

Robust and Accurate Non-Parametric Estimation of Reflectance using Basis Decomposition and Correction Functions

Supplementary Material

Tobias Nöll^{1,2}, Johannes Köhler^{1,2}, and Didier Stricker^{1,2}

¹ German Research Center for Artificial Intelligence, Kaiserslautern, Germany

{Tobias.Noell, Johannes.Koehler, Didier.Stricker}@dfki.de

² University of Kaiserslautern, Germany

Abstract. In the first part of this additional evaluation we show that our algorithm achieves a significantly lower perceptual error (error of 0.19) for a set of 16 newly measured materials when compared to other state-of-the-art basis decomposition methods (LC, error of 0.50; RBF, error of 0.38). Additionally, we show that the BRDF estimates computed by RBF and the scattered data interpolation technique PP are often physically implausible, in contrast to the results provided by our method. In the second part of this additional evaluation we analyze how the size of the generated basis affects the fitting accuracy of the final solution.

1 Experiments with Measured Data

1.1 Introduction


To demonstrate the applicability of our method on real objects we measured the reflectance of a chocolate Santa Claus that was wrapped in a colored aluminium foil showing spatially varying reflectance behavior. The 3D geometry of the object was scanned using a structured light setup. Reflectance samples were captured from 766 different sensor / light directions.

We then clustered the reflectance samples of each surface point according to their diffuse (i.e. average) color which results in BRDF samplings for 16 base materials. To obtain the underlying dense BRDFs we applied our method (with 10 iterations) on every of these material samplings independently.

In this supplementary material we evaluate how well our algorithm can describe the individual materials based on the raw reflectance measurements. We also evaluate and compare the performance of other state-of-the-art methods (LC [4, 6, 5, 1], RBF [7, 6] and PP [4, 3]) in terms of perceptual error and physical plausibility.

Table 1. Perceptual error comparison (in terms of ε) for all evaluated 16 measured materials. To give a visual impression of these numerical values the table cells are colored according to the error (from blue (0.0, low perceptual error), over green to red (1.0, high perceptual error)). Our method outperforms the other basis decomposition methods (LC [4, 6, 5, 1], RBF [7, 6]) significantly. Because PP [4, 3] reproduces each measured value exactly the error is 0.0 constantly. However, the BRDF estimations achieved by PP are physically implausible and not suitable for synthesizing images with novel viewing or lighting configurations (see text for details)

		Method			
		LC	RBF	PP	Our
Material	1	0.37	0.31	0.00	0.19
	2	0.49	0.35	0.00	0.21
	3	0.53	0.47	0.00	0.23
	4	0.51	0.29	0.00	0.18
	5	0.58	0.29	0.00	0.19
	6	0.54	0.41	0.00	0.19
	7	0.37	0.34	0.00	0.19
	8	0.45	0.38	0.00	0.19
	9	0.57	0.37	0.00	0.19
	10	0.48	0.36	0.00	0.18
	11	0.47	0.41	0.00	0.23
	12	0.56	0.40	0.00	0.18
	13	0.54	0.55	0.00	0.20
	14	0.49	0.39	0.00	0.18
	15	0.61	0.44	0.00	0.19
	16	0.49	0.31	0.00	0.15
Avg.		0.50	0.38	0.00	0.19



1.2 Evaluation

Perceptual Error We cannot provide the perceptual CIELAB error [2] of the estimated fittings (as we did for the Benchmark section (3.1) of the main paper), as no ground truth is available for this newly measured dataset. Nevertheless, we want to provide a meaningful error that indicates how well the fitting is in the *perceptual sense*.

In the main paper we have shown that iteratively applying a correction function minimizes the (perceptually meaningful) CIELAB error [2]. Thus, the more correction functions are applied to a current estimate, the more accurate the current estimate gets in the perceptual sense and the more the current correction function approaches 1 constantly. We conclude, that the closer a correction function is to a constant 1 function the better the current approximation is in the perceptual sense.

