Microfacet-based Photometric Stereo for Surfaces with Isotropic Reflectance

Lixiong Chen

Master of Science

School of Computer Science
McGill University
Montreal, Quebec, Canada

December 2017

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of Master of Science

Abstract

A precise, stable, and invertible model for surface reflectance is the key to the success of calibrated photometric stereo. Though models addressing low frequency reflectance have been proposed for a broad group of non-Lambertian surfaces, an effective solution directly targeting highly specular reflectance remains elusive. This thesis introduces an analytical isotropic microfacet-based reflectance model, based on which there is a physically interpretable approximation that can be tailored for highly specular surfaces. With this approximation, it shows that a surface recovery problem is essentially an ellipsoid of revolution fitting problem, and a fast, non-iterative and globally optimal solver is derived to attack the latter. Additionally, the introduced model also justifies the fact that, if specularity is not captured by any directional light, a very smooth surface can also appear to be diffusive. This fact leads to a design of an iterative solver for surface estimation with general isotropic reflectance, where the formulation handling specularity is taken as a special instance. Empirical results on images of both synthetic appearances and real objects can validate this model and demonstrate that the proposed solution can stably deliver the state-of-the-art performance.

Résumé

La clé du succès pour une stéréophotométrie calibrée réside dans un modèle de réflectance précis, stable et inversible. Bien que de nombreux modèles de réflectance à basse fréquence furent proposés pour un large groupe de surfaces non-Lambertienne, une solution efficace ciblant directement une réflectance hautement spéculaire reste hors de portée.

Ce mémoire présente un modèle analytique de réflectance isotrope à micro-facette, à partir duquel il existe une approximation physiquement interprétable permettant de cibler des surfaces hautement spéculaires. Avec cette approximation, le problème de restauration de surface est réduit à un problème d'ajustement sur un ellipsoïde de révolution résolu à l'aide d'un solveur non-itératif.

De plus, le modèle introduit justifie le fait qu'une surface très lisse puisse paraître d-iffuse si aucune specularité n'est capturée à partir des lumières directionnelles. De cette propriété découle la formulation d'un solveur itératif pour l'estimation de surface avec un-e réflectance isotopique générale. Dans cette formulation la gestion de la spécularité est considérée comme un cas particulier.

Les résultats obtenus sur les images d'objet capturées et les apparences synthétisées valident ce modèle et montre que la solution proposée est en accord avec les performances décrites dans l'état de l'art.

Contributions

Part of this thesis describes the research that I carried out as a visitor at the National Institute of Informatics (NII) in Japan, while I was enrolled in the M.Sc. program at McGill University. It includes an application of microfacet theory to analyze BRDFs and highly specular reflectance for photometric stereo. Some findings have been published in a proceedings paper [CZS+17], with four co-authors who have made the following contributions:

- Professor Imari Sato oversaw my research activities operated in her lab at NII.
- Professor Yinqiang Zheng helped me design the non-iterative solver introduced in Chapter 4.2.
- Dr. Boxin Shi shared with me the related source code to facilitate a fair comparison between our results and the ones published in the recent literature.
- Mr. Art Subpa-Asa assisted me to setup and carry out the experiment.

Acknowledgements

My deep gratitude goes to my thesis adviser professor Michael Langer for his guidance on my thesis work and for opening the door for me to the joyful research experience in computer vision. The days I spent at School of Computer Science of McGill has made a significant impact on a then software engineer when he was in search of his career goal.

My appreciation also extends to everyone I have worked with during my visit to the National Institute of Informatics in Tokyo. In particular, I must acknowledge the invaluable assistance from professor Imari Sato and professor Yinqiang Zheng, who not only created a resourceful environment for my research, but also shared with me their know-how that helped me pinpoint a research topic.

Moreover, I would like to thank professor Luc Devroye for connecting me with the excellent researchers in Japan, Mr.Art Subpa-Asa from Tokyo Institute of Technology for automating the laborious data collection tasks, professor Kaleem Siddiqi for his enlightening introductory course on computer vision, and Mr.Ilyes Kacher fomr ENSIMAG for preparing the French version of the abstract of this document.

I am indebted to my family, particularly my mother Ping Shi and my father Ming Chen. It is their constant understanding and trust that gives me the courage to go back to campus for a fresh start. Throughout the years, their unconditioned support has been invaluable to me.

Contents

List of Figures

1

Introduction

1.1 Overview

This thesis centers round calibrated photometric, a technique estimates the surface normals of a rigid object by observing its appearance under known but varying illuminations. In a nutshell, it can be considered as a problem of system identification where the input to the system is the lighting, the output is the appearance of the object captured in a set of images, and the system sits in between encodes the shape and reflectance of the object surface.

The system model is essential for its identification, as it describes how lights interact with the object surface and produce signals that can be captured by a fixated imaging sensor. Like other system identification problems, determining the system model involves locating the trade-off between complexity and accuracy: while too simple models fail to capture physical significance of the system, excessively complicated models make their analysis infeasible. For calibrated photometric stereo, it is about devising an effective surface reflectance model to be able to correctly interpret appearance signals critical for shape estimations.

Specularity happens to be this type of signal: it is omnipresent in object appearances, carrying strong shape cues, but tricky to deal with. On one hand, its highly nonlinear variation with respect to lighting requires complicated model to describe; on the other hand, it carries most of the reflectance power, hence ignoring its presence and fitting the appearance into a simpler model is error-prone. In fact, in real-world reflectance there is no clear boundary drawn between specular reflectance and its diffusive counterpart, hence dealing with these two components separately is a non-trivial task. In short, specular reflectance poses a major challenge for calibrated photometric stereo, so an effective photometric stereo with general surface reflectance is what this thesis focuses on.

1.2 Problem definition

This thesis investigates a problem with a slightly restrictive setting. Under directional light $\vec{l}$, the appearance I of a surface point observed from direction $\vec{v}$ is described as a product of shading signal $\vec{n}^\intercal\vec{l}$ and the reflectance signal specified by the isotropic Bidirectional Reflectance Distribution Function (BRDF) $\rho(\cdot)$:

$$I = \rho(\vec{v}, \vec{l}, \vec{n})\vec{n}^\intercal\vec{l} \tag{1.1}$$

where $\vec{n}$ is the unknown to be obtained. Following the discussion in Section 1.1, a tradeoff has to be located between the expressiveness and complexity of Equation 1.1. For example, Lambertian model [Woo80] allows straightforward normal recovery, but does apply to nonlinear reflectance well. In recent years, the photometric stereo has seen dramatic development in shape recovery techniques for surfaces made of a great variety of materials [CJ08, Geo03, GCHS10, IWMA14, TMDB16, WGS$^+$10]. Most know-how has been provided in extracting shape cues from low-frequency reflectance. However, how to properly handle highly specular surfaces, like the ones illustrated in Figure 1.1, remains to be a hard nut. This is because when specular reflectance become predominant, knowledge about inferring shape solely from specularities for photometric stereo is very limited. Consequently, many existing approaches choose to treat the specular signals as outliers, making pre-processing an inalienable part of their design.

On the contrary, how to render visually realistic specular/general appearances has been extensively studied by the computer graphics community. For example, the microfacet reflectance model [AS00, CT82, DWMG15, MB03, WMLT07] offer significant insight into the formation of specularity. Nevertheless, it is unclear how to apply these models directly to infer surface geometry, as analysis to warrant model invertibility and estimation accuracy

(a) chrome steel (b) metallic paint (c) two layer gold

Figure 1.1: Examples of highly specular reflectance targeted in this paper are pervasive in the real world, but effectively estimating the shape of the specular surfaces is a challenging task.

is not in place.

Recent studies [DHI+15, BSN16] have demonstrated that a typical isotropic BRDF can be compactly spanned as a product of several low dimensional functions, where the behavior of these functions is dictated by observations. Since low dimensionality implies stable invertibility, and data-driven approximation preserves expressiveness, apparently a BRDF which is able to extend this notion to surface normal parameterizations can be potential cure for photometric stereo.

1.3 Summary

This thesis expands the idea described in [CZS+17]. It aims to give a comprehensive investigation over the state-of-the-art microfacet theory involving the ellipsoid microfacet normal distribution function (a.k.a. ellipsoid NDF), and derive an analytical model serviceable for photometric stereo. Furthermore, a physical interpretable approximation that brings appealing algebraic properties to specular surface normal estimation is introduced. Essentially, with this approximation, one can identify that the calibrated photometric stereo problem boils down to an ellipsoid of revolution fitting problem, for which a fast, non-iterative and globally optimal solver targeting a system of polynomials is devised. Additionally, an iterative solver targeting general reflectance is implemented, in which the approximated formulation is taken as a special instance. More importantly, the proposed model shows that, when specularity is not fully capture, even a very smooth surface tends to exhibit strong diffusiveness. Empirical results on both synthetic and real images successfully justify the proposed theory and its model, as they demonstrate that the proposed solution can always deliver state-of-the-art performance directly from original measurements.

To sum up, the contributions of this thesis lie in

1. Deriving an analytical form based on microfacet theory and ellipsoid NDF for photometric stereo.
2. Developing a physically interpretable approximation for highly specular reflectance, which equates the problem of normal estimation with an ellipsoid of revolution fitting problem.
3. Designing a fast, non-iterative and globally optimal solver to stably obtain the normal of specular surfaces.
4. Describing a generalized solution for photometric stereo with general isotropic reflectance.

This thesis is organized as follows: Chapter 2 surveys the relevant existing literature, Chapter 3 presents the proposed analytical microfacet reflectance model, and its reduction for highly specular reflectance is derived in Section 3.5. Chapter 4 introduces a novel surface normal estimation algorithm applicable to general reflectance with specularity. Then in Chapter 5 the experiment results obtained from images of both synthetic objects and real objects are examined. Chapter 6 concludes this thesis and looks ahead to the possible future work.

2

Literature Review

This chapter reviews existing work in two separate but converging domains: (1) BRDF design for photo-realistic rendering; (2) Calibrated photometric stereo. The former studies appearance modeling of an object whose geometry is known, so the discussion is centered around modeling and acquisition; whereas since the latter infer the geometry of an object from its appearance, the focus is put on the development of algorithmic implementations.

2.1 BRDF modeling

BRDF accounts for the interaction between incident light and matter by evaluating the power distribution of reflected light in space. The finer the reflectance variation a BRDF is able to capture, the more desirable this BRDF is. BRDF models can be categorized into two types: (1) empirically derived reflectance model and (2) physical-based reflectance model.

2.1.1 Empirically derived reflectance model

Empirical models are devised directly from observations. Namely, they provide a simple and intuitive formulation to conform observers' perception of the reflection. This group of models are in general parameterized by few variables that are not required to carry physical interpretations.

Common empirical BRDF are expressed in terms of a set of geometrical angles: $\vec{v}^\intercal\vec{l}$, the angle made between the light $\vec{l}$ and the viewing angle $\vec{v}$ [Pho75]; $\vec{n}^\intercal\vec{h}$, the angle involving angular bisector $\vec{h}$ between light and view as well as the surface normal $\vec{n}$ [Bli77]. Together with $\vec{v}^{\intercal}\vec{n}$ and $\vec{l}^{\intercal}\vec{n}$, they serve as the fundamental building blocks for some more advanced models. A typical way to express the sharp specular reflectance is to superimposing higher order polynomials containing these terms [LW94, LFTG97]. Since the formulation is derived directly from perception, there is no physical constraint imposed to ensure energy conservation. Hence the visual effects are largely dependent on ad-hoc parameter tuning.

2.1.2 Physically-based reflectance model

The microfacet-based BRDFs represent a group of reflectance models that conform to law of physics. One common postulation made by these models is that the surface consists of randomly disposed facets, and the way each facet reflects the incident light is dictated by the law of electromagnetic, and depending on the size of the facet in contrast to the wavelength of the incident light, the aggregated reflectance can also be simulated by geometric optics [TS67, Smi67] or wave optics [Kaj85, DWMG15], and it can model smooth surfaces as well as rough ones [ON94].

The microfacet-based BRDF is mainly characterized by its normal distribution function (NDF), which specifies how microfacets are statistically distributed with respect to their orientations. Typical examples of NDF include Gaussian [War92] or Beckmann distribution [BS87], followed by more complicated expressions proposed recently [AS00, AP07, BS12]. Various types of NDFs have been applied to describe their targeted reflectance phenomenon, hence for rendering tasks NDF is left to designers as a design choice.

2.1.3 Data acquisition for measurement-based BRDF

A relevant line of research proposes to tabulate BRDF directly. After all, rendering applications do not need to sacrifice appearance subtleties for improved memory usage. Hence data acquisition advocates for dense sampling [GTHD03] turns out to be a task worth trying out [MB03, NDM05]. Moreover, later studies show that directly bookkeeping measured data is actually unnecessary, as real world BRDFs are naturally embedded in a low dimensional manifold [DWT+10, BÖK11], hence can be decomposed into several low dimensional representations [LRR04, BSN16]. In the language of microfacet theory, in a way these representations can be associated with the NDF, masking-shadowing function [Hei14] and the Fresnel term [Sch94] respectively.

For a factorable microfacet-based BRDF, NDF is found to carry dominating effect [AP07], and it can be well-parameterized. Through measurement fitting, parameters are to be determined. Among the proposed models, this thesis chooses to focus on the ellipsoid NDF [WDMG16]. In case of isotropic reflectance, it is analogous to the microfacet BRDF with GGX/Trowbridge-Reitz distribution [WMLT07, TR75]. It is worth noting that though energy conservation is not warranted for rendering, dense light distribution is always desirable in order to capture finer appearance [NLW+16]. Also, the sample material is assumed to have known and simple geometry [DWMG15] (*e.g.,* a flat surface), so the appearance sampling in such scenario exhibit some analogies to calibrated photometric stereo.

2.2 Photometric stereo under calibrated illumination

Photometric stereo can be considered as an inverse problem for rendering. It is a technique that exploit the appearance variations caused by illumination changes to recover surface orientation. Calibrated photometric stereo assumes that the illumination condition is known, and unlike shape from shading problem where a single image is put under investigation [Hor70], the variation of lighting conditions is sufficient in that the surface normal can be unambiguously determined [Woo80]. The number of required images, however, depends on the reflectance property as well as the shape of the surface. For example, for convex Lambertian surfaces, it is proved that the shape can be recovered using as few as three images [BK98] under directional lights. If attached shadow is considered, it is proved that appearances most lie in a 9-dimensional subspace [BJ03, Ram02].

If shiny objects are observed, low dimensional linear representations no longer apply. But if specular signals appear to be sparse, they could be discarded through outlier rejection [WGS+10, IWMA12, IWMA14]. The general reflectance could also be modeled by parametric BRDF models [TMDB16, Geo03, CJ08], or a composition of several simply parameterized components [NIK91a, GCHS10].

Recent approaches adopt non-parametric or semi-parametric formulation to handle a broader range of materials. Without explicit modeling of the BRDF, some general reflectance properties, such as similarity [HS05], isotropy [AK07], monotonicity [STMI12], and their combination with visibility [HMI10], are exploited to constrain surface orientation. The BRDFs can also be explicitly represented as a bivariate function [AZK08, WT13], a constrained bivariate regression [IA14] or a sparse dictionary-based representation [HS15]. Again, BRDF designed for photometric stereo is a delicate trade off between generality and complexity.

