

A Benchmark Dataset and Evaluation for Non-Lambertian and Uncalibrated Photometric Stereo

Boxin Shi^{1,2*} Zhe Wu^{2,3} Zhipeng Mo² Dinglong Duan³ Sai-Kit Yeung² Ping Tan⁴

¹Artificial Intelligence Research Center, National Institute of AIST

²Singapore University of Technology and Design

³National University of Singapore ⁴Simon Fraser University

Abstract

Recent progress on photometric stereo extends the technique to deal with general materials and unknown illumination conditions. However, due to the lack of suitable benchmark data with ground truth shapes (normals), quantitative comparison and evaluation is difficult to achieve. In this paper, we first survey and categorize existing methods using a photometric stereo taxonomy emphasizing on non-Lambertian and uncalibrated methods. We then introduce the ‘DiLiGenT’ photometric stereo image dataset with calibrated Directional Lightings, objects of General reflectance, and ‘ground Truth’ shapes (normals). Based on our dataset, we quantitatively evaluate state-of-the-art photometric stereo methods for general non-Lambertian materials and unknown lightings to analyze their strengths and limitations.

1. Introduction

Photometric stereo computes surface normals from radiance variation due to illumination changes. The classic algorithm [1] assumes the Lambertian surface, a fixed orthographic camera with linear radiometric response, and known directional illuminations to recover a per-pixel surface normal from at least three images. Many algorithms have been developed to extend it to surfaces with general reflectance [2], uncalibrated lighting [3], general environment illumination [4], nonlinear camera [5], multiple viewpoints [6], subpixel resolution [7], moving objects [8], and so on. However, unlike the studies in binocular and multi-view stereo [9], a standard taxonomy of photometric stereo algorithms is not yet available; benchmark dataset and evaluation, which greatly inspire future research directions, are available in binocular and multi-view stereo [9] but not in photometric stereo community.

*Part of this work was finished when B. Shi was working at Nanyang Technological University.

This paper aims to provide a taxonomy, a benchmark dataset and evaluation for photometric stereo algorithms. Unlike multi-view stereo, photometric stereo algorithms have a huge diversity in model assumptions and imaging conditions (e.g., single or multiple viewpoints, orthographic or perspective camera, directional or environmental lighting, etc.). Therefore, benchmarking photometric stereo is more feasible for a specific category of algorithms. We assume a fixed orthographic camera and directional lighting and focus on non-Lambertian and uncalibrated photometric stereo algorithms. Directional lighting and fixed orthographic camera is the most classical data capture setting [1]. Many other extensions (e.g., multi-view, environment lighting) can be built upon this setup. Under this setup, non-Lambertian reflectance and uncalibrated lighting directions are the two major difficulties for photometric stereo.

We first provide a taxonomy of photometric stereo methods, with an emphasis on recent non-Lambertian and uncalibrated algorithms, followed by a brief categorization for other methods with different lighting models, different numbers of input images, and several additional properties. We then present the ‘DiLiGenT’ photometric stereo image dataset which is captured under carefully calibrated directional lighting, with reference shapes for all objects from a high-end laser scanner provided. Our dataset contains objects of various materials and shapes, from ideally diffuse to highly reflective, from smooth to bristly. Finally, we quantitatively evaluate recently representative non-Lambertian [2, 10–16] and uncalibrated [3, 17–20] photometric stereo algorithms to understand their pros and cons and motivate future research on unsolved issues.

Our main contributions are threefold:

- An up-to-date survey/taxonomy on recent photometric stereo techniques with a special focus on non-Lambertian and uncalibrated methods¹;
- The first photometric stereo dataset with ‘ground truth’

¹We refer the readers to [21] and [22] for more comprehensive surveys on earlier and recent photometric stereo techniques.

shapes and objects of complex reflectance for single-view methods under directional lighting;

- A quantitative evaluation of recent algorithms for non-Lambertian and uncalibrated photometric stereo.

The dataset and evaluation results are available for download from <https://sites.google.com/site/photometricstereodata/>.

2. A Photometric Stereo Taxonomy

We categorize photometric stereo methods employing directional lightings and a fixed orthographic camera. With such a data capture setup, the main research challenges for photometric stereo is how to deal with non-Lambertian materials (Section 2.1) and unknown lighting conditions (Section 2.2). To be comprehensive, we also briefly review other algorithms with different lighting and camera models in Section 2.3 and number of images required for different algorithms in Section 2.4. We conclude this section by summarizing and labeling all surveyed methods in Section 2.5 and the ‘References’ section.

2.1. Non-Lambertian photometric stereo

The Lambert’s reflectance model is widely adopted in photometric stereo for its simplicity. However, the real world is full of non-Lambertian objects. Many photometric stereo algorithms have been developed to deal with non-Lambertian materials.

Non-Lambertian photometric stereo problems could be summarized as in the top row of Table 1. Here, $\mathbf{I}_{P \times Q}$ is the observation matrix of P points under Q different lighting conditions; $\mathbf{N}_{P \times 3}^T$ and $\mathbf{L}_{3 \times Q}$ are stacks of normal vector $\mathbf{n} \in \mathbb{R}^{3 \times 1}$ and lighting vector $\mathbf{l} \in \mathbb{R}^{3 \times 1}$ respectively; $\rho(\mathbf{n}, \mathbf{l})_{P \times Q}$ is the Bidirectional Reflectance Distribution Function (BRDF) values for each observation, which is a function of normal and lighting directions given viewing direction fixed at $\mathbf{v} = (0, 0, 1)^T$; and ‘ \circ ’ denotes the element-wise multiplication.

The Lambertian photometric stereo [1] (BASELINE) simply assumes $\rho(\mathbf{n}, \mathbf{l})$ is irrelevant of \mathbf{n} and \mathbf{l} , and it is an unknown constant scaling for each normal; attached shadow ($\max\{\cdot, 0\}$) can also be ignored. Given calibrated \mathbf{L} , the surface normal matrix \mathbf{N} is solved by linear least squares, and the reflectance values are approximated by normalizing each row of the estimated \mathbf{N} . The goal of non-Lambertian photometric stereo methods is to estimate \mathbf{N} when $\rho(\mathbf{n}, \mathbf{l})$ is a general unknown reflectance function.

Outlier rejection based methods assume non-Lambertian phenomena (such as shadow and specular highlight) are local and sparse, such that they can be detected and discarded as outliers. Earlier methods [23–25] select three optimal lights out of four where the surface appears mostly Lambertian to estimate normal. With more input images, this subset of Lambertian images can be

extracted in a more robust manner by graph cuts [26], RANSAC [27, 28], maximum-likelihood estimation [29], maximum feasible subsystem [30], taking the median values [31], or expectation maximization [32]. Instead of discarding non-Lambertian reflection as outliers, Zickler *et al.* [33] analyze in a subspace that is free from highlights.

Recent methods apply robust statistical techniques to reject outliers, by assuming a low-rank (Lambertian) observation matrix plus a sparse outlier matrix (shadow and specular) [10, 11] (WG10, IW12), as shown in Table 1. The difference that is WG10 aims at minimizing the rank of matrix, while IW12 employs sparse Bayesian regression to explicitly enforce the rank-3 constraint. Such rank minimization approaches have the advantages of robustness, since the only assumption to the outlier matrix is its sparsity regardless of the error magnitude and distribution.