Input to all algorithms is a sparse set of n measurements $(\theta_{h_i}, \theta_{d_i}, \phi_{d_i}, \rho_i, w_i)$, i.e. a BRDF was measured at $\rho(\theta_{h_i}, \theta_{d_i}, \phi_{d_i}) = \rho_i$. Each algorithm (LC, RBF, PP and our) then outputs a BRDF estimate $\varrho \approx \rho$ for this input. To describe the

error in the perceptual sense we compute a correction factor for each available raw measurement ρ_i as

$$\sigma_i = \frac{\rho_i}{\varrho(\theta_{hi}, \theta_{di}, \phi_{di})}. \quad (1)$$

We can then define an error metric that indicates how close the correction factors σ_i approach 1 by

$$\varepsilon_i = 1 - \begin{cases} \sigma_i & \text{if } \sigma_i \leq 1 \\ \frac{1}{\sigma_i} & \text{else} \end{cases}. \quad (2)$$

Finally, we compute the average error value ϱ as

$$\varepsilon = \frac{1}{n} \sum \varepsilon_i. \quad (3)$$

Based on our previous conclusions, the value ε then indicates the error of the current estimate in the perceptual sense (i.e. close to 0 is low perceptual error, close to 1 is high perceptual error).

For the experiments we compute a dense BRDF estimate ϱ for each of the 16 sparsely measured base materials using the respective algorithms (LC, RBF, PP and our) and compute the perceptual error ε for each of these estimation. The results for all experiments are presented in Table 1. Compared to the other basis decomposition methods LC (average $\varepsilon = 0.50$) and RBF (average $\varepsilon = 0.38$) our algorithm can achieve a significantly lower perceptual error (average $\varepsilon = 0.19$). We can thus conclude that our algorithm can significantly better describe the measured materials than the other evaluated basis decomposition method.

It is obvious that PP reproduces each measured value ρ_i exactly and thus has an error $\varepsilon = 0.0$ in all cases. Consequently ε is not expressive when comparing to PP. We have to empathize that this is not a drawback of how we defined the perceptual error ε and it will be the same for every differently defined error metric that compares to the raw measurements (since *all* measured values are represented *exactly* using PP). Because the perceptual error is not expressive in the case of PP we additionally compare the estimated solutions in terms of the physical plausibility.

Physical Plausibility An estimated BRDF should be physically plausible and be suitable for synthesizing images with novel viewing or lighting configurations. We thus additionally provide synthesized images of spheres (of the respective material estimations) under a natural illumination given by a HDR environment map (*Grace Cathedral*, courtesy of Paul Debevec, see Figure 7) and under a single light source illumination.

When comparing the synthesized images for each material the following observations can be made (see Figure 1):

1. The estimations achieved by PP are physically implausible. This is not surprising: First, the measurement noise including all outliers are exactly found in the resulting BRDF estimate. Also the outliers are not detected and interpolated within the whole BRDF domain. As a consequence many visual artifacts can be found in the synthesized sphere images.

