Uncalibrated Neural Inverse Rendering for Photometric Stereo of General Surfaces

Conference Paper**Author(s):**

Kaya, Berk; Kumar, Suryansh; [Porto de Oliveira, Carlos Eduardo](#) ; Ferrari, Vittorio; Van Gool, Luc

Publication date:

2021

Permanent link:

<https://doi.org/10.3929/ethz-b-000520037>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

Originally published in:

<https://doi.org/10.1109/CVPR46437.2021.00380>

Uncalibrated Neural Inverse Rendering for Photometric Stereo of General Surfaces

Berk Kaya¹ Suryansh Kumar¹ Carlos Oliveira¹ Vittorio Ferrari² Luc Van Gool^{1,3}
Computer Vision Lab, ETH Zürich¹, Google Research², KU Leuven³

Abstract

This paper presents an uncalibrated deep neural network framework for the photometric stereo problem. For training models to solve the problem, existing neural network-based methods either require exact light directions or ground-truth surface normals of the object or both. However, in practice, it is challenging to procure both of this information precisely, which restricts the broader adoption of photometric stereo algorithms for vision application. To bypass this difficulty, we propose an uncalibrated neural inverse rendering approach to this problem. Our method first estimates the light directions from the input images and then optimizes an image reconstruction loss to calculate the surface normals, bidirectional reflectance distribution function value, and depth. Additionally, our formulation explicitly models the concave and convex parts of a complex surface to consider the effects of interreflections in the image formation process. Extensive evaluation of the proposed method on the challenging subjects generally shows comparable or better results than the supervised and classical approaches.

1. Introduction

Since Woodham’s seminal work [69], the photometric stereo problem has become a popular choice to estimate an object’s surface normals from its light varying images. The formulation proposed in that paper assumes the Lambertian reflectance model of the object, and therefore, it does not apply to general objects with unknown reflectance property. While multiple-view geometry methods exist to achieve a similar goal [57, 20, 70, 76, 35, 24, 36, 37], photometric stereo is excellent at recovering fine details on the surface, like indentations, imprints, and even scratches. Of course, the solution proposed in Woodham’s paper has some unrealistic assumptions. Still, it is central to the development of several robust algorithms [71, 30, 55, 1, 22, 26] and also lies at the core of the current state-of-the-art deep photometric stereo methods [28, 65, 12, 10, 11, 42, 41, 27].

Generally, deep learning-based photometric stereo methods assume a calibrated setting, where all the light source

information is given both at the train and test time [28, 56, 12, 65]. Such methods attempt to learn an explicit relation between the reflectance map and the ground-truth surface normals. But, the exact estimation of light directions is a tedious process and requires expert skill for calibration. Motivated by that, Chen *et al.* [10, 11] recently proposed an uncalibrated photometric stereo method. Though it estimates light directions using image data, the proposed method requires ground-truth surface normals for training the neural network. Certainly, procuring ground-truth 3D surface geometry is difficult, if not impossible, which makes the acquisition task of correct surface normals strenuous. For 3D data acquisition, active sensors are mostly used, which is expensive and often needs post-processing of the data to remove noise and outliers. Hence, the necessity of ground-truth surface normals limits the usage of such an approach.

Further, most photometric stereo methods, including current deep-learning methods, assume that each surface point is illuminated only by the light source, which generally holds for a convex surface [49]. However, objects, mainly from ancient architectures, have complex geometric structures, where the shape may compose of convex, concave, and other fine geometric primitives (see Fig.1(a)). When illuminated under a varying light source, certain concave parts of the surface might reflect light onto other parts of the object, depending on its position. Surprisingly, this phenomenon of interreflections is often ignored in the modeling and formulation of a photometric stereo problem, despite its vital role in the object’s imaging [28, 65, 12, 10, 11].

In this work, we overcome the above shortcomings by proposing an uncalibrated neural inverse rendering network. We first estimate all the light source directions and intensities using image data. Computed light source information is then fed into the proposed neural inverse rendering network to estimate the surface normals. The idea is, those correct surface normals, when provided to the rendering equation, should reconstruct the input image as close as possible. Consequently, we can bypass the requirement of the ground-truth surface normals at train time. Unlike recent methods, we model the effects of both the light source and the interreflections for rendering the image. Although one

Figure 1: (a) Example showing the interreflection effect due to concave geometric structure. The light from the primary source hits the concave region of the surface that illuminates the other surface points which then act as a secondary light source. (b) Comparison of our approach against the classical and deep-learning methods on the Vase dataset which shows that it performs better than others. We used Mean Angular Error (MAE) metric to report the results.

can handle interreflection using classical methods [49, 9], the reflectance characteristics of different types of material are quite diverse. Hence, we want to leverage neural network’s powerful capability to learn complex reflectance behavior from the input image data.

For evaluation, we performed experiments on DiLiGenT dataset [62]. We noticed that the objects present with this dataset are not apt for studying interreflections. To that end, we proposed a novel dataset to study the behavior and effect of interreflections on the object’s imaging §5. We observed that ignoring interreflections can dramatically affect the accuracy of the surface normals estimate (see Fig 1(b)). To sum up, our paper makes the following contributions:

- This paper presents an uncalibrated deep photometric stereo method that does not require ground-truth surface normals at train time to solve photometric stereo.
- Our work considers the contribution of both the source light and interreflections in the image formation process. Consequently, our approach is more general and applicable to a wide range of objects.
- The proposed method leverages neural inverse rendering principles to infer the surface normals, depth, and spatially varying bidirectional reflectance distribution function (BRDF) values from input images. Our method generally provides comparable or better results than the classical [49, 2, 60, 73, 45, 52, 44] and the recent supervised uncalibrated deep learning methods [12, 15, 11].

2. Related Work

For comprehensive review on photometric stereo readers may refer to Herbot et al. [25], and Chen et al. [11] work.

1. Calibrated Photometric Stereo. The methods proposed under this setting assume that all the light source information is known for computing surface normals. Sev-

eral calibrated methods have been proposed to handle non-Lambertian surfaces [48, 72, 71, 47, 50, 31]. These methods assume non-Lambertian effects, such as specularities, are sparse and confined to a local region of the surface. So, they filter them before computing surface normals. For example, Wu et al. [71] proposed a rank minimization approach to robustify photometric stereo. Oh et al. [50] introduced a partial sum of singular values optimization algorithm for the low-rank normal matrix recovery. Other popular outlier rejection methods were based on RANSAC [48], Bayesian regression [30, 31], and expectation-maximization [72].

With the recent success of deep learning in many computer vision areas, several learning-based approaches have also emerged for the photometric stereo problem. Santo et al. [56] introduced a deep photometric stereo network (DPSN) that learns the mapping between the surface normals and the reflectance map. Ikehata [28] merged all pixel-wise information to an observation map and trained a network to perform per-pixel estimation of normals. In contrast, Taniai et al. [65] used a self-supervised framework to recover surface normals from input images. Yet, it uses the classical photometric equation that fails to model interreflections. Moreover, it uses Woodham’s method [69] to initialize the surface normals in their loss function which is not robust, and therefore, their trained network model is susceptible to noise and outliers.

2. Uncalibrated Photometric Stereo. These methods assume unknown light source information for solving photometric stereo. However, not knowing the light sources leads to an ambiguity i.e., there exists a set of surfaces under unknown distant light sources that can lead to identical images. Hence, the actual surface can be recovered up to a three-parameter ambiguity popularly known as Generalized Bas-Relief (GBR) ambiguity [5, 9]. Existing methods elim-

inate this ambiguity by making some additional assumptions in their proposed solution. Alldrin *et al.* [2] assumes bounded values on the GBR variables and resolves the ambiguity by minimizing the entropy of albedo distribution. Shi *et al.* [60] assumes at least four pixels with different normals but the same albedo. Papadimitri *et al.* [52] presents a closed-form solution by detecting local diffuse reflectance maxima (LDR). Other methods assume perspective projection [51], specularities [21, 18], low-rank [59], interreflections [9] or symmetry properties of BRDFs [64, 73, 44].

Apart from the traditional methods, Chen *et al.* [12] proposed a learning framework (UPS-FCN). This method bypasses the light estimation process and learns a direct mapping between the image and the surface normal. But, the knowledge of the light source would provide useful evidence about the surface normals, and therefore completely ignoring the light source data seems implausible. The self-calibrating deep photometric stereo networks work [10] recently introduced an initial lighting estimation stage (LC-Net) from images to overcome the problem with UPS-FCN. Recently, Chen *et al.* [13] also proposed a guided calibration network (GCNet) to overcome the limitations of LC-Net. Unlike existing uncalibrated deep-learning methods that rely heavily on ground-truth surface normals for training, our method can solve photometric stereo by using an image reconstruction term as a function of estimated surface normals. The goal is to let the network learn the image formation process and the complex reflectance model of the object via explicit interreflection modeling.

3. Photometric Stereo

Photometric stereo aims to recover the surface normals of an object from its multiple images captured under varying light illuminations. It assumes a unique point light source per image taken by a camera from a constant view direction $\mathbf{v}$ which is commonly assumed to be at $(0, 0, 1)^T$. Under such configuration, when a surface point $\mathbf{x}$ is illuminated by a distant point light source from direction ' $\mathbf{l}_s \in \mathbb{R}^{3 \times 1}$ ', the image intensity $X_s(\mathbf{x})$ measured by the camera due to s^{th} source in the view direction $\mathbf{v}$ is given by

$$X_s(\mathbf{x}) = e_s \cdot \rho(\mathbf{n}(\mathbf{x}), \mathbf{l}_s, \mathbf{v}) \cdot \zeta_a(\mathbf{n}(\mathbf{x}), \mathbf{l}_s) \cdot \zeta_c(\mathbf{x}) \quad (1)$$

Here, the camera projection model is assumed to be orthographic. The function $\rho(\mathbf{n}(\mathbf{x}), \mathbf{l}_s, \mathbf{v})$ gives the BRDF value, $\zeta_a(\mathbf{n}(\mathbf{x}), \mathbf{l}_s) = \max(\mathbf{n}(\mathbf{x})^T \mathbf{l}_s, 0)$ accounts for the attached shadow, and $\zeta_c(\mathbf{x}) \in \{0, 1\}$ assign 0 or 1 value to $\mathbf{x}$ depending on whether it lies in the cast shadow region or not. $e_s \in \mathbb{R}_+$ is a scalar for light intensity value, and $\mathbf{n}(\mathbf{x}) \in \mathbb{R}^{3 \times 1}$ is the surface normal vector at point $\mathbf{x}$. Eq:(1) is most-widely used photometric stereo formulation which generally works well in practise [9, 30, 28, 65, 13, 11].

