

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

A comprehensive introduction to photometric 3D-reconstruction

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Durou J. -D., Falcone M., Queau Y., Tozza S. (2020). A comprehensive introduction to photometric 3D-reconstruction. Cham : Springer Nature Switzerland AG [10.1007/978-3-030-51866-0_1].

Availability:

This version is available at: <https://hdl.handle.net/11585/844068> since: 2022-01-05

Published:

DOI: http://doi.org/10.1007/978-3-030-51866-0_1

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Durou, JD., Falcone, M., Quéau, Y., Tozza, S. (2020). A Comprehensive Introduction to Photometric 3D-Reconstruction. In: Durou, JD., Falcone, M., Quéau, Y., Tozza, S. (eds) Advances in Photometric 3D-Reconstruction. Advances in Computer Vision and Pattern Recognition. Springer, Cham

The final published version is available online at
https://dx.doi.org/10.1007/978-3-030-51866-0_1

Rights / License:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

A Comprehensive Introduction to Photometric 3D-reconstruction

Jean-Denis Durou, Maurizio Falcone, Yvain Quéau and Silvia Tozza

Abstract Photometric 3D-reconstruction techniques aim at inferring the geometry of a scene from one or several images, by inverting a physical model describing the image formation. This chapter presents an introductory overview of the main photometric 3D-reconstruction techniques which are shape-from-shading, photometric stereo and shape-from-polarisation.

1 Introduction

Inferring the 3D-shape of a scene is necessary in many applications such as quality control, augmented reality or medical diagnosis. Depending upon the requirements of the application, 3D-estimation can be carried out using a variety of technological solutions, from coordinate measuring machines to X-ray scanners. Over the last decades, digital cameras have become a reliable alternative to such sensors, as they represent a reasonable compromise between resolution and affordability. Given one or several 2D-images of a scene captured by a digital camera, the process of estimating its 3D-shape is called *3D-reconstruction*. It is a classic inverse problem in computer vision, which has been addressed in several ways.

Jean-Denis Durou
IRIT, UMR CNRS 5505, Université de Toulouse, Toulouse, France, e-mail: durou@irit.fr

Maurizio Falcone
Department of Mathematics, Sapienza University of Rome, Rome, Italy, e-mail: falcone@mat.uniroma1.it

Yvain Quéau
CNRS, GREYC Laboratory, ENSICAEN and University of Caen, France, e-mail: yvain.queau@ensicaen.fr

Silvia Tozza
Department of Mathematics and Applications "Renato Caccioppoli", Università degli Studi di Napoli Federico II, Naples, Italy, e-mail: silvia.tozza@unina.it

A 3D-model consists of a set of *geometric* (position, orientation, etc.) and *photometric* (color, texture, etc.) information. Knowing both these pieces of information allows to render synthetic images, by simulating the trajectory of the light rays from the sources to the camera, after reflection on the surface of the scene. 3D-scanning is the dual of rendering: one aims at a geometric and photometric characterisation of the scene’s surface by reversing the trajectory of the light rays. In fact 3D-scanning includes both the subproblems of 3D-reconstruction (estimating the scene’s geometry) and appearance estimation (estimating its photometric properties).

The various 3D-reconstruction techniques from digital cameras are grouped under the generic terminology *shape-from-X*, X indicating that shape estimation can be based on various clues (shadows, contours, etc.). The main shape-from-X techniques are presented in Table 1. In this table, they are classified according to the clue they are based on (photometric or geometric) and the number of images they require.

| | Geometric techniques | Photometric techniques |
|--------------|-----------------------------|------------------------------------|
| Single image | Structured light [39] | Shape-from-shading (SfS) [48] |
| | Shape-from-shadows [105] | |
| | Shape-from-contours [15] | |
| | Shape-from-texture [125] | |
| | Shape-from-template [10] | |
| Multi-images | Structure-from-motion [78] | Photometric stereo (PS) [129] |
| | Stereopsis [43] | Shape-from-polarisation (SfP) [96] |
| | Shape-from-silhouettes [44] | |
| | Shape-from-focus [81] | |

Table 1: Main shape-from-X techniques. Geometric techniques aim at identifying and analysing features. This presentation rather focuses on photometric techniques, which aim at inverting a physics-based image formation model.

Geometric shape-from-X techniques are built upon the identification of features in the image(s). On the other hand, photometric shape-from-X techniques are based on the analysis of the quantity of light received in each photosite of the camera’s sensor. Photometric 3D-reconstruction techniques indeed rely on a physics-based forward image formation model describing the interactions between light, matter and the camera, and aim at inverting this model in order to infer the geometry of the scene and, possibly, its photometric properties.

There exist out-of-the box solutions for geometric 3D-reconstruction, e.g. Microsoft Kinect (based on stereopsis or structured light, depending on the version), or the CMPMVS [54] or AliceVision [2] projects (based on structure-from-motion and stereopsis). On the contrary, there is a lack of such solutions for photometric techniques, which are usually rather viewed as “lab” reconstruction techniques because they rely on several assumptions on the acquisition setup. Still, they bear great promises in terms of level of geometric details which can be recovered, and of applicability to a wide range of materials.

The aim of this chapter is to present an overview of the three main photometric shape-from-X techniques: shape-from-shading, photometric stereo and shape-from-polarisation. We first review in Sect. 2 the shape-from-shading problem, which is a computer vision technique consisting in inferring geometry from a single image. Then, we discuss two techniques where multiple images are analysed under controlled incident or reflected lighting. In photometric stereo (Sect. 3), a series of images are acquired under varying incident lighting, which permits to estimate both the shape and the reflectance of the pictured surface. In shape-from-polarisation (Sect. 4), it is the state of polarisation of the reflected light which is analysed, by considering a series of images acquired with a controllable polarising filter attached to the camera. Sect. 5 eventually concludes this study by presenting the subsequent chapters of this volume.

2 Shape-from-shading

Inferring 3D-geometry from a single image of a shaded surface is a problem known as *shape-from-shading*. This technique was first developed in the seventies at MIT, under the impulse of Horn [48].

2.1 Non-differential SfS Models

Let us briefly outline the problem, attaching to the camera a 3D-coordinate system $Oxyz$, such that Oxy coincides with the image plane and Oz with the optical axis. Assuming orthographic projection, the visible part of the scene is, up to a scale factor, a graph $z = u(\mathbf{x})$, where $\mathbf{x} = [x, y]^T$ is an image point. The SfS problem can be modeled by the *image irradiance equation* [49]:

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

where $I(\mathbf{x})$ is the graylevel at point \mathbf{x} (in fact, $I(\mathbf{x})$ is the *irradiance* at point \mathbf{x} , but both quantities are proportional), and the *radiance* $R(\mathbf{n}(\mathbf{x}))$ gives the value of the light re-emitted by the surface as a function of its orientation, i.e., of the unit normal $\mathbf{n}(\mathbf{x})$ to the surface at the 3D-point $[x, y, u(\mathbf{x})]^T$ conjugate with \mathbf{x} (cf. Fig. 1). Assuming that, besides I , the radiance function R is also known, then solving Eq. (1) is a non-differential model of SfS, in which the unknown is the normal $\mathbf{n}(\mathbf{x})$.

Let us assume there is a unique light source at infinity, whose direction is characterised by the unit vector $\omega = [\omega_1, \omega_2, \omega_3]^T \in \mathbb{R}^3$, and whose intensity is denoted by $\psi(\mathbf{x})$. Let us also assume for simplicity that the surface is *Lambertian*, i.e., an ideal diffuse reflecting surface for which the apparent brightness is independent from the viewing angle. Then, R is written in such a way that Eq. (1) becomes:

$$I(\mathbf{x}) = r(\mathbf{x}) \psi(\mathbf{x}) \omega^T \mathbf{n}(\mathbf{x}). \quad (2)$$




Fig. 1: The surface is represented as a graph $z = u(\mathbf{x})$, where $\mathbf{x} = [x, y]^T$ is an image point in the reconstruction domain Ω . The normal at the surface point $[x, y, u(\mathbf{x})]^T$ conjugate with \mathbf{x} is denoted by $\mathbf{n}(\mathbf{x})$, and the incident light direction by ω .

In Eq. (2), $r(\mathbf{x})$ is the *reflectance* (or *albedo*), and the scalar product $\psi(\mathbf{x}) \omega^T \mathbf{n}(\mathbf{x})$ is called *shading*. This is another example of non-differential SfS model.

Equation (2) is fundamentally ill-posed, according to the *trompe-l'œil* principle, which is well illustrated by Adelson and Pentland's “workshop metaphor” [1] (cf. Fig. 2). If a painter, a light designer and a sculptor are asked to design an artwork explaining a given image $I(\mathbf{x})$, they may propose very different, but plausible, solutions. The painter will assume a planar surface and a uniform lighting, the changes in intensity being explained by changes in reflectance $r(\mathbf{x})$. The light designer may propose a sophisticated lighting configuration $\psi(\mathbf{x})$ placed in front of a planar surface with uniform reflectance. Eventually, the sculptor will assume lighting and reflectance are uniform, and explain the changes in intensity solely by the shading, which results from variations in the local orientation $\mathbf{n}(\mathbf{x})$ of the surface.

This last explanation, which comes down to inverting Eq. (2) in order to infer a 3D-shape, assuming everything is known but $\mathbf{n}(\mathbf{x})$, is precisely the shape-from-shading problem. So it is assumed that the reflectance is known, which is usually written $r(\mathbf{x}) \equiv 1$, and that the lighting is uniform, i.e., $\psi(\mathbf{x}) \equiv 1$.




Fig. 2: Adelson and Pentland's "workshop metaphor" [1]. To explain an image (a) in terms of reflectance, lighting and shape, (b) a painter, (c) a light designer and (d) a sculptor will design three different, but plausible, solutions. Inferring the shape (d) from a single image is the shape-from-shading problem.

