

Near Field Photometric Stereo with Point Light Sources*

Roberto Mecca[†], Aaron Wetzler[‡], Alfred M. Bruckstein[‡], and Ron Kimmel[‡]

Abstract. Shape recovery of an object based on shading variations resulting from different light sources has recently been reconsidered. Improvements have been made that allow for the photometric stereo approach to serve as a competitive alternative to other shape reconstruction methods. However, most photometric stereo methods tend to ignore factors that are relevant in practical applications. The setup considered in this paper tackles photometric stereo reconstruction in the case of a specific near-field imaging. This means that both the camera and the light sources are close to the imaged object, where close can be loosely considered as a setup having similar distances between lights, camera, and object. The known challenges that characterize the problem involve perspective viewing geometry, point light sources, and images that may include shadowed regions. Here, we pay special attention to the question of how to faithfully model these aspects and at the same time design an efficient and robust numerical solver. We present a mathematical formulation that integrates the above assumptions into a single coherent model based on quasi-linear PDEs. The well-posedness is proved showing uniqueness of a weak (i.e., Lipschitz continuous) solution. The surface reconstruction in our near-field scenario can then be executed efficiently in linear time. The merging strategy of the irradiance equations provided for each light source allows us to consider a characteristic expansion model which enables the direct computation of the surface. We evaluate several types of light attenuation models with a nonuniform albedo and noise on synthetic data. We also demonstrate the proposed method on surface reconstruction of real data using three images, each one taken with a different light source by a working prototype. We demonstrate the accuracy of the proposed method compared to other methods that ignore the near-field setup and assume distant, parallel beam light sources.

Key words. photometric stereo, quasi-linear PDEs, well-posedness, weak solution, perspective view, light attenuations, upwind scheme, semi-Lagrangian scheme

AMS subject classifications. 68T45, 35A02, 65N21, 65N12, 65M25

DOI. 10.1137/140968100

1. Perspective Shape from Shading and photometric stereo techniques. Since the seminal Shape from Shading (SfS) paper by Horn [15], new models have been introduced in order to extend the range of shape recovery problems that can be solved by SfS methods [28, 42, 12, 23, 43, 29, 22]. A particularly important direction of research has been the transition from the assumption of orthographic viewing geometry [21, 20, 17, 18, 4, 16, 19] to perspective cameras [39, 33, 32, 13].

Several papers reported using Perspective SfS (PSfS) methods applied to endoscopic image

*Received by the editors May 6, 2014; accepted for publication (in revised form) November 6, 2014; published electronically December 16, 2014. This work was supported by the Broadcom Foundation and by the European Community's FP7- ERC program, grant agreement 267414.

<http://www.siam.org/journals/siims/7-4/96810.html>

[†]Pattern Analysis & Computer Vision, Istituto Italiano di Tecnologia, 16163 Genova, Italy (mecca.roberto@gmail.com).

[‡]Department of Computer Science, Technion – Israel Institute of Technology, Haifa 32000, Israel (twerd@cs.technion.ac.il, freddy@cs.technion.ac.il, ron@cs.technion.ac.il).

analysis [43, 42, 28]. Okatani and Deguchi [28] introduced the perspective viewing shape estimation for objects located close to the camera, taking into account a realistic endoscopic model based on PDEs for close light source illumination (i.e., ideally attached to the optical center). Later, Prados and Faugeras [33], considering the same setup, studied a different parametrization, concluding that their formulation for the endoscopic perspective shape from the shading problem is well-posed. A recent paper by Breuß et al. [3] shows that there might be ambiguity in the Prados–Faugeras model. In fact, near-field endoscopic perspective shape reconstruction remains an open problem if only a single image is considered.

In this paper we deal with the case of calibrated photometric stereo, where multiple images of the same scene are taken under different illumination conditions (with known light positions) for which the camera, lights, and object are close to one another with comparable distances. For a recent overview on calibrated (and uncalibrated) photometric stereo methods, see [38] (and [30]). In our particular framework, the endoscopic perspective, Wu, Narasimhan, and Jaramaz [42] studied the multi-image endoscopic perspective problem by considering two light sources placed off the optical center. They develop a model based on the radiance information obtained by simultaneously illuminating an object with two different light sources. They then recover the surface by considering a single irradiance equation for the sum of Lambertian reflectance functions of the two different light sources. The use of this reflectance function results in a loss of information. In order to avoid this problem and issues related to an unknown albedo, they use a photometric calibration. Surface recovery is performed within a variational framework that involves high computational complexity compared to alternative direct methods [24]. The shape from an endoscopic perspective problem solved via a photometric stereo technique using n images (EPPS $_n$) was first addressed by Collins and Bartoli [7]. They solve the close-range photometric stereo with weak considerations on the parametrization of variables such as unknown surface height, light sources, and solving the problem of the light attenuation with an a priori light calibration procedure. Furthermore, they use a prior for a reflectance model learning, adding physical markers on the inspected object even when the surface is assumed to be Lambertian. In particular, their mathematical formulation is based on the usual double step procedure where an energy functional is minimized (which allows the computation of the surface derivatives), and only later is the surface recovered [10, 36, 1]. Moreover, their energy is based on the sum of Lambertian irradiance equations rather than using photometric ratios [41, 24, 5, 40] that lead to more suitable problems. For example, the most important feature of photometric ratios is to obtain independence from the albedo.

Recently, Parot et al. [31] studied the EPPS $_n$ problem by using a straightforward heuristic approach to photometric stereo. In their work, even if camera and lights are close to the inspected object, they assume orthographic viewing geometry, with uniform and unattenuated light directions calibrated by assuming *reasonable* distance between the object and the camera. The discrepancy with respect to the real physics they are working with is faced by filtering the directional gradients depending on the frequencies. The authors heuristically handle this by removing the lower frequencies and the DC components. Then, the resulting depth map is computed using a multigrid Poisson solver. The work describes purely qualitative results in the sense that they do not represent accurate reconstructions of the environment and instead are used as a qualitative tool for detecting lesions.

Here, we present a new mathematical formulation for the EPPS_n by considering the endoscopic perspective presented in [28] and using the perspective parametrization used in [25] based on nonlinear PDEs from photometric ratios. This extension introduces several realistic properties which result in a practical and mathematically interesting model. We formulate the model based on quasi-linear PDEs and develop an efficient finite difference upwind scheme and a semi-Lagrangian numerical scheme for the direct approximation of the surface, extending previous results obtained by Mecca et al. [26, 27]. This scheme is simple to implement, is parallelizable, and converges efficiently. The new model can also successfully handle images with missing data such as the case of shadows and occlusions. In this paper we do not focus on the problem of how to accurately detect occlusions or shadows [35], but instead show that our model successfully reconstructs surfaces in the presence of missing data. A significant advantage of our model is the direct computation of the three-dimensional (3D) surface without first explicitly computing surface normals and then performing integration. Nonlinearities resulting from our realistic lighting model yield a normal field that depends on the depth of the surface as well as the direction of the light source. However, this does not pose a problem in our framework because we avoid directly computing the surface normals and we do not use them in the numerical scheme.

In section 2 we recall a well-known parametrization for the surface under close camera observation. The new mathematical model is introduced in section 3, where we show the differential model, including two sources of nonlinear light attenuation. In section 4 we prove the well-posedness (unique weak solution) of the new differential formulation for the EPPS_2 that consists of a quasi-linear PDE. Section 5 introduces the upwind and semi-Lagrangian numerical schemes, and in section 6 we prove the convergence of the semi-Lagrangian scheme. Section 7 describes the theoretical formulation of the new differential approach for three images which can be easily extended to $n \geq 3$. Next, in section 8 experimental results on synthetic and real data are presented.

2. The endoscopic perspective setup and photometric stereo. In order to provide the necessary ingredients to understand the geometry behind the model, we start by considering the parametrization of the surface Σ (see Figure 1) up to an unknown function z from the image domain $\overline{\Omega}_p = \Omega_p \cup \partial\Omega_p$ to \mathbb{R} , such that

$$(2.1) \quad \mathbf{M}(x, y) = [\xi(x, y), \eta(x, y), \zeta(x, y)] := \left[-x \frac{z(x, y)}{f}, -y \frac{z(x, y)}{f}, z(x, y) \right].$$

Here, $f > 0$ is the focal length of the camera, $\zeta < -f < 0$, and the triple $[\xi(x, y), \eta(x, y), \zeta(x, y)] = [\xi, \eta, \zeta]$ comprises the real world coordinates (with respect to the image coordinates). This parametrization is based on the pinhole camera model and is due to the perspective viewing geometry, as seen in Figure 1, where the camera is placed at the origin (optical center) C of the coordinate system $C\xi\eta\zeta$ [33].

We recall that the unnormalized outgoing normal vector to the surface Σ is given by

$$(2.2) \quad \bar{\mathbf{n}}(x, y) = \frac{z}{f^2} [f \nabla z(x, y), z(x, y) + (x, y) \cdot \nabla z(x, y)],$$

and since the irradiance equation depends on the unit normal, we take into account the

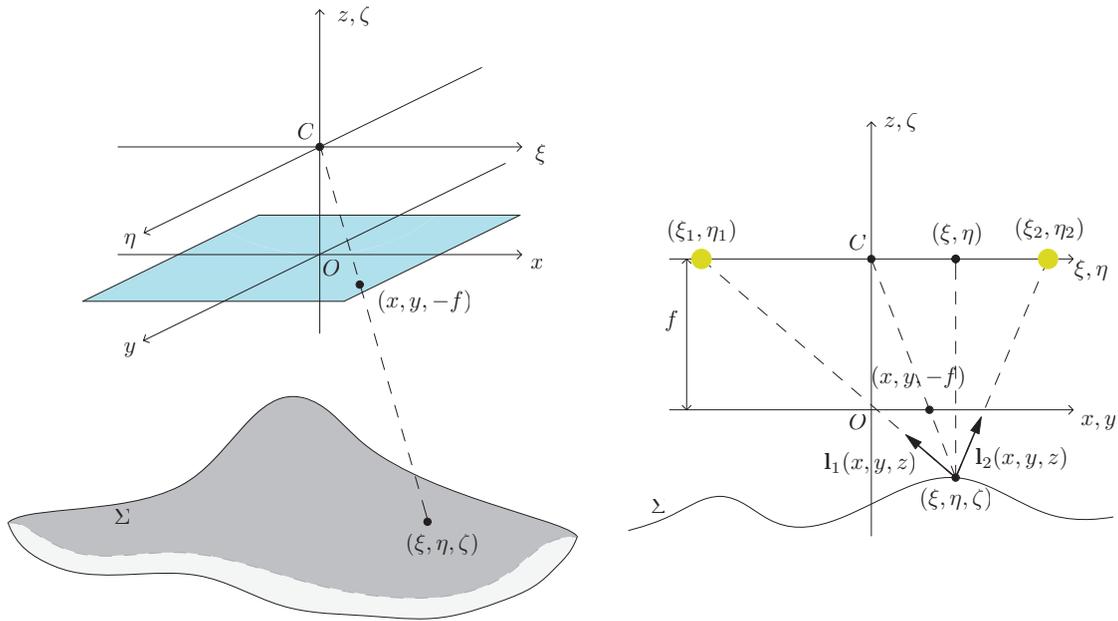


Figure 1. On the left is a diagram of the perspective view geometry. In the perspective world (i.e., the image coordinate system given by $O\xi\zeta$) the light observed at image point (x, y) comes from the real point $[\xi, \eta, \zeta]$ on the surface. On the right is a schematic section of the geometry on the left illustrating the positions of the lights placed on the optical plane. Clearly the light rays from nearby point light sources are not parallel.

following unit normal vector:

$$(2.3) \quad \mathbf{n}(x, y) = \frac{\bar{\mathbf{n}}(x, y)}{|\bar{\mathbf{n}}(x, y)|}.$$

We consider the well-known irradiance equation for Lambertian surfaces, given by the cosine law by the following inner product:

$$(2.4) \quad I(x, y) = \rho(x, y)(\mathbf{n}(x, y) \cdot \mathbf{l}(x, y, \dots)),$$

where $I : \bar{\Omega}_p \rightarrow [0, 1]$ is the image function, $\rho(x, y)$ is the unknown albedo, and $\mathbf{l}(x, y, \dots)$ is the light source direction. Let us emphasize that the light direction \mathbf{l} incident to the surface depends explicitly on the image points (x, y) and on other factors that we shall specify in the next section. The endoscopic formulation assumes close light source illumination, a property that will be considered in our model construction.

3. A new photometric stereo model for nearby light sources. In most papers dealing with PSfS, such as [33, 28], a single light source is placed at the optical center. This model is difficult to realize when considering a camera close to the inspected object. Since photometric stereo uses several light sources, we consider more realistic placements (ξ_j, η_j) on the optical plane, that is, $\zeta = 0$. We therefore define the light directions as

$$(3.1) \quad \mathbf{l}_j(x, y, z) = \frac{[-\frac{\xi_j f}{z} - x, -\frac{\eta_j f}{z} - y, f]}{\sqrt{(x + \frac{\xi_j f}{z})^2 + (y + \frac{\eta_j f}{z})^2 + f^2}} = \frac{\bar{\mathbf{l}}_j(x, y, z)}{q_j(x, y, z)},$$

where

$$(3.2) \quad q_j(x, y, z) = \sqrt{\left(x + \frac{\xi_j f}{z}\right)^2 + \left(y + \frac{\eta_j f}{z}\right)^2 + f^2}.$$

Let us emphasize that for this model, the light directions depend not only on the point (x, y) , but since they are displaced from the optical center, they also depend on z . This introduces a nonlinearity regardless of the derivatives of z . Now, in order to make the mathematical model more consistent to the actual physical aspects, we define the light attenuation function specifically set for two kinds of light attenuations. However, theoretically speaking the attenuation can be generalized to any behavior of light propagation.

3.1. Two types of light attenuation. We now consider two different kinds of light attenuation. The first factor is due to the reduction of light energy proportional to the inverse squared distance between the light source and object. The second factor of attenuation we describe is a result of a realistic directional lighting model of a real surface mount LED light. However, in principle, we could choose among many different continuous attenuation models to suit our needs.

3.1.1. Light attenuation due to distance. The standard way to model attenuation of the light intensity is to compute the distance between the light source and the surface explicitly. For this purpose, let us consider the functions $\mathbf{r}_1(x, y, z)$ and $\mathbf{r}_2(x, y, z)$ as the distances between the point of the surface $(\xi(x, y), \eta(x, y), \zeta(x, y))$ and the respective light source; see Figure 2.

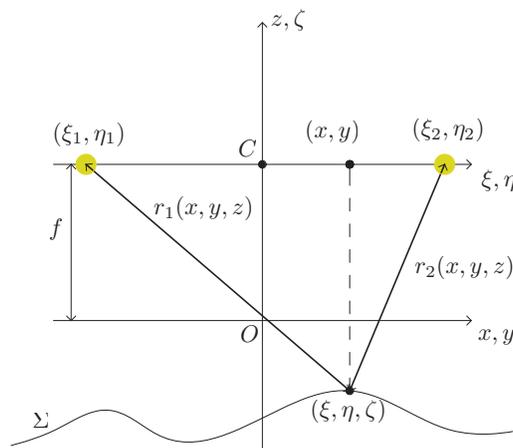


Figure 2. The distances $(\mathbf{r}_1$ and $\mathbf{r}_2)$ between the light sources and the surface Σ are taken into account as a dissipative factor of the light intensity.

In this case, since the light sources are shifted with respect to the origin, the distance

functions are given by

$$(3.3) \quad r_j(x, y, z) = \text{dist} \left((\xi_j, \eta_j, 0), \left(-x \frac{z}{f}, -y \frac{z}{f}, z \right) \right) \\ = \frac{z}{f} \sqrt{\left(x + \frac{\xi_j f}{z} \right)^2 + \left(y + \frac{\eta_j f}{z} \right)^2 + f^2} = \frac{z}{f} q_j(x, y, z).$$

The attenuation factor can then be written as r_j^{-2} .

3.1.2. Radial attenuation of the light. Many existing light sources are directional. That is to say, they are bright along a principle direction and become dimmer at angles further from the main direction. This behavior can be observed in Figure 3 and can be effectively modeled by multiplication with $\cos(\theta)^\mu$, where μ is the attenuation coefficient and is reminiscent of the specular model for surface reflectance.

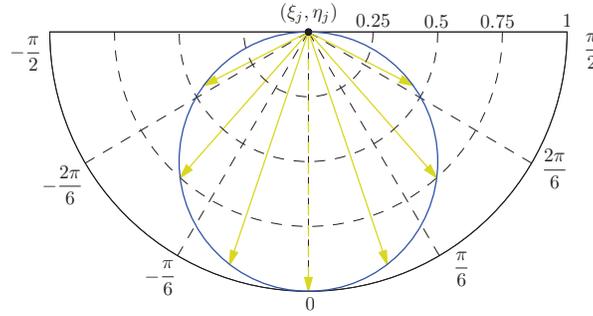


Figure 3. The intensity of light for a directional light source pointing downwards.

The attenuation factor is easily computable since

$$(3.4) \quad \cos(\theta)^\mu(x, y, z) = (\mathbf{l}_j(x, y, z) \cdot (0, 0, 1))^\mu = \frac{f^\mu}{q_j^\mu(x, y, z)},$$

where $\mathbf{l}_j(x, y, z)$ is the i th light source placed at (ξ_j, η_j) .

