



Some Illumination Models for Industrial Applications of Photometric Stereo

Yvain Quéau, Jean-Denis Durou

► To cite this version:

Yvain Quéau, Jean-Denis Durou. Some Illumination Models for Industrial Applications of Photometric Stereo. The International Conference on Quality Control by Artificial Vision 2015, Jun 2015, Le Creusot, France. pp.95341C, 10.1117/12.2182921 . hal-02118579

HAL Id: hal-02118579

<https://normandie-univ.hal.science/hal-02118579>

Submitted on 3 May 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Some Illumination Models for Industrial Applications of Photometric Stereo

Yvain Quéau and Jean-Denis Durou

Université de Toulouse, IRIT, UMR CNRS 5505, Toulouse, France

ABSTRACT

Among the possible sources of error in 3D-reconstruction using the photometric stereo technique, lighting modelling is often neglected, though it can create a dramatic large-scale bias. In this paper, after recalling the physical definition of a primary Lambertian source (isotropic lightings), we show how to derive a lighting model for several real-world scenarios, including directional lightings, nearby sources and extended planar illuminants. Finally, we show how to calibrate general spatially-varying lightings within a plane, in the case where explicitly modelling the lightings would be tedious.

Keywords: Photometric stereo, lighting, illumination models, nearby punctual sources, extended planar illuminants, light field sampling.

1. INTRODUCTION

Photometric stereo,¹ referred to as PS in the following, is a monocular 3D-reconstruction technique where orientation, depth and albedo clues are obtained by taking m pictures $I^1 \dots I^m$ of a scene under m different illuminations $\mathbf{s}^1 \dots \mathbf{s}^m$. This problem is usually formulated as an inverse problem, which requires an explicit model for the illuminations. While the BRDF models for the surface to reconstruct have attracted a lot of studies, the lighting models are often neglected. Yet, the choice of an appropriate lighting model is of major importance in PS, in order to avoid bias in the 3D-reconstructions.

Assuming Lambertian BRDF and neglecting shadows, the graylevel at pixel p is given by:

$$I(p) = \rho(\mathbf{x}_p) \mathbf{n}(\mathbf{x}_p) \cdot \mathbf{s}(\mathbf{x}_p) \quad (1)$$

where $\mathbf{x}_p = \pi^{-1}(p) \in \mathbb{R}^3$ is the point on the surface associated to pixel p through (orthographic or perspective) projection π , $\rho(\mathbf{x}_p)$ is the albedo at $\mathbf{x}_p$, $\mathbf{n}(\mathbf{x}_p)$ is the unit outward normal to the surface at $\mathbf{x}_p$, and $\mathbf{s}(\mathbf{x}_p)$ represents the total illumination vector at $\mathbf{x}_p$, in intensity and direction, oriented towards the light source. Modelling adequately $\mathbf{s}$, when this vector represents a primary lighting, is the target of the present study. Secondary reflections from the surface onto the surface are neglected in this work, though they can be recovered implicitly by refining the illumination vector $\mathbf{s}$ during the reconstruction process.²

Though Lambert's law is a simple model, which may seem quite unrealistic for real-world applications of PS, it is actually a rather realistic one for modelling the luminance of an extended primary source, since manufacturers of lighting systems (LEDs, LCD screens etc.) usually guarantee isotropy, as much as possible. After recalling the definition of a Lambertian primary source, we provide in the following three examples of such illuminants, and derive the associated image graylevel models. We eventually show that in the presence of more complex lightings, it is still possible to consider arbitrary illumination fields, using an appropriate calibration. These situations are summarised in Figure 1.

