Workshop on Material Appearance Modeling (2015)
H. Rushmeier and R. Klein (Editors)

Linear models for material BTFs and possible applications

D. den Brok[1], M. Weinmann[1] and R. Klein[1]

[1]University of Bonn, Germany

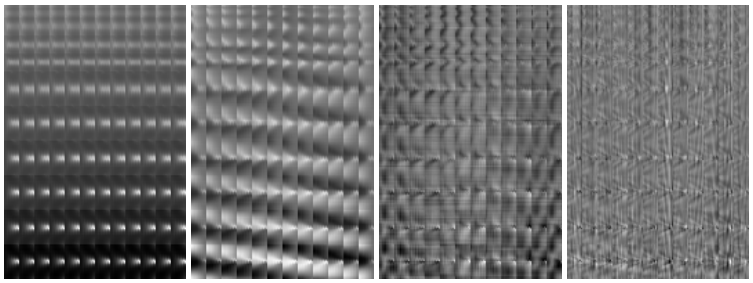

Figure 1: Selection of $\log(Y)$-channel eigen-ABRDFs computed on the BTFs of 11 different leathers.

Abstract

Due to the richness of real-world materials, arguably one of the biggest challenges in rendering is to come up with models that describe their appearance well. The image-based bidirectional texture function *(BTF) is known to be able to model many effects that are hard or impossible to reproduce with analytical reflectance models. This advantage comes at the price of demanding storage and acquisition requirements. In previous work, we have demonstrated that these requirements can be lifted to some extent by means of data-driven linear models. We give a more in-depth overview on our research on such models and summarize the applications we investigated so far, followed by an outlook on what might yet be achievable.*

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

1. Introduction

There is tremendous demand in many branches of science and industry for realistic reproduction and prediction of material appearance, be it in high-quality offline rendering, real-time or even interactive applications. In the latter two scenarios in particular, analytical reflectance models such as the spatially-varying BRDF often fail to accurately model the appearance of materials with noticeable non-local effects such as self-shadowing, interreflections, and subsurface scattering. Due to its image-based nature, the *bidirectional texture function* (BTF), first put to practical use by Dana et al. [DvGNK99], is capable of reproducing many such effects, even in real-time. In order to adequately capture high-frequency effects, however, a huge number of images is required, which makes dealing with BTFs a challenging task: for instance, methods that readily apply to 2D textures often become intractable in the case of BTFs. As a means to mititage this problem we investigated using linear models to reduce the BTFs' dimensionality. We found that a reduction of up to 94% can be achieved that way, which we believe to be sufficient to bring at least some previously intractable important problems when dealing with BTFs into reach. We further present some previously published applications for linear models in BTF acquisition and discuss a number of possible further applications.

DOI: 10.2312/mam.20151198

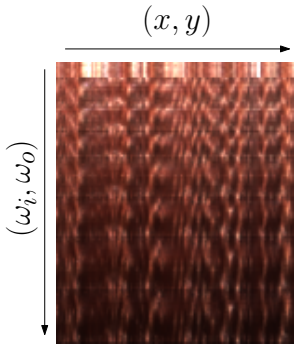

Figure 2: Representation of a discretized leather BTF as a matrix.

2. Bidirectional texture functions

A *bidirectional texture function* (BTF) is a function of the form

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

where $\omega_{i,o} \in \mathbf{R}^2$ denote the incoming and outgoing light directions, respectively, and $\mathbf{x} \in \mathbf{R}^2$ denotes the position on a parameterized surface V. In the case of material BTFs, V is typically flat; it does not need to coincide with the material's actual surface geometry. It is generally assumed that light sources are directional and have the same spectrum. Note that the function $\mathcal{B}(\mathbf{x}, -)$ is not BRDF-valued, as its values do not have some properties, such as conservation of energy, usually required of BRDFs. For this reason the term *apparent BRDF* (ABRDF) has been suggested by Wong et al. [WHON97] for this kind of functions. The values of the function $\mathcal{B}(-, \omega_i, \omega_o)$, however, are indeed 2D textures corresponding to specific pairs of incoming and outgoing light directions. During measurement, a finite discretization of the measured material's BTF $\mathcal{B}$ is obtained. After rectification of the acquired images, the discrete BTF has a natural representation as a matrix $\mathbf{B} \in \mathbf{R}^{n \times m}$ with the columns representing the m discrete ABRDFs, each entry corresponding to some pair (ω_i, ω_o) of incoming and outgoing angle, and the rows representing the n rectified textures (cf. Figure 2).

