

Photometric Stereo with Only Two Images: A Theoretical Study and Numerical Resolution[☆]

Yvain Quéau^{a,b}, Roberto Mecca^{c,d,1}, Jean-Denis Durou^a, Xavier Descombes^e

^a*IRIT, Université de Toulouse, France*

^b*Technical University Munich, Garching, Germany*

^c*Department of Engineering, University of Cambridge, United Kingdom*

^d*Department of Mathematics, University of Bologna, Italy*

^e*INRIA, Sophia Antipolis, France*

Abstract

This work tackles the problem of two-image photometric stereo. This problem constitutes the intermediate case between conventional photometric stereo with at least three images, which is well-posed, and shape-from-shading, which is ill-posed. We first provide a theoretical study of ambiguities arising in this intermediate case. Based on this study, we show that when the albedo is known, disambiguation can be formulated as a binary labelling problem, using integrability and a nonstationary Ising model. The resulting optimization problem is solved efficiently by resorting to the graph cut algorithm. These theoretical and numerical contributions are eventually validated in an application to three-image photometric stereo with shadows.

Keywords: 3D-reconstruction, shape-from-shading, photometric stereo, PDEs, numerical analysis, optimization, graph cut.

1. Introduction

In the computer vision field, 3D-shape reconstruction using digital images as input data has gained a growing importance. Interest in this task has increased even more since most mass digital devices have been equipped with cameras. Based on more than thirty years of research, such devices are

[☆]Author version. Publisher version's DOI: 10.1016/j.imavis.2016.11.006

Email address: `yvain.queau@tum.de` (Yvain Quéau)

¹Marie Curie fellow of the Istituto Nazionale di Alta Matematica, Italy.

potentially convertible into 3D-scanners without any hardware correction. Among all the photographic 3D-reconstruction techniques, we focus in this work on shape-from-shading (SFS) and photometric stereo (PS), which exploit shading information when one (SFS) or several (PS) sources illuminate the observed object. For a comprehensive overview on these techniques, see the reference book [1] by Horn and Brooks, but also [2] and [3, 4] for up-to-date surveys on SFS and PS, respectively.

Many articles enlightened the impossibility of avoiding any ambiguity while retrieving the shape from a single image, as in the SFS problem [5]. This impossibility arises from the difficulty met in distinguishing the concave from the convex surfaces. The most natural way to solve this problem is to use more than one image: Woodham showed in [6] that three is the minimum number of images to ensure well-posedness of the PS problem.

This work focuses on the intermediate case when only two images are taken into account (this specific situation will be referred to as PS2). Besides being particularly interesting for dedicated applications as single-day outdoor PS from sun light [7, 8], the PS2 problem can be seen as the degenerative case of lack of information from the three-source PS problem due to shadows [9]. In this view, we provide working tools aimed at solving the underlying ambiguities, derived from a theoretical study.

There exist a combinatorial number of normal fields which are solutions of the PS2 problem. Exhaustive search can be carried out among these normal fields, in order to find the one which best satisfies a smoothness constraint [10]. Alternatively, one may resort to the differential approach of PS, which implicitly enforces smoothness. A meaningful solution of the resulting PDE can be obtained either by specifying an explicit boundary condition [11], or by resorting to regularization [9]. Unfortunately, knowledge of the surface on the boundary is rarely available, and regularization techniques come along with parameters to tune, which might be tedious.

We put forward a new method for solving the PS2 problem, which is based on the non-differential approach of PS (normal estimation). We assume that the setup consists of orthographic viewing geometry, parallel and uniform lighting, as well as Lambertian reflectance. Under such assumptions, and provided that the albedo is known, it is shown in this paper that the number of possible solutions can be predicted beforehand. It is then demonstrated that exhaustive search of the “best” normal field can be recast as a binary labelling problem, efficiently solvable by resorting to the graph cut algorithm. In contrast with existing methods, the proposed one requires neither knowing

a boundary condition, nor tuning any parameter.

In the following, we first review the general equations of SFS, PS and PS2 in Section 2. In Section 3, we compare the differential and non-differential formulations of PS2. We show why the PS2 problem usually has a unique solution in Section 4. In Section 5, a practical graph cut-based [12] algorithm to compute this solution is introduced, and it is applied in Section 6 to the problem of three-source PS with shadows.

2. From Shape-from-shading to Photometric Stereo

2.1. Shape-from-shading

To fully describe the problem we are interested in, let us first recall some features of the SFS problem. We attach to the camera a 3D-Cartesian coordinate system xyz , so that xy coincides with the image plane and z with the optical axis. Under the assumption of orthographic projection, the visible part of a surface is a graph $z = u(x, y)$. It is well known that the SFS problem is modeled by the image irradiance equation [1]:

$$R(\mathbf{n}(x, y)) = \mathcal{I}(x, y) \quad (1)$$

where $\mathcal{I}(x, y)$ is the graylevel at image point (x, y) , and the reflectance function $R(\mathbf{n}(x, y))$ gives the value of the light re-emitted by the surface as a function of its orientation i.e., of the unit-length outgoing normal $\mathbf{n}(x, y)$ to the surface at surface point $[x, y, u(x, y)]^\top$. The unknown depth u has to be reconstructed on a compact domain $\Omega \subset \mathbb{R}^2$ called the reconstruction domain.

Let us consider a unique parallel and uniform light beam whose direction is indicated by the unit-length vector $\mathbf{s} = [s_1, s_2, s_3]^\top = [\tilde{\mathbf{s}}^\top, s_3]^\top \in \mathbb{R}^3$, and whose intensity is denoted by ψ . Assuming the observed object has purely diffuse reflection, and ignoring shadows, Eq. (1) can be written as follows:

$$\rho(x, y) \psi \mathbf{s} \cdot \mathbf{n}(x, y) = \mathcal{I}(x, y) \quad (2)$$

where $\rho(x, y) \in [0, 1]$ is the albedo.

In fact, this equality is nothing more than a relation of proportionality. Knowing that the vectors \mathbf{s} and $\mathbf{n}(x, y)$ have unit-length, and assuming that ψ is a constant factor, it seems justified to rewrite Eq. (2) as a real equality:

$$\rho(x, y) \mathbf{s} \cdot \mathbf{n}(x, y) = I(x, y) \quad (3)$$

where $I(x, y) \in [0, 1]$ should now be considered as the *normalized* graylevel.

Eq. (3) is a particular non-differential formulation (amongst many others) of the SFS problem. Once the normal field \mathbf{n} has been estimated, it has to be integrated. This means that the following equation in u has to be solved [13]:

$$\mathbf{n}(x, y) = \frac{1}{\sqrt{1 + \|\nabla u(x, y)\|^2}} [-\nabla u(x, y)^\top, 1]^\top \quad (4)$$

where $\nabla u(x, y) = [\partial_x u(x, y), \partial_y u(x, y)]^\top$ denotes the gradient of $u(x, y)$. From Eqs. (3) and (4), we get the following differential formulation of SFS:

$$\rho(x, y) \frac{-\tilde{\mathbf{s}} \cdot \nabla u(x, y) + s_3}{\sqrt{1 + \|\nabla u(x, y)\|^2}} = I(x, y) \quad (5)$$

which is a first-order nonlinear PDE of the Hamilton-Jacobi type. We refer the interested reader to the survey presented in [2] for a presentation of recent results on the eikonal equation, which follows from (5) when $\mathbf{s} = [0, 0, 1]^\top$.

2.2. Photometric Stereo

Even if \mathbf{s} is known, SFS is ill-posed without any additional knowledge on the surface to be reconstructed. In most papers on SFS, the albedo $\rho(x, y)$ is supposed to be known, but this is still not enough to make the problem well-posed. The simplest way to overcome SFS ill-posedness is to use $m \geq 2$ images taken from the same point of view, illuminated by m light sources (\mathbf{s}^i, ψ^i) , $i \in [1, m]$. This new problem is called photometric stereo (PS). The classical resolution of PS is based on a local estimate of the outgoing unit-length normal to the surface [6]. For a Lambertian surface, the non-differential formulation of PS consists in solving a system of m equations of type (3):

$$\rho(x, y) \mathbf{s}^i \cdot \mathbf{n}(x, y) = I^i(x, y), \quad i \in [1, m] \quad (6)$$

As for SFS, this formulation requires that Eq. (4) is solved afterwards. From a theoretical point of view, Eq. (4) admits a solution in u only if the estimated normal field is *integrable* [14] (cf. Section 4.5). Due to estimation errors, this is rarely the case in practice, hence projection of the estimated normal field on the space of integrable fields must be achieved, resorting for instance to Fourier analysis [14] or to variational methods [13]. Alternatively, one can directly try to estimate the “most integrable” normal field. This is the approach that is followed in Section 5.

Of course, a differential formulation of PS also exists, which aims at solving a system of m nonlinear PDEs of type (5):

$$\rho(x, y) \frac{-\tilde{\mathbf{s}}^i \cdot \nabla u(x, y) + s_3^i}{\sqrt{1 + \|\nabla u(x, y)\|^2}} = I^i(x, y), \quad i \in [1, m] \quad (7)$$

In the usual case, denoted PS3, where $m \geq 3$ non-coplanar distant calibrated light sources are used [6], System (6) reduces to a full-rank linear system in $\mathbf{m}(x, y) = \rho(x, y) \mathbf{n}(x, y) \in \mathbb{R}^3$. Solving this system has several advantages, compared to SFS: it is well-posed and can be locally solved, thus parallelized. Furthermore, the albedo no longer has to be known.

We may wonder whether the differential formulation (7) would really be pertinent for PS3. The main advantage of solving (7) is that integrability is implicitly ensured, unlike solving (6), knowing that the lack of integrability of \mathbf{n} complicates the resolution of Eq. (4) [14]. However, the problem (7) has two drawbacks: it is nonlinear and cannot be solved locally [15].

In this paper, we focus on the resolution of PS when the linear system (6) is not full-rank. In such cases, the non-differential formulation does not have as many advantages as for PS3, and the differential formulation might be worthwhile, at least because it is better-posed since the integrability constraint is implicitly satisfied. Indeed, practical solutions to the rank-deficient PS problem use this differential formulation. Yet, as discussed in Section 3.2, differential approaches have to resort either to a boundary condition (which is rarely available) or to regularization (which requires parameter tuning). The solution presented in Section 5, which is based on the non-differential formulation, involves neither boundary condition nor parameter tuning.

2.3. Scope of our Work: Photometric Stereo Using Two Images

The scope of our work is the intermediate case $m = 2$ of PS between $m = 1$ (SFS) and $m \geq 3$ (PS3), denoted PS2. Such a problem arises in real-time 3D-reconstruction of non-rigid objects [16]. This application can be carried out using PS3, an RGB sensor and three colored light sources. Yet, when one of the light source is occluded, only two of the color channels provide meaningful shading clues, and hence a PS2 problem must be solved. The PS2 problem is also very similar to the case $m \geq 3$ with coplanar light vectors \mathbf{s}^i . A concrete example to be mentioned is the case when a scene is illuminated by the sun [7]. In both cases, the linear system (6) is no longer full-rank in $\mathbf{m}(x, y) = \rho(x, y) \mathbf{n}(x, y)$.

