

Photometric Stereo using Constrained Bivariate Regression for General Isotropic Surfaces

Satoshi Ikehata*
The University of Tokyo, Japan
ikehata@hal.t.u-tokyo.ac.jp

Kiyoharu Aizawa
The University of Tokyo, Japan
aizawa@hal.t.u-tokyo.ac.jp

Abstract

This paper presents a photometric stereo method that is purely pixelwise and handles general isotropic surfaces in a stable manner. Following the recently proposed sum-of-lobes representation of the isotropic reflectance function, we constructed a constrained bivariate regression problem where the regression function is approximated by smooth, bivariate Bernstein polynomials. The unknown normal vector was separated from the unknown reflectance function by considering the inverse representation of the image formation process, and then we could accurately compute the unknown surface normals by solving a simple and efficient quadratic programming problem. Extensive evaluations that showed the state-of-the-art performance using both synthetic and real-world images were performed.

1. Introduction

Photometric stereo estimates the surface normals of an object from appearance variations under different lighting conditions. Since Woodham [20] first introduced the photometric stereo for Lambertian scenes, the extension to real-world objects which exhibit diverse appearances beyond Lambertian model has drawn significant interest.

Traditionally, certain parametric reflectance models are assumed to inversely solve the photometric stereo problem. One of the most popular classes assumes a basic Lambertian model but augmented with outlier detection for handling all non-Lambertian regions of the scene [21, 11]. This strategy is numerically stable and relatively robust to shadows and image noises, but complex reflections such as rough specularities can be highly disruptive. In contrast to the first class, a second class of methods treats non-Lambertian reflections as inliers using the nonlinear bidirectional reflectance distribution function (BRDF) [9, 16]. While these methods are more capable of handling a wide variety of objects includ-

ing rough surfaces, they may suffer from numerical instabilities derived from the complex nonlinear optimization.

Instead of explicitly modeling the parametric form of reflectance, the monotonicity property of reflectance has been recently integrated in the photometric stereo problem [2, 10, 17]. Chandraker and Ramamoorthi [6] show that an isotropic BRDF consists of a sum of lobes whose contribution to the reflected intensity decreases monotonically as the surface normal deviates away from the direction where the reflectance lobe is concentrated (*i.e.*, referred to as a *preferred direction*). Following this remark, the surface normal has been recovered utilizing the monotonicity of the reflectance function under the assumption that the number of lobes is one and its preferred directions are known (*e.g.*, the lighting direction in [10] and half vector in [2, 17]). While effective, these methods are highly disruptive when the assumption on the preferred direction is incorrect or the reflectance function is composed of two or more lobes. Furthermore, to our knowledge, simultaneous estimation of both azimuth and elevation angles has never been achieved by enforcing the monotonicity of a reflectance function that has a preferred direction that is different from the lighting vector (note that [2, 17] assume that the azimuth angle of the surface normal is known).

This paper presents a photometric stereo algorithm for accurate estimation of surface normals of a general isotropic scene by enforcing the monotonicity of a reflectance function with an unknown lobe number and preferred directions. For this purpose, the bivariate reflectance model is developed in Section 2 where pixelwise appearances are well approximated by a bivariate monotonic (and therefore invertible) smooth function of the dot products between the surface normal and the lighting direction, and between the lighting and viewing directions. We may then consider the inverse representation of the image formation process, where the unknown normal vector is separated from the unknown monotonic inverse reflectance function. By parameterizing the latter using the Bernstein polynomials [13], we obtain a set of constrained linear equations in both the surface normals and reflectance parameters, leading to a sim-

*This work was supported by the Grants-in-Aid for JSPS Fellows (No. 248615).

ple, quadratic programming problem.

The proposed framework benefits from an efficient pixelwise optimization that is easily amenable to parallel processing and does not require typical smoothness constraints for both object structure and reflectance, which could disrupt the recovery of fine details.

2. Photometric stereo using constrained bivariate regression

In this section, we formulate the photometric stereo as a constrained bivariate regression problem. Henceforth we rely on the following assumptions:

- (1) The relative position between the camera and the object is fixed across all images.
- (2) The object is illuminated by a point light source at infinity from varying and known directions.
- (3) The camera view is orthographic, and the radiometric response function is linear.

2.1. Problem Statement

Diverse appearances of real-world objects can be encoded by a BRDF (ρ) that relates the observed intensity I at a given point on the object to the associated surface normal $\mathbf{n} \in \mathbb{R}^3$, the incoming lighting direction $\mathbf{l} \in \mathbb{R}^3$, and the outgoing viewing direction $\mathbf{v} \in \mathbb{R}^3$ via

$$I = \rho(\mathbf{n}, \mathbf{l}, \mathbf{v}) \max(\mathbf{n}^T \mathbf{l}, 0), \quad (1)$$

where $\max(\mathbf{n}^T \mathbf{l}, 0)$ accounts for attached shadows. There is a problem for photometric stereo in recovering the surface normal \mathbf{n} of a scene by inversely solving Eq. (1) from a collection of m observations under different lighting conditions. Note that except for uncalibrated photometric stereo problems such as [8], \mathbf{l} and \mathbf{v} are usually known.