Correspondence:

Y. Quéau: yvain.queau@enseiht.fr

J.-D. Durou: durou@irit.fr




Figure 1. Overview of the illumination configurations considered in this paper. (a) In the isotropic punctual model (Section 4), the lighting direction in $\mathbf{x}_p$ depends on $\mathbf{x}_p$, which is precisely the unknown, and the lighting intensity is a decreasing function of $\|\mathbf{x}_S - \mathbf{x}_p\|$. Such a model is useful when dealing with LEDs. If the source is infinitely far away, a good approximation of this model is the directional model (Section 3). (b) In the extended planar anisotropic model (Section 5), several anisotropic sources, with principal lighting direction $\mathbf{n}_0$, are used together. This is a good model for LCD screens. (c) In the most general case (Section 6), it may become tedious to derive an explicit illumination model. Assuming the surface to reconstruct lies in the vicinity of a plane $\mathcal{P}$, it is easier to a priori sample the light field on $\mathcal{P}$.

2. LAMBERTIAN PRIMARY SOURCES

Before describing several explicit illumination models used in various PS applications, we recall hereafter the definition of a primary Lambertian light source,³ which is a realistic model for most real-world applications of PS.

Assume an elementary primary light source $d\mathcal{S}$ is located at $\mathbf{x}_S \in \mathbb{R}^3$, with infinitesimal surface element $d\Sigma_S$ and normal $\mathbf{n}(\mathbf{x}_S)$. The *emitted luminance* of the source in the direction $\mathbf{u}_e$ ($\|\mathbf{u}_e\| = 1$) is given by:

$$L_{\mathbf{x}_S}(\mathbf{u}_e) = \frac{d^2\Phi}{d\Sigma_S(\mathbf{n}(\mathbf{x}_S) \cdot \mathbf{u}_e) d\Omega} \quad (2)$$

where $d^2\Phi$ denotes the luminous flux emitted by the source inside an elementary solid angle $d\Omega$ around the direction $\mathbf{u}_e$, and $d\Sigma_S(\mathbf{n}(\mathbf{x}_S) \cdot \mathbf{u}_e)$ is equal to the apparent surface of the source seen from that direction. In the context of PS, $\mathbf{u}_e$ points towards an elementary scene surface $d\Sigma$ with normal $\mathbf{n}(\mathbf{x}_p)$, around the point $\mathbf{x}_p \in \mathbb{R}^3$. The solid angle $d\Omega$ is thus defined as:

$$d\Omega = \frac{d\Sigma(-\mathbf{u}_e \cdot \mathbf{n}(\mathbf{x}_p))}{\|\mathbf{x}_S - \mathbf{x}_p\|^2} \quad (3)$$

Hence, the *irradiance* at $\mathbf{x}_p$ due to $d\mathcal{S}$ reads, using (2) and (3):

$$dI = \frac{d^2\Phi}{d\Sigma} = L_{\mathbf{x}_S}(\mathbf{u}_e) \frac{d\Sigma_S(\mathbf{n}(\mathbf{x}_S) \cdot \mathbf{u}_e)(-\mathbf{u}_e \cdot \mathbf{n}(\mathbf{x}_p))}{\|\mathbf{x}_S - \mathbf{x}_p\|^2} \quad (4)$$

Considering $d\mathcal{S}$ as a *Lambertian primary source* means that $L_{\mathbf{x}_S}(\mathbf{u}_e)$ is assumed to be independent from $\mathbf{u}_e$. We will thus denote $L_{\mathbf{x}_S}(\mathbf{u}_e) = L(\mathbf{x}_S)$. This hypothesis is usually rather realistic, since, as stated earlier, manufacturers of lighting systems tend to guarantee this property of the emitted luminance, so as to improve users' comfort.

Introducing the vector $d\mathbf{s}_{\mathbf{x}_S}(\mathbf{x}_p)$ such that $dI = d\mathbf{s}_{\mathbf{x}_S}(\mathbf{x}_p) \cdot \mathbf{n}(\mathbf{x}_p)$, whose norm is called the *luminous flux density*, we obtain from (4):

$$d\mathbf{s}_{\mathbf{x}_S}(\mathbf{x}_p) = \frac{L(\mathbf{x}_S)d\Sigma_S}{\|\mathbf{x}_S - \mathbf{x}_p\|^2} (\mathbf{n}(\mathbf{x}_S) \cdot \mathbf{u}_e)(-\mathbf{u}_e) \quad (5)$$

In (5), the first factor represents the inverse of square falloff, the second represents the anisotropy of the primary source, and the third is the lighting direction. Let us emphasize that, though the luminance is assumed to be isotropic, the source itself is allowed to be anisotropic, since its apparent surface usually varies.

