Fast Method of Sparse Acquisition and Reconstruction of View and Illumination Dependent Datasets

Jiří Filip, Radomír Vávra

Institute of Information Theory and Automation of the AS CR, Prague, Czech Republic

Abstract

Although computer graphics uses measured view and illumination dependent data to achieve realistic digital reproduction of real-world material properties, the extent of their utilization is currently limited by a complicated acquisition process. Due to the high dimensionality of such data, the acquisition process is demanding on time and resources. Proposed is a method of approximate reconstruction of the data from a very sparse dataset, obtained quickly using inexpensive hardware. This method does not impose any restrictions on input datasets and can handle anisotropic, non-reciprocal view and illumination direction-dependent data. The method's performance was tested on a number of isotropic and anisotropic apparent BRDFs, and the results were encouraging. The method performs better than the uniform sampling of a comparable sample count and has three main benefits: the sparse data acquisition can be done quickly using inexpensive hardware, the measured material does not need to be extracted or removed from its environment, and the entire process of data reconstruction from the sparse samples is quick and reliable. Finally, the ease of sparse dataset acquisition was verified in measurement experiments with three materials, using a simple setup of a consumer camera and a single LED light. The proposed method has also shown promising performance when applied to sparse measurement and reconstruction of BTFs, mainly for samples with a lower surface height variation. Our approach demonstrates solid performance across a wide range of view and illumination dependent datasets, therefore creating a new opportunity for development of time and cost-effective portable acquisition setups.

Keywords: apparent BRDF, measurement, reconstruction, sparse sampling, portable setup, BTF

1. Introduction

View and illumination dependent data can be beneficial in many computer graphic applications, due to their ability to digitally represent the actual appearance of respective material. However, their measurement is costly and time-consuming, because standard acquisition procedures of such data often require lengthy measurements, or either a specific shape of the measured sample or a dedicated measurement setup. *Bidirectional reflectance distribution function* (BRDF) [25], *spatially varying BRDF* (SV- BRDF) and *bidirectional texture function* (BTF) [3] are examples of such data. While a four-dimensional BRDF describes distribution of energy reflected to the viewing direction when illuminated from a specific direction, a six-dimensional SVBRDF additionally captures the spatial dependency of reflectance across a material surface. While BRDF and SVBRDF impose restrictions on reciprocity, opacity and a range of sample height variations, the six-dimensional BTF generally does not fulfill these restrictions. This is due to local effects in a rough material structure such as occlusions, masking, subsurface scattering, and inter-reflections.

Therefore, individual BTF pixels are not regarded as BRDF but rather apparent BRDF (ABRDF). If we process individual color/spectral channels separately, the ABRDF can be represented by a four-dimensional function $ABRDF(\theta_i, \varphi_i, \theta_v, \varphi_v)$. ABRDF is the most general data representation of a reflectance of opaque materials dependent on local illumination $\mathbf{I}(\theta_i, \varphi_i)$ and view $\mathbf{V}(\theta_v, \varphi_v)$ directions; therefore, we focus on its proper acquisition and reconstruction in this paper. Its typical parameterization by elevation θ and azimuthal φ angles is shown in Fig. 1-left. A projection of the 4D ABRDF, representing dependence of view and illumination directions of a single pixel (BTF) or its average value (BRDF) by means of a 2D image, is shown in Fig. 1-right. Note that individual rectangles (an example is shown in red) represent 2D subspaces of 4D ABRDF at constant elevations (θ_i/θ_v). These subspaces are toroidal. That is data of the highest $\varphi \approx 2\pi$ are followed by data of the lowest $\varphi \approx 0$.

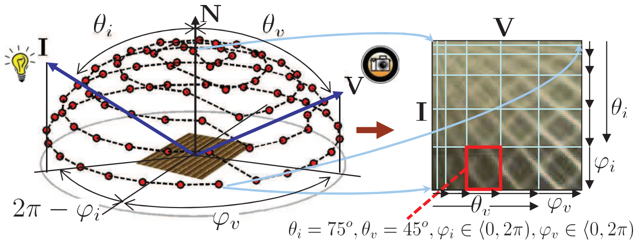

Figure 1: Parameterization (left) of view and illumination-dependent data of a single/average pixel (right).

Email addresses: `filipj@utia.cas.cz` (Jiří Filip), `vavra@utia.cas.cz` (Radomír Vávra)

Main contributions of this paper:

- a reconstruction method of the entire anisotropic ABRDF space from less than two hundred sparsely measured samples
- a practically verified novel method for intuitive and fast ABRDF acquisition and reconstruction using a consumer camera and LED light in under 10 minutes.

Main features of the proposed method:

- a correct reconstruction of non-reciprocal, energy non-conserving ABRDF data
- an arbitrarily dense sampling of specular highlights, without increasing measurement time
- no need for lengthy measurement using a dedicated and expensive measurement setup
- not necessary to process or 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.

This paper is structured as follows: Section 2 sets work in the context of related research. Section 3 explains the principle of the proposed method. Section 4 shows results of performed experiments. The method's limitations are discussed in Section 5, and pilot project results of the real data acquisition scenario are shown in Section 6. Section 7 shows experimental reconstruction of BTF samples, and Section 8 concludes the paper.

2. Prior Work

The proposed work relates to methods of BRDF or SVBRDF acquisition and interpolation from sparse samples.

Such data were initially captured by setups based on gonioreflectometers realizing a required four mechanical degrees of freedom (DOF) of camera/light/sample movement [12], [29]. Because measurement times were too long, setups were used which reduced the required number of DOF using parabolic mirrors [4], or kaleidoscope [10]. They allowed the capture of many viewing directions simultaneously; however a limited range of surface height or elevation angles resulted. Measurement time can also be reduced by using multiple lights and sensors simultaneously [20]; yet, high financial cost is associated with such a setup. Another group of fast acquisition methods reduces the number of DOF by using a known shape of the sample [19], [32], [17], [23], [13]. However, these approaches are often limited to isotropic BRDFs, or focus on representing sparsely sampled data using a parametric BRDF model. There is an existing statistical acquisition approach [22] allowing quick and economical measurement of ABRDF; however, it requires several samples of material with regular structure, positioned in different orientations with respect to the camera. Unfortunately, methods [22],[23] require a specific sample shape or placement coupled with its extraction from the original environment. Isotropic SVBRDF can also be estimated from photometric stereo using a parametric reflectance model [9] or bivariate BRDF [1]. Finally, it is possible to use portable setups measuring SVBRDF by matching sparsely locally measured isotropic BRDFs (using condenser lens optics) with sparsely measured global reflectance fields [6]. Another approach records SVBRDF from a single view using 1 DOF-moving linear light source and a set of known BRDF samples recorded simultaneously with the sample [27]. Recently, sparse SVBRDF measurement and reconstruction have been used based on the measurement of several images of known geometry illuminated by a circularly polarized light [8]. Although this method requires the capture of only four sample images, its usage is limited to flat and isotropic measurements and it requires a complex measurement setup.

View and illumination dependent data interpolation is often performed when sparse images of known geometry and illumination direction are recorded. Reflectance data collected from such images can be interpolated either by a parametric reflectance model [17], or in the form of isotropic BRDFs interpolated by means of three-dimensional radial basis functions [32] in achieving a reconstruction of SVBRDF. Alternatively, a 4D BRDF can be decomposed into simpler 1D and 2D components having physical meaning, to allow parametric editing of visual properties [16].

However, to the best of our knowledge, no measurement technique yet exists enabling rapid capture of anisotropic ABRDF using an consumer camera and light. Proposed is a method for fast, non-restricted anisotropic ABRDF space reconstruction from extremely sparse samples that can be measured in a few seconds by continual movement of the camera and light. Contrary to parametric BRDF models [24] or other simplified solutions (see Fig. 2), this method is capable of correctly reconstructing non-reciprocal, non-energy-conserving ABRDF data.