Resuming, both attenuation effects can be expressed by multiplication of the following factor:

$$(3.5) \quad a_j(x, y, z) = \frac{f^\mu}{r_j^2(x, y, z) q_j^\mu(x, y, z)} = \frac{f^{\mu+2}}{z^2 q_j^{\mu+2}(x, y, z)}.$$

3.2. EPPS₂ model. Our model for the SfEPPS₂ problem considers the following set of irradiance equations:

$$(3.6) \quad I_j(x, y) = \rho(x, y) \frac{\bar{\mathbf{l}}_j(x, y, z) \cdot \mathbf{n}(x, y)}{r_j^2(x, y, z) q_j^{\mu+1}(x, y, z)} f^\mu = \rho(x, y) \frac{\bar{\mathbf{l}}_j(x, y, z) \cdot \bar{\mathbf{n}}(x, y) f^{\mu+2}}{z^2 q_j^{\mu+3}(x, y, z) |\bar{\mathbf{n}}(x, y)|},$$

where

$$(3.7) \quad \bar{I}_j(x, y, z) \cdot \bar{\mathbf{n}}(x, y) = \frac{z}{f^2} \left(fz_x \left(-\frac{\xi_j f}{z} - x \right) + fz_y \left(-\frac{\eta_j f}{z} - y \right) + f(z + xz_x + yz_y) \right) = -\xi_j z_x - \eta_j z_y + \frac{z^2}{f}.$$

Now, solving the SfEPPS₂ from a mathematical point of view consists of solving the following system of nonlinear PDEs of Hamilton–Jacobi type (with Dirichlet boundary condition $g(x, y)$):

$$(3.8) \quad \begin{cases} I_1(x, y) = \rho(x, y) \frac{\bar{I}_1(x, y, z) \cdot \bar{\mathbf{n}}(x, y) f^{\mu+2}}{z^2 q_1^{\mu+3}(x, y, z) |\bar{\mathbf{n}}(x, y)|}, & (x, y) \in \Omega_p, \\ I_2(x, y) = \rho(x, y) \frac{\bar{I}_2(x, y, z) \cdot \bar{\mathbf{n}}(x, y) f^{\mu+2}}{z^2 q_2^{\mu+3}(x, y, z) |\bar{\mathbf{n}}(x, y)|}, & (x, y) \in \Omega_p, \\ z(x, y) = g(x, y), & (x, y) \in \partial\Omega_p. \end{cases}$$

Our strategy to solve such a problem is to merge the irradiance equations (3.8) by noting that the nonvanishing quantity $\frac{\rho(x, y) f^{\mu+2}}{|\bar{\mathbf{n}}(x, y)| z^2}$ is present in both equations. We merge such equations as follows:

$$\underbrace{\frac{I_1(x, y) q_1^{\mu+3}(x, y, z)}{\bar{I}_1(x, y, z) \cdot \bar{\mathbf{n}}(x, y)}}_{\text{first equation of (3.8)}} = \overbrace{\frac{\rho(x, y) f^{\mu+2}}{|\bar{\mathbf{n}}(x, y)| z^2}}^{\text{second equation of (3.8)}} = \frac{I_2(x, y) q_2^{\mu+3}(x, y, z)}{\bar{I}_2(x, y, z) \cdot \bar{\mathbf{n}}(x, y)},$$

getting

$$(3.9) \quad \frac{I_1(x, y) q_1^{\mu+3}(x, y, z)}{-\xi_1 z_x - \eta_1 z_y + \frac{z^2}{f}} = \frac{I_2(x, y) q_2^{\mu+3}(x, y, z)}{-\xi_2 z_x - \eta_2 z_y + \frac{z^2}{f}}.$$

Merging the irradiance equations allowed us to eliminate the nonlinearity with respect to the partial derivatives of z contained in $|\bar{\mathbf{n}}(x, y)|$. It makes the resulting problem completely independent of the albedo.

After some algebra, we obtain the following first-order quasi-linear PDE:

$$(3.10) \quad \begin{aligned} & \left(I_1(x, y) q_1^{\mu+3}(x, y, z) \xi_2 - I_2(x, y) q_2^{\mu+3}(x, y, z) \xi_1 \right) z_x \\ & + \left(I_1(x, y) q_1^{\mu+3}(x, y, z) \eta_2 - I_2(x, y) q_2^{\mu+3}(x, y, z) \eta_1 \right) z_y \\ & = \left(I_1(x, y) q_1^{\mu+3}(x, y, z) - I_2(x, y) q_2^{\mu+3}(x, y, z) \right) \frac{z^2}{f}, \end{aligned}$$

resumed as follows:

$$(3.11) \quad \begin{cases} \mathbf{b}(x, y, z) \cdot \nabla z(x, y) = s(x, y, z), & (x, y) \in \Omega_p, \\ z(x, y) = g(x, y), & (x, y) \in \partial\Omega_p, \end{cases}$$

where

$$(3.12) \quad \mathbf{b}(x, y, z) = \left[I_1(x, y)q_1^{\mu+3}(x, y, z)\xi_2 - I_2(x, y)q_2^{\mu+3}(x, y, z)\xi_1, \right.$$

$$(3.13) \quad \left. I_1(x, y)q_1^{\mu+3}(x, y, z)\eta_2 - I_2(x, y)q_2^{\mu+3}(x, y, z)\eta_1 \right],$$

and

$$(3.14) \quad s(x, y, z) = \left(I_1(x, y)q_1^{\mu+3}(x, y, z) - I_2(x, y)q_2^{\mu+3}(x, y, z) \right) \frac{z^2}{f}.$$

3.3. Handling light attenuations. In this section we emphasize the ability of our model to handle any type of light attenuation. That is, instead of considering the light attenuation terms

$$(3.15) \quad \frac{1}{r_j^2(x, y, z)} \quad \text{attenuation by distance,}$$

$$(3.16) \quad \frac{f^\mu}{q_j^\mu(x, y, z)} \quad \text{radial attenuation}$$

previously introduced, let us consider the function $a_j(x, y, z)$ as a general attenuation function of the i th light source that multiplies the irradiance equation as follows:

$$(3.17) \quad I_j(x, y) = \rho(x, y)a_j(x, y, z)\mathbf{l}_j(x, y, z) \cdot \mathbf{n}(x, y) = \rho(x, y)a_j(x, y, z) \frac{\bar{\mathbf{l}}_j(x, y, z) \cdot \bar{\mathbf{n}}(x, y)}{q_j(x, y, z)|\bar{\mathbf{n}}(x, y)|}.$$

By merging the two irradiance equations in (3.6) and by noting that $\frac{\rho(x, y)}{|\bar{\mathbf{n}}(x, y)|}$ is present in both of them, we get

$$(3.18) \quad \frac{I_1(x, y) \frac{q_1(x, y, z)}{a_1(x, y, z)}}{-f\xi_1 \frac{z_x}{z} - f\eta_1 \frac{z_y}{z} + z} = \frac{I_2(x, y) \frac{q_2(x, y, z)}{a_2(x, y, z)}}{-f\xi_2 \frac{z_x}{z} - f\eta_2 \frac{z_y}{z} + z}.$$

After some algebra we get the following first-order quasi-linear PDE:

$$(3.19) \quad \left(I_1(x, y) \frac{q_1(x, y, z)}{a_1(x, y, z)} \xi_2 - I_2(x, y) \frac{q_2(x, y, z)}{a_2(x, y, z)} \xi_1 \right) z_x \\ + \left(I_1(x, y) \frac{q_1(x, y, z)}{a_1(x, y, z)} \eta_2 - I_2(x, y) \frac{q_2(x, y, z)}{a_2(x, y, z)} \eta_1 \right) z_y \\ = \left(I_1(x, y) \frac{q_1(x, y, z)}{a_1(x, y, z)} - I_2(x, y) \frac{q_2(x, y, z)}{a_2(x, y, z)} \right) \frac{z^2}{f}.$$

The well-posedness results of the EPPS₂ model will apply to the most general light attenuation models.

4. Well-posedness of the new model. The first-order quasi-linear PDEs with the Dirichlet boundary condition (3.11) can generally be solved by *characteristic strip expansion* [24]. Focusing on this particular method, we prove the well-posedness of (3.11) by proving some nice properties of the vector field $\mathbf{b}(x, y, z) = [b_1(x, y, z), b_2(x, y, z)]$. In fact, this field describes the projected characteristic field since the complete characteristic system is usually formulated by

$$(4.1) \quad \begin{cases} \text{(a)} & \dot{x}_c(t) = b_1(x_c(t), y_c(t), z(t)), \\ \text{(b)} & \dot{y}_c(t) = b_2(x_c(t), y_c(t), z(t)), \\ \text{(c)} & \dot{z}(t) = s(x_c(t), y_c(t), z(t)), \end{cases}$$

where t is the parameterization variable and $z(t)$ is the value of z on the projected characteristic, that is, $z(t) = z(x_c(t), y_c(t))$.

The initial condition to integrate this system of ODEs is taken from the values of the function z known on the boundary $\Gamma_{in} := \{(x, y) \in \partial\Omega_p : \mathbf{b} \cdot \nu \leq 0\}$, where ν is the outgoing normal vector to the curve $\partial\Omega_p$. Then,

$$\begin{cases} \text{(a}_0) & x_c(0) = x_0, \\ \text{(b}_0) & y_c(0) = y_0, \\ \text{(c}_0) & z(x_c(0), y_c(0)) = g(x_0, y_0), \end{cases}$$

with $(x_0, y_0) \in \partial\Omega_p$.

Now, since singular points prevent the proper information propagation of the characteristic method, first we prove the following lemma, ensuring that such points do not exist for the differential problem (3.11). Let us start with the following property.

Lemma 4.1. *The vector field $\mathbf{b}(x, y, z)$ does not vanish in Ω_p , that is,*

$$(4.2) \quad |\mathbf{b}(x, y, z)| \neq 0 \quad \forall (x, y) \in \Omega_p, \quad \forall z \in C^1(\bar{\Omega}_p).$$

Proof. Let us prove the lemma by contradiction, assuming that there exist a point $(\tilde{x}, \tilde{y}) \in \Omega_p$ and a function $z \in C^1(\bar{\Omega}_p)$ such that

$$(4.3) \quad \begin{cases} b_1(\tilde{x}, \tilde{y}, \tilde{z}) = 0, \\ b_2(\tilde{x}, \tilde{y}, \tilde{z}) = 0, \end{cases}$$

where $\tilde{z} = z(\tilde{x}, \tilde{y})$. According to the irradiance equations (3.17), by omitting the dependence on $(\tilde{x}, \tilde{y}, \tilde{z})$ for readability, we get

$$(4.4) \quad \begin{cases} \tilde{\rho}_{a_1} \frac{\bar{\mathbf{l}}_1 \cdot \bar{\mathbf{n}} \tilde{q}_1}{\tilde{q}_1 |\bar{\mathbf{n}}| \tilde{a}_1} \xi_2 - \tilde{\rho}_{a_2} \frac{\bar{\mathbf{l}}_2 \cdot \bar{\mathbf{n}} \tilde{q}_2}{\tilde{q}_2 |\bar{\mathbf{n}}| \tilde{a}_2} \xi_1 = 0, \\ \tilde{\rho}_{a_1} \frac{\bar{\mathbf{l}}_1 \cdot \bar{\mathbf{n}} \tilde{q}_1}{\tilde{q}_1 |\bar{\mathbf{n}}| \tilde{a}_1} \eta_2 - \tilde{\rho}_{a_2} \frac{\bar{\mathbf{l}}_2 \cdot \bar{\mathbf{n}} \tilde{q}_2}{\tilde{q}_2 |\bar{\mathbf{n}}| \tilde{a}_2} \eta_1 = 0, \end{cases}$$

which can be simplified in the following equation:

$$(4.5) \quad \begin{cases} \bar{\mathbf{l}}_1 \cdot \bar{\mathbf{n}} \xi_2 - \bar{\mathbf{l}}_2 \cdot \bar{\mathbf{n}} \xi_1 = 0, \\ \bar{\mathbf{l}}_1 \cdot \bar{\mathbf{n}} \eta_2 - \bar{\mathbf{l}}_2 \cdot \bar{\mathbf{n}} \eta_1 = 0, \end{cases}$$

that is,

$$(4.6) \quad \begin{cases} \bar{\mathbf{n}} \cdot (\bar{\mathbf{I}}_1 \xi_2 - \bar{\mathbf{I}}_2 \xi_1) = 0, \\ \bar{\mathbf{n}} \cdot (\bar{\mathbf{I}}_1 \eta_2 - \bar{\mathbf{I}}_2 \eta_1) = 0, \end{cases}$$

which means that the vectors

$$(4.7) \quad \mathbf{u}(\tilde{x}, \tilde{y}, \tilde{z}) = \bar{\mathbf{I}}_1(\tilde{x}, \tilde{y}, \tilde{z})\xi_2 - \bar{\mathbf{I}}_2(\tilde{x}, \tilde{y}, \tilde{z})\xi_1$$

and

$$(4.8) \quad \mathbf{v}(\tilde{x}, \tilde{y}, \tilde{z}) = \bar{\mathbf{I}}_1(\tilde{x}, \tilde{y}, \tilde{z})\eta_2 - \bar{\mathbf{I}}_2(\tilde{x}, \tilde{y}, \tilde{z})\eta_1$$

have to be simultaneously orthogonal to the normal $\bar{\mathbf{n}}(\tilde{x}, \tilde{y})$. That is, \mathbf{u} and \mathbf{v} belong to the tangent plane of the surface. Let us write the system (4.6) explicitly; using (3.1) and (2.2) we have

$$(4.9) \quad \begin{cases} (\xi_2 - \xi_1)z^2 + f(\xi_1\eta_2 - \xi_2\eta_1)z_y = 0, \\ (\eta_1 - \eta_2)z^2 + f(\xi_1\eta_2 - \xi_2\eta_1)z_x = 0. \end{cases}$$

Let us span all the possible cases starting with the possibility of computing the partial derivatives as follows:

$$(4.10) \quad z_y = \frac{(\xi_1 - \xi_2)z^2}{(\xi_1\eta_2 - \xi_2\eta_1)f},$$

$$(4.11) \quad z_x = \frac{(\eta_1 - \eta_2)z^2}{(\xi_2\eta_1 - \xi_1\eta_2)f},$$

assuming $\xi_1\eta_2 - \xi_2\eta_1 \neq 0$. If we substitute such derivatives in (3.7), we find that the image functions I_1 and I_2 both take zero value, which means that the image data at point (\tilde{x}, \tilde{y}) vanishes in both images, which is obviously a contradiction since no information is available from the image data. On the other hand, considering $\xi_1\eta_2 - \xi_2\eta_1 = 0$, it follows from (4.9) that $(\xi_1, \eta_1) \equiv (\xi_2, \eta_2)$ (since $z < 0$), in contradiction to the photometric stereo definition. ■

This guarantees that the solution can be integrated along the characteristics from the system of ODEs (4.1) without being prevented by singular points.

4.1. Well-posedness for Lipschitz surfaces. In order to complete the theoretical analysis we extend the uniqueness results of the differential problem (3.19) in the case of a Lipschitz solution. We assume the points where it is not differentiable as the family of regular curves $(\gamma_1(t), \dots, \gamma_k(t))$, where t is the argument of the parametric representation.

It is clear that these curves on the surface *reflect* their presence on the images, making the functions $I_1(x, y)$ and $I_2(x, y)$ discontinuous along the same curves $(\gamma_1(t), \dots, \gamma_k(t))$. Now, since $\mathbf{b}(x, y, z)$ and $s(x, y, z)$ contain the image functions, if we consider our differential problem as an inverse problem of SfEP with the photometric stereo technique, searching for a weak solution (i.e., defined almost everywhere) implies a study of the quasi-linear PDE with discontinuous coefficients.

Our purpose is to prove the uniqueness of solution of (3.19) in the Lipschitz function space via the method of characteristics. The meaning of weak solution here is intended as a

combination of classical solutions, each defined on a different domain. The following theorem shows that these domains are then going to be patched together in such a way that, across the boundaries γ between domains on which there are discontinuities in some derivatives, (3.19) is satisfied.

Theorem 4.2. *Let $\gamma(t) \in \Omega_p$ be a regular curve of discontinuity for the function $\mathbf{b}(x, y, z)$ (and $s(x, y, z)$). Let $(\tilde{x}, \tilde{y}) \in \gamma(t)$, and let $\mathbf{n}^\gamma(\tilde{x}, \tilde{y})$ be the outgoing normal with respect to the set Ω_p^+ (see Figure 4); then we have*

$$(4.12) \quad \left(\lim_{\substack{(x,y) \rightarrow (\tilde{x}, \tilde{y}) \\ (x,y) \in \Omega_p^+}} \mathbf{b}(x, y, z(x, y)) \cdot \mathbf{n}^\gamma(\tilde{x}, \tilde{y}) \right) \left(\lim_{(x,y) \in \Omega_p^-} \mathbf{b}(x, y, z(x, y)) \cdot \mathbf{n}^\gamma(\tilde{x}, \tilde{y}) \right) \geq 0.$$

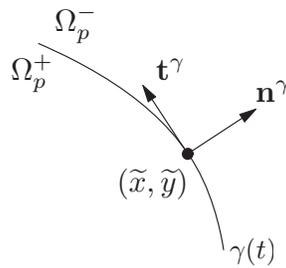


Figure 4. *The curve of discontinuity $\gamma(t)$ separates the image domain in two subsets Ω_p^+ and Ω_p^- where the classical solution is well defined. The vectors \mathbf{n}^γ and \mathbf{t}^γ are, respectively, the normal and the tangent vectors of $\gamma(t)$.*

Proof. Let us define the following quantities for the discontinuous images:

$$\begin{aligned} I_1^+ &:= \lim_{\substack{(x,y) \rightarrow (\tilde{x}, \tilde{y}) \\ (x,y) \in \Omega_p^+}} I_1(x, y), & I_1^- &:= \lim_{\substack{(x,y) \rightarrow (\tilde{x}, \tilde{y}) \\ (x,y) \in \Omega_p^-}} I_1(x, y), \\ I_2^+ &:= \lim_{\substack{(x,y) \rightarrow (\tilde{x}, \tilde{y}) \\ (x,y) \in \Omega_p^+}} I_2(x, y), & I_2^- &:= \lim_{\substack{(x,y) \rightarrow (\tilde{x}, \tilde{y}) \\ (x,y) \in \Omega_p^-}} I_2(x, y), \end{aligned}$$

and these other quantities considering that z is a continuous function (i.e., $\tilde{z} = z(\tilde{x}, \tilde{y})$)

$$\begin{aligned} \tilde{q}_1 &:= q_1(\tilde{x}, \tilde{y}, \tilde{z}), & \tilde{q}_2 &:= q_2(\tilde{x}, \tilde{y}, \tilde{z}), \\ \tilde{a}_1 &:= a_1(\tilde{x}, \tilde{y}, \tilde{z}), & \tilde{a}_2 &:= a_2(\tilde{x}, \tilde{y}, \tilde{z}), \end{aligned}$$

in order to work with the vector field $\mathbf{b}(x, y, z)$ in the neighborhood of (\tilde{x}, \tilde{y}) considering

$$(4.13) \quad \begin{aligned} \lim_{\substack{(x,y) \rightarrow (\tilde{x}, \tilde{y}) \\ (x,y) \in \Omega_p^+}} \mathbf{b}(x, y, z(x, y)) &= [b_1^+, b_2^+] = \left[I_1^+ \frac{\tilde{q}_1}{\tilde{a}_1} \xi_2 - I_2^+ \frac{\tilde{q}_2}{\tilde{a}_2} \xi_1, I_1^+ \frac{\tilde{q}_1}{\tilde{a}_1} \eta_2 - I_2^+ \frac{\tilde{q}_2}{\tilde{a}_2} \eta_1 \right], \\ \lim_{\substack{(x,y) \rightarrow (\tilde{x}, \tilde{y}) \\ (x,y) \in \Omega_p^-}} \mathbf{b}(x, y, z(x, y)) &= [b_1^-, b_2^-] = \left[I_1^- \frac{\tilde{q}_1}{\tilde{a}_1} \xi_2 - I_2^- \frac{\tilde{q}_2}{\tilde{a}_2} \xi_1, I_1^- \frac{\tilde{q}_1}{\tilde{a}_1} \eta_2 - I_2^- \frac{\tilde{q}_2}{\tilde{a}_2} \eta_1 \right]. \end{aligned}$$