Since there is no interference, the light flows additively, and thus a general expression for a light vector $\mathbf{s}(\mathbf{x}_p)$ modelling a Lambertian primary source can be obtained by integrating $d\mathbf{s}_{\mathbf{x}_S}(\mathbf{x}_p)$ over all the elementary sources $d\mathbf{S}$ which illuminate $\mathbf{x}_p$. This provides a natural framework for modelling any real-world situation, provided that the elementary sources can be explicitly modelled, which is the case in many industrial applications of PS, as for instance when dealing with distant illuminants (Section 3), nearby LEDs (Section 4) or LCD screens (Section 5). When modelling the elementary sources becomes difficult or impossible, it is better to avoid considering an explicit model for the lightings. In this case, we show how to sample the incident lighting onto a surface when this surface is well approximated by a plane $\mathcal{P}$ (Section 6).

In the following, we describe how to obtain a closed-form expression of the lightings in each of these cases. To illustrate how to use such models in industrial PS applications, we also provide some clues on how to invert the resulting image graylevel model deriving from (1), in order to simultaneously recover the albedo, the normals and the 3D-points associated to the surface.

3. UNCALIBRATED PS: DIRECTIONAL LIGHTING



Figure 2. 3D-reconstruction using uncalibrated PS. Left: the $m = 3$ input images, with unknown illuminations. Right: 3D-reconstruction, with estimated albedo warped onto the reconstructed surface.

If no information about the illumination is available, it is usually assumed that the lightings are directional. This hypothesis is for example systematically introduced in *uncalibrated* PS (unknown illuminations), so as to simplify the image graylevel model and to allow efficient resolution through matrix factorisation techniques.⁴

It is assumed in this case that in each image, a single source is emitting light from an infinitely far away unknown location $\mathbf{x}_S$ (extension to several distant illuminants, using spherical harmonic decomposition, was presented in Ref. 5). Then, $\mathbf{u}_e$ and $\mathbf{x}_S - \mathbf{x}_p$ in (5) can be considered to be the same for all the surface points $\mathbf{x}_p$. The elementary light vector $d\mathbf{s}_{\mathbf{x}_S}(\mathbf{x}_p)$ is thus independent from $\mathbf{x}_p$, which implies:

$$\mathbf{s}(\mathbf{x}_p) = \mathbf{s} \quad (6)$$

The m lightings $\mathbf{s}^i$ being assumed to be uniform, the mn equations obtained by writing Eq. (1) for each of the m illuminations and each of the n pixels can be re-arranged as a linear system:

$$\underbrace{\begin{bmatrix} I^1(\mathbf{x}_1) & \dots & I^m(\mathbf{x}_1) \\ \vdots & & \vdots \\ I^1(\mathbf{x}_n) & \dots & I^m(\mathbf{x}_n) \end{bmatrix}}_{\mathbf{I}} = \underbrace{\begin{bmatrix} \rho(\mathbf{x}_1)\mathbf{n}(\mathbf{x}_1)^\top \\ \vdots \\ \rho(\mathbf{x}_n)\mathbf{n}(\mathbf{x}_n)^\top \end{bmatrix}}_{\mathbf{M}} \underbrace{\begin{bmatrix} \mathbf{s}^1 & \dots & \mathbf{s}^m \end{bmatrix}}_{\mathbf{S}} \quad (7)$$