In combination with dichromatic model [Sha85], applications of photometric stereo can also process signals from multiple channels [SI94, MZKB05]. This is because by definition specularity directly encodes spectral information from incident light, whereas diffusive reflectance carries the information about the surface itself.

Benchmark evaluation [SWM+16] demonstrates that the state-of-the-art performance can be achieved with data containing less-specular observations [IA14, STMI14]. These approaches work well for a great diversity of real-world materials, but are challenged by specularity-dominant appearances. In contrast, the model described in this thesis aims to attack this challenge with theoretical support.

3

A Microfacet BRDF with Ellipsoid Normal Distribution Function

This chapter describes the basic notions of microfacet-based reflectance model, with a focus on a type of Normal Distribution Function (NDF) named Ellipsoid NDF. Microfacet reflectance model postulates that the surface is made up by a large collection of tiny facets, and the surface radiance is essentially a composition of microfacet reflections, where the radiance intensity can be evaluated [CT82, TS67, APS00] as

$$I(\vec{v}) = \int_{\Omega_+} \max(\vec{m} \cdot \vec{l}, 0) D(\vec{m}) G(\vec{l}, \vec{v}) \rho_{\vec{m}}(\vec{m}, \vec{l}, \vec{v}) \, d\vec{m} \tag{3.1}$$

As illustrated in Figure 3.1, Ω_+ denotes the the visible upper half sphere, $\rho_m(\vec{m}, \vec{l}, \vec{v})$ describes the reflectance of a specific microfacet with normal $\vec{m}$ under directional light $\vec{l}$ while being perceived along $\vec{v}$, $D(\vec{m})$ is the NDF tabulating the population of the microfacets of orientation $\vec{m}$, and $G(\vec{l}, \vec{v})$ is the masking-shadowing term ensuring power conservation. To model general reflectance, each microfacet can be effectively assumed to exhibit mirror reflection [WMLT07], namely, $\rho_m(\vec{m}, \vec{l}, \vec{v}) = F(\theta_d)\delta_{\vec{m}}(\vec{h})$ dictates that a microfacet contributes to the actual reflection only if its normal $\vec{m}$ and bisector $\vec{h} = \frac{\vec{l}+\vec{v}}{|\vec{l}+\vec{v}|}$ are perfectly aligned, and according to the Fresnel equations the amount of power it reflects is determined by the angle θ_d made by the normal and the incident light. Hence, Equation 3.1 can be rewritten as:

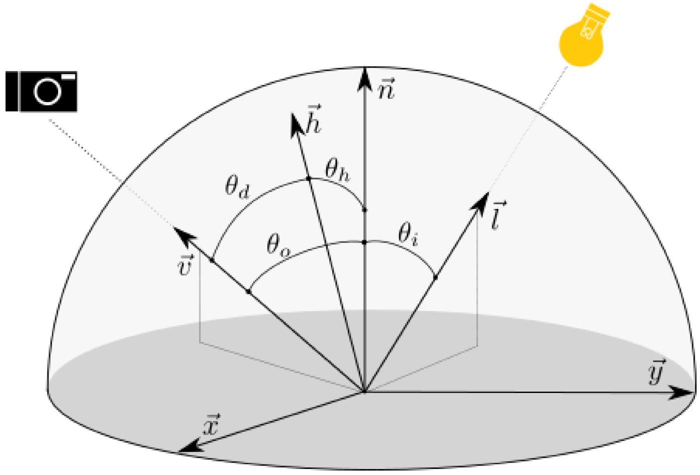

Figure 3.1: The coordinates in which BRDF is defined. By convention $\vec{n} = (0, 0, 1)$, and $\vec{v}$ and $\vec{l}$ are unit vectors that allow to orient arbitrarily above the positive half-sphere. This is in contrast to the typical setup for photometric stereo, where $\vec{v} = (0, 0, 1)$

$$I(\vec{v}, \vec{l}) = G(\vec{l}, \vec{v})D(\vec{h})F(\theta_d) \tag{3.2}$$

Essentially, the microfacet model is to built upon the construction of a Gauss map that parameterizes the microfacet in Euclidean Space $\mathbf{R}^3$ with the its normal $\vec{h}$, where the NDF evaluates its rate of the change over a unit sphere $\mathbf{S}^2$. In this regard, NDF is inherently the gaussian curvature of the surface that the Gauss Map applies to. For example, a planar surface with zero gaussian curvature leads to a Dirac delta NDF that only spikes along the normal of the plane. So, with identical setting given above, the NDF can also be implicitly defined as the inverse of the Gaussian curvature of the illusory surface covered by the microfacets. Moreover, recent study [WDMG16] demonstrates that the ellipsoidal microfacet arrangement and the general GGX NDF are equivalent. Whereas the success of the latter has been widely acknowledged in rendering, the following presents the appealing algebraic properties the former manifests for shape analysis.

3.1 Ellipsoid NDF for isotropic reflectance

Implicitly defining the NDF over an ellipsoid offers several algebraically appealing properties. As illustrated in Figure 3.2, if Ω_+ denotes an arbitrarily defined unit area of the physical surface under examination, and the microfacets can be geometrically translated to cover the upper half of a ellipsoid, then there exists a unique parametrization of surface point $\vec{p}$ by the surface normal $\frac{S\vec{p}}{|S\vec{p}|} = \vec{h}$, then the following is always satisfied:

$$\vec{p}^{\intercal} S \vec{p} = 1 \tag{3.3}$$

where S is a 3-by-3 matrix can always be re-scaled to normalize the RHS of the equation to 1, and it has following properties to characterize the shape of the ellipsoid:

1. S is symmetric and positively definite.
2. In the case of isotropic reflectance, S denotes an ellipsoid of revolution, so its eigenvalues satisfy that $\lambda_3 \geqslant \lambda_2 = \lambda_1 > 0$.
3. Correspondingly, the lengths of the major and the minor axes are $\frac{1}{\sqrt{\lambda_1}}$ and $\frac{1}{\sqrt{\lambda_3}}$ respectively.
4. The arrangement of the microfacetets has to be physically consistent with the surface geometry, so the minor axis is aligned with the surface normal. Namely, $S\vec{n} = \lambda_3\vec{n}$

Spectral theorem states that $S = \lambda_1 \vec{u}\vec{u}^{\intercal} + \lambda_2 \vec{v}\vec{v}^{\intercal} + \lambda_3 \vec{n}\vec{n}^{\intercal}$, where $\vec{u}$, $\vec{v}$, and $\vec{n}$ are its eigenvectors. Correspondingly, $S^{-1} = \frac{1}{\lambda_1}\vec{u}\vec{u}^{\intercal} + \frac{1}{\lambda_2}\vec{v}\vec{v}^{\intercal} + \frac{1}{\lambda_3}\vec{n}\vec{n}^{\intercal}$. Also, let $|S|$ denote the determinant of S, and $K_g = |S|(\vec{h}^{\intercal} S^{-1} \vec{h})^2$ denote the Gaussian curvature of the microfacet-parameterized ellipsoidal surface [Gol05], the ellipsoid NDF $D(\vec{h})$ for Isotropic Reflectance ($\lambda_1 = \lambda_2$) thus can be expressed as:

$$\begin{aligned} D(\vec{h}) = \frac{1}{K_g} &= \frac{1}{\lambda_1^2 \lambda_3 (\frac{(\vec{h}^{\intercal}\vec{u})^2 + (\vec{h}^{\intercal}\vec{v})^2}{\lambda_1} + \frac{(\vec{h}^{\intercal}\vec{n})^2}{\lambda_3})^2} \\ &= \frac{1}{\lambda_3 (1 - (\vec{h}^{\intercal}\vec{n})^2 + \frac{\lambda_1}{\lambda_3}(\vec{h}^{\intercal}\vec{n})^2)^2} \end{aligned} \tag{3.4}$$

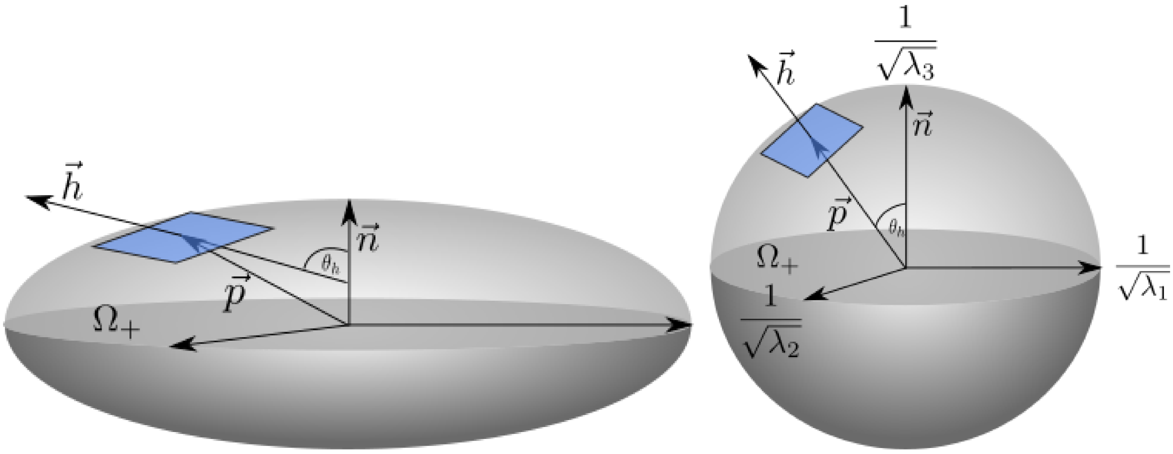

(a) specular surface (b) diffusive surface

Figure 3.2: The ellipsoid NDF describes that the microfacets can be re-arranged through translation to cover the upper surface of an ellipsoid. A "flatter" ellipsoid indicates that more microfacets are aligned with the surface, representing a smoother material.

3.2 The masking-shadowing function

The masking-shadowing function $G(\vec{l}, \vec{v})$ is introduced to impose a physical constraint that the visible and the illuminated area must not exceed the projected area along respectively the perceived direction $\vec{v}$ or the illumination direction $\vec{l}$ (Figure 3.3). In the case where $\vec{v}$ is fixed, $G(\vec{l}, \vec{v}) = G(\theta_i)$ has to satisfy the following for isotropic reflectance [Hei14]:

$$\vec{n}^{\intercal}\vec{l} = \int_{\Omega_+} \max(\vec{h} \cdot \vec{l}, 0) D(\vec{h}) G(\theta_i)\, d\vec{h} \tag{3.5}$$

where $\max(\vec{h} \cdot \vec{l}, 0)$ is to ensure the the microfacet lies in the shadow vanishes. Because in a typical photometric stereo setup a large population of the microfacets are illuminated by a moving directional light over the upper hemisphere, this highly nonlinear term is only significant when a light significantly deviates away from the normal. So, with the premise that lights are distributed sufficiently, it is possible to relax this expression by removing this operator. Therefore, by plugging the widely adopted Smith Microsurface Profile [Smi67] and the ellipsoid NDF with $\int(\vec{h} \cdot \vec{l}) D(\vec{h})\, d\vec{h} = \pi\sqrt{\vec{l}^{\intercal} S \vec{l} |S|^{-1}}$ [Vic96] into Equation 3.5, one

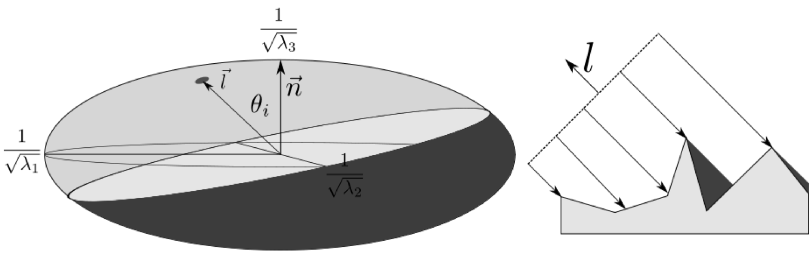

Figure 3.3: The shadowing function guarantees that the total area receiving illumination over a surface of unit area does not exceed $\vec{l}^\intercal\vec{n}$. In the proposed model, the restriction that the region has to be in the upper shpere Ω_+ is removed, so the entire intersected area is considered.

arrives at the following derivation for the shadowing function:

$$G(\vec{l}) = \frac{\vec{l}^\intercal\vec{n}}{\pi}\frac{\lambda_1\sqrt{\lambda_3}}{\sqrt{\vec{l}^\intercal S\vec{l}}} = \frac{\vec{l}^\intercal\vec{n}}{\pi}\frac{\lambda_1\sqrt{\lambda_3}}{\sqrt{\lambda_1(1-(\vec{l}^\intercal\vec{n})^2)+\lambda_3(\vec{l}^\intercal\vec{n})^2}} \tag{3.6}$$

3.3 The Fresnel term

In theory, the Fresnel term $F(\theta_d)$ only starts to vary dramatically as $\theta_d \rightarrow \frac{\pi}{2}$, as a result, in most cases it does not encode sufficient information for shape analysis unless both view and light are at the grazing angles, which only occasionally occurs when lights are located at numerous locations for photometric stereo. Figure 3.4 illustrates the result reported in [BSN16], in which the Fresnel term is evaluated for 100 materials available in [MB03]. It can be observed that, for a wide range of values chosen for θ_i (which is equivalent to θ_d in the microfacet setting), this term can be safely taken as a unknown constant.

3.4 A general reflectance model

Putting Equation 3.1, 3.4 and 3.6 together and letting $\lambda = \frac{\lambda_1}{\lambda_3}$ leads to:

$$I(\vec{l}) = C\frac{\lambda}{(1-(1-\lambda)(\vec{h}^\intercal\vec{n})^2)^2}\frac{\vec{l}^\intercal\vec{n}}{\sqrt{\lambda+(1-\lambda)(\vec{l}^\intercal\vec{n})^2}} \tag{3.7}$$

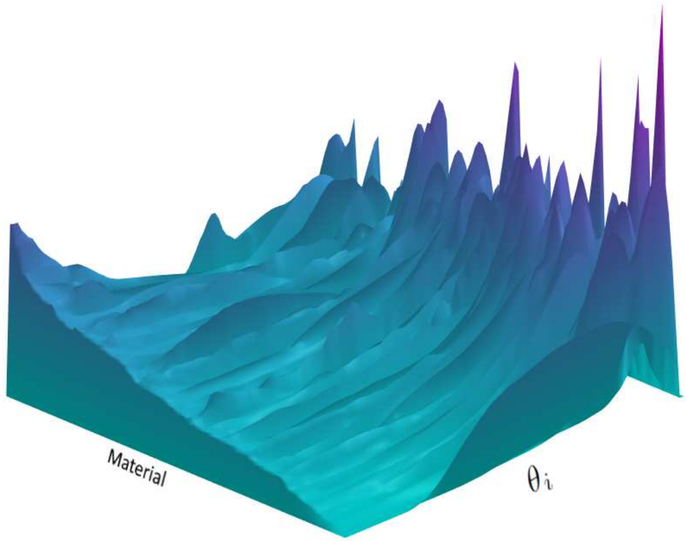

Figure 3.4: A plot of fresnel term obtained in [BSN16]. It shows that except for the extreme values taken by the incident angle θ_i, the fresnel term can be take as a material-specific constant.

where C is a unknown product subsuming the camera gain, the Fresnel term and $\frac{1}{\pi}$; more importantly, one should notice that λ is a factor independent of the geometry and illumination and a sole term characterizing the material's reflectance property. Being the square of the ratio of the minor axis length $\frac{1}{\sqrt{\lambda_3}}$ to the major axis length $\frac{1}{\sqrt{\lambda_1}}$, λ successfully decouples the pixel-wise material evaluation from the actual imaging process: since $\frac{1}{\sqrt{\lambda_1}}$ is the radius of a circular patch orthographically imaged to a specific pixel, λ is essentially the "normalized" shape descriptor of the ellipsoid, regardless how large the "volume" that the ellipsoid occupies, which is also depends on camera pose and light intensity.