1. Classical Photometric Stereo Model. It assumes a convex Lambertian surface model resulting in a constant

BRDF value across the whole surface. Additionally, the surface is considered to be illuminated only due to the light source. Under such assumptions, Eq:(1) becomes a linearly tractable problem and it is possible to recover the surface normals by solving a simple system of linear equations. Let all the n light source directions be denoted as $\mathbf{L} = [\mathbf{l}_1, \mathbf{l}_2, \dots, \mathbf{l}_n] \in \mathbb{R}^{3 \times n}$ and m unknown surface point normal be $\mathbf{N} = [\mathbf{n}(\mathbf{x}_1), \mathbf{n}(\mathbf{x}_2), \dots, \mathbf{n}(\mathbf{x}_m)] \in \mathbb{R}^{3 \times m}$. Using the notation, we can write Eq:(1) due to all the light sources and surface points compactly as

$$\mathbf{X}_s = \rho \mathbf{N}^T \mathbf{L} \quad (2)$$

where, $\mathbf{X}_s \in \mathbb{R}^{m \times n}$ is the matrix consisting of n images with m object pixels stacked as column vectors, and ρ is the constant albedo. The above system can be solved for the surface normals using the matrix pseudo-inverse approach under calibrated setting if $n \geq 3$ (i.e., at least three light sources are given in non-degenerate configuration).

2. Interreflection Model. In contrast to the classical photometric stereo, here, the total radiance at a point $\mathbf{x}$ on the surface is the sum of radiance due to light source s and the radiance due to interreflection from other surface points.

$$X(\mathbf{x}) = \underbrace{X_s(\mathbf{x})}_{\text{due to light source}} + \underbrace{\frac{\rho(\mathbf{x})}{\pi} \int_{\Omega} K(\mathbf{x}, \mathbf{x}') X(\mathbf{x}') d\mathbf{x}'}_{\text{due to interreflections}} \quad (3)$$

where, Ω represents the surface, $\mathbf{x}'$ is another surface point, and $d\mathbf{x}'$ the differential surface element at $\mathbf{x}'$. The value of the interreflection kernel ' K ' at $\mathbf{x}$ due to $\mathbf{x}'$ is defined as:

$$K(\mathbf{x}, \mathbf{x}') = \left(\frac{(\mathbf{n}(\mathbf{x})^T (-\mathbf{r})) \cdot (\mathbf{n}(\mathbf{x}')^T \mathbf{r}) \cdot V(\mathbf{x}, \mathbf{x}')}{(\mathbf{r}^T \mathbf{r})^2} \right) \quad (4)$$

The values of K , when measured for each surface element form a symmetric and positive semi-definite matrix. In Eq:(4), $V(\mathbf{x}, \mathbf{x}')$ captures the visibility. When $\mathbf{x}$ occludes $\mathbf{x}'$ or vice-versa then V is 0. Otherwise, V gives the orientation between the two points using the following expression:

$$V(\mathbf{x}, \mathbf{x}') = \left(\frac{\mathbf{n}(\mathbf{x})^T (-\mathbf{r}) + |\mathbf{n}(\mathbf{x})^T (-\mathbf{r})|}{2|\mathbf{n}(\mathbf{x})^T (-\mathbf{r})|} \right) \cdot \left(\frac{\mathbf{n}(\mathbf{x}')^T \mathbf{r} + |\mathbf{n}(\mathbf{x}')^T \mathbf{r}|}{2|\mathbf{n}(\mathbf{x}')^T \mathbf{r}|} \right) \quad (5)$$

where, $\mathbf{n}(\mathbf{x})$ and $\mathbf{n}(\mathbf{x}')$ are the surface normal at $\mathbf{x}$ and $\mathbf{x}'$, and $\mathbf{r} = \mathbf{x} - \mathbf{x}'$ is the vector from $\mathbf{x}'$ to $\mathbf{x}$. Substituting V and K in Eq:(3) gives an infinite sum over every infinitesimally small surface element (point) and therefore, it is not computationally easy to find a solution to $X(\mathbf{x})$ in its continuous form. Nevertheless, the solution to Eq:(3) is guaranteed to converge as $\rho(\mathbf{x}) < 1$ for a real surface. To practically implement the interreflection model, the object surface is

discretized into m facets [49]. Assuming the radiance and albedo values to be constant within each facet, then Eq:(3) for the i^{th} facet becomes $X_i = X_{si} + \frac{\rho_i}{\pi} \sum_{j=1, j \neq i}^m X_j K_{ij}$, where $X_i \in \mathbb{R}^{n \times 1}$ and ρ_i are the radiance and albedo of facet i . Considering the contribution of all the light sources for each facet, it can be compactly re-written as:

$$\mathbf{X} = \mathbf{X}_s + \mathbf{P}\mathbf{K}\mathbf{X}, \Rightarrow \mathbf{X} = (\mathbf{I} - \mathbf{P}\mathbf{K})^{-1}\mathbf{X}_s \quad (6)$$

where, $\mathbf{X} = [X_1, X_2, \dots, X_m]^T$ is the total radiance for all the facets, and $\mathbf{X}_s = [X_{s1}, X_{s2}, \dots, X_{sm}]^T$ is the light source contribution to the radiance of m facets. Furthermore, $\mathbf{P}$ is a diagonal matrix composed of albedo values and $\mathbf{K}$ is a $m \times m$ interreflection kernel matrix with $\text{diag}(\mathbf{K}) = 0$. Nayar *et al.* [49] proposed Eq:(20) to recover the surface normals for concave objects. The algorithm proposed to estimate surface normals using Eq:(20) first computes the pseudo surface normals by treating the object as directly illuminated by light sources. These pseudo surface normals are then used to iteratively update for the interreflection kernel and surface normals via depth map estimation step, until convergence. In the later part of the paper, we denote the normals estimated using Eq:(20) as $\mathbf{N}_{ny}$. The Nayar’s interreflection model assumes Lambertian surfaces and overlooks surfaces with unknown non-Lambertian properties.

4. Proposed Method

Given $\mathbf{X} = [X_1, X_2, \dots, X_n]$ a set of n input images and the object mask $\mathbf{O}$, we propose an uncalibrated photometric stereo method to estimate surface normals. Here, each image X_i is reshaped as a column vector and not a facet symbol as used in interreflection modeling. Even though the problem with unknown light directions gives rise to the bas-relief ambiguity [5], we leverage the potential of the deep neural networks to learn those source directions from the input image data using a light estimation network §4.1. The estimated light directions are used by the inverse rendering network §4.2 to infer the unknown BRDFs and surface normals using our proposed rendering equation. Our rendering approach explicitly utilizes the role of the light source and interreflections in the image reconstruction process.

4.1. Light Estimation Network

Given $\mathbf{X}$ and $\mathbf{O}$, the light estimation network predicts the light source intensities ($\mathbf{e}_i$ ’s) and direction vectors ($\mathbf{l}_i$ ’s). We can train such a network either by regressing the intensity values and the corresponding unit vector in the source’s direction or classifying intensity values into pre-defined angle-range bins. The latter choice seems reasonable as it is easier than regressing the exact direction and intensity values. Further, quantizing the continuous space of directions and intensities for classification makes the network robust to

Figure 2: (a) The estimated source directions are given by two parameters: $\phi \in [0, \pi]$ and $\theta \in [-\pi/2, \pi/2]$. (b) Illustration of surface reflectance. When light ray $\mathbf{l}_i$ hits a surface element, the specular component along the view-direction of the point $\mathbf{x}$ due to i^{th} source is given by $\mathbf{r}_{xi}$. Figure 2(b) geometry presentation is inspired by Keenan work [17].

small changes due to outliers or noise. Following that, we express the light source directions in the range $\phi \in [0, \pi]$ for azimuth angles and $\theta \in [-\pi/2, \pi/2]$ for elevation angles (Fig.2(a)). We divide the azimuth and elevation spaces into $K_d = 36$ classes. We classify azimuth and elevation separately, which reduces the problem’s dimensionality and leads to efficient computation. Similarly, we divide the light intensity range $[0.2, 2.0]$ into $K_e = 20$ classes [10].

We used seven feature extraction layers to extract image features for each input image separately, where each layer applies 3×3 convolution and LReLU activation [74]. The weights of the feature extraction layers are shared among all the input images. However, single image features cannot completely disambiguate the object geometry with the light source information. Therefore, we utilize multiple images to have a global implicit knowledge about the surface’s geometry and its reflectance property. We use image specific local features and combine them using a fusion layer to get a global representation of the image set via a max-pooling operation (Fig.3). The global feature representation with the image-specific features is then fed to a classifier. The classifier applies four layers of 3×3 convolution and LReLU activation [74] as well as two fully-connected layers to provide output softmax probability vectors for azimuth (K_d), elevation (K_d), and intensity (K_e). Similar to the feature extraction, the classifier weights are shared among each other. The output value with maximum probability is converted into a light direction vector $\mathbf{l}_i$ and scalar intensity $\mathbf{e}_i$.

Loss function for Light Estimation Network. The light estimation network is trained using a multi-class cross-entropy loss [10]. The total calibration loss $\mathcal{L}_{\text{calib}}$ is:

$$\mathcal{L}_{\text{calib}} = \mathcal{L}_{az} + \mathcal{L}_{el} + \mathcal{L}_{in} \quad (7)$$

Here $\mathcal{L}_{az}$, $\mathcal{L}_{el}$, and $\mathcal{L}_{in}$ are the loss terms for azimuth, elevation, and intensity respectively. We used synthetic Blobby and Sculpture datasets [12] to train the network. The light source labels from these datasets are used for supervision at the train time. The network is trained using the above loss for once and the same network is used at the test time for all other datasets §5.

