

Full BRDF Reconstruction Using CNNs from Partial Photometric Stereo-Light Field Data

Doris Antensteiner and Svorad Štolc
Austrian Institute of Technology
2444 Seibersdorf, Austria

{doris.antensteiner, svorad.stolc}@ait.ac.at

Abstract

The acquisition of partial BRDF measurements using light field cameras and several illumination directions raises critical questions regarding the accuracy of inferences based on that data. Therefore, we attempt to verify the quality of the reconstruction of a full BRDF using partial input data. A dataset that provides a densely sampled BRDF was used, both in viewing and illumination directions. We show the reconstruction of dense BRDFs when the viewing angles are limited to top central regions, while the illumination angles are not reduced and are positioned in the shape of a half sphere around the material object, these properties are characteristic of data provided by plenoptic cameras paired with a photometric light dome. The partial reconstruction of the dense BRDF out of data is achieved by utilizing convolutional neural networks. We obtain a competitive full reconstruction when up to 2/3 of the BRDF is unknown.

1. Introduction

Material appearances of structured materials such as metal, wood, or plastic can be uniquely characterized by the *bidirectional reflectance distribution function* (BRDF), as described by Nicodemus *et al.* [1]. The BRDF consists of four independent variables $(\theta_i, \phi_i; \theta_r, \phi_r)$, which denote the angles of the incident and reflected radiant flux. It maps the irradiance incident from a specific direction to its influence towards the reflected radiance of another direction. The *bidirectional texture distribution function* (BTF) was later introduced by Dana *et al.* [2], which describes a texture image parametrized by the viewing and illumination angle and therefore a spatially varying BRDF.

The acquisition of BRDFs can be tedious. Several databases were constructed for research and academic use, such as the CURET BRDF database [3] which provides sparse samples with 200 measurements. The MERL BRDF




Figure 1: Input data stack: (a) Ground truth BRDF, (b) reconstruction with input data dimensions $48 \times 48 \times u$ without the transposed channels with an SSIM=0.73 of the estimation (for the shown BRDF sample), (c) reconstruction with additional transposed input data added to the original input channels, which leads to input data dimensions $48 \times 48 \times 2u$ with an SSIM=0.88.

Database [4] was obtained with a light source moving in a circle at the same level as the camera, in respect to the surface normal. Since both illumination and camera position are restricted to a defined circle, it neither allows the simulation of viewing angles similar to plenoptic cameras nor the positioning of the illuminations in the shape of a half sphere. The BRDFs from BTFs Dataset from UTIA [5] provides densely sampled high precision BRDF measurements with independent camera and illumination positions, which allow both capturing anisotropic material behaviour as well as reducing the dimensionality for both, illumination and camera separately.

Using a very sparse BRDF for the task of image classification was discussed by Wang *et al.* [6], where a fixed traditional camera was placed at the top center of a light dome in order to acquire images. Each pixel observation is fit to a 2nd order hemispherical harmonic model. By Wang *et al.* [7] convolutional neural networks (CNNs) [8] were used to classify materials from plenoptic acquisitions without illumination variations. It was shown that material recognition can be improved through light field data, compared to

single image acquisitions. A compressive sensing approach was implemented by Zupanic and Soler [9], in order to reconstruct BRDFs from a single image with known normals and illumination directions for isotropic and spatially constant materials. Reconstructing the BRDF from a limited set of samples was addressed previously by Nielsen *et al.* [10], where the MERL database was used. The dimensionality was reduced based on a principal component analysis in order to retrieve the most influential regions.

This paper demonstrates the ability to reconstruct dense BRDFs for restricted observation angles on the scene, as occurring in plenoptic cameras, while the illumination sources can be triggered from a wide range of positions around the object. The goal is to reconstruct dense BRDFs out of sparse data, with restricted viewing angles concentrated around the top down view.