2.2 Differential SfS Models

Let us now turn to differential SfS models. Under orthographic projection, the normal is easily expressed as:

$$\mathbf{n}(\mathbf{x}) = \frac{1}{\sqrt{1 + p(\mathbf{x})^2 + q(\mathbf{x})^2}} [-p(\mathbf{x}), -q(\mathbf{x}), 1]^\top, \quad (3)$$

where:

$$p := \frac{\partial u}{\partial x} \quad \text{and} \quad q := \frac{\partial u}{\partial y}, \quad (4)$$

so that $\nabla u(\mathbf{x}) = [p(\mathbf{x}), q(\mathbf{x})]^\top$. It is easily deduced from Eqs. (2) and (3), assuming $r(\mathbf{x}) \equiv 1$ and $\psi(\mathbf{x}) \equiv 1$, that the following equation holds true for a general parallel lighting whose direction is characterised by $\omega = [\omega_1, \omega_2, \omega_3]^\top$:

$$I(\mathbf{x})\sqrt{1 + |\nabla u(\mathbf{x})|^2} + [\omega_1, \omega_2] \nabla u(\mathbf{x}) - \omega_3 = 0. \quad (5)$$

This is a first-order nonlinear partial differential equation (PDE) of Hamilton-Jacobi type, which constitutes an example of differential SfS model, in which the unknown is now the function u , called the *height map*. This equation has to be solved on a compact domain $\Omega \subset \mathbb{R}^2$, called the *reconstruction domain*.

The PDE which appears in most of the papers on SfS corresponds to a frontal lighting, i.e., $\omega = [0, 0, 1]^\top$. This assumption leads to the *eikonal equation*, which is a particular case arising from the differential SfS model (5):

$$|\nabla u(\mathbf{x})| = f(\mathbf{x}) := \sqrt{\frac{1}{I(\mathbf{x})^2} - 1}, \quad (6)$$

where the graylevel function I , which typically takes integer values between 0 and 255, is implicitly resampled to take real values in the range (0, 1).

Note that even in the most simple case (the eikonal equation) we get a nonlinear PDE of the first order, and the solutions are a priori non-differentiable and non-unique, even if we complement Eq. (6) with a Dirichlet boundary condition, i.e., $u = g$ on the boundary $\partial\Omega$ of Ω . Moreover, the right-hand side of the eikonal equation is not always defined because $I(\mathbf{x})$ can vanish at some points. A simple example is the hemisphere $z = \sqrt{1 - x^2 - y^2}$ under a parallel frontal lighting. In this case, $I(\mathbf{x}) = 0$ at the equator. However, in that simple situation the boundary condition $u = 0$ can help to solve the problem.

Under oblique light direction the same example becomes more difficult because there will be a *black shadow region* $\Omega_s \subset \Omega$ where $I(\mathbf{x}) \equiv 0$, and in that region the model has no information to reconstruct the surface. The boundary of Ω_s , which is not known *a priori*, is a curve where it would be difficult to impose boundary conditions in the numerical approximation. In general, the curve separating the region Ω_s will depend on the shape of the surface and on the light source direction ω . Note that in case of black shadows, the model is clearly unable to produce a reasonable surface approximation, because the information is missing. In this situation, one can follow a global approach avoiding to impose boundary conditions on $\partial\Omega_s$. This leads to the concept of *maximal solution*, where we solve the PDE on the whole domain with the standard Dirichlet boundary condition on $\partial\Omega$ (not on $\partial\Omega_s$), and recover a linear reconstruction on Ω_s (we refer to [34, 35, 17] for more details).

2.3 Ill-posedness of the SfS Models

The “workshop metaphor” illustrated in Fig. 2 is representative of the ill-posedness of SfS, because a posteriori estimating 3D-geometry from a single image is possible only if reflectance and lighting are known a priori. The reliability of these priors is of fundamental importance to guarantee that the solution of SfS is meaningful.

This is illustrated in Fig. 3, which shows how the assumptions leading to Eq. (6), i.e., a uniform reflectance ($r(\mathbf{x}) \equiv 1$) and a parallel uniform lighting ($\psi(\mathbf{x}) \equiv 1$ and $\omega = [0, 0, 1]^\top$), which is just a rough approximation in the case of Fig. 3-a, yield the erroneous interpretation of the scene 3D-shape shown in Fig. 3-b. However, this solution is an exact solution of Eq. (6), as shown by frontally relighting this uniformly white 3D-shape (cf. Fig. 3-c).

Even when reflectance and lighting are known, i.e., when $r(\mathbf{x}) \equiv 1$ and $\psi(\mathbf{x}) \equiv 1$, the non-differential model (2) of SfS remains ill-posed:

$$I(\mathbf{x}) = \omega^\top \mathbf{n}(\mathbf{x}). \quad (7)$$

Except for some sparse *singular points*, where $\mathbf{n}(\mathbf{x})$ points in the same direction as ω , there exists an infinity of surface normals explaining the graylevel in one pixel. It comes down from Eq. (7) that these normals $\mathbf{n}(\mathbf{x})$ form a revolution cone around the lighting direction ω . It is thus very difficult to *locally* solve SfS.




Fig. 3: Illustration of the importance of reflectance and lighting priors on the solution of SfS [25]. (a) A well-known graylevel image $I(\mathbf{x})$. (b) 3D-shape estimated by solving the SfS model (6), by (wrongly) assuming uniform reflectance $r(\mathbf{x}) \equiv 1$, and uniform frontal lighting, i.e., $\psi(\mathbf{x}) \equiv 1$ and $\omega = [0, 0, 1]^\top$. The 3D-reconstruction largely departs from the true geometry. Yet, by taking from above a picture of the uniformly white 3D-shape (b), using the camera's flash as single light source, one gets image (c), which resembles (a). 3D-shape (b) is thus a plausible explanation of (a): the bias comes from the inappropriate reflectance and lighting priors.

A simple example in dimension 1 is given by the surface $z = u_1(x) = 1 - x^2$ in the interval $(-1, 1)$ (cf. Fig. 4-a), under vertical lighting $\omega = [0, 1]^\top$, which satisfies an equation of the form (6) and the homogeneous boundary condition $u_1(1) = u_1(-1) = 0$. However, the function $u_2(x) = -u_1(x)$ in the same interval still satisfies this equation, since $|\nabla u(\mathbf{x})|$ is the same, and the same boundary condition. This example is an illustration of the famous *concave/convex ambiguity* of SfS.

Note that u_1 and u_2 are two differentiable solutions to the same problem. If we decide to accept also solutions which are differentiable only almost everywhere (a very natural choice in view of real applications), we suddenly have an infinite number of solutions which can be obtained just considering all the possible reflections of one of those solutions, e.g. u_1 , with respect to a horizontal axis located at the height $z = h$, where $h \in (0, 1)$. This is illustrated in Fig. 4-b, where three such solutions are exhibited. If one fixes the height at the singular point, then only one of these solutions can be accepted (we refer to [64] for this result), but such an additional knowledge is clearly not very realistic. A general theory for weak solutions of Hamilton-Jacobi equations (that includes the eikonal equation) has been developed in the last twenty years starting from the seminal paper by Crandall and Lions [24]. We refer the interested reader to the book [8] and the references therein.

Practical ways to reduce this ambiguity include resorting to more realistic models such as perspective camera [117] or near-lighting [89], however it has been shown that this remains insufficient to ensure well-posedness [16]. Recently, the introduction of an attenuation factor in the brightness equations relative to various perspective SfS models allowed to make the corresponding differential problems well-posed. In [18], a unified approach based on the theory of *viscosity solutions* has been proposed, showing that the brightness equations coming from different non-Lambertian reflectance models with the attenuation term admit a unique viscosity solution.




Fig. 4: (a) Example of the 1D-surface $z = u_1(x) = 1 - x^2$. (b) Under vertical lighting, three other solutions (amongst an infinity), which are differentiable almost everywhere, satisfy the same eikonal equation as u_1 . Pictures taken from [30].

2.4 Numerical Approximation

An important step towards the numerical solving of SfS was achieved when inverse problems in computer vision caught the attention of mathematicians [64]. Efficient numerical approaches were suggested to solve the eikonal equation (6), which is the SfS fundamental differential model relating the surface slope to the image graylevel. By construction, this equation can only be solved *globally*, therefore SfS ambiguities are reduced, in comparison with local approaches. They are not eliminated yet, because the concave/convex ambiguity remains.

An overview of the numerical methods for solving SfS can be found in [30, 134]. PDE-based methods (e.g., [34]) find a viscosity solution to the eikonal equation. Just to give an example let us consider the basic eikonal equation (6). A typical technique to solve it is to use a finite difference scheme. One example is the following iterative Lax-Friedrichs scheme which, in its simplest form, can be written as:

$$u_{i,j}^{(k+1)} = \frac{u_{i-1,j}^{(k)} + u_{i+1,j}^{(k)} + u_{i,j-1}^{(k)} + u_{i,j+1}^{(k)}}{4} - \frac{1}{2} \left(\sqrt{\left(\frac{u_{i+1,j}^{(k)} - u_{i-1,j}^{(k)}}{2} \right)^2 + \left(\frac{u_{i,j+1}^{(k)} - u_{i,j-1}^{(k)}}{2} \right)^2} - f_{i,j} \right), \quad (8)$$

where $f_{i,j}$ is the right-hand side of the eikonal equation (6) at the pixel (i, j) , $u_{i,j}$ is the height at this pixel, and the index k is the number of the iteration of the iterative scheme. The values $\{u_{i,j}^0\}$ represent an initial guess for the height, typically a constant value. Let us briefly explain the meaning of the iterative scheme: the first term is an average of four values around the pixel (i, j) , and inside the square root there are the centered finite difference approximations of the partial derivatives $\partial u / \partial x$ and $\partial u / \partial y$. In practice, several approximation schemes are available, e.g. finite difference as illustrated in [103, 86], semi-Lagrangian schemes [32, 33]. Most of the efficient schemes use up-wind approximations of the derivatives and additional

terms to control the diffusion in the scheme. Let us also mention that a fast-marching version for these methods allows to drastically reduce the CPU time for this type of algorithms [103, 26] and has been extensively applied in the area of image processing. Another delicate point is that the graylevel function I is typically a discontinuous function, so the approximation scheme should take into account this lack of regularity (a result in this direction is in [36]).

