

What An Image Reveals About Material Reflectance

Manmohan Chandraker
University of California, Berkeley

Ravi Ramamoorthi
University of California, Berkeley

Abstract

We derive precise conditions under which material reflectance properties may be estimated from a single image of a homogeneous curved surface (canonically a sphere), lit by a directional source. Based on the observation that light is reflected along certain (a priori unknown) preferred directions such as the half-angle, we propose a semiparametric BRDF abstraction that lies between purely parametric and purely data-driven models. Formulating BRDF estimation as a particular type of semiparametric regression, both the preferred directions and the form of BRDF variation along them can be estimated from data.

Our approach has significant theoretical, algorithmic and empirical benefits, lends insights into material behavior and enables novel applications. While it is well-known that fitting multi-lobe BRDFs may be ill-posed under certain conditions, prior to this work, precise results for the well-posedness of BRDF estimation had remained elusive. Since our BRDF representation is derived from physical intuition, but relies on data, we avoid pitfalls of both parametric (low generalizability) and non-parametric regression (low interpretability, curse of dimensionality). Finally, we discuss several applications such as single-image relighting, light source estimation and physically meaningful BRDF editing.

1. Introduction

Objects in the natural world exhibit diverse appearances due to the wide variety in underlying material reflectance. This is encoded by the bidirectional reflectance distribution function (BRDF), which relates incoming and outgoing directions of light transport. Parametric models of the BRDF such as Blinn-Phong and Torrance-Sparrow have been widely studied in the vision and graphics communities [1, 21]. While parametric models are inspired by physics and mimic material appearance reasonably well, there has been an emphasis in recent years on purely data-driven representations to achieve closer conformity to empirical observations [16, 20]. However, the high dimensionality of the BRDF space necessitates a large amount of data. More importantly, it is not clear under what conditions one may expect generic nonparametric methods to produce unambiguous results. Even for multi-lobe parametric models like Cook-Torrance [4], optimization is challenging [24] and it is unclear when an unambiguous fit can be achieved.

In this paper, we focus on BRDF estimation from a single image of a sphere, or more generally a curved surface of known geometry and homogeneous reflectance, lit by a directional source. Rather than purely parametric (potentially

inaccurate) or purely data-driven (potentially ill-posed for fitting), this paper suggests the viability of semiparametric approaches for estimating material reflectance properties from a single image. The empirical observation that motivates our approach is that, for most materials, light is reflected along certain preferred directions. These directions, as well as the behavior of light reflection about them, may vary with material type and must be estimated from the data. As a comparison, parametric models of the BRDF assume that both these directions (half-angle, back-scatter direction and so on), as well as the form of the distribution (Gaussian, Beckmann and so on) are precisely known. As opposed to general nonparametric regression approaches, we make the physically valid observation that the reflectance is a sum of (unknown) univariate non-linear functions, acting on projections of the surface normal on a few (unknown) directions.

Our approach has significant theoretical, algorithmic and practical benefits. While our formulation is physically valid and encompasses traditional parametric models, it makes minimal assumptions on the shape and orientation of the involved distributions. Yet, in Section 4, we show that a number of insights into BRDF estimation from a single image may be derived. For instance, we show that estimating reflectance functions with unknown lights or mixture models is inherently ill-posed. But when there is a backscatter in the source direction, the problem becomes well-posed. With known lights, we show that BRDFs with one or two specular lobes, oriented along unknown directions, may be estimated uniquely. To the best of our knowledge, this is the first set of results that exactly specifies the estimability of BRDFs.

Note that the uniqueness and ambiguity conditions that we outline are inherent in the BRDF estimation problem and independent of actual choice of estimation procedure. But from an algorithmic perspective, our semiparametric BRDF model is a natural fit to statistical data-fitting techniques like projection pursuit regression. Since the link functions acting on each projection are univariate, such methods are computationally cheaper than nonparametric regression.

In Section 6, we show that leveraging the physical structure of BRDFs lends insights which may be difficult to obtain with generic functional regression. Exploiting physical intuition also enables novel applications like single-image relighting, light source estimation and BRDF editing.

To summarize, our semiparametric outlook on material behavior leads to the following fundamental contributions:

- Precise conditions may be derived under which single-image BRDF estimation is well-posed.
- Only the most basic physical conditions are required to be satisfied, so closer fits to empirical data may be achieved than possible with parametric models.

- Physically-based modeling avoids the curse of dimensionality inherent in nonparametric estimation.
- Exploiting physical intuition lends insights into material behavior and enables novel applications.

2. Related Work

Parametric models of reflectance have a long history in computer graphics and vision [1, 21]. We refer the reader to recent empirical studies like [18] that compare various models. There has been prior work on estimating parametric reflectance models from single images, such as [2, 25]. As opposed to these methods, we do not explicitly define the functions that form our BRDF models.

Purely data-driven representations of reflectance have also been enabled by the availability of comprehensive BRDF databases [16]. More closely related to our work are nonparametric approaches such as [20], which estimate reflectance from a single image under environment lighting, by exploiting the bilinearity of the reflectance and illumination. But it is difficult to derive significant insight into material properties through nonparametric estimation of higher-dimensional functions and it is as yet unknown under what conditions that problem is ill-posed. In contrast, we use directional lights, estimate only univariate functions and derive precise conditions for well-posedness.

Semiparametric models of spatially varying BRDF for interactive editing have been proposed in [14], where the reflectance functions are unknown, but the directions are known. Our formulation may also be related at a high level to recent approaches like [17] that factor the BRDF into incoming and outgoing directions as precomputation and recombine them for rendering.

We derive inspiration from works on identifiability of additive models [9, 26], but our theory holds for the 2-sphere \mathcal{S}^2 , rather than \mathbb{R}^3 . Our algorithms belong to the class of projection pursuit regression [6], where the regression surface is modeled as a sum of general smooth functions of linear combinations of predictor variables (in our case, components of the surface normal). This can also be considered as a generalized version of independent components analysis [11] and blind source separation, with the number of components not restricted to be equal to the predictor dimension.