Outlier rejection based methods generally cannot deal with materials with broad and soft specular reflection, where the non-Lambertian outliers are dense. They often require more input images to make statistical analysis.

Analytic reflectance models are employed in some photometric stereo algorithms to model non-Lambertian reflectance. Instead of discarding specular reflection as outliers, they fit a nonlinear analytic BRDF to interpret all the observed data. Therefore, they have the advantage of exploiting all available data. Along this direction, various analytic BRDF models have been incorporated, including the Torrance-Sparrow model [34], the Ward model [35], a mixture of multiple Ward models [2, 5], and the specular spike [36].

Take the method in [2] (GC10) as an example. GC10 adopts the Ward model to explicitly represent $\rho(\mathbf{n}, \mathbf{l})$. Each Ward model involves several parameters *i.e.*, diffuse strength d , specular strength s , and surface roughness α , to describe the shininess of the surface, as summarized in Table 1. In order to handle spatially-varying BRDFs, GC10 linearly combines several $\rho_i(d_i, s_i, \alpha_i)$ (i -th basis materials) multiplied by a spatially-varying weight w_i . GC10 needs to solve a nonlinear system for $\{d, s, \alpha, w\}$ besides estimating surface normal \mathbf{n} , which is computationally expensive.

The drawback of such approaches is that the analytic models vary significantly from material to material, and each of them is limited to a narrow class of materials. Furthermore, these models are highly nonlinear. So such approaches in principal require complicated case by case analysis for different material classes.

General properties of a BRDF such as isotropy, reciprocity, and monotonicity are valid for a broader class of materials. Photometric stereo algorithms capitalizing on these general properties have the potential to deal with much broader types of materials.

Many real world materials are isotropic, which show

Table 1. Summary of non-Lambertian photometric stereo assumptions and formulations.

Solve \mathbf{N} from $\mathbf{I} = \max\{\boldsymbol{\rho}(\mathbf{n}, \mathbf{l}) \circ (\mathbf{N}^\top \mathbf{L}), 0\}$ by using different assumptions and constraints on $\boldsymbol{\rho}(\mathbf{n}, \mathbf{l})$ Notations: $\mathbf{h} = (\mathbf{l} + \mathbf{v})/\ \mathbf{l} + \mathbf{v}\ $, $\theta_h = \langle \mathbf{n}, \mathbf{h} \rangle = \arccos(\mathbf{n}^\top \mathbf{h})$, $\theta_d = \langle \mathbf{l}, \mathbf{h} \rangle = \arccos(\mathbf{l}^\top \mathbf{h})$	
BASELINE	$\boldsymbol{\rho}(\mathbf{n}, \mathbf{l}) \approx \mathbf{D}$, where each row of \mathbf{D} is a constant representing the albedo of a Lambertian surface
WG10	$\boldsymbol{\rho}(\mathbf{n}, \mathbf{l}) \approx \mathbf{D} + \mathbf{E}$, where \mathbf{E} is sparse and $\text{rank}(\mathbf{I})$ is minimized
IW12	$\boldsymbol{\rho}(\mathbf{n}, \mathbf{l}) \approx \mathbf{D} + \mathbf{E}$, where \mathbf{E} is sparse and $\text{rank}(\mathbf{I}) = 3$
GC10	$\boldsymbol{\rho}(\mathbf{n}, \mathbf{l}) \approx \sum_i \mathbf{w}_i \circ \boldsymbol{\rho}_i(d_i, s_i, \alpha_i)$, where $\boldsymbol{\rho}_i(d_i, s_i, \alpha_i) = \frac{d_i}{\pi} + \frac{s_i}{4\pi\alpha_i^2 \sqrt{(\mathbf{n}^\top \mathbf{l})(\mathbf{n}^\top \mathbf{v})}} \exp\left(\frac{(\mathbf{l} - \mathbf{l}/\mathbf{n}^\top \mathbf{h})}{\alpha_i^2}\right)$
AZ08	$\boldsymbol{\rho}(\mathbf{n}, \mathbf{l})$ is isotropic and depends only on (θ_h, θ_d)
ST12	$\boldsymbol{\rho}(\mathbf{n}, \mathbf{l})$ is isotropic, depends only on θ_h , and is monotonic about $\mathbf{n}^\top \mathbf{h}$
HM10	$\boldsymbol{\rho}(\mathbf{n}, \mathbf{l})$ is isotropic, monotonic about $\mathbf{n}^\top \mathbf{l}$, and $\boldsymbol{\rho}(\mathbf{n}, \mathbf{l}) = 0$ for $\mathbf{n}^\top \mathbf{l} \leq 0$
ST14	The low-frequency part of $\boldsymbol{\rho}(\mathbf{n}, \mathbf{l})$ is a bi-polynomial $A(\cos(\theta_h))B(\cos(\theta_d))$, where A, B are polynomials
IA14	$\boldsymbol{\rho}(\mathbf{n}, \mathbf{l}) \approx \sum_i \boldsymbol{\rho}_i(\mathbf{n}^\top \alpha_i)$, where $\alpha_i = (p_i \mathbf{l} + q_i \mathbf{v})/\ p_i \mathbf{l} + q_i \mathbf{v}\ $, p_i, q_i are nonnegative unknown values

structured BRDF values. The isotropy means equal reflectance values are observed if \mathbf{n} is symmetric about the plane spanned by \mathbf{v} and \mathbf{l} . By exploring isotropy, the azimuth angle of a surface normal can be computed [37], and ‘iso-depth’ and ‘iso-slope’ contours [37–39] can be derived. It has been proved in [38, 39] that these two set of contours determine the surface shape.

Many isotropic BRDFs can be well approximated by bivariate functions [40], which allows iterative estimation of surface normals and BRDF represented using a 2D discrete table [12] (AZ08). By further assuming this bivariate function is monotonic in one dimension, Shi *et al.* [13] (ST12) estimate the elevation angle of a surface normal by searching for the correct candidate showing monotonic BRDF values. Both AZ08 and ST12 are ‘two-step’ approaches that first estimate the azimuth angle using isotropic constraint [37] before the elevation angle estimation. Isotropy and monotonicity are further combined with visibility constraint in [14] (HM10) to deal with general materials whose BRDF consists of a single lobe. Some recent methods [15, 16] (ST14, IA14) develop novel bivariate BRDF representations to facilitate normal estimation. They use two thresholds, one to exclude shadows and the other one to exclude specular reflections, so the remaining reflectance values vary slowly. The specular threshold is necessary for ST14, but optional for IA14. ST14 models such low-frequency reflectance using a bi-polynomial representation, while IA14 models it as a sum of lobes with unknown center directions. Please refer to Table 1 for their analytic formulae and constraints on general BRDFs.

Handling anisotropic BRDF is a well-known difficult problem for photometric stereo. Holroyd *et al.* [41] utilize symmetric elliptical microfacet normal distribution function to solve this issue. This can also be solved by example-based method like [42], which captures example (known shape) and target (unknown shape) objects of the same material under varying lighting conditions. The example-based approach has been adopted to build a high-resolution micro-

geometry sensor [43].

The photometric stereo problem can also be solved by embedding all pixels to a hemisphere (*i.e.*, a 2D manifold representing all possible normal directions) while keeping appropriate distances between different pixels. Such distances are measured before embedding by similarity of pixel intensity profiles derived from analytic BRDFs [44], general isotropic BRDFs [20], or attached shadow codes [45] which in principal can even handle anisotropic BRDFs.

