

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

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. A common approach to non-parametric BRDF estimation is the approximation of the sparsely measured input using basis decomposition. In this paper we greatly improve the fitting accuracy of such methods by iteratively applying a novel correction function to an initial estimate. We also introduce a basis to efficiently represent such a function. Based on this general concept we propose an iterative algorithm that is able to explicitly identify and treat outliers in the input data. Our method is invariant to different error metrics which alleviates the error-prone choice of an appropriate one for the given input. We evaluate our method based on a large set of experiments generated from 100 real-world BRDFs and 16 newly measured materials. The experiments show that our method outperforms other evaluated state-of-the-art basis decomposition methods by an order of magnitude in the perceptual sense for outlier ratios up to 40%.

Keywords: Non-parametric BRDF estimation, reflectance, basis decomposition, correction function, error metric, sparse data, outliers

1 Introduction

How an object *appears* in reality is essentially determined by the complex interaction of its shape, its surface materials and the lighting environment it is currently observed in. Reproduction of this appearance is important in the context of various application domains such as advertisement, movie-production and cultural heritage preservation where realistic images of real objects need to be synthesized. This can be achieved by capturing and modeling the shape and surface material properties of the objects. In the last decades a considerable amount of progress has been made that allows to capture and model the 3D shape of an object precisely. This paper thus focuses on the robust and precise modeling of captured surface material properties only.

It is well known that the appearance properties for opaque materials are effectively described using the *Bidirectional Reflectance Distribution Function* (BRDF) [22]. For a specific material this function describes how much light from a specific *incident* direction is reflected to an *outgoing* direction. Various

types of BRDF measuring devices were proposed in the past – one device group for measuring isolated material probes (e.g. [18, 19]), another group allowing a combined acquisition of geometry and reflectance of general objects (e.g. [12, 14]). Regardless of the respective technical implementation all devices have in common that they can independently drive a light source and a sensor to different positions around the object of interest and thereby sample how the BRDF responds at these specific incident / outgoing directions. Densely sampling the BRDF is however intractable: As pointed out in [13], sampling at an angular resolution of 1 degree would already amount to a total number of 2×10^8 required measurements. Further, the BRDF is usually a complex function that can change drastically by several orders of magnitude for even small angular changes. This makes the extrapolation of a complete BRDF from a sparse set of measurements a complicated task that has not been completely solved yet. However, a *full* reconstruction that describes how the BRDF behaves under all configurations is essential for photo-realistic image generation.

In this paper we assume that we are provided with a sparse, irregularly sampled set of angular measurements. The task is to reconstruct the complete BRDF that accurately describes the sparsely measured behavior. This is particularly difficult for the following two properties of real measured data: First, even if the light and capturing directions were sampled in a regular way, the measurements themselves are usually irregular for curved objects, because the local coordinate system changes over the surface. Additionally, holes exist for areas that were occluded during acquisition. Second, measured reflectance data often contains outliers originating from imprecise calibration of the devices, imprecise reconstruction of the geometry, global illumination artifacts or self-shadowing. Our method is robust with respect to these practical considerations, achieves more accurate solutions than proposed state-of-the-art methods and to the best of our knowledge we are the first to explicitly consider and identify outliers during BRDF reconstruction (an overview of the problem is given in Figure 1).

1.1 Related Work

Parametric model fitting: Traditionally, parametric models – either empirical or based on a theoretical model – provide the BRDFs currently used in computer graphics. These range from ad-hoc models (e.g. Blinn-Phong [3], Lafortune [15], Ashikhmin [2], DSBRDF [23]) designed for efficiency, to physically derived descriptions either based on micro-facet theory (e.g. Ward [30], Cook-Torrance [4], Schlick [28]) or wave optics (e.g. He [11]). Independent of the derivation all parametric models are analytic functions that are well defined over the whole BRDF domain and only depend on a small set of meaningful parameters. Consequently, an obvious solution to the task would be to choose a parametric model and tune its parameters in a way that best fits the sparse set of actual measurements (in fact some models [15, 30] were designed especially with this purpose in mind). The ability of analytic functions to describe measured reflectance behavior was examined in [21]. As it turns out, reducing a rich set of measurements to only a small set of parameters can introduce significant errors: First, a specific

