Direct Shape Recovery from Photometric Stereo with Shadows

Roberto Mecca*, Aaron Wetzler*, Ron Kimmel* and Alfred Marcel Bruckstein*
*Department of Computer Science
Technion - Israel Institute of Technology,
Technion City, Haifa 32000, Israel
Email: {robertom,tweed,ron,freddy}@cs.technion.ac.il

***Abstract*—Reconstruction of 3D objects based on images is useful in many applications. One of the methods based on multi-image data is the Photometric Stereo technique relying on several photographs of the observed object from the same point of view, each one taken under a different illumination condition. The common approach is to estimate the gradient field of the surface by minimizing a functional, integrating the distance from the camera and thereby obtaining the geometry of the observed object. We propose an alternative method that consists of a novel differential approach for multi-image Photometric Stereo and permits a direct solution of a novel PDE based model without going through the gradient field while naturally dealing with shadowed regions. The mathematical well-posedness of the problem in terms of numerical stability yields a fast algorithm that efficiently converges, even for pictures of sizes in the order of several megapixels affected by noise.**

***Keywords*-Photometric Stereo, Partial Differential Equations, Shadows.**

I. Introduction

The classical computer vision topic of Shape from Shading (SfS) was recently revitalised by a series of research contributions driven in part by some interesting new applications [17], [20], [13]. In this context a number of multi-image depth recovery techniques, based on inverted shading models have been addressed in the literature [21]. Among these is the method of Photometric Stereo (PS) which involves capturing multiple images of an object from the same viewpoint and illuminated independently from different directions. It has gained popularity due to the feasibility of implementing controlled scene illumination systems. This method emerges as one of the most readily implementable depth recovery methods and recently new ideas have been introduced in order to solve the PS problems more efficiently, see e.g. [11].

A. The Shape-from-Shading problem

Let us briefly recall the simplest setting of the SfS problem model for the orthographic viewing of a Lambertian surface.

We consider an unknown surface $h(x,y) = (x,y,z(x,y))$, defined on the compact domain $\overline{\Omega} = \partial\Omega \cup \Omega$ and denote the gradient of z by $\nabla z = (\frac{\partial z}{\partial x}, \frac{\partial z}{\partial y})$. The basic Lambertian surface assumption yields an image formation model determined by the inner product between the light source direction ω and the outgoing normal to the surface $n(x,y)$, [8]:

$$R(n)(x,y) = \rho(x,y)\big(\omega \cdot n(x,y)\big) \tag{1}$$

where ρ is the surface albedo and $\omega = (\omega_1, \omega_2, \omega_3)$ (with $\omega_3 > 0$) is a unit vector pointing toward a far-away light source. The surface normal is given by:

$$n = \left(-\frac{\partial z}{\partial x}, -\frac{\partial z}{\partial y}, 1 \right) \frac{1}{\sqrt{1+|\nabla z|^2}}. \tag{2}$$

With this notation, from (1) we obtain the following non-linear model of image formation:

$$I(x,y) = \rho(x,y) \frac{-\frac{\partial z}{\partial x}(x,y)\omega_1 - \frac{\partial z}{\partial y}(x,y)\omega_2 + \omega_3}{\sqrt{1+|\nabla z(x,y)|^2}}. \tag{3}$$

It is clear that this partial differential equation does not allow the recovery of the correct surface $z(x,y)$ if additional information is not provided [14]. Indeed, we need sufficient further data, in the form of boundary conditions and also complete knowledge of the albedo to successfully recover $z(x,y)$ [15].

Most of the studies which have addressed the SfS problem using this formulation with various types of additional data work in two steps:

1) the estimation of the first derivatives of the surface $\frac{\partial z}{\partial x}$ and $\frac{\partial z}{\partial y}$, see e.g. [7] (usually via some minimization algorithms);
2) the recovery of the height from the gradient field over $\overline{\Omega}$, see e.g. [1] (by integration or by functional minimisation).

A two-step procedure introduces approximation errors which could be avoided if direct integration could be performed by considering (3) as a non-linear PDE of the Hamilton-Jacobi type. Such an approach would allow us to recover the function z without passing through its gradient field. In order to directly compute the height of the surface, several variational approaches have been studied using the Euler-Lagrange equations associated with functionals of $z(x,y)$ that involve the given image and iterate based on

This research was partly supported by European Community's FP7- ERC program, grant agreement no. 267414 and by Broadcom foundation.

some prior information on the surface [7]. Such PDEs have several drawbacks, most significant of which is the high order of the partial derivatives involved and the impossibility to have a solution with regularity less than a C^2 function [6].