2.2. Uncalibrated photometric stereo

Many photometric stereo algorithms assume known lighting conditions. Typically, this requires inserting additional calibration objects, such as a mirror sphere, into the scene during data capture. This process is tedious and the calibration sphere often causes interreflection on other scene objects. So uncalibrated methods are developed to automatically calibrate lighting conditions.

Most uncalibrated photometric stereo methods are Lambertian and based on the factorization technique proposed in [46]. The Lambertian assumption simplifies $\boldsymbol{\rho}(\mathbf{n}, \mathbf{l})$ as \mathbf{D} with each row of \mathbf{D} being a constant, so that $\mathbf{I} = \mathbf{S}^\top \mathbf{L}$ and \mathbf{S}^\top encodes the albedo-scaled normals. By applying SVD, the pseudo-normal $\tilde{\mathbf{S}}$, which is different from the true albedo-scaled normal \mathbf{S} by a 3×3 ambiguity matrix \mathbf{A} , can be decomposed, as shown in the top row of Table 2. Resolving this intrinsic shape-lighting ambiguity requires additional information. It is solved up to a rotation ambiguity, if six points with the same albedo or six lights of the same intensity can be identified [46]. If the surface is integrable [47] or the shadow boundary is observed [48], this ambiguity is reduced to the General Bas-Relief (GBR) ambiguity [49], which is denoted as the matrix \mathbf{G} with only three unknowns. Since then, many algorithms have been developed to further resolve this GBR ambiguity. There are methods relying on analyzing the specular spikes [50, 51], parametric specular reflection (*e.g.*, the Torrance-Sparrow model) [34], isotropic specular reflection [19, 52, 53], inter-

Table 2. Summary of uncalibrated photometric stereo constraints and solutions.

Solve \mathbf{N} from $\mathbf{I} = \max\{\boldsymbol{\rho}(\mathbf{n}, \mathbf{l}) \circ (\mathbf{N}^\top \mathbf{L}), 0\}$ when \mathbf{L} is unknown For Lambertian objects, $\mathbf{I} = \max\{\mathbf{D} \circ (\mathbf{N}^\top \mathbf{L}), 0\} = \mathbf{S}^\top \mathbf{L} = \tilde{\mathbf{S}}^\top \mathbf{A}^\top \mathbf{A}^{-1} \tilde{\mathbf{L}} = \hat{\mathbf{S}}^\top \mathbf{G}^\top \mathbf{G}^{-1} \hat{\mathbf{L}}$	
AM07	\mathbf{D} has only a few different albedos, <i>i.e.</i> , the rows of \mathbf{S} have only a few different lengths
SM10	Several surface points have equal albedo, <i>i.e.</i> , several rows of \mathbf{S} having equal length are identified
PF14	Several points with locally maximum intensity on a Lambertian surface, <i>i.e.</i> , points with $\mathbf{n} = \mathbf{l}$ are identified
WT13	$\boldsymbol{\rho}(\mathbf{n}, \mathbf{l}) \approx \mathbf{D} + \boldsymbol{\rho}_s(\theta_h, \theta_d)$, <i>i.e.</i> , the specular reflection depends only on $\{\theta_h, \theta_d\}$

reflection [54], a ring of light sources [55], and a perspective camera [56, 57].

We summarize some recent methods and their constraints for solving the ambiguity matrix \mathbf{G} in Table 2. Alldrin *et al.* [3] (AM07) assume the distribution of albedos in a natural image contains a few dominant values (*i.e.*, a few dominant colors of a scene), so that the GBR transform can be found by minimizing the entropy of this distribution. Shi *et al.* [17] (SM10) automatically identify pixels with the same albedo by chromaticity clustering to solve \mathbf{G} . Papadhimetri and Favaro [18] (PF14) propose to identify local diffuse reflectance (LDR) maxima, where the normal vector is coincident with the lighting direction. Wu and Tan [19] (WT13) assume an additive bivariate specular reflection to solve the GBR ambiguity. SM10 and PF14 need to select desired pixels (or lighting directions), but their solutions for \mathbf{G} are quite simple, *i.e.*, a linear problem for SM10 and a close-form solution for PF14. AM07 and WT13 seek the optimal three unknowns in \mathbf{G} by a brute-force search.

Manifold embedding based methods [20, 44, 45] can also work with uncalibrated lighting conditions. They first recover normals up to a global rotation ambiguity, which is removed by boundary or integrability constraint. In particular, the method in [20] (LM13) directly handles general isotropic BRDFs by mapping the distances of intensity profiles to angular differences of normals. Such a method cannot be described using the framework in Table 2.

2.3. Other categorization

To be comprehensive, we briefly review other photometric stereo algorithms beyond the classic setup of fixed orthographic camera and directional lighting.

Point lighting based methods adopt the more precise point light source model with spatially-varying lighting directions, while directional lighting assumes all pixels share the same lighting direction which is simpler but less accurate. The point lighting model is often used for a nearby light source, and the lighting intensity fall-off is considered in [58] for normal estimation. This model is adopted in [59] to build a handheld photometric stereo device.

General lighting (or environment lighting) model allows photometric stereo to work in less constrained lighting conditions. Environment lighting could be considered as a spherical function. Basri *et al.* [4] take spherical harmon-

ics representation and extend the factorization based uncalibrated photometric stereo method [46] to arbitrary lighting. Such a representation is followed by many methods [60–63] that study photometric stereo in general lighting. Environment lighting can also be represented as a sum of many directional lights and calibrated by a mirror sphere [64].

Perspective camera model is a more precise camera model. A perspective camera model will improve photometric stereo at the cost of nonlinear optimization even for Lambertian material [65]. Perspective model is also useful in building a handheld photometric stereo camera [59] and auto-calibrating the lighting conditions [56, 57].

Cameras with non-linear response distort the scene radiance values. While most of the methods apply a separate radiometric calibration to correct this distortion, some methods [5, 17, 63, 66] include automatic radiometric camera calibration. Reflectance monotonicity [14], symmetry [37], manifold embedding [44, 45], and example-based methods [42] are naturally unaffected by the nonlinear camera response. Thus, radiometric calibration is not required for these methods.

Multiple viewpoints provide useful constraints in various aspects. While single-view methods only recover a normal map, multi-view photometric stereo methods reconstruct a complete 3D model [6, 67–69]. The rough shape from multi-view stereo can calibrate directional lighting [6, 8, 68, 70], near point lighting [59], or general lighting [63] for Lambertian surfaces. Additional viewpoints also simplify non-Lambertian photometric stereo. For arbitrary isotropic surfaces, Zhou *et al.* [69] fuse the iso-depth contours reconstructed by [37] from multiple viewpoints to obtain a complete 3D model at high precision.

Object motion provides stronger constraint than multi-view data when illumination condition is fixed, since the angle between lighting and surface normal varies under object motion. The 3D shape can be recovered from object motion with as few as two frames [71]. With more frames, it can calibrate the directional lighting [8, 72] and general lighting [60] for Lambertian surfaces, and directional lighting for general BRDFs [73].

2.4. Number of input images

We further categorize photometric stereo methods according to their number of input images. The basic trend

is the more general assumption and complicated problem require more images as input.