3. Linear models for BTFs

3.1. Rationale

There exists a large repertoire of dimensionality reduction methods. One of the most simple such techniques is the singular value decomposition (SVD), which is known by the Eckart-Young theorem [EY36] to provide the best (in the least-squares sense) rank-k approximation $\mathbf{B} \approx \mathbf{UX}$ of a given matrix $\mathbf{B}$ for any value of k. Due to its simplicity, it cannot be expected to yield a dimensionality reduction on par with more sophisticated techniques. However, it is extremely easy to handle: its truncated version can be approximated efficiently in several ways, for instance using Roweis' EM-PCA [Row98], and most operations one would want to perform on a BTF can be performed just the same on the SVD coefficients. Moreover, if one such linear model exists that generalizes to many or even all BTFs, then their respective coefficients are defined over the same basis. That way, they can still be meaningfully combined or compared, for instance in distance computations in clustering or nearest-neighbor search, which become much faster due to the reduced dimensionality.

3.2. Experimental setup

We used our "Dome II" acquisition device, described in great detail by Schwartz et al. [SSW*14], to obtain BTFs for 8 pieces of (processed) wood, 8 pieces of leather, and 8 fabrics. The "Dome II" has 198 LEDs and 11 cameras, and the materials were evenly rotated about the z-axis 12 times to increase the sampling density, for a total extrinsic ABRDF dimensionality of $198 \times 11 \times 11 = 26136$ per color channel. There are several possible ways to reduce the data's intrinsic dimensionality even without using typical dimensionality reduction techniques; one could, for instance, simultaneously align all ABRDFs for maximum correlation as demonstrated by Müller et al. [MSK06]. Most such methods, however, alter the data to some degree; Müller et al.'s technique e.g. involves multilinear resampling which may change the shape and intensity of highlights. For the following analysis we therefore simply used rectified and radiometrically corrected but otherwise unprocessed BTFs.

We randomly selected 4 materials per class for testing and performed the same experiments successively four times, each time with a different material from each class. Numerical results were subsequently averaged to obtain the final results.

We first computed rank-512 approximations of all the training and test BTFs along with the reconstruction errors from rank-128 approximations of the test BTFs for future comparison. From the bases thus obtained, we train common bases for each material class, and finally for the entire database. All computations were performed on the $\log(Y)$-channel, which has by far the highest dimensionality among the three channels. We applied the logarithm to the Y-channel in order to reduce its dynamic range and, by extension, the method's sensitivity to outliers such as highlights.

3.3. Results

Figure 3 shows the average projection errors for both the per-material class basis and the global basis for the three material classes under consideration. For comparison, the average relative error when reconstructing from a truncated SVD of rank 128, which is a typical configuration for real-time applications, is shown as well.

Interestingly, for two of three classes – *wood* and *fabric* –

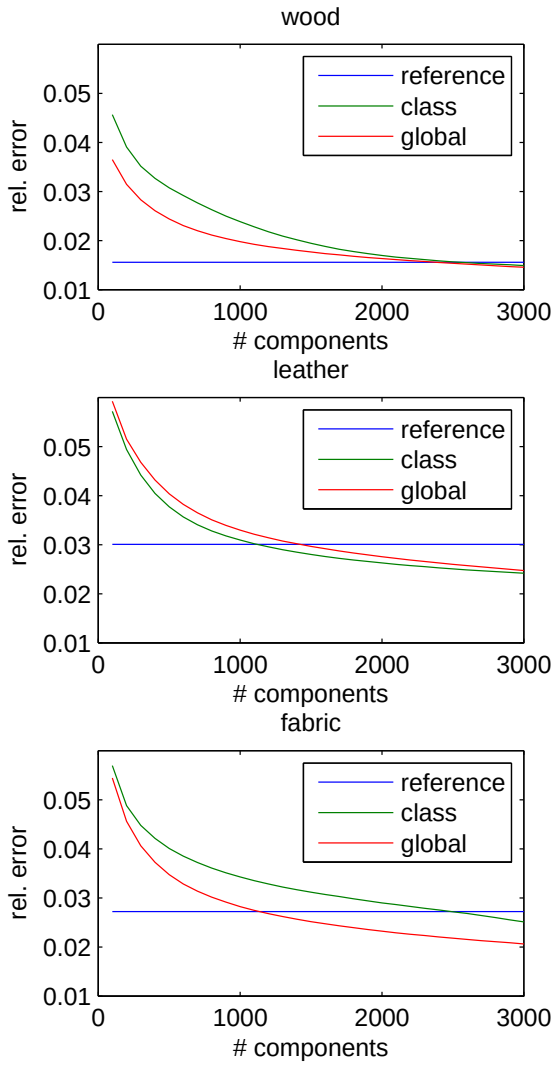

Figure 3: Average relative projection errors for various bases and material classes.