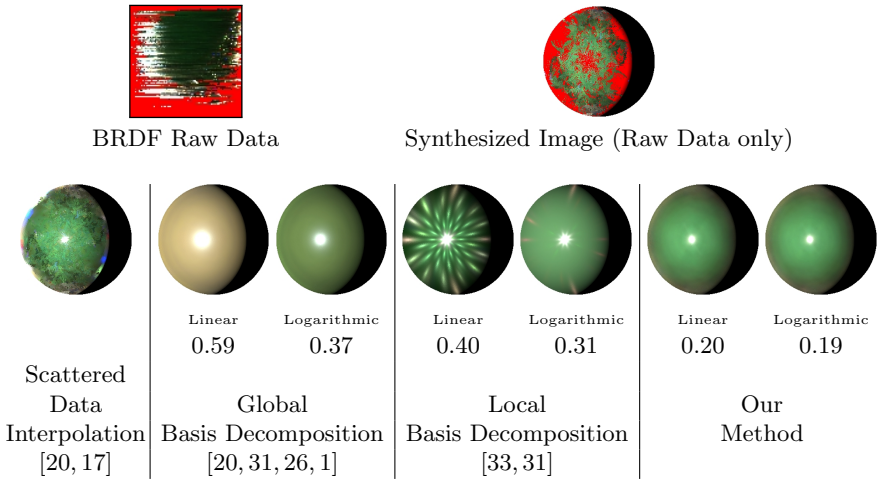


Fig. 1. Problem overview: First row (left) shows raw measured isotropic BRDF data (a 2D projection of the 3D BRDF table for visualization). Red corresponds to unmeasured areas. The right side of the first row shows a synthesized sphere based on the raw BRDF data. Note that for this specific camera / light configuration e.g. the specular lobe was never measured. A BRDF estimation algorithm tries to compute the full BRDF data based on the sparse raw input. The second row shows synthesized spheres using the estimates of different non-parametric algorithms (numerical values indicate the fitting error). The different state-of-the-art algorithms compute either physically implausible results (see scattered data interpolation and local basis decomposition) or non-accurate solutions (see global basis decomposition). Further, the estimate for the basis decomposition methods is highly dependent on the choice of an appropriate error metric for the current dataset (compare linear vs. logarithmic error metric) which can result in unpredictable solutions. Our method achieves accurate estimates that are additionally physically plausible. Additionally, our method is invariant with respect to different error metrics

reflectance model must be chosen *a priori*. However, each model was designed having a specific reflectance behavior in mind. In an extreme case there might not be *any* parameter configuration for the chosen model that describes the measured data well. Second, an error metric has to be chosen in order to mathematically define what the *best* parameter setting is. As pointed out in [9] this is non-trivial because choosing a parameter setting that minimizes a specific numerical error might not lead to a well fitting result in the perceptual sense. For BRDF fitting it is not yet clear which error metric is the most suitable to model human perception. Third, the models are usually highly nonlinear in their parameters which requires nonlinear optimization techniques for the parameter fitting. However, for these methods the quality of the fit is dependent on a good initial guess, and reaching a global minimum cannot be guaranteed. Because of these issues, research has shifted towards a non-parametric description of measured materials.

Non-parametric description: The methods in this category allow for a potentially greater accuracy and generality for describing measured reflectance behavior. The most general, non-parametric description of a BRDF is a regularly sampled multidimensional *table* where the available scattered data points are resampled to. To fill the unmeasured entries, it was proposed in [20, 17] to use the general purpose scattered data interpolation technique *Pull-Push* [10].

Most non-parametric techniques however describe the BRDF using basis decomposition. Hereby a BRDF ρ is approximated using a weighted linear combination of basis functions Ψ_i as

$$\rho(\mathbf{x}) \approx \sum_i \alpha_i \Psi_i(\mathbf{x}) . \quad (1)$$

The literature proposes a large set of suitable basis functions Ψ_i : Spherical harmonics [32, 25], Zernike polynomials [13], spherical wavelets [29, 16, 20], radial basis functions [33, 31], measured BRDFs [20, 31, 26, 1] and rational functions [24]. The basis operates either *globally* (i.e. each Ψ_i is non-zero in large areas of the parameter domain) or *locally* (i.e. each Ψ_i is non-zero only in a small area of the parameter domain).

In this paper we address two problems that exist for all basis decomposition methods in general:

1. The choice of a local or global basis is until today a choice of accuracy versus robustness (as also examined in [31]): Choosing a local basis will allow for a better matching of the measured input since the weighting factors α_i can be chosen more independently. However, outliers in the input will then directly propagate to the final solution. Choosing a global basis behaves reversely: While it is in general more robust with respect to outliers, the weighting factors are less independent, thus the approximation accuracy is degraded.
2. Previous methods are not robust with respect to different error metrics (see Figure 1): The weighting factors a_i are chosen in a way that the linear combination fits best to the measured input data. Thus, an appropriate error metric to be minimized must be chosen *a priori*. This generates the same problems as in the parametric case: The estimated solution quality is highly dependent on this choice and a minimization in the numerical sense is not necessarily a minimization in the perceptual sense.

We propose a non-parametric method that addresses these problems as follows: It operates globally, thus is very robust with respect to outliers in the input. Our key idea is then to avoid the inevitably reduced accuracy of a global basis by iteratively applying different corrections to the initial solution. We therefore introduce the novel concept of BRDF correction functions and provide a novel global basis that can be used to model these functions effectively. In order to converge to the best solution during this iterative process we explicitly consider and identify outliers in the measured data. To the best of our knowledge this was not done by any previously proposed method. Additionally, our method is intrinsically robust with respect to different error metrics due to its iterative nature: A specific error metric only changes the iteration sequence but the overall

estimate still converges to the same solution. We show in the evaluation section that our method has a higher robustness and accuracy than previous global and even local and tabular state-of-the-art methods.

1.2 Background

In its most general form a BRDF ρ can be represented as a 4-dimensional table $\rho(\theta_i, \phi_i, \theta_o, \phi_o)$ which is indexed by the spherical coordinates (θ, ϕ) of the incident and outgoing direction. This function then defines how much light from an incident direction i is reflected to an outgoing direction o .