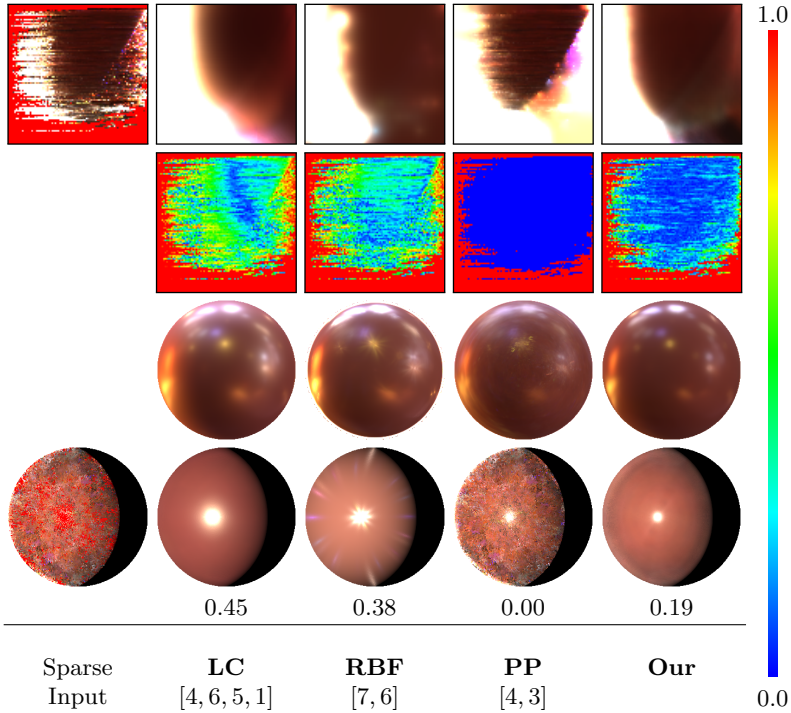


Fig. 1. Visual evaluation for measured material 8:

- Row 1 correspond to the BRDF data. The sparse input and outlier afflicted input can be seen in the first column (red indicates missing data). For visualization the 3D BRDF data was projected along ϕ_d (x-Axis $\hat{=}$ θ_h , y-Axis $\hat{=}$ θ_d). The other columns correspond to the estimations of the complete BRDF using the respective algorithm (LC [4, 6, 5, 1], RBF [7, 6], PP [4, 3], Our).
- Row 2 is a visualization of the perceptual error ε when compared to the original input (from blue (0.0, low perceptual error), over green to red (1.0, high perceptual error)).
- Row 3 shows a rendered sphere of the estimated material in a natural environment.
- Row 4 shows a rendered sphere of the estimated material under a single light source which shows the estimated BRDF values more precise. The sphere corresponding to the sparse input shows the available measured raw data for this specific camera / light configuration.
- Row 5 indicates the average perceptual error ε

2. Artifacts can also be found in the results achieved by RBF (e.g. star-like highlights). This is because RBF can not detect outliers and also tries to approximate these as well.
3. The estimations computed by LC are physically plausible because its basis consists of 100 densely measured BRDFs and is thus physically plausible itself. As a consequence no visual artifacts are found in the synthesized sphere images. However, the LC solutions have the largest perceptual error when compared to the other methods.

4. Our method achieves physically plausible results for two reasons: First, it is initialized based on a physically plausible result. Second, the basis used in each iteration step consists of 100 correction functions that were derived from densely measured materials and are thus also physically plausible. As a consequence, no artifacts can be observed in the synthesized sphere images. This is similar to LC, however, our method achieves a significantly lower perceptual error.

Additional examples for the estimated BRDFs can be found in Figure 2 - 5.

1.3 Additional Results

At this step we have already computed a number of 16 dense basis materials that the object consists of. To model the *overall* spatially varying reflectance behavior of the object we describe the reflectance of each surface point as a weighted combination of these base materials (similar to [3, 6]). We choose the weights using a simple non-negative least squares approach in order to remain physically plausible.

Choosing these weights accordingly is not in the scope of this paper and not evaluated further at this point. Nevertheless, we provide synthesized images of the whole object (using the basis materials computed by our algorithm) in Figure 6.

1.4 Conclusion

In this Section we provided additional evaluation for the case of 16 newly measured materials. We showed that our algorithm achieves a significantly lower perceptual error (error of 0.19) when compared to other state-of-the-art basis decomposition methods LC (error of 0.50) and RBF (error of 0.38). Additionally, we showed that the BRDF estimates computed by PP and RBF are often physically implausible. In contrast, our method provides physically plausible results that allow an artifact free synthesizing of images with novel viewing or lighting configurations.

