques are adapted for a fast measurement of the dataset by

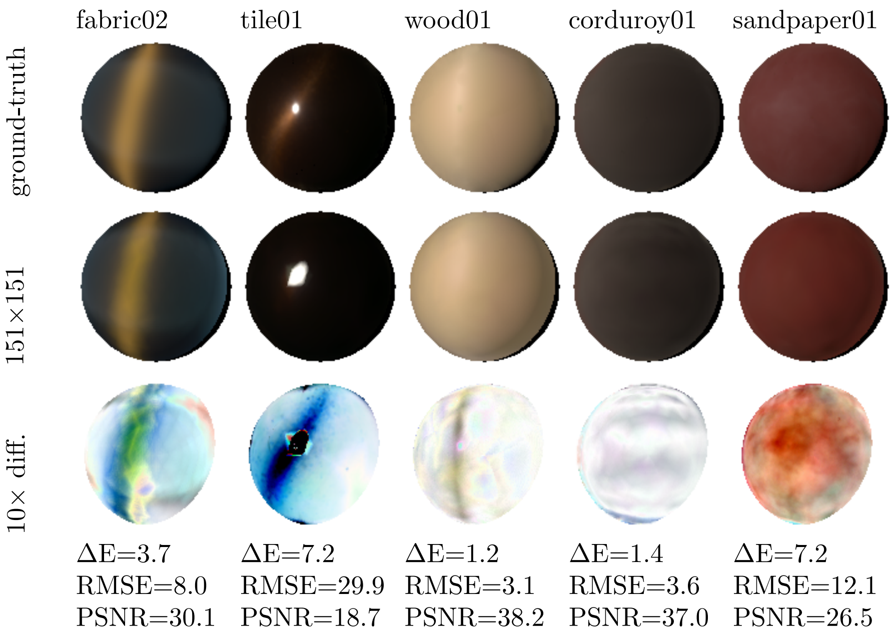

Figure 1.5: A comparison of BRDFs on spheres: 10,975 samples, one for each pixel (the first row) vs. uniformly-distributed measured and interpolated $151 \times 151 = 22,801$ samples (the second row). Below are the difference images scaled 10× (the third row) and the difference values in CIE ΔE / RMSE / PSNR [dB].

a simple inexpensive hardware.

Another consequence of the measurement process that is insensitive to the measured data is the resulting huge amount of acquired data. It is the most apparent in a case of the SVBRDF or the BTF, which comprise besides four angular dimensions, two additional spatial dimensions. In case of the sampling step of 1° in angular dimensions and the moderate spatial resolution of 1024×1024 pixels, the size of the data would be 6.6×10^{15} bytes, assuming three color channels and two-byte half-float data type.

Such big data cannot be directly used for fast rendering and have to be represented by an analytical model or substantially reduced, compressed and stored in a table. Although the analytical representation is very compact and more suitable for fast rendering, it often cannot represent reflectance of a real material in sufficient detail.

This effect can be seen in the work of Nauyoks et al. [NFM14], who fit analytical BRDF models to gradually increasing quantity of measured data. Even in this simple case of the isotropic BRDF, one can see that the more data a model has available the less precisely it can represent them due to a limited analytical formula. The problem would be even more obvious in the case of the anisotropic BRDF or even the unrestricted BTF. Therefore, the tabulated function should be used whenever high accuracy matters.

When we use the tabulated data, they should be very compact and a fast interpolation technique has to be employed to achieve a real-time rendering and a possibility to use many materials in one scene. One example of such an interpolation technique is the barycentric interpolation [Cox69]. Unfortunately, it cannot work with irregularly measured data; therefore, it cannot represent directionally localized effects of the BRDF. Hence, our main motivation is not only development of efficient techniques of adaptive measurement, but also development of a suitable interpolation technique that enables fast data reconstruction during a visualization.

1.3 Contributions of the Dissertation Thesis

The main contributions of this thesis can be divided into four categories as follows.

Analysis of Material Appearance – enables us to study anisotropy or structure of a material. We bring the following contributions:

1. An innovative, fast, and inexpensive image-based approach for detection of an extent of anisotropy, main axes of anisotropy and a width of the anisotropic highlights. The approach uses only an ellipsoidal reflector and a compact camera (see Sec. 3.1).

2. Methods for identification of the *critical viewing distance*, i.e., the distance where the structure of a material in a virtual environment becomes indistinguishable for a human observer, which are based on a psychophysical study (see Sec. 3.3).

Novel Methods for Measurement of Material Appearance – enable us to acquire data efficiently or to process them more precisely. Specifically, we offer the following:

1. A method for measurement of 2D BRDF subspaces of constant elevations that uses an ellipsoidal reflector (see Sec. 3.1.6).

2. A novel registration method of multi-view images (see Sec. 4.2).

3. A method for acquisition of the angularly-dense anisotropic BRDF that exploits variations of illumination and viewing directions across the surface of a measured material for fast collection of many samples (see Sec. 4.3).

4. An inexpensive setup for intuitive and fast approximative measurement of the BTF and the (A)BRDF, which consists of a consumer camera and LED light sources (see Chap. 6).

5. A reconstruction method of the entire anisotropic (A)BRDF or BTF space from less than 200 samples measured by the setup (see Chap. 6).

Adaptive Measurement of Material Appearance – enables us to acquire only important data, and consequently, save measurement time. Particularly, our contributions are as follows:

1. A method for estimation of a sampling pattern tailored to a measured material that uses an ellipsoidal reflector (see Sec. 3.1.5).

2. A concept for efficient measurement of material appearance based on 1D signals, which we call the BRDF slices, and which form a sparse 4D structure in the BRDF space (see Sec. 5.1).

3. A heuristic algorithm for adaptive measurement of values of the BRDF slices (see Sec. 5.2).

4. A reconstruction method for fast interpolation of any BRDF value from the BRDF slices (see Sec. 5.3).

5. A possible improvement of the reconstruction method which reduces some artifacts introduced by the former method, but can handle reconstruction of only 2D BRDF subspaces yet (see Sec. 5.7).

6. The optimal placement of the BRDF slices in the 4D BRDF space (see Sec. 5.4).

Data-Driven Adaptive Measurement – enables even more efficient acquisition of the BRDF, provided we have a sufficiently large database of BRDFs in sufficient resolution. In this area, we contribute by the following:

1. An algorithm that looks for near-optimal sampling patterns of materials in a database with the defined number of samples (see Sec. 7.1).

2. Two adaptive template-based methods for efficient measurement of new materials that use precomputed sampling patterns (see Sec. 7.1).

3. A method for minimal sampling of the anisotropic BRDF by only the very small number of samples, which are captured for several rotations of a measured material around its normal (see Sec. 7.2).

4. An extension of the minimal-sampling method to adaptive measurement of anisotropy of the material (see Sec. 7.2.6).

5. A study that analyzes ability of industrial multi-angle reflectometers to conveniently measure the anisotropic BRDF by the minimal-sampling method (see Sec. 7.2.7).

1.4 Structure of the Dissertation Thesis

The thesis is organized into 8 chapters and 6 appendices as follows:

Chapter 1 – *Introduction* – introduces a concept of adaptive measurement, explains theoretical background behind representations of material appearance and describes the motivation behind our efforts together with our goals. There is also a list of the main contributions of this thesis.

Chapter 2 – *Related Work* – provides a survey of setups and methods for acquisition of the BRDF and the BTF, lists papers that deal with the registration of the acquired data and introduces current methods of adaptive measurement and data interpolation. Also, there is a survey of analytical BRDF models related to the thesis, papers related to perception of materials in virtual environments and a list of publicly available datasets.

Chapter 3 – *Analysis of Material Appearance* – focuses on analysis of visual properties of materials. Directionally-dependent behavior, including effects of anisotropy, is analyzed by usage of an ellipsoidal reflector. We show that the reflector can be used for preparation of a sampling pattern tailored to a measured material, or it can even be used for measurement of some 2D BRDF subspaces of constant elevations. Moreover, the chapter summarizes our main contributions in development of a novel anisotropic BRDF model. Finally, we study human's ability to distinguish structure or texture of a material in a virtual environment in dependence on an observation distance and we propose a function for prediction of the distance where the structure or the texture becomes indistinguishable.

Chapter 4 – *Measurement Setup and Data Processing* – presents a measuring device, a gonioreflectometer with four degrees of freedom, which is used for measurement of the majority of the datasets used in this thesis. We also overview a method for the proper registration of captured photographs and describe a method for collection of a huge amount of BRDF samples from the usual number of photographs of a measured spatially-uniform material.

Chapter 5 – *Adaptive Measurement Based on One-Dimensional Slices* – introduces a method for adaptive measurement of the BRDF that uses a sparse 4D structure formed in the BRDF space by the 1D BRDF slices. Values along the BRDF slices are measured adaptively by a heuristic algorithm and then any BRDF value can be interpolated by provided equations. We study the optimal placement of the BRDF slices in the sparse 4D structure and finally, we outline a possible improvement of the interpolation method which reduces imperfections of the former method.

Chapter 6 – *Towards Rapid and Affordable Appearance Acquisition* – presents an inexpensive setup for the approximative acquisition of material appearance and its applications. The setup enables the rapid acquisition of a subset of BRDF and BTF datasets and besides its detailed description, we introduce a method for reconstruction of the complete dataset from the subset.

Chapter 7 – *Data-Driven Adaptive Measurement* – introduces approaches for adaptive measurement of material appearance that are based on a database of already measured materials. The first group of approaches works with the precomputed sampling patterns tailored to materials in the database when a new material is measured. The second looks for the minimal number of sampling directions and measures those directions for several rotations of a measured material around its normal. Then, the material is modeled by a linear basis computed on the database. When rotations of the material are adapted according to appearance of the material, it enables us to achieve even better results. Finally, we study the ability of industrial multi-angle reflectometers to conveniently measure material appearance by the proposed method.

Chapter 8 – *Conclusions* – summarizes all contributions and achievements of the thesis and concludes the thesis. Moreover, there are suggestions for future work.

Appendix A – *Published BRDF and BTF Datasets* – provides a survey of all measured datasets that were acquired during work on the thesis and that are publicly available. Moreover, we describe methods that were used for preprocessing of the datasets.

Appendix B – *A Study on Visual Perception of Material Structure* – describes materials, a virtual scene, a post-processing method of the obtained psychophysical data and a method for analysis of averaged responses used in Section 3.3.

Appendix C – *Registration of Multi-View Images of Planar Surfaces* – complements Section 4.2 and describes a standard approach for image registration, camera calibration and a method for iterative fitting of a position of the surface plane.

Appendix D – *BRDF Slices* – complements Chapter 5. First, equations for reconstruction of a BRDF value from the BRDF slices in a 2D subspace are rewritten to make transition into the 4D space possible and then equations for interpolation in the 4D space are provided. Next, we describe a dataset and regular sampling schemes used for our experiments. Finally, implementation of the reconstruction method on a graphics hardware is outlined.

Appendix E – *Applications of the Portable Setup* – demonstrates the speed and portability of the setup presented in Chapter 6 on the measurement of human skin and a dynamic process of desiccation of glue.

Appendix F – *Minimal Sampling for Effective Acquisition of the Anisotropic BRDFs* – provides derivation of the gradient of the objective function used in Section 7.2 and describes a method for a fast update of the condition number of a matrix used in Algorithm 7.4.

Chapter 2

Related Work

As the main objective of the thesis is a novel approach to adaptive BRDF and BTF acquisition using a limited sample set and their consequent reconstruction, the related areas of research are methods and setups for BRDF and BTF datasets acquisition (Sec. 2.1) and their registration (Sec. 2.2). Related adaptive measurement approaches and data interpolation techniques are summarized in Section 2.3, while survey of related analytical BRDF models is in Section 2.4. Section 2.5 summarizes previous work on visual perception of materials in virtual environments. Also, there is a brief survey of publicly available datasets suitable for research purposes (Sec. 2.6).

2.1 Acquisition Methods and Setups

Based on the way the four degrees of mechanical freedom are realized, BRDF and BTF acquisition setups can be divided into three categories. The first is based on gonioreflectometers where all combinations of illumination and viewing directions are achieved by the sequential mutual positioning of a light source, a sensor and a sample. The second category leverages mechanical complexity using some kind of mirror to capture multiple views in a single image or using multiple light sources and sensors (e.g., dome-based setups). The third category utilizes a prior knowledge of non-planar geometry of a sample to capture many illumination and viewing directions in a single image. A recent very detailed comparison of goniometer- and dome-based setups is available in [SSW*14].

As the BTF is a more general representation of material appearance than the SVBRDF and the BRDF, setups that allow acquisition of the BTF can be used for capturing of the SVBRDF and the BRDF as well. On the other hand, setups for the SVBRDF and BRDF acquisition can be less complex, e.g., as the SVBRDF is restricted by its definition to opaque and almost flat surfaces, its acquisition techniques make use of BRDF reciprocity. Another limitation of a majority of these approaches is that they capture only the isotropic SVBRDF and the isotropic BRDF, which is hardly the case for the most real-world materials.

Goniometer-Based Setups realize the required four degrees of freedom (DOFs) by the sequential mutual positioning of a light source, a sensor, and a sample during BRDF or BTF data acquisition, e.g., [MCS90, WSB*98, DvGNK99, KMBK03, SSK03, LFD*08, HLZ10, ELU11, HAH*12]. A very precise setup [HAH*12] based on the combination of a moving light source on a circular rail and a robotic arm that orients a sample has been recently used even for industrial measurements of paint pigments.

Mirror-Based Setups have been developed to reduce mechanical complexity as the angularly dense measurement using gonioreflectometers is very time demanding. These setups allow to capture multiple views in a single image. They can use kaleidoscopically arranged flat mirrors [HP03], hemispherical mirrors [War92], off-axis parabolic mirrors [Dan01, DW04, GAHO07], ellipsoidal mirrors [MSY07b, MSY09], or a combination of concave parabolic and custom-built mirrors [GHAO10]. Most of these setups place a measured material into the focal point of a mirror, share the optical and illumination axis using a coaxial pair of a camera and a projector, and thus allow to record many illumination or viewing directions in a single image. The main advantage of such an arrangement is elimination of any mechanical component in the measurement setup. Another advantage is the possibility to capture retro-reflections, which is difficult for goniometric acquisition techniques due to physical occlusions of a light source or a sensor. Principal disadvantages are a limited range of recorded elevation angles, a variable reflectance attenuation across elevations, or a low dynamic range of measurements when a projector is used as a source of illumination. The most setups place a material into a focal point of a mirror in such a way that the material is aligned with the axis of the mirror [Dan01, MSY07b, MTK*11]. This is not ideal for fast detection of anisotropy as it requires extraction and positioning of a sample inside the mirror. More applicable are setups that align the axis of a mirror with the normal of the surface [War92, MSY07a].

Approximate BRDF measurements can be also obtained using the condenser reflectometry that records many views at once [LDW*10]. Mukaigawa et al. [MSY09] measure the anisotropic BRDFs, and besides an ellipsoidal mirror, use a projector. Although the setup allows measurement in a couple of minutes, it is limited by a low dynamic range of the projector and low overall accuracy of the setup, as it relies on an accurate placement of a sample and uniform properties of the mirror. Another approach called the hemispherical confocal imaging uses an assortment of flat mirrors distributed over an ellipsoid [MTK*11]. It can be viewed as a set of virtual cameras and projectors positioned uniformly over the hemisphere.

Although a majority of the cited setups offer only a measurement of the BRDF, a solution for the rapid capture of the BTF has been proposed by Han and Perlin [HP03]. Their system comprises of a triangular tapered tube covered with mirrors, which reflects a kaleidoscopic mosaic of a material, which is positioned beneath the tube. A fixed camera views the kaleidoscopic image where individual triangular areas correspond to the surface as observed from different viewpoints. The sample is illuminated by a digital projector by covering individual triangles (i.e., generating illumination directions) and the projector uses a beam-splitter to share an optical path with the camera. The main advantage of this system is exclusion of any moving parts that results into an extremely fast measurement limited only by acquisition speed of the camera. However, resulting drawbacks are a limited spatial resolution (relative to resolution of the camera and the projector), a restricted dynamic range (the projector is used as a source of illumination), a greater cost of the setup due to the mirror and the projector, and a restricted size of a measured sample, which stems from an adequate size of the setup. When one wants to increase a size of the setup, there are greater demands on a size and quality of the mirror, and consequently, a greater cost of the setup. Also, requirements for geometrical accuracy of the system may limit its portability.

Setups Based on Multiple Light Sources and Sensors reduce measurement time by reduction of their mechanical positioning and by parallel acquisition [NZG04, MMS*05, MBK05, BEWW*08]. In the case of a BTF measurement, a financial cost of such a setup is high. On the other hand, due to reduction of measurement time to about two hours, capturing of time decaying materials like food is possible [SWR*11]. A light stage originally designed for capturing of human face is used by Gu et al. [GTR*06] for fast measurement of time-varying processes. The measurement takes only 30 seconds, however, only 16 viewing directions are recorded. While these solutions generally allow faster measurement, lower accuracy or a lower range of illumination or viewing directions often result.

Prior Knowledge-Based Setups exploit knowledge of the geometry or positioning of a sample to capture many illumination and viewing directions in a single image. These setups capture either isotropic BRDF data [MPBM03a, MPBM03b, HHA*10], anisotropic BRDF data [LL95, MWLT00, NDM05] or SVBRDF data [LLSS03, ZREB06]. In [MPBM03a, MPBM03b] a sphere is used for isotropic BRDF measurements, while [MWLT00, NDM05] use a cylinder for anisotropic BRDF measurements. Another image-based approach estimates the BRDF of an object from a single view and illumination [SHGW13] by extraction and compensation of sparse BRDF data, which are later fitted by an analytical isotropic model.

Accuracy of measurements in this category is often compromised due to inhomogeneity of a specimen and its shape imperfections. The approaches are often limited because they represent sparsely-sampled data by a parametric BRDF model. Challenges also result because some flexible materials would exhibit unwanted folds when wrapped, and the approaches are incompatible for usage with rigid materials such as woods or tiles.

The isotropic SVBRDF can also be estimated from photometric stereo using a parametric reflectance model [GCHS05] or the bi-variate BRDF [AZK08]. Wang et al. [WZT*08] capture a single-view 2D slice of the BRDF at each surface point and fit an anisotropic micro-facet BRDF model to the measured normal distribution function. Finally, it is possible to use portable setups that measure the SVBRDF by matching the sparsely-measured isotropic BRDFs (using condenser lens optics) with the sparsely-measured global reflectance fields [DWT*10]. Another approach records the SVBRDF from a single view using a moving linear light source and a set of materials with the known BRDF, which are recorded simultaneously with a measured material [RWS*11]. Acquisition of the sparse SVBRDF proposed in [GCP*10] is based on the measurement of several images of known geometry illuminated by a circularly polarized light. Although this method requires to capture only four images, its usage is limited to flat and isotropic materials and it requires a complex measurement setup.

Likely one of the first attempts of an efficient measurement of the anisotropic BRDF was presented by Lu et al. [LKK98], where the number of mechanical elements is reduced by placement of strips, which are cut from a measured velvet material, on a cylinder of a known geometry in different orientations. In their follow-up work [LKK00], the authors predict locations of the anisotropic reflections on cylinders, cones and spheres made of materials that contain tangential hairs or grooves. The former approach is further revived in [NDM05], where 20 strips of a measured material are attached to a cylinder. A light source moves around the rotating cylinder and a fixed camera captures its images. Its main limitations are long measurement time (16 hours) and significant noise in resulting data.

A statistical acquisition approach [ND06] allows for fast and inexpensive measurement of the BTF; however, it requires a large specimen of a regular material, which have uniform statistical properties and which is cut and positioned in different orientations with respect to a camera to capture several viewing directions simultaneously. The requirement of several strips of a material with the same statistical properties limits practical applicability of the approach to only spatially regular materials. Moreover, the need for extraction of a material from its original environment prevents many scenarios of a portable measurement.

The BTF can be recorded on a 3D object simultaneously with its shape to achieve higher angular density of illumination and viewing directions [MBK05, HLZ10, WSRK11, RSK12]. In such a case, application of the measured data to another object is very difficult. Holroyd et al. [HLZ10] use a system based on a spherical gantry, where each arm is fitted with a camera and a source of high-frequency, spatially-modulated light, which share a common focal point and an optical axis. The proposed measurement method exploits multi-view stereo, phase-based profilometry, and light descattering to avoid problems with registration of 2D-3D data, but it leverages a restrictive assumption about the BRDF, as is often done by related methods. Weinmann et al. [WSRK11] added several projectors into the setup [MMS*05] for a detailed acquisition of a 3D object. The projectors emit structured light used for unique identification of points on a surface of the object.

2.2 Registration of Acquired Data

Sattler et al. [SSK03] measure BTF data and register them using a projective transformation based on registration marks (see Sec. C.1). Such registration is sufficient as long as a resolution of captured images of a measured material is low enough and the plane which represents the surface of the material is close to the reference registration plane. Additionally, the problem of registration is not so pressing in this case because the authors apply compression to data of individual views separately. Thus, the multi-view correspondence does not affect performance of compression to the same degree as other compression methods do [KMBK03, MMK03].

The registration based on the registration marks properly aligns only those parts of a measured sample which lie close to the plane specified by the marks. Neubeck et al. [NZG04] were aware of that. Their work is the most relevant to ours (Sec. 4.2) as they propose to evaluate quality of BTF alignment using a function which computes an average Euclidean distance between reflectance values of those neighboring views that share the same illumination direction. They test several shifts of the surface plane and select the one for which this function is minimal. In contrast, our technique allows us to not only compensate for height misalignment, but also for a mutual rotation of the registration and sample planes, without the need for a repetitive measurement.

Müller et al. [MMS*05] use a setup with no moving parts. Therefore, positions of the image sensors are known in advance and registration can be done in sub-pixel accuracy without a need for registration marks. In another paper, Müller et al. [MSK06] propose an approach that attempts to align individual BTF pixels based on optimization techniques that reduce certain intra-variations in the data. This method rotates individual ABRDFs to achieve better global compression performance; therefore, it requires storage of an additional per-pixel rotation map. Nevertheless, in both cases, the accuracy of the measurement or the fit depends on an initial position of a calibration plane and its difference from the plane that represents the surface of a material.

Ruiters and Klein [RK09] published a technique which represents appearance of a material using a combination of a surface depth-map and a spatially-varying reflectance. The authors define a dense reference mesh and align its polygons to fit the original data the best in order to estimate a depth-map of near-flat surfaces. This technique can deal with materials that have variable height of a surface; however, our method (Sec. 4.2) is easier for implementation and it is computationally less expensive. We do not attempt to interpret surface depth (which might even be impossible for translucent materials), but to find an alignment that maximizes precision of registration of multi-view data.

2.3 Adaptive Measurement and Data Interpolation

Adaptive measurement attempts to measure an unknown function by the given number of samples that should be placed in such a way that reconstruction precision of the function is close to the best possible precision for the given number of samples. The first category of such approaches estimates a *fixed set* of directional samples using a linear combination of BRDFs from a database or their precomputed common basis and iterative search for a well-conditioned subset of the required number of directions. Then for any new material, only this subset of several hundred samples has to be measured and the respective BRDF is obtained by linear combination of BRDFs or basis vectors. Matusik et al. [MPBM03b] represent the MERL database using a wavelet basis or the linear combination of BRDFs. Later on, Ali et al. [ASOS12] use bivariate representation [RVZ08] to improve reconstruction of the BRDF using the linear basis. Similarly, Nöll et al. [NKS14] represent the MERL database using basis functions; however, deviations of their reconstruction from the reference are approximated by a basis of correction functions. Nielsen et al. [NJR15] present an approach to reconstruct the isotropic BRDFs by basis functions, which uses extremely sparse measurements. These approaches can optimize a set of appropriate directional samples based on information in a database; however, they cannot achieve any adaptivity towards measured materials. As a result, they cannot measure materials with features not present in the database precisely enough. Although there are approaches to a direct BRDF measurement in this representation from a single image given the know illumination environment [RVZ08], they are limited to a certain subset of isotropic materials.

Suboptimal placement of samples to represent a known function has been also extensively studied. Notice that this task is different to ours because in this case the function can be analyzed in advance and the samples can even be repositioned to find their (sub)optimal placement. Robinson and Ren [Rob95] apply a sparse one-dimensional yard-stick coding for compression of image data along ridges and valleys, which are obtained using a Laplacian-of-Gaussian image filter. They proposed a suboptimal relaxation algorithm for refinement of positions of samples. Sampling of known 1D functions based on Douglas-Peucker algorithm was introduced in [LRR05]. Computer graphics deals with a similar problem in the case of importance sampling.

The second category of approaches performs a stepwise data-driven measurement of an unknown BRDF. General methods of the adaptive sampling [TS96] have been extensively studied. Their application to the adaptive sampling of illumination is investigated in [FBLS07], yet only in two dimensions, which represents sampling of the reflectance field. An adaptive approach of an image-based BRDF measurement is proposed in [LLSS03] with planning of viewing and illumination directions based on minimization of uncertainty of parameters of an analytical model. Although one can use global methods for interpolation and prediction of errors [PTVF92], they have high computational demands, which become intractable when the number of samples exceed several thousand. Nauyoks et al. [NFM14] fit six isotropic BRDF models to the measured data. They iteratively add new samples by including illumination and viewing directions where the models differ the most. A disadvantage of the method is that if a BRDF model cannot fit measured data accurately, it has to be excluded from subsequent computation in order to avoid distraction of the result. Xu et al. [XSD*13] present anisotropic spherical Gaussians as a closed form representation of anisotropic behavior, which can be practically used for analytical modeling of the sparsely sampled BRDF. Recently, Fichet et al. [FSH16] introduced a method that fits a locally anisotropic analytical model to sparse measurements transformed to Fourier domain.

Methods of the adaptive measurement and data interpolation are closely related as our task is to achieve a good reconstruction of an unknown function by the optimal placement of novel samples based on previously measured values and the optimal placement depends on the chosen interpolation method. Good candidates for a useful interpolation method are global interpolation methods [PTVF92] based on radial basis functions (RBFs) or Kriging [JSW98]. They have often been used in the past for prediction of a black-box function in industry or for an estimation of profitable mining locations. Although, it seems that these methods might solve the problem of the adaptive sampling in measurement of material appearance, they have high computational demands that become intractable when the number of samples exceeds several thousands. Other examples of applicable interpolation methods are the barycentric interpolation [Cox69], thin plate splines [Boo89], or the Gaussian pyramid in the so-called pull-push algorithm [GGSC96]. Recently, Ward et al. [WKB14] reconstructed uniformly sampled BRDF measurements by RBFs interpolated using mass-transport solution. Zickler et al. [ZREB06] achieve a reconstruction of the SVBRDF in the form of the isotropic BRDFs interpolated by means of 3D RBFs. Alternatively, the 4D BRDF can be decomposed into simpler 1D and 2D components which have physical meaning, to allow parametric editing of visual properties [LBAD*06].

2.4 Analytical BRDF Models

To overcome relative sparsity and a noise present in a typical BRDF measurement, material appearance is often represented by analytical BRDF models. While their main advantage is the effective suppression of outliers and the ability to provide data for arbitrary directions, their complicated fitting process often remains in a local minimum and some of the appearance features are often not captured properly. These problems are even more distinct for anisotropic materials, which require a more sophisticated model.

We can roughly divide BRDF models into two categories. While empirical models compromise accuracy and physical plausibility in order to achieve a low number of parameters and thus faster evaluation, the physically derived models offer higher descriptive qualities, albeit at the cost of computational demands. A recent survey of the BRDF models can be found in [MU12]. Here we only overview models related to our work.

The most of empirical models represent the BRDF by means of a specific type of a reflectance function, which expresses a mutual relationship between the illumination direction ω_i, the viewing direction ω_v, and the normal vector $\mathbf{N}$ (see Fig. 1.1). The empirically derived models are usually based on a very simple formula with several adjustable parameters designed to fit a certain class of reflectance functions. The most common is a model of Phong [Pho75] and its more practical modification of Blinn [Bli77]. A generalization of the Phong model based on cosine lobes was introduced by Lafortune et al. [LFTG97]. A simplified, physically-plausible anisotropic reflectance model based on Gaussian distribution of micro-facets was presented by Ward [War92]. The model has the necessary bidirectional characteristics, and all four of its parameters have physical meaning; therefore, it can be fitted independently to measured BRDF data to produce a physically valid reflectance function that fulfills the reciprocity and the energy conservation. A model of Schlick [Sch94] stands halfway between empirical and theoretical models. Ashikhmin et al. [AS00] extended the Phong-based specular lobe, made this model anisotropic, and incorporated a Fresnel behavior while attempting to preserve the physical plausibility and computational simplicity of the initial model.

In contrast, physically motivated models are designed to represent some known physical phenomena; therefore, individual parameters or fitted functions are related to properties of real materials. A pioneer model of Torrance and Sparrow [TS67] represents a surface by the distribution of vertical V grooves, perfectly specular micro-facets. This model was later enhanced and introduced to the computer graphics community by Cook and Torrance [CT81] and further extended for more accurate and stable fitting by a shifted gamma micro-facet distribution in [BSH12]. The micro-facet models are typically based on a combination of Fresnel, facet-distribution, and shadowing/masking functions. It was shown [Hei14] that the distribution and masking functions can be effectively used for modeling of anisotropic behavior. Such an anisotropic extension of the Cook-Torrance model was recently introduced by Kurt et al. [KSKK10]. This model is based on a normalized microfacet distribution function, is physically plausible, and is designed to be more advantageous for data fitting and real-time rendering. Dupuy et al. [DHI*15] proposed an approach to the efficient fitting of distributions of facets based on solving an eigenvector problem, which relies solely on backscattering samples. Although this method provides a very convenient way for estimation of the properties and can fit the anisotropic BRDF in twenty seconds, its final accuracy depends, due to its principle, on the density of BRDF measurements in backscattering directions. Therefore, precision of fitted parameters is limited, especially for the complex anisotropic BRDFs.

The most complete and complex BRDF model was proposed by He et al. [HTSG91]. Although it accounts for many physical phenomenons involved in reflection of light on rough surfaces such as polarization, diffraction, conductivity, and subsurface scattering, it cannot represent the anisotropic BRDFs. Oren and Nayar [ON95] extended the model of Torrance and Sparrow [TS67] to model roughness of a surface using Lambertian facets. An efficient isotropic BRDF model based on a rational function was recently proposed by Pacanowski et al. [PCS*12]. The model uses reduced Rusinkiewicz's parameterization [Rus98], relies on a compact set of parameters, and its extension allows limited modeling of anisotropic effects.

Another physically motivated model of Sadeghi et al. [SBDDJ13] extends the original work of Poulin and Fournier [PF90]. It is based on the physics of light, which scatter in cylindrical threads of fibers, and achieves anisotropy by application of a user-defined weaving pattern of two threads together with parameters of their reflectance model. A drawback of this model is complicated application of parameter optimization techniques; therefore, derivation of appropriate parameters of individual threads is left completely to users and their experience.

Finally, let us remind the recent work of Raymond et al. [RGB*14] that exploits knowledge of behavior of the anisotropic reflections for manipulations with locations of anisotropic highlights by modifications of tangential fields on an arbitrary geometry.

2.5 Perception of Materials in Virtual Environments

As our work is related to the realistic representation of materials in a virtual environment, we provide here survey of previous work related to psychophysical analysis of visual perception of materials in virtual environments and to prediction of visibility of structural elements of materials.

Psychophysical Analysis of Materials in Virtual Environments – Although visual perception of material appearance represented by means of the isotropic BRDF has been thoroughly studied as a function of global illumination and 3D geometry [VLD07, RFWB07, KFB10], the perceptual effects of the BTF have not been studied in such detail to date. The psychophysical analysis of BTFs has been limited either to the verification of visual quality [MMK*06] or for guidance of data compression algorithms [FCGH08, FHC10]. Guthe et al. [GMSK09] apply model of the achromatic contrast sensitivity function [BK80] to achieve more efficient compression of the BTF without significant loss of a perceived fidelity. The recent work of Jarabo et al. [JWD*14] psychophysically analyses an effect of filtering (down- and over-sampling) of the BTFs in both spatial and directional domains.

A Prediction of Visibility of a Structure – Our work also closely relates to applied aspects of human visual system (HVS) and its models. The contrast sensitivity of HVS was extensively studied and modeled in the past [Wan95]. The spatiotemporal visibility of a texture is modeled in [BK80] by a product of spatial and temporal frequency response curves. These contrast sensitivity curves are approximated by the model based on two space-time separable Gaussian filters. Further, luminance and chromatic contrast of stimuli images can be estimated directly from pixel-wise cone responses to stimuli images (cone channels respectively) according to [WM97]. Alternatively, the image salience [PN04] can be used. This method predicts visual fixations by combination of first- and second-order image statistics, namely luminance and texture contrasts based on per-pixel spatial gradients over stimuli luminance represented by a Gaussian-Laplacian pyramid. Other approaches predict visual difference between two images either by a simple statistical model [WBSS04] or by means of more involved modeling of low-level visual perception [MKRH11].

2.6 Publicly Available Datasets

Despite the relatively high number of published approaches for measurement of the BRDF, there are very few datasets available to the research community at large. Since the well-known MERL database of Matusik et al. [MPBM03a], which comprises one hundred isotropic BRDFs, was published in 2003, only much later smaller numbers of samples were presented in other research projects. BRDFs of eight automotive paints are introduced in [GCG*05] and BRDFs of four paints, skin, ceramics, and a felt fabric are published in [MWLT00]. Ngan et al. [NDM05] perform fitting of all MERL datasets in the database by several analytical models. Also, they provide four additional anisotropic BRDF measurements acquired by their own image-based setup. Although the data have an azimuthal step of 2°, they sample the four-dimensional space relatively sparsely, and the data are very noisy due to the lower reliability of samples captured at higher elevations, which cause strong visual artifacts.

Dana et al. [DvGNK99] published the CUReT BTF database, which is the first BTF database and has a relatively sparse angular resolution. Koudelka et al. [KMBK03] introduce the BTF database with six materials in a moderate angular resolution. Further, Sattler et al. [SSK03] provide BTF datasets of the University of Bonn, which consists of six materials in a moderate angular resolution (81×81) and a rather lower spatial resolution (256×256). Later, measurements of four additional materials in a higher spatial resolution (800×800), and especially with HDR, were added. Multispectral measurements of four materials are introduced by Rump et al. [RSK10]. An extensive database of 84 materials in a higher angular resolution (151×151) was recently published by Weinmann et al. [WGK14]. Unfortunately, the datasets are provided compressed only.

Chapter 3

Analysis of Material Appearance

Prior to description of our efforts in area of adaptive measurement, we focus ourselves on analysis of visual properties of typical materials. Section 3.1 analyses anisotropic properties and describes reasons behind anisotropic behavior of fabric materials. We encompass contributions of the paper [A.12] that introduces an innovative, fast, and inexpensive image-based approach to detect an extent of anisotropy, its main axes and a width of the corresponding anisotropic highlights. The method does not rely on any moving parts and uses only an off-the-shelf ellipsoidal reflector with a compact camera. We analyze our findings obtained by a scan of micro-structure of a material, and we present how results correspond to a micro-structure of individual threads in a particular fabric. We show that knowledge on anisotropic behavior of a material can be effectively used in order to design a material-dependent sampling pattern so as the BRDF of a material could be measured much more precisely in the same amount of time using a common gonioreflectometer. Moreover, we summarize contributions of the paper [A.10] that extends usage of an ellipsoidal reflector to an approximative BRDF measurement.

Section 3.2 briefly summarizes the main contributions of the paper [A.3] that models anisotropic behavior of materials. The BRDF model is based on our experience and uses adaptive stencils of the anisotropic highlights to fit the behavior of a material. Besides the directionally-dependent behavior of materials, the spatially-dependent behavior is an important attribute of many materials. In Section 3.3 we summarize our efforts to study a human's ability to distinguish structure or texture of a material in virtual environments presented in [A.2]. Specifically, we focus ourselves on identification of the minimal distance where structure becomes indistinguishable for a human observer and on prediction of the distance for any material. The identified distance can be used, e.g., to accelerate rendering of virtual scenes.

3.1 Analysis of Materials using an Ellipsoidal Mirror

The most straightforward way to detect presence of anisotropy or its type is to mutually fix positions of a light source and a sensor while the measured sample rotates beneath

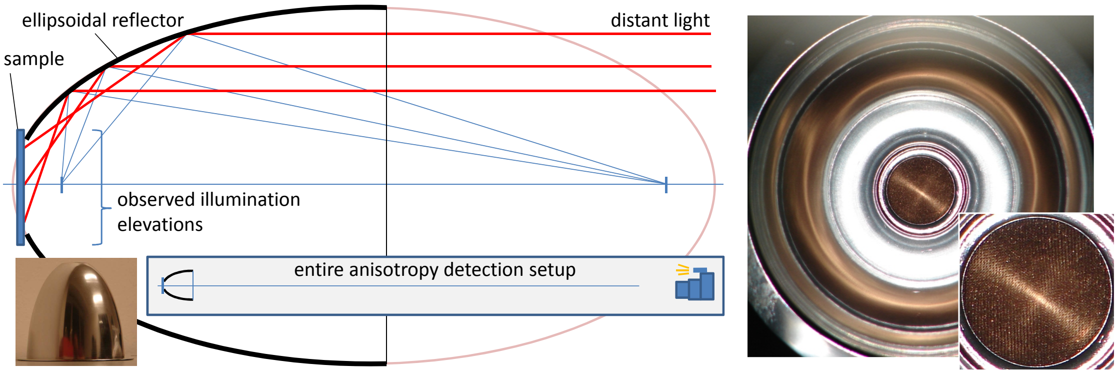

Figure 3.1: A scheme of the setup for detection of anisotropy that uses an ellipsoidal reflector (left) and an image captured by a camera (right).

them and to record outgoing reflectance. However, it is possible to do it more quickly and without any mechanical movement. Our anisotropy detection technique originates out of mirror-based setups; however, it is considerably simplified. It consists of an ellipsoidal reflector, with an opening at the narrowest part (see Fig. 3.1-bottom-left), which is placed on the measured material. This reflector is photographed by a compact camera, which have an optical axis approximately aligned with an axis of the reflector that also represents the normal of the measured surface.

In contrast to other setups, we do not attempt to record the BRDF but only its anisotropic behavior. Therefore, we do not sweep the mirror by a controlled ray of light from the projector but instead, we use the flash from our camera (Sony Cybershot H1). Similarly to [Dan01], we use aggregated illumination, which is however, achieved by a single shot of a flash. We do not use any special gantry and the image is taken from a distance of 1.5 m. This is done in order to capture the reflector (11 cm in diameter) approximately symmetrically, i.e., with the opening in the center of the reflector. Note that our configuration, as shown in the inset of Figure 3.1, is a reasonable approximation of the coaxial setup as the distance between a lens and a flash is 7.5 cm.

Such an omni-directional illumination lights the measured material from many azimuthal angles. As the material is not positioned in the focal point of the mirror, the elevation angles of illumination vary across the opening (see Fig. 3.1-right). This allows to capture directions of the anisotropic reflections. In the image of the material at the opening (26 mm wide) taken by the camera is each location illuminated from two elevations at azimuths 180 degrees apart. Contributions of these two elevations are summed in the captured photo. Mean elevation angle of illumination visible in the opening across the entire plane of the material is 50 degrees. When an isotropic material is analyzed, the captured image contains a circularly shaped peak in the center. This is caused by accumulation of reflected rays at that minute location. However, for an anisotropic material, we additionally observe a couple of triangle shaped highlighted areas, which run symmetrically from the center to the edges of the opening as shown in the first row of Figure 3.3. Their

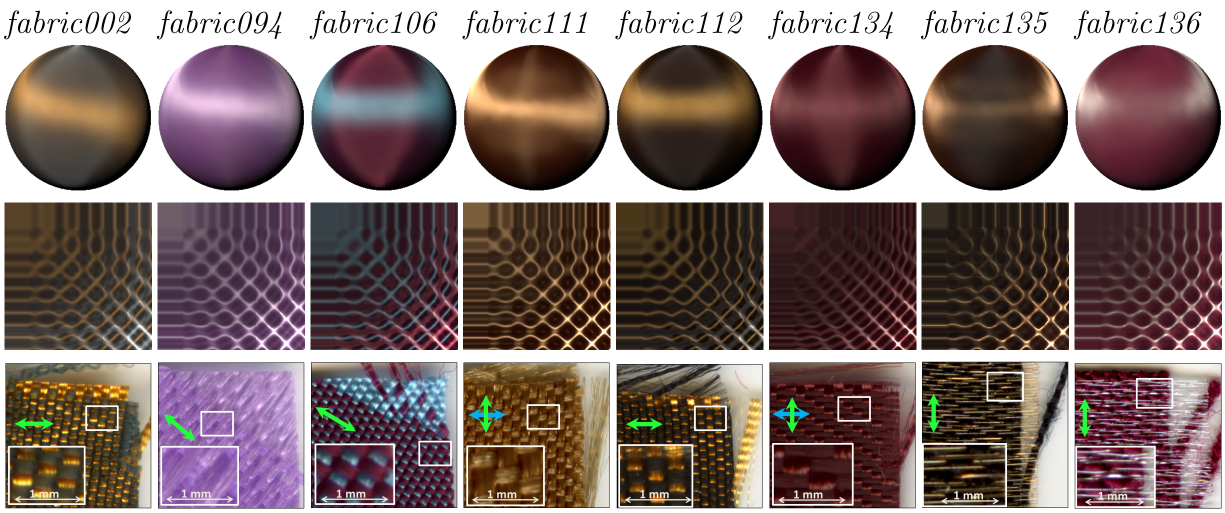

Figure 3.2: The materials rendered on a sphere for point-light illumination (the first row), their BRDFs (the second row). Detailed zoom into the micro-structure of the tested materials with marked primary (green) and secondary (blue) axes of anisotropy (the third row).

azimuthal orientation coincides with the direction(s) of the main axis(axes) of anisotropy, while their width corresponds to the width of the anisotropic highlight for the elevation of illumination $\theta_i = 50°$. Thus, when an unknown material is observed, we can easily determine the presence of anisotropy, its shape and direction.

3.1.1 Tested Materials

To test our approach, we use eight complex anisotropic fabric materials from the UTIA BRDF Database (see Appx. A) and we compare scans of their micro-structure with the detected direction of anisotropy recorded by our method. All tested materials consist of two (same or different) interwoven threads and all materials exhibit strong specific anisotropic behavior as shown on their rendering in the first row of Figure 3.2. The second row shows their BRDFs. The third row of the figure shows the detailed zoom into the micro-structure of the materials and reveals the threads and weaving pattern that influence anisotropic behavior of the materials. Primary and secondary (if applicable) axes of anisotropy are shown by green and blue arrows respectively.

The main axis of anisotropy of materials that consist of two identical threads (*fabric111*, *fabric134*) is determined by the weaving pattern. The main axis is orthogonal to threads that occupy the upper layer of the material. These materials have also distinctive secondary axis of anisotropy (shown as a vertical highlight in the renderings in Fig. 3.2). Other materials that consist of nylon fibers exhibit the highest reflectivity in the direction orthogonal to the direction of their upper threads (*fabric094*) or nylon fibers

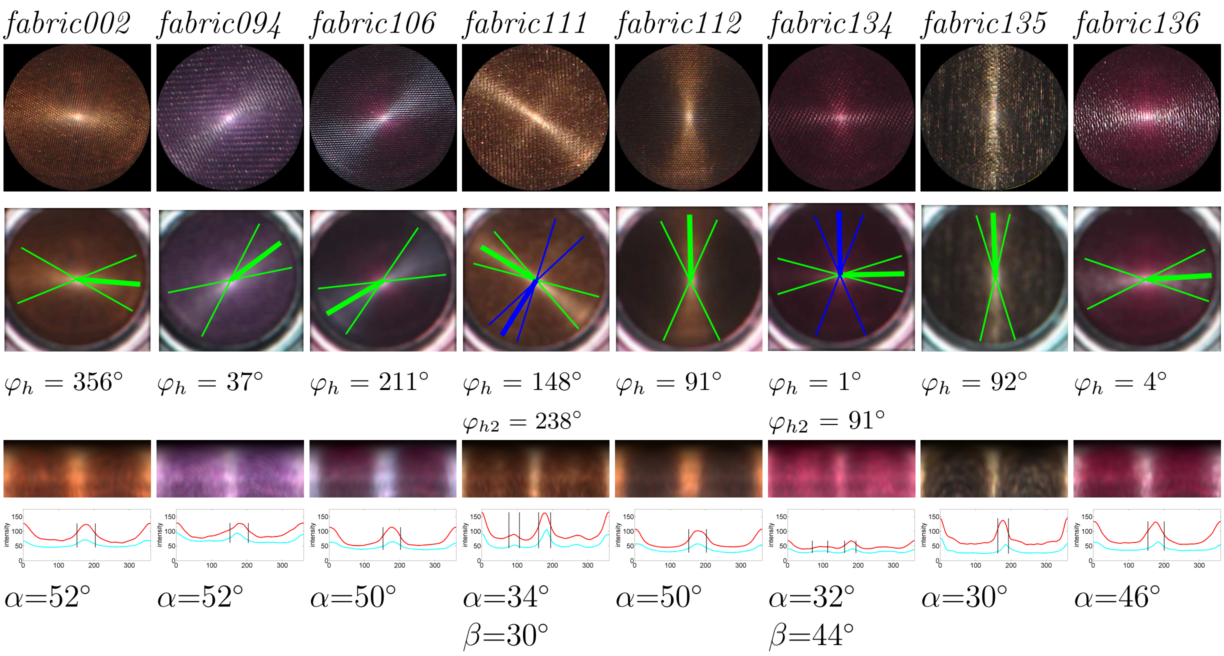

Figure 3.3: Detection of presence or type of anisotropy by the proposed method: the first row - the captured image of the opening of the reflector, the second row - its smoothed version with the detected anisotropic highlights (green - primary axis of anisotropy, blue - secondary axis of anisotropy), the third row - the radius/φ_h image with the clear vertical highlight, the fourth row - 1D plot of profile of the anisotropic highlight (red line) with an estimated width of the anisotropic highlights α, β. The reference BRDF values are shown as cyan lines.

(*fabric135*, *fabric136*). Finally, the materials that consist of similar threads of different colors (*fabric002*, *fabric106*, *fabric112*) exhibit an intricate color shifting, which depends on the actual illumination and viewing directions of the weaving structure. In fact, these materials have two anisotropic highlights; although the second one is very wide and difficult to be detected at elevations below 60°.

3.1.2 Results

The results obtained for the tested materials are shown in the first row of Figure 3.3. The first row shows the opening in the mirror as observed by the camera. It is apparent that this simple approach is capable of detecting the number of the anisotropic highlights, their colors, directions, and widths. Indeed, this image provides very descriptive information on the anisotropic behavior of a material without the necessity of lengthy measurement. Note the striking resemblance with renderings of the respective BRDFs in Figure 3.2.

The second row shows these images after filtering and smoothing by median and box filters. Additionally, these images illustrate the automatically detected axes of anisotropy (main as green, secondary as blue), defined by the azimuthal angle φ_h, as well as the estimated width of the anisotropic highlight. Each such image can be viewed as a polar plot which can be transformed to a Cartesian plot as shown in the third row. These Cartesian plots show reflectance from the edge to the center of the opening based on the azimuthal angle of illumination. In the Cartesian plots, the increase of intensity toward the center of the images is compensated for using an image of a Lambertian-like black paper and all images are aligned according to their main anisotropic highlight. Note that this is not necessary when only the lower part of the images is used. In these images, we can observe the width and the shape of the anisotropic highlights. One can average images in the third row along the vertical axis and obtain 1D profiles of the anisotropic highlights as shown in the fourth row (red line). Given the predefined threshold, we can estimate the width of the highlights. In validation of our method, we can observe very close resemblance of the averaged values (red line) to the known reference BRDF values, which correspond to $\theta_v = 0°$ and $\theta_i = 50°$ (cyan line).

3.1.3 Advantages

The main advantage of the proposed technique is fast measurement and processing, which takes 4.4 seconds using a single core of Intel Core i7. Another advantage is a low acquisition cost of the setup as only an of-the-shelf ellipsoidal mirror (price $7) and a basic compact camera with sufficient zoom and resolution are needed. A final notable advantage is the fact that the setup does not require any calibration procedure apart from attaching the reflector to the measured, nearly planar surface and taking its picture from a distance of 1.5 m with a forced flash in conditions of lower ambient lighting.

3.1.4 Limitations

The proposed anisotropy-detection approach is limited to nearly homogeneous and almost planar materials, i.e., without significant spatially-varying appearance whose texture or effects like shadows would visually mask the anisotropic highlights. Therefore, we experiment with different smooth fabric materials that avoid this limitation while still provide a wide range of available types of anisotropic appearance. Similarly, one can use this method for the analysis of any flat groove- or fibers-based type of anisotropy, which is present in, e.g., wooden, machined or polished surfaces.

As the mirror detects anisotropy at elevations around 50°, it cannot principally detect the anisotropic highlights which are not apparent at such an elevation. This is the case of materials *fabric002*, *fabric106*, *fabric112*, where a wide width of the secondary anisotropic highlight serves to make its distinction from the background reflectance quite difficult.

3.1.5 Application to the BRDF Sampling

The results of the proposed analysis can be applied to improve the sampling of the BRDF. Instead of simplistic uniform sampling, we suggest using the obtained information about a

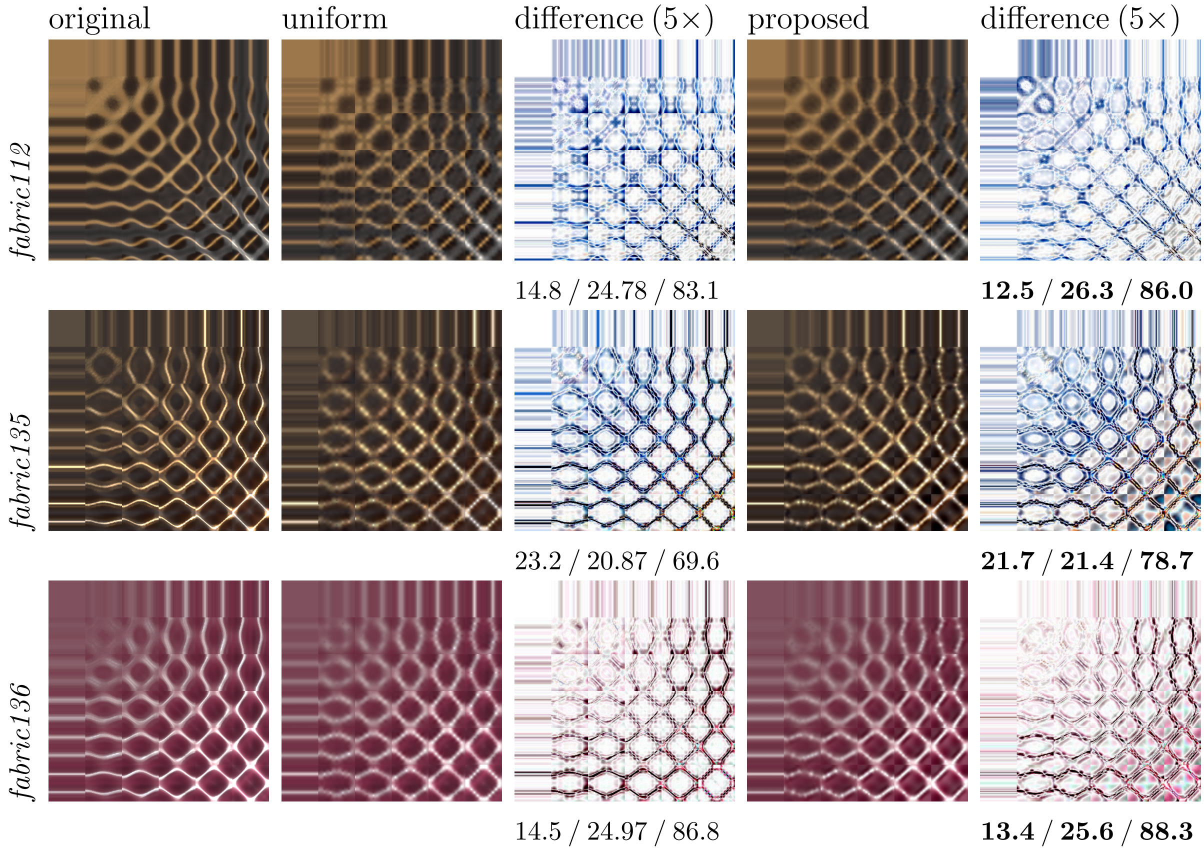

Figure 3.4: A comparison of the reference BRDF and its reconstructions, which use 1025 samples. Below are the difference values in RMSE / PSNR [dB] / VDP2.

position of the anisotropic highlight(s) and their shape to develop a custom-based sampling pattern. For our experiments, we use the UTIA BRDF Database, which provides angularly-dense BRDF measurements of three fabric materials: *fabric112*, *fabric135*, *fabric136* with the angular resolution of 2° (see Sec. A.3).

Figure 3.4 compares performance of reconstruction of the entire BRDF by the uniform sampling and the proposed sampling, both using 1025 samples. Our results, when compared to the uniform sampling, demonstrate a better smoothness of the reflectance function and continuity of the anisotropic highlights. We also compare the reconstructed and reference BRDFs using computational measures (RMSE, PSNR, VDP2) and we observe a notable visual improvement, when the detected information on material anisotropy is used for generation of a proper sampling pattern. Lower performance of the proposed method for *fabric135* is because of the used RBF interpolation, which undershoots the background reflectance near its contours due to a high dynamic range of the anisotropic highlights.

3.1.6 Using a Reflector for Acquisition of BRDF Subspaces

This section summarizes endeavors of the paper [A.10] that extends usage of affordable ellipsoidal reflectors to fast measurement of BRDF subspaces. We have experimented with parabolic reflectors too, but results are rather poor. Captured subspaces may serve for further analysis of anisotropic properties of the measured material, and as we believe, to tailor a sampling pattern for the measured material even better.

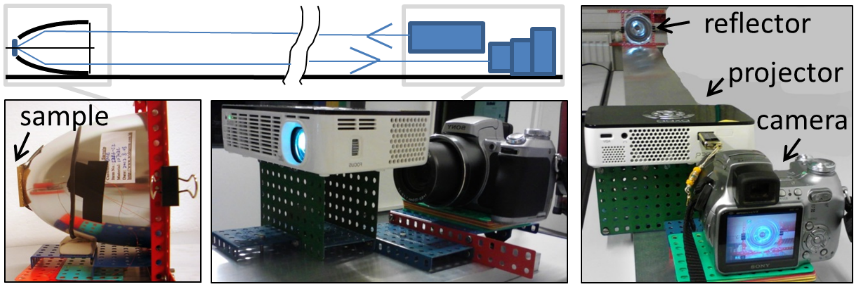

Figure 3.5: The setup for measurement of the anisotropic BRDF that uses an ellipsoidal reflector.

For measurement of the BRDF, we substitute flash illumination by a digital DSP projector AAXA P300 (luminous flux 300 lm), which illuminates the reflector by a controlled ray of light. The entire setup is shown in Figure 3.5. To guarantee alignment of the camera, the projector, and the reflector, we attach them to an optical bench. The lens of the camera is positioned closely to the emitting window of the projector to achieve almost identical optical paths. Distance of the camera and the projector from the reflector is 1.5 meters. The projector circularly sweeps interior of the reflector by a ray of light at various elevations and azimuths. The light is reflected from the material and from the inner body of the reflector back in the direction of illumination. For each azimuthal location of the

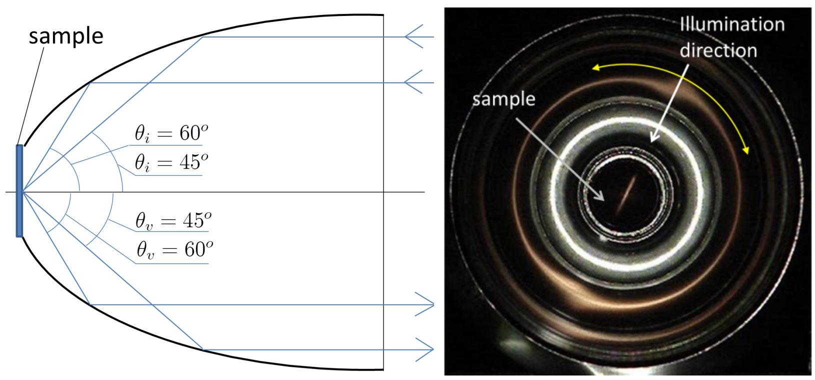

Figure 3.6: A principle of the BRDF measurement that uses an ellipsoidal reflector.

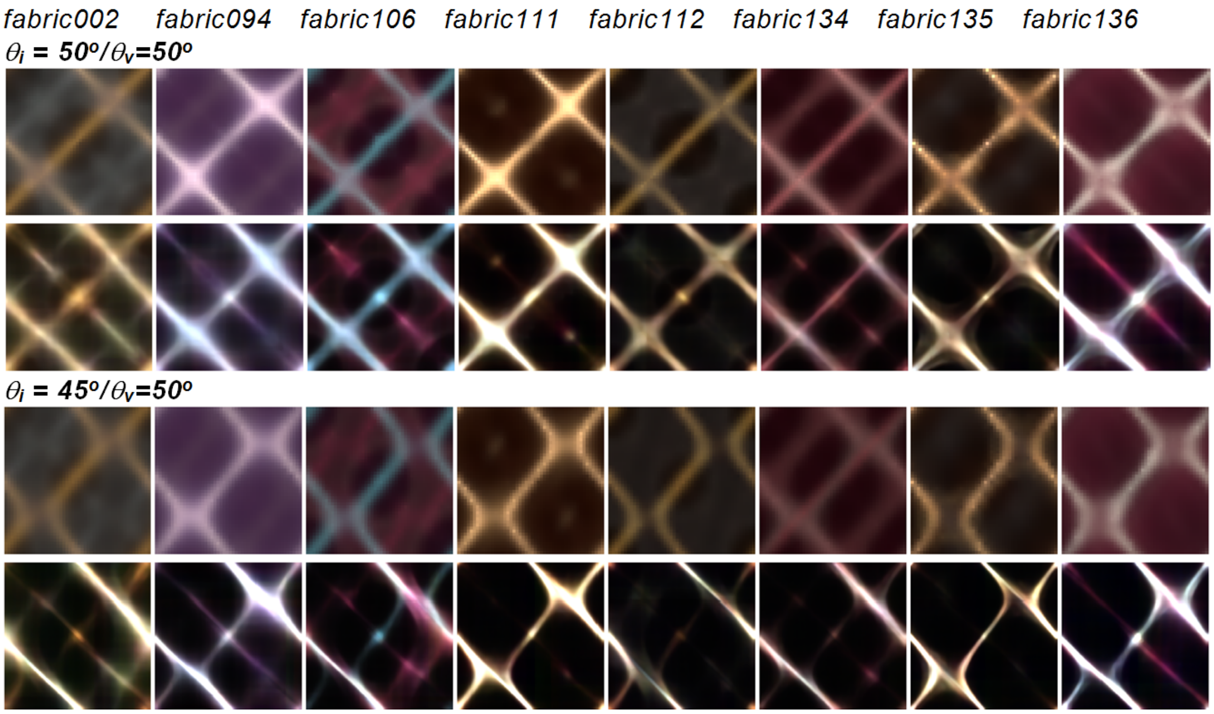

Figure 3.7: Comparison of BRDF subspaces: sparsely captured data using the UTIA gonioreflectometer (see Sec. 4.1, odd rows), ellipsoidal reflector measurements (even rows).

ray, a photo of the reflector is captured exhibiting appearance of the measured material for many azimuthal- and elevation-angle dependent viewing directions. The principle of the method is shown schematically together with illustration of one captured image in Figure 3.6.

Our setup allows acquisition of the 2D BRDF subspaces for elevation angles near 50° due to its non-ideal properties as shown in Figure 3.7. On the other hand, it can capture the main anisotropic features, including the retro-reflective highlights. Another disadvantage of our setup is a low dynamic range due to limited intensity of the projector and its further restriction by reflectivity of the mirror. Further, the lowest elevation angles are principally restricted to elevations around 30°; however, this drawback can be leveraged by a custom design of the reflector.

The main advantage of our technique is short acquisition time, which depends purely on frame-rate of a camera and the required azimuthal resolution. For an azimuthal step of 1° and acquisition frame-rate 30 fps, we finish measurement of a single BRDF subspace in 12 seconds. A high azimuthal resolution can be achieved and is limited only by the required measurement time. Another advantage of the setup is capability to capture back scattering reflections, which is impossible by the most of other acquisition techniques due to physical conflict of a light source and a sensor.

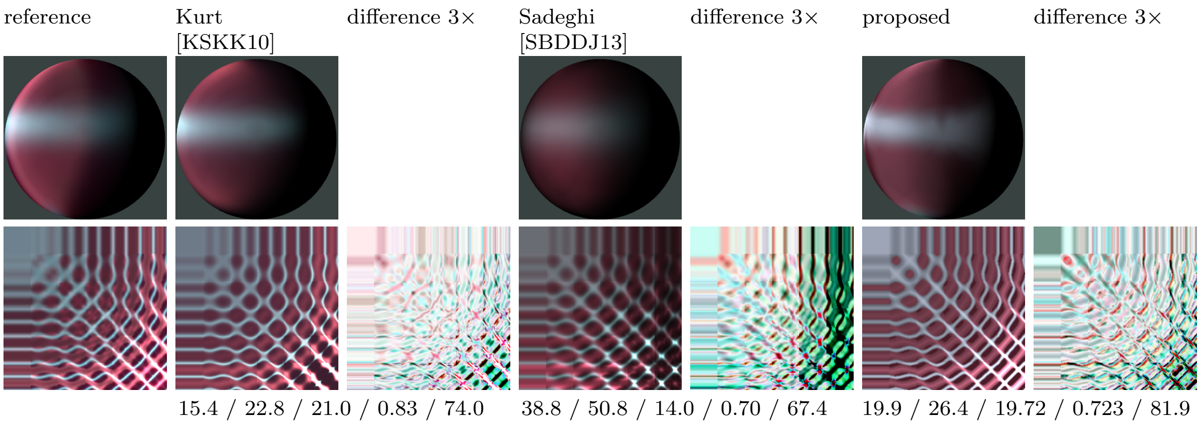

Figure 3.8: A comparison of the reference rendering and the reference BRDF with their reconstructions using three compared models for *fabric106*. Below are the difference images scaled 3× and the difference values in ΔE / RMSE / PSNR [dB] / SSIM / VDP2.

3.2 Modeling of Multi-Axial BRDF Anisotropy

Directionally-dependent appearance of anisotropic materials is widely represented by BRDF models. One of their drawbacks is complicated fitting of their parameters. In paper [A.3], we present a novel empirical anisotropic BRDF model that fits individual modes of anisotropy (i.e., unique highlights) independently using adaptive stencils of the anisotropic highlights. An example of the *fabric106* material, which have two distinct modes of anisotropy, is shown in Figure 1.4. Due to individual fitting of the modes, our approach does not need any time demanding and potentially unstable numerical optimization. We test performance on fourteen BRDFs that exhibit a wide range of anisotropic behavior. Our model achieves robust results while simultaneously minimize the need for user interaction during the fitting. Results of the *fabric106* are in Figure 3.8. For more results see [A.3]. Furthermore, due to the independent fitting of individual modes of anisotropy, the proposed model allows for an intuitive editing of anisotropic appearance and can handle more than two modes of anisotropy.