4.2. Inverse Rendering Network

To estimate an object surface normals from $\mathbf{X}$, we leverage neural networks' powerful capability to learn from data. The prime reason for that is, it is difficult to mathematically model the broad classes of BRDFs without any prior assumptions about the reflectance model [21, 16, 22]. Although there are methods to estimate BRDF values using its isotropic and low-frequency property [29, 61], it prohibits the modeling of unrestricted reflectance behavior of the material. Instead of such explicit modeling, we build on the idea of neural inverse rendering [65], where the BRDFs and surface normals are predicted during the image reconstruction process by the neural network. We go beyond Tani et al. [65] work by proposing an inverse rendering network that synthesizes the input images using a rendering equation that explicitly uses interreflections to infer surface normals.

(a) Surface Normal Modeling. We first convert $\mathbf{X}$ into a tensor $\mathcal{X} \in \mathbb{R}^{h \times w \times nc}$, where $h \times w$ denote the spatial dimensions, n is the number of images, and c is the number of color channels ($c = 1$ for grayscale and $c = 3$ for color images). $\mathcal{X}$ is then mapped to a global feature map Φ as follows:

$$\Phi = \xi_f(\mathcal{X}, \mathbf{O}, \Theta_f) \quad (8)$$

$\mathbf{O}$ is used to separate the object information from the background. ξ_f is a three layer feed-forward convolutional network with learnable parameter Θ_f . Each layer applies 3×3 convolution, batch-normalization [32] and ReLU activation [74] to extract global feature map Φ . In the next step, we use Φ to compute the surface normals. Let ξ_{n1} be the function that converts Φ into output normal map $\mathbf{N}_o$ via 3×3 convolution and L2-normalization operation.

$$\mathbf{N}_o = \xi_{n1}(\Phi, \Theta_{n1}) \quad (9)$$

Here, Θ_{n1} is the learnable parameter. We used the estimated $\mathbf{N}_o$ to compute $\mathbf{N}_{ny}$ using function ξ_{n2} .

$$\mathbf{N}_{ny} = \xi_{n2}(\mathbf{N}_o, \mathbf{P}, \mathbf{K}) \quad (10)$$

ξ_{n2} requires the interreflection kernel $\mathbf{K}$ and albedo matrix $\mathbf{P}$ as input. To calculate $\mathbf{K}$, we integrate the $\mathbf{N}_o$ over masked object pixel coordinates (x, y) to obtain the depth map [3, 63]. Afterward, the depth map is used to infer the kernel matrix $\mathbf{K}$ (see Eq:(4)). Once we have $\mathbf{K}$, we employ Eq:(20) to compute $\mathbf{N}_{ny}$. Later, $\mathbf{N}_{ny}$ is used in the rendering equation (Eq:(15)) for image reconstruction.

(b) Reflectance Modeling. For effective learning of BRDFs, it is important to model the specular component. To incorporate that, we feed a specular map along with the input image as a channel. Consider the specular-reflection direction $\mathbf{r}_{xi}$ at a surface element $\mathbf{x}$ with normal $\mathbf{n}_o(\mathbf{x})$ due to the i^{th} light source. We compute $\mathbf{r}_{xi}$ along the view-direction vector $\mathbf{v}$ using the following relation:

$$\mathbf{r}_{xi} = \mathbf{v}^T \left(2(\mathbf{n}_o(\mathbf{x})^T \mathbf{l}_i) \cdot \mathbf{n}_o(\mathbf{x}) - \mathbf{l}_i \right) \quad (11)$$

Here, $\|\mathbf{l}_i\|_2, \|\mathbf{n}_o(\mathbf{x})\|_2, \|\mathbf{r}_{xi}\|_2$ are 1 (see Fig.2(b)). Computing $\mathbf{r}_{xi}$ for all surface points provides the specular-reflection map $R_i \in \mathbb{R}^{h \times w \times 1}$. Concatenating $X_i \in \mathbb{R}^{h \times w \times c}$ with R_i across channel guides the network to learn complex BRDFs. Thus, we compute feature map S_i as:

$$S_i = f_{sp}(X_i \oplus R_i, \Theta_{sp}) \quad (12)$$

We used $\oplus$ to denote the concatenation operation. f_{sp} is a three-layer network where each layer applies 3×3 convolution, batch-normalization [32] and ReLU operations [74]. Although the feature map S_i models the actual specular component of a BRDF, it is computed using a single image observation X_i which has limited information. To enrich the feature, we concatenate it with the global features Φ (see Eq:(8)) and compute enhanced feature block Z_i .

$$Z_i = f_{lg}(S_i \oplus \Phi, \Theta_{lg}) \quad (13)$$

f_{lg} function applies 1×1 convolution, batch normalization [32] and ReLU operations [74] to estimate Z_i . Finally, we define the reflectance function f_r that blends the image specific features with Φ along with the specular component of the image to compute the reflectance map Ψ_i .

$$\Psi_i = f_r(Z_i, \Theta_{ri}) \quad (14)$$

The function f_r applies 3×3 convolution, batch normalization [32], ReLU operation [74] with an additional 3×3 convolution layer to compute Ψ_i . The predicted Ψ_i by the network contains the BRDFs and cast shadow information. The specular (Θ_{sp}), local-global (Θ_{lg}), and reflectance image (Θ_{ri}) parameters are learned over SGD iteration by the network. Details about the implementation of above functions, learning and testing strategy are described in §5.

(c) Rendering equation. Assuming photometric stereo setup, once we have the surface normals, reflectance map, and light source information, we render the input image using the following equation:

$$\tilde{X}_i = \Psi_i \odot (e_i \cdot \zeta_a(\mathbf{N}_{ny}, \mathbf{l}_i)) \quad (15)$$

Here, we explicitly model the effects of interreflections in the image formation. For a given source, Ψ_i encapsulates the BRDF values with the cast shadow information. Further, ζ_a is defined for the attached shadow. With a slight abuse of notation used in Eq:(1), ζ_a computes the inner product between a light source and the surface normal matrix for each pixel, and the maximum operation is done element-wise i.e., $\max(\mathbf{N}_{ny}^T \mathbf{l}_i, 0)$. $e_i \in \mathbb{R}_+$ is a scalar intensity value of the light source, and $\odot$ denotes the Hadamard product. Fig.3 shows the entire rendering network pipeline.

Figure 3: The proposed method consists of two networks. Light estimation network initially predicts the light source directions and intensities from input images. Then, neural inverse rendering network uses the images and the light source estimations to recover surface normals, depth and BRDF values.

Loss Function for Inverse Rendering Network. To train the proposed inverse rendering network, we use l_1 loss between the rendered images $\tilde{\mathbf{X}}$ and input images $\mathbf{X}$ on the masked pixels ($\mathbf{O}$). The network parameters are learned by minimizing the following loss using the SGD algorithm:

$$\mathcal{L}_{rec}(\mathbf{X}, \tilde{\mathbf{X}}) = \frac{1}{mnc} \sum_{i,c,\mathbf{x}} |X_{i,c}(\mathbf{x}) - \tilde{X}_{i,c}(\mathbf{x})| \quad (16)$$

Here, m is the number of pixels within $\mathbf{O}$ and n, c are the number of input images and color channels, respectively. The optimization of the above image reconstruction loss function seems reasonable; but, it may provide unstable behavior leading to inferior results. Therefore, we apply weak supervision to the network at the early stages of the optimization by adding a surface normal regularizer in the loss function using an initial normal estimate $\mathbf{N}_{init}$. Such a strategy guides the network for stable convergence behavior and a better solution to the surface normals. The total loss function is defined as:

$$\mathcal{L} = \mathcal{L}_{rec}(\mathbf{X}, \tilde{\mathbf{X}}) + \lambda_w \mathcal{L}_{weak}(\mathbf{N}_{ny}, \mathbf{N}_{init}) \quad (17)$$

where, function $\mathcal{L}_{weak}$ is defined as $\mathcal{L}_{weak}(\mathbf{N}_{ny}, \mathbf{N}_{init}) = \frac{1}{m} \sum_{\mathbf{x}} \|\mathbf{n}_{ny}(\mathbf{x}) - \mathbf{n}_{init}(\mathbf{x})\|_2^2$. Least-square solution of $\mathbf{N}$ in Eq:(2) can provide weak supervision to the network in the early stage of the optimization. However, such initialization may provide undesirable behavior at times. Therefore, we adhere to the robust optimization algorithm on photometric stereo (§5) to initialize the surface normal in Eq:(17).

5. Dataset Acquisition and Experiments

We performed evaluations of our method on DiLiGenT dataset [62]. DiLiGenT is a standard benchmark for photometric stereo, consisting of ten different real-world objects. Despite it provides surfaces of diverse reflectances, the subjects are not elegant for studying interreflections. Therefore, we propose a new dataset that is apt for analyzing such

complex imaging phenomena. The acquisition is performed using two different setups. In the first setup, we designed a physical dome system to capture the cultural artifacts. It is a 35cm hemispherical structure with 260 LEDs on the nodes for directed light projection, and with a camera on top, looking down vertically. The object under investigation lies at the center. Using it, we collected images of three historical artifacts (*Tablet1*, *Tablet2*, *Broken Pot*) with spatial resolution of 180×225 . Ground-truth normals are acquired using active sensors with post-refinements. We noted that it is onerous to capture 3D surfaces with high-precision. For this reason, we simulated the dome environment using Cinema 4D software with 100 light sources. Using this synthetic setup, we rendered images of three objects (*Vase*, *Golf-ball*, *Face*) with spatial resolution of 256×256 . Our dataset introduces new subjects with general reflectance property to initiate a broader adaptation of photometric stereo algorithm for extracting 3D surface information of real objects.

Implementation Details. Our method is implemented in PyTorch [54]. The light estimation network is trained using Blobby and Sculpture datasets [12] with Adam [34] optimizer and initial learning rate of 5×10^{-4} . We trained the model for 20 epochs with a batch size of 32. The learning rate is divided by two after every 5 epochs. Training of the neural inverse rendering network is not required as it learns the network parameters at the test time. However, the initialization of the network is crucial for stable learning.