This paper is organized as follows. In Sec. 2, we discuss the connection of light fields with BRDFs. The used BRDF dataset is described in Sec. 3. Our chosen regression CNN structure is introduced and discussed in Sec. 4. What follows clearly stated is the behaviour of the reconstruction from 1/6 to 5/6 of the dense known BRDF using CNNs in Sec. 5. There, we show quantitative and qualitative evaluations of our dense BRDF reconstruction experiments. A summary of the paper and our key results as well as an outlook to our further work is given in Sec. 6.

2. Light Field and BRDF

In order to measure BRDFs, usually gonireflectometers are utilized, where the word *gonio* is derived from the greek *gonía* which means angle and refers to the multiple direction capability of the device. Varying viewing and illumination angles allow the measurement of various incoming and outgoing light rays. For the measurement of anisotropic BRDFs all four dimensions have to be sequentially sampled [11].

Plenoptic cameras capture light fields by acquiring irradiance values from different viewing directions on scenes, using a lenslet array in front of the image sensor, as described in [12]. These light fields can be represented by two directional and two spatial dimensions. Thereby they provide 4D information of the flow of light through space in a static scene. Using multiple illumination sources allows both the estimation of local surface orientations (photometric stereo) as well as capturing additional reflectance properties.

Previously, a setup combining plenoptic acquisitions with photometric stereo, where the illumination sources were placed on a half sphere around the scene, was used in [13] in order to achieve an improved depth reconstructions of scene. Utilizing such a setup, comprising plenoptic cameras with several illumination sources, shows an ideal tool for measuring sparse BRDFs. Since the viewing direc-

tions are restricted by the cameras sensor size and position, one must address the question of the estimation of a dense BRDF from the data acquired from light field cameras.

For evaluating our reconstruction, a dense BRDF dataset is utilized which was captured with a gonireflectometer. We simulate the reduction of viewing angles by reducing the data accordingly. This dataset is described in the following section.

3. BRDF Dataset

The BRDFs from UTIA’s BTF Dataset [5] were used, which were measured with a high precision gonireflectometer [15], with an angular precision of 0.03 degrees and a spatial resolution of 1000 DPI. Data is then sampled at fixed azimuth and elevation angles for both the illumination directions and the camera position [16]. The dataset consists of 150 samples, comprising the materials of wood, fabric, scotchlite, car paint, carpet, corduroy, glitter, leather, paper, plaster, plastic, stone, wallpaper and plastic tape, as demonstrated in Fig. 2.

The data acquisition framework and data arrangement is shown in Fig. 3. The BRDF is described as a 4D function of $(\theta_i, \varphi_i; \theta_v, \varphi_v)$, where θ denotes the elevation- and φ the azimuth angle. These angles are defined for both the illumination direction I and the camera direction V in respect to the object normal N . The object normal points towards the top center of the acquisition device. The data contained in the dataset was interpolated in order to propagate information to missing parts of the BRDF subspace by an adapted swept surface technique (see [16]).

Limiting the cameras elevation angles towards the top center approximates capturing the scene with a plenoptic camera. Though it is obvious, that plenoptic cameras provide a more dense representation in a narrow angle range, we will be able to infer strong indications regarding the reconstruction abilities of BRDF data which is limited in its observation angles. Reducing the cameras elevation angles affect the material representation, as certain reflectance behaviour is unknown, as well as some of the representation of specular peaks and of anisotropic behaviour. Anisotropic patterns and specular peaks are visible in axial and diagonal slopes respectively, in a square $\theta_i \times \theta_v$ of the BRDF.

Each input feature stack for our CNN contains all azimuth variations (φ_i, φ_v) , where $\varphi_i = (0..2\pi)$ and $\varphi_v = (0..2\pi)$, which each cover 48 pixels at defined elevation angles (θ_i, θ_v) .

A description follows how the described dataset is used to approximate a dense BRDF from sparse samples.