A small number of images (at least three) are required for classic photometric stereo [1]. More input images help to enhance the system robustness, such as four- [23–25, 33] or five-light configurations [31]. Under general lighting, the first and second order spherical harmonics approximations require at least four [61] or nine images [4, 60].

Many photometric stereo algorithms take more images than the theoretical minimum number for better robustness. Typically 10 - 20 images are used for calibrated methods with shadow analysis [26], analytic BRDF models [2, 34, 35] and general lighting [64]. Most uncalibrated methods under directional lighting [3, 17–19, 28, 38, 50, 51, 54, 56] or near point lighting [48, 57] also take 10 - 20 input images.

A medium number of images (about 50 - 100) are required for many outlier rejection based methods [10, 11, 27, 29, 30]. Most methods handling general BRDFs also take about 50 - 100 images [12–16, 20, 37]. Multi-view methods naturally take more images since they need to capture images from different viewpoints. Lambertian multi-view photometric stereo [6, 59] often take 50 - 100 images to cover more than 36 viewpoints.

A large number of images (about 500 - 1000) are required for outdoor photometric stereo [5, 66] and manifold embedding based methods [44, 45], or to handle anisotropic BRDFs [41]. Some multi-view photometric stereo methods also take such a number of images to deal with dynamic shapes [67], intricate geometries [68], or non-Lambertian materials [69].

2.5. Category labels

We propose the following labels to categorize the papers reviewed in Section 2:

- Reflectance model: **L** – Lambert’s model, **R** – Robust Lambert’s model + outliers (usually solved by Robust methods), **A** – Analytic non-Lambertian model (*e.g.*, the Ward, Cook-Torrance model, *etc.*), and **B** – Basic General properties of BRDF (*e.g.*, symmetry, monotonicity, *etc.*).
- Lighting calibration requirement: **C** – Calibrated (known lighting conditions) and **U** – Uncalibrated (unknown lighting conditions).
- Lighting model: **D** – Directional lighting, **P** – Point lighting, and **G** – General (environment) lighting.
- Number of input images: **S** – Small, at least three and typically 10 - 20, **M** – Medium, typically 50 - 100, and **L** – Large, typically 500 - 1000.

Such labels provide an easy way to categorize a photometric stereo method. For example, the classic method [1] is indicated as [LCDS], *i.e.*, Lambert’s model, Calibrated, Directional lighting, with a Small number of input images. Other properties are labeled as: **PC**: Perspective Camera;

NL: Non-Linear camera response; **MV**: Multi-View setup; **OM**: Object Motion. For easy reference, we include such labels for all surveyed papers in the ‘References’ section.

3. Photometric Stereo Dataset

In this section, we introduce the ‘DiLiGenT’ dataset to evaluate photometric stereo algorithms under classic imaging condition (*i.e.*, fixed orthographic camera and directional lighting), with a special focus on non-Lambertian and uncalibrated methods. There are a few publicly available photometric stereo dataset^{2,3} without ground truth shape; the point feature dataset in [74] and multi-view stereo dataset in [75] also contain images under varying illumination and with scanned shapes provided. But none of them are suitable for evaluating non-Lambertian photometric stereo, because most of their objects are simple in reflectance and the number of different lighting conditions is small. This motivates us to create a new dataset with objects of various shapes and BRDFs, carefully calibrated lighting conditions, and ‘ground truth’ 3D shapes (normals).

The ‘DiLiGenT’ dataset contains ten objects, as shown in Figure 1. In terms of surface BRDFs, it covers materials that are mostly diffuse (CAT) or with a rough surface (POT1), with strong and sparse specular spikes (BALL, READING), with broad and soft specular lobes on uniform (BEAR, BUDDHA) and spatially-varying materials (POT2, GOBLET⁴), and with metallic paint on mostly uniform (COW) and spatially-varying (HARVEST) surfaces. In terms of surface shapes, we have the simple sphere (BALL), smoothly curved surfaces (BEAR, CAT, GOBLET, COW), smooth surfaces with local details (POT1, POT2), surfaces with complicated geometry (BUDDHA, READING), and delicate shapes with concave parts (HARVEST).

Data capture is performed using a Point Grey Grasshopper (GRAS-50S5C-C) camera that has a linear radiometric response function with a 50mm lens. All images have a resolution of 2448 × 2048. The objects are about 20cm in diameter and placed about 1.5m from the camera to approximate orthographic projection. To avoid interreflection, all data are captured in a dark room with everything covered by black cloth except the target object. Under each lighting condition, we take four images under different exposure times⁵ and combine them to obtain a single input HDR image. According to Section 2.4, most non-Lambertian methods require about 100 differently illuminated images, so we design our light sources as 96 white LED bulbs fixed on a

²<http://courses.cs.washington.edu/courses/csep576/05wi/projects/project3/project3.htm>

³<http://vision.ucsd.edu/~nalldrin/research/>

⁴This object is made from one material, but the rust on surface makes its BRDF spatially-varying. We exclude the concave interior of this object for evaluation due to its strong interreflection, which is not modeled in almost all existing photometric stereo algorithms.

⁵A typical exposure times setting is {20, 40, 70, 100}ms.

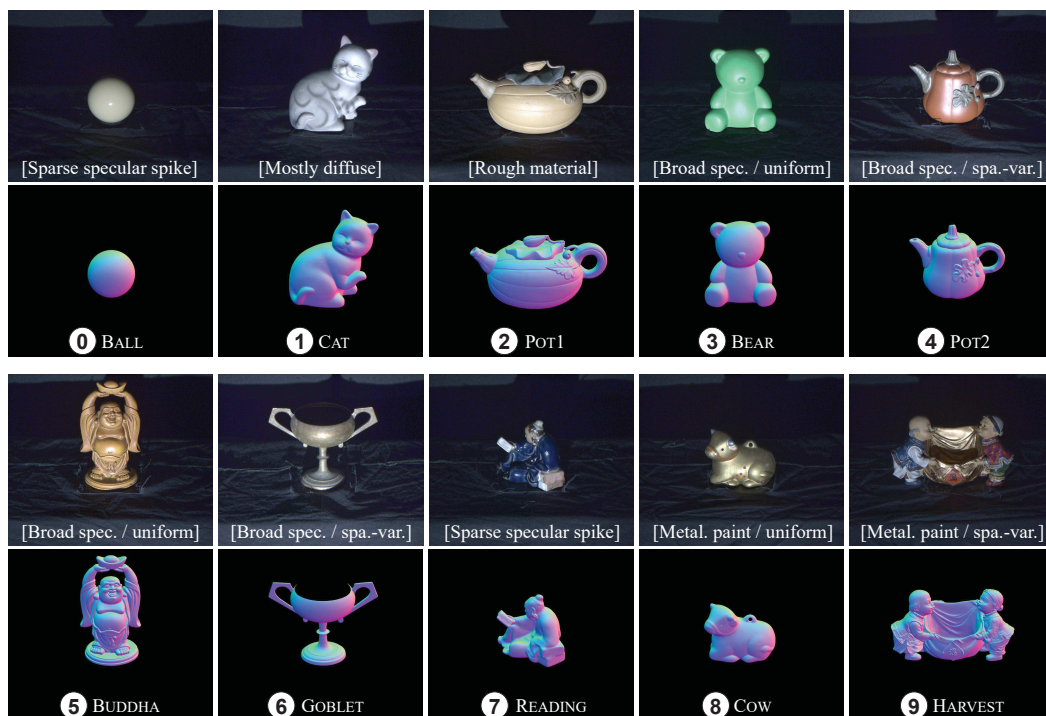


Figure 1. Photometric stereo images (tone-mapped HDR images) and ‘ground truth’ normals for ten objects in the ‘DiLiGenT’ dataset. The texts in brackets indicate the dominant reflectance properties of the objects. Please zoom in the electronic version for better details.

rectangular metal frame. The frame is divided into 12×8 regular samples spanning $74.6^\circ \times 51.4^\circ$, with one LED fixed at each of the grid points. We design the light frame in this way for its portability and efficiency in cost and space. The light frame is fixed with the camera at its center.