• **Initialization:** Our method uses an initial surface normals prior $\mathbf{N}_{init}$ (Eq:(17)) to warm up the rendering network and to initialize the interreflection kernel $\mathbf{K}$ values. Woodham’s classical method [69] is a conventional way to do so under given light sources. However, initialization using Woodham’s method is observed to provide a unstable network behavior leading to inferior results [65]. Therefore, for initialization, we propose to use partial sum of singular values optimization [50]. Let $\mathbf{X} \in \mathbb{R}^{m \times n}$, $\mathbf{L} \in \mathbb{R}^{3 \times n}$, $\mathbf{N} \in \mathbb{R}^{3 \times m}$, then Eq:(2) under Lambertian assumption with $\rho = 1$ can

Type	G.T. Normal	Methods↓ Dataset →	Ball	Cat	Pot1	Bear	Pot2	Buddha	Goblet	Reading	Cow	Harvest	Average
Classical	✗	Alldrin et al.(2007) [2]	7.27	31.45	18.37	16.81	49.16	32.81	46.54	53.65	54.72	61.70	37.25
Classical	✗	Shi et al.(2010) [60]	8.90	19.84	16.68	11.98	50.68	15.54	48.79	26.93	22.73	73.86	29.59
Classical	✗	Wu et al.(2013) [73]	4.39	36.55	9.39	6.42	14.52	13.19	20.57	58.96	19.75	55.51	23.93
Classical	✗	Lu et al.(2013) [45]	22.43	25.01	32.82	15.44	20.57	25.76	29.16	48.16	22.53	34.45	27.63
Classical	✗	Pap. et al.(2014) [52]	4.77	9.54	9.51	9.07	15.90	14.92	29.93	24.18	19.53	29.21	16.66
Classical	✗	Lu et al.(2017) [44]	9.30	12.60	12.40	10.90	15.70	19.00	18.30	22.30	15.00	28.00	16.30
NN-based	✓	Chen et al.(2018) [12]	6.62	14.68	13.98	11.23	14.19	15.87	20.72	23.26	11.91	27.79	16.02
NN-based	✓	Chen et al.(2018) [†] [12]	3.96	12.16	11.13	7.19	11.11	13.06	18.07	20.46	11.84	27.22	13.62
NN-based	✓	Chen et al.(2019) [10]	2.77	8.06	8.14	6.89	7.50	8.97	11.91	14.90	8.48	17.43	9.51
NN-based	✗	Ours	3.78	7.91	8.75	5.96	10.17	13.14	11.94	18.22	10.85	25.49	11.62

Table 1: Without using ground-truth light or surface normals of this dataset at train time, our method supplies results that is comparable to the recent state-of-the-art [10]. The 1st and 2nd best performing methods are colored in light-red and dark-red respectively. **G.T. Normal** column indicates the use of ground-truth normal at train time. Comparisons are done against well-known uncalibrated methods. [†] indicates the deeper version of the UPS-FCN model.

Figure 4: Qualitative results on DiLiGenT dataset using our method.

be written as $\mathbf{X} = \mathbf{N}^T \mathbf{L} + \mathbf{E}$. Here, $\mathbf{E} \in \mathbb{R}^{m \times n}$ is a matrix of outliers and assumed to be sparse [71]. Substituting $\mathbf{Z} = \mathbf{N}^T \mathbf{L}$, the normal estimation under low rank assumption can be formulated as a RPCA problem [71]. We know that RPCA performs the nuclear norm minimization of $\mathbf{Z}$ matrix which not only minimizes the rank but also the variance of $\mathbf{Z}$ within the target rank. Now, for the photometric stereo model, it is easy to infer that $\mathbf{N}$ lies in a rank-3 space. As the true rank for $\mathbf{Z}$ is known from its construction, we do not minimize the subspace variance within the target rank (K). We preserve the variance of information within the target rank while minimizing other singular values outside it via the following optimization:

$$\min_{\mathbf{Z}, \mathbf{E}} \|\mathbf{Z}\|_{r=K} + \lambda \|\mathbf{E}\|_1, \text{ subject to: } \mathbf{X} = \mathbf{Z} + \mathbf{E} \quad (18)$$

Eq:(22) is a well-studied problem and we solved it using ADMM [8, 50, 40]. We use the Augmented Lagrangian form of Eq:(22) to solve $\mathbf{Z}$, $\mathbf{E}$ for $K = 3$. The recovered solution is used to initialize the surface normal in Eq:(17). For detailed derivations, refer to supplementary material.

• **Testing:** For testing, we first feed the test images to the light estimation network to get source directions and intensities. For objects like *Vase*, where the cast shadows and interreflections play a vital role in the object’s imaging, light estimation network can have questionable behavior. So, we use the light source directions and intensities estimated from a calibration sphere for testing our synthetic objects. Once normal is initialized using our robust ap-

Figure 5: Qualitative results of our method on proposed dataset.

proach, we learn inverse rendering network’s parameters by minimizing $\mathcal{L}$ of Eq:(17). To compute $\mathcal{L}_{rec}$, we randomly sample 10% of the pixels in each iteration and compute it over these pixels to avoid local minimum. To provide weak-supervision, we set $\lambda_w = \mathcal{L}_{rec}(0, \mathbf{X})$ to balance the influence of $\mathcal{L}_{rec}$ and $\mathcal{L}_{weak}$ to network learning process. Note that λ_w is set to zero after 50 iterations to drop early stage weak-supervision. We perform 1000 iterations in total with initial learning rate of 8×10^{-4} . The learning rate is reduced by factor of 10 after 900 iterations for fine-tuning. Before feeding the images to the normal estimation network, we normalize them using a global scaling constant σ , i.e. the quadratic mean of pixel intensities $\mathbf{X}' = \mathbf{X}/(\sigma)$. During the learning of inverse rendering network, we repeatedly update the kernel $\mathbf{K}$ using $\mathbf{N}_o$ after every 100 iterations.

5.1. Evaluation, Ablation Study and Limitation

(a) **DiLiGenT Dataset.** Table(1) provides statistical comparison of our method against other uncalibrated methods on DiLiGenT benchmark. We used popular mean angular error (MAE) metric to report the results. It can be inferred that our method achieves competitive results on this benchmark with an average MAE of 11.62 degrees, achieving the second best performance overall without ground-truth surface normal supervision. On the contrary, the best performing method [10] uses ground-truth normals during training, and therefore, it performs better for objects like *Harvest*, where imaging is deeply affected by discontinuities.

Type	G.T. Normal	Methods ↓	Dataset →	Vase	Golf-ball	Face	Tablet 1	Tablet 2	Broken Pot	Average
Classical	✗	Nayar et al.(1991) [49]		28.82	11.30	13.97	19.14	16.34	19.43	18.17
NN-based	✓	Chen et al.(2018) [12]		35.79	36.14	48.47	19.16	10.69	24.45	29.12
NN-based	✓	Chen et al.(2019) [10]		49.36	31.61	13.81	16.00	15.11	18.34	24.04
NN-based	✗		Ours	19.91	11.04	13.43	12.37	13.12	18.55	14.74

Table 2: Comparison against recent uncalibrated deep photometric stereo methods and Nayar *et al.* [49] on our dataset. In contrast to our approach, Chen *et al.* [12] and Chen *et al.* [10] require ground-truth normal for training the network. We can observe that our method shows consistent behavior over a diverse dataset that is on average better than other methods. The two best-performing methods are shaded with light-red and dark-red color respectively.

Figure 6: Surface normal accuracy achieved w.r.t its initialization.

(b) **Our Dataset.** Table(2) compares our method with other deep uncalibrated methods on the proposed dataset. For completeness, we analyzed Nayar *et al.* [49] algorithm by using light sources data obtained using our approach. The results show that our method achieves the best performance overall. We observed that other deep learning methods cannot handle objects like *Vase* as they fail to model complex reflectance behavior. Similarly, Nayar *et al.* [49] results indicate that modeling interreflections alone is not sufficient. Since we not only model the effects of interreflections, but also the reflectance mapping associated with the geometry, our method consistently performs well.

(c) **Ablation Study.** For this study, we validate the importance of robust initialization and interreflection modeling.

- **Robust Initialization:** To show the effect of initialization, we consider three cases. First, we use classical approach [69] to initialize inverse rendering network. Second, we replace the classical method with our robust initialization strategy. In the final case, we remove the weak-supervision loss from our method. Fig.7 shows MAE and image reconstruction loss curve per learning iteration obtained on *Cow* dataset. The results indicate that robust initialization allows the network to converge faster as outliers are separated from the images at an initial stage. Fig.6 shows the MAE of surface normals during initialization as compared to the results obtained using our method.

- **Interreflection Modeling:** To demonstrate the effect of interreflection modeling, we remove the function ξ_{n2} in Eq.(10) and use N_o in image reconstruction as in classical rendering. Fig.7 provides learning curves with and without interreflection modeling. As expected, excluding the effect of interreflections inherently impacts the accuracy of the surface normals estimates even if the image reconstruction quality remains consistent. Hence, it is important to explicitly constrain the geometry information.

(d) **Limitations.** Discrete facets assumption of a continuous surface for computing depth and interreflection kernel

Figure 7: **Ablation Study:** We demonstrate the effect of robust initialization on *Cow* (top) and interreflection modeling on *Vase* (bottom).

may not be suitable where the surface is discontinuous in orientation, *e.g.*, surface with deep holes, concentric rings, *etc.* As a result, our method may fail on surfaces with very deep concavities and cases related to naturally occurring optical caustics. As a second limitation, the light estimation network may not resolve GBR ambiguity for all kinds of shapes. Presently, we did not witness such ambiguity with the light calibration network as it is trained to predict lights under non-GBR transformed surface material distribution.

6. Conclusion

From this work, we conclude that uncalibrated neural inverse rendering approach with explicit interreflection modeling enforces the network to model complex reflectance characteristics of objects with different material and geometry types. Without using ground-truth surface normals, we observed that our method could provide comparable or better results than the supervised approaches. And therefore, our work can enable 3D vision practitioners to opt for photometric stereo methods to study a broader range of geometric surfaces. That's said, image formation is a complex process, and additional explicit constraints based on the 3D surface geometry types, material, and light interaction behavior could further advance our work.

