Global Optimization for Estimating a BRDF with Multiple Specular Lobes*

Chanki Yu, Yongduek Seo and Sang Wook Lee
Dept. of Media Technology, Sogang University, Seoul, Korea
{ckyu, yndk, slee}@sogang.ac.kr

Abstract

This paper presents a global minimization framework for estimating analytical BRDF model parameters using the techniques of convex programming and branch and bound. Traditional local minimization suffers from local minima and requires a large number of initial conditions and supervision for successful results especially when a model is highly complex and nonlinear. We consider the Cook-Torrance model, a parametric model with the Gaussian-like Beckmann distributions for specular reflectances. Instead of optimizing the multiple parameters simultaneously, we search over all possible surface roughness values based on a branch-and-bound algorithm, and reduce the estimation problem to convex minimization with known fixed surface roughness. Our algorithm guarantees globally optimal solutions. Experiments have been carried out for isotropic surfaces to validate the method using the extensive high-precision measurements from the MERL BRDF database.

1. Introduction

Accurate descriptions of surface reflection have been a topic of research in computer vision and computer graphics. The Bidirectional Reflectance Distribution Function BRDF) is a four-dimensional function that defines how light is reflected at a surface, and a number of BRDFs have been developed to model real-world material surfaces [18].

Analytical models have been popular for their compact representations, and some of those are developed for physical plausibility [1] [4] [5] [10] [16] [19] [23] [24]. Since these models generally do not account for all the reflectance properties of all kinds of materials, approaches to representing reflectances on basis functions have recently been developed to describe reflectances from a wider range of material surfaces [2] [20] [21]. Despite all their comprehensiveness, however, the number of bases needs to be kept large to account for viewing and lighting variability and to maintain high frequency details. Once the analytical model is appropriately selected for a specific type of surface material and its parameters are accurately estimated, the model describes the reflection in a highly compact way.

For estimating a BRDF, reflections are measured under various viewing and illumination angles, and data is usually fitted to an analytical model using conventional (constrained) least-squares nonlinear minimization [13] [17] [24]. Highly nonlinear BRDFs that include multiple Gaussian-like functions can have a huge number of local minima. In traditional nonlinear optimization, the quality of the fit is dependent on a good initial guess. To make sure the optimization converges to a local minimum yielding a satisfactory result or hopefully the global minimum, the fitting quality of the result is visually/manually inspected and if necessary the optimization is restarted from a different set of initial guesses. However, even this supervised optimization does not necessarily guarantee the globally minimum solution.

While global optimization has recently been an area of active research in geometric vision [9] [11] [12], no explicit effort has been made in photometric vision. To our knowledge, no previous work has given globally optimal solutions to the BRDF estimation problems under photometrically meaningful L_∞ or L_2 cost functions. Our work focuses on the globally optimal parameter estimation of the Cook-Torrance model with multiple specular lobes using a set of photometric measurements with known surface orientations and viewing/lighting directions [5]. The Cook-Torrance model has been developed based on the geometrical optics and considered one of the most physically plausible models. Nonlinearity arises from the surface roughness parameter in the Beckmann distribution function, and the model becomes quite complex when multiple specular lobes are considered.

Our idea is based on the observation that the estimation problem becomes convex when the surface roughness parameter is known, in which case the global optimum can be found easily. Therefore, searching for the best

* This work was supported by the strategic technology development program of MCST/MKE/KEIT. [2008-F-030-01, Development of Full 3D Reconstruction Technology for Broadcasting Communication Fusion]

978-1-4244-6983-3/10/$26.00 ©2010 IEEE

roughness is the key procedure for the minimization. We employ a branch-and-bound algorithm for an efficient search over the whole interval of the roughness. It recursively bisects the interval into small sub-intervals, and for each of them it solves a convex feasibility problem which tests whether or not it may contain a better roughness inside. The algorithm is efficient because infeasible sub-intervals are discarded from the search space and not considered any more. Furthermore, our algorithm guarantees the global optimality. In this paper, we provide a mathematical analysis to derive the convex feasibility problem.

Recent advances in data acquisition technique and apparatus have resulted in several high quality BRDF databases available for research and academic use [6] [7] [8] [14] [15]. To validate our algorithm, we carried out experiments using the isotropic material data in the MERL database, a set of extensive, densely sampled, and high-precision HDR measurements [14] [15]. To demonstrate the limitations of the conventional approach, we also performed experiments with a local optimization algorithm and the results are compared with those with the presented method.

The rest of this paper is organized as follows. Section 2 describes the BRDF model that our work is based on. Section 3 states the problems for us to solve and Section 4 describes our brand and bound algorithm. Experimental results and discussions are presented in Section 5 and 6, respectively, and we conclude in Section 7.

2. BRDF Model

The Cook-Torrance is given as:

$$\frac{\rho}{\pi}(\boldsymbol{N}\cdot\boldsymbol{L})+\frac{F}{\pi}\frac{D}{(\boldsymbol{N}\cdot\boldsymbol{L})}\frac{G}{(\boldsymbol{N}\cdot\boldsymbol{V})}.$$

The diffuse reflectance is assumed to be Lambertian and ρ is the diffuse albedo. The Fresnel coefficient F is dependent on the light incident angle and the refractive index of the material which may vary along the wavelength λ. The unit vectors $\boldsymbol{N}$, $\boldsymbol{L}$ and $\boldsymbol{V}$ denote the surface normal, the illumination direction and the viewing direction, respectively. (See Figure 1 for the local geometry of reflection.)

