

疎な画像列からの複雑物体表面の推定に関する研究

沈 黎[†] 町田 貴史[‡] 竹村 治雄[‡]

[†]大阪大学 情報科学研究科

[‡]大阪大学 サイバーメディアセンター

E-mail: [†]shenli@lab.ime.cmc.osaka-u.ac.jp, [‡]{machida, takemura}@ime.cmc.osaka-u.ac.jp

あらずし 複雑な表面材質の推定は現実環境に対する光学的な整合性を解決するための重要な課題である。これに対する解決法としては様々な従来法が存在するが、物体表面上で任意に変化する反射特性の推定に膨大な数の画像列を必要とする問題があった。そこで本論文では、少ない画像枚数で複雑な材質の物体表面反射特性を推定する問題を解決する新しい手法を提案する。提案手法では、まず物体表面反射を分析し、拡散反射が低次元線形空間で近似できることを利用して、少ない画像から鏡面反射成分を抽出する。次に、効率よく非線形の鏡面反射成分の係数を推定することで、通常では悪条件問題となる部分を解決し、疎な画像列から複雑な表面特性を推定することが可能になる。

キーワード インパースレンダリング, 複雑な表面材質, 複雑照明条件

Spatial Appearance Acquisition of Real Objects with Sparse Images

Li Shen[†], Takashi Machida^{†‡}, Haruo Takemura^{†‡}

[†]Graduate School of Information Science and Technology, Osaka University

[‡]Cybermedia Center, Osaka University

E-mail: [†] shenli@lab.ime.cmc.osaka-u.ac.jp, [‡] {machida, takemura}@ime.cmc.osaka-u.ac.jp

Abstract Real-world objects are usually composed of different materials with complex textures. Acquiring spatial reflectance appearance of objects requires a dense set of images. Recovering spatially varying reflectance is usually considered as a numerically ill-conditioned problem for a very sparse set of images. In this paper, we address the problem, and propose the techniques for acquiring spatial appearances of real objects from sparse images. Spatially varying reflectance recovery is investigated under controlled lighting and general distant lighting conditions. Regarding controlled illumination conditions, we propose efficient algorithms to recover a complete object model, including both the shape and the spatial reflectance, from four single-light images. Regarding distant illumination conditions, we developed an algorithm to recover high-quality spatially varying BRDFs under arbitrary illumination using a sparse image set (6~10 images). The recovered spatially varying surface models can be used to create realistic renderings of the object under novel scenes.

Keyword Spatially varying reflectances, Reflectance recovery, Distant illumination

1. Introduction

The central problem in computer graphics is creating photorealistic images. In order to synthesize realistic images, physically-based rendering algorithms require accurate input models for geometry, reflective properties

and illumination. In recent years, there has been a significant interest in inverse rendering techniques to obtain realistic rendering attributes from a set of photographs. However, inverse rendering problems usually require extremely complex and costly processes. The difficulty of obtaining accu-

rate models of a scene has been a major limiting factor in realism of physically-based rendering today.

A major challenge is the acquisition of spatial reflectance appearance of objects. Real-world objects generally have spatially varying reflectance properties over surfaces. For a spatially varying material, recovery algorithms estimate the *bidirectional reflectance distribution function* (BRDF) at each point, demanding dense images and a costly fitting process. In this paper, we focus on inverse rendering problem for acquiring the spatial appearances of real objects.

1.1. Previous Work

1.1.1. Controlled illumination

Non-Lambertian BRDF is a complicated 4D function. A great deal of work on inverse rendering has been conducted in highly controlled lighting conditions which usually assume that illumination source is a single point source. Traditionally, BRDF is measured with a gonio-reflectionmeter system, positioning a point source and detector with respect to a flat material. A measure system (e.g. [12], [14], [15]) generally demands a very large set of measurements in order to capture high-frequency information, such as sharp highlights. The reconstruction of spatially varying BRDFs requires sampling 6D function, which would be tedious if not impossible. Some methods recover only a subset of the BRDF. For example, Debevec et al. [6] measured the 2D lighting-dependent variation (64*32 images) with fixed viewpoint and constructed a reflectance map for each point.

Most methods approximate reflectance using an analytic reflection model, thereby simplify the problem from recovering a 4D function to that of estimating a handful parameters. However, fitting processes still demand a large set of samples to robustly estimate the parameters. Sato et al. [19] used about 120 images to estimate spatial reflectance with Torrance-Sparrow model. The estimation of spatial reflectance from a small image set, which is in fact an ill-conditioned problem without any assumptions, is very challenging. Recently, Lensch et al. [13] described an acquisition process for spatially varying BRDFs from sparse image set (about 15-25). They partitioned the surface points into clusters corresponding to different BRDFs and built a set of basis BRDFs to describe the reflectance of each point. Their method however operates under the assumption that the surface is composed of a few different BRDFs, which cannot be applied to the surface with high-frequency variation in BRDF.

There are some methods to recover both the shape and the reflectance properties of an object using techniques based on Photometric Stereo. Hertzmann and Seitz [10] proposed the empirical methods for spatial materials recovery. Their approach requires a material segmentation and a reference object for each material to build look-up tables. Kay and Caelli [11] investigated the problem using an analytic reflection model without any assumptions. They applied nonlinear regression to a large number (thousands) of input images. Goldman et al. [9] developed a Photometric Stereo approach based on [13]. On the assumption that the surface is composed of

only one or two materials, they recovered the *spatially-varying* BRDF and the shape from a small number of images.

Coleman and Jain [5] has proposed a specularly detection method to avoid the ill-conditioned problem. Improvements were suggested by Solomon and Ikeuchi [20], and recently by Barsky and Petrou [2]. This approach is based on a strict constraint that no specular reflection spatially overlaps among the images. By treating highlights as outliers to the Lambertian model, the method excludes the highlighted pixel and recovers the surface normal with the others.

1.1.2. Complex illumination

Comparatively less work has been done under general illumination. Yu et al. [23] presented an inverse radiosity method to estimate reflectance by using 40 images taken at different viewpoints in a closed scene. High-frequency variation is allowed in the diffuse reflectance with the assumption that the specular reflectance was constant for a surface. Boivin and Gagalowicz [4] recover the uniform BRDFs of the surfaces in a scene using one image. Both methods recover BRDFs by an iterative estimation process with very costly radiosity computation.

Instead of conventional light transport computation, some methods (e.g. [1], [16], [18]) have been proposed to represent reflection in term of spherical harmonics. Westin et al. [22] proposed a method to reconstruct the BRDFs of non-Lambertian surfaces by estimating the coefficients of the spherical harmonics. However, their results shows that even the moderately complex BRDFs require large numbers of spherical harmonic basis functions to be represented correctly. To address the problem, Ramamoorthi and Hanrhan [17] presented a single-processing method of inverse rendering under distant illumination. They proposed an algorithm consisted of nested procedures with a dual angular and frequency-space representation. However, both [22] and [17] require a large number of images for the estimation of spatially varying reflectance.

We have briefly reviewed the previous work on spatial appearance acquisition. Little literature exists on spatial appearance recovery with sparse images, particularly under general illumination conditions. Previous methods demand dense input images, or constraints are imposed on materials and lighting conditions.