It was shown that the *Rusinkiewicz re-parametrization* [27] is much more suitable for basis decomposition because features in common BRDFs are aligned with the transformed coordinate axes. The BRDF is therefore usually expressed as a function based on the half-vector h between incoming and outgoing light direction and a difference vector d (please refer to [19, 20, 27] for a more detailed explanation). This yields

$$\rho(\theta_h, \phi_h, \theta_d, \phi_d) . \quad (2)$$

Several constraints can be formulated regarding the estimated solution which can reduce the dimensionality (i.e. the complexity) of the BRDF function: Many methods [20, 25, 13, 33, 31, 26, 24] (and also our) estimate a solution in the space of *isotropic* materials and such materials can already be represented using a 3-dimensional function $\rho(\theta_h, \theta_d, \phi_d)$. All opaque materials without “grain” are in fact isotropic, hence this establishes a valid and meaningful complexity reduction.

2 Method

To robustly initialize our method we use a global basis: It was examined in [20, 31, 26, 1] that novel BRDFs can be described using a basis of previously measured BRDFs. In [19] a publicly available database is introduced that contains 100 densely and precisely measured isotropic materials M_i that are used in [20, 31, 1] as basis functions.

Our initial estimate ϱ for a sparsely measured BRDF ρ can thus be written as a weighted linear combination

$$\varrho(\theta_h, \theta_d, \phi_d) = \sum_i \alpha_i M_i(\theta_h, \theta_d, \phi_d) \approx \rho(\theta_h, \theta_d, \phi_d) . \quad (3)$$

Our key idea is to formulate a BRDF correction function σ that represents the error of this initialization. There exist many different possibilities to do this, but we chose to represent σ by means of scaling factors, which yields

$$\sigma(\theta_h, \theta_d, \phi_d) = \frac{\rho(\theta_h, \theta_d, \phi_d)}{\varrho(\theta_h, \theta_d, \phi_d)} . \quad (4)$$

Using scaling factors is very reasonable because we will later only transform *similar* BRDFs, thus its values will be closely distributed around 1 for the whole

parameter domain. Consequently, a correction function (as defined in Equation 4) is usually of drastically lower complexity than a BRDF which may change rapidly. If the correction function σ is known, the initial estimate is corrected by

$$\rho(\theta_h, \theta_d, \phi_d) = \sigma(\theta_h, \theta_d, \phi_d) \varrho(\theta_h, \theta_d, \phi_d) . \quad (5)$$

It is clear that finding the *true* correction function σ is as hard as finding the *true* BRDF ρ . However, it will turn out later that finding a good *approximation* can be performed robustly and more efficiently than in the BRDF case. When applying such an approximate correction function on ϱ we can improve our initial estimate and iterate this process until convergence.

Our approach to approximate σ is to perform another basis decomposition using a set of suitable basis functions C_i , which yields

$$\sigma(\theta_h, \theta_d, \phi_d) \approx \sum_i \beta_i C_i(\theta_h, \theta_d, \phi_d) . \quad (6)$$

Our intuition was – analogous to the BRDF approximation – that novel correction functions can be well described using a basis of previously generated correction functions C_i .

To generate a large set of correction functions we again used the material database provided in [19]: For each BRDF M_i from this database, we compute an approximation ϱ using the remaining 99 materials as a basis. We then compute the corresponding correction function σ using Equation 4 (with $\rho = M_i$). Note that σ is well defined over the whole isotropic BRDF parameter domain.

For each correction function C_i generated this way, we observed two important characteristics:

1. The values of each correction function are distributed within a narrow range.
2. Each correction function itself is a relatively smooth function.

These characteristics are in sharp contrast to measured BRDF data whose minimal and maximal values often differ by several orders of magnitude and usually show a disproportionately steep increase close to the specular lobe. This justifies our concept of using BRDF corrections since it indicates that the *space of all correction functions* is a less complex space than the *space of all materials*.

To verify that C_i indeed defines an expressive basis we evaluated how well each correction function itself is described using the remaining 99 correction functions as a basis. Even though the functions are of very different appearance they are in general very well represented within this basis (see Figure 2) with an average deviation to the original scaling function of only 0.076 units. This indicates that the space of *all* correction functions is well described through our generated basis. In the supplementary material we additionally analyze how the size of the generated basis affects the fitting accuracy of the final solution.

2.1 Algorithm

Based on the previous considerations we outline the following algorithm that exploits the concept of correction functions to provide a robust and accurate BRDF estimate for a sparse set of material measurements:




Fig. 2. Our novel BRDF corrections basis (11 out of 100); 1st row: Color coded 2D slices (x-Axis $\hat{=}$ θ_h , y-Axis $\hat{=}$ θ_d , $\phi_d = 90^\circ$) of our generated 3D BRDF correction functions C_i for densely measured materials. 2nd row: Approximations of the upper correction functions using the remaining 99 correction functions as basis. Each of these functions is well approximated even though they have very different appearance

1. We compute a *dense* initial material estimate for the sparse input using M_i as basis functions.
2. We compute a correction factor for each value of this estimate where an input measurement is available. During this process outliers are identified and assigned with a low weight.
3. For this weighted sparse set of correction factors we can approximate the underlying *dense* correction function using our generated C_i as basis functions.
4. We improve our initial estimate by applying this correction function.
5. We iterate this procedure until convergence (go to step 2).

2.2 Basis Decomposition

In this subsection we demonstrate how we perform the basis decomposition used for the initialization and for each iteration step of our algorithm.