The geometrical attenuation factor G accounts for the shadowing and masking of microfacets and given as:

$$G=\boldsymbol{min}\left[1,\ \frac{2(\boldsymbol{N}\cdot\boldsymbol{H})(\boldsymbol{N}\cdot\boldsymbol{V})}{(\boldsymbol{V}\cdot\boldsymbol{H})},\ \frac{2(\boldsymbol{N}\cdot\boldsymbol{H})(\boldsymbol{N}\cdot\boldsymbol{L})}{(\boldsymbol{V}\cdot\boldsymbol{H})}\right],$$

where the unit vector $\boldsymbol{H}$ is the bisector of $\boldsymbol{L}$ and $\boldsymbol{V}$. The facet slope distribution function D represents the fraction of the facets that are oriented in the direction $\boldsymbol{H}$. Various facet

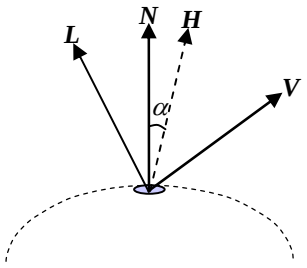

Figure 1: Local geometry of surface reflection

slope distribution functions have been considered by many investigators [5] [23] [24]. For rough surfaces, Cook and Torrance used the Beckmann distribution function [3] [5]:

$$D=\frac{1}{\sigma^2\cos^4\alpha}\exp\left(-\frac{\tan^2\alpha}{\sigma^2}\right),$$

where α is the angle between $\boldsymbol{N}$ and $\boldsymbol{H}$. The parameter σ is the root mean square of the microfacets and represents the surface roughness. Some surfaces have two or more scales of roughness, and can be modeled by using more than one distribution functions. In such cases, D is a weighted sum of the distribution functions, i.e.,

$$D=\sum_p w_p D(\sigma_p),$$

where σ_j is the surface roughness of the jth distribution and w_j is the weight of the jth distribution. The sum of the weights is 1 [5].

It should be noted that the distribution function D was developed for rough surfaces. When it is used for a smooth surface, its physical plausibility may diminish to some degree and the model serves merely as an empirical model. For our investigation presented in this paper, however, we use the same distribution function for all the surfaces in the MERL database as did in [17] where the one-lobe and two-lobe Cook-Torrance models show good performance compared to others.

As mentioned above, the Fresnel factor $F \in [0.1]$ is a complex function of the refractive index and light incident angle. Since it is almost constant for the incidence up to about 70 degrees, however, we excluded the BRDF data beyond 60 degrees of light incidence in our experiments and F is considered as a constant.

3. Problem Statement

We are interested in the classic problem of estimating parameters of an analytical BRDF model, in particular of the Cook-Torrance model, with known object geometry and photometric calibration. The main problem is given as follows.

Problem 1 (BRDF estimation) *Given an object of known shape, we want to estimate the parameters ρ, F, and σ of the Cook-Torrance model from a set of image-based photometric measurements with known viewing and illumination directions.*

With known object shape, we have $\boldsymbol{N}$, and with photometric calibration, we know $\boldsymbol{V}$, $\boldsymbol{L}$, $\boldsymbol{H}$ and α. The parameters we like to estimate are ρ, F and σ. The parametric reprojection equation of the one-lobe Cook-Torrance model for the ith measurement is given as:

$$\hat{I}_i(x,y,\sigma) = a_i x + b_i y \frac{1}{\sigma^2}\boldsymbol{exp}\left(-\frac{c_i}{\sigma^2}\right), \quad x \geq 0, y \geq 0, \sigma \geq 0.$$

where I_s is the incident illumination intensity at the surface, and

$$x = \rho I_s, \quad y = F I_s,$$
$$a_i = \frac{(\boldsymbol{N}\cdot\boldsymbol{L}_i)}{\pi}, \; b_i = \frac{G_i}{\pi(\boldsymbol{N}\cdot\boldsymbol{V}_i)\boldsymbol{cos}^4\alpha_i}, \; c_i = \boldsymbol{tan}^2\alpha_i .$$

For a distant point light source, I_s is constant but unknown. Thus, we actually estimate ρ and F up to scale, i.e.:

$$x = k\rho, \quad y = kF \quad \text{where } k = I_s.$$

Now **Problem 1** is to estimate x, y, σ with known a_is, b_is, and measured I_i s for i =1, 2,…, N.

The residual function for a measurement I_i is defined as:

$$\begin{aligned} e_i(x,y,\sigma) &= \left|\hat{I}_i(x,y,\sigma) - I_i\right| \\ &= \left|a_i x + b_i y \frac{1}{\sigma^2}\boldsymbol{exp}\left(-\frac{c_i}{\sigma^2}\right) - I_i\right|. \end{aligned} \tag{1}$$

We use the following vector notations for a set of errors, measurements and reprojections:

$$\boldsymbol{e} = \begin{bmatrix} e_1 & e_2 & \cdots & e_N \end{bmatrix}^T$$
$$\boldsymbol{I} = \begin{bmatrix} I_1 & I_2 & \cdots & I_N \end{bmatrix}^T$$
$$\hat{\boldsymbol{I}}(\sigma) = \begin{bmatrix} \hat{I}_1(\sigma) & \hat{I}_2(\sigma) & \cdots & \hat{I}_N(\sigma) \end{bmatrix}^T.$$

Note that $\hat{\boldsymbol{I}}$ is represented as a function of σ among other parameters. This will be used later in developing our branch-and-bound algorithm.