Calling $[n_1^\gamma, n_2^\gamma] = [n_1^\gamma(\tilde{x}, \tilde{y}), n_2^\gamma(\tilde{x}, \tilde{y})]$ the two coordinates of $\mathbf{n}^\gamma(\tilde{x}, \tilde{y})$ and replacing (4.13) in inequality (4.12), we obtain

$$(4.14) \quad b_1^+ b_1^- (n_1^\gamma)^2 + b_2^+ b_2^- (n_2^\gamma)^2 + (b_1^+ b_2^- + b_2^+ b_1^-) n_1^\gamma n_2^\gamma \geq 0.$$

By substitution of (4.13) into (4.14) and replacing the image functions with the respective reflectance functions (3.17), we have

$$(4.15) \quad \begin{aligned} & \left(\tilde{\rho}_{a_1} \frac{\bar{\mathbf{l}}_1 \cdot \bar{\mathbf{n}}^+}{\tilde{q}_1^2 |\bar{\mathbf{n}}^+|} \frac{\tilde{q}_1}{\tilde{a}_1} \xi_2 - \tilde{\rho}_{a_2} \frac{\bar{\mathbf{l}}_2 \cdot \bar{\mathbf{n}}^+}{\tilde{q}_2^2 |\bar{\mathbf{n}}^+|} \frac{\tilde{q}_2}{\tilde{a}_2} \xi_1 \right) \left(\tilde{\rho}_{a_1} \frac{\bar{\mathbf{l}}_1 \cdot \bar{\mathbf{n}}^-}{\tilde{q}_1^2 |\bar{\mathbf{n}}^-|} \frac{\tilde{q}_1}{\tilde{a}_1} \xi_2 - \tilde{\rho}_{a_2} \frac{\bar{\mathbf{l}}_2 \cdot \bar{\mathbf{n}}^-}{\tilde{q}_2^2 |\bar{\mathbf{n}}^-|} \frac{\tilde{q}_2}{\tilde{a}_2} \xi_1 \right) (n_1^\gamma)^2 \\ & + \left(\tilde{\rho}_{a_1} \frac{\bar{\mathbf{l}}_1 \cdot \bar{\mathbf{n}}^+}{\tilde{q}_1^2 |\bar{\mathbf{n}}^+|} \frac{\tilde{q}_1}{\tilde{a}_1} \eta_2 - \tilde{\rho}_{a_2} \frac{\bar{\mathbf{l}}_2 \cdot \bar{\mathbf{n}}^+}{\tilde{q}_2^2 |\bar{\mathbf{n}}^+|} \frac{\tilde{q}_2}{\tilde{a}_2} \eta_1 \right) \left(\tilde{\rho}_{a_1} \frac{\bar{\mathbf{l}}_1 \cdot \bar{\mathbf{n}}^-}{\tilde{q}_1^2 |\bar{\mathbf{n}}^-|} \frac{\tilde{q}_1}{\tilde{a}_1} \eta_2 - \tilde{\rho}_{a_2} \frac{\bar{\mathbf{l}}_2 \cdot \bar{\mathbf{n}}^-}{\tilde{q}_2^2 |\bar{\mathbf{n}}^-|} \frac{\tilde{q}_2}{\tilde{a}_2} \eta_1 \right) (n_2^\gamma)^2 \\ & + \left[\left(\tilde{\rho}_{a_1} \frac{\bar{\mathbf{l}}_1 \cdot \bar{\mathbf{n}}^+}{\tilde{q}_1^2 |\bar{\mathbf{n}}^+|} \frac{\tilde{q}_1}{\tilde{a}_1} \xi_2 - \tilde{\rho}_{a_2} \frac{\bar{\mathbf{l}}_2 \cdot \bar{\mathbf{n}}^+}{\tilde{q}_2^2 |\bar{\mathbf{n}}^+|} \frac{\tilde{q}_2}{\tilde{a}_2} \xi_1 \right) \left(\tilde{\rho}_{a_1} \frac{\bar{\mathbf{l}}_1 \cdot \bar{\mathbf{n}}^-}{\tilde{q}_1^2 |\bar{\mathbf{n}}^-|} \frac{\tilde{q}_1}{\tilde{a}_1} \eta_2 - \tilde{\rho}_{a_2} \frac{\bar{\mathbf{l}}_2 \cdot \bar{\mathbf{n}}^-}{\tilde{q}_2^2 |\bar{\mathbf{n}}^-|} \frac{\tilde{q}_2}{\tilde{a}_2} \eta_1 \right) \right. \\ & \left. + \left(\tilde{\rho}_{a_1} \frac{\bar{\mathbf{l}}_1 \cdot \bar{\mathbf{n}}^-}{\tilde{q}_1^2 |\bar{\mathbf{n}}^-|} \frac{\tilde{q}_1}{\tilde{a}_1} \xi_2 - \tilde{\rho}_{a_2} \frac{\bar{\mathbf{l}}_2 \cdot \bar{\mathbf{n}}^-}{\tilde{q}_2^2 |\bar{\mathbf{n}}^-|} \frac{\tilde{q}_2}{\tilde{a}_2} \xi_1 \right) \left(\tilde{\rho}_{a_1} \frac{\bar{\mathbf{l}}_1 \cdot \bar{\mathbf{n}}^+}{\tilde{q}_1^2 |\bar{\mathbf{n}}^+|} \frac{\tilde{q}_1}{\tilde{a}_1} \eta_2 - \tilde{\rho}_{a_2} \frac{\bar{\mathbf{l}}_2 \cdot \bar{\mathbf{n}}^+}{\tilde{q}_2^2 |\bar{\mathbf{n}}^+|} \frac{\tilde{q}_2}{\tilde{a}_2} \eta_1 \right) \right] n_1^\gamma n_2^\gamma \geq 0. \end{aligned}$$

Finally, we can write (4.15) as follows:

$$(4.16) \quad \begin{aligned} & (i_1^+ \xi_2 - i_2^+ \xi_2)(i_1^- \xi_2 - i_2^- \xi_1)(n_1^\gamma)^2 + (i_1^+ \eta_2 - i_2^+ \eta_1)(i_1^- \eta_2 - i_2^- \eta_1)(n_2^\gamma)^2 \\ & + ((i_1^+ \xi_2 - i_2^+ \xi_1)(i_1^- \eta_2 - i_2^- \eta_1) + (i_1^- \xi_2 - i_2^- \xi_1)(i_1^+ \eta_2 - i_2^+ \eta_1)) n_1^\gamma n_2^\gamma \geq 0, \end{aligned}$$

where

$$(4.17) \quad i_j^+ = \bar{\mathbf{l}}_j \cdot \bar{\mathbf{n}}^+ = \frac{\tilde{z}}{f} \left(-z_x^+ \frac{\xi_j f}{\tilde{z}} - z_y^+ \frac{\eta_j f}{\tilde{z}} + \tilde{z} \right),$$

$$(4.18) \quad i_j^- = \bar{\mathbf{l}}_j \cdot \bar{\mathbf{n}}^- = \frac{\tilde{z}}{f} \left(-z_x^- \frac{\xi_j f}{\tilde{z}} - z_y^- \frac{\eta_j f}{\tilde{z}} + \tilde{z} \right),$$

with $j = 1, 2$. We have

$$(4.19) \quad i_j^+ = i_j^- - \mathbf{\Delta} \cdot [\xi_j, \eta_j],$$

where $\mathbf{\Delta} \equiv [\delta_x, \delta_y] = [z_x^+ - z_x^-, z_y^+ - z_y^-]$. We can write (4.16) as

$$(4.20) \quad \begin{aligned} & \left[\left(i_1^- - \mathbf{\Delta} \cdot (\xi_1, \eta_1) \right) \xi_2 - \left(i_2^- - \mathbf{\Delta} \cdot (\xi_2, \eta_2) \right) \xi_1 \right] (i_1^- \xi_2 - i_2^- \xi_1) (n_1^\gamma)^2 \\ & + \left[\left(i_1^- - \mathbf{\Delta} \cdot (\xi_1, \eta_1) \right) \eta_2 - \left(i_2^- - \mathbf{\Delta} \cdot (\xi_2, \eta_2) \right) \eta_1 \right] (i_1^- \eta_2 - i_2^- \eta_1) (n_2^\gamma)^2 \\ & + \left\{ \left[\left(i_1^- - \mathbf{\Delta} \cdot (\xi_1, \eta_1) \right) \xi_2 - \left(i_2^- - \mathbf{\Delta} \cdot (\xi_2, \eta_2) \right) \xi_1 \right] (i_1^- \eta_2 - i_2^- \eta_1) \right. \\ & \left. + (i_1^- \xi_2 - i_2^- \xi_1) \left[\left(i_1^- - \mathbf{\Delta} \cdot (\xi_1, \eta_1) \right) \eta_2 - \left(i_2^- - \mathbf{\Delta} \cdot (\xi_2, \eta_2) \right) \eta_1 \right] \right\} n_1^\gamma n_2^\gamma \geq 0. \end{aligned}$$

After some algebra we get the following inequality:

$$(4.21) \quad [i_2^-(n_1^\gamma \xi_1 + n_2^\gamma \eta_1) - i_1^-(n_1^\gamma \xi_2 + n_2^\gamma \eta_2)] \\ \left[i_2^-(n_1^\gamma \xi_1 + n_2^\gamma \eta_1) - i_1^-(n_1^\gamma \xi_2 + n_2^\gamma \eta_2) + (\delta_y n_1^\gamma - \delta_x n_2^\gamma)(\xi_2 \eta_1 - \xi_1 \eta_2) \right] \geq 0.$$

In order to conclude the proof, let us recall that

$$(4.22) \quad \delta_y n_1^\gamma - \delta_x n_2^\gamma \equiv 0.$$

That is,

$$(4.23) \quad \delta_y n_1^\gamma - \delta_x n_2^\gamma = n_1^\gamma(z_y^+ - z_y^-) - n_2^\gamma(z_x^+ - z_x^-) = -z_x^+ n_2^\gamma + z_y^+ n_1^\gamma - (-z_x^- n_2^\gamma + z_y^- n_1^\gamma) \\ = \lim_{\substack{(x,y) \rightarrow (\tilde{x}, \tilde{y}) \\ (x,y) \in \Omega_p^+}} \nabla_{[-n_2^\gamma, n_1^\gamma]} z(x, t) - \lim_{\substack{(x,y) \rightarrow (\tilde{x}, \tilde{y}) \\ (x,y) \in \Omega_p^-}} \nabla_{[-n_2^\gamma, n_1^\gamma]} z(x, y) = \nabla_{\mathbf{t}^\gamma}^+ z - \nabla_{\mathbf{t}^\gamma}^- z = 0,$$

where we used the orthogonality of $[-n_2^\gamma, n_1^\gamma] = \mathbf{t}^\gamma$ with respect to the normal vector $[n_1^\gamma, n_2^\gamma]$ of the discontinuity curve $\gamma(t)$; see Figure 4. This means that we have to consider the gradient in the tangential direction with respect to $\gamma(t)$ for both sides Ω_p^+ and Ω_p^- . Now, since we are assuming Lipschitz continuity for the surface z , we have that these two derivatives are equal for all points $(x, y) \in \gamma(t)$.

Finally, we arrive from (4.21) to

$$(i_1^-(n_1^\gamma \xi_2 + n_2^\gamma \eta_2) - i_2^-(n_1^\gamma \xi_1 + n_2^\gamma \eta_1))^2 \geq 0,$$

which concludes the proof. ■

5. Numerical schemes. Next, we consider numerical methods that implement the proposed model. Some of the ideas behind the schemes we consider were proposed [24] for the SfPS problem where finite difference upwind schemes and semi-Lagrangian schemes are used for the forward (and backward in the appendix) approximation of the differential problem (3.11).

Let us start by considering a square domain Ω like the set $[a, b]^2$ (in particular considering in the numerical tests $[-1, 1]^2$) and with a uniform discretization space step $\Delta = (b - a)/n$, where n is the number of intervals that divide each side of the square (that is, $x_i = -1 + i\Delta_x$, $y_j = -1 + j\Delta_y$, with $i, j = 0, \dots, n$). We will denote by $\overline{\Omega}_p^d$ all the points of the lattice belonging to $\overline{\Omega}_p$, by Ω_p^d all the internal points, and by $\partial\Omega_p^d$ all the boundary points.

5.1. Forward numerical schemes. We recall the numerical schemes used for the forward approximation of (3.11) where the propagation of the information is considered as starting from the inflow part of the boundary Γ_{in} . We formulate the differential problem solved by the forward schemes as follows:

$$(5.1) \quad \begin{cases} \mathbf{b}(x, y, z) \cdot \nabla z(x, y) = s(x, y, z) & \text{a.e. } (x, y) \in \Omega_p, \\ z(x, y) = g(x, y) & \forall (x, y) \in \Gamma_{in}. \end{cases}$$

In order to simplify the notation, we shall denote $\mathbf{b}(x_i, y_j, z_{i,j})$ by $\mathbf{b}_{i,j} = [b_{i,j}^1, b_{i,j}^2]$ and $s(x_i, y_j)$ by $s_{i,j}$.

5.1.1. Forward upwind scheme. Let us consider the following implicit upwind scheme, obtained by adding vanishing viscosity:

$$(5.2) \quad b_{i,j}^1 \frac{Z_{i+1,j} - Z_{i-1,j}}{2\Delta_x} + b_{i,j}^2 \frac{Z_{i,j+1} - Z_{i,j-1}}{2\Delta_y} \\ = |b_{i,j}^1| \frac{\Delta_x}{2} \frac{Z_{i+1,j} - 2Z_{i,j} + Z_{i-1,j}}{\Delta_x^2} + |b_{i,j}^2| \frac{\Delta_y}{2} \frac{Z_{i,j+1} - 2Z_{i,j} + Z_{i,j-1}}{\Delta_y^2} + s_{i,j}$$

for $i, j = 1, \dots, n-1$. The artificial diffusion introduced in the right-hand side of (5.2) allows us to follow the vector field \mathbf{b} by considering the most appropriate discretization for the first derivative in order to track the characteristic lines [34, 37]. In particular, it is a consistent numerical scheme of order equal to one with respect to both Δ_x and Δ_y .

All possible simplifications can easily be seen by writing the scheme (5.2) as

$$(5.3) \quad (\text{UW}_F) \quad Z_{i,j}^{(k+1)} = \frac{|b_{i,j}^1(Z_{i,j}^{(k)})| Z_{i-\text{sgn}(b_{i,j}^1(Z_{i,j}^{(k)})),j}^{(k)} + |b_{i,j}^2(Z_{i,j}^{(k)})| Z_{i,j-\text{sgn}(b_{i,j}^2(Z_{i,j}^{(k)}))}^{(k)} + \Delta s_{i,j}(Z_{i,j}^{(k)})}{|b_{i,j}^1(Z_{i,j}^{(k)})| + |b_{i,j}^2(Z_{i,j}^{(k)})|},$$

where $\Delta_x = \Delta_y = \Delta$, and emphasizing the dependence of the functions \mathbf{b} and s on z . That is, $\mathbf{b}_{i,j}(Z_{i,j}^{(k)}) = [b_{i,j}^1(Z_{i,j}^{(k)}), b_{i,j}^2(Z_{i,j}^{(k)})] = [b_1(x_i, y_j, Z^{(k)}(x_i, y_j)), b_2(x_i, y_j, Z^{(k)}(x_i, y_j))]$ and $s_{i,j}(Z_{i,j}^{(k)}) = s(x_i, y_j, Z^{(k)}(x_i, y_j))$. The well-posedness of (5.3) is ensured by Lemma 4.1 since it guarantees that division by $|b_{i,j}^1(Z_{i,j}^{(k)})| + |b_{i,j}^2(Z_{i,j}^{(k)})|$ does not pose any difficulty.

5.1.2. Forward semi-Lagrangian scheme. A second type of numerical approach to solve problem (5.1) is presented in [24]. We derive it considering the following equivalent equation obtained by dividing both sides of (5.1) by the norm of $\mathbf{b}(x, y, z)$:

$$(5.4) \quad \nabla_{\mathbf{\Lambda}} z(x, y) = \frac{s(x, y, z)}{|\mathbf{b}(x, y, z)|} \quad \forall (x, y) \in \Omega_p,$$

with $\mathbf{\Lambda}(x, y, z) = (\Lambda_1(x, y, z), \Lambda_2(x, y, z)) = \frac{\mathbf{b}(x, y, z)}{|\mathbf{b}(x, y, z)|}$ and $\mathbf{\Lambda}_{i,j}(Z_{i,j}^{(k)}) = [\Lambda_{i,j}^1(Z_{i,j}^{(k)}), \Lambda_{i,j}^2(Z_{i,j}^{(k)})] = [\Lambda_1(x_i, y_j, Z^{(k)}(x_i, y_j)), \Lambda_2(x_i, y_j, Z^{(k)}(x_i, y_j))]$ in the discrete domain.

Considering a regular-uniform discretization $\overline{\Omega}_p^d$, as in the previous section, we can finally write the semi-Lagrangian scheme as

$$(5.5) \quad (\text{sLF}) \quad Z_{i,j}^{(k+1)} = Z^{(k)}(x_i - h\Lambda_{i,j}^1(Z_{i,j}^{(k)}), y_j - h\Lambda_{i,j}^2(Z_{i,j}^{(k)})) + \frac{s_{i,j}(Z_{i,j}^{(k)})}{|\mathbf{b}_{i,j}(Z_{i,j}^{(k)})|} h,$$

where the parameter h is greater than zero and assumed to be equal to the grid size Δ , in order to reach the best convergence rate [24].

Let us complete the dissertation concerning the numerical schemes solving (3.11) by considering the backward schemes derived in the appendix.

6. Convergence of the semi-Lagrangian schemes. In this section, we prove the convergence of the semi-Lagrangian numerical schemes giving, as known results, the order of consistency [9]. In fact, for both schemes (backward and forward) we have the following order of consistency:

$$(6.1) \quad \mathcal{O}(h) + \mathcal{O}\left(\frac{\Delta^2}{h}\right),$$

where the first part depends on the approximation of the directional derivative (stopped at order one using the Taylor expansion) and the second has in the numerator the order of the bilinear interpolation. We can obtain a global first-order consistency for $h = \Delta$ as repeated in [24].