Recently Chandraker and Ramamoorthi [6] have presented a semiparametric model of an isotropic BRDF that is represented as a sum of K different functions giving

$$\rho = \sum_{k=1}^K \rho_k(\mathbf{n}^T \boldsymbol{\alpha}_k). \quad (2)$$

Here ρ_k are (unknown) nonlinear functions, and $\boldsymbol{\alpha}_k$ (i.e., $\|\boldsymbol{\alpha}_k\| = 1$) are called *preferred directions*, along which ρ_k are concentrated. It is known that physically valid reflectance functions satisfy the following requirements.

- (L1) Monotonicity: $\rho_k' > 0$.
- (L2) Nonnegativity: $\rho_k \geq 0$.
- (L3) Passing thorough the origin: $\rho_k(0) = 0$.

Chandraker and Ramamoorthi [6] have shown that inversely solving Eq. (2) under known surface normals gives good estimates of a wide variety of isotropic BRDFs without suffering from the curse of dimensionality. Unfortunately, however, solving Eq. (2) directly is prohibitively dif-

ficult in the context of the photometric stereo problem because of the numerous unknown parameters, some of which are coincident in the same term (i.e., $\mathbf{n}, \boldsymbol{\alpha}_k, \rho_k$). Therefore, most of the conventional photometric stereo algorithms have assumed that the dominant preferred direction of the reflectance function is unique and known [10, 2, 17].

Instead of this approach, we only assume that the preferred direction ($\boldsymbol{\alpha}_k$) of each function (ρ_k) is lying on the plane spanned by the lighting and viewing directions as

$$\boldsymbol{\alpha}_k = \frac{p_k \mathbf{l} + q_k \mathbf{v}}{\|p_k \mathbf{l} + q_k \mathbf{v}\|}, \quad (3)$$

where p_k and q_k are nonnegative unknown values (i.e., $p_k \geq 0$ and $q_k \geq 0$). The degree of freedom of $\boldsymbol{\alpha}_k$ is actually 1 because $\|\boldsymbol{\alpha}_k\| = 1$.¹ Then, this assumption provides us following important result.

Theorem: Suppose there is no shadow at a surface point (i.e., $\forall i \mathbf{n}^T \mathbf{l}_i, \mathbf{l}_i^T \mathbf{v} \geq 0$ and $I_i \geq 0$, where i is the index of the light) and $\rho(\mathbf{n}, \mathbf{l}, \mathbf{v})$ in Eq. (1) has the form of Eq. (2), whose parameters satisfy the requirements of a physically valid BRDF (L1)–(L3) and Eq. (3). Then, it is guaranteed that there exists at least one continuous bivariate function $f(x, y) \forall x, y \in [0, 1]$, which satisfies $f \geq 0, \partial f / \partial x > 0, \partial f / \partial y \leq 0$ and $\forall i I_i = f(\mathbf{n}^T \mathbf{l}_i, \mathbf{l}_i^T \mathbf{v})$.

Proof: From Eq. (3), $\mathbf{n}^T \boldsymbol{\alpha}_k$ is transformed into

$$\mathbf{n}^T \boldsymbol{\alpha}_k = \frac{p_k \mathbf{n}^T \mathbf{l} + q_k \mathbf{n}^T \mathbf{v}}{\sqrt{p_k^2 + q_k^2 + 2p_k q_k \mathbf{l}^T \mathbf{v}}}. \quad (4)$$

Here we used $\|\mathbf{l}\| = \|\mathbf{v}\| = 1$. Eq. (4) illustrates that $\mathbf{n}^T \boldsymbol{\alpha}_k$ is non-decreasing for $\mathbf{n}^T \mathbf{l}$ with fixed $\mathbf{l}^T \mathbf{v}$ and non-increasing for $\mathbf{l}^T \mathbf{v}$ with fixed $\mathbf{n}^T \mathbf{l}$ since p, q are non-negative constant values and $\mathbf{n}^T \mathbf{v}$ is constant over different lightings. From (L1), it is guaranteed that each $\rho_k(\mathbf{n}^T \boldsymbol{\alpha}_k)$ is also non-decreasing/non-increasing for $\mathbf{n}^T \mathbf{l}$ and $\mathbf{l}^T \mathbf{v}$ when either of them is fixed. Integrating these results into Eq. (1) and Eq. (2), it is proved that I is monotonic increasing for $\mathbf{n}^T \mathbf{l}$ with fixed $\mathbf{l}^T \mathbf{v}$ and non-increasing for $\mathbf{l}^T \mathbf{v}$ with fixed $\mathbf{n}^T \mathbf{l}$, which implies we can always define continuous functions $f(x, y)$ which satisfy $f \geq 0, \partial f / \partial x > 0, \partial f / \partial y \leq 0$ and $\forall i I_i = f(\mathbf{n}^T \mathbf{l}_i, \mathbf{l}_i^T \mathbf{v})$ since $I_i \geq 0$ and $\forall i 0 \leq \mathbf{n}^T \mathbf{l}_i, \mathbf{l}_i^T \mathbf{v} \leq 1$.

We illustrate this theorem in Fig. 1-(a). We note that assuming that $f(x, y)$ is always passing through the y -axis (i.e., $f(0, y) = 0$) does not limit any kind of isotropic BRDF represented by Eq. (1) because $I = \mathbf{n}^T \mathbf{l} \rho$.