Given a set of measurements I_is, i =1, …, N, we want to find the best solution that minimizes the L_2 norm of residuals:

$$\begin{aligned} \min_{(x,y,\sigma)} \|\boldsymbol{e}\|_2 &= \min_{(x,y,\sigma)} \left(\sum_i e_i^2\right)^{\frac{1}{2}} \\ &= \min_{(x,y,\sigma)} \left(\sum_i \left(I_i - \hat{I}_i(x,y,\sigma)\right)^2\right)^{\frac{1}{2}}. \end{aligned}$$

The difficulty of finding the optimal solution results from the nonlinearity caused by the exponential term. However, one may easily observe that if we know σ, then the estimation becomes a convex optimization problem and we can find the global solution to this reduced problem. Let us suppose that we know $\sigma = \sigma_0$. Then, the estimation of x and y becomes a second-order cone programming problem:

Problem 2 (*Restricted problem for BRDF estimation*)

$$\begin{aligned} &\textit{minimize} \;\; t \\ &\textit{subject to} \;\; \left\|\boldsymbol{I} - \hat{\boldsymbol{I}}(x,y,\sigma_0)\right\|_2 \leq t, \\ &\qquad x \geq 0, \;\; y \geq 0, \end{aligned} \tag{2}$$

Here the number of variables is three including the auxiliary variable t and the number of constraints is three (one second-order constraint and two linear constraints). Note that this is a convex optimization problem that can easily be solved using a convex solver such as SeDuMi [22]. If this convex problem is solved for every value of the roughness, then the global solution can be obtained. It is of course impossible in practice to deal with an infinite number of surface roughness values. Our solution to overcome this impossibility is to efficiently search over the σ space using a branch-and-bound technique and take advantage of the convexity when σ is known. Below we outline our branch-and-bound strategy.

Starting with any σ_0 that can be found by any method at all, our algorithm computes the L_2 residual $\varepsilon_{\min}$, and takes it as the initially best residual. Then, the interval of σ domain is divided up into two sub-intervals. On each of the sub-intervals we determine whether there is a solution to the restricted optimization problem having cost less than $\varepsilon_{\min}$. This question is formulated as a *feasibility problem*. If the answer is negative on a sub-interval, it is excluded from further consideration. Otherwise, the algorithm evaluates the cost function for some σ value inside the sub-interval, and if this is less than $\varepsilon_{\min}$, it updates $\varepsilon_{\min}$, x and y. Then this feasible sub-interval is bisected into two smaller regions. This procedure of feasibility check and bisection is repeated until the length of sub-domains is short enough.

The coarse-to-fine sub-division and discarding infeasible sub-intervals result in an efficient search over the σ space. Next section presents a mathematical derivation of the feasibility problem which is the core of our BRDF estimation algorithm based on the branch-and-bound technique.

4. Branch-and-Bound BRDF estimation

1. An appropriate value σ_0 is determined and then the other parameters are computed. Then we have $\varepsilon_{\min}$, the residual corresponding to σ_0:

$$\varepsilon_{\min} = \left\| \boldsymbol{e}(\sigma_0) \right\|_2 .$$

2. The domain $\boldsymbol{S}$ of σ is subdivided into several sub-intervals S_j s in such a way that $\boldsymbol{S} = \cup_j S_j$ and $S_i \cap S_j = \varnothing$ for $i \neq j$, except at the boundary. This is an initialization step. In the following steps, feasible sub-domains are bisected.

Any sub-interval S is represented by its center position $\bar{\sigma}$ and the half length h of the interval. The lower bound and upper bound of the interval are denoted by σ_l and σ_u, respectively:

$$S = [\sigma_l \quad \sigma_u], \quad \sigma_l = \bar{\sigma} - h, \ \sigma_u = \bar{\sigma} + h$$

3. A convex optimization problem, called the *feasibility* problem, is solved to check whether the interval S_j contains a better solution than the current estimate.

4. Discard all the infeasible sub-intervals.

5. The restricted convex optimization problem (Equation 2) is solved for each of the feasible sub-intervals, and the best estimate is updated.

6. The feasible sub-intervals are bisected; Go to step 3.

We stop this branch-and-bound iteration when the half length h of the sub-intervals is small enough. As we show below, the feasibility of a sub-domain is dependent on two factors: the current best residual and the length of the domain $2h$.