In order to prove the convergence of the previous semi-Lagrangian schemes, let us write them using the fixed point operator formulation

$$(6.2) \quad z^{(k+1)} = T_h(z^{(k)})(x, y)$$

as follows:

$$(6.3) \quad \begin{aligned} T_h^b(z)(x, y) &= z((x, y) + h\Lambda(x, y, z(x, y))) - h \frac{s(x, y, z(x, y))}{|\mathbf{b}(x, y, z(x, y))|}, \\ T_h^f(z)(x, y) &= z((x, y) - h\Lambda(x, y, z(x, y))) + h \frac{s(x, y, z(x, y))}{|\mathbf{b}(x, y, z(x, y))|}, \end{aligned}$$

where the superscript letters (*b* and *f*) indicate the backward and the forward numerical schemes, respectively. Our aim is to prove the existence of a unique fixed point for our semi-Lagrangian schemes by solving that they belong to the asymptotically nonexpansive map family of differential operators.

Definition 6.1. Let K be a subset of a Banach space X . A transformation $T : K \rightarrow K$ is said to be asymptotically nonexpansive if for each $u, v \in K$,

$$(6.4) \quad |T^i(u) - T^i(v)| \leq k_i |u - v|,$$

where $\{k_i\}$ is a sequence of real numbers such that $\lim_{i \rightarrow \infty} k_i = 1$.

Furthermore, for asymptotically nonexpansive mapping it may be assumed that $k_i \geq 1$ and that $k_{i+1} \leq k_i$ for $i = 1, 2, \dots$

Our next goal is to prove the convergence of (6.3), for which we want to exploit the following main result of [11].

Theorem 6.2. Let K be a nonempty, closed, convex, and bounded subset of a uniformly convex Banach space X , and let $F : K \rightarrow K$ be asymptotically nonexpansive. Then F has a fixed point.

We are now able to prove the following.

Theorem 6.3. The numerical schemes defined by the operators (6.3) converge.

Proof. Without loss of generality, we start by restricting the image domain to $\overline{\Omega}_p = [0, 1]^2$. Let us focus the attention on the first operator T_h^b introducing some sets which will permit us to understand that this operator is not contractive but asymptotically nonexpansive. For the

moment, let us consider the case where the image functions are smooth in order to use the characteristic strip expansion method. The case where $\bar{\Omega}_p$ is divided by $(\gamma_1(t), \dots, \gamma_k(t))$ in a finite number of regions is a trivial extension of the regular one we explain now. We start with the following sets:

$$\tilde{\Omega}_{out}^h = \{(x^*, y^*) \in \mathbb{R}^2 \setminus \bar{\Omega}_p : (x^*, y^*) = (x, y) + h\mathbf{\Lambda}(x, y) \forall (x, y) \in \bar{\Omega}_p\}$$

and

$$\Omega_{out}^h = \tilde{\Omega}_{out}^h \cup \bar{\Omega}_p,$$

represented in Figure 5.

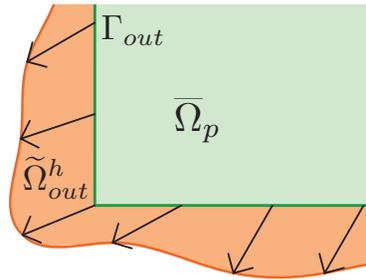


Figure 5. In green is one part of the set $\bar{\Omega}_p$. The set $\tilde{\Omega}_{out}^h$ is emphasized in orange. The arrows represent the vector field $h\mathbf{\Lambda}(x, y, z)$ belonging to the outgoing side of $\partial\Omega_p$, that is, Γ_{out} (marked with a darker color). By definition the length of the arrows is h , so the measure of $\tilde{\Omega}_{out}^h$ tends to zero when h goes to zero.

In order to exploit Theorem 6.2, let us consider as uniformly convex Banach space the L^2 space with respect to appropriate and smooth functions. We proceed by considering the following set of functions:

$$(6.5) \quad \bar{L}^2(\Omega_{out}^h) = \{f \in L^2(\Omega_{out}^h) : f|_{\tilde{\Omega}_{out}^h} = g\}.$$

Once we take $z, v \in \bar{L}^2(\Omega_{out}^h)$ the following relations hold:

$$\begin{aligned} & |T_h^b(z)(x, y) - T_h^b(v)(x, y)|_{L^2(\bar{\Omega}_p)} \\ &= \left| z((x, y) + h\mathbf{\Lambda}(x, y, z(x, y))) - h \frac{s(x, y, z(x, y))}{|\mathbf{b}(x, y, z(x, y))|} \right. \\ & \quad \left. - v((x, y) + h\mathbf{\Lambda}(x, y, v(x, y))) + h \frac{s(x, y, v(x, y))}{|\mathbf{b}(x, y, v(x, y))|} \right|_{L^2(\bar{\Omega}_p)} \\ (6.6) \quad & \leq \left| z((x, y) + h\mathbf{\Lambda}(x, y, z(x, y))) - v((x, y) + h\mathbf{\Lambda}(x, y, v(x, y))) \right|_{L^2(\bar{\Omega}_p)} \\ & \quad + h \left| \frac{s(x, y, v(x, y))}{|\mathbf{b}(x, y, v(x, y))|} - \frac{s(x, y, z(x, y))}{|\mathbf{b}(x, y, z(x, y))|} \right|_{L^2(\bar{\Omega}_p)}. \end{aligned}$$

Let us consider the previous two addends separately. The first can be manipulated to

$$\begin{aligned}
 & \int_{\overline{\Omega}_p} \left(z((x, y) + h\mathbf{\Lambda}(x, y)) - v((x, y) + h\mathbf{\Lambda}(x, y)) \right)^2 dx dy \\
 & \leq \int_{\Omega_{out}^h} \left(z(x, y) - v(x, y) \right)^2 dx dy \\
 (6.7) \quad & = \int_{\tilde{\Omega}_{out}^h} \left(z(x, y) - v(x, y) \right)^2 dx dy + \int_{\overline{\Omega}_p} \left(z(x, y) - v(x, y) \right)^2 dx dy \\
 & = \int_{\overline{\Omega}_p} \left(z(x, y) - v(x, y) \right)^2 dx dy
 \end{aligned}$$

since $z|_{\tilde{\Omega}_{out}^h} = v|_{\tilde{\Omega}_{out}^h} = g$. As for the second term, once we denote

$$(6.8) \quad k(x, y, z(x, y)) = \frac{s(x, y, z(x, y))}{|\mathbf{b}(x, y, z(x, y))|}$$

we exploit the regularity of the functions s and b in order to consider k with k_l as a Lipschitz constant. We obtain

$$(6.9) \quad h|l(x, y, z(x, y)) - k(x, y, v(x, y))|_{L^2(\overline{\Omega}_p)} \leq hk_l|z(x, y) - v(x, y)|_{L^2(\overline{\Omega}_p)}.$$

Then, we can continue the main inequality as follows:

$$(6.10) \quad (6.6) \leq (1 + hk_l)|z(x, y) - v(x, y)|_{L^2(\overline{\Omega}_p)},$$

where clearly

$$(6.11) \quad \lim_{h \rightarrow 0} 1 + hk_l = 1,$$

which allows us to use the result in [11]. Since Theorem 6.2 does not ensure the existence of a unique fixed point, the convergence to the right (and unique) solution of (A.2) is guaranteed by the consistency of (A.5).

The proof can be repeated with respect to T_h^f following the same steps as before since

$$(6.12) \quad T_h^f(z)(x, y) = T_{-h}^b(z)(x, y). \quad \blacksquare$$

The computational cost for every iteration of both schemes is considerable, taking into account that for every internal point of the grid we have to solve a linear system whose dimension is small and constant at that point. In the numerical tests we consider a bilinear interpolation for every pixel, and thus each linear system is of dimension 4.

The stopping criterion is based on the convergence of the sequence z^k through the Cauchy criterion with the infinity norm; that is, since $z^k \rightarrow z$, then the algorithm will stop when

$$(6.13) \quad |z^k - z^{k+1}|_{\Delta_\infty} = \max_{(x_i, y_j) \in \Omega_p^d} |z^k(x_i, y_j) - z^{k+1}(x_i, y_j)| < \varepsilon,$$

with ε chosen conveniently small. Even if it is not consistent with the proof of the convergence (where the L^2 norm was used), the discrete L^∞ norm allows us to achieve good results when used as a stopping criterion.

7. EPPS_n with initial point condition. The small distance between the camera, light sources, and the object increase the likelihood of acquiring images with shadows. This requires that a model is constructed in order to naturally handle such occlusions. Clearly, the presence of shadows in the images represents a loss of information that could compromise the shape recovery. The strategy we consider extends [25], where n images (with $n > 2$) are taken into account in order to compensate for a loss of information where shadows and occlusions occur.

If we have three images, then we can consider the set of unique image pairs and construct the system comprised of the following quasi-linear PDEs:

$$(7.1) \quad \mathbf{b}_{(m,l)}(x, y, z) \cdot \nabla z(x, y) = s_{(m,l)}(x, y, z) \quad \text{in } \Omega_p,$$

of the same type as (3.11), where

$$(7.2) \quad \mathbf{b}(x, y, z) = \left[I_m(x, y) \frac{q_m(x, y, z)}{a_m(x, y, z)} \xi_l - I_l(x, y) \frac{q_l(x, y, z)}{a_l(x, y, z)} \xi_m, \right.$$

$$(7.3) \quad \left. I_m(x, y) \frac{q_m(x, y, z)}{a_m(x, y, z)} \eta_l - I_l(x, y) \frac{q_l(x, y, z)}{a_l(x, y, z)} \eta_m \right],$$

and

$$(7.4) \quad s(x, y, z) = \frac{z^2}{f} \left(I_m(x, y) \frac{q_m(x, y, z)}{a_m(x, y, z)} - I_l(x, y) \frac{q_l(x, y, z)}{a_l(x, y, z)} \right).$$

Remark 1. Since the well-posedness of (3.11) is guaranteed by the boundary condition, the system of quasi-linear PDEs (7.1) is also well-posed only if the same boundary condition $z(x, y) = g(x, y)$ is given.

We can now describe our novel contribution, which is to ensure the well-posedness of the EPPS_n problem by exploiting the linearity of the basic differential formulation (3.11) and reducing it to a single PDE which can handle shadowed regions in a natural fashion. Since (3.11) does not lose the well-posedness if we multiply the equations by a function $w(x, y)$ on both sides (i.e., $\mathbf{b}(x, y, z)$ and $s(x, y, z)$), we are able to define the ingredients of a weighted EPPS_n model (W-EPPS_n) by considering the functions

$$(7.5) \quad \mathbf{b}_n^w(x, y, z) = \sum_{t \in \binom{[n]}{2}} w_t(x, y) \mathbf{b}_t(x, y, z)$$

and

$$(7.6) \quad s_n^w(x, y, z) = \sum_{t \in \binom{[n]}{2}} w_t(x, y) s_t(x, y, z),$$

where $\binom{[n]}{2}$ is the set of pairs of integer indices with no repetition. For example, if $n = 3$, we have $\binom{[3]}{2} = \{(1, 2), (1, 3), (2, 3)\}$.

The complete construction of the W-EPPS_n formulation is therefore

$$(7.7) \quad \begin{cases} \mathbf{b}_n^w(x, y, z) \cdot \nabla z(x, y) = s_n^w(x, y, z) & \text{a.e. } (x, y) \in \Omega_p, \\ z(x, y) = g(x, y) & \forall (x, y) \in \partial\Omega_p. \end{cases}$$

We next explain how shadows influence the definition of the weights (hence of \mathbf{b}_n^w and s_n^w). For the sake of clarity we denote W-EPPS_n as simply EPPS_n, where it implies that weights are used for controlling lights. A key observation is that it is possible to use weight-functions w_t that are vanishing while preserving the well-posedness of the problem. We are not concerned by the signs of the functions w_t ; instead we focus on the set of points where they vanish.

Let us observe that the well-posedness of the differential formulation is guaranteed for image pixels lit in at least two images and preserved if the same condition holds in the multi-image, weighted case. Since we want to exploit the photometric stereo technique, we assume that each pixel is illuminated in at least two images, thereby avoiding reduction to an EPPS₁ problem. Our goal is to consider the weights as switches able to locally nullify the involvement of an image pair in the summations (7.5) and (7.6) when the functions \mathbf{b}_t and s_t for that pair do not contain relevant information due to the presence of shadows in the images involved. Since no ambient light is assumed in our setup, we consider the point $(x, y) \in \overline{\Omega}_p$ shadowed in the i th image when $I_i(x, y) = 0$. Now, by using the Heaviside function, we can easily define the weights as

$$(7.8) \quad w_{(m,l)} = H(I_m(x, y))H(I_l(x, y)).$$

7.1. Removing the boundary condition almost completely. In this section we focus on the applicability of our model extending the EPPS_n model by assuming given data along the boundary $g(x, y)$. Clearly, such a hypothesis compromises the use of that model for many real applications. It is therefore important to find a way to solve the EPPS_n problem while removing the requirement for a priori knowledge of the boundary condition.

Researchers facing the SfS, PSfS, or the EPPS problems often adopt a two stage strategy where surface normals are computed over the domain and then are integrated in order to achieve surface recovery. Usually, the normal computation is realized by considering simplified assumptions where the light propagation is considered uniform [1, 31, 8] and by neglecting the realistic behavior of point light sources. For those cases, the model often consists of a linear system of irradiance equations. Let us emphasize that when close light sources are taken into account the usual linearity of the Lambertian reflectance equation is lost. In this particular case, since the light sources are not placed in the optical center, the normalization of the light direction and the further light attenuations introduce a nonlinearity with respect to z . In other words, given n images, it is very hard to compute the unknown z_x, z_y , and z just by solving the following system of irradiance equations:

$$(7.9) \quad \begin{cases} I_1 = \rho(x, y)a_1(x, y, z) \frac{\bar{\mathbf{l}}_1(x, y, z) \cdot \mathbf{n}(x, y)}{q_1(x, y, z)}, \\ I_2 = \rho(x, y)a_2(x, y, z) \frac{\bar{\mathbf{l}}_2(x, y, z) \cdot \mathbf{n}(x, y)}{q_2(x, y, z)}, \\ \vdots \\ I_n = \rho(x, y)a_n(x, y, z) \frac{\bar{\mathbf{l}}_n(x, y, z) \cdot \mathbf{n}(x, y)}{q_n(x, y, z)}. \end{cases}$$

The approach in this paper is to compute z without completely eliminating the nonlinearity. In fact, the quasi-linear PDE of (3.11) still contains a nonlinear component, but the problem of recovering z can be successfully achieved directly by solving a quasi-linear PDE.

Furthermore, partial images clearly represent a loss of information in the image set, and several authors have approached the problem of surface recovery with occlusions [14, 6]. However, we are not aware of attempts to overcome the problem of missing parts within the framework of the endoscopic problem.

We take the above issues into consideration and use them to help design a numerical strategy for reconstructing the surface using our model. The strategy involves selecting a single arbitrarily valued initial seed point within the reconstruction domain and robustly manipulating the path of the characteristics spreading from this point. We do this in order to numerically integrate the linear differential problem (7.7) so as to let the information travel in the most convenient directions for the whole domain. Our method allows us to use more than one source point if other known points are available within the domain.

7.2. Steering the characteristic field. A numerical strategy will depend on the direction in which information flows. Intuitively it is clear that if we have multiple light directions, then we should also be able to choose image pairs that have vector fields which allow different flow directions. Somewhat unintuitively, however, it turns out that we can use our unique formulation to linearly combine vector fields to choose the most convenient flow direction provided the fields are not parallel to each other. This allows us to manipulate the path along which the information travels. To formalize the idea we present the following result.

Theorem 7.1. *Let $\mathbf{b}_t(x, y, z)$ be the vector field of (7.2), where $t \in \binom{[n]}{2}$. Then, $\forall t_1, t_2 \in \binom{[n]}{2}$, $\forall (x, y) \in \Omega_p$, and $\forall z \in Lip(\overline{\Omega}_p)$, we have*

$$(7.10) \quad \mathbf{b}_{t_1}(x, y, z) \cdot \mathbf{b}_{t_2}(x, y, z) \neq \pm |\mathbf{b}_{t_1}(x, y, z)| |\mathbf{b}_{t_2}(x, y, z)|.$$

Proof. In order to avoid too many parameters, let us fix the indices t_1 and t_2 as (1, 2) and (1, 3), respectively. In order to prove that $\mathbf{b}_{(1,2)}$ and $\mathbf{b}_{(1,3)}$ are never parallel, we consider the contradiction, assuming that there exist a point $(\tilde{x}, \tilde{y}) \in \Omega_p$ and a Lipschitz function \tilde{z} such that

$$(7.11) \quad \mathbf{b}_{(1,2)}(\tilde{x}, \tilde{y}, \tilde{z}) \cdot \mathbf{b}_{(1,3)}(\tilde{x}, \tilde{y}, \tilde{z}) = \pm |\mathbf{b}_{(1,2)}(\tilde{x}, \tilde{y}, \tilde{z})| |\mathbf{b}_{(1,3)}(\tilde{x}, \tilde{y}, \tilde{z})|.$$

To improve readability we omit the dependence on $(\tilde{x}, \tilde{y}, \tilde{z})$. Now, by squaring both sides we have

$$(7.12) \quad (b_{(1,2)}^1 b_{(1,3)}^1 + b_{(1,2)}^2 b_{(1,3)}^2)^2 = ((b_{(1,2)}^1)^2 + (b_{(1,2)}^2)^2)((b_{(1,3)}^1)^2 + (b_{(1,3)}^2)^2).$$

Using the terms from (7.2) and eliminating the not null quantity $\frac{\rho}{|\bar{\mathbf{n}}|}$, this can be rewritten as follows:

$$(7.13) \quad [(\bar{\mathbf{I}}_1 \cdot \bar{\mathbf{n}}\xi_2 - \bar{\mathbf{I}}_2 \cdot \bar{\mathbf{n}}\xi_1)(\bar{\mathbf{I}}_1 \cdot \bar{\mathbf{n}}\xi_3 - \bar{\mathbf{I}}_3 \cdot \bar{\mathbf{n}}\xi_1) + (\bar{\mathbf{I}}_1 \cdot \bar{\mathbf{n}}\eta_2 - \bar{\mathbf{I}}_2 \cdot \bar{\mathbf{n}}\eta_1)(\bar{\mathbf{I}}_1 \cdot \bar{\mathbf{n}}\eta_3 - \bar{\mathbf{I}}_3 \cdot \bar{\mathbf{n}}\eta_1)]^2 \\ = [(\bar{\mathbf{I}}_1 \cdot \bar{\mathbf{n}}\xi_2 - \bar{\mathbf{I}}_2 \cdot \bar{\mathbf{n}}\xi_1)^2 + (\bar{\mathbf{I}}_1 \cdot \bar{\mathbf{n}}\eta_2 - \bar{\mathbf{I}}_2 \cdot \bar{\mathbf{n}}\eta_1)^2][(\bar{\mathbf{I}}_1 \cdot \bar{\mathbf{n}}\xi_3 - \bar{\mathbf{I}}_3 \cdot \bar{\mathbf{n}}\xi_1)^2 + (\bar{\mathbf{I}}_1 \cdot \bar{\mathbf{n}}\eta_3 - \bar{\mathbf{I}}_3 \cdot \bar{\mathbf{n}}\eta_1)^2].$$

After some algebra, we arrive at

$$(7.14) \quad (\xi_3(\eta_2 - \eta_1) + \xi_2(\eta_1 - \eta_3) + \xi_1(\eta_3 - \eta_2))^2 \bar{\mathbf{I}}_1 \cdot \bar{\mathbf{n}} = 0,$$

and assuming that I_1 has no missing parts in (\tilde{x}, \tilde{y}) (i.e., $\bar{\mathbf{I}}_1 \cdot \bar{\mathbf{n}} > 0$) we have that (7.14) is verified only when

$$(7.15) \quad \xi_3(\eta_2 - \eta_1) + \xi_2(\eta_1 - \eta_3) + \xi_1(\eta_3 - \eta_2) = 0.$$

This brings us to the final contradiction since it means that the light source has to be placed at collinear points in the optical plane, which is not our case. ■

In other words, this theorem states that two different vector fields \mathbf{b}_{t_1} and \mathbf{b}_{t_2} cannot be parallel. By using this theorem and adopting a similar fast marching strategy of [21, 27, 26] we steer the direction of the characteristics for the case when $n = 3$.

