

Open Archive TOULOUSE Archive Ouverte (OATAO)

OATAO is an open access repository that collects the work of Toulouse researchers and makes it freely available over the web where possible.

This is an author-deposited version published in : <http://oatao.univ-toulouse.fr/>
Eprints ID : 17235

The contribution was presented at BMVC 2016 :
<http://www.bmva.org/bmvc/2016/>

To cite this version : Logothetis, Fotios and Mecca, Roberto and QuØau, Yvain and Cipolla, Roberto *Near-Field Photometric Stereo in Ambient Light*. (2016) In: 27th British Machine Vision Conference (BMVC 2016), 19 September 2016 - 22 September 2016 (York, United Kingdom).

Any correspondence concerning this service should be sent to the repository administrator: staff-oatao@listes-diff.inp-toulouse.fr

Near-Field Photometric Stereo in Ambient Light

Fotios Logothetis¹

fl302@cam.ac.uk

Roberto Mecca^{1,2}

roberto.mecca@eng.cam.ac.uk

Yvain Quéau³

yvain.queau@enseeiht.fr

Roberto Cipolla¹

cipolla@eng.cam.ac.uk

¹ Department of Engineering

University of Cambridge

United Kingdom

² Department of Mathematics

University of Bologna

Italy

³ Université de Toulouse

France

Abstract

Shape recovery from shading information has recently regained importance due to the improvement towards making the Photometric Stereo technique more reliable in terms of appearance of reflective objects. However, although more advanced models have been lately proposed, 3D scanners based on this technology do not provide reliable reconstructions as long as the considered irradiance equation neglects any additive bias. Depending on the context, such bias assumes different physical meanings. For example, in murky water it is known as *saturated backscattered effect* or for acquisition in pure air medium it is known as *ambient light*. Although the theoretical part covers both cases, this work mostly focuses on the pure air acquisition case. Indeed, we present a new approach based on ratios of differences of images where an exhaustive set of physical features are tackled while dealing with Photometric Stereo acquisition with considerable importance for the ambient light. To the best of our knowledge, this is the first attempt to recover the shape from Photometric Stereo considering simultaneously perspective viewing geometry, non-linear light propagation, both specular and diffuse reflectance plus the additive bias of the ambient light. Proof of concept is provided by showing experimental results on synthetic and real data.

1 Introduction

Photometric Stereo (PS) can be seen as an evolution of the Shape from Shading (SfS) problem [12] where the ambiguity of retrieving the shape using shading information from a single image can be resolved by adding more images under different illumination conditions. There is an extensive literature on PS. Starting from the work by Woodham [32], PS has been studied for over three decades. However, until recently most approaches have relied on numerous simplifying assumptions such as orthographic viewing geometry, uniform directional lighting and the Lambertian model for the surface reflectance. In addition, the environment where the images are acquired is assumed dark, considering negligible ambient light.

On the contrary, in this work we introduce a PS method which simultaneously relaxes all the aforementioned assumptions. We deal with ambient light as an additional pixel-wise

component to the usual irradiance model (abusively referred to as BRDF in the following) depending on the light direction \mathbf{l} , the outgoing normal to the surface \mathbf{n} and the viewing direction \mathbf{v} as follows

$$I(x,y) = I_{\text{BRDF}}(\mathbf{l}, \mathbf{n}, \mathbf{v}) + A(x,y). \quad (1)$$

Our choice of tackling ambient light is motivated by two recent applications of PS: real time 3D-reconstruction [9], and PS in open environments (outdoor PS [14], or PS in a scattering medium such as murky water [22]). The usual approach to deal with ambient light consists of capturing an image without any artificial lights first, and then subtracting it from all the images. However, it is not possible to apply such a preprocessing in real time PS, where the shape information is extracted from a single RGB image of a commercial camera sensor [6, 9]. In addition, it is impossible to have a priori knowledge of the ambient light if it appears only when the artificial lights are on, as it is the case with back scattering in murky water [30, 31]. To deal with these issues, we introduce a framework for PS in ambient light which does not require any preprocessing step.