3.3 Predicting Visual Perception of Material Structure in Virtual Environments

Nowadays, photo-realistic renderings of real-world materials are widely used in computer graphics applications spanning from the motion picture industry, computer games, or paint industry to visual safety simulations or virtual prototyping in the automotive and aircraft industry, or architecture. Such the renderings rely on various representations of material appearance that can offer various levels of realism, but unfortunately also different data acquisition, storage, and related computational costs. In the most static virtual scenes,

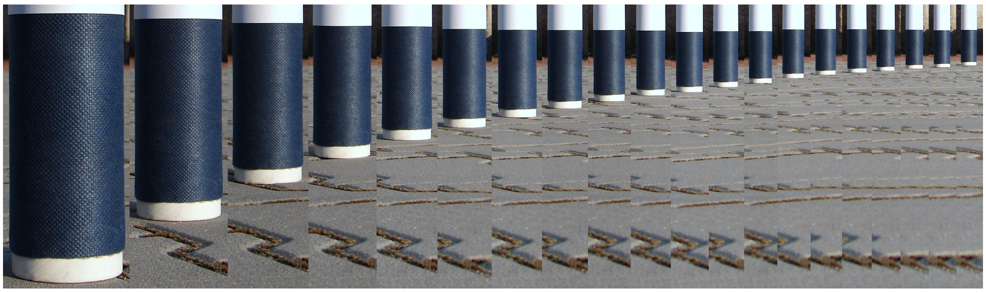

Figure 3.9: A photo-collage of a real non-woven fabric material observed at distances ranging from 0.6 m to 4.8 m (step 0.2 m).

the data storage and computational costs can be reduced by selection of the less demanding appearance representation that can still introduce material properties sufficiently at minimal costs.

The paper [A.2] summarized in this section focuses specifically on the analysis and modeling of the human's ability to distinguish material structure (or texture) in a virtual environment. We consider this information important to estimate a distance at which the material structure becomes visually indistinguishable. Figure 3.9 illustrates this effect on a set of real photographs of a material observed at distances in the range from 0.6 m to 4.8 m.

As we target our work on virtual environments, we analyze human visual perception of material structure directly on digitized representations of material appearance: the BRDF and the BTF. While the BRDF that uses moderate directional resolution occupies, e.g., $3 \times 100 \times 100 = 30,000$ float values, the BTF of the same angular and a very low spatial resolution 128×128 takes about $128 \times 128 \times 3 \times 100 \times 100 = 491,520,000$ values. This tremendous difference in the storage size implies higher memory and computational demands in related rendering applications.

While the BRDF imposes restrictions on reciprocity and opacity, the BTF does not fulfill these restrictions due to local effects in a rough material structure such as occlusions, masking, subsurface scattering, and inter-reflections. These effects constitute a great difference in realism as compared to the BRDF. However, if we move away from the observed surface, the spatial structure and local effects become less apparent. Finally at some viewing distance, we arrive at the appearance that is equivalent to the BRDF, i.e., it exhibits the appearance of a homogeneous, flat surface. This is common in the remote sensing where appearance of, e.g., foliage or urban areas, is represented by the BRDF due to a long viewing distance [QKM*00, SGS*02]. Therefore, we argue that all opaque materials can be represented by the BRDF when they are viewed from an appropriate distance without compromising their basic appearance properties. The main contributions of this section are:

- A psychophysical analysis of visibility of structure for common material categories, which identifies the distance where the structure becomes indistinguishable for the human observer. We call this distance the *critical viewing distance*.

- Identification of the factors that define visibility of material structure to allow its statistical prediction for an arbitrary material.
- Analysis of rendering speed gain obtained when our model of visibility of structure is applied for on-the-fly selection of appropriate representation of material appearance.

3.3.1 The Experiments

To perform three psychophysical studies described in this section, we use 25 materials out of categories often used in interior design (carpet, fabric, leather, and wood). Description of the used materials together with their classification to seven groups can be found in Section B.1. The groups are as follows: G1 - *fabric smooth* (4 materials), G2 - *fabric meso* (8 materials), G3 - *fabric rough* (3 materials), G4 - *leather meso* (3 materials), G5 - *leather rough* (1 material), G6 - *wood meso* (4 materials), G7 - *wood rough* (2 materials). We also include two artificial chessboard patterns into the experiment stimuli, whose widths are 2 mm and 5 mm. While the first one is used to check subjects' monitor resolution, the second one is used to filter subjects who either have impaired vision or do not understand the task.

The task of our experiments is to identify the *critical viewing distance* for all analyzed materials where the structure becomes visually indistinguishable. We alter the task into an analysis of the difference between the BRDF and the BTF as a function of viewing distance. For this purpose, we use test stimuli of receding cylinders as described more in detail in Section B.2. The upper part of the cylinder always displays the BRDF obtained by spatial averaging of its BTF, while the lower part always displays the BTF. Specifically, the subject's task is: *Enter the number of the most distant cylinder where you can still distinguish difference between upper and lower material.* Thus, the experiment type can be denoted as alternative-forced-choice (AFC).

Experiment A – Controlled, Static Stimuli – The scene in the first experiment consists of 20 receding cylinders (example stimulus image in Fig. 3.10-left) simulating distances from 0.6 to 4.4 m, i.e., the distance step between cylinders is 20 cm. Therefore, the experiment can be denoted as 20AFC.

Fourteen volunteer observers participated in the experiment. All had normal or corrected to normal vision and were naive with respect to the purpose of the experiment. Subjects evaluated 25 stimuli images, one per material. There was no time constraint to finish the task, and a typical session took about 5 minutes. The resolution of images was 1920×1080 pixels and no resampling was applied. The material renderings were precomputed using a ray-tracing (256 samples/pixel) with texture mip-mapping enabled. All stimuli were presented on a calibrated 27" ASUS LCD display VG27AH (60 Hz, resolution 1920×1080 pixels), and as the pixel-size was 0.311 mm, the final displayed resolution was 82 DPI. The experiment was performed under dim room lighting and participants were advised to view the screen at a distance of 0.6 m conforming to the designed scene geometry.

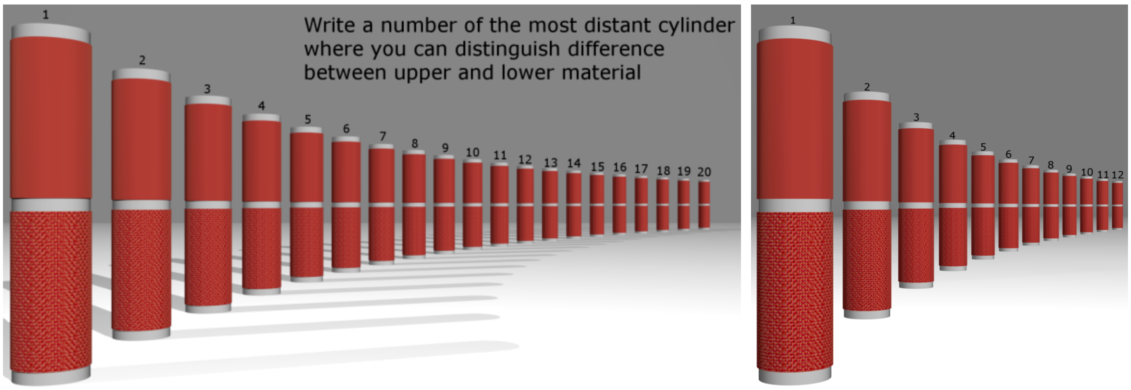

Figure 3.10: Left: An example of stimuli of the controlled *Experiment A* – 20 cylinders at distances ranging from 0.6 to 4.4 m (step 0.2 m). Right: An example of stimuli of the web-based *Experiment B* – 12 cylinders at distances ranging from 0.6 to 4.45 m (step 0.35 m).

Experiment B – Web-Based, Static Stimuli – To validate results of the controlled experiment and to generalize its conclusion for different displaying scenarios, we perform a second psychophysical study using a web-based testing interface. As the physical size of the scene has to be the same regardless of the size and DPI of a user's display device, we asked subjects to measure the width of a 500 pixels wide rectangle on the screen. This information allows us to compute DPI of the display and resize stimuli images to appropriate size. To avoid image blur due to down-sampling of the images on the observer's device, the images were precomputed in higher resolution of 2160×1868 pixels. As the size of a user's display device is often limited, we restricted the number of cylinders in the scene to 12; however, they span almost the same viewing distances (0.6 - 4.45 m) as distance between neighboring cylinders is 35 cm. Therefore, the experiment type can be denoted as 12AFC. This step allowed us to reduce a physical width of displayed stimuli images to around 29 cm, which is the width of a typical 13" screen (see the example stimulus image in Fig. 3.10-right).

The meaning of the stimuli images and the task for the subjects is the same as in *Experiment A*. The images were precomputed using the OpenGL rendering (256 samples/pixel) with a texture mip-mapping enabled. We applied high-quality filtering to suppress aliasing artifacts in rendering of the BTF. Forty volunteer observers participated in the experiment. Subjects evaluated 37 stimuli images, one per each material plus two additional variants for five anisotropic materials (material rotated for 45^o and 90^o) and two chessboard patterns. The subjects were strongly advised to observe the display from the distance of an outstretched arm (≈ 0.6 m) and to perform the experiment under dim lighting. The recorded average viewing parameters across all subjects are following: display resolution 1627×998 pixels, 96.7 DPI, pixel-size 0.279 mm.

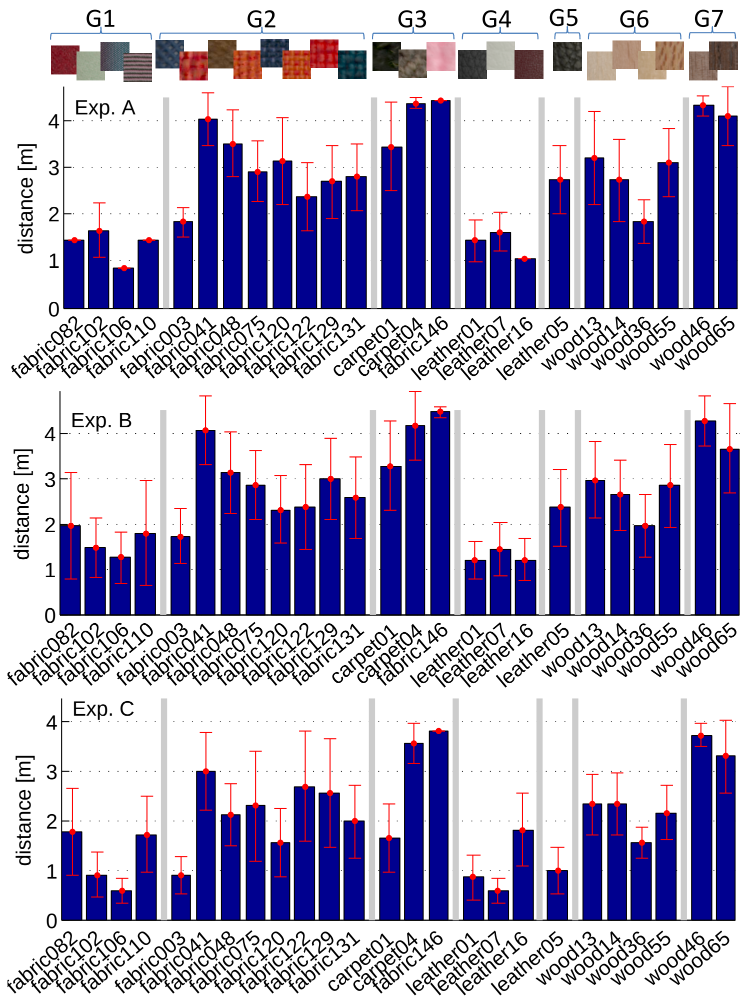

Figure 3.11: Perceptually estimated *critical viewing distances* where the average subject is still able to distinguish between the BRDF and the BTF. Included are results of all three experiments with standard deviations across all subjects.

Experiment C – Controlled, Dynamic Stimuli – Finally, we replicate 12AFC *Experiment B* in the controlled environment as defined in *Experiment A*. The only difference is that instead of a static scene we use rotating cylinders. The task for the subjects is the same as in the previous experiments, and stimuli images were computed using the OpenGL. Eight volunteer subjects participated in the experiment, and the majority of them participated in *Experiment A* too. Subjects evaluated 25 stimuli images, one per each material. Movies were stored uncompressed with 30 frames per second. To avoid visually distractive seams near the material overlap on the cylinder, which could produce artificial lower frequency content, we placed a patch over this seam along the whole cylinder.

3.3.2 Analysis of the Results

The first step in the analysis of the perceptual data is the filtering of unreliable results, whose description can be found in Section B.3. Next step is computation of the average *critical viewing distance* across all subjects. Figure 3.11 shows such data obtained from all three experiments across all 25 tested materials divided into seven groups as described in Section B.1. The error-bars represent standard deviation values computed across all responses.

Figure 3.12 depicts the estimated distances averaged across individual material categories. We observe that in the dynamic scene (*Experiment C*) the material structure becomes perceptually indistinguishable at closer distances (0.2 - 1.5 m closer) generally for all categories. The biggest difference is in groups that have rough surface structure (G3, G5, G7). We assume that the reason for this behavior may be in visual blurring of high-frequency features, present in these materials due to material motion. For more information see Section B.4, which confirms validity of web-based *Experiment B*, summarizes our findings about subjects' response time and discusses influence of anisotropy on the estimated distance.

For further analysis of subjects' sensitivity in structure recognition, we fit the data of each material group by a psychometric function [WH01]. For each observer, we record value 1 for those cylinders where the difference in structure is spotted, and 0 otherwise. The summarized data across all subjects are fitted by a Gumbel function (logaritmic Weibull function) using Palamedes psychophysical data analysis toolbox for Matlab (*http://www.palamedestoolbox.org/*). Figure 3.13 shows fitted functions for data of individual material categories from all experiments. We can observe close similarity of the results from the static stimuli experiments (A and B).

In accordance with Figure 3.12, the perceived structure first disappears for the category of *leather smooth* (G4), although a typical size of structural elements of leather is much larger when compared with smooth fabric. This can be explained by their low surface profile and thus limited contrast. Then follows the category of *fabric smooth* (G1) although its slope is considerably lower than the one of the leather, which is presumably due to higher variability of materials within this group. Similarly shaped and positioned psychometric functions are obtained for the groups of *fabric meso* (G2), *leather rough* (G5), and *wood*

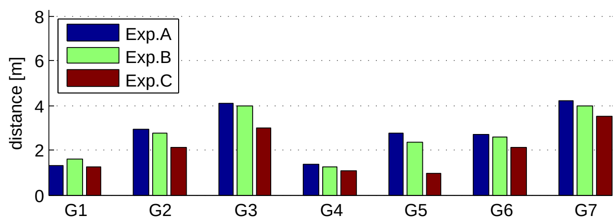

Figure 3.12: The distances averaged across the individual material groups obtained in the three psychophysical studies (A, B, and C).

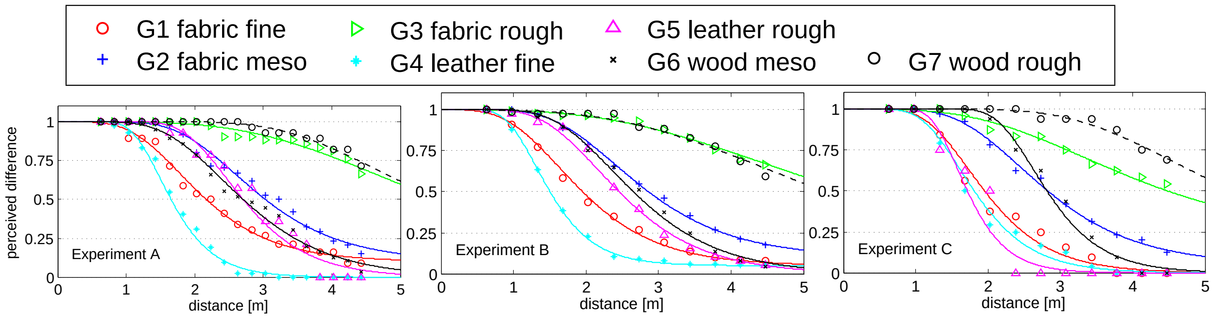

Figure 3.13: Psychometric functions that fit data obtained from all three psychophysical experiments (A, B, and C).

meso (G6). Finally, closely resembling functions are achieved for the groups of *fabric rough* (G3) and *wood rough* (G7), which mark completely different perceived distances.

The psychometric data from dynamic-stimuli *Experiment C* show slightly different behavior. Generally, the slants of the functions are steeper, which suggests a lower variance in subjects' responses. This is supported by the lower standard deviation values in *Experiment C* in Figure 3.11. The main difference is in the category of *leather rough* (G5), where its structure becomes significantly less apparent. Similarly, the structure of smooth wood materials (G6) is also less apparent. This behavior conforms with conclusions of Jarabo et al. [JWD*14] referring that for dynamic scenes the blur from the motion results to higher visual equivalence than in static scenes. On the other hand, for the category of *leather fine* (G4) the motion helps to recognize structure at a longer distance. We assume that this atypical behavior is caused by the flickering of fine visual features in the material structure that cannot be easily detected in static scenes. A discussion with the subjects after the experiment support this as they mentioned that they were not comparing the BRDF with the BTF, but they rather focused on the presence of the motion in the bottom part of a cylinder.

Finally, the estimated thresholds of the *critical viewing distance* for all three experiments and the individual material groups for significance level 25% and 50% are shown in Figure 3.14. The level 25% means that only 25% or less of observers spotted difference between materials on a cylinder at given distance. While the level 50% should provide values of a typical user, we consider the level 25% as a safe adjustment for applied rendering algorithms.

Relationship to the Contrast Sensitivity – As our work closely relates to the contrast sensitivity of the human visual system (HVS), we implemented the achromatic contrast sensitivity function (CSF) of Burbeck and Kelly [BK80]:

$$CSF(f_s, f_t) = 4\pi^2 f_s f_t \cdot e^{\frac{-4\pi(f_t+2f_s)}{45.9}} \cdot \left(6.1 + 7.3 \left|\log_{10} \frac{f_t}{3f_s}\right|^3\right)$$

and plot it as a function of a viewing distance that is present in our test scene. A temporal frequency f_t was set to 1 and a spatial frequency f_s in cycles-per-degree was obtained from

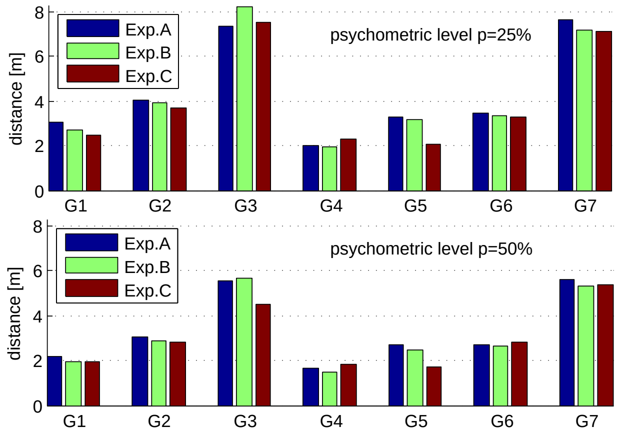

Figure 3.14: The distances obtained by psychometric functions with the thresholds at levels 25% (top) and 50 % (bottom). Results are shown for the individual material categories and all three psychophysical studies (A, B, and C).

the size of a material structural element. The result is shown in Figure 3.15. The CSF was evaluated for extreme and typical sizes of structural elements present in our materials (see Fig. B.2), i.e., 1 mm, 8 mm, and 3.4 mm.

When compared to the psychometric functions in Figure 3.13, it is apparent that although the CSF itself can predict sensitivity of the HVS to frequencies of typical materials relatively correctly, it is unable to predict reliably the sensitivity to materials of extreme sizes, where it tends to underestimate the *critical viewing distance*. This is probably due to lower texture contrast in appearance of the real-world materials, which is different from idealized sinusoidal-wave gratings assumed by the CSF as well as due to impact of illumination. Therefore, we focus on texture contrast features in our model of the *critical viewing distance*. Moreover, the tested spatiotemporal CSF shows that contrast sensitivity of the HVS decreases with an increasing temporal frequency, and thus guarantees the visual temporal coherence of renderings in applications relying on the estimated *critical viewing distances*.

Limitations of the Experiments – Due to limited DPI of a display, we filtered a certain portion of subjects' responses (see Sec. B.3) that correspond to more distant cylinders. This unfortunately slightly biases the *critical viewing distance* of materials with very fine structure (*fabric082*, *fabric106*, *fabric110*, *leather16*) towards ≈ 0.2 m lower distances. This filtering effect for these materials can be safely compensated in applications intended for standard screens with DPI around 100, where the users would face the same DPI limitation as subjects in the experiment.

When we compare the distances derived from psychometric scaling (Fig. 3.14) with simple averaging of the distances (Fig. 3.12), the psychophysically scaled variant provides higher perceived distances for more rough surfaces. This is due to a limited span of distances available in the experiment, which results in saturated subjects' responses for

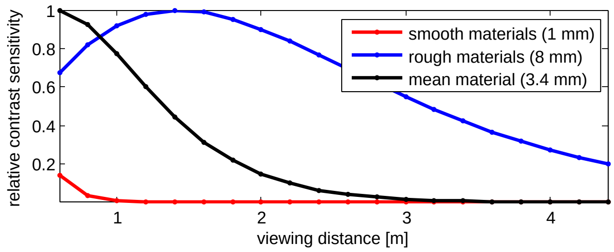

Figure 3.15: A relative contrast sensitivity as a function of viewing distances in our test scene for the materials with the smallest, the biggest, and the typical size of a structural element.

very rough materials, i.e., subjects were forced to select the last cylinder although they might be able to spot difference even further. Unfortunately, we could not increase a viewing distance due to a limited resolution of a screen, so a proper analysis of visibility of rough materials structure (a size of a structural element above 10 mm) becomes a subject of future endeavors.

Note that the type of motion in *Experiment C* is quite rare, as in many other scenes a geometry of an object combined with a movement would presumably visually mask these fine visual differences and thus increase the *critical viewing distance*.

3.3.3 Structure Visibility Prediction

This section compares obtained psychophysical results with the combination of standard image and computational metrics. The intention is to understand clues behind the visual perception of material structure. To this end, we analyze the correlation of psychophysical results with several promising features. We use data from *Experiment B* due to the highest number of subjects (40) as our golden standard.

The tested computational features are: (1) the average luminance of the cylinder, (2) the contrast approximated by the mean standard deviation of a difference between the BTF and the BRDF, (3) $1-\mathrm{SSIM}$ between the BTF and the BRDF, (4) $1-\mathrm{VDP2}/100$ between the BTF and the BRDF, (5) the luminance contrast C_L, (6) the texture contrast C_T, and (7) the saliency obtained as the linear combination of the last two features $C_L + 10 \cdot C_T$ [PN04]. We have also tested features based on luminance and chromatic contrast estimated from pixel-wise cone responses [WM97]; however, their performance is relatively poor. All the selected features are computed from the gray-scale bottom area of a cylinder in stimuli images from *Experiment B*, and values of most of them are obtained by the averaging of values across the image plane.

The correlation between the experiment and the features are shown in the left part of Table 3.1. The obtained values are not very high, which suggests that there is another element perceived by subjects not fully captured using the tested features. As an important factor impacting the recognition of material structure is its scale, we test the largest size

Table 3.1: Pearson correlations between results of *Experiment B* and the tested statistical features.

	Feature	ρ	p-Val		Feature	ρ	p-Val
1	luminace	0.067	0.749150	9	S·luminace	0.499	0.011103
2	contrast	0.583	0.002248	10	S·contrast	0.777	0.000005
3	1-SSIM	0.564	0.003326	11	S·(1-SSIM)	**0.878**	0.000000
4	1-VDP2/100	0.649	0.000452	12	S·(1-VDP2/100)	**0.808**	0.000001
5	lum. contrast	0.616	0.001037	13	S·lum. contrast	**0.941**	0.000000
6	tex. contrast	0.506	0.009898	14	S·tex. contrast	**0.934**	0.000000
7	saliency	0.591	0.001875	15	S·**saliency**	**0.943**	0.000000
8	structure size (S)	0.677	0.000203				

Table 3.2: Estimated linear parameters that fit three selected features to the data from *Experiment B*.

Feature	k_1 (std)	k_2 (std)
S·(1-SSIM)	1.582 (0.044)	1.748 (0.053)
S·(1-VDP2/100)	1.777 (0.033)	18.694 (0.384)
S·saliency	1.292 (0.023)	1.201 (0.015)

of a structural element as an additional feature (8) (see Fig. B.2). As expected, this gives the best correlation value so far. Therefore, we use the structure size to scale results of the individual tested features. The correlation values, shown in the right part of Table 3.1, show a very strong correlation, especially for scaled variants of the SSIM, the VDP2, and the image saliency.

While these features seem to be very promising candidates for prediction of the *critical viewing distance*, they are just proportional to experimental data and have very low values for materials where computed difference (SSIM, VDP2) or contrast are too low. Therefore, we suggest to adapt the values of the features (F) to the result of an experiment (E) using a simple linear model $E = k_1 + k_2F$. We use the robustfit function in Matlab for computation of parameters. To check stability of the parameters, we perform leave-one-out validation, i.e., a distance value for a particular material is obtained from parameters computed on values of all-but-this material. Obtained standard deviation values of such 25 sets of parameters are very low. The mean k_1, k_2 values and their standard deviations are shown in Table 3.2.

This analysis reveals that the most stable parameters are achieved for the image saliency feature. Another advantage of this feature is that, due to its Gaussian pyramid, it does not rely so much on the selection of a structure neighborhood size (SSIM) or information on a display and observer's distance (VDP2); see Table 3.3. Typical times needed for evaluation of individual features on our stimuli images are: 0.005 s for the SSIM, 0.270 s

Table 3.3: Parameters of individual tested methods.

SSIM	$C_1 = 0.01$, $C_2 = 0.03$, Gaussian window size 11 pixels, $\sigma = 1.5$
VDP2	screen size: 24", resolution 1900×1200, viewing dist. 0.6 m
saliency*	Gaussian window size 11 pixels, $\sigma_c = 0.5$, $\sigma_s = 2.5$, $\beta = 2$, $s = 5$

* our implementation in Matlab is available at http://staff.utia.cas.cz/filip/projects/15CGF.

for the VDP2, and 0.082 s for the saliency.

The predicted distances are shown in Figure 3.16 together with errors, where red bars show the proportion of underestimation of the correct *critical viewing distance*, while the green ones its overestimation. The best performance is achieved by the image saliency feature (c). When compared with distances obtained from *Experiment B* (Fig. 3.11), it is apparent that errors of our model are within standard deviations across subjects' responses. This supports our conclusion that the proposed image saliency feature [PN04] scaled by the structure size provides a promising and computationally reasonable model of visibility of material structure.

Finally, Figure 3.17 compares values of the *critical viewing distances* as obtained from *Experiment B* (E) and from the proposed prediction (P) for typical representatives of the individual tested material categories.

3.3.4 Applications

Although position of an observer with respect to viewed objects in virtual reality is not constrained in general, we assume that there exist typical constrained viewing scenarios. For instance, a car interior is typically viewed and rendered from driver's perspective. Similarly, interior of a room in virtual-walk-through applications is visualized from perspective of its visitor usually located near an entrance or a center of the room. In other words, we assume that our method can be beneficial in any application where constraints of the viewing distance can be imposed or assumed. As a consequence, it allows real-time rendering of more complex scenes (with more materials) using the same HW.

Even though we would assume a non-constrained viewing scenario where it is necessary to store the full BTF in a memory, our method can reduce computational cost related to a spatial interpolation of the BTF when the viewing distance exceeds the critical one. Although this interpolation is often supported by HW (depending on data structuring), its removal can save typically four BTF reconstructions (needed for the spatial interpolation) and substitute them by a single cheaper BRDF reconstruction. Although one can claim that our results can be substituted simply by a texture mip-mapping, we argue that our results are more restrictive as the estimated *critical viewing distance* is in majority cases shorter than the one corresponding to the lowest mip-mapping level (i.e., mapping the entire material structure onto a single pixel).

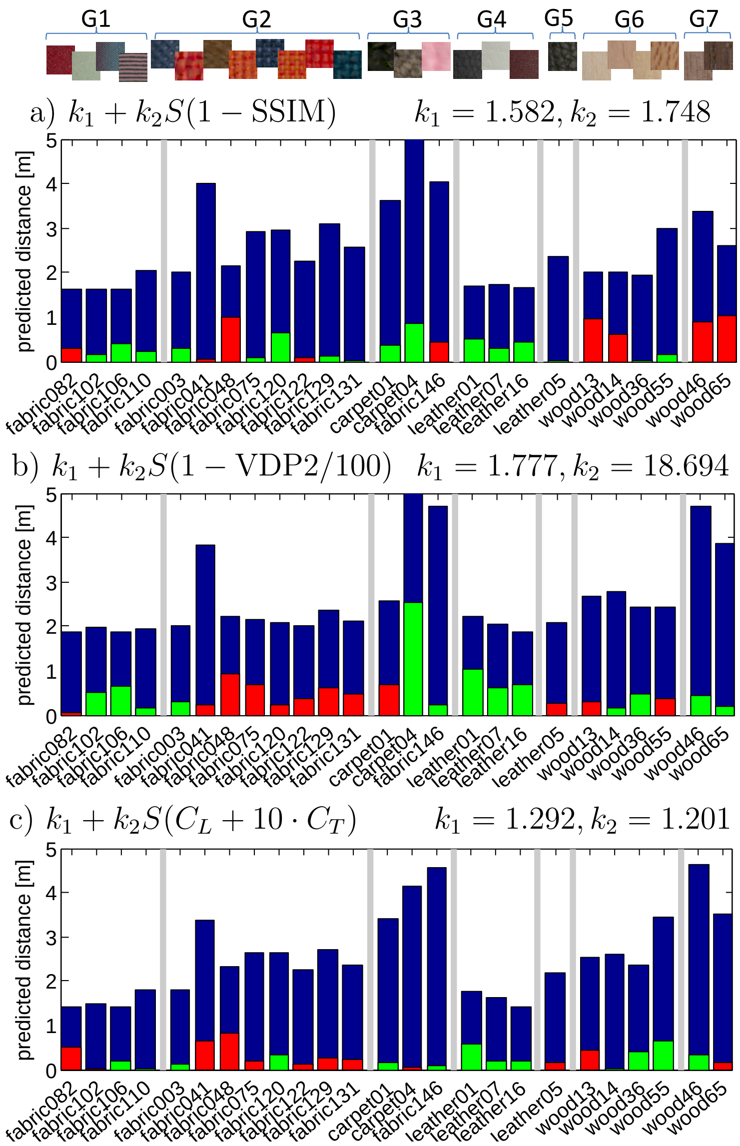

Figure 3.16: A comparison of the prediction of the *critical viewing distance* that uses the features based on (a) the SSIM, (b) the VDP2, (c) the image saliency. Red bars show the proportion of underestimation of the correct distance while the green ones its overestimation.

To support our claims, we created a virtual scene that consists of chairs in a room depicted in Figure 3.18. The chairs are organized in six rows, each consisting of twelve chairs. The upper cushioned parts of chairs are covered by fabric or leather materials while the bottom construction is covered by wood materials. The floor consists of two different carpets and walls are represented by wood (total 14 materials). The depth range of the room is 12 m (see Fig. 3.18-c) and materials are mapped on the objects in a physical scale.

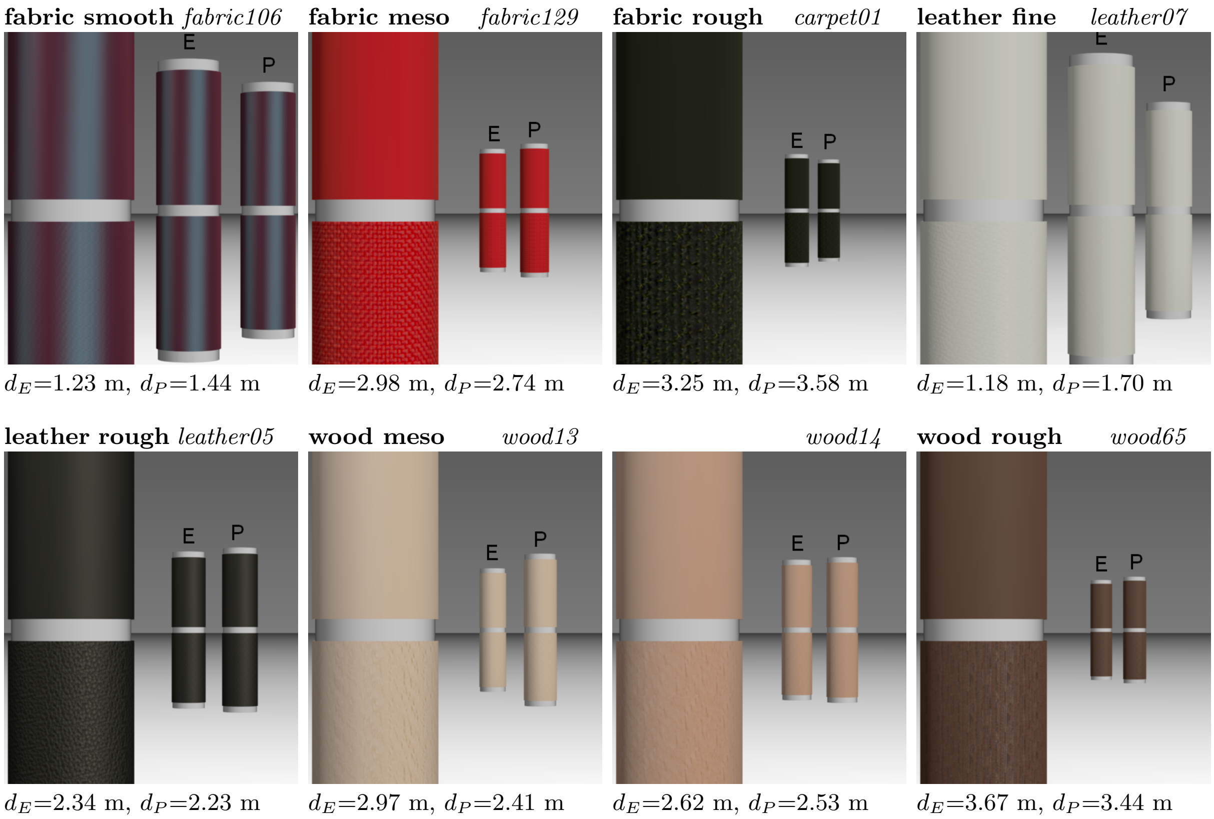

Figure 3.17: An example of the *critical viewing distance* that compares the 0.6 m distant reference cylinder (left) with cylinders in distances that correspond to the average subjects' responses from *Experiment B* (E) and in distances of the statistical predictor (P). To conform with the real scene geometry, the image should be zoomed so as the width of the left reference cylinder subtends 59 mm on the screen and the screen is assumed to be viewed from the distance of 0.6 m.

The scene is illuminated by two point-light sources. First, we rendered the scene in the full BTF representation as it is shown in Figure 3.18-a. Then, we implemented an OpenGL shader that switches rendering of individual pixels from the BTF to the BRDF where the viewing distance exceeds the predicted *critical viewing distance* d_E for a given material. The BTFs in the scene are filtered to remove aliasing artifacts. The result of the shader is shown in Figure 3.18-b and its 30× scaled difference image to the rendering of only the BTF is shown in Figure 3.18-d.

One can see that visual differences are negligible and the PSNR between the full BTF rendering and our approach is 58.6 dB as shown in Table 3.4, where is also highlighted the gain in computational speed. While the full BTF rendering of the scene in resolution of 1920×1080 pixels achieves 8.3 FPS, the BRDF rendering (as our upper bound) achieves 20.2 FPS, our method that combines both approaches based on the predicted *critical viewing distance* achieves 11.8 FPS. This represents 30% gain in speed when compared to the

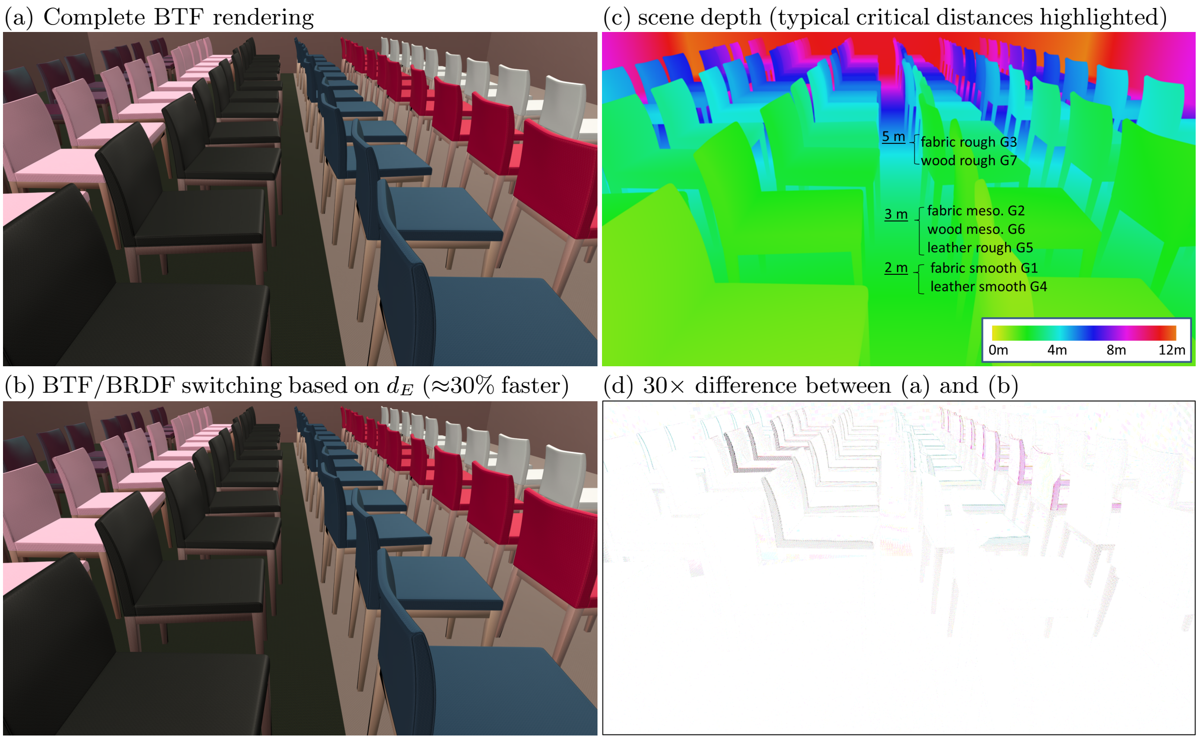

Figure 3.18: A scene that consists of six rows of chairs and 14 materials: (a) rendering using the BTF only, (b) switching between the BTF and the BRDF based on the proposed model of the *critical viewing distance* d_E, (c) range-map of the scene (depth 12 m), (d) 30× scaled difference between (a) and (b). The proposed approach achieves 30% faster rendering performance while maintain high visual fidelity to the full BTF rendering (PSNR 58.6 dB)

full BTF rendering without significant loss of perceived fidelity. Although these values will vary across different scenes, we consider this experiment as an important showcase application of the proposed *critical viewing distance*. We have not encountered any temporal discontinuities in our material appearance renderings.

All tests were performed on a PC with the Intel Core i5-2500 3.3GHz, 16GB RAM and the nVidia GeForce GTX 570. The relatively low speeds in Table 3.4 result from the high

Table 3.4: Measured rendering speed (in frames-per-second) of a scene that consists of 1920×1080 pixels with various material appearance representation approaches, and the PSNR of both approaches related to the reference full BTF method.

method	BTF	BRDF	proposed
FPS	8.3	20.2	11.8
PSNR [dB]	–	46.3	58.6

screen resolution (Full HD), where all pixels are evaluated. Further, as we do not use an analytical BRDF model, we have to interpolate the BRDF measurements (from 9 values using the barycentric interpolation). Finally, we use just unoptimized shader and data are reconstructed for each light source independently, which also prolongs the rendering time.

Another application of our findings can be an instrumental tool that would help interior designers to select materials according to their distance and intended appearance in the virtual environment. The *critical viewing distance* can create a map of the scene (similar to Fig. 3.18-c), e.g., it can highlight objects where the fabric structure would be visible from a given viewpoint.

3.4 Summary

A novel method for convenient and fast detection of strength of anisotropy, main anisotropy axes, and corresponding shape of highlights is introduced in Section 3.1. The proposed detection setup consists of an ellipsoidal mirror and a compact camera, which allow for almost a real-time analysis of material anisotropy without necessity to extract a sample from its environment. We show the application of parameters estimated using our method to improvement of sampling of an unknown BRDF. Furthermore, we show that the fast measurement of some subspaces of the anisotropic BRDF is possible due to inclusion of a projector to the setup. We believe that the proposed fast analysis of anisotropic appearance will be beneficial for a variety of tasks, ranging from improved acquisition of material appearance to automatic classification of materials or image-based material retrieval.

The knowledge on typical behavior of anisotropic materials helped us to design a novel empirical anisotropic BRDF model presented in Section 3.2. The model fits individual modes of anisotropy independently using adaptive stencils of the anisotropic highlights. Therefore, no time-demanding nor potentially unstable numerical optimization is needed. Moreover, an intuitive adjustment of appearance is possible and the model is not limited in terms of a number and an orientation of the anisotropic highlights.

The main objective of Section 3.3 is to analyze visibility of structure of materials in virtual environments as a function of viewing distance. Identification of the so-called *critical viewing distance* (i.e., a material-dependent maximal viewing distance where material structure or texture can be visually distinguished) is important in the decision whether to replace the BTF by the BRDF. An application of the *critical viewing distance* can save computational costs in situations where the BTF may be substituted by the BRDF without loss of visual fidelity. To this end, we estimated such distances for 25 interior materials in two static and one dynamic psychophysical studies with up to forty participants. We divided our materials into seven categories and identified the *critical viewing distances* for each of them. The *critical viewing distance* ranges from 2 m for smooth fabric and leather through 4 m for rough leather, moderately structured fabric, smooth wood to 8 m for rough fabric, carpet, and wood.

Moreover, we analyze a number of computational features to derive a reliable model of psychophysical results applicable as a predictor of the *critical viewing distance* for a new material. The best performance is obtained by the feature based on the image salience scaled by the size of the largest element in structure of a material.

Our results can benefit many applications that deal with rendering of complex virtual scenes. We show that application of the *critical viewing distance* can save rendering costs significantly. Further, it can serve as a useful tool for interior designers to help them select materials according to an intended appearance in a virtual environment. In future efforts, we plan to analyze the impact of an illumination environment and investigate whether it relates to the *critical viewing distance* or not.

Chapter **4**

Measurement Setup and Data Processing

This chapter presents a measuring device (Sec. 4.1) that we use for adaptive and reference measurements described in the following chapters, a practical registration method of the measured data (see Sec. 4.2), and an interesting method that overcomes the main disadvantage of gonioreflectometers, the long measurement time, by collecting many BRDF samples from each photograph (see Sec. 4.3).

4.1 Measurement Setup

At the Pattern Recognition department of UTIA, we have built a highly precise robotic gonioreflectometer (see Fig. 4.1). The device was first introduced in [A.8], and later on, reintroduced in [A.14]. It consists of independently controlled arms with a camera and a light source. Its parameters such as angular precision better than 0.03° or a spatial resolution up to 1071 DPI classify this gonioreflectometer to the state-of-the-art devices. The typical resolution of an area of interest is around 2000×2000 pixels and each of them is represented by means of at least a 16-bit floating point value for reasonable representation of visual information with the high dynamic range. The memory requirements for storage of a single material amount to 360 giga-bytes per color channel. Three RGB colour channels require about 1 tera-byte. More precise spectral measurements with moderate sampling of the visible spectrum (400-700 nm) further escalate the amount of data to 5 tera-bytes or even more. Such measurements allow to capture very fine meso-structure of individual entities which the material is made from.

Mechanical Construction – The setup consists of the measured sample held by a rotary stage and two independently controlled arms with a camera (one axis) and a light source (two axes) as shown in Figure 4.1. It allows for flexible and adaptive measurements of nearly arbitrary combinations of illumination and viewing directions. Although occlusion of view of the camera by the arm with the light source may occur, it can be analytically

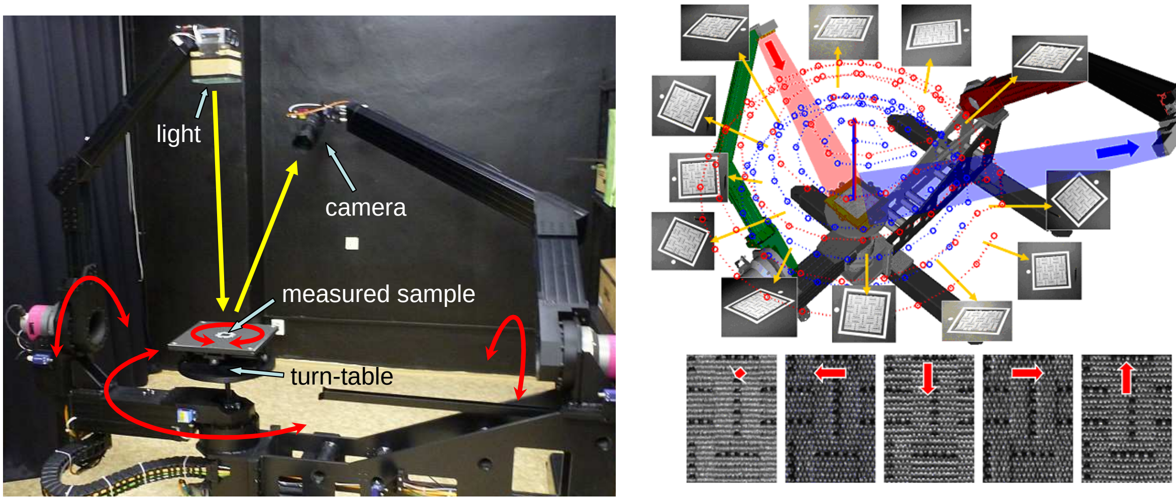

Figure 4.1: A photograph of the UTIA gonioreflectometer (left) and its schema (right). The total number of images is equal to every possible combination of red dots (view directions) and blue dots (illumination directions). Below is an example of how the visual appearance of a fabric is dependent upon illumination direction. The fabric is illuminated from above, right, top, left, bottom, respectively.

detected, and in the most cases an alternative configuration exists. Positioning accuracy of the arms is $\pm 0.03°$ across all the axes.

The Light Source – Due to advantages of color stability, long-lifetime and reasonable heat dissipation, we formerly used a custom build array of 11 LEDs Cree XML, each having flux 280 lm at 0.7 A (max. current 3 A). Each LED was equipped with its own optics producing a narrow and uniform beam of light. Moreover, a light diffuser can be attached with various types of matte glass to produce uniform illumination across the area of a sample. Currently, we use a single-area LED light source that allows current 3 A.

The Imaging Device – The industrial full-frame 16 Mpix RGB camera AVT Pike 1600C used in the setup contains a Kodak KAI-16000 CCD sensor. Resolution of the camera is 4872×3248 pixels, bit depth is 14 bits/pixel and the shortest integration time is 0.6 ms. Distance of the sensor from the sample is ≈ 2 m. By usage of different lenses, we can achieve spatial resolutions listed in Table 4.1.

Acquisition in the High Dynamic Range – is achieved by adaptive exposure times and variable intensity of illumination (through the current fed into LEDs); both are controlled remotely depending on the dynamic range of the measured sample. Due to the high bit-depth of the sensor, an extensive range of used exposure times and possibility to adapt intensity of illumination, the dynamic range of each photograph can be up to 28.4 EV.

System Calibration and Data Processing – The initial position of all the axes can be found using a spirit level, plummet, intersections of the axes of the device, etc. Non-uniformity of illumination over the target area is measured using a lux-meter, fitted

Table 4.1: A list of available lenses, resulting spatial resolutions, and maximal sizes of the measured samples.

Focal Length [mm]	DPI	μm/pixel	Sample Size [mm]
50	90.7	280	554×554
180	353.3	71.9	139×139
300	602.1	42.2	70×70
500 + extension tubes	1070.9	23.7	40×40

by a 2D polynom, and compensated from the photos. Defective pixels of the camera are detected and interpolated in the photos. Internal Bayer pattern in the raw float data is interpolated into RGB using a local linear interpolation. Vignetting of lenses is estimated from a photo of a white sheet and its non-uniformity is compensated. The colorimetric calibration 3×3 matrix is obtained by solving a set of linear equations, which relate the known and the measured color patches on the Xrite target in the CIE XYZ colorspace. Images taken in the OpenEXR floating point HDR format from different views are mutually pixel-wisely registered by means of detection of lines on a registration target that surrounds the measured material (see Sec. 4.2).

System Control – The complete measurement procedure is automatically controlled by a single server. The controlling application stores the list of the required measured positions, which can be adaptively modified during a measurement. Measured data are stored on a 20 terabytes disk-array and accessed via speedy 10 Gbit optical LAN. Mechanical positioning, exposure, and data transfer of 6,561 measurements typically take around 18 hours, mainly due to exposure times.

4.2 BTF Data Registration

In this section, we summarize the main contributions of the paper [A.15] that introduces a novel registration method of multi-view data. Processing of the acquired multi-view data consists of two steps: data registration and compression. Although the measured materials are planar, their rough structure often shows height variations causing significant variance of their appearance depending on illumination and viewing directions. The final appearance is affected by self-occlusions, shadows, inter-reflections, and subsurface-scattering. This is the reason why the features of the material are non-stationary and cannot be directly used for reliable feature-based registration. Due to this, we use registration marks placed on a reference plane, which allows the measured sample to be easily replaced. However, an orientation and a shift of the sample with respect to the registration plane are unknown (see Fig. 4.2). Although one might use a tilt/shift mechanical stage to fine-tune this misalignment manually, it is expensive and far less accurate than the proposed approach.

The registration based on the registration marks properly aligns only those parts of the measured sample which lie close to the plane specified by the marks (see Fig. 4.2-b). This

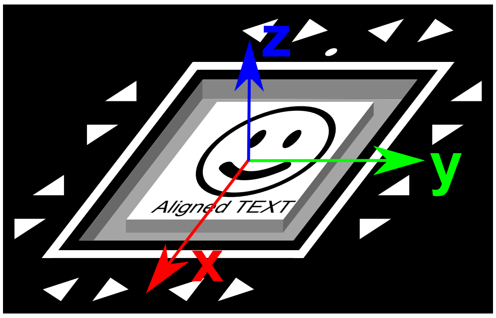

(a)

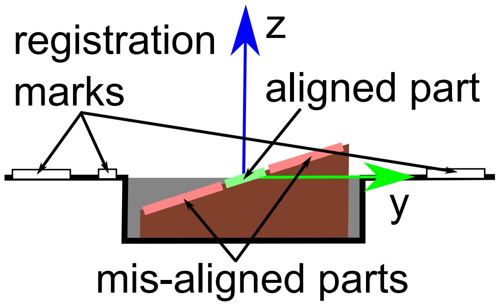

(b)

Figure 4.2: (a) A measured sample, the reference plane with the registration marks (the white frame and triangles) and the world coordinate system. (b) A cross-section of (a) shows aligned and misaligned parts of the surface of a measured material when the orientation of the reference plane and the measured surface differ.

is not an issue when distance between the reference plane and the surface plane is small and the slant difference is also negligible. Also, this might be less relevant when registered data are of low resolution or are used directly for a rendering. However, it has significant impact on the resulting visual quality if compression methods for multi-view data are used (e.g., all the classes of global factorization methods based on PCA [SSK03] or clustering [HFM10], etc.).

The main contribution of this section is a novel technique for registration of multi-view images of planar surfaces that aligns measured planar surfaces regardless of their slight height and slant differences from the reference registration plane. Such data alignment allows us to better exploit the power of compression techniques and to produce sharper images.

4.2.1 Overview of the Proposed Registration Method

A principle of image registration (Fig. 4.4-a) is usually as follows. Given a set of photos of the same planar surface, registration applies a projective transformation to all the photos so that the features of the planar surface are aligned across all of the transformed images. For more information on the standard registration approach see Section C.1. When the standard registration approach is used, the features which do not lie in the registered plane are not aligned after application of the projective transformation. The same pixel then corresponds to different physical points on the surface of the material (see Fig. 4.3-a,c). Therefore, we have to estimate the plane which represents the surface of the measured sample for appropriate registration (Fig. 4.3-b,d). Unfortunately, this plane (i.e., its offset and orientation with respect to the world coordinate system) cannot be determined accurately enough directly from the acquired photos.

We propose to find the position and the slant of the surface plane as follows. First, the

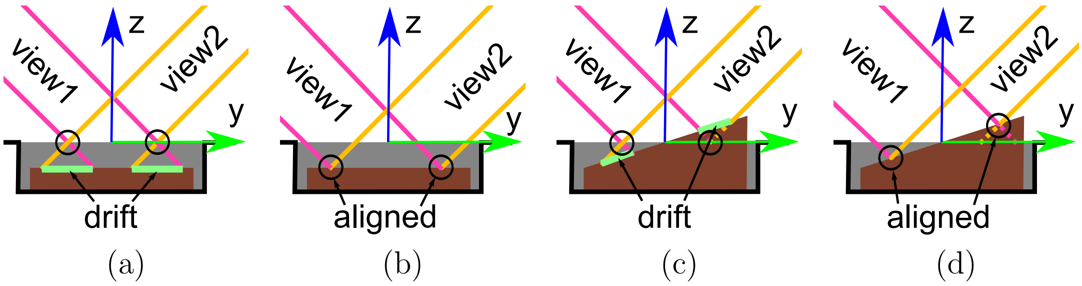

Figure 4.3: When the plane defined by the registration marks is chosen for registration, the same pixel corresponds to different points on the surface of a measured sample (a,c). When the appropriate plane on the surface of the sample is chosen, the drift disappears (b,d).

reference plane defined by the registration marks is taken. As we expect that the estimated plane which represents the surface of the material is close to the reference plane, a new hypothetical position of the estimated plane can be generated using a slight modification of the position of the reference plane by shifting it in a direction of its normal vector and by tilting it. Finally, the best estimation of the shift and the slant of the plane can be found by repetitive adjustment of the position and the slant, projective transformation of the photos, and evaluation of alignment of the surface features. A principle of the proposed registration method is expanded upon in Figure 4.4-b. The method consists of two main parts discussed in more detail in Appendix C. The first one is estimation of the camera intrinsic and extrinsic parameters (Sec. C.2) and the second is the iterative fitting of the position of the estimated plane (Sec. C.3).

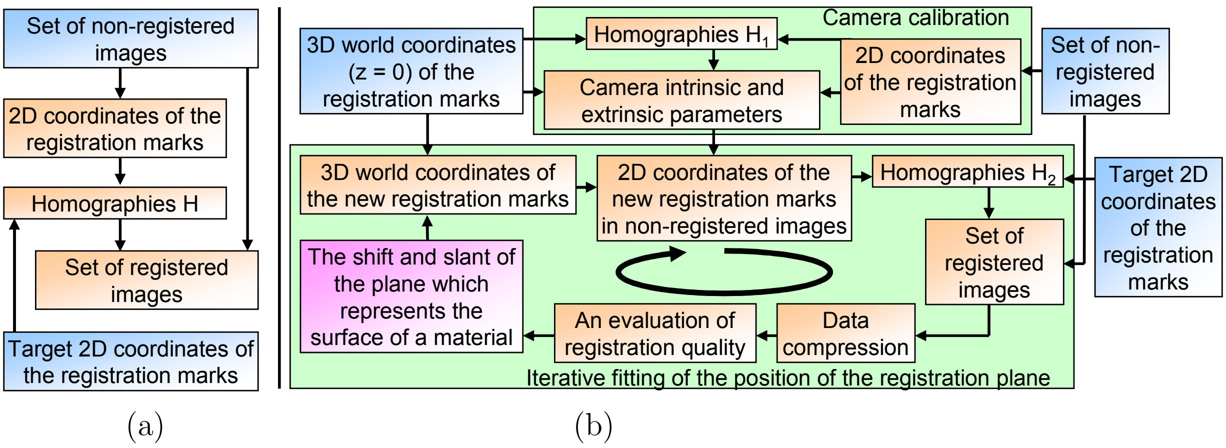

Figure 4.4: The standard image registration approach (a), and a scheme of the proposed method (b).

4.2.2 Results

This section illustrates performance of the method on two registration experiments using artificial and real data. In the experiments, we use the PCA compression of all registered images [KMBK03] and apply its data reconstruction error as an evaluation function of registration performance. All pixels selected from the individual BTF images are ordered into vectors and centered using the mean BTF image vector. All these vectors form a matrix, whose PCA is computed. The individual eigenvalues from the resulting diagonal matrix weight the importance of the resulting eigenvectors. A limited set of k eigenvectors is used to reconstruct the original n images, where $k << n$. The PCA-based methods are the most common in multi-view data compression; however, any other global BTF data compression technique would also benefit from the proposed approach.

In the first experiment, a flat paper printout is used, positioned approximately one millimeter below the reference registration plane. We took 80 different views on the plane which uniformly covered a hemisphere of viewing directions. The illumination direction was fixed in a direction opposite to the normal vector of the reference plane. When the standard registration approach is used, only features of the reference plane are aligned, while the misalignment in individual images causes that the mean image of all the registered images is blurred in the area of the measured sample (see Fig. 4.5-a). In contrast, when the proposed approach is applied, we obtain the mean image shown in Figure 4.5-b, where the desired surface is aligned well but the registration marks are blurred. The estimated surface plane is 1.24 millimeters bellow the reference plane and its normal vector is tilted for 0.29° with azimuth of 176°.

In the second experiment, the BTFs of five materials are registered using the standard and the proposed approaches. The BTFs of *wood01*, *fabric01*, *fabric02*, *fabric03*, and *leather01* are taken from the UTIA BTF Database (see Appx. A). The results of our

standard registration

proposed approach

Figure 4.5: An extreme example of the registration based on the reference plane (left) and based on the plane which represents the true surface of the sample found using the proposed method (right) when the shift of the surface plane differs considerably. Mean images of the 80 registered images are depicted here.

Table 4.2: The estimated values of the shift (z) and the slant (θ,φ) for the tested samples.

sample	shift z [mm]	elevation θ	azimuth φ
wood01	-0.49	1.15°	101°
fabric01	-0.20	0.15°	128°
fabric02	0.12	0.20°	260°
fabric03	-0.30	0.43°	323°
leather01	0.10	0.05°	354°

method are shown in Figure 4.7 and mark a considerable improvement over the standard registration approach. The compression of the data registered in a standard way (Fig. 4.7-b) leads to visualization that is less sharp in comparison with the non-compressed data (Fig. 4.7-a) that are considered to be the ground-truth. The compression after application of the proposed method leads to considerably sharper images (Fig. 4.7-c). Note that in both cases the same compressed parametric representation is used (50 PCA components, which allow a real-time rendering).

Processing of such a BTF sample that comprises $6,561$ images typically takes around five hours on 6 cores of Intel Xeon 2.7 GHz using our Matlab implementation. However, due to the massive size of datasets (415 GB), much of this time is consumed by disk data transfer operations. Note that a smaller visible area of non-aligned datasets (Fig. 4.7-b) is due to cropping of visual artifacts at borders of individual misaligned images. As the proposed alignment method re-projects original locations of registration marks, the registered images are slightly shifted and scaled. Therefore, their fair pixel-wise comparison (e.g., using RMSE, SSIM) with the original image is impossible. As there is no robust texture-similarity measure available, we performed a psychophysical experiment with 5 naive subjects who compare Figure 4.7-a with Figure 4.7-b and Figure 4.7-c in a random order. The (c) was always perceived as more visually similar to the (a) than the (b).

Table 4.2 shows estimated values of the shift and the slant of the estimated plane with respect to the reference plane. The images show that the more the plane of a sample deviates from the reference plane, the higher visual improvement is achieved as, e.g., for the sample of *wood01* in Figure 4.7. From the shown values, it is apparent that even when the sample is mechanically aligned with the registration plane as much as possible, the estimated differences are still relatively high. Finally, we remark that the visual effects of such misalignment are more pronounced if the resolution of captured images is higher.

Speed of the algorithm depends on the size of user-defined patches on the planar surface that are used for evaluation of registration quality, as well as on the number of multi-view images. Figure 4.6 shows execution time for a single iteration of the algorithm, which depends on the number of processed pixels and images. While the speed increases almost linearly with the number of pixels, it depends on the number of images n with $O(n^3)$ due to the PCA compression.

The proposed method is very robust. Its only obvious limitation is that it cannot

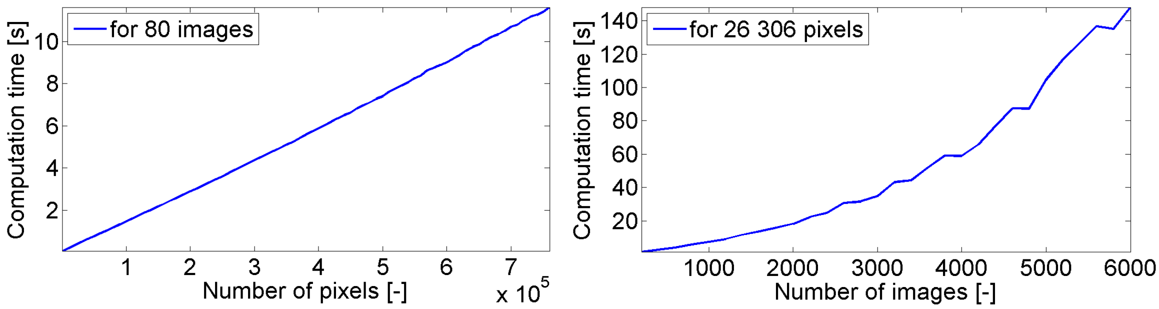

Figure 4.6: Computational time of one iteration of the method depends on the number of pixels used for quality evaluation (left) and on the number of processed images (right).

guarantee a correct alignment for surfaces that have wide height variations or several possible alignments (see, e.g., material *fabric03* in Fig. 4.7). However, even in such a case, the material is aligned in a way to minimize the compression error. Additionally, only pixels which belong to the required height can be selected by a user-defined mask and can be taken into account during the registration to achieve even better alignment for such materials.

4.3 Acquisition of the Dense Anisotropic BRDF

Our primary motivation for the angularly-dense BRDF measurement proposed in [A.13] and summarized in this section is the lack of those measurements suitable not only for fitting of analytical models and their development, but especially for development and evaluation of adaptive measurement techniques of anisotropic materials. Consequently, we use the highly precise gonioreflectometer-based setup (see Sec. 4.1) for the BRDF measurement. We propose to extend an idea of the isotropic BRDF measurement [RMS*08] based on capturing variable incoming and outgoing directions in images of a flat sample. We assume that the size of a measured material is comparable to its distance to a light source and a camera. Due to varying illumination and viewing directions across the surface of the material, a large variety of incoming and outgoing directions can be measured.

We show the way a gonioreflectometer can be combined with the image-based methodology to achieve a reasonable measurement time while still yielding very accurate dense BRDF data. Through interpolation, we achieve the uniform angular step of $2°$ in all four dimensions. We acquired and made publicly available BRDF measurements of three anisotropic fabric materials (see Sec. A.3).