1.2. Overview

Recovery of spatially varying materials from sparse images is considered as an ill-conditioned problem. In this paper, we address the problem, and describe the techniques for acquiring truly spatial appearances of real objects which are efficient and require only a very small image set.

Spatially varying BRDF recovery is investigated under controlled lighting and general distant lighting conditions, respectively. We show that the high-frequency varying term of the reflection consists of the diffuse reflection component. We demonstrate the high-frequency varying diffuse term can be removed with a small set of images since the diffuse reflection can be approximated by a low-dimensional linear subspace. This allows us to extract the low-frequency varying term of the reflection and address the

ill-conditioned problem for fitting process. The contributions of the paper are:

- We represent an approach to recover both the shape and the spatially varying BRDF from four images under controlled illumination.
- We represent an approach to recovered the spatially varying BRDF from a small set of images (6–10 images) under general distant illumination.

The rest of the paper is organized as follow: Section 2 explains the reflection model and discusses the reflection of a surface with spatially varying BRDF. Section 3 describes our method of surface recovery under controlled illumination conditions. Section 4 describes our method of inverse reflectometry under general illumination conditions. The effectiveness of our algorithms is supported by experimental results presented in Section 5. Conclusions are presented in Section 6.

2. Reflection Model

In this section, we will describe the reflection model of surfaces, and discuss the problems for estimating spatially varying BRDF from sparse images.

2.1. Dichromatic Reflection Model




Figure 2 Dichromatic reflection model.

An opaque dielectric material without multi-layer structure is composed by a supporting substrate into which colored pigment particles are embedded as shown in Figure 1. The reflection of opaque inhomogeneous dielectric objects can be modeled as a linear combination of diffuse and specular reflection, which known as the *dichromatic reflection model*. The light reflected directly at surface cause specular reflection. The refracted ray is diffused by the pigment particles and produces an isotropic distribution of scattered light resulting in diffuse reflection. The BRDF f at a location x can be expressed as

$$f(x; \bar{\omega}_i, \bar{\omega}_o) = f_d(x; \bar{\omega}_i, \bar{\omega}_o) + f_s(x; \bar{\omega}_i, \bar{\omega}_o), \quad (1)$$

where $\bar{\omega}_i$ and $\bar{\omega}_o$ are the incident and outgoing directions, respectively. $f_d(x; \bar{\omega}_i, \bar{\omega}_o)$ and $f_s(x; \bar{\omega}_i, \bar{\omega}_o)$ are the diffuse and specular BRDFs, respectively. In this paper, we denote color vectors by using a bold symbol, e.g. $f = \{f_r, f_g, f_b\}$. The diffuse BRDF can be approximated by:

$$f_d(x; \bar{\omega}_i, \bar{\omega}_o) = \frac{\rho_d(x)}{\pi}, \quad (2)$$




Figure 1 The 2D transfer functions of different points on a textured surface.

where the diffuse reflectance $\rho_d(x)$ depends on the pigments, which determines the diffuse color of the material. $f_s(x; \bar{\omega}_i, \bar{\omega}_o)$ is a complex nonlinear function. Truly spatially varying BRDFs vary in both the diffuse and the specular reflectances. The diffuse reflectance usually has high-frequency of variation. However, the specular reflectance only depends on the index of refraction and the surface facet conditions, which would not change high-frequently. Therefore, the specular reflectance properties have a slowly varying nature. In real world, objects are usually composed of one or a few of different materials¹.

2.2. Reflection on Spatial Reflectance Surfaces

We consider the incident illumination is distant illumination which allows us to use the same lighting function regardless of surface location. We consider curved objects, accounting for attached but not cast shadows. Also, interreflection will be ignored.

If the light source is a uniform point light source, the radiance at x in direction $\bar{\omega}_o$ can be represented by:

$$I(x; \bar{\omega}_o) = \hat{f}(x; \bar{\omega}_i', \bar{\omega}_o') \xi, \\ = \frac{\rho_d(x)}{\pi} \max[\bar{\omega}_i \cdot \bar{n}, 0] \xi + \underbrace{f_s(x; \bar{\omega}_i', \bar{\omega}_o') \max[\bar{\omega}_i \cdot \bar{n}, 0] \xi}_{I_{spec}} \quad (3)$$

where we have mixed local (primed) and global (unprimed) coordinates, $\hat{f}(x; \bar{\omega}_i', \bar{\omega}_o') \equiv f(x; \bar{\omega}_i', \bar{\omega}_o') \max[\bar{\omega}_i' \cdot \bar{n}, 0]$ is the transfer function, which absorbs the clamped cosine term; and ξ is the intensity of the point light. \bar{n} is the surface normal. Under general distant illumination, $I(x; \bar{\omega}_o)$ can be computed by integrating over all incoming direction of the hemisphere $\Omega_{2\pi}$ as:

$$I(x; \bar{\omega}_o) = \int_{\Omega_{2\pi}} L(\bar{\omega}_i) \hat{f}(x; \bar{\omega}_i', \bar{\omega}_o') d\bar{\omega}_i \\ = \int_{\Omega_{2\pi}} L(\bar{\omega}_i) \underbrace{\frac{\rho_d(x)}{\pi} \max[\bar{\omega}_i \cdot \bar{n}, 0]}_{I_{diff}} d\bar{\omega}_i \\ + \int_{\Omega_{2\pi}} L(\bar{\omega}_i) \underbrace{f_s(x; \bar{\omega}_i', \bar{\omega}_o') \max[\bar{\omega}_i \cdot \bar{n}, 0]}_{I_{spec}} d\bar{\omega}_i \quad (4)$$

¹ In this paper, we refer to a surface which has the same specular properties as a material.

where $L(\omega_i)$ is the incident radiance.

Let us first consider the case where a surface has only one material, i.e. the specular reflectance is identical for the surface. The transfer function of a point on the surface can be expressed by replacing $f_s(x; \bar{\omega}_i, \bar{\omega}_o)$ with $f_s(\bar{\omega}_i, \bar{\omega}_o)$ as:

$$\hat{f}(x; \bar{\omega}_i, \bar{\omega}_o) = \frac{\rho_d(x)}{\pi} \max[\bar{\omega}_i \cdot \bar{n}, 0] + f_s(\bar{\omega}_i, \bar{\omega}_o) \max[\bar{\omega}_i \cdot \bar{n}, 0], \quad (5)$$

Figure 1 illustrates the transfer functions in 2D at different points on a textured surface, where θ_h is the angle between \bar{n} and the half vector of $\bar{\omega}_i$ and $\bar{\omega}_o$. As shown in Figure 1, the transfer functions with respect to the diffuse reflectance component have large variations, and those with respect to the specular reflectance component have a complicated nonlinear characteristic. Hence, it is necessary for recovery algorithms to estimate the BRDF of each point on the surface. Recovering the reflectance with sparse images is numerically ill-conditioned.

For a spatial reflectance surface, the reflection function on the surface has high-frequency variations because of its diffuse reflectance component. With sparse images, if we can extract the specular term from the various BRDFs, the estimation problem will not be ill-conditioned because the specular BRDF (*SBRDF*) has a slowly varying nature.