Contribution We propose a new approach for the PS which allows 3D reconstructions that have data with a complete list of physical features that are not negligible when images are taken in outdoor scenes. We extend the model presented in [19] which includes perspective viewing geometry and point light source parameterization for diffuse and specular reflectance, by considering non-negligible ambient light as in Equation (1). In Sec. 2 we briefly review relevant works in the same field while Sec. 3 describes the mathematical modeling of the proposed approach. The numerical solver is presented in Sec. 4 and lastly, Sec. 5 shows the experimental validation of the proposed approach over synthetic and real data.

2 Related Works

The Lambertian assumption assumed in most of the literature regarding PS is not realistic for a large number of materials as shown in [23]. In order for diffuse reflection to be a reliable model for any surface, Wu et al. [33] introduced an image preprocessing based on low-rank matrix factorization. On the other hand, Johnson et al. [17, 18] developed an elastomer that is attached to the surface to change its BRDF. Another attempt to deal with more complicated reflectance functions is by Alldrin and Kriegman [2], who overcome the standard reflectance assumptions by realizing that for an isotropic material the BRDF will experience bilateral symmetry. Hertzmann and Seitz [10] removed the need for an explicit BRDF model by comparing object images with images of reference objects. Another strategy involves keeping the Lambertian assumption, while estimating the normal field in a robust manner, for instance using sparsity-enhancing estimators [15]. This is possible as long as the number of images is high enough, which is a valid assumption in outdoor PS approaches considering the sun as a moving light source [1, 11, 14, 28]. Indeed, works considering PS in open environments acquired hundreds of images in order to have redundant information on the geometry of the object(s), thus getting around the lack of properly modeling environmental interference. In order to eventually obtain a depth map, all these approaches must be followed by an integration step [8]. This final step can be avoided by resorting to a differential approach, as described in the next paragraph.

With the aim to overcome these difficulties, PDEs-based approaches, derived by combining the differential PS formulation with image ratios, were recently applied to more complicated models as shown by Mecca et al. [20]. The method of the image ratios has also

been used by Jacobs et al. [16] to show that for any two images there is always an object and two lighting conditions consistent with these images. It should be noted that the use of ratios makes the approach very sensitive to noise. As a result Chandraker et al. [7] require a very specialised set-up and present results on simple data sets only. To ensure robustness, variational approaches aimed at solving a system of PDEs independent from the albedo were recently proposed [19, 27, 29].

Yet, all these PS approaches assume ambient light free environments. The concept of ambient light basically covers all the unexplained secondary reflections between the surface and the environment, which may be either due to stray lights, inter-reflections or back scattering, and are thus difficult or impossible to model in closed-form. Such deviations induce a systematic additive bias of the linear model which can be modeled as (Equation 1) where $A(x,y)$ represents the unexplained bias that could be dependent on the light source \mathbf{l} itself. A relatively simple way to account for this bias is by using a spherical harmonics model for the lighting [3]. Such a model was used by Basri et al. [4] to tackle the problem of unknown and non directional lighting for a Lambertian object. Using spherical harmonics decomposition, it was shown that the surface normals lie on a low-dimensional space spanned by the principal components of the image. In particular, a four harmonics model may describe rather well the combination of directional lighting and “ambient” lighting. Such a model was recently used by Or-El *et al.* [24] in order to apply SfS under “natural” illumination. Yet, such spherical harmonics models prevent one from using closed form expressions for non-directional lightings, *e.g.* the pointwise source model accounting for radial propagation of light and inverse of squared distance light attenuation.

Instead, in this work, we assume that the ambient light consists of a non-uniform scalar field, which is independent from \mathbf{l} . Although it cannot be justified physically that mutual reflections, or backscatter lightings, are independent from the lighting, this approximation is numerically tractable, and was already successfully used in [22, 31, 34]. In the literature, such an ambient term was accounted for using two different strategies: offline calibration or online estimation. The former strategy simply consists of capturing an image in the “dark” (*i.e.* with no active lights turned on). This “ambient image” can then be subtracted from the PS images to create “ambient-free” images.