Acknowledgement. This work was funded by Focused Research Award from Google (CVL, ETH 2019-HE-318, 2019-HE-323). We thank Vincent Vanweddigen from KU Lueven for providing some datasets for our experiments.

References

- [1] Neil Alldrin, Todd Zickler, and David Kriegman. Photometric stereo with non-parametric and spatially-varying reflectance. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008.
- [2] Neil G Alldrin, Satya P Mallick, and David J Kriegman. Resolving the generalized bas-relief ambiguity by entropy minimization. In *2007 IEEE conference on computer vision and pattern recognition*, pages 1–7. IEEE, 2007.
- [3] Doris Antensteiner, Svorad Štolc, and Thomas Pock. A review of depth and normal fusion algorithms. *Sensors*, 18(2):431, 2018.
- [4] Louis-Philippe Asselin, Denis Laurendeau, and Jean-Francois Lalonde. Deep SVBRDF estimation on real materials. In *2020 International Conference on 3D Vision (3DV)*. IEEE, nov 2020.
- [5] Peter N Belhumeur, David J Kriegman, and Alan L Yuille. The bas-relief ambiguity. *International journal of computer vision*, 35(1):33–44, 1999.
- [6] Sai Bi, Zexiang Xu, Kalyan Sunkavalli, David Kriegman, and Ravi Ramamoorthi. Deep 3d capture: Geometry and reflectance from sparse multi-view images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5960–5969, 2020.
- [7] James F Blinn. Models of light reflection for computer synthesized pictures. In *Proceedings of the 4th annual conference on Computer graphics and interactive techniques*, pages 192–198, 1977.
- [8] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.
- [9] Manmohan Krishna Chandraker, Fredrik Kahl, and David J Kriegman. Reflections on the generalized bas-relief ambiguity. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR’05)*, volume 1, pages 788–795. IEEE, 2005.
- [10] Guanying Chen, Kai Han, Boxin Shi, Yasuyuki Matsushita, and Kwan-Yee K Wong. Self-calibrating deep photometric stereo networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8739–8747, 2019.
- [11] Guanying Chen, Kai Han, Boxin Shi, Yasuyuki Matsushita, and Kwan-Yee Kenneth Wong. Deep photometric stereo for non-lambertian surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [12] Guanying Chen, Kai Han, and Kwan-Yee K Wong. Ps-fcn: A flexible learning framework for photometric stereo. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–18, 2018.
- [13] Guanying Chen, Michael Waechter, Boxin Shi, Kwan-Yee K Wong, and Yasuyuki Matsushita. What is learned in deep uncalibrated photometric stereo? In *European Conference on Computer Vision*, 2020.
- [14] Lixiong Chen, Yinqiang Zheng, Boxin Shi, Art Subpa-asa, and Imari Sato. A microfacet-based model for photometric stereo with general isotropic reflectance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [15] Zhang Chen, Anpei Chen, Guli Zhang, Chengyuan Wang, Yu Ji, Kiriakos N Kutulakos, and Jingyi Yu. A neural rendering framework for free-viewpoint relighting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5599–5610, 2020.
- [16] Hin-Shun Chung and Jiaya Jia. Efficient photometric stereo on glossy surfaces with wide specular lobes. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008.
- [17] Keenan Crane. *Conformal Geometry Processing*. PhD thesis, Caltech, June 2013.
- [18] Ondrej Drbohlav and M Chaniler. Can two specular pixels calibrate photometric stereo? In *Tenth IEEE International Conference on Computer Vision (ICCV’05) Volume 1*, volume 2, pages 1850–1857. IEEE, 2005.
- [19] Kenji Enomoto, Michael Waechter, Kiriakos N Kutulakos, and Yasuyuki Matsushita. Photometric stereo via discrete hypothesis-and-test search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2311–2319, 2020.
- [20] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multiview stereopsis. *IEEE transactions on pattern analysis and machine intelligence*, 32(8):1362–1376, 2009.
- [21] Athinodoros S Georgiades. Incorporating the torrance and sparrow model of reflectance in uncalibrated photometric stereo. In *ICCV*, pages 816–823. IEEE, 2003.
- [22] 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, 2009.
- [23] Elaine T Hale, Wotao Yin, and Yin Zhang. Fixed-point continuation for l_1 -minimization: Methodology and convergence. *SIAM Journal on Optimization*, 19(3):1107–1130, 2008.
- [24] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003.
- [25] Steffen Herbot and Christian Wöhler. An introduction to image-based 3d surface reconstruction and a survey of photometric stereo methods. *3D Research*, 2(3):4, 2011.
- [26] Tomoaki Higo, Yasuyuki Matsushita, and Katsushi Ikeuchi. Consensus photometric stereo. In *2010 IEEE computer society conference on computer vision and pattern recognition*, pages 1157–1164. IEEE, 2010.
- [27] Santo Hiroaki, Michael Waechter, and Yasuyuki Matsushita. Deep near-light photometric stereo for spatially varying reflectances. In *European Conference on Computer Vision*, 2020.
- [28] Satoshi Ikehata. Cnn-ps: Cnn-based photometric stereo for general non-convex surfaces. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–18, 2018.
- [29] Satoshi Ikehata and Kiyoharu Aizawa. Photometric stereo using constrained bivariate regression for general isotropic

- surfaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2179–2186, 2014.
- [30] Satoshi Ikehata, David Wipf, Yasuyuki Matsushita, and Kiyoharu Aizawa. Robust photometric stereo using sparse regression. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 318–325. IEEE, 2012.
- [31] 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.
- [32] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015.
- [33] Micah K Johnson and Edward H Adelson. Shape estimation in natural illumination. In *CVPR 2011*, pages 2553–2560. IEEE, 2011.
- [34] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [35] Suryansh Kumar. Jumping manifolds: Geometry aware dense non-rigid structure from motion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5346–5355, 2019.
- [36] Suryansh Kumar, Yuchao Dai, and Hongdong Li. Monocular dense 3d reconstruction of a complex dynamic scene from two perspective frames. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4649–4657, 2017.
- [37] Suryansh Kumar, Yuchao Dai, and Hongdong Li. Superpixel soup: Monocular dense 3d reconstruction of a complex dynamic scene. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [38] Junxuan Li, Antonio Robles-Kelly, Shaodi You, and Yasuyuki Matsushita. Learning to minify photometric stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7568–7576, 2019.
- [39] Min Li, Zhenglong Zhou, Zhe Wu, Boxin Shi, Changyu Diao, and Ping Tan. Multi-view photometric stereo: A robust solution and benchmark dataset for spatially varying isotropic materials. *IEEE Transactions on Image Processing*, 29:4159–4173, 2020.
- [40] Zhouchen Lin, Minming Chen, and Yi Ma. The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices. *arXiv preprint arXiv:1009.5055*, 2010.
- [41] Fotios Logothetis, Ignas Budvytis, Roberto Mecca, and Roberto Cipolla. A CNN based approach for the near-field photometric stereo problem. In *31st British Machine Vision Conference 2020, BMVC 2020, Virtual Event, UK, September 7-10, 2020*. BMVA Press, 2020.
- [42] Fotios Logothetis, Ignas Budvytis, Roberto Mecca, and Roberto Cipolla. Px-net: Simple, efficient pixel-wise training of photometric stereo networks. *arXiv preprint arXiv:2008.04933*, 2020.
- [43] Fotios Logothetis, Roberto Mecca, and Roberto Cipolla. A differential volumetric approach to multi-view photometric stereo. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1052–1061, 2019.
- [44] Feng Lu, Xiaowu Chen, Imari Sato, and Yoichi Sato. Symps: Brdf symmetry guided photometric stereo for shape and light source estimation. *IEEE transactions on pattern analysis and machine intelligence*, 40(1):221–234, 2017.
- [45] Feng Lu, Yasuyuki Matsushita, Imari Sato, Takahiro Okabe, and Yoichi Sato. Uncalibrated photometric stereo for unknown isotropic reflectances. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1490–1497, 2013.
- [46] Wojciech Matusik. *A data-driven reflectance model*. PhD thesis, Massachusetts Institute of Technology, 2003.
- [47] Daisuke Miyazaki, Kenji Hara, and Katsushi Ikeuchi. Median photometric stereo as applied to the segonko tumulus and museum objects. *International Journal of Computer Vision*, 86(2-3):229, 2010.
- [48] Yasuhiro Mukaigawa, Yasunori Ishii, and Takeshi Shikunaga. Analysis of photometric factors based on photometric linearization. *JOSA A*, 24(10):3326–3334, 2007.
- [49] Shree K Nayar, Katsushi Ikeuchi, and Takeo Kanade. Shape from interreflections. *International Journal of Computer Vision*, 6(3):173–195, 1991.
- [50] Tae-Hyun Oh, Hyeonwoo Kim, Yu-Wing Tai, Jean-Charles Bazin, and In So Kweon. Partial sum minimization of singular values in rpca for low-level vision. In *Proceedings of the IEEE international conference on computer vision*, pages 145–152, 2013.
- [51] Thoma Papadhimetri and Paolo Favaro. A new perspective on uncalibrated photometric stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1474–1481, 2013.
- [52] Thoma Papadhimetri and Paolo Favaro. A closed-form, consistent and robust solution to uncalibrated photometric stereo via local diffuse reflectance maxima. *International journal of computer vision*, 107(2):139–154, 2014.
- [53] Jaesik Park, Sudipta N Sinha, Yasuyuki Matsushita, Yu-Wing Tai, and In So Kweon. Multiview photometric stereo using planar mesh parameterization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1161–1168, 2013.
- [54] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017.
- [55] Yvain Quéau, Tao Wu, François Lauze, Jean-Denis Durou, and Daniel Cremers. A non-convex variational approach to photometric stereo under inaccurate lighting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 99–108, 2017.
- [56] Hiroaki Santo, Masaki Samejima, Yusuke Sugano, Boxin Shi, and Yasuyuki Matsushita. Deep photometric stereo network. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 501–509, 2017.