We assume that we are initially provided with a sparse set of n measurements $(\theta_{h_i}, \theta_{d_i}, \phi_{d_i}, \rho_i, w_i)$, i.e. a function was measured at $\rho(\theta_{h_i}, \theta_{d_i}, \phi_{d_i}) = \rho_i$ with a confidence of w_i (w_i could e.g. be provided by the measurement device or set to 1 if unknown). Given an error metric ε we compute the *best* representation of this function by a dense basis $\Psi = \{\Psi_1, \dots, \Psi_m\}$ as

$$\begin{pmatrix} \varepsilon(\rho_1) \\ \vdots \\ \varepsilon(\rho_n) \end{pmatrix} \approx \begin{bmatrix} \varepsilon(\Psi_1(\theta_{h_1}, \theta_{d_1}, \phi_{d_1})) & \dots & \varepsilon(\Psi_m(\theta_{h_1}, \theta_{d_1}, \phi_{d_1})) \\ \vdots & & \vdots \\ \varepsilon(\Psi_1(\theta_{h_n}, \theta_{d_n}, \phi_{d_n})) & \dots & \varepsilon(\Psi_m(\theta_{h_n}, \theta_{d_n}, \phi_{d_n})) \end{bmatrix} \begin{pmatrix} \alpha_1 \\ \vdots \\ \alpha_m \end{pmatrix}, \quad (7)$$

abbreviated as $b \approx Ax$.

We can write the squared error of this basis decomposition as

$$\begin{aligned} d^T d &= (Ax - b)^T (Ax - b) = (x^T A^T - b^T)(Ax - b) \\ &= x^T A^T Ax - 2b^T Ax + b^T b. \end{aligned} \quad (8)$$

To include the given confidence weighting we define $W = \text{diag}(w_1, \dots, w_n)$ and extend Equation 8 as

$$d^T d = x^T A^T W Ax - 2b^T W Ax + b^T W b. \quad (9)$$

We want to minimize the squared error, thus we can drop the constant term $b^T W b$ and define $Q = A^T W A$ and $c^T = b^T W A$. This leads to the canonical form of a *quadratic programming* (QP) problem, which is

$$\begin{aligned} & \underset{x}{\text{minimize}} && J(x) = \frac{1}{2} \mathbf{x}^T Q \mathbf{x} + c^T \mathbf{x} , \\ & \text{subject to} && D \mathbf{x} \leq \mathbf{d} \text{ (inequality constraint) and} \\ & && E \mathbf{x} = \mathbf{e} \text{ (equality constraint) .} \end{aligned} \quad (10)$$

We can solve this efficiently, e.g. using an online active set strategy [7] implemented in [6]. We set $D = -I$ and $d, E, e = 0$ to constrain $x \geq 0$ for a basis decomposition with only positive factors. The final solution is then given as

$$\rho(\theta_h, \theta_d, \phi_d) \approx \varepsilon^{-1} \left(\sum_{i=1}^m \alpha_i \varepsilon(\Psi_i(\theta_h, \theta_d, \phi_d)) \right) . \quad (11)$$

2.3 BRDF Correction

If we are provided with a dense BRDF estimate ϱ for a sparse set of n measurements $(\theta_{h_i}, \theta_{d_i}, \phi_{d_i}, \rho_i, w_i)$ we may generate a sparse set of correction factors $(\theta_{h_i}, \theta_{d_i}, \phi_{d_i}, \sigma_i, v_i)$ by setting

$$\sigma_i = \frac{\rho_i}{\varrho(\theta_{h_i}, \theta_{d_i}, \phi_{d_i})} . \quad (12)$$

Using our novel basis functions C_i , the underlying *dense* correction function σ may be generated from this sparse set. We can then use σ to correct the current BRDF estimate (Equation 5).

However, very relevant in this context are outliers in the measured BRDF values ρ_i . These will result in wrong scaling factors that will affect the quality of σ . In the worst case the generated correction function could even degrade the current estimate. As a result our method would diverge. It is thus extremely important to identify and treat all outliers accordingly. An intrinsic feature of our algorithm is that at each iteration step ϱ is already a robust, non-parametric approximation of the sparsely measured BRDF ρ . By comparing each measured input value with its corresponding estimate $\varrho_i = \varrho(\theta_{h_i}, \theta_{d_i}, \phi_{d_i})$ we can thus robustly estimate outliers even in the non-parametric case and assign a low weight v_i to them before estimating the dense correction function. We set

$$v_i = w_i e^{-\gamma \frac{|\rho_i - \varrho_i|}{\varrho_i}} , \quad (13)$$

i.e. we assign an exponentially lower weight to the correction factor the more ϱ_i and ρ_i differ. γ is the only parameter of our algorithm and needs to be chosen appropriately depending on the dynamic range of the input values and the estimated outlier ratio.

Note that there might be more elaborate heuristics for detecting outliers. However, we found this simple one to work quite robustly. During any of the experiments performed for the evaluation we never experienced our method diverging (though we cannot guarantee this formally).

2.4 Implementation

Before our algorithm can be applied, we need to compute the new basis C_i once. This can be done as explained in Section 2. We used a logarithmic error metric $\varepsilon(x) = \ln(1 + x)$ and $W = I$ for the basis decomposition (Subsection 2.2). The *same* basis C_i is used for all iterations of our algorithm.