3. Surface Recovery under Controlled Illumination

In the section, we explain how to resolve the problem under controlled illumination. We first describe a method to extract the specular reflection from an input image set. Then, we introduce an algorithm to recover the shape and spatially varying BRDF of an object from the images. Our input is four image of an object taken under different single point sources at a fixed viewpoint. The light sources are known.

3.1. Specular Reflection Component Extraction

Under single point illumination, we can solve the problem with four images. We denote the k^{th} illuminant by the vector $\bar{\zeta}^k \equiv \zeta^k \bar{\omega}_i^k$, where ζ^k is the illumination irradiance and $\bar{\omega}_i^k$ is the incident direction. Since any three of the incident illumination directions are linearly independent, the four illuminants can be expressed as follows:

$$A^1 \bar{\zeta}^1 + A^2 \bar{\zeta}^2 + A^3 \bar{\zeta}^3 + A^4 \bar{\zeta}^4 = \mathbf{0}, \quad (6)$$

or as:

$$\bar{A} \cdot [\bar{\zeta}] = \mathbf{0}, \quad (7)$$

where $[\bar{\zeta}] = (\bar{\zeta}^1 \cdots \bar{\zeta}^4)^T$ is the illumination matrix, the coefficient vector $\bar{A} = (A^1 \cdots A^4)$ can be computed directly from the known illumination for each color channel. Because, for an unshadowed point, the diffuse reflection is a linear function of its incident illumination, we can derive from Eq. 7 that any non-shadowed Lambertian set of irradiances $\bar{I}_{diff} = (I_{diff}^1 \cdots I_{diff}^4)$ should satisfy:

$$\bar{A} \cdot \bar{I}_{diff}(x) = \frac{\rho_d(x)}{\pi} \bar{A} \cdot [\bar{\zeta}] \cdot \bar{n} = \mathbf{0}. \quad (8)$$




Figure 3 Extraction of specular reflection components. Left: Four input images of a textured ball. Right: Result of $\bar{A} \cdot \bar{I}$.

Therefore, in accordance with Eq. 3, we can remove the diffuse terms of a non-shadowed irradiance set $\bar{I} = (I^1 \cdots I^4)$ by multiplying $\bar{I}(x)$ by the coefficient \bar{A} as:

$$\bar{A} \cdot \bar{I}(x) = \bar{A} \cdot (\bar{I}_{diff}(x) + \bar{I}_{spec}(x)) = \bar{A} \cdot \bar{I}_{spec}(x). \quad (9)$$

Thus, we can directly extract the specular components for the unshadowed points without any complex and fragile separating process. Figure 3 shows that after multiplying the four images by \bar{A} , the diffuse reflection terms are offset, where \bar{A} depends on only the light set.

With the specular reflection components extracted, the problem of data inadequacy is resolved. We can robustly recover the slowly varying specular reflectance with the unshadowed points when the surface geometry is known. Consequently, we can address the ill-conditioned problem for spatially varying BRDF with the sparse images.

3.2. Estimation of the Surface Model

In this section, we introduce an algorithm to recover a full surface model with four images. We first describe how to use the chromatic points to recover the specular reflectance parameters and then explain the algorithm used to recover a surface model by using the specular reflectance.

In order to recover the surface geometry from the color information of pixels, we assume that the surface is satisfied with the NIR model [1]. We assume that the specular reflectance remains invariable for a surface, while allowing that the diffuse reflectance varies arbitrarily over a surface. We will explain how to cope with varying specular BRDFs in Section 4.

3.2.1. Recovering Specular Reflectance with Chromatic Points

According to the NIR model, the specular reflection component has the same spectral distribution. The irradiance $I(x, \bar{\omega}_o)$ can be expressed by rewriting Eq. 3 as:

$$I(x, \bar{\omega}_o) = \underbrace{\frac{\rho_d(x)}{\pi} \max[\bar{\omega}_i \cdot \bar{n}, 0]}_{I_{diff}} \bar{\zeta} + \underbrace{f_s(x; \bar{\omega}_i, \bar{\omega}_o)}_{I_{spec}} \max[\bar{\omega}_i \cdot \bar{n}, 0] \bar{\zeta} \quad (10)$$

If the chromaticity of a pixel is different from that of the illumination, the geometric information can be directly obtained from these pixels by using the cues of color information.

We first consider the case that the color of the k^{th} illumination $\xi^k = \{\xi_r^k, \xi_g^k, \xi_b^k\}$ is pure white, i.e. $\xi_r^k = \xi_g^k = \xi_b^k \equiv \xi_o^k$. According to Eq. 10, we can simply remove the specular term from the irradiance I^k using:

$$\begin{aligned} i^k(x) &= I_u^k(x) - I_v^k(x) = \xi_o^k b(x) \max[\bar{\omega}_i^k \cdot \bar{n}, 0] \\ b(x) &= \rho_{d,u}(x) - \rho_{d,v}(x), \end{aligned} \quad (11)$$

where $\rho_d = \{\rho_{d,r}, \rho_{d,g}, \rho_{d,b}\}$ is the diffuse reflectance in RGB space, u is the maximum color channel of the irradiance $I(x)$, and v is the minimum channel, i.e. $I_{uv} = \max/\min\{I_r, I_g, I_b\}$. Eq. 11 can be extended to the general case of a known illumination to obtain the following:

$$\begin{aligned} i^k(x) &= I_u^k(x) - \frac{\xi_u^k}{\xi_v^k} I_v^k(x) = \xi_o^k b(x) \max[\bar{\omega}_i^k \cdot \bar{n}, 0], \\ b(x) &= \rho_{d,u}(x) - \rho_{d,v}(x), \end{aligned} \quad (12)$$

where the intensity $i^k(x)$ has identical geometrical information to $I_{uv}(x)$ under the k^{th} illumination. Using $\bar{i} = (i^1 \dots i^4)^T$ obtained from Eq. 12, the surface normal for an unshadowed point can be recovered by:

$$\begin{aligned} b(x)\bar{n} &= ([\xi_u^1 \dots \xi_u^4]^T [\xi_u^k])^{-1} [\xi_u^k]^T \bar{i}, \\ [\xi_u^k] &\equiv (\xi_u^1 \dots \xi_u^4)^T \quad \bar{\xi}_u^k = \xi_u^k \bar{\omega}_i^k, \end{aligned} \quad (13)$$

where $[\xi_u^k]$ is the illumination matrix of channel u and $b(x)$ is a scale.

3.2.2. Chromatic Point Selection

Rather than all of the chromatic pixels, we only use the reliable pixels, for which the intensities of the maximum and minimum channels are distinctly different. If the difference between the two channels is small, the variation in irradiance due to geometrical affection becomes sensitive to noise in the images.

Eq. 9 and Eq. 13 are only satisfied when the four irradiances of a point are non-shadowed. Since, for any non-shadowed irradiance set, $\bar{A}_u \cdot \bar{i}(x)$ should be 0, where \bar{A}_u is the coefficient vector of channel u . In the algorithm, we detect shadows and noise by an appropriate threshold of $\bar{A}_u \cdot \bar{i}(x)$: if $\bar{A}_u \cdot \bar{i}(x)$ exceeds a threshold, we will not chose the point.