2 Effect of the Basis Size


We also analyzed how the size of the basis affects the fitting accuracy of the final solution. We therefore repeated the benchmark in the main paper (see Subsection 3.1 of the main paper) with an outlier ratio of 0.0 and data ratio of 1.0. However, this time we also controlled the size of the basis (i.e. the number of basis functions used during basis decomposition). For each of the 100 test materials we chose random subsets of the original basis functions M_i and C_i of fixed sizes n . In this scenario $n = 99$ corresponds to the solutions calculated in the original benchmark of the main paper, $n < 99$ corresponds to solutions calculated by our algorithm using bases of reduced sizes. The mean CIELAB

Table 2. Measured CIELAB error (color coded) for different basis sizes n and methods

	Method			20.0 0.0
	RBF	LC	Our	
758	4.11			
99		4.88	0.28	
80		5.01	0.35	
60		5.26	0.43	
40		6.43	0.67	
20		7.54	1.22	
10		10.29	2.68	

Table 3. The measures indicate that the expected fitting error of our method in relation to n can be predicted as ε_n

Basis size (n)	Perceptual Error	
	Measured	Predicted
99	0.28	0.28
80	0.35	0.35
60	0.43	0.47
40	0.67	0.70
20	1.22	1.41
10	2.68	2.82



error [2] for varying values of n is provided in Table 2. Note that our algorithm together with bases of sizes $n = 10$ still provides solutions with a 2 times lower perceptual error than LC with $n = 99$ or RBF with $n = 758$ basis functions.

The measurements also indicate (see Table 3) that the expected fitting error of our method in relation to n can be predicted as

$$\varepsilon_n = \varepsilon \frac{99}{n}, \quad (4)$$

where ε corresponds to the original CIELAB error using the original bases with $n = 99$ M_i and C_i . This means if only half of the basis functions are used, the fitting error of our algorithm can be expected to be twice as high. To the contrary, using our algorithm together with a basis having double the amount of functions can be expected to have a two times lower perceptual error for an estimated solution.

3 Acknowledgment

Our work has been partially funded by the project DENSITY (01IW12001) and the Google Research Award received by Didier Stricker, University of Kaiserslautern.

References

1. Ali, M.A., Sato, I., Okabe, T., Sato, Y.: Toward efficient acquisition of brdfs with fewer samples. In: Proceedings of the 11th Asian Conference on Computer Vision - Volume Part IV. pp. 54–67. ACCV'12, Springer-Verlag, Berlin, Heidelberg (2013), http://dx.doi.org/10.1007/978-3-642-37447-0_5
2. Fairchild, M.D.: Color appearance models. John Wiley & Sons (2013)
3. Lawrence, J., Ben-Artzi, A., DeCoro, C., Matusik, W., Pfister, H., Ramamoorthi, R., Rusinkiewicz, S.: Inverse shade trees for non-parametric material representation and editing. In: ACM SIGGRAPH 2006 Papers. pp. 735–745. SIGGRAPH '06, ACM, New York, NY, USA (2006), <http://doi.acm.org/10.1145/1179352.1141949>
4. Matusik, W., Pfister, H., Brand, M., McMillan, L.: Efficient isotropic brdf measurement. In: Proceedings of the 14th Eurographics Workshop on Rendering. pp. 241–247. EGRW '03, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland (2003), <http://dl.acm.org/citation.cfm?id=882404.882439>
5. Ren, P., Wang, J., Snyder, J., Tong, X., Guo, B.: Pocket reflectometry. In: ACM SIGGRAPH 2011 Papers. pp. 45:1–45:10. SIGGRAPH '11, ACM, New York, NY, USA (2011), <http://doi.acm.org/10.1145/1964921.1964940>
6. Weistroffer, R.P., Walcott, K.R., Humphreys, G., Lawrence, J.: Efficient basis decomposition for scattered reflectance data. In: Proceedings of the 18th Eurographics Conference on Rendering Techniques. pp. 207–218. EGSR'07, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland (2007), <http://dx.doi.org/10.2312/EGWR/EGSR07/207-218>
7. Zickler, T., Enrique, S., Ramamoorthi, R., Belhumeur, P.: Reflectance sharing: Image-based rendering from a sparse set of images. In: Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques. pp. 253–264. EGSR'05, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland (2005), <http://dx.doi.org/10.2312/EGWR/EGSR05/253-264>




Fig. 2. Visual evaluation for measured material 1 and 3 (see caption of Figure 1 for explanation)