When provided with a sparse set of measurements $(\theta_{hi}, \theta_{di}, \phi_{di}, \rho_i, w_i)$ our algorithm performs the following steps (individually per RGB color channel) to estimate the complete BRDF ρ :

1. Initialize a dense BRDF ϱ for the sparse input using basis decomposition with

$$\Psi = \{M_1, \dots, M_{100}\} \text{ and } \varepsilon(x) = \ln(1 + x) . \quad (14)$$

2. Generate a sparse set of correction factors $(\theta_{hi}, \theta_{di}, \phi_{di}, \sigma_i, v_i)$ with:

$$\sigma_i = \frac{\rho_i}{\varrho(\theta_{hi}, \theta_{di}, \phi_{di})} \text{ and } v_i = w_i e^{-\gamma \frac{|\rho_i - \varrho_i|}{\varrho_i}} \quad (\varrho_i = \varrho(\theta_{hi}, \theta_{di}, \phi_{di})) . \quad (15)$$

3. Using these correction factors, estimate a dense BRDF correction function σ using basis decomposition with

$$\Psi = \{C_1, \dots, C_{100}\} \text{ and } \varepsilon(x) = x . \quad (16)$$

4. Correct the current BRDF estimate ϱ using

$$\varrho(\theta_h, \theta_d, \phi_d) := \sigma(\theta_h, \theta_d, \phi_d) \varrho(\theta_h, \theta_d, \phi_d) . \quad (17)$$

5. Stop if the maximum iteration count has been reached or σ is 1 *almost* everywhere, otherwise go to 2.

3 Evaluation

We first demonstrate how our method is executed based on an example. We assume a sparse and irregular sampling of an isotropic BRDF to be given as input which was obtained by capturing a specular red material under different viewing and lighting directions. Our task is to extrapolate this sparse input to a consistent tabular dense BRDF representation suitable for image synthesis. We assume that only 10% of the discretized BRDF domain has been sampled at random positions. However, due to inevitable inaccuracies in the geometry reconstruction as well as in the sensor and light calibration, the data set exhibits a significant amount of outliers: In 40% of the cases a BRDF measurement value does not belong to the provided calibrated pair of incident and outgoing direction but to a random value of the BRDF.

The high ratio of outliers makes this dataset extremely challenging for existing non-parametric methods. In Figure 3 we show the solutions provided by different types of non-parametric state-of-the-art methods:




Fig. 3. Evaluation of different BRDF estimation techniques: The top row shows a 2D slice of the underlying 3D BRDF data (x -Axis $\hat{=}$ θ_h , y -Axis $\hat{=}$ θ_d , $\phi_d = 90^\circ$), while the center row shows a rendered sphere of this material in a natural environment. The bottom as well as the numerical values indicate the image difference in the CIELAB [5] space of the current sphere with respect to the ground truth (from blue (0.0), over green to red (20.0 and higher)). All evaluated method types (global, local, tabular) could not provide a satisfactory solution. Moreover, the solution quality also depends highly on the chosen error metric (e.g. linear $\varepsilon_{\text{lin}}(x) = x$, logarithmic $\varepsilon_{\text{log}}(x) = \ln(1 + x)$)

– *Global basis decomposition*: Using a linear combination of measured materials (as proposed in [20, 31, 26, 1]) performs most physically meaningful with respect to the outliers. The reason for this is that each basis function is *globally* defined over the whole BRDF domain, thus physically non-meaningful outliers cannot be part of the solution. However, the outliers do affect the global solution quality, leading to a bad overall fit for in example.

– *Local basis decomposition*: In this class of methods the solution is provided by a basis decomposition where each basis function only has a *local* support on the BRDF domain (such as radial basis functions as proposed in [33, 31]). These methods have a larger flexibility and allow for a potentially better fitting. However, outliers directly contribute to each basis function making the overall fitting less robust and less physically plausible.

– *Tabular representation*: This representation provides the largest flexibility because each measured value is *directly* represented. The missing values are completed using scattered data interpolation techniques (such as *Pull-Push* [10] as proposed in [20, 17]). However, because the outliers are also interpolated over the BRDF domain this results in a physically implausible solution in this example.

All these evaluated method types cannot provide a satisfactory solution for our exemplary dataset. Also the final solution quality is highly dependent on

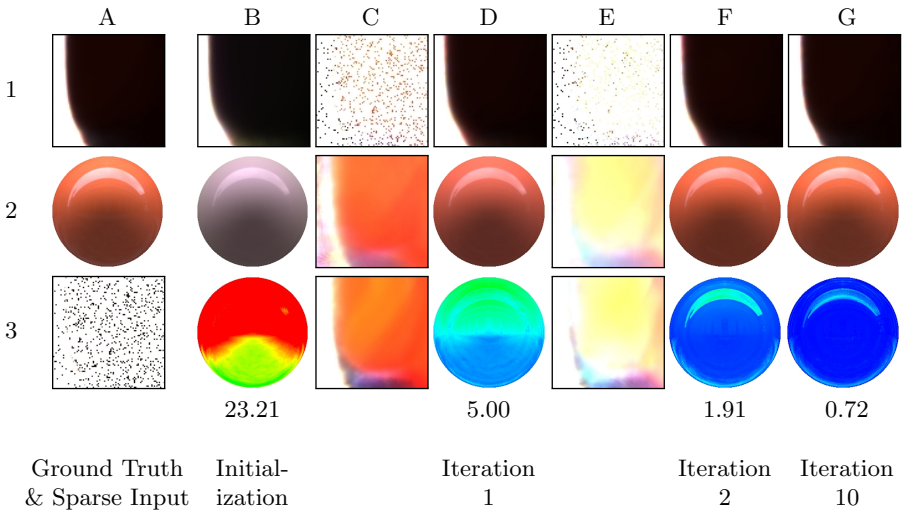


Fig. 4. Based on the sparse input (A3) our method is initialized by a linear combination of measured BRDFs (B1), the error of this initialization is shown in B2-3. Using this initialization we compute a sparse set of correction factors (C1) for which we estimate a dense correction function (C2). Note the similarity to the ground truth correction function (C3) even in the significant presence of outliers. Column D shows the corrected initial estimate, the error is already drastically reduced (D2-3) (error scale as in Figure 3). Columns E-F correspond to C-D for the second iteration. The last column shows the estimated solution after 10 iterations with almost no perceptual difference