4. BRDF Estimation

This section introduces our regression neural network and discusses the depth of the network structure. Using re-

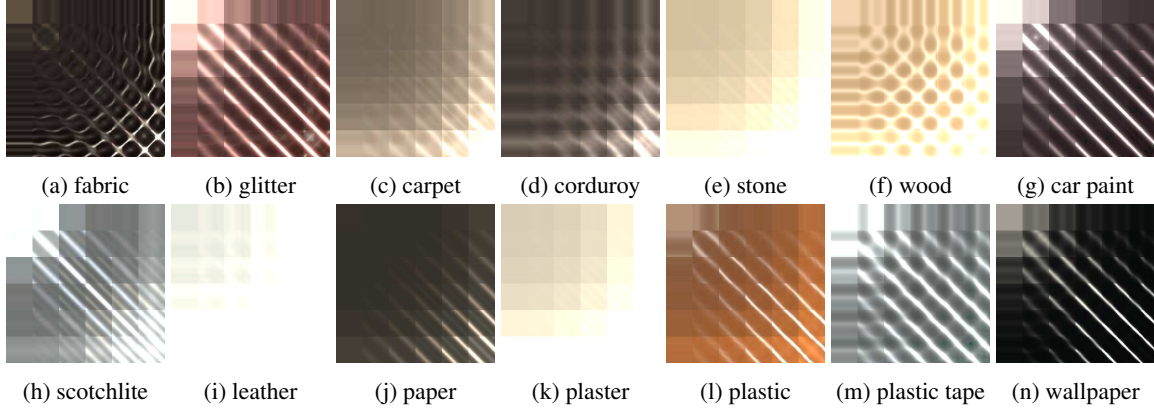


Figure 2: Examples of measured BRDFs of materials as contained in the UTIA dataset [14].




Figure 3: Acquisition setup and data arrangement from [14]. This illustration is courtesy of UTIA.

gression, one can estimate those parts of the BRDF which are unknown due to a restricted elevation viewing angle on the scene.

4.1. Network Structure

We consider the problem of estimating a full BRDF consisting of all measured views and illumination angles, which will further be referred to as a dense BRDF. This estimation is inferred from BRDF slices that are limited in their viewing directions. More specifically, we reduce the dimension θ_v from 15° to 75° , in step sizes of 15° . In the representation in Fig. 3 this is equal to removing 1 to 5 columns with each a width of 48 pixels, starting from the right side of the image. These unknown parts are estimated from the residual known data. For this regression problem, we make use of a CNN. The network can be represented as a function $f(x, w)$ that transforms the input data into a stack of output feature maps:

$$f(x, w) = f_{c4} \left(f_{r3} \left(f_{c3} \left(f_{r2} \left(f_{c2} \left(f_{r1} \left(f_{c1}(x, w_{c1}) \right), w_{c2} \right) \right), w_{c3} \right) \right), w_{c4} \right) \quad (1)$$

with a function relation

$$f : \mathbb{R}^{48 \times 48 \times (2u)} \rightarrow \mathbb{R}^{48 \times 48 \times (36-u)}. \quad (2)$$

The input data x of the dimension $48 \times 48 \times (2u)$ leads to a number of $2u$ input maps, where the number of input channels $2u$ contains $u = \{6, 12, 18, 24, 30\}$ BRDF slices of defined elevation angles (θ_i, θ_v) as well as their transposed counterparts. The benefit of that input data structure is depicted in Fig. 1, which shows an improved reconstruction when using enriched input data. Transposing the input channels supports the convolutional neural network to follow the Helmholtz reciprocity, which improves the reconstruction qualitatively and quantitatively (see Fig. 5), especially of the top right regions. This is described in more detail in Sec. 4.3

The input data is then mapped to $n = 36 - u$ output channels, using convolutional layers f_{c*} , followed by rectified linear units (ReLU) f_{r*} . The convolutional layers f_{c2} , f_{c3} and f_{c4} are 1×1 convolutional layers. The networks weight parameters $w = (w_{c1}, w_{c2}, w_{c3}, w_{c4})$ are initialized with normally distributed random numbers and learned from the input data in order to solve the problem of reconstructing the unknown BRDF regions. The 1×1 convolutional layers followed by ReLU layers benefit the reconstruction of the BRDF. We train the model to minimize a quadratic loss function $\sum_i^N (l_i - f(x_i, w))^2$, where l is the ground truth BRDF value and $f(x_i, w)$ the estimation for the input x_i .

