Vision, Modeling, and Visualization (2015)
D. Bommes, T. Ritschel, and T. Schultz (Eds.)

Fast multiplexed acquisition of high-dynamic-range material appearance

D. den Brok,[1] H. C. Steinhausen[1] and R. Klein[1]

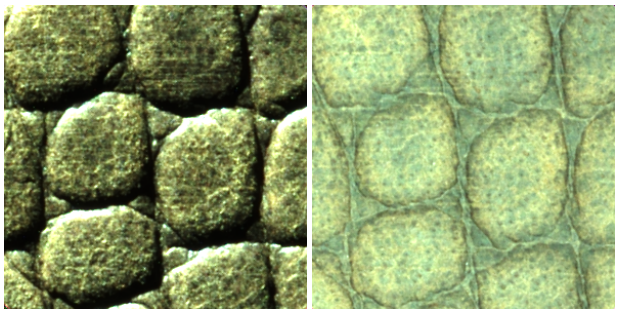

Figure 1: *Left: leather sample under single-light illumination. Right: same sample under multiplexed illumination.*

Abstract

There is tremendous demand for digital representations that allow for materials to be re-lit under as wide a range of illumination scenarios as possible. It is therefore desirable to capture the entire dynamic range of a material's appearance. This process can require excessive shutter times for many materials that reflect only small amounts of light for certain lighting and viewing directions, for instance in the presence of low albedo or self-shadowing. The problem is amplified in the case of image-based appearance models such as the bidirectional texture function (BTF), where possibly many thousands of images are required to accurately sample high-frequency details in the angular domain. We propose to capture material BTFs with their dynamic range compressed by multiplexed illumination. We demonstrate that the signal-dependent noise associated with demultiplexing can be mitigated by means of an existing database of low-noise material BTFs. Moreover, we investigate a method to quickly create a suitable database from scratch.

Categories and Subject Descriptors (according to ACM CCS): I.4.1 [Image Processing and Computer Vision]: Digitization and Image Capture—Reflectance

1. Introduction

In many applications it is desirable to be able to predict a given material's appearance under illumination that differs vastly from the controlled illumination that is typically employed when obtaining a material's digital representation. It is therefore important to capture as much of the dynamic range of material's reflectance as possible. The process of doing so, however, may require exposure series of many separate steps with shutter times that can range from only a few milliseconds up to minutes, for instance in the presence of low albedo or self-shadowing. In the case of image-based appearance models such as the *bidirectional texture function* (BTF), where the number of single images contributes strongly to the quality of reproductions in renderings, the resulting resource requirements can become excessive: Due to the many and possibly long shutter times, acquiring a single material's BTF may take from hours to even days. Moreover, a BTF typically comprises many thousands of (HDR) images, which during acquisition is multiplied by the number of (LDR) images per exposure series. This can easily lead to terabytes of data and increases the time spent on post-processing.

We propose to acquire material BTFs using *illumination multiplexing*: the material sample is illuminated not by each

DOI: 10.2312/vmv.20151270

single light source individually, but with a series of patterns of many light sources at once. That way, the dynamic range of a material's reflectance is compressed, because less shadows and highlights appear due to the many different simultaneous illumination angles, and the material is exposed to a larger amount of light. A BTF can thus be obtained with shorter exposure series and greatly reduced shutter times. Recovery of the desired images under single-light illumination amounts to solving an appropriate linear system, exploiting the linearity of the superposition of light. In the presence of signal-dependent noise, this *demultiplexing* process may yield intolerably noisy images, in particular if the material's dynamic range is wide.

We demonstrate that this effect can be mitigated by projecting the noisy BTF onto a linear subspace spanned by an existing database of traditionally measured material BTFs. In order to make this projection more robust, we heuristically identify possibly problematic images to be treated as missing values during projection.

As a suitable database may not always be available, we investigate a method of speeding up the process of creating one by using a combination of single-light and multiplexed illumination. We show that a database with material BTFs thus obtained performs not much worse in de-noising but is much quicker to create.

We evaluate our method on a number of real-world materials in a camera dome setup and show that it is possible to reduce acquisition times by about 75–95%, often down to our setup's physical limits, while maintaining a satisfactory reproduction quality.