On the other hand, many optimisation-based methods have been proposed to compute the normal field \mathbf{n} . Under orthographic projection, (3) shows that \mathbf{n} just depends on p and q as defined in (4). Therefore, there exists a function \mathcal{R} such that $\mathcal{R}(p(\mathbf{x}), q(\mathbf{x})) := R(\mathbf{n}(\mathbf{x}))$. From this and Eq. (1), the following least-squares variational model of SfS is derived (robust estimators have also been used [130]):

$$\min_{p, q: \Omega \rightarrow \mathbb{R}} \int_{\Omega} \left| I(\mathbf{x}) - \mathcal{R}(p(\mathbf{x}), q(\mathbf{x})) \right|^2 d\mathbf{x}. \quad (9)$$

As already said in the previous subsection, this problem is clearly ill-posed. Nevertheless, if u is of class C^2 , p and q are two non-independent functions since, according to Schwarz's theorem, $\partial p / \partial y = \partial q / \partial x$. For numerical reasons [49], this hard constraint is usually replaced by a quadratic regularisation term weighted by a hyper-parameter $\lambda > 0$, which gives the following better-posed problem than (9):

$$\min_{p, q: \Omega \rightarrow \mathbb{R}} \int_{\Omega} \left| I(\mathbf{x}) - \mathcal{R}(p(\mathbf{x}), q(\mathbf{x})) \right|^2 d\mathbf{x} + \lambda \int_{\Omega} \left| \frac{\partial p}{\partial y}(\mathbf{x}) - \frac{\partial q}{\partial x}(\mathbf{x}) \right|^2 d\mathbf{x}. \quad (10)$$

Another regularisation term has been extensively used, since it is easier to discretise [62]:

$$\min_{p, q: \Omega \rightarrow \mathbb{R}} \int_{\Omega} \left| I(\mathbf{x}) - \mathcal{R}(p(\mathbf{x}), q(\mathbf{x})) \right|^2 d\mathbf{x} + \lambda \int_{\Omega} \left[|\nabla p(\mathbf{x})|^2 + |\nabla q(\mathbf{x})|^2 \right] d\mathbf{x}. \quad (11)$$

Typical optimisation methods are descent methods. For instance, the Euler-Lagrange equations derived from (11) are written (dependencies on \mathbf{x} are omitted):

$$\left[I - \mathcal{R}(p, q) \right] \frac{\partial \mathcal{R}}{\partial p}(p, q) + \lambda \Delta p = 0 \quad \text{and} \quad \left[I - \mathcal{R}(p, q) \right] \frac{\partial \mathcal{R}}{\partial q}(p, q) + \lambda \Delta q = 0. \quad (12)$$

Using the classical discrete approximation of the Laplacian Δp at pixel (i, j) :

$$\Delta p_{i,j} \approx \frac{p_{i+1,j} + p_{i-1,j} + p_{i,j+1} + p_{i,j-1}}{4} - p_{i,j}, \quad (13)$$

the following iterative scheme for solving (11) comes down from (12) and (13) [53]:

$$\begin{cases} p_{i,j}^{(k+1)} = \bar{p}_{i,j}^{(k)} + \frac{1}{\lambda} \left[I_{i,j} - \mathcal{R}(p_{i,j}^{(k)}, q_{i,j}^{(k)}) \right] \frac{\partial \mathcal{R}}{\partial p}(p_{i,j}^{(k)}, q_{i,j}^{(k)}), \\ q_{i,j}^{(k+1)} = \bar{q}_{i,j}^{(k)} + \frac{1}{\lambda} \left[I_{i,j} - \mathcal{R}(p_{i,j}^{(k)}, q_{i,j}^{(k)}) \right] \frac{\partial \mathcal{R}}{\partial q}(p_{i,j}^{(k)}, q_{i,j}^{(k)}), \end{cases} \quad (14)$$

where \bar{p} denotes the local average of p , $p^{(0)}$ and $q^{(0)}$ are given initial conditions, and the index k is the number of the iteration.

To avoid divergence for such schemes [31], it has been proposed to directly minimise the functional in (11), using conjugate gradient descent [115, 62] or line search [29], but the approximate solution is typically a local minimum. A way to overcome this limitation is to use a global optimisation, e.g. simulated annealing [27]. Finally, to decrease the CPU time, it has been dealt with multi-resolution [115].

Even if some optimisation-based methods aim to directly solve the SfS problem in the height u , as for instance [23] where a parametric model with few parameters is used, most of them first compute a normal field \mathbf{n} . Once the components p and q of the normal (cf. Eq. (4)) have been computed, it remains to *integrate* them into a height map. Several methods can be used for this task, depending on the application's requirements in terms of speed, robustness to noise in the estimated normal field, and preservation of discontinuities [91]. For instance, a standard solution for the recovery of a smooth height map consists in considering the quadratic variational problem:

$$\min_{u: \Omega \rightarrow \mathbb{R}} \int_{\Omega} \left| \nabla u(\mathbf{x}) - \begin{bmatrix} p(\mathbf{x}) \\ q(\mathbf{x}) \end{bmatrix} \right|^2 d\mathbf{x}, \quad (15)$$

which can be solved, e.g., using Fourier analysis [37], discrete sine or cosine transform [109] or iterative methods [7], depending upon the shape of Ω and the boundary conditions.

2.5 Applications of SfS

The natural application of SfS is the 3D-reconstruction of a scene from a single image. However, in real-world settings the assumptions formulated above on reflectance and lighting are too restrictive. Therefore, efforts have recently been devoted to move beyond the assumptions of Lambertian reflectance [56, 118, 119] and controlled illumination [55, 94]. In such works, reflectance and lighting are allowed to take a more general form, yet they still must be calibrated. To remove this limitation, additional priors must be introduced, as it is common in the field of intrinsic image decomposition where the reflectance is often assumed to be piecewise smooth [9]. Alternatively, deep learning techniques can be employed to simultaneously estimate shape, reflectance and lighting, provided that the object to reconstruct resembles those in the learning database [102]. In the absence of such priors, SfS can be combined with another 3D-reconstruction technique: the latter provides a coarse prior on geometry, whose details are then refined using SfS. In this view, SfS has been combined with shape-from-texture [123], structure-from-motion [38], multi-view stereopsis [61, 68, 72], or depth sensors [41, 85].

An alternative strategy to resolve the ambiguities of SfS consists in using additional images taken under varying lighting. This approach, which is called photometric stereo, will be discussed in the next section.

3 Photometric Stereo

The photometric stereo technique, first developed by Woodham [129], is an extension of SfS which considers several images acquired under the same viewing angle, but various lighting conditions.

3.1 Well-posedness of PS

One may reasonably hope that shape inference by PS will be better-posed, in comparison with the single-image case of SfS. Indeed, 3D-shape and Lambertian reflectance can be exactly and uniquely determined from a set of three images taken under non-coplanar, uniform, calibrated directional lighting. This is easily shown by considering a system of $m \geq 3$ image irradiance equations such as (2), obtained under illumination with uniform intensity $\psi(\mathbf{x}) \equiv 1$, but varying direction characterised by vectors ω_i , $i \in \{1, \dots, m\}$:

$$I_i(\mathbf{x}) = r(\mathbf{x}) \omega_i^\top \mathbf{n}(\mathbf{x}), \quad i \in \{1, \dots, m\}. \quad (16)$$

This system of equations comes down to a linear system of m equations in $\mathbf{m}(\mathbf{x}) := r(\mathbf{x}) \mathbf{n}(\mathbf{x})$. Provided that $m = 3$ and the three illumination vectors ω_i , $i \in \{1, 2, 3\}$, are non-coplanar, there exists a unique solution $\mathbf{m}(\mathbf{x})$ of this system, from which the albedo can be extracted as $r(\mathbf{x}) = |\mathbf{m}(\mathbf{x})|$ and the surface normal as $\mathbf{n}(\mathbf{x}) = \frac{\mathbf{m}(\mathbf{x})}{|\mathbf{m}(\mathbf{x})|}$. When $m > 3$, an approximate solution of the system can be estimated as long as the m illumination vectors remain non-coplanar. An example of result obtained with this approach on a banknote is presented in Fig. 5. It illustrates well the unique ability of PS both to estimate fine-scale geometric details, and to estimate the reflectance.

PS can however be ill-posed in two particular scenarios. Firstly, when lighting is unknown (*uncalibrated PS*), the local estimation of surface normals is under-constrained. As in SfS, the problem must be reformulated globally, and the integrability constraint must be imposed [133]. But even then, a low-frequency ambiguity known as the *generalised bas-relief ambiguity* remains [12]: it is necessary to introduce additional priors, see [106] for an overview of existing uncalibrated photometric stereo approaches, and [19] for a modern solution based on deep learning. Another situation where PS is ill-posed is when only two images are considered [93]: in each pixel there exist two possible normals explaining the pair of graylevels, even with known reflectance and lighting. Again, integrability must be imposed in order to limit the ambiguities [84].




Fig. 5: Photometric stereo-based 3D-reconstruction of a 10 euro banknote. From a set of images captured under varying lighting (left), PS infers both the surface geometry (top-right, we show the RGB-coded estimated normals and the 3D-shape obtained by integration of the normals), as well as its reflectance (bottom-right, we show the estimated albedo).

3.2 Numerical Solving of PS

In the previous subsection we described a simple strategy to estimate the surface normals by photometric stereo. The knowledge of surface normals is however not sufficient to fully characterise the geometry of the pictured scene. To obtain a complete 3D-representation, the normals must then be integrated into a height map. We have already discussed this integration problem in Sect. 2.4, and we refer the reader to [91] for a comprehensive overview.

With this pipeline, one first estimates the surface normals, and then integrates them into a height map. This strategy is however suboptimal, since any error in the normal estimation step will propagate during the subsequent normal integration one. An alternative strategy is to reformulate (16) as a system of partial differential equations in the unknown height map u , and directly estimate u . For instance, one may consider the ratio of two equations such as (16), for $i \neq j$, while replacing the surface normal $\mathbf{n}(\mathbf{x})$ by its definition (3). This yields the following PDE:

$$[I_i(\mathbf{x}) \omega_j - I_j(\mathbf{x}) \omega_i]^\top \begin{bmatrix} -\nabla u(\mathbf{x}) \\ 1 \end{bmatrix} = 0, \quad \forall \mathbf{x} \in \Omega, \quad (17)$$

which is linear in ∇u , and is independent from the reflectance r . It can be solved, e.g., using a finite difference upwind scheme or semi-Lagrangian methods [69]. When more than a single pair of images is considered, the joint approximate solving of the system of equations such as (17), obtained for every pair $\{i, j\}$, can be formulated as a variational problem:

$$\min_{u: \Omega \rightarrow \mathbb{R}} \sum_{i < j} \int_{\Omega} \left| [I_i(\mathbf{x}) \boldsymbol{\omega}_j - I_j(\mathbf{x}) \boldsymbol{\omega}_i]^\top \begin{bmatrix} -\nabla u(\mathbf{x}) \\ 1 \end{bmatrix} \right|^2 d\mathbf{x}. \quad (18)$$

Such an approach, initially proposed in [92, 110], also easily extends to more elaborate camera or reflectance models [70].