The other strategy consists in estimating $A(x,y)$ along with the surface characteristics. The classic pixelwise PS estimation is simple to extend: the new recovery problem is another linear system of diffuse irradiance equations accounting $A(x,y)$ as unknown whose solution is uniquely defined as long as the number of images is at least 4 and the rank of the system is 4. The fourth component of the unknown vector is the ambient light, and its first three components represent the surface normals scaled by the albedo. This is essentially the same strategy as the approach of Yuille *et al.* in [34], although they extend this procedure to the *uncalibrated* PS problem. However, this approach relies on the restrictive assumptions of directional lighting as well as a Lambertian surface reflection, and it is expected to perform poorly if these assumptions are violated.

Hence, all existing methods to handle ambient light in PS lack flexibility. Thus, there is still a need for a technique that can be applied to more general scenarios. Our aim in this paper is to fill this gap providing an approach able to deal with complex non-linear phenomena occurring while considering open environment.

3 Dynamic Ambient Light Removal

PS is generally modeled considering restrictive assumptions in order to deal with a reasonable amount of non-linearity describing the image reflection model. Usually, physical based

BRDF models are very complicated to invert as required for shape reconstruction. With the aim of considering non-linearities reasonable to deal with, we use the irradiance equation introduced by Mecca and Quéau [19] which tackles most of the physical effects involved in the image formation model such as perspective viewing geometry, radial propagation, light attenuation and both diffuse and specular reflection by unifying the Lambertian shading model [13] with the Blinn-Phong one [5, 26]. In this work we propose a PS modeling with the following irradiance equation for the i^{th} light source placed at point (ξ_i, η_i, ζ_i) :

$$I_i(x, y) = \rho(x, y) a_i(x, y, z) (\bar{\mathbf{n}}(x, y) \cdot \bar{\mathbf{h}}_i(x, y, z))^{c(x, y)} + A(x, y) \quad (2)$$

where ρ is the albedo, c is a material property in $(0, 1]$ determining the shininess and a_i is the light attenuation parametrized as follows

$$a_i(x, y, z) = \frac{\phi_i a_{ai}(x, y, z)}{|\mathbf{l}_i(x, y, z)|^2} \quad (3)$$

where

$$a_{ai}(x, y, z) = (\bar{\mathbf{l}}_i(x, y, z) \cdot \bar{\mathbf{p}}_i)^\mu \quad (4)$$

takes inspiration from [20], $\bar{\mathbf{p}}_i$ is the principal lighting direction of the source (i.e. its orientation), ϕ_i is the illuminance of the source and

$$\mathbf{l}_i(x, y, z) = \mathbf{X}(x, y) - (\xi_i(x, y), \eta_i(x, y), \zeta_i(x, y)) \quad (5)$$

is the light direction for the surface point \mathbf{X} , assuming the light position (ξ_i, η_i, ζ_i) to be known. For the real world experiments, the parameter $\mu \geq 0$ is given by the datasheet of the LEDs used in the setup.

Using the notation $\bar{\cdot}$ for a general unit vector, we define the normal vector $\bar{\mathbf{n}} = \frac{\mathbf{n}}{|\mathbf{n}|}$ parametrized according the notation provided in [25]:

$$\mathbf{n}(x, y) = \left(\nabla z(x, y), -\frac{f+z(x, y)}{f} - \frac{(x, y) \cdot \nabla z(x, y)}{f} \right), \quad (6)$$

and the \mathbf{h}_i vector is defined as:

$$\mathbf{h}_i(x, y, z) = \bar{\mathbf{l}}_i(x, y, z) + \min \left\{ 1, \frac{|1 - c(x, y)|}{\varepsilon} \right\} \bar{\mathbf{v}}(x, y) \quad (7)$$

where ε is a material-related parameter (we took $\varepsilon = 0.01$ in our experiments directly following [19]). We consider the definition of the other quantities from [21] having the following viewer direction

$$\mathbf{v}(x, y) = (x, y, -f). \quad (8)$$

We remark that in our irradiance equation, the ambient light $A(x, y)$ is considered as a pixel-wise unknown of the problem. By assuming that it is independent from \mathbf{l}_i , it can be cancelled out by considering the ratios of image differences, as described hereafter. With the aim to consider a readable ratio of images, we avoid as much as possible to write the dependencies of the functions from now on and we preliminary manipulate the irradiance equation as follows:

$$(I_i - A)^c = (\rho a_i)^c \bar{\mathbf{n}} \cdot \bar{\mathbf{h}}_i. \quad (9)$$