2. Related work

2.1. Multiplexed illumination

Illumination multiplexing belongs to the wider field of plenoptic multiplexing. An extensive theoretical introduction to the topic has been provided by Hartwit and Sloane [HS79]. Essential for the present article is their proof of the optimality of Hadamard patterns and their binary siblings with respect to a number of measures. They also discuss the various noise sources in optical systems and their influence on demultiplexing on a high level.

Wenger et al. [WGT*05] propose multiplexing for the purpose of capturing time-varying light fields of human faces. However, they observed an intolerable amount of noise when using Hadamard patterns that they were unable to reduce to a tolerable level through simple filtering.

Schechner et al. [SNB07] provide a careful analysis of illumination multiplexing using digital photo cameras and various types of light sources, paying great attention to the possible kinds of noise. They arrive at a formula as a criterion whether multiplexing is beneficial given the setup's relevant intrinsic parameters.

Ratner et al. [RS07] demonstrate an optimization method which produces illumination patterns that take the noise characteristics of a given setup into account. They assume a one-dimensional affine noise model and a nearly diffuse scene, which does not lend itself well to BTF acquisition. Furthermore, they show that in the presence of overexposure it is preferable to reduce the number of light sources instead of shutter times.

Mitra et al. [MCV14] take this even further and compute illumination patterns using an optimization based on image priors. In contrast to previous methods, their method is able to handle large amounts of light, but it still relies on the assumption of a one-dimensional affine noise model. They show how to extend their method to low-resolution light fields, which is, however, computationally very expensive and therefore likely prohibitive in the case of BTFs.

As far as the authors are aware, den Brok et al. [dBSHK15] are the first to use multiplexed illumination in the context of BTF acquisition. They, too, use a database of previously acquired materials for the purpose of de-noising; they do, however, not provide a means to quickly bootstrap such a database, and we shall demonstrate that their de-noising scheme can be improved upon.

2.2. Sparse acquisition

We briefly review a number of articles on sparse acquisition, even though not directly related, because its purpose, too, is to reduce acquisition times.

Matusik et al. [MPBM03] demonstrate two sparse reconstruction methods for measured isotropic BRDFs: In the first method, they determine a set of basis wavelets they use to reconstruct previously unseen BRDFs with approximately 1.5 million samples from approximately half that amount of measurements. In the second, they reconstruct fully resolved measured BRDFs from 800 out of the original 1.5 million samples using their entire training BRDF database as a linear model. Drawbacks are slightly increased reconstruction errors and the required availability of the BRDF database. They do not investigate how well their methods generalize to more general appearance models.

Koudelka et al. [KMBK03] use single per-material linear models for apparent BRDFs for the purpose of BTF compression.

Peers et al. [PML*09] introduce *compressed sensing* [Don06] to the acquisition of *reflectance fields*, assuming both 2D *outgoing* (here: fixed viewing direction) and *incident light fields*. Their algorithm uses a hierarchical, multiresolution Haar wavelet basis that takes spatial coherence into account. It is unclear how to extend their approach to BTFs, where multiple viewing directions come into play, and the typically very limited number of light sources counteracts the advantage of compressed sensing. We expect shot noise to become a problem in this scenario as well.

Dong et al. [DWT*10] reconstruct a material's SVBRDF from a sparse measurement using a manifold constructed from analytical BRDFs fit to fully measured BRDFs of manually selected representative points on the material's surface. The algorithm is unlikely to scale to BTFs because of the typically much higher intrinsic dimensionality of the manifold of per-texel reflectance distribution functions. A generalization to previously unseen materials is not obvious, albeit conceivable.

Marwah et al. [MWBR13] use sparsity-based methods related to compressed sensing in order to sparsely acquire 4D light fields with an angular resolution of 5×5. They compute a dictionary of what they call *light field atoms* – 11×11 spatial light field patches which allow for a sparse representation of natural light fields. It is unlikely that such a dictionary exists in the case of BTFs, as their dimensionality is much higher.

Following Matusik et al. [MPBM03], den Brok et al. [dBSHK14] demonstrate that there exist linear models that lend themselves to sparse reconstruction of BTFs. In order to account for non-local effects, they propose to fit the models to small BTF patches, manually clustered by material class to constrain the dimensionality of that data. That way, they are able to reconstruct BTFs with the acquisition setup's full resolution from 6% of the total number of sample images.

Miandji et al. [MKU15] introduce a novel compressed sensing framework they demonstrate to work for 2D images, videos, and even 4D light fields. They use 2D patches from training data and show how to convert the problem of 2D sparse signal recovery to an equivalent problem in 1D. It would be interesting to investigate if their method could be lifted to 6D BTFs.