4.3.1 Acquisition Approach

The image of the measured material is taken $8,505$-times. Where $6,561$ images represent uniform sampling of the hemisphere above the material surface using 81×81 viewing and

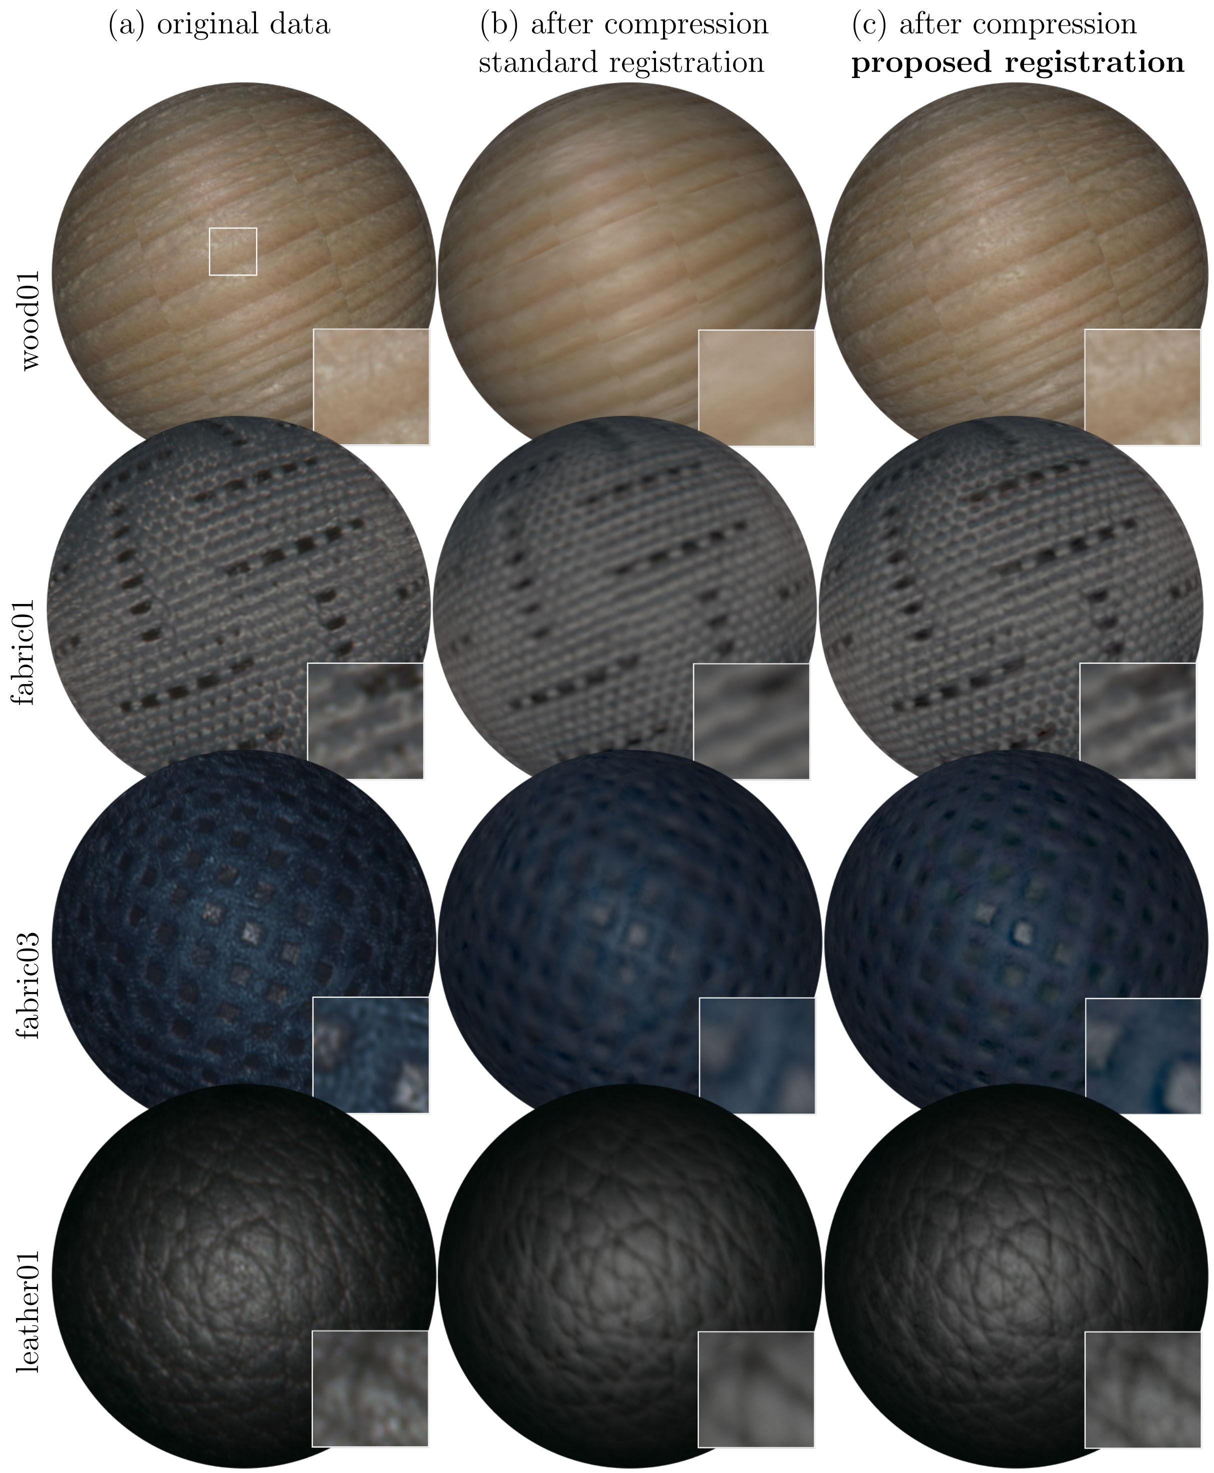

Figure 4.7: Visualization of the BTF: (a) a rendering using all 6, 561 images, (b,c) renderings from a compressed representation using only 50 eigen-images without and with application of the proposed method, respectively.

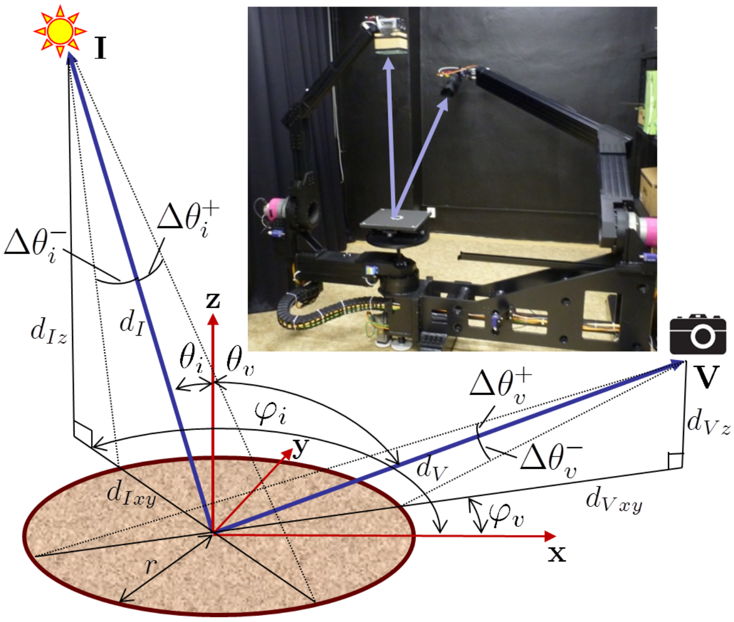

Figure 4.8: The setup and a scheme of the BRDF measurement based on variations of viewing and illumination directions across the measured surface.

illumination directions uniformly covering the hemisphere [SSK03]. The remaining $1,944 = 81 \times 24$ samples represent additional oblique illumination directions at the elevation angle $\theta_i = 85°$ with the azimuthal step of 15°.

The measurement principle is based on the fact that an image of a flat surface of the size relative to its distance to a light source and a camera captures a variety of illumination and viewing directions. This is due to varying elevation and azimuthal angles across the surface of the measured material, as shown for elevation angles in Figure 4.8.

Our measurement target consists of a flat desk with registration marks and the measured material as illustrated in Figure 4.9. The measured fabric samples were flattened and stretched using pins. We use the Sobel filter for detection of edges of the marks and the Hough transformation for detection of lines. Intersections of the lines are used for computation of the homography and rectification of the images. The effective radius of the material is $r = 277$ mm, and we sample it in the rectified images using a uniform grid of samples 14 mm apart (see Fig. 4.9). All samples are bined to a matrix of dimensions $|\theta_i| \times |\theta_v| \times |\varphi_i| \times |\varphi_v| = 44 \times 44 \times 180 \times 180$, i.e., with the angular step of 2° and the maximal elevation is 86°. Sampled values are computed as medians from neighborhood of a size of 14×14 mm (50×50 pixels) to suppress imperfections in geometry of the material (folds, textile defects, etc.). The resulting number of samples balances well their overall density to their visual contribution in the final data. The slight spatial non-homogenity is the main source of noise in the data; however, it can be effectively filtered by computation of the median of samples that are present in each bin.

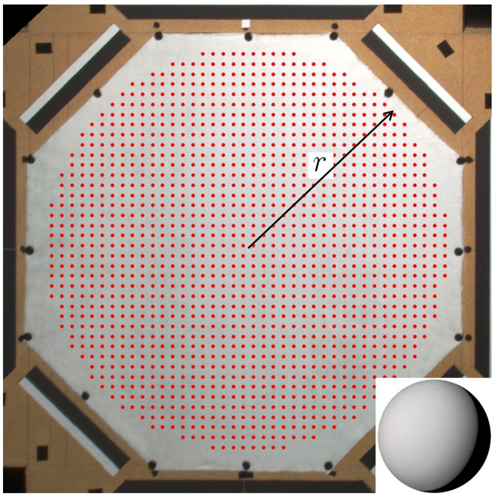

Figure 4.9: A measured material of the width of 700 mm, with uniform distribution of sampling points 14 mm apart within a circle of the radius of 277 mm. Bottom-right: An inset image shows rendering of the reference material used for compensation of non-uniformity of illumination.

4.3.2 Measured Data Evaluation

To test our approach, we measure three highly anisotropic materials: *fabric112*, *fabric135*, *fabric136* (see Sec. A.3). We analyze quality of the measured datasets in several ways. Specifically, we compare our data with: photographs of the measured materials, a sparse uniform sampling, results of two state-of-the-art BRDF models, and the ground-truth data in an applied situation.

Comparison with Photographs – A comparison of the measured BRDF with real

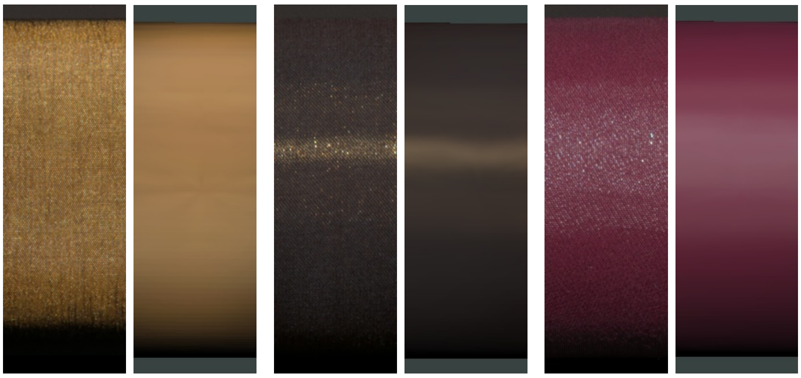

Figure 4.10: A comparison of photos of the measured materials (left) and renderings of their BRDF (right) on a cylinder illuminated from above a camera.

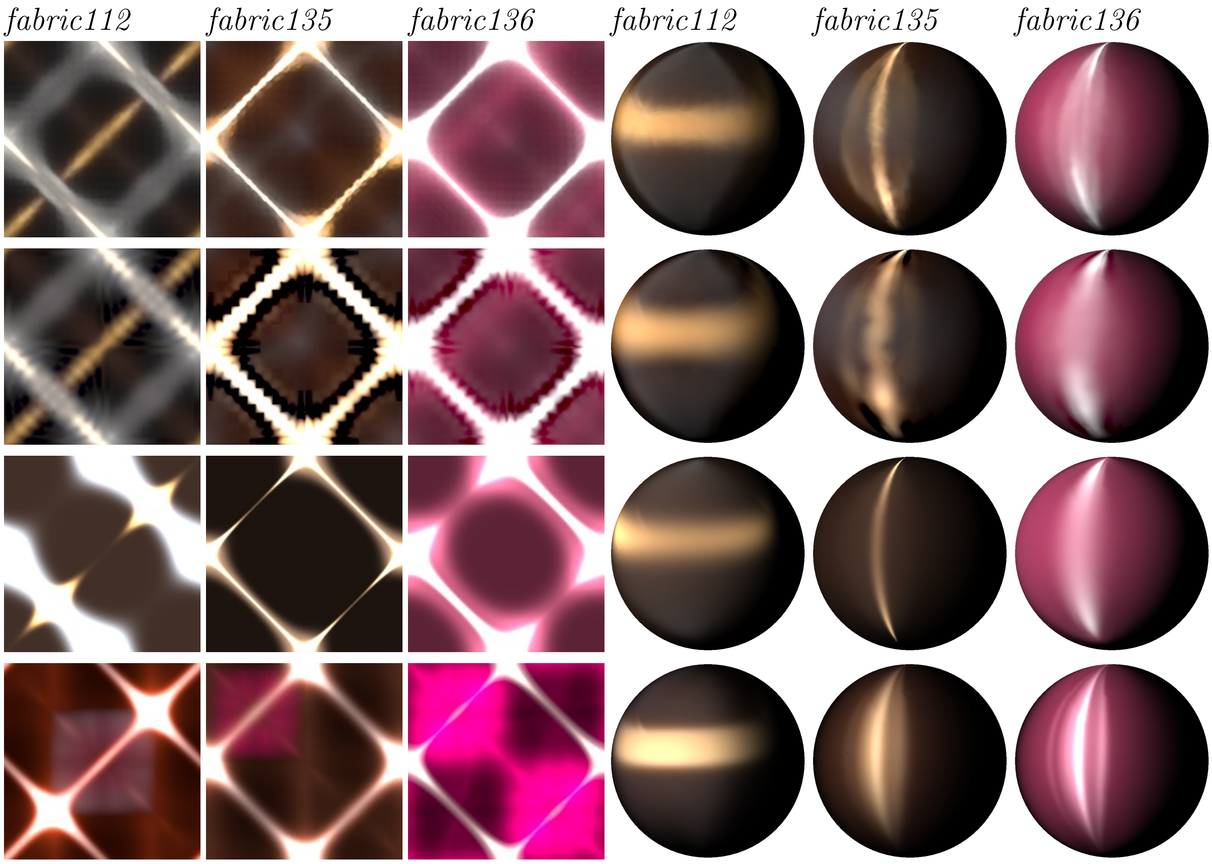

Figure 4.11: A BRDF subspace ($\theta_i = 74°$, $\theta_v = 74°$) and a rendering of the measured dense data (first row), the sparsely (81×81) measured data (second row), the BRDF model of Kurt et al. [KSKK10] (third row) and the BRDF model of Sadeghi et al. [SBDDJ13] (fourth row).

appearance of the measured fabrics is achieved by bending the fabrics along a cylinder and photographing them. Then, the measured datasets are rendered on a cylinder given the same conditions. Results shown in Figure 4.10 illustrate a good agreement between the photographs and the measured BRDFs. One has to take into account that the structure that is present in the photographs cannot by preserved by the BRDF.

Comparison with a Sparse Uniform Sampling – One might suggest that the sparse uniform sampling of the BRDF space would perform similarly; therefore, we take only samples from the centers of the measured images. In such a case, the hemisphere over the material is sampled by 81 illumination × 81 viewing directions resulting in $6,561$ samples distributed uniformly over the hemisphere. When missing samples are interpolated using the RBFs up to a density of 2°, we can observe considerable differences, especially in improper reproduction of high frequency features, e.g., width and continuity of the specular highlights, as it is clearly visible in the second row of Figure 4.11.

Fitting by BRDF Models – We test two recent BRDF models. First, we fit two lobes

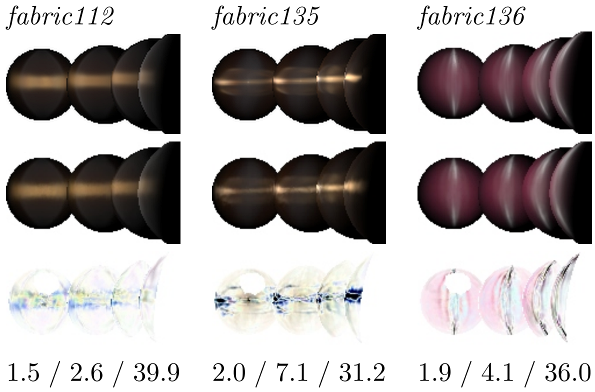

Figure 4.12: A comparison of the ground-truth data (first row) with renderings obtained using the dense BRDF measurements (second row). Below are the difference values in CIE ΔE / RMSE / PSNR [dB] and the difference images (third row) are scaled 10×.

of the Kurt et al. model [KSKK10] and results are shown in the third row of Figure 4.11. Second, we use Sadeghi et al. model [SBDDJ13]. The results of our manual fitting of thread shapes and parameters are shown in the fourth row of Figure 4.11. In our case, the last two materials contain tiny transparent nylon fibers instead of a cylindrical thread of fibers. Therefore, it is difficult to reproduce them reliably by the model. Although application of both models yields significant errors for the observed subspace of rather extreme elevations, the renderings on a sphere produce appearance similar to our measurements. These results underline the importance of the accurate and dense BRDF measurements for BRDF related research as fitting of the analytical models always sacrifice a certain portion of apparent details.

Comparison with the Ground-Truth Data – We compare the data in an applied situation using the 3D scene described in Section 5.5.3 and depicted in Figure 5.10. We measured exact values of 6, 195 pixels of the scene and we compare them to the renderings using the measured BRDFs. Results are shown in Figure 4.12. From the accompanying error values and difference images, we can conclude that the data are reasonably precise. The main source of error is a slight difference in a hue and a weak noise in the data.

4.4 Summary

This chapter presents the measurement setup and data processing and measurement approaches beneficial to methods of adaptive measurement described in the following chapters. First, the setup, its parameters and properties are summarized in Section 4.1. Our setup consists of a gonioreflectometer that realizes four DOFs, a LED light source and a camera with an interchangeable lens. Adaptation of exposure time and intensity of illumination allows us to achieve extensive dynamic range and a measurement can be finished within 18 hours.

Section 4.2 describes the registration method introduced in the paper [A.15] that focuses on the correct registration of multi-view images of planar surfaces. As the reference plane and the sample plane might be misaligned, our aim is to find this misalignment. When it is found and compensated from the measured dataset during the registration stage, better alignment of material features is achieved. Quality of the registration is verified by a reconstruction error of the compressed data. Consequently, the proposed approach allows more efficient application of compression approaches on multi-view data, i.e., we obtain sharper images using the same size of compressed parametric representation. The proposed method is robust, easy to implement and computationally efficient.

An innovative procedure for image-based measurement of the angularly-dense anisotropic BRDF that uses a gonioreflectometer introduced in [A.13] is summarized in Section 4.3. We sample the BRDF space uniformly while each sample represents an image of a circular material probe of the size comparable to the distance from the material to a light source and a camera. Due to the uniformity of a material and variations of illumination and viewing directions across its surface, we are able to collect many BRDF samples from each image.

Our approach allows for the acquisition of nearly three million samples from $8,505$ images. Therefore, it overcomes the main limitation of a gonioreflectometer based setup - the long measurement time, as the measurement is completed within twenty hours. We believe that due to the high accuracy of the acquisition process, our publicly available measurements represent unique anisotropic data useful as a reference for further BRDF-oriented research.

Chapter **5**

Adaptive Measurement Based on One-Dimensional Slices

In this chapter, we introduce a method of adaptive measurement of the BRDF of an arbitrary material in the high angular resolution. The method does not rely on a database of already measured materials and yet can provide precise results using a limited number of samples. The method is well scalable and can be used in every application where precision is important and might be helpful for acquisition of a BRDF database crucial for other methods. The chapter builds upon our paper [A.14] that presents our first efforts with adaptive measurement based on one-dimensional slices and extends its contributions.

Our method is based on our findings of the typical behavior of the BRDFs (Sec. 3.1). We assume that each two-dimensional BRDF subspace of given θ_i, θ_v and its important features can be captured by several diagonal and anti-diagonal cross-sections (so-called *slices*), as depicted in Figure 5.1. Sampling the BRDF subspace only along the slices saves a huge amount of samples. Another saving is done by selection of only certain subspaces (e.g., by limitation of values of θ_i, θ_v to multiples of 15°, see Fig. 5.2-left). Then, elevation-dependent behavior is captured by another type of slices (see Fig. 5.2-right). The slices form a sparse 4D structure in the BRDF space. Its ability to capture important features of the BRDF depends on its density: the more dense structure, the more precise reconstruction can be achieved; however, more samples are needed. Although all the values along the structure itself present substantial reduction of BRDF samples when compared to the dense regular sampling, we need to reduce the number of measured values even more to make the method practical. That is done by the sparse but adaptive sampling along individual slices, i.e., one-dimensional signals.

To our best knowledge, adaptive measurement of an unknown one-dimensional signal is still an open problem. Therefore, after a detailed description of the slices in Section 5.1, we focus ourselves on this problem in Section 5.2 where we propose an enhanced heuristic for the problem which was first introduced in [A.14]. Section 5.3 encompasses an interpolation method for reconstruction of the BRDF from the slices in a 2D BRDF subspace proposed in [A.14] and Section D.1 describes its extension for reconstruction from the

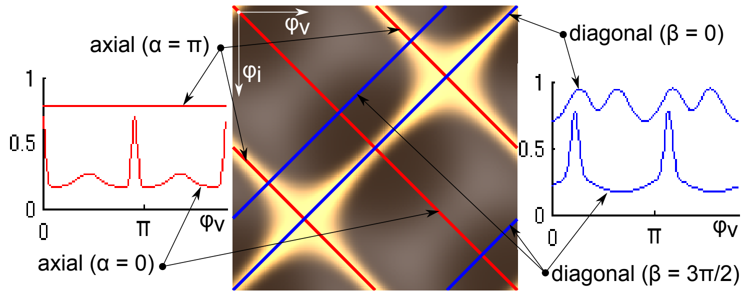

Figure 5.1: A schema of the axial (red, perpendicular to the anisotropic highlights) and the diagonal (blue, perpendicular to the specular highlights) slices placed in the 2D subspace ($\theta_i = 60°$, $\theta_v = 60°$) and their function values. Note that the subspace is periodical.

4D structure. Then, a study on the optimal placement of the slices into the sparse 4D structure is presented in Section 5.4. While Section 5.5 compares results of the proposed method with other methods, Section 5.6 discusses advantages and limitations of the proposed method and confirms validity of the study on the optimal placement. Instructions for implementation of the proposed reconstruction method on a graphics hardware are provided in Section D.4. Next, Section 5.7 encompasses the new reconstruction method presented in the paper [A.11]. The new method enables us to remove the main disadvantage of the reconstruction method described in Section 5.3 – inaccurate reconstruction of the features not perpendicular to the slices. Although the new method is very promising, it brings new challenges and it needs substantial improvement to replace the former method. Finally, Section 5.8 summarizes the chapter.

5.1 Introduction to One-Dimensional BRDF Slices

As is explained in Section 1.1, the BRDF is the four-dimensional vector-valued function of elevation angles θ_i, θ_v and azimuthal angles φ_i, φ_v. To enable effective sampling of the entire BRDF, we propose four types of one-dimensional so-called *slices*. Two of them, *axial* and *diagonal* slices, were introduced in [Fil12], where only one slice of each type per the 2D BRDF subspace is used. We propose to use up to dozens of slices per subspace to capture the most of subtle details of the BRDF. Moreover, we suggest the extension of the concept to additional *horizontal* and *vertical* slices.

Axial and Diagonal Slices – take place in the 2D BRDF subspaces defined by fixed θ_i and θ_v. These slices are designed to optimize capture of the specular reflections and the anisotropic reflections. The axial slices s_A (red in all figures) are perpendicular to the anisotropic highlights and the diagonal slices s_D (blue in all figures) are perpendicular to the specular highlights (see Fig. 5.1):

$$s_{A,\theta_i\theta_v\alpha}(\varphi_v) = f_r(\theta_i, \theta_v, \varphi_i = \varphi_v - \alpha, \varphi_v), \tag{5.1}$$
$$s_{D,\theta_i\theta_v\beta}(\varphi_v) = f_r(\theta_i, \theta_v, \varphi_i = \beta - \varphi_v, \varphi_v), \tag{5.2}$$

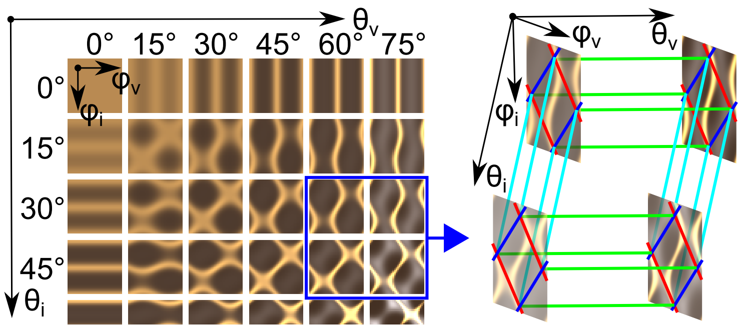

Figure 5.2: A preview of the BRDF with the highlighted axial (red), diagonal (blue), horizontal (green) and vertical (cyan) slices. Note that the BRDF is a four-dimensional continuous function and that its axes are orthogonal.

where α or β determine a position of the slice in the 2D subspace chosen by elevation angles θ_i, θ_v. Each subspace is typically measured using several axial and several diagonal slices to accurately capture all reflections and their shape (see Fig. 5.1).

Horizontal and Vertical Slices – are the 1D subspaces of the BRDF. They are designed to capture a change in reflectance values when the viewing elevation angle θ_v (horizontal slice, green) or the illumination elevation angle θ_i (vertical slice, cyan) is changed while other parameters are fixed as show in Figure 5.2-right:

$$s_{H,\theta_i\varphi_i\varphi_v}(\theta_v) = f_r(\theta_i, \theta_v, \varphi_i, \varphi_v),$$
$$s_{V,\theta_v\varphi_i\varphi_v}(\theta_i) = f_r(\theta_i, \theta_v, \varphi_i, \varphi_v).$$

Each horizontal or vertical slice passes through an intersection of the axial and diagonal slices at all sampled elevations (e.g., 0^o, 15^o, 30^o, 45^o, 60^o, 75^o). Note that due to Helmholtz reciprocity, values of the corresponding horizontal and vertical slices are equal, i.e., $s_{H,\theta_i\varphi_i\varphi_v}(\theta_v) = s_{V,\theta_i\varphi_v\varphi_i}(\theta_v)$. So only, e.g., the horizontal slices need to be measured.

All the four types of slices represent the sparse 4D structure in the 4D BRDF space that effectively captures the main visual features of the measured BRDF (see Fig. 5.2-right).

5.2 Adaptive Sampling Along the Slices

Each slice can be interpreted as an unknown one-dimensional signal that we need to measure and reconstruct. In case of the axial or diagonal slice, the signal is periodical with the period of 360°. Any slice can be sampled uniformly with a defined step (e.g., 1°) or adaptively decreasing the number of samples on the one hand and increasing reconstruction accuracy in areas with the high variance of the signal on the other. As the behavior of the signal is unknown, the adaptive algorithm can work with already measured samples only; adding new samples in areas where it can improve accuracy of the reconstructed signal.

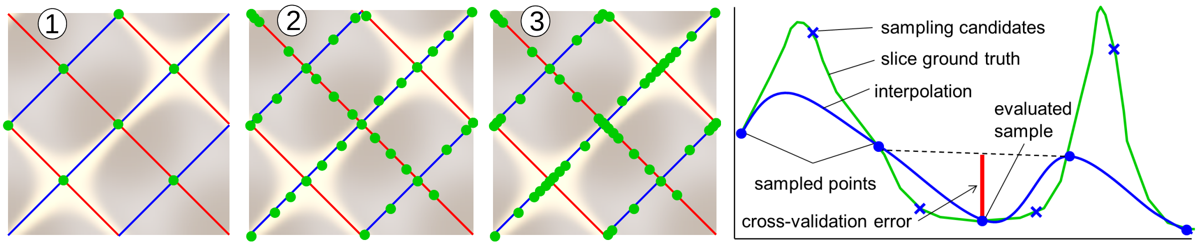

Figure 5.3: A description of individual steps of the adaptive sampling algorithm (left), and a description of the cross-validation procedure (right).

When the samples are taken in proper directions and their count is sufficient, every value of the slice can be interpolated precisely enough using, e.g., piece-wise cubic splines.

We propose heuristic Algorithm 5.1 that enables a very efficient adaptive sampling of the slices using a given count of samples n. The algorithm is based on a simple assumption. If the value of a sample can be predicted well by the neighboring samples, the neighborhood of the sample can probably be predicted well too and there is no need to place new samples there. On the other hand, if the value of any sample cannot be predicted well by the neighboring samples, there is possibility that neighborhood of the sample cannot be predicted well even by the sample itself together with the neighbors. Therefore, we should place new samples there.

The algorithm consists of three steps as shown in Figure 5.3. In the first step, the algorithm samples the signal only in the directions of the intersections of the slices to collect at least some information on the signals. Then, several iterations of adaptive measurement are performed relying on Adaptive Sampling Iteration (ASI) Algorithm 5.2, which adaptively selects and measures n_i sampling candidates.

The count of samples n_i in each iteration depends on parameters k, p_1, p_2. The parameter k defines the ratio of the count of samples measured in the second step of the Algorithm 5.1 versus the count of samples taken in the third step. We have found experimentally that reasonable choice is $k = 0.9$ expressing the ratio $90\% : 10\%$. The parameters p_1 and p_2 define how many iterations are performed in each step and our default choice is $p_1 = p_2 = 5$. Using these parameters, the algorithm forces to divide even those intervals on the slices where the error is lower in the second step as there could be some hidden variation of the signal. In the third step, only the intervals where the high cross-validation error prevails are divided. Although the algorithm is not very sensitive to the parameters, one should not decrease the total number of iterations $p_1 + p_2$ too much as intervals on the slices could not be divided sufficiently then. On the contrary, using too many iterations (up to addition of only one new sample during an iteration) is not recommended as the algorithm would be too focused on the specular highlights and might overlook variations of the signal in unexplored intervals.

Algorithm 5.2 demands the table T of already measured directions and their values and the count of samples n_i that will be identified and measured by the algorithm. The output of the algorithm is the appended table T. The algorithm itself consists of three steps. First, the leave-one-out cross-validation error is evaluated in each already measured sample. The

Algorithm 5.1 Adaptive Sampling Along the Slices

Input: n, k, p_1, p_2.

Output: T - a table of measured directions and their values.

1: Measure values of all n_0 samples at all intersections of the BRDF slices creating T.
2: p_1-times append T performing $T = \text{ASI}(T, [n - n_0] \cdot k/p_1)$.
3: p_2-times append T performing $T = \text{ASI}(T, [n - n_0] \cdot [1\text{-}k]/p_2)$.

Algorithm 5.2 Adaptive Sampling Iteration (ASI), $T = \text{ASI}(T, n_i)$

1: Evaluate the leave-one-out cross-validation error for each direction in T (for each slice independently), i.e., interpolate the value of a sample in T based on neighboring samples and evaluate the error by comparison with the true value of the sample.
2: Make an ordered list of the potential directions (see crosses in Fig. 5.3) in descending value of their weight. The potential direction is the one in the middle of each two neighboring already measured directions in T. Its weight is the maximal error value of the neighboring directions.
3: Measure values of the first n_i directions in the list and append them to T.

error is evaluated independently for each slice, i.e., the sample at the intersection of the slices has several error values. The evaluation for each sample is done by exclusion of the evaluated sample from the dataset, the linear interpolation of a value in the location of the evaluated sample from the neighboring samples, and computation of a distance of the interpolated value and the actual value of the evaluated sample (see Fig. 5.3-right). As a distance measure, we use the maximum difference over all color channels. Then, in the second step of Algorithm 5.2, a list of all directions where new measurements could be performed is prepared. Each direction has assigned a weight equal to the maximum of cross-validation errors of its already measured neighbors (see Fig. 5.3-right). The list is then sorted in descending manner, values of n_i first directions from the list are measured, and the table T is appended by the newly measured directions and their values.

5.3 Reconstruction from the Sparse Structure

To reconstruct an arbitrary point in the four-dimensional BRDF space from the slices, we adapted a two-dimensional swept surface technique that uses two cross-section and two profile curves and we constrain the resulting minimal value. We tested also polynomial and spline interpolations, but they require predefined rank and introduce disturbing artifacts.

Therefore, in the 4D BRDF space, the demanded data are restored by a multi-linear interpolation of values of the slices. This processing is done separately for each color channel. For clarity, the following equations describe only a bilinear interpolation in the 2D subspace in a form that enables illustrative visualization (see Fig. 5.4). In Section D.1, we rewrite the equations to make transition to four dimensions possible and finally, we provide there equations for reconstruction of the desired value in the 4D space.

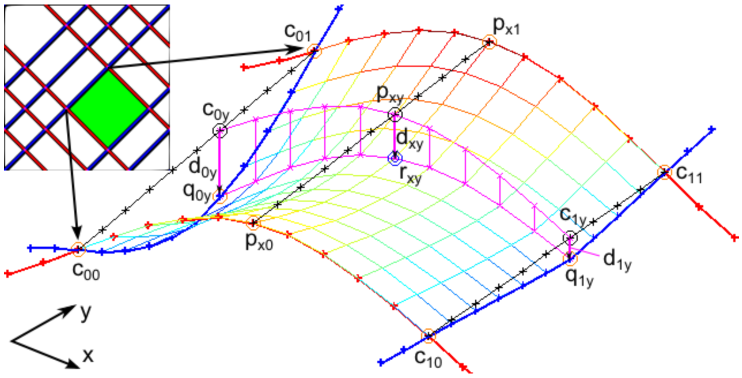

Figure 5.4: A scheme of the 2D interpolation estimating interior points of a general rectangle from values on a border and at its corners.

For purposes of clarification, edges of a reconstructed rectangle, which are formed by the axial and diagonal slices, are normalized to unitary length and the axial slices are parallel to axis x and the diagonal slices are parallel to axis y (see Fig. 5.4). The interpolation requires a knowledge of values of eight points to interpolate the desired value r_{xy} at (x, y) inside the square. The first four of them are values at the corners of the square (the intersections of the slices) $c_{\bar{x}\bar{y}}$, where $\bar{x}, \bar{y} \in \{0, 1\}$. The next two values p_{x0}, p_{x1} are on the axial slices and the last two values q_{0y}, q_{1y} are on the diagonal slices.

First, the values p_{x0}, p_{x1} are linearly interpolated along axis y yielding the value:

$$p_{xy} = (1-y) \cdot p_{x0} + y \cdot p_{x1} \ .$$

To compute the value r_{xy}, the value p_{xy} has to be compensated for a height difference introduced by the linear interpolation and values of the diagonal slices. Therefore, the values c_{0y} and c_{1y} are also linearly interpolated along axis y from c_{00}, c_{01} and c_{10}, c_{11}:

$$\begin{aligned} c_{0y} &= (1-y) \cdot c_{00} + y \cdot c_{01} \ , \\ c_{1y} &= (1-y) \cdot c_{10} + y \cdot c_{11} \ . \end{aligned}$$

The differences d_{0y} and d_{1y} between the linear interpolations of the corner values and values of the x-aligned slices are computed as $d_{0y} = q_{0y} - c_{0y}$ and $d_{1y} = q_{1y} - c_{1y}$. Finally, the height difference in the point (x, y) is obtained by the linear interpolation of the differences in the axis x:

$$d_{xy} = (1-x) \cdot d_{0y} + x \cdot d_{1y} \ .$$

The final value r_{xy} is a sum of the value p_{xy} and the difference d_{xy} and its minimal value is constrained by the minimal value of p_{x0}, p_{x1}, q_{0y}, q_{1y} ensuring non-negativity of the reconstructed function:

$$\begin{aligned} r_{xy} &= max(p_{xy} + d_{xy}, min_{xy}) \ , \\ min_{xy} &= min(p_{x0}, p_{x1}, q_{0y}, q_{1y}) \ . \end{aligned}$$

5.4 A Study on the Optimal Placement of the Slices

Though values along individual slices are measured adaptively, a position of the slices in the BRDF space must be known in advance. To investigate which placement of the slices is optimal with respect to a given count of samples, we perform the following study. Due to the insufficient number of available anisotropic BRDF measurements of high angular density and high accuracy, we perform the study on ten materials represented by an analytical BRDF model (see Sec. D.2).

5.4.1 Placement Theory

The axial and diagonal slices are placed into 2D subspaces selected by θ_i and θ_v. The most straightforward approach is to deploy the slices regularly into the subspaces using $\alpha, \beta \in \{k \cdot a_s; k \in \{0, 1, \ldots, a_n - 1\}, a_n = 2\pi/a_s\}$, where a_n is a count of the slices of one type in one subspace and a_s is the *azimuthal step* between two consecutive slices. Only the a_s parameter has to be chosen. Let us denote that the position of the axial and diagonal slices might be chosen arbitrarily, but for simplicity, we limited ourselves to the regular distribution of both types of slices by one common step parameter.

Similarly, 2D subspaces are selected by one parameter called *elevation step* e_s as $\theta_i, \theta_v \in \{k \cdot e_s, \theta_{max}; k \in \{0, 1, \ldots, e_n - 2\}, e_n = \lceil \theta_{max}/e_s \rceil + 1\}$, where e_n is a count of elevations involved and θ_{max} is the maximal elevation. Measured subspaces are those resulting from all combinations of θ_i, θ_v. Finally, the horizontal and vertical slices are placed through all intersections of the axial and diagonal slices (see Fig. 5.2).

5.4.2 An Analysis of the Optimal Placement

We perform an experiment to determine the optimal positions of the slices based on a user-defined count of samples n. We use all the ten BRDFs of the ten materials (Sec. D.2) and compute reconstructions of these BRDFs for various combinations of the *azimuthal step* a_s, the *elevation step* e_s and the count of samples n. To make the precalculation computationally feasible, we restrict a resolution of the reconstructed BRDFs to 2^o and we use the maximal elevation angle $\theta_{max} = 80^o$ due to the unstable fitting of the BRDF by the analytical model for high elevation angles. Values of the *azimuthal step* parameter are $a_s \in \{12°, 20°, 36°, 60°, 180°\}$ and values of the *elevation step* parameter are $e_s \in \{6°, 8°, 10°, 12°, 14°, 16°, 20°, 28°\}$. In total, we perform 20,646 simulated measurements and reconstructions of BRDFs, which consume over 20 days of computational time using four cores of Intel Xeon E5-2643 3.3GHz.

To evaluate quality of the reconstructed BRDFs for the tested values of the parameters, we compute the MRE between the reference BRDF $f_r(\theta_i, \theta_v, \varphi_i, \varphi_v)$ and its reconstruction $f'_r(\theta_i, \theta_v, \varphi_i, \varphi_v)$. We compute graphs of the MRE as a function of the number of adaptive samples n. As a result, there are $|a_s| \times |e_s| \times |m| = 5 \times 8 \times 10 = 400$ graphs $err_{m,a_s,e_s}(n)$, where m stands for one of 10 materials. Each graph captures the reconstruction error for a large range of samples n. In Figure 5.5 are plotted some of those graphs averaged

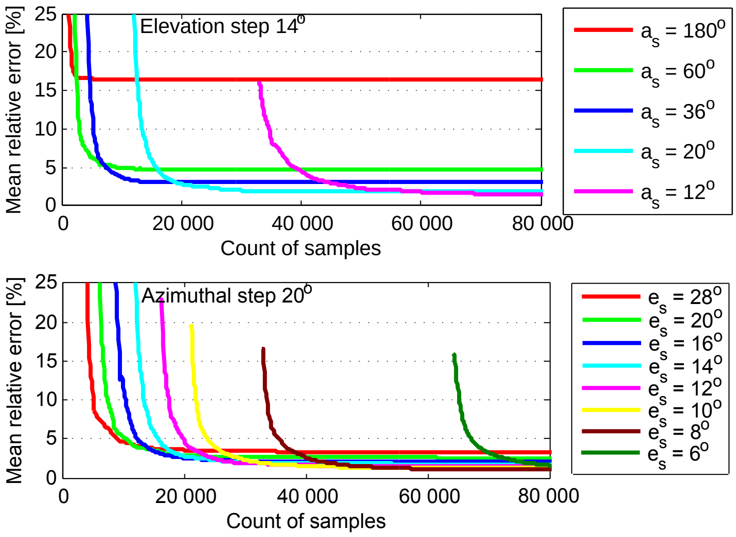

(a)

(b)

Figure 5.5: The MRE of the reconstructed BRDF as a function of the count of samples averaged across all materials for five values of the *azimuthal step* with the fixed value of $e_s = 14°$ (a) and for eight values of the *elevation step* with the fixed value of $a_s = 20°$ (b).

across all the materials. The first group of graphs (a) shows progress of the error for the fixed parameter $e_s = 14^o$ and various values of the a_s parameter; while, the second group (b) shows the progress for the fixed parameter $a_s = 20^o$ and various values of the e_s parameter. Note that the fast convergence of individual graphs confirms efficiency of the adaptive sampling algorithm.

From the graphs, it could be possible to conclusively select the best combination of the a_s, e_s parameters for a selected count of samples n relative to a given material m. Unfortunately, these selections are not unique across different materials as they exhibit individual behavior with respect to changes in the azimuthal and elevation angles. Therefore, we estimate the optimal values $\widehat{a_s}, \widehat{e_s}$ of the parameters for a given count of samples n so as

Table 5.1: The optimal values of the a_s, e_s parameters depending on the demanded count of samples.

$n \leq$	e_s	a_s	$n \leq$	e_s	a_s
667	28°	180°	17 096	12°	36°
932	20°	180°	20 969	10°	36°
1 034	16°	180°	22 291	8°	36°
1 060	28°	60°	33 879	12°	20°
2 272	20°	60°	38 735	10°	20°
3 230	16°	60°	79 469	8°	20°
4 928	14°	60°	184 655	6°	20°
5 645	16°	36°	∞	6°	12°
9 660	14°	36°			

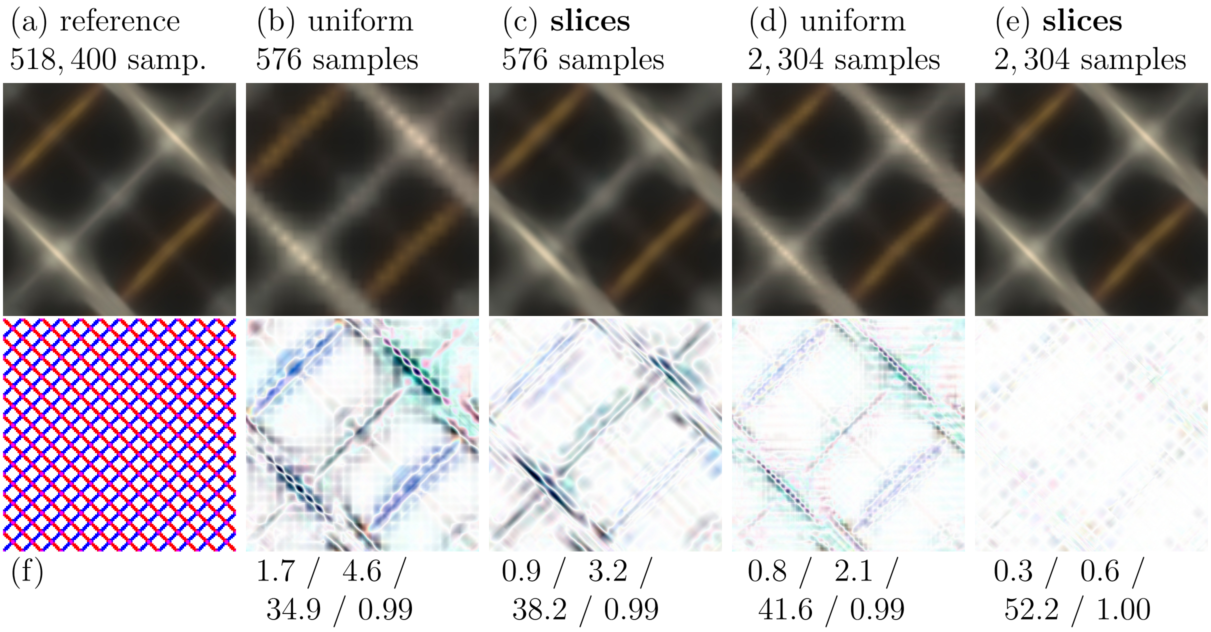

Figure 5.6: The reference BRDF subspace (a) compared with the barycentric interpolations of the same size from two densities of the uniform sampling (b,d), and the proposed method using the same counts of samples (c,e). Below are the difference images scaled 10× and the difference values in CIE ΔE / RMSE / PSNR [dB] / SSIM. Bottom-left: The used uniform placement of the measured slices (f).

the sum of the errors across all the materials relative to the achievable error is minimized:

$$(\widehat{a_s}, \widehat{e_s})(n) = \underset{(a'_s \in a_s, e'_s \in e_s)}{\arg\min} \sum_{m=1}^{|m|} \frac{err_{m,a'_s,e'_s}(n)}{\underset{(a''_s \in a_s, e''_s \in e_s)}{\min} err_{m,a''_s,e''_s}(n)}. \tag{5.3}$$

The resulting optimal values of the a_s, e_s parameters are summarized in Table 5.1; when a new material is measured, one should select appropriate values of the parameters for the demanded count of samples.

5.5 Results

This section presents results of the proposed method. First, due to lack of reliable data, we investigate performance of the axial and diagonal slices only in two-dimensional BRDF subspaces. Then, we use synthetic data generated by a BRDF model to test the method in the four-dimensional BRDF space. To evaluate performance of the method on real data we use a 3D scene (Fig. 5.10) and measure all the data needed to visualize the scene by the gonioreflectometer (Sec. 4.1). Finally, the contribution of the horizontal and vertical slices included on the top of the axial and diagonal slices is evaluated.

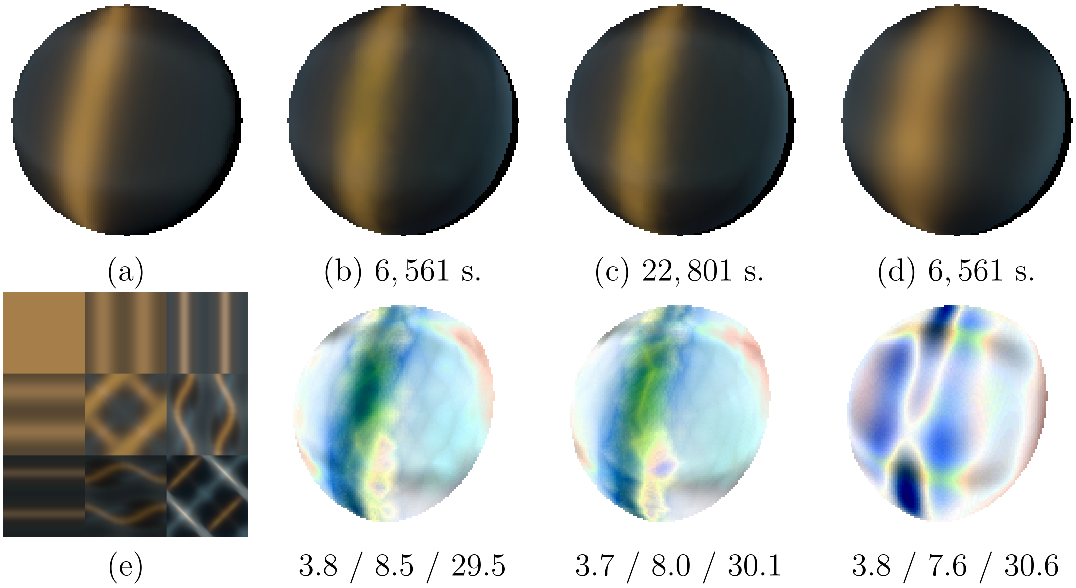

Figure 5.7: The per-pixel ground-truth measurements on a sphere (a) compared with the barycentrically interpolated uniform measurements (b,c) and with the proposed method (d). Below are the difference values in CIE ΔE / RMSE / PSNR [dB] as well as the 10× scaled difference images. Bottom-left: The subspaces measured by the method (e).

5.5.1 Performance of the Axial and Diagonal Slices

Performance of the proposed method using only the axial and diagonal slices is compared here with the conventional uniform sampling; both methods use the same count of samples. We perform this experiment first due to unavailability of the accurate dense reference data owing to their extremely time-demanding measurement. Therefore, we use only a single material described in Section A.2 for our experiment. However, the selected material has one of the strongest anisotropic behavior we have seen so far.

First, we focus ourselves to test behavior of the material in one subspace only. We selected the subspace $\theta_i = \theta_v = 75°$, which was very precisely measured using the UTIA gonioreflectometer (Sec. 4.1) and its azimuthal resolution is 0.5° (see Sec. A.2). Figure 5.6 shows the densely measured reference BRDF subspace (a), the uniform sampling by means of 24×24 samples interpolated to the same azimuthal resolution as the reference using the barycentric interpolation (b) and 48×48 samples interpolated in the same manner (d). Figure 5.6 (c) and (e) shows performance of the proposed interpolation method in the suggested representation using 12 axial and 12 diagonal slices (f). The 10× scaled difference images (second row of Fig. 5.6) show that the proposed method is able to achieve a significantly better reconstruction of the original data using the same number of samples.

Furthermore, Figure 5.7 shows a comparison of the ground-truth measurements on a sphere (a), with the rendering using the barycentrically interpolated 81×81 (b) and 151×151 (c) uniform samples. The result of the proposed method using 6, 561 adaptive

Table 5.2: A decrease of the MRE due to usage of the proposed method instead of the barycentric interpolation (fist row) or the RBF interpolation (second row). Average values over 30 sampling schemes.

	Brush. alum	Purple satin	Red velvet	Yellow satin	fabric 002	fabric 041	fabric 112	fabric 135	fabric 139	wood 01	mean
prop. vs. bar.	5.68×	5.88×	26.40×	2.56×	6.68×	6.80×	5.06×	2.04×	4.53×	9.24×	7.49×
prop. vs. RBF	6.97×	2.21×	10.58×	1.68×	2.79×	4.48×	1.45×	1.11×	2.28×	5.07×	3.86×

samples is shown in (d). We measured nine subspaces at elevations $0°, 30°, 75°$ and their combinations and interpolated data at missing elevations using the four-dimensional Krig interpolation of the spherical angles $(\theta_i, \varphi_i, \theta_v, \varphi_v)$ represented in a $0 \approx 2\pi$ continuity-preserving parameterization [HFM10].

5.5.2 Simulated Measurement Experiment

To evaluate performance of the proposed method using all types of slices, the reliable BRDF data are needed. As a sufficient number of the ground-truth anisotropic BRDF measurements of the high angular density and high accuracy is not available so far, we perform a simulated measurement experiment using ten materials represented by an analytical BRDF model (see Sec. D.2). Thus, we can easily and quickly obtain a BRDF value of any direction and results of the experiment are not influenced by errors caused by a measurement process.

We evaluate performance of the proposed method in comparison with the uniformly distributed samples, which are taken at directions according to one of the thirty sampling schemes we designed (see Sec. D.3). These schemes produce in total from $n = 435$ to $n = 354,061$ reciprocal samples. Values of the samples are interpolated using the barycentric [Cox69] or the RBF [PTVF92] interpolation. Notice that the second method is global while the first one is local and is therefore suitable for fast rendering on a GPU. Both methods compute results separately in each color channel. We interpolate the BRDF to

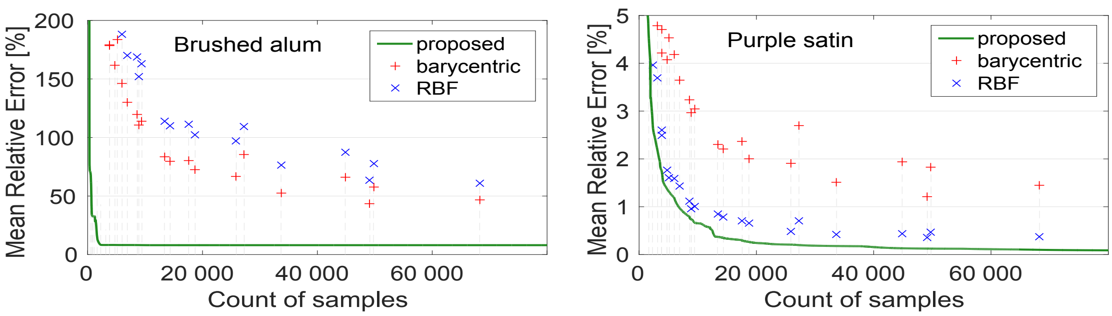

Figure 5.8: The MRE as a function of the count of samples for two materials.

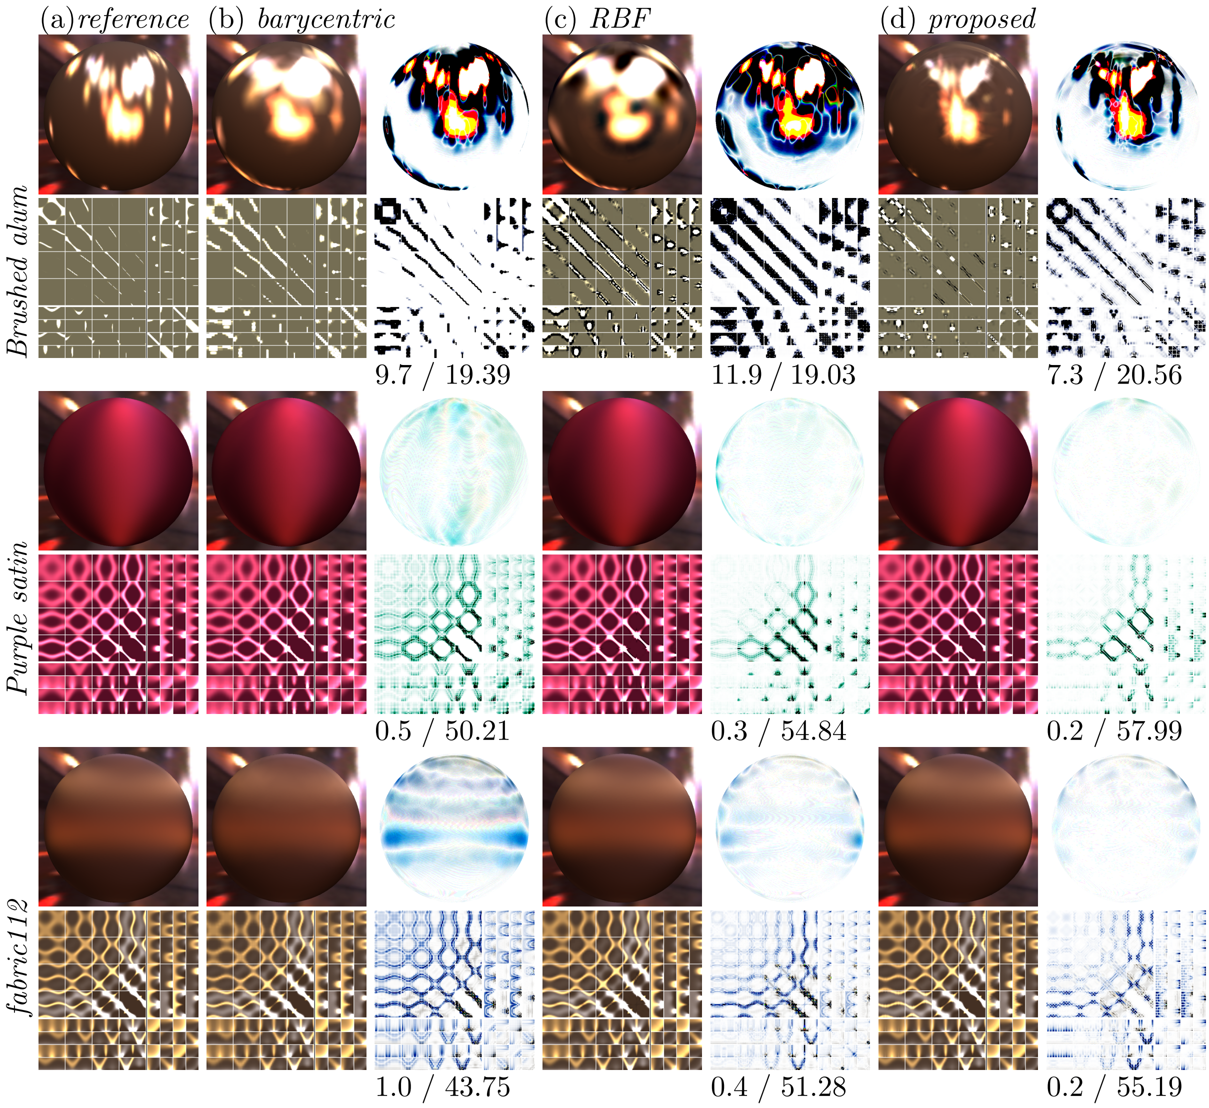

Figure 5.9: The BRDF rendering on a sphere in the *grace* environment. The reference BRDF (a) is compared with its reconstruction from 8,911 samples using (b) the barycentric interpolation, (c) the RBF interpolation, and (d) the proposed method. Difference images are scaled 10× and below are the difference values in CIE ΔE / PSNR [dB].

a four-dimensional array using a uniform step of 2^o and the maximal elevation is 80^o, i.e., dimensions of the array are $|\theta_i| \times |\theta_v| \times |\varphi_i| \times |\varphi_v| = 41 \times 41 \times 180 \times 180$.

To evaluate quality of the reconstructed BRDF we compute the MRE between the reference BRDF $f_r(\theta_i, \theta_v, \varphi_i, \varphi_v)$ and its reconstruction $f'_r(\theta_i, \theta_v, \varphi_i, \varphi_v)$ in $N = 3 \times 41 \times 41 \times 180 \times 180 = 163,393,200$ data points. The average decrease of the MRE due to usage of the proposed method instead of the barycentric interpolation or the RBF interpolation across all counts of samples of the thirty sampling schemes is shown in Table 5.2. Note that

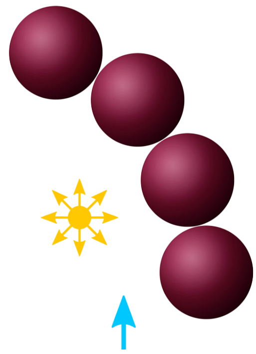

Figure 5.10: A schema of a 3D scene of Figure 5.11.

the error values are evaluated only for the discrete number of samples as the barycentric and RBF interpolations operate on the 30 predefined sampling schemes. The average MRE (over all the materials and schemes) of the barycentric interpolation is almost 7.5-times as high as for the BRDF slices. The RBF achieves better performance, but its average MRE is still almost 3.9-times higher than the MRE of the BRDF slices.

Figure 5.8 shows the progression of the reconstruction error as a function of the number of samples for all tested methods with two materials. The convergence of the proposed algorithm to low MRE values as a function of the count of samples on the tested dataset is very fast (see green line in Fig. 5.8). While the RBF method performs well for lower numbers of samples, the proposed method has a superior performance relative to higher numbers of samples. When the count of samples is over 5,000, our method achieves high quality results that are significantly better than those achieved using uniform interpolation methods. In Figure 5.9, we show a comparison of all three methods with the reference rendering for three materials and 8,911 reciprocal samples. We use the *grace* environment represented by means of 256 lights. Achieved results show that our method provides the best reproduction of the specular reflections and the anisotropic highlights.

5.5.3 Practical Measurement Experiment

The previous experiment was performed using the reference BRDFs represented by an analytical model. However, such an approach sacrifices some visual features of the original reflectance behavior that cannot be reliably represented by the model. Therefore, we perform practical BRDF measurement experiments collecting 8,911 and 18,721 samples. First, we record these samples uniformly (sampling schemes 14 and 19, see Sec. D.3) and interpolate them using the barycentric and RBF interpolations. Then, we record the same count of samples adaptively using the proposed method. Altogether, we measure four datasets for each of the three materials (*fabric112*, *fabric135*, *fabric136*).

All measurements are performed using the UTIA gonioreflectometer (Sec. 4.1) that allows for the placement of an almost arbitrary combination of illumination and viewing

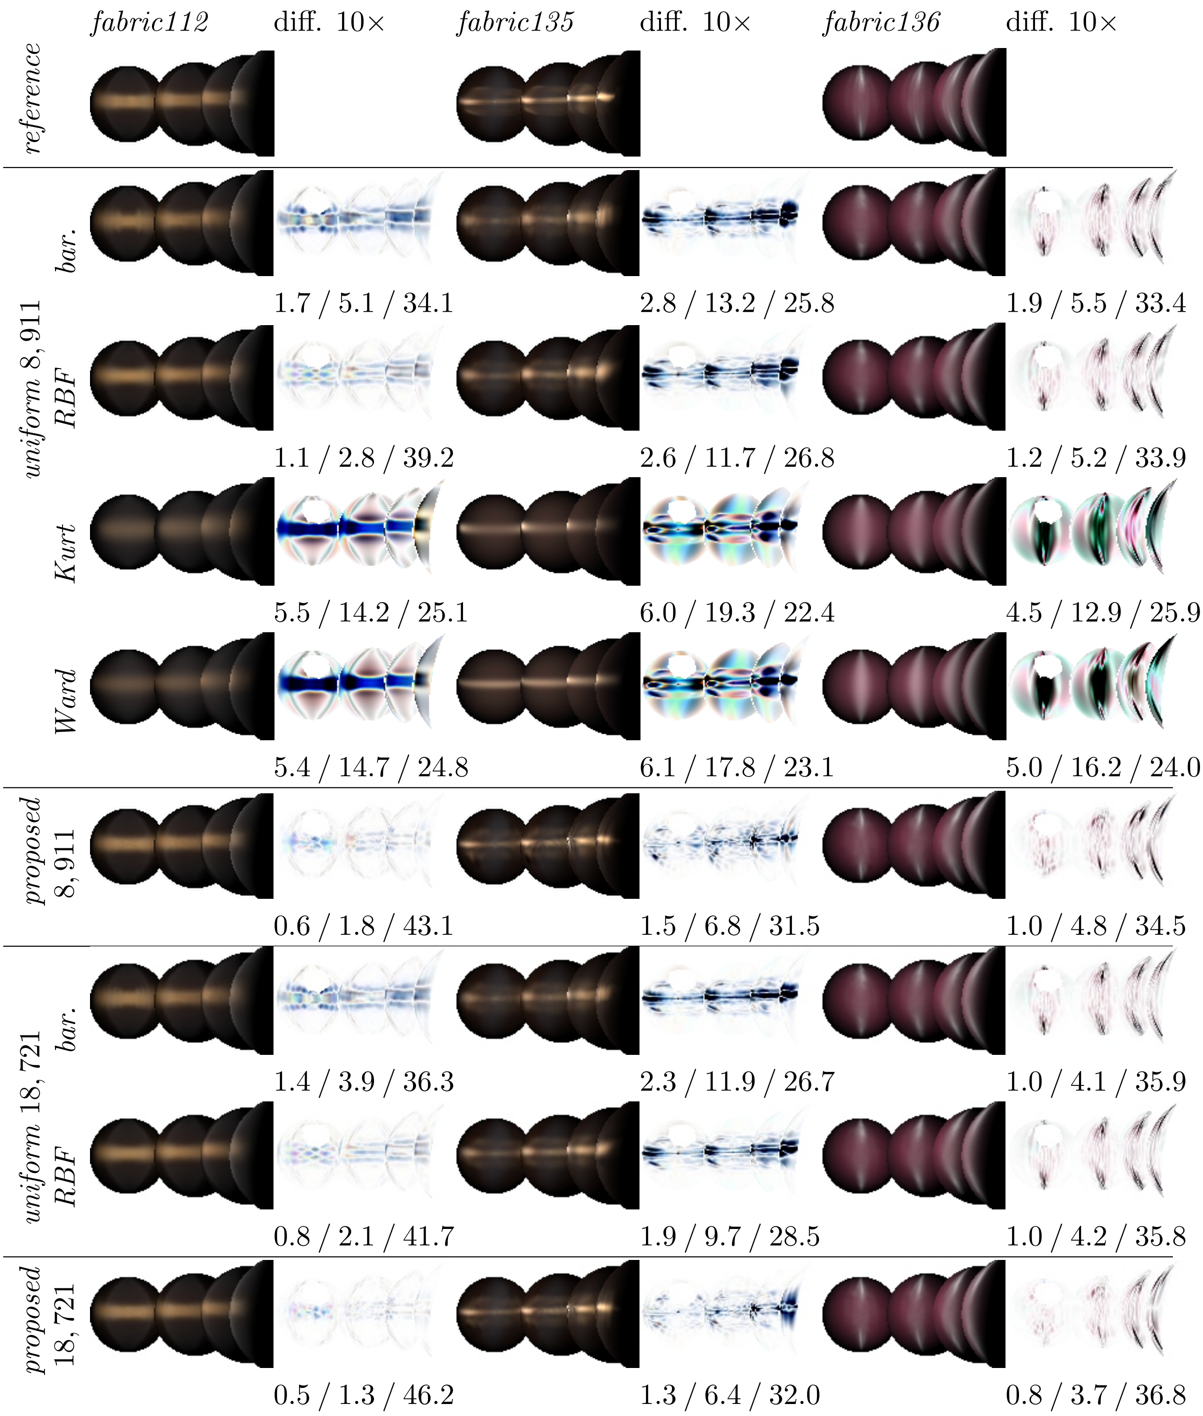

Figure 5.11: All reachable directions in the virtual scene (see Fig. 5.10) measured and compared with the barycentric and RBF interpolations of unifomly measured data, fits of two analytical models ([KSKK10, War92]), and the proposed adaptive measurement. $8,911$ and $18,721$ samples are used, the difference values are in CIE ΔE / RMSE / PSNR [dB].