Following this theorem, we formulate the photometric stereo as a constrained bivariate regression problem whose

¹Note that this assumption does not violate most existing BRDF models such as [7, 19, 15, 12, 3].

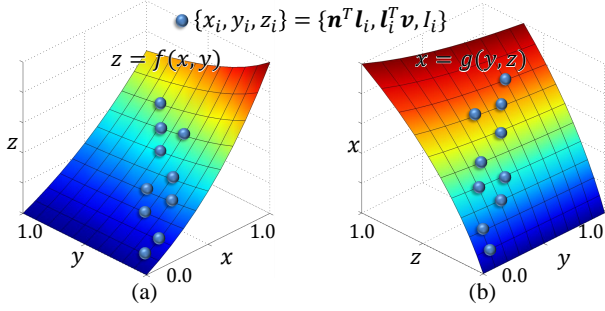


Figure 1. (a) Collections of $\{x, y, z\} = \{\mathbf{n}^T \mathbf{l}_i, \mathbf{l}_i^T \mathbf{v}, I_i\}$ are lying on a continuous function of $z = f(x, y)$ which satisfies $\partial f / \partial x > 0$, $\partial f / \partial y \leq 0$ and $f(0, y) = 0$. (b) 3-d points lying on f are also lying on a inverse function ($x = g(y, z)$).

goal is to recover an unknown surface normal \mathbf{n} and a continuous bivariate function f from a collection of lighting directions \mathbf{l}_i and associated appearances I_i ($i = 1, \dots, m$), which satisfy the following equations and constraints:

$$I_i = f(\mathbf{n}^T \mathbf{l}_i, \mathbf{l}_i^T \mathbf{v}) \quad i = 1, \dots, m. \quad (5)$$

(L4) Monotonicity (x): $\partial f / \partial x > 0$.

(L5) Monotonicity (y): $\partial f / \partial y \leq 0$.

(L6) Nonnegativity: $f \geq 0$.

(L7) Passing through y-axis: $f(0, y) = 0$.

We call Eq. (5) the *forward bivariate reflectance model*. The major benefit of this formulation is that we do not need to explicitly approximate the number of lobes K and their preferred directions α_k . However, there is a critical issue, namely the coincidence of unknown parameters \mathbf{n} and f in the same term. We overcome this difficulty by a convenient, inverse representation of the imaging model applied through a constrained bivariate regression framework.

2.2. Inverse Bivariate Reflectance Model

Strict monotonicity of $f(x, y)$ **(L4)** guarantees the unique existence of the function giving $x = g(y, f(x, y)) = g(y, z)$, which obeys the following requirements.

(L8) Monotonicity (y): $\partial g / \partial y \geq 0$.

(L9) Monotonicity (z): $\partial g / \partial z > 0$.

(L10) Nonnegativity: $g \geq 0$.

(L11) Passing through the y-axis: $g(y, 0) = 0$.

The proof, which has been omitted for brevity, is obvious by seeing Fig. 1-(b). From the definition, each 3-D point of $\{x, y, z\} = \{\mathbf{n}^T \mathbf{l}_i, \mathbf{l}_i^T \mathbf{v}, I_i\}$ ($i = 1, \dots, m$) lying on f is also lying on g as follows.

$$\mathbf{n}^T \mathbf{l}_i = g(\mathbf{l}_i^T \mathbf{v}, I_i) \quad i = 1, \dots, m. \quad (6)$$

In contradiction to Eq. (5), we call Eq. (6) the *inverse bivariate reflectance model*. Our goal is now updated to recover the surface normal \mathbf{n} and a continuous bivariate function g

with some shape restrictions **(L8)–(L11)**. The fundamental advantage of Eq. (6) is that unknown variables of \mathbf{n} and g are separated, which contributes to simplifying the problem.

Constraints on g limit the solution space of Eq. (6), but there are still multiple feasible solutions of a pair of \mathbf{n} and g because $\{\mathbf{l}_i^T \mathbf{v}, I_i\}$ are sparsely distributed on the valid range of $\{y, z\}$. To reduce the inherent ambiguity of the problem, we further assume a parametric model of the inverse bivariate reflectance function $g(y, z)$. Given that the left-hand side of Eq. (6) is linear in the unknown normal vector \mathbf{n} , for computational simplicity we would like to impose a similar linearity on the right-hand side in our parameterized representation of $g(\mathbf{l}_i^T \mathbf{v}, I_i)$ (we have omitted subscripts for simplicity). For this purpose, we then choose to express $g(\mathbf{l}_i^T \mathbf{v}, I_i)$ as a summation over p known nonlinear basis functions $g_k(\mathbf{l}_i^T \mathbf{v}, I_i)$ weighted by an unknown coefficient vector $\beta \triangleq [\beta_1, \dots, \beta_p]^T$, leading to the representation

$$g(\mathbf{l}_i^T \mathbf{v}, I_i) = \sum_{k=1}^p \beta_k g_k(\mathbf{l}_i^T \mathbf{v}, I_i). \quad (7)$$

While nonlinear in $\mathbf{l}_i^T \mathbf{v}$ and I_i , $g(\mathbf{l}_i^T \mathbf{v}, I_i)$ is clearly linear in β . We need to choose g_k carefully because estimating a multivariate regression function subject to shape restrictions with compact support is challenging and usually very time-consuming [18]. Here we adopt the bivariate Bernstein polynomials [13], where the shape-restricted regression function estimation is shown to be the solution of a quadratic programming problem [5, 18], making it computationally attractive. Furthermore, the Bernstein polynomials approximation naturally selects smooth functions with little computational effort, unlike other nonparametric regression functions (*e.g.*, smoothing spline [4]), which implicitly enforces the smoothness of BRDF as in [2].