Considering the Binomial expansion of the left hand side:

$$(I_i - A)^c = I_i^c - cAI_i^{c-1} + \frac{c(c-1)}{2}A^2I_i^{c-2} + \dots \quad (10)$$

and truncating this expansion to the first order, we obtain after some algebra:

$$I_i - cA \approx \rho^c \gamma_i \frac{\mathbf{n}}{|\mathbf{n}|} \cdot \bar{\mathbf{h}}_i \quad (11)$$

where $\gamma_i = \frac{a_i^c}{I_i^{c-1}}$. We remark that, in the pure diffuse case (i.e. $c = 1$), there is no need for Binomial expansion and truncation. This makes the model more accurate in the framework where most of the research about PS focuses on. From a practical point of view, as ambient light only makes up a small portion of the total reflected light, the error committed by considering a first order truncation is reasonably negligible. For example, if $\frac{I}{A} = 4$ and $c = 0.5$ then the third term of the expansion is 6.25% of the second. Now, we consider two pairs of irradiance equations, namely the i^{th} , j^{th} and the q^{th} , r^{th} in order to let the ambient light cancel out together with the albedo while taking ratio as follows:

$$\frac{I_i - cA - I_j + cA}{I_q - cA - I_r + cA} \approx \frac{\frac{\rho^c}{|\mathbf{n}|} [\gamma_i \mathbf{n} \cdot \bar{\mathbf{h}}_i - \gamma_j \mathbf{n} \cdot \bar{\mathbf{h}}_j]}{\frac{\rho^c}{|\mathbf{n}|} [\gamma_q \mathbf{n} \cdot \bar{\mathbf{h}}_q - \gamma_r \mathbf{n} \cdot \bar{\mathbf{h}}_r]} \quad (12)$$

that is

$$\frac{I_i - I_j}{I_q - I_r} \approx \frac{\gamma_i \mathbf{n} \cdot \bar{\mathbf{h}}_i - \gamma_j \mathbf{n} \cdot \bar{\mathbf{h}}_j}{\gamma_q \mathbf{n} \cdot \bar{\mathbf{h}}_q - \gamma_r \mathbf{n} \cdot \bar{\mathbf{h}}_r}. \quad (13)$$

By substituting the parameterization of the normal from Equation 6, Equation 13 yields the following quasi-linear PDE:

$$b_1(x, y, z) \frac{\partial z}{\partial x} + b_2(x, y, z) \frac{\partial z}{\partial y} = s(x, y, z) \quad (14)$$

with:

$$b_1 = (I_q - I_r) \left[\gamma_i \left(\bar{h}_i^{-1} - \frac{x}{f} \bar{h}_i^{-3} \right) - \gamma_j \left(\bar{h}_j^{-1} - \frac{x}{f} \bar{h}_j^{-3} \right) \right] - (I_i - I_j) \left[\gamma_q \left(\bar{h}_q^{-1} - \frac{x}{f} \bar{h}_q^{-3} \right) - \gamma_r \left(\bar{h}_r^{-1} - \frac{x}{f} \bar{h}_r^{-3} \right) \right], \quad (15)$$

$$b_2 = (I_q - I_r) \left[\gamma_i \left(\bar{h}_i^{-2} - \frac{y}{f} \bar{h}_i^{-3} \right) - \gamma_j \left(\bar{h}_j^{-2} - \frac{y}{f} \bar{h}_j^{-3} \right) \right] - (I_i - I_j) \left[\gamma_q \left(\bar{h}_q^{-2} - \frac{y}{f} \bar{h}_q^{-3} \right) - \gamma_r \left(\bar{h}_r^{-2} - \frac{y}{f} \bar{h}_r^{-3} \right) \right], \quad (16)$$

$$s = \frac{f+z}{f} \left[(I_q - I_r) \left(\gamma_i \bar{h}_i^3 - \gamma_j \bar{h}_j^3 \right) - (I_i - I_j) \left(\gamma_q \bar{h}_q^3 - \gamma_r \bar{h}_r^3 \right) \right]. \quad (17)$$

As a final note, it is possible to increase the accuracy of the truncation: as differences of pairs of images cancel out the first order remainder of the expansion, it is not hard to see how weighted quadruples of images can cancel up to the second order remainder of the expansion. Furthermore, the same approach can be applied to Taylor expansions of more physical BRDFs and this is left as a future work.