Table 5.3: The MRE [%] depending on types of slices: (a,c) all types of slices are used (horizontal, vertical, axial, and diagonal), and (b,d) only the axial and diagonal slices are used. The MREs are for 8,911 samples (a,b) and for 18,721 samples (c,d).

	Brush. alum	Purple satin	Red velvet	Yellow satin	fabric 002	fabric 041	fabric 112	fabric 135	fabric 139	wood 01	mean
(a) all slices	24.0	0.8	0.1	10.2	0.5	0.2	1.3	3.9	0.5	0.2	4.2
(b) axial & diag.	68.5	1.6	2.1	10.7	1.8	1.1	2.8	3.9	0.9	1.1	9.5
(c) all slices	19.6	0.5	0.1	5.6	0.3	0.1	0.9	2.5	0.3	0.1	3.0
(d) axial & diag.	31.7	0.9	1.3	6.8	1.2	0.8	1.9	2.6	0.6	0.7	4.9

directions with high angular accuracy. Acquisition time using the adaptive method is about 10% longer in comparison to the uniform sampling of the same count of samples due to the data processing overhead as required by the adaptive algorithm. It takes approximately 25 hours for 8,911 samples and 51 hours for 18,721 samples.

We compare the data in an applied situation using a 3D scene that comprises four spheres illuminated by a single point-light source (see Fig. 5.10). Therefore, many combinations of illumination and viewing directions are covered providing a comprehensive preview of properties of the materials. The rendered images are divided into a sparse raster with only 6,195 occupied pixels representing the directions that are reachable by the gonioreflectometer; therefore, only this count of BRDF values is measured for three different materials. The entire scene is then rendered using those pixels, which we call *control samples*, and we use them as our reference. Note that pixels representing directions unreachable by the gonioreflectometer due to occlusion of view of the camera by the light source are not included into the *control samples* (resulting into white spot in the difference images in Fig. 5.11 on the most-left sphere). Their value for purposes of visualization is interpolated from the regular measurements.

Then, we render the same scene using the values obtained by the three compared methods and evaluate their results at all *control samples*. Figure 5.11 shows the reference scene side-by-side its reconstruction using all tested methods as well as results of two tested analytical anisotropic models of Kurt et al. [KSKK10] and Ward [War92]. From the accompanying error values, we conclude that our method achieves by far the best performance on the real BRDF data.

5.5.4 Contribution of the Horizontal and Vertical Slices

Last, we evaluate the contribution of the horizontal and vertical slices included on the top of the axial and diagonal slices. Table 5.3 shows the performance gain, which demonstrates that usage of the additional slices is important especially when considering the lower count of samples where with, e.g., 8,911 samples we achieve twice the lower MRE averaged across materials. Note that when only the axial and diagonal slices are used, they are sampled more densely as the same count of samples is used (and they do not need to cover the horizontal and vertical slices in this case). The general recommendation based on the results is to sample along the horizontal and vertical slices whenever a measurement setup allows that.

5.6 Discussion

This section discusses advantages and limitations of the proposed method and quotes timings related to the used interpolation methods. Furthermore, we discuss reliability of the proposed approach that finds the optimal placement of the slices.

5.6.1 Advantages of the Method

In contrast to competitive methods, the main advantage of the proposed adaptive sampling is its faster decrease of reconstruction errors in correlation with increasing number of samples (especially for their lower counts, see Fig. 5.5). Thus, on average we achieve almost four-times lower relative errors given the same count of samples and between two- and five-times less samples necessary to reach the same reconstruction errors. Additionally, the proposed adaptive sampling method operates along the one-dimensional slices in the BRDF space that allows data acquisition by a continuous movement of a light source and a sensor. It is especially beneficial for accurate gonioreflectometer-based measuring devices.

5.6.2 Limitations of the Method

What is common to all tested methods is their decreasing improvement of the performance relative to the increasing number of samples approaching an asymptotic value. A reason for this behavior is lack of samples at locations representing the specular and anisotropic

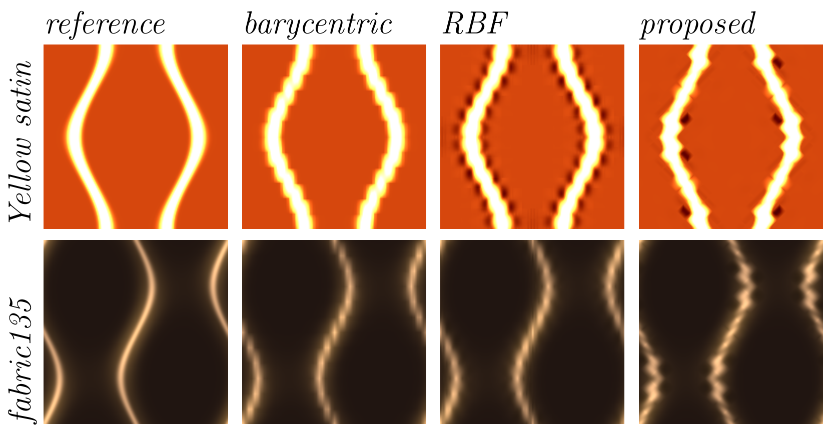

Figure 5.12: An example of a limitation of the proposed interpolation method when highlights are not perpendicular to the axial or diagonal slices ($\theta_i = 30^o$, $\theta_v = 75^o$).

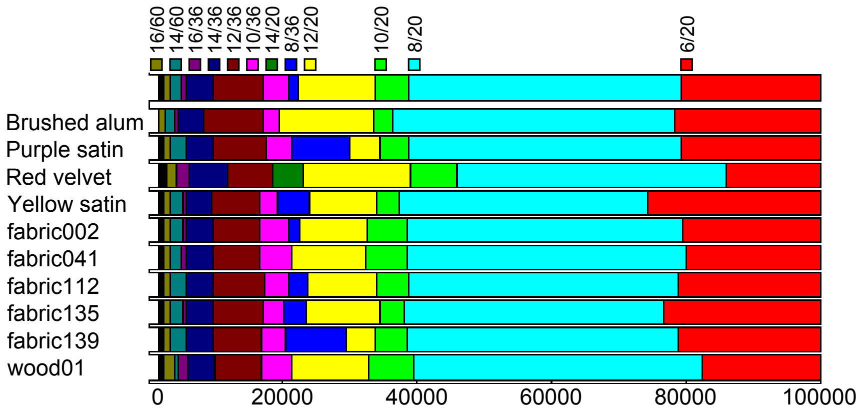

Figure 5.13: The first line: The color coded optimal placement of the slices in terms of e_s/a_s for a various count of samples. Other lines: The stability test of the optimal placement of the slices; at each line, one material is left out from computation of the optimal placement.

highlights. They are not sampled properly by the uniform sampling for the RBF and barycentric methods or by the uniform positioning of the slices in BRDF subspaces by the proposed method. We believe that the development of a better parameterization in combination with an adaptive placement of the slices would further improve the performance.

As the proposed method is designed to represent, in particular, features that are perpendicular to the slices, it can suffer from an improper representation of curved highlights or visual features occurring between the slices. Such behavior is typical for subspaces with a high difference between the viewing and illumination elevation angles as shown in Figure 5.12. We believe that these errors can be suppressed by introduction of an elevation-angle dependent interpolation method (see Sec. 5.7).

5.6.3 Timings

Reconstruction of the BRDF subspace (720×720 directions) measured adaptively by 24 slices takes typically 1-2 seconds regardless the count of samples using a non-optimized Matlab implementation. Reconstruction of the entire BRDF space with a uniform step of 2^o requires about 45 seconds using the barycentric interpolation; whereas, the RBF interpolation takes about 22 minutes, both regardless the number of samples. Reconstruction of the whole BRDF array by the proposed method using our Matlab implementation typically takes 4 minutes. All timings are obtained using a single core of Intel Xeon E5-2643 3.3 GHz.

5.6.4 Stability of the Placement of the Slices

To test stability of the proposed placement optimization, we perform the following experiment. We compute the optimal values of the parameters $|m|$-times using Equation 5.3;

Table 5.4: An increase of the MRE using the optimal placement of the slices instead of the best possible placement.

Brush. alum	Purple satin	Red velvet	Yellow satin	fabric 002	fabric 041	fabric 112	fabric 135	fabric 139	wood 01	mean
1.31×	1.19×	1.49×	1.32×	1.02×	1.03×	1.05×	1.27×	1.11×	1.17×	1.20×

each time with one material left out of the computation. As a consequence, there are $|m|$ different results, and we compare them with the results in Table 5.1. This is graphically illustrated in Figure 5.13, where the first line corresponds to the computation for all $|m|$ materials, and each other line corresponds to the computation with one material left out.

As one can see, boundaries of the individual regions (representing combinations of parameters e_s/a_s) vary slightly but not dramatically. When we leave *Red velvet* out, there is a green region (14/20) which does not appear on other lines substituting purple (10/36) and blue (8/36) regions. When we leave *Brushed alum*, *fabric041* or *wood01* out, the blue region (8/36) is missing. Note that usage of the parameters of neighboring regions brings only slightly worse reconstruction errors.

We evaluate an increase of the MRE due to usage of the optimal placement of the BRDF slices according to Table 5.1 instead of usage of the best placement for the tested material (which is unknown for a newly measured material). Evaluation is performed across a broad count ranging from 55 samples (resulting from the count of intersections of the slices in the sparsest configuration $e_s = 28^o/a_s = 180^o$) to 354,061 samples (resulting from the densest tested uniform distribution of samples across a hemisphere, see Sec. D.3). Results are summarized in Table 5.4. While the maximal increase of the error is almost 1.5-times, the average MRE across all the materials is only 1.2-times worse than if we use the best possible placement of the slices. We conclude that usage of the optimal placement according to Equation 5.3 is a good practice in the measurement of an unknown material.

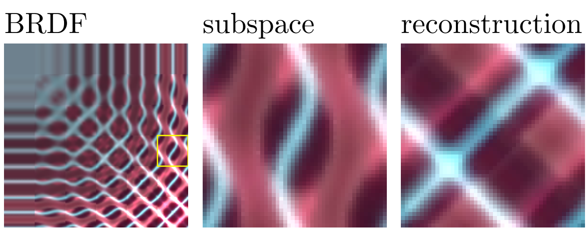

Figure 5.14: A BRDF subspace ($\theta_i = 30°, \theta_v = 75°$, middle) reconstructed from only one axial and one diagonal slice using the method proposed in Section 5.3 (right).

5.7 Improvement of the Reconstruction Method

As shown in Section 5.6.2, the proposed interpolation method introduced in Section 5.3 cannot handle reconstruction of features that are not perpendicular to the slices. This is especially evident at subspaces with a significant difference between illumination and viewing elevations, i.e., $|\theta_i - \theta_v| \geq 45^o$ as the anisotropic highlights at such elevations are not perpendicular to the slices as shown in Figure 5.14. Therefore, in this section, we aim to improve the reconstruction method using our findings of the typical behavior of anisotropic materials (Sec. 3.1).

To emphasize potential of the improved method and to keep things simple, we use only one axial and one diagonal slice per subspace. The proposed method is based on interpolation along the predicted direction of the anisotropic highlights (see Fig. 5.15-right). Then, the highlights are interpolated correctly without disturbing artifacts.

5.7.1 Improved Interpolation Method

Our interpolation method is based on the observation that BRDF values of an azimuthal subspace are almost constant for the constant value of the azimuthal angle φ_h of the half vector H. Therefore, a value of any reconstructed point P in azimuthal subspace (φ_i, φ_v) at the given elevations (θ_i, θ_v) can be interpolated from a value of point A on the diagonal slice and a value of point B on the axial slice as shown in Figure 5.15-right. Note that the azimuthal angle φ_h must be the same at P, A, B.

Equations to find coordinates of both points (A, B) in azimuthal space (φ_i, φ_v) can be derived as follows. First, let us remind basic equations for conversions between the spherical and Cartesian coordinates:

$$\begin{aligned} x &= \sin(\theta) \cdot \cos(\varphi), \qquad \theta = \arccos(z) \\ y &= \sin(\theta) \cdot \sin(\varphi), \qquad \varphi = \arctan\left(\frac{y}{x}\right) \\ z &= \cos(\theta) \ . \end{aligned}$$

From now on, the superscript will denote a point (P, A, B) and the subscript will, as usual, denote a direction (illumination, view, half). Elevation angles are fixed across the subspace where interpolation is performed, i.e., $\theta_i^A = \theta_i^B = \theta_i^P, \theta_v^A = \theta_v^B = \theta_v^P$. The azimuthal angle φ_h of the stencil intersecting the reconstructed point P can be obtained as:

$$\varphi_h^P = \arctan\left(\frac{y_i^P + y_v^P}{x_i^P + x_v^P}\right) \ ,$$

To find the point $B(\varphi_i^B, \varphi_v^B)$ on the axial slice we create a system of equations:

$$\begin{aligned} \varphi_i^B &= \varphi_v^B - \alpha, \\ \varphi_h^B &= \arctan\left(\frac{y_i^B + y_v^B}{x_i^B + x_v^B}\right) \ , \end{aligned} \tag{5.4}$$

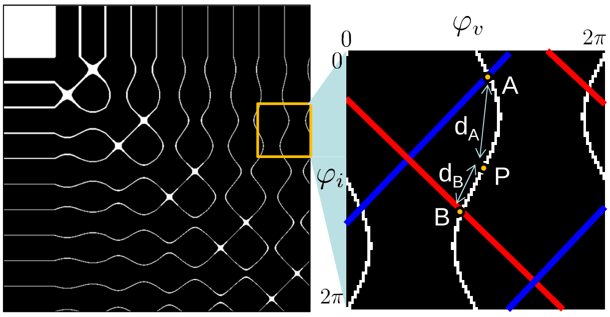

Figure 5.15: The interpolation along an anisotropic stencil that intersects the point P.

where α is a parameter of the axial slice (see Eq. 5.1). Similar equations can be derived for the point $A(\varphi_i^A, \varphi_v^A)$ on the diagonal slice:

$$\begin{aligned} \varphi_i^A &= \beta - \varphi_v^A, \\ \varphi_h^A &= \arctan\left(\frac{y_i^A + y_v^A}{x_i^A + x_v^A}\right) \quad , \end{aligned} \tag{5.5}$$

where β is a parameter of the diagonal slice (see Eq. 5.2). Based on $\varphi_h^A = \varphi_h^B = \varphi_h^P$ the system of equations can be created. The solution is:

$$\begin{aligned} \varphi_v^A &= \arctan\left(\frac{C_1 + C_3(\beta) + C_4(\beta)}{C_2 + C_5(\beta) - C_6(\beta)}\right), \\ \varphi_v^B &= \arctan\left(\frac{C_1 + C_3(\alpha) - C_4(\alpha)}{C_2 - C_5(\alpha) - C_6(\alpha)}\right) \quad , \end{aligned}$$

where

$$\begin{aligned} C_1 &= \sin(\theta_v^P) \cdot \sin(\varphi_h^P), \\ C_2 &= \sin(\theta_v^P) \cdot \cos(\varphi_h^P), \\ C_3(\nu) &= \sin(\theta_i^P) \cdot \cos(\nu) \cdot \sin(\varphi_h^P), \\ C_4(\nu) &= \sin(\theta_i^P) \cdot \sin(\nu) \cdot \cos(\varphi_h^P), \\ C_5(\nu) &= \sin(\theta_i^P) \cdot \cos(\nu) \cdot \cos(\varphi_h^P), \\ C_6(\nu) &= \sin(\theta_i^P) \cdot \sin(\nu) \cdot \sin(\varphi_h^P). \end{aligned}$$

As the viewing and illumination azimuthal angles are coupled by means of an initial difference between azimuthal angles prior to the measurement (α for the axial slice, β for the diagonal slice), we can also obtain a corresponding illumination azimuths φ_i for the axial and diagonal slices using Equations 5.4 and 5.5, respectively.

Finally, we can reconstruct the unknown reflectance value at the point P using a linear combination of the reflectance values at the points A, B weighted by the Euclidean distances:

$$f_r(\theta_i, \theta_v, \varphi_i, \varphi_v) = s_{A,\theta_i\theta_v\alpha}(\varphi_v^B)\frac{d_A}{d_A + d_B} + s_{D,\theta_i\theta_v\beta}(\varphi_v^A)\frac{d_B}{d_A + d_B} \quad ,$$

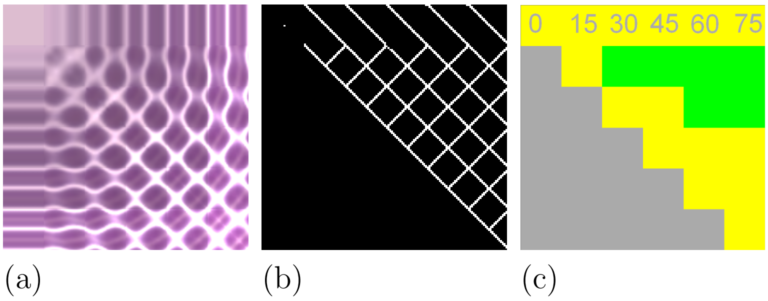

Figure 5.16: An example BRDF (a), a mask showing a distribution of $1,555$ samples in two slices per subspace (b), and a distribution of subspaces reconstructed by the original approach (yellow) and by the proposed approach (green) (c).

where d_A, d_B are Euclidean distances of the point P to the points A, B in the azimuthal space, respectively (see Fig. 5.15-right).

5.7.2 Results and Discussion

As our reconstruction method is currently beneficial only to anisotropic BRDFs, we analyzed the variance along the axial slices in all BRDF subspaces of the 150 BRDFs in the database (Sec. A.4) in order to detect only those having significant anisotropic behavior. Out of these 150 BRDFs, we eventually selected 37 anisotropic BRDFs to test our reconstruction method. See an example of the anisotropic BRDF of *fabric094* in Figure 5.16-a.

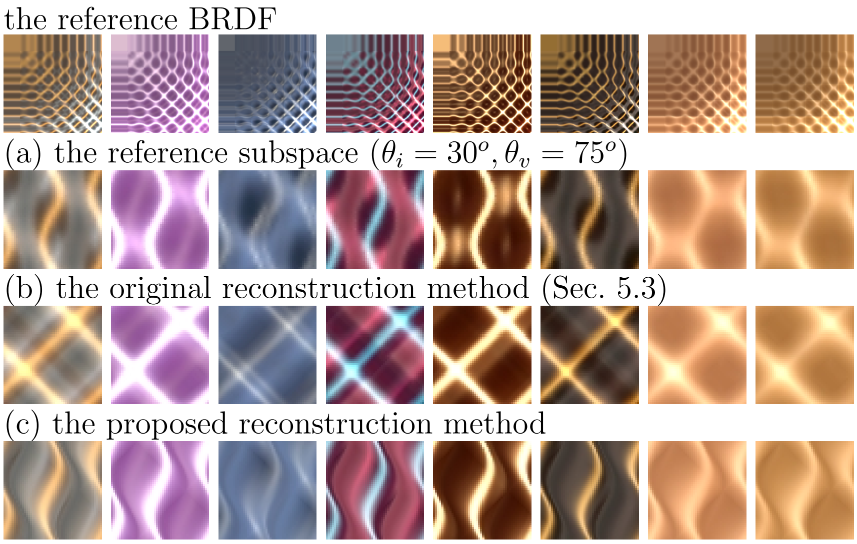

Figure 5.17: A comparison of the BRDF subspace $\theta_i = 30^o, \theta_v = 75^o$ (a) reconstructed using the original (b) and the proposed (c) methods.

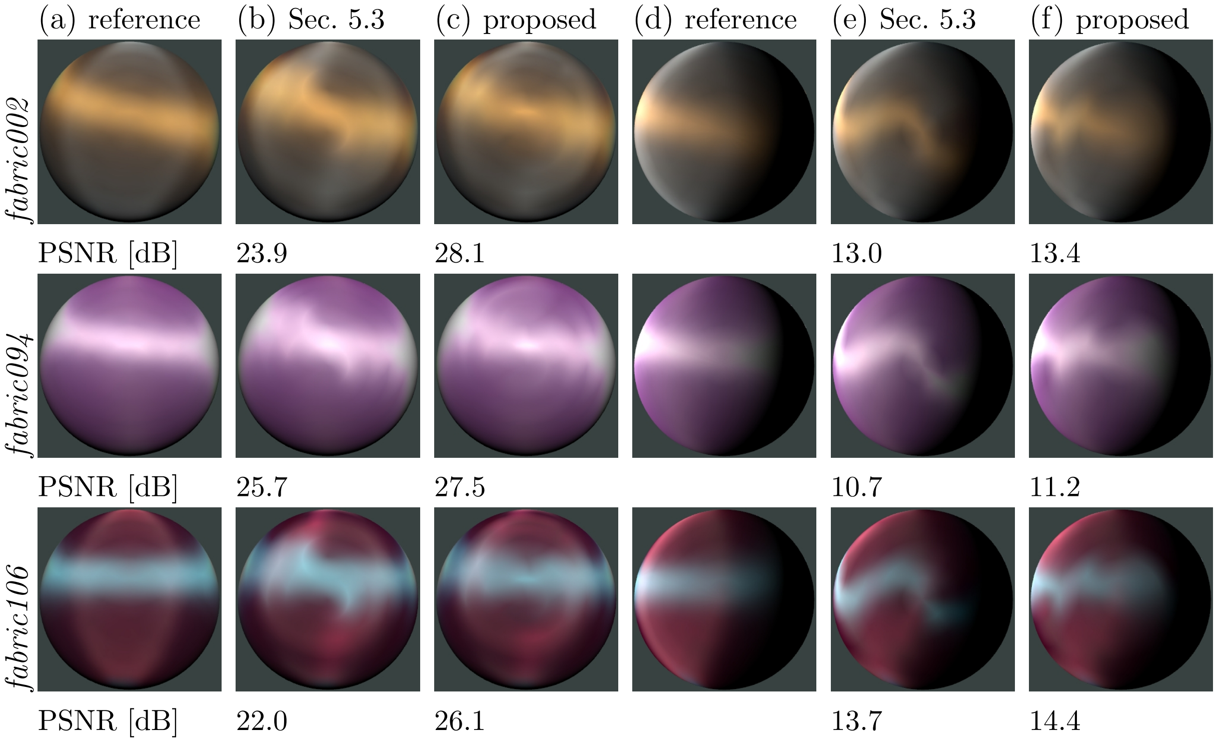

Figure 5.18: A comparison of the BRDFs on spheres for two positions of point-light sources: (a,d) the reference, (b,e) the method of Section 5.3, (c,f) the proposed method.

The new reconstruction method does not work properly when the viewing and illumination elevations are approximately equal. Therefore, in this case we use the original interpolation approach (Sec. 5.3). For the experiment, we settled at using the new method for elevations as denoted by green in Figure 5.16-c and using the original method at elevations as denoted by yellow. On the other hand, the original approach works well in this case as both specular and anisotropic highlights are perpendicular to the slices.

Performance of the proposed reconstruction method as compared to the original method (Sec. 5.3) for the subspaces ($\theta_i = 30^o, \theta_v = 75^o$) of eight anisotropic BRDFs are shown in Figure 5.17. One can observe an inaccurate reconstruction of the subspaces by means of the original method (b), while the proposed approach (c), when compared to the reference (a), works reasonably well. Results of the reconstructed BRDFs rendered on spheres illuminated from front or left respectively are shown in Figure 5.18. Again, the new approach (columns c and f) demonstrates an improvement over the original method (columns b and e). Three of the tested fabric materials illuminated by the *grace* environmental illumination are shown in Figure 5.19. These results also demonstrate a considerable quality gain achieved by the new method when compared to the former one. Finally, Figure 5.20 shows a graph comparing the performance of both methods across all 37 tested anisotropic BRDFs. One can observe a steady reconstruction improvement of the proposed method over the original solution.

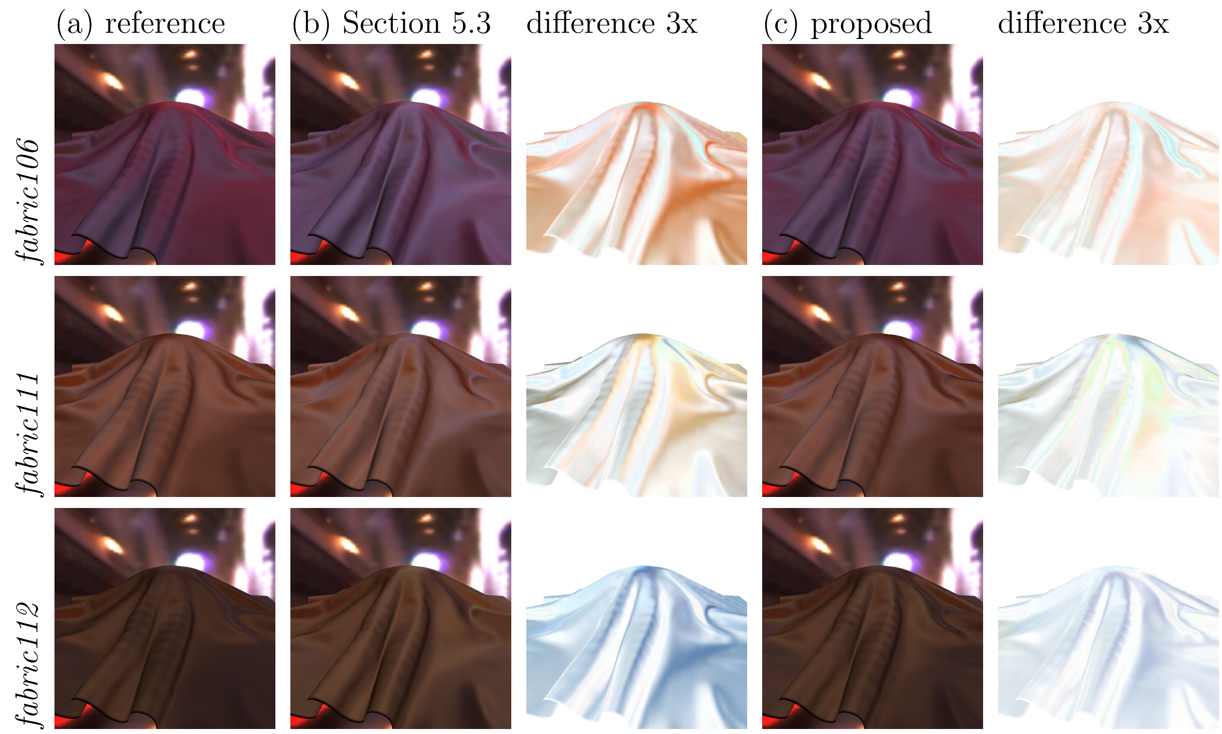

Figure 5.19: Renderings of three fabric materials in the *grace* illumination environment: (a) the reference BRDF, (b) the original reconstruction method, (c) the proposed method.

As the results suggest, our method has currently several limitations. It does not work properly if the material is isotropic or nearly isotropic, and it cannot handle subspaces where $\theta_i = \theta_v$ due to singularities in the equations (undefined φ_h if $\theta_h = 0$). We believe that this issues can be handled in future by a more elaborate design of the equations. Also, the method needs to be extended to handle more than two slices per subspace and to provide reliable reconstructions in the 4D dataspace.

As no demanding computation is involved, the computational costs of the proposed reconstruction method are comparable with the original one and the reconstruction of the 2D BRDF subspace takes approximately 1 second using our MATLAB implementation.

5.8 Summary

This chapter deals with the efficient sampling and reconstruction of the anisotropic BRDFs from a predefined number of samples. First, we sample the four-dimensional dataspace by means of the one-dimensional slices while adapting a density of the samples along the slices according to properties of a given material. To find the optimal placement of the slices in the 4D BRDF space, we perform the extensive study on ten materials and we provide parameters that enables us to chose the optimal placement based on the required number

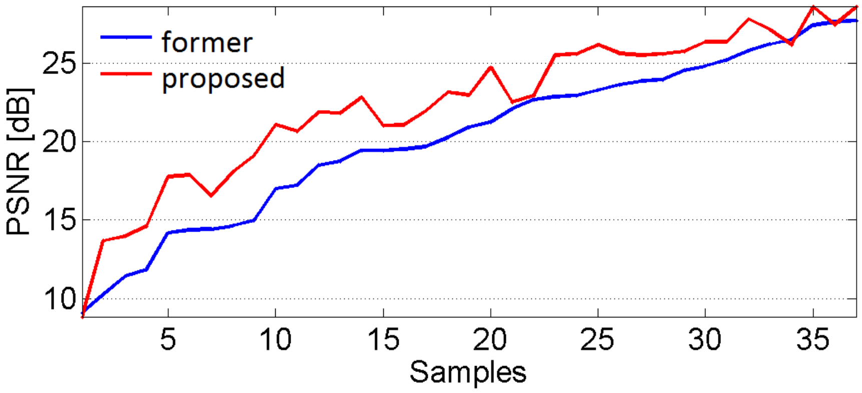

Figure 5.20: The performance of our method across 37 anisotropic BRDFs.

of measured samples. Then, we present the interpolation method to reconstruct the 4D BRDF space that is very fast, easy to implement, and allows for a GPU implementation as well (see Sec. D.4).

We compare the performance of our method with the barycentric and RBF interpolation approaches as applied to many uniform sampling schemes with a variable count of samples. First, we use only the axial and diagonal slices at nine subspaces of one material. The results of the proposed method are superior to those of conventional uniform methods. Then, we test our method in the 4D BRDF space by the simulated measurement experiment using an analytical anisotropic BRDF model as a source of reference data. The proposed method achieves on average almost four-times lower reconstruction errors than the competing approaches given the same number of samples. Alternatively, for a given reconstruction error, it requires between two- to five-times less samples. In the real measurement of the anisotropic BRDFs, our method achieves a twice lower reconstruction error than other approaches and our further experiments suggest that the captured BRDF data belong among the best publicly available anisotropic BRDFs (see Sec. A.5). Finally, we evaluate the contribution of the horizontal and vertical slices included on the top of the axial and diagonal slices. As they help to achieve twice the lower errors, we recommend to sample along the horizontal and vertical slices whenever a measurement setup allows that.

Our analysis shows that the reconstruction error nearly does not depend on a used type of the adaptive sampling along the slices as the convergence to low error rates is already very fast; therefore, the interpolation method (Sec. 5.3) remains the main source of errors. Hence, we propose a new interpolation method based on observation that BRDF values in a 2D subspace are almost constant for the constant value of the azimuthal angle φ_h. As a consequence, the method interpolates values along predicted locations of the anisotropic highlights. Although the new method performs better than the former method at some subspaces, it is not elaborate enough to handle subspaces where $\theta_i = \theta_v$, to perform well in case of an isotropic or a nearly isotropic material, and it does not enable us to reconstruct the whole 4D BRDF space yet. Therefore, improvement of the new method is subject of our future work.

Chapter 6

Towards Rapid and Affordable Appearance Acquisition

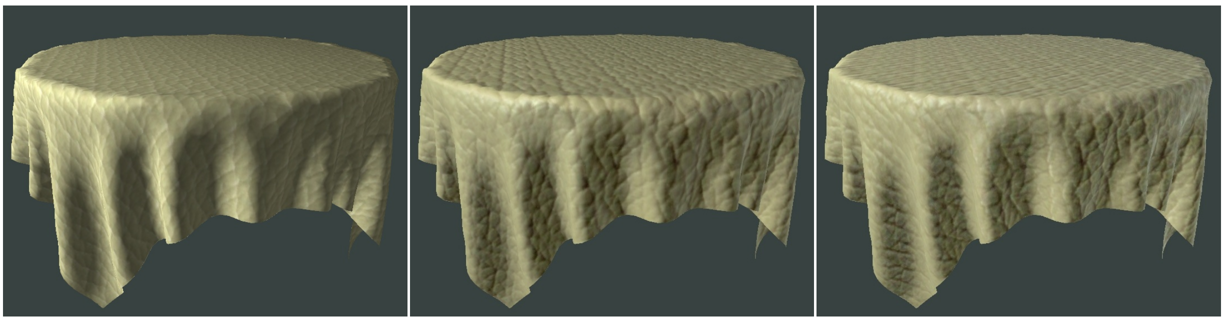

Figure 6.1: A comparison of renderings based on a single texture modulated by an analytical BRDF model [KSKK10] (left), the reference BTF measurement using 6,561 images (middle), and the proposed measurement and reconstruction using 168 images (right).

The acquisition of BRDF or BTF data is usually time- and resources-intensive due to the high dimensionality of the data. This results in expensive, complex measurement setups and excessively long measurement times. We describe an approximate BRDF, ABRDF and BTF acquisition setup proposed in [A.4] based on a simple, affordable mechanical gantry containing a consumer camera and two LED light sources. It allows for the rapid acquisition of a subset of the (A)BRDF or BTF dataset in six minutes and the fully automatic reconstruction of the entire approximate dataset in one hour by the measurement of a very limited set of images of a material by shooting several video sequences.

The method, first introduced in [A.6], is based on acquisition of one axial and one diagonal slice (see Sec. 5.1) per the selected subspace (see Fig. 6.2). As the slices are perpendicular to the anisotropic and specular highlights, they bear enough information to approximate the azimuthally-dependent behavior of the ABRDF. The main advantage

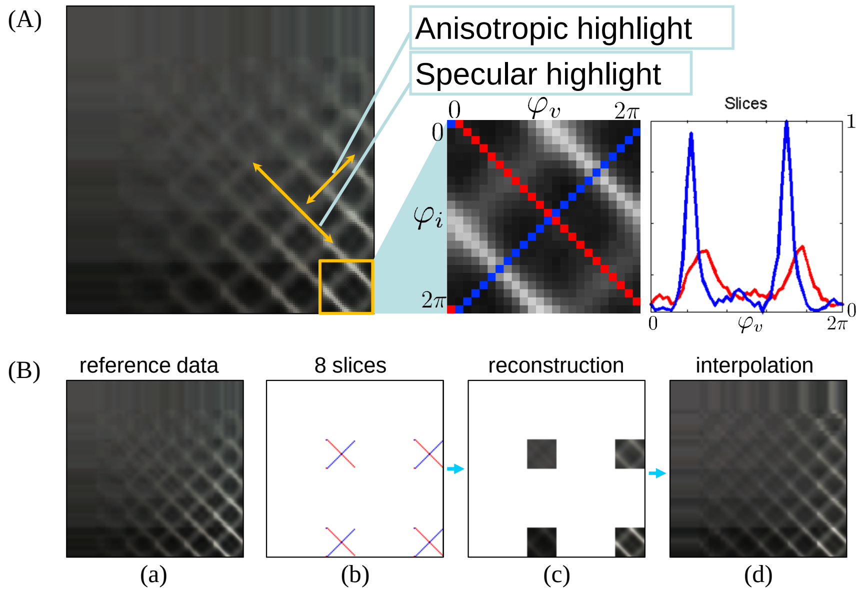

Figure 6.2: (**A**) A schema of an ABRDF subspace with two slices, and (**B**) a principle of the entire ABRDF reconstruction from four recorded subspaces: (a) the reference, (b) sampling by eight slices, (c) reconstruction of four subspaces, (d) interpolation of missing data.

of the slices is that they can be captured by a continuous rotation of arms holding a camera and a light source (see Fig. 6.3). That significantly reduces acquisition time of the appearance data and allows us to build a simple and affordable acquisition setup.

Our setup does not impose any additional restrictions on the measured type of material or its properties (namely in terms of isotropy, reciprocity, bilateral symmetry, opacity) when compared to other setups. Speed of the setup is demonstrated on the measurement of human skin and the measurement and modeling of the time-varying process: dessication of glue (see Appx. E). As it allows for the fast, inexpensive acquisition of the approximate BTF, this method can be beneficial to visualization applications demanding less accuracy where utilization of the BTF has previously been limited (see Fig. 6.1).

Generally, our setup is capable to measure several axial and diagonal slices per subspace. Nevertheless, we stick with the simplest scenario of only one axial and one diagonal slice per subspace to emphasize potential of the method. Similarly, it is possible to use more light sources and elevations of a camera, but again, we stick with the simplest reasonable scenario of two light sources and two positions of a camera. In Section 6.4.4, we investigate benefits of additional light sources and cameras as well as benefits of additional slices per subspace.

This chapter is structured as follows. First, Section 6.1 describes the method of reconstruction of the 4D dataspace from a very limited set of samples. Then, Section 6.2 describes the proposed inexpensive setup and procedures of data processing. Results of

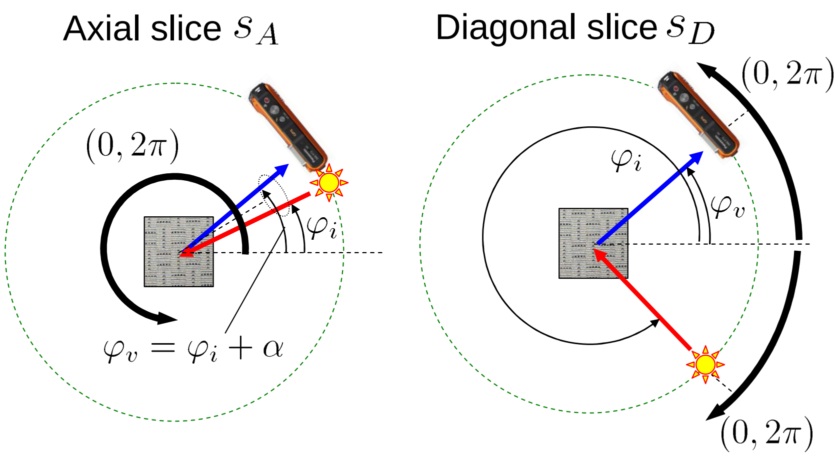

Figure 6.3: The principle of the proposed measurement of the axial slice (left) and the diagonal slice (right) at fixed elevations θ_i, θ_v.

the approach are presented in Section 6.3, Section 6.4 discuses advantages and limitations of the proposed method and setup and analyses possibilities of their improvement, and Section 6.5 summarizes the chapter.
The main contributions of this chapter are:

- A reconstruction method of the entire anisotropic (A)BRDF or BTF space from less than two hundred sparsely measured samples.

- A setup for intuitive and fast (A)BRDF and BTF acquisition that is composed of a consumer camera and LED light sources.

The main features of the proposed method are:

- A correct reconstruction of the non-reciprocal, energy non-conserving ABRDF.

- An arbitrarily dense sampling of the specular highlights, without increasing measurement time.

- Fast measurement using an inexpensive measurement setup.

- It is unnecessary to extract the measured sample from its environment.

- Contrary to analytical BRDF models, this method requires neither a lengthy fitting procedure nor guessing at initialization values.

6.1 The Proposed Method

A principle of the method is sketched in Figure 6.2B. First, ABRDF samples are measured at predefined locations at given elevations by means of eight slices (see Sec. 6.1.1 and Fig. 6.2B-b). Then, four sampled periodical ABRDF subspaces are reconstructed from the values of the slices (see Sec. 6.1.2 and Fig. 6.2B-c) and finally, the remaining values at unmeasured elevations are interpolated (see Sec. 6.1.3 and Fig. 6.2B-d). All pixels of the BTF (i.e., ABRDFs) are processed separately in this way.

6.1.1 Acquisition of the Slices

Because two of our major concerns are simplicity and speed of the acquisition process, we suggest taking samples by the continuous movement of a light source and a camera around the sample at fixed elevations. By doing so, samples can be taken at an arbitrary density, limited only by speed of the movement of a camera and its frame-rate. Each subspace of azimuthal angles (φ_i, φ_v) is sampled by means of two perpendicular slices (see Sec. 5.1 and Fig. 6.5-a), which differ in the direction of the mutual movement of a camera and a light source (see Fig. 6.3). In principle, the slices are orthogonal to the most prominent features (the specular reflections and the anisotropic reflections) of many materials. These features are often constant in the direction perpendicular to the slices and thus can be effectively represented by their marginal values (see Fig. 6.4).

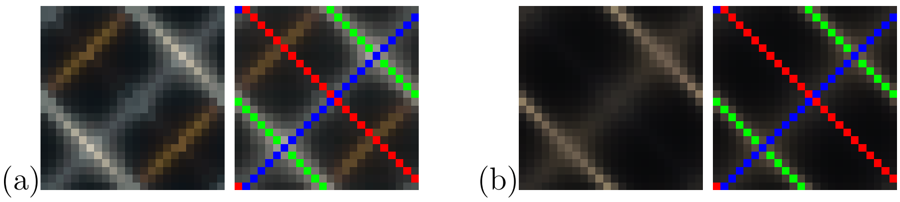

Figure 6.4: Examples of periodical ABRDF subspaces for derivation of the optimal placement of the axial slices (green and red alternatives).

We focus first on an analysis of the ABRDF subspace that have the highest elevations, in which the illumination and view dependent effects are the most pronounced (see the first row of Fig. 6.7). We analyze the optimal placement of the axial and diagonal slices across many ABRDFs. While the placement of the diagonal slice can be arbitrary, the highest variance along the axial slice is achieved near the specular highlight ($\alpha = 180°$, green dots in Fig. 6.4). Consequently, this is – the most likely – the best placement of the axial slice. However, such a placement might omit vital color or luminance information in some parts of the subspace. For example, it would completely miss yellow anisotropic features as in Figure 6.4-a or dark parts as in Figure 6.4-b. Therefore, we use the slice with the second highest variance ($\alpha = 15°$, red dots in Fig. 6.4) to capture the most of the visual features of the ABRDF subspace.

6.1.2 Reconstruction from the Slices

Reconstruction of a periodical ABRDF subspace is performed for elevation angles at which the slices are captured (see Fig. 6.2B-c). Section 5.3 provides a method for reconstruction of a subspace from an arbitrary number of slices. Here we provide an alternative method which can reconstruct a subspace from one axial and one diagonal slice only. The method is easier to implement than the method of Section 5.3 and often achieves slightly better results for one axial and one diagonal slice. It can be explained as a combination of two sets of marginal values as shown in Figure 6.5.

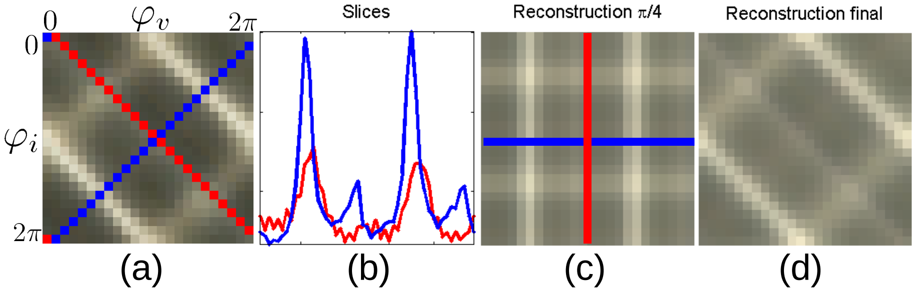

Figure 6.5: Reconstruction of a periodical ABRDF subspace from two slices at fixed elevations: (a) the reference data with placement of the slices, (b) the data profiles of the slices, (c) the reconstruction from the slices ($\frac{\pi}{4}$ rotated), (d) the final reconstruction.

Reconstruction of a direction (φ_i, φ_v) in the ABRDF subspace (θ_i, θ_v) starts with combining contribution of the axial (s_A) and diagonal (s_D) slices. We tested their sum and product; however, the latter improperly enhances the locations at intersection of the specular and anisotropic highlights as shown in Figure 6.6-b. Note that the sum preserves the specular highlights, which are less affected by the anisotropic highlights (see Fig. 6.6-c).

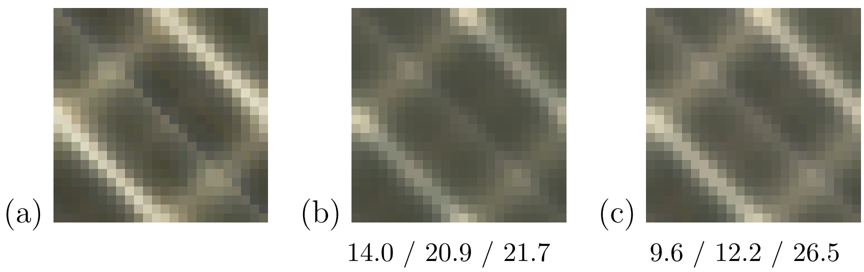

Figure 6.6: The reference subspace (a), and its reconstruction using the product (b) and the sum of the slices (c). Below are the difference values in CIE ΔE / RMSE / PSNR [dB].

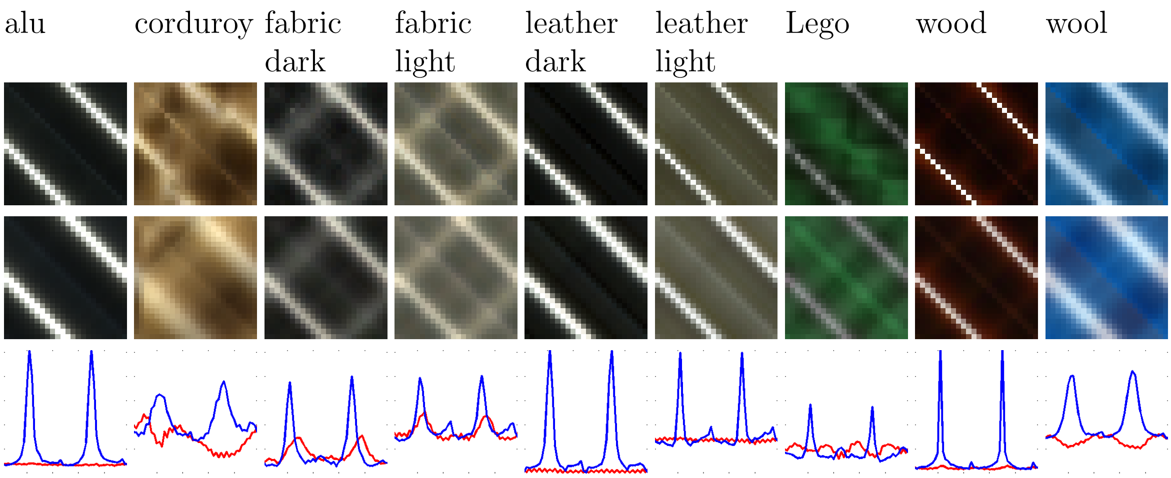

Figure 6.7: A comparison of an ABRDF subspace (the first row) with its reconstruction (the second row) from the axial (red) and diagonal (blue) slices (the third row).

Therefore, we use the sum of the slices in our reconstruction procedure:

$$\begin{aligned} v_{\theta_i\theta_v}(\varphi_i, \varphi_v) &= s_{A,\theta_i\theta_v\alpha}(\varphi_{i,R}) + s_{D,\theta_i\theta_v\beta}(\varphi_{v,R}) \ , \\ \begin{bmatrix} \varphi_{i,R} \\ \varphi_{v,R} \end{bmatrix} &= \begin{bmatrix} \cos(\pi/4) & -\sin(\pi/4) \\ \sin(\pi/4) & \cos(\pi/4) \end{bmatrix} \begin{bmatrix} \varphi_i \\ \varphi_v \end{bmatrix} \ . \end{aligned}$$

Note that the original azimuths φ_i, φ_v are rotated for $\pi/4$ (Fig. 6.5-c) to account for the slant of the slices with respect to the φ_i, φ_v coordinate system (Fig. 6.5-a). Finally, the summed value v is mapped to the dynamic range of the original slices:

$$\begin{aligned} f_r(\theta_i, \theta_v, \varphi_i, \varphi_v) &= v_{\theta_i\theta_v}(\varphi_i, \varphi_v) \cdot (M - m) + m \ , \\ m &= \min(s_{A,\theta_i\theta_v\alpha} \cup s_{D,\theta_i\theta_v\beta}) \ , \\ M &= \max(s_{A,\theta_i\theta_v\alpha} \cup s_{D,\theta_i\theta_v\beta}) \ . \end{aligned}$$

Since the axial slice has always a constant value for isotropic materials, the slices do not have to be combined and the reconstruction can be performed using the diagonal slice alone as:

$$f_r(\theta_i, \theta_v, \varphi_i, \varphi_v) = s_{D,\theta_i\theta_v}(\varphi_{v,R}) \ .$$

Figure 6.7 shows reconstruction of the anisotropic ABRDF subspace at elevations $\theta_i = 75°, \theta_v = 75°$ from two slices and proves the ability of the proposed approach to represent a variety of anisotropic materials.

6.1.3 Interpolation of Missing Values

Now, we know how to sample and reconstruct the four subspaces at the following elevations: $(\theta_i, \theta_v) = \{(30^o, 30^o), (30^o, 75^o), (75^o, 30^o), (75^o, 75^o)\}$ (see Fig. 6.2B-c). However,

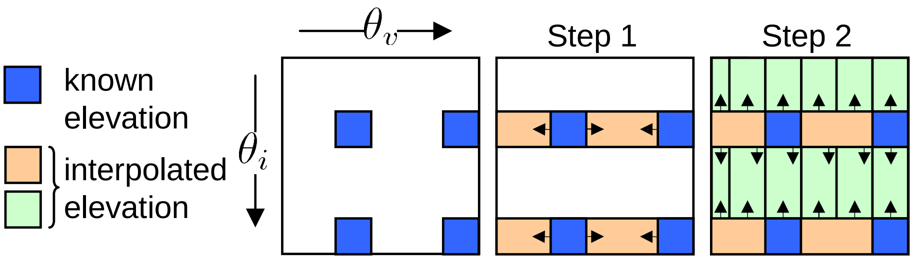

Figure 6.8: Steps of the interpolation method B.

data for the remaining subspaces are still unknown and have to be estimated. The BRDF parametric models, e.g., [NDM05], cannot be used to solve this problem because they impose restrictions on data properties (reciprocity, energy conservation, etc.), require more samples, a different distribution of samples, and lengthy fitting. Therefore, we propose the following two interpolation approaches.

Method A – In the first one, the interpolation is performed by means of the four-dimensional RBFs [PTVF92] computed separately in each color channel. We tested several parameterizations of illumination and viewing directions, e.g., $[\theta_i, \varphi_i, \theta_v, \varphi_v]$, $[\theta_h, \varphi_h, \theta_d, \varphi_d]$ from [Rus98] and finally, we use the parameterization according to [HFM10] applied to both illumination and viewing directions $[\alpha_i, \beta_i, \alpha_v, \beta_v]$. This parameterization has the lowest reconstruction error due to alignment of the specular highlights to the 0° value of the angle β_i.

Method B – Due to relatively high computational demands of the method A, we developed a faster hybrid linear interpolation constrained by a reflectance model. This interpolation consists of two steps shown in Figure 6.8. First, we interpolate data at different viewing and constant illumination elevations. Then, the remaining illumination elevations are filled. In each interpolation step, the average ABRDF value for each unknown elevation is approximated by fitting a simplified monospectral one-lobe Lafortune model

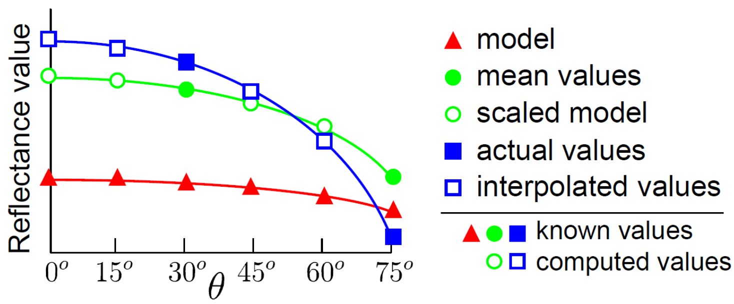

Figure 6.9: An interpolation of values at unmeasured elevations.

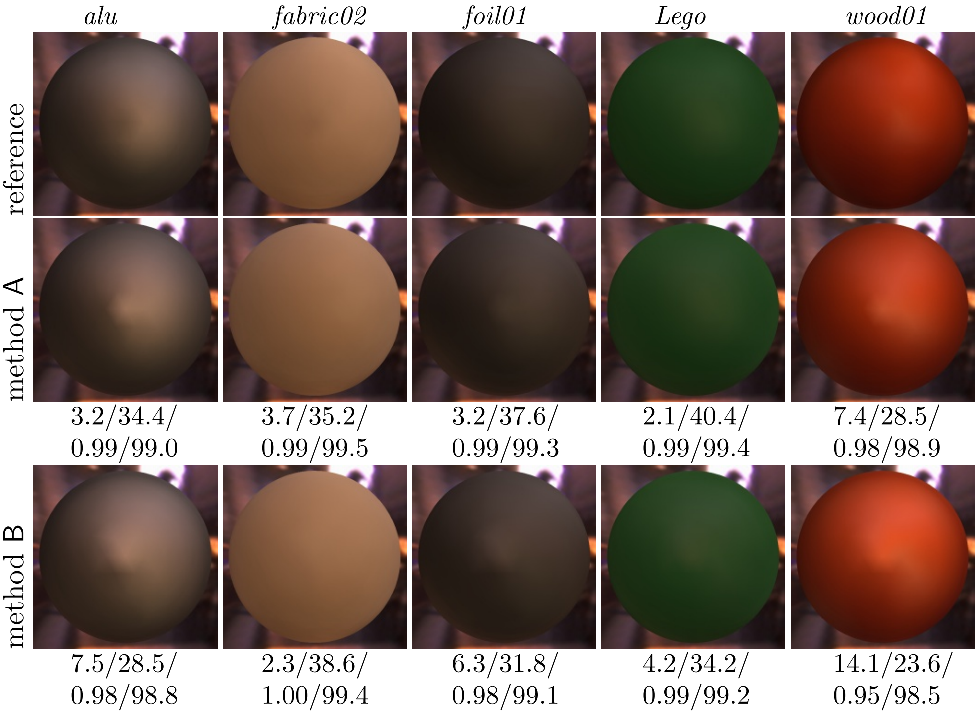

Figure 6.10: Performance of the interpolation methods in the *grace* illumination environment. Below are the difference values in CIE ΔE / PSNR [dB] / SSIM / VDP2.

[LFTG97] with parameters k, α to known values:

$$f_r(k, \alpha) = k(\cos\theta_i \cdot \cos\theta_v)^\alpha \quad . \tag{6.1}$$

Initialization of k, α is constant during all experiments. A value obtained from the model at elevation θ is scaled by the mean values of the slices and then used to obtain an interpolation weight. These weights are then applied to a linear interpolation of missing elevations from values at known elevations as shown in Figure 6.9. Elevations lower than 30° are extrapolated using predictions of the scaled model (Eq. 6.1). This procedure is performed over all azimuths as shown in Figure 6.8.

The method A is approximately three-times computationally more intensive than the method B and provides better results in the most cases. While the method A allows arbitrarily dense sampling, even for originally unmeasured azimuths, the method B reconstructs the ABRDF data in their original azimuthal sampling. The results in Figure 6.10 show the major visual differences between both proposed interpolation methods in the *grace* illumination environment [Deb98].

We also tested a modification of the method B denoted as BD that uses a displacement interpolation [BvdPPH11]. Its principle is based on solving the generalized mass transport optimization problem. As this method cannot extrapolate, the low elevation areas are reconstructed using the method B. The method BD gives better results than the method

B and comparable results to the method A. However, it is about thirty times slower than the method A when interpolation of the whole subspace is applied. When viewing or illumination direction is fixed, its speed is comparable to the method A.

6.2 Acquisition Setup and Data Processing

This section describes a fast, practical measurement setup to capture the slices in the (A)BRDF or BTF space using a consumer camera and LED point-light sources. First, mechanical construction of the setup is described. Then, information related to process of the measurement is provided and finally, data processing and calibration procedures are explained.

6.2.1 Acquisition Setup

The setup that realizes measurements of the slices is shown in Figure 6.11. It consists of a mechanical gantry holding two arms that rotate synchronously in either the same or the opposite direction. The gantry is built using the Merkur toy construction set (*http://www.merkurtoys.cz*). We use a single DC motor (4.5 V) that runs at a constant speed. Usage of additional gears guarantees accurate synchronization of the arms and allows us to switch a direction of the mutual rotation of the arms (see Fig. 6.11-top-right).

One of the arms holds two LED Cree XLamp XM-L light sources with the 20° frosted optics (see Fig. 6.11-bottom-right). Weight of the arms is balanced and we use a rotary slip ring to avoid clumsy wires of the LED power supply (0.7 A/3 V). The second arm has two positions for attachment of the Panasonic Lumix DMC-FT3 camera. One advantage of this camera is that it does not have a protruding lens, which could possibly block the arm bearing the light sources. During its movement, the camera records appearance of the

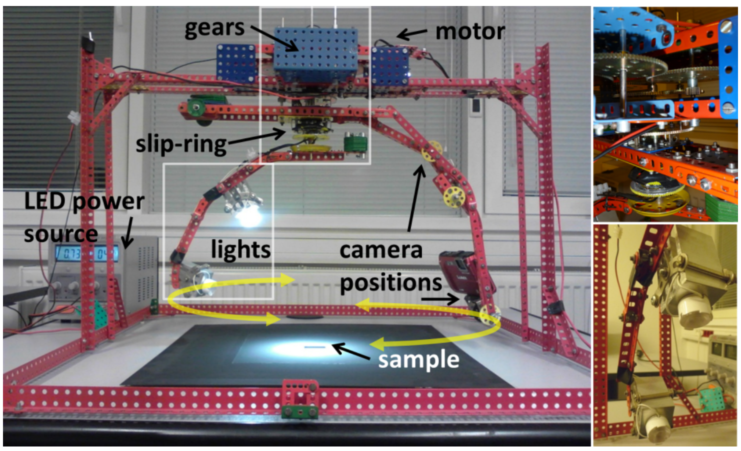

Figure 6.11: The proposed measuring device.

material as a video sequence at a resolution of 1280×720 pixels. Elevations of the LEDs and the camera at both positions are fixed at 30° and 75°.

The setup can be constructed in 10 hours, for less than $350, and can hold almost any compact camera. Dimensions of the setup are 0.6 × 0.6 × 0.4 m and its weight is 6 kg. The frame holding the arms can be folded down on the supporting platform, so the size of the setup can even be reduced to allow easy transportation and the usage for a field measurement of materials that cannot be removed from their environment.

6.2.2 Measurement

A material is placed below the setup in the axis of rotation of the arms (see Fig. 6.11). While data of the axial slice are measured by the mutually fixed rotation of the light source and the sensor around the material (see Fig. 6.3-left), data of the diagonal slice are obtained by the mutually opposite rotation of the light source and the sensor (see Fig. 6.3-right). In both cases, the camera and the light source travel full circle around the material and return to their initial positions.

The camera records appearance of a material as a video sequence at a resolution of 1280 × 720 pixels. Both axial and diagonal slices are recorded for two different elevations of the camera and the light source; therefore, eight slices are measured at elevations (θ_i, θ_v) $= \{(30°, 30°), (30°, 75°), (75°, 30°), (75°, 75°)\}$ as shown in Figure 6.2B-b. Recording of a single slice takes 30 seconds and the entire set of eight slices is captured in six minutes – during that time the LEDs are switched four times, gears twice, and the position of the camera once. Note that the fully automatic device assembled with two cameras would need only two minutes.

Zoom and white balance of the camera are fixed during the recording to allow calibrations described in Section 6.2.3. The image stabilizer of the camera is switched on and the global exposure level is set to $-\frac{2}{3}$ EV to avoid overexposure.

Due to a short distance of the camera and the LEDs from the material, we limited the size of the measured material to 30 × 30 mm, which is sufficient for a wide range of materials. However, shadows cast by the registration frame (see Fig. 6.12) limits the effective size to about 80% of the area. Although a larger size of the material can be used, we limited it due to a large span of illumination and viewing directions across the material, which might possibly cause a spatial non-uniformity in the captured images. Finally, we cut the smallest possible repetitive tiles found near the center of the material. A white border with a white dot is attached around the material (see Fig. 6.12) to detect orientation of the camera with respect to the material and for registration purposes.

6.2.3 Data Processing and Calibrations

Basic processing steps applied to the measured data are outlined in Figure 6.12 and explained in more detail here.

Frame Filtering and Registration – Eight slices are recorded as videos at a frame rate of 30 fps, i.e., we have 900 frames per slice. As the camera provides the M-JPEG

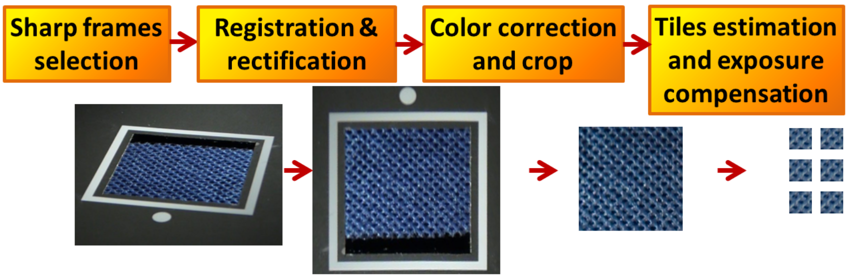

Figure 6.12: The processing pipeline of each image.

non-interleaved variable bit-rate format storing all recorded frames (i.e., not only the key-frames), the error introduced by a video coding is negligible compared to the ABRDF reconstruction error.

From each of the eight video sequences, 24 frames are extracted corresponding to sampling of azimuthal angles (φ_i, φ_v) every 15°. However, as not all of the frames are sharp due to the motion blur, we search in the neighborhood of ± 4 frames for the sharpest image minimizing the edge-based blur metric [MDWE02]. As this approach might miss intensity at the specular reflections, we additionally scan the entire sequence for the two brightest frames, which are sampled always. To capture even the very sharp specular highlights, four additional samples 1° apart from the specular reflections are recorded as well. This leaves us with a set of 28 frames for each diagonal slice and 24 frames for each axial slice, i.e., 208 frames in total. Note that in the practical measurement of the diagonal slice, up to two frames are removed due to an occlusion of view of the camera by the LEDs; therefore, the number of eventually selected frames is slightly lower. A subset of the recorded frames is used for camera calibration (see Sec. C.2) and further for compensation of a geometric distortion of all frames.

The frame registration itself is performed in two steps. First, all images are registered based on detection of the white border placed around the material using the Hough transform and computation of the homography between the detected and desired coordinates of corners. Second, as the plane of the measured material is usually at least 0.1 mm below the registration plane defined by the white border, we detect and compensate for this shift and eventual slant using the iterative method described in Section 4.2. Finally, the registered images are cropped to a size of 300 × 300 pixels, yielding a resolution of 340 DPI. If the sequences were recorded in the Full HD resolution (1900 × 1080 pixels), the BTF resolution would approach 500 DPI.

Exposure and Light Non-Uniformity Compensation – Unfortunately, the most compact cameras adapt their exposure depending on the amount of light coming from the scene, which is also true for the used camera. On a positive note, this enables us to capture as much information as possible, even using a limited dynamic range of the sensor (8bits/channel). However, the information about exposure throughout the video sequence could not be retrieved from an EXIF header as is possible for still photos.

Another related problem is the spatial non-uniformity of illumination, which is caused by the limited distance of LEDs from the material. Therefore, we compensate for exposure fluctuations as well as for the spatial non-uniformity of illumination in the original images I using the intensity of the spatially-uniform black material with the known BRDF at the locations beyond the white border (see Fig. 6.12).

A compensation image C is computed for each frame by the linear interpolation of intensity of the black material. First, the images I, originally measured in locations of the slices, are resampled to the azimuthally uniform grid producing the images I_G. A pixel of the monospectral correction image K represents the value in the center of the image divided by a value interpolated from measured intensities behind the white border, i.e., this value includes the known reference BRDF of the black material f_r divided by the compensation image C for the corresponding angles:

$$K_i(x, y) = \frac{1}{3} \sum_{\lambda \in \{R,G,B\}} \frac{f_{r\lambda}(\xi(i))}{C_i(x, y, \lambda)} \ ,$$

where $i = 1 \dots m$ is the number of the image in the slice that consists of m images, $\lambda \in \{R, G, B\}$ is a color channel and ξ is a known mapping function $\xi(i) \rightarrow (\theta_i, \varphi_i, \theta_v, \varphi_v)$. The compensated image I_C is obtained as:

$$I_{C,i}(x, y, \lambda) = I_{G,i}(x, y, \lambda) K_i(x, y) \ .$$

The reference BRDF of the black material f_r was measured by the UTIA gonioreflectometer (see Sec. 4.1).