Lighting condition calibration includes calibrating light source intensities and directions. The lighting intensity is calibrated using a Macbeth white balance chart, which is a good approximation of a uniform Lambertian surface. We calibrate the lighting directions by treating each LED as a point light source and fitting all positions of the 2D grid of lights together. Specifically, we capture an image of a mirror sphere with all LEDs turned on. The 3D position of the mirror sphere is determined from the camera intrinsics and the radius of sphere. So the reflections of the LEDs in the mirror sphere provide 3D light rays where the LED bulbs should lie on. We measure the physical dimension of the LED grid frame, and fit an optimal rigid transformation to align it to those 3D light rays. Finally, the lighting directions are calculated by subtracting the center of the mirror sphere from 3D positions of the LED bulbs. Note though we assume directional lighting in this paper, our data can also be used for photometric stereo methods with point light assumption by directly using the positions of point light sources.

‘Ground truth’ shapes are scanned with a Rexcan CS+ scanner (precision up to $0.01mm$). Each 3D mesh is

stitched from 162 scans of different rotation and swing poses of the turntable, which are then registered using the eZScan 7 software affiliated with the scanner. The registration error is reported to be lower than $0.01mm$. Manual refinement such as hole filling and smoothing are applied using Autodesk Meshmixer 2.8 only when it is really necessary. The shape-to-image registration is estimated using the mutual information method [76] in Meshlab. We try multiple manual initializations, and select the alignment with the best visual consistency. The alignment is further manually verified by comparing the normal map rendered from the scanned shape and the normal map estimated from photometric stereo to ensure subpixel precision.

Evaluation criteria is based on the statistics of angular error. For each pixel, the angular error is calculated as $\arccos(\mathbf{n}_0^\top \mathbf{n})$ in degrees, where \mathbf{n}_0 and \mathbf{n} are ‘ground truth’ and estimated normals respectively. In addition to the mean angular error, which is a commonly adopted metric in papers surveyed here, we also calculate the minimum, maximum, median, the first quartile, and the third quartile of angular errors for each estimated normal map.

4. Quantitative Evaluation

4.1. Evaluation for non-Lambertian methods

We evaluate and compare non-Lambertian photometric stereo methods summarized in Table 1, *i.e.*, BASELINE

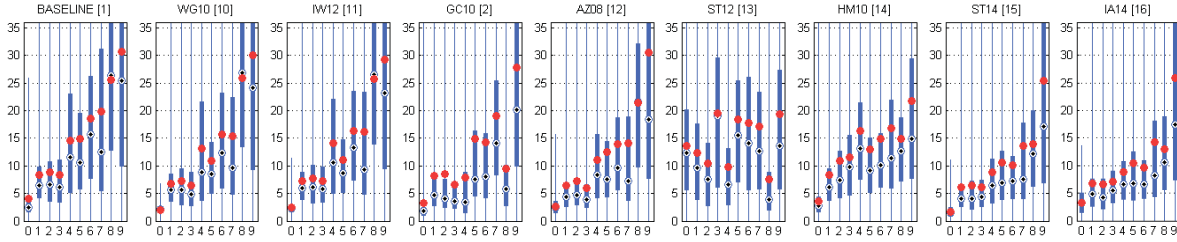


Figure 2. Benchmark results for calibrated non-Lambertian photometric stereo. Each subplot shows the results by one evaluated method using all ten data; the X -axis is the ID of the data according to Figure 1, and the Y -axis is the angular error in degrees. The statistics of angular errors are displayed using the box-and-whisker plot: The red dot indicates the mean value, the black dot is the median, the lower and upper bound of the blue box indicate the first and third quartile values, and the vertical blue lines indicate the minimum and maximum errors.

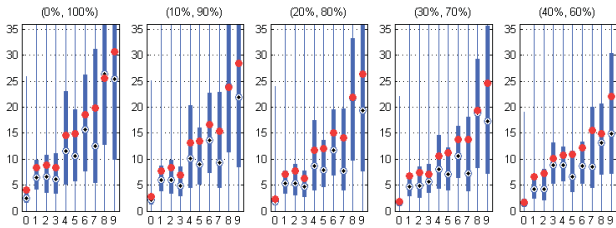


Figure 3. Benchmark results for classic photometric stereo (BASELINE) with position thresholds (T_{low} , T_{high}). Please refer to Figure 2 caption for more description.

[1], WG10 [10], IW12 [11], GC10 [2], AZ08 [12], ST12 [13], HM10 [14], ST14 [15], and IA14 [16]. We choose these methods because they cover most categories of non-Lambertian methods and achieve state-of-the-art performance. For all evaluated methods, we use the parameters provided in the original codes or suggested by the original papers. The evaluation results are summarized in Figure 2. Estimated normal maps and difference maps w.r.t. ‘ground truth’ are in the supplementary material.

The BASELINE method has the largest errors in most data. When sparse outliers could be efficiently removed, WG10 and IW12 show improvement on objects including a dominant Lambertian component (BALL, CAT, POT1). But they cannot handle broad and soft specularly well (POT2, COW). Adopting a mixture of the Ward model, GC10 produces superior results on objects with broad and soft specularity (POT2, COW). Its performance drops on objects with complicated spatially-varying BRDFs (GOBLET). Further, the mixture model fitting might suffer from some local minima as indicated by the relatively large errors in some objects (BUDDHA, READING). The bivariate BRDF constraint from AZ08 consistently achieves moderate accuracy on all data, and performs better when the highlight is sparser or weaker (BALL, BEAR). ST12 suffers from large error when the BRDF is not 1D monotonic (BALL, BEAR), but achieves best results among all evaluated methods on one challenging example (COW), where its assumption is satisfied. HM10 works reasonably well for monotonic reflectance described by a single lobe BRDF (GOB-

LET), but not for the mixture of diffuse and specular components (BEAR, POT2). In general, the two most recent methods ST14 and IA14 perform best for most data. All evaluated methods show large errors in concave areas (HARVEST), due to frequent cast shadow and interreflection.

A position threshold method: Non-Lambertian phenomena such as shadows and specular highlights are the main problems for calibrated photometric stereo. Since shadows and highlights are often associated with dark and bright pixels, we could simply sort pixels by their intensities and discard shadows and highlights as bottom and top ranked observations respectively. The classic photometric stereo (BASELINE) can be safely applied once the non-Lambertian observations are largely discarded. Such a strategy has been used in [2, 15, 16] as initialization. We experiment this simple position threshold method using the ‘DiLi-GenT’ dataset. The position thresholds for shadow and highlight are chosen among $T_{low} = \{0\%, 10\%, \dots, 40\%$ and $T_{high} = \{100\%, 90\%, \dots, 60\%$ respectively. The benchmark results are summarized in Figure 3. It is interesting to note that a narrower interval produces more accurate normal estimates on most of the data. In particular, the tightest thresholding ($T_{low} = 40\%$, $T_{high} = 60\%$) produces comparable results to state-of-the-art methods ST14 [15] and IA14 [16].