The PS2 problem has been addressed in few papers. Onn and Bruckstein prove in [10] that the determination of the normal is a priori ambiguous, and demonstrate how to eliminate this ambiguity using the integrability constraint. Yang et al. study the problem in the particular case of convex objects [17]. In [7], Sato and Ikeuchi use the resolution method designed by Onn and Bruckstein to solve the PS problem using $m \geq 3$ images under solar illumination, which partially brings us back to the PS2 problem [18]. In [19], Kozera reaches the same conclusions as Onn and Bruckstein, by making an analytical resolution of the differential formulation. From 1995 and for over ten years, only Ikeda addressed the PS2 problem [20]. However, he essentially considers the second image as a means to more accurately solve the SFS problem. More recently, the problem of outdoor PS has been reexplored in [8, 21, 22]. Finally, let us quote a preliminary version [11] of our work.

One could wonder whether the PS2 problem is not purely formal, considering that the actual trend is to deal with far more than two images, as in [23] where a full video is used, or in [21, 22] where timelapse images sequences, acquired over several years, are considered. Apart from the intrinsic interest we can find in studying the number of solutions to the PS2 problem, PS in the presence of shadows has been studied in several recent papers [9, 24, 25, 26, 27, 28, 29]. When $m > 3$, the graylevels which lie inside shadows are considered as outliers and left out of the estimation [24, 25, 26] or dealt with in a robust estimation process such as the Expectation-Maximization algorithm [27]. Most recent techniques assume that such outliers are sparsely distributed in the images, and hence design sparsity-enhancing algorithms [28, 29]. When $m = 3$, Hernández et al. show in [9] that implementing a specific treatment for the twice-lit points i.e., for those where the PS2 problem may arise, improves the 3D-reconstruction accuracy. The application of our study will precisely be to improve the 3D-reconstruction in such shadow areas (see Section 6).

3. Photometric Stereo Using Two Images: A Theoretical Study

3.1. Non-differential Formulation of the PS2 Problem

Taking $m = 2$ in (6), the non-differential formulation of PS2 is written as:

$$\begin{cases} \rho(x, y) [s_1^1 n_1(x, y) + s_2^1 n_2(x, y) + s_3^1 n_3(x, y)] = I^1(x, y) \\ \rho(x, y) [s_1^2 n_1(x, y) + s_2^2 n_2(x, y) + s_3^2 n_3(x, y)] = I^2(x, y) \\ n_1(x, y)^2 + n_2(x, y)^2 + n_3(x, y)^2 = 1 \end{cases} \quad (8)$$

If $\rho(x, y)$ is assumed to be known, the problem (8) consists at each point (x, y) in a nonlinear system of three equations with three unknowns (n_1, n_2, n_3) . The nonlinearity of the third equation could give rise to a non-unique solution, as observed by Ikeuchi and Horn in [30]: “Naturally, the above nonlinear equations [...] may have more than one solution, in which case additional information (such as a third image) may be needed to find a unique answer”. Indeed, most of the time, (8) does not have a unique but two solutions.

To solve (8), we only consider the set \mathcal{S} of twice-lit unit-length normals \mathbf{n} i.e., those where $\mathbf{s}^1 \cdot \mathbf{n} > 0$ and $\mathbf{s}^2 \cdot \mathbf{n} > 0$. This set is an open part of Gaussian sphere \mathcal{G} limited by two planes containing the origin (see Fig. 1-a):

$$\begin{cases} \pi^1 : & s_1^1 n_1 + s_2^1 n_2 + s_3^1 n_3 = 0 \\ \pi^2 : & s_1^2 n_1 + s_2^2 n_2 + s_3^2 n_3 = 0 \end{cases} \quad (9)$$

Planes π^1 and π^2 are orthogonal, respectively, to \mathbf{s}^1 and \mathbf{s}^2 . At each twice-lit point (x, y) characterized by the graylevels $I^1(x, y)$ and $I^2(x, y)$, and by the supposedly known albedo $\rho(x, y) \neq 0$, System (8) admits two solutions, denoted $\mathbf{n}^+(x, y)$ and $\mathbf{n}^-(x, y)$, which are the intersections of \mathcal{S} and of two planes which are obtained by translating π^1 and π^2 of $I^1(x, y)/\rho(x, y)$ and $I^2(x, y)/\rho(x, y)$ in the directions of \mathbf{s}^1 and \mathbf{s}^2 , respectively. See Fig. 1-b for a geometrical interpretation of $\mathbf{n}^+(x, y)$ and $\mathbf{n}^-(x, y)$. The analytical formulae for them will be derived in Section 5.1. Let us denote by π the plane supported by \mathbf{s}^1 and \mathbf{s}^2 and containing the origin of \mathcal{G} . It is obvious that $\mathbf{n}^+(x, y)$ and $\mathbf{n}^-(x, y)$ are symmetric with respect to π , for any $(x, y) \in \Omega$.

If $I^1(x, y)$ and $I^2(x, y)$ exactly match the Lambertian model (3), the non-differential problem (8) admits either two solutions, or one if $\mathbf{n}^+(x, y) = \mathbf{n}^-(x, y)$. But, if this model is not exactly satisfied, which may happen with real images, there may be some points (x, y) without any exact solution. The way to effectively handle these different cases will be specified in Section 5.1.




Figure 1: (a) The set \mathcal{S} of twice-lit normals, emphasized in red, is a part of Gaussian sphere \mathcal{G} limited by the planes π^1 and π^2 , which are orthogonal to light vectors \mathbf{s}^1 and \mathbf{s}^2 . (b) An example where the system (8) admits two solutions in \mathbf{n} (marked in red).

3.2. Differential Formulations of the PS2 Problem

Taking $m = 2$ in (7), and adding a Dirichlet boundary condition, a first differential formulation of the PS2 problem is written as two nonlinear PDEs:

$$\begin{cases} \rho(x, y) \frac{-\tilde{\mathbf{s}}^1 \cdot \nabla u(x, y) + s_3^1}{\sqrt{1 + \|\nabla u(x, y)\|^2}} = I^1(x, y) & \text{a.e. } (x, y) \in \Omega \\ \rho(x, y) \frac{-\tilde{\mathbf{s}}^2 \cdot \nabla u(x, y) + s_3^2}{\sqrt{1 + \|\nabla u(x, y)\|^2}} = I^2(x, y) & \text{a.e. } (x, y) \in \Omega \\ u(x, y) = g(x, y) & \forall (x, y) \in \partial\Omega \end{cases} \quad (10)$$

Besides its nonlinearity, the main drawback of Problem (10) concerns the need for a boundary condition: the function $g(x, y)$, taken in the space of Lipschitz functions, represents a piece of information which is rarely available. On the other hand, since a common factor $\rho(x, y)/\sqrt{1 + \|\nabla u(x, y)\|^2}$ occurs in both PDEs of (10), we can combine them in order to simultaneously eliminate the

albedo $\rho(x, y)$ and the nonlinearity (we only have to suppose that $\rho(x, y) \neq 0$):

$$[I^2(x, y) \tilde{\mathbf{s}}^1 - I^1(x, y) \tilde{\mathbf{s}}^2] \cdot \nabla u(x, y) = I^2(x, y) s_3^1 - I^1(x, y) s_3^2 \quad (11)$$

Considering the same boundary condition as in (10), we deduce:

$$\begin{cases} \mathbf{b}(x, y) \cdot \nabla u(x, y) = f(x, y) & \text{a.e. } (x, y) \in \Omega \\ u(x, y) = g(x, y) & \forall (x, y) \in \partial\Omega \end{cases} \quad (12)$$

where:

$$\begin{cases} \mathbf{b}(x, y) = I^2(x, y) \tilde{\mathbf{s}}^1 - I^1(x, y) \tilde{\mathbf{s}}^2 \\ f(x, y) = I^2(x, y) s_3^1 - I^1(x, y) s_3^2 \end{cases} \quad (13)$$

This second differential formulation of the PS2 problem allows us to propagate boundary information $g(x, y)$ across Ω through vector field $\mathbf{b}(x, y)$. Indeed, if $\mathbf{b}(x, y)$ and $f(x, y)$ are two bounded (but not necessarily continuous) functions defined by (13), and if $g(x, y)$ is a Lipschitz function, then (12) admits a unique Lipschitz solution $u(x, y)$ [15]. It is worth emphasizing that this implies that the differential formulations of PS2 allow one to reconstruct surfaces that are differentiable almost everywhere. In practice, this means that surfaces having sharp structures could be retrieved. In contrast, using the non-differential approach, the surface must be assumed to be \mathcal{C}^1 , so that the normal is defined everywhere. If edges or depth discontinuities are in fact present, they must be adequately handled during the integration stage [13].

In order to ensure robustness to noise, propagation schemes used in [11] may be advantageously replaced by variational methods, recasting the linear PDE (12) as the following optimization problem:

$$\begin{cases} \min_{u: \Omega \rightarrow \mathbb{R}} \iint_{\Omega} [\mathbf{b}(x, y) \cdot \nabla u(x, y) - f(x, y)]^2 dx dy \\ \text{s.t. } u(x, y) = g(x, y), \forall (x, y) \in \partial\Omega \end{cases} \quad (14)$$

When $m \geq 3$ images are available, it was shown in [31] that variational models such as (14) may be considered in several more difficult situations such as color PS, PS with pointwise sources or perspective PS, even without a boundary condition. Yet, when $m = 2$, explicit knowledge of the function g is required in order to ensure that the characteristics are not reconstructed independently. An alternative consists in using anisotropic regularization, in order to “couple”

the 3D-reconstructions of the different characteristics, by ensuring smoothness along directions that are not tangent to the characteristics. A suitable variational model, which was introduced by Hernández et al. in [9], is written as follows:

$$\min_{u: \Omega \rightarrow \mathbb{R}} \iint_{\Omega} \left\{ [\mathbf{b}(x, y) \cdot \nabla u(x, y) - f(x, y)]^2 + \alpha_1 [\mathbf{b}_{\perp}(x, y) \cdot \nabla u(x, y)]^2 + \alpha_2 [\mathbf{b}_{\perp}(x, y)^{\top} \mathcal{H}(u)(x, y) \mathbf{b}_{\perp}(x, y)]^2 \right\} dx dy \quad (15)$$

where the field \mathbf{b}_{\perp} is perpendicular to the characteristic curves and $\mathcal{H}(u)$ is the Hessian matrix of u . The parameters α_1 and α_2 must be tuned appropriately, in order to ensure that the regularization is sufficient, yet avoiding over-smoothing the solution (cf. Section 5.4).

3.3. Example of PS2 Problem

In order to question the consistency between the different formulations of the PS2 problem, let us take the example of a plane surface $u(x, y) = x$ illuminated by light vectors $\mathbf{s}^1 = [0, 0, 1]^{\top}$ and $\mathbf{s}^2 = \frac{1}{2} [1, 1, \sqrt{2}]^{\top}$. If $\rho \equiv 1$, the images of this surface using Model (6) are uniform, with graylevels $I^1(x, y) = \frac{1}{\sqrt{2}}$ and $I^2(x, y) = \frac{\sqrt{2}-1}{2\sqrt{2}}$. Let us first solve this PS2 example using the non-differential formulation (8), which is written as:

$$\begin{cases} n_3(x, y) = \frac{1}{\sqrt{2}} \\ \frac{n_1(x, y) + n_2(x, y) + \sqrt{2}n_3(x, y)}{2} = \frac{\sqrt{2}-1}{2\sqrt{2}} \\ n_1(x, y)^2 + n_2(x, y)^2 + n_3(x, y)^2 = 1 \end{cases} \quad (16)$$

This system admits two solutions $\mathbf{n}^+ = \frac{1}{\sqrt{2}} [-1, 0, 1]^{\top}$ and $\mathbf{n}^- = \frac{1}{\sqrt{2}} [0, -1, 1]^{\top}$, independently of (x, y) . If the surface to be reconstructed is supposed to be differentiable everywhere, there are only two acceptable normal fields $\mathbf{n}(x, y) = \mathbf{n}^+$ or $\mathbf{n}(x, y) = \mathbf{n}^-$. It follows from (4) that there are two possible values for $\nabla u(x, y)$:

$$\frac{-1}{n_3^+(x, y)} \begin{bmatrix} n_1^+(x, y) \\ n_2^+(x, y) \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix} ; \quad \frac{-1}{n_3^-(x, y)} \begin{bmatrix} n_1^-(x, y) \\ n_2^-(x, y) \end{bmatrix} = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \quad (17)$$

Imposing $u(0, 0) = 0$ to fix the integration constants, we finally obtain two solutions $u^+(x, y) = x$ and $u^-(x, y) = y$. Fortunately, one of these solutions is the genuine surface.