3. Representation

Notation: We will denote unit vectors as α and unnormalized vectors as $\tilde{\alpha}$. Unless stated otherwise, a vector is assumed to be a 3-vector. The i -th component of a vector is denoted as α_i , not to be confused with α_i , which is a 3-vector. Functions are denoted as $f(\cdot)$ and unless stated otherwise, they are univariate and defined on non-negative reals, that is, $f : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. We reserve the notations \mathbf{n} , \mathbf{s} , \mathbf{v} and \mathbf{h} for, respectively, the directions of the surface normal, light source, camera and half-angle bisector of the source and camera (all unit vectors). We denote the unit vector with the i -th component equal to 1 as \mathbf{e}_i .

Setup: We will henceforth assume that we have available a single image of an object of known geometry (say, a sphere), illuminated by a single directional point light source. Note that dependence of the BRDF exclusively on the source and view may be ignored for estimation in a single-image

setup, since it remains constant over the whole hemisphere of visible surface normals. (Thus, for instance, the dependence of the geometric term of microfacet models on $\mathbf{h}^\top \mathbf{v}$ is often inconsequential for some applications.)

Now, it is reasonable to assume that the reflectance is a function of the interaction of the surface normal with various (*a priori* unknown) directions, whose exact form is determined by intrinsic material properties. Note that while this setup restricts us to a slice of the BRDF, we do estimate complete reflectance functions allowing synthesis of images with novel lighting and viewpoints.

BRDF Abstraction: It is also a common observation that a BRDF is composed of a sum of “lobes” – that is, there exist certain preferred directions along which the reflectance function is “concentrated”. As the normal deviates away from such a direction, the contribution of that lobe to the reflected intensity diminishes monotonically. Thus, the BRDF can be intuitively represented as a sum of univariate functions:

$$\rho(\mathbf{n}) = \sum_{i=1}^K f_i(\alpha_i^\top \mathbf{n}), \quad (1)$$

where α_i are some directions, that is, $\|\alpha_i\| = 1$. The goal of BRDF estimation now becomes determining the directions α_i and the functions f_i .

In parametric models of the BRDF, α_i usually correspond to \mathbf{s} , \mathbf{v} and \mathbf{h} . Examples of such reflectance functions include Blinn-Phong (where $\alpha_1 = \mathbf{s}$ and $\alpha_2 = \mathbf{h}$, with $f_1 = \text{constant}$ and $f_2(t) = t^\mu$) or simplified Torrance-Sparrow models (where $f_2(t) = \exp(-\mu t^2)$). Our formulation will discover any such dependences, but also allows for the existence of other significant directions governed by material behavior (see Section 6).

In some cases, the BRDF may be represented as a product of monotonic functions [17]. Then, the logarithm of the BRDF may be written as a linear combination of monotonic functions. Examples also include simplified versions of the Torrance-Sparrow model, which model the BRDF as a Gaussian distribution around the half-angle, or recent statistics-based extensions such as mixtures of hemispherical exponential power distributions [19]. The Lafortune model [13] also fits multiple lobes to measured data, but unlike this work, assumes that their parametric forms are known.

One may also interpret the model in (1) as a generalization of standard methods like principal components analysis that seek linear structure in data. Instead, we are interested in understanding a non-linear structure, with the important restriction that the underlying functions are constant on certain hyperplanes. For this reason, the functions f_i are also referred to as ridge functions.

For a physically valid BRDF, the functions $f_i(\cdot)$ must satisfy certain technical conditions:

- (L1) The domain of each f_i is the real closed interval $[0, 1]$.
- (L2) Each f_i is non-negative, that is, $f_i \geq 0$.
- (L3) Each f_i is strictly monotonic, that is, $f_i' > 0$.
- (L4) Each f_i passes through the origin, that is, $f_i(0) = 0$.

Note that (L2), (L3) and (L4) together mean $f_i(y) > 0$ for any $y \in (0, 1] \subset \mathbb{R}$. Indeed, most aspects of our theory hold true for this weaker condition than (L3), but we continue to assume monotonicity to match physical intuition. Unless

explicitly stated, for the sake of brevity, we will refer to a “strictly monotonic” function as just “monotonic”.

It is known in functional analysis that any “nice” function can be represented to arbitrary precision using a model such as (1), thus, our modeling is always valid [5]. So, it is not surprising that reflectance functions may also be represented as such. However, the intuitive property of reflectance functions that we seek to exploit, namely monotonic behavior along preferred projections, has two important consequences.

- Low computational burden (Section 5): since the reflectance data is already aligned along certain directions, the number of ridge functions required, K , is very small, typically, two or three.
- Uniqueness (Section 4): physical requirements place strong constraints on the model that guarantee uniqueness, regardless of the actual estimation method.

4. Uniqueness of BRDF Estimation

In this section, we derive precise conditions when the general BRDF of (1) can be unambiguously estimated from data, using a single image. Note that we must determine the directions α_i and the functions $f_i(\cdot)$ to estimate the BRDF.