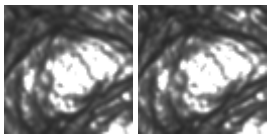

Figure 4: Reconstructed Y-channel textures of a leather. Left: from rank-128 SVD. Right: from projection to global ABRDF basis with 768 entries.

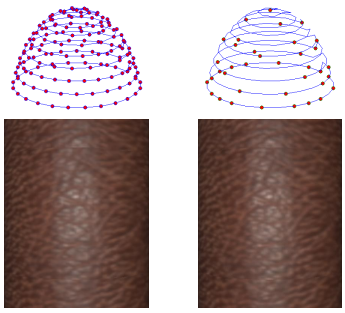

Figure 5: Demonstration of sparse BTF acquisition. Left: Excerpt from reference rendering along with the full sampling. Right: Sparse sampling used to produce the excerpt from a rendering of a sparse reconstruction shown below.

the global ABRDF basis performs slightly better in numerical terms than the class bases, and only very slightly worse in the case of *leather*. The number of basis ABRDFs required to achieve a reconstruction error comparable to what is typically used in real-time rendering varies from approximately 1500 to 2200, which amounts to a dimensionality reduction of up to 94%. By visual inspection we found 1000 basis ABRDFs to be typically sufficient as well; see Figure 4 for an example.

Note that while this number may still seem big, and indeed would make real-time rendering very demanding if not impossible in terms of hardware requirements, a factorized representation using a smaller basis tailored for the material of concern can be easily produced from the common basis and the corresponding coefficients:

Let $\mathbf{B} \approx \mathbf{UX}$ be a representation of a BTF as a matrix $\mathbf{B}$, factorized over an orthonormal basis $\mathbf{U}$. Let $\mathbf{X} = \mathbf{U_X}\Sigma_\mathbf{X}\mathbf{V_X}^T$ be the (exact) singular value decomposition of the coefficient matrix $\mathbf{X}$. Then $\mathbf{B} \approx (\mathbf{UU_X})\Sigma_\mathbf{X}\mathbf{V_X}^T$ is a truncated SVD of the projection of $\mathbf{B}$ onto the span of $\mathbf{U}$, and as the rank of $\mathbf{B}$ is assumed to be much smaller than that of $\mathbf{U}$, it can be truncated further with negligible impact. The exact SVD of the coefficient matrix $\mathbf{X}$ can be obtained quite quickly, because its first dimension is always bounded by the rank of $\mathbf{U}$, and the SVD can be computed on its transpose if necessary.

4. Applications

4.1. Sparse acquisition

Our first attempt at exploiting linear bases for BTFs [dBSHK14] drew much inspiration from Matusik et al.'s work on isotropic BRDFs [MPBM03]. The general idea is to obtain a small subset of the full number of individual images only, and subsequently filling the missing values by means of the model:

$$\mathbf{B} \approx \mathbf{U}(\mathbf{MU})^\dagger \widetilde{\mathbf{B}}, \tag{1}$$

Figure 6: Demonstration of denoising of demultiplexed BTFs. Texture taken from a leather BTF. Left: demultiplexed. Center: ground truth. Right: denoised. Note how even fine structural detail, completely covered by noise in the demultiplexed image, is recovered.

where $\mathbf{B}$ is the full BTF to be reconstructed from the sparsely measured BTF $\widetilde{\mathbf{B}}$, $\mathbf{U}$ the linear basis, $\mathbf{M}$ a binary measurement matrix determining which lighting and viewing angles to sample, and $(\mathbf{MU})^{\dagger}$ denotes the Moore-Penrose pseudo-inverse of $\mathbf{MU}$. Matusik et al. provide an algorithm to obtain $\mathbf{M}$ that optimizes the condition $\kappa(\mathbf{M})$.

This approach potentially drastically reduces both acquisition time and storage requirements. Note in particular that, as explained in Section 3, the uncompressed full BTF never needs to be stored anywhere. It turned out that without appropriate regularization, Matusik et al.'s method provides no real benefit when applied to ABRDFs because the much higher intrinsic dimensionality of the space of ABRDFs causes the number of required sample images to be inconveniently big. We solved this problem by using models fit to clusters of 3×3 BTF patches which we call *apparent BTFs*. Preliminary unpublished results suggest that a regularization scheme exists that could lead to similar performance with models fit to ABRDFs.