In the uncalibrated photometric stereo context, both $\mathbf{M}$ and $\mathbf{S}$ are unknown. Knowing that $\mathbf{I}$ should be of rank 3, these matrices can be simultaneously estimated by third-order truncated singular value decomposition (SVD). Yet, such a decomposition is ambiguous: $\mathbf{MS} = (\mathbf{MA})(\mathbf{A}^{-1}\mathbf{S})$ for any invertible 3×3 matrix $\mathbf{A}$. This linear ambiguity can be reduced to a *generalized bas-relief* ambiguity through discrete integrability enforcement.⁶ This final ambiguity is a 3-parameters ambiguity, which can be solved for instance through the introduction of a prior on the albedo,⁷ on the presence of diffuse maxima⁸ or on surface total variation.⁹ Eventually, the 3D-points can be estimated by integrating the normals.¹⁰ An example of 3D-reconstruction, using the disambiguation algorithm from Ref. 9, is shown in Figure 2.

4. LED-ILLUMINATED PS: ISOTROPIC PUNCTUAL SOURCES




Figure 3. 3D-reconstruction of a polystyrene dummy using PS with isotropic punctual sources. Left: 3 of the $m = 10$ input images. In each image, a single nearby LED with known position and intensity is used to illuminate the scene. Right: 3D-reconstruction.

The directional assumption above is not realistic when dealing with nearby sources, such as LEDs at short distance. A punctual model (Figure 1-a), with known position $\mathbf{x}_S$ (extension to unknown $\mathbf{x}_S$ was recently proposed in Ref. 11), is more appropriate than a directional model. In this case, the light direction $\mathbf{u}_e$ writes:

$$\mathbf{u}_e = \frac{\mathbf{x}_p - \mathbf{x}_S}{\|\mathbf{x}_p - \mathbf{x}_S\|} \quad (8)$$

Assume the primary source consists in a uniform Lambertian sphere such that $L(\mathbf{x}_S) = L_0$. If the sphere is small, for a given $\mathbf{x}_p$, $\mathbf{x}_S - \mathbf{x}_p$ and $\mathbf{u}_e$ are the same for all points $\mathbf{x}_S$, and by integrating over the hemisphere visible from $\mathbf{x}_p$:

$$\mathbf{s}(\mathbf{x}_p) = -\frac{\mathbf{u}_e L_0}{\|\mathbf{x}_S - \mathbf{x}_p\|^2} \iint_{\text{hemisphere}} (\mathbf{n}(\mathbf{x}_S) \cdot \mathbf{u}_e) d\Sigma_S \quad (9)$$

which we can write, using (8):

$$\mathbf{s}(\mathbf{x}_p) = s_0 \frac{\mathbf{x}_S - \mathbf{x}_p}{\|\mathbf{x}_S - \mathbf{x}_p\|^3} \quad (10)$$

with:

$$s_0 = L_0 \iint_{\text{hemisphere}} (\mathbf{n}(\mathbf{x}_S) \cdot \mathbf{u}_e) d\Sigma_S = \pi R^2 L_0 \quad (11)$$

where R is the radius of the sphere. We finally obtain:

$$I(p) = s_0 \rho(\mathbf{x}_p) \mathbf{n}(\mathbf{x}_p) \cdot \frac{\mathbf{x}_S - \mathbf{x}_p}{\|\mathbf{x}_S - \mathbf{x}_p\|^3} \quad (12)$$

where the quantities ρ , $\mathbf{n}$ and $\mathbf{x}_p$ are all unknown. Since the image graylevel model (12) explicitly depends on $\mathbf{x}_p$, unlike in the case of Section 3, it requires a careful minimisation.

Assuming the dependency of $\mathbf{n}$ in $\mathbf{x}_p$ is explicitly known (*i.e.*, the camera is calibrated), the authors of Ref. 12 propose to recover directly $\mathbf{x}_p$ (the albedo and the normals are eliminated from the set of unknowns) using a PDE-based method which propagates information from the boundary (in this case, the value of the depth on this boundary is assumed to be known) according to a semi-Lagrangian scheme.