4.1. 1-D BRDF

Let us begin with the simplest case of a 1-D BRDF. Examples of 1-D BRDFs include a Lambertian reflectance, or an arbitrary isotropic BRDF with a colocated source and sensor. Given a single image of, say, a sphere, with 1D reflectance $f(\alpha^\top \mathbf{n})$, we show that both $f(\cdot)$ and α can be determined:

Proposition 1. *Given that the BRDF is 1-D, it can be uniquely determined using a single image.*

Proof. Assume to the contrary that there exist β and monotonic $g(\cdot)$, distinct from α and $f(\cdot)$, such that for all $\mathbf{n} \in \mathcal{S}^2$

$$f(\alpha^\top \mathbf{n}) = g(\beta^\top \mathbf{n}). \quad (2)$$

For some constant $c \in [0, 1)$, consider the circle on the Gauss sphere given by $\Omega_\alpha = \{\mathbf{n} \in \mathcal{S}^2 : \alpha^\top \mathbf{n} = c\}$. Clearly, $f(\alpha^\top \mathbf{n}) = f(c) = \text{constant}$ on Ω_α . Since $\beta \neq \alpha$, there must exist $\mathbf{n}_1, \mathbf{n}_2 \in \Omega_\alpha$, such that $\beta^\top \mathbf{n}_1 \neq \beta^\top \mathbf{n}_2$. Since $g(\cdot)$ is a monotonic function, it must be true that $g(\mathbf{n}_1^\top \beta) \neq g(\mathbf{n}_2^\top \beta)$, whereby (2) contradicts the fact that $f(\cdot)$ is constant on Ω_α . \square

Thus, by a suitable estimation procedure, one may unambiguously determine α and $f(\cdot)$ to estimate the BRDF. While the above result will be utilized in the subsequent sections, note that it does not preclude the fact that a 1-D BRDF might be representable as a sum of many functions. That is established in the following result:

Proposition 2. *A 1-D BRDF is uniquely determinable.*

Proof. Assume that there exist β_i , $i = 1, \dots, K$, for some $K \geq 1$, such that the BRDF has an alternate expression

$$f(\alpha^\top \mathbf{n}) = \sum_{i=1}^K g_i(\beta_i^\top \mathbf{n}). \quad (3)$$

Note that non-uniqueness requires at least one $\beta_i \neq \alpha$. Let \mathbf{A} be an invertible 3×3 transformation whose first column

is α . For $\mathbf{x} = \mathbf{A}^\top \mathbf{n}$ and $\gamma_i = \mathbf{A}^{-1} \beta_i$, the above relation transforms to $f(x_1) = \sum_{i=1}^K g_i(\gamma_i^\top \mathbf{x})$.

Let \mathbf{e}_i be the unit vector with the i -th component equal to 1. For $\mathbf{x} = \mathbf{e}_2$ and $\mathbf{x} = \mathbf{e}_3$, we note that the LHS is $f(0) = 0$. Since $g_i(y) = 0$ only for $y = 0$, it must be true that $\gamma_i^\top \mathbf{e}_2 = \gamma_i^\top \mathbf{e}_3 = 0$. That is, $\gamma_i = \mathbf{e}_1$, for all $i = 1, \dots, K$. Thus, the decomposition (3) reduces to the trivial one, $f(x_1) = \sum_{i=1}^K g_i(x_1)$, which is a contradiction of the assumption that at least one β_i is not the same as α . Thus, a 1-D BRDF is uniquely determinable. \square

4.2. Some Basic Restrictions

In this section, we present some basic restrictions that must be satisfied by BRDFs of form (1) for them to be uniquely estimable. The remarks in this section are well-known for general semiparametric regression [5, 9, 26], but for BRDF estimation, they are also reasonable from a physical point of view.

We begin by noting that, for any vectors $\tilde{\alpha}', \tilde{\beta}'$ such that $\tilde{\alpha}' + \tilde{\beta}' = \alpha + \beta$, we have $\tilde{\alpha}'^\top \mathbf{n} + \tilde{\beta}'^\top \mathbf{n} = \alpha^\top \mathbf{n} + \beta^\top \mathbf{n}$. Thus, we remark:

Remark 1. *For a BRDF of form (1) to be estimable, at most one term can be diffuse.*

The following two conditions must also hold:

Remark 2. *For a BRDF of form (1) to be estimable, at most one of the f_i can be quadratic.*

Remark 3. *A BRDF of form (1) might not be estimable when the directions α_i are linearly dependent.*

We refer the reader to [9] for the simple proofs, but note that Remark 2 is an important restriction, since it is common in prior work to approximate BRDFs as Gaussian mixture models. In our theory, when considering logarithms, the behavior of at most one lobe may be modeled as Gaussian. Further, it follows from Remark 3 that BRDFs guaranteed to be estimable can be at most 3-lobe. Again, note that this condition is required only to guarantee determinability of the BRDF. It is possible for the directions to be linearly dependent while the BRDF remains uniquely estimable.

Next, we analyze the special case of dichromatic BRDFs similar to half-angle models like Blinn-Phong or Torrance-Sparrow (but with unknown directions and functions).

4.3. 1-lobe (Or Dichromatic) BRDFs

In many cases, such as a Blinn-Phong BRDF, we know that the diffuse component is a linear function of the surface normal. Thus, such BRDFs may be written as

$$\rho(\mathbf{n}) = \tilde{\alpha}^\top \mathbf{n} + f(\beta^\top \mathbf{n}) \quad (4)$$

where $\tilde{\alpha} \in \mathbb{R}^3$, $\beta \in \mathcal{S}^2$ and f is some monotonic, non-linear function. Further, $\tilde{\alpha}$ and β are linearly independent, else the model reduces to a 1-D BRDF.

In practice, we know that α corresponds to the light source direction. It is empirically known that the specular lobe points away from the light direction, so it is valid to assume linear independence of $\tilde{\alpha}$ and β . From Proposition 1, we know that f has to be a non-linear function for the

model to be uniquely defined. Again, this is a physically valid assumption, since the specular lobe of a half-angle BRDF model is known to be non-linear (that is, behaves in a more sharply-peaked fashion than the diffuse component).

Proposition 3. *For an unknown light source, the model in (4) is ambiguous.*

This can be easily seen by noting that

$$\tilde{\alpha}^\top \mathbf{n} + f(\beta^\top \mathbf{n}) = \tilde{\gamma}^\top \mathbf{n} + g(\beta^\top \mathbf{n}), \quad (5)$$

where $g(\beta^\top \mathbf{n}) = k\beta^\top \mathbf{n} + f(\beta^\top \mathbf{n})$ and $\tilde{\gamma} = \tilde{\alpha} - k\beta$. However, note that the direction β may still be uniquely recoverable.