A network architecture is used as shown in Fig. 4, the input consists of the known BRDF parts and their transposed versions.

4.2. Deep Neural Networks

For the regression task, we tested deeper convolutional network architectures, where we used the structure of the first 3×3 convolutional layer f_{c1} up to a number of 5 times. Multiple insertions of f_{c1} , each followed by the ReLU f_{r1} , are used to gain a deeper network structure. The network is depicted in Fig. 4. The depicted structure contains one 3×3 convolutional layer f_{c1} , while more insertions of that layer




Figure 4: CNN network structure.

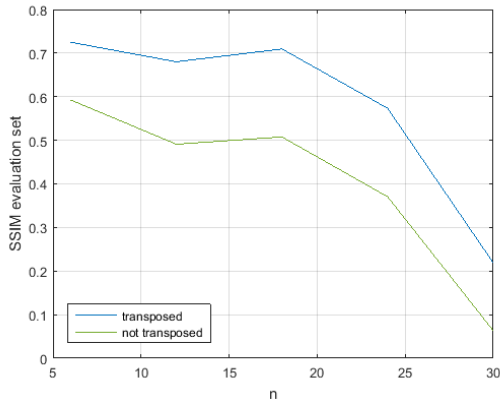


Figure 5: Comparison of the SSIM performance of our proposed network with transposed input data added to the original and the original input data (not transposed), both evaluated at epoch $e = 1300$.

demonstrate the performance of deeper networks. While the first convolutional layer is sensitive to local features as e.g. edge structures, deeper layers represent more complex combinations of lower level local features. For each network of increased depth, we observed the accuracy on the evaluation set. In Fig. 6, the performances of the networks are compared at several epochs e up to $e = 3000$. A performance drop is visible with each additional 3×3 convolutional layer, with a structural similarity (SSIM, described in detail in Sec. 5.1). Therefore, a deeper network of that type would not support us with our current reconstruction task. We chose a network with one 3×3 convolutional layer, as depicted in Fig. 4, as it shows superior performance compared to deeper network structures for the used dataset and architecture type.

Other kernel sizes were tested during our experiments. Where sizes up to 7×7 didn't improve the result in any significant way, a fully connected network showed to be infeasible due to memory limitations. The latter was addressed

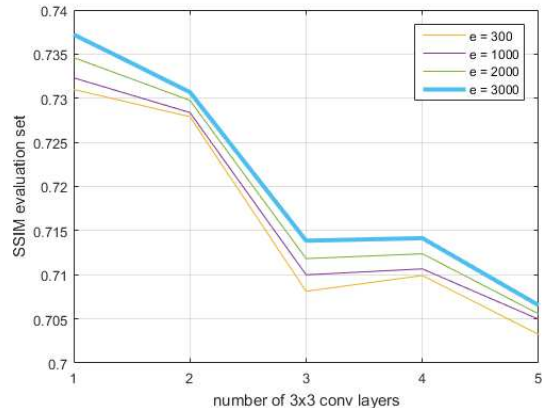


Figure 6: Evaluation of deep CNN structures up to epoch 3000 on the evaluation set, with $u = 30$, where the structure of the first 3×3 convolutional layer is used up to 5 times by additional insertions. The training set behaves in a similar way. The network was trained twice up to $e = 3000$, the second result showed the same tendency. Additionally we tested the mean absolute error (MAE) performance, which showed a comparable behaviour. Therefore the second result and the MAE are not displayed.

by enriching the input by the transposed data.