Algebraically, λ also plays a central role in identifying the type of surface reflectance. Since microfacet arrangement has to be consistent with the surface geometry, it is expected to have $\frac{1}{\lambda_3} \leqslant \frac{1}{\lambda_1}$, hence $\lambda \in (0, 1]$. When $\lambda \to 1$ results in a sphere, $I(\vec{l}) \to C\sqrt{\lambda}\vec{l}^{\intercal}\vec{n}$, which corresponds to the ideal diffusive case, because microfacets are arranged along an arbitrary direction with equal probability (Figure 3.2b). On contrary, when $\lambda \to 0$, $G(\vec{l}) \to 1$, the material reflectance becomes more conspicuous till only specularities are present. While the former has been extensively studied in the existing literature, here the latter is elaborated.

3.5 Physics driven approximation for specular reflectance

$\lambda \to 0$ leads to a description for perfect mirror reflection, where correspondingly $G(\vec{l}) \to 1$. In this case the surface radiance has to be evaluated under two scenarios:

1. $\vec{h}^\intercal \vec{n} = 1$. By Equation 3.1, $I(\vec{l}) \to \int_{\Omega_+} D(\vec{h})\, d\vec{h} \to C \int_{\Omega_+} \delta_{\vec{n}}(\vec{h})\, d\vec{h} = C$, where $\delta(\vec{h})$ is the dirac delta function describing the infinite impulse due to $\frac{1}{\lambda}$.
2. $\vec{h}^\intercal \vec{n} \neq 1$. $I(\vec{l}) = \frac{C\lambda}{(1-(\vec{h}^\intercal \vec{n})^2)^2} \to 0$, which is a direct simplification from Equation 3.7.

In the first case the light is directly observed through the ideal mirror reflection whereas no diffused radiance can be captured when $\vec{h}$ falls off from $\vec{n}$ in the second case.

Therefore, by letting λ take a sufficiently small value, one obtains a reflectance function for highly specular materials:

$$I(\vec{l}) \approx \frac{C\lambda}{(1-(1-\lambda)(\vec{h}^\intercal \vec{n})^2)^2} \tag{3.8}$$

which can be rearranged into

$$\sqrt{\frac{I(\vec{l})}{C\lambda}}(1-(\vec{h}^\intercal \vec{n})^2 + \lambda(\vec{h}^\intercal \vec{n})^2) \approx 1 \tag{3.9}$$

and by defining $\hat{S} = \vec{u}\vec{u}^\intercal + \vec{v}\vec{v}^\intercal + \lambda\vec{n}\vec{n}^\intercal$, Equation 3.9 can be further simplified as

$$((\frac{I(\vec{l})}{C\lambda})^{\frac{1}{4}}\vec{h})^T \hat{S}((\frac{I(\vec{l})}{C\lambda})^{\frac{1}{4}}\vec{h}) \approx 1 \tag{3.10}$$

which is essentially the standard equation for an ellipsoid of revolution $\hat{S}$ centered at the origin, where the original ellipsoid S and $\hat{S}$ are co-axial. Fitting an ellipsoid requires at least 4 points on its surface, which can be easily satisfied for photometric stereo. After all, as the directional light relocates, distinct appearances can be obtained except for nearly perfect mirror reflection, for which only the appearance of the illuminant can be directly seen from one specific location. This analysis rules out this extreme case.

Moreover, it is reasonable to expect this approximation to become less accurate when applied to diffusive materials. As discussed in Section 3.4, in diffusive cases $\lambda \to 1$, the ellipsoid degenerates to a sphere without elongation (see Figure 3.2b). Algebraically, this means that Equation 3.10 can be satisfied by a set of non-unique $\hat{S}$. Fortunately, the value of λ itself serves as a good measure for estimation confidence, so one can always safely "roll back" to the exiting solvers implemented for low-frequency reflectance if large values for λ are detected.

4

Algorithm

This chapter develops an algorithmic toolset for microfacet-based photometric stereo based on the model introduced in chapter 3. The toolset consists of three components: (1) λ-estimator for material extraction; (2) a polynomial-based $\vec{n}$-estimator for specular reflectance; and (3) an iterative joint $\vec{n}$-λ estimator for photometric stereo with general reflectance, which subsumes the first two components.

4.1 Estimating λ

When λ is estimated along, letting $x = \sqrt{\lambda + (1-\lambda)(\vec{l}^{\intercal}\vec{n})^2}$ transforms Equation 3.7 to a system of single-variable polynomials:

$$p(x) = C\lambda(x)(\vec{l}^{\intercal}\vec{n}) - I(\vec{l})(1 - (1-\lambda(x))(\vec{h}^{\intercal}\vec{n})^2)^2 x = 0 \tag{4.1}$$

where $\lambda(x) = \frac{x^2 - (\vec{l}^{\intercal}\vec{n})^2}{1-(\vec{l}^{\intercal}\vec{n})^2}$. Because $\vec{l}^{\intercal}\vec{n}$ and $\vec{h}^{\intercal}\vec{n}$ are parameters, $\lambda(x)$ is rational, so is $p(x)$. Since this has to hold for all k distinct lightings, with i denoting the observation index, solving this system of polynomials can be treated as a single-variable optimization problem, where the candidate solution has to satisfy the following:

$$\frac{\partial}{\partial x} \sum_{1 \leqslant i \leqslant k} p_i(x)^2 = \sum_{1 \leqslant i \leqslant k} \frac{\partial}{\partial x} p_i(x)^2 = 0 \tag{4.2}$$

which can solved using the corresponding companion matrix [HJ90].

4.2 Estimating surface normal with highly specular reflectance

Given the illumination direction $\vec{l}$ (equivalently the half vector $\vec{h}$) and the image radiance $I(\vec{l})$, photometric stereo seeks to recover the surface normal $\vec{n}$. In this scenario, this boils down to fitting a unknown surface of an ellipsoid of revolution in $\mathbf{R}^3$ using points $\vec{h}$ on a spherical surface $\mathbf{S}^2$ whose lengths are re-scaled by $\left(\frac{I(\vec{l})}{C\lambda}\right)^{\frac{1}{4}}$. After the ellipsoid is determined, by detecting its elongation the surface normal can also be obtained.

Fitting an ellipsoid can be formulated into an energy minimization problem, and one is able to retrieve the global minimum from the solutions to a system of polynomials. For simplification let $P = \sqrt{I(\vec{l})}$ and $\omega = \sqrt{\frac{1}{C\lambda}}$. To get around the unit norm constraint on $\vec{n}$, one can also let $\vec{\hat{n}} = \sqrt{(1-\lambda)\omega}\vec{n}$, hence for each of the k observations, Equation 3.9 can be rewritten as

$$P_i \left(\omega - \vec{\hat{n}}^T \vec{h_i}\vec{h_i}^\intercal \vec{\hat{n}}\right) = 1, i = 1, 2, \cdots, k, \tag{4.3}$$

where again i is the observation index.

By averaging all k equations, one obtains $\bar{P} = \frac{1}{k}\sum_{i=1}^{k} P$, $\bar{H} = \frac{1}{k}\sum_{i=1}^{k} P_i\vec{h_i}\vec{h_i}^\intercal$, and

$$\omega = \frac{1 + \vec{\hat{n}}^T \bar{H} \vec{\hat{n}}}{\bar{P}}, \tag{4.4}$$

Moreover, combining Equation 4.4 into Equation 4.3 leads to

$$\vec{\hat{n}}^T \left(P_i\vec{h_i}\vec{h_i}^\intercal - P_i\frac{\bar{H}}{\bar{P}}\right)\vec{\hat{n}} = \frac{P_i}{\bar{P}} - 1, \tag{4.5}$$

a quadratic polynomial with respect to $\hat{n} = [\hat{n_1}, \hat{n_2}, \hat{n_3}]^T$. Therefore, all k equations in Equation 4.5 can be organized into the matrix form

$$Mx = M[\hat{n}_1^2, \hat{n}_1\hat{n}_2, \hat{n}_1\hat{n}_3, \hat{n}_2^2, \hat{n}_2\hat{n}_3, \hat{n}_3^2]^T = b, \tag{4.6}$$

so $M \in \mathbb{R}^{k\times 6} and b \in \mathbb{R}^k$ are established.

However, due to measurement noise and model precision, one should not enforce the equality in Equation 4.6 to hold strictly, instead it is more desirable to find the optimal $\hat{n} = [\hat{n_1}, \hat{n_2}, \hat{n_3}]^T$ that minimizes the following energy function

$$f(\hat{n}) = \|Mx - b\|_2^2 = x^T M^T Mx - 2b^T Mx + b^T b. \tag{4.7}$$

Since the cost function in Equation 4.7 is nonconvex, one has to find its global minimizer by retrieving all its stationary points. Specifically, it is feasible to solve the three-variable cubic equations defined by the partial derivatives as

$$\frac{\partial f}{\hat{n}_1} = 0, \frac{\partial f}{\hat{n}_2} = 0, \frac{\partial f}{\hat{n}_3} = 0. \tag{4.8}$$

which is a three-variable cubic polynomial system that has 27 solutions. Since the system is homogeneous, the solutions are positive-negative symmetric. Therefore, one only needs to examine 13 independent solutions. These facts motivate a solver based on the symmetric Gröbner basis [LÅ16]. Finally, λ and C can be determined consecutively using the length of $\vec{\hat{n}}$.

4.3 An iterative normal estimator for surfaces with general reflectance

As discussed in 3.4, a value between 0 and 1 assigned to λ fully characterizes the reflectance of a material. However, highly specular reflection may not be observed from smooth surfaces under sparse light distributions. The existing approaches [WGS+10, SWM+16] take advantage of this spatially varying notion of surface reflectance, and choose to reject specular signals through thresholding over the pixel intensity. On contrary, this section describes a solver that directly attacks the general reflectance model described in Equation 3.7. By investigating a different type of approximation in response to spatially varying illuminations, it successfully gets around with the need to pre-process the inputs.

As opposed to the approximation discussed in Section 3.5 which is physically-driven by material smoothness, an perception-based approximation with respect to the varying lighting direction exists to reveal the essence of treating the specular signals as outliers. According to the shape of the specular lobe, the lighting directions with respect to a surface point can be partitioned into two groups: (1) $\vec{h}^{\intercal}\vec{n} \rightarrow 0$ and (2) $\vec{h}^{\intercal}\vec{n} \rightarrow 1$. Because $\vec{v}^{\intercal}\vec{n} \geqslant 0$, $\vec{h}^{\intercal}\vec{n} \rightarrow 0$ also implies that $\vec{l}^{\intercal}\vec{n} \rightarrow 0$. Then one can obtain the first-order approximation for the Taylor expansion of Equation 3.7 as follows:

$$\begin{aligned} I(\vec{l}) &= C\frac{\lambda}{(1-(1-\lambda)(\vec{h}^{\intercal}\vec{n})^2)^2}\frac{\vec{l}^{\intercal}\vec{n}}{\sqrt{1+(1-\lambda)((\vec{l}^{\intercal}\vec{n})^2-1)}} \\ &= C\lambda\big(1+\mathcal{O}((\vec{h}^{\intercal}\vec{n})^2)\big)\vec{l}^{\intercal}\vec{n}\big(1+\mathcal{O}((\vec{h}^{\intercal}\vec{n})^2)\big) \\ &\approx C\lambda\vec{l}^{\intercal}\vec{n} \end{aligned} \tag{4.9}$$

which is consistent with the assumption made in [WGS+10] that many readings of non-specular reflectance, regardless the surface material they are measured from, obey the Lambert's Law.

Therefore, combining with λ, there exists three scenarios:

- $\lambda \rightarrow 1$ and $\vec{h}^{\intercal}\vec{n} \rightarrow 1$: the specularity is guaranteed to be observed. $I(\vec{l})$ can be approximated by Equation 3.8.
- $\lambda \rightarrow 1$ and $\vec{h}^{\intercal}\vec{n} \rightarrow 0$: the specularity is unlikely to be detected even though the surface can be very smooth. $I(\vec{l})$ can be approximated by Equation 4.9.
- $\lambda \rightarrow 0$: no specular reflectance is to be captured in all cases. $I(\vec{l})$ can be approximated by $C\sqrt{\lambda}\vec{l}^{\intercal}\vec{n}$.

It is worth noting that, though carry distinct physical significance, the second and the third case only differ in scale. More importantly, this approximation indicates that when incomplete observations are made from a smooth surface which is supposed to produce specularity, by directly applying the proposed model one may obtain an estimated λ larger than its actual value, but the surface normal $\vec{n}$ stays invariant. Hence an ideal solver should always correctly estimate surface geometry regardless the light distribution.

Numerically, it is reasonable to expect the actual solution to satisfy one of the approximations discussed above, in addition to complying with the original model described in Equation 3.7. Moreover, as discussed in Section 3.5, when the non-iterative solver produces inaccurate estimation for diffusive observations, its estimation on λ can be somehow reliable. As a result, one can devise a two-stage iterative solver as follows:

Step 1 Apply the approximation for specular reflectance (Section 4.2) and the Lambertian model independently to the observations. Let $\vec{x}_s, \lambda_s$ and $\vec{x}_d, \lambda_d$ be the solution obtained for each case respectively.

Step 2 Plug in $\vec{x}_s, \lambda = 0$ and $\vec{x}_d, \lambda = \lambda_s$ as the initial estimate, solve the Equation 3.7 directly using nonlinear least square solver. Let $\vec{x}_s^*$ and $\vec{x}_d^*$ be the corresponding solutions.

Step 3 The final solution $\vec{x}^* \in \{\vec{x}_s^*, \vec{x}_d^*\}$ is the one results in less fitting error.

It is worth noting that Step 1 finishes in polynomial time for both models and is non-iterative. Step 2 applies nonlinear least square twice hence it completes iteratively. All subproblems are formulated on a system of equations involving all observations, and operations on pre-processing are avoided.

5

Experiment

This chapter describes the experiments designed to validate the proposed model and illustrates the effectiveness of the corresponding algorithms. Since inter-reflections and shadowing are not considered, sphere serves as a desirable sample object for evaluating its normal map. For quantitative examination, PBRT [PJH16] is used with MERL BRDF [MB03] to synthesize a sphere under numerous directional lights. Spiral points [RSZ94] are used to approximate uniform distribution. Aside from synthetic images, appearances real-world objects are also experimented. Typical samples include USC "Light Stage Data Gallery" [ECJ+06], UCSD photometric stereo data set [AZK08] and recently published "DiLiGent" benchmark data set [SWM+16]. These image sets contain objects with various reflectance properties, and "DiLiGent" image set also provides the ground truth of the normal maps.

In the experiments, the results produced by the proposed method are compared against the results produced by Nonlinear Least Square (NLS),Cook-Torrance Model [CT82] ,Ward Model [War92], Lafortune model [LFTG97], Biquadratic [STMI14] (Biquad), Constrained Bivariate Regression (CBR) [IA14], where the latter two are known to provide the state of the art performance to date. Also, these two approaches represent two distinctive groups of methodologies: while Biquad applies a relatively constrained but fine-tuned parametric model, CBR adopts a non-parametric form to describe general reflectance.