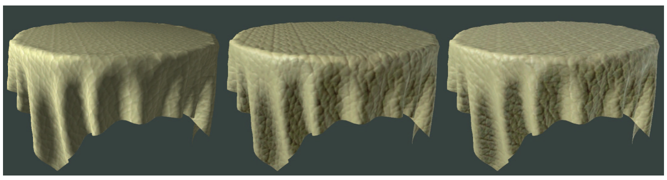

Figure 2: A comparison of renderings based on a single texture modulated by an analytical BRDF model [14] (left), reference BTF measurements using 6561 images (middle), and the proposed sparse data selection and reconstruction using 168 images (right).

Although the principle of the method has been outlined in [7], this paper provides additional thorough reasoning of the method's functionality. In addition, an introduction to a novel interpolation technique for missing elevations, results of real appearance acquisition, and the method's application on sparse reconstruction of BTF datasets are included.

3. The Proposed Reconstruction Method

A robust and sparse acquisition of general view and illumination dependent appearances is a tricky task. While a position of specular highlights can be expected near the mirror reflection, the location of anisotropic highlights is unknown. It depends on the local macro-geometry as well as on the micro-geometry of the measured surface. In our work we look for a very sparse set of illumination/view measurement points. They should allow a visually tolerable reconstruction of material reflectance, as well as a quick measurement of the sparse dataset using simple inexpensive hardware. There would be no need to preprocess the measured material sample or remove it from its environment.

Standard angularly uniform or even adaptive sampling strategies require many samples to preserve high frequencies in the data. On the other hand, employing analytical BRDF models imposes restrictions on data reciprocity and requires lengthy fitting, etc. Therefore, we analyzed a typical ABRDF and employed this knowledge to capture and reconstruct its behavior using a small set of measurements. This analysis has shown that it is most effective to place samples perpendicularly to specular highlights in a subspace of view/illumination azimuths. This is done in such a way that the samples form slices in the subspace and can be easily measured by horizontal movement of the light/camera around the measured sample. As the appearance of the azimuthal subspaces often depends on elevation angles, to create a more precise approximation we suggest sampling four combinations of view/illumination elevations.

A principle of the proposed method [7] is explained in Fig. 3. First, the material's ABRDF (Fig. 3-a) is sparsely measured in

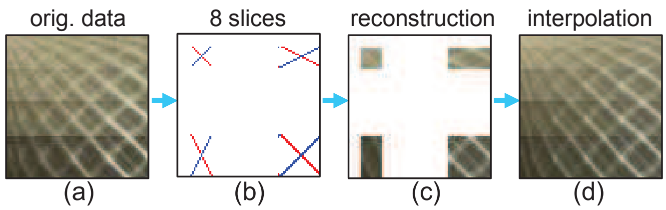

Figure 3: Example of ABRDF reconstruction: (a) original, (b) sparse-sampling using 8 slices, (c) reconstructions of elevations where the slices were measured, (d) missing data interpolation.

four subspaces by means of eight slices (Fig. 3-b), then the missing values in these subspaces are reconstructed from the values of the slices (Fig. 3-c), and lastly the remaining values at non-measured elevations are interpolated (Fig. 3-d).

3.1. Acquisition of Slices

Because one of our major concerns is the simplicity and speed of the acquisition process, we suggest taking samples by the continuous movement of light/camera around the sample at fixed elevations. By doing so, samples can be taken at an arbitrary density, limited only by camera movement speed and frame-rate. Each subspace of azimuthal angles φ_i/φ_v is sampled by means of two perpendicular slices (see Fig. 5-a), which differ in the direction of mutual movement of camera and light. In principle, the slices are, for a majority of the materials, orthogonal to their most prominent features: a specular reflection and an anisotropic reflection (see Fig. 4). These features are often constant in the direction perpendicular to the slices and thus can be effectively represented by their marginal values.

The slice aligned with the direction of the specular highlights is called *axial slice* s_A (red), i.e., $\varphi_v - \varphi_i = \alpha$ holds for azimuthal angles. The axial slice records the material's anisotropic properties (mutual positions of the light and camera are fixed while the sample rotates (see Fig. 14-left)), i.e., its value is almost a constant for near-isotropic samples.

The slice perpendicular to the highlights is called *diagonal slice* s_D (blue), i.e., $\varphi_i + \varphi_v = 2\pi$ holds for azimuthal angles. The diagonal slice captures the shape of the specular peaks (light and camera travel in mutually opposite directions over the sample (see Fig. 14-right)).

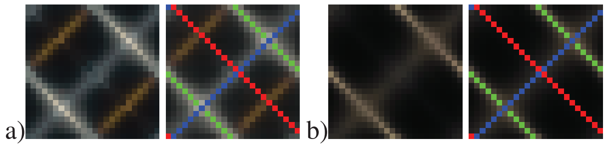

Figure 4: Examples of ABRDF toroidal subspaces for derivation of optimal placement of the axial slices (green and red alternatives).

We focused first on analysis of the ABRDF subspace having the highest elevations, in which the illumination and view dependent effects are the most pronounced. See in the first row of Fig. 7. The azimuthal difference of light and camera during axial slice measurement – α influences the placement of slices in the ABRDF toroidal subspace; therefore, we analyzed optimal placement of axial and diagonal slices across a number of ABRDFs. The study has shown that while the placement of a diagonal slice can be arbitrary, the highest variance along axial slices is achieved near the specular highlight. Consequently, this is – most likely – the best placement of the axial slice $\alpha = 180^o$ (green dots in Fig. 4). However, such a placement might omit vital color/luminance information in some parts of the subspace. For example, it would completely miss yellow anisotropic features as in Fig. 4-a or dark parts as in Fig. 4-b. Therefore, we used the slice with the second highest variance $\alpha = 15^o$ (red dots in Fig. 4) to eliminate occlusion of the camera with the light and capture most visual features of the ABRDF subspace (see the shift of the red axial slice from image diagonal in Fig. 5-a).

For experimental purposes the slices can be taken from the measured ABRDF (Fig. 5-a) as

$$
\begin{aligned}
s_{A,\theta_i\theta_v}(\varphi_i) &= ABRDF(\theta_i, \theta_v, \varphi_i, \varphi_v = \varphi_i + \alpha) \ , \\
s_{D,\theta_i\theta_v}(\varphi_v) &= ABRDF(\theta_i, \theta_v, \varphi_i = 2\pi - \varphi_v, \varphi_v) \ .
\end{aligned}
\tag{1}
$$

3.2. Reconstruction from slices

ABRDF toroidal subspace reconstruction is performed for elevation angles at which the slices were captured. It can be explained as a combination of two slices (i.e., sets of marginal values) as shown in Fig. 5. The reconstruction of point $\widehat{ABRDF}(\theta_i, \theta_v, \varphi_i, \varphi_v)$ in ABRDF subspace starts with combining contributions of the s_A and s_D slices. We tested their sum and product; however, the latter improperly enhanced the locations at intersections of the specular and anisotropic highlights as shown in Fig. 6-b. Note that the sum of slice contributions (see Fig. 6-c) preserves specular highlights, which are less affected by the anisotropic highlights. Therefore, we finally used the sum of slices in our reconstruction procedure:

$$
v_{\theta_i\theta_v}(\varphi_i, \varphi_v) = s_{A,\theta_i\theta_v}(\varphi_{i,R}) + s_{D,\theta_i\theta_v}(\varphi_{v,R}) \ , \tag{2}
$$

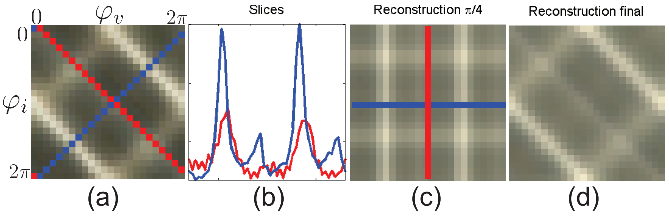

Figure 5: Reconstruction of a toroidal ABRDF subset from two slices at fixed elevations: (a) reference data with slice placements, (b) data profiles in the slices, (c) reconstruction from slices ($\frac{\pi}{4}$ rotated), (d) final reconstruction.

$$\begin{bmatrix} \varphi_{i,R} \\ \varphi_{v,R} \end{bmatrix} = \begin{bmatrix} \cos(\pi/4) & -\sin(\pi/4) \\ \sin(\pi/4) & \cos(\pi/4) \end{bmatrix} \begin{bmatrix} \varphi_i \\ \varphi_v \end{bmatrix} .$$

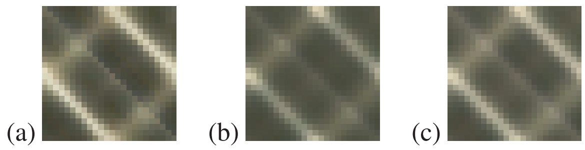

14.0 / 20.9 / 21.7 9.6 / 12.2 / 26.5

Figure 6: Original subspace (a), and its reconstructions using product of slices (b) and sum of slices (c). Below are the difference values in: CIE Δ E / RMSE / PSNR[dB].