Proposition 4. *For a known light source direction and strength, the model in (4) is not ambiguous.*

Proof. Suppose the light source direction is α and strength c_1 . Let us assume to the contrary that the model is ambiguous. Then, there exist directions β, γ and univariate monotonic functions f, g , not all of them identical, such that

$$\tilde{\alpha}^\top \mathbf{n} + f(\beta^\top \mathbf{n}) = \tilde{\alpha}^\top \mathbf{n} + g(\gamma^\top \mathbf{n}). \quad (6)$$

But this violates Proposition 1, so the model is unambiguous. \square

In practice, it is easy to determine the light source direction, but not the strength. However, it can be shown that it only suffices to know the light source direction in order for the model in (4) to be uniquely defined.

Proposition 5. *For a light source at a known direction, but of unknown strength, the model in (4) is not ambiguous.*

Proof. Suppose the light source direction is α . Then, let us assume to the contrary that the model is ambiguous. Then, there exist constants c_1, c_2 , directions β, γ and univariate monotonic functions f, g , not all of them identical, such that

$$c_1 \alpha^\top \mathbf{n} + f(\beta^\top \mathbf{n}) = c_2 \alpha^\top \mathbf{n} + g(\gamma^\top \mathbf{n}). \quad (7)$$

Suppose $c_1 = c_2$. Then, we must have $f(\beta^\top \mathbf{n}) = g(\gamma^\top \mathbf{n})$, where at least one of $\beta \neq \gamma$ and $f \neq g$ must be true. But that will violate Proposition 1, thus, we must have $c_1 \neq c_2$.

Without loss of generality, let $c = (c_1 - c_2) > 0$. Then,

$$c \alpha^\top \mathbf{n} + f(\beta^\top \mathbf{n}) = g(\gamma^\top \mathbf{n}). \quad (8)$$

Consider the vector $\mathbf{n}_0 \perp \text{span}(\alpha, \beta)$. Then, we must have $\alpha^\top \mathbf{n}_0 = 0$ and $\beta^\top \mathbf{n}_0 = 0$. So, $f(\beta^\top \mathbf{n}_0) = 0$ and from (8), we must have $g(\gamma^\top \mathbf{n}_0) = 0$. Since g is zero only at the origin, this is only possible if $\gamma^\top \mathbf{n}_0 = 0$, that is, $\gamma \in \text{span}(\alpha, \beta)$. Thus, we may write $c\alpha = a\beta - b\gamma$, whereby (8) becomes

$$a\beta^\top \mathbf{n} + f(\beta^\top \mathbf{n}) = b\gamma^\top \mathbf{n} + g(\gamma^\top \mathbf{n}). \quad (9)$$

Now, consider the transformation of variables given by $\mathbf{x} = \mathbf{A}^\top \mathbf{n}$, where $\mathbf{A} = [\beta, \gamma, \mathbf{e}_3]$. Recall our notation, where $\mathbf{e}_3 = (0, 0, 1)^\top$. Then, we may rewrite (9) as

$$ax_1 + f(x_1) = bx_2 + g(x_2). \quad (10)$$

Since x_1 and x_2 are independent variables, this is only possible if $f(y) = -ay + c_0$ and $g(y) = -by + c_0$, for some constant c_0 . This is a contradiction, since we know that f is a non-linear function. Thus, the model in (4) is uniquely determined when the source direction is known. \square

4.4. General 2-lobe BRDFs

A general 2-lobe BRDF has the form

$$\rho(\mathbf{n}) = f_1(\alpha_1^\top \mathbf{n}) + f_2(\alpha_2^\top \mathbf{n}), \quad (11)$$

where α_1, α_2 are linearly independent and both f_1, f_2 are non-linear, non-quadratic functions. Note that we have already considered the cases where these functions are linear or quadratic in the previous section.

Also, recall that all the functions we are considering correspond to physical BRDFs, so they are constrained to be non-negative, passing through the origin and non-zero elsewhere.

In order to prove the uniqueness conditions for estimation of a 2-lobe BRDF, following [26], we use an auxiliary lemma from probability theory.

Lemma 1 (Khatri and Rao). *For linearly independent vectors $\alpha_i \in \mathbb{R}^n$, any unknown functions $f_i : \mathbb{R} \rightarrow \mathbb{R}$ and $g_j : \mathbb{R} \rightarrow \mathbb{R}$ that satisfy the relation*

$$\sum_{i=1}^m f_i(\alpha_i^\top \mathbf{x}) = \sum_{j=1}^n g_j(x_j) \quad (12)$$

over the domain $\{\mathbf{x} \in \mathbb{R}^n : |x_j| < r, j = 1, \dots, n\}$, for some $r > 0$, must be polynomials of degree at most 2.

For a proof, we refer the reader to [12]. Notice the domain of the equation in the above lemma, which is an open subset of \mathbb{R}^n . Such generality of the domain is often an important requirement for functional analysis. In contrast, for the problems that are encountered in BRDF estimation, the domain of definition for the surface normals is the closed set \mathcal{S}^2 with no interior, so unlike [26], the above result cannot be applied in a straightforward manner.

Yet, given the above lemma, we can prove the following:

Proposition 6. *A general 2-lobe BRDF of the form in (11) is always uniquely determined by a single image.*

Proof. Assume to the contrary that there exist functions g_i and directions β_i , where $1 \leq i \leq K$, such that

$$f_1(\alpha_1^\top \mathbf{n}) + f_2(\alpha_2^\top \mathbf{n}) = \sum_{i=1}^K g_i(\beta_i^\top \mathbf{n}). \quad (13)$$

Since α_1, α_2 are linearly independent, there exists a 3×3 matrix \mathbf{A} , such that $\mathbf{A}^{-1} [\alpha_1, \alpha_2] = [\mathbf{e}_1, \mathbf{e}_2]$, where \mathbf{e}_i is the unit vector with the i -th component equal to 1. Let us define $\mathbf{x} = \mathbf{A}^\top \mathbf{n}$ and $\gamma_i = \mathbf{A}^{-1} \beta_i$, for $i = 1, \dots, K$.