We perform the following steps:

1. Fix the exact depth value z for a point. In our case it shall usually be towards the center of the image domain, adding all of that point's neighbors to a list of pixels to be visited.
2. Traverse the list of pixels to be visited, and update the value of z for each one by the scheme (5.5) (or (5.3)) derived in section 5.
3. For each newly visited pixel, add its unvisited neighbors to the list of pixels to be visited.
4. In the case of shadow set regions of the scene which are illuminated by only two or less lights, we can change the wavefront propagation direction in order to *surround* the shadow sets (i.e., computing the boundary condition changed the images). If there are exactly two lights, we can solve the appropriate equation in (7.1) provided that the characteristics enable the flow of information from a direction in which it already exists.
5. The above steps are repeated until the L^∞ discrete norm of the difference between the last two elements of the approximating sequence is smaller than a stopping threshold.

For the case where $n > 3$ we iterate over unique triplets of images for each pixel until one of the triplets results in characteristics that permit information flow in the correct direction; otherwise the pixel is skipped.

We remark that in this near-field setup, the convergence of the previous algorithm is not as straightforward as for [27, 26] because here the vector field $\mathbf{b}_{(m,l)}(x, y, z)$ depends on z ; i.e., they are unknown. However, this does not impede the convergence of the scheme.

8. Experimental results. The method we describe has a number of factors that influence the speed and quality of the reconstruction. We now describe the full experimental procedure we followed while investigating the properties of the new model. The numerical schemes were all implemented in MATLAB MEX files in unoptimized C++ using OpenMP for the parallelization and executed in MATLAB using a 2013 Dell Precision M6700 with an Intel i7 CPU clocked at 3GHz and 32GB of RAM.

8.1. Synthetic cases: Eve.

8.1.1. Setup. Synthetic images are useful for testing and accurately quantifying the different properties of a photometric stereo model provided they can faithfully approximate real world data. We therefore first consider a realistic shape obtained from a mannequin head where real depth data has been obtained from a structured light range scanner. This is stored as a height field over a grid with 4-neighbor connectivity. The mannequin face (known as Eve) is approximately 70cm from the virtual camera center. The virtual camera is a perspective pinhole camera (see, for example, [44]) with no simulated lens distortion and has intrinsic parameters $f_x = f_y = 960$, $c_x = 120$, $c_y = 160$. These values are chosen to match the parameters of the structured light scanner that captured the depth data. All depth pixels which are background or beyond 1m are set to NaN to indicate that they are not part of the mesh. Each pixel is also assigned an albedo which we restrict to be between 0 and 1. For the experiments with Eve we use either a uniform albedo of 1 or a synthetically generated albedo based on Perlin noise. We define four virtual light sources which lie on the camera plane $z = 0$ and are positioned at 90° intervals at a radius of either 4cm or 15cm depending on the experiment being performed. Each light source is defined by its direction and the nonlinear light attenuation coefficient μ , as described in section 3.1.1, which we vary between 1 and 100 for Eve. We compute the surface normals at every point on the grid and using the known light positions and attenuation directly synthesize four radiance images using the full lighting model in (3.6). Each of the four images is synthesized with a single light source activated. The resulting images are then scaled to have values between 0 and 255 according to the smallest and largest pixel values across all the images. At this point, zero mean Gaussian noise of strength σ ranging from 0% to 5% of 255 is appended, and finally the images are converted to unsigned 8 bit gray scale. The whole process aims to fully simulate a set of real digital images which match the assumed lighting conditions of the proposed model and are designed to be a better physical approximation than classical photometric stereo assumptions.

Using the above setup we performed a number of different experiments. In all the experiments on Eve, the central pixel was initialized with the ground truth depth value and the method described in section 7.2 was used together with the backward upwind scheme of section A.1.

8.1.2. EPPS₃ without nuisance factors. Figure 6 shows the three images used to solve the EPPS₃ problem in which we consider a uniform albedo and light attenuation with $\mu = 1$. To the right of the input images the angle error map between the ground truth normals and those reconstructed using our proposed model is displayed with a heat map scaling (blue is low and red is high) between 0 and 30 degrees. Areas near to fast changing surface gradients can be seen to be more prone to error in reconstruction of surface normals. However, the effect of this on the final result is quantitatively small because the normals are not directly integrated. The final reconstruction can be seen on the bottom left of the figure. The same model textured with the distance error map indicates which areas have suffered a loss in reconstruction fidelity. The error map measures the Euclidean distance in millimeters between the ground truth vertex locations in R^3 and the reconstructed vertices where the heat map scale runs between 0mm and 6mm. This experiment involves no nuisance factors and demonstrates the simplest reconstruction scenario with three images. This is the only experiment performed with three

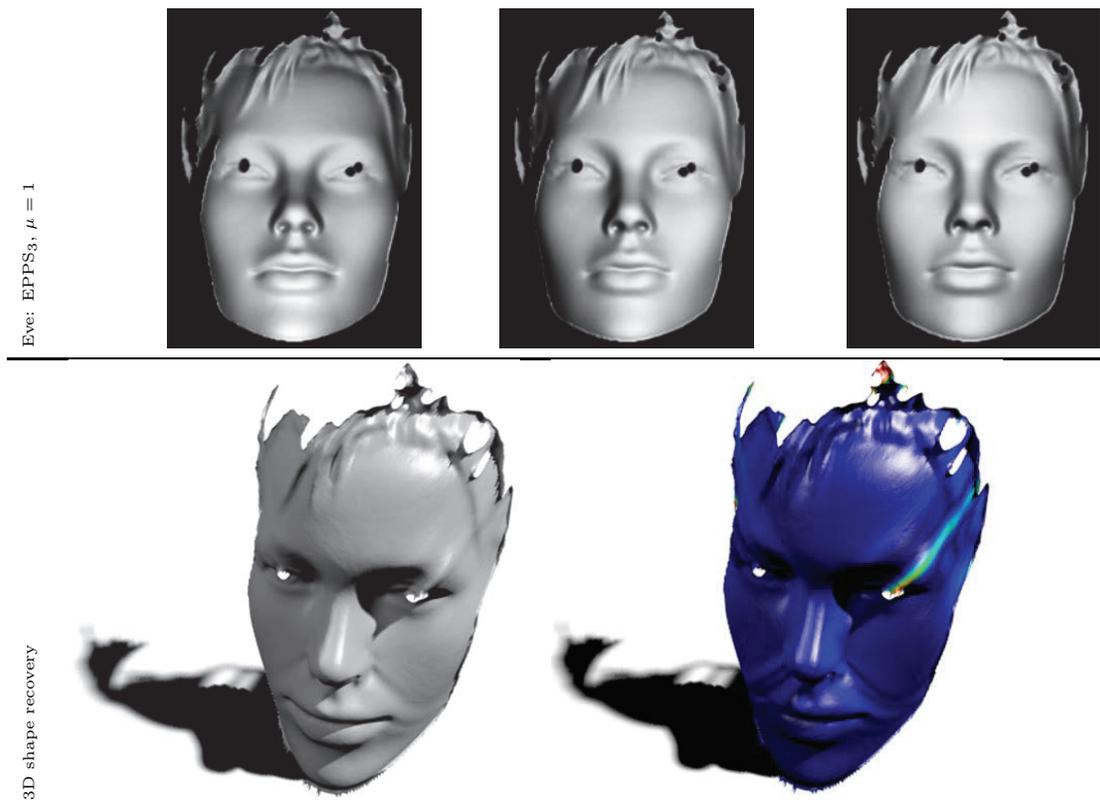


Figure 6. Solving the $EPPS_3$ problem. The first row shows the synthetic input data I_1 , I_2 , and I_3 which are synthesized radiance images of Eve without noise and light attenuation fixed by $\mu = 1$ and three lights placed at 150mm from the focal point. In the second row, the recovered surface is shown on the left and the same surface is shown on the right but textured by the Euclidean error at each point with the heat map scaled from blue to red representing 0mm to 6mm. $MSE = 0.46mm^2$.

images. All other experiments used four images.

Although our method does not explicitly require the surface normals, we can compute them once our scheme has converged. We show the ground truth normals compared to the reconstructed normals in Figure 7. The most likely places for errors to occur are at points with high variance or surface curvature. It can be seen in the figure that errors partially propagate from the point at which they occur in the direction of the processing wavefront but do not necessarily become worse further from the point of distortion. Figure 8 shows a similar result.

8.1.3. $EPPS_4$ with noise. A more informative experiment was performed that demonstrates the effects of injecting artificial noise into the system and solving the $EPPS_4$ problem. Using four images improves the likelihood of having at least two images which observe any given point in the scene. Unlike many other photometric stereo experiments in the literature we do not inject noise into the normals because this is not realistic and we neither directly extract nor use the normals. Instead the noise is directly added to the image itself, as described in the setup, which is a far more realistic approach to synthesis. Most digitally captured

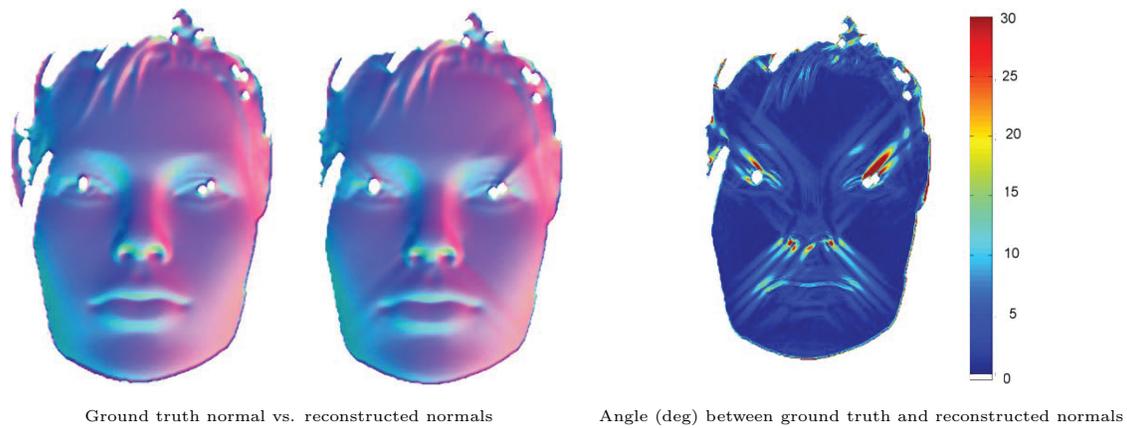


Figure 7. This figure illustrates the ground truth normals on the left, the reconstructed normals in the middle as produced by our method after the surface has been recovered, and the angular difference between the two with an error map scaling of between 0° and 30° . The standard method of converting normal coordinates to RGB is used to display the normals as color images.



Figure 8. The shape reconstructions for the $EPPS_4$ problem for Eve is displayed here showing 0%, 2%, and 5% from left to right. The error maps are textured on to the reconstructed surfaces to help illustrate the magnitude and locations of the errors. The MSE for each instance is 0.51mm^2 , 2.20mm^2 , and 10.1mm^2 from left to right.

images include some level of sensor noise during capture, especially in low light conditions, which can often be assumed to be approximately Gaussian. Figure 8 shows the influence of the noise by texture mapping the Euclidean error maps scaled between 0mm and 6mm onto the reconstructions of image sets with 0%, 2%, and 5% additive Gaussian noise. The reconstruction MSE is 0.50mm^2 , 2.19mm^2 , and 9.98mm^2 , respectively. It is clear that the more noise in the system, the more susceptible the method becomes to propagating errors, as can be seen from the error lines streaking outwards from the central seed point. However, despite up to 5% noise the significant features and characteristics of Eve's face remain intact.

8.1.4. Missing data and nonuniform albedo. As described in section 7 our proposed method handles shadows and occlusions by switching off the information from the missing

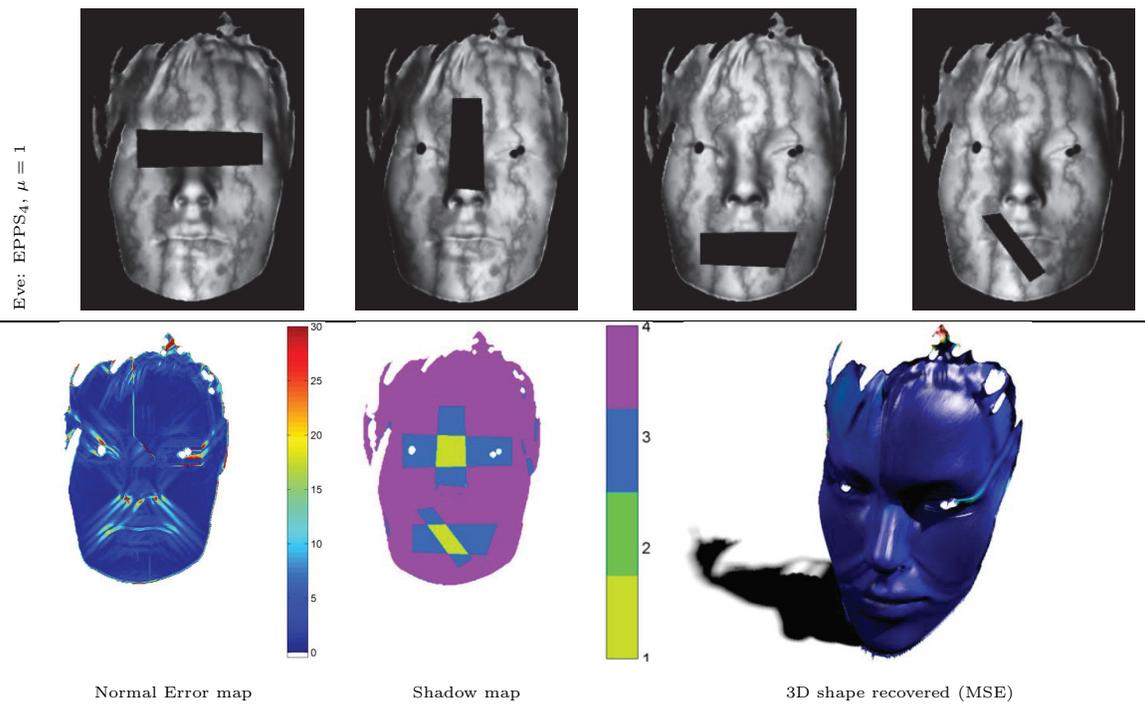


Figure 9. In the first row are the images for the EPPS₄ problem for Eve considering light attenuation with $\mu = 1$, no noise, and missing parts artificially added. In the second row, the left image shows the normal error map in the same way as Figure 7. The middle image illustrates the number of lights visible at each pixel, and the right image is a render of the resulting 3D reconstruction with the Euclidean error map textured on. $MSE = 0.52\text{mm}^2$.

pixels in the affected images. We demonstrate the results of artificially deleting portions of the four input images that have a totally nonuniform albedo. The images can be seen in Figure 9 on the top row. Solving the EPPS₄ problem in this scenario results in an MSE of 0.53mm^2 which is comparable to the previous noiseless result of 0.50mm^2 without the additional nuisance factors. The impact of the missing data is to change the path that the numerical scheme follows. The sequence of the set of independent pixels that are processed at each stage is a moving processing wavefront, and this can be seen in Figure 10. This illustrates the mechanism by which the EPPS₄ problem is solved. In areas where there are three or more illuminated pixels there is no impedance to the propagation of the active wavefront. However, for pixels with only two illumination sources the characteristic has to allow data to come from the right direction, which has the effect of redirecting parts of the advancing processing wavefront.