4 Variational resolution

In order to ensure robustness to noise (camera thermal noise, quantization) and to outliers (shadows, depth discontinuities, specularities not consistent with the Blinn-Phong model), we follow the variational approach from [19]. This approach solves the system of quasilinear PDEs (13) in an approximate manner, in the sense of the algebraic error (L^1 norm). Denoting

by $\mathbf{b}^d = [b_1^d, b_2^d]^\top$ and s^d the quantities in (15), (16) and (17) associated to the d -th ratio of differences, this variational problem reads:

$$\min_z \sum_d \left\| \mathbf{b}^d \cdot \nabla z - s^d \right\|_{L^1(\Omega)} + \mu \|z - z_0\|_{L^2(\Omega)} \quad (18)$$

where $\|\cdot\|_{L^1(\Omega)}$ is the L^1 norm over Ω , ∇z is the gradient of z , z_0 is some prior depth estimate and $\mu > 0$ controls the influence of this regularization term, which basically ensures numerical stability (in our experiments, we used $\mu = 10^{-6}$ and a constant value for z_0).

In order to apply the ADMM scheme proposed in [19], the b_1^d , b_2^d and s^d must be known, which is not the case when $c \neq 1$ or if the lighting is not directional. Hence, we embed the resolution of the variational problem (Equation (18)) into a fixed point algorithm, where these coefficients are iteratively refined. As described in [19], by embedding the ADMM scheme in a fixed point algorithm, the depth update comes down to solving a linear system of equations, which can be solved, e.g., by means of Gauss-Seidel iterations. We note that the variation problem is discretised with finite differences and that no boundary condition is required ([27]). The photometric stereo problem is summarized in Algorithm 1.

Algorithm 1: Ambient PS

Input : A sequence of images with light sources and rough mean distance z_0

Output: Depth map z , shininess parameter c

Initialize $z_k = z_0, c_k = 1$;

while $|z_{k+1} - z_k| > 10^{-4} \times |z_k|$ **do**

 calculate fields $a_k(z_k, c_k), \mathbf{h}_k(z_k, c_k)$;

 calculate fields $\mathbf{b}(z_k, a_k, \mathbf{h}_k), s(z_k, a_k, \mathbf{h}_k)$;

 Solve variational problem for z_{k+1} ;

 Re-calculate fields $a_k(z_{k+1}, c_k), \mathbf{h}_k(z_{k+1}, c_k)$;

 Estimate $c_{k+1}(z_{k+1}, a_k, \mathbf{h}_k)$

end

The update of the fields $a, \mathbf{h}, \mathbf{b}$ and s comes in a straightforward manner from their definition. The update of the shininess parameter c (final step in Algorithm 1) is less straightforward, and hence it is detailed hereafter. Starting again from the irradiance Equation (2) (before considering the Taylor expansion) and taking a ratio of the difference of equations i, j to q, r , we get:

$$\frac{I_i - I_j}{I_q - I_r} = \frac{a_i(\mathbf{n} \cdot \bar{\mathbf{h}}_i)^{\frac{1}{c}} - a_j(\mathbf{n} \cdot \bar{\mathbf{h}}_j)^{\frac{1}{c}}}{a_q(\mathbf{n} \cdot \bar{\mathbf{h}}_q)^{\frac{1}{c}} - a_r(\mathbf{n} \cdot \bar{\mathbf{h}}_r)^{\frac{1}{c}}}. \quad (19)$$