Nevertheless, this ratio-based approach does not provide the reflectance, contrarily to the simple pipeline presented in the previous subsection. Moreover, solving the linearised partial differential equations (17) is not equivalent to solving the original equations (16): for instance, Gaussian noise on the images turns into Cauchy noise on the ratios, making least-squares inference suboptimal. Thus, the joint recovery of height and reflectance by variational inversion of the image irradiance equations (16) has also been explored. For example, plugging the definition (3) of $\mathbf{n}(\mathbf{x})$ into (16), the joint estimation of height and reflectance in a least-squares sense leads to

$$\min_{u, r: \Omega \rightarrow \mathbb{R}} \sum_{i=1}^m \int_{\Omega} \left| I_i(\mathbf{x}) - r(\mathbf{x}) \boldsymbol{\omega}_i^\top \begin{bmatrix} -\nabla u(\mathbf{x}) / \sqrt{1 + |\nabla u(\mathbf{x})|^2} \\ 1 / \sqrt{1 + |\nabla u(\mathbf{x})|^2} \end{bmatrix} \right|^2 d\mathbf{x}, \quad (19)$$

which can be solved, e.g., using alternating reweighted least-squares [95].

3.3 PS with Non-trivial Reflectance or Lighting

The surface has been assumed Lambertian in our models, and lighting has been assumed directional but those assumptions are difficult to satisfy in real-world scenarios. An important feature of PS, in comparison with SfS, is that the redundancy provided by the multiple images enables relaxing such assumptions. Indeed, shadows or off-Lambertian effects such as specularities can be coped with by solving PS in a robust manner, for instance by resorting to sparse regression which treats such effects as outliers to the Lambertian model [52, 95]. Other ways to deal with off-Lambertian effects include inverting a reflectance model which is more sophisticated than Lambert's [70, 107], or pre-processing the images according to a low-rank prior [131]. Let us also mention data-driven methods, which either compare the intensity variations with those observed on a reference object with known shape [46], or resort to a deep neural network trained on large dataset [98].

Another direction of research on PS is the study of more realistic lighting models, in order to simplify the acquisition of data. Methods have for instance been developed to handle images acquired under nearby point light illumination [71], which finds a natural application in LED-based photometric stereo [90]. This permits to build a simple acquisition setup based on cheap hardware. Extended light sources have also been considered, which permits for instance to use the screen of a LCD display as light source [21]. Eventually, other approaches have considered the case of natural illumination in order to bring PS outdoor [11], and numerical solving methods based on variational principles [42] or deep learning [47] have recently been suggested.

3.4 Combining PS and other 3D-reconstruction Methods

A criticism which is frequently formulated against PS is that it excels with the recovery of high-frequency geometric details, yet it is prone to a low-frequency bias which may distort the global geometry [82]. In fact, such a bias usually comes from a contradiction between the assumptions behind the image formation model and the actual experiments, e.g. assuming a directional light source instead of a nearby point light one. Therefore, the methods discussed in the previous subsection provide a remedy to such a bias.

On the other hand, it is sometimes simpler from a practical perspective to stick to the simplest assumptions, and rather remove the low-frequency bias by coupling PS with another 3D-reconstruction method such as shape-from-silhouette [122], multi-view stereopsis [63] or depth sensing [87]. In such works, PS provides the fine-scale geometric details, which are combined with the gross geometry provided by the alternative technique.

Another interesting application of PS is 3D-reconstruction from a single shot, which can be achieved by using a multichannel camera coupled with monochromatic colored light sources which are simultaneously turned on: each channel can then be viewed as a graylevel image obtained under a single light source. This idea, which dates back from the nineties [60], has more recently been applied to the real-time 3D-reconstruction of deformable surfaces by combining PS with optical flow [44] or scene flow [40].

3.5 Applications of PS

The ability of PS to estimate both the fine-scale geometric details and the reflectance of the surface has proven useful in many applications. Here, we briefly highlight a few of them. For instance, PS can be used to infer 3D-models for augmented reality, which can be very helpful for computer-aided surgery using laparoscopy [22]. Another medical application of PS is the characterisation of the melanoma's shape and color, as proposed in [114]. Besides medical applications, PS has been extensively used in the field of quality control, e.g. for the inspection of defects on metallic surfaces [113]. Also, let us mention Reflectance Transform Imaging (RTI) techniques, which are based on PS principles, and allow one to interactively relight the pictured surfaces. Such an approach finds a natural application in the field of cultural heritage, see the recent survey [88] for an overview. Finally, Chapter 7 in the present volume addresses a novel application, which is the estimation of facial aging.

4 Shape-from-polarisation

Another problem belonging to the shape-from-X class is the shape-from-polarisation one. The goal is the same, i.e., recover the 3D-shape of the object, but starting from a different input data, given by polarisation information.




Fig. 6: (a) Polarimetric capture, and (b-c-d) decomposition into polarisation images, from captured data of a piece of fruit. Pictures taken and adapted from [112].

4.1 Description and Generation of a Polarisation Image

When unpolarised light is reflected by a surface, it becomes partially polarised [128]. This applies to both specular [96] and diffuse [4] reflections caused by subsurface scattering. Using a linear polarising filter placed in front of a camera, a sequence of $m \geq 3$ images (cf. Fig. 6-a) is captured by rotating the filter under varying polariser angle ϑ_j , $j \in \{1, \dots, m\}$. The measured brightness at each pixel \mathbf{x} varies in accordance to the *transmitted radiance sinusoid* corresponding to

$$i_{\vartheta_j}(\mathbf{x}) = \frac{I_{\max}(\mathbf{x}) + I_{\min}(\mathbf{x})}{2} + \frac{I_{\max}(\mathbf{x}) - I_{\min}(\mathbf{x})}{2} \cos[2\vartheta_j - 2\phi(\mathbf{x})], \quad (20)$$

where $\phi(\mathbf{x})$ is the phase angle, I_{\max} the maximum measured pixel brightness, and I_{\min} the minimum one.

A *polarisation image* (cf. Figs. 6-b-c-d), i.e., the full set of polarisation data for a given object or scene, can be obtained by decomposing the sinusoid at every pixel into three separate components [127]: the *phase angle*, $\phi(\mathbf{x})$, the *unpolarised intensity*, $i_{\text{un}}(\mathbf{x})$, and the *degree of polarisation*, $\rho(\mathbf{x})$, where

$$i_{\text{un}}(\mathbf{x}) = \frac{I_{\text{max}}(\mathbf{x}) + I_{\text{min}}(\mathbf{x})}{2} \quad \text{and} \quad \rho(\mathbf{x}) = \frac{I_{\text{max}}(\mathbf{x}) - I_{\text{min}}(\mathbf{x})}{I_{\text{max}}(\mathbf{x}) + I_{\text{min}}(\mathbf{x})}. \quad (21)$$

The phase angle $\phi(\mathbf{x})$ is directly related to the angle of the linearly polarised component of the reflected light and can be defined as the angle of maximum or minimum transmission. Since polarisers cannot distinguish between two angles separated by π radians, the range of initially acquired phase measurements is $[0, \pi)$. Therefore, there is a π ambiguity, since two maxima in pixel brightness are found as the polariser is rotated through 2π . The unpolarised image $i_{\text{un}}(\mathbf{x})$ is simply the image that would be obtained using a standard camera. The degree of polarisation $\rho(\mathbf{x})$ can be defined in terms of refractive index and zenith angle of the surface normal [128], but the explicit formula is different depending on the polarisation model used, as we will see in Sect. 4.2 below.

These quantities can be estimated from the captured image sequence using different methods, e.g. the Levenberg-Marquardt nonlinear curve fitting algorithm [4], linear methods [50] or following the procedure suggested by Wolff in [127] for the specific case of $m = 3$, $\vartheta \in \{0, \frac{\pi}{4}, \frac{\pi}{2}\}$.

4.2 Diffuse and Specular Polarisation Models

A polarisation image provides information on the azimuth and zenith angles of the normal, and, hence, a constraint on the surface normal direction at each pixel. The exact nature of the constraint depends on the polarisation model used.

Using a diffuse polarisation model, the phase angle $\phi(\mathbf{x})$ is the polariser angle ϑ_j at which I_{max} is observed. It determines the azimuth angle $\alpha(\mathbf{x}) \in [0, 2\pi[$ of the surface normal up to a π ambiguity: $\alpha(\mathbf{x}) = \phi(\mathbf{x})$ or $\phi(\mathbf{x}) + \pi$. The degree of polarisation $\rho_d(\mathbf{x})$, on the other hand, is related to the zenith angle $\theta(\mathbf{x}) \in [0, \frac{\pi}{2}]$ of the normal in viewer-centered coordinates (i.e., the angle between the normal and viewer) as follows:

$$\rho_d(\mathbf{x}) = \frac{\sin^2 \theta(\mathbf{x}) \left(\eta - \frac{1}{\eta}\right)^2}{4 \cos \theta(\mathbf{x}) \sqrt{\eta^2 - \sin^2 \theta(\mathbf{x})} - \sin^2 \theta(\mathbf{x}) \left(\eta + \frac{1}{\eta}\right)^2 + 2\eta^2 + 2}, \quad (22)$$

where η is the refractive index (in general, η is unknown, but for most dielectrics typical values range between 1.4 and 1.6, hence an accurate estimate of geometry can be obtained without a precise estimate of η [4]).