3. Background

3.1. Bidirectional texture functions

BTFs have been used in practice first by Dana et al. [DvGNK99]. Like spatially-varying BRDFs, they are 6-dimensional functions of the form

$$\mathcal{B}(\mathbf{x}, \omega_i, \omega_o),$$

where $\omega_{i,o} \in \mathbf{R}^2$ denote the directions of incoming and outgoing light, respectively, and $\mathbf{x} \in \mathbf{R}^2$ denotes the position on a parameterized surface V which, in the case of material BTFs, V is typically assumed flat; it does not need to coincide with the material's actual surface geometry. Light sources are usually assumed to be directional and have the same spectrum, which means that effects such as phosphorescene, fluorescence and subsurface scattering cannot be accounted for accurately.

Note that the values of the function $\mathcal{B}(\mathbf{x}, -)$ are not BRDFs in the strict sense: they typically do not adhere to Helmholtz reciprocity and conservation of energy and, in contrast to BRDFs, are therefore capable of capturing non-local effects such as interreflections and self-shadowing. When V does not coincide with the material's actual surface, they also describe parallax effects. For these reasons, the term *apparent BRDF* (ABRDF) has been established by Wong et al. [WHON97] for this kind of functions. The values of the function $\mathcal{B}(-, \omega_i, \omega_o)$, in contrast, are straightforward 2D textures corresponding to individual pairs of incoming and outgoing light directions.

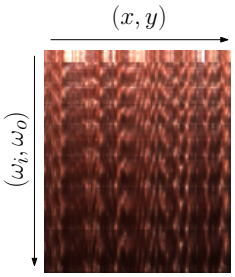

Figure 2: *Representation of a discretized BTF as a matrix.*

Discrete BTFs have a natural representation as a matrix $\mathbf{B} \in \mathbf{R}^{n \times m}$, where the columns represent the $w \times h$ discrete ABRDFs, each entry of which corresponds to a pair (ω_i, ω_o) of incoming and outgoing light direction, and the rows correspond to the n rectified 2D textures, where n is the number of distinct pairs of incoming and outgoing light directions (cf. Fig. 2). Measured BTFs are assumed to be arranged like this in the following. For a very detailed overview of BTF acquisition devices, we refer the reader to the recent survey by Schwartz et al. [SSW*14].

3.2. Illumination multiplexing

Multiplexed measurement in general is based on the observation that by using series of appropriate patterns, the amount of *signal-independent* noise in the demultiplexed measurements will actually be lower than without multiplexing. In the special case of illumination multiplexing, the resulting linear system to be solved is

$$\mathbf{M} \cdot \mathbf{I}_{\text{single}} = \mathbf{I}_{\text{multiplexed}}$$

where $\mathbf{M} \in \mathbf{Z}^{n \times n}$, $\mathbf{I}_{\text{single}} \in \mathbf{R}^{n \times (w \cdot h)}$ with each row a $w \times h$ image of the scene lit by an individual light source, and $\mathbf{I}_{\text{multiplexed}} \in \mathbf{R}^{n \times (w \cdot h)}$ where each row is a $w \times h$ image of the scene lit by an individual illumination pattern determined by the corresponding row of $\mathbf{M}$.

In setups where light sources only have two states, "on" and "off", $\mathbf{M}$ needs to be binary. It has been shown that a class of binary matrices called *S-matrices* is at least very close to minimizing the average mean square error, and that the associated signal-to-signal-independent-noise ratio is increased by a factor of $\sqrt{n}/2$ for large n [HS79]. It has, however, also been shown that the presence of signal-dependent

noise such as photon noise strongly counteracts this advantage [HS79, SNB07].

A straight-forward way to obtain S-matrices $\mathbf{M} \in \mathbf{N}^{n \times n}$ of a given order $n = 4p+3$ for n, $p \in \mathbf{N}$ prime, is the *quadratic residue construction* [HS79]: Let Q be the set of *quadratic residues* of $0, \ldots, \frac{n-1}{2}$ modulo n. Then

$$m_{1j} = \begin{cases} 1, & j \in Q \\ 0, & \text{else} \end{cases}$$

The remaining $n-1$ rows of $\mathbf{M}$ are obtained by cyclically shifting the preceding row to the left by one.