Note that the original azimuths φ_i, φ_v had to be rotated for $\pi/4$ (Fig. 5-c) to account for the slant of slices with respect to the φ_i, φ_v coordinate system (Fig. 5-a). Finally, the summed value v is mapped to a dynamic range of original slices

$$\widehat{ABRDF}(\theta_i, \theta_v, \varphi_i, \varphi_v) = v_{\theta_i\theta_v}(\varphi_i, \varphi_v) \cdot (M - m) + m \quad , \quad (3)$$

$$m = \min(s_{A,\theta_i\theta_v} \cup s_{D,\theta_i\theta_v}) \qquad M = \max(s_{A,\theta_i\theta_v} \cup s_{D,\theta_i\theta_v}) \quad .(4)$$

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

$$\widehat{ABRDF}(\theta_i, \theta_v, \varphi_i, \varphi_v) = s_{D,\theta_i\theta_v}(\varphi_{v,R}) \quad . \quad (5)$$

Fig. 7 shows a reconstruction of anisotropic ABRDF subspace at elevations $\theta_i/\theta_v = 75^o/75^o$ (the second row) from two slices (the third row) and prove the ability of the proposed approach to represent a variety of anisotropic materials.

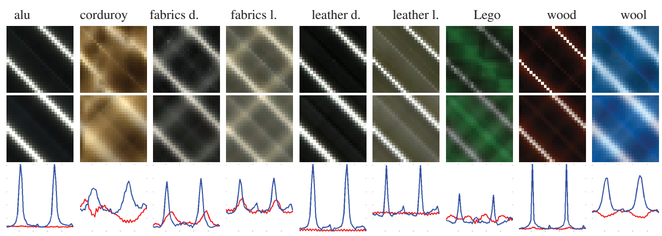

Figure 7: Comparison of the material's ABRDF toroidal subspace at elevation 75^o (the first row), with its reconstruction (the second row) from the axial (red) and diagonal (blue) slices (the third row).

3.3. Interpolation of missing values

At this point, sparse acquisition and reconstruction of the ABRDF subspace has been explained. However, the selection of elevations at which the slices are measured significantly influences the final ABRDF reconstruction. Therefore, we performed an experiment with two measured ABRDFs (isotropic specular and diffuse anisotropic material) in order to find the proper combination of two elevations at which the four subspaces should be measured. We tested six different combinations of elevations. Only samples from these illumination/view elevations were used for the entire ABRDF interpolation using radial basis functions [26]. The average RMSE differences between ground-truth ABRDF data and interpolation shown in Fig. 8 suggest that the combination of $\theta = 45^o/75^o$ provides the lowest reconstruction error. Therefore, we have chosen

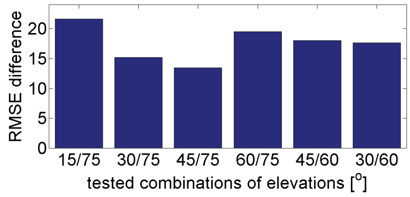

Figure 8: ABRDF reconstruction error (RMSE) for different elevation combinations used for selection of four measured subspaces.

the highest elevation angles $\theta_i = \theta_v = 75^o$, where the specular reflections are the most intensive (see first row of Fig. 12). The lower elevation angles were decreased to $\theta_i = \theta_v = 30^o$ (the second best choice from Fig. 8) for better representation of material appearance at orthogonal viewing and illumination directions, which are the most visually salient. More than these four subspaces can be used at the expense of more camera/light elevations; however, this would increase the number of samples and the complexity of their measurement. Finally the four subspaces at the following elevations were sampled: $\theta_i/\theta_v = 30^o/30^o, 75^o/75^o, 30^o/75^o, 75^o/30^o$,).

However, data for the remaining subspaces are still unknown and have to be estimated. The BRDF parametric models, e.g., [23], cannot be used to solve this problem because they impose restrictions on data properties (reciprocity, energy conservation, etc.), require many more samples or a different distribution of samples, and lengthy fitting. They also depend on initial values. We tried to fit measured samples using the anisotropic parametric BRDF model [14]. Due to a low number of samples and their distribution we were unable to find a stable parameter fit for most of the tested ABRDFs. Moreover, these models are not designed to handle non-reciprocal ABRDF data. Therefore, we tested the following two interpolation approaches:

Method A – In the first one, the interpolation was performed by means of the four-dimensional radial basis functions [26] 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 [28], and finally used parameterization according to [11], applied to both illumination and view directions $[\alpha_i, \beta_i, \alpha_v, \beta_v]$. This parameterization has shown the lowest reconstruction error due to alignment of specular highlights 0^o value of angle β_i.

Method B – Due to relatively high computational demands of Method A, we developed a faster hybrid linear interpolation constrained by a reflectance model. This interpolation consists of two steps shown in Fig. 9. First we interpolate data at different viewing and constant illumination elevations. Then the remaining illumination elevations are filled. In each interpo-

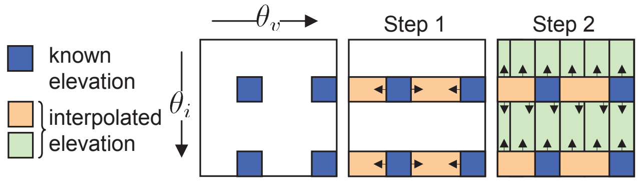

Figure 9: Steps of the interpolation approach A.

lation step, the average ABRDF value for each unknown elevation was approximated by fitting a simplified monospectral one-lobe Lafortune model [15] with parameters k,α to known slice values

$$f_r(k, \alpha) = k(\cos\theta_i \cdot \cos\theta_v)^{\alpha} \quad . \tag{6}$$

Initialization of k, α was constant during all experiments. Values obtained from the model at elevations θ are scaled by the mean values of the slices and then used for obtaining interpolation weights. These weights are then applied for a linear interpolation of missing elevations from slice values at known elevations as shown in Fig. 10. Elevations lower than 30^o are extrapolated using the scaled model's (6) predictions. This procedure is performed over all azimuthal directions as shown in Fig. 9.

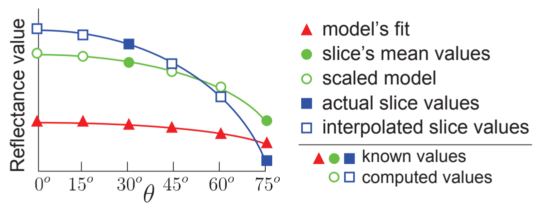

Figure 10: An interpolation of non-measured elevation values.

Method A is approximately three times more computationally intensive than Method B and provides better results in most cases. While Method A allows arbitrarily dense sampling, even for originally unmeasured azimuthal directions, Method B reconstructs ABRDF data in their original azimuthal sampling. The results in Fig. 11 show the major visual differences between both proposed interpolation methods in two illumination environments.

We also tested modification of this step-wise interpolation of subspaces (Fig. 9) using a displacement interpolation (denoted as BD) method [2]. Compared to the weighted linear interpolation (Method B), its principle is based on solving the generalized mass transport optimization problem. As this method cannot extrapolate, the low elevation areas were reconstructed using Method B. Method BD gives better results than Method B and comparable results to Method A. The whole subspace interpolation is about thirty times slower than Method A, while when view or illumination direction is fixed, i.e., individual lines in subspaces are interpolated separately, its speed compares to Method A. For this reason we have not used Method BD further in the paper.