In the framework of classical PDEs for single image data and known albedo there exists a well known direct approach to SfS which uses level sets [9], [3]. Its primary drawback is the need to know a-priori the albedo which is a key assumption in order to make this method work. Here however we shall consider a more realistic scenario where $\rho(x, y)$ is an unknown function and surface reconstruction is performed in a multi-image framework.

B. Photometric Stereo SfS setup

There are several ways to collect information about the surface in order to counteract the ill-posedness of (1) and to achieve the reconstruction of the object under observation. Photometric Stereo can be formulated to exploit the information contained in several images of the scene taken from the same viewpoint under different light source directions [19]. Let n be the number of images where each one corresponds to different light sources used in the PS setting, and let us denote PS_n to be the associated depth recovery problem. In this paper we derive the differential formulation of PS_n which relies on the previous setup of the basic orthographic SfS problem (1) and therefore we study the following non-linear system of PDEs:

$$I_i(x, y) = \rho(x, y)\frac{-\nabla z(x, y) \cdot \widetilde{\omega}^i + \omega_3^i}{\sqrt{1 + |\nabla z(x, y)|^2}}, \tag{4}$$

where $i = 1, \ldots, n$ and I_i is the image obtained by lighting up the surface h using the i-th light source ω^i.

The classical two-step procedure for PS_2 has been extensively investigated and in particular is analysed in [15], [10]. These works highlight the fact that it is not possible to recover ∇z locally even if the albedo is known, since an ambiguity remains. The local ambiguity is shown to be removable in most cases by global integrability constraints. A completely different way to view PS_2 is presented in [11]. This formulation exploits the property of the *photometric ratio* introduced in [18]. In other terms, the common non-linearity $\sqrt{1 + |\nabla z(x, y)|^2}$ is removed by division of the image irradiance equations and in [11] is proved that an equivalent differential problem

$$\begin{cases} b(x, y) \cdot \nabla z(x, y) = f(x, y), & \text{a.e.}(x, y) \in \Omega; \\ z(x, y) = g(x, y) & \forall (x, y) \in \partial\Omega, \end{cases} \tag{5}$$

admits a unique Lipschitz solution (i.e. differentiable almost everywhere). Here $g(x, y)$ is the boundary condition and we have that

$$b(x, y) = (I_2(x, y)\omega_1' - I_1(x, y)\omega_1'', I_2(x, y)\omega_2' - I_1(x, y)\omega_2'') \tag{6}$$

is the characteristics field and

$$f(x, y) = I_2(x, y)\omega_3' - I_1(x, y)\omega_3''. \tag{7}$$

In [11], the uniqueness of the solution has been proved in absence of shadows. In particular, the proof is based on full knowledge of the boundary condition $g(x, y)$. But this is not a reasonable assumption as this information can not be obtained just by taking pictures of the object we want to reconstruct.

In this paper we exploit the well-posed differential formulation (5) presented in [11] in order to use it for the PS_3 problem in the presence of shadows which add further complexity to the surface recovery. Our new approach uses the full PS_3 formulation where all the data is available (i.e. the regions where we have three images with no shadows) and employs the PS_2 problem for parts of Ω where there is shadow in one of the captured images.

Note that our hypotheses are slightly weaker than the ones assumed in [5], which addressed the same problem, but considers a two step procedure and regularization terms for smooth surfaces. We will focus here on the direct computation of the height z by solving a system of linear PDEs.

In Section II we introduce the theoretical formulation of the new differential approach defining how the problem with $n > 3$ can be addressed starting from a complete formulation of the problem with three images. Using this formulation we will show how it is possible to overcome the problem of handling shadowed regions. We also focus on important mathematical details by sketching the proof of uniqueness of the weak solution for our new PDE using the characteristic strip expansion method. In Section III we solve the problem of PS_3 when there are shadowed regions and an unknown albedo. Section IV is devoted to the explanation of the numerical schemes used in the numerical test. They are essentially based on an up-wind and semi-Lagrangian scheme to follow the propagation of information on Ω. In Section V the numerical results on a realistic synthetic test are presented in which shadows are present.

II. Direct Surface Reconstruction Using Multiple Images and Shadows

When using multiple light sources we have an additional requirement on the lighting directions. A further assumption needs be made when more than two images are taken into account. By considering the reflectance equation (1) as a linear function with respect to the light source vector, we are constrained to consider only non-coplanar light sources. This inconvenience has been studied in [12] with respect to the PS_3 problem. It is well known that an image obtained with a linearly dependent light source with respect to the other images (in the sense of (1)) does not add any additional information [16], [2].