On synthetic images, three sets of experiments are performed under six lighting conditions (*i.e.,* 60, 96, 100, 150, 250, 500 lights, Figure 5.1): (1) estimating only λ assuming the geometry of the sphere is given; (2) feasibility study for the non-iterative solver designed for specular reflectance; (3) estimating the surface normal of the sphere using the iterative approach for general reflectance. The distribution of the 96 lights is non-uniform and simulates the lighting adopted for “DiLiGent” benchmark data set. The distribution of the 100 lights is randomly generated to create an adversary illumination condition. For CBR, N_1 is set to 2 and N_2 is set to 4, with “retroreflective” on. In other parametric models, surface normal $\vec{n}$ is directly estimated using NLS.

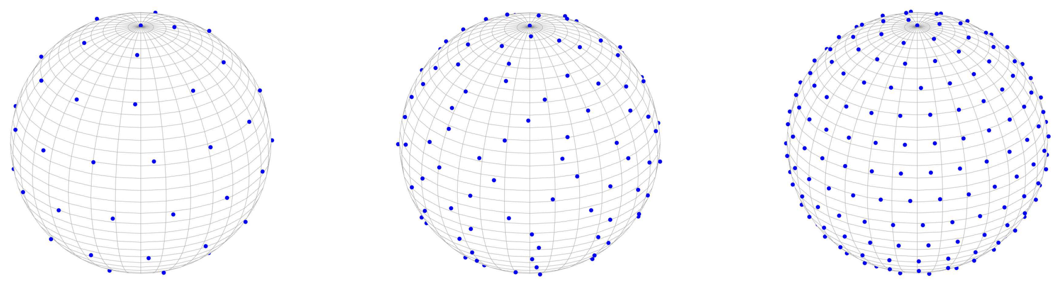

(a) 60 lights (b) 150 lights (c) 250 lights

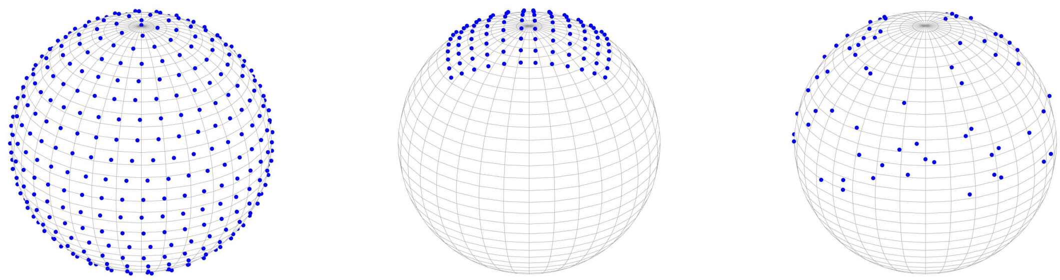

(d) 500 lights (e) 96 lights (f) 100 lights

Figure 5.1: Distribution of lights with various densities and patterns. The positive Z-axis is pointing upward. Hence the lights located in the bottom are likely to contribute less to appearance formation.

5.1 Model validation

Figure 5.2 compares the performance of the proposed model on fitting the appearance of a sphere against the performance of some well-known BRDFs for rendering as well as for photometric stereo. The metric measures the mean fitting discrepancy between the estimated appearance and the actual observation. As one may expect, since the geometry of the object is already specified, the BRDFs for rendering takes an advantage in that they adopt complicated models to capture the fine-grained appearance details.

Under uniformly distributed lights (Figure 5.2b and Figure 5.2c), the denser the the light, the more accurate appearance is generated. Also, the performance gain is more significant for specular appearances. This is mainly because specularity carries more energy conveyed by the appearance. Once it is captured by light, its approximation by the model determines the overall evaluation for model effectiveness. Consequently, since partially distributed lights (Figure 5.1e) is less likely to capture specularity completely, in such setting low-frequency reflectance model ([STMI14]) delivers improved results (Figure5.2a).

Overall speaking, the proposed ellipsoid-based model outputs a stable and accurate appearance estimation for all 100 materials.

5.2 Non-iterative solver for specular reflectance

5.2.1 Evaluation on synthetic images

Figure 5.3a and 5.3c present the angular estimation errors in degrees over the 100 materials in MERL. It can be seen that the proposed solution targeting on specular reflectance produces a trace complementary to those produced by the other three approaches targeting on low-frequency reflectance. Two conclusions can be drawn from these plots:

1. All methods perform inferiorly over a specific set of materials, and this division performance over the materials is almost independent of light density;
2. The existing approaches outperform on surface with diffusive reflectance, and the proposed method delivers better performance for specular surfaces with a few ex-

ceptions (Section 5.2.4), which is consistent with the prediction made in Section 3.4.

5.2.2 Evaluation on images of real objects

The proposed solver is also applied to two images sets of real scenes, "helmet side right" from "Light Stage Data Gallery" and "DiLigent". The "helmet" image set contains specular appearances captured under 253 directional lights, and the "DiLigent" data set that mainly represents diffusive materials.

Figure 5.4 visualizes the estimated normals with their respective $+x$,$+y$ and $+z$ components, with (1) lights with positive z-values only and (2) all 253 lights. No qualitative comparison is made because ground truth is unavailable, but qualitatively the proposed method appears to present reasonable and consistent results, indicating its stability. Besides, the convex shape of the model is clearly illustrated by $+x$,$+y$ and $+z$ components together. It is also worth noting that CBR delivers much more reasonable result when only lights with $+z$ are selected (Figure 5.5). In general, light distribution makes a major impact on estimation accuracy. It is reported that existing approaches shall perform better with properly adjusted "position thresholds" [SWM+16], which is to be examined in Section 5.3.3.

5.2.3 Detection of diffusive reflectance

Among the ten models in "DiLiGenT", "ball", "reading","cow" and "harvest" represent relatively more specular materials. Figure 5.6 compares the estimation error produced by CBR, Biquad and the proposed approach, together with the median value of λ obtained for each model. Here it is assumed that each material is made of homogenous material so a rough cross-pixel analysis is permitted. It is interesting to observe that except for "ball" and "reading", the lower the λ value detected, the better the performance the proposed method delivers. The inferior performance on "ball" is due to limited light distribution as to be discussed in Section 5.2.4, and "reading" is inaccurately estimated because surface non-convexity has caused a significant amount of specular inter-reflections.

However, it should be noted that though λ can correctly indicate the smoothness of the surface, a fail-safe "switch" that allows one to roll back to the existing solutions for low-frequency reflectance remains absent. Specifically, whether an arbitrary λ between 0 and 1 represent a physical specular surface or diffusive surface is unclear. So, as pointed out in Section 4.3, an effective solver addressing general reflectance should not only rely on λ, but also consider the spatial variation between light and the surface geometry.

5.2.4 Impact of light density

The distribution density of lights affects the accuracy of the proposed solver. Figure 5.7 compares the accuracy obtained under various illumination densities: 500, 250, 150, 60 lights, respectively. The main observation made is that the proposed solution produces stable output for various lighting densities, but the denser the distribution, the higher the accuracy is achieved. More importantly, the gain due to a denser light distribution is more significant for specular surfaces, whose labels are highlighted in red. For example, "specular green phenolic", as indicated in Figure 5.8, exhibit extremely localized specularity with large estimation error. This is because when light distribution is insufficient, many pixels (*e.g.,* point B) do not exhibit specular property at all, as compared with the pixels (*e.g.,* point A) covered by specularities, they are less accurately estimated as a consequence of mis-fitting the model. Though uncommon in a photometric stereo setup, it is expected to see that with a denser light distribution, estimation on the specular materials highlighted in red in Figure 5.3 shall continue to improve.

On the other hand, the light that directly generates the specularity is disfavored by the methods introduced in [WGS+10, SWM+16]. Specifically, by intentionally ignoring the specular appearances, the remaining set of reflectance is thought of as if its observed from a diffusive surface. Theoretically, lights of an arbitrarily dense distribution in their limit shall make specular appearance unavoidable, and in such case a model that handles specularity should excel in robustness and stability.

5.3 Iterative solver for general surface reflectance

As described in Section 4.3, the iterative solver for general reflectance is a wrapper around the non-iterative solver for specular reflectance, with restriction of $\lambda \to 0$ removed. In the experiment, the camera response is assumed to be linear, and the input is taken as-is except that pixels with zero-readings are ignored. Also, no pre-processing is applied. The iteration is dictated by the standard NLS algorithm, with constraint that $\lambda \in [0, 1]$.

5.3.1 Evaluation on synthetic images

Figure 5.9 compares the performance of the proposed general model with the estimations of a sphere respectively provided by Ward, Torrance, *BVR*,and Biquad reflectance model, under all 6 lighting conditions (Figure 5.1). It can be observed that the proposed method outperforms its peers in all situations, including irregularly distributed lights (Figure 5.9c). The proposed method successfully recovers the shape of the sphere of all 100 materials in MERL database.

5.3.2 Impact of light distribution

Figure 5.10 contrasts the performance of the proposed method under various conditions. It is not surprising to see that the worst case arises for the arbitrarily distributed 100 lights, the second most inferior performance is obtained when lights are non-uniformly distributed (e.g. 96 lights, Figure 5.1e). This is caused by insufficient observations made available for model fitting: the parameters are better learned if dense and sufficiently distributed observations are made. For uniformly distributed lights, as expected, the denser the light distribution, the more accurate the estimation.

5.3.3 Evaluation on real images

Figure 5.11 presents the recovered normal map and their decomposition along $+x$,$+y$,and $+z$ axes respectively. Since each object is convex and symmetric, the symmetry of the decomposed illustration indicates that reasonable results are obtained for these three objects even though their reflectance properties vary.

Figure 5.12 demonstrates the normal map estimated by Biquad, BVR and the proposed general methods for “DiLiGent” image set. Here Biquad and BVR are taking preprocessed inputs (“position threshold” (PT) [SWM+16]) in order to achieve the state-of-the-art performance for calibrated photometric stereo. A quantitative comparison is also presented in Figure 5.13. The performances delivered by these methods can be summarized as follows:

- The proposed method for general reflectance is comparable with the state-of-the-art approaches in any situation.
- The proposed method delivers much more accurate estimation on "cow" and "harvest", which are mainly covered by specular surfaces.
- Inferior estimations are made for highly non-convex shapes. For example, Part of the estimated result for "reading" is inaccurate because many specular reflections caused by inter-reflections; the estimation error for "buddha" is mainly due to the cast shadow in the back of the scene.

Figure 5.13 also totals the estimation error when the proposed method is applied in tandem with thresholding. As discussed in Section 5.2.4, thresholding is essentially equivalent to selecting a subset of lights for analysis. Hence as expected, doing so makes no impact on the performance of the proposed approach.

5.4 Summary

To sum up, the proposed microfacet-based model captures the general isotropic appearance well, and stands out in describing the specular reflectance. Accordingly, a non-iterative solver targeting specular reflectance and an iterative solver for general reflectance that takes the former as a special instance together address photometric stereo successfully, delivering results that are comparable with the state-of-the-art performance. Since the proposed methods are insensitive to light distribution, these solvers can handle more materials and offer accurate estimations with improved stability and robustness.

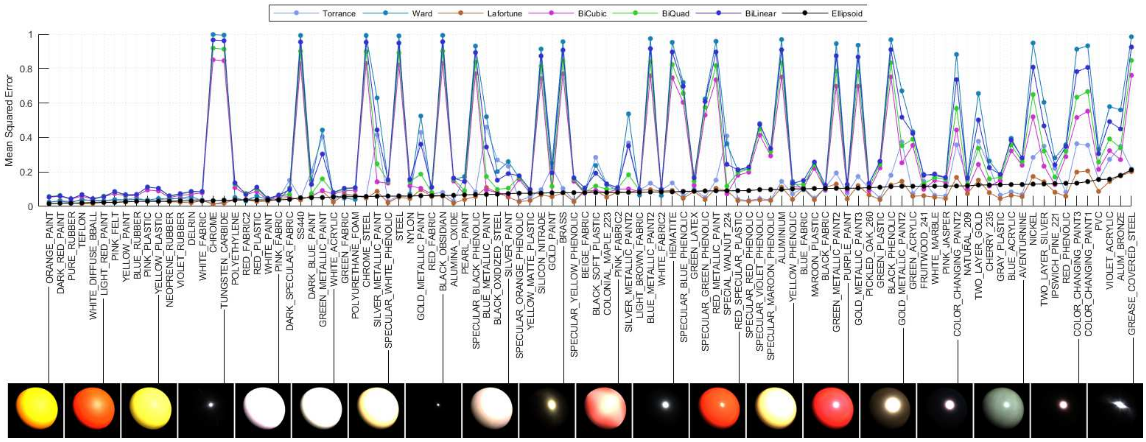

(a) Mean squared estimation error on sampled MERL appearance using 96 lights.

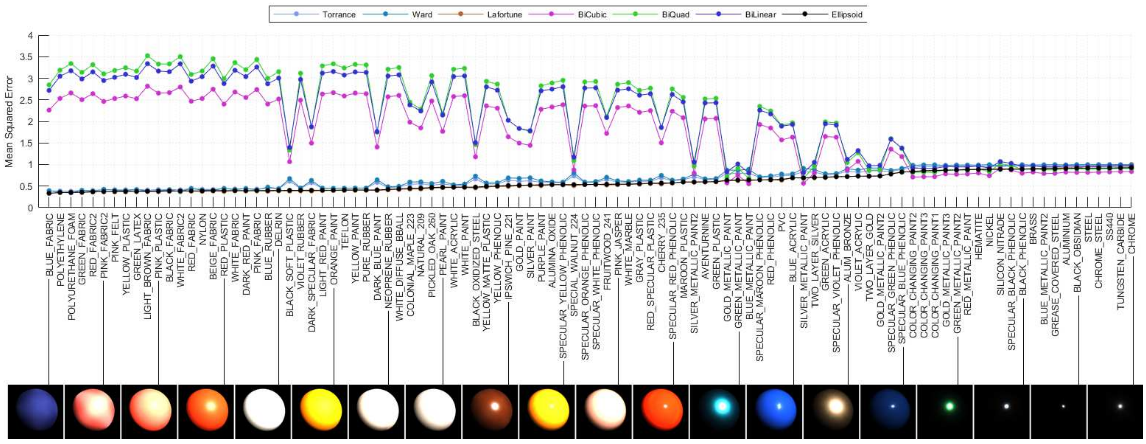

(b) Mean squared estimation error on sampled MERL appearance using 100 lights.