Computation of the Viewing and Illumination Directions – When all images are registered and compensated, their corresponding illumination and viewing directions can be identified. The viewing direction (θ_v, φ_v) is obtained using the extrinsic parameters of the camera given the known camera calibration and coordinates of the corners of the white border. As the viewing direction is known and the arm that holds the light sources is mechanically coupled with the arm that holds the camera (Fig. 6.11), the illumination azimuth can be computed as $\varphi_i = \varphi_v - \alpha$ for the axial slice and $\varphi_i = \beta - \varphi_v$ for the diagonal slice. The elevation θ_i is known as the vertical positions of the LEDs are fixed in the setup (Fig. 6.11).

Colorimetric Calibration – The images are further colorimetrically compensated using a transformation matrix that relates measured and known color values of the Color-Gauge Micro Target (35 × 41 mm) in the least-square sense [Joh96]. This target has colors of the same specification as the X-Rite ColorChecker Classic target, however, its physical size is reduced.

Reconstruction of the Entire BTF Space – Since direct use of images of the material for BTF rendering would exhibit distractive seams on textured surfaces, we employ an image-tiling approach to find seamless tiles [SH05]. Finally, the compensated and tiled images of the material measured in locations of the slices are used for reconstruction of the remaining directions. Due to a lower computational complexity, we usually use the method B described in Section 6.1.3.

Timings – A typical timeframe for data processing using a desktop PC with the Intel i7-3610QM 2.3 GHz is as follows: video sequences decoding, 1 minute; obtaining sharp frames, 15 minutes; image registration and plane alignment, 15 minutes; image-tiling, 5 minutes – in total 36 minutes. As reconstruction of a single BTF pixel from the slices takes 0.2 seconds, reconstruction of a BTF tile of the size of 128×128 pixels takes under one hour using a single core or 17 minutes using four cores. To summarize, our method allows to measure a material and reconstruct its BTF in one hour.

6.3 Results

In this section, we thoroughly test our method and setup on a series of both synthetic and practical experiments. In the synthetic experiments, we use various datasets and simulate measurement by taking samples in only those directions that are reachable by the proposed setup. In the practical experiments, we perform actual measurements of real materials. In both cases, we compare data reconstructed by the proposed method with the reference data and report the results.

First, we perform synthetic tests on the isotropic BRDFs and the anisotropic ABRDFs and then, we do practical measurement of the anisotropic ABRDFs. Also, we compare look-and-feel of renderings of the measured materials with actual photographs of the materials. Finally, we perform both synthetic and practical experiments on the BTFs. Results of the latter are evaluated by two psychophysical experiments.

6.3.1 Synthetic Test on the Isotropic BRDFs

In the first experiment, we test our method on reconstruction of 55 isotropic BRDFs from the MERL BRDF Database [MPBM03a] resampled to 81×81 directions [SSK03]. The advantage of the isotropic reconstruction is that only four diagonal slices have to be obtained. In our case, we collected 84 samples in total for each material (12 samples in the subspace where $\theta_i = 30°, \theta_v = 30°$ and 24 samples in each other). We leave the experiment without figure and denote only the mean reconstruction errors of all 55 BRDFs which are: CIE ΔE = 9.1, RMSE = 15.7, and PSNR = 24.9 dB.

6.3.2 Synthetic Test on the Anisotropic ABRDFs

In the second experiment, ten BTF samples (nine from the BTF database of the University of Bonn[1] and one from the Volumetric Surface Texture Database[2]) are used (aluminium, corduroy, dark and light fabrics, dark and light leatherettes, lacquered wood, knitted wool, upholstery fabric Proposte, and Lego). These materials represent a challenging dataset to test the proposed method. Due to their rough structure and often translucent properties, they exhibit anisotropic effects of occlusions, masking, and subsurface scattering.

[1] http://btf.cs.uni-bonn.de/

[2] http://vision.ucsd.edu/kriegman-grp/research/vst/

All BTF pixels were averaged to obtain the average ABRDF of the material (first row of Fig. 6.13). The results of the complete reconstruction of the original ABRDFs from 168 samples (12 + 12 samples in the subspace where $\theta_i = 30°, \theta_v = 30°$ and 24 + 24 samples in each other) are shown in Figure 6.13. Together with difference images (10× scaled) and reconstruction errors in terms of CIE ΔE, RMSE, PSNR, these results show that even a very sparse set of measured values can provide promising reconstruction of such challenging anisotropic datasets.

Although on average both interpolation approaches performed similarly, the difference images in Figure 6.13 show that the global interpolation method A estimates incorrect values for elevations between the two sampled elevations (30^o and 75^o). On the other hand, the method B gives, due to a lack of global knowledge, the worst estimation for extrapolated elevations, i.e., the elevations smaller than 30^o. Finally, the displacement interpolation (the method BD) scores similarly as the method A.

To validate contribution of our approach, we compare its performance with the uniform sampling of a similar count of samples. For this purpose, hemispheres of illumination and viewing directions are sampled nearly-uniformly by means of 13 × 13 samples, producing a total of 169 samples. Then, the missing values in the ABRDF space are interpolated from these samples by means of the four-dimensional RBFs [PTVF92]. The interpolation is computed separately in each color channel, and $0 \approx 2\pi$ discontinuity is avoided by usage of the onion-like parameterization of illumination and viewing directions [HFM10].

A comparison of ten reconstructed ABRDFs in Figure 6.13 shows that the proposed reconstruction method has better performance than the interpolation from the uniform samples, mainly near the specular highlights, as confirmed by the objective criterion values shown below the reconstructions. On average, the proposed reconstruction provides 1.4 lower CIE ΔE values, 3.2 lower RMSE values and 1.8 dB higher PSNR values across the ten tested ABRDFs. Moreover, the data acquisition process using our method is considerably faster and less demanding on hardware.

6.3.3 Practical Test on the Anisotropic ABRDFs

Three anisotropic fabric materials (as shown in Fig. 6.14) are used to perform the practical test. Data for all eight slices are obtained by the setup and the data processing described in Section 6.2 and reconstruction of the ABRDF space described in Section 6.1 is performed. Figure 6.15 compares the reference ABRDF measurements of the materials with their reconstruction from the 168 *reference* samples using the method B and with the reconstruction using 192 measured samples (24 + 24 samples in each subspace) obtained by the proposed setup and interpolated by the method B. Figure 6.15 also compares our method with a uniform measurement using 196 *reference* samples (14×14 samples nearly-uniformly distributed over the hemispheres), while the remaining directions are interpolated by the RBFs. Note that, while the visual performance of the uniform measurement might look similarly, the complexity of its measurement is considerably higher in comparison with the proposed measurement approach.

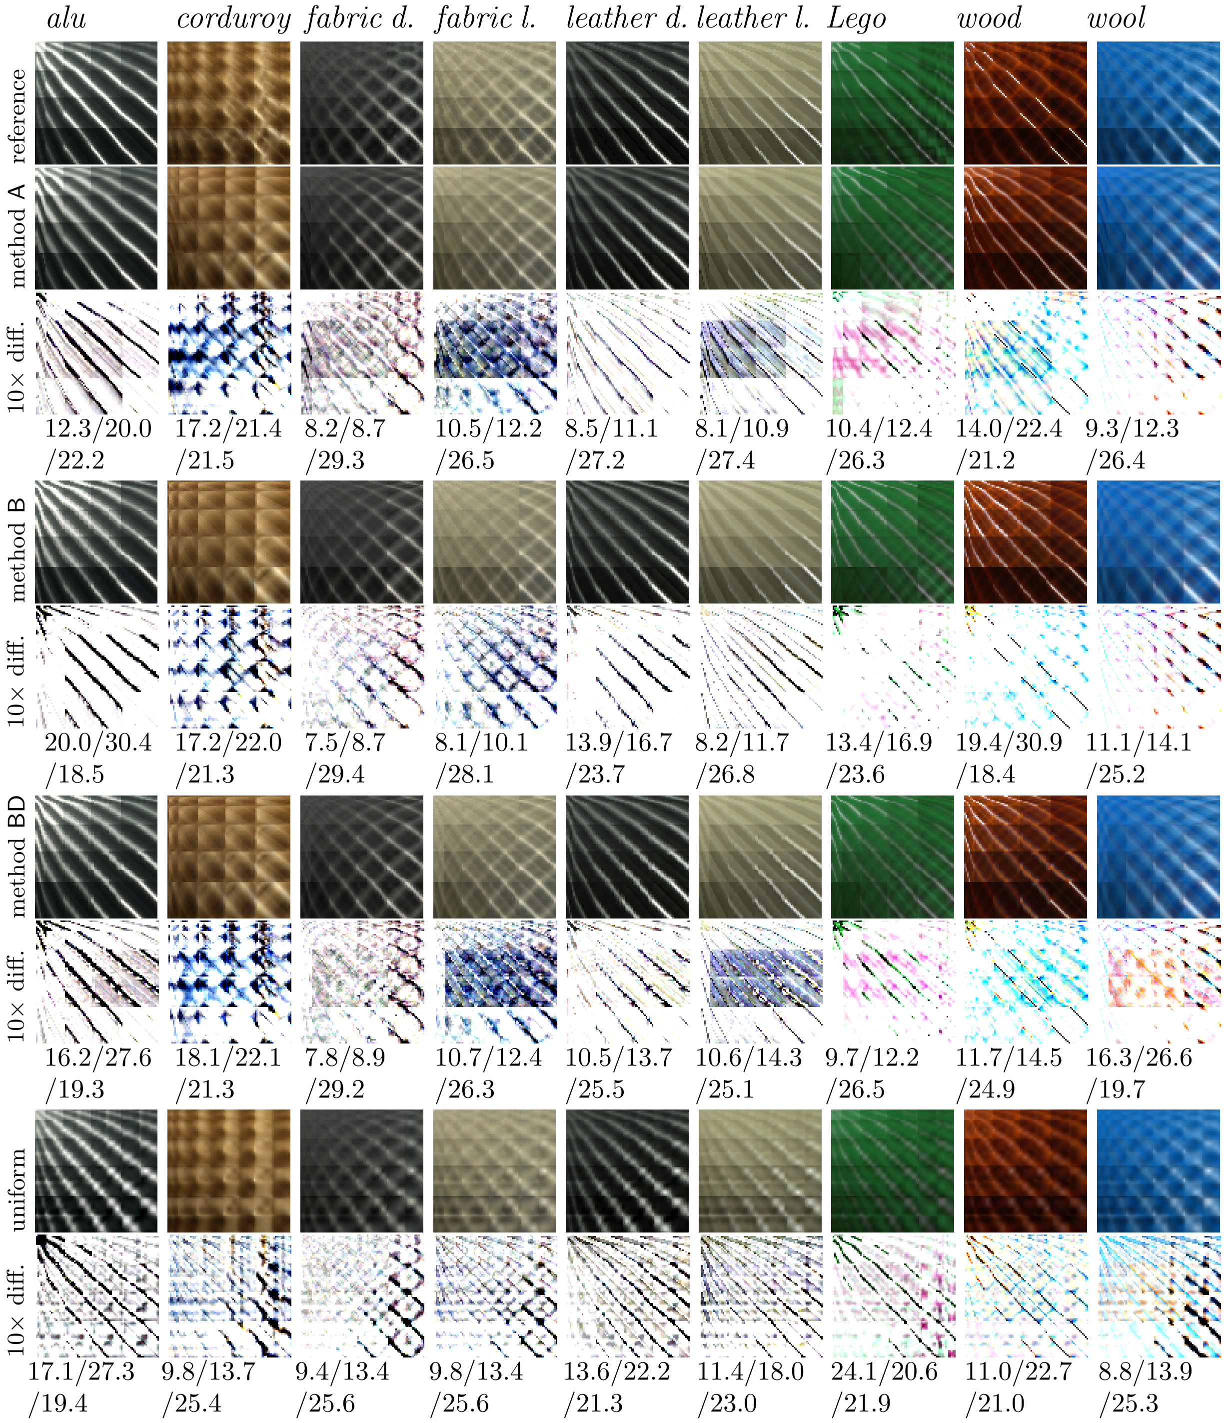

Figure 6.13: A comparison of the reference BRDFs and their reconstruction from the 168 samples by the method A, the method B, the method BD, and from the 169 uniformly-distributed samples, respectively. Below are the 10× scaled difference images and the difference values in CIE ΔE / RMSE / PSNR [dB].

Figure 6.14: Three anisotropic fabric materials whose ABRDFs were measured using the proposed setup.

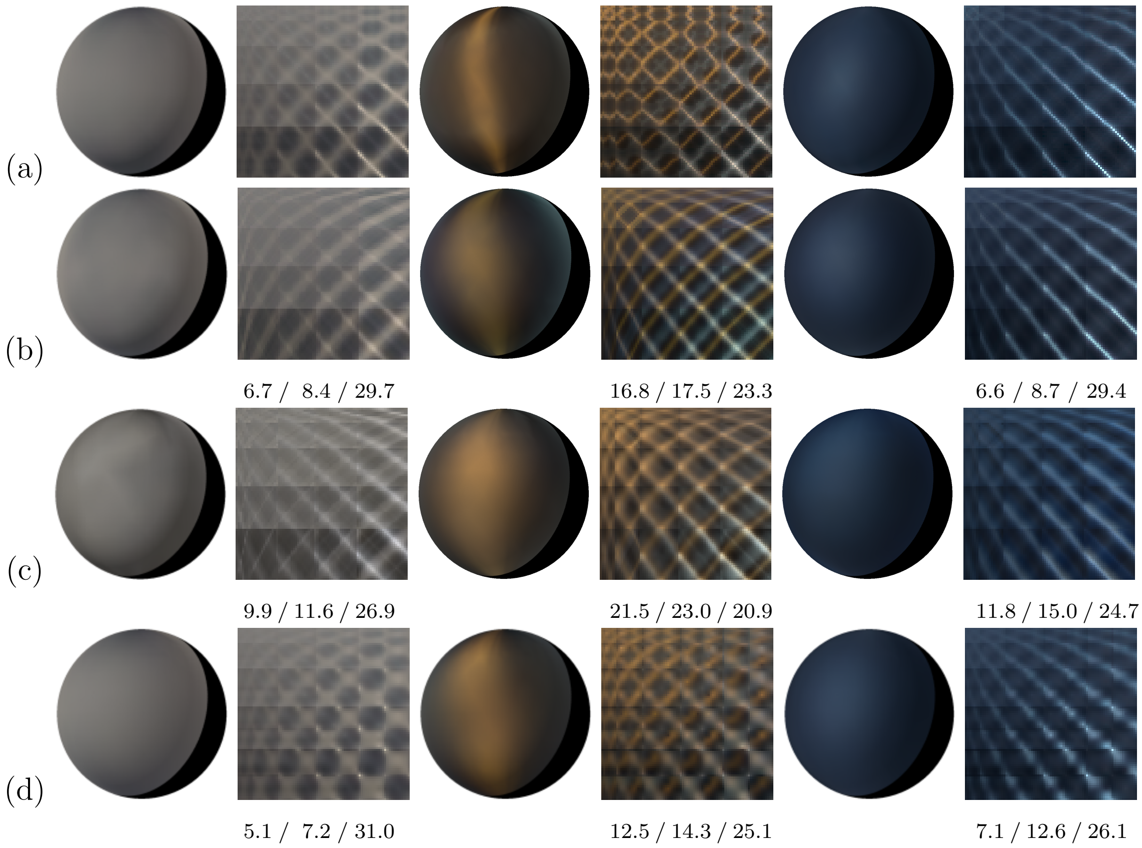

Figure 6.15: The reference measurement of the ABRDF (a), compared to its reconstruction from the 168 reference samples (b), from the 192 samples measured by the proposed setup (c), and from the 196 uniformly-distributed samples (d). Below are the difference values in CIE ΔE / RMSE / PSNR [dB].

6.3.4 Practical Evaluation on Cylinders

Figure 6.16 compares photographs of *fabric02* and *fabric03* attached to a cylinder with renderings on a cylinder using the reference ABRDFs and the ABRDFs obtained by measurement of the 192 samples by the proposed setup. The results for several illumination

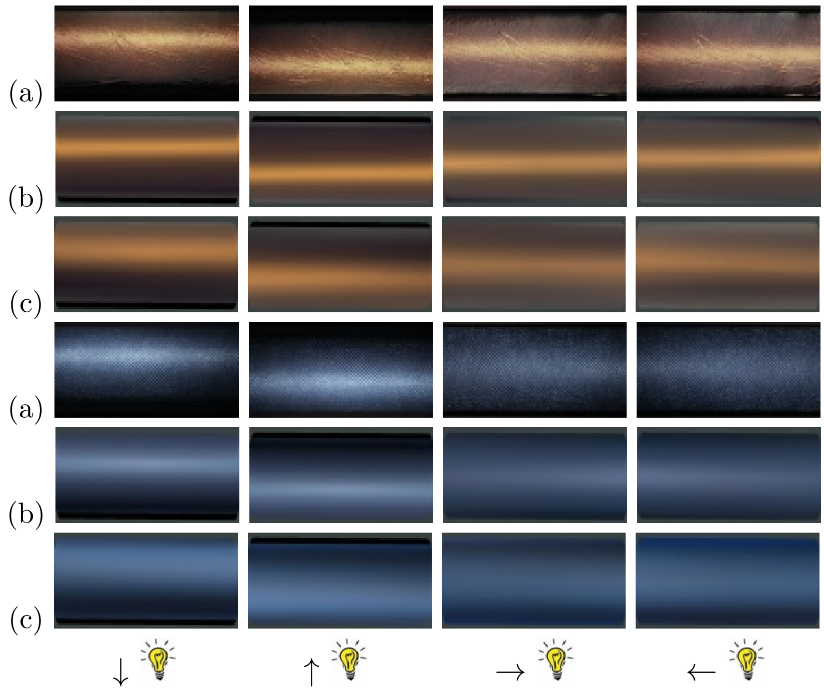

Figure 6.16: Photographs of *fabric02* and *fabric03* on a cylinder illuminated from top, bottom, left, and right (a) compared with renderings using the reference ABRDFs (b), and the ABRDFs measured and reconstructed by the proposed method (c).

conditions confirm that even the proposed setup for approximative measurement can record ABRDFs with reasonable accuracy in comparison with the reference measurements.

6.3.5 Synthetic Test on the BTFs

As acquisition and reconstruction of spatially-varying datasets is a straightforward extension of the proposed method, we test performance of the method on ten BTFs of a resolution of $81 \times 81 = 6,561$ images as described in Section 6.3.2. Only the 168 images corresponding to the eight slices are selected from the BTFs and used for the pixel-wise reconstruction of the remaining images using the proposed method.

Renderings of the original data together with results of the proposed method for the point-light illumination are shown side-by-side in Figure 6.17. From the results, it is apparent that for materials with lower height variations, there is a close match to the original data. The apparent deviations from the original data for materials having higher height variations (e.g., *corduroy* and *Lego* materials shown in Fig. 6.18) are caused mainly by incorrect preservation of a geometry of structural elements due to the very sparse measurement of the azimuthal space as well as the measurement of four elevations only. While the former produces a geometrical deformation of the structure, the latter causes its blur and an improper extrapolation of highlights for low elevations.

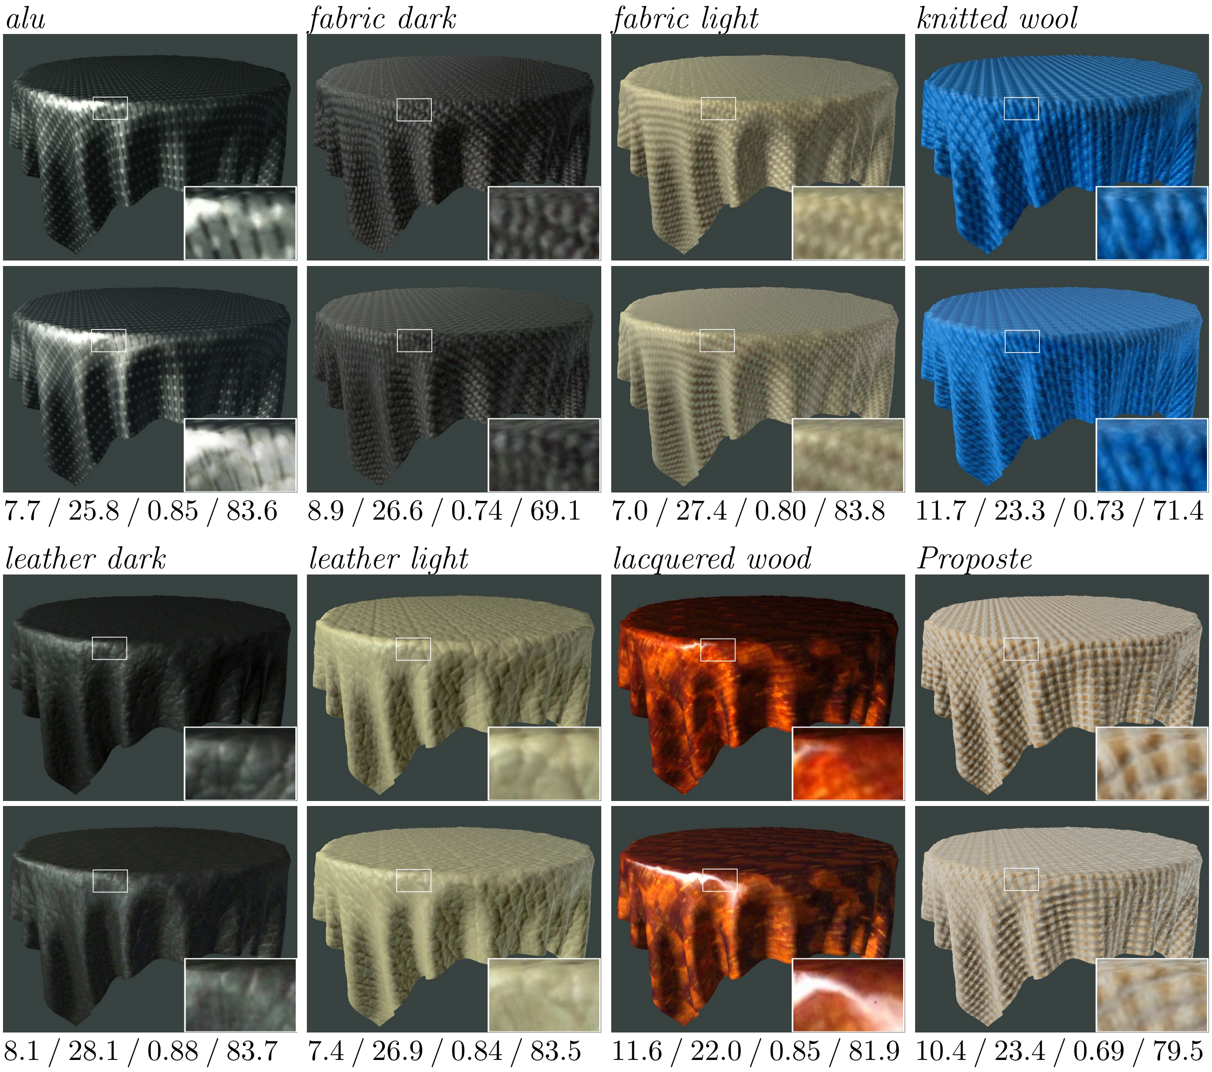

Figure 6.17: A comparison of the BTF rendering from the full dataset of 6,561 images (the first row), with its reconstruction from the 168 images only (the second row) in the single point-light illumination. Below are the difference values in CIE ΔE / PSNR [dB] / SSIM / VDP2.

Even though the reconstruction from a small number of samples is not accurate in terms of correct shading of structural elements for some materials, the method correctly captures the look-and-feel of spatially-varying appearance for nearly-flat materials, e.g., for *fabric dark*, *fabric light*, and *leather light*. However, in comparison with approaches to measure and represent the SVBRDF, the proposed method is not limited to restrictions imposed by the BRDF itself. Therefore, it may be found useful for quick, low-cost, and fairly accurate acquisition and reconstruction of the BTF of many materials having a limited height variation, e.g., fabric and leather.

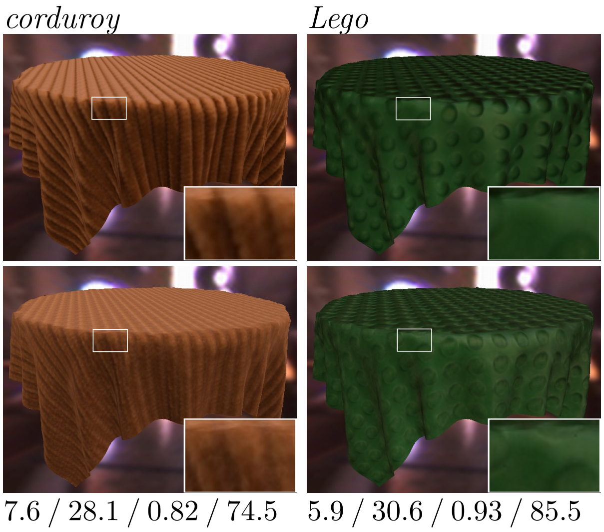

Figure 6.18: A failure case of the proposed reconstruction for rough materials: *corduroy* and *Lego*. Below are the difference values in CIE ΔE / PSNR [dB] / SSIM / VDP2.

6.3.6 Practical Test on the BTFs

We use seven materials to test performance of our method on the practical measurements of the BTFs. We selected different types of materials as illustrated in Figure 6.19: non-woven fabric (*fabric03*), upholstery with a height profile (*fabric04*), woven fabric (*fabric38*), corduroy-like upholstery (*fabric78*), artificial leather (*leather01*), raw wood (*wood01*), and rough sandpaper (*sandpaper01*); the most of them highly anisotropic as is clear from their spatially averaged ABRDFs shown in the second row.

As the reference measurements acquired by the UTIA gonioreflectometer (Sec. 4.1) have

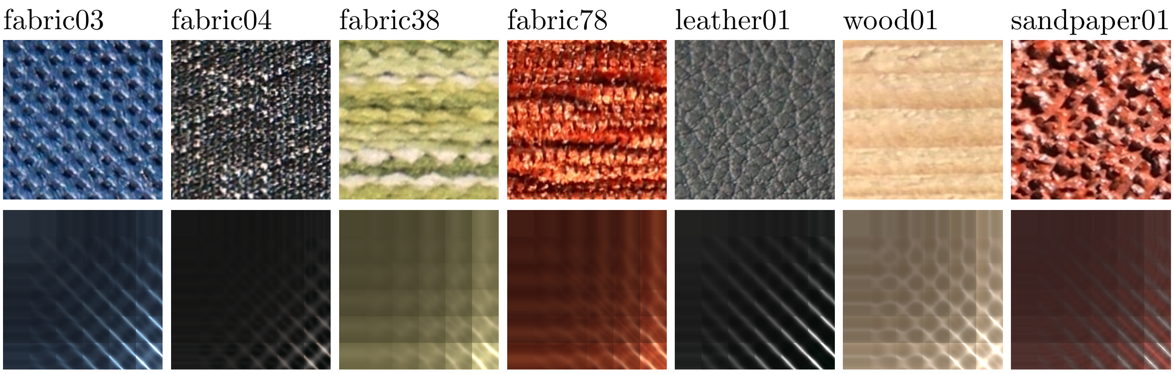

Figure 6.19: The measured materials on an area of 15 × 15 mm (the first row) and their average ABRDFs (the second row).

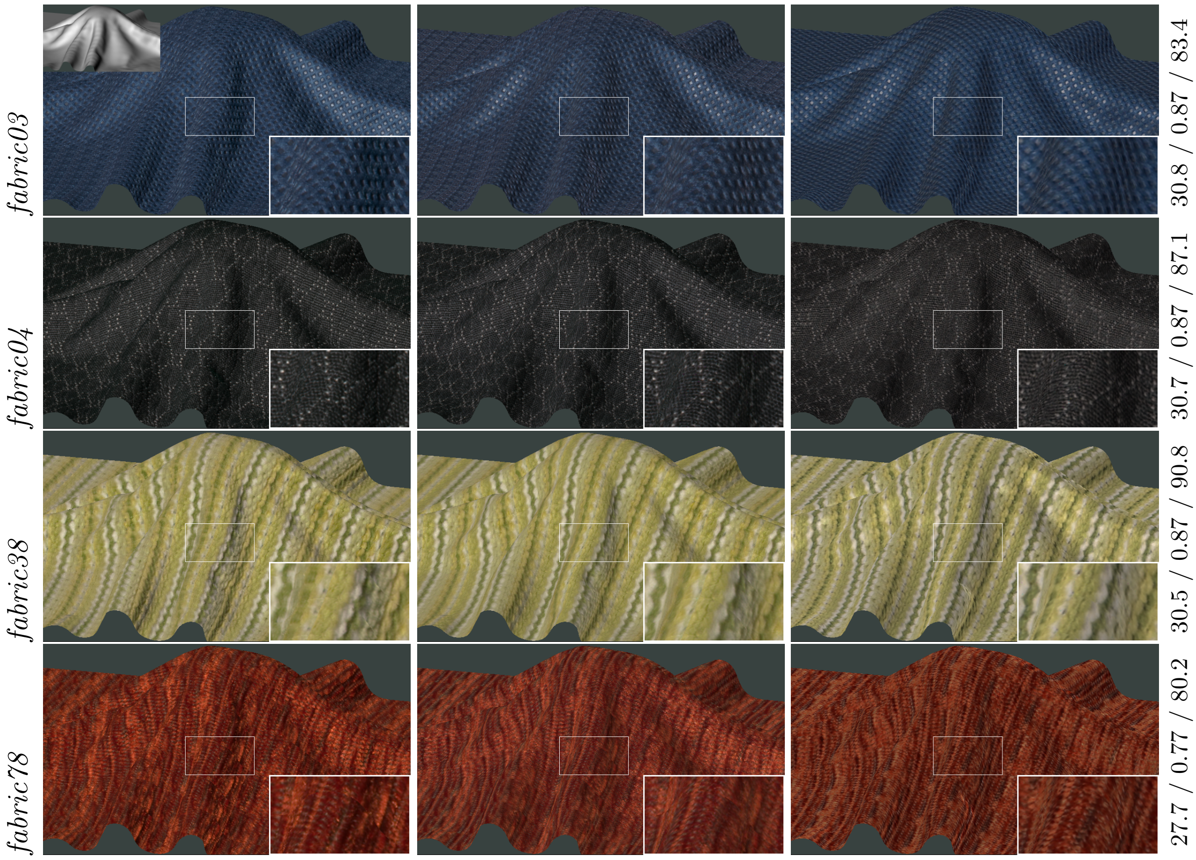

Figure 6.20: A comparison of renderings: (left) using the entire reference BTF datasets (6, 561 images), (middle) using the reconstructed BTFs from the 208 *reference* images, (right) using the reconstructed BTFs from the 208 *measured* images. At the end of each row are PSNR [dB] / SSIM / VDP2 values between images in the first two columns.

a higher resolution (353 or 1071 DPI) than the data captured by the Panasonic Lumix camera (340 DPI), we down-sampled the reference data to match the lower resolution. We also attempt to cut similar tiles [SH05] from both the reference data and the captured BTF datasets to achieve a fair comparison of our measurements with the reference BTF data.

A comparison of the reference BTF data with the BTF datasets measured by our setup and reconstructed by the proposed method is shown in Figure 6.20 (fabrics) and Figure 6.21 (leather, wood, sandpaper). For the purpose of distinction between differences introduced by the reconstruction procedure and those resulting from the proposed acquisition technique, the reference measurements (the left column) are compared with two types of results. The first one (the middle column) reconstructs the BTF from the subset of *reference* measurements (208 images, i.e., 24 + 28 samples in each subspace) while the second one (the right column) reconstructs the BTF from the same set of 208 images; however, measured by the proposed setup.

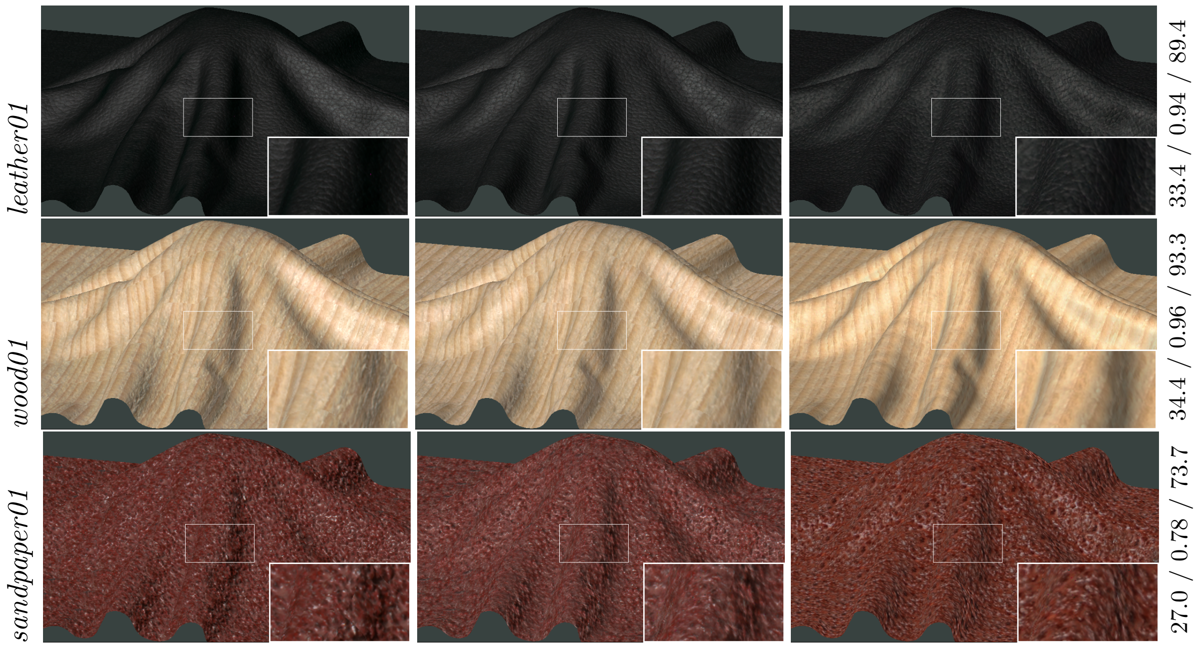

Figure 6.21: A comparison of renderings: (left) using the entire reference BTF datasets (6,561 images), (middle) using the reconstructed BTFs from the 208 *reference* images, (right) using the reconstructed BTFs from the 208 *measured* images. At the end of each row are PSNR [dB] / SSIM / VDP2 values between images in the first two columns.

These images show a close resemblance of the both results to the reference data. The main difference between them is in color hue caused by differences in the acquisition processes as discussed in Section 6.4. Smoother appearance of the material *wood01* results from variations in the amount of sanding that has been performed on the surface of the raw wood. Note that, although we use the same materials as the reference measurement system, we cannot achieve a pixel-wise alignment between the reference and the proposed measurements. Therefore, the accuracy of the proposed method cannot be assessed by any of the standard pixel-wise quality evaluation metrics. This type of comparison is possible only between the data in the first two columns of Figures 6.20 and 6.21.

Therefore, to objectively evaluate our method, we ran two psychophysical experiments. In the first one, we evaluate a visual error introduced by means of measurement of a small subset of directions compared to the complete BTF dataset. The second experiment compares performance of the reference and proposed setups recording the same small set of images; in other words, a visual quality trade-off occurred by using the proposed instead of the reference setup.

Experiment 1 – Eight naive subjects compared differences between the reference and reconstructed BTFs. Subjects were shown two rendered images on a calibrated screen together with actual specimens of materials as shown in Figure 6.19. The question was:

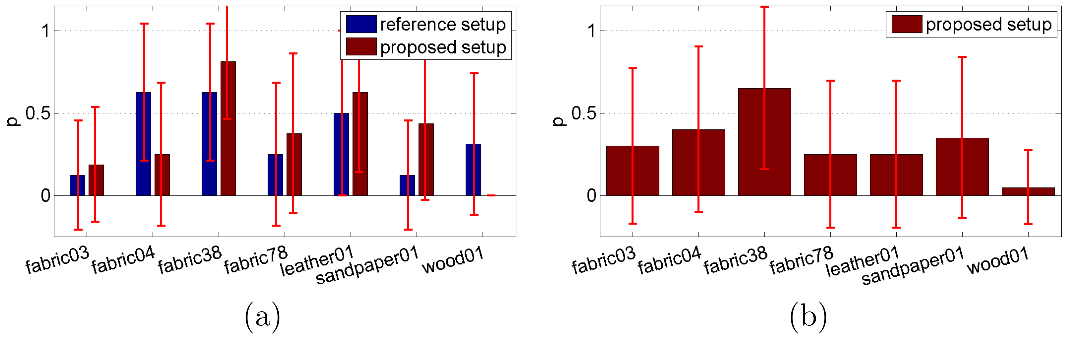

Figure 6.22: Results of psychophysical experiments: (a) *Experiment 1* – probability p that the image rendered using the BTF dataset reconstructed from the 208 images captured by the *reference* setup (blue) or the proposed setup (red) looks more realistic than the image rendered using the whole BTF dataset, (b) *Experiment 2* – probability p that the image rendered using the BTF dataset reconstructed from the 208 images captured by the proposed setup is more similar to the image rendered using the whole BTF dataset than the image rendered using the BTF dataset reconstructed from the 208 images captured by the *reference* setup.

Which of the two images looks more realistic? When we compare the renderings that use the complete dataset with its subset that consists of the 208 images (the first *vs.* the second column in Figures 6.20 and 6.21), on average 37% of our subjects prefer those that use the subset. When we compare renderings that use the complete dataset with our measurements (the first *vs.* the third column in Figures 6.20 and 6.21), 38% of our subjects prefer those from the proposed setup. The results for individual materials are shown in Figure 6.22-a. Interestingly, for materials *fabric38*, *fabric78*, and *sandpaper01* the proposed setup scores much better than the reference one. Note that the difference for material *wood01* is a result of a slightly different specimen of the same material, with a smoother surface finishing.

Experiment 2 – To determine an error resulting from the proposed simplified acquisition technique when compared to the reference measurement of the 208 images, we ran another perceptual study with 20 subjects. For each tested material, they were shown the reference rendering using the complete measurement of 6,561 BTF images and were asked to choose between the two renderings using 208 images only (obtained using the proposed and reference setups). An example stimulus image is shown in Figure 6.23; subjects were asked: *Select the image more similar to the reference image above.* We used a website-based interface, and image couples were ordered randomly. Subjects were encouraged to consider an overall appearance instead of focusing on local details. To avoid the latter, we reduced the size of the reference image. Analysis of responses revealed that on average, only 32% of subjects considered the results from the proposed setup more similar to the reference data (see Fig. 6.22-b), while 68% of subjects preferred data from the reference setup. When material *wood01* is excluded (due to differences in specimens as explained above) the ratio becomes 38% : 62%. The lower preference of our setup is probably caused by the similarity of the reference dataset to its subset (both sharing the same acquisition

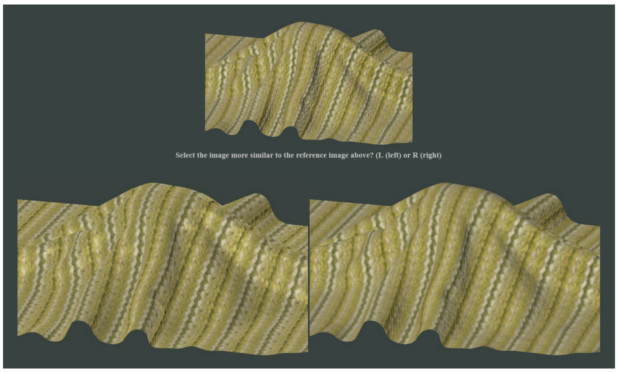

Figure 6.23: An example of a stimulus image from the second psychophysical experiment.

process, i.e., color hue, texture tiles). Even so, the result is encouraging, especially when considering the relatively high standard deviation of the responses of the subjects shown as errorbars in Figure 6.22-b.

In summation, although results invariably depend on the measured material, we found that the average subject is able to recognize differences between the complete dataset and its subset as well as between the data from the reference and proposed setups. However, high standard deviations in both experiments suggest that feedback significantly depends on personal preference rather than unequivocal visual differences resulting from usage of the simplified setup. Therefore, we believe that even though the reconstruction from a small subset of images is not always physically accurate (in terms of proper shading of structural elements), the performed perceptual study confirms that our setup captures the overall look-and-feel of a given material.

6.4 Discussion

This section discusses advantages and limitations of the proposed method, setup, and data processing procedures. Then, we analyze a possible improvement of the setup in terms of usage of more slices per subspace or measurement of additional elevations.

6.4.1 Advantages

A notable advantage of our setup is its ability to quickly measure almost any flat or slightly rough material without necessity to extract it from its environment. Compared to approaches to measure and represent the SVBRDF, the proposed method is not limited by the restrictions imposed by the BRDF properties (reciprocity, energy conservation). Therefore, it can be found useful for fast and inexpensive approximate acquisition of the BTF of many materials. As the measurement can be done in 6 minutes and the entire BTF can be reconstructed in one hour, our method can be used for fast and inexpensive measurements of such materials as human skin (see Appx. E), precious planar cultural heritage artifacts, e.g., coins, engravings, fabrics, etc.

6.4.2 Limitations of the Proposed Method

The limitations of the proposed method are threefold. First, the method restores the reflectance at given elevations only from two orthogonal slices that measure only a small subset of the azimuthal subspace and as such, can omit some reflectance features (see *corduroy* in Fig. 6.7), or results in a slightly different appearance of color or brightness of the reconstructed data. To avoid this problem, the azimuthal subspace can be measured by additional slices at the cost of longer acquisition times (see Sec. 6.4.4).

Second, the interpolation method expects monotonicity of reflectance values across different illumination and viewing elevations. However, this condition is rarely invalid and no such behavior was experienced with any of the tested materials. The accuracy of the method can be further improved in this respect by taking additional samples at different elevations (see Sec. 6.4.4).

Finally, highlights of extremely specular materials are not always represented accurately enough (see Fig. 6.17) mainly due to an insufficient sampling of azimuthal angles (step 15^o) in original datasets used in the synthetic experiments. Note that density of the sampling across the specular highlights in the diagonal slices can be arbitrarily increased to provide a better results without prolongation of the measurement.

Note that the proposed method does not fit any analytical model to the measured data and as such it is sensitive to noise in the measurement process. However, since the measurement procedure is fast and simple, this noise can be effectively suppressed by repeated measurement of the slices and by computation of median values.

6.4.3 Limitations of the Proposed Setup and Data Processing

Although there is not any restriction imposed on the measured materials when compared with other setups capturing the BTF datasets, the results show the following general limitations of the proposed setup and data processing approach:

- *Lower sharpness of structural elements* – results from a geometrical deformation of features of a structure, which is caused by mechanical vibrations during the measurement. Another reason for the lower contrast is a low dynamic range of the sensor, where certain details are lost after the exposure compensation (e.g., white dots in *fabric04*); therefore, a sensor with a dynamic range over 8 bits/channel would help.

- *Color hue differences* – are caused by a different dynamic range and spectral response of the RGB sensors used for the reference and our acquisition of the data, and by different calibration targets used. Another source of these differences can be slight color variations across the specimen (e.g., *sandpaper01*).

- *Visible repeatable seams* – are caused by tiling with the aid of only a single BTF tile and by less than ideal compensation of non-uniformity of the illumination, and are apparent for materials represented by a very large tile.

- *A limited size of a material* – A larger size is not a severe limitation of our setup. We can scan larger areas sequentially due to a high speed of the measurement; and this approach does not compromise portability of the setup. Alternatively, the setup can be built in a larger size for only minor additional costs for a price of lower DPI.

6.4.4 An Analysis of Placement of the Slices

To achieve the best possible reconstruction of the ABRDFs, we analyze influence of: (a) selection of elevations at which the slices are measured, (b) the number of used elevations, and (c) the number of slices per subspace. In our analysis, we use four ABRDFs of highly anisotropic materials from the UTIA BTF Database (Appx. A). Each material was measured as the BTF with ten elevations (7.5° apart) and 24 × 24 azimuthal samples per subspace. Finally, all pixel values were averaged to obtain the ABRDFs (see Fig. 6.25-top).

First, we test six different combinations of two elevations to find the optimal configuration (see Fig. 6.24). Only samples from the resulting four subspaces at given elevations are used for reconstruction of the entire ABRDF. The average RMSE differences between the ground-truth ABRDF data and the interpolated data suggest that the combination of 45°/75° provides the lowest reconstruction error. Therefore, we have chosen the highest elevation angles $\theta_i = \theta_v = 75°$, where the specular reflections are the most intensive (see the first row of Fig. 6.13). The lower elevation angles were decreased to $\theta_i = \theta_v = 30°$ (the second best choice from Fig. 6.24) for better representation of appearance of a material at orthogonal viewing and illumination directions, which are the most visually salient.

Further, we analyze the benefit of additional elevations. Specifically, we use the following four configurations: two elevations (30° and 75°, used in this chapter), three elevations (30°, 52.5°, 75°), four elevations (30°, 45°, 60°, 75°), and five elevations (15°, 30°, 45°, 60°,

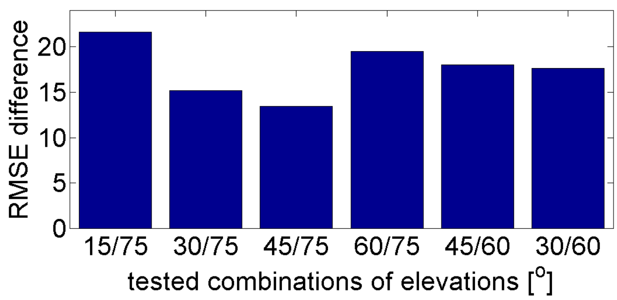

Figure 6.24: A RMSE of reconstruction of the ABRDF for different combinations of two elevations used for selection of four measured subspaces.

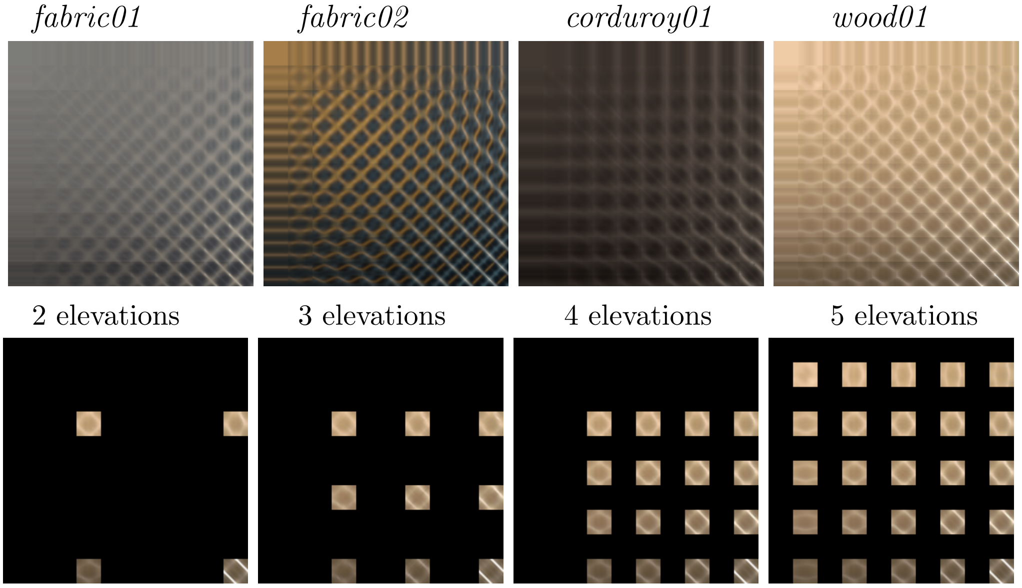

Figure 6.25: Tested BRDFs of fabrics and wood (top), and examples of the sampled subspaces for four tested configurations (bottom).

75°) as shown in Figure 6.25-bottom. We evaluate the benefit of the additional elevations together with the benefit of additional slices per subspace as described below.

Reconstruction of a subspace can be performed either from only two slices by means of the method proposed in Section 6.1 or from an arbitrary number of slices by means of more general interpolation approach (Sec. 5.3). Our analysis reveals that the main quality bottleneck resides in reconstruction of a subspace by two slices only. Therefore, significant improvement can be achieved by increase of the number of measured slices.

Figure 6.26-a shows PSNR and SSIM errors of the reconstruction method (Sec. 5.3) averaged across the tested ABRDFs for the linearly increasing number of slices. The dashed lines represent the upper error bounds of reconstruction (using all 24×24 samples per subspace). Already 4 axial and 4 diagonal (4 + 4) slices uniformly distributed across a subspace deliver a reasonably good reconstruction compared to mere 1 + 1 slices (the reconstruction using the method proposed in Section 6.1 is shown as crosses). The reconstruction error also depends on the number of sampled elevations (represented by charts of different colors), where the quality gain starts to be significant when a subspace is sampled by more than 3 + 3 slices. The main quality contribution is gained using additional middle elevation between 30° and 75° and first of all, additional low elevation 15°. We also found that usage of the five elevations (15°, 30°, 45°, 60°, 75°) instead of only four elevations (15°, 30°, 52.5°, 75°; not shown in graphs) provides improvement lower than 0.5%.

Graphs in Figure 6.26-b mark almost a linear increase of the count of samples and time with the number of slices recorded per subspace. For instance, four cameras can record material appearance for each light source in parallel and the measurement time for 8 slices per subspace (e.g., 4 + 4, 2 + 6) can be around 7 minutes. Note that this time includes

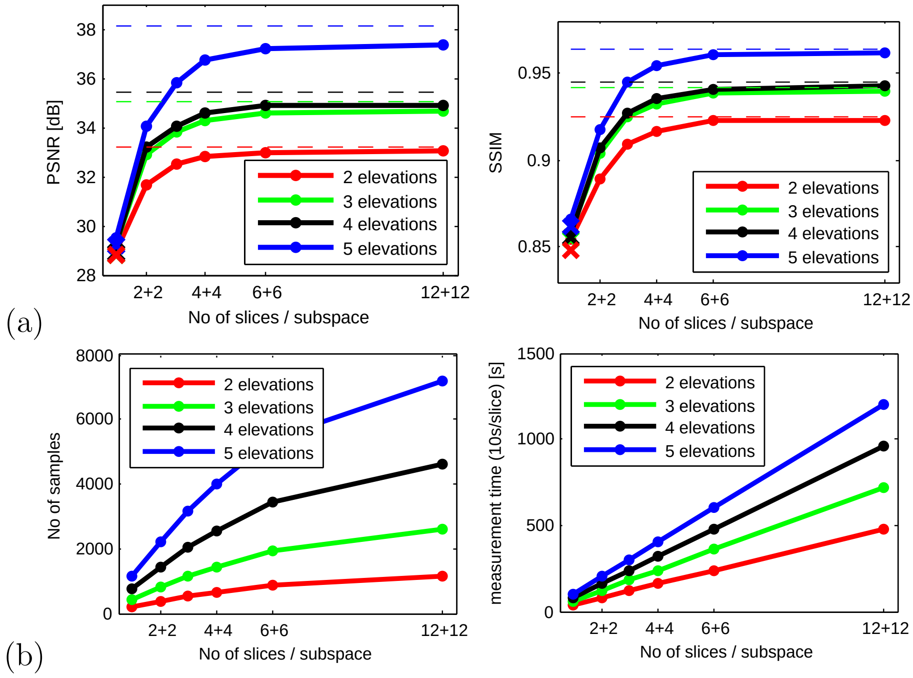

Figure 6.26: (a) Errors of reconstruction of the entire BRDF (left: PSNR, right: SSIM), (b) the total number of samples (left) and the measurement time (right) as functions of the number of sampled elevations and slices per subspace.

overhead of positioning of the arms and assumes measurement time of 10 s per slice, which we believe can be achieved with a steadier gantry and sufficient intensity of the illumination (see [A.9]).

Finally, we theoretically analyze the configuration of our setup comparable with other BTF system [MBK05] in terms of the measurement time *vs.* the number of taken images. The other setup is one of the fastest setups and collects and processes $151 \times 151 = 22,801$ images within 120 minutes. Our setup that would use 10 cameras and 10 light sources could measure 1,600 slices within 30 minutes assuming time of 10 s per slice and process them in the remaining 90 minutes. Total of the 1,600 slices at ten viewing and ten illumination elevations result in 16 slices per subspace, e.g., a configuration of $8 + 8$ slices. The azimuthal density is restricted only by the video frame rate, e.g., for a step of 15° it is 24 samples per slice, and we obtain 38,400 images. When we compare the reference ABRDFs measured at a resolution of 151×151 directions with the simulated measurement and reconstruction using the proposed configuration, we obtain PSNR 50.1 dB (averaged across all four materials). Similarly for $4 + 4$ slices per subspace (19,199 images) we obtain PSNR 44.2 dB and for $4 + 6$ slices per subspace (24,209 images) we obtain PSNR 46.8 dB. Although a setup in this configuration would still be less expensive than the setup with 151 cameras [MBK05], we consider it reasonable to limit the number of cameras and slices, allowing a much higher speed of the BTF acquisition while maintaining the required visual quality.

6.5 Summary

We present a fast and inexpensive setup for acquisition of material appearance in the form of the approximate ABRDF or BTF. The proposed acquisition setup is based purely on consumer hardware and is easy to build. Acquisition and subsequent fully automatic reconstruction of the entire dataset is fast and computationally non-intensive. The measurement process records appearance of a material using eight video sequences, from which only 208 frames are taken to approximate the entire ABRDF or BTF of the measured material. Although different materials require different numbers or distributions of samples, we fixed the number of samples in our experiments due to the diffuse character of the tested materials.

Performance of the method was thoroughly tested on the isotropic BRDFs, anisotropic ABRDFs, and BTFs with encouraging results. Our experiments show that retrieval of the samples and the consequent reconstruction of the complete dataset take about one hour. Reconstruction of the BTF datasets show that the method can be a reasonably accurate alternative to a lengthy measurement of the entire datasets, especially for materials having a smaller height variation. Moreover, we demonstrate flexibility of the acquisition process on examples of human skin and measurement of dynamic processes (see Appx. E).

We show that accuracy of the measurement can be significantly increased by addition of new cameras and light sources together with sampling by more slices per subspace. Our analysis reveals that the reasonable trade-off between the visual quality and the measurement time can be achieved by usage of four cameras and four light sources placed at elevations 15°, 30°, 52.5°, 75°, recording at least six slices per subspace. The measurement using this configuration of the setup would take 7 minutes. One also must consider the increased price of such the setup that requires four cameras.

Although the presented setup has certain limitations and does not capture exact detailed appearance for some materials, its speed, simplicity, and portability makes it superior to alternative approaches such as the bump/displacement mapping or the parametric BRDF modeling. This approximate method can be therefore utilized in less accuracy-demanding applications. One option for further improvement would be adaptive selection of samples from the video (along the slices) based on the overall reflectance of a material as shown in Section 5.2. Since digital reproduction of look-and-feel of a material can be created inexpensively, it could be particularly useful in the fields of computer gaming, film and digital presentations of e-commerce. We believe that our setup has, due to its scalability, high application potential for development of simple, inexpensive, and portable systems for the BTF measurement with performance adapted to the requested visual quality or speed.

Chapter 7

Data-Driven Adaptive Measurement

This chapter summarizes contributions of the papers [A.5, A.1] that describe methods for measurement of the anisotropic BRDF which use a database of already measured materials. The methods either use sampling patterns learned on the database when we measure a new material or they even model a measured material by a linear basis computed on the database. Then, the new material can be measured efficiently with only the small number of samples.

First, we describe a method for computation of sampling patterns tailored to materials in a database and two methods that reuse the patterns for measurement (Sec. 7.1). Then is explained a principle of a method that enables effective measurement of anisotropic materials by only few samples taken for several rotations of a material (Sec. 7.2). Finally, Section 7.3 summarizes the chapter.

7.1 Template-Based Measurement

In this section, we describe a method of adaptive sampling proposed in the paper [A.5] that is based on an analysis of the measured BRDFs. We build upon the unique database of 150 BRDFs described in Section A.4, which represents a wide range of materials; the majority exhibiting anisotropic behavior. For each BRDF in the database, a unique sampling pattern is precomputed, given a predefined count of samples. Further, template-based methods are introduced based on reusage of the precomputed sampling patterns for measurement of a new material. We compare the template-based methods with a measurement by the nearly-uniformly distributed samples, with a linear-combination-based method [MPBM03b], with a linear-basis-based method [ASOS12], and with the adaptive BRDF slices (Chap. 5). Finally, we discuss advantages and disadvantages of the proposed methods.

7.1.1 Approximation of the Optimal Sampling

For the sake of comparison of the tested sparse-sampling methods, we look for the reference adaptive sampling of the BRDF for the required number of samples. We build on a heuristic algorithm introduced by Robinson and Ren [RR95] that refines positions of samples based on a cross-validation principle, and we propose Algorithm 7.1 that iteratively adds and shifts samples while minimizing an error between the reconstructed and ground-truth BRDFs. The error of the reconstructed BRDF f_r' with respect to the reference BRDF f_r is evaluated by the modified PSNR in n_s samples as:

$$PSNR = 10 \cdot \log_{10} \left(\frac{n_s}{\sum_{i=1}^{n_s} [f_r(\xi(i)) - f_r'(\xi(i))]^2} \right) ,$$

where ξ is a known mapping function $\xi(i) \to (\theta_i, \varphi_i, \theta_v, \varphi_v)$ that maps the number of a sample to its directions.

Algorithm 7.1 Reference Sampling Estimation (*adaptive*)

1: Select one sample in each subspace of the BRDF so as the value of the sample is the closest to the mean of the subspace.
2: Reconstruct missing BRDF values in each subspace.
3: Select the subspace with the lowest PSNR.
4: Add a new sample to the location in the subspace where the difference between the reconstructed value and the ground-truth value is maximal.
5: Goto step 2 until the requisite number of samples is reached.
6: For each sample (if the sample is not the only one in the subspace):
   Remove the sample and perform steps 2 – 4.
7: Repeat step 6 multiple times.

Note that this approach (denoted as *adaptive*) does not guarantee to find the optimal sampling, and some of the tested sampling techniques show even better performance sometimes; however, we believe that it can serve as a reasonable approximation of the optimal sampling. We use thin plate splines for interpolation of data from the samples. The reconstruction error based on the samples estimated using this algorithm is shown as a blue line in Figure 7.1.

7.1.2 Proposed Template-Based Sampling

Here we estimate the proper sampling of unknown BRDFs given the predefined number of samples and using the template-based sampling method. The method is based on reusage of precomputed sampling patterns (Alg. 7.1) for a more efficient BRDF measurement of an unknown material.

In contrast to the method introduced above, the sampling of an unknown material is more difficult because we do not have a prior knowledge of directional behavior of the material. Therefore, at least minimal information must be collected, prior to selection of a sampling strategy. Alignment of the main axis of anisotropy (see Sec. A.4.2) of an unknown material requires the initial measurement of 48 samples along the axial slice (see Sec. 5.1). This number of samples is then subtracted from the available number of samples. Once anisotropy of a material is aligned, we measure additional 42 samples and look for the most similar material in the database based on the 90 samples. Then, we use its sampling pattern precomputed by Algorithm 7.1 to measure the material. Note that the best results are achieved when the samples are taken at the subspace at $\theta_i = 60°, \theta_v = 60°$; however, usage of another subspace does not significantly impact the performance. The entire sampling process (denoted as *method1*) is described in Algorithm 7.2.

Algorithm 7.2 Template-Based *Method1*

1: Record additional 42 samples adaptively along four diagonal slices in the subspace at $\theta_i = 60°, \theta_v = 60°$.
2: Select the most similar (L1 norm) material in the database based on the 90 samples.
3: Measure the BRDF using the pattern of the selected material.

Alternatively, we test also a variant of the method (denoted as *method2*) that iteratively adds new samples and can change the chosen similar material if there is another more similar one based on the samples collected so far as described in Algorithm 7.3.

Both approaches directly compare brightness values $Y_U \in \mathbb{R}^{n_s \times 1}$ of an unknown measured material with brightness values $Y_D \in \mathbb{R}^{n_s \times 1}$ of a material in the database using L1 norm in n_s locations of already taken samples. The brightness values Y_U, Y_D are obtained as a weighted average of the RGB values as in their conversion to the YUV color space. To remove influence of brightness of the compared materials, we use least squares regression to identify linear parameters a, k for vectors Y_U, Y_D. The query material is then compared with all $N = 150$ BRDFs to find the one most resembling the sparse measurements using:

$$\hat{i} = \underset{i=1 \ldots N}{\arg\min} \|Y_U - (a_i + k_i \cdot Y_{Di})\| \ .$$

Algorithm 7.3 Template-Based *Method2*

1: Measure 16 sparse samples (4×4) per subspace.
2: Select the most similar (L1 norm) material in the database based on already taken samples.
3: Reconstruct the BRDF of the *selected* material using the available samples.
4: Measure new sample at a location where the difference between the reconstructed value and the ground-truth value is maximal.
5: Goto step 2 until the requisite number of samples is taken.

7.1.3 Overview of other approaches

We compare results of the following approaches with results achieved by the proposed methods. The compared techniques can be split into two categories. The first one comprises methods that estimate the fixed sampling pattern valid for any BRDF: the uniform

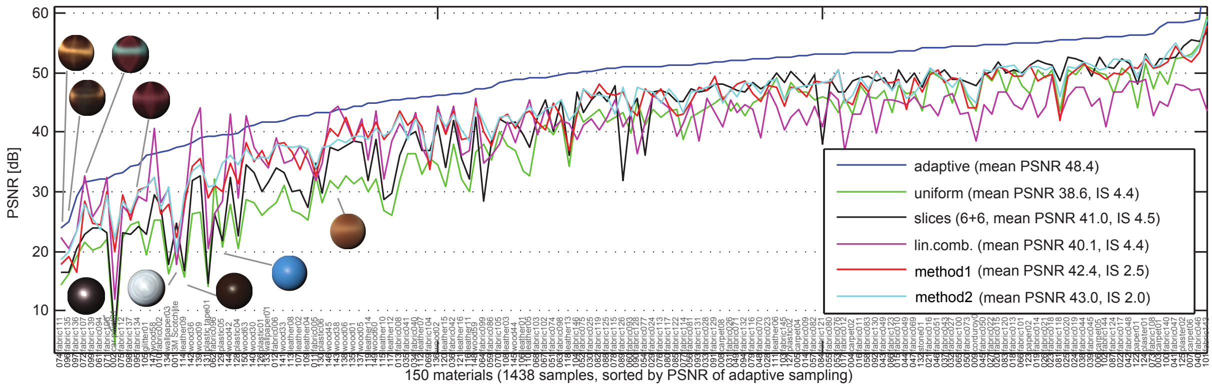

Figure 7.1: Fidelity of reconstruction of all 150 BRDFs using the estimated reference sampling (blue) and the tested approaches for the fixed number of 1438 samples. The legend shows PSNR [dB] and IS values averaged across all BRDFs.

sampling (*uniform*), the linear combination of BRDFs (*lin.comb.*), and the global adaptive sampling (*global*). The second one includes methods that estimate the sampling pattern of each measured material individually: the BRDF slices (*slices*) and the proposed template-based methods (*method1*, *method2*).

The uniform sampling places samples to each BRDF subspace by means of regular matrices with resolution (e.g., 4×4 samples) tuned independently in a way to achieve the lowest PSNR across the entire database for the required count of samples, i.e., each subspace is sampled by the optimal number of samples so that the error across the entire database is minimal for the given total number of samples.

We use Algorithm 7.1 for derivation of the fixed sampling pattern tailored to the entire BRDF database (denoted as *global*), i.e., we use all 150 BRDFs instead of a single BRDF to tune placement of samples. Consequently, both the *uniform* and *global* methods are tuner to our database, but the *global* method has greater freedom in placement of the samples; therefore, we believe that it can achieve better results.