3.2.3. Estimation of SBRDF

For a surface with many chromatic points, we can obtain a sufficient number of samples with the recovered surface normals to estimate the SBRDF robustly. Many different BRDF models with different strengths and weakness have been proposed. We denote the *SBRDF model* by $f_s(\bar{v}; \bar{\omega}_i, \bar{\omega}_o)$, where \bar{v} is the fit parameter set of the SBRDF, and define the specular pixels $\hat{I}(x) \equiv \bar{A} \cdot \bar{i}(x)$. The square error between a given SBRDF and a specular measurement is computed as follows:

$$\begin{aligned} error(\hat{I}(x), f_s(\bar{v})) \\ = \sum_{c=r,g,b} \left(\hat{I}_c(x) - \sum_{k=1}^4 A_c^k f_{s,c}(\bar{v}_c; \bar{\omega}_i^k, \bar{\omega}_o^k) \bar{\xi}_c^k \cdot \bar{n} \right)^2. \end{aligned} \quad (14)$$

The SBRDF then can be computed using a nonlinear minimization algorithm:




Figure 4 SBRDF recovery with the specular image and the selected chromatic points.

$$f_s(\bar{v}) = \arg \min_{x \in M} \sum error(\hat{I}(x), f_s(\bar{v})), \quad (15)$$

where M are the selected chromatic points on the surface. In Figure 4, the plot on the left shows the $\hat{I}(x)$ of the selected chromatic points, where the x and y axes are the elevation and azimuth angles of the surface normals, respectively. The plot on the right shows the reconstructed distribution of $\hat{I}(x)$ with the recovered SBRDF, which is estimated using a Levenberg-Marguardt optimizing algorithm. In practice, because of camera noise, we recommend the least-median-of-squares technique, rather than the least-squares technique, for estimating real images.

The advantage of the proposed method is that without any separating process, we can directly derive and fit the specular components to the most suitable model. In contrast to methods that fit available pixels to a full reflection model, which simply treats specular components as outliers in estimation, the specular reflectance can be robustly estimated using the proposed method.

3.2.4. Recovering the Full Surface Model

The estimation of the surface parameters at local points is no longer an ill-conditioned problem when the SBRDF has been obtained.

The specular reflection for the chromatic points can be computed with known surface normal. By subtracting the specular reflection component from the irradiance, the diffuse reflection component can be retrieved. Consequently, the diffuse reflectance can be computed by Eq. 3:

$$\begin{aligned} \rho_d(x) &= \\ \arg \min_{k=1}^4 & \left(I^k(x) - I_{spec}^k(x) - \frac{\rho_d(x)}{\pi} \max[\xi^k \cdot \bar{n}, 0] \right)^2. \end{aligned} \quad (16)$$

For the points, of which neither the diffuse reflectance nor the surface normal is known, estimating these parameters from the set of pixel intensities is equivalent to resolving the nonlinear equations of Eq. 3. We first recover the surface normal, which should satisfy the following:

$$\begin{aligned} ([\xi]^T [\xi])^{-1} [\xi]^T (\bar{I}(x) - \bar{I}_{spec}(x)) \times \bar{n} &= 0, \\ I_{spec}^k(x) &= f_s(\bar{v}; \bar{\omega}_i^k, \bar{\omega}_o^k) \max[\bar{\xi}^k \cdot \bar{n}, 0]. \end{aligned} \quad (17)$$

The surface normal can then be computed using a minimization algorithm:

$$\bar{n} = \arg \min \left(([\xi]^T [\xi])^{-1} [\xi]^T (\bar{I}(x) - \bar{I}_{spec}(x)) \times \bar{n} \right) \quad (18)$$

Subsequently, the diffuse reflectance can be recovered with Eq. 16. Given the normals, the surface shape can be reconstructed by using the method of Frankot and Chellappa [8] for a continuous surface.

4. Inverse Reflectometry under Complex Illumination

In this section, we explain the spatially varying BRDF recovery under general distant illumination. We first introduce a method to extract the specular reflection component under general illumination. Then, we describe an algorithm to recover a truly spatially varying BRDF model of a surface which consists of multiple materials.

The input of our algorithm consists of the object geometry, a small set of images taken at a fixed viewpoint under different illumination conditions, and the environment maps of the illuminations. We consider real materials with spatial variation in both the diffuse and the specular reflectance. We assume that the illumination has high-frequency components in order to be able to reconstruct BRDF perfectly.

4.1. Specular Reflection Component Extraction

According to the work of [1] and [18], Lambertian reflection on a surface can be accurately approximated by the low-dimensional linear subspace constructed by spherical harmonic functions. We denote spherical harmonics by $Y_m(l=0,1,\dots; m=-l,\dots,l)$. Then, the illumination can be expressed by spherical harmonics as $L(\bar{\omega}_l) = \sum_{l=0}^{\infty} \sum_{m=-l}^{m=l} L_{lm} Y_{lm}(\bar{\omega}_l)$. The diffuse reflection term of Eq. 4 then is given by:

$$I_{diff}(x) = \frac{\rho_d(x)}{\pi} \sum_{l=0}^{\infty} \sum_{m=-l}^{m=l} \alpha_l L_{lm} Y_{lm}(\bar{n}), \quad (19)$$

where the coefficient α_l vanishes for odd values of $l > 1$, and even terms fall off very rapidly as $O(l^{2.5})$ (see, e.g. [1]). Therefore, the diffuse reflection can be approximated with just a few low order harmonics:

$$I_{diff}(x) \approx \frac{\rho_d(x)}{\pi} \sum_{l=0}^{\Gamma} \sum_{m=-l}^{m=l} \alpha_l L_{lm} Y_{lm}(\bar{n}), \quad (20)$$

We define $\hat{Y}_k(x, \bar{n}) = (\rho_d(x)/\pi) \alpha_l Y_{lm}(\bar{n})$ and $L_k = L_{lm}$, where $k = l \times (l+1) + m + 1$. Eq. 20 can be written as:

$$I_{diff}(x) \approx \sum_{k=1}^{\infty} L_k \hat{Y}_k(x, \bar{n}), \quad N = (\Gamma+1)^2. \quad (21)$$

We take $N+1$ images under different lightings, and use L^j to denote the j^{th} illumination. $\bar{\ell}^j = (L^j_1 \dots L^j_N)$ is defined as the vector of its N low harmonics coefficients which can be computed with the known illumination. The diffuse reflection component under j^{th} illumination then can be expressed by rewriting Eq. 21 as:

$$I^j_{diff}(x) \approx \bar{\ell}^j \cdot (\hat{Y}_1(x, \bar{n}) \dots \hat{Y}_N(x, \bar{n}))^T. \quad (22)$$

Assuming that the N low order harmonics terms of any N incident lightings are linear independent, the $\bar{\ell}$ s of the $N+1$ lightings can be expressed as follows:

$$A^1 \bar{\ell}^1 + \dots + A^{N+1} \bar{\ell}^{N+1} = \mathbf{0}, \quad (23)$$

or written as:

$$\bar{A}_{1 \times (N+1)} \cdot [\bar{\ell}]_{(N+1) \times N} = \mathbf{0}, \quad (24)$$

where the coefficient vector $\bar{A} = (A^1 \dots A^{N+1})$ can be computed from the illumination harmonic coefficient matrix $[\bar{\ell}] = (\bar{\ell}^1 \dots \bar{\ell}^{N+1})^T$. Since the diffuse reflection can be linear represented by the basis of the harmonic function, we can derive from Eq. 22 that any diffuse reflection set of intensities $\bar{I}_{diff} = (I^1_{diff} \dots I^{N+1}_{diff})^T$ should satisfy:

$$\bar{A} \cdot \bar{I}_{diff}(x) \approx \bar{A} \cdot [\bar{\ell}] \cdot (\hat{Y}_1(x, \bar{n}) \dots \hat{Y}_N(x, \bar{n}))^T = \mathbf{0}. \quad (25)$$