Now let us develop the feasibility problem. Plainly, it is cast as follows:

Problem 3 *Do there exist x, y and* $\sigma \in S$ *such that*

$$\left\| \boldsymbol{e} \right\|_2 \ \leq \ \varepsilon_{\min}.$$

Unfortunately, it is not easy to give an answer to this question. Instead, we want to consider an alternative but equivalent problem [9]:

Problem 4 *Do there exist x and y such that*

$$\left\| \boldsymbol{e} \right\|_2 < \varepsilon_{\min} + \varepsilon_S \ , \qquad (3)$$

where ε_S *is a bound due to the variation of* σ *in S.*

Now let us derive a form of the bound ε_S. First, we consider the case when the sub-interval S contains a roughness that yields a smaller residual. Let σ_{opt} be one of the roughness values in S. Then, the following inequalities hold:

$$\begin{aligned} \left\| \boldsymbol{e}(\bar{\sigma}) \right\| &= \left\| \boldsymbol{I} - \hat{\boldsymbol{I}}(\bar{\sigma}) \right\| \\ &\leq \left\| \boldsymbol{I} - \hat{\boldsymbol{I}}(\sigma_{\text{opt}}) \right\| + \left\| \hat{\boldsymbol{I}}(\sigma_{\text{opt}}) - \hat{\boldsymbol{I}}(\bar{\sigma}) \right\| \\ &\leq \varepsilon_{\min} + \left\| \hat{\boldsymbol{I}}(\sigma_{\text{opt}}) - \hat{\boldsymbol{I}}(\bar{\sigma}) \right\|. \end{aligned}$$

The first inequality is due to the triangle inequality. Let us examine the second term on the right hand side as a function of σ_{opt}:

$$\begin{aligned} &\left\| \hat{\boldsymbol{I}}(\sigma_{\text{opt}}) - \hat{\boldsymbol{I}}(\bar{\sigma}) \right\|_2 \\ &= \left(\sum_i \left(yb_i \frac{1}{\sigma_{\text{opt}}^2} \boldsymbol{exp}\left\{ -\frac{c_i}{\sigma_{\text{opt}}^2} \right\} - yb_i \frac{1}{\bar{\sigma}^2} \boldsymbol{exp}\left\{ -\frac{c_i}{\bar{\sigma}^2} \right\} \right)^2 \right)^{\frac{1}{2}}. \end{aligned}$$

We may find an upper bound of this norm function by replacing σ_{opt} by one of the two boundary points of S due to the characteristic of this function. By finding maximizers $\sigma_{i,*}$ (either $\sigma_{i,u}$ or $\sigma_{i,l}$) of this function, we obtain an inequality:

$$\left\| \boldsymbol{I} - \hat{\boldsymbol{I}}(\bar{\sigma}) \right\| < \varepsilon_{\min} + y\varepsilon_\sigma \ ,$$

where

$$\varepsilon_\sigma = \left(\sum_i \left(b_i \frac{1}{\sigma_{i,*}^2} \boldsymbol{exp}\left\{ -\frac{c_i}{\sigma_{i,*}^2} \right\} - b_i \frac{1}{\bar{\sigma}^2} \boldsymbol{exp}\left\{ -\frac{c_i}{\bar{\sigma}^2} \right\} \right)^2 \right)^{\frac{1}{2}}.$$

It should be noted that the function

$$\frac{1}{\sigma^2} \boldsymbol{exp}\left\{ -\frac{c_i}{\sigma^2} \right\}$$

monotonically increases for $0 < \sigma < \boldsymbol{tan}\alpha$ and monotonically decreases for $\boldsymbol{tan}\alpha < \sigma$. Thus we make sure that no sub-interval S is set up across $\boldsymbol{tan}\alpha$.

We recast the feasibility problem (3) as follows: Find x and y subject to**:**

$$\left\| \boldsymbol{I} - \hat{\boldsymbol{I}}(\bar{\sigma}) \right\| < \varepsilon_{\boldsymbol{min}} + y\varepsilon_\sigma, \qquad x \geq 0, \ y \geq 0,$$

and if there is no such solution, report it.

Note that this is a second-order cone problem whose parametric formulation is given as:

Problem 5 (Feasibility problem)

$$\begin{aligned} &\textit{minimize} \quad s \\ &\textit{subject to} \quad \left\| \boldsymbol{I} - \hat{\boldsymbol{I}}(x, y, \bar{\sigma}) \right\|_2 < \varepsilon_{\boldsymbol{min}} + y\varepsilon_\sigma + s \\ &\qquad\qquad x \geq 0 \, , \ y \geq 0 \end{aligned}$$

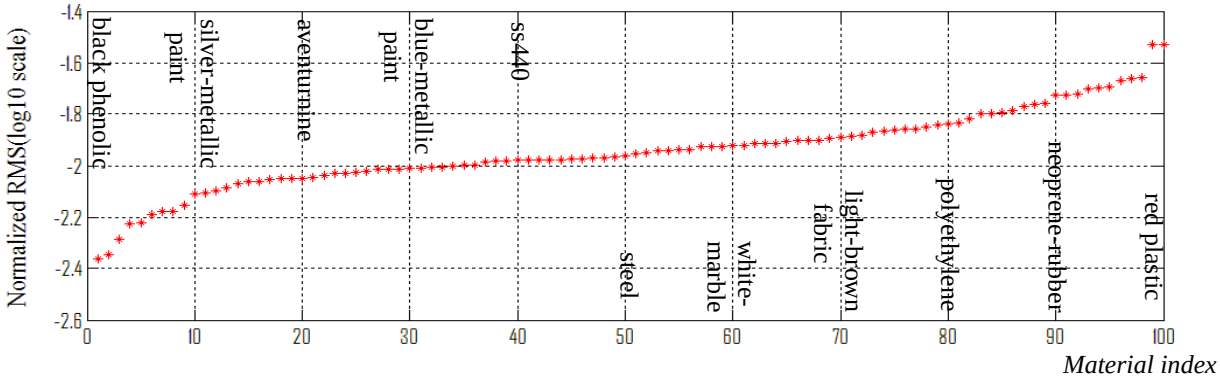

Figure 2 : The normalized *RMS* reprojection errors (logarithmic scale) from the estimated Cook-Torrance model for the MERL isotropic BRDF dataset of 100 materials.