We consider the particular choice $\mathbf{A} = [\alpha_1, \alpha_2, \mathbf{e}_3]$. Then, if \mathcal{Q} is the domain of \mathbf{x} , we note that $\mathbf{e}_3 \in \mathcal{Q}$. Now, we can rewrite (13) as

$$f_1(x_1) + f_2(x_2) = \sum_{i=1}^K g_i(\gamma_i^\top \mathbf{x}). \quad (14)$$

Substitute $\mathbf{x} = \mathbf{e}_3$. Then, the LHS is $f_1(0) + f_2(0) = 0$. For the RHS to be also 0, since the g_i are non-negative functions, it must be true that $\gamma_i^\top \mathbf{e}_3 = 0$ for all $i = 1, \dots, K$. Thus, $\gamma_i \in \text{span}(\mathbf{e}_1, \mathbf{e}_2)$ and we may express $\gamma_i = [p_i, q_i, 0]^\top$, for all $i = 1, \dots, K$. Now, we can rewrite (14) as

$$f_1(x_1) + f_2(x_2) = \sum_{i=1}^K g_i(p_i x_1 + q_i x_2). \quad (15)$$

Let $\mathcal{Q}' \subset \mathbb{R}^2$ be the domain of $[n_1, n_2]^\top$ and \mathcal{Q}'' be the domain of $[x_1, x_2]^\top$. Define

$$\mathcal{B}_r(y_1, y_2) = \{(y_1, y_2)^\top \in \mathbb{R}^2 : |y_1| < r, |y_2| < r\}. \quad (16)$$

Since $\mathcal{B}_r(n_1, n_2) \subset \mathcal{Q}'$ for any $r < \frac{1}{\sqrt{2}}$, it must be true that

there exists an $r' > 0$, such that $\mathcal{B}_{r'}(x_1, x_2) \subseteq \mathcal{Q}''$. Then, in accordance with the Khatri-Rao Lemma, all of f_j, g_i , for $j = 1, 2$ and $i = 1, \dots, K$ are restricted to be quadratic functions. But this contradicts the condition that the f_j are neither linear nor quadratic, so the general 2-lobe BRDF in (11) is uniquely determined by a single image. \square

K-lobe BRDF, $K \geq 3$: Note that we have already established that for $K > 3$, uniqueness of BRDF estimation cannot be guaranteed. It remains an open question whether estimation of BRDF models with $K = 3$ lobes is well-posed. In practice, it is known that projection pursuit with 3 or more ridge functions is often ill-conditioned [10]. However those results are also known only for convex subsets of \mathbb{R}^3 . It has been observed that fitting 3-lobe BRDFs can be unstable, so we conjecture that similar ill-conditioning results may also exist for estimation on \mathcal{S}^2 .

5. Estimation Algorithms

We now indicate the wide suite of algorithms available for single-image BRDF estimation within our framework. While a complete comparison is beyond the scope of this paper, these methods are well-studied in the statistics community and we will point the reader to the relevant references.

5.1. 1-D and Phong-Like BRDFs

Nonparametric methods may be used to directly estimate the projection direction α in a 1-D BRDF $f(\alpha^\top \mathbf{n})$, followed by a LOESS regression [3]. But it is also relatively inexpensive to search over the sphere for the projection direction. We adopt the latter approach in a coarse-to-fine implementation. Projecting the covariates along each putative α , we fit a cubic spline to minimize the energy:

$$\min \sum_{i=1}^N (y_i - s(x_i))^2 + \lambda \int s''(x)^2 dx. \quad (17)$$

The smoothing parameter is chosen by cross-validation. General cubic splines yield a good fit to data, but may not be monotonic. It is well-known that sufficient conditions for monotonicity of a cubic spline within each knot-interval are guaranteed by a system of seven linear constraints [7]. Additional terms that penalize second derivative discontinuities can be included to ensure C^2 smoothness [23].

5.2. K -lobe BRDFs, $K > 1$

In some cases, one may have prior knowledge of the preferred directions along which the reflectance lobes are centered, which reduces the problem to estimating a generalized additive model [8]. When such prior knowledge is not available, one approach might be to alternate between solving for the projection directions and a neural-net regression to estimate the link functions as sums of sigmoids.




Figure 1. With a colocated source and camera, we have a 1-D BRDF. For gold-metallic-paint material, the estimated projection direction is found to be within 0.7 degrees of the viewing direction, $(0, 0, 1)^\top$. The reflectance curves obtained in the red and green channels are shown (the curve for blue channel is similar). The blue points represent data and the red curve is the estimated fit.

We use a more general, well-known method to simultaneously estimate the directions and functions that constitute a BRDF of the form (1), namely projection pursuit [6], which is an application of the successive refinement principle to solve regression problems of the form:

$$\min_{\alpha, f_i} \|\rho(\mathbf{n}) - \sum_i f_i(\alpha_i^\top \mathbf{n})\|^2. \quad (18)$$

The algorithm builds each term in the model, by selecting a direction that, for a corresponding smoothed representation of residuals, maximizes a suitable figure of merit, such as the fraction of unexplained variance. That is, if $r_j = y_j - \sum_{i=1}^k f_i(\alpha_i^\top \mathbf{n}_j)$ are the residuals after k terms have been fitted, the $(k + 1)$ -th direction, α_{k+1} , is chosen as

$$\alpha_{k+1} = \arg \min \frac{\sum_j \|r_j - \sum_{i=1}^{k+1} f_i(\alpha_i^\top \mathbf{n})\|^2}{\sum_j r_j^2}. \quad (19)$$

We refer the reader to [10] for an introduction and [22] for greater details on the implementation. As is standard practice in projection pursuit analysis, we first fit a greater number of terms (5, in our experiments) than required, then sequentially drop the least effective term followed by refitting. We use a projection pursuit implementation in the R language for statistical computing, employing the super-smoother of [6].