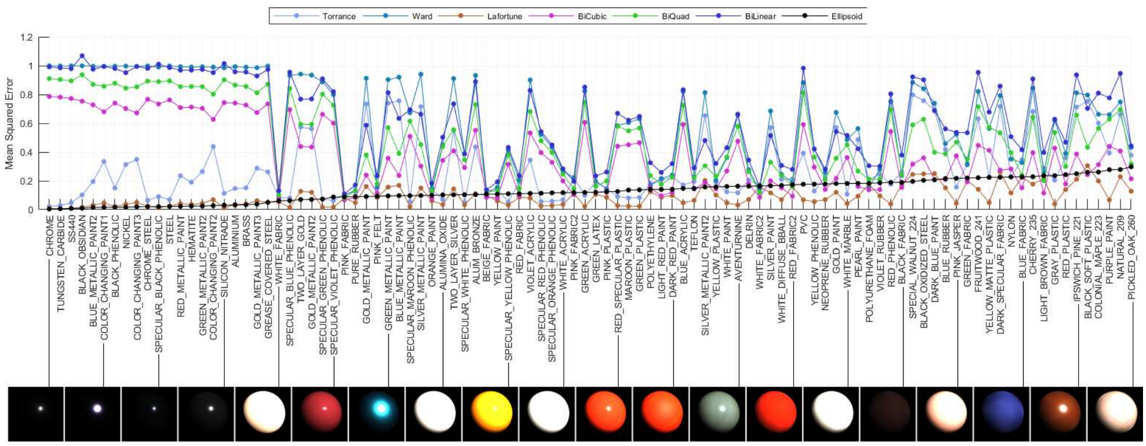

(c) Mean squared estimation error on sampled MERL appearance using 150 lights.

Figure 5.2: Performance comparison on appearance fitting under various lightings.

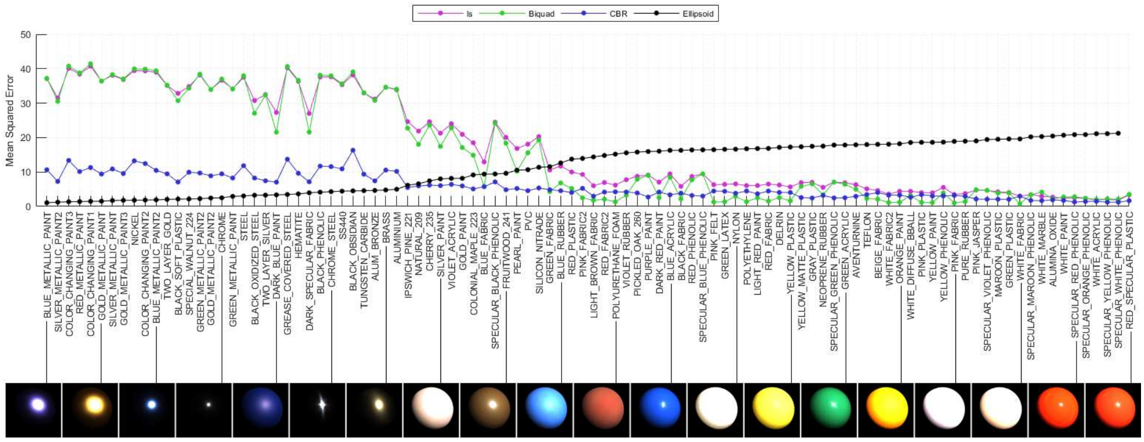

(a) Mean estimated angular error in degrees produced by various methods using 60 lights.

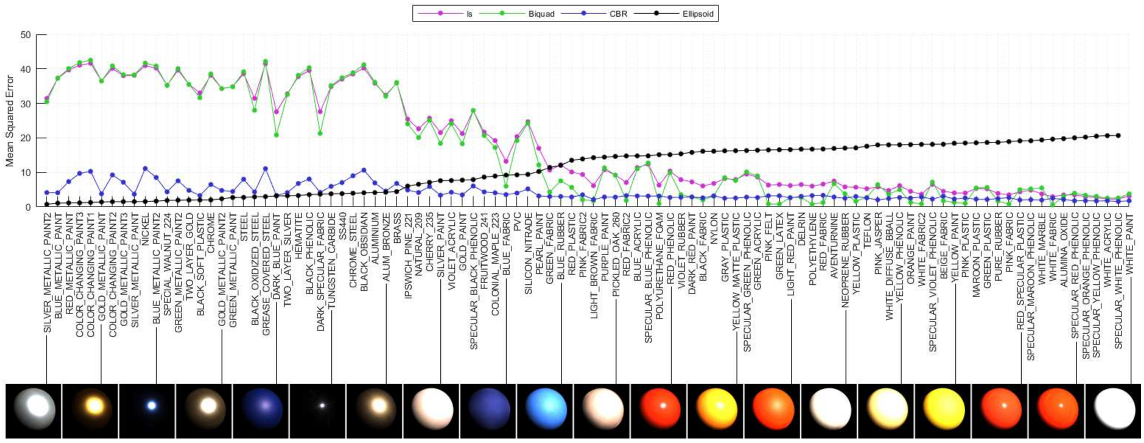

(b) Mean estimated angular error in degrees produced by various methods using 150 lights.

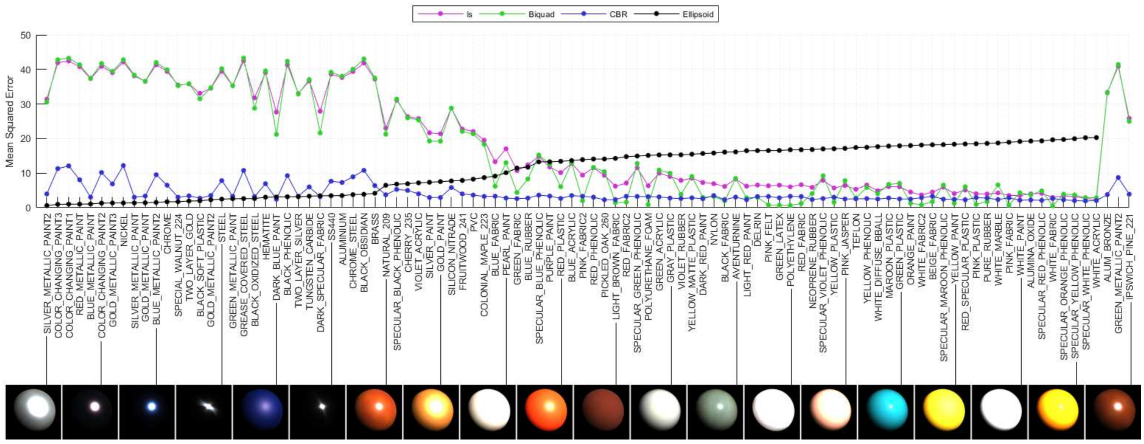

(c) Mean estimated angular error in degrees produced by various methods using 250 lights.

Figure 5.3: Performance comparison on shape recovery with specular reflectance under various lightings.

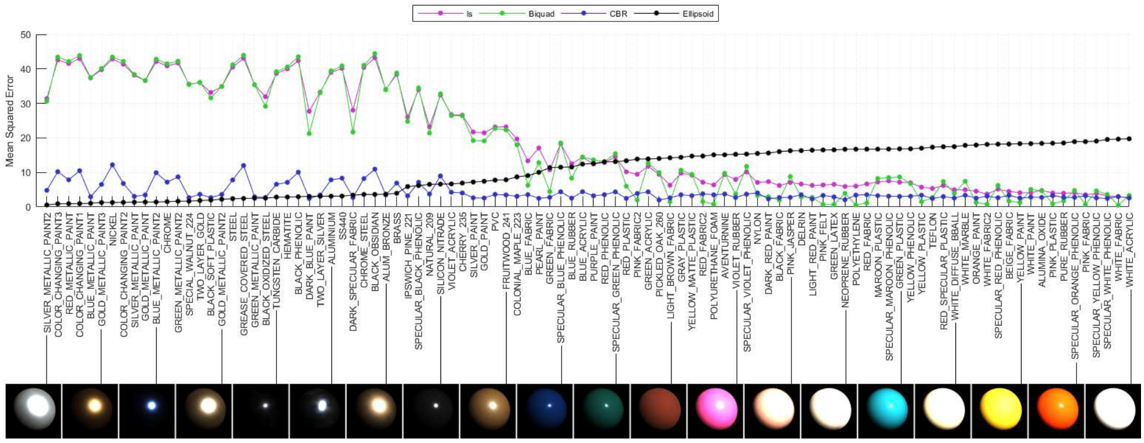

(d) Mean estimation angular error in degrees produced by various methods using 500 lights.

Figure 5.3: Performance comparison on shape recovery with specular reflectance under various lightings.

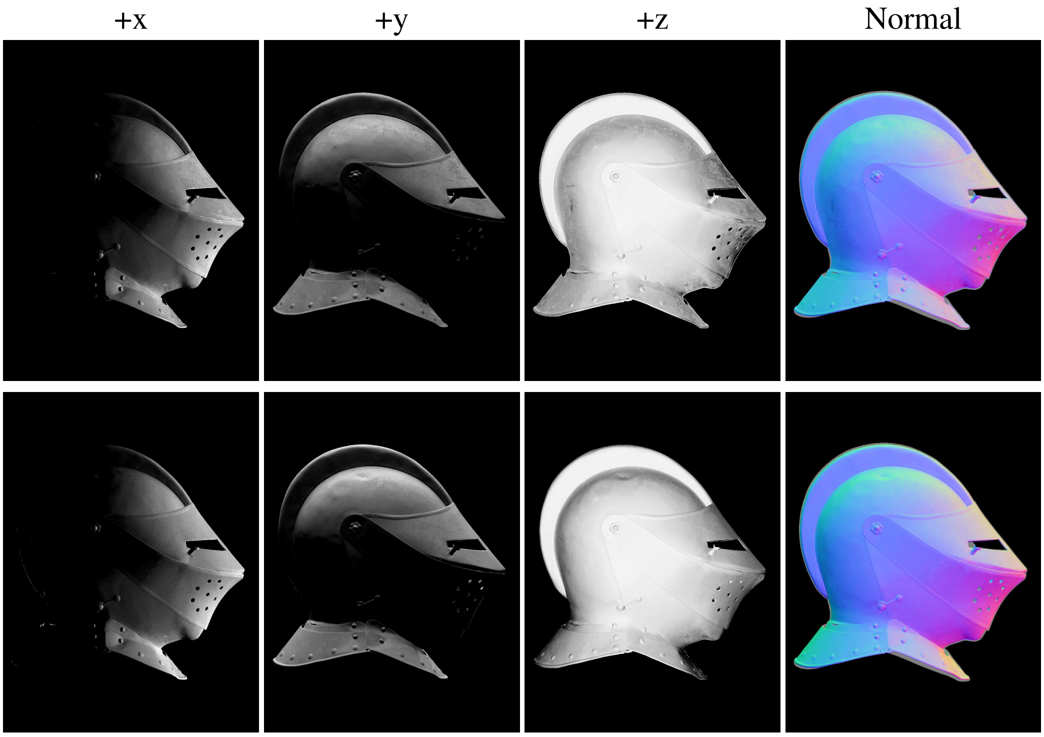

Figure 5.4: Normal maps obtained by the proposed method for “helmet front left”. Row 1: with $+z$ lights. Row 2: with 253 lights.

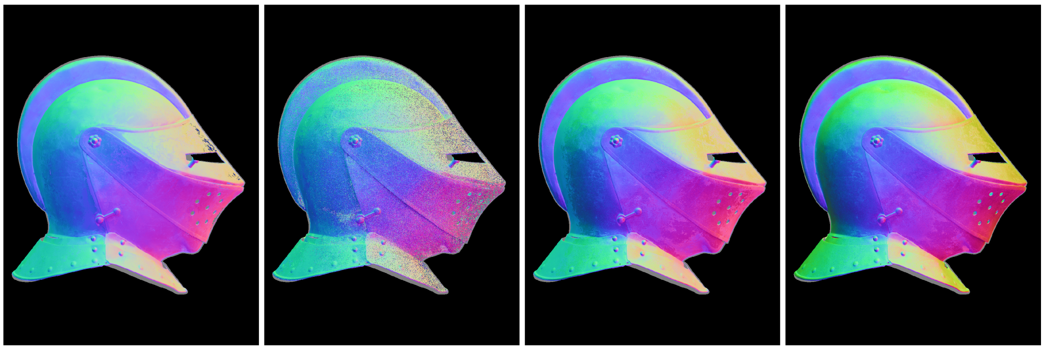

Figure 5.5: Normal maps obtained by CBR and Ellipsoid. From left to right: CBR with $+z$ lights and with 253 lights; Ellipsoid with $+z$ lights and with 253 lights;

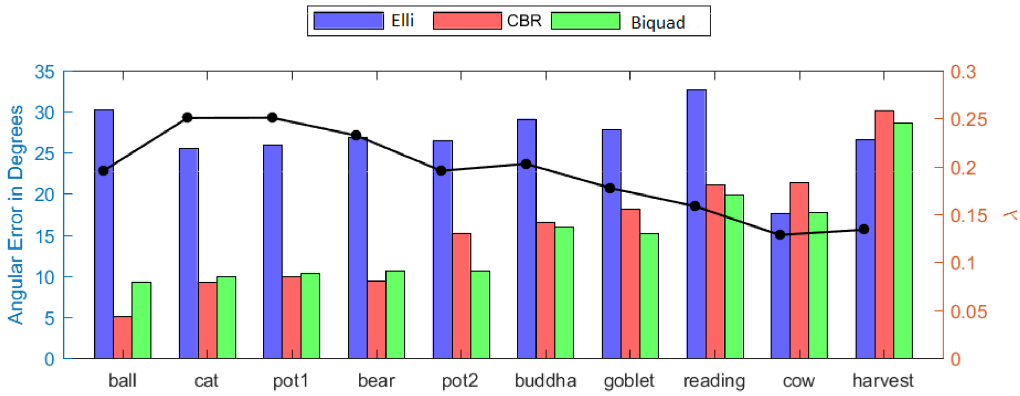

Figure 5.6: Average estimation error in degrees produced by the three methods.The black dotted line indicates the corresponding λ obtained for each model. The proposed model predicts that the trace of λ should be consistent with the trace of estimations error for homogeneous materials.

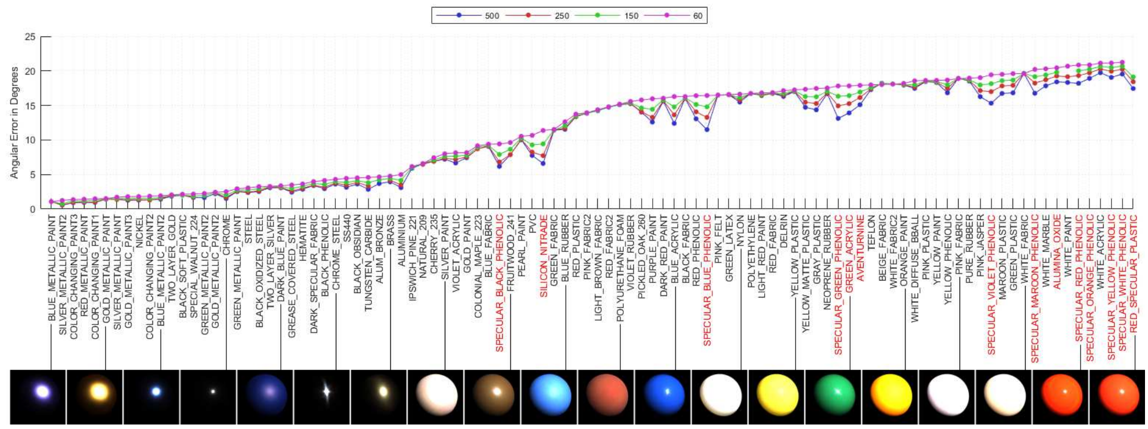

Figure 5.7: Mean estimated angular error in degrees produced by the proposed ellipsoid specular method under various lighting conditions.