- [57] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4104–4113, 2016.
- [58] Soumyadip Sengupta, Angjoo Kanazawa, Carlos D Castillo, and David W Jacobs. Sfsnet: Learning shape, reflectance and illuminance of faces in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6296–6305, 2018.
- [59] Soumyadip Sengupta, Hao Zhou, Walter Forkel, Ronen Basri, Tom Goldstein, and David Jacobs. Solving uncalibrated photometric stereo using fewer images by jointly optimizing low-rank matrix completion and integrability. *Journal of Mathematical Imaging and Vision*, 60(4):563–575, 2018.
- [60] Boxin Shi, Yasuyuki Matsushita, Yichen Wei, Chao Xu, and Ping Tan. Self-calibrating photometric stereo. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1118–1125. IEEE, 2010.
- [61] 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, 2013.
- [62] Boxin Shi, Zhe Wu, Zhipeng Mo, Dinglong Duan, Sai-Kit Yeung, and Ping Tan. A benchmark dataset and evaluation for non-lambertian and uncalibrated photometric stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3707–3716, 2016.
- [63] Richard Szeliski. *Computer vision: algorithms and applications*. Springer Science & Business Media, 2010.
- [64] Ping Tan, Satya P Mallick, Long Quan, David J Kriegman, and Todd Zickler. Isotropy, reciprocity and the generalized bas-relief ambiguity. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007.
- [65] Tatsunori Taniai and Takanori Maehara. Neural inverse rendering for general reflectance photometric stereo. In *International Conference on Machine Learning (ICML)*, pages 4857–4866, 2018.
- [66] Xueying Wang, Yudong Guo, Bailin Deng, and Juyong Zhang. Lightweight photometric stereo for facial details recovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 740–749, 2020.
- [67] Xi Wang, Zhenxiong Jian, and Mingjun Ren. Non-lambertian photometric stereo network based on inverse reflectance model with collocated light. *IEEE Transactions on Image Processing*, 29:6032–6042, 2020.
- [68] Olivia Wiles and Andrew Zisserman. SilNet : Single- and multi-view reconstruction by learning from silhouettes. In *Proceedings of the British Machine Vision Conference 2017*. British Machine Vision Association, 2017.
- [69] Robert J Woodham. Photometric method for determining surface orientation from multiple images. *Optical engineering*, 19(1):191139, 1980.
- [70] Changchang Wu, Sameer Agarwal, Brian Curless, and Steven M Seitz. Schematic surface reconstruction. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1498–1505. IEEE, 2012.
- [71] Lun Wu, Arvind Ganesh, Boxin Shi, Yasuyuki Matsushita, Yongtian Wang, and Yi Ma. Robust photometric stereo via low-rank matrix completion and recovery. In *Asian Conference on Computer Vision*, pages 703–717. Springer, 2010.
- [72] Tai-Pang Wu and Chi-Keung Tang. Photometric stereo via expectation maximization. *IEEE transactions on pattern analysis and machine intelligence*, 32(3):546–560, 2009.
- [73] Zhe Wu and Ping Tan. Calibrating photometric stereo by holistic reflectance symmetry analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1498–1505, 2013.
- [74] Bing Xu, Naiyan Wang, Tianqi Chen, and Mu Li. Empirical evaluation of rectified activations in convolutional network. *arXiv preprint arXiv:1505.00853*, 2015.
- [75] Zhuokun Yao, Kun Li, Ying Fu, Haofeng Hu, and Boxin Shi. Gps-net: Graph-based photometric stereo network. *Advances in Neural Information Processing Systems*, 33, 2020.
- [76] Enliang Zheng and Changchang Wu. Structure from motion using structure-less resection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2075–2083, 2015.
- [77] Qian Zheng, Yiming Jia, Boxin Shi, Xudong Jiang, Ling-Yu Duan, and Alex C Kot. Spline-net: Sparse photometric stereo through lighting interpolation and normal estimation networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8549–8558, 2019.

[Supplementary Material] Uncalibrated Neural Inverse Rendering for Photometric Stereo of General Surfaces

Abstract

In our supplementary material, we first present a few case studies to analyze our method’s effectiveness. Next, we give a detailed description of our coding implementation for training and testing the neural network outlined in the main paper. Formally, this report includes the coding platform details —both hardware and software, with train and test time observed across different datasets. Further, mathematical derivations of our robust initialization and specular-reflectance map formulations are supplied. Finally, we analyze the light estimation performance and discuss the possible future extensions of our method. Besides, our supplementary material includes a short video clip that illustrates the image acquisition setup and visual results.

A. Case Study

This section provides the observation on the case study that we conducted for our proposed method. It is done to analyze the behavior of our method under different possible variations in our experimental setup. Such a study can help us understand the behavior, pros, and cons of our approach.

Case Study 1: *What if we use ground-truth light as input to inverse rendering network instead of relying on light estimation network?*

This case study investigates the reliability of our method. To conduct this experiment, we supplied ground-truth light source directions and intensities as input to the inverse rendering network and robust initialization. The goal is to study the expected deviation in the accuracy of surface normals when ground-truth light sources information is used, compared to the light calibration network. Table (3) compares our method’s performance with recent deep calibrated photometric stereo methods on our proposed dataset. The results show that our inverse rendering method achieves the best performance in the calibrated setting, although it does not use a training dataset like other deep-learning-based methods. Additionally, we observed that the CNNPS model proposed by Ikehata [28] which performs per-pixel estimation using observation maps, may not provide accurate surface normals for interreflecting surfaces such as the *Vase* and the *Broken Pot*. Hence, we conclude that extracting information by utilizing the surface geometry is crucial for solving photometric stereo since all surface points affect each other.

Moreover, in Table (3), we show the comparison of our method’s performance under calibrated and uncalibrated settings. Our method achieves 12.68° MAE on average, using ground-truth light as input. At the same time, it reaches

Figure 8: **Failure case:** Qualitative results on the Rose dataset.

an average MAE of 14.74° utilizing the information of the light source obtained from the light estimation network. The difference between these two scores is 2.06 degrees, which indicates that the gap between the calibrated and uncalibrated settings is not substantial. Accordingly, we can conclude that our method is robust to the variations in the estimated lighting. Further, we observed that our method performs better with the network estimated light sources information in the categories like *Golf-ball*, *Face*. Hence, based on that observation, we can conclude that the availability of ground-truth calibration data is not a strict requirement for achieving better surface normals estimates in photometric stereo for all kinds of surface geometry.

Case Study 2: *What if we use noisy images?*

Photometric stereo uses a camera acquisition setup, and this implies that noise due to imaging is inevitable. This case study aims to investigate the behavior of our method on different noise levels. To study such a behavior, we synthesized images by adding noise to the images of our proposed dataset. Fig.9 compares the performance of our method under different noise levels. For this case study, we used zero-mean Gaussian noise with different standard deviations ($\sigma=0.05$, $\sigma=0.1$, $\sigma=0.2$). The quantitative results indicate that increasing the noise generally degrades the performance. We observed that the behavior under different noise levels varies among the subjects.

Case Study 3: *Photometric stereo on concentric surfaces with deep concavities and large surface discontinuity.*

To study our photometric stereo method’s boundary-condition, we took a complex geometric structure with concentric surfaces, deep-concavities, and large discontinuities for investigation. Accordingly, we synthesized the *Rose* dataset using the same dome-settings outlined in the main paper. Fig.8 shows the qualitative results obtained on this dataset. Our method achieves 60.82 degrees of MAE on this particular example. We observed that our approach could not handle this complex geometry because the surface is highly discontinuous with excessive gaps between the leaves. The scene is also affected by occlusions and cast shadows, and therefore, modeling the interreflections for this case seems very difficult.

Though our method applies to a broad range of objects,

Type	G.T. Normal	Methods, Dataset →	Vase	Golf-ball	Face	Tablet 1	Tablet 2	Broken Pot	Average Performance
NN-based	✓	Ikehata (2018)[28]	34.00	14.96	16.61	16.64	12.32	18.31	18.81
NN-based	✓	Chen et al.(PS-FCN)(2018)[12]	27.11	15.99	16.17	10.23	5.79	8.68	14.00
NN-based	✗	Ours (Ground-truth light/ calibrated)	16.40	14.23	14.24	10.77	4.49	15.92	12.68
NN-based	✗	Ours (Estimated light/ uncalibrated)	19.91	11.04	13.43	12.37	13.12	18.55	14.74
		Diff. in MAE (Ours(Est)-Ours(GT))	+3.51	-3.19	-0.81	+1.60	+8.63	+2.63	+2.06

Table 3: Comparison of recent deep **calibrated** photometric stereo methods Ikehata [28] and Chen *et al.* [12] (PS-FCN) against our method under **uncalibrated** and **calibrated** setting. For testing our method under the calibrated setting, we evaluate the performances assuming that ground-truth light source directions and intensities are available. Note that Chen *et al.* [12] and Ikehata [28] additionally uses ground-truth surface normals for training, in contrast to our method. The last row shows the difference between our method results when used under uncalibrated and calibrated setting respectively. We can see that the average difference in MAE between the two settings of our method is not significant.

our interreflection modeling is inspired by Nayar *et al.* [49] formulation, which may not hold for all kinds of surfaces. The interreflection modeling computes depth from the normal map under the continuous surface assumption, which fails in this case study. Furthermore, it models continuous surfaces with discrete facets. Due to such limitations, our method may not be suitable for concentric surfaces with deep concavities and large discontinuities. In such cases, the interreflection effect is very complicated, and our approach may disappoint to model such complex light phenomena.

B. Coding Details

This section provides a detailed description of our source code implementation. We start by introducing the light estimation network’s training phase. Then we focus on the testing phase, where the inverse rendering network is optimized to estimate the surface normals, depth, and BRDF values. Finally, we present details on training and testing run-times.