6. Experiments, Insights and Applications

In this section, we verify our theory on real data from the publicly available MERL database [16]. Along the way, we highlight two of the main advantages of our theory – that a semiparametric approach lends insights into material behavior, as well as towards the problem of BRDF estimation itself. We also show auxiliary practical benefits such as single-image relighting, BRDF editing and simultaneous estimation of light source and reflectance.

1-D BRDF: We begin with the simplest case of 1-D BRDFs, which are proved to be uniquely determined by a single image in Proposition 2. In Figure 1, we consider the input image of a gold-metallic-paint sphere, with a colocated source and camera. We search over all directions and show that there exists one along which the covariates can be projected, so that a one-dimensional curve fit explains most of the variance. Moreover, this estimated direction is actually the camera direction, which explains why a univariate function well-approximates the data.




Figure 2. Ambiguity in 1-lobe model with unknown light source. Note the rise in the specular red curve, corresponding to the half angle, even for low values of $n^T h$. This is because the diffuse and specular terms interact in (5).

Plastics: Shiny plastics are well-approximated by a 1-lobe model, centered around the half-angle. First, we verify Proposition 3 that the estimation is ambiguous with unknown source direction. This is shown for the `green-plastic` and `gray-plastic` materials in Figure 2.

Application 1 - Light source estimation: Even though the reflectance curves are ambiguous in Figure 2, we note that the direction of non-linear dependence (β in Proposition 3) can be accurately recovered. Thus, knowing that the material is expected to follow a 1-lobe half-angle model, we may estimate the light source direction. This is of significant practical interest, since light source directions are hard to calibrate and even an error of one or two degrees can introduce substantial errors in BRDF measurements. For the `green-plastic` and `gray-plastic` materials in Figure 2, the estimated light source directions are $[1.004, 1.006, 1]^T$ and $[1.004, 1.008, 1]^T$, which are 0.19 and 0.14 degrees away from the “calibrated” direction $[1, 1, 1]^T$.

Application 2 - Relighting: With known light source direction, from Proposition 5, the reflectance curves as well as the direction of projection may be estimated from data. In Figure 3, we use the method of Section 5.1 to estimate both β and $f(\cdot)$ in (4) for the `violet-acrylic` material, thereby verifying Proposition 5. We observe that, as expected, the shape of the reflectance curves remain constant across light source directions. This constancy, together with the fact that the estimated directions are physically meaningful, allows us to predict appearance under novel lighting.

Figure 4 shows the relighting application for the `specular-yellow-phenolic` material. As a compar-




Figure 3. For a `violet-acrylic` sphere, no significant change in the error is found between one and two projection directions. So, we fit a 1-lobe model. The estimated projection direction is found to be within 3 degrees of the half-angle. The top row shows an input image under light source $(1, 1, 1)^T$ and the bottom row under $(-1, 1, 1)^T$. The estimated curves in each case are found to be nearly identical. We can use this fact to predict the appearance under novel lighting configurations.




Figure 4. Relighting from a new light source direction for the `specular-yellow-phenolic` material. The prediction closely matches the ground truth appearance, the error shown is scaled 3 times for visualization (brighter represents higher error).




Figure 5. Given input from source at $(1, 1, 1)^T$, in (b) we predict from a new light source direction $(-1, 1, 1)^T$. The ground truth is shown in (c). Note the accuracy of our estimate in both the specular highlight and far from it, relative to a parametric fit (d).

son, we also show the inferior prediction using a simplified Torrance-Sparrow model for `green-metallic-paint` in Figure 5. For a `red-phenolic` material, Figure 6 shows the 1-term PPR fit and relighting of a complex geometry.

Paints: In Figure 7(c), we show an input image for the `ipswich-pine-221` material in the MERL database, imaged under an oblique light source $(10, 10, 1)^T$. The esti-

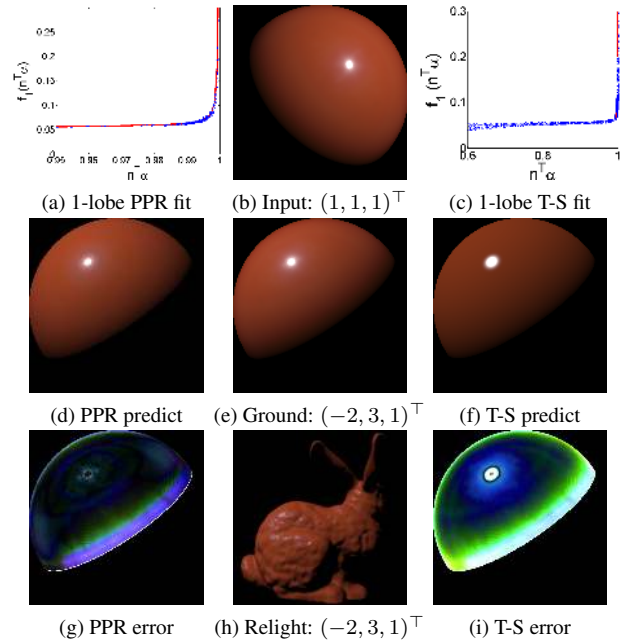


Figure 6. Given an image of a `red-phenolic` sphere with lighting direction $(1, 1, 1)^T$ (b), a 1-lobe simplified Torrance-Sparrow fit (c) shows slightly greater variance than a 1-term PPR (a). The Torrance-Sparrow prediction of the appearance under a light source at $(-2, 3, 1)^T$ is reasonable (f). But using the method of Section 5.1, the prediction (d) clearly matches the ground truth (e) better, both in the highlight and diffuse regions. The corresponding errors are shown in (i) and (g), scaled 3 times for visualization, with brighter colors representing higher error. The predicted image for a complex geometry is shown in (h).