Bivariate Bernstein polynomials [13] are composed of multiple basis functions of the form

$$b_{k_1, k_2}(x_1, x_2, N_1, N_2) = b_{k_1}(x_1, N_1) b_{k_2}(x_2, N_2), \quad (8)$$

$$b_{k_i}(x_i, N_i) = \binom{N_i}{k_i} x_i^{k_i} (1 - x_i)^{N_i - k_i} \quad (i = 1, 2),$$

where $0 \leq x_i \leq 1$ and N_i is the order of the polynomial as for x_i , which will be chosen as a function of the sample size m (*e.g.*, $N_i = o(m^{\gamma_i})$ with $\gamma_i > 0$ suitably chosen via the popular V-fold cross-validation method as shown in [18]). We transform I_i ($i = 1, \dots, m$) to lie in the unit $[0, 1]$ via a simple linear equation as $I_i \leftarrow I_i / \max(I)$. Note that $\mathbf{l}_i^T \mathbf{v}$ naturally lies in $[0, 1]$ as we only consider the case $\mathbf{v} = [0, 0, 1]^T$ and $l_z > 0$. Then, the bivariate Bernstein polynomials approximation of g is represented as

$$x = g(y, z) = \beta^T \mathbf{b}_{N_y, N_z}(y, z)$$

$$= \sum_{k_y=0}^{N_y} \sum_{k_z=0}^{N_z} \beta_{k_y, k_z} b_{k_y, k_z}(y, z, N_y, N_z), \quad (9)$$

where $\mathbf{b}_{N_y, N_z} \triangleq [b_{0,0}, \dots, b_{N_y, N_z}]^T \in \mathbb{R}^{(N_y+1)(N_z+1) \times 1}$ and $\boldsymbol{\beta} \triangleq [\beta_{0,0}, \dots, \beta_{N_y, N_z}]^T \in \mathbb{R}^{(N_y+1)(N_z+1) \times 1}$.

Unlike the B-splines procedure (which may require quadratic constraints on the coefficients) [4], shape restrictions (e.g., monotonicity, nonnegativity) on Eq. (9) are easily encoded via linear constraints, that is $A\boldsymbol{\beta} \geq 0$ and $C\boldsymbol{\beta} = 0$, where A, C are shape restriction matrices. Following [18], the shape restriction matrices required for our problem are defined as follows.

(1) **Monotonicity:** $\partial g / \partial y \geq 0$ and $\partial g / \partial z \geq 0$. **(L8), (L9)**
The first-order partial derivatives of g with respect to y in Eq. (9) can be represented as

$$\begin{aligned} & \partial g(y, z) / \partial y \\ &= N_y \sum_{k_z=0}^{N_z} \sum_{k_y=0}^{N_y-1} (\beta_{k_{y+1}, k_z} - \beta_{k_y, k_z}) b_{k_y, k_z}(y, z, N_y - 1, N_z). \end{aligned} \quad (10)$$

From the definition in Eq. (8), it is easy to show that all Bernstein basis polynomials are nonnegative with respect to $0 \leq y, z \leq 1$. Hence, the non-decreasing constraint (i.e., $\partial g / \partial y \geq 0$) is simply achieved by enforcing $\beta_{k_{y+1}, k_z} \geq \beta_{k_y, k_z}$. The non-decreasing constraint with respect to z (i.e., $\partial g / \partial z \geq 0$) is also achieved in the same manner. The restriction matrix for linear constraint $A_{mono}\boldsymbol{\beta} \geq 0$ is represented as $A_{mono} = [A_y^T A_z^T]^T$, which is composed of submatrices $A_y \in \mathbb{R}^{N_y(N_z+1) \times (N_y+1)(N_z+1)}$ and $A_z \in \mathbb{R}^{N_z(N_y+1) \times (N_y+1)(N_z+1)}$, where A_r ensures the monotonicity of the function with respect to r .² Note that the strict monotonicity constraint $\partial g / \partial z > 0$ was eased to $\partial g / \partial z \geq 0$ for computational simplicity.

(2) **Nonnegativity:** $g \geq 0$. **(L10)**

The nonnegativity of g is guaranteed when $\forall i \beta_i \geq 0$. Hence, the restriction matrix for a linear constraint $A_{nonneg}\boldsymbol{\beta} \geq 0$ is as $A_{nonneg} \triangleq \text{diag}([1, \dots, 1]) \in \mathbb{R}^{(N_y+1)(N_z+1) \times (N_y+1)(N_z+1)}$.

(3) **Passing through the y -axis:** $g(y, 0) = 0$. **(L11)**

From the definition in Eq. (8), $b_{k_y, k_z}(y, 0) = 0$ for all $k_z \neq 0$. Therefore $g(y, 0) = \sum_{k_y=0}^{N_y} \beta_{k_y, 0} b_{k_y, 0}(y, 0, N_y, N_z)$ becomes zero for all y when $\forall k_y \beta_{k_y, 0} = 0$. This constraint is encoded via a linear constraint $C\boldsymbol{\beta} = 0$, where $C \in \mathbb{R}^{(N_y+1)(N_z+1) \times (N_y+1)(N_z+1)} \triangleq \text{diag}([1, 0, \dots, 1, 0, \dots, 1, \dots])$ with N_z of 0 between 1.