Now, let us write the first differential formulation (10), supposing $\rho \equiv 1$:

$$\begin{cases} \frac{1}{\sqrt{1 + \|\nabla u(x, y)\|^2}} = \frac{1}{\sqrt{2}} \\ \frac{-\partial_x u(x, y) - \partial_y u(x, y) + \sqrt{2}}{2\sqrt{1 + \|\nabla u(x, y)\|^2}} = \frac{\sqrt{2} - 1}{2\sqrt{2}} \end{cases} \quad (18)$$

Note that no boundary condition is available. System (18) is equivalent to:

$$\begin{cases} \partial_x u(x, y) + \partial_y u(x, y) = 1 \\ \|\nabla u(x, y)\|^2 = 1 \end{cases} \quad (19)$$

By replacing $\|\nabla u(x, y)\|^2$ with $\partial_x u(x, y)^2 + \partial_y u(x, y)^2$, we quickly find that (19) has the same solutions (17) in $\nabla u(x, y)$ as (16). Both formulations are thus consistent, in the sense that they provide the same global solutions.

It follows from the definitions (13) that:

$$\mathbf{b}(x, y) = -\frac{1}{2\sqrt{2}}[1, 1]^\top \quad ; \quad f(x, y) = -\frac{1}{2\sqrt{2}} \quad (20)$$

In the absence of a boundary condition, the second differential formulation (12) reduces to:

$$\partial_x u(x, y) + \partial_y u(x, y) = 1 \quad (21)$$

Unsurprisingly, this PDE is the first equation of (19), which admits an infinity of solutions. For instance, all functions $u(x, y) = a x + (1 - a) y + w(x - y)$, for any $a \in \mathbb{R}$ and any differentiable function $w : \mathbb{R} \rightarrow \mathbb{R}$, are solutions to (21). How can it be explained that this formulation does not lead to the same conclusion as the others? In fact, without a boundary condition, differential formulation (12) is a necessary but insufficient condition. For better constraint, a boundary condition is required (but rarely available).

To conclude, there is obviously no inconsistency between the non-differential formulation and the first differential formulation of the PS2 problem. Hence, both these formulations should allow us to predict the same number of global solutions. Let us examine in more detail this issue.

4. PS2 Problem: Predicting the Number of Solutions

At any twice-lit point $(x, y) \in \Omega$, we know from Section 3.1 that the PS2 problem admits one or two solutions in $\mathbf{n}(x, y)$. Of course, in order to predict the number of normal fields, the points where the normal can be determined unambiguously are of primary importance. Such *singular points* have been studied in detail to solve the SFS problem [30].

4.1. Singular Points

As already observed in Section 3.1, the first situation where normal uniqueness can be proved is reached when both solutions $\mathbf{n}^+(x, y)$ and $\mathbf{n}^-(x, y)$ to the problem (8) coincide. In this case, the unique solution is inside plane π . The set \mathcal{S}_R of such normals is a geodesic on \mathcal{G} . It is the intersection between \mathcal{S} and π (see Fig. 2-a). A second type of singular point occurs when one solution to (8) points towards the viewer (\mathcal{S}_G set in Fig. 2-b), while the other one points away (\mathcal{S}_Y set in Fig. 2-b). In this case, the ambiguities are easily eliminated by choosing the “visible” normal. Let Ω_R and Ω_G be the singular points sets where the normal belongs to \mathcal{S}_R or \mathcal{S}_G , respectively.




Figure 2: (a) Red geodesic \mathcal{S}_R is the intersection between \mathcal{S} and π . Each normal pointing to \mathcal{S}_R is known without ambiguity. (b) This also holds true for each normal pointing to \mathcal{S}_G , since the other possible normal points towards the non-visible part \mathcal{S}_Y of \mathcal{S} .

Since a PS2 problem comes down to a pair of SFS problems (with the constraint that the camera pose is unique), one could wonder why the singular points of each SFS problem are not taken into account. These points are such that $\mathbf{n} = \mathbf{s}^1$ in the first image, or $\mathbf{n} = \mathbf{s}^2$ in the second. However, both these values of \mathbf{n} are inside \mathcal{S}_R , since they support plane π . This shows us that the singular points of the PS2 problem include those of each SFS subproblem.

The normal field is continuous if the surface is supposed to be C^1 . Under such an assumption, is it possible to propagate the knowledge of the normal in a singular point to its non-singular neighbours? The answer to this question depends on which type of singular points is referred to. Let $P_R \in \Omega_R$ be a singular point of the first type i.e., one whose normal is inside \mathcal{S}_R . The normal in a non-singular neighbour \bar{P}_R of P_R can lie on both sides of \mathcal{S}_R . The subsets of \mathcal{S} which are above and below the geodesic \mathcal{S}_R are respectively called \mathcal{S}_U and \mathcal{S}_B , see Fig. 3. In other words, there is a remaining ambiguity on the normal in \bar{P}_R . Now, let $P_G \in \Omega_G$ be a singular point of the second type, whose normal is inside \mathcal{S}_G . In any non-singular neighbour \bar{P}_G of P_G , we can infer from the normal field continuity that the normal is inside \mathcal{S}_U . To conclude, the normal is unambiguously known in all points connected to Ω_G inside $\Omega \setminus \Omega_R$, which can be considered as supplementary singular points.




Figure 3: (a) Set \mathcal{S}_B is colored in pink. It is the subset of \mathcal{S} bounded by equator E and geodesic \mathcal{S}_R . (b) Gaussian sphere \mathcal{G} is seen from the camera point of view (i.e. z direction). Orange set \mathcal{S}_U is the subset of \mathcal{S} located between \mathcal{S}_G and \mathcal{S}_R .

4.2. Using the Singular Points to Construct a Boundary Condition

As an illustration, let us calculate a pair of images of the smooth surface depicted in Fig. 4-a, supposing $\rho \equiv 1$. The light directions are given, respectively, by $(\theta_1, \phi_1) = (60^\circ, 17^\circ)$ and $(\theta_2, \phi_2) = (135^\circ, 17^\circ)$, which avoids shadows (see Figs. 4-b and 4-c).

The sets of singular points Ω_R and Ω_G , which are numerically estimated from the estimation of \mathbf{n}^+ and \mathbf{n}^- described in Section 5, are superimposed to these images. In each of the five connected parts of $\Omega \setminus \Omega_R$, there are two solutions: one is contained in \mathcal{S}_U , the other in \mathcal{S}_B . This gives rise to $2^5 = 32$ continuous normal fields. Nevertheless, a difference between our count and that put forward by Onn and Bruckstein [10] comes from the detection of Ω_G : since the region $\Omega \setminus \Omega_R$ containing Ω_G is determined, only the other four regions are ambiguous, leading eventually to $2^4 = 16$ possible normal fields.

The solution u can be calculated over each connected singular points set, up to a constant of integration. Knowing that differential formulation (12) needs a boundary condition to be well-posed, this is a simple way to construct one. Since $\partial\Omega$ is connected to Ω_G inside $\Omega \setminus \Omega_R$ (see Figs. 4-b and 4-c), the solution can be univocally calculated along $\partial\Omega$. According to [15], this allows us to predict a unique solution, which is of course the original surface shown in Fig. 4-a. Based on this example, we could conclude that the non-differential and differential approaches to the PS2 resolution are complementary, in order to predict the number of solutions. Namely, the boundary condition required by the latter is provided by the former.




Figure 4: (b-c) A pair of 256×256 synthetic images (stored in 32 bits) of the smooth surface (a), such that all the points are twice-lit, over which Ω_R (in red) and Ω_G (in green) are superimposed.

4.3. A Possible Remaining Ambiguity

The conclusion of the previous section does not always hold true. For instance, in a case such as the example of Section 3.3, there is no singular point (all the points have the same normal). Let us show another counterexample originally exhibited by Kozera in [32]. The surface represented by equation $z = xy$, with uniform albedo $\rho \equiv 1$, illuminated by light vectors $\mathbf{s}^1 = [s, s, c]^\top$ and $\mathbf{s}^2 = [-s, -s, c]^\top$, where $s = \sqrt{2}/2 \sin \phi$ and $c = \cos \phi$, for a given $\phi \in]0, \pi/2[$, is characterized by the following graylevel functions:

$$\begin{cases} I^1(x, y) = \frac{s(-x - y) + c}{\sqrt{1 + x^2 + y^2}} \\ I^2(x, y) = \frac{-s(-x - y) + c}{\sqrt{1 + x^2 + y^2}} \end{cases} \quad (22)$$

The two first equations of the non-differential problem (8) are rewritten as:

$$\begin{cases} s[n_1(x, y) + n_2(x, y)] + c n_3(x, y) = \frac{s(-x - y) + c}{\sqrt{1 + x^2 + y^2}} \\ -s[n_1(x, y) + n_2(x, y)] + c n_3(x, y) = \frac{-s(-x - y) + c}{\sqrt{1 + x^2 + y^2}} \end{cases} \quad (23)$$

Since by definition, s and c are nonzero, System (23) is equivalent to:

$$\begin{cases} n_1(x, y) + n_2(x, y) = \frac{-x - y}{\sqrt{1 + x^2 + y^2}} \\ n_3(x, y) = \frac{1}{\sqrt{1 + x^2 + y^2}} \end{cases} \quad (24)$$

Using (24), the third equation of (8) can be rewritten as:

$$n_1(x, y)^2 + \frac{x + y}{\sqrt{1 + x^2 + y^2}} n_1(x, y) + \frac{xy}{1 + x^2 + y^2} = 0 \quad (25)$$

It is easy to check that this second-order equation always admits two real solutions in $n_1(x, y)$, which come down to a unique solution when $y = x$.

These solutions give rise to two possible normals at each point $(x, y) \in \Omega$:

$$\mathbf{n}^+(x, y) = \frac{1}{\sqrt{1+x^2+y^2}} \begin{bmatrix} -y \\ -x \\ 1 \end{bmatrix} ; \mathbf{n}^-(x, y) = \frac{1}{\sqrt{1+x^2+y^2}} \begin{bmatrix} -x \\ -y \\ 1 \end{bmatrix} \quad (26)$$

We deduce from (26) and (4) two possible values for $\nabla u(x, y)$:

$$\frac{-1}{n_3^+(x, y)} \begin{bmatrix} n_1^+(x, y) \\ n_2^+(x, y) \end{bmatrix} = \begin{bmatrix} y \\ x \end{bmatrix} ; \frac{-1}{n_3^-(x, y)} \begin{bmatrix} n_1^-(x, y) \\ n_2^-(x, y) \end{bmatrix} = \begin{bmatrix} x \\ y \end{bmatrix} \quad (27)$$

Both these vector fields are easily integrated, which provides us with two solutions $u^+(x, y) = xy$ and $u^-(x, y) = (x^2 + y^2)/2$, up to two additive constants. That is to say, there is a remaining ambiguity.