If the number of light sources is not prime of the required form, which will likely be the case in setups not constructed with multiplexing in mind, one can choose a suitable higher order and truncate the resulting S-matrix such that the number of columns equals the number of light sources. The resulting linear system can then be solved in the least-squares sense [HS79].

4. Proposed method

4.1. Thresholded de-noising

As we wish to achieve the maximum gain with respect to compression of dynamic range and reduction of shutter times, we chose to use the highest S-matrix order applicable to our acquisition setup and, by extension, the largest amount of demultiplexing noise. In order to mitigate that noise, we assume the availability of a linear model $\mathbf{U}$ derived from a database $\mathbf{D}$ of traditionally measured BTFs by means of a truncated SVD

$$\mathbf{D} \approx \mathbf{U}\Sigma\mathbf{V}^t$$

of some rank $k \ll n$, which has been shown to be an optimal rank-k approximation in the least-squares sense [EY36]. Such models are known to generalize to materials $\mathbf{B}$ that do not belong to the particular database [dBSHK14]; i.e.

$$\min_{\mathbf{C_B}} ||\mathbf{U}\mathbf{C_B} - \mathbf{B}|| < \varepsilon.$$

Den Brok et al. propose to de-noise by simply projecting a demultiplexed BTF, obtained as described in Sec. 3.2, onto the corresponding subspace via

$$\mathbf{B}_{\text{denoised}} = \mathbf{U} \cdot (\mathbf{U}^T \cdot \mathbf{B}_{\text{demultiplexed}})$$

We observed that demultiplexing produces a number of very noisy outlier textures that may have an undesirable impact on the projection. We therefore introduce one intermediate step to make this projection more robust: For the entire BTF, we compute per-texture variances and do not take the BTF's textures that correspond to some percentile of the variances into consideration during projection. Let $\widetilde{\mathbf{U}}$ and $\widetilde{\mathbf{B}}_{\text{demultiplexed}}$ be the linear model and demultiplexed BTF, respectively, with the rows corresponding to the selected textures removed. The equation thus becomes

$$\mathbf{B}_{\text{denoised}} = \mathbf{U} \cdot (\widetilde{\mathbf{U}}^{\dagger} \cdot \widetilde{\mathbf{B}}_{\text{demultiplexed}}),$$

where † denotes the Moore-Penrose pseudo-inverse. We tried various deciles with normalized and unnormalized data and found the 9th decile of the latter to perform best.

4.2. Database bootstrapping

A drawback of the the outlined method certainly is the assumed availability of a suitable database. We therefore propose a bootstrapping scheme that helps constructing a suitable database faster: Based on the observation that the amount of demultiplexing noise depends on both dynamic range and number of light sources, the idea is to divide the hemisphere of light sources into quadrants of approximately equal size, such that one quadrant contains the light sources directly opposite the cameras. As the direct reflections and Fresnel effect usually contribute most to the dynamic range, the images for this quadrant are best captured using single-light illumination. The images for the remaining three quadrants can be successively captured using appropriate S-matrix patterns. Unfortunately, this approach is somewhat tied to acquisition setup paradigms where such a division of the hemisphere is actually possible. Different means would need to be found for other paradigms.

5. Results

5.1. Experimental setup

We chose 12 materials from 3 classes – *cloth*, *leather*, and *wood* – each, of which we used the same four materials per class as den Brok et al. for performance evaluation. For the purpose of evaluating our database bootstrapping scheme we further selected four of the remaining leathers.

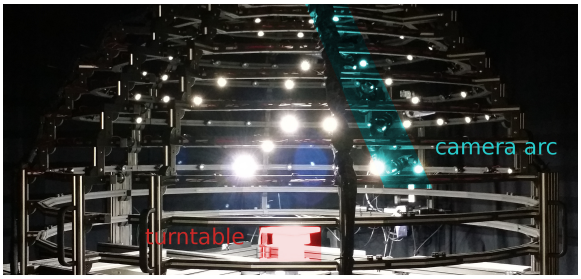

Figure 3: *Our measurement device illuminating a material sample with an S-matrix pattern.*

The measurement device we used in our experiments is a camera dome with 11 industrial-grade cameras with a maximum frame-rate of 8 Hz and 198 LED light sources, 10 of which are placed opposing the lower-most 10 cameras to produce direct reflections. The material sample is placed on a turntable that is rotated during acquisition to increase the setup's angular resolution. For an extensive description of the setup, cf. Schwartz et al. [SSWK13]. The quadratic residue construction requires the S-matrix order to be prime

of the form $4p+3$, p prime. The smallest S-matrix order of that form greater or equal the number of LEDs in our setup is $n = 199$, which we used in our experiments.