Given a solution s of this problem, we have the following two cases:

a) If $s \leq 0$, then the problem is feasible. The sub-interval S contains a better (or at least equivalent) solution inside.

b) If $s > 0$, then the problem is infeasible. The sub-interval S can be excluded for further consideration.

When the problem is found to be feasible, we solve the Problem 2 and check whether $\bar{\sigma}$ yields a smaller L_2 error norm. If it does, then the total residual $\varepsilon_{\min}$ and the optimal solution $(\hat{x}_{\bar{\sigma}}, \hat{y}_{\bar{\sigma}})$ are updated.

Our branch-and-bound algorithm can be easily extended to the multiple-lobe Cook-Torrance model. For the two-lobe model, the residual function for the i-th measurement is given as:

$$e_i(x,y,\sigma) = \left| a_i x + \sum_{p=1}^{2} b_i y_p \frac{1}{\sigma_p^2} \exp\left(-\frac{c_i}{\sigma_p^2}\right) - I_i \right|,$$

where

$$x = k\rho, \quad y_1 = k w_1 F, \quad y_2 = k w_2 F, \quad k = I_s, \quad (w_1 + w_2 = 1)$$
$$a_i = \frac{(\mathbf{N}\cdot\mathbf{L}_i)}{\pi}, \quad b_i = \frac{G_i}{\pi(\mathbf{N}\cdot\mathbf{V}_i)\cos^4\alpha_i}, \quad c_i = \tan^2\alpha_i .$$

We have again a restricted convex second-order cone programming problem if we know the two roughness parameters σ_1 and σ_2. The branch-and-bound search is performed now over the two dimensional (σ_1, σ_2) space. From the triangle inequality for the residual function, we can derive the following two-lobe second-order cone programming problem:

$$\begin{aligned} &\text{minimize} \quad t \\ &\text{subject to} \quad \left\| \mathbf{I} - \hat{\mathbf{I}}(x,y,\bar{\sigma}) \right\|_2 < \varepsilon_{\min} + y_1\varepsilon_{\sigma_1} + y_2\varepsilon_{\sigma_2} + t, \\ &\qquad x \geq 0,\ y_1 \geq 0,\ y_2 \geq 0 \end{aligned}$$

5. Experimental Results

To validate our approach, we have carried out experiments on the MERL isotropic BRDF database. It contains reflections from a set of 100 objects that includes near diffuse objects (e.g. fabrics and paints), glossy objects (e.g. paints, metals and wood), mirror-like objects (e.g., metals and plastics). For a single common σ, we estimated the three sets of x and y parameters in each color channel independently. Therefore, the solutions are obtained for the seven parameters: x_r, x_g, x_b, y_r, y_g, y_b and σ.

One Specular Lobe

For the one-specular-lobe model, we performed branch and bound (BnB) on σ over the range of 10^{-12} ~6.0, and the shortest sub-interval after the final bisection was preset to 2^{-11}. This limits the number of bisectioning to 10. Many of the materials in the dataset can be described reasonably well with the one-specular-lobe model. Figure 2 shows a plot of *RMS* reprojection errors from the estimated one-specular-lobe model in logarithmic scale for all the 100 materials. Each of the *rgb* channels is normalized by its maximum value before their *RMS* reprojection errors are computed. The materials are sorted in ascending order of the *RMS* errors. We used the data with light incidence of up to 60 degrees to keep the Fresnel factor constant. For each object, a total of 84 sampled data are used (21 lighting direction for 4 views). The error curve shows the similar pattern to that shown in [17]. However, the different normalization factors do not allow us to compare the numerical values directly.

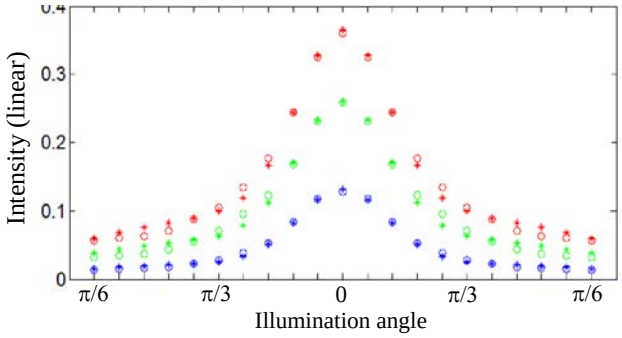

Figure 3: One-specular-lobe estimation: measurements and reprojections for the gold paint.

Figure 3 shows the measurements and reprojections for the gold paint which has the 33st smallest *RMS* error. The parameters are estimated using the set of 21 data with the viewing direction $(\theta, \phi) = (0, 0)$ and the illumination angles that range from (0, 0) to $(\pi/3, \pi)$. The circles denote the measurements and the asterisks show the estimated values in *rgb*. The notations for data plots for Figures 3, 4, 5, 6 and 7 are shown in Table 1.