On the other hand, Ω_G is empty, but it is easily deduced from (26) that Ω_R is the straight line $y = x$. The solution can therefore be calculated along this line, up to a constant. Using a similar rationale as in Section 4.2, we should conclude that the solution is unique, which would contradict the previous result. This contradiction is easily explained: the prediction of [15] holds true only if the solution is known on a curve which is not a characteristic. As we will see in the next section, this condition is precisely not valid in this case.

4.4. Using the Characteristics to Predict the Number of Solutions

In the previous example, each vector $\mathbf{b}(x, y)$ defined in Eq. (13) is parallel to $[1, 1]^\top$. The characteristics are thus the straight lines represented by equations $y = x + \eta$, $\eta \in \mathbb{R}$, including the set Ω_R . Depth u can be univocally calculated along each characteristic, up to a constant. This uniqueness result is not contradicted by the previous two-fold ambiguity since, for any $\eta \in \mathbb{R}$:

$$u^+(x, x + \eta) - u^-(x, x + \eta) = x(x + \eta) - \frac{x^2 + (x + \eta)^2}{2} = -\frac{\eta^2}{2} \quad (28)$$

is independent from x .

If $u(x, y)$ is known at one point of each characteristic, we have a better understanding why Problem (12) has a unique solution. Following this rationale, all functions of the following form seem to be solutions to the previous example:

$$u(x, y) = u^+(x, y) + v(y - x) \quad (29)$$

provided that v is a scalar function such that $v(y - x)$ is constant along each

characteristic. In fact, any function v is not acceptable because, as already noted, differential formulation (12) is a necessary but insufficient condition, in the absence of boundary condition. From (29), we deduce:

$$\nabla u(x, y) = \begin{bmatrix} y - v'(y - x) \\ x + v'(y - x) \end{bmatrix} \quad (30)$$

Eq. (7) tells us that the surface $z = u(x, y)$ of albedo $\rho \equiv 1$ is a solution to the previous example only if:

$$\begin{cases} \frac{-s[y - v'(y - x)] - s[x + v'(y - x)] + c}{\sqrt{1 + [y - v'(y - x)]^2 + [x + v'(y - x)]^2}} = I^1(x, y) \\ \frac{s[y - v'(y - x)] + s[x + v'(y - x)] + c}{\sqrt{1 + [y - v'(y - x)]^2 + [x + v'(y - x)]^2}} = I^2(x, y) \end{cases} \quad (31)$$

Using (22), we easily find the only two solutions in $v'(y - x)$ to System (31):

$$\begin{cases} v'_1(y - x) = 0 \\ v'_2(y - x) = y - x \end{cases} \implies \begin{cases} v_1(y - x) = K_1, \quad K_1 \in \mathbb{R} \\ v_2(y - x) = \frac{(y - x)^2}{2} + K_2, \quad K_2 \in \mathbb{R} \end{cases} \quad (32)$$

Plugging (32) into (29), we eventually obtain the two following solutions:

$$\begin{cases} u_1(x, y) = u^+(x, y) + K_1 \\ u_2(x, y) = u^+(x, y) + \frac{(y - x)^2}{2} + K_2 = u^-(x, y) + K_2 \end{cases} \quad (33)$$

This result confirms that there are only two analytical solutions u^+ and u^- , up to the constants K_1 and K_2 .

4.5. Integrability Constraint

Even if vector field $[p, q]^\top = [-n_1/n_3, -n_2/n_3]^\top$ is easily calculated from a normal field, there is no guarantee that this vector field is *integrable* i.e., that it satisfies the integrability constraint [33]:

$$\frac{\partial p}{\partial y} = \frac{\partial q}{\partial x} \quad (34)$$

whereas this is required, if the surface is supposed to be at least C^2 , in order to be sure that equation $\nabla u = [p, q]^\top$ has a solution in u . Onn and Bruckstein note in [10] that “most of the time”, in each connected part \mathcal{P} of $\Omega \setminus \Omega_R$, one of the two possible normal fields may be discarded since it is not integrable. To decide, they rely on a criterion deduced from (34):

$$\iint_{(x,y) \in \mathcal{P}} \left[\frac{\partial p}{\partial y}(x, y) - \frac{\partial q}{\partial x}(x, y) \right]^2 dx dy = 0 \quad (35)$$

They also characterize the “rare cases” where (35) is satisfied by more than one normal field, in which case the PS2 problem admits several solutions. In fact, the examples of Sections 3.3 and 4.3 are such “rare cases”.

Let us now return to the example of Fig. 4. Using the non-differential formulation, we found $2^4 = 16$ possible normal fields. A more complex rationale based on the topology of sets Ω_R and Ω_G allowed us to predict a unique solution. Consequently, among the sixteen normal fields, only one is integrable. However, a prediction based on topology is difficult to extend to the discrete framework, since the notion of continuity will be lost.

Indeed, let us calculate the sets Ω_R and Ω_G when the images are quantized. It is clear from Fig. 5 that the topology of Ω_R is very sensitive to quantization noise. It is thus no longer possible to predict the number of solutions using the rationale of Section 4.2 from such a fragmented set Ω_R .




Figure 5: Two same images as in Fig. 4, quantized using 256 levels (8 bits). While this is not the case of Ω_G (right), Ω_R becomes fragmented (left).

We show in the next section that it is possible to efficiently find the most integrable normal field, and hence to eliminate the ambiguities of the PS2 problem, without knowledge of a boundary condition [11] nor parameter tuning [9].

5. Photometric Stereo Using Two Images: A Numerical Resolution

We know from Section 3.1 that the PS2 problem admits at most two solutions in \mathbf{n} at each point $(x, y) \in \Omega$ (see Fig. 6). We can thus a priori construct at most $2^{|\Omega|}$ different discrete normal fields, where $|\Omega|$ denotes the number of pixels inside Ω . Yet, except in some rare cases [10], only one of these candidates is integrable. Let us show how to efficiently find this most integrable normal field.




Figure 6: Top: RGB-encoded normal fields (\mathbf{n} is the ground truth, \mathbf{n}^+ and \mathbf{n}^- are the normal fields estimated by (42)), using the quantized dataset presented in Fig. 5. Bottom: 3D-shapes obtained integrating the normals [13]. The normal fields \mathbf{n}^+ and \mathbf{n}^- are only two solutions among $2^{|\Omega|}$, since any combination of both these normal fields is also plausible.

5.1. Estimating the Candidate Normal Fields

Problem (8) has two solutions \mathbf{n}^+ and \mathbf{n}^- at each point. Let us first show how to express these solutions using a purely algebraic method.

At each point $(x, y) \in \Omega$, the two first equations of Problem (8) write (the dependencies in (x, y) are omitted):

$$\underbrace{\begin{bmatrix} \mathbf{s}^{1\top} \\ \mathbf{s}^{2\top} \end{bmatrix}}_{\mathbf{S}^\top \in \mathbb{R}^{2 \times 3}} \mathbf{n} = \underbrace{\begin{bmatrix} \frac{I^1}{\rho} \\ \frac{I^2}{\rho} \\ \frac{\rho}{\rho} \end{bmatrix}}_{\mathbf{i} \in \mathbb{R}^2} \quad (36)$$

With a view to solving Eq. (36) in the least-squares sense in order to handle quantization noise, we reformulate problem (8) as follows:

$$\begin{cases} \min_{\mathbf{n}} \|\mathbf{S}^\top \mathbf{n} - \mathbf{i}\|^2 \\ \text{s.t. } \|\mathbf{n}\| = 1 \end{cases} \quad (37)$$

The singular value decomposition (SVD) of \mathbf{S}^\top can be written as:

$$\mathbf{S}^\top = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top = \sigma_1 \mathbf{u}_1 \mathbf{v}_1^\top + \sigma_2 \mathbf{u}_2 \mathbf{v}_2^\top \quad (38)$$

In (38), $\sigma_1 \geq \sigma_2 > 0$ are the pair of strictly positive singular values of \mathbf{S} (\mathbf{S} has rank 2, since \mathbf{s}^1 and \mathbf{s}^2 are non-collinear), \mathbf{u}_1 and \mathbf{u}_2 are orthonormal vectors of \mathbb{R}^2 , and \mathbf{v}_1 and \mathbf{v}_2 are orthonormal vectors of \mathbb{R}^3 . It can be shown (Theorem 5.5.1 in [34]) that the minimum-norm solution of the least-squares problem resulting from (36):

$$\begin{cases} \mathbf{n}_0 = \underset{\mathbf{n} \in \mathcal{E}}{\operatorname{argmin}} \|\mathbf{n}\| \\ \text{s.t. } \mathcal{E} = \{\mathbf{n} \in \mathbb{R}^3, \|\mathbf{S}^\top \mathbf{n} - \mathbf{i}\|^2 \leq \|\mathbf{S}^\top \mathbf{m} - \mathbf{i}\|^2 \forall \mathbf{m} \in \mathbb{R}^3\} \end{cases} \quad (39)$$

is written as:

$$\mathbf{n}_0 = \frac{\mathbf{u}_1^\top \mathbf{i}}{\sigma_1} \mathbf{v}_1 + \frac{\mathbf{u}_2^\top \mathbf{i}}{\sigma_2} \mathbf{v}_2 \quad (40)$$

Let us now solve Problem (37). Three cases can occur:

1. If $\|\mathbf{n}_0\| = 1$, then $\mathbf{n}^+ = \mathbf{n}^- = \mathbf{n}_0$ is the only solution to (37), and the objective is null. This corresponds to the singular points in Ω_R , cf. Section 4.1.
2. If $\|\mathbf{n}_0\| < 1$, then neglecting the unit-length constraint, there exists a set of solutions with null objective to the optimization problem (37), which is written as:

$$\mathbf{n}_0 + \mathbb{R}\mathbf{v}_3 \quad (41)$$

where \mathbf{v}_3 is a unit-length vector of the kernel of \mathbf{S}^\top (this kernel has dimension 1, according to the rank theorem). Two of these solutions

have unit-length, which are the only two solutions to Problem (37)²:

$$\begin{cases} \mathbf{n}^+ = \mathbf{n}_0 + \sqrt{1 - \|\mathbf{n}_0\|^2} \mathbf{v}_3 \\ \mathbf{n}^- = \mathbf{n}_0 - \sqrt{1 - \|\mathbf{n}_0\|^2} \mathbf{v}_3 \end{cases} \quad (42)$$

3. Finally, if $\|\mathbf{n}_0\| > 1$, there is no solution with null objective to problem (37). It is tempting to choose the approximate solution $\mathbf{n}^+ = \mathbf{n}^- = \frac{\mathbf{n}_0}{\|\mathbf{n}_0\|}$, but this is not the real solution to (37). It is proven in [34] that this solution writes as:

$$\mathbf{n}(\lambda_0) = \frac{\sigma_1 \mathbf{u}_1^\top \mathbf{i}}{\sigma_1^2 + \lambda_0} \mathbf{v}_1 + \frac{\sigma_2 \mathbf{u}_2^\top \mathbf{i}}{\sigma_2^2 + \lambda_0} \mathbf{v}_2 \quad (43)$$

which reduces to \mathbf{n}_0 when $\lambda_0 = 0$, yet λ_0 is in fact the unique positive solution in λ to the following secular equation:

$$\left(\frac{\sigma_1 \mathbf{u}_1^\top \mathbf{i}}{\sigma_1^2 + \lambda} \right)^2 + \left(\frac{\sigma_2 \mathbf{u}_2^\top \mathbf{i}}{\sigma_2^2 + \lambda} \right)^2 - 1 = 0 \quad (44)$$

Eq. (44) could be rewritten under the form of an algebraic equation of degree 4 in λ , and solved using the Ferrari-Cardan formulae. Instead, we used in our implementation a Newton method. However, it should be reminded that an equation of this type has to be solved at each pixel such that $\|\mathbf{n}_0\| > 1$. If the targetted application has real-time requirements [16], $\mathbf{n}_0/\|\mathbf{n}_0\|$ may be preferred as a fast approximation of the solution to (37), although it is not the exact solution.