Therefore, in accordance with Eq. 4, we can remove the diffuse reflection component by multiplying the irradiance set $\bar{I}(x) = (I^1 \dots I^{N+1})^T$ by the offset vector \bar{A} as follows:

$$\bar{A} \cdot \bar{I}(x) = \bar{A} \cdot (\bar{I}_{diff}(x) + \bar{I}_{spec}(x)) \approx \bar{A} \cdot \bar{I}_{spec}(x). \quad (26)$$

As a result, we can directly derive the specular component by multiplying the sparse image set by \bar{A} , without any complex and fragile separating process. The diffuse reflection terms which have high-frequency variations are offset, and we obtain an image which includes only the specular reflection component of the measurements. We refer to the image as the specular image. For a convex object without cast shadows and interreflections, the offset coefficient \bar{A} only depends on the lighting set.

With the specular reflection component extracted, we can resolve the problem of data inadequacy for sparse images. For a surface with uniform specular properties, we can robustly recover the specular reflectance from the specular image with all the points over the surface. In the case of a surface with multiple materials, since specular reflectance properties have slowly varying nature, we can obtain enough specular irradiance samples of a material to estimate its specular reflectance. Consequently, we can address the ill-conditioned problem for BRDF recovery with sparse images.

4.1.1. How many images at least should we require?

According to the results of [1] and [16], more than 99% of the energy can be captured by $\Gamma = 2$ ($N = 9$), i.e. diffuse reflection can be accurately approximated by the 9D linear subspace constructed by spherical harmonic functions. A straightforward method is to use 10 images under different lighting conditions, and to compute the offset vector \bar{A} with the spherical harmonic coefficient matrix $[\bar{\ell}]$ of the illuminations. However, the 9D approximation is held for the whole sphere of possible normal. At a fixed viewpoint, we have only the upper hemisphere of normals which should produce an approximation with lower dimension as Ramamroothi has discussed in [16]. In this case, we can compute an optimal basis of the diffuse reflection space for a specific object under the light set with SVD method, and obtain \bar{A} by using the basis. Empirically, diffuse reflection under a fixed viewpoint lies very close to a 5D subspace, i.e. at least 6 images are required.

4.2. Fitting Process for Multiple Materials

In this section we will explain how we estimate the BRDFs for a surface composed of multiple materials.

4.3. SBRDF Model

We use the Lafortune model [12] in our work. We consider only isotropic materials in the clustering process. The Lafortune model for isotropic surfaces is expressed as:

$$f_s(\bar{\omega}_i', \bar{\omega}_o') = \sum_l (C_{xy,l}(\omega_{lx}'\omega_{ox}' + \omega_{ly}'\omega_{oy}') + C_{z,l}\omega_{oz}')^{\gamma_l} \quad (27)$$

Where $\bar{\omega}_i'$ and $\bar{\omega}_o'$ are the incident and outgoing directions in the surface local coordinate systems, respectively. $C_{xy,l}$, $C_{z,l}$ and γ_l are the parameters for lobe l of the SBRDF.

4.4. Clustering Algorithm

For the points of the same material, they should have an identical SBRDF. We employ a cluster process inspired by [13] to cluster the specular measurements under the distant illumination condition. In our method, however, each cluster corresponds to a material i.e. a SBRDF. Since specular reflectance properties have a slowly varying nature, it is rational to generate clusters corresponding to SBRDF. In contrast, [13] is according to BRDF, which can not be applied for highly colorful surfaces.

The clustering algorithm is shown in Figure 6. We define the specular pixels $\tilde{I} \equiv \bar{A} \cdot \bar{I}$, the lighting $\tilde{L} \equiv \bar{A} \cdot \bar{L}$, and the SBRDF $\tilde{f} = f_s$. The square error between a given SBRDF \tilde{f} and a specular measurement is computed as follow:

$$error_{\tilde{f}}(x) = \sum_{c=r,g,b} (\tilde{I}_c(x) - \int_{\Omega} \tilde{L}_c(\bar{\omega}_i) \tilde{f}_c(\bar{\omega}_i, \bar{\omega}_o) \max[\bar{\omega}_i \cdot \bar{n}, 0] d\bar{\omega}_i)^2 \quad (28)$$

We value the quality of the approximation in a cluster Λ , using the relative square error, defined as:

$$error(\Lambda_j) = \sum_{x \in \Lambda_j} error_{\tilde{f}_j}(x) / \sum_{x \in \Lambda_j} \sum_{c=r,g,b} \tilde{I}_c(x)^2 \quad (29)$$

We first estimate a SBRDF with an initial cluster consisting of all specular pixels. Then we go into the splitting step to split the cluster into two new clusters each with a different SBRDF model.

4.4.1. Splitting

We denote a SBRDF model by $\tilde{f}(\bar{v}; \bar{\omega}_i, \bar{\omega}_o)$, where \bar{v} is the parameter set of the model. We need to compute the direction of maximum




Figure 6 Split process. The initial SBRDF is split into two new SBRDFs.




Figure 5 The reflectance parameters of a one-lobe Lafortune BRDF approximation for all 100 samples in the MIT/MERL database.

variance in the parameter space in order to generate two new SBRDFs to split the parameters space. Instead of complex calculations, we propose a new method to generate the SBRDFs by taking advantage of the form of the Lafortune model.

One-lobe Lafortune model is used in the splitting step which is sufficient for the cluster splitting procedure. The \bar{v} of the Lafortune model is (C_{xy}, C_z, γ) , where γ is the specular exponent, C_{xy} and C_z indicate the off-specularity of the lobe. As shown in Figure 5, C_{xy} and C_z remain close to the value of one for most materials, especially when γ is big. On the other hand, γ changes much with different materials. Consequently, we can consider generally γ has the maximum variance in the parameter space. Therefore, we construct the two initial SBRDFs of the splitting step as:

$$\tilde{f}_1^0(C_{xy}, C_z, \gamma + 1; \bar{\omega}_i, \bar{\omega}_o), \text{ and } \tilde{f}_2^0(C_{xy}, C_z, \gamma - 1; \bar{\omega}_i, \bar{\omega}_o).$$

After the initialization, the specular pixels from the original cluster Λ are assigned into the new clusters Λ_1^k and Λ_2^k according to their distance $error_{\tilde{f}_i}(x)$ to the two SBRDFs \tilde{f}_1^k and \tilde{f}_2^k . Then, new SBRDF parameters of \tilde{f}_1^{k+1} and \tilde{f}_2^{k+1} are fitted again to the specular pixels in each the cluster Λ_1^k and Λ_2^k . We iterate the reclustering and fitting until the resulting BRDFs and clusters have converged. The superscript k denotes the iteration number.