2.3. Solution Method

By substituting Eq. (9) into the inverse bivariate reflectance model, Eq. (6) becomes

$$\mathbf{n}^T \mathbf{l}_i = \boldsymbol{\beta}^T \mathbf{b}_{N_y, N_z}(\mathbf{l}_i^T \mathbf{v}, I_i) \quad i = 1, \dots, m, \quad (11)$$

²The concrete form of the shape restriction matrix is included in the supplementary material.

where the coefficients of Bernstein polynomials ($\boldsymbol{\beta}$) are restricted via the following equation:

$$A\boldsymbol{\beta} = \begin{bmatrix} A_{mono} \\ A_{nonneg} \end{bmatrix} \boldsymbol{\beta} \geq 0, \quad C\boldsymbol{\beta} = 0. \quad (12)$$

Collecting variations of observations at the same pixel under different lighting directions, Eq. (11) can be merged into following linear problem:

$$L^T \mathbf{n} = B^T \boldsymbol{\beta}. \quad (13)$$

Here, $B \triangleq [\mathbf{b}_{N_y, N_z}(\mathbf{l}_1^T \mathbf{v}, I_1), \dots, \mathbf{b}_{N_y, N_z}(\mathbf{l}_m^T \mathbf{v}, I_m)]$ and $L \triangleq [\mathbf{l}_1, \dots, \mathbf{l}_m]$. By merging unknown variables ($\mathbf{n}, \boldsymbol{\beta}$), this problem is transformed into

$$P\mathbf{x} = [L \quad -B]^T \mathbf{x} = 0, \quad (14)$$

where $\mathbf{x} \triangleq [n_x, n_y, n_z, \beta_{0,0}, \dots, \beta_{N_y, N_z}]^T$ and n_x, n_y, n_z are the three elements of the surface normal. Without loss of generality, we may avoid the degenerate $\mathbf{x} = 0$ solution to Eq. (14) by constraining $\sum_i x_i = 1$, which implies $\mathbf{c}^T \mathbf{x} = 1$ where $\mathbf{c} = [1, \dots, 1]^T$.

Given the appearance variations (I_1, I_2, \dots, I_m) under different known lighting conditions ($\mathbf{l}_1, \mathbf{l}_2, \dots, \mathbf{l}_m$), the optimal surface normal (\mathbf{n}) and model parameters ($\boldsymbol{\beta}$) are recovered by solving the constrained linear problem:

$$\min_{\mathbf{x}} \|\mathbf{P}\mathbf{x}\|_2^2, \quad \text{s.t. } \tilde{A}\mathbf{x} \geq 0 \text{ and } \tilde{C}\mathbf{x} = 0, \quad (15)$$

where $\tilde{A} \triangleq [\mathbf{0} \quad A]$ and $\tilde{C} \triangleq \begin{bmatrix} \mathbf{c}^T \\ \mathbf{0} \quad C \end{bmatrix}$. Note that Eq. (15) can be effectively solved by the general quadratic programming.

3. Handling Retroreflective Materials

If the reflectance of a target object obeys Eq. (2), our method reasonably recovers the surface normal of the object by solving Eq. (15). However, one limitation of our method is that this assumption is not satisfied in the presence of retroreflections that are often observed on rough surfaces because Eq. (2) does not have the ability to represent this kind of reflection. In the presence of retroreflections, our surface normal estimation fails because of the violation of **(L5)** by the behavior of retroreflections in that the power of reflections increases as $\mathbf{l}^T \mathbf{v}$ increases. While it may limit available materials, fortunately we have found that our method practically handles retroreflective materials by simply reversing the direction of the monotonicity constraint on $\mathbf{l}^T \mathbf{v}$ in Eq. (15) (i.e., use **(L8)**' $\partial g / \partial y \leq 0$ instead of **(L8)**) because the retroreflections do not affect the monotonicity for $\mathbf{n}^T \mathbf{l}$ while they unify the direction of the monotonicity for $\mathbf{l}^T \mathbf{v}$ over the BRDF space. The problem of course is that we do not know whether the material is retroreflective or not. To overcome this difficulty,

we present a practical approach for handling both non- and retroreflective materials. The important observation is that when we incorrectly constrain the problem, the regression usually fails. Therefore, after we have the regression outputs under both constraints, we can judge which constraint was optimal by examining regression errors.

However, we have empirically found that comparing regression errors in Eq. (15) does not work since the flexible Bernstein polynomials are generally well fitted to observations even though the constraint was not correct; instead, we compute the following linear regression error E for choosing the optimal solution:

$$\tilde{a} = \arg \min_a \sum_{i=1}^m \|\hat{\mathbf{n}}^T \mathbf{l}_i - a I_i\|_2^2, \quad (16)$$

$$E = \sum_{i=1}^m \|\hat{\mathbf{n}}^T \mathbf{l}_i - \tilde{a} I_i\|_2^2. \quad (17)$$

Here, $\hat{\mathbf{n}}$ is a recovered surface normal by solving Eq. (15) under the monotonicity constraint for $\mathbf{l}^T \mathbf{v}$ in either of two directions. We simply choose the direction for which E is smaller. This strategy is very simple but very efficient as we will show in Section 4.