Expanding out Equation (19), we get:

$$(I_i - I_j)a_q(\bar{\mathbf{n}} \cdot \bar{\mathbf{h}}_i)^{\frac{1}{c}} - (I_i - I_j)a_r(\bar{\mathbf{n}} \cdot \bar{\mathbf{h}}_j)^{\frac{1}{c}} - (I_r - I_q)a_i(\bar{\mathbf{n}} \cdot \bar{\mathbf{h}}_q)^{\frac{1}{c}} + (I_r - I_q)a_j(\bar{\mathbf{n}} \cdot \bar{\mathbf{h}}_r)^{\frac{1}{c}} = 0 \quad (20)$$

which is of the form $\sum_{n=1}^4 D_n M_n^x = 0$, with D_n, M_n known and $x = \frac{1}{c}$. Equation 20 cannot be solved analytically; it can, however, be solved numerically with the Newton-Raphson method. In addition, too small (< 0.1) or too large (> 1) values of c are treated as outliers and are not updated ($c_{k+1} = c_k$). Finally, the new c map is convolved with a Gaussian filter of size 21 pixels and standard deviation $\sigma = 8$ pixels (those values were found experimentally to provide the best results). This step, which reduces the effect of outliers, is consistent with the assumption that the material distribution is smooth.

5 Experimental Part

We evaluated our algorithm on both synthetic and real data. Synthetic data were made with OpenGL directly implementing Equation 2. The resolution of all the synthetic data was set to 800x600 pixels in order to match our real camera. For synthetic model, we took "bimba" from the AIM@Shape Repository.

We compare our method against Yuille *et al.* [34], Ikehata *et al.* [15] and Mecca and Quéau [19]. For [15, 19], we used their original code whereas for [34] we used our implementation. In addition, we used our variational solver (see Section 4) in order to numerically integrate the normal maps produced by [15, 34].

5.1 Synthetic Data

First of all, we generated synthetic data (Figure 1) with a simplified image formation model including orthographic viewing geometry, directional illumination and Lambertian reflection. This was done in order to have a fair comparison with Yuille *et al.* [34] and Ikehata *et al.* [15]. We used the "Lena" image for the albedo. The ambient light was set to grow linearly from bottom-left to up-right to up to 45% of the maximum intensity value.




Figure 1: 2 out of the 24 images of our simplified synthetic dataset rendered with orthographic geometry, directional lighting and Lambertian reflection.

The reconstructions are shown in Figure 2. Quantitative evaluation is performed by comparing the generated normal maps with the ground truth (Figure 3). For [19] and our method the normals are calculated using finite differences of the depth map.

Our approach clearly outperforms [15, 19] as these methods suffer from the additive bias of the ambient light which has a flattening effect on the reconstruction. We also have a slightly smaller mean normal error than [34] (8.5° vs 11°) and this is probably because of the use of the robust variational solver. In addition, the error is very slightly increasing as a function of c (Figure 4) as expected from the truncation of the Binomial series.

Finally the robustness of our method to more realistic effects such as perspective viewing geometry and near point source as well as Gaussian noise is shown in Figure 5.

5.2 Real Data

We acquired real data with two different setups. The first setup consists of a FL3-U3-20E4C camera of Point Grey Research placed in the center of a 5cm base containing SHARP LED modules, MINIZENIGATA series. We used 24 LEDs arranged in 2 concentric rings of radii 3cm and 5cm respectively. This setup was used to obtain a marble Buddha statue (6(a)) and a shiny plastic head (6(b)). A second setup, consisting in a series of 8 LEDs and a Canon EOS 7D camera, was used to acquire images of a plaster Arlequin mask (6(c)) and of print of teeth (6(d)). For both setups, the camera intrinsics were calibrated using Matlab's computer vision toolbox, the position of the LEDs was calibrated using a mirror sphere, and their intensities were calibrated using a diffuse sphere. The anisotropy parameter was set to $\mu = 1$, consistently with the manufacturer's specification.





Figure 2: Comparison of several PS methods on the dataset of Figure 1. Note that 2(b) and 2(c) have a substantial flattening due to the ambient light.