4.3. Enriched Input Data

The input data is enriched to help the network with the reconstruction task. In order to provide a higher variety in the input data, each BRDF slice is transposed and added to the original BRDF data. This is reasonable since the Helmholtz reciprocity holds for BRDFs of most realistic materials [17]. The principle describes that the BRDF value would be unmodified, when swapping the camera and illumination positions: $BRDF(\theta_i, \varphi_i; \theta_v, \varphi_v) = BRDF(\theta_v, \varphi_v; \theta_i, \varphi_i)$.

The results for both enriched and non-enriched input data are illustrated in Fig. 5, with output channels

$n = \{6, 12, 18, 24, 30\}$. Using a transposed input additionally to the original input channels allows a more realistic reconstruction, especially in regions where the camera and illumination positions are swapped and either of both is represented in the input data stack.

4.4. Training

The CNN architecture, which we described in Sec. 4, was trained by learning its weights using back propagation with the 150 BRDF examples from the dataset described in Sec.3, which was randomly divided into a training and an evaluation set. The probability of a sample to belong to the training set is 80%. The input data was enriched as described in Sec. 4.3. Data augmentation, which is frequently used in CNNs, was not implemented because of the strict reflectance behaviour of BRDFs.

5. Experimental Results

In this section the qualitative and quantitative evaluation performances of the results of our dense BRDF reconstruction are demonstrated.

5.1. Quantitative Evaluation

Performance evaluations are pursued for defined numbers of known elevation angles in the following order of $u = \{30, 24, 18, 12, 6\}$, which corresponds to maximum viewing elevation angles $\theta_v = \{60^\circ, 45^\circ, 30^\circ, 15^\circ, 0^\circ\}$. Fig. 7 shows the the reconstruction using three different evaluation methods, namely the structural similarity (SSIM), the mean absolute error (MAE), and a percentage of correct reconstruction compared to the total range of the ground truth square. In our evaluation representation, the x -axis denotes the number of unknown estimated BRDF data $n = 36 - u$ and the y -axis shows the performance measurements.

The SSIM is shown in Fig. 7a and Fig. 7b for the evaluation and training set respectively. We compute the SSIM as suggested by Wang *et al.* [18], while taking the mean value over our N estimated slices for all pixel values, as follows:

$$SSIM = \frac{1}{N \cdot P} \sum_{p=1}^P \sum_{i=1}^N (I_i(p))^\alpha \cdot (c_i(p))^\beta \cdot (s_i(p))^\gamma, \quad (3)$$

with the pixel position $p := (x, y)$ and a maximum pixel index P . While I is denoting the luminance, c the contrast and s a structural term, where weights α, β, γ are set to 1. The SSIM covers values from -1 to 1 , while a value of 1 describes two identical images.

The MAE is computed by the mean absolute differences between our estimation $f(p)$ and the ground truth value l

for each slice i and pixel position p :

$$MAE = \frac{1}{N \cdot P} \sum_{p=1}^P \sum_{i=1}^N |l_i(p) - f_i(p)|. \quad (4)$$

MAE evaluation results are shown in Fig. 7c. Here, a higher value represents a less accurate matching result.

The percentage rate PR of correct classification is defined by the distance from our estimation to the ground truth value in relation to the full value range of a specific slice:

$$PR = 100 \cdot \left(1 - \frac{1}{N \cdot P} \sum_{p=1}^P \sum_{i=1}^N \frac{|l_i(p) - f_i(p)|}{l_{i_{max}} - l_{i_{min}}} \right), \quad (5)$$

where the maximum value over each slice is defined by $l_{i_{max}} = \max_{p \in P} l_i(p)$ and the minimum value by $l_{i_{min}} = \min_{p \in P} l_i(p)$ for each slice $i \in N$. The PR evaluation results are shown in Fig. 7d. A value of 100% corresponds to a perfect match, while a value of 0% would mean that each pixel in our estimation is on average wrong by the full value range of that slice.