4.4.2. Reclustering

Having the new clusters and the new SBRDFs, we can re-group all the specular pixels and estimate an exact full Lafortune model for each cluster. This global reclustering is repeated until the SBRDFs are convergent. Clustering process will finish when $\sum error(\Lambda_i)$ is smaller than some ϵ , or the desired number of materials is reached. Otherwise, the cluster with the largest relative square error is chose to split the next.

4.5. Recovering the Full BRDF

With the specular properties obtained, we can recover the diffuse reflectance at each point from the given data. By subtracting the specular component from the pixel intensity, the diffuse component $I_{diff}^k(x)$ is obtainable by using Eq.4. According to Eq.22, the diffuse reflectance then can be computed by:

$$\begin{aligned} \rho_d(x) &= \arg \min error(\rho_d(x)) \\ &= \arg \min \sum_{k=1}^{N+1} (I_{diff}^k(x) - \frac{\rho_d(x)}{\pi} \sum_{l=0}^2 \sum_{m=-l}^{m=l} \alpha_l I_{lm}^k Y_{lm}(\vec{n}))^2. \end{aligned} \quad (30)$$

In our algorithm, we remove the two brightest $I_{diff}^k(x)$ and two darkest one which carry the large errors from the set, and use the middle 5 of $I_{diff}^k(x)$ to compute $\rho_d(x)$ with Eq.30.

5. Results

In this section, we show our experimental results of our implements of the proposed methods. The objects in our experiments have highly textured and colorful surfaces, which are very different from those used in the previous study [13]. In the experiments, a DELL PWS530 (CPU: Pentium4 Xeon 2.2 GHz \times 2; Memory: 2GB) was used.

5.1. Results of Surface Recovery under Controlled Lighting

We first show the results of our Photometric Stereo method under con-

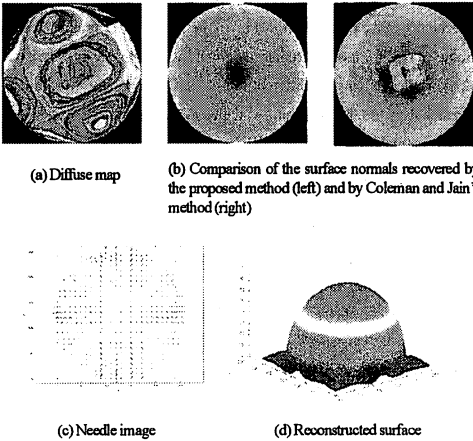


Figure 8 Results for a colored ball.



Figure 7 Comparison of an original image to an image rendered with recovered parameters. L: Original image. M: Rendered image. R: Difference

trolled lighting with four images.

5.1.1. Results for Simulated Objects

We tested the proposed algorithm on a number of simulated objects, which allowed us to verify the accuracy of the algorithm and compare the obtained results to ground truth data.


We used the RADIANCE rendering system [21] to produce four input photographs of an object under a point source from different directions. The simulated objects were rendered with the isotropic Ward's reflection model:

$$f(x, \vec{\omega}_i, \vec{\omega}_o) = \frac{\rho_d(x)}{\pi} + \rho_s \cdot \frac{\exp(-\tan^2 \theta_h / \alpha^2)}{4\pi\alpha^2 \sqrt{\cos\theta_i \cos\theta_r}}, \quad (1)$$

where θ_i , θ_o and θ_h are the geometrical angles, $\rho_d(x)$ is the diffuse reflectance, and ρ_s and α are the specular reflectance parameters. In other words, Ward's SBRDF can be expressed as $f_s((\rho_s, \alpha): \vec{\omega}_i, \vec{\omega}_o)$.

First, the complete set of results for a textured ball is presented. One of the four input images is shown in Figure 10(left). The SBRDF was robustly estimated from the chromatic points as described in Section 3.2. With the specular reflectance, the diffuse reflectances and the surface normals are recovered. The total running time is approximately 15 seconds, where the image size is 256*256. Figure 8(a) and (c) show the diffuse albedo map and the needle map, respectively. Figure 8(d) shows the surface shape recovered from the normals. The specular reflectance parameters are listed in Table 1. In Figure 8(b), the proposed method is compared with that of Coleman and Jain, which uses the same input images. The method of Coleman and Jain failed to recover the surface normals for the locations at which highlights overlap. In addition, some errors occurred in highlight/shadow detection. In contrast, the surface normals were recovered correctly by the proposed method.

We also tested the proposed algorithm with a marble ball and a synthetic orange. While the surface of the marble has complex variation in diffuse reflectance, the surface of the orange has high-frequency spatial variation in the surface normal. Figure 10 shows the recovered diffuse reflectance map for the marble ball and the normal map for the orange. The original and recovered BRDF parameters of these objects are listed in Table 1. The algorithm successfully recovered the surface parameters for both objects.



(a) L: Marble ball, R: diffuse reflectance (b) L: Orange ball, R: normal map.

Figure 10 Results for a marble ball and an orange.

Table 1 Comparison between true and recovered BRDF parameters of the simulated objects.

		ρ_d	ρ_s	α
Marble ball	True	/	.100	.056
	Estimated	/	.098	.056
Textured ball	True	/	.050	.015
	Estimated	/	.056	.015
Orange ball	True	.600, .200, .010	.050	.010
	Estimated	.589, .195, .010	.053	.014

Having recovered the BRDF and the surface normal, we can render objects under new views and lightings. Figure 7 shows the rendering result compared to the original result. Figure 7(right) shows the difference between the rendering and the original images. The errors in the difference image occur on the edges of different-colored segments as a result of the integration process of the rendering system.

5.1.2. Results for a Real Scene


We also tested the proposed method on real objects. The images were captured using a CCD camera (Canon EOS-20D) with no gamma correction, the parameters of which were computed by the Calibration Toolbox of Matlab. Four PLP (60 W) halogen lamps were set approximately 2 meters from the objects (diameter < 20 cm), which were used as point sources. The illuminations were measured by an illuminometer (Konica Minolta CL-200).

An apple made of wax is used in this experiment. In order to avoid saturated pixels, we used the input images in a high-dynamic range radiance model. Four high dynamic range radiance images were acquired from two exposures for the apple. Figure 9(a) shows two of the input images for the apple, respectively. The overlapped highlights among the input images can be seen in both cases.

The diffuse reflectance map and the surface shape recovered by the proposed method are shown in Figure 9(b) and (d), respectively. Figure 9(c) shows the rendering results using the recovered surface model under the same conditions as the input images. There is qualitative agreement, with the highlights similar in area, brightness and position over the surfaces, which verifies the correctness of the estimated surface parameters. The rendering results under the novel lightings and views are shown in Figure 9 (e).




(a) Real images



(c) Rendered (same view, same lighting)

(d) Surface 3D shape



(e) Rendered (novel view, novel lighting). Left: local illumination. Right: global illumination

Figure 9 Results of an apple

5.2. Results of Inverse Reflectometry under Complex Lighting

We tested the proposed approach for recovering BRDF under general illumination on simulated objects which allowed us to verify the accuracy of the algorithm.