Another possibility, discussed in Ref. 11, is to alternate estimations of $\mathbf{m} = \rho \mathbf{n}$ with $\mathbf{x}_p$ fixed, and estimations of $\mathbf{x}_p$ with $\mathbf{m}$ known, through normal field integration. The first update is a simple linear least-squares problem:

$$\mathbf{m}^{k+1}(\mathbf{x}_p^k) = \underset{\mathbf{m}}{\operatorname{argmin}} \sum_i \left(I^i(p) - s_0^i \mathbf{m} \cdot \frac{\mathbf{x}_S^i - \mathbf{x}_p^k}{\|\mathbf{x}_S^i - \mathbf{x}_p^k\|^3} \right)^2 \quad (13)$$

Assuming orthographic camera, the second update reads:

$$\begin{cases} \mathbf{x}_p^{k+1} = [p, z^{k+1}(p)]^\top \\ \text{s.t. } z^{k+1} = \underset{z}{\operatorname{argmin}} \left\| \nabla z - [-\mathbf{m}_1^{k+1}(\mathbf{x}_p^k)/\mathbf{m}_3^{k+1}(\mathbf{x}_p^k), -\mathbf{m}_2^{k+1}(\mathbf{x}_p^k)/\mathbf{m}_3^{k+1}(\mathbf{x}_p^k)]^\top \right\|_{L^2}^2 \end{cases} \quad (14)$$

which can be solved using numerous approaches.¹⁰ Eventually, one may recover the albedo and the normals by $\rho = \|\mathbf{m}\|$ and $\mathbf{n} = \mathbf{m}/\|\mathbf{m}\|$. Unlike in the differential approach of PS,¹² extension of such a two-steps approach to perspective projection is trivial, since this only changes the datum in the normal integration problem, as discussed for example in Ref. 10. To our knowledge, there is no theoretical proof of convergence for this scheme (unlike in Ref. 12), but usually a few iterations suffice to reach acceptable results.¹¹ In Figure 3, we show some images of a dummy, illuminated by nearby LEDs, as well as a 3D-reconstruction obtained after 10 iterations of the alternating optimisation scheme.

5. LCD-SCREEN ILLUMINANTS: EXTENDED PLANAR ANISOTROPIC SOURCES



Figure 4. 3D-reconstruction using PS with extended anisotropic light sources. Left: 3 of the $m = 4$ input images, obtained by displaying white rectangular patterns on the screen of a laptop, and capturing the images using the laptop's integrated webcam. Right: 3D-reconstruction (the eyes were manually removed from the reconstruction domain). In this experiment, we used the RGB values instead of the graylevels, which allowed us to recover RGB albedos, as explained in Ref. 13.

Another interesting PS application is 3D-reconstruction using a LCD screen as light source, as proposed for instance in Ref. 14. The elementary source $d\mathcal{S}$ in this case is planar, with normal $\mathbf{n}_0$ (see Figure 1-b). From (5) and (8), we obtain:

$$ds_{\mathbf{x}_S}(\mathbf{x}_p) = L(\mathbf{x}_S) d\Sigma_S \frac{\mathbf{n}_0 \cdot (\mathbf{x}_p - \mathbf{x}_S)}{\|\mathbf{x}_S - \mathbf{x}_p\|^4} (\mathbf{x}_S - \mathbf{x}_p) \quad (15)$$

where L is expected to be proportional to the color level (in each channel) of the image displayed on the screen, and $\mathbf{n}_0 \cdot (\mathbf{x}_p - \mathbf{x}_S)$ corresponds to a cosine-like anisotropy factor, which was experimentally validated in Ref. 15, and is also often used in punctual models.¹² This term is usually introduced without physical motivation: in fact, it is directly derived from the application of Lambert's law to the primary illuminant.