B.1. Training Details

As our inverse rendering network optimizes its learnable parameters at the test time, we apply a training stage only to the light estimation network. For training the network, we used Blobby and Sculpture datasets that are introduced by Chen *et al.* [12]. This dataset is created by using 3D geometries of Blobby [33], and Sculpture [68] shape datasets and combining them with different material BRDFs taken from MERL dataset [46]. In total, the complete dataset contains 85212 subjects. For each subject, there exist 64 renderings with different light source directions. The intensity of the light sources is kept constant during the whole data generation process. To simulate different intensities during training, image intensity values are randomly generated in the range of $[0.2, 2]$, and these intensity values are used to scale the image data linearly. In each training iteration, the input data is perturbed in the range of $[-0.025, 0.025]$ for augmentation.

The light estimation network is a multiple-input multiple-output (MIMO) system which requires images of the same object captured under different illumination conditions (see Fig.10). The core idea is that all input images

Figure 9: The performance of our method against different noise levels. We used zero-mean Gaussian noise ($\mu = 0$) with different standard deviations (σ). We observed that increasing the noise level generally degrades the performance. Still, the behavior under different noise levels varies among the subjects as the performance depends on the signal-to-noise ratio of the images.

have the same surface, and having more images helps the network extract better global features. During training, we use 32 images of the same object for global feature extraction. Note that all of the images are used for feature extraction at test time to achieve the best performance from the network.

B.2. Testing Details

Given a set of test images $\mathbf{X}$ and object mask $\mathbf{O}$, we first use the light estimation network to have light source directions and intensities. However, the light estimation network operates on 128×128 images because it uses fully connected layers for classification, and these layers process only fixed-length vectors. Consequently, we scale the input images into the resolution of 128×128 before feeding them to the network. We apply this pre-processing step only for the light estimation network and use the original image size for all other operations during testing.

Once we obtain the light source directions and intensities, we apply the robust initialization algorithm to get an initial surface normal matrix $\mathbf{N}_{init}$. It also provides an albedo map that is transformed into $\mathbf{P} \in \mathbb{R}^{m \times m}$ which is required for interreflection modeling. Details about the robust initialization method are explained and derived in §C.1.

After the robust initialization process, we start the optimization of our inverse rendering framework. First, we

Figure 10: **Architecture of the light estimation network.** The network first extracts features from the input images separately using feature extraction layers (purple). Then, the extracted image-specific features (light-green) are fused with max-pooling operation to obtain a global representation of the entire scene (dark-green). Finally, all image-specific and global features are used in classifier network where convolution (brown) and fully-connected (blue) layers are used to predict light intensity values (e_i 's) and direction vectors (l_i 's).

initialize all the network parameters (Θ_f , Θ_{n1} , Θ_{sp} , Θ_{lg} , Θ_{ri}) which correspond to the weights of the convolution operations. In this step, we initialize the weights randomly by sampling from a Gaussian distribution with zero mean and 0.02 variance. We perform 1000 iterations in total using Adam optimizer [34] with an initial learning rate of 8×10^{-4} . The learning rate is reduced by a factor of 10 after 900 iterations for fine-tuning. We observed that setting these hyperparameters may result in convergence problems in our dataset. For this reason, we set the initial learning rate of the estimation branch (ξ_f and ξ_{n1}) to 8×10^{-5} while experimenting on our dataset. We also inject Gaussian noise with zero mean and 0.1 variance to the images before feeding them to f_{sp} for image reconstruction. We observed that this prohibits the network from generating degenerate solutions. At every 100 iterations, we update the depth and the interreflection kernel matrix entries using the normal estimation N_o .

(a) Depth: To compute the depth from normals, we use a gradient-based method with surface orientation constraint [3]. Given the surface normals, we first compute a gradient field $\hat{G} \in \mathbb{R}^{h \times w \times 2}$ where h and w are the spatial dimensions. The idea is that the gradient field computed from surface normal map and the estimated depth $D \in \mathbb{R}^{h \times w}$ should be consistent, *i.e.*, $\nabla D \approx \hat{G}$. That corresponds to an overdetermined system of linear equations and is solved by minimizing the following objective function *i.e.*, Eq:(19) using the least-squares approach

$$\min_D \|\nabla D - \hat{G}\|^2 \quad (19)$$

(b) Interreflection Modeling: To consider the effect of interreflection during the image reconstruction process, we define the function ξ_{n2} which uses the estimated normal $N_o \in \mathbb{R}^{3 \times m}$, albedo matrix $P \in \mathbb{R}^{m \times m}$ and the inter-

reflection kernel $K \in \mathbb{R}^{m \times m}$. Given all these components, Nayar *et al.* [49] relates the observed radiance (X) and the radiance due to primary light source (X_s) as follows:

$$X = (I - PK)^{-1} X_s \quad (20)$$

Assuming the surface shows Lambertian reflectance property, we model the radiance in terms of facet matrices as follows:

$$X = F_{ny} L, \quad X_s = F L, \quad \Rightarrow F_{ny} = (I - PK)^{-1} F \quad (21)$$

Here $F_{ny} \in \mathbb{R}^{m \times 3}$ and $F \in \mathbb{R}^{m \times 3}$ are the facet matrices which contain surface normals N_{ny} and N_o scaled with local reflectance value. We use Eq:(21) to obtain F_{ny} and normalize each row to unit vector to obtain N_{ny} .

The computation of the interreflection kernel K has the complexity of $\mathcal{O}(n^2)$ where n is the number of facets. Therefore, treating each pixel as a facet limits the application of our method. To approximate the effect of interreflections, we downsample the normal maps with the factor of 4 and calculated the kernel values accordingly. After the normal is updated, we scale it to the original size managing the image details appropriately.

B.3. Timing Details

Our framework is implemented in Python using PyTorch version 1.1.0. Table (4) provides the light estimation network's training time and the inference time of neural inverse rendering network on two datasets separately.

C. Mathematical Derivations

Here, we supply the mathematical derivation pertaining to the initialization of the surface normals to the inverse ren-

Figure 11: We present visual results of our method on all of the DiLiGenT categories. The bottom row demonstrates the angular error maps obtained from our estimations and ground-truth normals.

dering network. For completion, we also supplied the well-known deviation of reflection vector §C.2.

C.1. Robust Initialization

Our surface normals initialization procedure aims at recovering the low rank matrix $\mathbf{Z} \in \mathbb{R}^{m \times n}$ from the image matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$ such that $\mathbf{X} = \mathbf{Z} + \mathbf{E}$ where $\mathbf{E} \in \mathbb{R}^{m \times n}$ is the matrix of outliers. Here, we assume that the low-rank matrix follows the classical photometric stereo model ($\mathbf{Z} = \mathbf{N}^T \mathbf{L}$) and the outlier matrix $\mathbf{E}$ is sparse in its distribution. Since it is known by definition that $\mathbf{Z}$ spans a rank-3 space, it can be formulated as a standard RPCA problem [71]. However, we know that RPCA formulation performs the nuclear norm minimization of $\mathbf{Z}$ matrix which not only minimizes the rank but also the variance of $\mathbf{Z}$ within the target rank. Now, for the photometric stereo model, it is easy to infer that $\mathbf{N}$ lies in a rank 3 space. As the true rank for $\mathbf{Z}$ is known from its mathematical construction, we do not want to minimize the subspace variance within the target range. Nevertheless, this strict constraint is difficult to meet due to the complex imaging model, and therefore, we encourage to preserve the variance of information within the target range while minimizing the other singular values outside the target rank (K). So, we minimize the partial sum of the singular values which are outside the target rank with the following optimization as follows:

$$\underset{\mathbf{Z}, \mathbf{E}}{\text{minimize}} \|\mathbf{Z}\|_{r=K} + \lambda \|\mathbf{E}\|_1, \quad \text{subject to: } \mathbf{X} = \mathbf{Z} + \mathbf{E} \quad (22)$$

	GPU	Time
Training of Light Estimation Network	Titan X Pascal (12GB)	≈ 22 hours
Inference on DiLiGenT	GeForce GTX TITAN X (12GB)	53.41 ± 41.57 min per subject
Inference on our Dataset	GeForce GTX TITAN X (12GB)	29.08 ± 15.99 min per subject

Table 4: Measured training and testing time with respect to the utilized hardware. For our dataset, we have 100 to 260 images per subject and the DiLiGenT dataset has 96 images per subject. Note: Deep photometric stereo method processes a set of images rather than one image for estimating normals.

The Augmented Lagrangian function of Eq:(22) can be written as follows:

$$\mathcal{L}(\mathbf{Z}, \mathbf{E}, \mathbf{Y}) = \|\mathbf{Z}\|_{r=K} + \lambda \|\mathbf{E}\|_1 + \frac{\mu}{2} \|\mathbf{X} - \mathbf{Z} - \mathbf{E}\|_F^2 + \langle \mathbf{Y}, \mathbf{X} - \mathbf{Z} - \mathbf{E} \rangle \quad (23)$$

Here, μ is a positive scalar and $\mathbf{Y} \in \mathbb{R}^{m \times n}$ is the estimate of the Lagrange multiplier. As minimizing this function is challenging, we solve it by utilizing the alternating direction method of multipliers (ADMM)[8, 50, 40]. Accordingly, the optimization problem in Eq:(23) can be divided into sub-problems, where $\mathbf{Z}$, $\mathbf{E}$ and $\mathbf{Y}$ are updated alternatively while keeping the other variables fixed.

1. Solution to $\mathbf{Z}$:

$$\mathbf{Z}^* = \underset{\mathbf{Z}}{\text{argmin}} \|\mathbf{Z}\|_{r=K} + \frac{\mu_k}{2} \|\mathbf{Z} - (\mathbf{X} - \mathbf{E}_k + \mu_k^{-1} \mathbf{Y}_k)\|_F^2 \quad (24)$$

The solution to Eq:(24) sub-problem at k^{th} iteration is given by $\mathbf{Z}_k = \mathcal{P}_{K, \mu_k^{-1}}[\mathbf{X} - \mathbf{E}_k + \mu_k^{-1} \mathbf{Y}_k]$ where, $\mathcal{P}_{K, \tau}[\mathbf{M}] = \mathbf{U}_M(\Sigma_{\mathbf{M}_1} + \mathcal{S}_\tau[\Sigma_{\mathbf{M}_2}])\mathbf{V}_M^T$ is the partial singular value thresholding operator [50] and $\mathcal{S}_\tau[x] = \text{sign}(x) \max(|x| - \tau, 0)$ is the soft-thresholding

operator [23]. Here, $\mathbf{U}_M, \mathbf{V}_M$ are the singular vector of matrix $\mathbf{M}$ and $\Sigma_{\mathbf{M}_1} = \text{diag}(\sigma_1, \sigma_2, \dots, \sigma_K, 0, 0)$, $\Sigma_{\mathbf{M}_2} = \text{diag}(0, 0, \dots, \sigma_{K+1}, \dots, \sigma_N)$.