Although both global (A) and local (B) interpolation approaches provide similar visual quality (see more discussion in Section 4), we believe their performance can be further improved, e.g., using estimated height-map as an interpolation constraint.

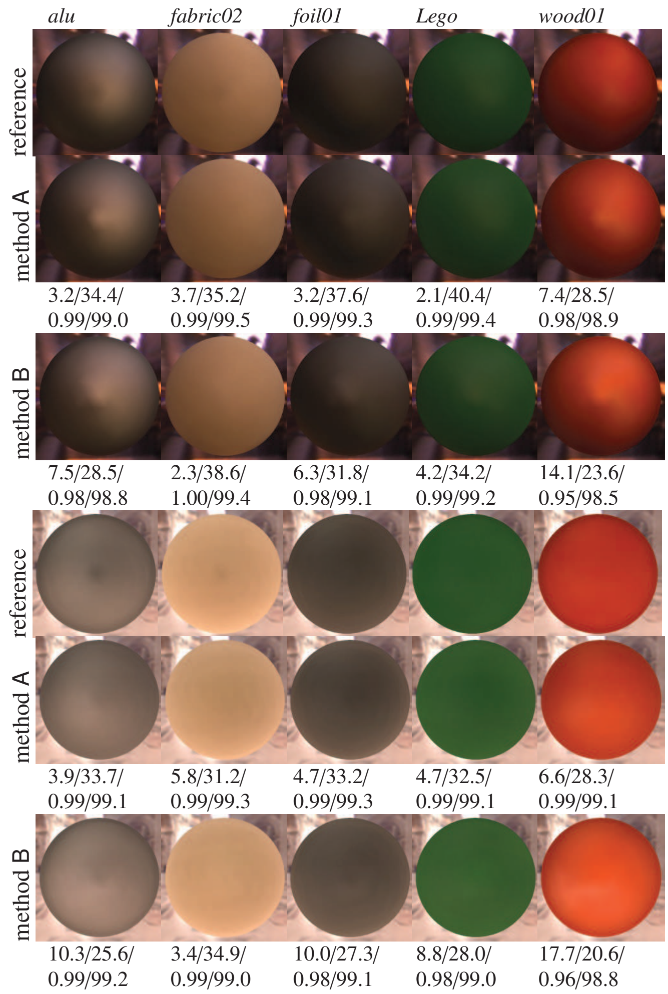

Figure 11: A comparison of the interpolation methods performance in *grace* and *st.peters* illumination environments [5] . Below are the CIE ΔE / PSNR[dB] / SSIM / VDP2 difference values.

4. Results of Simulated Measurement

In this section we show results of sparse reconstruction experiments performed on isotropic BRDF and anisotropic ABRDF data. The data served as a source of sparse sampling and were simultaneously used for evaluating reconstruction quality of the method.

Generally, the angular resolution used in all experiments in this paper was 81 view $\times$81 illumination directions (6561 values) [29] distributed uniformly over the hemisphere (Fig. 1) To sample this resolution in the slices of the proposed method we need only 168 samples to obtain information sufficient for data reconstruction.

In the first experiment, we tested our method on reconstruction of 55 isotropic BRDF samples (resampled to 81×81 directions) from the MERL BRDF database [19]. The advantage of isotropic reconstruction is that only four diagonal slices s_D

have to be obtained (in our case 84 samples instead of the 168 needed for anisotropic data). Mean reconstruction errors of all 55 BRDFs (8bits/channel) were: CIE ΔE=9.1, RMSE=15.7, and PSNR= 24.9 [7].

In the second experiment, ten BTF samples (nine from Bonn University BTF database[1] and one from Volumetric Surface Texture Database[2]) were used (aluminum profile, corduroy, dark and light fabrics, dark and light leatherettes, lacquered wood, knitted wool, upholstery fabric Proposte, and Lego). These materials, due to their rough structure and often non-opaque properties, exhibit anisotropic effects of occlusions, masking, subsurface scattering and therefore represent a challenging dataset to test the proposed method. All BTF pixels were averaged to obtain the average ABRDF of the material (first row of Fig. 12). The results of the complete reconstruction of original ABRDFs from 168 sparse samples are shown in Fig. 12. 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 Fig. 12 show that the global interpolation Method (A) estimated incorrect values for elevations between the two sampled elevation values (30^o and 75^o). On the other hand, Method (B) gives, due to a lack of global knowledge, the worst estimation for extrapolated elevations. That is, elevations smaller than 30^o as represented by the first few rows/columns in the images. Finally, the displacement interpolation – Method BD scored similarly to Method A. The reconstruction and interpolation of a single ABRDF from 168 samples take ≈1 second using interpolation Method A, ≈0.3 second using Method B, and ≈1 second using Method BD on Intel Xeon 2.7GHz (using 3 cores).

To validate the contribution of our method, we compared its reconstruction performance using 168 samples with the uniform sampling of a similar samples count. For this purpose, hemispheres of illumination and viewing directions were sampled by means of 13×13 samples, producing a total of 169 samples. Then the missing values in the ABRDF space were interpolated from these sparse samples by means of four-dimensional radial basis functions [26] (Method A). The interpolation was computed separately in each color channel, and $0 \approx 2\pi$ discontinuity has been avoided using the onion parameterization of illumination and view directions [11]. Comparison of ten interpolated ABRDFs (see Fig. 13) has shown that the proposed reconstruction method has a better performance than the interpolation from uniform samples, mainly near specular highlights, as confirmed by the objective criterion values shown below the reconstructions. On average, the proposed reconstruction provides 1.4 and 3.2 lower ΔE / RMSE values and 1.8 higher PSNR value across ten tested ABRDFs. Moreover, the data acquisition process using our method is considerably faster and less demanding on hardware as shown in Section 6.

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

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

5. Limitations

The limitations of the proposed method are threefold. First, since the method restores reflectance at given elevations only from two orthogonal slices, it cannot reliably capture features that are not orthogonal to the slices (see second example of *corduroy* in Fig. 7). It must also be noted that the proposed slices represent a very sparse sampling of the azimuthal subspace and as such, can omit some reflectance features, resulting in a slightly different color/brightness appearance of the reconstructed data. To avoid this problem, the azimuthal subspace can be sampled by additional slices at the cost of slightly longer acquisition times. Second, the interpolation step of the algorithm expects monotonicity of reflectance values across different illumination and view elevations. However, this condition is rarely invalid and no such behavior was experienced with any of the tested materials. The method's accuracy can be further improved in this respect by taking more slices at different elevations. Finally, highlights of extremely specular samples are not always represented accurately enough (see Fig. 19) mainly due to an insufficient angular sampling of azimuthal angles (step 15^o) in original datasets used in the experiments. Note that the sampling density along specular highlights in diagonal slices can be arbitrarily increased to provide a better match of specular highlights of a high-dynamic-range within the model without increasing the measurement time.

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 measuring the slices several times and computing the measurements' median values.

6. Sparse ABRDF Data Measurement

This section describes a practical experiment of capturing sparse ABRDF samples using a consumer camera and a LED point-light source and is followed by a complete ABRDF reconstruction from such measurements.