By integrating (15) over a part $\mathcal{S}$ of the screen:

$$\mathbf{s}(\mathbf{x}_p) = \iint_{\mathcal{S}} L(\mathbf{x}_S) \frac{\mathbf{n}_0 \cdot (\mathbf{x}_p - \mathbf{x}_S)}{\|\mathbf{x}_S - \mathbf{x}_p\|^4} (\mathbf{x}_S - \mathbf{x}_p) d\Sigma_S \quad (16)$$

and finally:

$$I(p) = \rho(\mathbf{x}_p) \mathbf{n}(\mathbf{x}_p) \cdot \iint_{\mathcal{S}} L(\mathbf{x}_S) \frac{\mathbf{n}_0 \cdot (\mathbf{x}_p - \mathbf{x}_S)}{\|\mathbf{x}_S - \mathbf{x}_p\|^4} (\mathbf{x}_S - \mathbf{x}_p) d\Sigma_S \quad (17)$$

Remark that the integral can be either approached by considering a discrete sum of the contributions of each pixel of the screen, or exactly computed if a simple model for the distribution of the luminance is available.¹⁴

The albedo, the normals and the 3D-points can be eventually recovered using the same alternating optimisation strategy as described in Section 4: by fixing the value of the integral in (17), we get a linear system in $\mathbf{m} = \rho \mathbf{n}$, while the 3D-points $\mathbf{x}_p$ can be updated by integrating the normals. A 3D-reconstruction result, using this strategy, is shown in Figure 4.

6. PS WITH ARBITRARY PRIMARY ILLUMINANTS: DENSE LIGHT FLOW SAMPLING ON A PLANE

In all the examples above, an explicit model is derived from a priori knowledge of the illumination conditions. In some applications, this prior is not available. For instance, as shown in Figure 1-c, the light reaching a point on the scene surface may be composed of a primary lighting coming directly from the source and of secondary lightings due to reflections of the primary light beams on some reflection device (secondary reflections from the surface onto the surface are neglected here), for example if LEDs are located in front of a reflecting system: using an explicit lighting model would be quasi impossible in such a complex case.

Yet, it is always possible to consider that the lighting is an arbitrary vector field $\mathbf{s} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$, and to try to sample the values of this vector field during the calibration process. This is particularly adapted to the case of almost planar objects: in this case the light variations in the z -direction can be neglected, because the surface to be reconstructed lies in vicinity of a plane $\mathcal{P}$ of equation $z = \text{constant}$. In this case, the values of the light field only depends on the 2D position p : $\mathbf{s}(\mathbf{x}_p) = \mathbf{s}(p)$. As a consequence, direct recovery of the normals and the albedo can be obtained without using an iterative scheme, provided the values of $\mathbf{s}(p)$ are a priori measured.

This is exactly the approach followed in Ref. 16, where images of a flat reference object with known albedo are used to compensate for the non-directionality of the lightings, by a priori dividing the images by those of the reference object (a similar approach is considered in Ref. 17 in the uncalibrated context). The underlying idea in such algorithms is to consider the light direction is uniform, but light intensity is not. By a priori capturing these spatially-varying intensities and appropriately compensating the input images, the reconstruction problem is turned into a standard directional PS problem.

More generally, we propose to sample the values of both the light direction and the light intensity by introducing a planar grid of calibration patterns *in lieu* of the plane $\mathcal{P}$. Standard calibration of lighting can be applied to all these patterns, before interpolating these samples so as to obtain a dense calibration of the light field $\mathbf{s}(p)$.

We then find ourselves in the standard calibrated PS case, which does not require any iterative procedure, since the model is linear w.r.t. $m = \rho \mathbf{n}$. Eventually, the 3D-points can be estimated by integrating the estimated normals. In Figure 5, we show the results of applying this approach to images captured using a dermoscopy device (a high-focal camera usually intended to observe the skin). Microstructures of a painted (so as to remove the specularities) 2-cents euro coin are revealed, while the 3D-reconstruction also exhibits an accurate 3D-reconstruction of low-frequencies, thanks to the appropriate dense calibration of the lightings.



Figure 5. 3D-reconstruction using PS with complex primary lightings. Left: 3 of the $m = 15$ input images. In each image, the scene is illuminated by a combination of primary and secondary illuminants (LEDs and their reflections inside the device), which is sampled by introducing a planar calibration grid, before interpolation. Right: 3D-reconstruction. Both high- and low-frequency details are finely recovered.