Instead, using a specular polarisation model, the azimuth angle of the surface normal is perpendicular to the phase of the specular polarisation [97] leading to a $\frac{\pi}{2}$ shift, so that the azimuth angle corresponds to polariser angle ϑ_j at which I_{\min} is observed: $\alpha(\mathbf{x}) = \phi(\mathbf{x}) \pm \frac{\pi}{2}$. Regarding the degree of polarisation $\rho_s(\mathbf{x})$, it relates to the zenith angle according to

$$\rho_s(\mathbf{x}) = \frac{2 \sin^2 \theta(\mathbf{x}) \cos \theta(\mathbf{x}) \sqrt{\eta^2 - \sin^2 \theta(\mathbf{x})}}{\eta^2 - \sin^2 \theta(\mathbf{x}) - \eta^2 \sin^2 \theta(\mathbf{x}) + 2 \sin^4 \theta(\mathbf{x})}, \quad (23)$$

and in that case the dependency of the degree of polarisation ρ_s on η is weaker than in the diffuse case.

4.3 3D-shape Recovery using Polarisation Information

The phase angle $\phi(\mathbf{x})$ (cf. Fig. 6-b) and the degree of polarisation $\rho(\mathbf{x})$ (cf. Fig. 6-d) of reflected light conveys information about the surface orientation through information on zenith and azimuth angles and, therefore, provides a cue for 3D-shape recovery.

There are nice and attractive properties to the SFP cue: it requires only a single viewpoint and a single illumination condition, it is invariant to illumination direction and surface albedo, and it provides information about both the zenith and azimuth angle of the surface normal. Unfortunately, the polarisation information alone restricts the surface normal at each pixel to two possible directions, providing in such a way only ambiguous estimates of the surface orientation.

SfP methods can be categorised into three groups:

1. Methods which use only polarisation information (cf. Sect. 4.3.1). They are passive since, typically, a polarisation image is obtained by capturing a sequence of images in which a linear polarising filter is rotated in front of the camera (possibly with unknown rotation angles [100]). These methods can be considered “single shot” methods by using custom CCD cameras configured for polarisation imaging¹ or by mounting the polarisation filter on a CMOS sensor in order to acquire polarisation information in real-time²).
2. Methods which combine polarisation with shading cues (cf. Sect. 4.3.2).
3. Methods which combine a polarisation image with an additional cue (cf. Sect. 4.3.3) such as stereo, multispectral measurements, an RGBD sensor or active polarised illumination.

SfP methods can also be classified according to the polarisation model (dielectric versus metal, diffuse [77, 4, 50], specular [79] or hybrid models [116]) and whether they compute shape in the surface normal or surface height domain.

¹ <http://www.fluxdata.com/products/fd-1665p-imaging-polarimeter>

² https://www.ricoh.com/technology/tech/051_polarization

4.3.1 Resolution using only Polarisation Information

The earliest work focused on capture, decomposition and visualisation of polarisation images was by Wolff in the nineties [127], even if older works on shape recovery by polarisation information exist since 1962 [108]. Both Atkinson and Hancock [4] and Miyazaki et al. [77] disambiguated the polarisation normals via propagation from the boundary under an assumption of global convexity. Huynh et al. [50] also disambiguated polarisation normals with a global convexity assumption, estimating the refractive index in addition. These works used a diffuse polarisation model whereas Morel et al. [79] used a specular polarisation model for metals. Recently, Taamazyan et al. [116] introduced a mixed diffuse/specular polarisation model. All of these methods estimate surface normals which must then be integrated into a height map. Moreover, since they rely entirely on the weak shape cue provided by polarisation and do not enforce integrability, the results are extremely sensitive to noise.

4.3.2 Polarisation and Shading Cues

A polarisation image also contains an unpolarised intensity channel (cf. Fig. 6-c), which provides a shading cue. Mahmoud et al. [67] used a shape-from-shading cue assuming known light source direction, known albedo and Lambertian reflectance, in order to disambiguate the polarisation normals. Atkinson and Hancock [6] used calibrated, three-source Lambertian photometric stereo for disambiguation but avoiding an assumption of known albedo. Smith et al. [111] showed how to express polarisation and shading constraints directly in terms of surface height, leading to a robust and efficient linear least-squares solution. They also showed how to estimate the illumination, up to a binary ambiguity, making the method uncalibrated. However, they require known or uniform albedo. This requirement was afterwards relaxed in [112], where spatially-varying albedo was estimated from a single polarisation image, assuming known illumination and strong smoothness assumptions. In [120] variants of the aforementioned method have been exploited by introducing additional constraints which arise when a second light source is considered, allowing to relax the uniform albedo assumption even under unknown lighting. In this work, albedo-invariant or phase-invariant formulations were proposed. Another differential approach has been proposed in [65], where the geometry of the object is described through its level-sets for both diffuse and specular reflections. Ngo et al. [83] derived constraints which allowed surface normals, light directions and refractive index to be estimated from polarisation images under varying lighting. However, this approach requires at least four light sources.

4.3.3 Combining Polarisation with other Cues

In order to solve the ambiguities generated by models using only polarisation information, some attempts have been done combining SfP with other cues. In addition to photometric cues (from SfS or PS), auxiliary geometric information can be considered. Stereo cues has been combined with polarisation to obtain surface orientation information since the nineties [126]. Rahmann et al. [96] proposed to reconstruct specular surfaces taking polarisation images from multiple views. The reconstruction problem is solved by an optimisation scheme where the surface geometry is modelled by a set of hierarchical basis functions. Atkinson et al. [3, 5] refined estimates of the surface normal to establish correspondences between two views of an object, extracting surface patches from each view. Multi-view stereo (MVS) and polarisation have also been adopted for transparent and specular objects [73, 76], and a polarimetric MVS method applied to objects with mixed polarisation models is proposed in [28]. With respect to this last paper, which is offline and needs a manual preparation, Yang et al. proposed in [132] a fully automatic approach to produce a height map in real-time using two views. More than two views have been used in [20]. Space carving [75, 76] or RGBD sensors [57, 58] have been employed to obtain initial 3D-shape, from which the ambiguities in SfP are resolved. Zhu et al. [135] used polarisation and an RGBD stereo pair to disambiguate the polarisation surface normal estimates using a higher order graphical model. Cameras with multiple spectral bands [74, 51] could be useful for disambiguating and estimating the refractive index of the surface.

4.4 An Example of Numerical Resolution for Shape Recovery

In this section we want to give an example of numerical resolution of the SfP problem, either by following a non-differential approach, which considers as unknowns the partial derivatives p and q as defined in Eq. (4), or by solving a linear differential system directly in the height u .

We assume orthographic projection and directional illumination. We consider only the diffuse polarisation model, hence the degree of polarisation is defined as in (22), and the object we want to recover is composed by dielectric (i.e., non metallic) materials. Moreover, the refractive index η is supposed to be a known constant, and interreflections are neglected. In order to estimate the phase angle $\phi(\mathbf{x})$ and the degree of diffuse polarisation $\rho_d(\mathbf{x})$ at each point, we fit the data to the transmitted radiance sinusoid (20) following one of the aforementioned methods, e.g. the idea by Wolff [127]. The zenith angle $\theta(\mathbf{x})$ of the surface normal can be obtained from Eq. (22) arriving to:

$$\cos \theta(\mathbf{x}) = \mathbf{n}(\mathbf{x}) \cdot \mathbf{v} = f(\rho(\mathbf{x}), \eta) = \sqrt{\frac{2\rho + 2\eta^2\rho - 2\eta^2 + \eta^4 + \rho^2 + 4\eta^2\rho^2 - \eta^4\rho^2 - 4\eta^3\rho\sqrt{-(\rho-1)(\rho+1)} + 1}{\eta^4\rho^2 + 2\eta^4\rho + \eta^4 + 6\eta^2\rho^2 + 4\eta^2\rho - 2\eta^2 + \rho^2 + 2\rho + 1}}, \quad (24)$$

where we have denoted ρ_d simply by ρ and we have dropped the dependency of ρ on \mathbf{x} for readability. The normal vector defined in (3) can be written in terms of azimuth and zenith angles as

$$\mathbf{n}(\mathbf{x}) = \begin{bmatrix} \cos \alpha(\mathbf{x}) \sin \theta(\mathbf{x}) \\ \sin \alpha(\mathbf{x}) \sin \theta(\mathbf{x}) \\ \cos \theta(\mathbf{x}) \end{bmatrix}. \quad (25)$$

Remembering that the phase angle $\phi(\mathbf{x})$ determines the azimuth angle $\alpha(\mathbf{x})$ of the normal up to a π ambiguity ($\alpha(\mathbf{x}) = \phi(\mathbf{x})$ or $\alpha = \phi(\mathbf{x}) + \pi$), the normal vector can be estimated up to an ambiguity. Several attempts have been done in order to disambiguate the azimuth angle, as explained in Sect. 4.3.1. Once the surface normal has been estimated, by integration we can recover the height, which is our real and final unknown to be found. Again, we refer the interested reader to Sect. 2.4 and the survey [91] for some discussion on the integration problem.

As an alternative, we can solve the problem directly in the unknown height following a differential approach, starting again from a single polarisation image, but using also the unpolarised intensity quantity, which is the image obtained using a standard camera for the SfS problem. For example, let us assume Lambertian reflectance, known illumination and uniform albedo that is factored into the light source vector ω . The shading constraint coming from the unpolarised intensity channel of a polarisation image reads as (cf. Eq. (5)):

$$i_{\text{un}}(\mathbf{x}) = \frac{-\omega_1 p(\mathbf{x}) - \omega_2 q(\mathbf{x}) + \omega_3}{\sqrt{1 + p(\mathbf{x})^2 + q(\mathbf{x})^2}}. \quad (26)$$

Since we are working in a viewer-centered coordinate system, with the viewer $\mathbf{v} = [0, 0, 1]^\top$, Eq. (24) simplifies to $n_3(\mathbf{x}) = f(\rho(\mathbf{x}), \eta)$, which can be expressed in terms of the surface gradient as

$$f(\rho(\mathbf{x}), \eta) = \frac{1}{\sqrt{1 + p(\mathbf{x})^2 + q(\mathbf{x})^2}}. \quad (27)$$

Now, by using the image ratio technique commonly applied also in PS-SfS problems [119], taking a ratio between (26) and (27), the nonlinear normalisation factor vanishes, yielding the following linear equation in the surface gradient:

$$\frac{i_{\text{un}}(\mathbf{x})}{f(\rho(\mathbf{x}), \eta)} = -\omega_1 p(\mathbf{x}) - \omega_2 q(\mathbf{x}) + \omega_3. \quad (28)$$

Instead of disambiguating the polarisation normals at each pixel locally, as illustrated before following a non-differential approach, here we express the azimuth ambiguity

as a collinearity condition which is satisfied by either of the two possible azimuth angles. In this way, we postpone resolution of the ambiguity until surface height is computed, solving the azimuthal ambiguities in a globally optimal way.