8.1.5. Full set of nuisance factors. In Figure 11 we solve the EPPS₄ problem showing the results of a scenario in which all nuisance factors are present. We include a nonuniform albedo, overlapped missing data, and 2% noise. We exaggerate the directional light attenuation by setting $\mu = 100$. The 3D reconstruction with the error map textured on shows that attenuation of light reduces the quality of the recovered shape due to the decreasing amount of information toward the edges. This is also visible in the error map for the normals, with normals toward

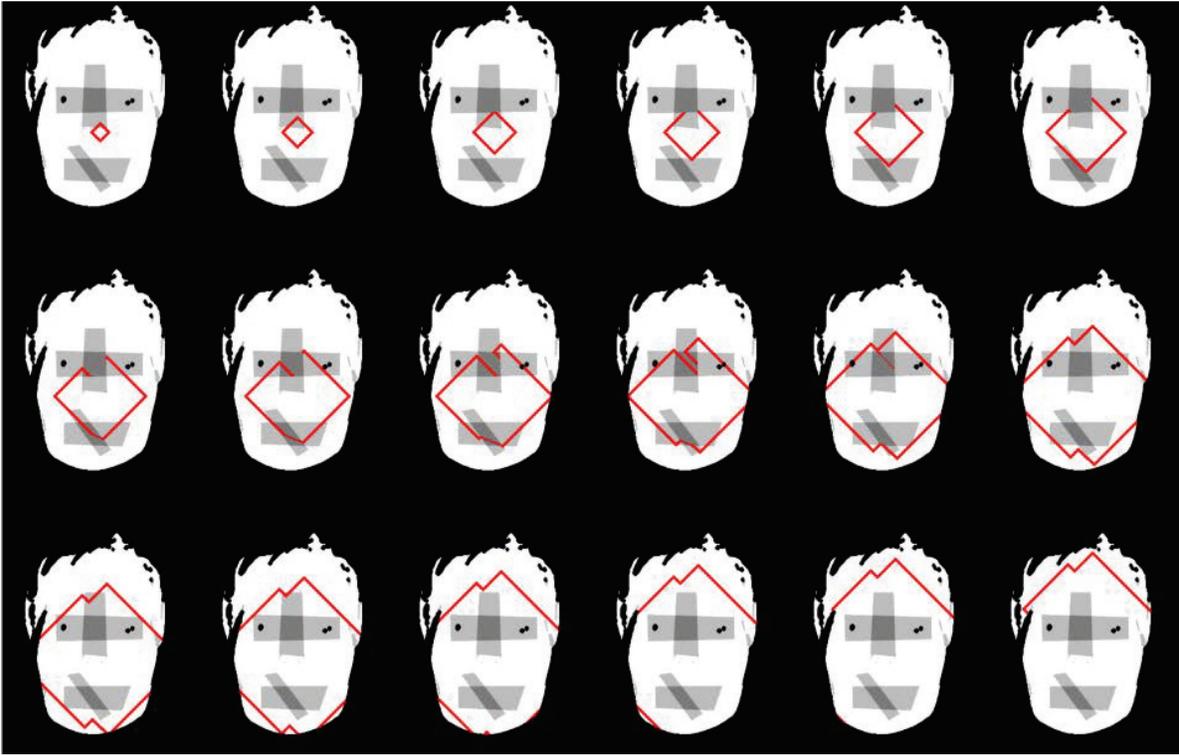


Figure 10. At each stage in our method a number of pixels sit in the processing queue. These pixels shown in red are updated independently to every other pixel in the queue, and this image displays a sequence of the set of pixels processed at different stages during one iteration of our upwind scheme. The sequence appears as a wavefront and moves along the characteristics of the PDE in (7.7). The image used is the shadow map from Figure 9 rendered in grayscale to help highlight the wavefront propagation.

the edges of the image being more than 30° different from the ground truth. Figure 12 shows the albedo recovered under this exaggerated difficult setup and the error map of the difference with the ground truth albedo. Although the resulting surface is somewhat deformed, both the albedo and shape retain much of the structure of the ground truth shape.

8.2. Synthetic cases: AbsPeaks.

8.2.1. Setup. A second synthetic surface is used for testing the proposed method when realistic synthetic shadows appear in the input images as well as performing timing tests for different sized images. The main purpose is comparison to other methods. We generate the surface by taking the absolute value of the well-known MATLAB function *peaks*, scaling it by 0.1 and offsetting it by 5 units along the z -axis of a coordinate system centered at the focal point of a virtual camera. The virtual camera is modeled in the same way as in section 8.1 but with intrinsic parameters $f_x = f_y = N, c_x = N/2, c_y = N/2$. Here, N is the width of the image in pixels and each synthesized image is square in size. This simulates a camera with a 90° field of view in both directions. Again, a synthetic albedo is generated and light sources are defined and placed as in section 8.1. Images are generated of size 256×256 pixels. Here, the radial distance of each light source from the focal point is between 3 and 20 units, depending on

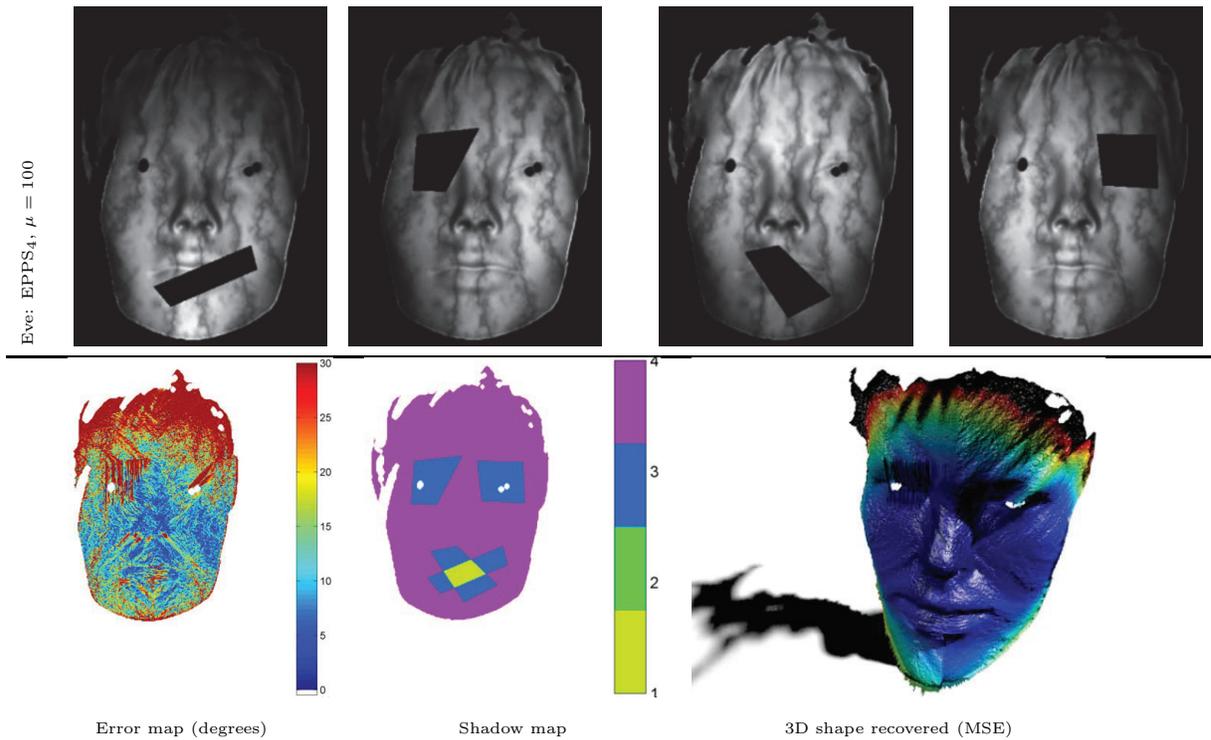


Figure 11. Visualizing the effect of the full set of nuisance factors. In the first row are the images for the $EFPS_4$ problem for Eve considering light attenuation with $\mu = 100$, 2% noise, and artificially deleted data. In the second row, the left image shows the normal error map in the same way as Figure 7. The middle image illustrates the number of lights visible at each pixel, and the right image is a render of the resulting 3D reconstruction with the Euclidean error map textured on. $MSE = 288\text{mm}^2$.

the experiment being performed. By setting the light sources at a distance of 20 units, which is large relative to the size and position of AbsPeaks, we are actually approximating parallel ray distant light sources that are commonly assumed in photometric stereo methods. At the other end, by positioning the light sources close to the focal point at a radius of 3 units, we can achieve an accurate simulation of non-parallel-ray realistic light sources and positioning. Irradiance images will then be nonlinearly dependent on light source location and orientation as well as surface depth and location.

8.2.2. Realistic shadows. One issue, when simulating light sources very close to a virtual object, occurs when the normal of the surface at any point observed by the virtual camera is at an angle of more than 90° from the angle of incidence of a light ray. In this case the irradiance value from (7.7) will be negative. Double precision floating point is initially used for the synthesized images so that negative irradiance, although physically impossible, can be stored. The correct course of action is to set these negative values to be missing data and convert the images to 8-bit grayscale images, as described in section 8.1. This is the case in Figure 13, which shows the set of significantly overlapped shadows and light attenuation where $\mu = 5$ for light sources placed at a radial distance of 10 units. The MSE for the reconstruction is $3.72\text{E-}2$ units². Once again this demonstrates that our formulation enables

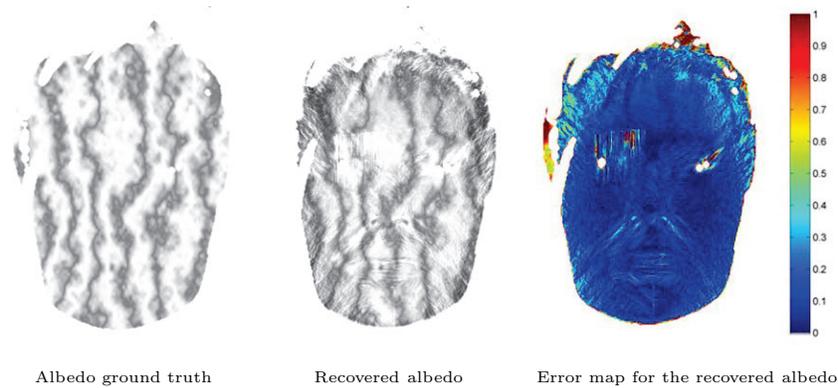
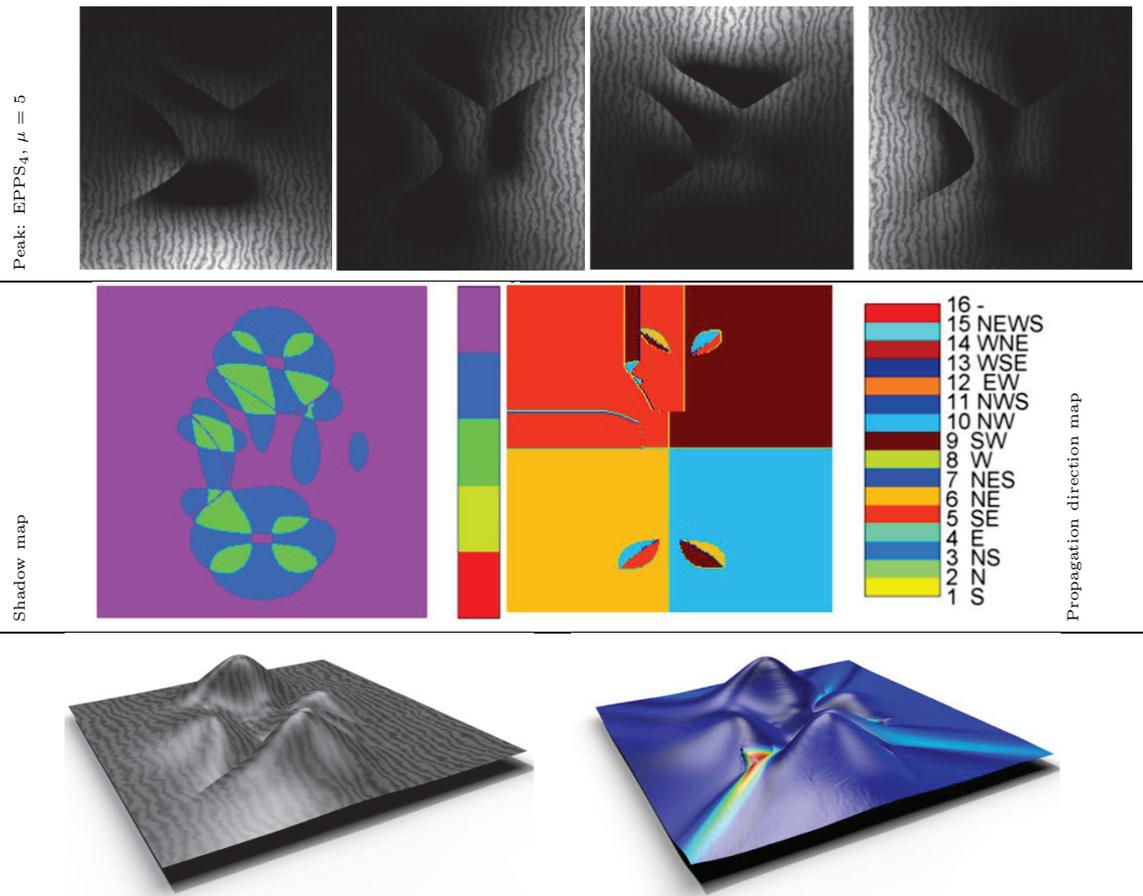


Figure 12. The left and middle images show the ground truth synthetic albedo and the reconstructed albedo, respectively, from the experiment in section 8.1.5. The image on the right is an error map of the difference between the ground truth and reconstructed albedos.

endoscopic photometric stereo even for shadowed regions and non-parallel-ray light sources. We do not explicitly simulate the case of cast shadows and instead consider them to be areas of missing data from in one or more images.

8.2.3. Comparison to other methods. Let us compare our shape recovery approach to other methods which have not necessarily been designed to handle shadows. We therefore allow all irradiance values and retain the data in the double precision floating point format without first converting to 8-bit grayscale. As a result there are no areas with missing data in our synthesized images even though the images themselves may not be physically realizable. The first row of Figure 14 shows the synthesized images from the AbsPeaks setup with $\mu = -2$, which completely turns off attenuation effects for both distance and nonuniform lighting. The point light sources are positioned 3 units from the focal point, which, as we have mentioned, is in contrast to the classic assumption of distant light sources. The reasoning behind this placement is to demonstrate the deformation of the shape recovered by other methods when the model of nearby light sources is not taken into account. We compare our approach with the Poisson solver of [1, 8], which are surface integration methods and whose authors make their code available. The surface normals provided to these two methods are first computed by assuming distant parallel ray light sources where the light direction is computed by averaging between the light source position and an approximate distance from the object (as done in [31]). [1] requires the use of full boundary conditions, whereas [8] assumes so-called natural boundary conditions which require no additional information. Furthermore, we remark that [1] actually compares several methods of 3D shape recovery from the normal field, but all of them are prevented from providing the correct shape in our scenario since the normal field is *deformed* due to neglected physical effects. In the bottom row of Figure 14 from left to right one can see the ground truth surface, our reconstruction, the reconstruction using [1], and the reconstruction using [8]. Even in this simple scenario, the other methods produce deformed reconstructions, whereas our method faithfully reconstructs the surface with an MSE of $3.27\text{E-}2$ units². Figure 15 illustrates that the reconstruction by the other two methods becomes considerably worse when we add attenuation of $\mu = 1$, whereas our result remains



Approximate 3D shape with recovered albedo textured on 3D shape recovered with error map textured on

Figure 13. In the first row are the images for the $EPPS_4$ problem for *AbsPeaks* considering light attenuation with $\mu = 5$, no noise, and missing parts artificially added. In the second row, the left image shows the number of lights visible at each pixel. The right image shows the direction from which the information flowed for every pixel. So, for example, any pixel labeled SW was updated from the pixel southwest of it. If a pixel is colored NES, then it was updated by the average of the two pixels to its northeast and southeast. The shadowed image results in a change of direction of information flow and processing wavefront propagation, and hence the update locations are not uniform. In the third row, the left image is a render of the resulting 3D reconstruction with its reconstructed albedo textured on to it. The right image is the same shape but with its error map textured on. $MSE = 3.75E-4 \text{ units}^2$.

consistent with an MSE of $3.80E-2 \text{ units}^2$. The shape recovered has a radial deformation which increases moving far away from the projection of the optical center. This is because of the effect of realistic light propagation that gets stronger when moving far from the center.

Finally, we want to compare our method with a modern approach aimed at getting shape, illumination, and reflectance from shading (SIRFS) based on the machine learning procedure [2]. By a default setting, it uses only one input image and the model is required to be appropriately trained in advance. Furthermore, SIRFS has many parameters which have to be tuned and that drastically change the final shape. Even if there are several technical aspects

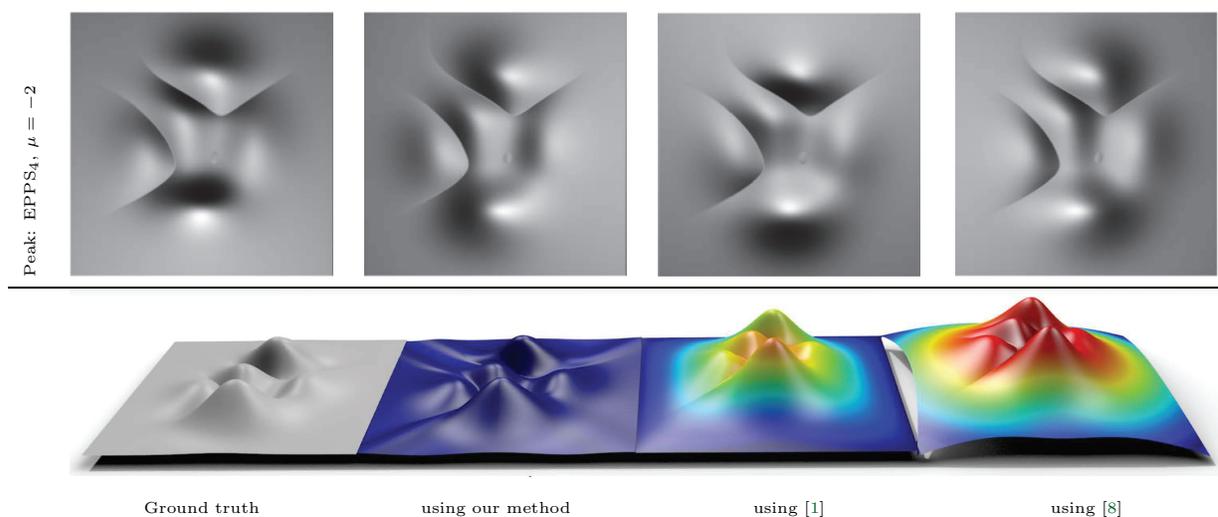


Figure 14. In the first row are images from *AbsPeaks* with $\mu = -2$ and lights positioned 3 units from the focal point. In the second row, from left to right: ground truth shape, reconstruction using our method $MSE = 3.29E-4units^2$, the reconstruction using [1] $MSE = 0.53units^2$, the reconstruction using [8] $MSE = 0.94units^2$. All reconstructed shapes have their error maps textured on.

which differ from our approach, the goal of this comparison is to show that the scenario we intended to face has many critical features that prevent other methods from working. Instead of having as the main goal the 3D shape recovery *in the wild*, our method takes into account a specific setup reliably obtained with the use of a prototype. Figure 16 shows the shape reconstructed by SIRFS considering a surface with a nonuniform albedo and light attenuated. Let us remark that we used SIRFS as it is provided¹ without retraining the model or changing parameters. With no need to compute the error with respect to the ground truth, the actual difference is clear from the left image of Figure 7.