7. CONCLUSION AND PERSPECTIVES

In this paper, we studied several primary illumination models, from the simplest to the most general, in order to adapt the photometric stereo technique to various situations impelled by industrial applications. We showed that in most cases, an explicit lighting model can be derived from Lambert’s law, offering a natural framework for inverting the image graylevel model. In future work, we will further investigate on the dense light flow sampling we eventually described, so as to propose a strategy for calibrating the light fields in 3 dimensions.

ACKNOWLEDGMENTS

Part of this work was funded by Toulouse Tech Transfer, during technological transfers between IRIT and the Fitting Box company (for LEDs-based PS), and between IRIT and the Pixience company (for general lightings).

REFERENCES

- [1] Woodham, R. J., “Photometric method for determining surface orientation from multiple images,” *Optical engineering* **19**(1), 191139–191139 (1980).
- [2] Quéau, Y., Lauze, F., and Durou, J.-D., “A ll-tv algorithm for robust perspective photometric stereo with spatially-varying lightings,” in [*Fifth International Conference on Scale Space and Variational Methods in Computer Vision*], To appear (2015).
- [3] Horn, B. K. P., [*Robot vision*], MIT press (1986).
- [4] Yuille, A. L., Snow, D., Epstein, R., and Belhumeur, P. N., “Determining generative models of objects under varying illumination: Shape and albedo from multiple images using SVD and integrability,” *International Journal of Computer Vision* **35**(3), 203–222 (1999).
- [5] Basri, R., Jacobs, D., and Kemelmacher, I., “Photometric stereo with general, unknown lighting,” *International Journal of Computer Vision* **72**(3), 239–257 (2007).
- [6] Belhumeur, P., Kriegman, D., and Yuille, A. L., “The bas-relief ambiguity,” *International Journal of Computer Vision* **35**(1), 33–44 (1999).
- [7] Alldrin, N. G., Mallick, S. P., and Kriegman, D. J., “Resolving the generalized bas-relief ambiguity by entropy minimization,” in [*Computer Vision and Pattern Recognition, IEEE Conference on*], (2007).
- [8] Favaro, P. and Papadimitri, T., “A Closed-Form Solution to Uncalibrated Photometric Stereo via Diffuse Maxima,” in [*Computer Vision and Pattern Recognition, IEEE Conference on*], (2012).
- [9] Quéau, Y., Lauze, F., and Durou, J.-D., “Solving uncalibrated photometric stereo using total variation,” *Journal of Mathematical Imaging and Vision* , To appear (2015).
- [10] Durou, J.-D., Aujol, J.-F., and Courteille, F., “Integrating the normal field of a surface in the presence of discontinuities,” in [*Energy Minimization Methods in Computer Vision and Pattern Recognition*], *Lecture Notes in Computer Science* **5681**, 261–273 (2009).
- [11] Papadimitri, T. and Favaro, P., “Uncalibrated near-light photometric stereo,” in [*Proceedings of the British Machine Vision Conference*], (2014).
- [12] Mecca, R., Wetzler, A., Bruckstein, A. M., and Kimmel, R., “Near field photometric stereo with point light sources,” *SIAM Journal on Imaging Sciences* **7**(4), 2732–2770 (2014).
- [13] Barsky, S. and Petrou, M., “The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on* **25**(10), 1239–1252 (2003).
- [14] Clark, J. J., “Photometric stereo using lcd displays,” *Image and Vision Computing* **28**(4), 704–714 (2010).
- [15] Funk, N. and Yang, Y.-H., “Using a raster display for photometric stereo,” in [*Fourth Canadian Conference on Computer and Robot Vision*], (2007).
- [16] Sun, J., Smith, M., Smith, L., and Farooq, A., “Sampling light field for photometric stereo,” *International Journal of Computer Theory and Engineering* **5**(1), 14–18 (2013).
- [17] Angelopoulou, M. and Petrou, M., “Uncalibrated flatfielding and illumination vector estimation for photometric stereo face reconstruction,” *Machine Vision and Applications* **25**(5), 1317–1332 (2014).