Mutual movement of arms with camera and light with respect to the sample being measured is controlled manually as is shown in Fig. 14. The axial slice s_A data (left) are measured

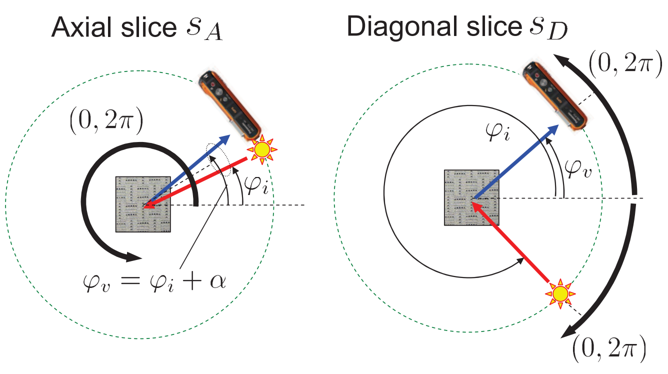

Figure 14: The proposed ABRDF measurement setup at fixed elevation angles θ_i/θ_v.

using rotation of the fixed light and sensor around the sample,

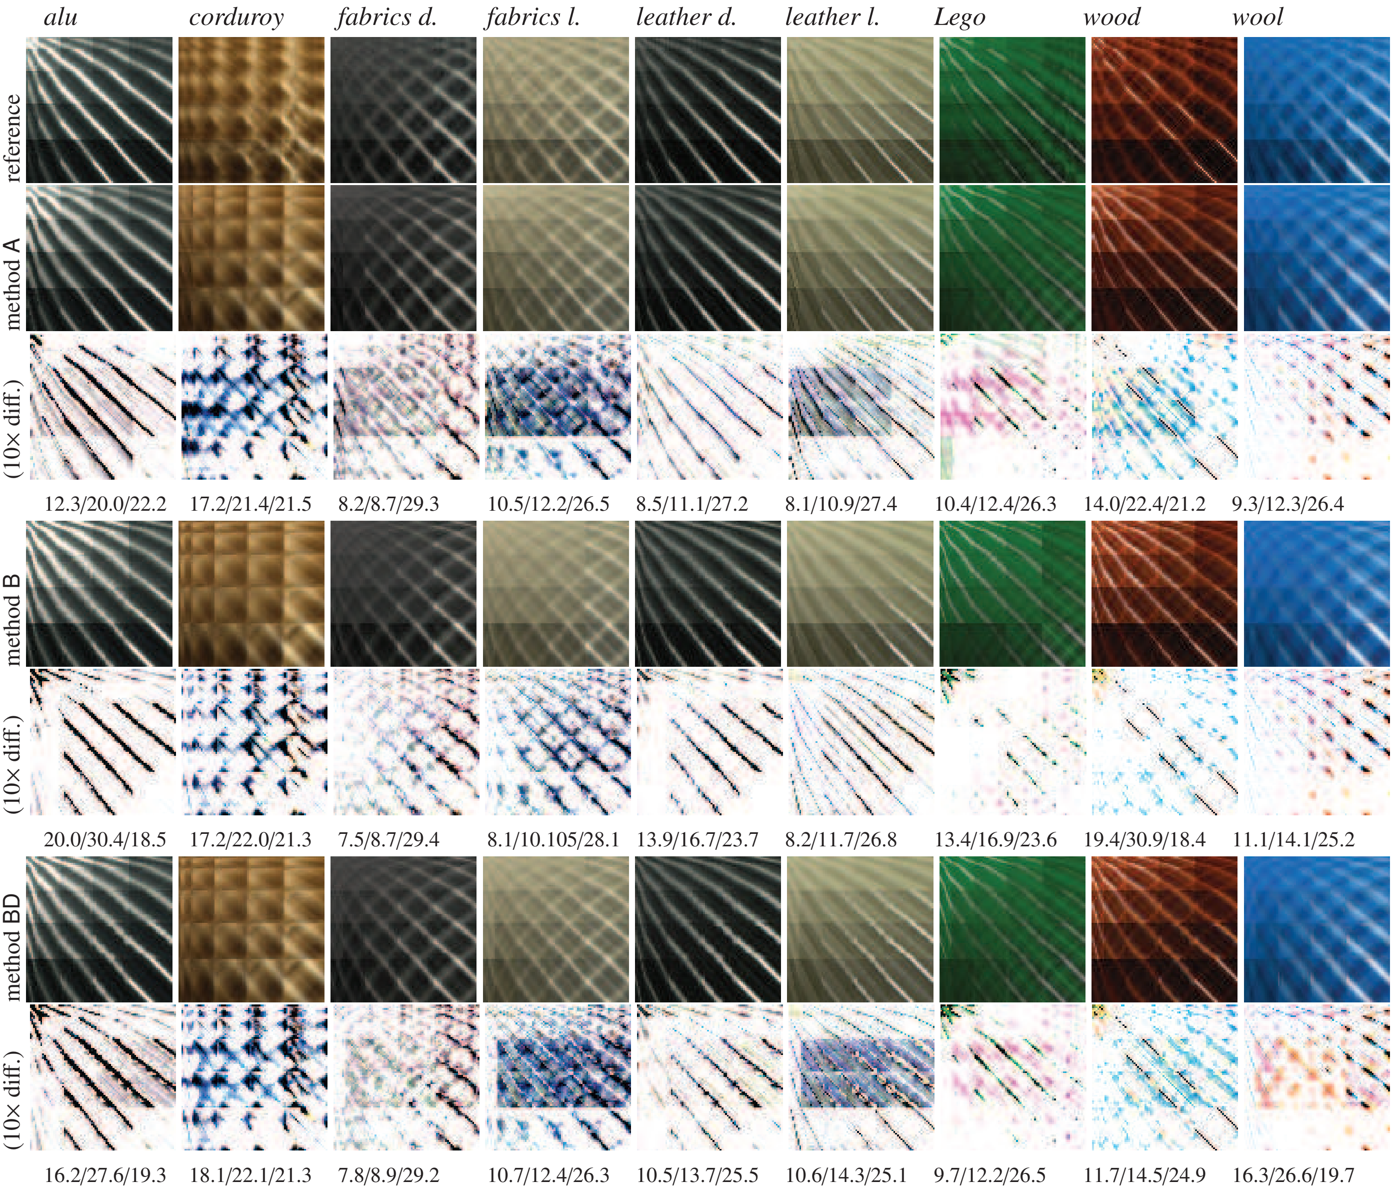

Figure 12: Comparison of the material's BRDF (the first row), and interpolation of missing values by means of method A (the second row), method B (the third row), and its modification method BD (the fourth row) respectively. Below are 10× difference images and global difference values in CIE ΔE / RMSE / PSNR[dB].

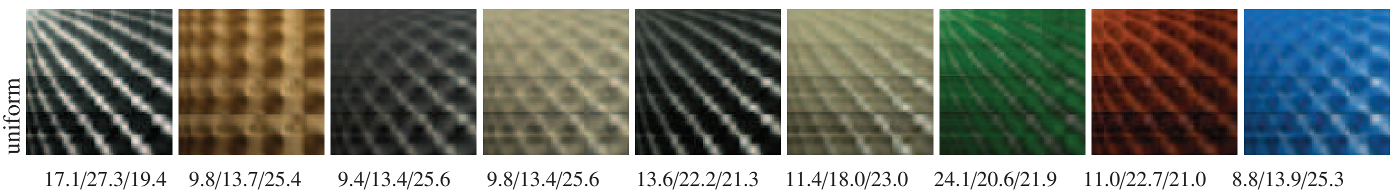

Figure 13: The performance of interpolation from 169 sparse uniform samples (13 samples per hemisphere). Below are the difference values in: CIE ΔE / RMSE / PSNR[dB].

while the diagonal slice s_D data (right) are obtained by mutually opposite movements of the light and sensor in respect to the sample. Both the camera and light travel full circle around the sample and return to the initial position.

Our acquisition setup consisted of the Panasonic camera Lumix DMC-FT3 and light using high-power LED Cree XLamp XM-L with 20^o frosted optics (Fig. 15-a). To achieve the required synchronous movement of light and camera, we constructed a frame with two arms using a Merkur toy[3] construction set shown in Fig. 15-b. During its movement, the camera records the material sample appearance as a video sequence at a resolution of 1280×720 pixels, and the elevation angles of the camera and light are kept constant using the setup Fig. 15-b. Both s_A and s_D slices are recorded for two different ele-

[3]http://www.merkur.cz/

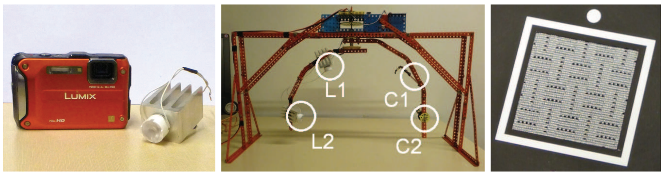

(a) (b) (c)

Figure 15: Data acquisition equipment (a) with its fixating frame (b), and measured sample with registration borders for calibration (c).

vations of the camera (C1, C2) and light (L1, L2); therefore, eight slices are measured approximately at elevations θ_i/θ_v = $[30^o/30^o, 30^o/75^o, 75^o/30^o, 75^o/75^o]$ as shown in Fig. 3-b. Recording of the slices took less than 10 minutes. From each of the eight video sequences, 24 frames were extracted corresponding to sampling of azimuthal angles φ_i/φ_v every 15^o. This resulted in a total of 192 samples being obtained. The number differs from 168 samples used in the reconstructions in Section 4, because this time all elevations were covered by the same number of samples. The effective number of samples is always slightly lower than 192, as some of the frames are removed due to occlusion of the material by the arm with light. Note that the method's principle allows adaptive density of the samples (frames) along the slices to also record extremely narrow specular highlights.