8.2.4. Run times and effect of image size on MSE. Using the *AbsPeaks* setup we can easily vary the size of the images and observe how this affects both the resulting reconstruction error as well as the run time. In Figure 17 the left graph was constructed by running the reconstruction for one iteration for a set of different image sizes. The reconstruction was performed on different numbers of Intel i7 cores to illustrate the almost linear speedup that we obtained through straightforward multicore parallelization. The parallelization is done by using OpenMP to schedule processing of independent pixels on the processing wavefront to separate cores. The graph on the right of Figure 17 shows that the MSE reduces as a function of image size in the case when there is little to no noise. This is expected behavior because the the order of consistency of the method is one. However, as noise is introduced the amount of noise as measured by the MSE indicates the lower bound that is reached after convergence. By injecting noise into the synthetic images, the final reconstruction accuracy is fundamentally limited. For both real and synthetic images the experiments were run until convergence. It is useful to note that the actual convergence relative to the MSE stopped

¹<http://www.cs.berkeley.edu/~barron/>

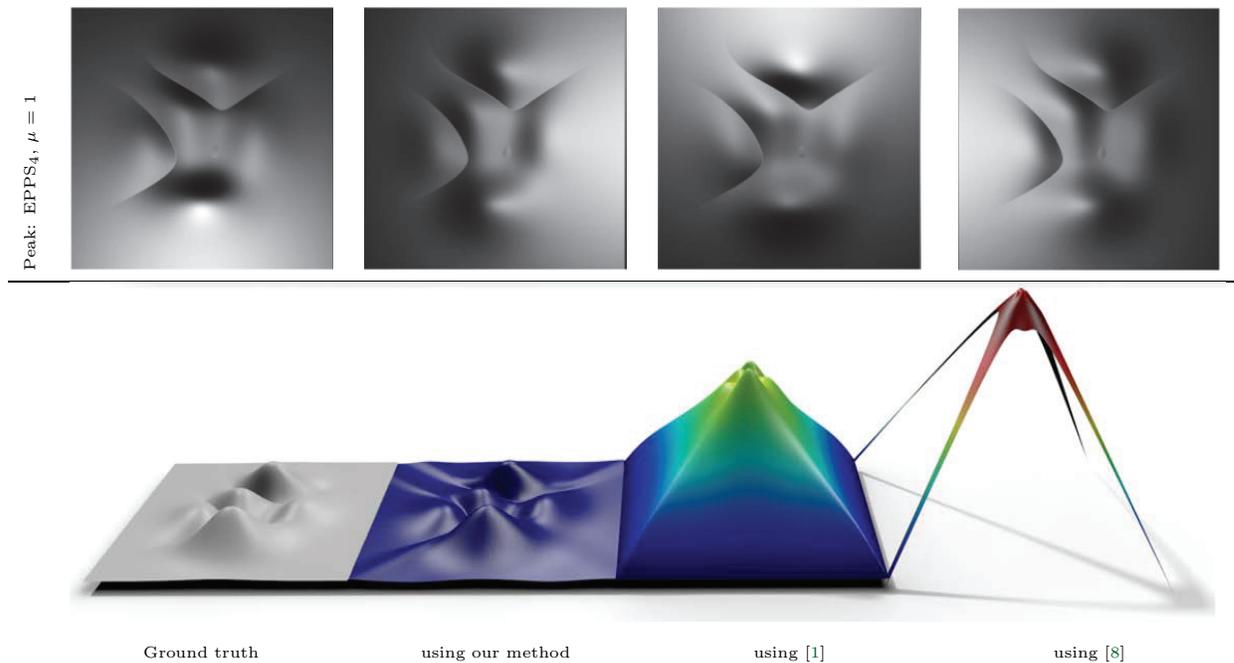


Figure 15. In the first row are images from *AbsPeaks* with $\mu = 1$ and lights positioned 3 units from the focal point. In the second row, from left to right: ground truth shape, reconstruction using our method $MSE = 3.82E-4units^2$, the reconstruction using [1] $MSE = 5.19units^2$, the reconstruction using [8] $MSE = 21.5units^2$. All reconstructed shapes have their error maps textured on.

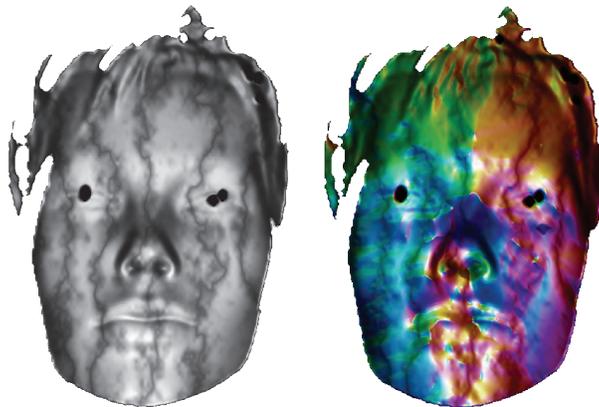


Figure 16. The image on the left is the input data with very low light attenuation and a nonuniform albedo. The one on the right is the normal field computed by SIRFS.

improving after three iterations of our upwind scheme. This means that in a time-sensitive use scenario the algorithm need not run to convergence and instead can be restricted to run only for three iterations.

Finally, we resume the synthetic tests shown in Figures 9 and 11 and Figures 13 and 14, respectively, in Tables 1 and 2 performed using the semi-Lagrangian scheme. Due to the

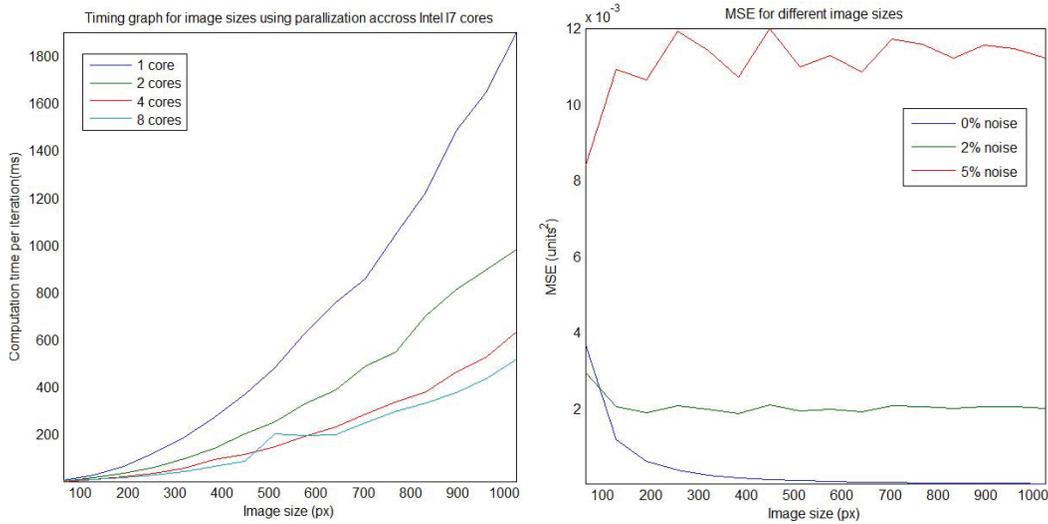


Figure 17. The left graphs show the speedup due to parallel implementation of the algorithm for different image width sizes. The current implementation has overheads in setting up each iteration which have not yet been parallelized, which is the reason that using 8 cores does not currently infer a significant advantage over 4 cores. The right graph shows the MSE convergence of the semi-Lagrangian numerical scheme for the $EPPS_4$ problem with light attenuation $\mu = 1$, no noise, 2% noise, and 5% noise.

Table 1

The values of this table explain how (in precision and in time) the semi-Lagrangian scheme converges for Eve and the gain obtained by parallel implementation of the algorithm.

Eve		$EPPS_4, \mu = 1$		$EPPS_4, \mu = 100$		
Δ	Cores	MSE s-L	time (sec)	MSE s-L	time (sec)	
640×480	4	0.506019	0.586100	0.530944	0.597518	
	8	0.506019	0.528541	0.530944	0.558816	
1280×960	4	0.627532	2.519946	0.687804	2.587541	
	8	0.627532	2.246315	0.687804	2.103126	
2% noise	640×480	4	2.201825	0.533312	4.659086	0.609900
	8	2.201825	0.502544	4.659086	0.451640	
1280×960	4	4.383240	2.576125	7.981745	2.603389	
	8	4.383240	1.855340	7.981745	2.113353	
5% noise	640×480	4	10.090009	0.530891	23.250882	0.624731
	8	10.090009	0.429353	23.250882	0.428937	
1280×960	4	17.338734	2.575009	35.542887	2.589999	
	8	17.338734	1.891218	35.542887	1.827679	

presence of occlusion regions, the parallelization does not provide a linear decreasing error with respect to the available cores. However, this implementation is a preliminary step for future computational advancement on a more suitable platform as GPUs. Regarding the MSE errors, of course, as long as the noise increases, it increases as well. However, we note that MSE increases also when the size of the input images increase since the error is mostly due to propagation of information accumulated during the characteristic strip expansion.

Table 2

The values of this table explain how (in precision and in time) the semi-Lagrangian scheme converges for Peak and the gain obtained by parallel implementation of the algorithm.

Peak		EPPS ₄ , $\mu = -2$		EPPS ₄ , $\mu = 5$	
Δ	Cores	MSE s-L	time (sec)	MSE s-L	time (sec)
512×512	4	0.000115	1.211298	0.000202	1.639071
	8	0.000115	0.910084	0.000202	1.149819
1024×1024	4	0.000033	4.377678	0.000060	6.576134
	8	0.000033	4.265500	0.000060	5.272947
512×512	4	0.001934	1.095136	0.551087	2.789768
	8	0.001934	0.749802	0.551087	2.008740
1024×1024	4	0.002010	4.330678	0.544797	10.715179
	8	0.002010	4.388909	0.544797	7.065750
512×512	4	0.010990	1.061619	1.123475	2.656035
	8	0.010990	0.748017	1.123475	1.765348
1024×1024	4	0.011230	4.259067	1.109509	10.750346
	8	0.011230	4.527862	1.109509	8.924932

8.3. Real cases. In the last scenario we use a low cost endoscopic camera synchronized to four individually controlled low power white SMD light sources. The camera has an approximately 90° field-of-view lens. The intrinsic camera parameters are $f_x = f_y = f = 635$, $c_x = 325$, and $c_y = 244$. The lights are attached to a 3D printed surface that is designed to allow the camera to sit in the same plane as the lights. They are positioned using callipers, and the assumption that placing them face flat side down would fix their beam orientation to be 90° to the surface proved to be sufficiently robust. We capture three objects which are approximately Lambertian, namely a set of small artificial stairs, the tip of a thumb, and the palm of a hand. The radial positioning of the light sources was adjusted for the size of the objects being observed, where we used 40mm, 28mm, and 12mm for the hand, stairs, and thumb, respectively. Our experimental setup permits us to measure the distance to a point on the surface to use as an initial seed point as required by the integration strategy. In a more practical setup we could use a single static laser point or laser line attached to the endoscope and calibrated to automatically provide a set of true depth coordinates for use as seed points. The point on each object observed by the central pixel in the camera frames was measured by hand in order to define the only required boundary condition for the method. These distances were 220mm, 120mm, and 35mm for the hand, stairs, and thumb, respectively. The SMD light sources have a real radial attenuation parameter of approximately $\mu = 1$. The endoscopic camera and SMD lights can be seen in Figure 18. An TI Launchpad was used to control the lights via serial communication in MATLAB.

The images were taken in a dark environment, and a calibration image was acquired without any lights activated. This was then subtracted from all subsequent images to account for ambient lighting. After this any pixels with a value less than 20 were marked as shadow. The intrinsic camera parameters for our low cost endoscopic camera were found using the Bouguet calibration toolbox,² and all images had the effect of lens distortion removed using the same toolbox. The processed images for each object can be seen in Figure 19 with a

²http://www.vision.caltech.edu/bouguetj/calib_doc/

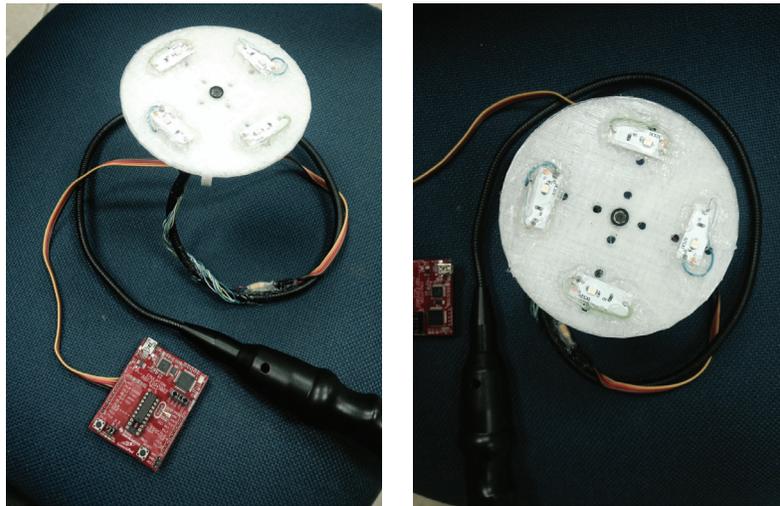


Figure 18. These two images show photographs of the endoscopic camera with the SMD LEDs mounted 90° apart on a 3D printed surface at a radial distance of approximately 40mm from the camera focal point. The TI launchpad is the red board which is visible. The simple endoscopic webcam has a field of view of around 90° .

rendering of a novel view of the same object based on our reconstruction. It is interesting to note that our method was able to successfully reconstruct the thumb and the ridges of the thumb-prints are clearly visible in the 3D render. Furthermore, despite the specularly and nonuniform albedo present in the step images, the reconstruction is unaffected and preserves the straight lines of the edges of the steps. The hand model is also interesting because it illustrates that our method can handle real world scenarios in which there is a level of noise in the captured images.

Using the results from the proposed method we can also resynthesize the original captured images to show how they would look under the conditions determined by the reconstruction process. This could in principle be used in an additional optimization process as a data term to drive the reconstruction to optimally match the initial images, but this has not been implemented here and remains an open avenue for future research. The process of resynthesis is also interesting because it allows us to reconstruct the original images in areas where data is missing, as can be seen in the first row of Figure 20. We measure reconstruction accuracy using $PSNR = 10\log_{10}(255^2/E[(I_{est} - I)^2])$. This measures how similar an image is to a disturbed version of the same image and is most commonly used to show the effects of image denoising algorithms. We do not measure the discrepancy in areas where data cannot be reconstructed. The resynthesized images shown are of a high quality according to their high $PSNR$ values.

9. Conclusions. An efficient model for shape reconstruction from the $EPPS_n$ problem was proposed. The main goal of the proposed model is to define a new differential formulation based on a quasi-linear PDE, where the well-posedness holds even in the presence of images with missing parts. We have shown that PDEs provide a strong way of modeling near-field photometric stereo that can be used to approximate completely general lighting scenarios.

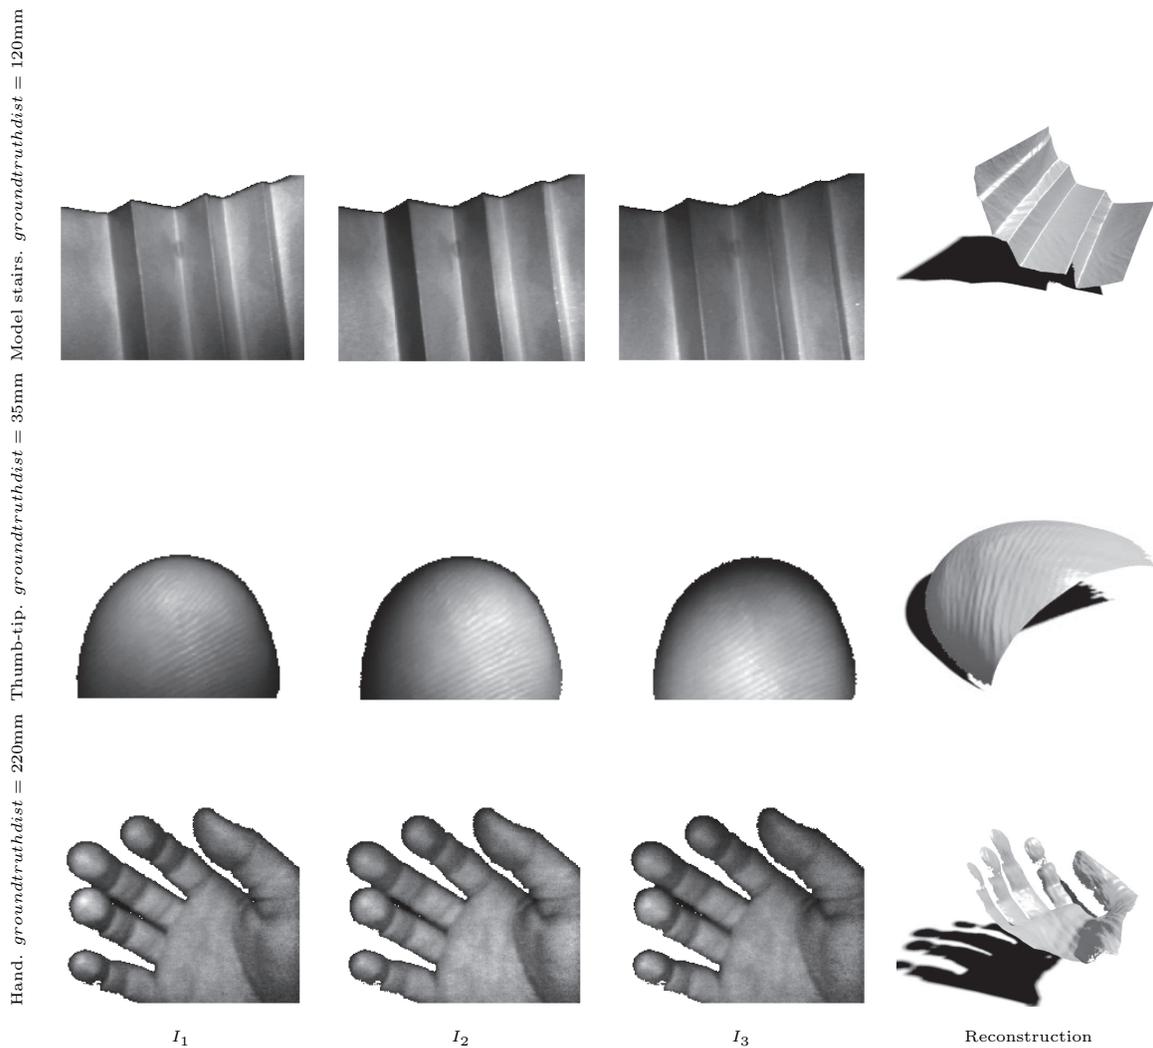


Figure 19. Novel views from reconstructions of real objects imaged with the experimental setup described in section 8.3. The captured images are shown on the left of each reconstruction. They all have the same size 640×480 and provide the relative 3D shape shown on the respective right column in less than half a second.