4. Experimental Results

In this section, we evaluate our method on synthetic and real image data. All experiments were performed on an Intel Core i7-2640M (2.80GHz, single thread) machine with 8GB RAM and were implemented in MATLAB. For the quantitative evaluation, we generated 32-bit HDR images of a sphere (256×256) with foreground masks under different BRDF settings: (A) common physical or phenomenological BRDF (Cook-Torrance [7], Ward [19], Lafortune [12], Oren-Nayar [15] and Ashikhmin-Shirley [3]) and (B) the measured MERL BRDF database [14]. Lighting directions were randomly selected from a hemisphere with the object placed at the center. Additionally, for the third dataset, denoted (C), we used real images for qualitatively evaluating our method in practical situations. For each dataset, shadows were removed via simple thresholding as in other works such as [16] (Thresholds for shadow removal were fixed over algorithms but varied over objects manually). Because ground truth surface normals are provided in datasets (A) and (B), we quantitatively evaluated our method by the angular error between the recovered normal map and the ground truth when using these datasets.

4.1. Evaluation with Synthesized BRDF

We evaluated the performance using the synthesized images in dataset (A) generated under 100 different lightings using five common BRDFs³. Here we compared

³Details of BRDF models are described in the supplementary material.

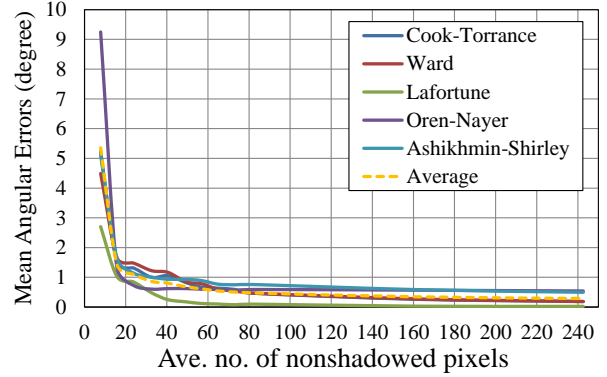


Figure 3. Experimental results of dataset (A) with a varying number of images. For a fair evaluation, we display the average number of nonshadowed pixels that join our algorithm on the x -axis instead of the number of images.




Figure 4. Comparison with [17]. The values in parentheses are the radian values added to the true azimuth angle map. The noisy azimuth maps were used in [17]. The sign was chosen randomly.




Figure 5. Evaluation of computational time. We illustrate the per-pixel computational time for each combination of the number of Bernstein basis functions and lightings.

our method with the standard Lambertian least-squares-regression-based approach [20] (LS) and a recent Lambertian sparse-regression-based approach implemented with the sparse Bayesian learning [11] (SBL) (λ is fixed by




Figure 2. Experimental results of dataset (A) with five different BRDFs with varying frequencies in observations.

10^{-6}). Our method was also compared with a recent parametric non-Lambertian photometric stereo method with the biquadratic reflectance model [16] (BQ). In this experiment, we fixed (N_y, N_z) in Eq. (11) by (1, 5) to examine the robustness of our method against these parameters.

Evaluation with varying frequencies

We first evaluated our algorithm using observations with varying frequency to present the flexibility of our model in comparison with models used in previous works (*e.g.*, BQ[16]) that only work for low-frequency reflectance values. Here high-frequency specularities were discarded by using the nonshadowed pixels with intensities that were ranked below the $T_{low}\%$ ($10 \leq T_{low} \leq 100$).

The result is illustrated in Fig. 2. Overall, we observed that our method outperformed other algorithms for almost all frequencies for the Lafortune, Ashikhmin–Shirley, and Oren–Nayar models, and performed competitively with SBL for the Cook–Torrance and Ward models. Interestingly, because of our inverse reflectance model and the simple $l^T v$ constraint selection strategy, our method works well for the Oren–Nayar model which exhibits strong retroreflective reflections, and it also works well when all the frequencies are included unlike other methods⁴. Unfortunately, however, our method did not work when the number of images was very small, as will be discussed below.

Valid number of input images

We also evaluated our algorithm using a varying number of images to find the valid number required for effective recovery. The results are displayed in Fig. 3. We observed that the minimum number of images required to make the algo-

⁴The result of BQ is similar as LS since the optimization procedure of BQ is nonlinear and strongly affected by the initial estimation by LS.

rithm work was around 20 and more than 60 were required for stable reconstruction because our method could suffer from the over-fitting when the number of the input images was very small.

Comparison with another monotonicity-based approach

In this section, we compared our method with the recent elevation-angle estimation algorithm [17] assuming that the dominant reflectance lobe was pointing at the half-vector direction. We generated a ground truth azimuth angle map and azimuth angle maps with some radian errors because [17] requires an azimuth angle map as input whereas our method simultaneously recovers all elements in the normal. The radian errors varied from 0.05 to 0.3 were added to the ground truth azimuth angle values, where the sign of the error was chosen randomly.