5.2. Disambiguating the Problem by Graph Cut

The case $\|\mathbf{n}_0\| < 1$ being by far the most frequent, we can actually build almost $2^{|\Omega|}$ normal fields, which are all solutions to Problem (8). In [35], it is proposed to better constrain the problem assuming that the normals are distributed according to a Laplace law, but this assumption is hard to justify. As discussed in Section 4.5, we would rather advise using the integrability constraint of the normal field, which is much less restrictive (the surface is simply assumed C^2 , at least piecewise). We will now describe the practical resolution of the problem (8) based on this constraint.

²The expressions (42) a posteriori explain the signification of the superscripts $+$ and $-$.

Finding the “most integrable” normal field amounts to finding the one such that the integrability constraint (34) is “best” approximated over Ω . This can be formulated as the variational problem (35). Yet, such a variational problem does not account for the fact that we know explicitly the $2^{|\Omega|}$ possible normal fields. An exhaustive search of the most integrable normal field could be preferred, but it may be computationally infeasible on large data. Instead, we suggest to recast the task of finding the most integrable normal field as a binary labelling problem whose solution can be found efficiently, by means of the graph cut algorithm [12].

Let us attribute to each pixel a label $l \in \{+, -\}$ indicating the normal \mathbf{n}^+ or \mathbf{n}^- , and let us denote by $[p^l, q^l]^\top = [-n_1^l/n_3^l, -n_2^l/n_3^l]^\top$ the corresponding discrete approximation of the surface gradient ∇u . The “optimal” labelling $l : \Omega \rightarrow \{+, -\}$ is the one which makes the normal field the “most integrable”. Thus, we now have to solve the following discrete version of the variational problem (35) over Ω :

$$\min_l \sum_{(x,y) \in \Omega} \left[\frac{\partial p^l}{\partial y}(x, y) - \frac{\partial q^l}{\partial x}(x, y) \right]^2 \quad (45)$$

In order to discretize the space derivatives of Problem (45) using finite differences of order 1, we should consider the four possible *cliques* families of order 3:

$$\begin{cases} \mathcal{C}_1^3 = \{ \{(x, y), (x-1, y), (x, y-1)\} \in \Omega^3 \} \\ \mathcal{C}_2^3 = \{ \{(x, y), (x+1, y), (x, y-1)\} \in \Omega^3 \} \\ \mathcal{C}_3^3 = \{ \{(x, y), (x-1, y), (x, y+1)\} \in \Omega^3 \} \\ \mathcal{C}_4^3 = \{ \{(x, y), (x+1, y), (x, y+1)\} \in \Omega^3 \} \end{cases} \quad (46)$$

If we denote (l_1, l_2, l_3) the labels of each clique’s three pixels, in the same order they are defined in (46), Problem (45) can be rewritten as:

$$\min_l \sum_{c^3 \in \mathcal{C}^3} V_{c^3}^{\text{int}}(l_1, l_2, l_3) \quad (47)$$

where $\mathcal{C}^3 = \cup_{i=1}^4 \mathcal{C}_i^3$, and *potential* $V_{c^3}^{\text{int}}(l_1, l_2, l_3)$ gives the local integrability for the current clique c^3 and the current labeling l . For example, if $c_1^3 \in \mathcal{C}_1^3$:

$$V_{c_1^3}^{\text{int}}(l_1, l_2, l_3) = \left[(p^{l_1}(x, y) - p^{l_3}(x, y-1)) - (q^{l_1}(x, y) - q^{l_2}(x-1, y)) \right]^2 \quad (48)$$

Problem (47) is a labeling problem where the local potential depends on the current pixel and on two of its neighbours. Such combinatorial optimization problems have been studied in [36], where it has been proven that the graph cut algorithm [12] can be used to minimize the energy $\sum_{c^3 \in \mathcal{C}^3} V_{c^3}^{\text{int}}(l_1, l_2, l_3)$, provided that its *regularity* (sub-modularity) is ensured, which means here:

$$\begin{cases} V_{c^3}^{\text{int}}(+, +, l_3) + V_{c^3}^{\text{int}}(-, -, l_3) \leq V_{c^3}^{\text{int}}(+, -, l_3) + V_{c^3}^{\text{int}}(-, +, l_3) \\ V_{c^3}^{\text{int}}(+, l_2, +) + V_{c^3}^{\text{int}}(-, l_2, -) \leq V_{c^3}^{\text{int}}(+, l_2, -) + V_{c^3}^{\text{int}}(-, l_2, +) \\ V_{c^3}^{\text{int}}(l_1, +, +) + V_{c^3}^{\text{int}}(l_1, -, -) \leq V_{c^3}^{\text{int}}(l_1, +, -) + V_{c^3}^{\text{int}}(l_1, -, +) \end{cases} \quad (49)$$

for any $c^3 \in \mathcal{C}^3$ and any $(l_1, l_2, l_3) \in \{+, -\}^3$. Of course, these inequalities have no reason to be satisfied.

5.3. Ensuring the Regularity Condition

To ensure the regularity condition, we modify the problem (47) by introducing a regularization term of the Ising type:

$$\min_l \sum_{c^3 \in \mathcal{C}^3} V_{c^3}^{\text{int}}(l_1, l_2, l_3) + \sum_{c^2 \in \mathcal{C}^2} V_{c^2}^{\text{Ising}}(l_1, l_2) \quad (50)$$

where $\mathcal{C}^2 = \cup_{i=1}^4 \mathcal{C}_i^2$ gathers the four following sets of cliques of order 2:

$$\begin{cases} \mathcal{C}_1^2 = \{ \{(x, y), (x-1, y)\} \in \Omega^2 \} \\ \mathcal{C}_2^2 = \{ \{(x, y), (x, y-1)\} \in \Omega^2 \} \\ \mathcal{C}_3^2 = \{ \{(x, y), (x-1, y-1)\} \in \Omega^2 \} \\ \mathcal{C}_4^2 = \{ \{(x, y), (x+1, y-1)\} \in \Omega^2 \} \end{cases} \quad (51)$$

and $V_{c^2}^{\text{Ising}}(l_1, l_2)$ is defined as follows, where (l_1, l_2) denote the labels of the two pixels of each clique $c^2 \in \mathcal{C}^2$, in the same order they are defined in (51):

$$V_{c^2}^{\text{Ising}}(l_1, l_2) = \beta_{c^2} \delta(l_1 \neq l_2) \quad (52)$$

In Eq. (52), β_{c^2} is a positive or zero *local* coefficient i.e., such a coefficient must be fixed for each $c^2 \in \mathcal{C}^2$. Enforcing the regularity condition for (50), we deduce a lower bound for each β_{c^2} . For example, if $c_3^2 = \{(x, y), (x-1, y-1)\} \in \mathcal{C}_3^2$:

$$\beta_{c_3^2} \geq \frac{1}{2} \max \left\{ 0, \max_{(l_3, l_4) \in \{+, -\}^2} \left\{ \Delta V_{c_3^2}^{\text{int}}(l_3), \Delta V_{c_3^2}^{\text{int}}(l_4) \right\} \right\} \quad (53)$$

where $c_3^3 = c_3^2 \cup (x, y - 1)$, $c_2^3 = c_3^2 \cup (x - 1, y)$, and (l_3, l_4) denote the labels of $(x, y - 1)$ and $(x - 1, y)$, respectively. Finally, for $j \in \{2, 3\}$ and $k \in \{+, -\}$:

$$\Delta V_{c_j^3}^{\text{int}}(k) = V_{c_j^3}^{\text{int}}(k, +, +) + V_{c_j^3}^{\text{int}}(k, -, -) - V_{c_j^3}^{\text{int}}(k, +, -) - V_{c_j^3}^{\text{int}}(k, -, +) \quad (54)$$

From a Markovian point of view, our approach consists in using non-stationary Ising models. The use of such models corresponds to a piecewise uniform prior on the labeling l . We know from Section 4 that there exist a finite number of connected areas over which the solutions \mathbf{n}^+ and \mathbf{n}^- are different (these areas are bounded by the set Ω_R). If these normal fields are estimated as indicated in Section 5.1, they will be continuous inside each area, so the optimal labeling should change only along their boundaries. The Ising prior is thus consistent with the theoretical analysis conducted in Section 4.

Yet, in order not to bias the 3D-reconstruction which should rather be guided by integrability, this prior should have as little influence as possible: according to inequalities like (53), it is possible to predict, for each clique $c^2 \in \mathcal{C}^2$, the smallest value of coefficient β_{c^2} to ensure the regularity condition. These coefficients should therefore not be considered as parameters, which is an advantage. This allows us to limit the energy regularization, only in order to ensure regularity in the sense of Kolmogorov [36], thus avoiding oversmoothing. When coefficients β_{c^2} are not fixed to their minimum values, the labeling is oversmoothed, which biases the 3D-reconstruction (see Fig. 7). Our approach thus avoids the difficulty of tuning the regularization parameters, which is a problem with existing methods (cf. Figs. 8 and 9).

5.4. An Efficient Method for Solving the PS2 Problem

To sum up, the method of resolution of the PS2 problem that we recommend comprises two stages:

1. The calculation of normal fields \mathbf{n}^+ and \mathbf{n}^- as indicated in Section 5.1.
2. The disambiguation of the problem using the integrability criterion, an 8-connected Ising model, the minimum values of the coefficients β_{c^2} like for instance (53), and the graph cut algorithm.

Both steps can be efficiently conducted: the initial solutions \mathbf{n}^+ and \mathbf{n}^- are explicit, and optimization by graph cut is very efficient³.

³It is stated in [12] that “the running time is nearly linear in practice”.




Figure 7: The oversmoothing of the labeling obtained using high values of coefficients β_{c^2} , on the dataset from Fig. 6. From left to right: obtained labeling (black for \mathbf{n}^+ , white for \mathbf{n}^-), XOR map between the estimated labels and the ground truth one (τ indicates the percentage of wrong labels), estimated normal field, and absolute angular error in degrees (MAE is the mean angular error). The choice $\beta_{c^2} = \beta_{\min}$ indicates minimal coefficients (cf. (53)).

The 3D-reconstruction shown in the first row of Fig. 7 is obtained applying this method to the example of Fig. 6. The normal field is very similar to that of the ground truth: when the albedo is known and in the absence of shadow, our method is able to recover almost exactly the genuine normals.

In addition to being parameter-free and not requiring a boundary condition, Figs. 8 and 9 show us that the proposed method is able to confine estimation errors, which may be useful if a shadow is present in one (or both) of the images. This is in contrast with existing methods, for which an estimation error in one pixel may propagate to its neighbors.