Open problems: High-quality photometric stereo for complicated BRDFs (COW), spatially-varying materials (GOBLET), and concave shapes (READING, HARVEST) is still challenging. Methods exploiting general reflectance properties such as symmetry or monotonicity often require far more input images than the BASELINE. Non-Lambertian photometric stereo also involves extensive computation, and all evaluated methods here take much more memory and computation resources than the BASELINE. Computationally efficient non-Lambertian photometric stereo methods using a small number of images are still missing. The unavoidable global illumination effect deteriorates all evaluated methods, and how to integrate global illumination removal (e.g., using structured light pattern [77]) to photometric stereo is an open problem.

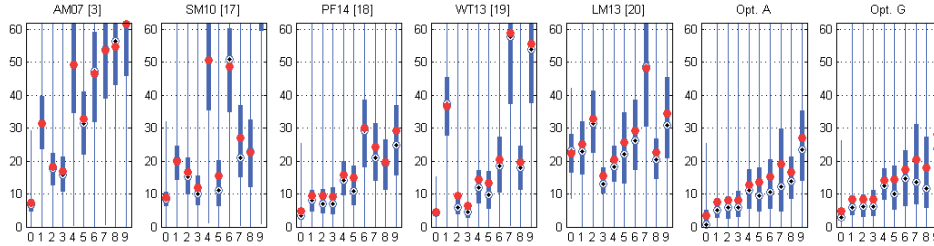


Figure 4. Benchmark results for uncalibrated photometric stereo. Please refer to Figure 2 caption for more description.

4.2. Evaluation for uncalibrated methods

We evaluate and compare recent uncalibrated methods summarized in Table 2, *i.e.*, AM07 [3], SM10 [17], PF14 [18], and WT13 [19], plus a manifold embedding based method LM13 [20]. These methods so far are evaluated by their consistency with calibrated methods. Using the ‘DiLiGenT’ dataset with ‘ground truth’ shapes, we can directly evaluate the accuracy in surface normal directions. To fairly evaluate and compare AM07, SM10, and PF14, we start from the same GBR-distorted normal map that is generated from the code of PF14. Unlike these three methods, WT13 applies diffuse-specular separation [78] beforehand and LM13 is not specific to GBR distortion, so we simply take their complete system for evaluation. We summarize the evaluation results in Figure 4.

The uncalibrated methods produce much larger errors than the calibrated ones especially for non-Lambertian materials. AM07 and SM10 only produce reasonable results on near Lambertian objects (BALL, POT1, BEAR). AM07 assumes piecewise constant albedo, and is sensitive to objects with smoothly varying texture (CAT). SM10 requires albedo based segmentation and hence fails on objects with high-frequency material changes (GOBLET). Though PF14 is designed for Lambertian surfaces, it shows excellent robustness and tolerance to non-Lambertian objects (POT2, BUDDHA, COW). Its average error for all test data is the smallest. The specular reflection based method WT13 outperforms other methods on materials with broad and soft highlights (BEAR, POT2, BUDDHA, COW). Its performance deteriorates when the specular reflection is weak (CAT) or spatially-varying (GOBLET, HARVEST) and it suffers from catastrophic failure when diffuse-specular separation is poor (READING). The performance of LM13 is poor in our evaluation, partially due to its requirement on uniformly distributed directional lighting.

We further evaluate the linearly- and GBR-distorted normals, which are inputs to auto-calibration. We solve the optimal linear transformation \mathbf{A} and GBR transformation \mathbf{G} by registering the distorted normals to the ‘ground truth’ normals. The mean angular errors between the registered and ‘ground truth’ normals are reported in Figure 4 as Opt. A and Opt. G for linearly- and GBR-distorted normals respectively. Both results show surprisingly large residual er-

rors, which suggests a more robust method is needed for recovering the pseudo-normal $\hat{\mathbf{S}}$ (up to \mathbf{A}) and $\hat{\mathbf{S}}$ (up to \mathbf{G}).

Open problems: Uncalibrated photometric stereo is still a very difficult problem for real-world objects and a major bottleneck is from the widely adopted factorization step [46] which is designed for Lambertian materials. Pre-processing of the data using outlier rejection methods [10, 11] can only solve the problem in a limited scale, since they need to assume the existence of a dominant Lambertian component. LM13 presents a successful trial to avoid this factorization for general BRDFs, but it is largely limited by the requirement of uniformly distributed lighting. Exploring general BRDF properties and how to accurately estimate the pseudo-normal for non-Lambertian objects could be interesting future topics. In terms of resolving the GBR ambiguity, it might be interesting to combine the strength of diffuse reflection based method (*e.g.*, PF14) and specular reflection based method (*e.g.*, WT13).

5. Conclusion

We survey recent photometric stereo methods according to our taxonomy focusing on non-Lambertian and uncalibrated photometric stereo algorithms and present photometric stereo images dataset with carefully calibrated lighting conditions as well as scanned and registered ‘ground truth’ shapes (normals). Eight recent non-Lambertian photometric stereo methods and five uncalibrated methods are quantitatively evaluated and compared using our data, which reveals interesting future research directions.

Acknowledgement

This paper is based on results obtained from a project commissioned by the New Energy and Industrial Technology Development Organization (NEDO). S.-K. Yeung is supported by SUTD-MIT International Design Center Grant IDG31300106, Singapore MOE Academic Research Fund MOE2013-T2-1-159. We also acknowledge the support of the SUTD Digital Manufacturing and Design (DManD) Centre which is supported by the Singapore National Research Foundation (NRF). P. Tan is supported by the NSERC Discovery Grant 31-611664 and the NSERC Discovery Accelerator Supplement 31-611663.

References

* The capital letters in brackets are photometric stereo category labels defined in Section 2.5.