A. Weighted Photometric Stereo for multiple images with shadows

If we have three linearly independent images obtained by individually shining three parallel-ray light sources in non-coplanar directions then we can consider the set of unique image pairs and have the following system of linear PDEs:

$$\begin{cases} b^{(1,2)}(x,y)\cdot\nabla z(x,y) = f^{(1,2)}(x,y) \\ b^{(1,3)}(x,y)\cdot\nabla z(x,y) = f^{(1,3)}(x,y) \\ b^{(2,3)}(x,y)\cdot\nabla z(x,y) = f^{(2,3)}(x,y) \end{cases} \tag{8}$$

of the same type as (5), where

$$\begin{aligned} b^{(h,k)}(x,y) = (I_k(x,y)\omega_1^h - I_h(x,y)\omega_1^k, \\ I_k(x,y)\omega_2^h - I_h(x,y)\omega_2^k) \end{aligned} \tag{9}$$

and

$$f^{(h,k)}(x,y) = I_k(x,y)\omega_3^h - I_h(x,y)\omega_3^k \tag{10}$$

with (h,k) being the combination of two of the first three natural integers without repetitions.

We can now describe our novel contribution which is to ensure the well posedness of the PS$_n$ problem by exploiting the linearity of the basic differential formulation (5) and reducing it to a single PDE which can handle shadowed regions in a natural fashion. Since (5) does not lose the well-posedness if we multiply the equations by a function $q(x,y)$ on both sides (i.e. $b(x,y)$ and $f(x,y)$), we are able to define the ingredients of a weighted PS$_n$ model (W-PS$_n$) by considering the functions

$$b_n^w(x,y) = \sum_{p\in\binom{[n]}{2}} q_p(x,y)b^p(x,y) \tag{11}$$

and

$$f_n^w(x,y) = \sum_{p\in\binom{[n]}{2}} q_p(x,y)f^p(x,y) \tag{12}$$

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-PS$_n$ formulation is therefore

$$\begin{cases} b_n^w(x,y)\cdot\nabla z(x,y) = f_n^w(x,y), & \text{a.e. } (x,y)\in\Omega \\ z(x,y) = g(x,y) & \forall (x,y)\in\partial\Omega. \end{cases} \tag{13}$$

We next explain how shadows influence the definition of the weights (hence of b_n^w and f_n^w).

A key observation we can make is that it is possible to use weight-functions q_p that are vanishing while preserving the well-posedness of the problem. It is obvious that we do not care about the signs of the functions q_p; the main importance is related to the set of points where they are null.

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 a PS$_1$ problem. Our aim is to consider the weights as switches able to locally nullify the involvement of an image pair in the summations (11) and (12) when the functions b^p and f^p 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 set-up, we consider the point $(x,y)\in\overline{\Omega}$ 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 follows:

$$q_{(h,k)} = H(I_h(x,y))H(I_k(x,y)). \tag{14}$$

These functions are not continuous and therefore in order not to complicate (13) by adding such discontinuous functions, we consider a regularization of (14). By assuming the shadows to be open sets, we can regularize the weights using *cutoff functions* as constructed in [4].

B. Uniqueness of the weak solution of W-PS$_3$

In order to complete the theoretical analysis we will extend the uniqueness results of the differential problem (13) in the case of a weak solution. Discussion of discontinuities and multiple objects is beyond the scope of this paper. Our purpose is to prove the uniqueness of the solution of (13) in the Lipschitz function space via the method of characteristics. The meaning of a weak solution here is intended as a combination of classical solutions, each defined on a different domain. These domains are then going to be patched together in such a way that, across the boundaries between domains on which there are discontinuities in some derivatives, the equation (13) is satisfied. Let us recall that the points where the surface z is not differentiable are the same where the functions b_n^w and f_n^w are discontinuous (jump discontinuity) [11].

Since the complete proof of well-posedness of the W-PS$_3$ model can not be shown due to lack of space, we sketch it in the following two steps:

1) (**R1**), the absence of a critical point for the projected characteristic field, i.e. $b_3^w(x,y)\neq(0,0)$;
2) (**R2**), the propagation of the information from the boundary is not prevented between two sets separated by discontinuity.

III. W-PS$_n$ With No Boundary Condition

In Section II we extended the PS$_n$ model by supposing knowledge of the boundary condition $g(x,y)$. Clearly such a hypothesis compromises the use of that model in most real applications. It is therefore important to find a way to solve the PS$_n$ problem without requiring knowledge of the boundary condition.