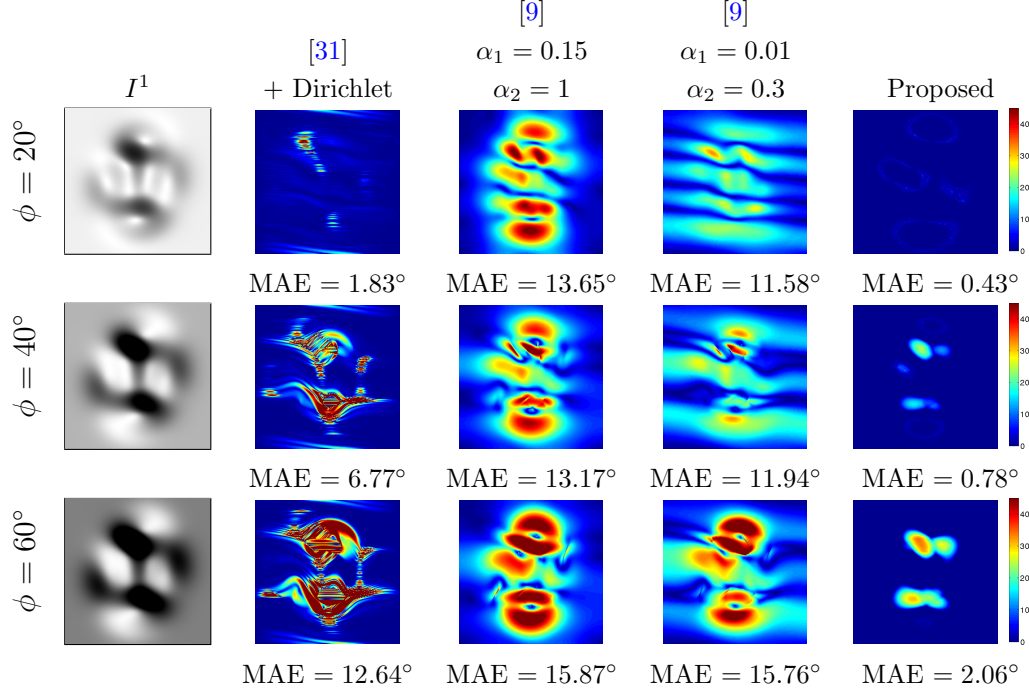


Figure 8: Angular errors on the normals. The dataset from Fig. 5 was considered while increasing the zenithal angles ϕ for the lightings, in order to create stronger shadowing effects. Our method confines the errors in the shadow areas, while requiring neither a boundary condition nor parameter tuning (α_1 and α_2 are defined in Eq. (15)).




Figure 9: 3D-reconstructions corresponding to the second row in Fig. 8.

Yet, our approach has one major weakness with respect to existing methods. Indeed, we assume that the albedo is known. In contrast, since existing methods [9, 31] are based on image ratios, knowledge of the albedo is not required. As a consequence, their methods are robust to unpredicted albedo variations, while ours is not. As it is illustrated in Figs. 10 and 11, this represents an important limitation of our method. To conclude, our method yields overall more accurate 3D-reconstructions than existing ones, and is more flexible (neither need for boundary condition nor parameter tuning), but only when the albedo is known. This restriction is also recurrent in shape-from-shading, where albedo estimation must be carried out beforehand using interpolation techniques [37, 38], or within the 3D-reconstruction process by introducing priors on the albedo and the shape [39].

For sake of completeness, well-posedness of the PS2 problem with unknown albedo is briefly discussed hereafter.




Figure 10: The 256×256 images were created with ρ everywhere equal to 1 except on a diagonal band with width d_ρ pixels where it was set to 0.9. Assuming (wrongly) that $\rho \equiv 1$ induces a bias with our method, while existing methods are independent from the albedo.

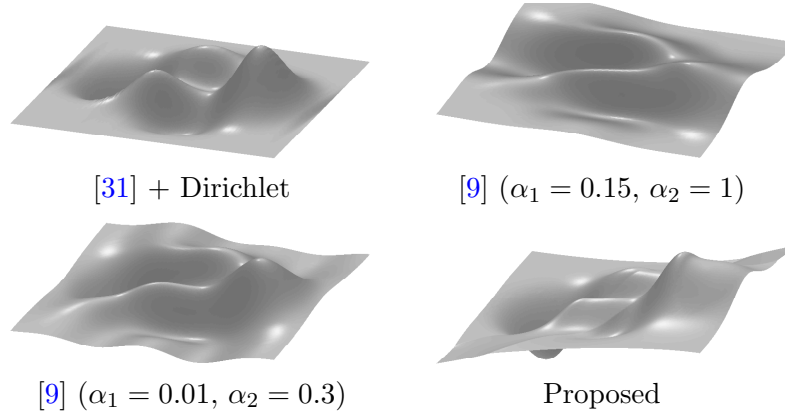


Figure 11: 3D-reconstructions corresponding to the second row in Fig. 10.

5.5. PS2 Problem with Unknown Albedo

At this stage, it is interesting to quote a variant of the PS2 problem, where the surface to be reconstructed has an unknown albedo. The presence of a supplementary unknown $\rho(x, y)$ at each point $(x, y) \in \Omega$ could drastically complicate the problem. A main feature of the PS3 problem is that the albedo can be univocally estimated [6]. This is not the case for SFS since, even if the albedo is known, the problem is usually ill-posed [40]. We will see that the PS2 problem with unknown albedo is more similar to SFS than to PS3.

Let us first give a geometric interpretation of (11). Since $\mathbf{s} = [\tilde{\mathbf{s}}^\top, s_3]^\top$ and $\mathbf{n}(x, y)$ is parallel to $[-\nabla u(x, y)^\top, 1]^\top$, this equation is rewritten as:

$$[I^2(x, y) \mathbf{s}^1 - I^1(x, y) \mathbf{s}^2] \cdot \mathbf{n}(x, y) = 0 \quad (55)$$

Eq. (55) can be directly derived from the two first equations of (8). It means that $\mathbf{n}(x, y)$ lies within a plane $\chi(x, y)$ which is orthogonal to the vector $I^2(x, y) \mathbf{s}^1 - I^1(x, y) \mathbf{s}^2$. Naturally, $\chi(x, y)$ is orthogonal to the plane π as well, which is supported by the vectors \mathbf{s}^1 and \mathbf{s}^2 . It is noteworthy that Eq. (55) holds true for any albedo value $\rho(x, y)$, because Eq. (11) has been derived from Problem (10) by elimination of $\rho(x, y)$.

We know from Section 3.1 that, for any known albedo value $\rho(x, y)$, Problem (8) has two solutions $\mathbf{n}^+(x, y)$ and $\mathbf{n}^-(x, y)$ which are symmetric with respect to π . The geometric interpretation of Eq. (55) means that for a fixed pair of graylevels $(I^1(x, y), I^2(x, y))$, different values of $\rho(x, y)$ will

provide different values of $\mathbf{n}^+(x, y)$ and $\mathbf{n}^-(x, y)$ which all lie within $\chi(x, y)$. Fig. 12 shows us that $\mathbf{n}^+(x, y)$ and $\mathbf{n}^-(x, y)$ move towards plane π as $\rho(x, y)$ decreases from 1 towards a limit value $\rho_{\inf}(x, y)$ which corresponds to the limit case $\mathbf{n}^+(x, y) = \mathbf{n}^-(x, y)$. Hence, at each (x, y) , a range of values $[\rho_{\inf}(x, y), 1]$ are feasible for $\rho(x, y)$. If $\mathbf{s}^1 \neq \mathbf{s}^2$, it can be shown that this limit is:

$$\rho_{\inf}(x, y) = \sqrt{\frac{I^1(x, y)^2 + I^2(x, y)^2 - 2 I^1(x, y) I^2(x, y) (\mathbf{s}^1 \cdot \mathbf{s}^2)}{1 - (\mathbf{s}^1 \cdot \mathbf{s}^2)^2}} \quad (56)$$

Hence, the PS2 problem with unknown albedo is ill-posed. However, it is worth underlining that, when a boundary condition is available, Problem (12) can be solved without any knowledge of $\rho(x, y)$. Moreover, $\rho(x, y)$ can be a posteriori calculated in this case, using any of the PDEs of (10).




Figure 12: Normals $\mathbf{n}^+(x, y)$ and $\mathbf{n}^-(x, y)$ move towards π as $\rho(x, y)$ decreases from 1 towards a minimum value $\rho_{\inf}(x, y)$ which corresponds to the limit case $\mathbf{n}^+(x, y) = \mathbf{n}^-(x, y)$.

As an example, let us reconsider the example of Section 3.3, under the assumption that $\rho(x, y)$ is unknown. The non-differential problem (16) becomes:

$$\begin{cases} \rho(x, y) n_3(x, y) = \frac{1}{\sqrt{2}} \\ \rho(x, y) \frac{n_1(x, y) + n_2(x, y) + \sqrt{2} n_3(x, y)}{2} = \frac{\sqrt{2} - 1}{2\sqrt{2}} \\ n_1(x, y)^2 + n_2(x, y)^2 + n_3(x, y)^2 = 1 \end{cases} \quad (57)$$

For any given value of $\rho(x, y)$, the two first equations of (57), which can be rewritten as $n_3(x, y) = \frac{1}{\sqrt{2}\rho(x, y)}$ and $n_1(x, y) + n_2(x, y) = \frac{-1}{\sqrt{2}\rho(x, y)}$, admit an infinity of solutions:

$$\mathbf{n}(x, y) = \frac{1}{\rho(x, y)} \left\{ \frac{1}{\sqrt{2}} \begin{bmatrix} -1 \\ 0 \\ 1 \end{bmatrix} + t \begin{bmatrix} -1 \\ 1 \\ 0 \end{bmatrix} \right\} \quad (58)$$

which depend on a real parameter t . Replacing the expression (58) of $\mathbf{n}(x, y)$ in the third equation of (57), we find after some algebra the following result:

$$\mathbf{n}(x, y) = \frac{1}{2\sqrt{2}\rho(x, y)} \begin{bmatrix} -1 - \epsilon \sqrt{4\rho(x, y)^2 - 3} \\ -1 + \epsilon \sqrt{4\rho(x, y)^2 - 3} \\ 2 \end{bmatrix} \quad (59)$$

where $\epsilon = \pm 1$. In (59), the albedo value $\rho(x, y)$ can be arbitrarily chosen, provided that $4\rho(x, y)^2 - 3 \geq 0$. This means that $\rho(x, y) \geq \sqrt{3}/2$, which is actually the limit value $\rho_{\inf}(x, y)$ given in (56). Clearly, this problem is ill-posed since at each point $(x, y) \in \Omega$, for each value $\rho(x, y) \in [\sqrt{3}/2, 1]$, there are two possible normals, due to the two possible values of ϵ in (59).

To conclude, with the non-differential approach that we follow in this article, the PS2 problem can be solved unambiguously only if the albedo is known beforehand. Nevertheless, let us now show that solving the PS2 problem in an efficient way is not a purely formal challenge.

6. An Application: Improving Three-source Photometric Stereo

We have seen in Section 5 that the accuracy of PS2 strongly depends on the presence of shadows. The simplest way to ensure robustness to shadows is to consider a third image, i.e. the PS3 problem. By placing the lights appropriately, one can ensure that each surface point is lit in at least two out of the three images, and resort to a combination of the PS2 and the PS3 techniques. This can improve a lot the accuracy of real-time PS based on color photometric stereo [16].

6.1. The Recurrent Problem of Shadows

The three photographs of the first row of Fig. 13, which are available on the web⁴, show a plaster bust of Beethoven illuminated by three non-coplanar, parallel and uniform light beams. Since the light vectors are provided, these real data are particularly well adapted to the PS3 technique (see Section 2.2). In addition, let us remark that no boundary condition is available.