- [1] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering* 19(1):139–144, 1980, [LCDS]. 1, 2, 5, 7
- [2] D. B. Goldman, B. Curless, A. Hertzmann, and S. M. Seitz. Shape and spatially-varying BRDFs from photometric stereo. *IEEE TPAMI* 32(6):1060–1071, 2010, [ACDS]. 1, 2, 5, 7
- [3] N. G. Alldrin, S. P. Mallick, and D. J. Kriegman. Resolving the generalized bas-relief ambiguity by entropy minimization. In *Proc. CVPR*, 2007, [LUDS]. 1, 4, 5, 8
- [4] R. Basri, D. Jacobs, and I. Kemelmacher. Photometric stereo with general, unknown lighting. *IJCV* 72(3):239–257, 2007, [LUGS]. 1, 4, 5
- [5] J. Ackermann, F. Langguth, S. Fuhrmann, and M. Goesele. Photometric stereo for outdoor webcams. In *Proc. CVPR*, 2012, [ACDM-NL]. 1, 2, 4, 5
- [6] C. Hernández, G. Vogiatzis, and R. Cipolla. Multiview photometric stereo. *IEEE TPAMI* 30(3):548–554, 2008, [LUDM-MV]. 1, 4, 5
- [7] P. Tan, S. Lin, and L. Quan. Subpixel photometric stereo. *IEEE TPAMI* 30(8):1460–1471, 2008, [ACDM]. 1
- [8] L. Zhang, B. Curless, A. Hertzmann, and S. M. Seitz. Shape and motion under varying illumination: Unifying structure from motion, photometric stereo, and multiview stereo. In *Proc. ICCV*, 2003, [LUDL-MV-OM]. 1, 4
- [9] S. M. Seitz, B. Curless, J. Diebel, D. Scharstein, and R. Szeliski. A comparison and evaluation of multi-view stereo reconstruction algorithms. In *Proc. CVPR*, 2006. 1
- [10] L. Wu, A. Ganesh, B. Shi, Y. Matsushita, Y. Wang, and Y. Ma. Robust photometric stereo via low-rank matrix completion and recovery. In *Proc. ACCV*, 2010, [RCDM]. 1, 2, 5, 7, 8
- [11] S. Ikehata, D. Wipf, Y. Matsushita, and K. Aizawa. Robust photometric stereo using sparse regression. In *Proc. CVPR*, 2012, [RCDM]. 2, 5, 7, 8
- [12] N. G. Alldrin, T. Zickler, and D. J. Kriegman. Photometric stereo with non-parametric and spatially-varying reflectance. In *Proc. CVPR*, 2008, [BCDM]. 3, 5, 7
- [13] B. Shi, P. Tan, Y. Matsushita, and K. Ikeuchi. Elevation angle from reflectance monotonicity: Photometric stereo for general isotropic reflectances. In *Proc. ECCV*, 2012, [BCDM]. 3, 7
- [14] T. Higo, Y. Matsushita, and K. Ikeuchi. Consensus photometric stereo. In *Proc. CVPR*, 2010, [BCDM-NL]. 3, 4, 7
- [15] B. Shi, P. Tan, Y. Matsushita, and K. Ikeuchi. Bi-polynomial modeling of low-frequency reflectances. *IEEE TPAMI* 36(6):1078–1091, 2014, [BCDM]. 3, 7
- [16] S. Ikehata and K. Aizawa. Photometric stereo using constrained bivariate regression for general isotropic surfaces. In *Proc. CVPR*, 2014, [BCDM]. 1, 3, 5, 7
- [17] B. Shi, Y. Matsushita, Y. Wei, C. Xu, and P. Tan. Self-calibrating photometric stereo. In *Proc. CVPR*, 2010, [LUDS-NL]. 1, 4, 5, 8
- [18] T. Papadhimetri and P. Favaro. A closed-form, consistent and robust solution to uncalibrated photometric stereo via local diffuse reflectance maxima. *IJCV* 107(2):139–154, 2014, [LUDS]. 4, 8
- [19] Z. Wu and P. Tan. Calibrating photometric stereo by holistic reflectance symmetry analysis. In *Proc. CVPR*, 2013, [BUDS]. 3, 4, 5, 8
- [20] F. Lu, Y. Matsushita, I. Sato, T. Okabe, and Y. Sato. Uncalibrated photometric stereo for unknown isotropic reflectances. In *Proc. CVPR*, 2013, [BUDM]. 1, 3, 4, 5, 8
- [21] S. Herbot and C. Wöhler. An introduction to image-based 3D surface reconstruction and a survey of photometric stereo methods. *3D Research* 2(3):1–17, 2011. 1
- [22] J. Ackermann and M. Goesele. A survey of photometric stereo techniques. *FnT CGV* 9(3-4):149–254, 2015. 1
- [23] E. N. Coleman and R. Jain. Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry. *CGIP* 18(4):309–328, 1982, [RCDS]. 2, 5
- [24] F. Solomon and K. Ikeuchi. Extracting the shape and roughness of specular lobe objects using four light photometric stereo. *IEEE TPAMI* 18(4):449–454, 1996, [RCDS].
- [25] S. Barsky and M. Petrou. The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows. *IEEE TPAMI* 25(10):1239–1252, 2003, [RCDS]. 2, 5
- [26] M. Chandraker, S. Agarwal, and D. J. Kriegman. Shadow-Cuts: Photometric stereo with shadows. In *Proc. CVPR*, 2007, [RCDS]. 2, 5
- [27] Y. Mukaigawa, Y. Ishii, and T. Shakunaga. Analysis of photometric factors based on photometric linearization. *JOSA A* 24(10):3326–3334, 2007, [RCDM]. 2, 5
- [28] K. Sunkavalli, T. Zickler, and H. Pfister. Visibility subspaces: Uncalibrated photometric stereo with shadows. In *Proc. ECCV*, 2010, [RUDS]. 2, 5
- [29] F. Verbiest and L. Van Gool. Photometric stereo with coherent outlier handling and confidence estimation. In *Proc. CVPR*, 2008, [RCDM]. 2, 5
- [30] C. Yu, Y. Seo, and S. W. Lee. Photometric stereo from maximum feasible lambertian reflections. In *Proc. ECCV*, 2010, [RCDM]. 2, 5
- [31] D. Miyazaki, K. Hara, and K. Ikeuchi. Median photometric stereo as applied to the segonko tumulus and museum objects. *IJCV* 86(2):229–242, 2010, [RCDS]. 2, 5
- [32] T.-P. Wu and C.-K. Tang. Photometric stereo via expectation maximization. *IEEE TPAMI* 32(3):546–560, 2010, [RCDL]. 2
- [33] T. Zickler, S. P. Mallick, D. J. Kriegman, and P. N. Belhumeur. Color subspaces as photometric invariants. *IJCV* 79(1):13–30, 2008, [RCDS]. 2, 5
- [34] A. S. Georghiadis. Incorporating the Torrance and Sparrow model of reflectance in uncalibrated photometric stereo. In *Proc. ICCV*, 2003, [AUDS]. 2, 3, 5
- [35] H. S. Chung and J. Jia. Efficient photometric stereo on glossy surfaces with wide specular lobes. In *Proc. CVPR*, 2008, [ACDS]. 2, 5
- [36] S.-K. Yeung, T.-P. Wu, C.-K. Tang, T. F. Chan, and S. Osher. Normal Estimation of a Transparent Object Using a Video. *IEEE TPAMI* 37(4):890–897, 2015, [ACDL]. 2
- [37] N. G. Alldrin and D. J. Kriegman. Toward reconstructing surfaces with arbitrary isotropic reflectance: A stratified photometric stereo approach. In *Proc. ICCV*, 2007, [BCDM-NL]. 3, 4, 5