In order to speed up the acquisition of the ground truth data, the single-light measurements were obtained using a camera gain of 10 dB, while the multiplexed measurements were obtained with a camera gain of 0 dB. Additionally, we measured 4 small patches (approximately 1 cm $\times$ 1 cm) of different materials of the same class at once. We argue that, even in practive, this is a reasonable trade-off to be made: If required, larger BTFs for each single sample can be produced quickly by measuring them once more using the proposed method, or by extrapolation as described e.g. by Steinhausen et al. [SdBHK14].

The materials are placed on the turntable, which during measurement is rotated 12 times by 30°. For each turntable position, the cameras take pictures of the material lit by each of the LEDs separately using several exposure times determined manually before the actual acquisition. The test materials were additionally measured using S-matrix patterns of order $n = 199$ before proceeding to the next turntable position. In the case of leather, the four leathers selected for evaluation of the database bootstrapping scheme were additionally measured using S-matrix patterns of orders $n_{\{1,2,3\}} = 47$ for each of the three of the illumination hemisphere's quadrants not sampled using single-light illumination.

After measurement, we combine the low dynamic range (LDR) raw images into high dynamic range (HDR) images and subsequently demosaic and rectify them. We use LED calibration data in order to account for variations in the LEDs' spectra. For further details on the various post-processing steps, cf. Schwartz et al. [SSW*14].

Finally, the resulting images are cropped to minimize influence from the materials' surroundings and arranged and stored as matrices, which in the case of multiplexed measurements are then demultiplexed. For rendering we resample the BTFs in the angular domain such that the light and view hemispheres are the same for each texel.

5.2. De-noising

We compute the linear models intended for de-noising seperately from the per-color channel ABRDFs of each of the three material classes, leaving out one test material at a time. We use the log(Y) U/Y V/Y color space, which lends itself well to least-squares fitting [MPBM03, dBSHK14].

For our experiments we used 768 basis vectors for log(Y) and 128 for both U/Y and V/Y channels, which is sufficient for the material classes under consideration and allows for better comparability with the method proposed by den Brok et al. [dBSHK15].

Tab. 1 shows relative errors

$$\varepsilon = \frac{\|\mathbf{B}_{\text{reconstructed}} - \mathbf{B}_{\text{reference}}\|_{\text{F}}}{\|\mathbf{B}_{\text{reference}}\|_{\text{F}}}$$

for BTFs reconstructed in the denoted ways. Errors are computed on the log(Y) channel to approximate human perception. We found the general tendencies to be the same for the other color channels.

In the case of wood, the relative errors are relatively small, and indeed there are no obvious differences between the renderings of ground truth, demultiplexed and denoised BTFs, even though the relative error is reduced significantly by denoising (cf. Fig. 4, 1st row). Curiously, this is the only material class where the proposed method performs worse than the state of the art in purely numerical terms. We have yet to investigate what the reason might be. Possibly the error introduced by sparse reconstruction exceeds and hence does not counteract the low demultiplexing error. A higher threshold might mitigate this problem.

In contrast, Cloth 1 and Leather 2 exhibit annoying artifacts for grazing viewing angles, as seen at the cylinder's borders. The BTF de-noised using den Brok et al.'s method looks much more plausible than the demultiplexed BTF, but it still exhibits minor artefacts which are close to unnoticable in the BTF with our method (cf. Fig. 4, 2nd and 3rd row).

Demultiplexed Leather 4, however, exhibits so much noise (yielding a relative error of over 27%) that both denoising strategies break down (cf. Fig. 4, 4th row), even though the relative error is more than halved by den Brok et al's method and reduced even further by ours. Note that the intensity at grazing angles in the result produced using den Brok et al.'s method is much too low. Our method does not suffer from this problem; it fails, however, at reducing the demultiplexing noise to an acceptable level for this particular material.

As can be seen in Tab. 2, the larger amount of light reaching the material samples on each particular image makes it possible to reduce total acquisition times by about 75–95%, even though the ground truth data was obtained with a much higher camera gain. In the case of Leather 1–4 and Wood 1–4, our setup's minimum acquitision time is reached. For all materials, 4 different shutter times were necessary under single-light illumination, whereas with multiplexed illumination, only a single exposure step was needed in the case of leather and wood, and two in the case of cloth, the latter because the cloths exhibited significantly different albedos. As a result, storage requirements for the raw measured data are reduced by 50–75% from originally 0.5–1.5 TB.