O	*r* channel measured	*	*r* channel estimated
O	*g* channel measured	*	*g* channel estimated
O	*b* channel measured	*	*b* channel estimated

Table 1: Notations for data plots for Figures 3 ~ 7.

Two Specular Lobes

Many material surfaces cannot be well represented with only one specular lobe [17] [5] [13]. The BRDF estimation with the one-specular-lobe model using the data from the two-layer gold shows noticeable errors as can be seen in Figure 4 (logarithmic scale). The two-layer gold exhibits the 34st smallest *RMS* error (Figure 2). We have applied our two-lobe BnB algorithm that searches the two dimensional (σ_1, σ_2) space to the two-layer gold data, and the results are shown in Figure 5. The reprojection errors are substantially reduced.

We conducted experiments on the red plastic that exhibits the largest RMS errors. Figure 6 shows the measurements and reprojections estimated from the single lobe model for the four different views (linear scale). On the other hand, the results shown in Figure 7 are obtained using the two- specular-lobe BnB algorithm.

To confirm global optimality, we compared the results from the BnB algorithm with those from a two dimensional brute force search. For both σ_1 and σ_2, the smallest sampling interval is set to 2^{-11} and a brute force search is

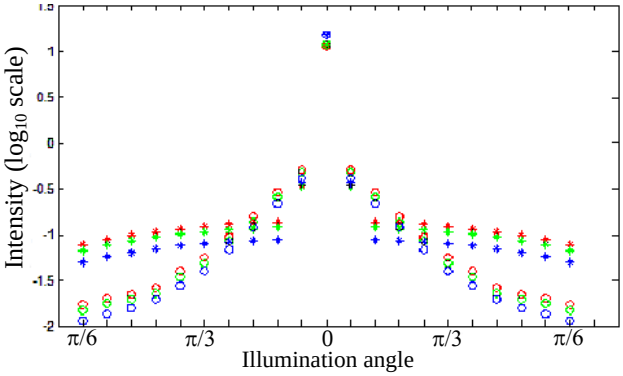

Figure 4: One-specular-lobe estimation: measurements and reprojections for the two-layer gold.

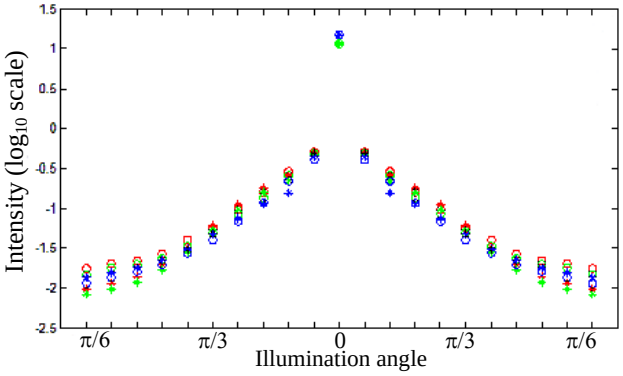

Figure 5: Two-specular-lobe estimation: measurements and reprojections for the two-layer gold.

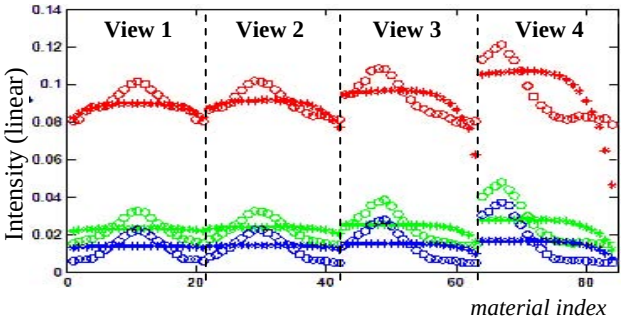

Figure 6: One-specular-lobe estimation: measurements and reprojections for the red plastic.

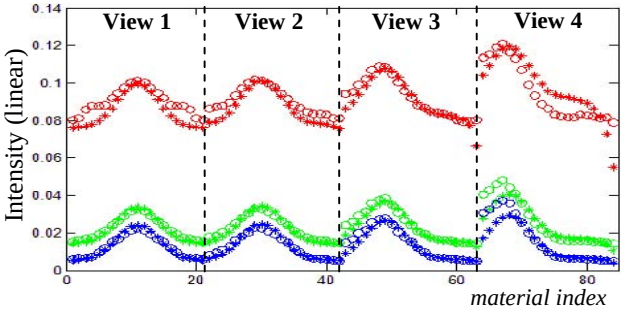

Figure 7: Two-specular-lobe estimation: measurements and reprojections for the red plastic.

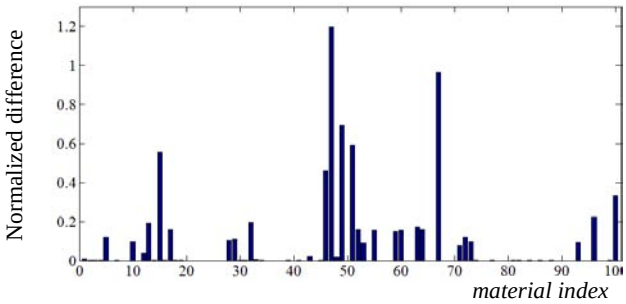

Figure 8 : Normalized differences between the residuals from the one-specular-lobe BnB and two-specular-lobe BnB algorithms. The materials are sorted in alphabetical order.