The method based on the linear combination of BRDFs [MPBM03b] and the method based on the linear combination of basis functions derived from BRDFs [ASOS12] provide nearly similar results; therefore, due to lower computational demands, we use the former method. The BRDF slices (Chap. 5) sample each subspace by means of 6 axial and 6 diagonal slices and samples along the slices are taken adaptively.

7.1.4 Results

Here we present results of the proposed methods together with results of other approaches described above. Methods *global*, *method1*, *method2*, and *lin.comb.* allocate the first 48 samples for alignment of anisotropy of the BRDF using a method presented in Section A.4.2. Results of *method1* and *method2* are computed with the query material ex-

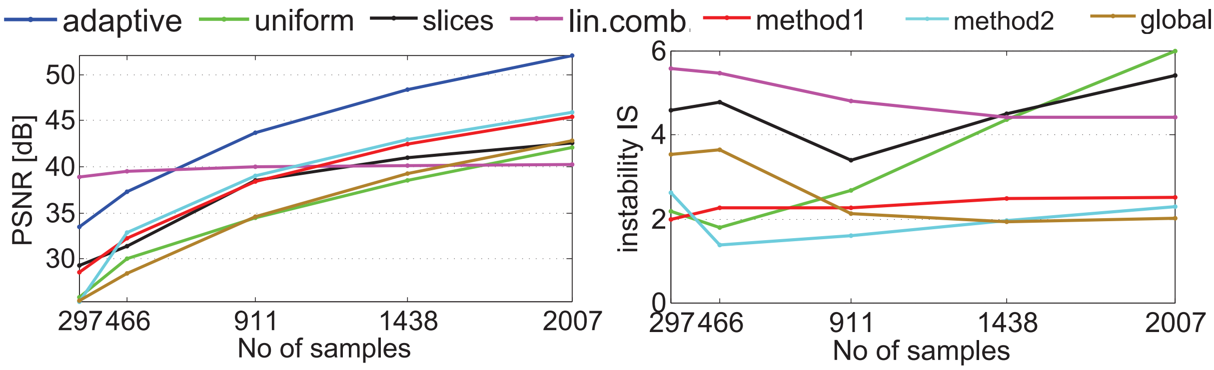

(a) (b)

Figure 7.2: The average PSNR (a) and IS (b) of the compared methods dependent on the number of samples.

cluded from the database. Similarly, the query material should be excluded from computation for methods *lin.comb.* and *global*, but computational times would be too long. Our experiments on several materials has shown that impact of not exclusion of the query material is negligible.

Results of the performed experiments for 1438 samples and all 150 BRDFs are shown in Figure 7.1. The BRDFs are sorted according to the PSNR error of the reference *adaptive* method. From the graph, it is apparent that none of the tested methods is ideal for sparse sampling of all tested BRDFs. As expected, adaptive methods show better performance than fixed-sampling methods. For BRDFs where the reference *adaptive* method obtains the worst PSNR (left side of the graph), the *lin.comb.* method works the best, while for the materials that can be represented in higher PSNR, the *slices* dominate. Although the *lin.comb.* approach shows superior performance for some of the difficult materials, it has the worst stability across the BRDFs and is unable to achieve better results than about 45 dB (regardless the count of samples).

The template-based *method1* and *method2* mark steady and the most stable performance across all materials. Some of the tested methods indicate superior performance for some BRDFs yet perform poorly on some others but still keeping the same reconstruction PSNR as in more stable approaches. Therefore, we compute the instability of performance (IS) as the standard deviation of the difference between the PSNR of the *adaptive* method and of the tested method. A legend of Figure 7.1 shows the PSNR and IS values averaged across all BRDFs.

We also analyze performance of all methods for the variable number of samples. The reference *adaptive* sampling as well as a sampling by the tested methods was computed for 297, 466, 911, 1438 and 2007 samples. The PSNR and IS values averaged across all materials as functions of the number of samples are shown in Figure 7.2.

In Figure 7.2-a, we can observe almost constant performance of the *lin.comb.* method regardless the number of samples. The method is superior for a very small set of samples

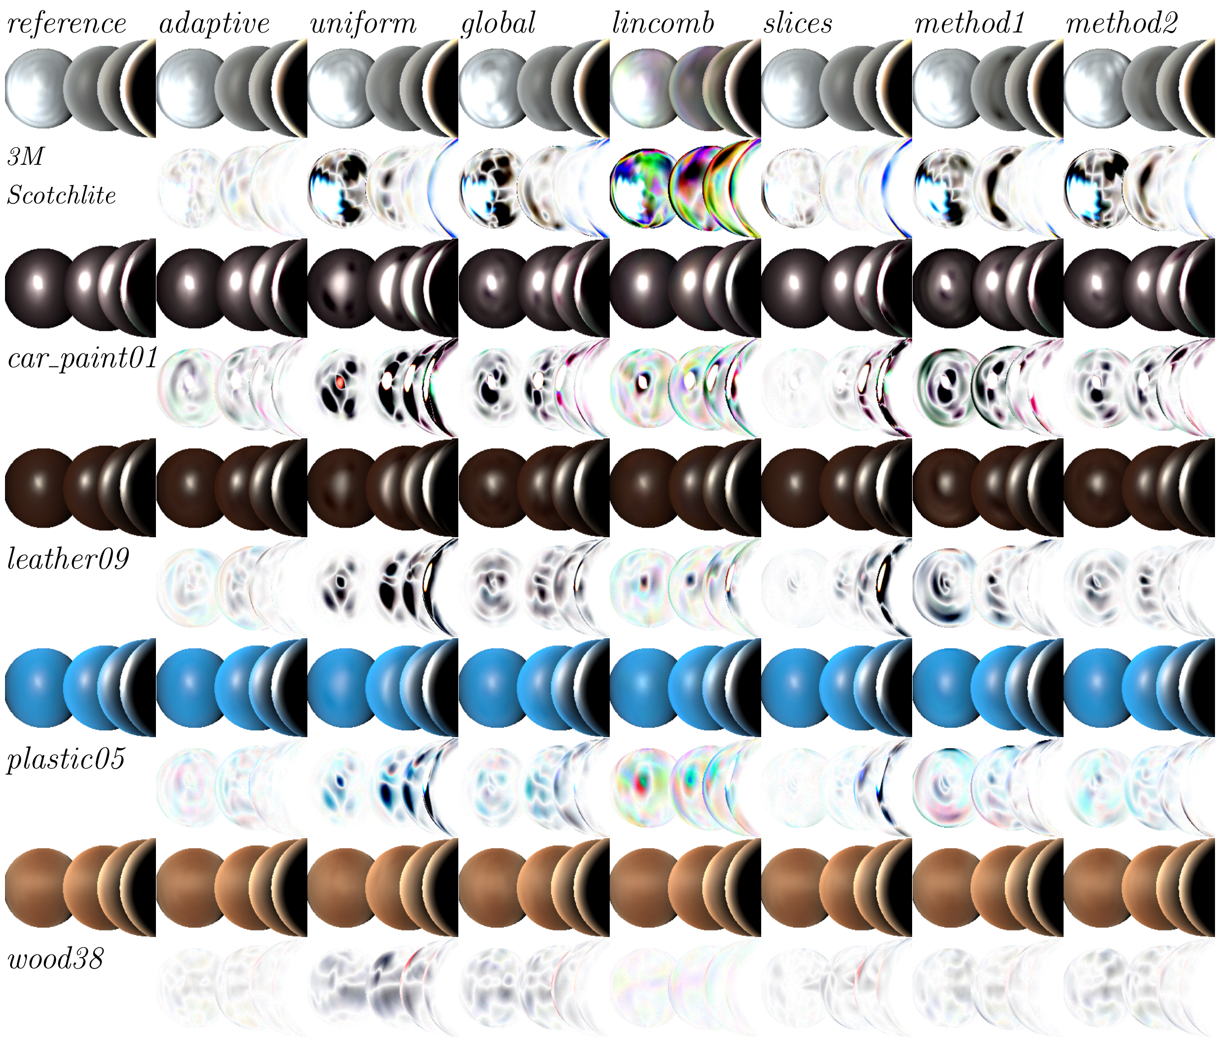

Figure 7.3: Renderings of the nearly isotropic BRDFs on four spheres: the reference vs. reconstructions using 1438 samples obtained by the seven tested methods. The difference images are scaled 10×.

(under 1000), however, it is limited to a certain quality (about 40 dB) due to a limited descriptive power of the finite number of linear bases. When more samples are used, all other techniques perform better. The highest reconstruction quality can be achieved using the proposed template-based approaches, while the worst results are achieved using the *uniform* and *global* approaches.

The lowest IS of the tested methods in Figure 7.2-b is obtained using the proposed template-based techniques, which have almost constant dependability across all numbers of samples. This is in contrast to the other tested approaches that mark higher or more variable IS with the increasing number of samples.

Finally, Figures 7.3 and 7.4 show renderings for challenging, nearly isotropic and an-

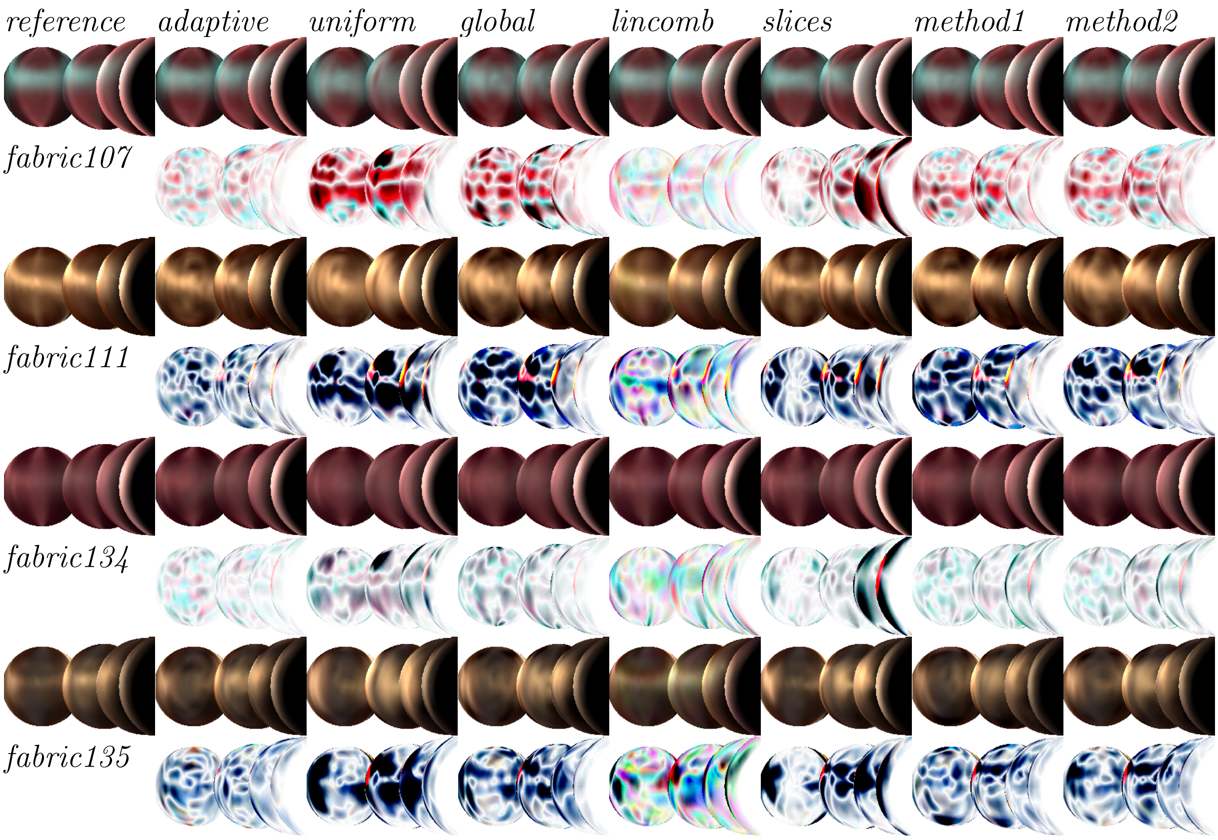

Figure 7.4: Renderings of the highly anisotropic BRDFs of fabrics on four spheres: the reference vs. reconstructions using 1438 samples obtained by the seven tested methods. The difference images are scaled 10×.

isotropic BRDFs selected from our database. A rendered scene is very similar to the 3D scene depicted in Figure 5.10 and described in Section 5.5.3. We fixed the number of samples to 1438 and compare renderings obtained by all the tested approaches with the reference BRDFs. The 10× scaled difference images reveal that the *uniform* and *global* approaches perform the worst, while the *lin.comb.* method generally works well at the cost of increased color error; especially for materials whose representation using a limited basis is difficult. Similarly, the *slices* perform very well with an exception of the grazing angles. Performance of the proposed methods is well balanced across all the materials.

7.1.5 Discussion

To test hard limits of the template-based methods, we consider sampling patterns of all materials in the database and use the one that provides the lowest error when reconstructing the known BRDF of one material excluded from the database. This test shows that overall performance gain is mere 1 dB, and thus proves that the proposed selection of the pattern is correct without much room for improvement.

An advantage of the fixed-sampling methods is a universal sampling applicable for any measured material, albeit at the cost of lower reconstruction quality. An advantage of *method1* over its variant *method2* and the BRDF *slices* is that it does not require any on-the-fly analysis during the measurement. Main factors affecting an error in process of selection of the sampling pattern are possibly the limited scope of the BRDF database (some materials of unique reflectance are present only once) and the number and distribution of initial samples.

Timings – The proposed *adaptive* and template-based techniques are suboptimal approaches only. It would be extremely computationally-demanding to find the optimal sampling pattern. Computational times for identification of 1438 samples are as follows: the *adaptive* method needs 70 minutes for each BRDF, the *global* method needs 140 minutes altogether, the *method1* needs 2 seconds for each material and the *method2* needs about 2 minutes per material using a single core of the Intel Xeon E5-2643 3.3 GHz. Identification of the fixed pattern by the *lin.comb.* method that uses 33 iterations takes about one day, and the BRDF *slices* can select required samples in 8 seconds for each BRDF. Of course, these times would increase if the measurement process were to be included.

7.2 Minimal Sampling of the Anisotropic BRDFs

This section encompasses description of the method proposed in the paper [A.1] that deals with the effective acquisition of the anisotropic BRDFs and uses the very limited count of samples. This goal is achieved by measurement of only few sampling directions for several rotations of the material around its normal, i.e., the measurement process is decomposed into several sparse scans taken at regular azimuthal intervals. Similarly to [NJR15], our data reconstruction relies on precomputed information on appearance of BRDFs in a database; however, our method extends the former method to handle the general anisotropic BRDFs. The key principle is decomposition of the training anisotropic BRDFs into a set of isotropic slices in half-difference parameterization [Rus98], which are obtained for a constant φ_h. These isotropic slices are further used as basis functions that describe arbitrary anisotropic behavior without need of its lengthy measurement by regular sampling in the four-dimensional space.

Further, we propose significant acceleration of the method for identification of a sparse set of samples presented by Nielsen et al. [NJR15]. In case of the anisotropic BRDF, the samples are measured for each azimuthal scan. The measured BRDF is then reconstructed by a linear combination of isotropic slices constrained by the measured samples. While the current methods to capture the anisotropic appearance require demanding sampling of many different incoming and outgoing directions, the proposed method allows a more practical measurement using a predefined set of fixed sampling directions in combination with rotation of a material. Typically less than 10 sampling directions in combination with 8 rotations of a material (80 samples taken together) suffice for reasonable reconstruction

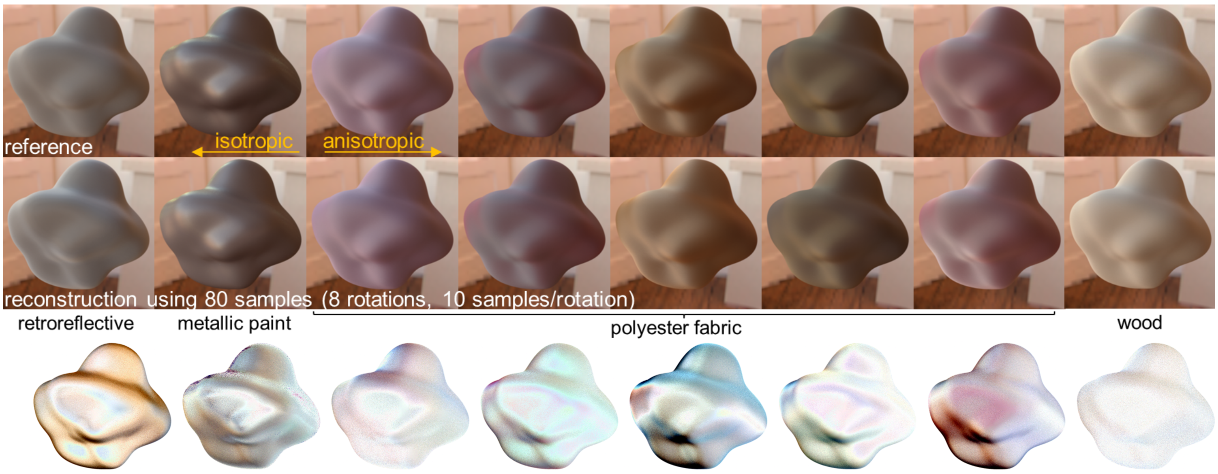

Figure 7.5: BRDFs reconstructed by the proposed method from only 80 samples. The last row shows the 10× scaled difference images.

of the main anisotropic features. Visualizations of eight reconstructed BRDFs from only 10 samples per each of 8 rotations are shown in Figure 7.5.

The main contributions of this section are:

- A novel, effective method for convenient and fast measurement of the anisotropic BRDF and its reconstruction.
- Extension of the proposed method to adaptive sampling of anisotropy of a material.
- Significant acceleration of the method presented by Nielsen et al. [NJR15] that identifies a sparse set of samples.
- A study that analyzes ability of the industrial multi-angle reflectometers to conveniently measure the anisotropic BRDF by the proposed method.

7.2.1 Formation of the Linear Basis

As the number of publicly available anisotropic BRDFs is limited, we use the database of 150 BRDFs described in Section A.4. The database is provided in the standard parameterization; therefore, we convert it to the half-difference parameterization with the uniform sampling step of 5^o in all dimensions. We verified that, due to the initial resolution of the database, this sampling step is sufficient to capture all visual features in the original BRDFs. The benefit of the non-linear sampling along θ_h proposed by Matusik et al. [MPBM03a] is negligible in this case as the database does not contain any extremely specular material. Each anisotropic BRDF is also decomposed into $n_{\varphi_h} = 360^o/5^o = 72$

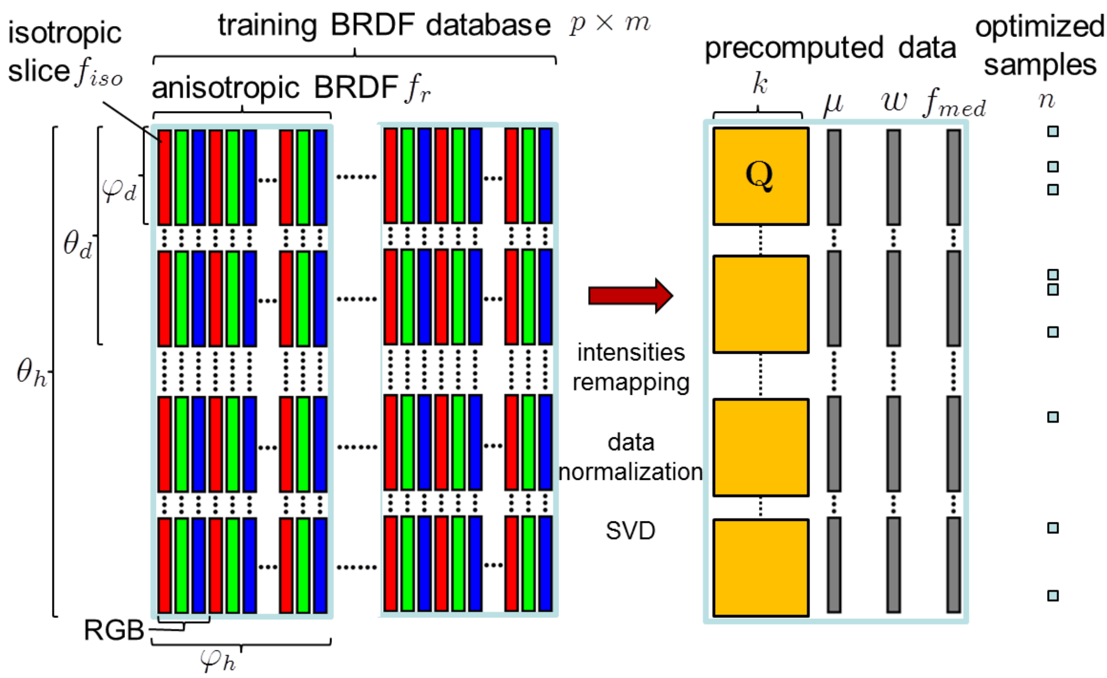

Figure 7.6: An arrangement of the BRDF database and the precomputed data necessary for reconstruction of the BRDF.

isotropic slices f_{iso} per color channel. As we treat color channels independently, we obtain $3 \times 72 = 216$ isotropic slices per each material.

Values of these isotropic BRDFs f_{iso} depending on parameters φ_d, θ_d, θ_h are vectorized and stacked together to form the matrix $\boldsymbol{A} \in \mathbb{R}^{p \times m}$ whose structure is depicted in Figure 7.6 and where $p = n_{\varphi_d} \times n_{\theta_d} \times n_{\theta_h} = 36 \times 18 \times 18 = 11,664$ is the number of rows and $m = 3 \times n_{\varphi_h} \times 150 = 32,400$ is the number of columns as there are 150 materials in the BRDF database.

Now, we can remap intensities of the data as suggested in [NJR15]. First, we compute the median isotropic BRDF f_{med} across all m columns of $\boldsymbol{A}$ and the cosine-weight factor $\boldsymbol{w} = \max\{\cos\theta_i \cos\theta_v, \epsilon\}$ that compensates for extreme grazing-angle values and where $\epsilon = 0.001$ is a constant guaranteeing numerical stability. Then, we remap the data forming the matrix $\boldsymbol{X} \in \mathbb{R}^{p \times m}$, whose each element is computed as:

$$X_{a,b} = \ln\left(\frac{A_{a,b} \cdot w_a + \epsilon}{f_{med,a} \cdot w_a + \epsilon}\right), \forall a \in \{1, \ldots, p\}, \forall b \in \{1, \ldots, m\}. \tag{7.1}$$

Next, we subtract the mean $\boldsymbol{\mu}$ over columns of the matrix $\boldsymbol{X}$ introducing the matrix

$$\boldsymbol{Y} = \boldsymbol{X} - \boldsymbol{\mu} \cdot \boldsymbol{J}^{1,m} \quad , \tag{7.2}$$

where $\boldsymbol{J}^{a,b}$ is a matrix of ones of the size of $a \times b$, replicating the column vector $\boldsymbol{\mu}$ m-times. Finally, the linear basis is derived by application of the singular value decomposition (SVD) to the matrix

$$\boldsymbol{Y} = \boldsymbol{U\Sigma V}^T \ ,$$

where columns of the matrix $\boldsymbol{U\Sigma}$ are the principal components of the data. Note that the decomposition can be performed based on eigenvectors of $\boldsymbol{Y}^T\boldsymbol{Y}$ or $\boldsymbol{Y}\boldsymbol{Y}^T$ to save computational time. We keep here the same notation as [NJR15] and define the matrix $\boldsymbol{Q} \in \mathcal{R}^{p\times k}$ of a limited count of the principal components $\boldsymbol{Q} = \boldsymbol{U}\Sigma$, where $k \leq m$ is the number of the applied principal components.

7.2.2 Data Reconstruction

Our goal is to reconstruct the missing elements of a measured BRDF given the small number of samples n for each of r rotations. For simplicity, assume $r = n_{\varphi_h} = 72$, i.e., we are given values for the same number of rotations as in the database. Later, we describe how to manage situations where $r < n_{\varphi_h}$.

Let $\tilde{\boldsymbol{A}} \in \mathbb{R}^{n\times r}$ be a matrix of known values of a reconstructed BRDF, let $\tilde{f}_{med} \in \mathbb{R}^n$ be a vector of corresponding median values, let $\tilde{\boldsymbol{w}} \in \mathbb{R}^n$ be a vector of corresponding weights, and let $\tilde{\boldsymbol{\mu}} \in \mathbb{R}^n$ be a vector of corresponding mean values. Applying those to the matrix $\tilde{\boldsymbol{A}}$ as in Equations 7.1 and 7.2 respectively, we obtain the matrix $\tilde{\boldsymbol{Y}} \in \mathbb{R}^{n\times r}$. Let $\tilde{\boldsymbol{Q}} \in \mathbb{R}^{n\times k}$ be the corresponding rows of the principal components in $\boldsymbol{Q}$. Then, the matrix $\boldsymbol{C} \in \mathbb{R}^{k\times r}$ of coefficients of the linear combination of the principal components that best models the observed data is obtained by:

$$\begin{aligned}
\tilde{\boldsymbol{Y}} &= \tilde{\boldsymbol{Q}}\boldsymbol{C} \\
\boldsymbol{C} &= \arg\min_{\boldsymbol{C}} E(\boldsymbol{C}) \\
E(\boldsymbol{C}) &= \|\tilde{\boldsymbol{Y}} - \tilde{\boldsymbol{Q}}\boldsymbol{C}\|_F^2 \\
\boldsymbol{C} &= (\tilde{\boldsymbol{Q}}^T\tilde{\boldsymbol{Q}})^{-1}\tilde{\boldsymbol{Q}}\tilde{\boldsymbol{Y}}.
\end{aligned}$$

Notice that in this case, individual columns of $\mathbf{C}$ can be computed separately. The full remapped BRDF, $\bar{\boldsymbol{X}} \in \mathbb{R}^{p\times r}$, is then reconstructed by usage of the full principal components:

$$\bar{\boldsymbol{X}} \ = \ \boldsymbol{Q}\cdot\boldsymbol{C} + \boldsymbol{\mu}\cdot\boldsymbol{J}^{1,r}. \tag{7.3}$$

Then, we obtain the full reconstructed BRDF, $\bar{\boldsymbol{A}} \in \mathbb{R}^{p\times r}$, by application of the inverse mapping of Equation 7.1 to each element:

$$\bar{A}_{a,b} = \frac{e^{\bar{X}_{a,b}}(f_{med,a}\cdot w_a + \epsilon) - \epsilon}{w_a}, \forall a \in \{1,\ldots,p\}, \forall b \in \{1,\ldots,r\}.$$

Nielsen et al. [NJR15] noticed that the least squares solution above usually results in over-fitted results, deviating significantly from the ground-truth. Therefore, they proposed a biased solution based on the work of Blanz et al. [BMVS04], who claim that $\|\boldsymbol{C}\|_F^2$ is proportional to the unlikelihood of a reconstruction. By means of introduction of the

Lagrange multiplier η in conjunction with the Frobenius norm of $\boldsymbol{C}$, it is possible to favor reconstructions closer to the observed distribution of the BRDFs:

$$\begin{aligned} E(\boldsymbol{C}) &= \|\tilde{\boldsymbol{Y}} - \tilde{\boldsymbol{Q}}\boldsymbol{C}\|_F^2 + \eta\|\boldsymbol{C}\|_F^2 \\ \boldsymbol{C} &= (\tilde{\boldsymbol{Q}}^T\tilde{\boldsymbol{Q}} + \eta\boldsymbol{I})^{-1}\tilde{\boldsymbol{Q}}\tilde{\boldsymbol{Y}}, \end{aligned} \tag{7.4}$$

where $\boldsymbol{I}$ is the identity matrix.

When this method is applied for reconstruction of anisotropic data, we have to keep in mind that a value of the BRDF in the specular reflection ($\theta_h = 0$) does not change when only φ_h change (as it is undefined there). Therefore, the function values must meet in the specular reflections and we can write that if $\theta_{h,a} = \theta_{h,b} = 0$ and $\theta_{d,a} = \theta_{d,b}$ and $\varphi_{d,a} = \varphi_{d,b}$ then $f_r(\theta_{h,a}, \varphi_{h,a}, \theta_{d,a}, \varphi_{d,a}) = f_r(\theta_{h,b}, \varphi_{h,b}, \theta_{d,b}, \varphi_{d,b})$ must hold $\forall a, b \in \{1, \ldots, p\}$. In other words, all isotropic slices $\bar{f}_{iso}$ of the reconstructed material belonging to the same color channel must have the same value on the rows where $\theta_h = 0$, i.e., $\bar{A}_{c,a} = \bar{A}_{c,b}, \forall c \in \{1, \ldots, p\}$ where $\theta_h = 0, \forall a, b \in \{1, \ldots, r\}$. This condition is very important in our model as data are tabulated and modeled independently of φ_h or else there are perceptible artifacts around the specular reflections in the reconstructed BRDFs.

Let $\tilde{\tilde{\boldsymbol{Q}}} \in \mathbb{R}^{s\times k}$ be a matrix of the rows of the principal components in $\boldsymbol{Q}$ satisfying condition above, i.e., the rows where $\theta_h = 0$, where s is the number of the rows satisfying the condition (in our case $s = n_{\varphi_d} \times n_{\theta_d} = 36 \times 18 = 648$). To satisfy the condition, we introduce another Lagrange multiplier λ in conjunction with the Frobenius norm of the residue after subtraction of its mean from reconstruction of the involved rows:

$$\begin{aligned} E(\boldsymbol{C}) &= \|\tilde{\boldsymbol{Y}} - \tilde{\boldsymbol{Q}}\boldsymbol{C}\|_F^2 + \eta\|\boldsymbol{C}\|_F^2 + \lambda\|\tilde{\tilde{\boldsymbol{Q}}}\boldsymbol{C} - \boldsymbol{\nu}\cdot\boldsymbol{J}^{1,r}\|_F^2 \\ &= \|\tilde{\boldsymbol{Y}} - \tilde{\boldsymbol{Q}}\boldsymbol{C}\|_F^2 + \eta\|\boldsymbol{C}\|_F^2 + \lambda\|\tilde{\tilde{\boldsymbol{Q}}}\boldsymbol{C} - \frac{1}{r}\cdot\tilde{\tilde{\boldsymbol{Q}}}\boldsymbol{C}\boldsymbol{J}^{r,r}\|_F^2 \\ &= \|\tilde{\boldsymbol{Y}} - \tilde{\boldsymbol{Q}}\boldsymbol{C}\|_F^2 + \eta\|\boldsymbol{C}\|_F^2 + \lambda\|\tilde{\tilde{\boldsymbol{Q}}}\boldsymbol{C}\boldsymbol{M}\|_F^2, \end{aligned} \tag{7.5}$$

where $\boldsymbol{\nu}$ is a column vector of the mean values of the rows of a matrix $\tilde{\tilde{\boldsymbol{Q}}}\boldsymbol{C}$, and $\boldsymbol{M} = \boldsymbol{I} - \frac{1}{r}\cdot\boldsymbol{J}^{r,r}$. Equation 7.5 does not have a close-form solution and must be optimized numerically. During the optimization, we can at least use its gradient, whose detailed derivation is in Section F.1:

$$\nabla E(\boldsymbol{C}) = -2\tilde{\boldsymbol{Q}}^T\boldsymbol{Y} + 2\tilde{\boldsymbol{Q}}^T\tilde{\boldsymbol{Q}}\boldsymbol{C} + 2\eta\boldsymbol{C} + 2\lambda\tilde{\tilde{\boldsymbol{Q}}}^T\tilde{\tilde{\boldsymbol{Q}}}\boldsymbol{C}\boldsymbol{M}\boldsymbol{M}.$$

In our implementation, we first estimate $\boldsymbol{C}$ using Equation 7.4 and then we perform an unconstrained minimization (Matlab function `fminunc`) to find a local optimum of Equation 7.5. We found that the method works the best for parameter values $\eta = 40$ and $\lambda = 0.1$ on our dataset. As we remove color information from data during the formation of the matrix $\boldsymbol{A}$ (see Sec. 7.2.1 and Fig. 7.6), coefficients $\boldsymbol{C}$ for each color channel must be computed separately.

Now, we describe how to handle situations where the number of rotations r of the measured material around its normal is smaller than that of original data n_{φ_h}. There are three possible solutions. The first relies on interpolation of the measured data to the required resolution by, e.g., Piecewise Cubic Hermite Interpolating Polynomials (PCHIP) producing periodic interpolated functions. The second approach is based on a similar interpolation of the linear coefficients $\boldsymbol{C}$. Although the second approach is faster, as it evaluates the significantly lower number of variables in Equation 7.5, it produces slightly worse results. Moreover, the second approach cannot be used when we do not measure all sampling directions in each rotation of a material. We use the second approach as it is much faster, unless for adaptive sampling (Sec. 7.2.6) where the first one is used. The third approach is a simple interpolation of the reconstructed BRDF $\bar{\boldsymbol{A}}$.

7.2.3 Optimization of Sampling Directions

To identify the optimal sampling directions, we follow an approach of Matusik et al. [MPBM03b] improved by Nielsen et al. [NJR15]. Let us recall that a measured BRDF is reconstructed by Equation 7.3 using the matrix $\boldsymbol{C}$ of the coefficients of the linear combination of the principal components $\boldsymbol{Q}$. Computation of the coefficients $\boldsymbol{C}$ by means of Equation 7.5 is highly dependent on a subset of rows of the principal components $\tilde{\boldsymbol{Q}}$ that corresponds to the selected sampling directions. To minimize sensitivity to errors in the modeling, we need to select the subset with the minimal condition number κ for the given count of samples n. The condition number $\kappa(\tilde{\boldsymbol{Q}})$ is defined as the ratio of the maximal and minimal singular values of $\tilde{\boldsymbol{Q}}$. Similarly to Nielsen et al. [NJR15], we exploit the fact that we deal with the three dimensional BRDF array $(\varphi_d, \theta_d, \theta_h)$ whose each location corresponds to one row in $\boldsymbol{Q}$. We slightly modified approach of Nielsen et al. [NJR15] and propose Algorithm 7.4.

We use the neighborhood of the size $h = 7$. Using the method of a fast update of

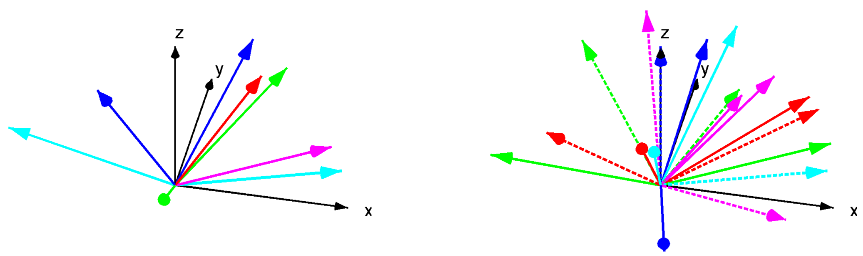

5 samples		
θ_h	θ_d	φ_d
30	0	130
0	30	40
35	60	75
0	75	0
65	0	0

10 samples					
θ_h	θ_d	φ_d	θ_h	θ_d	φ_d
35	50	55	15	60	35
25	70	95	0	0	0
0	75	15	0	30	165
15	40	80	75	0	0
35	10	65	30	50	140

Figure 7.7: Optimized sampling directions for $n = 5$ and $n = 10$ samples per rotation. All values are in degrees.

Algorithm 7.4 Optimization of Sampling Directions

1: Pick n random sampling locations in $(\varphi_d, \theta_d, \theta_h)$, i.e., pick n random rows of $\boldsymbol{Q}$ forming $\tilde{\boldsymbol{Q}}$. Based on our experience, even random selection with an extremely high $\kappa(\tilde{\boldsymbol{Q}})$ might lead to a good solution. Save $\tilde{\boldsymbol{Q}}$ for a later usage in step 4.

2: Randomly choose one of the n rows and exclude it from the matrix $\tilde{\boldsymbol{Q}}$ forming the matrix $\bar{\boldsymbol{Q}}$. For each location in $(\varphi_d, \theta_d, \theta_h)$ in a neighborhood of the excluded location of a size $h \times h \times h$ evaluate the condition numbers efficiently using the method described in Section F.2 based on eigenvectors and eigenvalues of $\bar{\boldsymbol{Q}}^T\bar{\boldsymbol{Q}}$. Accept the location with the lowest κ if it is lower than the current κ.

3: Repeat step 2 until convergence.

4: Restore the original $\tilde{\boldsymbol{Q}}$ from step 1 and repeat from step 2; 10-times.

5: Repeat from step 1; 10-times.

singular values when adding one row to a matrix described in Section F.2, we can search the neighborhood of a current location very quickly; we can find a better solution in a given amount of time. When compared to the standard approach for computation of the condition numbers, the speed gain of the proposed method is $2\times$ for $n = 5$, $4\times$ for $n = 20$, and $15\times$ for $n = 50$.

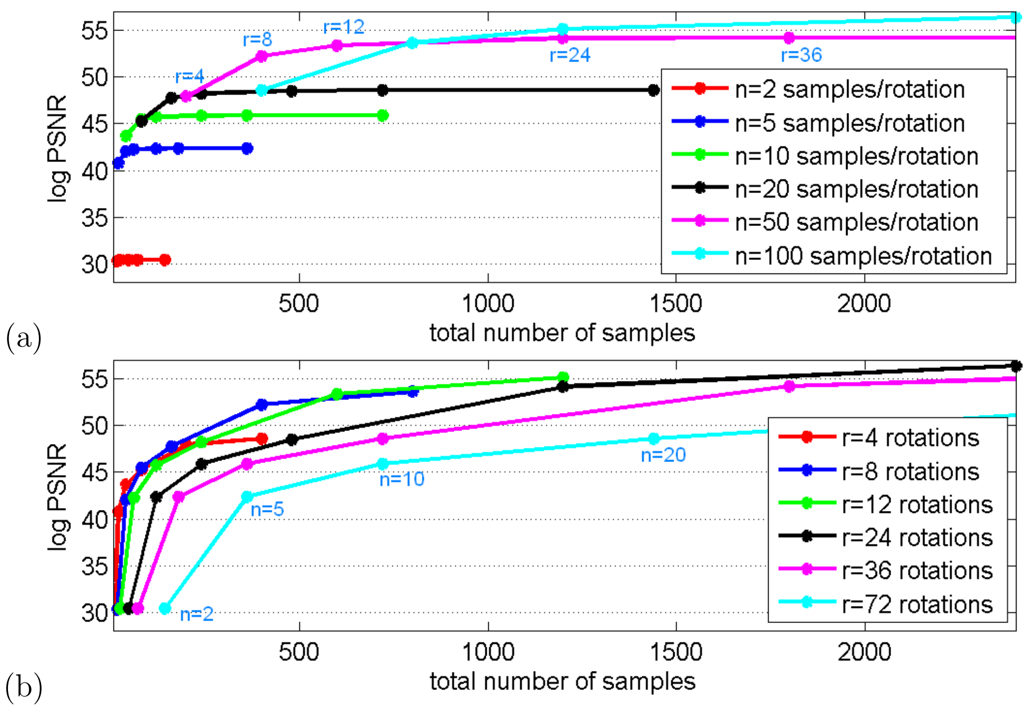

Figure 7.8: The reconstruction error (Log PSNR) as a function of the total number of samples for the given number of samples per rotation n and the given number of rotations r.

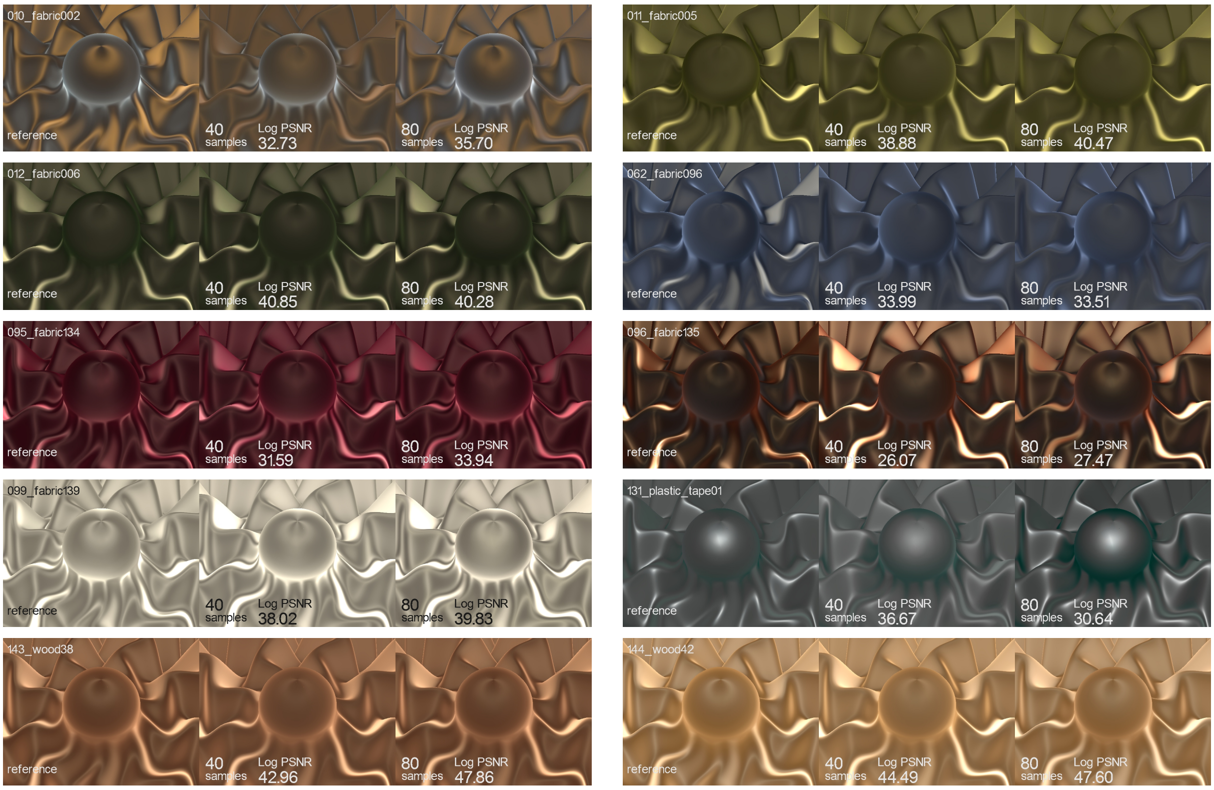

Figure 7.9: Results of the proposed method on selected anisotropic materials for configurations using: 40 samples, i.e., $n = 5$, $r = 8$ (middle), and 80 samples, i.e., $n = 10$, $r = 8$ (right). Included are values of the Log PSNR error.

The optimized directions for $n = 5, 10$ are shown in Figure 7.7. Optionally, we can even search the whole space and place the excluded sampling location to the optimal position with respect to the other sampling locations. Unfortunately, this approach does not lead to better solutions than the algorithm described above and is rather slower. Codes for formation of the linear basis, data reconstruction, and optimization of sampling directions in Matlab are provided at *http://staff.utia.cas.cz/filip/projects/16PG/*.

7.2.4 Results

Here we present results of the proposed method on the database of 150 anisotropic BRDFs (Sec. A.4). The errors are evaluated using the logarithmic PSNR (Log PSNR) HDR measure presented by Aydin et al. [AMS08]. To avoid sharing of the training and testing data, we apply the leave-one-out validation, i.e., we exclude each tested BRDF from the linear basis.

Figure 7.8 shows the error averaged across all materials from the database as a function of the total number of samples. The outlines in the graph (a) depict the given number of

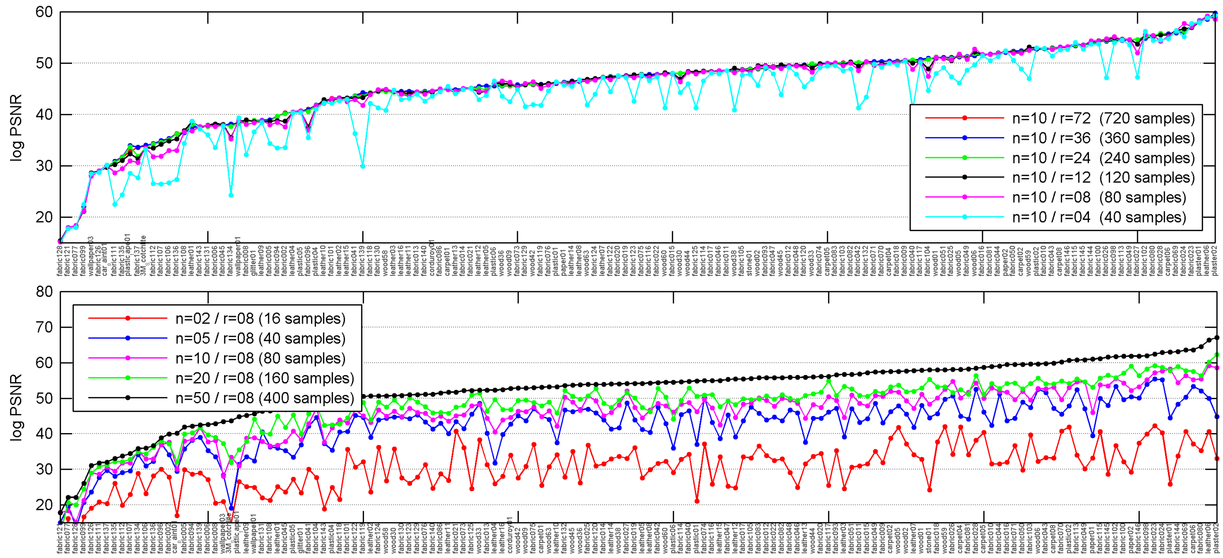

Figure 7.10: The reconstruction error for the 150 tested materials: for the fixed number of samples per rotation $n = 10$ (top), for the fixed number of rotations $r = 8$ (bottom).

samples per rotations $n = 2, 5, 10, 20, 50, 100$, while in the graph (b) they show the variable number of rotations $r = 4, 8, 12, 24, 36, 72$. From the graphs, we can draw a conclusion that in general, n has greater impact on the reconstruction error than r, where only a minor improvement is achieved for more that $r = 12$ rotations. One can observe rapid decrease of the error when eight or more rotations are used. Similarly, the error drops for five to ten samples per rotation; whereas for more, the improvement is approximately linear to the number of samples. Therefore, the efficient method should decompose measurement into 8 rotations and the number of samples taken for each rotation should be selected up on accuracy requirements starting from $n = 5$. Visual results for selected materials for configurations $n = 5$, $r = 8$ (40 samples) and $n = 10$, $r = 8$ (80 samples) are shown for an environmental illumination in Figure 7.5 and for a point-light illumination in Figure 7.9.

Figure 7.10 presents the reconstruction errors across all 150 tested materials from the database for 10 samples per rotation and the various number of rotations, and 8 rotations and the various number of samples per rotation. Here again, one can clearly observe that the impact of the number of samples per rotation is more significant.

For visual comparison, we use the surface optimized to high coverage of illumination and viewing directions proposed by Havran et al. [HFM16]. We render the surface for different combinations of parameters $n = 5, 10, 20$ and $r = 4, 8, 12$. Figures 7.11 and 7.12 show, accordingly to the previous results, the significant performance gain when more than four rotations and 10 samples per rotation are used.

Finally, we compare our method with other recent approaches to adaptive sampling of anisotropic BRDFs presented in Section 7.1. Figure 7.13 illustrates that the proposed method performs better then the template-based approaches (*method1*, *method2*) when the total number of samples drops under 500. One can observe that the method based on the

Figure 7.11: Renderings of the reconstructed BRDFs for one anisotropic material and influence of the number of samples per rotation n (columns), and the number of rotations r (rows). Included are values of the Log PSNR error and the total number of samples.

linear combination of BRDFs (*lin.comb.*) achieves slightly higher PSNR values. Note that our method uses the regularization constraint (see Eq. 7.4) that avoids data over-fitting at the cost of slightly lower PSNR values. Thus, PSNR values of *lin.comb.* might not always correspond to the visual performance.

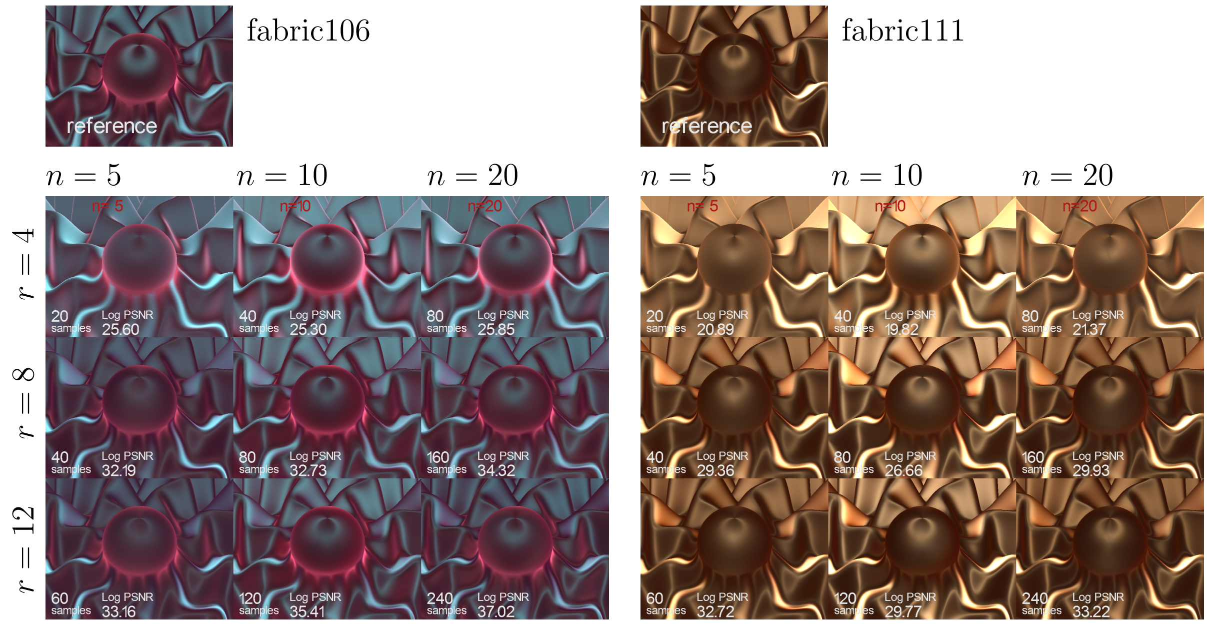

Figure 7.12: Renderings of the reconstructed BRDFs for two anisotropic materials and influence of the number of samples per rotation n (columns), and the number of rotations r (rows). Included are values of the Log PSNR error and the total number of samples.

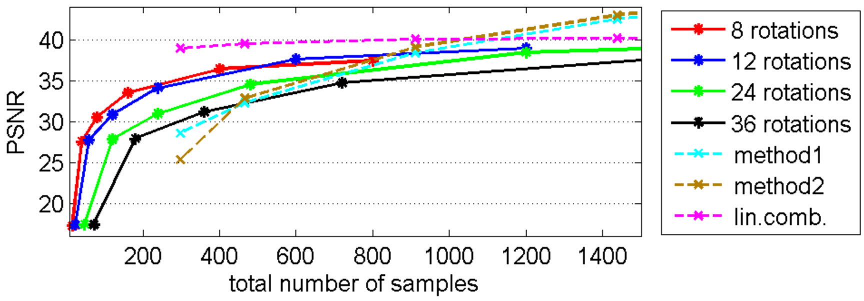

Figure 7.13: The reconstruction error (PSNR) as a function of the total number of samples for the proposed method and methods presented in Section 7.1.

7.2.5 Discussion

Here we discuss several topics related to the proposed method and the database used for our experiments.

Contribution of the Proposed Constraint – the original optimization approach produces disruptive anisotropic artifacts running from the singularity point ($\theta_h = 0$) and interfering with the anisotropic highlights. Figure 7.14 compares contribution of the proposed constrained optimization according to Equation 7.5 when compared to the original approach (Eq. 7.4) on a sphere. While the original approach achieves overall a higher PSNR

(a) the original optimization [NJR15]

(b) application of the proposed regularization constraint

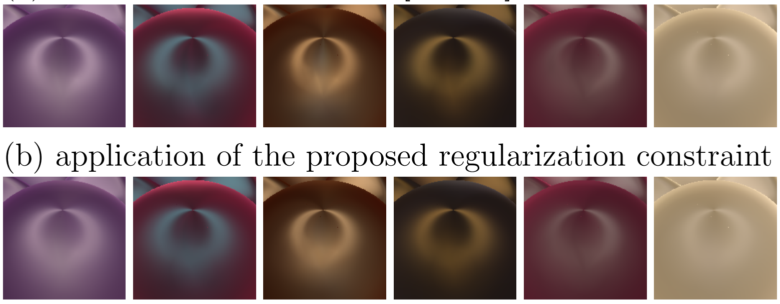

Figure 7.14: A comparison of two optimization approaches: (a) the original optimization according to Equation 7.4, (b) the proposed constrained optimization according to Equation 7.5. Results are for 360 samples, i.e., $n = 10$, $r = 36$.

than the constrained approach, it introduces the aforementioned anisotropic artifacts.

Linear Basis Formed by the Isotropic Slices *vs.* by the Full BRDFs – alternatively to our approach, one can optimize sampling directions over the four-dimensional BRDF space. Although the idea seems appealing, it has several drawbacks. First, our method, in comparison with other adaptive-sampling approaches (Sec. 7.1), does not require any azimuthal alignment (Sec. A.4.2) of a material during the measurement process. This saves up to 48 samples that were previously needed for the alignment. Second, this approach would require a much bigger database of anisotropic BRDFs as the datavectors would have higher dimensionality. The insufficient number of datasets results in various color artifacts present in the reconstructed BRDFs (see the *lin.comb.* method in Fig. 7.3 and 7.4). Third, usage of the full BRDFs would not allow the proposed fast measurement of the anisotropic BRDF by a simple rotation of a material.

Limits of the Training Dataset – As any machine learning method, our approach is limited mainly by descriptiveness of the training dataset. Therefore, a collection of anisotropic materials of various appearance behavior is necessary. An example of the insufficient training data is the highly retro-reflective material *3M_Scotchlite*, where performance of our method is not ideal (the first material in Fig. 7.5) as it is the only material of this kind in the database. We test sensitivity of selection of BRDFs used for formation of the linear basis to the reconstruction error, i.e., we select a subset of the isotropic slices that are the most similar to the reconstructed isotropic slice. We use a technique based on similarity of the captured samples of the query material and the isotropic slices stored in the basis. Although there is a noticeable effect of this filtering, the final performance gain is negligible when we consider increased computational costs related to such customized formation of the linear basis. One may consider also addition of the MERL isotropic BRDFs [MPBM03a] to expand our basis. Our experience shows that usage of the isotropic BRDFs as the only data is insufficient as such data miss unique features occurring in anisotropic BRDFs. On the other hand, their addition to the basis might improve the performance, although one have to carefully resample any added BRDF into the resolution of the basis.

Accuracy of the Reconstruction – In the applied context of the measurement process, one may require information how well the available linear basis represents the measured material. To this end, we evaluate the reconstruction error only in the measured samples and compare its correlation with the error evaluated across the entire BRDF. The correlation depends on a combination of n and r; however, above $n \geq 10$ and $r \geq 8$ the Pearson correlation coefficient is over 0.9. Thus, relation of the errors over the 150 tested materials can be approximately represented by a single scaling factor 0.8, i.e., one should multiply the error obtained in sampling locations to estimate the reconstruction error of the entire BRDF.

Timings – our method was implemented in MATLAB and all processing was done on the Intel Xeon E5-2643 3.3 GHz. Computation of the linear basis that consists of 32,400 isotropic slices took approximately four minutes. Identification of 20 optimized directions took three minutes. Note that these two steps have to be performed only once, while the reconstruction of the anisotropic BRDF from the measured samples for the configuration $n = 20$, $r = 12$ takes about 1.5 seconds and for $n = 20$, $r = 72$ takes about 14 seconds.

7.2.6 Adaptive Measurement of Anisotropy

Our technique captures anisotropic appearance by a set of regularly spaced rotations of the measured material. Another performance gain can be achieved by a coarse regular sampling combined with a further refinement by adaptive placement of other rotations. This method can be easily applied once the remotely controlled rotary stage is used. The method enables us to decrease the error of the measurement while preserving or even decreasing time of the measurement. The process of the adaptive measurement is described in Algorithm 7.5.

Figure 7.15 shows dependency of the total number of samples and the reconstruction error on the adaptive threshold e for the fixed number of initial rotations $r = 8$. The horizontal dashed lines depict results of non-adaptive regular sampling of anisotropy using $r = 8$ rotations. The best trade-off between the number of samples and the reconstruction error is achieved when threshold $e = 0.1$ is used.

Algorithm 7.5 Adaptive Measurement of Anisotropy

1: Capture the BRDF by the sparse azimuthal sampling (e.g., 8 rotations).
2: For each captured sample respectively:
    Evaluate the reconstruction error by interpolation of its value from the captured samples in its neighborhood by, e.g., PCHIP.
3: If the interpolated and actual values of the sample differ more than e percent, plan to measure two sampling directions:
    The middle point between the sample and its left neighbor.
    The middle point between the sample and its right neighbor.
4: Capture all planed sampling directions and repeat from step 2 until all errors are lower than e percent.

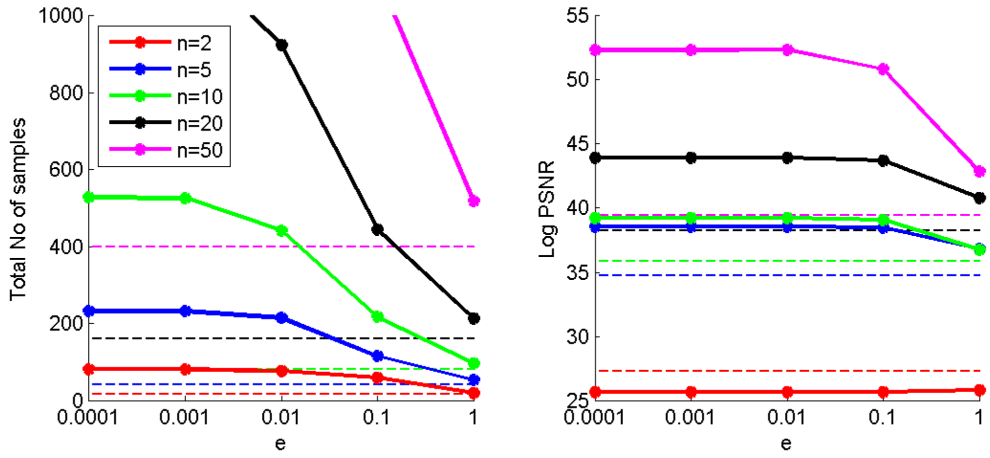

Figure 7.15: The total number of samples and the reconstruction error averaged over six highly anisotropic materials (see Fig. 7.5) as functions of the adaptive threshold e for the fixed number of initial rotations $r = 8$.

Figure 7.16 illustrates the visual fidelity gain achieved by the adaptive placement of the rotations. One can observe significant improvement while only 70 samples placed adaptively are used instead of 80 samples placed over uniformly distributed rotations. The adaptive sampling along φ_h is more technically demanding, but when the remotely controlled rotary stage is used, one can achieve better results while a setup can be simpler and less expensive.

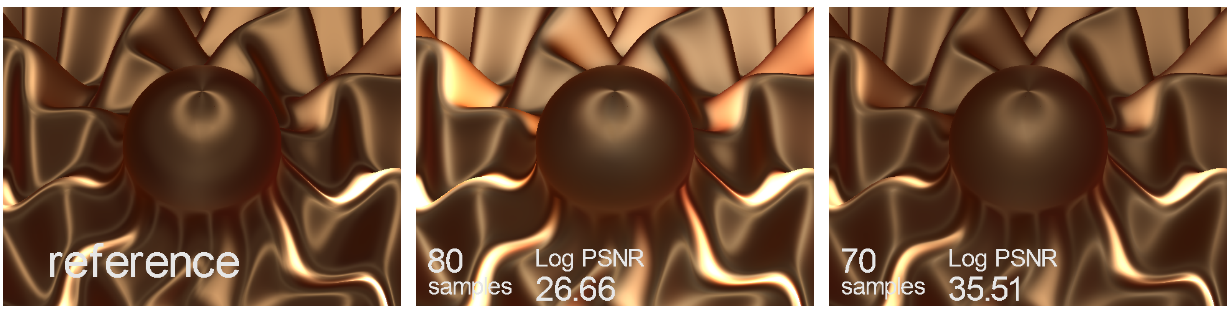

(a) (b) non-adaptive (c) adaptive

Figure 7.16: The benefit of the adaptive measurement of the highly anisotropic material *fabric111*: (a) the reference, (b) the non-adaptive measurement using 80 samples ($n = 10$, $r = 8$), (c) the adaptive measurement using 70 samples ($n = 5$, $e = 1.0$, the initial number of rotations $r = 8$).

7.2.7 Measurement by Multi-Angle Reflectometers

The proposed method allows to capture anisotropic appearance by the very low number of bidirectional pairs of light sources and sensors. As these pairs are taken identically for each rotation of a material, one can develop a relatively simple setup for measurement of material appearance. Note that anisotropy can be captured either by rotation of a device above a material or by rotation of a material below a device.

Before development of any new setup, one can validate current industrial standards used for measurement of appearance [WM01] and test current devices used in industry for multi-angle reflectance measurements. To this end, we test the bidirectional pairs that correspond to four common industrial devices: MA68 and MA98 by X-rite, BYK-mac by Gardner, and MultiFX10 by Datacolor as described in more detail in [PCVMV13]. These devices take from 5 to 11 in-plane samples to identify unique reflectance properties of materials as shown in Table 7.1. Only MA98 uses eight additional out-of-plane bidirectional pairs. Notice that the first five in-plane samples, stemming from ASTM and DIN standards [WM01], are captured by all evaluated devices.

In the context of our work, one can directly use pairs specified by a device for reconstruction of the BRDF instead of the optimized bidirectional pairs (Sec. 7.2.3). Figure 7.17 shows the results of the devices averaged across the 150 BRDFs, and their comparison with the results of various numbers of the optimized directions ($n = 5, 10, 20$) as functions of the number of rotations. The graph shows performance of the setups comparable with the proposed sampling for the low number of samples (MA68, BYK-mac); however, with the increasing number of samples, the best reconstruction quality is obtained by means of the optimized directions. The best trade-off between the reconstruction error and the number of rotations is obtained at eight rotations. However, when we focus on highly anisotropic materials, it is apparent that the optimized directions give better results both visually and in terms of the reconstruction error as shown in Figure 7.18.

These results demonstrate that although the current multi-angle devices are designed

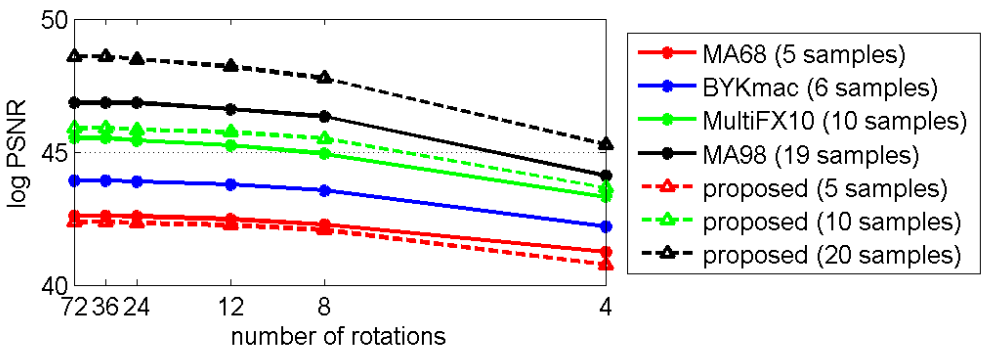

Figure 7.17: The reconstruction error as a function of the number of rotations for $n = 5, 10, 20$ optimized samples and for the various number of samples taken by the industrial devices.

Table 7.1: The geometries sampled by the industrial devices in the standard BRDF parametrization. All values are in degrees.