Figure 7. For paints such as *ipswich-pine-221* under oblique illumination, we expect significant backscatter. Compare the variance in the 1-lobe fit in (g) with that of highly specular materials in Figures 3 and 4. A projection pursuit assuming unknown light source converges to directions very close to the half-angle and light source (h). Both (d) and (e) are rescaled to 3 times their actual values for better visualization. Please see the text for details.




Figure 8. For a *natural-209* sphere illuminated from $(25, 25, 1)^T$ (a), 2-term PPR yields projection directions as \mathbf{h} and \mathbf{s} . Average relative error is 2.61%. This is used to relight simple (c) and complex (d) geometries under novel lighting $(-1, 1, 1)^T$.

mated best projection direction, assuming a 1-lobe BRDF with known light source, is found to be close to the half-angle. But as seen in Figure 7(g), the data lies far from a 1-D distribution (compare to the fits for sharply specular plastic materials in Figures 3 and 4). The mean relative error (d) in the estimated 1-lobe fit (a) is found to be 4.76%.

This high error is expected for paint materials similar to wood lacquer, which may show noticeable backscatter towards the light source. Incidentally, this phenomenon is also the motivation for BRDF models used in Neumann-Neumann shaders. Now, since we expect a non-linearity in the light source direction, we can assume that the light source direction is unknown and perform a projection pursuit regression to find the best-fitting 2-lobe BRDF. The best projection directions are found to be within 2 degrees of the ground truth half-angle and the light source. The estimated reflectance curves are shown in Figure 7(h). The mean relative error (e) in the estimated 2-lobe fit (b) is found to be 2.14%. Finally, using the 2-lobe fit, we can also predict the

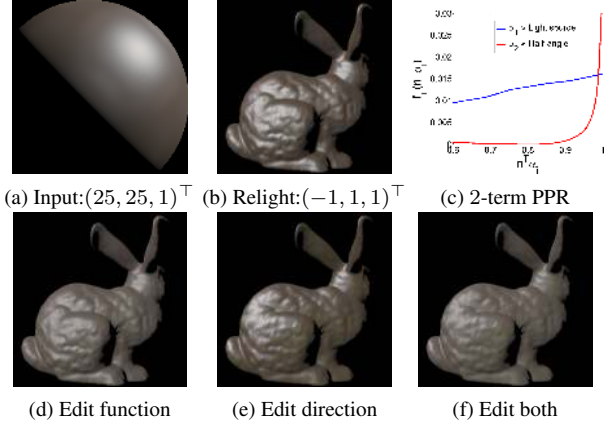


Figure 9. Examples of editing the appearance of objects with complex geometry. An advantage of our approach is that both the directions and functions that determine the BRDF can be edited to create novel appearances. Using an image of a *special-walnut-224* sphere under oblique lighting $(25, 25, 1)^T$ (a), a 2-lobe BRDF is estimated (c). An object with complex geometry is relighted from a novel direction $(-1, 1, 1)^T$ (b). The glossiness in (b) is reduced by editing the functions along the half-angle to half their value (d). The direction of the gloss in (b) is shifted from $\mathbf{s} + \mathbf{v}$ to an arbitrary new direction $\mathbf{s} + 0.3\mathbf{v}$ (e). Both the direction and the strength of gloss can be simultaneously edited (f).

image accurately from a very different light source direction, $(-1, 1, 1)^T$, as shown in Figures 7(f) and (i).

Thus, in this example, we have derived insight into the behavior of certain paints, verified our theory in Proposition 6 and provided another demonstration of simultaneous light source and reflectance estimation. A relighting example on simple and complex geometries for the *natural-209* material, whose BRDF shows similar characteristics as *ipswich-pine-221*, is shown in Figure 8.

Application 3 - BRDF Editing: An important practical advantage of our approach, compared to traditional methods, is that we simultaneously estimate both the directions and functions that determine the BRDF. Conversely, in a BRDF editing application, we can control both the directions along which the BRDF displays interesting behavior, as well as the behavior of the BRDF along those directions. A simple example where we independently edit the functions and directions of the BRDF for the *special-walnut-224* material is shown in Figure 9 (d) and (e), respectively. In Figure 9(f), we show simultaneous editing of the functions and directions of the material reflectance.

Synthetic and Natural Fibers: The most significant deviations from half-angle models were found in synthetic fibers like polyethylene and natural ones like fabrics. Again, this matches physical intuition, since reflectance from fibers is complex [15]. Figure 10(a) shows a projection of the polyethylene data along the half-angle. Rather, the most significant direction estimated by 2-D projection pursuit was close to $2\mathbf{s} + \mathbf{v}$ (and a second direction close to \mathbf{v}). The corresponding fit is shown in Figure 10(b). For the *blue-fabric* material, shown in Figure 11, the most significant direction was close to $4\mathbf{s} + 3\mathbf{v}$. In all cases, two lobes were necessary to match the observed images.




Figure 10. (a) For synthetic fiber materials like polyethylene, a projection along the half-angle is not the best one. (b) 1-D projection pursuit estimates the best fit direction to be $2s + v$. (c) Another term is found to exist aligned with the viewing direction. Average relative error is 1.57%.