5.3. Database bootstrapping

We evaluated the performance of our database bootstrapping scheme in two steps. First, we compared the projection errors when projecting onto linear bases computed from the ground-truth and the bootstrapping measurements, respectively (cf. Tab. 3). We found projection errors to be only slightly bigger for the bootstrapping measurements. It is thus also to be expected that de-noising performance will

Class	$\varepsilon_{\text{demult}}$					$\varepsilon_{\text{denoised}}$					$\varepsilon_{\text{threshold}}$				
	1	2	3	4	∅	1	2	3	4	∅	1	2	3	4	∅
Cloth	16.2	12.2	10.5	11.3	12.6	7.5	6.2	5.8	5.6	6.3	5.3	4.5	4.3	4.7	4.7
Leather	17.7	20.1	23.0	27.4	22.1	8.8	8.3	9.4	10.5	9.3	6.3	7.5	7.8	7.6	7.3
Wood	7.3	7.6	9.3	7.2	7.9	3.6	3.7	4.2	3.7	3.8	4.9	5.1	4.2	4.3	4.6

Table 1: *Comparison of log(Y) relative errors [%], demultiplexed and denoised.*

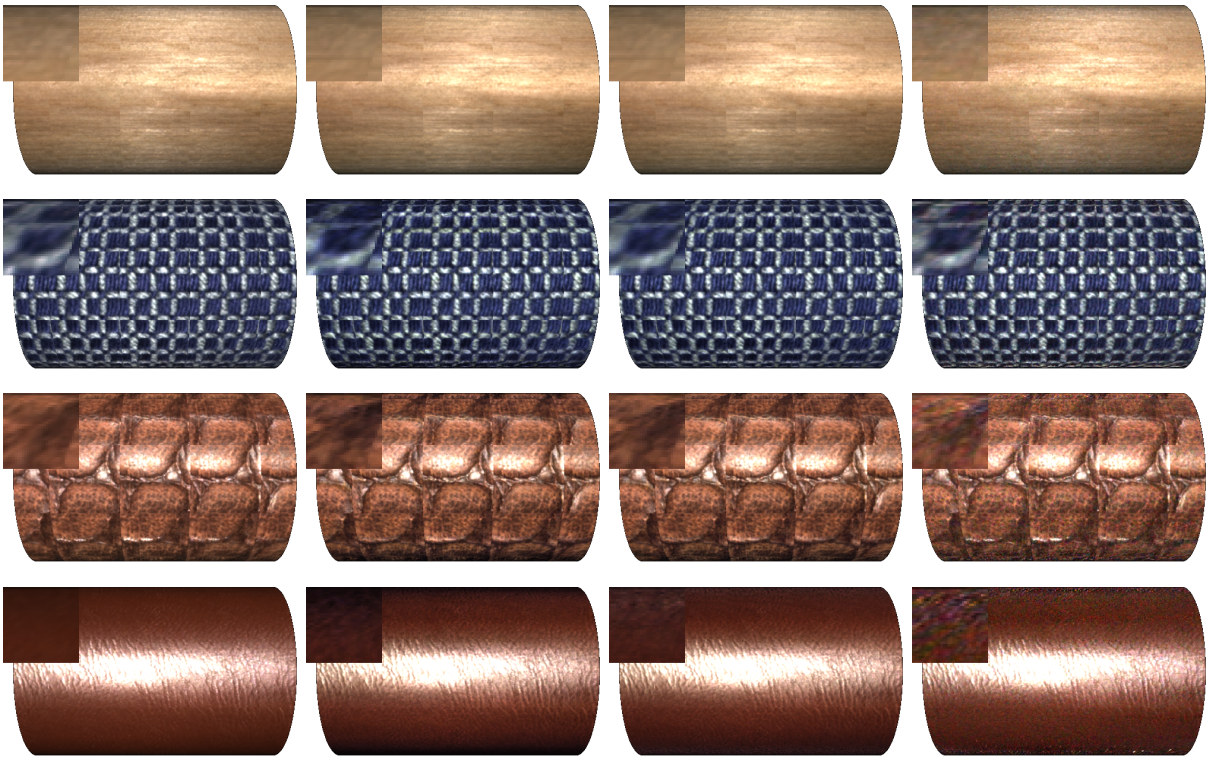

(a) *Reference.* **(b)** *[dBSHK15]* **(c)** *Proposed method.* **(d)** *Demultiplexed.*