More in detail, for the diffuse case we require that the projection of the surface normal into the image plane Oxy , $[n_1(\mathbf{x}), n_2(\mathbf{x})]^\top$, is collinear with a vector pointing in the phase angle direction, $[\sin \phi(\mathbf{x}), \cos \phi(\mathbf{x})]^\top$. This requirement translates into the following condition:

$$\mathbf{n}(\mathbf{x})^\top [\cos \phi(\mathbf{x}), -\sin \phi(\mathbf{x}), 0]^\top = 0. \quad (29)$$

By rewriting $\mathbf{n}(\mathbf{x})$ in terms of the surface gradient, noting that the nonlinear normalisation term is always non-null, we obtain from Eq. (29) a second linear equation in the surface gradient:

$$-p(\mathbf{x}) \cos \phi(\mathbf{x}) + q(\mathbf{x}) \sin \phi(\mathbf{x}) = 0. \quad (30)$$

At this point, after approximating the surface gradient, e.g. by using finite differences, we arrive to a linear system of equations in terms of the unknown surface height, which can be solved using linear least-squares. For stability reasons, priors on convexity and smoothness can be added to the linear system. For more information on this idea and for details on the implementation, we refer the interested reader to [111, 112].

4.5 Applications

The polarisation state of light reflected by a surface provides a cue on the material properties of the surface and, via a relationship with surface orientation, the 3D-shape. Polarisation has been used for several applications since the nineties, including early work on material segmentation [128] and diffuse/specular reflectance separation [80]. In recent years, there has been an increasing interest in using polarisation information for 3D-shape estimation [57, 83, 111, 116]. Nice applications include polarised laparoscopy [45] or in general biomedical applications [121]. In addition to the use of polarisation information for 3D-reconstruction, recently several other applications are using polarisation for different tasks. For example, for image segmentation [104], robot dynamic navigation [13, 14], image enhancement [99, 100], and reflection separation by a deep learning approach [66], which simplifies previous works requiring three images from different polariser angles [59, 101, 124]. For more details on possible applications, we refer the interested reader to Chapter 6 of this volume.

5 A Short Presentation of this Volume

As we said in the introduction, the volume contains several contributions which represent recent trends in 3D-reconstruction via photometric techniques. Here is an overview of the chapters.

In Chapter 2, Breuss and Yarahmadi focus on a more realistic shape-from-shading model than that we described in Section 2, where perspective projection is considered. A comprehensive state-of-the-art of perspective SfS (PSfS) is carried out. The case of a Lambertian surface illuminated either by a parallel and uniform luminous flux, or by a nearby point light source, is more specifically addressed. Finally, a comparative study is carried out between two methods of resolution of the PSfS problem under directional lighting, both of which are based on the fast-marching algorithm.

In Chapter 3, Or-El et al. tackle the problem of refining the depth map provided by RGBD sensors, by applying shape-from-shading techniques. The authors propose three ways to solve this problem. First, by a model-based approach effective for Lambertian surfaces, which refines the depth map by a SfS strategy applied to the RGB image. Then, they extend this approach to specular objects, using a Phong-type model and the IR image with the attached (near) light source. Lastly, a deep learning-based solution is proposed.

In Chapter 4, Gallardo et al. tackle the problem of the 3D-reconstruction of deformable surfaces using non-rigid structure from motion and shading. The authors propose an optimisation-based strategy, which aims at finding the geometry (parameterised by vertices) and reflectance (parameterised by a finite set of albedo values) which minimise a cost function combining a shape-from-shading term and a structure-from-motion one. Additional terms are also included in the cost function: a contour boundary one, a smoothness one and a quasi-isometry one. The resulting non-convex optimisation problem is addressed by a careful heuristical initialisation followed by an iterative, Gauss-Newton-based refinement over all variables in a multi-scale fashion. The proposition is evaluated both qualitatively and quantitatively against the state-of-the-art.

In Chapter 5, Brahimi et al. present a theoretical contribution on the well-posedness of uncalibrated photometric stereo under general illumination. In particular they prove that there is no ambiguity for the perspective model if lighting is represented by first-order spherical harmonics. In the process of establishing their main result they also provide a comprehensive survey of the available results regarding the well-posedness of several photometric stereo problems and they examine in detail the case of the orthographic projection. For this problem they prove that, even in the case of spherical harmonics, the concave/convex ambiguity still persists. They conclude with some numerical experiments.

Chapter 6, authored by Shi et al., represents a concise survey on SfP. After an introduction, the authors briefly recall the Fresnel theory, which is the theoretical basis of polarisation imaging. The process for the formation of a polarisation image is described, giving also details on the data acquisition. The authors discuss the estimation of azimuth and zenith angles of the normal for surfaces with different reflectance properties (specular, diffuse, and mixed polarisation). Then, the combination of SfP with auxiliary information is explored, e.g., geometric cues, spectral cues, photometric cues, and deep learning. Moreover, applications which can benefit from polarisation information, in addition to the 3D-shape recovery, are presented. The chapter ends with a discussion on problems still open.

Finally, Chapter 7 by Dahlan et al. addresses the problem of facial aging estimation, using light scattering photometry. It is shown that the roughness parameter of several BRDF models is correlated with the age. Therefore, facial aging estimation can be carried out by fitting a BRDF model to an input image. In this work, geometry estimation is carried out using photometric stereo, by resorting to an illumination dome. Then, given the estimated normals, an image with frontal lighting is used to infer the BRDF parameters. Various experiments are carried out to study whether these estimated parameters correlate with age and it is shown that this is the case for the roughness parameter. Several tests on real images are illustrated and analysed.

References

1. E. H. Adelson and A. P. Pentland. The perception of shading and reflectance. In *Perception as Bayesian inference*, pages 409–423. Cambridge University Press, 1996.
2. AliceVision. <https://alicevision.org/>.
3. G. A. Atkinson and E. R. Hancock. Multi-view surface reconstruction using polarization. In *Proceedings of the IEEE International Conference on Computer Vision*, volume 1, pages 309–316, 2005.
4. G. A. Atkinson and E. R. Hancock. Recovery of surface orientation from diffuse polarization. *IEEE Transactions on Image Processing*, 15(6):1653–1664, 2006.
5. G. A. Atkinson and E. R. Hancock. Shape estimation using polarization and shading from two views. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(11):2001–2017, 2007.
6. G. A. Atkinson and E. R. Hancock. Surface reconstruction using polarization and photometric stereo. In *Proceedings of the International Conference on Computer Analysis of Images and Patterns*, pages 466–473, 2007.
7. M. Bähr, M. Breuß, Y. Quéau, A. S. Boroujerdi, and J.-D. Durou. Fast and accurate surface normal integration on non-rectangular domains. *Computational Visual Media*, 3(2):107–129, 2017.
8. G. Barles. *Solutions de viscosité des équations de Hamilton-Jacobi*, volume 17 of *Mathématiques et Applications*. Springer-Verlag, 1994.
9. J. T. Barron and J. Malik. Shape, Illumination, and Reflectance from Shading. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(8):1670–1687, 2015.
10. A. Bartoli, Y. Gérard, F. Chadebecq, T. Collins, and D. Pizarro. Shape-from-Template. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(10):2099–2118, 2015.
11. R. Basri, D. Jacobs, and I. Kemelmacher. Photometric stereo with general, unknown lighting. *International Journal of Computer Vision*, 72(3):239–257, 2007.

12. P. N. Belhumeur, D. J. Kriegman, and A. L. Yuille. The Bas-Relief Ambiguity. *International Journal of Computer Vision*, 35(1):33–44, 1999.
13. K. Berger, R. Voorhies, and L. Matthies. Incorporating polarization in stereo vision-based 3D perception of non-Lambertian scenes. In *Unmanned Systems Technology XVIII*, volume 9837 of *Proceedings of the SPIE*, 2016.
14. K. Berger, R. Voorhies, and L. H. Matthies. Depth from stereo polarization in specular scenes for urban robotics. In *Proceedings of the International Conference on Robotics and Automation*, pages 1966–1973, 2017.
15. M. Brady and A. L. Yuille. An Extremum Principle for Shape from Contour. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(3):288–301, 1984.
16. M. Breuß, E. Cristiani, J.-D. Durou, M. Falcone, and O. Vogel. Perspective Shape from Shading: Ambiguity Analysis and Numerical Approximations. *SIAM Journal on Imaging Sciences*, 5(1):311–342, 2012.
17. F. Camilli and L. Grüne. Numerical approximation of the maximal solutions for a class of degenerate Hamilton-Jacobi equations. *SIAM Journal on Numerical Analysis*, 38(5):1540–1560, 2000.
18. F. Camilli and S. Tozza. A Unified Approach to the Well-Posedness of Some Non-Lambertian Models in Shape-from-Shading Theory. *SIAM Journal on Imaging Sciences*, 10(1):26–46, 2017.
19. G. Chen, K. Han, B. Shi, Y. Matsushita, and K.-Y. K. Wong. Self-calibrating deep photometric stereo networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8739–8747, 2019.
20. L. Chen, Y. Zheng, A. Subpa-Asa, and I. Sato. Polarimetric Three-View Geometry. In *Proceedings of the European Conference on Computer Vision*, pages 20–36, 2018.
21. J. J. Clark. Photometric stereo using LCD displays. *Image and Vision Computing*, 28(4):704–714, 2010.
22. T. Collins and A. Bartoli. 3D-reconstruction in laparoscopy with close-range photometric stereo. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, volume 7511 of *Lecture Notes in Computer Science*, pages 634–642, 2012.
23. F. Courteille, J.-D. Durou, and G. Morin. A Global Solution to the SFS Problem Using B-spline Surface and Simulated Annealing. In *Proceedings of the International Conference on Pattern Recognition (volume II)*, pages 332–335, 2006.
24. M. G. Crandall and P.-L. Lions. Viscosity solutions of Hamilton-Jacobi equations. *Transactions of the American Mathematical Society*, 277(1):1–42, 1983.
25. E. Cristiani. 3D Printers: A new challenge for mathematical modeling, 2014. arXiv preprint arXiv:1409.1714.
26. E. Cristiani and M. Falcone. Fast semi-Lagrangian schemes for the eikonal equation and applications. *SIAM Journal on Numerical Analysis*, 45(5):1979–2011, 2007.
27. A. Crouzil, X. Descombes, and J.-D. Durou. A multiresolution approach for shape from shading coupling deterministic and stochastic optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(11):1416–1421, 2003.
28. Z. Cui, J. Gu, B. Shi, P. Tan, and J. Kautz. Polarimetric Multi-View Stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1558–1567, 2017.
29. P. Daniel and J.-D. Durou. From Deterministic to Stochastic Methods for Shape From Shading. In *Proceedings of the Asian Conference on Computer Vision*, pages 187–192, 2000.
30. J.-D. Durou, M. Falcone, and M. Sagona. Numerical methods for shape-from-shading: A new survey with benchmarks. *Computer Vision and Image Understanding*, 109(1):22–43, 2008.
31. J.-D. Durou and H. Maître. On Convergence in the Methods of Strat and of Smith for Shape from Shading. *International Journal of Computer Vision*, 17(3):273–289, 1996.
32. M. Falcone and R. Ferretti. *Semi-Lagrangian approximation schemes for linear and Hamilton-Jacobi equations*. Society for Industrial and Applied Mathematics, First edition, 2014.
33. M. Falcone and R. Ferretti. Numerical methods for Hamilton-Jacobi type equations. In *Handbook of numerical methods for hyperbolic problems*, volume 17 of *Handbook of Numerical Analysis*, pages 603–626. Elsevier, 2016.