- [38] P. Tan, L. Quan, and T. Zickler. The geometry of reflectance symmetries. *IEEE TPAMI* 33(12):2506–2520, 2011, [BUDS]. 3, 5
- [39] M. Chandraker, J. Bai, and R. Ramamoorthi. On differential photometric reconstruction for unknown, isotropic BRDFs. *IEEE TPAMI* 35(12):2941–2955, 2013, [BUDS]. 3
- [40] F. Romeiro, Y. Vasilyev, and T. Zickler. Passive reflectometry. In *Proc. ECCV*, 2008. 3
- [41] M. Holroyd, J. Lawrence, G. Humphreys, and T. Zickler. A photometric approach for estimating normals and tangents. *ACM TOG (Proc. SIGGRAPH Asia)* 27(5):133, 2008, [BCDL]. 3, 5
- [42] A. Hertzmann and S. M. Seitz. Example-based photometric stereo: Shape reconstruction with general, varying BRDFs. *IEEE TPAMI* 27(8):1254–1264, 2005, [BUDS-NL]. 3, 4
- [43] M. K. Johnson, F. Cole, A. Raj, and E. H. Adelson. Microgeometry capture using an elastomeric sensor. *ACM TOG (Proc. SIGGRAPH)* 30(4):46:1–46:8, July 2011, [BCPS]. 3
- [44] I. Sato, T. Okabe, Q. Yu, and Y. Sato. Shape reconstruction based on similarity in radiance changes under varying illumination. In *Proc. ICCV*, 2007, [AUDL-NL]. 3, 4, 5
- [45] T. Okabe, I. Sato, and Y. Sato. Attached shadow coding: Estimating surface normals from shadows under unknown reflectance and lighting conditions. In *Proc. ICCV*, 2009, [BUDL-NL]. 3, 4, 5
- [46] H. Hayakawa. Photometric stereo under a light source with arbitrary motion. *JOSA A* 11(11):3079, 1994, [LUDS]. 3, 4, 8
- [47] A. L. Yuille and D. Snow. Shape and albedo from multiple images using integrability. In *Proc. CVPR*, 1997, [LUDS]. 3
- [48] D. J. Kriegman and P. N. Belhumeur. What shadows reveal about object structure. *JOSA A* 18(8), 2001, [LUPS]. 3, 5
- [49] P. N. Belhumeur, D. J. Kriegman, and A. L. Yuille. The bas-relief ambiguity. *IJCV* 35(1):33–44, 1999, [LUDS]. 3
- [50] O. Drbohlav and R. Sara. Specularities reduce ambiguity of uncalibrated photometric stereo. In *Proc. ECCV*, 2002, [AUDS]. 3, 5
- [51] O. Drbohlav and M. Chantier. Can two specular pixels calibrate photometric stereo? In *Proc. ICCV*, 2005, [AUDS]. 3, 5
- [52] P. Tan, S. Mallick, L. Quan, D. Kriegman, and T. Zickler. Isotropy, reciprocity and the generalized bas-relief ambiguity. In *Proc. CVPR*, 2007, [BUDM]. 3
- [53] P. Tan and T. Zickler. A projective framework for radiometric image analysis. In *Proc. CVPR*, 2009, [BUDM]. 3
- [54] M. Chandraker, F. Kahl, and D. J. Kriegman. Reflections on the generalized bas-relief ambiguity. In *Proc. CVPR*, 2005, [LUDS]. 4, 5
- [55] Z. Zhou and P. Tan. Ring-light photometric stereo. In *Proc. ECCV*, 2010, [LUDS]. 4
- [56] T. Papadhimetri and P. Favaro. A new perspective on uncalibrated photometric stereo. In *Proc. CVPR*, 2013, [LUDS-PC]. 4, 5
- [57] T. Papadhimetri and P. Favaro. Uncalibrated near-light photometric stereo. In *Proc. BMVC*, 2014, [LUPS-PC]. 4, 5
- [58] Y. Iwahori, H. Sugie, and N. Ishii. Reconstructing shape from shading images under point light source illumination. In *Proc. ICPR*, 1990, [LCPS]. 4
- [59] T. Higo, Y. Matsushita, N. Joshi, and K. Ikeuchi. A hand-held photometric stereo camera for 3-D modeling. In *Proc. ICCV*, 2009, [LUPM-PC-MV]. 4, 5
- [60] D. Simakov, D. Frolova, and R. Basri. Dense shape reconstruction of a moving object under arbitrary, unknown lighting. In *Proc. ICCV*, 2003, [LUGS-OM]. 4, 5
- [61] C. Chen and C. Chen. The 4-source photometric stereo under general unknown lighting. In *Proc. ECCV*, 2006, [LUGS]. 5
- [62] L. Shen and P. Tan. Photometric stereo and weather estimation using internet images. In *Proc. CVPR*, 2009, [LUGM].
- [63] B. Shi, K. Inose, Y. Matsushita, P. Tan, S.-K. Yeung, and K. Ikeuchi. Photometric stereo using Internet images. In *Proc. 3DV*, 2014, [LUGM-NL]. 4
- [64] L.-F. Yu, S.-K. Yeung, Y.-W. Tai, D. Terzopoulos, and T. F. Chan. Outdoor photometric stereo. In *Proc. ICCP*, 2013, [LCGS]. 4, 5
- [65] A. Tankus and N. Kiryati. Photometric stereo under perspective projection. In *Proc. ICCV*, 2005, [LCDS-PC]. 4
- [66] A. Abrams, C. Hawley, and R. Pless. Heliometric stereo: Shape from sun position. In *Proc. ECCV*, 2012, [LCDL-NL]. 4, 5
- [67] D. Vlasic, P. Peers, I. Baran, P. Debevec, J. Popović, S. Rusinkiewicz, and W. Matusik. Dynamic shape capture using multi-view photometric stereo. *ACM TOG (Proc. SIGGRAPH Asia)* 28(5):174, 2009, [LCDL-MV]. 4, 5
- [68] J. Park, S. N. Sinha, Y. Matsushita, Y. W. Tai, and I. S. Kweon. Multiview photometric stereo using planar mesh parameterization. In *Proc. ICCV*, 2013, [LUDL-MV]. 4, 5
- [69] Z. Zhou, Z. Wu, and P. Tan. Multi-view photometric stereo with spatially varying isotropic materials. In *Proc. CVPR*, 2013, [BCDL-MV]. 4, 5
- [70] N. Joshi and D. J. Kriegman. Shape from varying illumination and viewpoint. In *Proc. ICCV*, 2007, [LUDS-MV]. 4
- [71] R. Basri and D. Frolova. A two-frame theory of motion, lighting and shape. In *Proc. CVPR*, 2008, [LCDS-OM]. 4
- [72] J. Lim, J. Ho, M. H. Yang, and D. J. Kriegman. Passive photometric stereo from motion. In *Proc. ICCV*, 2005, [LUDL-OM]. 4
- [73] M. Chandraker, D. Reddy, Y. Wang, and R. Ramamoorthi. What object motion reveals about shape with unknown BRDF and lighting. In *Proc. CVPR*, 2013, [BUDS-PC-OM]. 4
- [74] H. Aanæs, A. L. Dahl, and K. S. Pedersen. Interesting interest points: A comparative study of interest point performance on a unique data set. *IJCV* 97:18–35, 2012. 5
- [75] R. Jensen, A. Dahl, G. Vogiatzis, E. Tola, and H. Aanæs. Large scale multi-view stereopsis evaluation. In *Proc. CVPR*, 2014. 5
- [76] M. Corsini, M. Dellepiane, F. Ponchio, and R. Scopigno. Image-to-geometry registration: A mutual information method exploiting illumination-related geometric properties. *CGF* 28(7):1755–1764, 2009. 6
- [77] S. K. Nayar, G. Krishnan, M. D. Grossberg, and R. Raskar. Fast separation of direct and global components of a scene using high frequency illumination. *ACM TOG (Proc. SIGGRAPH)* 25(3):935–944, 2006. 7
- [78] Y. Sato and K. Ikeuchi. Temporal-color space analysis of reflection. *JOSA A* 11(11):2990–3002, 1994. 8