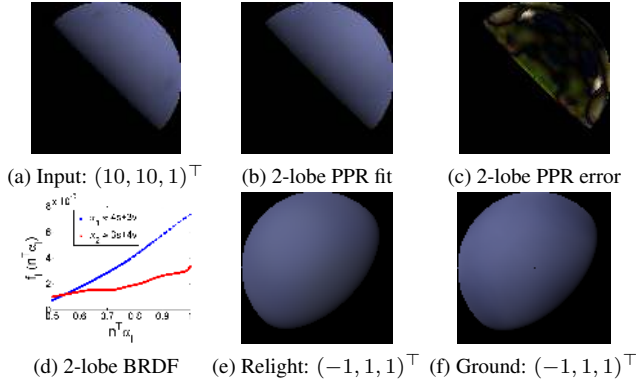


Figure 11. Given images of a blue-fabric sphere, under light source $(10, 10, 1)^T$ (a), a 2-term PPR closely fits the data (b). The corresponding error, scaled 3 times for visualization, is shown in (c). Average relative error is 2.05%. The two significant directions of the 2-lobe BRDF are $\alpha_1 = 4s + 3v$ and $\alpha_2 = 3s + 4v$ (d). The directions and curves obtained from (a) are used to predict the appearance under a new light source at $(-1, 1, 1)^T$ (e), which closely matches the ground truth (f).

7. Discussion and Future Work

In this paper, we have specified exact conditions for BRDF estimability using a single image under directional lighting. Our semiparametric BRDF model is motivated by physically valid, but minimal requirements on the form of reflectance functions. Unlike parametric models, we achieve good fits to data and unlike nonparametric ones, our estimates are easy to interpret. We leverage powerful tools from semiparametric regression analysis to develop general algorithms, derive significant insights into material behavior and present novel applications such as single-image relighting, lighting estimation and BRDF editing.

Looking ahead, it is natural to extend our models to handle anisotropy, where tangent directions are also important in determining reflectance. Relating our theoretical discoveries to the dual problem of shape reconstruction from a sampling of light source directions is also an interesting avenue. A drawback of projection pursuit regression is its tendency to get stuck in local minima, so it is also important to develop better estimation algorithms. Our models retain predictive power along with physical interpretability, so they are well-suited to drive machine learning algorithms for material classification and scene interpretation, which are directions we seek to explore in future work.

Acknowledgments This work is funded by ONR PECASE N00014-09-1-0741 and generous support from Intel, Adobe, NVIDIA and Pixar.

References

- [1] J. F. Blinn and M. E. Newell. Texture and reflection in computer generated images. *Comm. ACM*, 19:542–547, 1976.
- [2] S. Boivin and A. Gagalowicz. Image-based rendering of diffuse, specular and glossy surfaces from a single image. *SIGGRAPH*, 2001.
- [3] W. Cleveland and S. Devlin. Robust locally weighted regression and smoothing scatterplots. *JASA*, 74(368):829–836, 1979.
- [4] R. Cook and K. Torrance. A reflectance model for computer graphics. *ACM ToG*, 15(4):307–316, 1981.
- [5] P. Diaconis and M. Shahshahani. On nonlinear functions of linear combinations. *SIAM J. Sci. Comp.*, 5:175–191, 1984.
- [6] J. Friedman and W. Stuetzle. Projection pursuit regression. *JASA*, 76(376):817–823, 1981.
- [7] F. Fritsch and R. Carlson. Monotone piecewise cubic interpolation. *SIAM Num. Anal.*, 17(2):238–246, 1980.
- [8] T. Hastie and R. Tibshirani. Generalized additive models. *Stat. Sci.*, 1(3):297–318, 1986.
- [9] J. Horowitz. *Semiparametric methods in econometrics*. Springer, NY, 1998.
- [10] P. Huber. Projection pursuit. *Ann. Stat.*, 13(2):435–475, 1985.
- [11] A. Hyvarinen and J. Karhunen. *Independent Components Analysis*. Wiley, NY, 2001.
- [12] C. Khatri and C. Rao. Solutions to some functional equations and their applications to characterization of probability distributions. *Sankhya A*, 30(2):167–180, 1968.
- [13] E. P. F. LaFortune, S.-C. Foo, K. E. Torrance, and D. P. Greenberg. Non-linear approximation of reflectance functions. In *SIGGRAPH*, pages 117–126, 1997.
- [14] J. Lawrence, A. Ben-Artzi, C. Decoro, W. Matusik, H. Pfister, R. Ramamoorthi, and S. Rusinkiewicz. Inverse shade trees for non-parametric material representation and editing. *ACM ToG (SIGGRAPH)*, pages 735–745, 2006.
- [15] S. R. Marschner, H. W. Jensen, M. Cammarano, S. Worley, and P. Hanrahan. Light scattering from human hair fibers. *ACM ToG*, 22:780–791, 2003.
- [16] W. Matusik, H. Pfister, M. Brand, and L. McMillan. A data-driven reflectance model. *ACM Transactions on Graphics*, 22(3):759–769, July 2003.
- [17] M. D. McCool, J. Ang, and A. Ahmad. Homomorphic factorization of BRDFs for high-performance rendering. In *SIGGRAPH*, pages 171–178, 2001.
- [18] A. Ngan, F. Durand, and W. Matusik. Experimental analysis of BRDF models. In *EGSR*, 2005.
- [19] K. Nishino. Directional statistics BRDF model. In *ICCV*, 2009.
- [20] F. Romeiro and T. Zickler. Blind reflectometry. In *ECCV*, 2010.
- [21] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces. *JOSA*, 57:1105–1112, 1967.
- [22] W. Venables and B. Ripley. *Modern Applied Statistics with S*. Springer, 2002.
- [23] G. Wolberg and I. Alf. An energy-minimization framework for monotonic cubic spline interpolation. *JCAM*, 143:145–188, 2002.
- [24] C. Yu, Y. Seo, and S. Lee. Global optimization for estimating a BRDF with multiple specular lobes. In *CVPR*, 2010.
- [25] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: recovering reflectance models of real scenes from photographs. *SIGGRAPH*, 1999.
- [26] M. Yuan. On the identifiability of additive index models, 2010.