Three anisotropic fabric materials (30×30 mm) were used as test samples, as shown in Fig. 16. A white border was attached around the sample to help detect camera orientation in respect to the sample coordinate space and for sample registration Fig. 15-c). Because of this, we first calibrated the camera [31]. Unfortunately, the used low-end camera adapts its exposure depending on the amount of light coming from the scene. On the other hand, this feature enables us to capture as much information as possible, even using a limited dynamic range of the camera's sensor (8bits/color). Since the information about exposure throughout the video sequence could not be retrieved from an EXIF header as is possible for still photos, we used the reference BRDF data of dark material surrounding the sample to compensate for exposure of each image. That is, we compensated color values of the sample using the black part of the sample holder (near the white borders as shown in Fig. 15-c) and its reference measurements.

Subsequent processing was then performed for each image. Camera viewing angles θ_v/φ_v were obtained from camera extrinsic parameters, given the known camera calibration and corner points of the white borders. Coordinates of these points were obtained from the image registration based on the camera calibration. When the viewing angles were known, the illumination azimuth angle was computed as: $\varphi_i = \varphi_v - \alpha$ for the axial slice s_A, and $\varphi_i = 2\pi - \varphi_v$ for the diagonal slice s_D. The elevation angles θ_i were estimated from the slant of the light during measurement. Finally, the slice's ABRDF value from each image was obtained as the average of RGB values near the sample's center, and colorimetrically calibrated. The non-optimized data processing described above took approximately 10 minutes to perform over all selected images. The reference ABRDF measurements of the black target and materials are obtained from the UTIA BTF database[4] and have the same angular resolution as BTF Database Bonn [29].

When all of the selected images were processed in this way and data for all eight slices were obtained, the ABRDF space reconstruction described in Sections 3.2 and 3.3 was performed. Figure 17 compares reference ABRDF measurements of the material (a) with their reconstruction from the 192 sparse reference samples using Method B (b), and with a reconstruction using 192 sparse measurements obtained by the proposed setup and interpolation Method B (c). The last column (d) of Figure 17 compares our method with a uniform sampling using 196 samples (14^2), while the remaining samples are interpolated using Method A (compare with column (b)). Note that, while the visual performance of the uniform sampling might look similar, the complexity of its measurement is considerably higher in comparison with the proposed measurement approach.

fabric01 fabric02 fabric03

Figure 16: Three anisotropic fabric samples whose ABRDFs were measured using the proposed setup.

Finally we took photographs of the *fabric02* and *fabric03* materials attached on a cylinder (a) and compared them with renderings on a cylinder using their reference BRDFs (b) and BRDFs captured by the proposed setup (c). The results for different illumination conditions are shown in Fig. 18 and confirm that even the proposed approximate measurement setup can record BRDFs with reasonable accuracy, in comparison to the reference measurements.

The reconstruction results from our preliminary measurements (Figure 17-c) are encouraging and we believe that they convey the idea of ABRDF capturing speed and simplicity without the need for dedicated and thus costly devices.

A notable advantage of our setup and the proposed sampling pattern is its ability to quickly measure any flat samples without needing to extract them from their environment, and thus it can be used for fast and inexpensive measurements of such samples as human skin and precious cultural heritage objects.

7. Experimental BTF Reconstruction

As the acquisition and reconstruction of spatially-varying datasets is a straightforward extension of the proposed sparse sampling and reconstruction method, we tested the method's performance on ten BTF samples of angular resolution 81 × 81 = 6561 images as described in Section 4. Only 168 images (corresponding to the eight data slices) were selected from the BTF samples and used for pixel-wise reconstruction of the remaining images using the proposed method.

[4]http://btf.utia.cas.cz

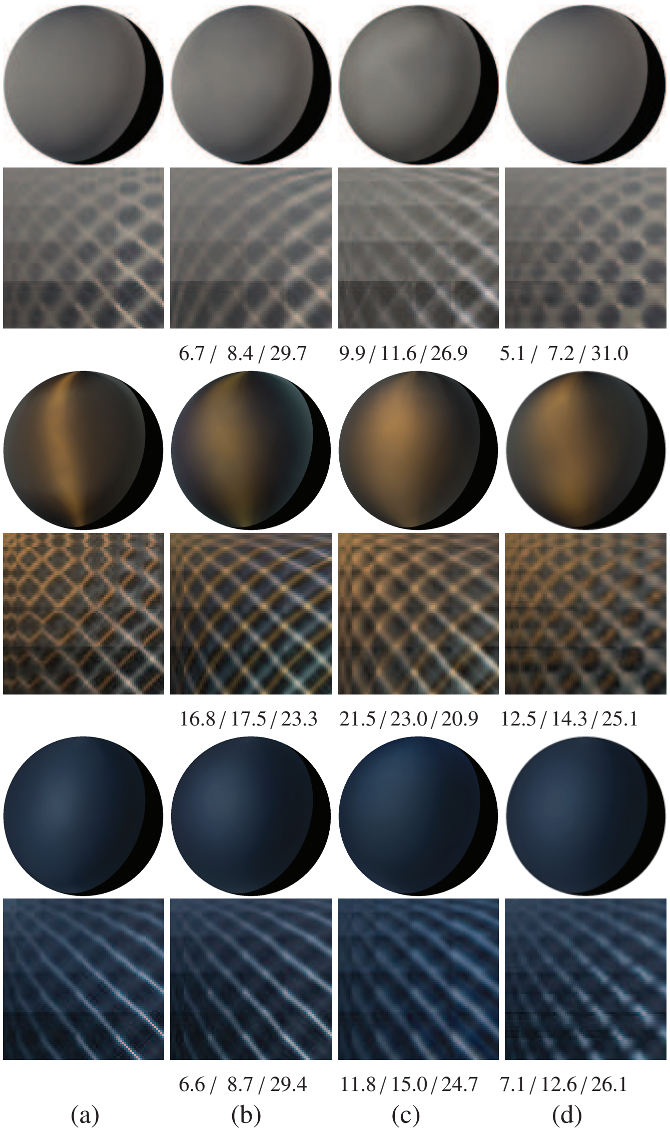

Figure 17: ABRDF reference measurement (a), compared to reconstruction from 168 sparse reference measurements only using Method B (b), and reconstruction from the proposed measurement procedure using 192 samples and interpolation Method B (c), uniform interpolation using 196 samples (d). Below are the difference values in: CIE ΔE / RMSE / PSNR[dB].

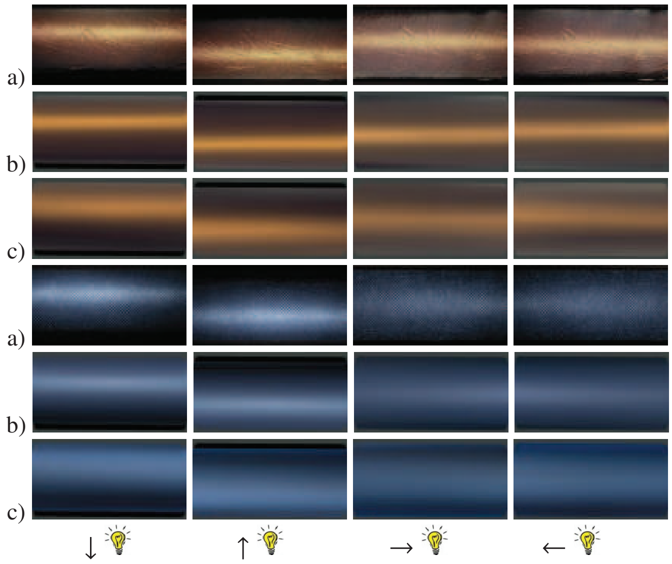

Figure 18: Photographs of the *fabric01* and *fabric02* samples on a cylinder illuminated from top, bottom, left, and right (a) compared with renderings using reference ABRDF (b), and sparsely measured and reconstructed ABRDF (c).