Figure 4: *Renderings of Wood 1, Cloth 1, Leather 2, and Leather 4.*

Class	shutter times [ms]		acquisition times [h]		rel. Δ
	single	multiplexed	single	multiplexed	
Cloth 1 – 4	150, 608.2, 2466.2, 10000	15, 45	10.8	1.8	−83%
Leather 1 – 4	10, 144.2, 2080.1, 30000	30	23.6	1.2	−95%
Wood 1 – 4	10, 60.4, 364.4, 2200	20	4.4	1.2	−75%

Table 2: *Comparison of shutter [ms] and acquisition times [h], single light vs. multiplexed.*

be slightly worse, which our experiments confirmed. However, both numerically and perceptually, we found it still to be better overall than that of den Brok et al.'s method, and the resulting BTFs are perceptually close to the those produced with the traditionally obtained basis (cf. Tab. 3 & Fig. 5). The acquisition itself took approximately 9 hours, as opposed to almost 24 hours for the ground-truth.

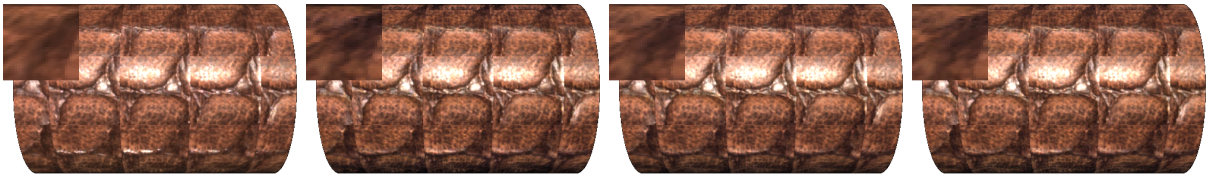

(a) *Reference.* **(b)** *[dBSHK15]* **(c)** *Bootstrapped basis.* **(d)** *Original basis.*

Figure 5: *Bootstrapping: renderings of Leather 4.*

	1	2	3	4	∅
$\varepsilon_{\text{proj}}$ ground-truth	8.4	7.6	8.6	9.3	8.5
$\varepsilon_{\text{proj}}$ bootstrapped	8.7	8.0	9.0	9.5	8.8
$\varepsilon_{\text{denoise}}$ ground-truth	6.3	7.5	7.8	7.3	7.3
$\varepsilon_{\text{denoise}}$ bootstrapped	7.0	8.3	8.8	8.5	8.2

Table 3: *Comparison of projection and reconstruction errors for Leather 1–4 with a bootstrapped database.*

6. Conclusion & Future work

We demonstrated the feasibility of illumination multiplexing in the context of BTFs, supported by using linear models derived from an existing database of material BTFs as a prior for a de-noising method that perceptually and, most of the time, numerically outperforms the state of the art. We found that using illumination multiplexing enables both dramatically reduced dynamic ranges and shutter times. As a result, storage requirements for raw measurement data could be reduced by up to 75%, and total acquisition times by up to 95%, even reaching the limits of our acquisition setup.

Moreover, we presented a "bootstrapping" method that allows for faster creation of a database suitable for the purpose of de-noising by using a hybrid approach where the quadrant of the illumination hemisphere likely to cause noise is sampled using single-light illumination, and the remaining quadrants are separately sampled using multiplexed illumination. That way, acquisition time could be reduced by about 63%. We found linear models derived from a database such obtained to perform not much worse than linear models derived from our ground-truth database.

It seems worthwhile to determine a proper noise model for our camera dome setup to use with the method proposed by Mitra et al. [MCV14] (cf. Sec. 2.1 for more details), which might help decrease the amount of demultiplexing noise even further.

Acknowledgements

This work was developed in the X-Rite Graduate School on Digital Material Appearance at the University of Bonn.