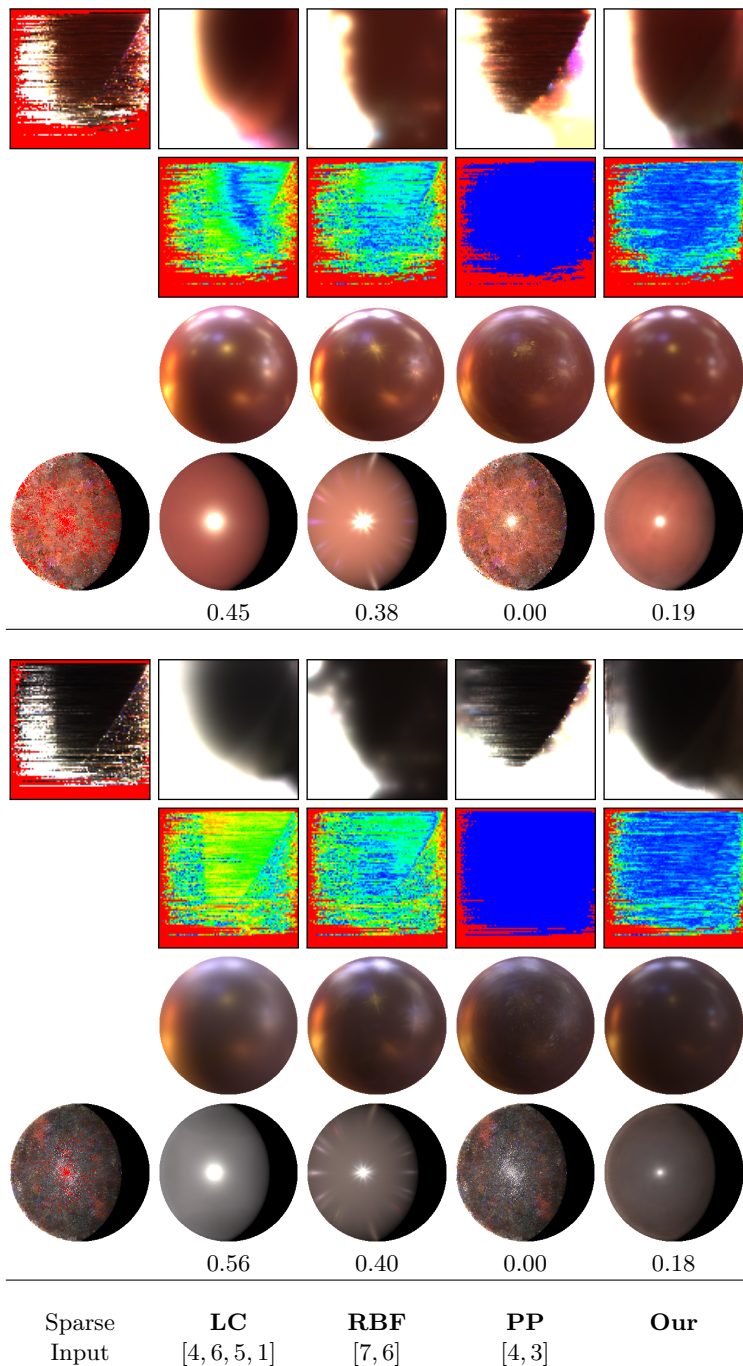


Fig. 4. Visual evaluation for measured material 8 and 12 (see caption of Figure 1 for explanation)



Fig. 6. Synthesized images of the measured object under novel viewing and illumination conditions (1st row – Grace Cathedral, 2nd row – St. Peter’s Basilica, 3rd row – Uffizi Gallery)




Fig. 7. The HDR environment maps that define the natural illumination environment (courtesy of Paul Debevec) for the synthesized sphere images in the main paper and in this supplementary material:

- Grace Cathedral (left, used for the Benchmark section (3.1) of the main paper and for all sphere images in this supplementary material)
- St. Peter's Basilica (middle, used for Figure 2 of the main paper)
- Uffizi Gallery (right, used for Figure 4 and 5 of the main paper)