7.1. Results

Renderings of the original data with results of the proposed reconstruction methods for point-light and environment illumination is shown side-by-side in Figures 19 and 20, respectively. All differences are objectively compared using CIE ΔE, PSNR[dB], SSIM [30], and VDP2 [18] metrics. From the results it is apparent that for samples with lower height variations, there is a close match to the original data. The apparent deviations from the original data for materials having higher surface height variations are caused mainly by the incorrect geometry preservation of structural elements.

7.2. Limitations

Although there are not any restrictions imposed on view and illumination dependent datasets, the results have shown that the BTF reconstruction is incorrect for those materials which have a wide range of surface height variation, e.g., *corduroy* and *Lego* samples shown in Fig. 21. This is caused partly by very sparse sampling of the azimuthal space, as well as by interpolation of the data at missing elevations. While the former produces geometrical deformation of the structure's features, the latter causes their blur as well as improper highlights extrapolation for low elevation angles. Even though the reconstruction from sparse samples for such materials is not accurate in terms of correct shading of structure elements, the method correctly captures the look-and-feel of the material's spatially-varying appearance for nearly-flat samples, e.g., for *fabric dark*, *fabric light*, and *leather light* samples. However, in comparison with the SVBRDF measurement and representation approaches, the proposed method is not limited to restrictions imposed by BRDF itself. Therefore, it may be found useful for quick, low-cost, and fairly accurate acquisition and BTF reconstruction of many materials having a limited height variation, e.g., fabric and leather.

The time of BTF data reconstruction depends only on its spatial resolution, since individual pixels are regarded as independent ABRDFs. Due to huge sizes of datasets, only repetitive BTF tiles were used. While reconstruction of a single pixel took ≈ 1 second, the non-optimized reconstruction of a BTF tile of size 128^2 took 4.5 hours using the 3 cores of the Intel Xeon 2.7 GHz. Therefore, using optimized multi-core CPU's implementation processing times of less than one hour can be easily achieved.

Note that the proposed sparse acquisition and reconstruction method is complementary to BTF compression methods. In always processing an entire BTF dataset, any of these methods can be applied to compress the reconstructed data. By its sparse measurement, our method can achieve a compression ratio 1:39; however, in terms of reconstruction quality and compression ratio, it cannot compete with BTF compression approaches, as seen in the local PCA method [21] (using 5 clus-

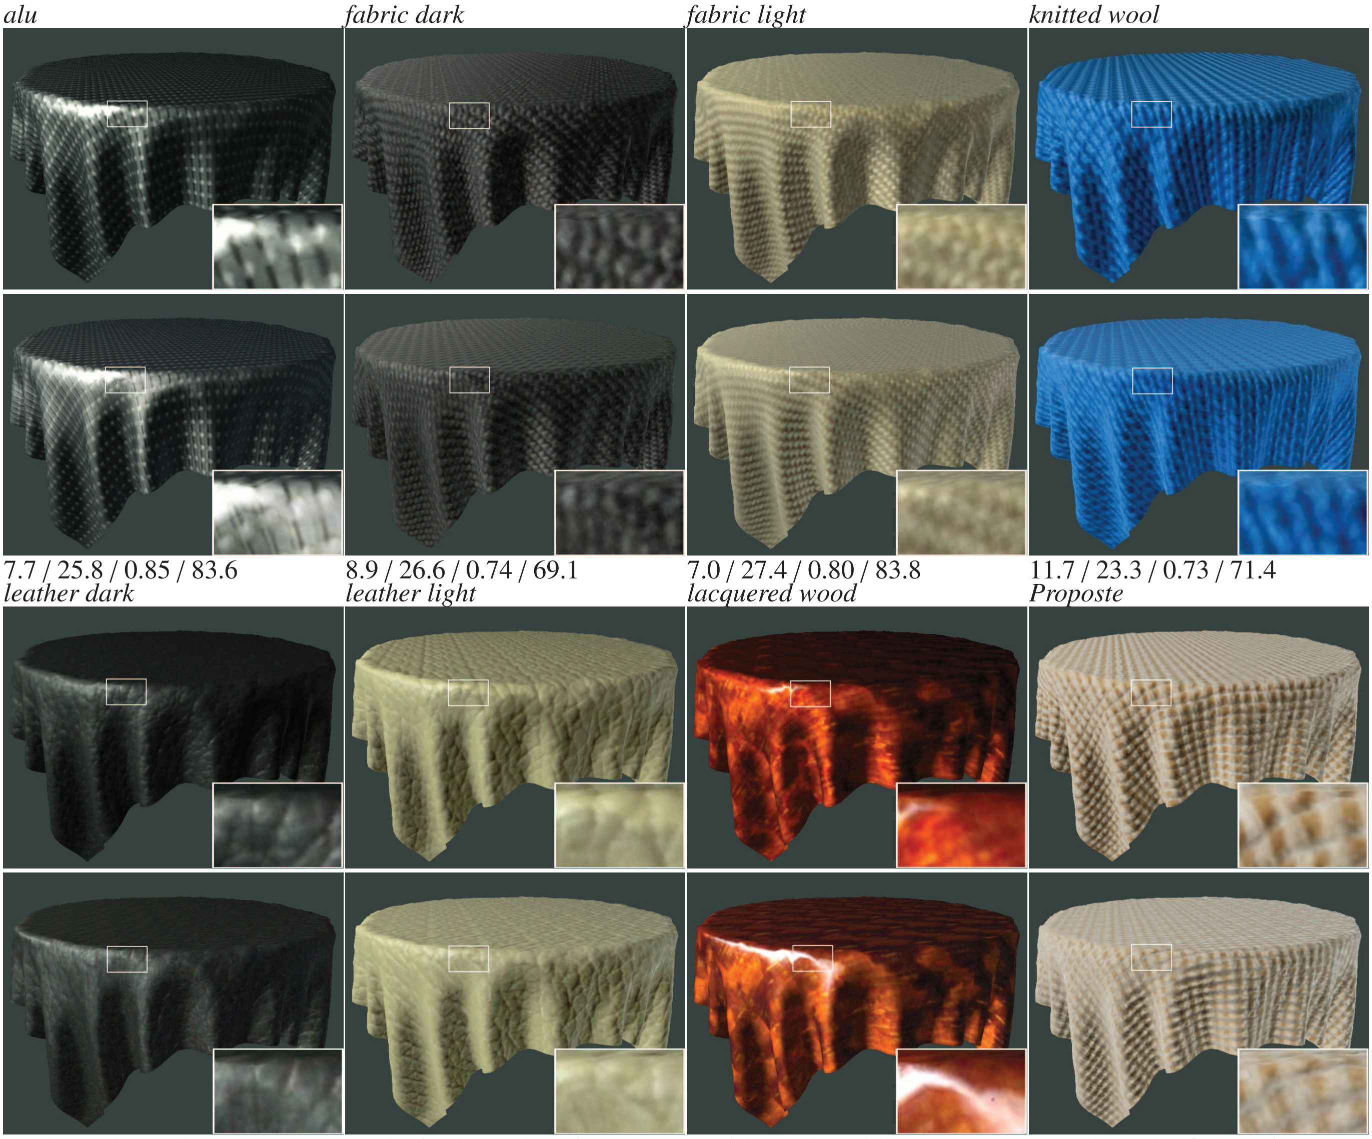

Figure 19: A comparison of BTF rendering from the full dataset of 6561 images (the first row), with its reconstruction from only 168 images (the second row) in single point-light illumination. Below are the CIE ΔE / PSNR[dB] / SSIM / VDP2 difference values.

ters, 5 components) in Tab. 1 (compare with our reconstruction errors in Fig. 19). The variable compression ratio of the local PCA method is due to the variable size of the BTF tile used.