Solving at each point $(x, y) \in \Omega$ a linear system of type (6), then integrating the normals using [13], we indeed obtain a “satisfactory” 3D-shape (see Fig. 14), but this is contradicted by the estimated albedo (see Fig. 13-d), which should be uniform since the material is homogeneous. We also note that the points where the albedo estimate is biased lie inside the shadows.

The problem of dealing with shadows in PS is well-known. Because they constitute an unavoidable departure from the Lambertian model (2), shadow graylevels are usually considered as outliers. In this view, most contributions [24, 25, 26, 27, 28, 29] assume that $m > 3$ images are available, which is not the case here. Only [9] considers the case $m = 3$, and thus “two-source photometric stereo [...] in the presence of shadows”.

Based on a simple shadow detection (we used the same graph cut-based approach as in [41]), Fig. 13-e shows how Ω can be split into four subsets: Ω_3 , $\Omega_2^{1,2}$, $\Omega_2^{1,3}$ and $\Omega_2^{2,3}$, with straightforward notations ($\Omega_2^{1,3}$ is empty in this example). A PS3/PS2 combination can then be considered: the PS2 technique can be used over the pixels lit in only two images, and the PS3 technique elsewhere. The PS3 solution is also used for the points shadowed in more than one image, for stability reasons, yet it would probably be possible to define three other subsets Ω_1^1 , Ω_1^2 and Ω_1^3 and resort to SfS over these sets.

⁴<http://www.ece.ncsu.edu/imaging/Archives/ImageDataBase/Industrial/>




Figure 13: (a-b-c) Three photographs of a plaster bust of Beethoven. (d) Albedo estimated using the PS3 technique, which is biased in the shadow areas. (e) Partition of Ω : Ω_3 (white), $\Omega_2^{1,2}$ (red), and $\Omega_2^{2,3}$ (blue). PS2 is applied in the Ω_2 sets, and PS3 elsewhere.



Figure 14: 3D-shape reconstructed from the three photographs of the top line of Fig. 13, using different techniques. The PS2 technique is applied to the pair of images (I^1, I^2) .

To apply the method described in Section 5.4 to the PS3/PS2 combination, we first need to estimate the albedo. Knowing that it is uniform since the material is homogeneous, this can be carried out by evaluating the histogram peak of the estimated albedo inside the set Ω_3 , which can be considered as the “real” albedo: we obtained the value $\rho = 0.74$.

As shown in Fig. 14, using only two images yields biased results because PS2 is not robust to shadows. Yet, as shown in Fig. 15, combining the PS3 and PS2 techniques can improve a lot the accuracy of the 3D-reconstruction. In addition, it seems that using the proposed PS2 framework inside this combination improves the results, in comparison with state-of-the-art.

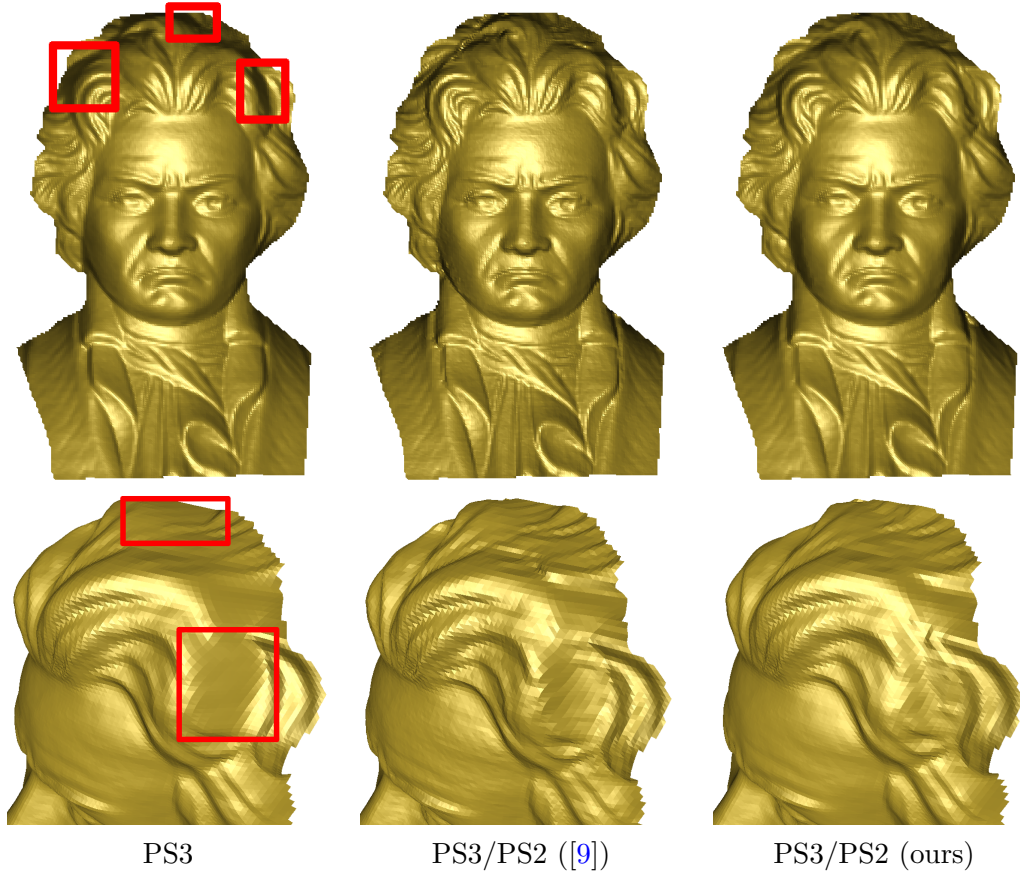


Figure 15: Frontal (top) and side (bottom) views of the 3D-reconstructions obtained using PS3, and two versions of the combination PS3/PS2. Ours is able to recover fine-scale details more accurately in the shadow areas.

In order to quantitatively assess the accuracy of the proposed PS2 framework, we extracted three images from the dataset [42], which contains 96 images in total, and for which ground truth normals are available. Our PS2 method was applied only over a shadow area, indicated in red, where the albedo is approximately uniform. The proposed method is able to handle shadows even when they are located near the boundaries, which is a known failure case of the differential approach in the absence of a boundary condition [9].




Figure 16: First row: Three images of a Lambertian object, and the $\Omega_3/\Omega_2^{2,3}$ partition. The area indicated in red lies almost entirely in the shadow in the first image. Second row: angular error (in degrees) obtained using, from left to right, the PS3 method (MAE = 11.84°), the PS2/PS3 combination from [9] (MAE = 11.27°), and the proposed one (MAE = 11.57°). The differential approach, which ensures smoothness, is globally more satisfactory than the non-differential one. Third row: 3D-reconstruction results. Fourth row: close-up on the estimated normals over the area indicated in red. The MAE are, respectively, 20.79° , 25.49° and 16.73° . Our method achieves the best 3D-reconstruction in the shadow area.

6.2. Color Photometric Stereo

Hernández et al. argue the following in [9]: “Using photometric stereo on just three images may seem like an unreasonably hard restriction. There is, however, a particular situation when only three images are available. This technique is known as color photometric stereo”. Indeed, the most straightforward application of PS2 is color photometric stereo, a technique where three light sources with different colors and positions are used to simultaneously provide three (graylevel) images of the surface under three different illuminations.

This idea is due to Kontsevich et al., who show in [43] how to reconstruct the 3D-shape of a white painted scene, illuminated by $m = 3$ color light sources. Indeed, the number of channels of a standard color image is three. Considering each channel as a graylevel image, a single RGB image is enough to apply the PS3 technique. A deformable scene such as a face can therefore be reconstructed, even if the person is not standing still: Hernández et al. show in [16] “how multispectral lighting allows one to essentially capture three images (each with a different light direction) in a single snapshot, thus making per-frame photometric reconstruction possible”. An example of such an RGB image extracted from a video sequence⁵ is shown in Fig. 17. Let us nevertheless point out that the albedo map should be the same in each channel. In practice, this requires that the scene is made-up. Since applying make-up also ensures that the albedo is uniform, it can be estimated by evaluating the histogram peak.




Figure 17: (a) An RGB image of a face illuminated by three directional, non-coplanar, color light sources. (b-c-d) Decomposition of the image (a) in three channels (red, green, blue). (e) Partition of Ω into sets Ω_3 (white), $\Omega_2^{R,B}$ (green), and $\Omega_2^{G,B}$ (blue).

⁵<http://mi.eng.cam.ac.uk/research/projects/VideoNormals/>

We can apply the PS3 technique to the image of Fig. 17-a but, without any specific treatment, the result is biased around the nose (see Fig. 18), since the shadow renders the red channel unusable in this area (see Fig. 17-b).

As in Section 6.1, the histogram of the albedo allows us to estimate the (uniform) albedo, and therefore to use the PS2 framework within the PS3/PS2 combination. In comparison with PS3, the result is greatly improved, even if artifacts are visible at the junction between the Ω_3 and Ω_2 areas. The proposed method even qualitatively outperforms the state of the art [9].




Figure 18: 3D-reconstructions obtained from the single RGB image of Fig. 17-a. Only our method is able to restore the nostril area. The artifacts around the shadow area boundary are probably due to the overlapping spectra of the light sources: these artifacts are also visible in the images of Fig. 17-b and Fig. 17-c, although they should not appear.

7. Conclusion and Future Prospects

In this paper, we first compared in Section 3 the non-differential formulation of the PS2 problem with two differential formulations. This comparison led us to provide in Section 4 a theoretical study of PS2. Building upon this study, we chose the non-differential formulation of PS2 for its numerical

resolution, which is detailed in Section 5. With this approach, two normals can explain the graylevels at each point. Our main contribution consists in reformulating the 3D-reconstruction problem as a binary labelling problem, which is solved by means of the graph cut algorithm. The criterion to be optimized is integrability, and a non-stationary Ising term must be added to satisfy the regularity condition.

The straightforward application of this contribution is color photometric stereo, which is explored in Section 6. Our method is validated on RGB images extracted from a video sequence. It improves the 3D-reconstruction fineness, while being quite fast. The real-time requirement expressed in [16] is thus not compromised. We moreover assert that this approach is more justified than regularization schemes. Regularization should not be used for disambiguating the 3D-reconstruction problem, but rather for smoothing the residual noise. In this view, future extensions of our work may include an additional regularization designed to ensure more robustness. Yet, designing computationally-efficient algorithms for regularizing normal fields is not an easy task, because of the unit-length constraint [44]. Recent developments on half-quadratic algorithms designed for manifold-valued images [45] may constitute an interesting option.

Nevertheless, the assumption that the albedo is known greatly limits the applicability of our study. As future work we plan to extend it to the PS2 problem with unknown albedo. A solution to this more complex problem would be, assuming the albedo known and uniform, that is to say $\rho^0 \equiv 1$, to get a first estimate u^0 of the 3D-shape using the method described in Section 5.1. It seems then possible to solve the problem, using the constraint $u \approx u^0$ as an additional regularization, to deduce a new estimate ρ^1 , and then to loop. A natural perspective of this extension is to cope with images taken by a webcam on a sunny day, as in [8], since the sun rays move within a plane and the linear system (6) is not full-rank.

Finally, throughout this article, we made the hypothesis of an orthographic camera. The proposed method of resolution of the PS2 problem could be used as such for a perspective camera, although the constraint of integrability is more complicated [46], as this only serves to disambiguate a combinatorial optimization problem. Obviously, the perspective should be taken into account in the final step, when the normals are integrated [13].