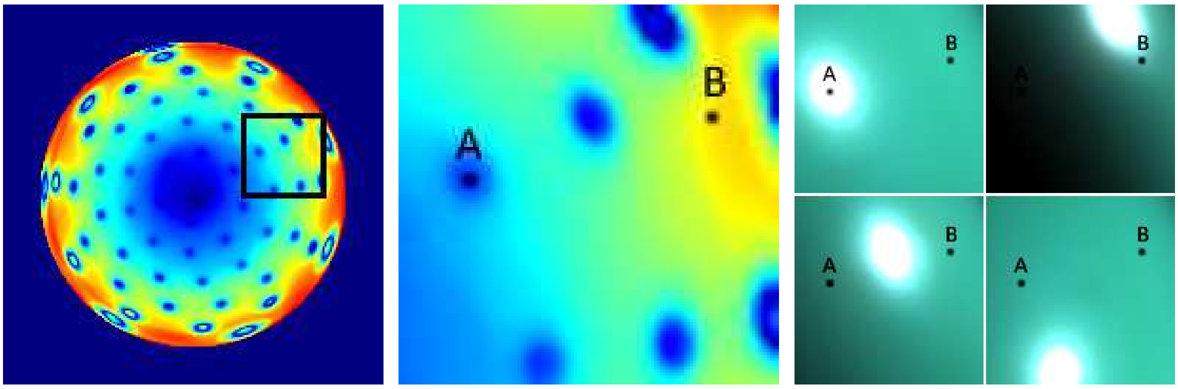

Figure 5.8: Some materials exhibit highly localized specularity, so its appearance is sensitive to light density. In terms of model fitting, the pixels that carry specular signals (*e.g.,* point A) are more likely to be correctly estimated by the proposed solver than those do not (*e.g.,* point B). Left: Estimation error for specular green phenolic. Middle: A closer view over the region showing both accurate and inaccurate estimations. Right: Appearances produced by four distinct lights. Point B is "by-passed" by all lights so it does not carry specular signals.

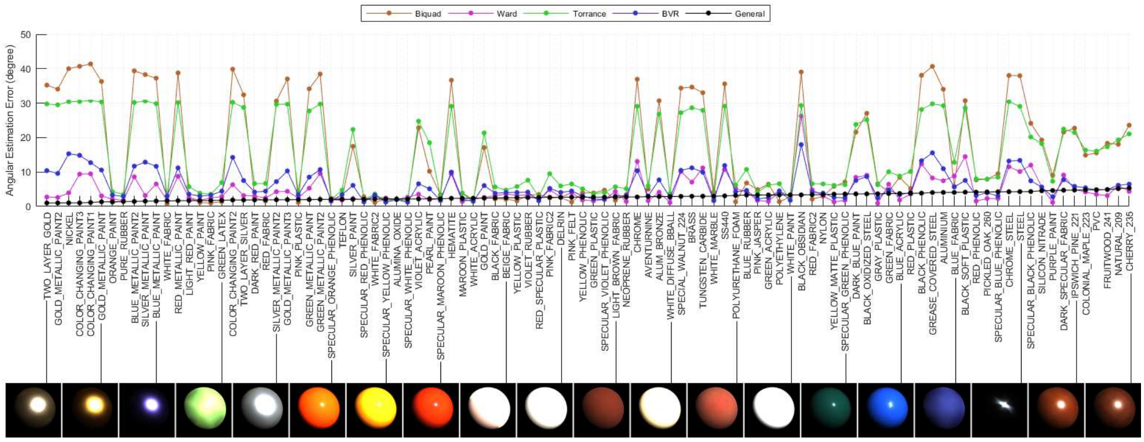

(a) Mean estimated angular error in degrees produced by various methods using 60 lights.

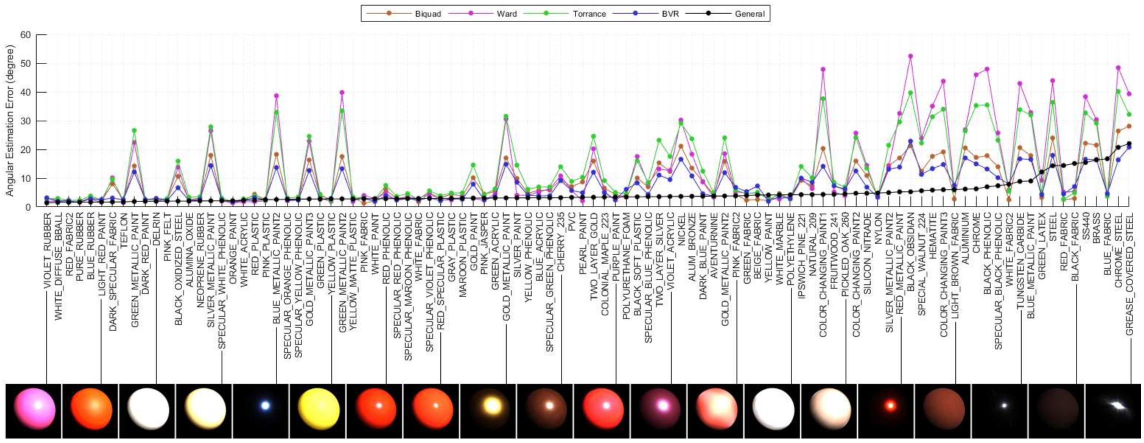

(b) Mean estimated angular error in degrees produced by various methods using 96 lights.

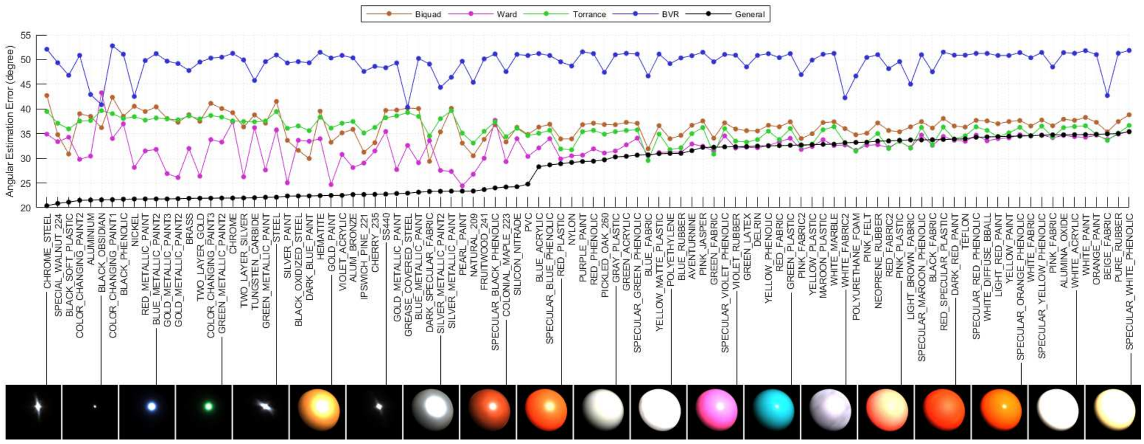

(c) Mean estimated angular error in degrees produced by various methods using 100 lights.

Figure 5.9: Performance comparison on shape recovery with general reflectance under various lightings.

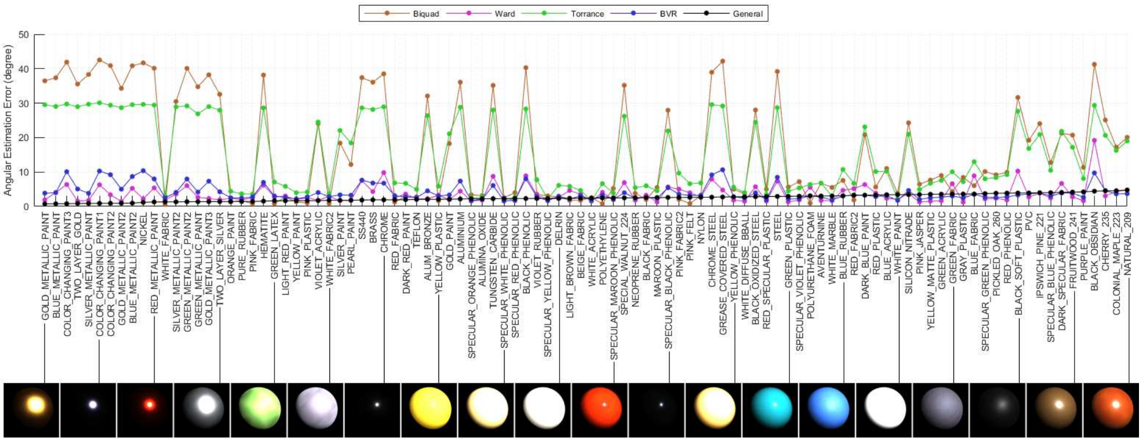

(d) Mean estimated angular error in degrees produced by various methods using 150 lights.

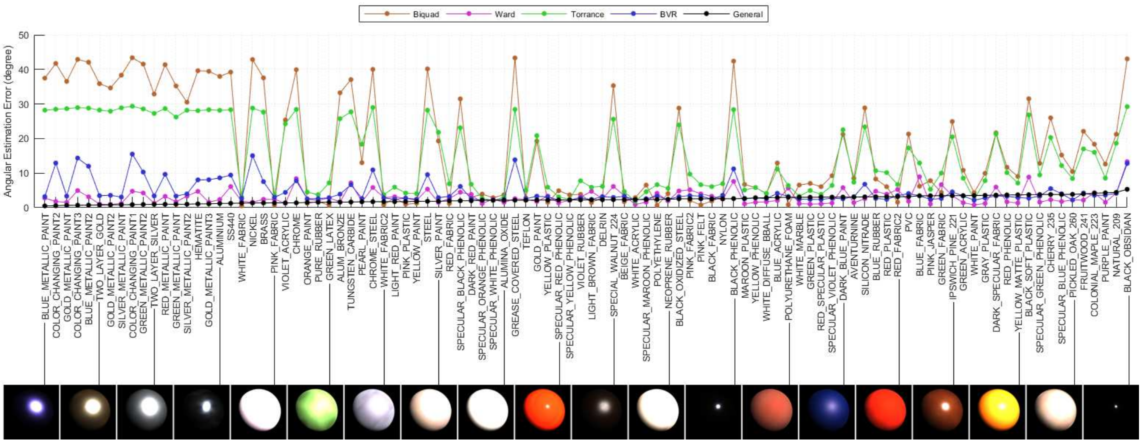

(e) Mean estimated angular error in degrees produced by various methods using 250 lights.

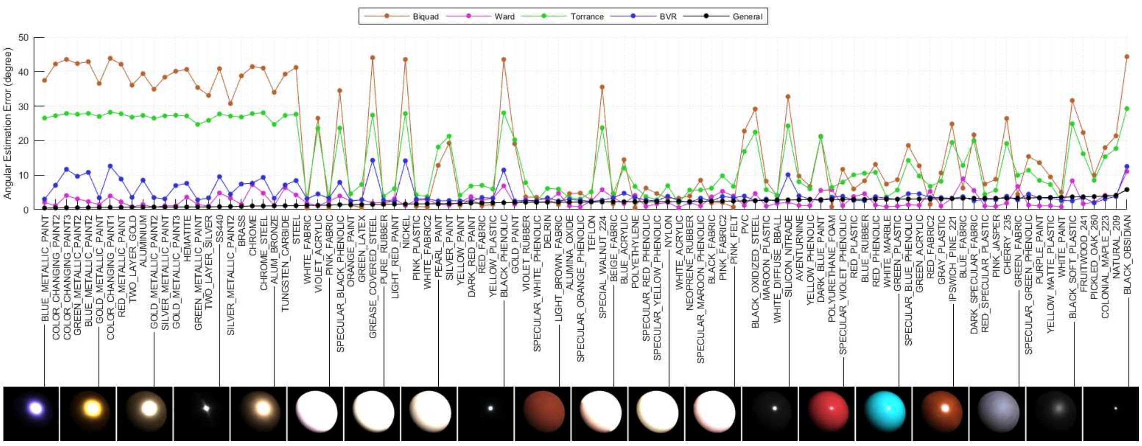

(f) Mean estimated angular error in degrees produced by various methods using 500 lights.

Figure 5.9: Performance comparison on shape recovery with general reflectance under various lightings.

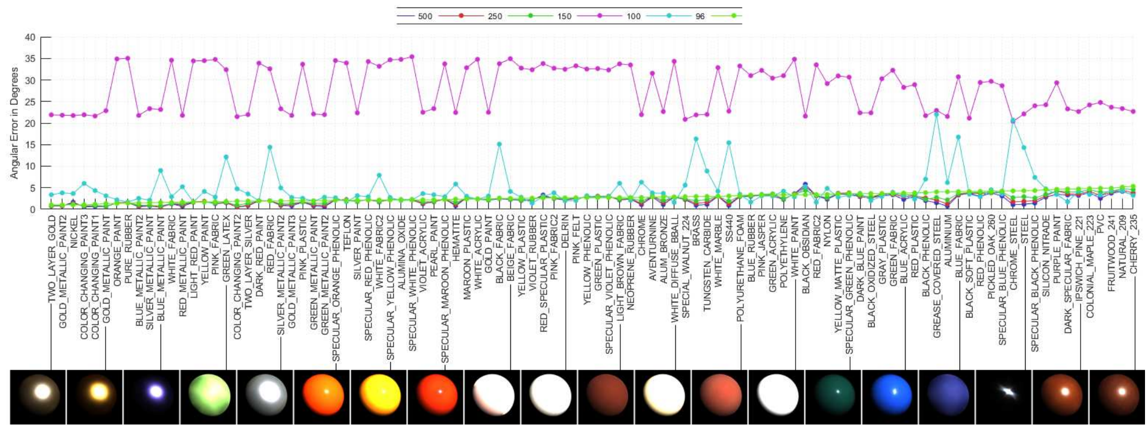

Figure 5.10: Mean estimated angular error in degrees produced by the proposed ellipsoid-based general method under various lighting conditions.

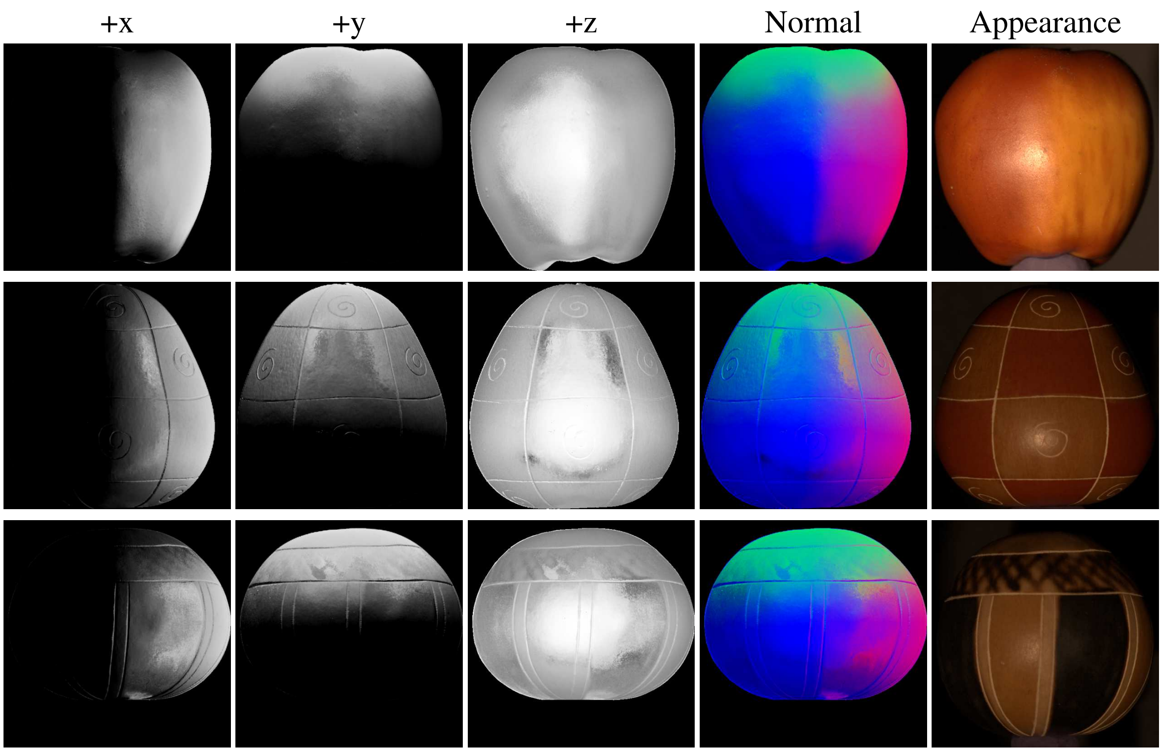

Figure 5.11: Estimated normal maps obtained by the proposed method for general reflectance applied to “apple” (row 1), “gourd1” (row 2) and “gourd2” (row 3).