2. Solution to E:

$$\mathbf{E}^* = \underset{\mathbf{E}}{\text{argmin}} \lambda \|\mathbf{E}\|_1 + \frac{\mu_k}{2} \|\mathbf{E} - (\mathbf{X} - \mathbf{Z}_{k+1} + \mu_k^{-1} \mathbf{Y}_k)\|_F^2 \quad (25)$$

The solution to Eq:(25) sub-problem at k^{th} iteration is given by $\mathbf{E}_k = \mathcal{S}_{\lambda \mu_k^{-1}}[\mathbf{X} - \mathbf{Z}_{k+1} + \mu_k^{-1} \mathbf{Y}_k]$ where, $\mathcal{S}_\tau[x] = \text{sign}(x) \max(|x| - \tau, 0)$ is a soft-thresholding operator [23]. For proof of convergence and theoretical analysis of partial singular value thresholding operator kindly refer to Oh *et al.* [50] work. We solve for $\mathbf{Z}, \mathbf{E}$ using ADMM until convergence for $K = 3$ and use the obtained surface normals for initializing the loss function of inverse rendering network.

3. Solution to Y: The variable $\mathbf{Y}$ is updated as follows over the iteration:

$$\mathbf{Y}_{k+1} = \mathbf{Y}_k + \mu_k(\mathbf{X} - \mathbf{Z}_{k+1} - \mathbf{E}_{k+1}) \quad (26)$$

For more details on the implementation kindly refer to Oh *et al.* [50] method.

C.2. Derivation of Specular-Reflection Equation 11 in the Main Paper

For completion, we derive Equation 11 of the main paper that is used to compute the specular-reflection map $R_i \in \mathbb{R}^{h \times w \times 1}$ for each image. To compute it, we first compute $\mathbf{r}_{xi}$ for each point $\mathbf{x}$ that is the direction vector with the highest specular component using the following well-known relation; assuming $\mathbf{l}_i$, and $\mathbf{n}_o$ as unit length vectors:

$$\begin{aligned} \mathbf{r}_{xi} + \mathbf{l}_i &= 2\cos(\theta) \cdot \mathbf{n}_o(\mathbf{x}); \quad \mathbf{n}_o(\mathbf{x})^T \mathbf{l}_i = \cos(\theta) \\ \mathbf{r}_{xi} &= 2(\mathbf{n}_o(\mathbf{x})^T \mathbf{l}_i) \mathbf{n}_o(\mathbf{x}) - \mathbf{l}_i \end{aligned} \quad (27)$$

Here, $\mathbf{r}_{xi}$ is also a unit length vector (see Fig.12). The component of specular reflection in the view-direction $\mathbf{v} = (0, 0, 1)^T$ of the point $\mathbf{x}$ due to i^{th} light is computed as:

$$\mathbf{r}_{xi} = \mathbf{v}^T \left(2(\mathbf{n}_o(\mathbf{x})^T \mathbf{l}_i) \mathbf{n}_o(\mathbf{x}) - \mathbf{l}_i \right) \quad (28)$$

The above relation show that the specular highlights are strongest if the normal $\mathbf{n}_o(\mathbf{x})$ is closest to $\mathbf{r}_{xi}$. Performing this operation for each point gives us the specular-reflection map $\mathbf{R}_i$.

D. Statistical Analysis of Estimated Light Source Directions

We aim to investigate the source directions' behavior predicted by the light estimation network (Fig.10). For that

Figure 12: Illustration of surface reflectance. When light ray $\mathbf{l}_i$ hits a surface element, the specular component along the view-direction of the point $\mathbf{x}$ due to i^{th} source is given by $\mathbf{r}_{xi}$. This presentation of 3D geometry is inspired by Keenan work [17].

purpose, we use a well-known setup used for light calibration, *i.e.*, a calibration sphere. Our renderings from the calibration sphere (see Fig.13(a)) has specular highlights and attached shadows, which provide useful cues for the light estimation network. Figures 13(b)-13(d) illustrate the x , y and z components of the estimated light source direction and ground-truth with respect to the images. We measured the MAE between these vectors as 6.31 degrees. We also observed that the x and y components match well with the ground-truth values. On the other hand, we observed fluctuations on the z component where the values slightly deviate from the ground-truth in a specific pattern. One possible explanation for this observation is that the network has a bias such that its behavior changes in the different regions of the lighting space. Since we generated the data by moving the light source on a circular pattern around z -axis, Fig. 13(d) also follows a similar pattern with the same frequency with x and y components' curves.

E. More Qualitative Results Comparison on our Dataset

Here, we present qualitative results on all of the categories of our proposed dataset. Figure 14 to Figure 19 compares the output normal maps of our method with other baselines. Note that our implementation of Nayar *et al.* [49] uses Woodham's classical photometric stereo [69] to calculate the pseudo surface and updates the normals with the interreflection modeling for 15 iterations. Even though the Nayar *et al.* [49] interreflection algorithm is not theoretically guaranteed to converge for all surfaces, it gives a stable response on our dataset. We initialized Nayar's algorithm using the same predicted light sources of our method for a fair comparison.

The results show that our method achieves the best results overall, both qualitatively and quantitatively. We observed that other deep learning networks [12, 10] may fail to remove the surface ambiguity in challenging subjects. This

Figure 13: Light source directions obtained from the calibration sphere (a) using the light estimation network. We demonstrate the x, y and z components of the light direction vectors (b-d). The mean angular error between the ground-truth and estimated light directions is 6.31 degrees.

Ball	Cat	Pot1	Bear	Pot2	Buddha	Goblet	Reading
985	2808	3601	2585	2193	2787	1636	1723
Cow	Harvest	Vase	Golf-ball	Face	Tablet 1	Tablet 2	Broken Pot
1651	3582	1280	468	435	1610	437	1046

Table 5: Number of facets per subject used for our experiments.

is because these networks require supervised training with ground-truth normals, and their performance depends on the content of the training dataset. On the other hand, the results show that Nayar *et al.* [49] performs much better on challenging concave shapes. However, it cannot model specularities and cast shadows. On the other hand, our method can model these non-Lambertian effects with the reflectance mapping, and therefore, it performs better than Nayar *et al.* in all the tested categories.

Lastly, we provide the reflectance map obtained using our method on the proposed dataset. Figure 20 and Figure 21 show the reflectance map obtained using our method on the synthetic and real sequence respectively.

F. Some General Comments

Q1: Influence of complex texture on the light estimation. Indeed, surface texture is can be important for light estimation. However, the present benchmark datasets *i.e.*, DiLiGenT is composed of textureless subjects, and therefore, our focus was to perform surface reconstruction on textureless objects.

Q2: Nayar interreflection model vs. Monte Carlo: Monte Carlo method can provide more photo-realistic renderings. However, such an approach is again expensive, requires analytic BRDF models, and a sophisticated sampling strategy for computation, which can make the pipeline better, but more involved. So, we favored Nayar’s method and used reflectance maps to handle non-Lambertian effects.

Q3: Number of parameters for the normal estimation network and interreflection kernel computation: The inverse rendering network has ≈ 3.7 million parameters (12.3 MB). The interreflection kernel is generally sparse, and efficient software are available to handle large-sized sparse matrices. Table (5) provides the number of facets used for our experiments to calculate the interreflection kernel.

Figure 14: Qualitative comparison on the **Vase** scene. Here, it is obvious that previous deep learning based methods fail to handle the concavity of the subject. In contrast, our method works reasonably well showing the competence of our modeling procedure.

Figure 15: Qualitative comparison on the **Golf-ball** scene. Although deep learning based methods perform well smooth objects, they cannot handle fine structures and indentations.

G. Other Possible Future Extension

Our proposed method enables the application of photometric stereo on a broader range of objects. Yet, we think that there are possible future directions to extend it. Firstly, our method is generally a two-stage framework that utilizes a light estimation network and inverse rendering network in separate phases during inference. As an extension of our work, we aim to combine those stages in an end-to-end framework where light, surface normals, and reflectance values are estimated simultaneously. Secondly, our method uses a physical rendering equation for image reconstruction that is not sufficient for modeling all physical interactions between the object and the light. We believe that an improved rendering equation with additional physical constraints will allow better normal estimates. In addition to that, our method utilizes a specular-reflectance

Figure 16: Qualitative comparison on the **Face** scene. Although Nayar *et al.* [49] models interreflections, it cannot handle cast shadows. Therefore, it performs poorly on regions surrounding the eyes and the nose where cast shadows are effective. Here, we also observe that Chen *et al.* [12] can not estimate accurately for higher slant angles.

Figure 17: Qualitative comparison on the **Tablet1** scene. This subject has a complicated geometry involving cuneiform and reliefs. Apart from these fine structures, the object can be treated as a composite surface which has a large concavity in the middle part.

Figure 18: Qualitative comparison on the **Tablet2** scene. Similar to *Tablet1*, this subject also contains reliefs and cuneiform scripts. Since the overall geometry is approximately flat, all methods perform comparable on this category.

Figure 19: Qualitative comparison on the **Broken Pot** scene.

map inspired by the Phong reflectance model. Using other sophisticated variants of specular-reflectance map such as the Blinn-Phong reflection model [7] may further advance our approach. Finally, we observed that our method is very

Figure 20: Reflectance maps obtained with our method from **Vase**, **Golf-ball** and **Face** categories.

Figure 21: Reflectance maps obtained with our method from **Tablet1**, **Tablet2** and **Broken Pot** categories.

convenient for practical usage as it doesn't require ground-truth normals for supervised training. However, it could be possible to improve performance by utilizing training data in a similar framework.