References

- [1] B. K. P. Horn, M. J. Brooks (Eds.), Shape from Shading, MIT Press, 1989. [2](#), [3](#)
- [2] J.-D. Durou, M. Falcone, M. Sagona, Numerical Methods for Shape-from-shading: A New Survey with Benchmarks, Computer Vision and Image Understanding 109 (1) (2008) 22–43. [2](#), [4](#)
- [3] G. Vogiatzis, C. Hernández, Practical 3D Reconstruction Based on Photometric Stereo, in: R. Cipolla, S. Battiato, G. M. Farinella (Eds.), Computer Vision: Detection, Recognition and Reconstruction, Springer-Verlag, 2010, Ch. 12. [2](#)
- [4] J. Ackermann, M. Goesele, A survey of photometric stereo techniques, Foundations and Trends in Computer Graphics and Vision 9 (3-4) (2015) 149–254. [2](#)
- [5] E. Prados, O. Faugeras, F. Camilli, Shape from Shading: a well-posed problem?, Rapport de Recherche 5297, Institut National de Recherche en Informatique et en Automatique, Sophia Antipolis, France (Aug. 2004). [2](#)
- [6] R. J. Woodham, Photometric Method for Determining Surface Orientation from Multiple Images, Optical Engineering 19 (1) (1980) 139–144. [2](#), [4](#), [5](#), [28](#)
- [7] Y. Sato, K. Ikeuchi, Reflectance analysis under solar illumination, in: Proceedings of the Workshop on Physics-Based Modeling in Computer Vision, Cambridge, Massachusetts, USA, 1995, pp. 180–187. [2](#), [5](#), [6](#)
- [8] J. Jung, J.-Y. Lee, I. S. Kweon, One-day outdoor photometric stereo via skylight estimation, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Boston, Massachusetts, USA, 2015, pp. 4521–4529. [2](#), [6](#), [37](#)
- [9] C. Hernández, G. Vogiatzis, R. Cipolla, Overcoming shadows in 3-source photometric stereo, IEEE Transactions on Pattern Analysis and Machine Intelligence 33 (2) (2011) 419–426. [2](#), [6](#), [10](#), [18](#), [26](#), [27](#), [28](#), [31](#), [32](#), [33](#), [34](#), [35](#), [36](#)

- [10] R. Onn, A. M. Bruckstein, Integrability Disambiguates Surface Recovery in Two-Image Photometric Stereo, *International Journal of Computer Vision* 5 (1) (1990) 105–113. [2](#), [6](#), [14](#), [18](#), [19](#)
- [11] R. Mecca, J.-D. Durou, Unambiguous Photometric Stereo Using Two Images, in: *Proceedings of the 16th International Conference on Image Analysis and Processing*, Vol. 6978 of *Lecture Notes in Computer Science*, Ravenna, Italy, 2011, pp. 286–295. [2](#), [6](#), [9](#), [18](#)
- [12] Y. Boykov, O. Veksler, R. Zabih, Fast approximate energy minimization via graph cuts, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 23 (11) (2001) 1222–1239. [3](#), [22](#), [23](#), [24](#)
- [13] J.-D. Durou, J.-F. Aujol, F. Courteille, Integration of a Normal Field in the Presence of Discontinuities, in: *Proceedings of the 7th International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*, Vol. 5681 of *Lecture Notes in Computer Science*, Bonn, Germany, 2009, pp. 261–273. [4](#), [9](#), [19](#), [31](#), [37](#)
- [14] R. T. Frankot, R. Chellappa, A Method for Enforcing Integrability in Shape from Shading Algorithms, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 10 (4) (1988) 439–451. [4](#), [5](#)
- [15] R. Mecca, M. Falcone, Uniqueness and Approximation of a Photometric Shape-from-Shading Model, *SIAM Journal on Imaging Sciences* 6 (1) (2013) 616–659. [5](#), [9](#), [14](#), [16](#)
- [16] C. Hernández, G. Vogiatzis, G. J. Brostow, B. Stenger, R. Cipolla, Non-rigid Photometric Stereo with Colored Lights, in: *Proceedings of the 11th IEEE International Conference on Computer Vision*, Rio de Janeiro, Brazil, 2007. [5](#), [21](#), [31](#), [35](#), [37](#)
- [17] J. Yang, N. Ohnishi, N. Sugie, Two image photometric stereo method, in: *Intelligent Robots and Computer Vision XI: Biological, Neural Net, and 3D Methods*, Vol. 1826 of *Proceedings of the International Society for Optical Engineering*, Boston, Massachusetts, USA, 1992, pp. 452–463. [6](#)
- [18] R. J. Woodham, Photometric Stereo: A Reflectance Map Technique for Determining Surface Orientation from a Single View, in: *Proceedings of*

- the 22nd SPIE Annual Technical Symposium, Vol. 155 of Proceedings of the International Society for Optical Engineering, San Diego, California, USA, 1978, pp. 136–143. [6](#)
- [19] R. Kozera, On Shape Recovery from Two Shading Patterns, *International Journal of Pattern Recognition and Artificial Intelligence* 6 (4) (1993) 673–698. [6](#)
 - [20] O. Ikeda, A Robust Shape-from-shading Algorithm using two Images and Control of Boundary Conditions, in: *Proceedings of the IEEE International Conference on Image Processing (volume I)*, Barcelona, Spain, 2003, pp. 405–408. [6](#)
 - [21] A. Abrams, C. Hawley, R. Pless, Heliometric Stereo: Shape from Sun Position, in: *Proceedings of the 12th European Conference on Computer Vision (volume II)*, Vol. 7573 of *Lecture Notes in Computer Science*, Firenze, Italy, 2012, pp. 357–370. [6](#)
 - [22] J. Ackermann, F. Langguth, S. Fuhrmann, M. Goesele, Photometric Stereo for Outdoor Webcams, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Providence, Rhode Island, USA, 2012, pp. 262–269. [6](#)
 - [23] S. J. Koppal, S. G. Narasimhan, Appearance Derivatives for Isonormal Clustering of Scenes, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 31 (8) (2009) 1375–1385. [6](#)
 - [24] S. Barsky, M. Petrou, The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25 (10) (2003) 1239–1252. [6](#), [31](#)
 - [25] M. K. Chandraker, S. Agarwal, D. J. Kriegman, ShadowCuts: Photometric Stereo with Shadows, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2nd Workshop Towards Benchmarking Automated Calibration, Orientation, and Surface Reconstruction from Images*, Minneapolis, Minnesota, USA, 2007. [6](#), [31](#)
 - [26] J. Sun, M. Smith, L. Smith, S. Midha, J. Bamber, Object surface recovery using a multi-light photometric stereo technique for non-Lambertian sur-

- faces subject to shadows and specularities, *Image and Vision Computing* 25 (7) (2007) 1050–1057. [6](#), [31](#)
- [27] F. Verbiest, L. Van Gool, Photometric stereo with coherent outlier handling and confidence estimation, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Anchorage, Alaska, USA, 2008. [6](#), [31](#)
 - [28] L. Wu, A. Ganesh, B. Shi, Y. Matsushita, Y. Wang, Y. Ma, Robust Photometric Stereo via Low-Rank Matrix Completion and Recovery, in: *Proceedings of the 10th Asian Conference on Computer Vision*, Vol. 6494 of *Lecture Notes in Computer Science*, Queenstown, New Zealand, 2010, pp. 703–717. [6](#), [31](#)
 - [29] S. Ikehata, D. Wipf, Y. Matsushita, K. Aizawa, Photometric Stereo Using Sparse Bayesian Regression for General Diffuse Surfaces, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36 (9) (2014) 1816–1831. [6](#), [31](#)
 - [30] K. Ikeuchi, B. K. P. Horn, Numerical Shape from Shading and Occluding Boundaries, *Artificial Intelligence* 17 (1–3) (1981) 141–184. [7](#), [12](#)
 - [31] Y. Quéau, R. Mecca, J.-D. Durou, Unbiased photometric stereo for colored surfaces: A variational approach, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Las Vegas, USA, 2016, pp. 3707–3716. [9](#), [26](#), [27](#), [28](#)
 - [32] R. Kozera, Existence and Uniqueness in Photometric Stereo, *Applied Mathematics and Computation* 44 (1) (1991) 1–103. [15](#)
 - [33] B. K. P. Horn, M. J. Brooks, The Variational Approach to Shape From Shading, *Computer Vision, Graphics, and Image Processing* 33 (2) (1986) 174–208. [17](#)
 - [34] G. H. Golub, C. F. Van Loan, *Matrix computations*, 4th Edition, Johns Hopkins University Press, 2012, Ch. 6.2.1, pp. 313–314. [20](#), [21](#)
 - [35] W. L. Woo, S. S. Dlay, 3D shape restoration using sparse representation and separation of illumination effects, *Signal Processing* 103 (2014) 258–272. [21](#)

- [36] V. Kolmogorov, R. Zabih, What Energy Functions Can Be Minimized via Graph Cuts?, *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26 (2) (2004) 147–159. [23](#), [24](#)
- [37] F. Courteille, A. Crouzil, J.-D. Durou, P. Gurdjos, Towards shape from shading under realistic photographic conditions, in: *Proceedings of the 17th International Conference on Pattern Recognition*, Cambridge, UK, 2004, pp. 277–280. [27](#)
- [38] Vogel, O. and Valgaerts, L. and Breuß, M. and Weickert, J. and Valgaerts, L., Making Shape from Shading Work for Real-World Images, in: *Proceedings of the 31st symposium of the German Association for Pattern Recognition*, 2009, pp. 191–200. [27](#)
- [39] J. T. Barron, J. Malik, High-Frequency Shape and Albedo from Shading using Natural Image Statistics, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Colorado Springs, Colorado, USA, 2011, pp. 2521–2528. [27](#)
- [40] J. Oliensis, Uniqueness in Shape from Shading, *International Journal of Computer Vision* 6 (2) (1991) 75–104. [28](#)
- [41] Hernández, C. and Vogiatzis, G. and Cipolla, R., Shadows in Three-Source Photometric Stereo, in: *Proceedings of the 11th European Conference on Computer Vision*, Vol. 5302 of *Lecture Notes in Computer Science*, Marseille, France, 2008, pp. 290–303. [31](#)
- [42] B. Shi, Z. Wu, Z. Mo, D. Duan, S.-K. Yeung, P. 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*, Las Vegas, USA, 2016, pp. 4359–4368. [34](#)
- [43] L. L. Kontsevich, A. P. Petrov, I. S. Vergelskaya, Reconstruction of shape from shading in color images, *Journal of the Optical Society of America - Part A: Optics, Image Science, and Vision* 11 (3) (1994) 1047–1052. [35](#)
- [44] Zeisl, B. and Zach, C. and Pollefeys, M., Variational Regularization and Fusion of Surface Normal Maps, in: *Proceedings of the 2nd International Conference on 3D Vision*, Vol. 1, Tokyo, Japan, 2014, pp. 601–608. [37](#)

- [45] Bergmann, R. and Chan, R. H. and Hielscher, R. and Persch, J. and Steidl, G., Restoration of manifold-valued images by half-quadratic minimization, *Inverse Problems and Imaging* 10 (2) (2016) 281–304. [37](#)
- [46] T. Papadhimetri, P. Favaro, A New Perspective on Uncalibrated Photometric Stereo, in: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, Portland, Oregon, USA, 2013, pp. 1474–1481. [37](#)