Deep convolutional layers were evaluated, where we used the structure of the first 3×3 convolutional layer f_{c1} up to 5 times. We observed the network for each increased depth up to epoch $e = 3000$ (see Fig. 6). Using our specific network type and input data, deeper structures showed a drop in performance with a range from the best result ($e = 3000$) of $SSIM = 0.737$ (one 3×3 convolutional layer) to $SSIM = 0.707$ (five 3×3 convolutional layers). The steepest drop was observed between the use of two and three of these layers. Therefore, a deeper network of that type would not support us with our current reconstruction task. Therefore, we chose a network with one 3×3 convolutional layer.

We tested the reconstruction with our trained networks for $u = 6, 12, 18, 24, 30$ at several iterations, up to the epoch $e = 2000$. The used dataset consists of 150 samples and was split in a training and evaluation set by random choice with a probability of 80% for a sample to belong to the training set. We show the performance on our evaluation set at several iterations as well as the SSIM on the training set (see Fig. 7). Notably, a higher number of iterations shows better performance, which indicates, that our network didn't overfit on the training data. The performance at iteration $e = 2000$ is numerically displayed in Tab. 1, using the SSIM, MAE and PR performance measurements. Competitive reconstruction results are shown when 1/6 ($n = 6$) to 3/6 ($n = 18$) of the BRDF is estimated, with an SSIM performance on the evaluation set of 0.742, 0.699 and 0.730 respectively. These results are represented in column 2 and 3 in Fig. 8, where even complex structures at a high angle θ_i and θ_v , shown at the bottom right of the BRDF structures, are well reconstructed.

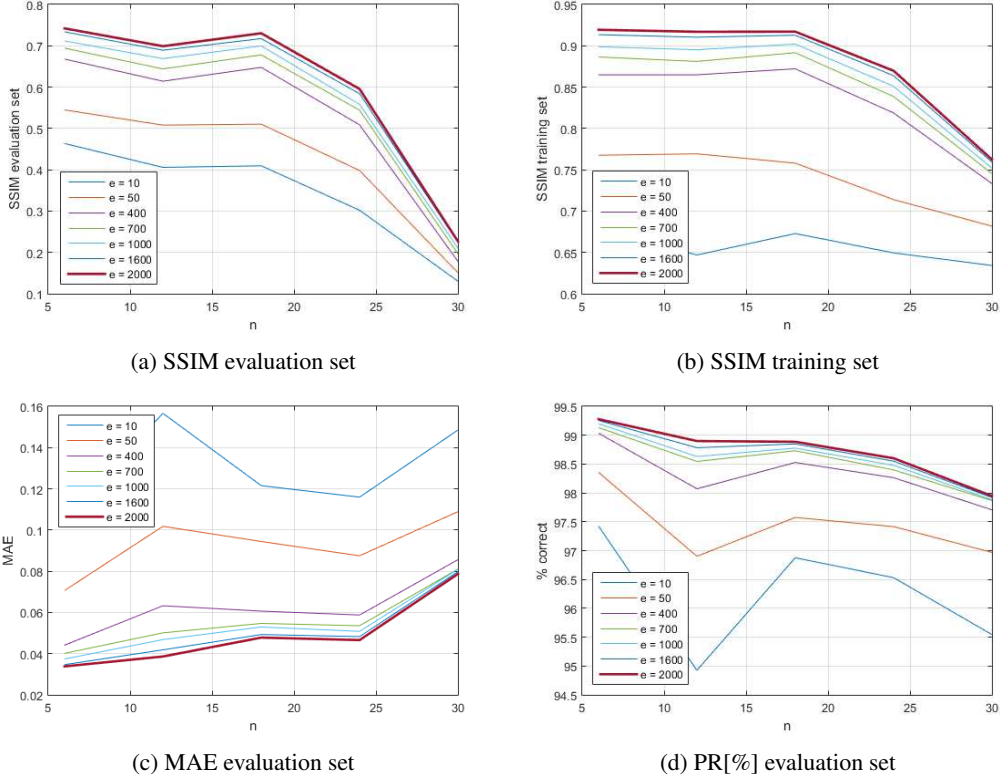


Figure 7: Performance evaluation from epoch $e = 10$ to 2000. The SSIM rate on the evaluation set is shown in (a), (b) depicts the SSIM training error. The MAE is shown in (c), and (d) demonstrates the correct classification in % compared to the full range in each BRDF slice.