Table 1: A reconstruction error and compression ratio of LPCA compression method.

material	ΔE / PSNR[dB] / SSIM / VDP2				C.R.	(tile)
alu	3.5	34.6	0.99	93.8	18.2	(21×26)
corduroy	2.7	37.5	0.97	93.6	55.0	(36×46)
fabric d.	4.6	32.2	0.94	86.2	16.1	(21×23)
fabric l.	8.6	26.3	0.95	89.4	10.4	(19×23)
leather d.	4.3	33.3	0.97	91.0	263.0	(93×86)
leather l.	10.8	24.3	0.95	87.7	193.0	(74×79)
l. wood	11.3	23.5	0.92	81.0	627.8	(137×142)
k. wool	4.3	32.1	0.96	93.0	20.8	(25×25)

8. Conclusions

A novel method of sparse measurement and reconstruction of view and illumination dependent datasets has been proposed. The proposed sparse sampling of illumination and viewing directions allows for intuitive continuous measurement by a consumer camera and LED light. The reconstruction from such sparse data does not impose any restrictions on input data and allows reliable approximation of anisotropic non-reciprocal view and illumination dependent datasets. Additionally, this method can provide arbitrarily dense data reconstruction of both incoming and outgoing directions. The method's performance was tested on isotropic BRDFs and anisotropic apparent BRDFs

Figure 20: A comparison of BTF rendering from the full dataset of 6561 images (the first row), with its reconstruction from only 168 images (the second row) in *grace* environment illumination [5]. Below are the CIE ΔE / PSNR[dB] / SSIM / VDP2 difference values.

with encouraging results. Our pilot ABRDF measurement experiments have shown that retrieval of sparse samples and the consequent reconstruction of the complete dataset take less than half an hour. Experimental sparse reconstruction of BTF datasets has shown that the method can be a reasonably accurate alternative to lengthy measurement, especially for samples having a smaller height variation. The ease of data acquisition and visual quality of the reconstruction using this method makes it superior to alternative approaches such as bump/displacement mapping or parametric BRDF modeling. Because of the simplicity of data acquisition and reconstruction, this approximate method can be utilized in less accuracy-demanding applications. Since digital reproduction of a material's appearance look-and-feel can be created inexpensively, it could be particularly useful in the fields of computer gaming, film and digital presentations of e-commerce.

In summation, we believe this research will contribute to future development of simple, inexpensive, and portable acquisition setups of illumination and view dependent data.

Acknowledgments

This work has been supported by the Czech Science Foundation grant GAP 103/11/0335 and EC Marie Curie European Reintegration Grant No 239294.

References

[1] N. G. Alldrin, T. Zickler, and D. Kriegman. Photometric stereo with non-parametric and spatially-varying reflectance. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2008.

[2] Nicolas Bonneel, Michiel van de Panne, Sylvain Paris, and Wolfgang Heidrich. Displacement interpolation using lagrangian mass transport. *ACM Transactions on Graphics*, 30(6), December 2011.

[3] K.J. Dana, B. van Ginneken, S.K. Nayar, and J.J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Transactions on Graphics*, 18(1):1–34, 1999.

corduroy *Lego*

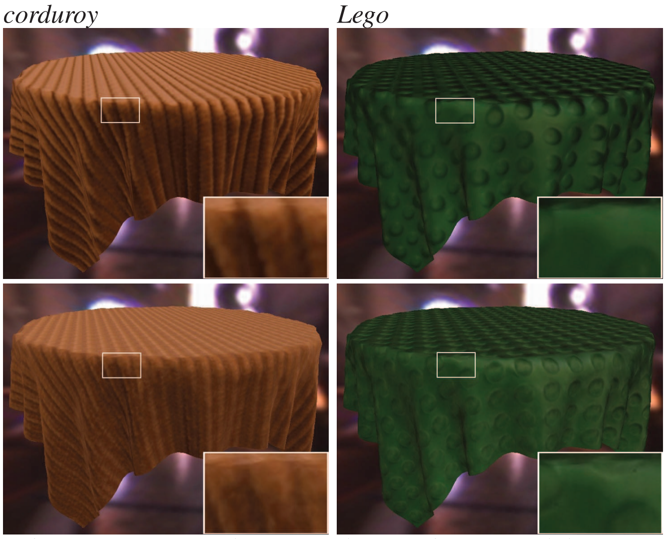

7.6 / 28.1 / 0.82 / 74.5 5.9 / 30.6 / 0.93 / 85.5

Figure 21: A failure case of the proposed BTF reconstruction for rough materials: *corduroy* and *Lego*. Below are the CIE ΔE / PSNR[dB] / SSIM / VDP2 difference values.

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

[5] Paul Debevec. 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*, SIGGRAPH '98, pages 189–198, 1998.

[6] Yue Dong, Jiaping Wang, Xin Tong, John Snyder, Yanxiang Lan, Moshe Ben-Ezra, and Baining Guo. Manifold bootstrapping for SVBRDF capture. *ACM Transactions on Graphics*, 29(4):98:1–98:10, 2010.

[7] J. Filip. Restoring illumination and view dependent data from sparse samples. In *Proceedings of the 21th International Conference on Pattern Recognition, ICPR 2012*, pages 1391–1394, November 2012.

[8] A. Ghosh, T. Chen, P. Peers, C.A. Wilson, and P. Debevec. Circularly polarized spherical illumination reflectometry. *ACM Transactions on Graphics*, 29(6):162:1–162:12, 2010.

[9] D. B. Goldman, B. Curless, A. Hertzmann, and S. M. Seitz. Shape and spatially-varying BRDFs from photometric stereo. In *Proceedings of ICCV*, ICCV '05, pages 341–348, 2005.

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

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

[12] M. Holroyd, J. Lawrence, and T. Zickler. A coaxial optical scanner for synchronous acquisition of 3D geometry and surface reflectance. *ACM Transactions on Graphics*, 29(4):99:1–99:12, 2010.

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

[14] M. Kurt, L. Szirmay-Kalos, and J. Křivánek. An anisotropic BRDF model for fitting and monte carlo rendering. *SIGGRAPH Comput. Graph.*, 44:3:1–3:15, February 2010.

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

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

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

[18] Rafat Mantiuk, Kil Joong Kim, Allan G. Rempel, and Wolfgang Heidrich. Hdr-vdp-2: a calibrated visual metric for visibility and quality predictions in all luminance conditions. *ACM Transactions on Graphics*, 30(4):40:1–40:14, 2011.

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

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

[21] G. Müller, J. Meseth, and R. Klein. Compression and real-time rendering of measured BTFs using local PCA. In *Vision, Modeling and Visualisation*, pages 271–280, 2003.

[22] A. Ngan and F. Durand. Statistical acquisition of texture appearance. *17th Eurographics Symposium on Rendering*, pages 31–40, August 2006.

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

[24] A. Ngan, F. Durand, and W. Matusik. Image-driven navigation of analytical BRDF models. *Proceedings of the Eurographics Symposium on Rendering*, pages 399–407, 2006.

[25] F.E. Nicodemus, J.C. Richmond, J.J. Hsia, I.W. Ginsburg, and T. Limperis. Geometrical considerations and nomenclature for reflectance. *NBS Monograph 160, National Bureau of Standards, U.S. Department of Commerce*, pages 1–52, 1977.

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

[27] Peiran Ren, Jiaping Wang, John Snyder, Xin Tong, and Baining Guo. Pocket reflectometry. *ACM Transactions on Graphics*, 30(4):45:1–45:10, 2011.

[28] S.M. Rusinkiewicz. A new change of variables for efficient BRDF representation. In *Rendering techniques' 98*, pages 11–22, 1998.

[29] M. Sattler, R. Sarlette, and R. Klein. Efficient and realistic visualization of cloth. In *Eurographics Symposium on Rendering 2003*, pages 167–178, 2003.

[30] Z. Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.

[31] Z. Zhengyou. Flexible camera calibration by viewing a plane from unknown orientations. In *In proceedings of ICCV*, pages 666–673, 1999.

[32] T. Zickler, R. Ramamoorthi, S. Enrique, and P. N. Belhumeur. Reflectance sharing: Predicting appearance from a sparse set of images of a known shape. *IEEE PAMI*, 28(8):1287–1302, 2006.