MA68				MA98				BYK-mac				Multi FX10			
θ_i	φ_i	θ_v	φ_v	θ_i	φ_i	θ_v	φ_v	θ_i	φ_i	θ_v	φ_v	θ_i	φ_i	θ_v	φ_v
in-plane samples															
45	0	0	0	45	0	0	0	45	0	0	0	45	0	0	0
45	0	30	0	45	0	30	0	45	0	30	0	45	0	30	0
45	0	65	0	45	0	65	0	45	0	65	0	45	0	65	0
45	0	20	180	45	0	20	180	45	0	20	180	45	0	20	180
45	0	30	180	45	0	30	180	45	0	30	180	45	0	30	180
				45	0	60	180	45	0	60	180	45	0	60	180
				15	0	30	0					15	0	0	180
				15	0	65	0					15	0	30	180
				15	0	20	180					65	0	50	180
				15	0	30	180					65	0	80	180
				15	0	60	180								
out-of-plane samples															
				15	0	50	33								
				15	0	50	327								
				15	0	45	90								
				15	0	45	270								
				45	0	50	33								
				45	0	50	327								
				45	0	45	90								
				45	0	45	270								
5 samples				19 samples				6 samples				10 samples			

primarily to discriminate materials and their visual properties, they can be, when combined with a rotary stage, readily used for measurement of the main features of the anisotropic BRDFs. However, whenever higher accuracy is needed, one should resort to a custom build apparatus with ten or more bidirectional pairs. Still, such a device could capture and reconstruct the anisotropic BRDF in several minutes.

7.3 Summary

This chapter deals with methods based on analysis of a BRDF database and which use gained knowledge for fast measurement of new materials. Introduced methods either look for the optimal sampling of materials in a database, which is then applied to sampling of a measured material, or even combine materials in a database to model a measured material.

First, we introduce an algorithm for a reference adaptive sampling. This heuristic algorithm adaptively places the required number of samples to find the suboptimal sampling patterns. Then, we present a template-based algorithm that selects a proper precomputed

MA68 MA98 BYC-mac MultiFX10 proposed n=5 proposed n=10 reference

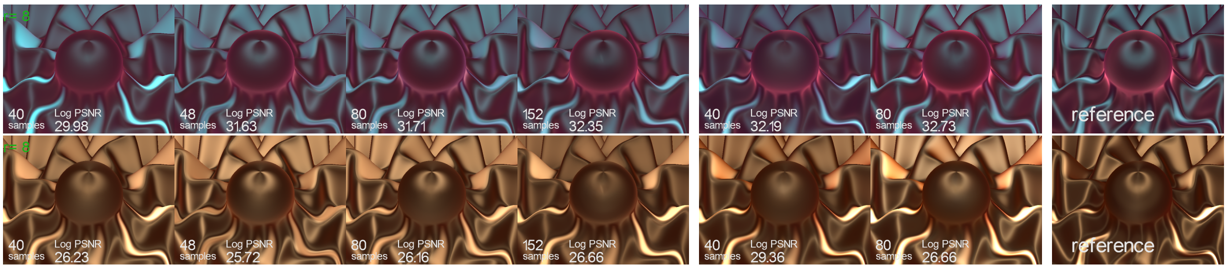

Figure 7.18: Renderings of the reconstructed BRDFs measured using the bidirectional pairs of the four industrial devices and the proposed optimized sampling using $n = 5, 10$ samples per rotation. The number of rotations $r = 8$. Included are values of the Log PSNR error and the total number of samples.

sampling pattern for an unknown material based on 90 samples only. The second proposed template-based method allows to change the pattern during the course of the measurement to achieve even better results.

A method for the convenient acquisition of the anisotropic BRDFs, which relies on the extremely low number of samples, is outlined in Section 7.2. The method allows decomposition of the measurement process into independent measurements for several rotations of a material. At each rotation, we take between five to ten samples at the optimized directions common for all rotations. We show that below one hundred samples are sufficient to capture the main features of even highly anisotropic materials. To further improve performance of our method, we extend it to the adaptive sampling of anisotropy. Finally, we show that our method can directly utilize current industrial multi-angle reflectometers for the convenient measurement of approximate anisotropic BRDFs.

We compare our template-based sampling techniques as well as the minimal-sampling method of Section 7.2 with the uniform sampling, the linear combination of BRDFs, and the BRDF slices (Chap. 5). Our experiments show that none of the tested methods performs perfectly across all BRDFs from our database. If less than thousand samples are needed, the *lin.comb.* method together with the minimal-sampling method work the best, while our method enables the much more convenient measurement of samples. In all other cases, the BRDF slices or the proposed template-based methods dominate. The best stability of reconstruction performance across tested BRDFs is achieved by the template-based methods.

The results show that the current adaptive approaches are able to achieve a low reconstruction error using fewer than 2,000 samples when reconstructing the anisotropic BRDF. The minimal-sampling method even suggests that below 100 samples are sufficient for the approximate reconstruction of the anisotropic BRDF when a database contains materials with appearance similar to appearance of a measured material.

Chapter **8**

Conclusions

In this chapter, we summarize the main contributions and achievements of the thesis as divided into five categories. Moreover, we propose possible extensions and improvements of the presented work.

Analysis of Material Appearance – is an important subject, which can help us better understand the directionally-dependent behavior of measured materials and we also analyze visibility of structure or texture of materials in virtual environments. Particularly, we introduce the innovative method and setup for the rapid detection of strength of anisotropy, main anisotropy axes, and corresponding shape of highlights. The analyzed material does not need to be extracted from its environment. Retrieved properties of the material can be used to prepare a sampling pattern that enables the efficient measurement of the material. Also, azimuthal subspaces of the material can be directly measured if we include a projector in the setup. Gathered knowledge on the typical behavior of anisotropic materials helped us design a novel empirical anisotropic BRDF model. Moreover, we study human's ability to distinguish structure or texture of a material in a virtual environment in dependence on an observation distance. We perform three psychophysical experiments and identify the *critical viewing distance* for each of 25 materials. Then, we look for computational features that can predict this distance for any material. Application of the distance saves computational costs if material appearance can be represented by the BRDF instead of the BTF without a loss of visual fidelity.

Novel Approaches to Measurement of Material Appearance – enable us to acquire data more efficiently than previous approaches or process data acquired by conventional methods more precisely. Specifically, we present a registration method that achieves the correct alignment of multi-view images of a planar surface, which is applicable on data acquired by both current and novel methods. Hence, compressed images are much sharper, or alternatively, we can use a more compact representation to achieve the same quality. The proposed method is robust, easy to implement and computationally efficient. Next, an image-based measurement of angularly-dense BRDF data is proposed. We exploit a wide span of illumination and viewing directions in each image due to the large size of a homogeneous specimen of a measured material and collect nearly three million

BRDF samples from only 8,505 images. So, the method overcomes the main limitation of a goniometric setup – the long measurement time.

Adaptive Measurement of Material Appearance – can be done by the proposed method that does not rely on any database of already measured materials. Therefore, it can be used for creation of the database of BRDFs in a selected resolution, which might then serve as the basis for other methods. Our method is based on a concept of so-called BRDF slices, which form a 4D structure in the BRDF space. While distribution of the BRDF slices in the structure must be given in advance, samples along the slices are captured adaptively by the proposed heuristic algorithm. Our experiments reveal that the algorithm exhibits fast convergence. We performed an extensive study to determine the optimal placement of the slices with respect to the demanded number of samples, which is highly correlated with the measurement time. Moreover, we provide equations for fast reconstruction of an arbitrary value in the BRDF space from values on the BRDF slices and their implementation on a graphics hardware. Our experiments show that the proposed method achieves almost four-times lower reconstruction errors in comparison with the standard regular sampling of the BRDF space in combination with the reconstruction performed by the RBFs and more than seven-times lower errors when the reconstruction of the standard sampling is performed by the barycentric interpolation, which enables real-time reconstruction on a GPU like our method but unlike the RBFs. Alternatively, we can say that our method requires from two- to five-times less samples to achieve the same error as other methods. Finally, we propose a reconstruction method that further improves our method; however, it works only in 2D subspaces of the BRDF yet.

Application of Our Experience – with the BRDF slices into practical measurement of material appearance brings the novel inexpensive setup, which is composed of a consumer hardware and enables the rapid acquisition of samples. Our psychophysical studies show that precision of the proposed approximative reconstruction of the entire (A)BRDF and BTF datasets is sufficient. We demonstrate the flexibility of the acquisition method by the rapid measurement of human skin on an author's hand. Moreover, we measure a dynamic process and present its results in a form of the TVBTF. Our analysis reveals that the setup might be easily improved by the usage of four cameras and four light sources, which provide the best trade off between accuracy, measurement time, and cost of the setup.

Data-Driven Adaptive Measurement – is done by the proposed methods that rely on a BRDF database of already measured materials. The template-based methods precompute the suboptimal sampling pattern for each material in the database. When we measure a new material, it is captured by the pattern of the most similar material. Alternatively, we iteratively select the most similar material during the process and use its characteristics for adaptive measurement. Next, we present the method for minimal sampling of the anisotropic BRDF, which enables convenient acquisition by the very low number of samples. Specifically, only a few samples of a material are acquired for several of its rotations around its normal, which is sufficient for precise reconstruction. It is even possible to use current industrial multi-angle reflectometers directly for the measurement. To achieve even better results, anisotropy can be measured adaptively.

All proposed methods of adaptive measurement achieve outstanding results when compared with other approaches. When we want to quickly acquire the appearance of a material, we should use the proposed minimal-sampling method to achieve optimal results. If we are willing to spend more time on measurement and to collect more than several hundred samples of the material, we can achieve even better results by the template-based techniques. If a database of sufficient size and resolution is not available, we can use our method based on the BRDF slices, which also achieves superior results to current methods.

During our research, we acquired many datasets, which are now publicly available for research purposes. Among others, we provide six BTF datasets in a standard resolution, three angularly-dense BRDF datasets, a unique database of 150 anisotropic BRDFs, and twelve angularly-dense BRDF datasets measured by both the standard methods and the proposed adaptive methods.

Altogether, we present methods for the analysis of material appearance, a method for the correct registration of captured images, an image-based method for angularly-dense measurement of the BRDF, setups for convenient measurement of material appearance, and foremost, methods for the adaptive measurement of material appearance that outperform other techniques in precision or speed of measurement of the anisotropic BRDFs.

8.1 Future Work

The author of the dissertation thesis suggests exploring the following:

- An extension of the improved method for reconstruction of values form the BRDF slices proposed in Section 5.7 into higher dimensions might make the method based on the BRDF slices introduced in Chapter 5 even better. The extension should be done in a way similar to the extension of the former equations for reconstruction of BRDF values in two-dimensions presented in Section 5.3, which is described in Section D.1. However, this extension has to be designed very carefully due to singularities in the half-difference parameterization.
- Our study on the optimal placement of the BRDF slices (see Sec. 5.4) might be extended in future work to a more precise measurement of new materials. We would like to identify a small number of samples that enable us to classify materials into groups with the similar directional behavior. Then for each group independently, we can find the optimal placement of the BRDF slices and we can even deploy the slices irregularly into the space. Consequently, the measurement might be even more efficient.
- Usage of the BRDF slices for measurement of the BTF data is possible as shown in Chapter 6; however, it requires further detailed analysis in future work.
- Our report on adaptive measurement of 2D BRDF subspaces [A.16] suggests that it is possible to extend the heuristic algorithm for adaptive measurement of 1D signals (Sec. 5.2) into adaptive measurement of 2D BRDF subspaces, and consequently, also into adaptive measurement of the 4D BRDF space. We determined that it is

important to measure many samples at each iteration of an algorithm, similarly as in [FBLS07], to suppress oversampling of locations with greater variations of values. Alternatively, we can suppress oversampling by introduction of some sort of penalization for oversampled locations.

- The optimization approach for the minimal-sampling method (Sec. 7.2.3) selects the directions that are measured for each rotation of a material according to the condition number of a subset of the principal components. According to our observation, this criterion is not ideal, because it does not reflect actual descriptiveness of the subset and even the subset defined by the X-Rite MA68 reflectometer achieves better results than our subset while its condition number is much worse. Therefore, it would be interesting to find another criterion that better correlates with the reconstruction error. Particularly, we plan to try the error function proposed in [XNY*16].

Bibliography

[AMS08] AYDIN T. O., MANTIUK R., SEIDEL H.-P.: Extending quality metrics to full dynamic range images. In *Human Vision and Electronic Imaging XIII* (San Jose, USA, January 2008), Proceedings of SPIE, pp. 6806–10.

[AS00] ASHIKHMIN M., SHIRLEY P.: An anisotropic phong light reflection model. *Journal of Graphics Tools 5*, 2 (2000), 25–32.

[ASOS12] ALI M. A., SATO I., OKABE T., SATO Y.: Toward efficient acquisition of BRDFs with fewer samples. In *ACCV* (2012), pp. 54–67.

[AZK08] ALLDRIN N. G., ZICKLER T., KRIEGMAN D.: Photometric stereo with non-parametric and spatially-varying reflectance. In *CVPR* (2008), pp. 2505–2511.

[BEWW*08] BEN-EZRA M., WANG J., WILBURN B., LI X., MA L.: An LED-only BRDF measurement device. In *CVPR* (2008), pp. 1–8.

[BK80] BURBECK C. A., KELLY D.: Spatiotemporal characteristics of visual mechanisms: Excitatory-inhibitory model. *Journal of the Optical Society of America 70*, 9 (1980), 1121–1126.

[Bli77] BLINN J.: Models of light reflection for computer synthesized pictures. *SIGGRAPH Comput. Graph. 11* (July 1977), 192–198.

[BMVS04] BLANZ V., MEHL A., VETTER T., SEIDEL H. P.: A statistical method for robust 3D surface reconstruction from sparse data. In *Proceedings of 2nd International Symposium on 3D Data Processing, Visualization and Transmission* (2004), pp. 293–300.

[BNS78] BUNCH J. R., NIELSEN C. P., SORENSEN D. C.: Rank-one modification of the symmetric eigenproblem. *Numerische Mathematik 31*, 1 (1978), 31–48.

[Boo89] BOOKSTEIN F. L.: Principal warps: thin-plate splines and the decomposition of deformations. *IEEE TPAMI 11*, 6 (1989), 567–585.

[BSH12] BAGHER MAHDI M., SOLER C., HOLZSCHUCH N.: Accurate fitting of measured reflectances using a shifted gamma micro-facet distribution. *Computer Graphics Forum 31*, 4 (2012), 1509–1518.

[BvdPPH11] BONNEEL N., VAN DE PANNE M., PARIS S., HEIDRICH W.: Displacement interpolation using lagrangian mass transport. *ACM Transactions on Graphics 30*, 6 (December 2011), 158:1–158:12.

[Cox69] COXETER H. S. M.: *Introduction to Geometry.* New York: Wiley, 1969.

[CT81] COOK R., TORRANCE K.: A reflectance model for computer graphics. *ACM SIGGRAPH 1981, ACM Press 15*, 3 (August 1981), 307–316.

[Dan01] DANA K.: BRDF/BTF measurement device. In *Proceedings of Eighth IEEE International Conference on Computer Vision. ICCV 2001* (July 2001), vol. 2, pp. 460–466.

[Deb98] DEBEVEC P.: Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proceedings of the 25th annual conference on Computer graphics and interactive techniques* (1998), SIGGRAPH '98, pp. 189–198.

[DHI*15] DUPUY J., HEITZ E., IEHL J.-C., POULIN P., OSTROMOUKHOV V.: Extracting microfacet-based BRDF parameters from arbitrary materials with power iterations. In *Proceedings of the 26th Eurographics Symposium on Rendering* (Aire-la-Ville, Switzerland, Switzerland, 2015), EGSR '15, Eurographics Association, pp. 21–30.

[DvGNK99] DANA K., VAN GINNEKEN B., NAYAR S., KOENDERINK J.: Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics 18*, 1 (1999), 1–34.

[DW04] DANA K., WANG J.: Device for convenient measurement of spatially varying bidirectional reflectance. *Journal of Optical Society of America 21*, 1 (2004), 1–12.

[DWT*10] DONG Y., WANG J., TONG X., SNYDER J., LAN Y., BEN-EZRA M., GUO B.: Manifold bootstrapping for SVBRDF capture. *ACM Transactions on Graphics 29*, 4 (July 2010), 98:1–98:10.

[ELU11] EILERTSEN G., LARSSON P., UNGER J.: A versatile material reflectance measurement system for use in production. In *SIGRAD* (2011).

[FBLS07] Fuchs M., Blanz V., Lensch H. P., Seidel H.-P.: Adaptive sampling of reflectance fields. *ACM Transactions on Graphics 26*, 2 (June 2007), 1–18.

[FCGH08] Filip J., Chantler M., Green P., Haindl M.: A psychophysically validated metric for bidirectional texture data reduction. *ACM Transactions on Graphics 27*, 5 (December 2008), 138.

[FH09] Filip J., Haindl M.: Bidirectional texture function modeling: A state of the art survey. *IEEE Trans. on Pattern Analysis and Machine Intelligence 31*, 11 (October 2009), 1921–1940.

[FHC10] Filip J., Haindl M., Chantler M.: Gaze-motivated compression of illumination and view dependent textures. In *ICPR* (August 2010), pp. 862–864.

[Fil12] Filip J.: Restoring illumination and view dependent data from sparse samples. In *ICPR* (November 2012), pp. 1391–1394.

[FSH16] Fichet A., Sato I., Holzschuch N.: Capturing Spatially Varying Anisotropic Reflectance Parameters using Fourier Analysis. In *Graphics Interface 2016* (Victoria, BC, Canada, June 2016).

[GAHO07] Ghosh A., Achutha S., Heidrich W., O'Toole M.: BRDF acquisition with basis illumination. In *IEEE 11th International Conference on Computer Vision, ICCV 2007* (October 2007), pp. 1–8.

[GCG*05] Günther J., Chen T., Goesele M., Wald I., Seidel H.-P.: Efficient acquisition and realistic rendering of car paint. In *Proceedings of 10th International Fall Workshop - Vision, Modeling, and Visualization (VMV) 2005* (November 2005), Akademische Verlagsgesellschaft Aka GmbH, pp. 487–494.

[GCHS05] Goldman D. B., Curless B., Hertzmann A., Seitz S. M.: Shape and spatially-varying BRDFs from photometric stereo. In *Proceedings of ICCV* (2005), pp. 341–348.

[GCP*10] Ghosh A., Chen T., Peers P., Wilson C., Debevec P.: Circularly polarized spherical illumination reflectometry. In *ACM Transactions on Graphics* (2010), vol. 29, ACM, pp. 162:1–162:12.

[GGSC96] Gortler S. J., Grzeszczuk R., Szeliski R., Cohen M. F.: The lumigraph. In *Proceedings of SIGGRAPH 1996* (1996), pp. 43–54.

[GHAO10] Ghosh A., Heidrich W., Achutha S., OToole M.: A basis illumination approach to BRDF measurement. *International Journal of Computer Vision 90*, 2 (2010), 183–197.

[GMSK09] Guthe M., Müller G., Schneider M., Klein R.: BTF-CIELab: A perceptual difference measure for quality assessment and compression of BTFs. *Comput. Graph. Forum 28*, 1 (2009), 101–113.

[Gol73] Golub G. H.: Some modified matrix eigenvalue problems. *Siam Review 15*, 2 (1973), 318–334.

[GTR*06] Gu J., Tu C.-I., Ramamoorthi R., Belhumeur P., Matusik W., Nayar S.: Time-varying surface appearance: acquisition, modeling and rendering. *ACM Transactions on Graphics 25*, 3 (July 2006), 762–771.

[HAH*12] Höpe A., Atamas T., Hünerhoff D., Teichert S., Hauer K.-O.: Argon^3: 3D appearance robot-based gonioreflectometer at PTB. *Review of Scientific Instruments 83*, 4 (2012), 045102.

[Hei14] Heitz E.: Understanding the masking-shadowing function in microfacet-based BRDFs. *Journal of Computer Graphics Techniques (JCGT) 3*, 2 (June 2014), 32–91.

[HF13] Haindl M., Filip J.: *Visual Texture*. Advances in Computer Vision and Pattern Recognition. Springer-Verlag, London, 2013.

[HFM10] Havran V., Filip J., Myszkowski K.: Bidirectional texture function compression based on multi-level vector quantization. *Computer Graphics Forum 29*, 1 (January 2010), 175–190.

[HFM16] Havran V., Filip J., Myszkowski K.: Perceptually motivated BRDF comparison using single image. *Computer Graphics Forum (Proceedings of EGSR) 35* (2016), 1–12.

[HHA*10] Hullin M. B., Hanika J., Ajdin B., Seidel H.-P., Kautz J., Lensch H. P. A.: Acquisition and analysis of bispectral bidirectional reflectance and reradiation distribution functions. *ACM Transactions on Graphics 29*, 4 (2010), 1–7.

[HLZ10] Holroyd M., Lawrence J., Zickler T.: A coaxial optical scanner for synchronous acquisition of 3D geometry and surface reflectance. *ACM Transactions on Graphics 29* (July 2010), 99:1–99:12.

[HP03] Han J. Y., Perlin K.: Measuring bidirectional texture reflectance with a kaleidoscope. *ACM Transactions on Graphics 22*, 3 (2003), 741–748.

[HTSG91] He X. D., Torrance K. E., Sillion F. X., Greenberg D. P.: A comprehensive physical model for light reflection. In *Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques* (New York, NY, USA, 1991), SIGGRAPH '91, ACM, pp. 175–186.

[HZ00] Hartley R. I., Zisserman A.: *Multiple View Geometry in Computer Vision*. Cambridge University Press, ISBN: 0521623049, 2000.

[Joh96] Johnson T.: Methods for characterizing colour scanners and digital cameras. *Displays 16*, 4 (1996), 183 – 191.

[JSW98] Jones D. R., Schonlau M., Welch W. J.: Efficient global optimization of expensive black-box functions. *Journal of Global Optimization 13*, 4 (December 1998), 455–492.

[JWD*14] Jarabo A., Wu H., Dorsey J., Rushmeier H., Gutierrez D.: Effects of approximate filtering on the appearance of bidirectional texture functions. *IEEE Transactions on Visualization and Computer Graphics 20*, 6 (June 2014), 880–892.

[KFB10] Křivánek J., Ferwerda J., Bala K.: Effects of global illumination approximations on material appearance. *ACM Transactions on Graphics 29*, 4 (2010), 112.

[KMBK03] Koudelka M., Magda S., Belhumeur P., Kriegman D.: Acquisition, compression, and synthesis of bidirectional texture functions. In *Texture 2003* (2003), pp. 47–52.

[KSKK10] Kurt M., Szirmay-Kalos L., Křivánek J.: An anisotropic BRDF model for fitting and Monte Carlo rendering. *SIGGRAPH Computer Graphics 44* (February 2010), 3:1–3:15.

[LBAD*06] Lawrence J., Ben-Artzi A., DeCoro C., Matusik W., Pfister H., Ramamoorthi R., Rusinkiewicz S.: Inverse shade trees for non-parametric material representation and editing. *ACM Transactions on Graphics 25*, 3 (2006), 735–745.

[LDW*10] Lan Y., Dong Y., Wang J., Tong X., Guo B.: Condenser-based instant reflectometry. *Computer Graphics Forum 29*, 7 (2010), 2091–2098.

[LFD*08] Leloup F. B., Forment S., Dutré P., Pointer M. R., Hanselaer P.: Design of an instrument for measuring the spectral bidirectional scatter distribution function. *Applied Optics 47*, 29 (2008), 5454–5467.

[LFTG97] Lafortune E. P., Foo S. C., Torrance K. E., Greenberg D. P.: Non-linear approximation of reflectance functions. *Computer Graphics Forum 31* (1997), 117–126.

[LKK98] Lu R., Koenderink J. J., Kappers A. M. L.: Optical properties (bidirectional reflectance distribution functions) of velvet. *Applied Optics 37* (1998), 5974–5984.

[LKK00] Lu R., Koenderink J. J., Kappers A. M.: Specularities on surfaces with tangential hairs or grooves. *Computer Vision and Image Understanding 78*, 3 (2000), 320–335.

[LL95] Lu J., Little J.: Reflectance function estimation and shape recovery from image sequence of a rotating object. In *Proceedings of the Fifth International Conference on Computer Vision* (Washington, DC, USA, 1995), ICCV '95, IEEE Computer Society, pp. 80–86.

[LLSS03] Lensch H. P., Lang J., Sá A. M., Seidel H.-P.: Planned sampling of spatially varying BRDFs. *Computer Graphics Forum 22*, 3 (2003), 473–482.

[LMM*10] Langenbucher T., Merzbach S., Möller D., Ochmann S., Vock R., Warnecke W., Zschippig M.: Time-varying BTFs. In *Central European Seminar on Computer Graphics for Students (CESCG)* (2010).

[LRR05] Lawrence J., Rusinkiewicz S., Ramamoorthi R.: Adaptive numerical cumulative distribution functions for efficient importance sampling. In *Eurographics Symposium on Rendering* (2005), pp. 11–20.

[MBK05] Müller G., Bendels G., Klein R.: Rapid synchronous acquisition of geometry and BTF for cultural heritage artefacts. In *The 6th International Symposium on Virtual Reality, Archaeology and Cultural Heritage (VAST)* (2005), pp. 13–20.

[MCS90] Murray-Coleman J., Smith A.: The automated measurement of BRDFs and their application to luminaire modeling. *Journal of the Illuminating Engineering Society 19* (1990), 87–99.

[MDWE02] Marziliano P., Dufaux F., Winkler S., Ebrahimi T.: A no-reference perceptual blur metric. In *ICIP 2002* (2002), vol. 3, IEEE, pp. III–57.

[MKRH11] Mantiuk R., Kim K. J., Rempel A. G., Heidrich W.: Hdr-vdp-2: a calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM Transactions on Graphics 30*, 4 (2011), 40:1–40:14.

[MMK03] Müller G., Meseth J., Klein R.: Compression and real-time rendering of measured BTFs using local PCA. In *Vision, Modeling and Visualisation* (November 2003), pp. 271–280.

[MMK*06] Meseth J., Müller G., Klein R., Röder F., Arnold M.: Verification of rendering quality from measured BTFs. In *Third Symposium on Applied Perception in Graphics and Visualization (APGV)* (July 2006), vol. 153, pp. 127–134.

[MMS*05] Müller G., Meseth J., Sattler M., Sarlette R., Klein R.: Acquisition, synthesis and rendering of bidirectional texture functions. *Computer Graphics Forum 24*, 1 (2005), 83–109.

[MPBM03a] Matusik W., Pfister H., Brand M., McMillan L.: A data-driven reflectance model. *ACM Transactions on Graphics 22*, 3 (2003), 759–769.

[MPBM03b] Matusik W., Pfister H., Brand M., McMillan L.: Efficient isotropic BRDF measurement. In *Proceedings of the 14th Eurographics Workshop on Rendering* (2003), EGRW '03, pp. 241–247.

[MSK06] Müller G., Sarlette R., Klein R.: Data-driven local coordinate systems for image-based rendering. *Computer Graphics Forum 25*, 3 (2006).

[MSY07a] Mukaigawa Y., Sumino K., Yagi Y.: High-speed measurement of BRDF using an ellipsoidal mirror and a projector. In *IEEE Conference on Computer Vision and Pattern Recognition, 2007. CVPR '07.* (June 2007), pp. 1–8.

[MSY07b] Mukaigawa Y., Sumino K., Yagi Y.: Multiplexed illumination for measuring BRDF using an ellipsoidal mirror and a projector. In *Proceedings of the 8th Asian conference on Computer vision - Volume Part II* (Berlin, Heidelberg, 2007), ACCV'07, Springer-Verlag, pp. 246–257.

[MSY09] Mukaigawa Y., Sumino K., Yagi Y.: Rapid BRDF measurement using an ellipsoidal mirror and a projector. *IPSJ Trans. on Computer Vision and Appl. 1* (January 2009), 21–32.

[MTK*11] Mukaigawa Y., Tagawa S., Kim J., Raskar R., Matsushita Y., Yagi Y.: Hemispherical confocal imaging using turtleback reflector. In *Computer Vision ACCV 2010*, vol. 6492 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2011, pp. 336–349.

[MU12] Montes R., Ureña C.: *An overview of BRDF models.* Tech. rep., 2012.

[Mül09] Müller G.: *Data-Driven Methods for Compression and Editing of Spatially Varying Appearance.* Dissertation, Universität Bonn, December 2009.

[MWLT00] Marschner S. R., Westin S. H., Lafortune E. P. F., Torrance K. E.: Image-based bidirectional reflectance distribution function measurement. *Applied Optics 39* (June 2000), 2592–2600.

[ND06] Ngan A., Durand F.: Statistical acquisition of texture appearance. *Proceedings of the Eurographics Symposium on Rendering* (August 2006), 31–40.

[NDM05] Ngan A., Durand F., Matusik W.: Experimental analysis of BRDF models. *Eurographics Symposium on Rendering 2005 2* (2005), 117–126.

[NFM14] Nauyoks S. E., Freda S., Marciniak M. A.: Dynamic data driven bidirectional reflectance distribution function measurement system. *Proc. SPIE 9205* (2014), 920502–920502–12.

[NJR15] Nielsen J. B., Jensen H. W., Ramamoorthi R.: On Optimal, Minimal BRDF Sampling for Reflectance Acquisition. *ACM Transactions on Graphics 34*, 6 (2015), 186:1–186:11.

[NKS14] Nöll T., Köhler J., Stricker D.: Robust and accurate non-parametric estimation of reflectance using basis decomposition and correction functions. In *ECCV 2014 (LNCS 8690)* (2014), pp. 376–391.

[NRH*77] Nicodemus F., Richmond J., Hsia J., Ginsburg I., Limperis T.: Geometrical considerations and nomenclature for reflectance. *NBS Monograph 160, National Bureau of Standards, U.S. Department of Commerce* (1977), 1–52.

[NZG04] Neubeck A., Zalesny A., Gool L. V.: Viewpoint consistent texture synthesis. In *Proceedings of the 3D Data Processing, Visualization, and Transmission, 2nd International Symposium* (2004), pp. 388–395.

[ON95] Oren M., Nayar S. K.: Generalization of the lambertian model and implications for machine vision. *International Journal of Computer Vision 14*, 3 (1995), 227–251.

[PCS*12] Pacanowski R., Celis O. S., Schlick C., Granier X., Poulin P., Cuyt A.: Rational BRDF. *IEEE Transactions on Visualization and Computer Graphics 18*, 11 (2012), 1824–1835.

[PCVMV13] Perales E., Chorro E., Viqueira V., Martínez-Verdú F. M.: Reproducibility comparison among multiangle spectrophotometers. *Color Research & Applic. 38*, 3 (2013), 160–167.

[PF90] Poulin P., Fournier A.: A model for anisotropic reflection. *SIGGRAPH Comput. Graph. 24*, 4 (September 1990), 273–282.

[Pho75] Phong B.-T.: Illumination for computer generated images. *Communications of the ACM 18*, 6 (June 1975), 311–317.

[Pir67] Pirenne M.: *Vision and the Eye.* Chapman and Hall, London, 1967.

[PN04] Parkhurst D. J., Niebur E.: Texture contrast attracts overt visual attention in natural scenes. *European Journal of Neuroscience 19*, 3 (2004), 783–789.

[PTVF92] Press W. H., Teukolsky S. A., Vetterling W. T., Flannery B. P.: *Numerical Recipes in C: The Art of Scientific Computing.* Cambridge University Press, 1992.

[QKM*00] Qi J., Kerr Y., Moran M., Weltz M., Huete A., Sorooshian S., Bryant R.: Leaf area index estimates using remotely sensed data and BRDF models in a semiarid region. *Remote sensing of environment 73*, 1 (2000), 18–30.

[RFWB07] Ramanarayanan G., Ferwerda J., Walter B., Bala K.: Visual equivalence: Towards a new standard for image fidelity. *ACM Transactions on Graphics 26*, 3 (2007).

[RGB*14] Raymond B., Guennebaud G., Barla P., Pacanowski R., Granier X.: Optimizing BRDF orientations for the manipulation of anisotropic highlights. In *Computer Graphics Forum* (2014), vol. 33, Wiley Online Library, pp. 313–321.

[RK09] Ruiters R., Klein R.: Heightfield and spatially varying BRDF reconstruction for materials with interreflections. *Computer Graphics Forum 28*, 2 (2009), 513–522.

[RMS*08] Rump M., Müller G., Sarlette R., Koch D., Klein R.: Photorealistic rendering of metallic car paint from image-based measurements. *Computer Graphics Forum 27*, 2 (2008), 527–536.

[Rob95] Robinson J.: Image coding with ridge and valley primitives. *IEEE Transactions on Communications 43*, 6 (1995), 2095–2102.

[RR95] Robinson J. A., Ren M. S.: Data-dependent sampling of two-dimensional signals. *Multidimensional Systems and Signal Processing 6* (1995), 89–111.

[RSK10] Rump M., Sarlette R., Klein R.: Groundtruth data for multispectral bidirectional texture functions. In *CGIV 2010* (June 2010), Society for Imaging Science and Technology, pp. 326–330.

[RSK12] Ruiters R., Schwartz C., Klein R.: Data driven surface reflectance from sparse and irregular samples. *Computer Graphics Forum (Proc. of Eurographics) 31*, 2 (May 2012), 315–324.

[Rus98] Rusinkiewicz S.: A new change of variables for efficient BRDF representation. In *Rendering techniques' 98* (1998), pp. 11–22.

[RVZ08] Romeiro F., Vasilyev Y., Zickler T.: Passive reflectometry. In *Proceedings of the 10th European Conference on Computer Vision: Part IV* (2008), ECCV '08, pp. 859–872.

[RWS*11] Ren P., Wang J., Snyder J., Tong X., Guo B.: Pocket reflectometry. *ACM Transactions on Graphics 30*, 4 (August 2011), 45:1–45:10.

[SAS05] Stark M., Arvo J., Smits B.: Barycentric parameterizations for isotropic BRDFs. *IEEE TVCG 11*, 2 (March-April 2005), 126 –138.

[SBDDJ13] Sadeghi I., Bisker O., De Deken J., Jensen H. W.: A practical microcylinder appearance model for cloth rendering. *ACM Transactions on Graphics 32*, 2 (April 2013), 14:1–14:12.

[Sch94] Schlick C.: An inexpensive BRDF model for physically-based rendering. *Computer Graphics Forum 13*, 3 (1994), 233–246.

[SGS*02] Schaaf C. B., Gao F., Strahler A. H., Lucht W., Li X., Tsang T., Strugnell N. C., Zhang X., Jin Y., Muller J.-P., et al.: First operational BRDF, albedo nadir reflectance products from MODIS. *Remote sensing of Environment 83*, 1 (2002), 135–148.

[SH05] Somol P., Haindl M.: Novel path search algorithm for image stitching and advanced texture tiling. In *WSCG 2005* (2005), pp. 155–218.

[SHGW13] Schick E., Herbort S., Grumpe A., Wöhler C.: Single view single light multispectral object segmentation. In *WSCG* (2013), pp. 171–178.

[SSK03] Sattler M., Sarlette R., Klein R.: Efficient and realistic visualization of cloth. In *Eurographics Symposium on Rendering* (2003), pp. 167–178.

[SSW*14] Schwartz C., Sarlette R., Weinmann M., Rump M., Klein R.: Design and implementation of practical bidirectional texture function measurement devices focusing on the developments at the University of Bonn. *Sensors 14*, 5 (2014).

[SWR*11] Schwartz C., Weinmann M., Ruiters R., Zinke A., Sarlette R., Klein R.: Capturing shape and reflectance of food. In *SIGGRAPH Asia 2011 Sketches* (2011), pp. 28:1–28:2.

[Sze06] Szeliski R.: Image alignment and stitching: A tutorial. *Found. Trends. Comput. Graph. Vis. 2*, 1 (January 2006), 1–104.

[TS67] Torrance K., Sparrow E.: Theory for off-specular reflection from rough surfaces. *Journal of the Optical Society of America 57*, 9 (February 1967), 1105–1114.

[TS96] Thompson S. K., Seber G. A. F.: *Adaptive Sampling.* John Wiley and sons, New York, 1996.

[VLD07] Vangorp P., Laurijssen J., Dutré P.: The influence of shape on the perception of material reflectance. *ACM Transactions on Graphics 26*, 3 (July 2007).

[Wan95] Wandell B. A.: *Foundations of vision.* Sinauer Associates, 1995.

[War92] Ward G.: Measuring and modeling anisotropic reflection. *Computer Graphics 26*, 2 (July 1992), 265 – 272.

[WBSS04] Wang Z., Bovik A., Sheikh H., Simoncelli E.: Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing 13*, 4 (2004), 600–612.

[WGK14] Weinmann M., Gall J., Klein R.: Material classification based on training data synthesized using a BTF database. In *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part III* (2014), Springer International Publishing, pp. 156–171.

[WH01] Wichmann F. A., Hill N. J.: The psychometric function: I. fitting, sampling, and goodness of fit. *Perception & Psychophysics 63*, 8 (2001), 1293–1313.

[WKB14] Ward G., Kurt M., Bonneel N.: Reducing anisotropic BSDF measurement to common practice. In *Proceedings of the 2nd Eurographics Workshop on Material Appearance Modeling: Issues and Acquisition* (2014), MAM '14, pp. 5–8.

[WM97] Webster M. A., Mollon J.: Adaptation and the color statistics of natural images. *Vision Research 37*, 23 (1997), 3283 – 3298.

[WM01] Westlund H., Meyer G.: *Applying appearance standards to light reflection models.* Association for Computing Machinery, 2001, pp. 501–510.

[WSB*98] White D. R., Saunders P., Bonsey S. J., van de Ven J., Edgar H.: Reflectometer for measuring the bidirectional reflectance of rough surfaces. *Journal of Applied Optics 37* (June 1998), 3450–3454.

[WSRK11] Weinmann M., Schwartz C., Ruiters R., Klein R.: A multi-camera, multi-projector super-resolution framework for structured light. In *International Conference on 3DIMPVT, 2011* (2011), pp. 397–404.

[WZT*08] Wang J., Zhao S., Tong X., Snyder J., Guo B.: Modeling anisotropic surface reflectance with example-based microfacet synthesis. *ACM Transactions on Graphics 27*, 3 (2008), 41:1–41:9.

[XNY*16] Xu Z., Nielsen J., Yu J., Jensen H., Ramamoorthi R.: Minimal BRDF sampling for two-shot near-field reflectance acquisition. *ACM Transactions on Graphics 35*, 6 (2016).

[XSD*13] Xu K., Sun W.-L., Dong Z., Zhao D.-Y., Wu R.-D., Hu S.-M.: Anisotropic spherical gaussians. *ACM Transactions on Graphics 32*, 6 (2013), 209:1–209:11.

[Zha99] Zhang Z.: Flexible camera calibration by viewing a plane from unknown orientations. In *Proceedings of ICCV* (1999), vol. 1, pp. 666–673.

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

Reviewed Publications of the Author Relevant to the Thesis

[A.1] Radomír Vávra and Jiří Filip. *Minimal Sampling for Effective Acquisition of Anisotropic BRDFs.* Computer Graphics Forum **35**:7, Proceedings of Pacific Graphics 2016, pp. 299-309, 2016.

[A.2] Jiří Filip, Radomír Vávra, Michal Havlíček and Mikuláš Krupička. *Predicting Visual Perception of Material Structure in Virtual Environments.* To appear in Computer Graphics Forum, 2015.

[A.3] Jiří Filip, Michal Havlíček and Radomír Vávra. *Adaptive Highlights Stencils for Modeling of Multi-Axial BRDF Anisotropy.* The Visual Computer, pp. 1-11, Springer, 2015.

The paper has been cited in:

- Guarnera D., Guarnera G.C., Ghosh A., Denk C., Glencross M. *BRDF Representation and Acquisition.* Computer Graphics Forum **35**:2, pp. 625-650, 2016.
- Guarnera D., Guarnera G.C., Ghosh A., Hall I. and Glencross M. *Capturing and Representing BRDFs for Virtual Reality.* SIGGRAPH Asia Courses, 2016.

[A.4] Jiří Filip, Radomír Vávra and Mikuláš Krupička. *Rapid Material Appearance Acquisition Using Consumer Hardware.* Sensors **14**:10, pp. 19785-19805, 2014.

The paper has been cited in:

- Steinhausen H.C., Martín R., Brok D., Hullin M.B. and Klein R. *Extrapolation of Bidirectional Texture Functions using Texture Synthesis Guided by Photometric Normals.* IS&T/SPIE Electronic Imaging, pp. 93980A-93980A, International Society for Optics and Photonics, 2015.

- Brok D., Steinhausen H.C., Hullin M.B. and Klein R. *Multiplexed Acquisition of Bidirectional Texture Functions for Materials*. IS&T/SPIE Electronic Imaging, pp. 93980F–93980F, International Society for Optics and Photonics, 2015.
- Steinhausen H.C., Brok D., Hullin M.B. and Klein R. *Extrapolating Large-Scale Material BTFs under Cross-Device Constraints*. Vision, Modeling & Visualization, pp. 143-150, The Eurographics Association, 2015.
- Guarnera D., Guarnera G.C., Ghosh A., Denk C. and Glencross M. *BRDF Representation and Acquisition*. Computer Graphics Forum **35**:2, pp. 625-650, 2016.
- Guarnera D., Guarnera G.C., Ghosh A., Hall I. and Glencross M. *Capturing and Representing BRDFs for Virtual Reality*. SIGGRAPH Asia Courses, 2016.
- Dong W. and Shen H.-L. *Sparse Acquisition and Reconstruction of Bidirectional Texture Functions*. Zhejiang Daxue Xuebao (Gongxue Ban)/Journal of Zhejiang University (Engineering Science) **50**:12, pp. 2364-2370, 2016.
- Dong W., Shen H.-L., Du X., Shao S.-J. and Xin J.H. *Spectral Bidirectional Texture Function Reconstruction by Fusing Multiple-Color and Spectral Images*. Applied Optics **55**:36, pp. 10400-10408, 2016.
- Zhang X., Yang L., Cui S. and Wu T. *YUV-Based Material Modeling from a Single Image*. In 9th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI), Datong, pp. 801-805, 2016.
- Zhang X., Cui S., Cui H., Yang L. and Wu T. *Material Appearance Acquisition from a Single Image*. Proceedings of SPIE - The International Society for Optical Engineering, volume 10322, 2017.
- Hošek J., Havran V., Němcová Š., Bittner J. and Čáp J. *Optomechanical Design of a Portable Compact Bidirectional Texture Function Measurement Instrument*. Applied Optics **56**:4, pp. 1183-1193, 2017.
- Havran V., Hošek J., Němcová Š., Čáp J. and Bittner J. *Lightdrum–Portable Light Stage for Accurate BTF Measurement on Site*. Sensors **17**:3, 423, 2017.

[A.5] Jiří Filip and Radomír Vávra. *Template-Based Sampling of Anisotropic BRDFs*. Computer Graphics Forum **33**:7, Proceedings of Pacific Graphics 2014, pp. 91-99, 2014.

The paper has been cited in:

- Dupuy J., Heitz E., Iehl J.C., Poulin P. and Ostromoukhov V. *Extracting Microfacet-Based BRDF Parameters from Arbitrary Materials with Power Iterations*. Computer Graphics Forum **34**:4, pp. 21-30, 2015.

- Macdonald L.W. *Realistic Visualisation of Cultural Heritage Objects*. Dissertation Thesis, University College London, 2015.
- Bonneel N., Peyré G. and Cuturi M. *Wasserstein Barycentric Coordinates: Histogram Regression Using Optimal Transport*. ACM Transactions on Graphics **35**:4, Proceedings of SIGGRAPH 2016, pp. 71:1-71:10, 2016.
- Yu C. and Lee S.W. *Reliable Estimation of Anisotropic BRDF*. Electronics Letters **52**:12, pp. 1024-1026, 2016.
- Natola F., Ntouskos V., Pirri F. and Sanzari M. *Single Image Object Modeling Based on BRDF and r-surfaces Learning*. In proceedings of the 2016 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pp. 4414-4423, 2016.
- Guarnera D., Guarnera G.C., Ghosh A., Denk C., Glencross M. *BRDF Representation and Acquisition*. Computer Graphics Forum **35**:2, pp. 625-650, 2016.
- Guarnera D., Guarnera G.C., Ghosh A., Hall I. and Glencross M. *Capturing and Representing BRDFs for Virtual Reality*. SIGGRAPH Asia Courses, 2016.
- Serrano A., Gutierrez D., Myszkowski K., Seidel H. P. and Masia B. *An Intuitive Control Space for Material Appearance*. ACM Transactions on Graphics **35**:6, pp. 186:1-186:12, 2016.
- Xue J., Zhang H., Dana K. and Nishino K. *Differential Angular Imaging for Material Recognition*. CoRR, 2016.
- Zubiaga C.J. *Real-Time 2D Manipulation of Plausible 3D Appearance using Shading and Geometry Buffers*. Dissertation Thesis, Université de Bordeaux, 2016.

[A.6] Jiří Filip and Radomír Vávra. *Fast Method of Sparse Acquisition and Reconstruction of View and Illumination Dependent Datasets.* Computers & Graphics **37**:5, pp. 376-388, Elsevier, 2013.

The paper has been cited in:

- Dong W. and Shen H.-L. *Sparse Acquisition and Reconstruction of Bidirectional Texture Functions*. Zhejiang Daxue Xuebao (Gongxue Ban)/Journal of Zhejiang University (Engineering Science) **50**:12, pp. 2364-2370, 2016.

[A.7] Michal Haindl, Matěj Sedláček and Radomír Vávra. *Virtual Reconstruction of Cultural Heritage Artifacts*. ERCIM News, No. 103, pp. 12-13, 2015.

[A.8] Michal Haindl, Jiří Filip and Radomír Vávra. *Digital Material Appearance: the Curse of Tera-Bytes*. ERCIM News, No. 90, pp. 49-50, 2012.

The paper has been cited in:

- Yengui M.Y. and Poulin P. *Simulated Bidirectional Texture Functions with Silhouette Details*. Proceedings of Graphics Interface 2013, pp. 45-54, Canadian Information Processing Society, 2013.
- Schwartz C., Sarlette R., Weinmann M. and Klein R. *DOME II: a Parallelized BTF Acquisition System*. Proceedings of the Eurographics 2013 Workshop on Material Appearance Modeling: Issues and Acquisition, pp. 25-31, Eurographics Association, 2013.
- Dong W. and Shen H.-L. *Sparse Acquisition and Reconstruction of Bidirectional Texture Functions*. Zhejiang Daxue Xuebao (Gongxue Ban)/Journal of Zhejiang University (Engineering Science) **50**:12, pp. 2364-2370, 2016.
- Guarnera D., Guarnera G.C., Ghosh A., Denk C., Glencross M. *BRDF Representation and Acquisition*. Computer Graphics Forum **35**:2, pp. 625-650, 2016.
- Guarnera D., Guarnera G.C., Ghosh A., Hall I. and Glencross M. *Capturing and Representing BRDFs for Virtual Reality*. SIGGRAPH Asia Courses, 2016.
- Zhang X., Yang L., Cui S. and Wu T. *YUV-Based Material Modeling from a Single Image*. In 9th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI), Datong, pp. 801-805, 2016.
- Zhang X., Cui S., Cui H., Yang L. and Wu T. *Material Appearance Acquisition from a Single Image*. Proceedings of SPIE - The International Society for Optical Engineering, volume 10322, 2017.

[A.9] Jiří Filip, Radomír Vávra and Michal Haindl. *Capturing Material Visualization Data Using Goniometers*. In proceedings of the 4th CIE Expert Symposium on Colour and Visual Appearance, pp. 121-127, 2016.

[A.10] Jiří Filip and Radomír Vávra. *Using Reflectors for Analysis and Acquisition of Anisotropic BRDF*. In proceedings of the 4th CIE Expert Symposium on Colour and Visual Appearance, pp. 155-162, 2016.

[A.11] Radomír Vávra and Jiří Filip. *BRDF Interpolation using Anisotropic Stencils*. In Proceedings of IS&T International Symposium on Electronic Imaging **2016**:9, pp. 1-6, Society for Imaging Science and Technology, 2016.

[A.12] Jiří Filip and Radomír Vávra. *Anisotropic Materials Appearance Analysis Using Ellipsoidal Mirror*. In Proceedings of IS&T/SPIE Conference on Measuring, Modeling, and Reproducing Material Appearance, pp. 93980P:1-93980P:9, 2015.

[A.13] Jiří Filip, Radomír Vávra and Michal Havlíček. *Effective Acquisition of Dense Anisotropic BRDF*. In Proceedings of the 22nd International Conference on Pattern Recognition, ICPR, pp. 2047-2052, IEEE Computer Society, 2014.

The paper has been cited in:

- Raymond B., Guennebaud G. and Barla P. *Multi-Scale Rendering of Scratched Materials using a Structured SV-BRDF Model.* ACM Transactions on Graphics **35**:4, pp. 57:1-57:11, 2016.
- Fichet A., Sato I. and Holzschuch N. *Capturing Spatially Varying Anisotropic Reflectance Parameters using Fourier Analysis.* Graphics Interface 2016.

[A.14] Jiří Filip, Radomír Vávra, Michal Haindl, Pavel Žid, Mikuláš Krupička and Vlastimil Havran. *BRDF Slices: Accurate Adaptive Anisotropic Appearance Acquisition.* In Proceedings of the 2013 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR, pp. 1468-1473, IEEE Computer Society, 2013.

The paper has been cited in:

- Schwartz C., Sarlette R., Weinmann M., Rump M. and Klein R. *Design and Implementation of Practical Bidirectional Texture Function Measurement Devices Focusing on the Developments at the University of Bonn.* Sensors **14**:5, pp. 7753-7819, Multidisciplinary Digital Publishing Institute, 2014.
- Inoshita C. *Shape from Scattering: Shape Estimation Based on Light Transport Analysis in Translucent Objects.* Dissertation Thesis, Osaka University, 2015.
- Weinmann M. and Klein R. *Advances in Geometry and Reflectance Acquisition.* Course Notes, SIGGRAPH Asia 2015 Courses, pp. 1:1-1:71, 2015.
- Dana K.J. *Capturing Computational Appearance: More than Meets the Eye.* IEEE Signal Processing Magazine **33**:5, pp. 70-80, 2016.
- Raymond B. *Contrôle de l'apparence des matériaux anisotropes.* Dissertation Thesis, Université de Bordeaux, 2016.

[A.15] Radomír Vávra and Jiří Filip. *Registration of Multi-View Images of Planar Surfaces.* 11th Asian Conference on Computer Vision, ACCV 2012, Revised Selected Papers, pp. 497-509, Springer Berlin Heidelberg, 2013.

The paper has been cited in:

- Klein J. *Correction of Longitudinal Aberrations in Goniometric Measurement with a Multispectral Camera.* 19th Workshop Farbbildverarbeitung, pp. 31-41, 2013.
- Klein J. and Schmücking G. *Analysis of Aberrations and Pixel Information in Goniometric Multispectral Imaging.* IS&T/SPIE Electronic Imaging, pp. 901806-901806, International Society for Optics and Photonics, 2014.
- Havran V. *Geometric accuracy analysis of stationary BTF gonioreflectometers.* Workshop on Material Appearance Modeling, MAM 2015, pp. 25-29, 2015.
- Havran V. *Practical Experiences with Using Autocollimator for Surface Reflectance Measurement.* Workshop on Material Appearance Modeling, MAM 2016, 2016.

- Havran V., Hošek J., Němcová Š., Čáp J. and Bittner J. *Lightdrum–Portable Light Stage for Accurate BTF Measurement on Site*. Sensors **17**:3, 423, 2017.

Remaining Publications of the Author Relevant to the Thesis

[A.16] Radomír Vávra, Jiří Filip and Petr Somol. *A Comparison of Adaptive Sampling and Interpolation of 2D BRDF Subspaces.* Research Report 2339. UTIA CAS, Prague 2013.

Appendix **A**

Published BRDF and BTF Datasets

During last years, we measured and made several datasets available for research purposes. The datasets are available at *http://btf.utia.cas.cz*. In this appendix, we describe these datasets and related data processing approaches. First, there are six BTF datasets (see Sec. A.1). Next, we published very precise BRDF measurement of few subspaces of one fabric material (see Sec. A.2). As we had found those precise and angularly-dense BRDF data very useful for research of adaptive measurement, we measured three more fabric materials using an approach described in Section 4.3. The measured materials are described in Section A.3. Additionally, we published a database of 150 anisotropic BRDFs measured in rather lower angular resolution (see Sec. A.4) and we measured three fabric materials using the BRDF slices described in Chapter 5 (see Sec. A.5).

A.1 BTF Datasets

Using the UTIA gonioreflectometer described in Chapter 4, we measured and made available six BTFs of interesting materials (see Tab. A.1). Their renderings on a sphere are depicted in Figure A.1. The directional resolution of the samples is 81 illumination $\times$ 81 viewing directions and the spatial resolution is about 512×512 pixels (see Tab. A.1), while the spatial density is 1071 DPI. The dynamic range of the datasets is up to 21 EV and data are provided in the OpenEXR format in the CIE XYZ colorspace and also in the PNG format in the sRGB colorspace. The ABRDFs created by averaging pixels of the BTFs are provided as well.

A.2 Anisotropic BRDF Subspaces

To make research of adaptive measurement possible for even those who do not have access to a gonioreflectometer, we measured and published the very angularly-dense measurement of one fabric material; fabric02. However, as a dense measurement of the whole four-dimensional angular space would be impossible, we had decided to measure only sev-

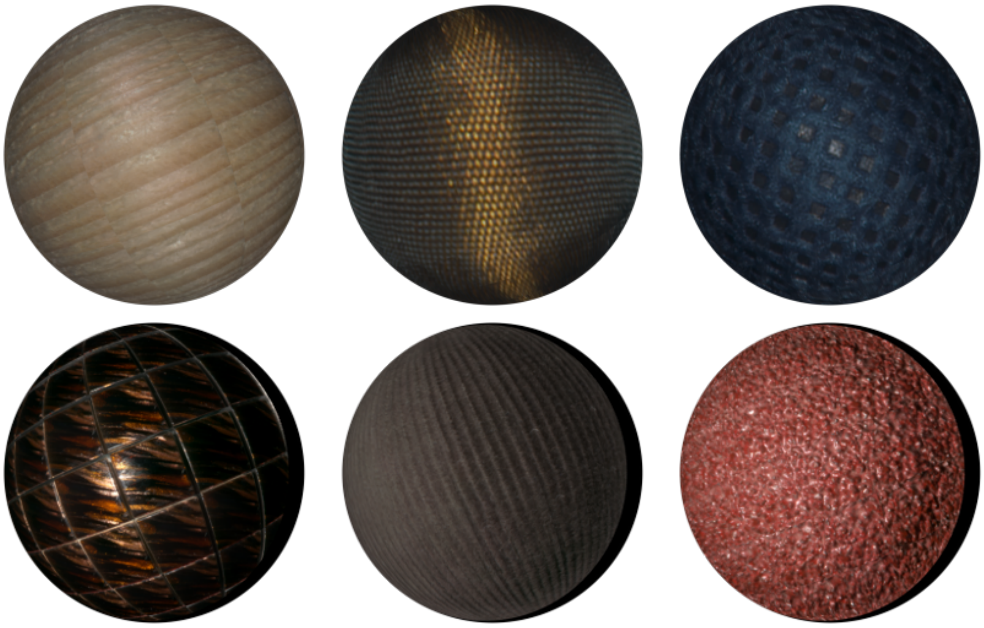

Figure A.1: Renderings of six BTF datasets on spheres for point-light illumination.

Table A.1: List of the materials whose BTF datasets are measured.

material	description	spatial resolution
wood01	raw spruce wood	681×717
fabric02	woven fabric, two threads	512×512
fabric03	non-woven fabric	512×512
tile01	bathroom wall tile	877×865
corduroy01	corduroy fabric	545×553
sandpaper01	raw sandpaper	489×477

eral two-dimensional subspaces. We selected three elevation angles $\theta = \{0°, 30°, 75°\}$ of illumination and viewing directions and measured nine subspaces resulting from all the combinations. Each subspace is measured using the angular resolution of 0.5° in both azimuthal angles.

We selected a fabric material consisting of two perpendicular, finely interwoven, yellow and gray fibers (see Fig. A.2-a). This material provides an intricate golden appearance with a strong anisotropic behavior as shown in example images of various illumination and viewing conditions in Figure A.2-b. The resulting data were used for research of adaptive measurement using the BRDF slices (see Chap. 5). The data are shown in Figure A.3.

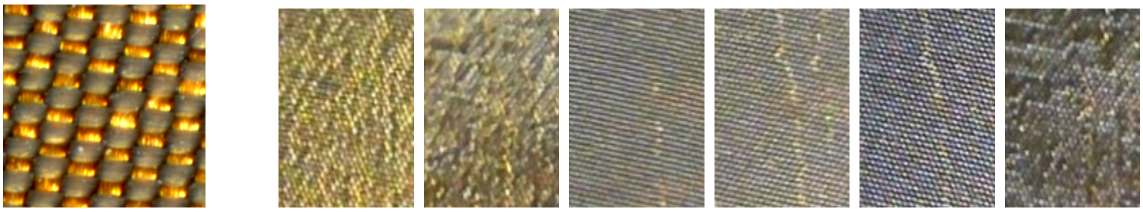

(a) (b)

Figure A.2: A high resolution scan of the fabric material (a), appearance of the material for different illumination and viewing conditions (b).

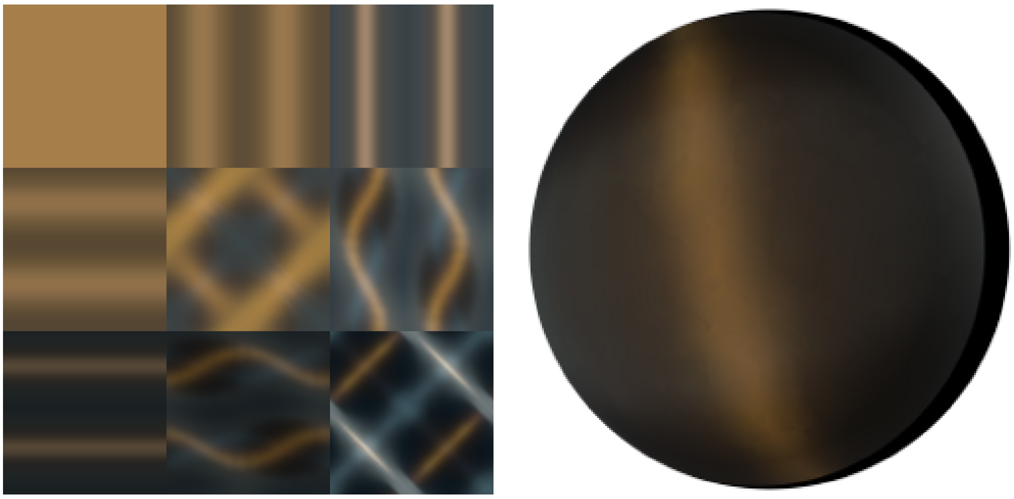

Figure A.3: The dense BRDF of the anisotropic fabric material (left) and its rendering on a sphere (right).

A.3 Dense Anisotropic BRDF Datasets

We measured three highly anisotropic materials as shown in Figure A.4. The material *fabric112* consists of two types of threads (gold and dark-green), which produce distinct anisotropic behavior. The threads are interwoven in a plain pattern so that the proportion of visibility of the threads is about 40% vs. 60%. The remaining two satin-like materials (*fabric135* and *fabric136*) consist of a combination of a diffuse dark-brown/red thread and tiny transparent nylon fibers that cover almost 40% of the surface. Real material specimens can be provided upon request. The rendering of the measured BRDF in an environmental illumination represented by 512 directional lights is shown in Figure A.5.

Each BRDF dataset is stored in an array of dimensions $|\theta_i| \times |\theta_v| \times |\varphi_i| \times |\varphi_v| = 44 \times 44 \times 180 \times 180$, i.e., with the angular step of $2°$ and the maximal elevation angle is $86°$. Size of each BRDF in OpenEXR HDR format is 700 MB. As the size of the dataspace is enormous,

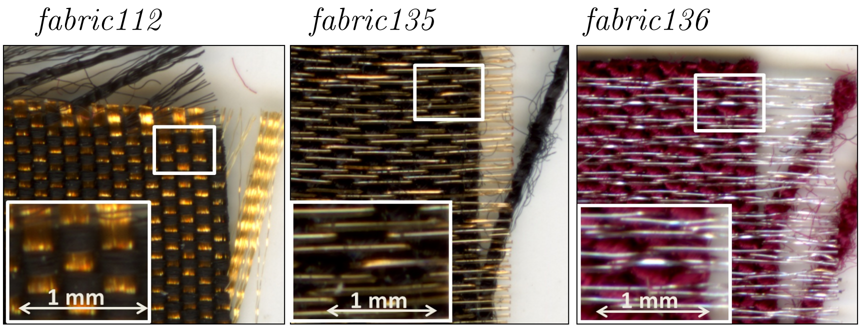

Figure A.4: Detail of the structure of the measured materials.

Figure A.5: Renderings of the measured dense BRDFs in an environmental illumination.

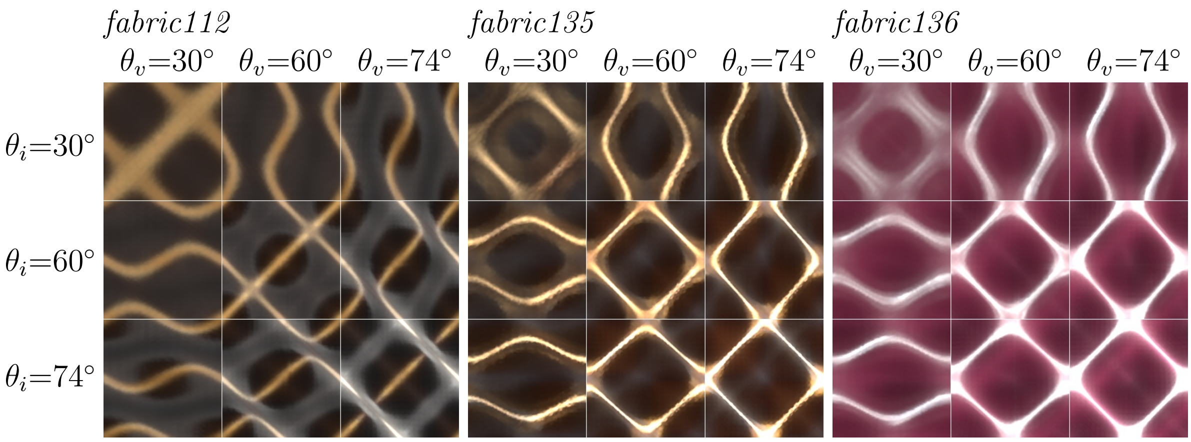

Figure A.6: Azimuth subspaces for nine different combinations of illumination and viewing directions of the three measured materials.

we show only nine azimuthal (φ_i, φ_v) subspaces in Figure A.6. The nine subspaces result from all the combinations of defined elevation angles $\theta_i, \theta_v \in \{30°, 60°, 74°\}$.

The materials *fabric135* and *fabric136* exhibit the strong anisotropic highlights when azimuth of illumination is perpendicular to nylon fibers that reflect light off their smooth surface (see Fig. A.4). The specular highlights are of comparable intensity.

The material *fabric112* exhibits a more complicated type of anisotropy (see Fig. A.4). While both specular and anisotropic highlights are comparable at low elevations (top-left subspace), with an increasing elevation two types of the anisotropic highlights appear: the yellow one for the gold thread and the gray one for the dark thread. Due to mutual occlusion of both threads the intensity of the specular highlights at high elevations depend on orientation of the material owing to a camera and a light source and its color is determined by the prevailing dark thread.

A.4 The Database of 150 Anisotropic BRDFs

We introduce a new BRDF database that contains 150 measurements of a wide range of materials (see Fig. A.7). Contrary to the majority of the past research we deal with full four-dimensional anisotropic BRDF measurements. In addition to fabrics, we measured leather, plastic, wood, paper, plaster, etc. As seen in Table A.2, we focused ourselves on anisotropic materials, which can be categorized into groups and summarized according to strength of anisotropy, the number of the anisotropic highlights, and presence of two distinct colors.