the chosen error metric. The various steps that our algorithm executes for this example are illustrated and explained in Figure 4. Already after a single iteration we can achieve a solution that is significantly better than those of all other evaluated methods. After a few more iterations our estimate is perceptually almost indistinguishable from the original BRDF.

3.1 Benchmark

To give a representative evaluation how our algorithm performs when compared to other state-of-the-art methods we performed a large number of experiments and compared the results. The material database in [19] provides 100 densely measured materials which served as a ground truth for our experiments. From each material we generated several sparse input sets with different characteristics. These differed in the data ratio (i.e. the number of randomly chosen samples from the ground truth BRDF table) and the outlier ratio (i.e. the probability for each sample to have consistent information). We assigned to each outlier a randomly chosen reflectance from the ground truth BRDF table (which is actually a more difficult problem than assigning completely random outlier values).

Similar to the previously given example we evaluated representative methods from each solution type in the literature: For global basis decomposition we chose

Table 1. Mean CIELAB [5] error of all evaluated methods under varying outlier and data ratio. To give a visual impression of these numerical values the table cells are colored using the same colors as for the error visualizations in Figure 3 and 4

		Data ratio							Data ratio				
Our		1.0	0.7	0.5	0.3	0.1	LC		1.0	0.7	0.5	0.3	0.1
Outlier ratio	0.00	0.28	0.29	0.29	0.29	0.31	Outlier ratio	0.00	4.88	4.89	4.91	4.96	5.05
	0.20	0.36	0.37	0.44	0.41	0.49		0.20	6.32	6.58	6.76	7.57	7.59
	0.40	0.73	0.72	0.75	0.78	0.92		0.40	7.97	8.22	7.93	8.32	9.03
	0.60	2.06	2.10	2.05	2.23	2.31		0.60	9.64	9.74	9.46	9.59	9.53
	0.80	4.86	4.89	4.83	4.81	4.83		0.80	11.33	11.16	11.05	11.05	10.61
RBF		1.0	0.7	0.5	0.3	0.1	PP		1.0	0.7	0.5	0.3	0.1
Outlier ratio	0.00	4.11	4.11	4.13	4.15	4.17	Outlier ratio	0.00	0.00	0.09	0.20	0.46	2.14
	0.20	6.85	6.78	6.96	7.12	7.74		0.20	17.32	15.59	17.28	14.00	12.30
	0.40	9.74	9.98	9.79	10.41	11.04		0.40	26.12	26.38	24.80	24.35	22.90
	0.60	12.49	12.63	12.49	12.56	13.45		0.60	33.32	34.90	32.71	32.13	28.22
	0.80	15.18	15.32	15.28	15.16	15.92		0.80	39.86	41.70	40.75	38.94	34.09

a linear combination (LC) of measured materials (as proposed in [20, 31, 26, 1]) with a logarithmic error metric. For local basis decomposition we chose radial basis functions (RBF) as proposed in [33, 31] with a logarithmic error metric. We distributed 758 RBF centers for each color channel as proposed in [31] on a radially re-parametrized BRDF domain which is more suitable for RBF interpolation [33]. For a tabular representation we used the *Pull-Push* (PP) method [10] as proposed in [20, 17] for the scattered data interpolation. Finally, we initialized our method with the LC solution and performed 10 iterations. In the outlier free case we set the parameter $\gamma = 0$, otherwise to 3 (for an outlier ratio of 0.2) or to 6 (for larger outlier ratios). For LC and our algorithm we removed the current material from the underlying database before the methods were executed. All steps of our algorithm (including the estimation of the correction basis) were therefore executed without any remaining traces of the current material.

As demonstrated in [8], surface reflectance properties are clearer and better comparable when objects are viewed under real-world illuminations. For each estimated solution we thus synthesized a sphere of this material under a natural illumination given by a HDR environment map (*Grace Cathedral*, courtesy of Paul Debevec, see supplementary material). We compared each rendered sphere with the ground truth sphere in the perceptually uniform CIELAB [5] space.

The mean CIELAB [5] error values for all experiments are given in Table 1. The evaluation confirms the observations sketched in the previously given example: Compared to the local RBF method the global LC method is more robust if the sparse input contains many outliers. However, the local RBF method has an overall larger flexibility and thereby allows to better approximate the data (for small outlier ratios). The tabular PP method performs well in the outlier free case, however its quality decreases drastically even for very small outlier ratios, since outliers are also interpolated within the table.