(a) [34], Mean error:11° (b) [15], Mean error:16.7° (c) [19], Mean error:17.2° (d) Ours, Mean error:8.5°

Figure 3: Quantitative evaluation of several PS methods on the dataset of Figure 1. Most errors are observed around the occlusion boundaries as expected.

The algorithm was implemented in Matlab and was tested on a i7 processor at 2.4GHz, with 16GB RAM. As selecting all possible quadruples of images is computationally intractable ($\binom{24}{4} = 10626$), for the first dataset we only considered differences of diametrically opposite images. This maximizes photometric parallax and reduces the number of quadruples to 32 (8 pairs from outer ring getting divided with 4 pairs from the inner ring). The CPU time on our datasets was a few minutes. To create challenging conditions, the relevant amount of ambient light was varied between the different datasets, with the Buddha being illuminated more than the Mask and the Teeth. Finally, the shiny head was placed next to an open window (the reflection of the window is clearly visible and it is creating some artifacts in the reconstructions at the second row of Figure 7) which makes the ambient light uncontrolled.

Acknowledgments

Roberto Mecca was supported through a Marie Curie fellowship of the “Istituto Nazionale di Alta Matematica”, Italy.




Figure 4: Comparison of several PS methods under different ambient light and shininess parameter. At high ambient levels, a considerable portion of the pixels get saturated, and this effect dominates the error.




Figure 5: Synthetic data with realistic assumptions including perspective deformation, near light sources, and moderate specularity ($c=0.4$); as well as 2% additive Gaussian noise.




Figure 6: Two samples from each object and the respective ambient light in the last row.




Figure 7: Evaluation of PS methods on the datasets of Figure 5.2

6 Conclusion and Further Work

In this work we tackled the problem of PS under ambient light as well as an extensive set of additional realistic assumption (perspective view geometry, non-linear light propagation, specular reflection). A new approach based on ratios of image differences was presented that is able to remove any additive bias on images. The problem is then expressed as a quasi-linear PDE and is solved through a robust variational optimizer performing L^1 minimization. Experiments on synthetic and real data verify that our approach achieves good reconstructions under significant ambient light, specular highlights and perspective deformation.

As a challenging future work, we foresee the applicability of our technique involving Binomial expansion, from the BRDF presented in [19] to other more physical and complicated BRDFs.

Finally, as most of the calculations are performed on a per-pixel basis, the implementation can easily be ported to a GPU allowing for a real time acquisition. This is coherently matching with the proposed ambient light removal approach towards making the PS a reliable technique capable of real time shape reconstruction in open environment.

References

- [1] A. Abrams, C. Hawley, and R. Pless. Heliometric stereo: Shape from sun position. In *ECCV*, 2012.
- [2] N. G. Aldrin and D. J. Kriegman. Toward Reconstructing Surfaces With Arbitrary Isotropic Reflectance : A Stratified Photometric Stereo Approach. In *ICCV*, 2007.
- [3] R. Basri and D. W. Jacobs. Lambertian reflectance and linear subspaces. In *ICCV*, 2001.
- [4] R. Basri, D. Jacobs, and I. Kemelmacher. Photometric Stereo with General, Unknown Lighting. *IJCV*, (3):239–257, 2007.
- [5] J. F. Blinn. Models of light reflection for computer synthesized pictures. In *SIGGRAPH*, 1977.
- [6] G. J. Brostow, C. Hernández, G. Vogiatzis, B. Stenger, and R. Cipolla. Video normals from colored lights. *PAMI*, 33(10):2104–2114, 2011.
- [7] M. Chandraker, J. Bai, and R. Ramamoorthi. On differential photometric reconstruction for unknown, isotropic brdfs. *PAMI*, 35(12):2941–2955, 2013.
- [8] Durou J. D., Aujol J.-F., and Courteille F. Integrating the normal field of a surface in the presence of discontinuities. In *EMMVCVPR*, 2009.
- [9] C. Hernández, G. Vogiatzis, G. J. Brostow, B. Stenger, and R. Cipolla. Non-rigid photometric stereo with colored lights. In *ICCV*, 2007.
- [10] A. Hertzmann and S.M. Seitz. Example-based photometric stereo: shape reconstruction with general, varying brdfs. *PAMI*, 27(8):1254–1264, 2005.
- [11] Y. Hold-Geoffroy, J. Zhang, P. F. U. Gotardo, and J.-F. Lalonde. X-hour Outdoor Photometric Stereo. In *3DV*, 2015.
- [12] B. K. P. Horn. *Shape from Shading: a Method for Obtaining the Shape of a Smooth Opaque Object from One View*. PhD thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1970.
- [13] B. K. P. Horn. Obtaining shape from shading information. *The Psychology of Computer Vision*, pages 115–155, 1975.
- [14] C.-H. Hung, T.-P. Wu, Y. Matsushita, L. Xu, J. Jia, and C.-K. Tang. Photometric Stereo in the Wild. In *WACV*, 2015.
- [15] S. Ikehata, D. Wipf, Y. Matsushita, and K. Aizawa. Robust photometric stereo using sparse regression. In *CVPR*, 2012.
- [16] D. W. Jacobs, P. N. Belhumeur, and R. Basri. Comparing images under variable illumination. In *CVPR*, 1998.
- [17] M. K. Johnson and E. H. Adelson. Retrographic sensing for the measurement of surface texture and shape. In *CVPR*, 2009.