To solve this problem we design a numerical strategy which involves selecting a single arbitrarily valued initial seed point within the reconstruction domain and robustly manipulating the path of the characteristics. We do this in order to numerically integrate the linear differential problem (13) so as to let the information travel in the most convenient directions for the whole domain.

A. Controlling the characteristic field

On the way to defining a numerical strategy we will need to manipulate the path along which the information travels. To do this we will exploit the following result:

Theorem *Let $b^p(x,y)$ be the vector field of (9) where $p \in \binom{[n]}{2}$. Then, $\forall p_1, p_2 \in \binom{[n]}{2}$ and $\forall (x,y) \in \Omega$ we have:*

$$b^{p_1}(x,y) \cdot b^{p_2}(x,y) \neq \pm |b^{p_1}(x,y)||b^{p_2}(x,y)|. \quad (15)$$

Proof

In order not to involve too many parameters, let us fix the indices p_1 and p_2 as $(1,2)$ and $(1,3)$ respectively. In order to prove that $b^{(1,2)}$ and $b^{(1,3)}$ are never parallel, we consider the contradiction assuming that there exists a point $(\widetilde{x},\widetilde{y}) \in \Omega$ such that:

$$b^{(1,2)}(\widetilde{x},\widetilde{y}) \cdot b^{(1,3)}(\widetilde{x},\widetilde{y}) = \pm |b^{(1,2)}(\widetilde{x},\widetilde{y})||b^{(1,3)}(\widetilde{x},\widetilde{y})|. \quad (16)$$

For the sake of clarity we omit the dependence on $(\widetilde{x},\widetilde{y})$. Now, by squaring both sides we have:

$$\begin{aligned}\left[b_1^{(1,2)} b_1^{(1,3)} + b_2^{(1,2)} b_2^{(1,3)}\right]^2 = \\ \left[\left(b_1^{(1,2)}\right)^2 + \left(b_2^{(1,2)}\right)^2\right]\left[\left(b_1^{(1,3)}\right)^2 + \left(b_2^{(1,3)}\right)^2\right] \quad (17)\end{aligned}$$

and by writing $b^{(1,2)}$ and $b^{(1,3)}$ explicitly from (9) we get:

$$\begin{aligned}&\left[I_2 I_3 (\omega_1')^2 - I_1 I_2 \omega_1' \omega_1''' - I_1 I_3 \omega_1' \omega_1'' + I_1^2 \omega_1'' \omega_1''' + \right. \\ &\left. I_2 I_3 (\omega_2')^2 - I_1 I_2 \omega_2' \omega_2''' - I_1 I_3 \omega_2' \omega_2'' + I_1^2 \omega_2'' \omega_2'''\right]^2 = \\ &\left[(I_2\omega_1' - I_1\omega_1'')^2 + (I_2\omega_2' - I_1\omega_2'')^2\right] \\ &\left[(I_3\omega_1' - I_1\omega_1''')^2 + (I_3\omega_2' - I_1\omega_2''')^2\right] \quad (18)\end{aligned}$$

Now, let us write the reflectance function (3) by simplifying the notation as follows:

$$I_j = \rho(x,y) \frac{i_j(x,y)}{\sqrt{1 + |\nabla z(x,y)|^2}}, \quad j = 1,2,3 \quad (19)$$

and substitute them in (18). We note that the quantity

$$\frac{\rho(x,y)}{\sqrt{1 + |\nabla z(x,y)|^2}} \quad (20)$$

is non vanishing and therefore can be eliminated from both sides of (18). Finally we can write (18) as follows:

$$\begin{aligned}&\left[i_2 i_3 (\omega_1')^2 - i_1 i_2 \omega_1' \omega_1''' - i_1 i_3 \omega_1' \omega_1'' + i_1^2 \omega_1'' \omega_1''' + \right. \\ &\left. i_2 i_3 (\omega_2')^2 - i_1 i_2 \omega_2' \omega_2''' - i_1 i_3 \omega_2' \omega_2'' + i_1^2 \omega_2'' \omega_2'''\right]^2 = \\ &\left[(i_2\omega_1' - i_1\omega_1'')^2 + (i_2\omega_2' - i_1\omega_2'')^2\right] \\ &\left[(i_3\omega_1' - i_1\omega_1''')^2 + (i_3\omega_2' - i_1\omega_2''')^2\right] \quad (21)\end{aligned}$$

and after some algebraic manipulation we get:

$$\begin{aligned}i_1(&\omega_1'''\omega_2''\omega_3' - \omega_1''\omega_2'''\omega_3' - \omega_1'''\omega_2'\omega_3'' + \\ &\omega_1'\omega_2'''\omega_3'' + \omega_1''\omega_2'\omega_3''' - \omega_1'\omega_2''\omega_3''') = 0. \quad (22)\end{aligned}$$