References

[dBSHK14] DEN BROK D., STEINHAUSEN H. C., HULLIN M. B., KLEIN R.: Patch-based sparse reconstruction of material BTFs. *Journal of WSCG 22*, 2 (June 2014), 83–90. 3, 4, 5

[dBSHK15] DEN BROK D., STEINHAUSEN C., HULLIN M. B., KLEIN R.: Multiplexed acquisition of bidirectional texture functions for materials. In *Measuring, Modeling, and Reproducing Material Appearance II (SPIE 9398)* (Feb. 2015), vol. 9398. 2, 5, 6, 7

[Don06] DONOHO D. L.: Compressed sensing. *IEEE Trans. Inform. Theory 52* (2006), 1289–1306. 2

[DvGNK99] DANA K. J., VAN GINNEKEN B., NAYAR S. K., KOENDERINK J. J.: Reflectance and texture of real-world surfaces. *ACM Trans. Graph. 18*, 1 (Jan. 1999), 1–34. 3

[DWT*10] DONG Y., WANG J., TONG X., SNYDER J., LAN Y., BEN-EZRA M., GUO B.: Manifold bootstrapping for SVBRDF capture. *ACM Trans. Graph. 29*, 4 (July 2010), 98:1–98:10. 3

[EY36] ECKART C., YOUNG G.: The approximation of one matrix by another of lower rank. *Psychometrika 1*, 3 (1936), 211–218. 4

[HS79] HARWIT M., SLOANE N.: *Hadamard transform optics*. Academic Press, 1979. 2, 3, 4

[KMBK03] KOUDELKA M. L., MAGDA S., BELHUMEUR P. N., KRIEGMAN D. J.: Acquisition, compression, and synthesis of bidirectional texture functions. In *ICCV Workshop on Texture Analysis and Synthesis* (2003). 2

[MCV14] MITRA K., COSSAIRT O., VEERARAGHAVAN A.: Can we beat hadamard multiplexing? data driven design and analysis for computational imaging systems. In *Computational Photography (ICCP), 2014 IEEE International Conference on* (May 2014), pp. 1–9. 2, 7

[MKU15] MIANDJI E., KRONANDER J., UNGER J.: Compressive image reconstruction in reduced union of subspaces. In *Eurographics 2015* (May 2015). 3

[MPBM03] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: A data-driven reflectance model. *ACM Trans. Graph. 22*, 3 (July 2003), 759–769. 2, 3, 5

[MWBR13] MARWAH K., WETZSTEIN G., BANDO Y., RASKAR R.: Compressive light field photography using overcomplete dictionaries and optimized projections. *ACM Trans. Graph. 32*, 4 (July 2013), 46:1–46:12. 3

[PML*09] PEERS P., MAHAJAN D. K., LAMOND B., GHOSH A., MATUSIK W., RAMAMOORTHI R., DEBEVEC P.: Compressive light transport sensing. *ACM Trans. Graph. 28*, 1 (Feb. 2009), 3:1–3:18. 2

[RS07] RATNER N., SCHECHNER Y.: Illumination multiplexing within fundamental limits. In *Computer Vision and Pattern*

Recognition, 2007. CVPR '07. IEEE Conference on (June 2007), pp. 1–8. 2

[SdBHK14] STEINHAUSEN H. C., DEN BROK D., HULLIN M. B., KLEIN R.: Acquiring bidirectional texture functions for large-scale material samples. *Journal of WSCG 22*, 2 (June 2014), 73–82. 5

[SNB07] SCHECHNER Y. Y., NAYAR S. K., BELHUMEUR P. N.: Multiplexing for optimal lighting. *IEEE Trans. Pattern Anal. Mach. Intell. 29*, 8 (Aug. 2007), 1339–1354. 2, 4

[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 (Apr. 2014). 3, 5

[SSWK13] SCHWARTZ C., SARLETTE R., WEINMANN M., KLEIN R.: DOME II: A parallelized BTF acquisition system. In *Eurographics Workshop on Material Appearance Modeling: Issues and Acquisition* (June 2013), Eurographics Association, pp. 25–31. 4

[WGT*05] WENGER A., GARDNER A., TCHOU C., UNGER J., HAWKINS T., DEBEVEC P.: Performance relighting and reflectance transformation with time-multiplexed illumination. *ACM Trans. Graph. 24*, 3 (July 2005), 756–764. 2

[WHON97] WONG T.-T., HENG P.-A., OR S.-H., NG W.-Y.: Image-based rendering with controllable illumination. In *Proceedings of the Eurographics Workshop on Rendering Techniques '97* (London, UK, UK, 1997), Springer-Verlag, pp. 13–22. 3