- [18] M. K. Johnson, F. Cole, A. Raj, and E. H. Adelson. Microgeometry capture using an elastomeric sensor. In *SIGGRAPH*, 2011.
- [19] R. Mecca and Y. Quéau. Unifying diffuse and specular reflections for the photometric stereo problem. In *WACV*, 2016.
- [20] R. Mecca, A. Wetzler, A. Bruckstein, and R. Kimmel. Near field photometric stereo with point light sources. *SIAM J. Imag. Sci.*, 7(4):2732–2770, 2014.
- [21] R. Mecca, E. Rodolà, and D. Cremers. Realistic photometric stereo using partial differential irradiance equation ratios. *Computers & Graphics*, 51:8–16, 2015.
- [22] Z. Murez, T. Treibitz, R. Ramamoorthi, and D. Kriegman. Photometric Stereo in a Scattering Medium. In *ICCV*, 2015.
- [23] A. Ngan, F. Durand, and W. Matusik. Experimental analysis of brdf models. In *EGSR*, 2005.
- [24] R. Or-el, G. Rosman, A. Wetzler, R. Kimmel, and A. M. Bruckstein. RGBD-Fusion: Real-Time High Precision Depth Recovery. In *CVPR*, 2015.
- [25] T. Papadimitri and P. Favaro. A new perspective on uncalibrated photometric stereo. In *CVPR*, 2013.
- [26] B. T. Phong. Illumination for computer generated pictures. *Commun. ACM*, 18(6):311–317, 1975.
- [27] Y. Quéau, R. Mecca, and J. D. Durou. Unbiased photometric stereo for colored surfaces: A variational approach. In *CVPR*, 2016.
- [28] F. Shen, K. Sunkavalli, N. Bonneel, S. Rusinkiewicz, H. Pfister, and X. Tong. Time-Lapse Photometric Stereo and Applications. *Computer Graphics Forum*, 33(7):359–367, 2014.
- [29] W. Smith and F. Fang. Height from photometric ratio with model-based light source selection. *CVIU*, 145:128 – 138, 2016.
- [30] T. Treibitz and Y. Y. Schechner. Active polarization descattering. *PAMI*, 31(3):385–399, 2009.
- [31] C. Tsiotsios, M. E. Angelopoulou, T.-K. Kim, and A. J. Davison. Backscatter compensated photometric stereo with 3 sources. In *CVPR*, 2014.
- [32] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):134–144, 1980.
- [33] L. Wu, A. Ganesh, B. Shi, Y. Matsushita, Y. Wang, and Y. Ma. Robust photometric stereo via low-rank matrix completion and recovery. In *ACCV*, 2010.
- [34] A. L. Yuille, D. Snow, R. Epstein, and P. N. Belhumeur. Determining generative models of objects under varying illumination: Shape and albedo from multiple images using SVD and integrability. *IJCV*, 35(3):203–222, 1999.