5.2.1. Results for a surface with one material

Firstly, we tested the proposed algorithm on a highly textured sphere with one material by using 9D spherical harmonic approximations. We produced 10 input object images under different illuminations by using the environment maps in [7]. The images are rendered with an isotropic two-lobe Lafortune model. The input to our algorithm consists of the surfaces geometry, and the 10 sets of the object images and the corresponding environment maps. The input sets are shown in Figure 11.

We first computed the 9 lighting coefficients for each illumination condition. Sequentially, the offset vector \vec{A} from the lighting coefficient set. Then, the input image set $[\vec{I}] = (\vec{I}_{upper} \dots \vec{I}_{at_pixel})^T$ was multiplied by \vec{A} in order to remove the diffuse reflection components in the images, where \vec{I} denotes all the pixel sets in the images. Figure 11(right) shows the specular image $\vec{A} \cdot [\vec{I}]$ and the specular environment map $\vec{A} \cdot [\vec{L}]$, respectively. We can see the varying diffuse reflections have been offset in with the simple linear computation. $\vec{A} \cdot [\vec{I}]$ was used as the

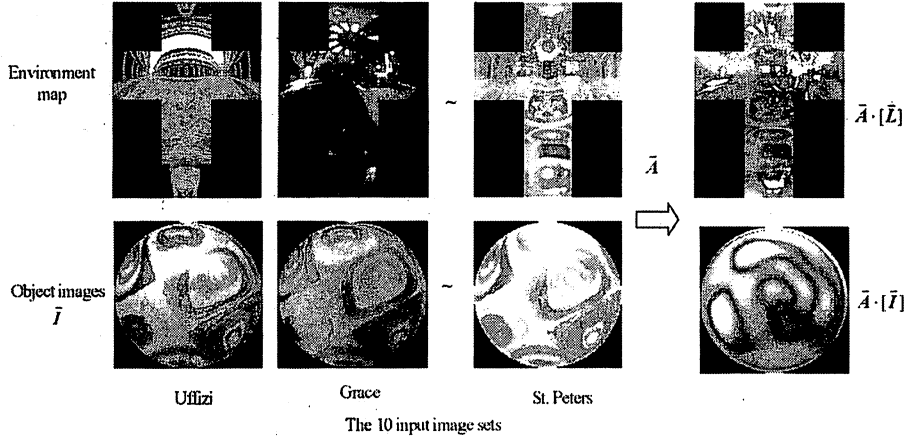


Figure 11 The specular reflection component extraction with 10 images by using 9D spherical harmonic approximation of the diffuse reflection field.

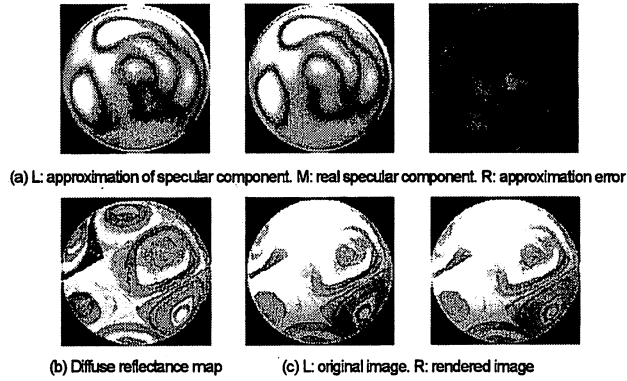


Figure 12 Results for a surface with one material

Table 2. Comparison of SH (9D) and SVD (5D). The first column stands for the quality of approximation for the diffuse reflection component. The other columns stand for the parameters of a two-lobe Lafortune model.

Method	$\frac{(\bar{A} \cdot \bar{I}_{spec})^2}{(\bar{A} \cdot \bar{I})^2}$	C_{xy1} (r, g, b)	C_{z1} (r, g, b)	N_1 (r, g, b)	C_{xy2} (r, g, b)	C_{z2} (r, g, b)	N_2 (r, g, b)
SH	99.89%	(-1.119, -1.121, -1.123)	(1.012, 1.013, 1.017)	(15.90, 15.43, 15.32)	(-1.053, -1.064, -1.084)	(0.710, 0.672, 0.630)	(117.7, 91.43, 66.43)
SVD	99.46%	(-1.129, -1.120, -1.122)	(1.014, 1.016, 1.019)	(14.06, 14.95, 15.31)	(-1.081, -1.047, -1.062)	(0.544, 0.781, 0.715)	(89.02, 114.5, 89.53)
Truth	/	(-1.118, -1.118, -1.120)	(1.013, 1.015, 1.019)	(15.87, 15.65, 15.46)	(-1.053, -1.064, -1.084)	(0.695, 0.662, 0.627)	(111.2, 88.92, 65.22)




Figure 13. Results for a surface with multiple materials

Table 3. Recovered BRDF parameters of the surface with multiple materials

Material		C_{xy}		C_z		N	
		Est.	Tru.	Est.	Tru.	Est.	Tru.
A	r	-1.052	-1.056	0.883	0.881	28.29	29.50
	g	-1.054		0.883		28.48	
	b	-1.055		0.896		34.96	
B	r	-1.033	-1.034	0.960	1.007	239.9	235.9
	g	-1.034		0.959		238.2	
	b	-1.034		0.959		236.3	
C	r	-1.007	-1.007	1.007	1.007	621.9	622.0
	g	-1.007		1.007		621.8	
	b	-1.007		1.007		622.4	

approximation of $\bar{A} \cdot [\bar{I}_{rec}]$ to recover the SBRDF. Figure 12(a) compares the extracted component e is the real specular component $\bar{A} \cdot [\bar{I}]$ to the real specular component $\bar{A} \cdot [\bar{I}_{rec}]$. The difference between $\bar{A} \cdot [\bar{I}]$ and $\bar{A} \cdot [\bar{I}_{rec}]$ is shown in Figure 12(a right). The SBRDF was estimated from the specular reflection image and the environment map by using a Levenberg-Margardt optimizing algorithm, which is shown in Table 2. And the diffuse reflectances (Figure 12b) were recovered, sequentially. Figure 12c shows the rendering result with the recovered parameters compared to the original one.

We also test the proposed algorithm using SVD approximation from 6 images. A 5D optimal basis was computed from the surface normals and the illumination conditions with SVD. The comparison of the estimating results with the two approximation methods is presented in Table 2. As Table 2 indicates, the reflectance model can be recovered successfully by using the proposed algorithm with both the approximation methods. The estimations of γ_2 for the higher frequency reflection lobe have larger errors comparing with the others. 9D spherical harmonic basis shows a better approximation than 5D SVD.

5.2.2. Results for a surface with multiple materials

We tested the proposed algorithm on a highly textured sphere with three different materials to verify the efficiency of the proposed clustering algorithm. We used 9D spherical harmonic approximation for extracting the specular component. One of the ten input object images is shown in Figure 13(d right). The original distribution pattern of the three materials is shown in Figure 13 (b left), where the red, yellow and blue parts denote the different materials. The result of $\bar{A} \cdot [\bar{I}]$ is shown in Figure 13(a). We inputted the value of 3 as the maximum cluster number. After the proposed fitting process had been finished, we obtained a compact expression for the spatially varying BRDF model of the surface including a diffuse reflectance map shown in Figure 13(c), a set of clusters shown in Figure 13 (b right) and the corresponding set of the SBRDFs represented in Table 3, which shows that the proposed method can efficiently recover the whole BRDF model from the sparse image set. Figure 13(d) represents the rendering result by using the recovered reflectance model.