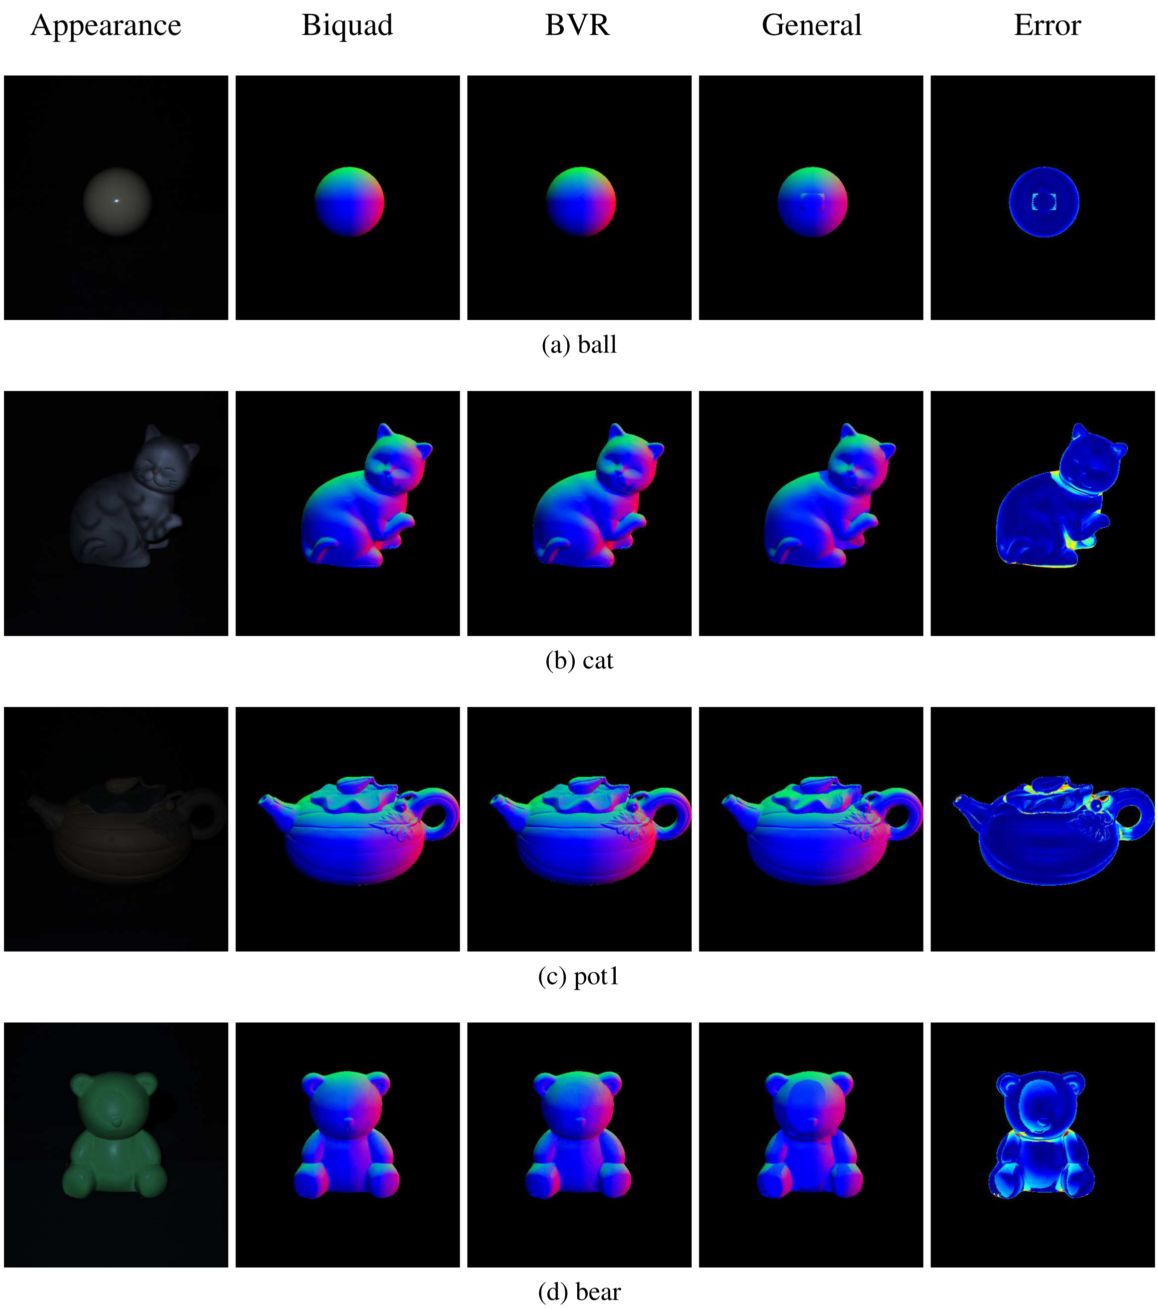

Figure 5.12: Estimated Normal Map obtained for “DiLiGent” image set using Biquad, BVR, and the proposed method for general reflectance.

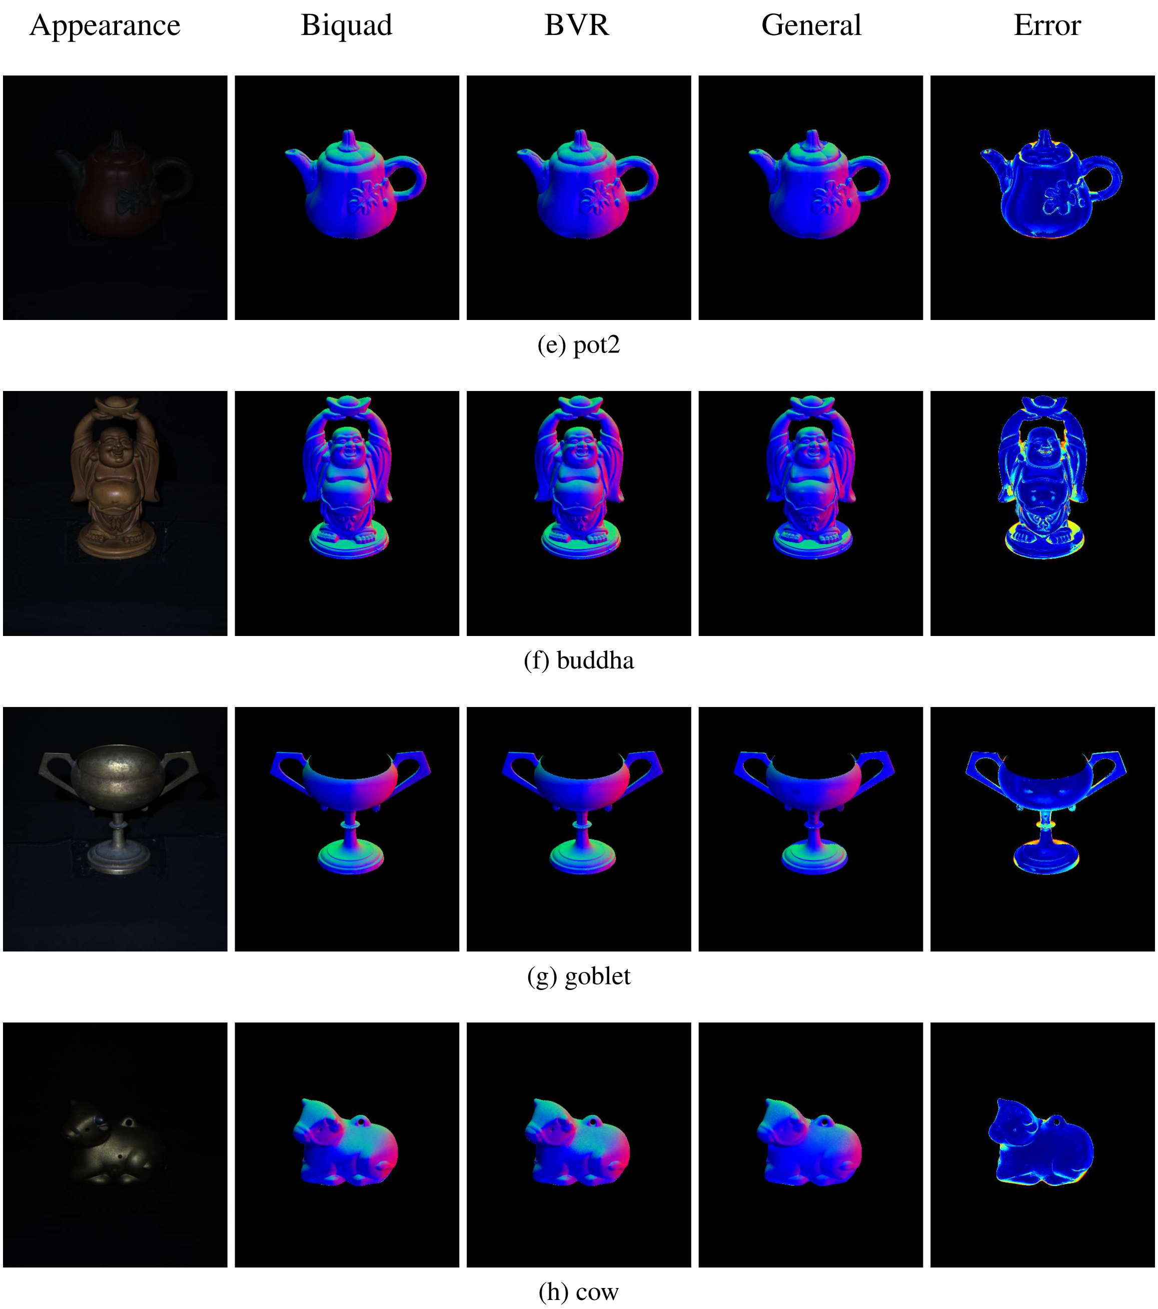

Figure 5.12: Estimated Normal Map obtained for “DiLiGent” image set using Biquad, BVR, and the proposed method for general reflectance.

Appearance Biquad BVR General Error

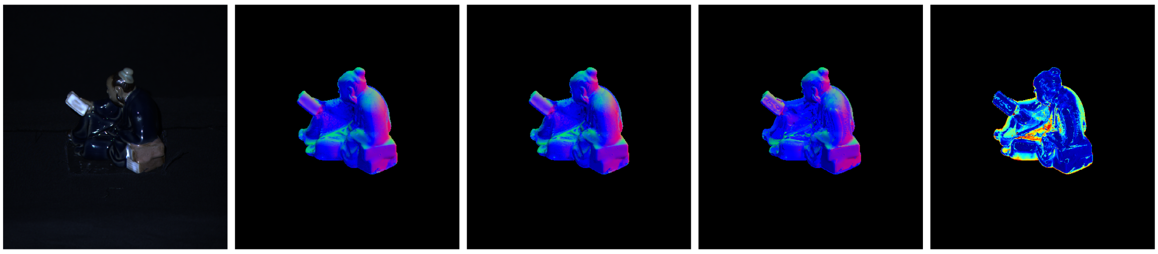

(i) reading

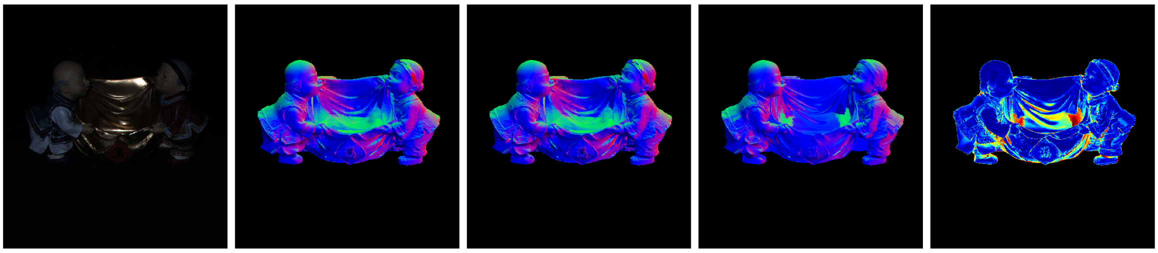

(j) harvest

Figure 5.12: Estimated Normal Map obtained for “DiLiGent” image set using Biquad, BVR, and the proposed method for general reflectance.

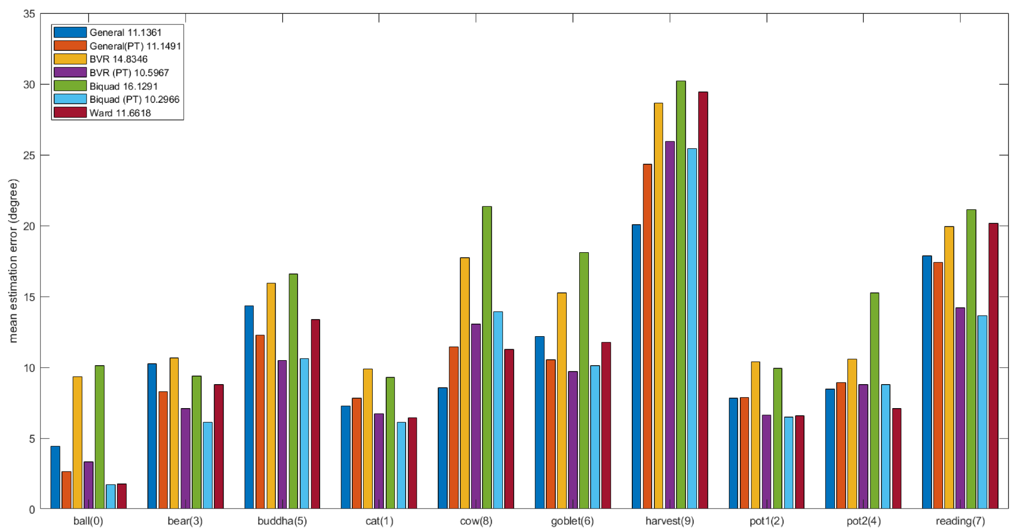

Figure 5.13: Mean estimation error in degrees producecd by the proposed method for general reflectance, BVR, Biquad and Ward model. PT indiates position threshold is applied. It can be observed that the proposed method does not rely on thresholding to obtain accurate estimation.