The model we presented overcomes the limitations of more classical approaches since the lighting is modeled realistically as a nearby source with fully general illumination. As far as we can tell this currently represents state of the art in physical modeling for any photometric stereo method. We are currently investigating even more general models for nonlinearities such as specular effects and general BRDF functions. Our current experiments demonstrate that the model gracefully deals with nonlinear light attenuations as well as nonuniform surface albedo and missing data. Furthermore, the suggested method handles real world surfaces and produces quantitatively faithful surface reconstructions for nearby objects. The method is highly parallelizable, and future work will attempt an implementation on a kilo-core GPU to demonstrate that real time endoscopic shape from photometric stereo is possible. Current

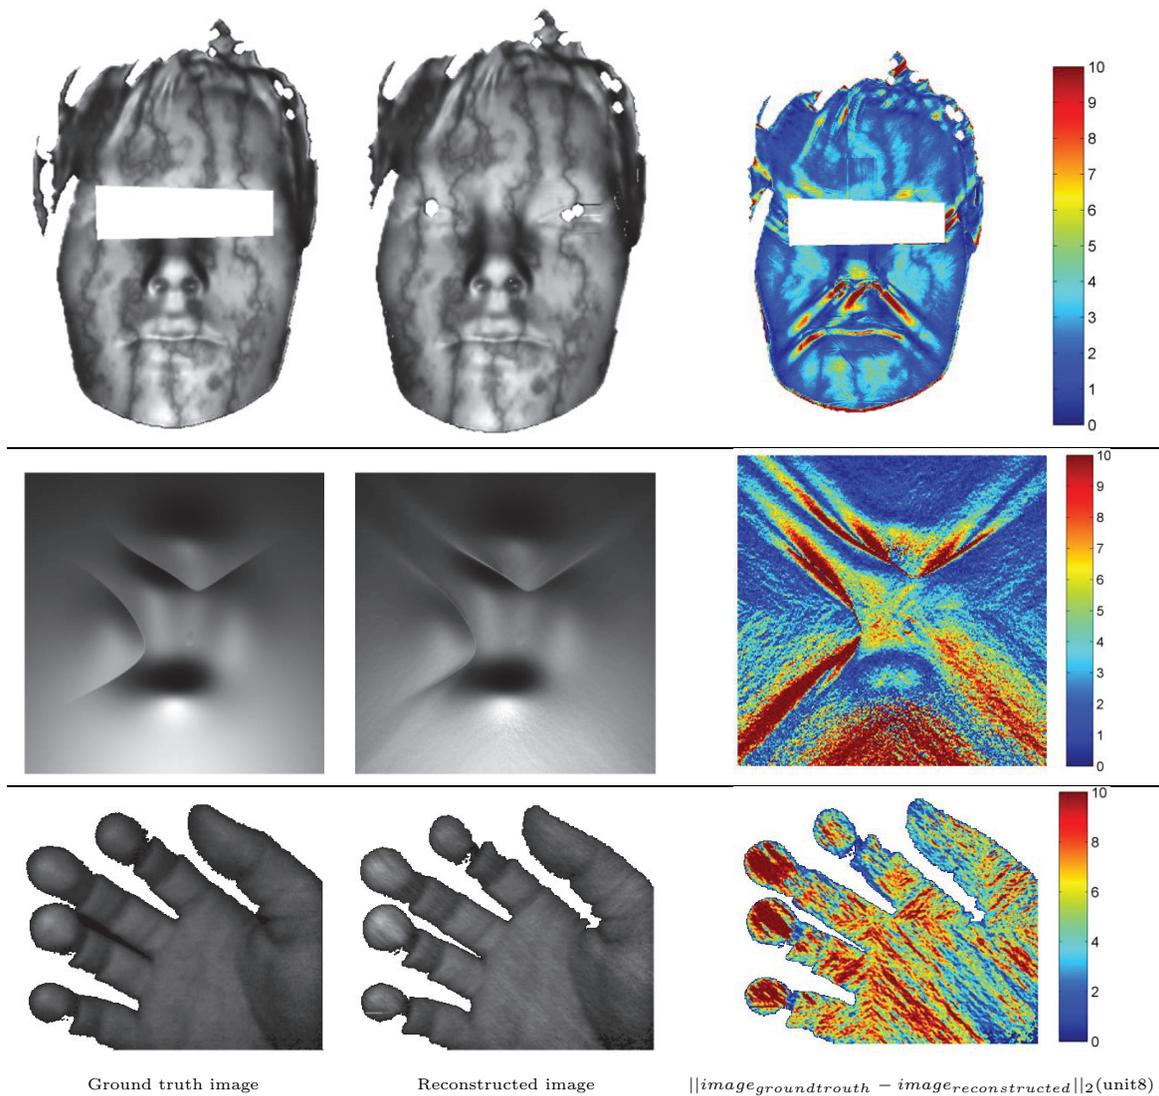


Figure 20. This figure compares the ground truth images to their resynthesized counterparts. In each row we display the first captured or synthesized image on the left. The middle image is the resynthesized image based on the proposed method, and the error map on the right shows the difference in intensity between the ground truth image and the resynthesized image. In the first row are Eve experiments with no noise, a nonuniform albedo, $\mu = 1$, lights positioned at 150mm, missing data. PSNR = 34.3dB. In the second row are AbsPeaks experiments with 2% noise, a uniform albedo, $\mu = 1$, lights positioned at a radius of 3 units. PSNR = 32.0dB. In the third row is a real hand experiment as described in section 8.3. PSNR = 30.8dB.

ongoing work includes shrinking the scanning head and performing tracking so that reconstructed depth maps can be fused into a larger surface. Another issue is that of obtaining at least one known ground truth point (assuming a single connected surface). This has not been implemented, but we are investigating using a laser dot or line calibrated to the camera.

Appendix. Backward numerical schemes. The backward numerical schemes are based on the approximation of the surface propagating the information stored on the outflowing parts of the boundary

$$(A.1) \quad \Gamma_{out} = \partial\Omega \setminus \Gamma_{in}.$$

The formulation of these schemes can be obtained by considering the equivalent problem

$$(A.2) \quad \begin{cases} -\mathbf{b}(x, y, z) \cdot \nabla z(x, y) = -s(x, y, z) & \text{a.e. } (x, y) \in \Omega_p, \\ z(x, y) = g(x, y) & \forall (x, y) \in \Gamma_{out} \end{cases}$$

and repeating the arguments used for the forward cases.

A.1. Backward upwind scheme. Using the same construction of the forward upwind scheme, we introduce an artificial diffusion that guarantees the ability to follow the orientation of the characteristics, this time in the opposite direction. The implicit numerical scheme

$$(A.3) \quad -b_{i,j}^1 \frac{Z_{i+1,j} - Z_{i-1,j}}{2\Delta} - b_{i,j}^2 \frac{Z_{i,j+1} - Z_{i,j-1}}{2\Delta} \\ = |b_{i,j}^1| \frac{Z_{i+1,j} - 2Z_{i,j} + Z_{i-1,j}}{2\Delta} + |b_{i,j}^2| \frac{Z_{i,j+1} - 2Z_{i,j} + Z_{i,j-1}}{2\Delta} - s_{i,j}$$

for $i, j = 1, \dots, n-1$ can be rewritten as

$$(A.4) \quad (UW_B) \quad Z_{i,j}^{(k+1)} = \frac{|b_{i,j}^1(Z_{i,j}^{(k)})| Z_{i+\text{sgn}(b_{i,j}^1(Z_{i,j}^{(k)}),j)}^{(k)} + |b_{i,j}^2(Z_{i,j}^{(k)})| Z_{i,j+\text{sgn}(b_{i,j}^2(Z_{i,j}^{(k)}))}^{(k)} - \Delta s_{i,j}(Z_{i,j}^{(k)})}{|b_{i,j}^1(Z_{i,j}^{(k)})| + |b_{i,j}^2(Z_{i,j}^{(k)})|}.$$

A.2. Backward semi-Lagrangian scheme. Once we have (A.2), we obtain the backward scheme as in the forward one, that is,

$$(A.5) \quad (sLB) \quad Z_{i,j}^{(k+1)} = Z^{(k)}(x_i + h\Lambda_{i,j}^1(Z_{i,j}^{(k)}), y_j + h\Lambda_{i,j}^2(Z_{i,j}^{(k)})) - \frac{s_{i,j}(Z_{i,j}^{(k)})}{|\mathbf{b}_{i,j}(Z_{i,j}^{(k)})|} h.$$

REFERENCES

- [1] A. AGRAWAL, R. RASKAR, AND R. CHELLAPPA, *What is the range of surface reconstructions from a gradient field?*, in Proceedings of the 9th European Conference on Computer Vision: Part I, Springer-Verlag, Berlin, Heidelberg, 2006, pp. 578–591.
- [2] J. T. BARRON AND J. MALIK, *Shape, Illumination, and Reflectance from Shading*, Tech. Report UCB/EECS-2013-117, EECS, UC Berkeley, Berkeley, CA, 2013.
- [3] M. BREUß, E. CRISTIANI, J. D. DUROU, M. FALCONE, AND O. VOGEL, *Perspective shape from shading: Ambiguity analysis and numerical approximations*, SIAM J. Imaging Sci., 5 (2012), pp. 311–342.
- [4] A. M. BRUCKSTEIN, *On shape from shading*, Comput. Vis. Graph. Image Process., 44 (1988), pp. 139–154.
- [5] M. CHANDRAKER, J. BAI, AND R. RAMAMOORTHY, *On differential photometric reconstruction for unknown, isotropic BRDFs*, IEEE Trans. Pattern Anal. Mach. Intell., 35 (2012), pp. 2941–2955.
- [6] M. K. CHANDRAKER, S. AGARWAL, AND D. J. KRIEGMAN, *ShadowCuts: Photometric stereo with shadows*, in Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2007.

- [7] T. COLLINS AND A. BARTOLI, *3d reconstruction in laparoscopy with close-range photometric stereo*, in Medical Image Computing and Computer-Assisted Intervention – MICCAI 2012, Lecture Notes in Comput. Sci. 7511, N. Ayache, H. Delingette, P. Golland, and K. Mori, eds., Springer, Berlin, Heidelberg, 2012, pp. 634–642.
- [8] J.-D. DUROU, J.-F. AUJOL, AND F. COURTEILLE, *Integrating the normal field of a surface in the presence of discontinuities*, in Proceedings of the 7th International Conference on Energy Minimization Methods in Computer Vision and Pattern Recognition, 2009, pp. 261–273.
- [9] M. FALCONE AND R. FERRETTI, *Semi-Lagrangian Approximation Schemes for Linear and Hamilton–Jacobi Equations*, SIAM, Philadelphia, 2014.
- [10] R.T. FRANKOT AND R. CHELLAPPA, *A method for enforcing integrability in shape from shading algorithms*, IEEE Trans. Pattern Anal. Mach. Intell., 10 (1988), pp. 439–451.
- [11] K. GOEBEL AND W. A. KIRK, *A fixed point theorem for asymptotically nonexpansive mappings*, Proc. Amer. Math. Soc., 35 (1972), pp. 171–174.
- [12] B. W. HAPKE, *A theoretical photometric function for the lunar surface*, J. Geophys. Res., 68 (1963), pp. 4571–4586.
- [13] J. K. HASEGAWA AND C. L. TOZZI, *Shape from shading with perspective projection and camera calibration*, Computers and Graphics, 20 (1996), pp. 351–364.
- [14] C. HERNÁNDEZ, G. VOGIATZIS, AND R. CIPOLLA, *Shadows in three-source photometric stereo*, in Proceedings of the 10th European Conference on Computer Vision: Part I, Springer-Verlag, Berlin, Heidelberg, 2008, pp. 290–303.
- [15] B. K. P. HORN, *Obtaining shape from shading information*, in The Psychology of Computer Vision, P. H. Winston, ed., McGraw–Hill, New York, 1975, pp. 115–155.
- [16] B. K. P. HORN, *Height and gradient from shading*, Int. J. Comput. Vis., 5 (1990), pp. 37–75.
- [17] B. K. P. HORN AND M. J. BROOKS, *The variational approach to shape from shading*, Comput. Vis. Graph. Image Process., 33 (1986), pp. 174–208.
- [18] B. K. P. HORN, R. S. SZELISKI, AND A. L. YUILLE, *Impossible shaded images*, IEEE Trans. Pattern Anal. Mach. Intell., 15 (1993), pp. 166–169.
- [19] I. HOROVITZ AND N. KIRYATI, *Depth from gradient fields and control points: Bias correction in photometric stereo*, Image Vis. Comput., 22 (2004), pp. 681–694.
- [20] R. KIMMEL AND A. M. BRUCKSTEIN, *Tracking level sets by level sets: A method for solving the shape from shading problem*, Comput. Vis. Image Understand., 62 (1995), pp. 47–58.
- [21] R. KIMMEL AND J. A. SETHIAN, *Optimal algorithm for shape from shading and path planning*, J. Math. Imaging Vision, 14 (2001), pp. 237–244.
- [22] R. KIMMEL AND I. YAVNEH, *An algebraic multigrid approach for image analysis*, SIAM J. Sci. Comput., 24 (2003), pp. 1218–1231.
- [23] S. LEE AND M. BRADY, *Integrating stereo and photometric stereo to monitor the development of glaucoma*, Image Vis. Comput., 9 (1991), pp. 39–44.
- [24] R. MECCA AND M. FALCONE, *Uniqueness and approximation of a photometric shape-from-shading model*, SIAM J. Imaging Sci., 6 (2013), pp. 616–659.
- [25] R. MECCA, G. ROSMAN, R. KIMMEL, AND A. M. BRUCKSTEIN, *Perspective photometric stereo with shadows*, in Proceedings of the Fourth International Conference on Scale Space and Variational Methods in Computer Vision, 2013, pp. 258–269.
- [26] R. MECCA, A. TANKUS, A. WETZLER, AND A. M. BRUCKSTEIN, *A direct differential approach to photometric stereo with perspective viewing*, SIAM J. Imaging Sci., 7 (2014), pp. 579–612.
- [27] R. MECCA, A. WETZLER, R. KIMMEL, AND A.M. BRUCKSTEIN, *Direct shape recovery from photometric stereo with shadows*, in Proceedings of the International Conference on 3D Vision, 2013, pp. 382–389.
- [28] T. OKATANI AND K. DEGUCHI, *Shape reconstruction from an endoscope image by shape from shading technique for a point light source at the projection center*, Comput. Vis. Image Understand., 66 (1997), pp. 119–131.
- [29] R. ONN AND A. M. BRUCKSTEIN, *Integrability disambiguates surface recovery in two-image photometric stereo*, Int. J. Comput. Vis., 5 (1990), pp. 105–113.
- [30] T. PAPADHIMITRI AND P. FAVARO, *A new perspective on uncalibrated photometric stereo*, in Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition, 2013.
- [31] V. PAROT, D. LIM, G. GONZALEZ, G. TRAVERSO, N. S. NISHIOKA, B. J. VAKOC, AND N. J. DURR,

- Photometric stereo endoscopy*, J. Biomed. Opt., 18 (2013), 076017.
- [32] E. PRADOS, F. CAMILLI, AND O. FAUGERAS, *A viscosity solution method for shape-from-shading without image boundary data*, M2AN Math. Model. Numer. Anal., 40 (2006), pp. 393–412.
- [33] E. PRADOS AND O. D. FAUGERAS, *Shape from shading: A well-posed problem?*, in Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition, Vol. 2, 2005, pp. 870–877.
- [34] A. QUARTERONI AND A. VALLI, *Numerical Approximation of Partial Differential Equations*, Springer, Berlin, 1994.
- [35] A. SANIN, C. SANDERSON, AND B. C. LOVELL, *Shadow detection: A survey and comparative evaluation of recent methods*, Pattern Recognition, 45 (2012), pp. 1684–1695.
- [36] T. SIMCHONY, R. CHELLAPPA, AND M. SHAO, *Direct analytical methods for solving Poisson equations in computer vision problems*, IEEE Trans. Pattern Anal. Mach. Intell., 12 (1990), pp. 435–446.
- [37] J. C. STRICKWERDA, *Finite Difference Schemes and Partial Differential Equations*, Wadsworth Brooks/Cole, Belmont, CA, 1989.
- [38] K. SUNKAVALLI, T. ZICKLER, AND H. PFISTER, *Visibility subspaces: Uncalibrated photometric stereo with shadows*, in Proceedings of the 11th European Conference on Computer Vision: Part II, Springer-Verlag, Berlin, Heidelberg, 2010, pp. 251–264.
- [39] A. TANKUS, N. A. SOCHEN, AND Y. YESHURUN, *Shape-from-shading under perspective projection*, Int. J. Comput. Vis., 63 (2005), pp. 21–43.
- [40] G. VOGIATZIS AND C. HERNANDEZ, *Practical 3d reconstruction based on photometric stereo*, in Computer Vision, Stud. Comput. Intell. 285, R. Cipolla, S. Battiato, and G. M. Farinella, eds., Springer, Berlin, Heidelberg, 2010, pp. 313–345.
- [41] L. B. WOLFF AND E. ANGELOPOULOU, *3-d stereo using photometric ratios*, in Proceedings of the 3rd European Conference on Computer Vision, Springer, Berlin, 1994, pp. 247–258.
- [42] C. WU, S. G. NARASIMHAN, AND B. JARAMAZ, *A multi-image shape-from-shading framework for near-lighting perspective endoscopes*, Int. J. Comput. Vis., 86 (2010), pp. 211–228.
- [43] S. YEUNG, H. TSUI, AND A. YIM, *Global shape from shading for an endoscope image*, in Medical Image Computing and Computer-Assisted Intervention – MICCAI 1999, Lecture Notes in Comput. Sci. 1679, C. Taylor and A. Colchester, eds., Springer, Berlin, 1999, pp. 318–327.
- [44] Z. ZHANG, *A flexible new technique for camera calibration*, IEEE Trans. Pattern Anal. Mach. Intell., 22 (2000), pp. 1330–1334.