The method was evaluated using the material database presented by Weinmann et al. [WGK14]. We found it to perform well for many simple (e.g. processed wood) to moderately complex materials (e.g. leather) with as little as 6 % of the original number of sample images; cf. Figure 5 for an example. It breaks down, however, for very specular materials (e.g. wall tiles) and materials with significant surface structure (e.g. carpet).

4.2. Multiplexed acquisition

As another approach to faster BTF acquisition, we investigated ways to reduce the exposure times necessary to capture the full dynamic range of materials' reflectance [dBSHK15]. An obvious solution is to increase the amount of light that hits the material sample. In a fixed acquisition setup with multiple light sources this increase can be gained by using illumination patterns composed of multiple light sources. Due to the linearity of the superposition of light the desired images under directional illumination can be obtained by solving a linear system

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

where $\mathbf{M}$ is a binary measurement matrix composed of illumination patterns, $\mathbf{I}_{\text{single}}$ a matrix with each row a $w \times h$ image of the scene lit by an individual light source, and $\mathbf{I}_{\text{multiplexed}}$ a matrix where each row is a $w \times h$ image of the scene lit by an individual illumination pattern. Suitable binary illumination patterns that are optimal in a certain sense, known as *S-matrices*, can be derived from the theory of Hadamard codes [HS79].

Under the assumption that demultiplexing noise does not lie in the subspace spanned by these bases, denoising can now be performed by simply projecting a demultiplexed BTF onto the corresponding basis:

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

We evaluated the method using the BTFs analyzed in Section 3. It performs well for moderate amounts of noise, even if they would cause intolerable artifacts in renderings; cf. Figure 6 for an example. Unfortunately, we encountered a number of leathers and fabrics that exhibit very large amounts of demultiplexing noise impossible to mitigate by means of the proposed method. In most cases, however, adequate results could be achieved in 75% to 95% less time than using traditional single light illumination, even though a higher camera gain was used in the traditional scenario.

4.3. Outlook

It is fair to say that when it comes to BTFs there are still more unsolved problems than there are solved ones. Among the most important ones are problems like meaningful editing, super-resolution, classification, and acquisition of BTFs of very large material samples. We firmly believe that linear models make at least some of these problems tractable and shall investigate this in the future.

5. Conclusion & Future work

We presented a short analysis of linear models for material BTFs, along with some applications from previously published research. It turned out that linear models provide a simple means to significantly reduce the dimensionality of BTFs while retaining many desirable properties such as ease of use, comparability, etc. We found that in terms of numerical error, a model trained on a wide range of materials performs better than models trained on manually selected material classes.

A lot of questions are still open: How does model performance change with the number of materials used for training? How can we achieve scale invariance, and do we even have to? Does automatic clustering improve the performance of class-wise bases? We hope to be able to answer some of these and other questions in future work.

6. Acknowledgements

The research leading to these results was partially funded by X-Rite Inc. via the Graduate School on Digital Material Appearance at Bonn University, and the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 323567 (Harvest4D); 2013-2016.

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

[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. 4

[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. URL: http://doi.acm.org/10.1145/300776.300778, doi:10.1145/300776.300778. 1

[EY36] ECKART C., YOUNG G.: The approximation of one matrix by another of lower rank. *Psychometrika 1*, 3 (1936), 211–218. URL: http://dx.doi.org/10.1007/BF02288367, doi:10.1007/BF02288367. 2

[HS79] HARWIT M., SLOANE N.: *Hadamard transform optics*. Academic Press, 1979. URL: http://books.google.de/books?id=UAy1AAAAIAAJ. 4

[MPBM03] MATUSIK W., PFISTER H., BRAND M., MCMILLAN L.: Efficient isotropic BRDF measurement. In *Proceedings of the 14th Eurographics Workshop on Rendering* (Aire-la-Ville, Switzerland, Switzerland, 2003), EGRW '03, Eurographics Association, pp. 241–247. URL: http://dl.acm.org/citation.cfm?id=882404.882439. 3

[MSK06] MÜLLER G., SARLETTE R., KLEIN R.: Data-driven local coordinate systems for image-based rendering. *Computer Graphics Forum 25*, 3 (Sept. 2006). 2

[Row98] ROWEIS S.: EM algorithms for PCA and SPCA. In *Proceedings of the 1997 Conference on Advances in Neural Information Processing Systems 10* (Cambridge, MA, USA, 1998), NIPS '97, MIT Press, pp. 626–632. URL: http://dl.acm.org/citation.cfm?id=302528.302762. 2

[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). URL: http://www.mdpi.com/1424-8220/14/5/7753, doi:10.3390/s140507753. 2

[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. 4

[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. URL: http://dl.acm.org/citation.cfm?id=647651.731971. 2