34. M. Falcone and M. Sagona. An algorithm for the global solution of the shape-from-shading model. In *International Conference on Image Analysis and Processing*, volume 1310 of *Lecture Notes in Computer Science*, pages 596–603, 1997.
35. M. Falcone and M. Sagona. A scheme for the shape-from-shading model with “black shadows”. In *Numerical mathematics and advanced applications*, pages 503–512. Springer, 2003.
36. A. Festa and M. Falcone. An Approximation Scheme for an Eikonal Equation with Discontinuous Coefficient. *SIAM Journal on Numerical Analysis*, 52(1):236–257, 2014.
37. R. T. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(4):439–451, 1988.
38. M. Gallardo, T. Collins, and A. Bartoli. Dense non-rigid structure-from-motion and shading with unknown albedos. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3884–3892, 2017.
39. J. Geng. Structured-light 3D surface imaging: A tutorial. *Advances in Optics and Photonics*, 3(2):128–160, 2011.
40. P. F. U. Gotardo, T. Simon, Y. Sheikh, and I. Matthews. Photogeometric Scene Flow for High-Detail Dynamic 3D Reconstruction. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 846–854, 2015.
41. B. Haefner, Y. Quéau, T. Möllenhoff, and D. Cremers. Fight ill-posedness with ill-posedness: Single-shot variational depth super-resolution from shading. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 164–174, 2018.
42. B. Haefner, Z. Ye, M. Gao, T. Wu, Y. Quéau, and D. Cremers. Variational uncalibrated photometric stereo under general lighting. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8539–8548, 2019.
43. R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision*. Cambridge University Press, Second edition, 2004.
44. C. Hernández. *Stereo and Silhouette Fusion for 3D Object Modeling from Uncalibrated Images Under Circular Motion*. Thèse de doctorat, École Nationale Supérieure des Télécommunications, 2004.
45. S. E. M. Herrera, A. Malti, O. Morel, and A. Bartoli. Shape-from-Polarization in laparoscopy. In *Proceedings of the International Symposium on Biomedical Imaging*, pages 1412–1415, 2013.
46. A. Hertzmann and S. M. Seitz. Example-based photometric stereo: Shape reconstruction with general, varying BRDFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1254–1264, 2005.
47. Y. Hold-Geoffroy, P. Gotardo, and J.-F. Lalonde. Single Day Outdoor Photometric Stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
48. B. K. P. Horn. *Shape from Shading: a Method for Obtaining the Shape of a Smooth Opaque Object from One View*. PhD Thesis, MIT, 1970.
49. B. K. P. Horn and M. J. Brooks. The variational approach to shape from shading. *Computer Vision, Graphics, and Image Processing*, 33(2):174–208, 1986.
50. C. P. Huynh, A. Robles-Kelly, and E. R. Hancock. Shape and refractive index recovery from single-view polarisation images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1229–1236, 2010.
51. C. P. Huynh, A. Robles-Kelly, and E. R. Hancock. Shape and refractive index from single-view spectro-polarimetric images. *International Journal of Computer Vision*, 101(1):64–94, 2013.
52. S. Ikehata, D. Wipf, Y. Matsushita, and K. Aizawa. Robust photometric stereo using sparse regression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 318–325, 2012.
53. K. Ikeuchi and B. K. P. Horn. Numerical Shape from Shading and Occluding Boundaries. *Artificial Intelligence*, 17(1–3):141–184, 1981.
54. M. Jancosek and T. Pajdla. Multi-view reconstruction preserving weakly-supported surfaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3121–3128, 2011.

55. M. K. Johnson and E. H. Adelson. Shape estimation in natural illumination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2553–2560, 2011.
56. Y.-C. Ju, S. Tozza, M. Breuß, A. Bruhn, and A. Kleefeld. Generalised Perspective Shape from Shading with Oren-Nayar Reflectance. In *Proceedings of the British Machine Vision Conference*, 2013.
57. A. Kadambi, V. Taamazyan, B. Shi, and R. Raskar. Polarized 3D: High-Quality Depth Sensing with Polarization Cues. In *Proceedings of the IEEE International Conference on Computer Vision*, 2015.
58. A. Kadambi, V. Taamazyan, B. Shi, and R. Raskar. Depth Sensing Using Geometrically Constrained Polarization Normals. *International Journal of Computer Vision*, 125:34–51, 2017.
59. N. Kong, Y.-W. Tai, and J. S. Shin. A physically-based approach to reflection separation: from physical modeling to constrained optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(2):209–221, 2013.
60. L. L. Kontsevich, A. P. Petrov, and I. S. Vergelskaya. Reconstruction of shape from shading in color images. *Journal of the Optical Society of America A*, 11(3):1047–1052, 1994.
61. F. Langguth, K. Sunkavalli, S. Hadap, and M. Goesele. Shading-aware multi-view stereo. In *Proceedings of the European Conference on Computer Vision*, pages 469–485, 2016.
62. Y. G. Leclerc and A. F. Bobick. The Direct Computation of Height from Shading. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 552–558, 1991.
63. M. Li, Z. Zhou, Z. Wu, B. Shi, C. Diao, and P. Tan. Multi-View Photometric Stereo: A Robust Solution and Benchmark Dataset for Spatially Varying Isotropic Materials. *IEEE Transactions on Image Processing*, 2020.
64. P.-L. Lions, E. Rouy, and A. Tourin. Shape-from-shading, viscosity solutions and edges. *Numerische Mathematik*, 64(1):323–353, 1993.
65. F. Logothetis, R. Mecca, F. Sgallari, and R. Cipolla. A Differential Approach to Shape from Polarisation: A Level-Set Characterisation. *International Journal of Computer Vision*, 127(11–12):1680–1693, 2019.
66. Y. Lyu, Z. Cui, S. Li, M. Pollefeys, and B. Shi. Reflection Separation using a pair of Unpolarized and Polarized Images. In *Proceedings of the Annual Conference on Neural Information Processing Systems*, 2019.
67. A. H. Mahmoud, M. T. El-Melegy, and A. A. Farag. Direct method for shape recovery from polarization and shading. In *Proceedings of the IEEE International Conference on Image Processing*, pages 1769–1772, 2012.
68. D. Maurer, Y.-C. Ju, M. Breuß, and A. Bruhn. Combining shape from shading and stereo: A joint variational method for estimating depth, illumination and albedo. *International Journal of Computer Vision*, 126(12):1342–1366, 2018.
69. R. Mecca and M. Falcone. Uniqueness and approximation of a photometric shape-from-shading model. *SIAM Journal on Imaging Sciences*, 6(1):616–659, 2013.
70. R. Mecca, Y. Quéau, F. Logothetis, and R. Cipolla. A single-lobe photometric stereo approach for heterogeneous material. *SIAM Journal on Imaging Sciences*, 9(4):1858–1888, 2016.
71. R. Mecca, A. Wetzler, A. Bruckstein, and R. Kimmel. Near field photometric stereo with point light sources. *SIAM Journal on Imaging Sciences*, 7(4):2732–2770, 2014.
72. J. Mélou, Y. Quéau, F. Castan, and J.-D. Durou. A splitting-based algorithm for multi-view stereopsis of textureless objects. In *Proceedings of the International Conference on Scale Space and Variational Methods in Computer Vision*, pages 51–63, 2019.
73. D. Miyazaki, M. Kagesawa, and K. Ikeuchi. Transparent surface modeling from a pair of polarization images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(1):73–82, 2004.
74. D. Miyazaki, M. Saito, Y. Sato, and K. Ikeuchi. Determining surface orientations of transparent objects based on polarization degrees in visible and infrared wavelengths. *Journal of the Optical Society of America A*, 19(4):687–694, 2002.