Table 2. Mean CIELAB [5] error for different error metric. In contrast to LC or RBF our method performs robust for different choices of error metrics

Our			LC			RBF		
ε_{lin}	$\varepsilon_{\text{root}}$	ε_{log}	ε_{lin}	$\varepsilon_{\text{root}}$	ε_{log}	ε_{lin}	$\varepsilon_{\text{root}}$	ε_{log}
0.44	0.31	0.28	13.39	5.05	4.88	24.14	3.32	4.11

The experiments indicate that our method performs significantly better than all other evaluated method types: In the outlier free case the proposed method has a precision comparable to the tabular PP method. As less input data is available we can even outperform this purely tabular representation. We expect this is because our method implicitly uses physically meaningful BRDF domain knowledge by its underlying measured material and correction database. If compared to the LC and RBF method our approach has a significantly lower error value in all experiments (an order of magnitude lower perceptual error for outlier ratios up to 40%). Additionally, our method is almost invariant to increasing sparsity and behaves robust for increased outlier ratios.

We also evaluated how the individual methods based on basis decomposition performed for different error metrics. The choice of the error metric can have a significant impact on the estimation quality. As also pointed out in [21] it is difficult in practice to make a single choice that consistently performs well for all materials. To quantify which effect this choice can have, we compared the results of all methods using different error metrics. For this we chose the following three metrics that were proposed by researchers before: linear $\varepsilon_{\text{lin}}(x) = x$, square root $\varepsilon_{\text{root}}(x) = \sqrt{x}$ and logarithmic $\varepsilon_{\text{log}}(x) = \ln(1 + x)$. These metrics were then used for LC, RBF and in the initialization stage of our algorithm. We only considered the case of full data and zero outlier ratio (which allows a more consistent comparison without any randomness in the generated input). The results of this experiment are shown in Table 2. It can be seen that our algorithm performs robust with respect to different error metrics, whereas the outcome of the other basis decomposition methods is highly dependent on this choice. This indicates that *any* meaningful metric can be used along with our algorithm which alleviates the choice of an appropriate one. This choice then only affects the iteration sequence but the overall method converges to the same solution.

Regarding the runtime our algorithm operates highly efficient, requiring only a few seconds per iteration on our hardware (3 GHz CPU, 24 GB RAM).

We additionally evaluated our method based on 16 newly measured materials (see Figure 5). This is important because real measured material samples can show different characteristics than our synthetically generated experiments. We found that our method allows to represent the measured data significantly better (in the numerical sense and in terms of physical plausibility). Summarized, for these newly measured materials our algorithm achieves a lower mean perceptual error (error of 0.19) compared to LC (error of 0.50) and RBF (error of 0.38). Please see the supplementary material for a detailed report.




Fig. 5. Exemplary fitting result for 1 of the 16 newly measured materials. Left: Raw measured isotropic BRDF data (a 2D projection of the 3D BRDF table for visualization). Red corresponds to unmeasured areas. Right: Synthesized spheres using the computed estimates under point light illumination (numerical values indicate the fitting error). In contrast to our method the different state-of-the-art algorithms compute either physically implausible results (PP, RBF) or non-accurate solutions (LC)

4 Conclusion

We introduced a novel, non-parametric method that can be used for accurate and robust reflectance data fitting. It is based on the novel concept of BRDF correction functions that are iteratively applied to improve an initial BRDF estimate. Moreover, we introduced a new basis that allows to efficiently represent these correction functions. During execution the resulting algorithm explicitly identifies and excludes outliers in the input data. This is challenging (in particular for *non-parametric* methods) and was not handled by any previously proposed method that we are aware of.

We evaluated and compared our method using a large number of experiments and showed that it performs very robust with respect to sparse and outlier afflicted input – in contrast to other evaluated methods. We further showed that our method is invariant to different error metrics. This alleviates the choice of an appropriate error metric for a given input which was challenging until now.

The main limitation of our method is its current exclusive applicability to isotropic materials. The only reason for this is that there is currently no large enough database of measured anisotropic materials available. We have successfully tested our method in the isotropic (3D) case and we see no particular conceptual limitation why it should not be adaptable for the general anisotropic (4D) case once such data is available. It may be possible to synthesize such a database by using a variety of physically based anisotropic BRDF models with varying parameters. However, we have not verified this yet and it is subject to future work.