The results are illustrated in Fig. 4. We observed that when true azimuth angles were given, [17] outperformed our method in the Cook–Torrance and Ward datasets. However, as the amount of errors increased, the differences became smaller and finally our method outperformed [17]. As we expected, [17] did not work for the Lafortune and Ashikhmin–Shirley models because these models violate their assumption.

Evaluation of computational time

Here, we examined the computational time required for our computation. We tried various combinations of N_y and N_z and m in Eq. (9) and solved our optimization problem Eq. (11) using the *lsqlin* function in MATLAB.

The evaluation results are illustrated in Fig. 5. Here we present a per-pixel computational time to solve *one* optimization problem. Therefore, the actual computational time was twice that in the figure as we applied our algorithm

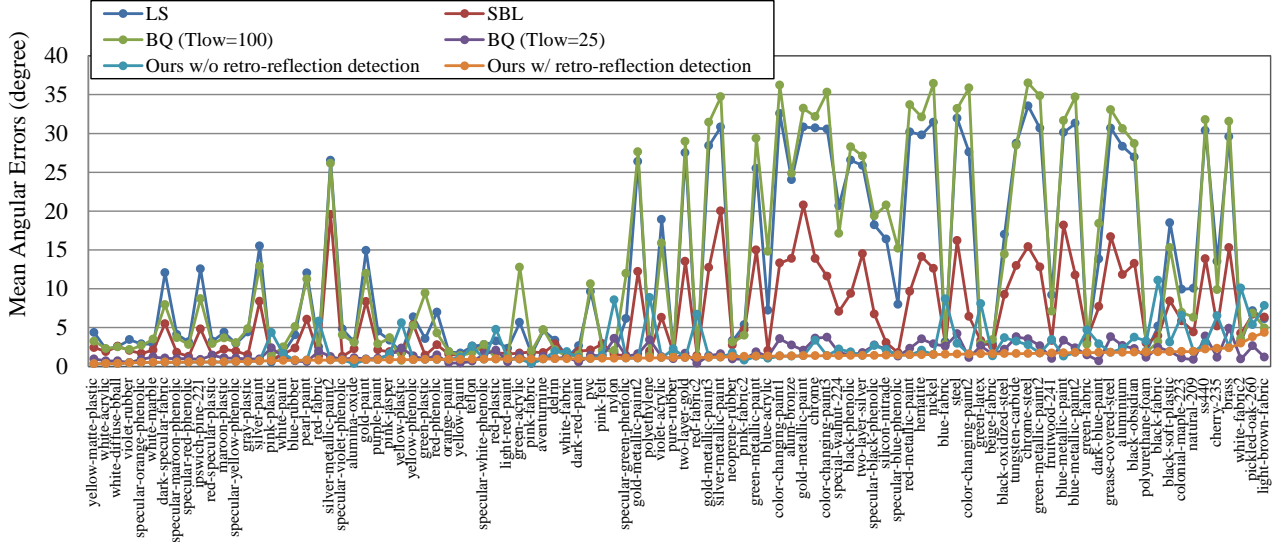


Figure 6. Comparison between different methods with dataset (B). Results are aligned in ascending order of the mean angular error of our method with the retroreflection detection algorithm.

twice to distinguish retroreflective materials. We observed that the computational complexity depends on the number of basis functions rather than the number of lightings.

4.2. Evaluation with Measured BRDF

Here we evaluated the performance of our method to the dataset (B). We generated images under 300 different lightings for 100 different materials from the MERL BRDF database [14]. In this experiment, our method was also compared with LS [20], SBL [11], and BQ [16]. LS, SBL, and our method used all the frequencies in observations (*i.e.*, $T_{low} = 100$) whereas only BQ was performed with both $T_{low} = 25$ and $T_{low} = 100$ as that model was originally designed to represent the low-frequency observations. In this experiment, we fixed $N_y = 3$ and $N_z = 5$ and performed our method with/without the retroreflection detection algorithm described in Section 3 to verify the effectiveness of this process.

The results are illustrated in Fig. 6. We observed that our method with our efficient retroreflection detection outperformed other algorithms for most of the materials. BQ ($T_{low} = 25$) is more effective for some materials, but we emphasize that our method was capable of handling all frequencies in observations because of our flexible reflectance model, while BQ ($T_{low} = 100$) does not work for most materials. We also observed that the angular errors of our method without retroreflection detection were relatively larger than that of our method with retroreflection detection for materials exhibiting strong retroreflections (*e.g.*, MERL fabrics), which indicated that our retroreflection detection algorithm worked quite well for those materials. Finally, the

average angular errors over 100 materials were 12.5 (LS), 6.2 (SBL), 13.1 (BQ, $T_{low} = 100$), 1.7 (BQ, $T_{low} = 25$), 2.4 (our method without retroreflection detection) and 1.2 (our method with retroreflection detection), respectively.

4.3. Qualitative Evaluation with Real Images

We also evaluated our algorithm using real images: (1) 100 images of *two-face* (this dataset is from [21]), (2) 100 images of *doraemon*, and (3) 44 images of *fatguy* (these two datasets are from [11]). In this experiment, we compared our algorithm with LS [20], SBL [11] (λ was fixed by 10^{-1}) and BQ [16] (for $T_{low} = 25$ and 100). Note that our method used a retro-reflective detection scheme and we fixed (N_y, N_z) by (1, 5). The threshold of the shadow removal for each dataset was chosen manually but fixed over algorithms.