75. D. Miyazaki, T. Shigetomi, M. Baba, R. Furukawa, S. Hiura, and N. Asada. Polarization-based surface normal estimation of black specular objects from multiple viewpoints. In *Proceedings of the International Conference on 3D Imaging, Modeling, Processing, Visualization and Transmission*, pages 104–111, 2012.
76. D. Miyazaki, T. Shigetomi, M. Baba, R. Furukawa, S. Hiura, and N. Asada. Surface normal estimation of black specular objects from multiview polarization images. *Optical Engineering*, 56(4):041303, 2016.
77. D. Miyazaki, R. T. Tan, K. Hara, and K. Ikeuchi. Polarization-based inverse rendering from a single view. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 982–987, 2003.
78. T. Moons, L. Van Gool, and M. Vergauwen. 3D Reconstruction from Multiple Images Part 1: Principles. *Foundations and Trends in Computer Graphics and Vision*, 4(4):287–404, 2008.
79. O. Morel, F. Meriaudeau, C. Stolz, and P. Gorria. Polarization imaging applied to 3D reconstruction of specular metallic surfaces. In *Machine Vision Applications in Industrial Inspection XIII*, volume 5679 of *Proceedings of the SPIE*, pages 178–186, 2005.
80. S. Nayar, X. Fang, and T. Boulton. Separation of reflection components using color and polarization. *International Journal of Computer Vision*, 21(3):163–186, 1997.
81. S. K. Nayar and Y. Nakagawa. Shape from focus. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(8):824–831, 1994.
82. D. Nehab, S. Rusinkiewicz, J. Davis, and R. Ramamoorthi. Efficiently combining positions and normals for precise 3D geometry. *ACM Transactions on Graphics*, 24(3):536–543, 2005.
83. T. T. Ngo, H. Nagahara, and R. Taniguchi. Shape and Light Directions from Shading and Polarization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2310–2318, 2015.
84. R. Onn and A. M. Bruckstein. Integrability disambiguates surface recovery in two-image photometric stereo. *International Journal of Computer Vision*, 5(1):105–113, 1990.
85. R. Or-El, G. Rosman, A. Wetzler, R. Kimmel, and A. Bruckstein. RGBD-Fusion: Real-Time High Precision Depth Recovery. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5407–5416, 2015.
86. S. Osher and R. Fedkiw. *Level set methods and dynamic implicit surfaces*, volume 153 of *Applied Mathematical Sciences*. Springer-Verlag, 2003.
87. S. Peng, B. Haefner, Y. Quéau, and D. Cremers. Depth super-resolution meets uncalibrated photometric stereo. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 2961–2968, 2017.
88. R. Pintus, T. G. Dulecha, I. Ciortan, E. Gobbetti, and A. Giachetti. State-of-the-art in Multi-Light Image Collections for Surface Visualization and Analysis. *Computer Graphics Forum*, 38(3):909–934, 2019.
89. E. Prados and O. Faugeras. Shape from Shading: A Well-Posed Problem? In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 870–877, 2005.
90. Y. Quéau, B. Durix, T. Wu, D. Cremers, F. Lauze, and J.-D. Durou. LED-based photometric stereo: Modeling, calibration and numerical solution. *Journal of Mathematical Imaging and Vision*, 60(3):313–340, 2018.
91. Y. Quéau, J.-D. Durou, and J.-F. Aujol. Normal integration: A survey. *Journal of Mathematical Imaging and Vision*, 60(4):576–593, 2018.
92. Y. Quéau, R. Mecca, and J.-D. Durou. Unbiased photometric stereo for colored surfaces: A variational approach. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4359–4368, 2016.
93. Y. Quéau, R. Mecca, J.-D. Durou, and X. Descombes. Photometric stereo with only two images: A theoretical study and numerical resolution. *Image and Vision Computing*, 57:175–191, 2017.
94. Y. Quéau, J. Mérou, F. Castan, D. Cremers, and J.-D. Durou. A variational approach to shape-from-shading under natural illumination. In *Proceedings of the International Workshop on Energy Minimization Methods in Computer Vision and Pattern Recognition*, pages 342–357, 2017.

95. Y. Quéau, T. Wu, F. Lauze, J.-D. Durou, and D. Cremers. A non-convex variational approach to photometric stereo under inaccurate lighting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 99–108, 2017.
96. S. Rahmann and N. Canterakis. Reconstruction of specular surfaces using polarization imaging. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, 2001.
97. A. Robles-Kelly and C. P. Huynh. *Imaging Spectroscopy for Scene Analysis*. Springer-Verlag, 2013.
98. H. Santo, M. Samejima, Y. Sugano, B. Shi, and Y. Matsushita. Deep photometric stereo network. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 501–509, 2017.
99. Y. Y. Schechner. Inversion by P^4 : Polarization-picture post-processing. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 366(1565):638–648, 2011.
100. Y. Y. Schechner. Self-calibrating imaging polarimetry. In *Proceedings of the IEEE International Conference on Computational Photography*, 2015.
101. Y. Y. Schechner, J. Shamir, and N. Kiryati. Polarization and statistical analysis of scenes containing a semireflector. *Journal of the Optical Society of America A*, 17(2):276–284, 2000.
102. S. Sengupta, A. Kanazawa, C. D. Castillo, and D. W. Jacobs. SfSNet: Learning Shape, Reflectance and Illuminance of Faces in the Wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6296–6305, 2018.
103. J. A. Sethian. *Level set methods and fast marching methods*, volume 3 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, Second edition, 1999.
104. A. E. R. Shabayek, C. Demonceaux, O. Morel, and D. Fofi. Vision based UAV attitude estimation: Progress and insights. *Journal of Intelligent & Robotic Systems*, 65(1–4):295–308, 2012.
105. S. A. Shafer and T. Kanade. Using shadows in finding surface orientations. *Computer Vision, Graphics, and Image Processing*, 22(1):145–176, 1983.
106. B. Shi, Z. Mo, Z. Wu, D. Duan, S. Yeung, and P. Tan. A Benchmark Dataset and Evaluation for Non-Lambertian and Uncalibrated Photometric Stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2):271–284, 2019.
107. B. Shi, P. Tan, Y. Matsushita, and K. Ikeuchi. Bi-polynomial modeling of low-frequency reflectances. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(6):1078–1091, 2013.
108. W. A. Shurcliff. *Polarized light, production and use*. Harvard University Press, 1962.
109. T. Simchony, R. Chellappa, and M. Shao. Direct analytical methods for solving Poisson equations in computer vision problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(5):435–446, 1990.
110. W. A. P. Smith and F. Fang. Height from photometric ratio with model-based light source selection. *Computer Vision and Image Understanding*, 145:128–138, 2016.
111. W. A. P. Smith, R. Ramamoorthi, and S. Tozza. Linear depth estimation from an uncalibrated, monocular polarisation image. In *European Conference on Computer Vision*, volume 9912 of *Lecture Notes in Computer Science*, pages 109–125, 2016.
112. W. A. P. Smith, R. Ramamoorthi, and S. Tozza. Height-from-Polarisation with Unknown Lighting or Albedo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41:2875–2888, 2019.
113. D. Soukup and R. Huber-Mörk. Convolutional neural networks for steel surface defect detection from photometric stereo images. In *International Symposium on Visual Computing*, volume 8887 of *Lecture Notes in Computer Science*, pages 668–677, 2014.
114. J. Sun, M. Smith, L. Smith, L. Coutts, R. Dabis, C. Harland, and J. Bamber. Reflectance of human skin using colour photometric stereo: with particular application to pigmented lesion analysis. *Skin Research and Technology*, 14(2):173–179, 2008.
115. R. Szeliski. Fast Shape from Shading. *Computer Vision, Graphics, and Image Processing: Image Understanding*, 53(2):129–153, 1991.

116. V. Taamazyan, A. Kadambi, and R. Raskar. Shape from Mixed Polarization. *arXiv preprint arXiv:1605.02066*, 2016.
117. A. Tankus, N. Sochen, and Y. Yeshurun. A new perspective [on] shape-from-shading. In *Proceedings of the IEEE International Conference on Computer Vision*, volume 2, pages 862–869, 2003.
118. S. Tozza and M. Falcone. Analysis and approximation of some shape-from-shading models for non-Lambertian surfaces. *Journal of Mathematical Imaging and Vision*, 55(2):153–178, 2016.
119. S. Tozza, R. Mecca, M. Duocastella, and A. Del Bue. Direct Differential Photometric Stereo Shape Recovery of Diffuse and Specular Surfaces. *Journal of Mathematical Imaging and Vision*, 56(1):57–76, 2016.
120. S. Tozza, W. A. P. Smith, D. Zhu, R. Ramamoorthi, and E. R. Hancock. Linear Differential Constraints for Photo-polarimetric Height Estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, 2017.
121. V. V. Tuchin, L. Wang, and D. A. Zimnyakov. *Optical polarization in biomedical applications*. Springer Science & Business Media, 2006.
122. G. Vogiatzis, C. Hernández, and R. Cipolla. Reconstruction in the round using photometric normals and silhouettes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 1847–1854, 2006.
123. R. White and D. Forsyth. Combining cues: Shape from shading and texture. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 2, pages 1809–1816, 2006.
124. P. Wieschollek, O. Gallo, J. Gu, and J. Kautz. Separating Reflection and Transmission Images in the Wild. In *Proceedings of the European Conference on Computer Vision*, pages 89–104, 2018.
125. A. P. Witkin. Recovering surface shape and orientation from texture. *Artificial Intelligence*, 17(1–3):17–45, 1981.
126. L. B. Wolff. Surface orientation from two camera stereo with polarizers. In *Optics, Illumination, and Image Sensing for Machine Vision IV*, volume 1194 of *Proceedings of the SPIE*, pages 287–298, 1990.
127. L. B. Wolff. Polarization vision: A new sensory approach to image understanding. *Image and Vision Computing*, 15(2):81–93, 1997.
128. L. B. Wolff and T. E. Boulton. Constraining object features using a polarization reflectance model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):635–657, 1991.
129. R. J. Woodham. Photometric Method For Determining Surface Orientation From Multiple Images. *Optical Engineering*, 19(1):134–144, 1980.
130. P. L. Worthington and E. R. Hancock. Needle map recovery using robust regularizers. *Image and Vision Computing*, 17(8):545–557, 1999.
131. L. Wu, A. Ganesh, B. Shi, Y. Matsushita, Y. Wang, and Y. Ma. Robust photometric stereo via low-rank matrix completion and recovery. In *Proceedings of the Asian Conference on Computer Vision*, pages 703–717, 2010.
132. L. Yang, F. Tan, A. Li, Z. Cui, Y. Furukawa, and P. Tan. Polarimetric Dense Monocular SLAM. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3857–3866, 2018.
133. A. L. Yuille, D. Snow, R. Epstein, and P. N. Belhumeur. Determining Generative Models of Objects Under Varying Illumination: Shape and Albedo from Multiple Images Using SVD and Integrability. *International Journal of Computer Vision*, 35(3):203–222, 1999.
134. R. Zhang, P.-S. Tsai, J. E. Cryer, and M. Shah. Shape-from-shading: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(8):690–706, 1999.
135. D. Zhu and W. A. P. Smith. Depth From a Polarisation + RGB Stereo Pair. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7586–7595, 2019.