type	set	n=6	n=12	n=18	n=24	n=30
SSIM	eval	0.742	0.699	0.730	0.596	0.227
MAE	eval	0.034	0.039	0.047	0.047	0.079
PR [%]	eval	99.28	98.90	98.89	98.60	97.94
SSIM	train	0.919	0.917	0.917	0.870	0.762
MAE	train	0.024	0.028	0.030	0.037	0.067
PR [%]	train	99.26	99.11	99.05	98.88	98.09

Table 1: Error evaluation at epoch 2000 of the evaluation set, for $n = \{6, 12, 18, 24, 30\}$.

5.2. Qualitative Evaluation

In order to evaluate the qualitative performance of our reconstruction, we show results from the introduced CNN structure, trained with input channels of the size $48 \times 48 \times (2u)$ (see Fig. 8). In terms of qualitative visual similarity, the second column ($u = 30$) with $n = 6$ unknown slices up to the the fourth column ($u = 18$) with $n = 18$ unknown slices, are very close to the ground truth. The last column ($u = 6$) with $n = 30$ unknown slices only a general tendency was learned with difficulties to reconstruct detailed structures. A major improvement was achieved by

adding transposed image data to the input channels from the used dataset (see Fig. 1).

6. Summary and Outlook

We introduced the reconstruction of dense BRDFs from limited observation angles, as characteristic for plenoptic image acquisitions. Using light field cameras, the scene is captured from different viewing angles, which are concentrated around a main viewing direction. Using a half sphere structure, various illumination sources can be placed around the scene at several illumination angles. The ability to infer




Figure 8: BRDF reconstructions where 5/6 to 1/6 of the dense BRDF are known. Columns from left to right: ground truth, $n = \{6, 12, 18, 24, 30\}$ unknown elements. The estimated parts of the BRDF data are framed with a red colour.

from thereby acquired sparse BRDF data to a more dense representation is of high interest.

In this work we showed the performance evaluation of several convolutional neural networks on BRDF data as well as the reconstruction ability and limitations. We showed the estimation of dense BRDF structures at five different sparsity levels on the UTIA dataset, where up to 5/6 of the BRDF structure was unknown.

Most notably, we gained an improved reconstruction of dense BRDF values from sparse BRDF data of low angles θ_i and high angles θ_v , which are represented towards the top right edge in Fig. 8, by using transposed images additionally in the input data. This means that the propagation of the Helmholtz reciprocity from the input layers to the output was supported.

This type of CNN architecture has strong local connections and therefore cannot capture complex behaviours such as the transpose, which is used to follow the Helmholtz reciprocity. We believe that other types of networks can improve the results in regions where the locality relation is weakened. Additional refinements can be achieved by further enriching the input data by useful structural information.

Future work will cover more sophisticated network structures for the reconstruction of the dense BRDF as well as partly reusing pre-learned networks of different topologies. U-Shaped networks were previously utilized for the depth estimation from light field data by Heber *et al.* [19], we will work on an adapted version of such a network structure with a weakened local orientation dependency. Our experiments will be additionally extended to more real-world objects for two different setups, namely a plenoptic camera with a light dome and an industrial acquisition setup with a multi-line-scan camera. For ground truth evaluations we will render BRDFs with our camera setups and reconstruct more complete BRDF data additionally.

7. Acknowledgements

This work is supported by the research initiative Mobile Vision with funding from the Austrian Federal Ministry of Science, Research and Economy and the Austrian Institute of Technology.

References

[1] F. Nicodemus, J. Richmond, J. Hsia, I. Ginsberg, and T. Limperis, “Geometrical considerations and nomenclature for reflectance,” 1977, nBS Monograph 160, U.S. Department of Commerce, National Bureau of Standards.

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

[3] CURET, “BRDF database,” <http://www.cs.columbia.edu/CAVE/curet/>, last viewed 2. March 2017.