5 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. Ashikhmin, M., Shirley, P.: An anisotropic phong brdf model. *J. Graph. Tools* 5(2), 25–32 (Feb 2000), <http://dx.doi.org/10.1080/10867651.2000.10487522>
3. Blinn, J.F.: Models of light reflection for computer synthesized pictures. In: Proceedings of the 4th Annual Conference on Computer Graphics and Interactive Techniques. pp. 192–198. SIGGRAPH '77, ACM, New York, NY, USA (1977), <http://doi.acm.org/10.1145/563858.563893>
4. Cook, R.L., Torrance, K.E.: A reflectance model for computer graphics. *ACM Trans. Graph.* 1(1), 7–24 (Jan 1982), <http://doi.acm.org/10.1145/357290.357293>
5. Fairchild, M.D.: Color appearance models. John Wiley & Sons (2013)
6. Ferreau, H.J.: qpOASES Library for Online Active Set Strategy. <http://set.kuleuven.be/optec/Software/qpOASES-OPTEC/>, accessed: 2014-02-13
7. Ferreau, H.J., Bock, H.G., Diehl, M.: An online active set strategy to overcome the limitations of explicit mpc. *International Journal of Robust and Nonlinear Control* 18(8), 816–830 (2008), <http://dx.doi.org/10.1002/rnc.1251>
8. Fleming, R.W., Dror, R.O., Adelson, E.H.: Real-world illumination and the perception of surface reflectance properties. *Journal of Vision* 3(5), 3 (2003)
9. Fores, A., Ferwerda, J., Gu, J.: Toward a perceptually based metric for brdf modeling. In: Twentieth Color and Imaging Conference. Los Angeles, California, USA. pp. 142–148 (November 2012)
10. Gortler, S.J., Grzeszczuk, R., Szeliski, R., Cohen, M.F.: The lumigraph. In: Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques. pp. 43–54. SIGGRAPH '96, ACM, New York, NY, USA (1996), <http://doi.acm.org/10.1145/237170.237200>
11. He, X.D., Torrance, K.E., Sillion, F.X., Greenberg, D.P.: A comprehensive physical model for light reflection. In: Proceedings of the 18th Annual Conference on Computer Graphics and Interactive Techniques. pp. 175–186. SIGGRAPH '91, ACM, New York, NY, USA (1991), <http://doi.acm.org/10.1145/122718.122738>
12. Holroyd, M., Lawrence, J., Zickler, T.: A coaxial optical scanner for synchronous acquisition of 3d geometry and surface reflectance. *ACM Trans. Graph.* 29, 99:1–99:12 (July 2010)
13. Koenderink, J.J., Doorn, A.J.V.: Phenomenological description of bidirectional surface reflection. *JOSA A* 15, 2903–2912 (1998)
14. Köhler, J., Nöll, T., Reis, G., Stricker, D.: A full-spherical device for simultaneous geometry and reflectance acquisition. In: Applications of Computer Vision (WACV), 2013 IEEE Workshop on. pp. 355–362 (2013)
15. Lafortune, E.P., Foo, S.C., Torrance, K.E., Greenberg, D.P.: Non-linear approximation of reflectance functions. In: SIGGRAPH. pp. 117–126 (1997), <http://dblp.uni-trier.de/db/conf/siggraph/siggraph1997.html#LafortuneFTG97>
16. Lalonde, P., Fournier, A.: A wavelet representation of reflectance functions. *IEEE Transactions on Visualization and Computer Graphics* 3(4), 329–336 (Oct 1997), <http://dx.doi.org/10.1109/2945.646236>
17. 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>
18. Marschner, S.R., Westin, S.H., Lafortune, E.P.F., Torrance, K.E.: Image-based bidirectional reflectance distribution function measurement. *Applied Optics* 39, 2592–2600 (2000)
 19. Matusik, W., Pfister, H., Brand, M., McMillan, L.: A data-driven reflectance model. In: *ACM SIGGRAPH 2003 Papers*. pp. 759–769. SIGGRAPH '03, ACM, New York, NY, USA (2003), <http://doi.acm.org/10.1145/1201775.882343>
 20. 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>
 21. Ngan, A., Durand, F., Matusik, W.: Experimental analysis of brdf models. In: *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques*. pp. 117–126. EGSR'05, Eurographics Association, Aire-la-Ville, Switzerland (2005), <http://dx.doi.org/10.2312/EGWR/EGSR05/117-126>
 22. Nicodemus, F.E.: Directional reflectance and emissivity of an opaque surface. *Appl. Opt.* 4(7), 767–775 (Jul 1965), <http://ao.osa.org/abstract.cfm?URI=ao-4-7-767>
 23. Nishino, K.: Directional statistics brdf model. In: *Computer Vision, 2009 IEEE 12th International Conference on*. pp. 476–483 (Sept 2009)
 24. Pacanowski, R., Salazar-Celis, O., Schlick, C., Granier, X., Pierre, P., Annie, C.: Rational BRDF. *IEEE Transactions on Visualization and Computer Graphics* 18(11), 1824–1835 (Feb 2012), <http://hal.inria.fr/hal-00678885>
 25. Ramamoorthi, R., Hanrahan, P.: A signal-processing framework for inverse rendering. In: *Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*. pp. 117–128. SIGGRAPH '01, ACM, New York, NY, USA (2001), <http://doi.acm.org/10.1145/383259.383271>
 26. 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>
 27. Rusinkiewicz, S.: A new change of variables for efficient brdf representation. In: *Rendering Techniques*. pp. 11–22 (1998)
 28. Schlick, C.: An inexpensive brdf model for physically-based rendering. *Comput. Graph. Forum* 13(3), 233–246 (1994)
 29. Schröder, P., Sweldens, W.: Spherical wavelets: Efficiently representing functions on the sphere. In: *Proceedings of the 22Nd Annual Conference on Computer Graphics and Interactive Techniques*. pp. 161–172. SIGGRAPH '95, ACM, New York, NY, USA (1995), <http://doi.acm.org/10.1145/218380.218439>
 30. Ward, G.J.: Measuring and modeling anisotropic reflection. In: *Proceedings of the 19th Annual Conference on Computer Graphics and Interactive Techniques*. pp. 265–272. SIGGRAPH '92, ACM, New York, NY, USA (1992), <http://doi.acm.org/10.1145/133994.134078>
 31. 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>
 32. Westin, S.H., Arvo, J.R., Torrance, K.E.: Predicting reflectance functions from complex surfaces. In: *Proceedings of the 19th Annual Conference on Computer*

- Graphics and Interactive Techniques. pp. 255–264. SIGGRAPH '92, ACM, New York, NY, USA (1992), <http://doi.acm.org/10.1145/133994.134075>
33. 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>