6. Conclusions

The main goal of this paper is the appearance acquisition for the real objects, in particular the objects with spastically varying BRDFs.

The paper addressed the problem for spatial reflectance recovery using a very sparse image set. By taking advantage of the low-dimensional characteristic of diffuse reflection functions, we can removed the high-frequency varying terms from the BRDFs by using a sparse set of images with a linear combination method. Since the specular properties have a slowly varying nature, we can robustly estimate the spatial specular reflectance from the observations, and acquire a high-quality appearance of spatially varying BRDF. We proposed a set of techniques and algorithms to address the ill-conditioned problem for sparse images in spatial reflectance recovery. In the following, we briefly summarize our algorithms.

First, we proposed a four-source *Photometric Stereo* method under controlled lighting. We assume the diffuse reflectance varies arbitrarily over the surface while the specular reflectance remains identical. (The assumption of specular reflectance is removed in Section 4). Since the unshadowed diffuse reflection field is a 3D linear space, we can use four single-light images to remove the varying diffuse term with a linear combination. Then, the identical specular reflectance can be robustly recovered from all of the unshadowed points. And the diffuse reflectance can be recovered at each point subsequently. In previous work, inverse reflectometry methods either exploit the diffuse measurements of points by changing view pointing, or use a large number of images at fixed viewpoint to estimate directly. Our method provides a new recovery approach for observations at a fixed viewpoint.

We then discuss the spatial reflectance recovery for real objects under general distant illumination. Variations are allowed in both the diffuse and the specular reflectance. The truly spatial reflectance recovery with very sparse images has not been achieved under general illumination by any previous techniques. Under general distant illumination, the complete Lambertian reflection field can be approximated by a 9D linear subspace. Thus, the ill-conditioned problem for sparse images can be addressed by using an approximation of low-dimensional subspace (6–10). And spatial SBRDFs can be recovered by an efficient cluster fitting process. To estimate the specular reflectance, we only need to fit SBRDF to a specular image, which substantially simplifies the fitting process. As a result, we can effectively and robustly recover a spatial varying BRDF model under arbitrary distant illumination.

Efficient acquisition of high-quality models is made possible by our techniques since they demand only a very small image set. The idea of removing the low-dimensional reflection terms might be more generally applied to inverse rendering, which can speed up many current inverse rendering processes.

Reference

- [1] R. Bajcsy, S.W. Lee, and A. Leonardis, "Detection of Diffuse and Specular Interface Reflections by Color Image Segmentation," *Int'l J. Computer Vision*, vol. 17, no. 3, pp. 249-272, 1996.
- [2] S. Barsky and M. Petrou, "The 4-Source Photometric Stereo Technique for Three-Dimensional Surfaces in the Presence of Highlights and Shadows," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 25, no. 10, pp. 1239-1252, 2003.
- [3] R. Basri and D. Jacobs, "Lambertian Reflectance and Linear Subspaces," *Proc. 8th IEEE Int'l Conf. Computer Vision*, vol. 2, pp. 383-390, July 2001.
- [4] S. Boivin and A. Gagalowicz, "Image-based Rendering of Diffuse, Specular and Glossy Surfaces from a Single Image," *Proc. SIGGRAPH '01*, pp. 107-116, 2001.
- [5] E.N. Coleman and R. Jain, "Obtaining 3-Dimensional Shape of Textured and Specular Surfaces Using Four-Source Photometry," *Computer Graphics and Image Processing*, vol. 18, pp. 309-328, 1982.
- [6] P.E. Debevec, T. Hawkins, C. Tchou, H.P. Duiker and W. Sarokin, "Acquiring the Reflectance Field of a Human Face." *Proc. SIGGRAPH '00*, pp. 145-156, 2000.
- [7] P.E. Debevec and D. Lemmo, "Acquiring Light Probe Image Gallery," <http://www.debevec.org/Probes/>
- [8] R.T. Frankot and R. Chellappa, "A method for enforcing integrability in shape from shading algorithms," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 10, pp. 439-451, 1988.
- [9] D. Goldman, B. Curless, A. Hertzmann and S. Seitz, "Shape and Spatially-Varying BRDFs From Photometric Stereo," *Proc. of ICCV 2005*, pp. 341-348.
- [10] A. Hertzmann, S.M. Seitz, "Shape and Materials by Example: A Photometric Stereo Approach", *Proc. of CVPR 2003*.
- [11] G. Kay and T. Caelli, "Estimation the Parameters of an Illumination Model Using Photometric Stereo," *Graphical Models and Image Processing*, vol. 57, no. 5, pp. 365-388, Sept. 1995.
- [12] P. F. LaFortune, S. Foo, E. Torrance, and P. Greenberg, "Nonlinear Approximation of Reflectance Functions," *Proc. 24th Conf. Computer Graphics and Interactive Techniques*, pp. 117-126, 1997.
- [13] H.P.A. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.P. Seidel, "Image-Based Reconstruction of Spatially Varying Materials," *Proc. 12th Eurographics Workshop on Rendering*, pp. 104-115, 2001.
- [14] S.R. Marschner, S.H. Westin, E.P.F. LaFortune, and K.E. Torrance, "Image-based BRDF Measurement," *Applied Optics*, vol. 39, no. 16, pp. 2592-2600, 2000.
- [15] W. Matusik, H. Pfister, M. Brand, and L. McMillan: A Datadriven, "Reflectance Model," *Proc. SIGGRAPH '03*, pp. 759-769, 2003.
- [16] R. Ramamoorthi and P. Hanrahan, "On the Relationship between Radiance and Irradiance: Determining the Illumination from Images of a Convex Lambertian Object," *J. Optical Society of America A*, vol. 18, no. 10, pp. 2448-2459, 2001.
- [17] R. Ramamoorthi and P. Hanrahan, "A Signal-processing Framework for Inverse Rendering," *Proc. SIGGRAPH '01*, pp. 117-128, 2001.
- [18] R. Ramamoorthi, "Analytic PCA Construction for Theoretical Analysis of Lighting Variability in Images of a Lambertian Object," *J. IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 24, no. 10, pp. 1322-1333, 2002.
- [19] Y. Sato, M. Wheeler, and K. Ikeuchi, "Object Shape and Reflectance and Reflectance Modeling from Observation," *Proc. SIGGRAPH '97*, pp. 379-387, 1997.
- [20] F. Solomon and K. Ikeuchi, "Extracting the shape and roughness of specular lobe objects using four light photometric stereo," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 18, no. 4, pp. 449-454, 1996.
- [21] G.J. Ward, "The RADIANCE lighting simulation and rendering system", *Proc. SIGGRAPH '94*, pp. 459-472.
- [22] S. Westin, R. Arvo and E. Torrance: Predicting Reflectance Functions from Complex Surfaces," *Proc. SIGGRAPH '92*, 1992.
- [23] Y. Yu, P. Debevec, J. Malik, and T. Hawkins, "Inverse Global Illumination: Recovering Reflectance Models of Real Scenes from Photographs," *Proc. SIGGRAPH '99*, pp. 215-225, 1999.