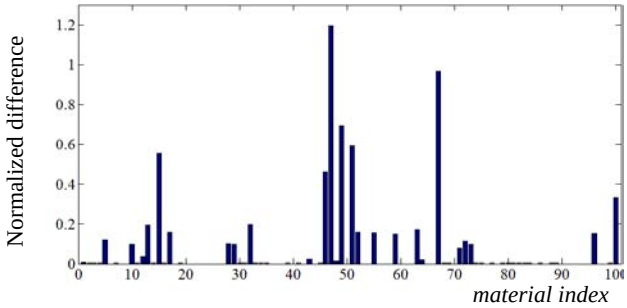

Figure 9: Normalized differences between the residuals from the two-specular-lobe L-M (100 trials) and BnB algorithms. The materials are sorted in alphabetical order.

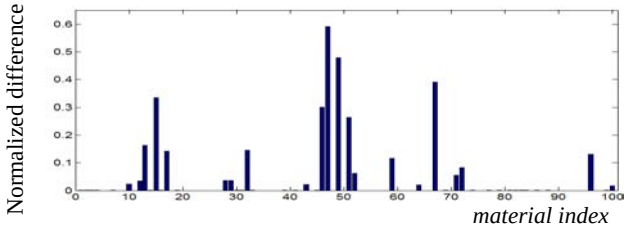

Figure 10: Normalized differences between the residuals from the two-specular-lobe L-M (5000 trials) and BnB algorithms. The materials are sorted in alphabetical order.

performed over the range of 10^{-12} ~6.0. The results are identical to those from the BnB algorithm.

To compare the model fitting errors, the difference between the reprojection errors $E_{BnB\text{-}1}$ (one-specular-lobe model) and $E_{BnB\text{-}2}$ (two-specular-lobe model) are normalized with respect to $E_{BnB\text{-}2}$, i.e.:

$$d_{BnB\text{-}1\text{-}2} = \frac{E_{BnB\text{-}1} - E_{BnB\text{-}2}}{E_{BnB\text{-}2}},$$

and the differences are shown for the 100 materials in Figure 8. The residuals are the same for 47 materials and the fitting quality with the two-specular-model improves for 53 materials.

***Local Optimization versus BnB* (*Two Specular Lobes*)**

For comparison, we carried out experiments on the same data set with conventional local optimization based on the Levenberg-Marquardt (L-M) algorithm. For each of the materials, the residuals $E_{BnB\text{-}2}$ from our BnB algorithm and $E_{L\text{-}M}$ from the L-M method are computed, and their difference is normalized with respect to $E_{BnB\text{-}2}$ as follows:

$$d_{L\text{-}M\text{-}BnB} = \frac{E_{L\text{-}M} - E_{BnB\text{-}2}}{E_{BnB\text{-}2}},$$

To compute $E_{L\text{-}M}$, we applied the L-M optimization a hundred times with different initial guesses for each of the materials, and selected the smallest residual as $E_{L\text{-}M}$. We found that data from some materials often fit the one-specular-lobe model better that the two-specular-lobe model with the local minimization approach, hence the residual from better fitting is taken as $E_{L\text{-}M}$, i.e.:

$$E_{L\text{-}M} = \boldsymbol{min}\{E_{L\text{-}M\text{-}1},\ E_{L\text{-}M\text{-}2}\},$$

where $E_{L\text{-}M\text{-}1}$ and $E_{L\text{-}M\text{-}2}$ denote the residuals from the one- and two-specular-lobe models, respectively. The local minimization was supervised such that the initial conditions are randomly chosen but unreasonably large or small values are excluded. Figure 9 shows the normalized residual differences. Even with the supervision, the L-M method with 100 trials found global optima only for 51 out of 100 materials.

We increased the number of initial guesses to 5000 to improve the results from the local minimization, and Figure 10 shows the normalized residual differences. Compared to the results from the one-specular-lobe model, the fitting quality degrades for 20 among 100 materials and improves for 52 materials with the two-specular-lobe model. It is the same for 28 materials with the one- and two-specular-lobe models. Even with 5000 trials, the L-M algorithm reaches the global optima only for 60 materials. We were not able to compare the computation time fairly since we implemented the L-M algorithm using a numerical tool based on C/C++ [25] and the BnB algorithm in Matlab [22]. However, the L-M method with 5000 trials takes a substantially longer time than the BnB method without any guarantee of global optimality.

6. Discussion

A common practice of BRDF fitting with a conventional local minimization algorithm is to run the algorithm with a predetermined number of trials and an error bound. Even with elaborate supervision and a huge number of initial guess, however, it is not always possible to find globally optimal solutions to the problem of fitting a BRDF model with multiple specular lobes. The experimental results show that our unsupervised global

optimization guarantees global optimality even for a complex BRDF and thus eliminates uncertainty in BRDF estimation.

In addition to the Cook-Torrance model, we have also developed a global method for the distribution function of the Torrance-Sparrow model [23]:

$$D = c\exp\left(-\alpha^2/\sigma^2\right),$$

and conducted experiments. However, we have not found any meaningful difference between the two models in terms of fitting quality. For some materials, the Torrance-Sparrow model works slightly better. Perhaps the accuracy in data measurement is not high enough to account for the subtle difference in physical modeling.