The experimental results are illustrated in Fig. 7. Here we show both recovered surface normal maps and surface meshes reconstructed by a poisson solver [1] with a fixed scale. We observe that our method succeeded to estimate smoother and more reasonable normal maps and surface meshes. We also observe that BQ ($T_{low} = 25$) worked poor for those datasets since shadows could not be completely removed by a simple thresholding therefore the low-frequency component in the observation was not reliable. In contrast to that, our method performed well since our method could account all observations without discarding the informative high-frequency component.

5. Conclusion

In this paper, we have proposed the constrained bivariate regression-based photometric stereo, which worked for




Figure 7. Experimental results using real data (*two-face*, *doraemon* and *fatguy*). We illustrate (a) example images of target objects, and normal maps recovered by (b) LS, (d) SBL, (f) BQ ($T_{low} = 100$), (h) BQ ($T_{low} = 25$) and (j) Ours (with a retro-reflective detection). We also show surface meshes generated from normal maps in (c), (e), (g), (i) and (k).

various kinds of isotropic surfaces by exploiting various conditions shared among physically valid BRDFs. Our detailed experimental results have shown the state-of-the-art performance of our method for both synthetic and real data. The current limitation is that the proposed method enforces a global monotonic constraint with regard to $l^T v$, which might not be true for surfaces with both the retro and specular reflections though they are rarely observed in the real world. Another limitation is that we assumed that shadows were discarded from images in advance, and this might be impractical in real scenes. To ease this condition, we are interested in incorporating the data cleansing scheme in a similar manner as [11].

References

- [1] A. Agrawal, R. Raskar, and R. Chellappa. What is the range of surface reconstructions from a gradient field? In *Proc. ECCV*, 2006. 7
- [2] N. Alldrin, T. Zickler, and D. Kriegman. Photometric stereo with non-parametric and spatially-varying reflectance. In *Proc. CVPR*, 2008. 1, 2, 3
- [3] M. Ashikhmin and P. Shirley. An isotropic phong brdf model. *Journal on Graphics Tools*, 5(2):25–32, 2000. 2, 5
- [4] K. Bollaerts, P. H. Eilers, and I. Mechelen. Simple and multiple p-splines regression with shape constraints. *British J. Math. Statist. Psych.*, 59(2):451–469, 2006. 3, 4
- [5] P. Chak. Semi-nonparametric estimation with bernstein polynomials. *Economics Letters*, 89(2):153–156, 2005. 3
- [6] M. Chandraker and R. Ramamoorthi. What an image reveals about material reflectance. In *Proc. ICCV*, 2011. 1, 2
- [7] R. Cook and K. Torrance. A reflectance model for computer graphics. *ACM Trans. on Graph.*, 15(4):307–316, 1981. 2, 5
- [8] P. Favaro and T. Papadimitri. A closed-form solution to uncalibrated photometric stereo via diffuse maxima. In *Proc. CVPR*, 2012. 2
- [9] D. Goldman, B. Curless, A. Hertzmann, and S. Seitz. Shape and spatially-varying brdfs from photometric stereo. In *Proc. ICCV*, October 2005. 1
- [10] T. Higo, Y. Matsushita, and K. Ikeuchi. Consensus photometric stereo. In *Proc. CVPR*, 2010. 1, 2
- [11] S. Ikehata, D. Wipf, Y. Matsushita, and K. Aizawa. Robust photometric stereo using sparse regression. In *Proc. CVPR*, 2012. 1, 5, 7, 8
- [12] E. Lafortune, S.-C. Foo, K. Torrance, and D. Greenberg. Non-linear approximation of reflectance functions. In *Proc. ACM SIGGRAPH*, 1997. 2, 5
- [13] G. G. Lorentz. *Bernstein Polynomials*. Chelsea Publishing Company, New York, 1986. 1, 3
- [14] W. Matusik, H. Pfister, M. Brand, and L. McMillan. A data-driven reflectance model. *ACM Trans. on Graph.*, 22(3):759–769, 2003. 5, 7
- [15] M. Oren and S. Nayar. Generalization of lambert’s reflectance model. In *In Proc. of the 21st annual conference on Computer graphics and interactive techniques*, 1994. 2, 5
- [16] B. Shi, P. Tan, Y. Matsushita, and K. Ikeuchi. A biquadratic reflectance model for radiometric image analysis. In *Proc. CVPR*, 2012. 1, 5, 6, 7
- [17] B. Shi, P. Tan, Y. Matsushita, and K. Ikeuchi. Elevation angle from reflectance monotonicity. In *Proc. ECCV*, 2012. 1, 2, 5, 6
- [18] J. Wang. *Shape restricted nonparametric regression with Bernstein polynomials*. PhD thesis, North Carolina State University, 2012. 3, 4
- [19] G. Ward. Measuring and modeling anisotropic reflection. *Computer Graphics*, 26(2):265–272, 1992. 2, 5
- [20] P. Woodham. Photometric method for determining surface orientation from multiple images. *Opt. Engg.*, 19(1):139–144, 1980. 1, 5, 7
- [21] L. Wu, A. Ganesh, B. Shi, Y. Matsushita, Y. Wang, and Y. Ma. Robust photometric stereo via low-rank matrix completion and recovery. In *Proc. ACCV*, 2010. 1, 7