6 Conclusion

This thesis has described a microfacet-based isotropic reflectance model for photometric stereo. By identifying the appealing algebraic properties of its ellipsoid NDF, one is able to derive a precise representation for general isotropic reflectance. Since microfacet-based theory is physically driven, accordingly a physically interpretable approximation that is particularly serviceable for specular reflectance analysis is derived. In particular, this approximation establishes a connection between the estimation for specular surface normal and fitting an ellipsoid of revolution, where the latter can be described by a system of polynomials that can be solved by a fast, non-iterative and globally optimal solver.

Moreover, the model has also been used to reason about the necessity of treating specular pixels as outliers. Its spatially dependent approximation shows that, most non-specular observations coincidentally comply with the Lambert's model, justifying the need for thresholding by some existing approaches. But for the proposed case, an iterative solver that extends the non-iterative implementation has been proposed to handle the generalized case. Extensive experiments are performed on images of both synthetic and real objects to prove its effectiveness. Since no pre-processing is required by the proposed approach, its stability and robustness is ensured even under adversary light distribution.

On the other hand, a microfacet model that directly addresses specularity for photometric stereo invites new challenges. First, specularity may also occur due to inter-reflections on concave shapes, in such case it is misleading for shape inference. Second, if cast shadow is captured with ineligible dark noise, low intensity pixels may also produce inaccurate estimations. Third, BRDFs in their essence does not favor subsurface scattering. However, it is worth noting that these sources of inaccuracies are mainly due to high frequency light (directional light in particular), so under smart illumination these components can be extracted and eliminated [NKGR06]. Because BRDFs in general describe the appearance formation under all types of lighting, their applications under illumination pattern other than directional light may augment the toolset for photometric stereo on non-convex shapes [NIK91b].

In a broad sense, this thesis has illustrated an idea of applying microfacet-based theory to photometric stereo, which relies on the ellipsoid NDF that preserves both expressiveness and numerical invertibility. One should reasonably envision that the ellipsoid NDF might be one such model that has surfaced among the others that are still awaiting investigation. More importantly, since NDF nowadays is directly measured from a sample object with known regular geometry, acquiring the ability to recover surface orientation and to extract BRDF simultaneously is an interesting part of the follow-up work.

Bibliography

[AK07] Neil G Alldrin and David J Kriegman. Toward reconstructing surfaces with arbitrary isotropic reflectance: A stratified photometric stereo approach. In *Proc. of Internatoinal Conference on Computer Vision*, pages 1–8, 2007.

[AP07] Michael Ashikhmin and Simon Premoze. Distribution-based brdfs. *Unpublished Technical Report, University of Utah*, 2:6, 2007.

[APS00] Michael Ashikmin, Simon Premože, and Peter Shirley. A microfacet-based brdf generator. In *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pages 65–74, 2000.

[AS00] Michael Ashikhmin and Peter Shirley. An anisotropic phong brdf model. *Journal of graphics tools*, 5(2):25–32, 2000.

[AZK08] Neil Alldrin, Todd Zickler, and David Kriegman. Photometric stereo with non-parametric and spatially-varying reflectance. In *Proc. of Computer Vision and Pattern Recognition*, pages 1–8, 2008.

[BJ03] Ronen Basri and David W Jacobs. Lambertian reflectance and linear subspaces. *IEEE transactions on pattern analysis and machine intelligence*, 25(2):218–233, 2003.

[BK98] Peter N Belhumeur and David J Kriegman. What is the set of images of an object under all possible illumination conditions? *International Journal of Computer Vision*, 28(3):245–260, 1998.

[Bli77] James F Blinn. Models of light reflection for computer synthesized pictures. In *ACM SIGGRAPH Computer Graphics*, volume 11, pages 192–198, 1977.

[BÖK11] Ahmet Bilgili, Aydn Öztürk, and Murat Kurt. A general brdf representation based on tensor decomposition. In *Computer Graphics Forum*, volume 30, pages 2427–2439. Wiley Online Library, 2011.

[BS87] Petr Beckmann and Andre Spizzichino. The scattering of electromagnetic waves from rough surfaces. *Norwood, MA, Artech House, Inc., 1987, 511 p.*, 1987.

[BS12] Brent Burley and Walt Disney Animation Studios. Physically-based shading at disney. In *ACM SIGGRAPH*, pages 1–7, 2012.

[BSN16] Mahdi M Bagher, John Snyder, and Derek Nowrouzezahrai. A non-parametric factor microfacet model for isotropic brdfs. *ACM Transactions on Graphics (TOG)*, 35(5):159, 2016.

[CJ08] Hin-Shun Chung and Jiaya Jia. Efficient photometric stereo on glossy surfaces with wide specular lobes. In *Proc. of Computer Vision and Pattern Recognition*, pages 1–8, 2008.

[CT82] Robert L Cook and Kenneth E. Torrance. A reflectance model for computer graphics. *ACM Transactions on Graphics*, 1(1):7–24, 1982.

[CZS$^+$17] Lixiong Chen, Yinqiang Zheng, Boxin Shi, Art Subpa-Asa, and Imari Sato. A microfacet-based reflectance model for photometric stereo with highly specular surfaces. In *Computer Vision, 2017. ICCV 2017. IEEE International Conference on*. IEEE, 2017.

[DHI$^+$15] Jonathan Dupuy, Eric Heitz, Jean-Claude Iehl, Pierre Poulin, and Victor Ostromoukhov. Extracting microfacet-based brdf parameters from arbitrary materials with power iterations. In *Computer Graphics Forum*, volume 34, pages 21–30, 2015.

[DWMG15] Zhao Dong, Bruce Walter, Steve Marschner, and Donald P Greenberg. Predicting appearance from measured microgeometry of metal surfaces. *ACM Transactions on Graphics (TOG)*, 35(1):9, 2015.

[DWT+10] Yue Dong, Jiaping Wang, Xin Tong, John Snyder, Yanxiang Lan, Moshe Ben-Ezra, and Baining Guo. Manifold bootstrapping for svbrdf capture. In *ACM Transactions on Graphics (Proc. of ACM SIGGRAPH)*, volume 29, page 98, 2010.

[ECJ+06] Per Einarsson, Charles-Felix Chabert, Andrew Jones, Wan-Chun Ma, Bruce Lamond, Tim Hawkins, Mark T Bolas, Sebastian Sylwan, and Paul E Debevec. Relighting human locomotion with flowed reflectance fields. *Rendering techniques*, 2006:17th, 2006.

[GCHS10] Dan B Goldman, Brian Curless, Aaron Hertzmann, and Steven M Seitz. Shape and spatially-varying brdfs from photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(6):1060–1071, 2010.

[Geo03] Athinodoros S Georghiades. Incorporating the torrance and sparrow model of reflectance in uncalibrated photometric stereo. In *Proc. of Computer Vision and Pattern Recognition*, pages 816–823, 2003.

[Gol05] Ron Goldman. Curvature formulas for implicit curves and surfaces. *Computer Aided Geometric Design*, 22(7):632–658, 2005.

[GTHD03] Andrew Gardner, Chris Tchou, Tim Hawkins, and Paul Debevec. Linear light source reflectometry. In *ACM Transactions on Graphics (TOG)*, volume 22, pages 749–758, 2003.

[Hei14] Eric Heitz. Understanding the masking-shadowing function in microfacet-based brdfs. *Journal of Computer Graphics Techniques*, 3(2):32–91, 2014.

[HJ90] Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge university press, 1990.

[HMI10] Tomoaki Higo, Yasuyuki Matsushita, and Katsushi Ikeuchi. Consensus photometric stereo. In *Proc. of Computer Vision and Pattern Recognition*, 2010.

[Hor70] Berthold KP Horn. Shape from shading: A method for obtaining the shape of a smooth opaque object from one view. 1970.

[HS05] Aaron Hertzmann and Steven M Seitz. Example-based photometric stereo: Shape reconstruction with general, varying brdfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1254–1264, 2005.

[HS15] Zhuo Hui and Aswin C Sankaranarayanan. A dictionary-based approach for estimating shape and spatially-varying reflectance. In *Proc. International Conference on Computational Photography (ICCP)*, pages 1–9, 2015.

[IA14] Satoshi Ikehata and Kiyoharu Aizawa. Photometric stereo using constrained bivariate regression for general isotropic surfaces. In *Proc. of Computer Vision and Pattern Recognition*, pages 2179–2186, 2014.

[IWMA12] Satoshi Ikehata, David Wipf, Yasuyuki Matsushita, and Kiyoharu Aizawa. Robust photometric stereo using sparse regression. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 318–325. IEEE, 2012.

[IWMA14] Satoshi Ikehata, David Wipf, Yasuyuki Matsushita, and Kiyoharu Aizawa. Photometric stereo using sparse bayesian regression for general diffuse surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(9):1816–1831, 2014.

[Kaj85] James T. Kajiya. Anisotropic reflection models. *SIGGRAPH Comput. Graph.*, 19(3):15–21, July 1985.

[LÅ16] Viktor Larsson and Kalle Åström. Uncovering symmetries in polynomial systems. In *Proc. of European Conference on Computer Vision*, pages 252–267. Springer, 2016.

[LFTG97] Eric PF Lafortune, Sing-Choong Foo, Kenneth E Torrance, and Donald P Greenberg. Non-linear approximation of reflectance functions. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 117–126, 1997.

[LRR04] Jason Lawrence, Szymon Rusinkiewicz, and Ravi Ramamoorthi. Efficient brdf importance sampling using a factored representation. In *ACM Transactions on Graphics (TOG)*, volume 23, pages 496–505. ACM, 2004.

[LW94] Eric P Lafortune and Yves D Willems. Using the modified phong reflectance model for physically based rendering. 1994.

[MB03] Wojciech Matusik and Matt Brand. A data-driven reflectance model. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)*, 22(3):759–769, 2003.

[MZKB05] Satya P Mallick, Todd E Zickler, David J Kriegman, and Peter N Belhumeur. Beyond lambert: Reconstructing specular surfaces using color. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, volume 2, pages 619–626. Ieee, 2005.

[NDM05] Addy Ngan, Frédo Durand, and Wojciech Matusik. Experimental analysis of brdf models. *Rendering Techniques*, 2005(16th):2, 2005.

[NIK91a] S. K. Nayar, K. Ikeuchi, and T. Kanade. Surface reflection: physical and geometrical perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):611–634, Jul 1991.

[NIK91b] Shree K Nayar, Katsushi Ikeuchi, and Takeo Kanade. Shape from interreflections. *International Journal of Computer Vision*, 6(3):173–195, 1991.

[NKGR06] Shree K Nayar, Gurunandan Krishnan, Michael D Grossberg, and Ramesh Raskar. Fast separation of direct and global components of a scene using high frequency illumination. In *ACM Transactions on Graphics (TOG)*, volume 25, pages 935–944. ACM, 2006.

[NLW+16] Giljoo Nam, Joo Ho Lee, Hongzhi Wu, Diego Gutierrez, and Min H Kim. Simultaneous acquisition of microscale reflectance and normals. *ACM TRANSACTIONS ON GRAPHICS*, 35(6), 2016.

[ON94] Michael Oren and Shree K Nayar. Generalization of lambert's reflectance model. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 239–246, 1994.

[Pho75] Bui Tuong Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, 1975.

[PJH16] Matt Pharr, Wenzel Jakob, and Greg Humphreys. *Physically based rendering: From theory to implementation*. Morgan Kaufmann, 2016.

[Ram02] Ravi Ramamoorthi. Analytic pca construction for theoretical analysis of lighting variability in images of a lambertian object. *IEEE transactions on pattern analysis and machine intelligence*, 24(10):1322–1333, 2002.

[RSZ94] Evguenii A Rakhmanov, EB Saff, and YM Zhou. Minimal discrete energy on the sphere. *Math. Res. Lett*, 1(6):647–662, 1994.

[Sch94] Christophe Schlick. An inexpensive brdf model for physically-based rendering. In *Computer graphics forum*, volume 13, pages 233–246. Wiley Online Library, 1994.

[Sha85] Steven A Shafer. Using color to separate reflection components. *Color Research & Application*, 10(4):210–218, 1985.

[SI94] Yoichi Sato and Katsushi Ikeuchi. Temporal-color space analysis of reflection. *JOSA A*, 11(11):2990–3002, 1994.

[Smi67] B Smith. Geometrical shadowing of a random rough surface. *IEEE Transactions on Antennas and Propagation*, 15(5):668–671, 1967.

[STMI12] Boxin Shi, Ping Tan, Yasuyuki Matsushita, and Katsushi Ikeuchi. Elevation angle from reflectance monotonicity: Photometric stereo for general isotropic reflectances. In *Proc. of European Conference on Computer Vision*, pages 455–468, 2012.

[STMI14] Boxin Shi, Ping Tan, Yasuyuki Matsushita, and Katsushi Ikeuchi. Bi-polynomial modeling of low-frequency reflectances. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(6):1078–1091, 2014.

[SWM+16] Boxin Shi, Zhe Wu, Zhipeng Mo, Dinglong Duan, Sai-Kit Yeung Yeung, and Ping Tan. A benchmark dataset and evaluation for non-lambertian and uncalibrated photometric stereo. In *Proc. of Computer Vision and Pattern Recognition*, 2016.

[TMDB16] S. Tozza, R. Mecca, M. Duocastella, and A. Del Bue. Direct differential photometric stereo shape recovery of diffuse and specular surfaces. *Journal of Mathematical Imaging and Vision*, 56:57–76, 2016.

[TR75] T. S. Trowbridge and K. P. Reitz. Average irregularity representation of a rough surface for ray reflection. *Journal of the Optical Society of America*, 65(5):531–536, 1975.

[TS67] Kenneth E Torrance and Ephraim M Sparrow. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America*, 57(9):1105–1112, 1967.

[Vic96] GT Vickers. The projected areas of ellipsoids and cylinders. *Powder technology*, 86(2):195–200, 1996.

[War92] Gregory J Ward. Measuring and modeling anisotropic reflection. *ACM SIGGRAPH Computer Graphics*, 26(2):265–272, 1992.

[WDMG16] Bruce Walter, Zhao Dong, Steve Marschner, and Donald P Greenberg. The ellipsoid normal distribution function. *Supplementary material*, http://www.cs.cornell.edu/Projects/metalappearance/, 2016.

[WGS+10] Lun Wu, Arvind Ganesh, Boxin Shi, Yasuyuki Matsushita, Yongtian Wang, and Yi Ma. Robust photometric stereo via low-rank matrix completion and recovery. In *Proc. of Asian Conference on Computer Vision*, pages 703–717, 2010.

[WMLT07] Bruce Walter, Stephen R Marschner, Hongsong Li, and Kenneth E Torrance. Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques*, pages 195–206, 2007.

[Woo80] Robert J Woodham. Photometric method for determining surface orientation from multiple images. *Optical engineering*, 19(1):191139–191139, 1980.

[WT13] Zhe Wu and Ping Tan. Calibrating photometric stereo by holistic reflectance symmetry analysis. In *Proc. of Computer Vision and Pattern Recognition*, pages 1498–1505, 2013.

Acronyms

BRDF	Bidirectional Reflectance Distribution Function
CBR	Constrained Bivariate Regression
NLS	Nonlinear Least Square
NDF	Normal Distribution Function
PBR	Physically Based Rendering
PT	Position Threshold