We aim at collecting as many types of anisotropic behaviors as possible. To achieve high variability we collect BRDFs of nearly homogeneous samples of leather, plastic, smooth fabric. Naturally, fabric materials present the highest variability of appearance due to individual thread-based weaving patterns; therefore, they constitute the biggest group of the measured materials. We avoided materials with macroscopic features. The most materials exhibit anisotropy caused by their small-scale geometry.

First, all images taken for all combinations of viewing and illumination directions are pixel-wisely registered and rectified. Then, the representative area of the real size of 11×11 mm is cropped and its values are spatially averaged. One of the advantages of spatial averaging in BRDF acquisition is suppression of a noise in the data. An example of a typical mean measurement noise can be observed by measurement of two specimens of the same corduroy fabric. When we compute a difference between their captured samples (prior to any processing) we obtain the PSNR of 79.03 dB.

All data are measured using the UTIA gonioreflectometric device (Sec. 4.1). The uniform directional resolution of $81 \times 81 = 6,561$ directions [SSK03] of our BRDFs is relatively sparse; however, for the majority of the materials it is sufficient to capture their behavior as the highlights are not too narrow.

Measurements at all elevations have the same level of reliability. All BRDFs are recorded as float values and the four dimensions of the anisotropic BRDF are stored as 2D images divided into the azimuthal BRDF subspaces for fixed elevation angles. Because the measurement of over six thousand directions takes around eighteen hours, we perform batch measurements of nine materials simultaneously, each spanning a region of 3×3 cm, where varying directions across patches are compensated. As we record all combinations

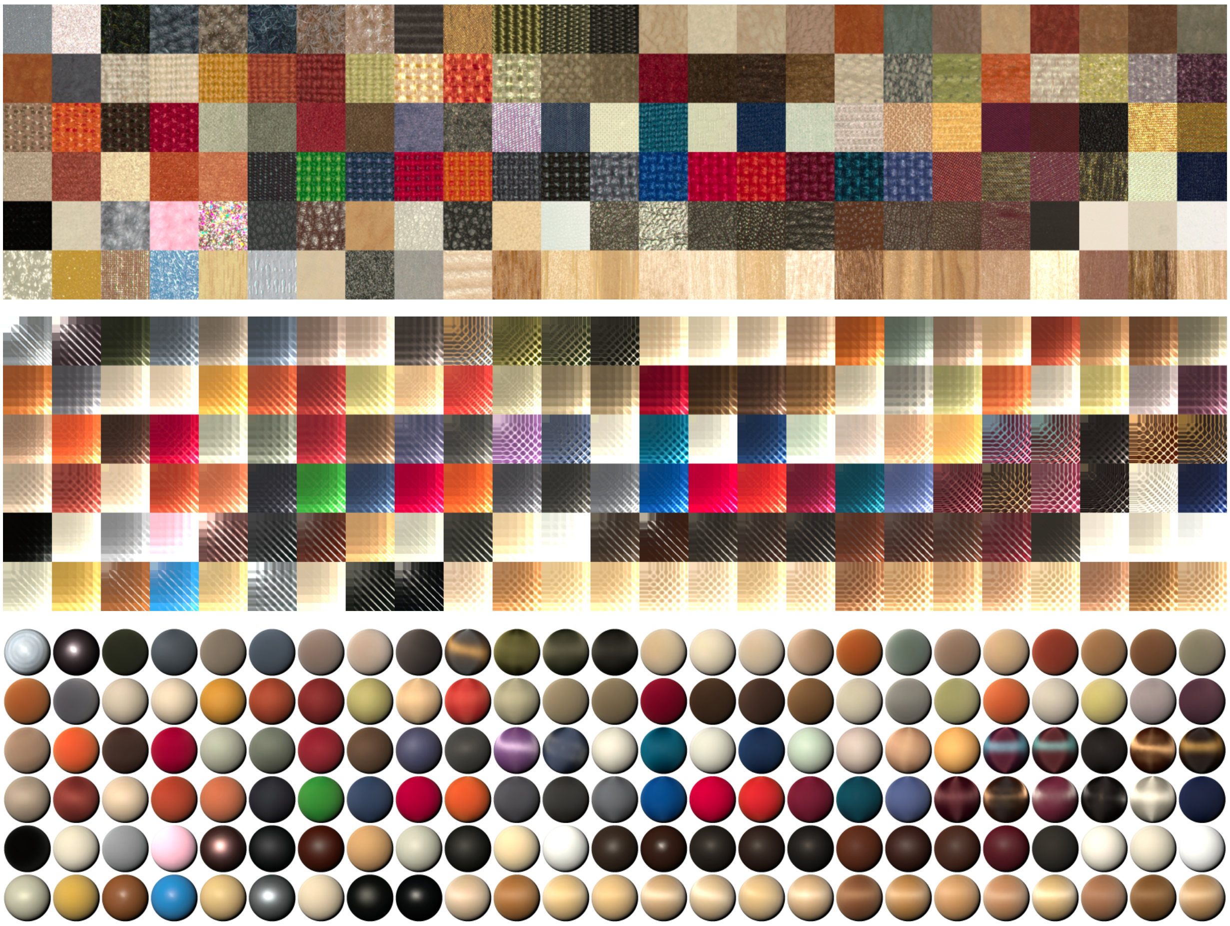

Figure A.7: Arrangements of textures (top), measured BRDFs (middle) and renderings on a sphere (bottom) of 150 materials. A spatial area of each materials is 11×11 mm. Each BRDF consists of 6×48 illumination directions and 6×48 viewing directions.

of illumination and viewing directions, we compensate foreshortening of the illuminated area in relation to an increasing elevation and finally, we enforce reciprocity, which helps to suppress a possible source of a noise in the measurement process.

The BRDF measurement for the hemispherically-uniform sampling of 81×81 directions results in a variable count of samples at different elevations. In order to obtain the constant count of samples needed for further analysis, we use an interpolation based on the radial basis functions [PTVF92] as explained in Section A.4.1. As a result, we obtain $6 \times 48 \times 6 \times 48 = 288 \times 288 = 82,944$ directions. All raw and interpolated measured BRDFs are available in the OpenEXR format and samples of some materials may be provided up on request.

Table A.2: Categorization of the measured materials. We report strength of anisotropy, the number of the anisotropic highlights, and presence of two distinct colors in the BRDF.

material	*total*	*anisotropy*		*anis. HLs*		*two*
category	*mat.*	*strong*	*weak*	*one*	*two*	*colors*
fabric	**96**	21	51	41	31	6
wood	**16**	16	-	16	-	-
leather	**16**	-	2	2	-	-
carpet	**6**	-	2	2	-	-
plastic	**6**	-	1	1	-	-
paper	**2**	-	-	-	-	-
wallpaper	**2**	-	-	-	-	-
plaster	**2**	-	-	-	-	-
others	**4**	-	1	1	-	-
total	**150**	37	57	63	31	6

A.4.1 Uniform Resampling and Interpolation of the BRDF

As the number of measured samples is insufficient for a proper sampling of more specular BRDFs, we perform interpolation of each azimuthal subspace to the resolution of 48 × 48 directions (angular resolution 7.5°). In contrast to the MERL database [MPBM03a], the number of specular materials in our database is limited; therefore, we could afford a significantly lower sampling density. To improve the continuity of highlights we detect positions of the specular and anisotropic highlights and perform the linear interpolation just over values along the highlights. Obtained data are then added to measured samples and used for the final interpolation based on the two-dimensional radial basis functions [PTVF92]. As a result, we obtain the smoother specular and anisotropic highlights instead of discontinuous ones as shown in Figure A.8.

Although the interpolation artificially modifies raw BRDF measurements, it allows us to obtain a database of relatively dense measurements for further analysis. We believe

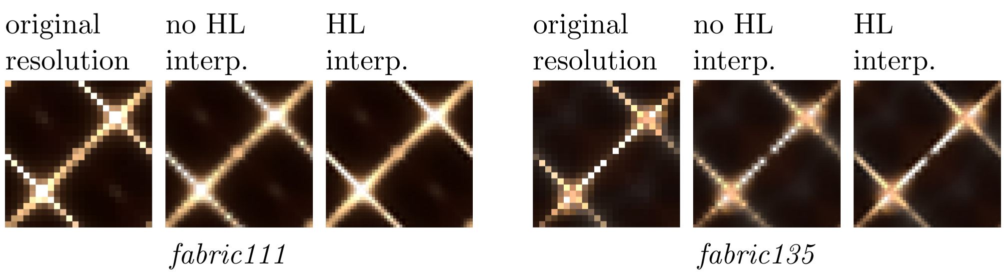

Figure A.8: The interpolation of two BRDFs – effect of the linear interpolation of the specular and anisotropic highlights (HL).

that our modifications are physically reasoned as the original measurements are altered to preserve a typical smooth behavior of the specular highlights. Each measured BRDF is finally represented by $1 + 5 \times 48 + (5 \times 48 \times 5 \times 48)/2 + (5 \times 48)/2 = 29,161$ reciprocal samples.

A.4.2 Alignment of the Anisotropic Highlights

Contrary to analysis of the isotropic BRDF where locations of the specular highlights can be predicted, the analysis of the anisotropic BRDF is more challenging. The number and location of the anisotropic highlights are unknown and depend entirely on the initial orientation of the measured material and on its properties. Therefore, before different anisotropic BRDFs can be compared, they must be first aligned according to the main anisotropic axis.

Wang et al. [WZT*08] solved a related problem of alignment and merging of normal distribution functions corresponding to different surface locations of the same material. In contrast to their approach, we tackle a more general problem of the BRDF alignment of materials that have completely different types of anisotropic behavior.

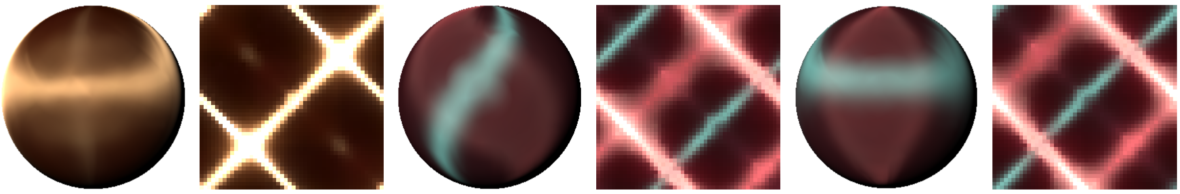

template misaligned aligned

Figure A.9: The anisotropic alignment of a subspace (elevations 75°/75°) of the material *fabric107* according to the template defined by the material *fabric111*.

The procedure is straightforward when we operate in a subspace of azimuthal angles (see Fig. A.9). As the subspace is circularly symmetric the circular shift of the subspace in diagonal direction has the same effect as the change of initial rotation of the measured material prior to the measurement. As the correct detection of the main anisotropic highlight is crucial for our experiments with the BRDF sampling (see Chap. 7), we test several approaches. The best results are obtained by an approach that uses a shiny anisotropic material (*fabric111*) with the distinct specular and anisotropic highlights as a template T. A query material B is aligned using shift $\hat{s}$ to this template by minimizing L1 norm between the query and the template subspaces across all possible circular shifts:

$$\hat{s} = \arg\min_{s=0\ldots\pi} \sum_{\varphi_i=0}^{2\pi} \sum_{\varphi_v=0}^{2\pi} |T(\varphi_i, \varphi_v) - B(\mathsf{circshift}(\varphi_i, \varphi_v, -s))| \ .$$

A.5 The Adaptively Measured Anisotropic BRDFs

Chapter 5 deals with adaptive measurement of the BRDF using the BRDF slices. During our experiments, we measured several datasets of three anisotropic materials; *fabric112*, *fabric135* and *fabric136*. For detailed description of the materials see Section A.3.

Each material was measured uniformly by $8,911$ samples, adaptively by $8,911$ samples, uniformly by $18,721$ samples and adaptively by $18,721$ samples. Together with the data and lists of measured directions, we provide four-dimensional arrays filled with interpolated values. Dimensions of the arrays are $|\theta_i| \times |\theta_v| \times |\varphi_i| \times |\varphi_v| = 41 \times 41 \times 180 \times 180$, i.e., the angular step is $2°$ and the maximal elevation angle is $80°$. Uniformly measured data are interpolated using barycentric weights [Cox69] or RBFs [PTVF92]. Adaptively measured data are interpolated using the proposed reconstruction method described in Section 5.3.

Moreover, we measured ground-truth data for each material to evaluate precision of the datasets, which are also provided. We compare data in an applied situation using the 3D scene described in Section 5.5.3 and depicted in Figure 5.10. Figure 5.11 depicts the scene rendered using the datasets together with difference values.

Appendix **B**

A Study on Visual Perception of Material Structure

This appendix complements Section 3.3 that summarizes endeavors of the paper [A.2] that deals with prediction of visibility of structural elements in virtual environments. While Section B.1 describes 25 materials used in the research and divides them into seven groups, Section B.2 introduces a virtual scene consisting of cylinders covered with the materials. Next, Section B.3 overviews a post-processing of the obtained psychophysical data and finally, Section B.4 confirms validity of the web-based *Experiment B*, summarizes our findings about response times of tested subjects, and discusses influence of anisotropy on the estimated distance.

B.1 Tested Materials

We test 25 materials out of categories often used in interior design (carpet, fabric, leather, and wood). Their BTFs were measured by the UTIA gonioreflectometer (see Sec. 4.1) at an angular resolution of 81×81 directions, and a spatial resolution of 350 DPI. Examples of materials are shown in Figure B.1. For the sake of seamless coverage of the tested objects in required spatial resolution, we estimated the single seamless repeatable tile using the approach proposed in [SH05]. Sizes of tiles range from 80 to 300 pixels.

As some of the tiles were not entirely spatially uniform in hue and luminance, we carefully removed their lowest frequency components in Fourier space. This step avoids visually distractive repetitions, that could introduce low frequencies that are not present in the original material. The approximate BRDFs of the materials were obtained by spatial averaging across the area of a tile. We consider this approximation as reasonably accurate as the size of the biggest tile is 22 mm, which is negligible when compared to the distances between the sample and a light source or a camera which are 1.3 m and 2.1 m, respectively.

Next, we visually estimated the physical size of the biggest structural elements of the tested materials. The results in mm are shown in Figure B.2. We further analyzed the

Figure B.1: An overview of 25 materials used in our study (listed alphabetically) and their distribution into seven analyzed groups.

height of the materials by integration of normals resulting from the over-determined photometric stereo using 80 illumination directions (frontal illumination was removed to reduce presence of the specular reflections). Although one could argue that we can compute variance of distribution of slopes of local facets to determine the size of material structure, we consider this approach unreliable, especially for fabric materials with translucent mesostructure. For the sake of further analysis, we divide materials into seven categories based on the type of a material, and the size and the height of its structure:

G1: *fabric smooth* group comprises materials *fabric082*, *fabric102*, *fabric106*, and *fabric110*. This group represents mostly the apparel fabric materials with a fine structure where the typical size of the biggest structural element is below 1 mm.

G2: *fabric meso* group comprises materials *fabric003*, *fabric041*, *fabric048*, *fabric075*, *fabric120*, *fabric122*, *fabric129*, and *fabric131*. This group represents the typical upholstery materials (used on, e.g., office chairs) where the size of the biggest structural element is below 3 mm.

G3: *fabric rough* group comprises materials *carpet01*, *carpet04*, and *fabric146*. It represents very rough fabric materials: office carpet and a crocheted sweater with the structure size above 4 mm.

G4: *leather meso* group comprises materials *leather01*, *leather07* (both synthetic), and

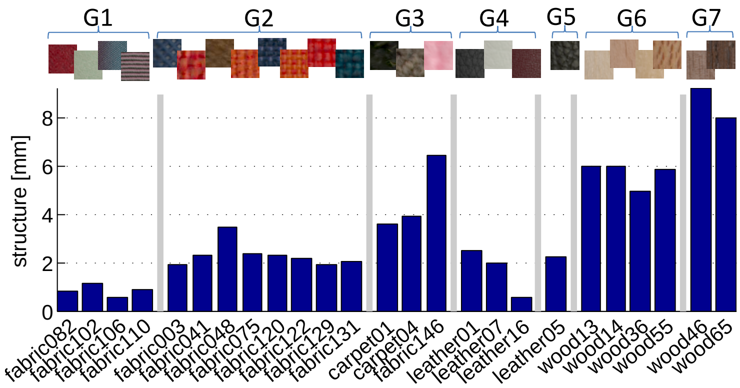

Figure B.2: Categorization of the studied materials into seven groups based on the size of their biggest structural element.

leather16 (genuine). This category comprises the relatively flat leather materials (often used in car interior design or upholstery) with the structure size up to 2 mm.
G5: *leather rough* group comprises single rough genuine leather material *leather05*. We designed a special category for this material due to its higher surface height that sets its appearance apart from the leather materials in the previous group.
G6: *wood meso* group includes *wood13* (beech), *wood14* (steamed beech), *wood36* (european lime), and *wood55* (plane). These wood materials exhibit smooth surfaces, smaller structural elements (below 6 mm), and less contrasted appearance.
G7: *wood rough* group includes material *wood46* (american walnut) and *wood65* (wenge). This group represents contrasting wood samples with characteristic and long vertical grooves of the size above 8 mm.

B.2 Test Scene

As the basic notion of our study is the psychophysical detection of the critical viewing distance where the structure becomes visually indistinguishable, we altered the task into an analysis of the difference between the BRDF and the BTF as a function of viewing distance. Therefore, we generated a virtual scene consisting of a set of consequently receding cylinders as shown in Figure B.3. The upper part of the cylinders shows the BRDF while the lower part shows the BTF of the same material. We opted for illumination by means of two directional light sources; one from the right (intensity 1.0) and the second from the left (intensity 0.3). Such configuration guarantees visibility of the structure on all visible areas of cylinders and preserves contrasting shadowing in the material structure that would be smoothed by a typical environmental illumination.

We selected a cylinder due to its clearly defined texture mapping without the need for warping or cutting of the mapped material. The perspective projection of the scene

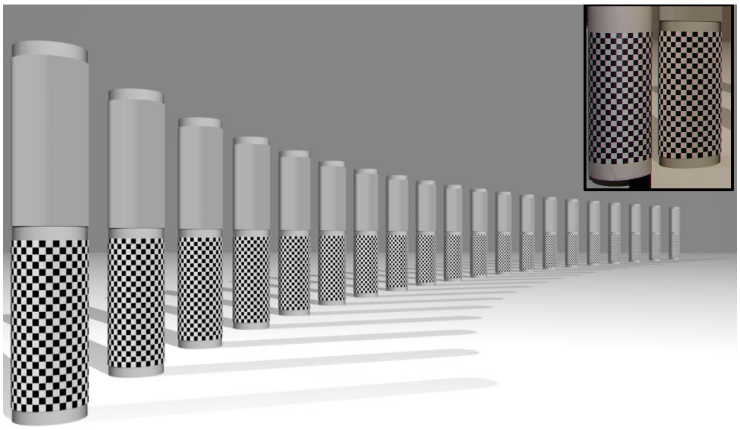

Figure B.3: An example of the virtual scene consisting of a set of consequently receding cylinders showing the BRDF and the BTF of the same material. Top-right: a side-by-side comparison of a real cylinder (left) with its representation on the screen (right). Both have the same width and a checkerboard pattern whose size is 5 mm.

is calibrated so as its appearance resembles the same physical template. See the set of photos of a physical cylinder of the defined geometry in Figure B.4. The distance of the virtual camera is set at 0.6 m from the orthogonal plane where the first cylinder appears. Assuming the full-frame sensor of the height of 24 mm, we set the camera focal distance to 45 mm, which corresponds to human vision. This translates to a vertical viewing angle of approximately 30^o.

The diameter of all cylinders is 59 mm, and the height of a material on its surface is 120 mm; their viewing distances range from 0.6 m to 4.4 m. A side-by-side comparison of a real 59 mm wide cylinder with its rendering observed from a distance of 0.6 m is shown in an inset at the top-right corner of Figure B.3.

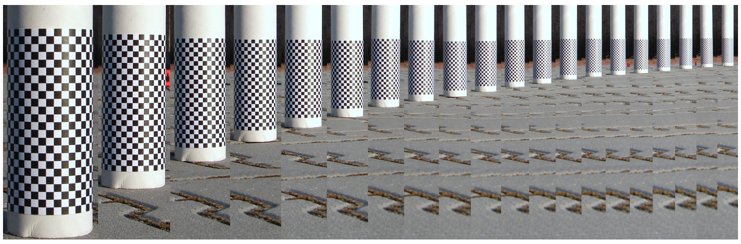

Figure B.4: A validation of the virtual scene using a set of photos of the physical cylinder from the same distances as in Figure B.3.

B.3 Filtering of Responses

The main limitation of our experiments using a virtual scene instead of a real scene lie in the limited DPI of a display device. This may potentially limit subjects in observing differences of materials that have too small size of the smallest structural element with regard to the DPI, especially on low-DPI displays. Therefore, we resorted to the filtering of the subjects' responses. First, we identify the most distant cylinder where structural elements are displayed on more than one pixels. The subject's response is recorded only if the number of the selected cylinder is lower than the number of the identified cylinder. In other words, this guarantees that the subject still had a chance to see the cylinder where more than one pixel reproduced the biggest structural element although she selected the one a step closer. This filtering removes 9.7% (34 out of 350) responses in *Experiment A*, and 8.5% (85 out of 1000) responses in *Experiment B*. The filtering takes effect especially for the materials with a very small size of the structural element, namely *fabric075*, *fabric106*, and *leather16*. Despite the filtering, we always obtained more than 30% valid responses for each material.

Further, we check the user's vision by the artificial chessboard patterns with the size of 2 mm and 5 mm. The visual angle of the 2 mm square on the most distant cylinder subtends 0.026^o, which is still above the resolution of a naked eye, which is typically around 0.7 arc-minutes = 0.012^o [Pir67]. In our experiment, over 90% of subjects were able to distinguish differences even on the last cylinder. Finally, we removed the subjects who were not able to spot a 5 mm pattern on the most distant cylinder as they probably either suffer from impaired vision or did not understand the task in the study (1 subject). Also, users who entered the measure of the calibration rectangle incorrectly in cm, and thus obtained an incorrect size of the stimuli images, were excluded from the study (2 subjects).

B.4 Analysis of Averaged Responses

Although we observe higher standard deviation values in data from *Experiment B* in Figure 3.11 (which was expected due to its uncontrolled nature), there is an apparent correlation between results of all experiments as shown in Table B.1. The high correlation of *Experiment B* with the other two validate results of this web-based study.

Along with responses, we also recorded the subjects' observing time for individual stimuli images. The average time span is between 8 and 20 seconds, while the longest

Table B.1: Pearson correlations between results of the experiments.

Correlation	ρ	p-Val
Experiment A vs. *Experiment B*	0.966	0.000000
Experiment A vs. *Experiment C*	0.811	0.000001
Experiment B vs. *Experiment C*	0.889	0.000000

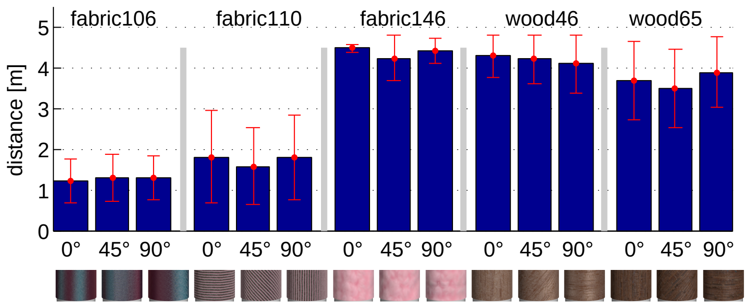

Figure B.5: Impact of the anisotropic behavior on the distance where a structure becomes indistinguishable.

ones belong to materials that have the very fine structure and thus are more difficult to be distinguished. Surprisingly, longer observation times were also recorded for wood materials (except *wood46*), where the structure is much more apparent but misses sufficient contrast. The shortest times were recorded for rough materials where most users quickly spotted differences across all cylinders, e.g., *fabric146* or *wood65*.

Finally, we analyze the estimated distance as a function of material anisotropy (*Experiment B*). The results for three fabric and two wood materials, each having three different orientations over the cylinder (0^o, 45^o, 90^o), are given in Figure B.5. Differences within individual materials are below 0.5 m, which is presumably due to the constant size, and thus visibility, of material structure regardless of changes in overall appearance.

Appendix C

Registration of Multi-View Images of Planar Surfaces

This appendix complements description of the registration method that was first introduced in [A.15] and is summarized in Section 4.2. While the standard registration approach is outlined in Section C.1, Section C.2 summarizes a method of calibration of a camera. Finally, the proposed iterative method is described in detail in Section C.3.

C.1 The Standard Image Registration Approach

A principle of the image registration is usually as follows. Given a set of photos of the same planar surface, the registration applies the projective transformation to all the photos so that the features of the planar surface are aligned across all of the transformed images. In the standard registration approach depicted in Figure 4.4-a, registration marks are placed around the photographed planar surface of a material. First, their 2D coordinates are found in all of the photos. Then, projective transformation matrices projecting these points to the desired target coordinates are computed. Finally, all the photos are transformed using these projective matrices. All the registered images have the same coordinate system.

The projective transformation, also called the *homography*, is a 2D coordinate transformation preserving straight lines (see [Sze06] for a survey of 2D coordinate transformations). Given a photo of a planar surface we want to transform together with an orthonormal coordinate system (u, v) of the photo, $\mathbf{m} = [u, v, 1]^T$ denotes an augmented point in this system in the planar surface. An augmented 2D point $\mathbf{m'} = [u', v', 1]^T$ in a new orthonormal coordinate system (u', v') into which we transform can be computed as $s\mathbf{m'} = \mathbf{Hm}$, where $\mathbf{H}$ is the 3×3 homography matrix and s is an arbitrary scalar. The homography can be computed if we know coordinates of at least four corresponding points in source and target images and it is defined uniquely up to a scale factor. If there are more than four such points and they are not perfectly corresponding, the homography has to be computed in a least-square sense (e.g., [HZ00]).

C.2 Camera Calibration

The calibration of a camera is a process where we look for a 3×3 matrix $\mathbf{A}$ of the intrinsic parameters of the camera, and the extrinsic parameters of the camera, which consist of a 3×3 rotation matrix $\mathbf{R}$ and a translation vector $\mathbf{t}$. While the intrinsic parameters $\mathbf{A}$ do not change as long as the internal setup of the camera does not change (e.g., focal length), the extrinsic parameters change when the camera moves, i.e., all the photos of the measured planar sample should have the same intrinsic parameters but the corresponding extrinsic parameters can be different. Using the parameters of the camera, we can project an augmented 3D point $M = [x, y, z, 1]^T$ from the world coordinate system into an image:

$$s\mathbf{m} = \mathbf{A} \left[\ \mathbf{R}\,\mathbf{t} \ \right] M, \tag{C.1}$$

where $\mathbf{m} = [u, v, 1]^T$ denotes an augmented 2D point and s is an arbitrary scalar. The usual pinhole camera model is assumed.

In a case where we work with extensive view- and illumination-dependent data (e.g., the BTF), the calibration procedure should start with selection of their representative subset. Although the calibration can be determined more precisely if all the photos are used, the estimation process would take a very long time if there were hundreds or even thousands of images. Therefore, we recommend working with images of one surface light field only, i.e., images where illumination directions are fixed while viewing directions change.

Next, the registration marks are found for all images in the subset. Without loss of generality, we define the world coordinate system so that the reference plane is on $z = 0$ (see Fig. 4.2). Spatial coordinates of the marks (x, y) should correspond to their real positions in a natural system of units, i.e., millimeters. Now, projective transformation matrices $\mathbf{H}_1$ projecting points from the reference plane to photos of the plane are computed based on the coordinates of the registration marks. From C.1, we have:

$$s \begin{bmatrix} u \\ v \\ 1 \end{bmatrix} = \mathbf{A} \left[\ \mathbf{r}_1\,\mathbf{r}_2\,\mathbf{r}_3\,\mathbf{t} \ \right] \begin{bmatrix} x \\ y \\ 0 \\ 1 \end{bmatrix} = \mathbf{A} \left[\ \mathbf{r}_1\,\mathbf{r}_2\,\mathbf{t} \ \right] \begin{bmatrix} x \\ y \\ 1 \end{bmatrix} = \mathbf{H}_1 \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}, \tag{C.2}$$

where $\mathbf{r}_i$ are column vectors of the rotation matrix $\mathbf{R} = [\mathbf{r}_1\,\mathbf{r}_2\,\mathbf{r}_3]$. Using the knowledge that column vectors of $\mathbf{R}$ are orthonormal and the matrix $\mathbf{A}$ is upper triangular, the camera intrinsic and extrinsic parameters can be derived from the homographies $\mathbf{H}_1$. For a detailed explanation, we refer the reader to Zhang's paper [Zha99]. We use the Camera Calibration Toolbox for Matlab (http://www.vision.caltech.edu/bouguetj/calib_doc/), which implements Zhang's work.

C.3 Iterative Fitting of the Position of the Surface Plane

We can define an almost arbitrary position of the expected plane which represents the surface of a material in the world coordinate system just by setting new z-coordinates of the registration marks. If we project the estimated 3D coordinates of the new registration marks back to non-registered photos, we obtain new 2D coordinates of the registration marks in the coordinate systems of the photos. We can now register the images by specification of target coordinates of the registration marks and by computation of homography matrices $\mathbf{H}_2$ (see Fig. 4.4-b). The latter should not be confused with the homography matrices $\mathbf{H}_1$ mentioned above, which project points in the reference registration marks plane of the world coordinate system to the non-registered images. In contrast, these new homography matrices $\mathbf{H}_2$ project points from the non-registered photos to the registered images as homography matrices $\mathbf{H}$ in the standard registration approach (Fig. 4.4-a).

Therefore, we suggest a novel iterative method for estimation of the shift and the slant of the surface plane, projective transformation of the images and evaluation of the alignment of surface features. The method is depicted in Figure 4.4-b. As we look for the optimal vertical shift and the optimal slant of the surface plane, three parameters have to be found: a z-coordinate of an auxiliary point $P = [0, 0, z]^T$ which lies in the surface plane, an elevation θ of the normal vector of the surface, and an azimuthal angle φ of the normal vector. Therefore, the search state-space is three-dimensional and execution of the transformation, compression and evaluation can be computationally demanding. We can quickly find at least the local minimum by alternating between estimation of individual parameters. As there can be significant variations of height of the surface, there may be more than one good position of the surface plane.

Our goal is to provide as accurate visualization of the measured material as possible. As quality of visualization relies mostly on visual quality after data compression, evaluation functions which estimate the alignment of surface features should reflect properties of the selected compression method. Therefore, we use the error of a compression technique as an objective quality measure. As we work with only a subset of all of the photos, the compression as well as its error evaluation can be done quickly enough to be practical. An ideal position of the plane is the one where the compression error is minimal.

One iteration of modification of the reference plane, transformation of the data, and evaluation of the compression error takes about one second depending on the count of pixels and images. In order to avoid local minima the search space is sampled uniformly, alternating between estimation of the three parameters (shift z, elevation θ, and azimuth φ) with the following step sizes: $\Delta z = 0.1$ mm (range [-2, 2] mm), $\Delta\theta = 0.1°$ (range [0, 3°]), $\Delta\varphi = 10°$ (range [0, 360°]), and then refined near the global minimum (step sizes: $\Delta z = 0.01$ mm, $\Delta\theta = 0.01°$, $\Delta\varphi = 1°$). Typically, around 800 iterations are necessary to find the proper orientation and shift of the surface plane.

Appendix D

BRDF Slices

This appendix complements the description of the BRDF slices provided in Chapter 5. First, an extension of the reconstruction method is encompassed in Section D.1. Then, Section D.2 presents the dataset used to study the optimal placement of the slices and also to evaluate the performance of the method together with regular sampling schemes presented in Section D.3. Finally, Section D.4 provides instructions for implementation of the method on a graphics hardware.

D.1 Reconstruction from the Sparse Structure

This section complements Section 5.3 and rewrites the equations for interpolation in the 2D BRDF subspace introduced there to make the transition to the 4D space possible. Then, equations for interpolation of any value in the BRDF space from the sparse 4D structure (in the form of four types of the BRDF slices) are provided.

It can be proven that the change of interpolation order of x, y axes in equations of Section 5.3 yields the same result. Also, the equations can be rewritten as follows to make the transition into four dimensions possible. First, we apply the bilinear interpolation of the corner values introducing:

$$c_{xy} = (1-x) \cdot [(1-y) \cdot c_{00} + y \cdot c_{01}] + x \cdot [(1-y) \cdot c_{10} + y \cdot c_{11}] \quad .$$

Values of the slices are interpolated linearly and the differences are computed as:

$$\begin{aligned} p_{xy} &= (1-y) \cdot p_{x0} + y \cdot p_{x1} \quad , \\ q_{xy} &= (1-x) \cdot q_{0y} + x \cdot q_{1y} \quad , \\ \Delta p_{xy} &= p_{xy} - c_{xy} \quad , \\ \Delta q_{xy} &= q_{xy} - c_{xy} \quad . \end{aligned}$$

The final value of the reconstructed function is:

$$r_{xy} = max(c_{xy} + \Delta p_{xy} + \Delta q_{xy}, min_{xy}) = max(p_{xy} + q_{xy} - c_{xy}, min_{xy}) \quad .$$

The extension into four dimensions is straightforward. We apply a multi-linear interpolation to all sixteen corners of the 4D hyper-cube $c_{\bar{x}\bar{y}\bar{z}\bar{w}}$, where $\bar{x}, \bar{y}, \bar{z}, \bar{w} \in \{0, 1\}$:

$$
\begin{aligned}
c_{xyzw} = (1-x) \cdot \Big\{ (1-y) \cdot \{(1-z) &\cdot [(1-w) \cdot c_{0000} + w \cdot c_{0001}] \\
+z &\cdot [(1-w) \cdot c_{0010} + w \cdot c_{0011}]\} \\
+y \cdot \{(1-z) &\cdot [(1-w) \cdot c_{0100} + w \cdot c_{0101}] \\
+z &\cdot [(1-w) \cdot c_{0110} + w \cdot c_{0111}]\} \Big\} \\
+x \cdot \Big\{ (1-y) \cdot \{(1-z) &\cdot [(1-w) \cdot c_{1000} + w \cdot c_{1001}] \\
+z &\cdot [(1-w) \cdot c_{1010} + w \cdot c_{1011}]\} \\
+y \cdot \{(1-z) &\cdot [(1-w) \cdot c_{1100} + w \cdot c_{1101}] \\
+z &\cdot [(1-w) \cdot c_{1110} + w \cdot c_{1111}]\} \Big\} \ .
\end{aligned}
$$

Then, values of the slices are interpolated in three remaining dimensions (see Fig. 5.2). Values of the eight axial slices $p_{x\bar{y}\bar{z}\bar{w}}$, where $\bar{y}, \bar{z}, \bar{w} \in \{0, 1\}$ are interpolated as:

$$
\begin{aligned}
p_{xyzw} = (1-y) \ &\cdot \{(1-z) \cdot [(1-w) \cdot p_{x000} + w \cdot p_{x001}] + z \cdot [(1-w) \cdot p_{x010} + w \cdot p_{x011}]\} \\
+y \ &\cdot \{(1-z) \cdot [(1-w) \cdot p_{x100} + w \cdot p_{x101}] + z \cdot [(1-w) \cdot p_{x110} + w \cdot p_{x111}]\} \ ,
\end{aligned}
$$

values of the eight diagonal slices $q_{\bar{x}y\bar{z}\bar{w}}$, where $\bar{x}, \bar{z}, \bar{w} \in \{0, 1\}$ are interpolated as:

$$
\begin{aligned}
q_{xyzw} = (1-x) \ &\cdot \{(1-z) \cdot [(1-w) \cdot q_{0y00} + w \cdot q_{0y01}] + z \cdot [(1-w) \cdot q_{0y10} + w \cdot q_{0y11}]\} \\
+x \ &\cdot \{(1-z) \cdot [(1-w) \cdot q_{1y00} + w \cdot q_{1y01}] + z \cdot [(1-w) \cdot q_{1y10} + w \cdot q_{1y11}]\} \ ,
\end{aligned}
$$

values of the eight horizontal slices $s_{\bar{x}\bar{y}z\bar{w}}$, where $\bar{x}, \bar{y}, \bar{w} \in \{0, 1\}$ are interpolated as:

$$
\begin{aligned}
s_{xyzw} = (1-x) \ &\cdot \{(1-y) \cdot [(1-w) \cdot s_{00z0} + w \cdot s_{00z1}] + y \cdot [(1-w) \cdot s_{01z0} + w \cdot s_{01z1}]\} \\
+x \ &\cdot \{(1-y) \cdot [(1-w) \cdot s_{10z0} + w \cdot s_{10z1}] + y \cdot [(1-w) \cdot s_{11z0} + w \cdot s_{11z1}]\} \ ,
\end{aligned}
$$

and values of the eight vertical slices $t_{\bar{x}\bar{y}\bar{z}w}$, where $\bar{x}, \bar{y}, \bar{z} \in \{0, 1\}$ are interpolated as:

$$
\begin{aligned}
t_{xyzw} = (1-x) \ &\cdot \{(1-y) \cdot [(1-z) \cdot t_{000w} + z \cdot t_{001w}] + y \cdot [(1-z) \cdot t_{010w} + z \cdot t_{011w}]\} \\
+x \ &\cdot \{(1-y) \cdot [(1-z) \cdot t_{100w} + z \cdot t_{101w}] + y \cdot [(1-z) \cdot t_{110w} + z \cdot t_{111w}]\} \ .
\end{aligned}
$$

Final value of reconstructed function is:

$$
\begin{aligned}
r_{xyzw} &= max(p_{xyzw} + q_{xyzw} + s_{xyzw} + t_{xyzw} - 3 \cdot c_{xyzw}, min_{xyzw}) \ , \\
min_{xyzw} &= min(min_p, min_q, min_s, min_t) \ , \\
min_p &= min(p_{x000}, p_{x001}, p_{x010}, p_{x011}, p_{x100}, p_{x101}, p_{x110}, p_{x111}) \ , \\
min_q &= min(q_{0y00}, q_{0y01}, q_{0y10}, q_{0y11}, q_{1y00}, q_{1y01}, q_{1y10}, q_{1y11}) \ , \\
min_s &= min(s_{00z0}, s_{00z1}, s_{01z0}, s_{01z1}, s_{10z0}, s_{10z1}, s_{11z0}, s_{11z1}) \ , \\
min_t &= min(t_{000w}, t_{001w}, t_{010w}, t_{011w}, t_{100w}, t_{101w}, t_{110w}, t_{111w}) \ .
\end{aligned}
$$

D.2 Test Dataset

To generate our test dataset, we use high-quality fits of an analytical BRDF model of Kurt et al. [KSKK10] as it is one of the state-of-the-art models for the anisotropic BRDFs. We selected ten materials, four of them measured in [NDM05]; fitted in [KSKK10], and the remaining materials come from our own measurements (Sec. A.4). All of the materials are anisotropic and include fabrics, brushed aluminum, and raw wood. Their fitted BRDFs are shown in Figure D.1 and parameters are included in Table D.1.

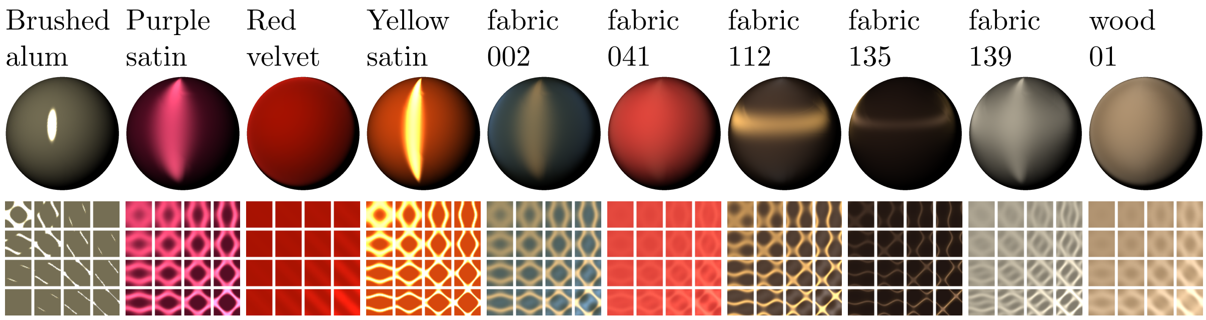

Figure D.1: Ten tested materials represented by the BRDF model [KSKK10] rendered on spheres for one point-light source and a preview of their BRDF.

Table D.1: Parameters of the BRDF model of Kurt et al. [KSKK10] for ten materials.

		Brushed alum.	Purple satin	Red velvet	Yellow satin	fabric002	fabric041	fabric112	fabric135	fabric139	wood01
	k_{dr}	0.0036	0.0026	0.0048	0.0066	0.0805	0.2565	0.0838	0.0395	0.1607	0.1548
	k_{dg}	0.0034	0.0004	0.0005	0.0022	0.0945	0.0703	0.0596	0.0258	0.153	0.1219
	k_{db}	0.0026	0.0011	0	0.0004	0.081	0.0613	0.0434	0.0205	0.135	0.0836
lobe 1	k_{sr}	0.0115	0.1404	0.1938	0.0542	0.0798	0.0211	0.1419	0.014	0.0483	0.0768
	k_{sg}	0.0105	0.0522	0.0333	0.0345	0.1205	0.019	0.1078	0.011	0.0461	0.0741
	k_{sb}	0.0075	0.0711	0.0267	0.0131	0.1796	0.0146	0.0521	0.0087	0.0414	0.0706
	f_0	0.999	0.055	0.041	0.207	0.3431	1	1	1	1	1
	m_x	0.035	0.339	2.337	0.129	1.3508	0.3665	1.2101	1.2025	0.2948	0.5672
	m_y	0.129	1.256	2.644	1.084	0.8121	1.1075	0.2139	0.1029	1.3215	0.7619
	α	0.005	0	0	0.197	0.8359	0.4984	0	0.5043	0.5753	0.5051
lobe 2	k_{sr}					0.0798	0.0211	0.1419	0.014	0.0483	0.0768
	k_{sg}					0.1205	0.019	0.1078	0.011	0.0461	0.0741
	k_{sb}					0.1796	0.0146	0.0521	0.0087	0.0414	0.0706
	f_0					1	0.1222	1	0.3078	1	1
	m_x					0.3168	1.1864	0.7046	2.038	1.3281	1.7007
	m_y					1.0061	0.7247	1.2292	0.3867	0.3677	1.923
	α					0.0905	0.8356	0.976	0.5723	0.667	0.2285

D.3 Regular sampling schemes

For the sake of comparison of the proposed adaptive method with the commonly used regular uniform sampling, we designed thirty regular sampling schemes with a variable sample density. The samples are distributed nearly uniformly across a hemisphere ranging from $\tilde{n} = 29$ to $\tilde{n} = 841$ directions. These schemes produce in total from $n = 435$ to $n = 354,061$ reciprocal samples ($n = \frac{\tilde{n}\cdot(\tilde{n}+1)}{2}$, see examples in Fig. D.2 and Tab. D.2).

01: 29 directions 435 samples | **14**: 133 directions 8,911 samples | **19**: 193 directions 18,721 samples | **30**: 841 directions 354,061 samples

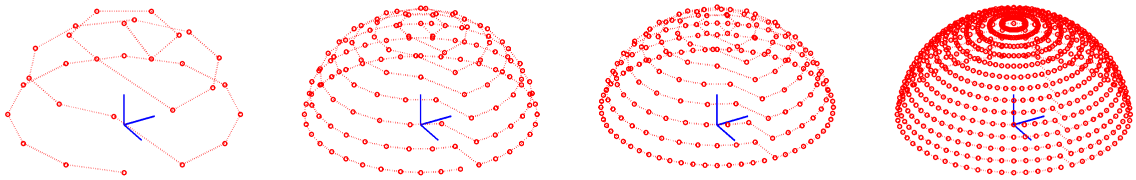

Figure D.2: Four from total thirty sampling schemes used for the uniform sampling of the hemisphere together with the barycentric and RBF based methods.

Table D.2: The measured directions of four from total thirty sampling schemes. Each selected elevation has defined an azimuthal step resulting in a certain count of samples.

scheme 1			scheme 14			scheme 19			scheme 30		
θ [°]	$\Delta\varphi$ [°]	#	θ [°]	$\Delta\varphi$ [°]	#	θ [°]	$\Delta\varphi$ [°]	#	θ [°]	$\Delta\varphi$ [°]	#
0	360	1	0	360	1	0	360	1	0	360	1
28	60	6	12	60	6	10	60	6	6	6	60
56	36	10	24	36	10	20	36	10	12	6	60
80	30	12	36	30	12	30	30	12	18	6	60
			48	20	18	40	20	18	24	6	60
			60	18	20	50	18	20	30	6	60
			72	12	30	60	12	30	36	6	60
			80	10	36	70	10	36	42	6	60
						80	6	60	48	6	60
									54	6	60
									60	6	60
									66	6	60
									72	6	60
									78	6	60
									80	6	60
sum		29	sum		133	sum		193	sum		841
illum.×view		841	illum.×view		17689	illum.×view		37249	illum.×view		707281
reciprocal		435	reciprocal		8911	reciprocal		18721	reciprocal		354061

D.4 Graphics Hardware Implementation

Another important feature that sets our method apart from the RBF interpolation is its feasibility for very fast interpolation on a graphics hardware (GPU). Although the barycentric interpolation allows for this as well, it does not support an adaptive non-uniform placement of the samples and thus cannot achieve the high reconstruction performance.

To reconstruct the BRDF on a GPU using our method, one need to store data in four textures. The first texture contains values at intersections of the slices, the second contains pre-interpolated values of the axial slices, the third contains pre-interpolated values of the diagonal slices and the fourth texture contains pre-interpolated values of the horizontal slices (identical to the vertical slices). The first three textures can be easily stacked together. Note that the interpolation along the slices has to be precomputed to avoid intricate lookups for measured values. Despite that, the data are still very compact. Table D.3 compares the size of the data structures required by the tested methods using $8,911$ and $18,721$ reciprocal samples.

Table D.3: The GPU storage space required by the barycentric and proposed interpolation methods for $8,911$ samples ($e_s = 14^o/a_s = 36^o$) and $18,721$ samples ($e_s = 10^o/a_s = 36^o$).

method	$8,911$ samples	$18,721$ samples
barycentric	3.5 MB (scheme 14)	3.6 MB (scheme 19)
proposed	3.3 MB	5.1 MB

Rendering speed on the GPU is 270 FPS for the barycentric interpolation and 130 FPS for the proposed method regardless of the complexity of a 3D scene as the majority of the computation is performed by the fragment shader rendering window of size 800×800 pixels. The timings are obtained using the nVidia GeForce GTX 570 graphics card. The slower performance of the proposed approach results from the higher complexity of the BRDF reconstruction from the slices, which is in contrast to implementation of the barycentric interpolation using fast indexing in precomputed cube-maps stored as textures. The visual performance of the GPU rendering is identical to that of the CPU rendering shown in Figure 5.9, i.e., the proposed method clearly provides superior visual quality to the barycentric interpolation.

Let us describe the content of the four textures in detail. The first texture containing values at intersections of the slices consists of $\frac{e_n(e_n+1)}{2}$ areas, where e_n is the count of elevations (see Fig. 5.2). Note that due to Helmholtz reciprocity we do not need to store values for subspaces where $\theta_i > \theta_v$. Each area covers $(a_n + 1) \times (2a_n + 1)$ pixels, where a_n is the count of slices of one type (axial or diagonal) in a subspace and for simplicity we assume that the number of the axial slices is equal to the number of the diagonal slices. Then, $2a_n$ is the number of intersections in each subspace. An additional row and column in each area is there due to easier interpolation of the values, which can be performed implicitly by the GPU during texture value lookup and explicit interpolation has to be performed in two remaining dimensions (θ_v, θ_i) only.

The second texture containing values of the axial slices also consists of $\frac{e_n(e_n+1)}{2}$ areas. Width of each area is $a_n + 1$ as one axial slice is repeated for easier interpolation and height of the texture depends on a selected resolution (e.g., for resolution 1^o the height is 360 pixels). As interpolation in two dimensions is provided by the GPU during texture value lookup, only interpolation in two other dimensions has to be implemented. The third texture covers the same area as the second texture, contains values of the diagonal slices and is interpolated in the same manner.

The last texture contains values of all horizontal slices. There are $2a_n^2$ horizontal slices for each measured illumination elevation. Therefore, with padding for easier interpolation, width of the texture is $a_n(2a_n + 1)$ and height is, e.g., $90 \cdot e_n$ for a resolution of 1^o. Interpolation in two dimensions is, again, performed during texture value lookup by the GPU and only two remaining dimensions has to be interpolated explicitly.

Appendix **E**

Applications of the Portable Setup

This appendix complements Chapter 6 to demonstrate the speed and portability of our acquisition procedure (using 2 illumination and 2 viewing elevations). First, we recorded and reconstructed the BTF of skin on the author's hand (Fig. E.1). Although there is present a certain blur caused by slight movement of subject's hand during the five minute measurement, note that such a measurement would not be possible, due to long measurement times or a setup design, by any BTF acquisition method presented so far. Compared to the acquisition setup introduced by Gu et al. [GTR*06], our approach is not constrained by the BRDF properties, i.e., we can also record rougher, non-opaque materials. Also, our setup is based on inexpensive components and provides a higher number of viewing directions; however, it has a slightly longer acquisition time.

The speed of our setup also allows us, contrary to other setups [LMM*10], to record time-varying BTF of dynamic processes at a sampling rate of up to 12 measurements per hour). We measured desiccation of glue placed on the underside of leather in form of a single dot as shown in the first row of Figure E.2 (measurements are replicated using an

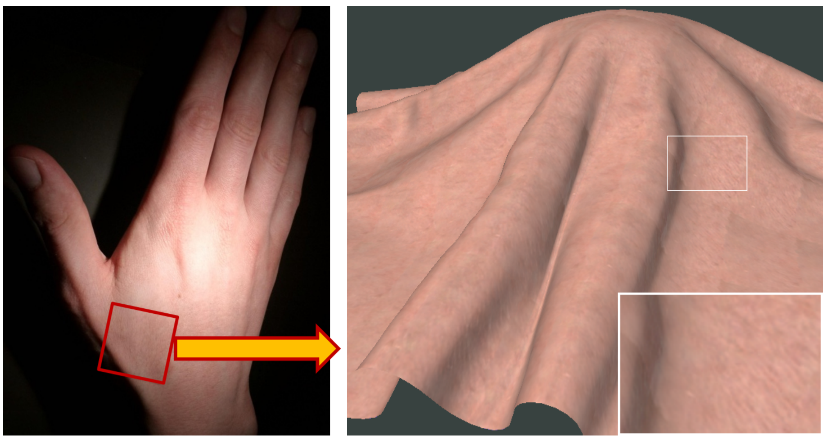

Figure E.1: The measured BTF of human skin.

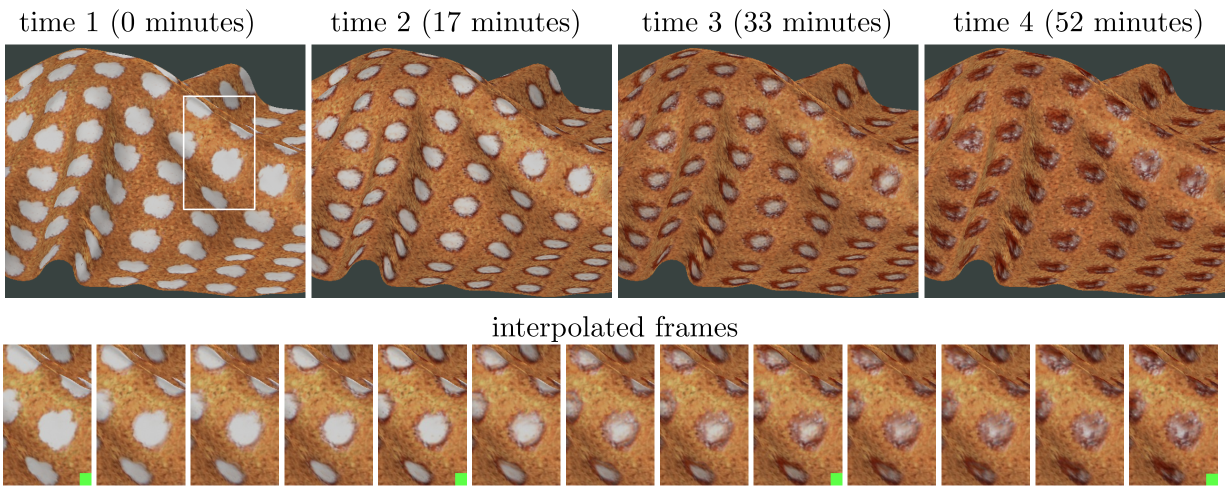

Figure E.2: Four measurements of the time-varying BTF of desiccation of glue (the first row). A complete sequence rendered from interpolated BTFs based on the four measured (marked as green) and the nine interpolated BTF time-frames (the second row).

image-tiling approach to show illumination- and view-dependent appearance of the dot). We performed four measurements, i.e., recorded the BTF of the process every 15 minutes in a span of one hour. Typical duration of a single measurement was 5 minutes.

Then, we tested three interpolation methods: the linear interpolation (applied to TVBTF in [LMM*10]), the space-time factorization (STAF) [GTR*06], and the displacement interpolation [BvdPPH11] as shown in Figure E.3. As expected, the linear interpolation is fastest, but introduces a blur at the regions where the area of the glue shrinks. The STAF algorithm performs significantly better, however, it is sensitive to careful setting of initial conditions as it performs the per-pixel non-linear fitting. Finally, we test the displacement interpolation (DI) solving a generalized mass transport optimization problem.

In contrast to the STAF algorithm, which models spatio-temporal dynamics of the entire time-sequence, the DI algorithm performs a weighted interpolation between only two images. Due to its stability, we resorted to the DI algorithm and applied it to interpolation

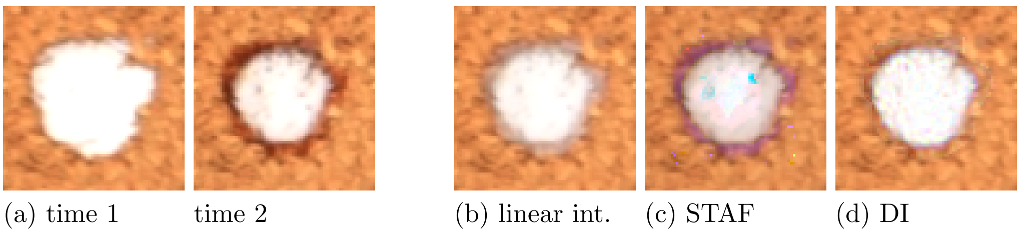

Figure E.3: Interpolation of the time-frames 1 and 2 (a) ($\Delta t = 17$ min) using: (b) the linear interpolation, (c) the STAF algorithm, and (d) the displacement interpolation.

between two sets of 208 images representing two time-frames of the TVBTF. The second row in Figure E.2 shows the obtained smooth transitions between the measured time-frames illustrating the illumination- and view-dependent effect of desiccation of glue over one hour (the measured time-frames are marked by green squares).

Appendix **F**

Minimal Sampling for Effective Acquisition of the Anisotropic BRDFs

This appendix provides detailed derivation of the gradient of the objective function introduced in Section 7.2.2 and describes a method for a fast update of the condition number used in Algorithm 7.4.

F.1 Gradient of the Objective Function

Unfortunately, the objective function of Equation 7.5 does not have the close-form solution. At least, we can compute its gradient to accelerate its numerical solution. Detailed derivation of the gradient is provided here:

$$\begin{aligned}
E(\boldsymbol{C}) &= \|\tilde{\boldsymbol{Y}} - \tilde{\boldsymbol{Q}}\boldsymbol{C}\|_F^2 + \eta\|\boldsymbol{C}\|_F^2 + \lambda\|\tilde{\tilde{\boldsymbol{Q}}}\boldsymbol{C}\boldsymbol{M}\|_F^2, \\
\nabla E(\boldsymbol{C}) &= \nabla\|\tilde{\boldsymbol{Y}} - \tilde{\boldsymbol{Q}}\boldsymbol{C}\|_F^2 + \eta\nabla\|\boldsymbol{C}\|_F^2 + \lambda\nabla\|\tilde{\tilde{\boldsymbol{Q}}}\boldsymbol{C}\boldsymbol{M}\|_F^2 \\
&= \nabla\,\mathrm{Tr}[(\tilde{\boldsymbol{Y}} - \tilde{\boldsymbol{Q}}\boldsymbol{C})^T(\tilde{\boldsymbol{Y}} - \tilde{\boldsymbol{Q}}\boldsymbol{C})] + \eta\nabla\,\mathrm{Tr}(\boldsymbol{C}^T\boldsymbol{C}) + \lambda\nabla\,\mathrm{Tr}[(\tilde{\tilde{\boldsymbol{Q}}}\boldsymbol{C}\boldsymbol{M})^T(\tilde{\tilde{\boldsymbol{Q}}}\boldsymbol{C}\boldsymbol{M})] \\
&= \nabla\,\mathrm{Tr}(\tilde{\boldsymbol{Y}}^T\tilde{\boldsymbol{Y}}) - \nabla\,\mathrm{Tr}(\boldsymbol{C}^T\tilde{\boldsymbol{Q}}^T\tilde{\boldsymbol{Y}}) - \nabla\,\mathrm{Tr}(\tilde{\boldsymbol{Y}}^T\tilde{\boldsymbol{Q}}\boldsymbol{C}) + \nabla\,\mathrm{Tr}(\boldsymbol{C}^T\tilde{\boldsymbol{Q}}^T\tilde{\boldsymbol{Q}}\boldsymbol{C}) + \\
&\quad\ \eta\nabla\,\mathrm{Tr}(\boldsymbol{C}^T\boldsymbol{C}) + \lambda\nabla\,\mathrm{Tr}(\boldsymbol{M}\boldsymbol{C}^T\tilde{\tilde{\boldsymbol{Q}}}^T\tilde{\tilde{\boldsymbol{Q}}}\boldsymbol{C}\boldsymbol{M}) \\
&= -2\tilde{\boldsymbol{Q}}^T\boldsymbol{Y} + 2\tilde{\boldsymbol{Q}}^T\tilde{\boldsymbol{Q}}\boldsymbol{C} + 2\eta\boldsymbol{C} + 2\lambda\tilde{\tilde{\boldsymbol{Q}}}^T\tilde{\tilde{\boldsymbol{Q}}}\boldsymbol{C}\boldsymbol{M}\boldsymbol{M}.
\end{aligned}$$

F.2 Fast Update of the Condition Number

Let $\tilde{\boldsymbol{Q}} \in \mathbb{R}^{n\times k}$ be the selected rows of the principal components in $\boldsymbol{Q}$, whose condition number we look for. The condition number is defined as $\kappa(\tilde{\boldsymbol{Q}}) = \frac{\sigma_{max}(\tilde{\boldsymbol{Q}})}{\sigma_{min}(\tilde{\boldsymbol{Q}})}$, i.e., the ratio of the maximum and minimum singular values of $\tilde{\boldsymbol{Q}}$. Note that the singular values of $\tilde{\boldsymbol{Q}}$ are equal to the square roots of the eigenvalues of the matrix $\tilde{\boldsymbol{Q}}^T\tilde{\boldsymbol{Q}}$. Therefore, we can exploit methods for a fast update of the eigenvalues when a rank-one modification of the matrix is performed [BNS78]. Let $\bar{\boldsymbol{Q}} \in \mathbb{R}^{n-1\times k}$ be the desired matrix $\tilde{\boldsymbol{Q}}$ with one row removed. Specifically, without loss of generality, we can write:

$$\tilde{\boldsymbol{Q}} = \begin{bmatrix} \bar{\boldsymbol{Q}} \\ \boldsymbol{q}^T \end{bmatrix}, \quad \tilde{\boldsymbol{Q}}^T\tilde{\boldsymbol{Q}} = \begin{bmatrix} \bar{\boldsymbol{Q}}^T \boldsymbol{q} \end{bmatrix} \begin{bmatrix} \bar{\boldsymbol{Q}} \\ \boldsymbol{q}^T \end{bmatrix} = \bar{\boldsymbol{Q}}^T\bar{\boldsymbol{Q}} + \boldsymbol{q}\boldsymbol{q}^T,$$

where $\boldsymbol{q}^T \in \mathbb{R}^{1\times k}$ is a row of $\boldsymbol{Q}$, which we want to add to the current matrix $\bar{\boldsymbol{Q}}$. Obviously, we want to select such a $\boldsymbol{q}^T$, that $\kappa(\tilde{\boldsymbol{Q}})$ is minimum.

Let $\rho = \|\boldsymbol{q}\|_2^2$ and $\boldsymbol{b} = \frac{1}{\sqrt{\rho}}\boldsymbol{q}$. Then $\tilde{\boldsymbol{Q}}^T\tilde{\boldsymbol{Q}} = \bar{\boldsymbol{Q}}^T\bar{\boldsymbol{Q}} + \rho\boldsymbol{b}\boldsymbol{b}^T$ and we can once solve the symmetric eigenproblem of $\bar{\boldsymbol{Q}}^T\bar{\boldsymbol{Q}} = \boldsymbol{V}\boldsymbol{D}\boldsymbol{V}^T$, where $\boldsymbol{D} \in \mathbb{R}^{k\times k}$ is a diagonal matrix of eigenvalues and $\boldsymbol{V} \in \mathbb{R}^{k\times k}$ is a matrix of corresponding eigenvectors. Now, we can write $\bar{\boldsymbol{Q}}^T\bar{\boldsymbol{Q}} + \rho\boldsymbol{b}\boldsymbol{b}^T = \boldsymbol{V}(\boldsymbol{D} + \rho\boldsymbol{z}\boldsymbol{z}^T)\boldsymbol{V}^T$, where $\boldsymbol{z} = \boldsymbol{V}^T\boldsymbol{b}$. Let $\boldsymbol{C} = \boldsymbol{D} + \rho\boldsymbol{z}\boldsymbol{z}^T = \boldsymbol{X}\tilde{\boldsymbol{D}}\boldsymbol{X}^T$. Then $\tilde{\boldsymbol{Q}}^T\tilde{\boldsymbol{Q}} = \tilde{\boldsymbol{V}}\tilde{\boldsymbol{D}}\tilde{\boldsymbol{V}}^T$, where $\tilde{\boldsymbol{V}} = \boldsymbol{V}\boldsymbol{X}$, but we are only interested in the maximum and minimum eigenvalue of $\tilde{\boldsymbol{D}}$ to compute $\kappa(\tilde{\boldsymbol{Q}})$.

Let $d_1 \leq d_2 \leq \cdots \leq d_k$ be the eigenvalues of $\boldsymbol{D}$ and let $\tilde{d}_1 =\leq \tilde{d}_2 \leq \cdots \leq \tilde{d}_k$ be the eigenvalues of $\boldsymbol{C}$ (and therefore also of $\tilde{\boldsymbol{D}}$). Then holds $d_1 \leq \tilde{d}_1 \leq d_2$ and $d_k \leq \tilde{d}_k \leq d_k + \rho$. Golub [Gol73] shows that this eigenvalues can be computed by finding zeros of the secular equation:

$$w(\tilde{d}_i) = 1 + \rho \sum_{j=1}^{k} \frac{z_j^2}{d_j - \tilde{d}_i} \quad ,$$

where $\boldsymbol{z} = [z_1, \ldots, z_k]^T$. The equation can be solved numerically by the bisection method. First, we set the lover bound of an interval to d_1 and the upper bound to d_2 and we find $\tilde{d}_1$ in the interval where $w(\tilde{d}_1) \leq \varepsilon$ for a sufficiently small ε. Second, we set the lover bound to d_k and the upper bound to $d_k + \rho$ and perform the bisection method once again to find $\tilde{d}_k$ in the interval where $w(\tilde{d}_k) \leq \varepsilon$. In praxis, we perform 30 subdivisions of an interval to achieve very precise results.

The condition number of the updated matrix $\tilde{\boldsymbol{Q}}$ is finally computed as $\kappa(\tilde{\boldsymbol{Q}}) = \frac{\sqrt{\tilde{d}_k}}{\sqrt{\tilde{d}_1}}$.