In our work, the L_2 error norm is minimized by the second-order cone programming. The L_∞ error norm can also be minimized yet with more computationally efficient linear programming. We have developed a method of feasibility check for the L_∞ cost function and performed experiments. However, the results are rather similar to those presented in the previous section for the same database. Our future work includes the development of global optimization algorithms based on the L_1 cost function for the robust estimation of BRDF from data with outliers. We are also interested in developing global optimization methods for other models than the Cook-Torrance and Torrance-Sparrow models, for instance, the Ashikhmin-Shirly, Lafortune, He, and Ward models [1][13][10][24].

7. Conclusion

We have developed a global optimization method for the estimation of the one- and two-lobe Cook-Torrance model parameters. To the best of our knowledge, this is the first approach that claims global optimality in the estimation of BRDF with multiple specular lobes. For the highly nonlinear BRDF function, we divide the estimation problem into the branch-and-bound search for the surface roughness parameters and the convex programming for others. Our method does not require any supervision, and we have demonstrated that it performs noticeably better than a carefully supervised traditional optimization method for the isotropic data in the MERL database.

References

[1] M. Ashikhmin, S. Premośe, and P. Shirley. A microfacet-based BRDF generator. *Proc. SIGGRAPH*, pages 65–74, 2000.

[2] R. Basri and D. W. Jacobs. Lambertian reflectance and linear subspaces. *IEEE PAMI*, 25(2):218–233, Feb. 2003.

[3] P. Beckmann and A. Spizzichino. The Scattering of Electromagnetic Waves from Rough Surfaces, MacMillan, pp.I-33, 70-98, 1963.

[4] Blinn, J. .Models of Light Reflection for Computer Synthesized Pictures. *Proc. SIGGRAPH*, 1977.

[5] R. Cook and K. Torrance. A Reflectance Model for Computer Graphics. *ACM Trans. on Graphics*, 15(4):307-316, 1981.

[6] Cornell. Light Measurement Laboratory. Web page. http://www.graphics.cornell.edu/online/measurements/

[7] CUReT. Columbia-Utrech Reflectance and Texture. Web page. http://www1.cs.columbia.edu/CAVE//software/curet/

[8] K.J. Dana, B. Van Ginneken, S.K. Nayar, and J.J. Koenderink. Reflectance and texture of real world surfaces. *ACM Trans. on Graphics* 1(18):1-34, 1999.

[9] R. Hartley and F. Kahl. Global Optimization through Rotation Space Search. *IJCV*, 82(1):64-79, 2009.

[10] X. He, K. Torrance, F. Sillion, and D. Greenberg. A Comprehensive Physical Model for Light Reflection, *Proc. SIGGRAPH*, 1991.

[11] F. Kahl. Multiple view geometry and the L_∞-norm. *Proc. ICCV*, pages 1002–1009, 2005.

[12] O. Ke and T. Kanade,. Quasiconvex optimization for robust geometric reconstruction. *IEEE Trans. PAMI*, 29(10): 1834–1847, 2007

[13] E. Lafortune, S.C. Foo, K. Torrance, and D. Greenberg,. Non-Linear Approximation of Reflectance Functions, *Proc. SIGGRAPH*, pages 117–126, 1997.

[14] W. Matusik, H. Pfister, M. Brand and L. McMillan. A Data-Driven Reflectance Model. *ACM Trans. on Graphics*, 22(3):759-769. 2003

[15] MERL database. Web page. http://www.merl.com/brdf/.

[16] S. K. Nayar, K. Ikeuchi, and T. Kanade, "Surface reflection: physical and geometrical perspectives," *IEEE Trans. on PAMI*, 13(7): 611-634, 1991.

[17] A. Ngan, F. Durand, and W. Matusik. Experimental analysis of BRDF models. *Proc. Eurographics Symposium on Rendering*, pages 117–226, 2005.

[18] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis. Geometric considerations and nomenclature for reflectance. *NBS Monograph*, 160, 1977.

[19] M. Oren and S.K. Nayar. Generalization of Lambert's Reflectance Model, *Proc. SIGGRAPH*, 1994.

[20] R. Ramamoorthi and P. Hanrahan. Frequency space environment map rendering. *Proc. SIGGRAPH* '02, pages 517–526, 2002.

[21] I. Sato, T. Okabe, Y. Sato, and K. Ikeuchi. Appearance sampling for obtaining a set of basis images for variable illumination. *Proc. of ICCV*, pages 800–807, 2003.

[22] J. Sturm, Using SeDuMi 1.02, a Matlab toolbox for optimization over symmetric cones. *Optimization Methods and Software* 11-12 : 625–653, 1999.

[23] K. E. Torrance and E. M. Sparrow, "Theory for off-specular reflection from roughened surface," *Journal of Optical Society of America*, 57: 1105-1114, 1967.

[24] G. J. Ward. Measuring and modeling anisotropic reflection. *Proc. SIGGRAPH*, pages 265–272, 1992.

[25] M.I.A. Lourakis. levmar : Levenberg-Marquardt nonlinear least squares algorithms in C/C++, Web page. http://www.ics.forth.gr/_lourakis/levmar.