[4] W. Matusik, H. Pfister, M. Brand, and L. McMillan, “A data-driven reflectance model,” vol. 22, no. 3, Jul. 2003, pp. 759–769.

[5] M. Haindl, J. Filip, and R. Vávra, “Digital material appearance: the curse of tera-bytes,” no. 90, 2012, pp. 49–50.

[6] O. Wang, P. Gunawardane, S. Scher, and J. Davis, “Material classification using BRDF slices,” in *In: Proc. CVPR*, 2009.

[7] T. Wang, J. Zhu, H. Ebi, M. Chandraker, A. A. Efros, and R. Ramamoorthi, “A 4d light-field dataset and CNN architectures for material recognition,” vol. abs/1608.06985, 2016. [Online]. Available: <http://arxiv.org/abs/1608.06985>

[8] Y. LeCun, P. Haffner, L. Bottou, and Y. Bengio, “Object recognition with gradient-based learning,” in *Shape, Contour and Grouping in Computer Vision*. London, UK, UK: Springer-Verlag, 1999, pp. 319–. [Online]. Available: <http://dl.acm.org/citation.cfm?id=646469.691875>

[9] B. Zupancic and C. Soler, “Sparse BRDF approximation using compressive sensing,” in *SIGGRAPH Asia 2013 Posters*, ser. SA ’13. New York, NY, USA: ACM, 2013, pp. 42:1–42:1. [Online]. Available: <http://doi.acm.org/10.1145/2542302.2542352>

[10] J. B. Nielsen, H. W. Jensen, and R. Ramamoorthi, “On optimal, minimal BRDF sampling for reflectance acquisition,” vol. 34, no. 6. New York, NY, USA: ACM, Oct. 2015, pp. 186:1–186:11. [Online]. Available: <http://doi.acm.org/10.1145/2816795.2818085>

[11] J. Filip, R. Vvra, and M. Havlicek, “Effective acquisition of dense anisotropic BRDF,” in *2014 22nd International Conference on Pattern Recognition*, Aug 2014, pp. 2047–2052.

[12] R. Ng, M. Levoy, M. Brédif, G. Duval, M. Horowitz, and P. Hanrahan, “Light field photography with a hand-held plenoptic camera,” vol. 2, no. 11, 2005, pp. 1–11.

[13] D. Antensteiner, S. Štolc, and R. Huber-Mörk, “Depth estimation with light field and photometric stereo data using energy minimization,” in *Progress in Pattern Recognition, Image Analysis, Computer Vision, and Applications: 21st Iberoamerican Congress, CIARP 2016, Lima, Peru, November 8–11, 2016, Proceedings, C.* ”Beltrán-Castañón, I. Nyström, and F. Famili, Eds. Springer International Publishing, 2017, pp. 175–183.

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

[15] M. Haindl, J. Filip, and R. Vávra, “Advanced textural representation of materials appearance,” no. 90, 2011, pp. 1:1–1:84.

[16] J. Filip, R. Vvra, M. Haindl, P. Id, M. Krupika, and V. Havran, “BRDF slices: Accurate adaptive anisotropic appearance acquisition,” in *2013 IEEE Conference on Computer Vision and Pattern Recognition*, June 2013, pp. 1468–1473.

- [17] B. Duvenhage, K. Bouatouch, and D. G. Kourie, “Numerical verification of bidirectional reflectance distribution functions for physical plausibility,” in *Proceedings of the South African Institute for Computer Scientists and Information Technologists Conference*, ser. SAICSIT '13. New York, NY, USA: ACM, 2013, pp. 200–208. [Online]. Available: <http://doi.acm.org/10.1145/2513456.2513499>
- [18] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, “Image quality assessment: from error visibility to structural similarity,” vol. 13, no. 4, April 2004, pp. 600–612.
- [19] S. Heber, W. Yu, and T. Pock, “U-shaped networks for shape from light field,” in *2016 British Machine Vision Conference BMVC*, 2016.