Assuming that we are considering a non shadowed point for the first image (i.e. $i_1 > 0$), we have that (22) is satisfied only if the light sources are collinear since it is equivalent to the following:

$$\det \begin{pmatrix} \omega_1' & \omega_2' & \omega_3' \\ \omega_1'' & \omega_2'' & \omega_3'' \\ \omega_1''' & \omega_2''' & \omega_3''' \end{pmatrix} = 0 \quad (23)$$

which is in contradiction with the photometric stereo assumption.

□

In other words this Theorem says that two different vector fields b^{p_1} and b^{p_2} can not be parallel. We will use this fact to control the direction of the characteristics for the case when $n = 3$ which can be easily generalised to any number of images. Let us introduce the set of three-lighted pixels $\mathcal{M}$ through its indicator function as follows:

$$\mathbb{1}_{\mathcal{M}}(x,y) = H(I_1(x,y))H(I_2(x,y))H(I_3(x,y)). \quad (24)$$

We would like to control the summation functions q_p in (11) and (12) such that we can take a linear weighted combination as given by the following two indexes:

$$(p_1^*, p_2^*) = \textit{argmin}_{(p_1,p_2)\in\binom{[3]}{2}} |b^{p_1} \cdot b^{p_2}| \quad (25)$$

with the aim to span the set of all possible directions of the derivatives by using the two least ill-conditioned directions $b^{p_1^*}$ and $b^{p_2^*}$. This is permitted to us through the linearity of the basic differential formulation (5).

Since $q_p(x,y) \equiv 1$, $\forall (x,y) \in \mathcal{M}$ and $\forall p \in \binom{[3]}{2}$, we can sum the equations in (8) as follows:

$$\left(\alpha b^{p_1^*} + \beta b^{p_2^*}\right) \cdot \nabla z = \alpha f^{p_1^*} + \beta f^{p_2^*} \quad (26)$$

where α and β are real coefficients.

We can now fully control the direction in which we compute the first derivatives of z independently at any pixel provided no shadows are involved. We will choose (α, β) in order to control the characteristic direction at a pixel to be in the most favourable direction as required by some integration strategy. On an axis aligned discretization grid we define eight primary directions of integration which can be seen in Figure 1.

We recognize these as the integration directions resulting from the possible locations accessible by the numerical schemes derived in the next section. Once we have chosen a particular direction d at a point $(x,y) \in \mathcal{M}$ we can compute values for (α, β) so that

$$\begin{pmatrix} b_1^{p_1^*}(x,y) & b_1^{p_2^*}(x,y) \\ b_2^{p_1^*}(x,y) & b_2^{p_2^*}(x,y) \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} d_1(x,y) \\ d_2(x,y) \end{pmatrix}. \quad (27)$$

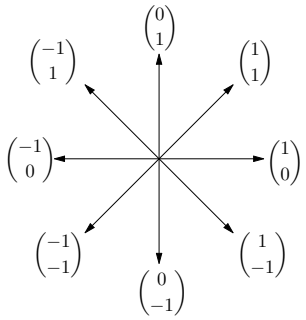

Figure 1. Here are shown all the main direction useful for the integration strategy on an axis aligned grid.

Provided that the light sources used to produce the three images are non-coplanar then Theorem (III-A) guarantees that the vectors $b^{p_1^*}$ and $b^{p_2^*}$ are not parallel and thus this 2×2 system is guaranteed to have a unique solution. Furthermore by virtue of (25) we have ensured that it is well conditioned. Thus for each point $(x, y) \in \mathcal{M}$ we can compute coefficients (α, β) such that

$$\left(\alpha b^{p_1^*} + \beta b^{p_2^*}\right) \cdot \nabla z = \alpha f^{p_1^*} + \beta f^{p_2^*}, \tag{28}$$

which is now a pixel specific version of (13).

B. Integration strategy for W-PS$_3$

We aim to minimize the accumulation error during the numerical integration and therefore we define a strategy in order to make clear how the surface can be reconstructed without the use of a-priori boundary conditions.

We have the following steps inspired by the wavefront expansion principle of fast marching methods over flat domains:

1) fix an arbitrary value to z for a point towards the center of the image domain not in shadow, namely $(\bar{x}, \bar{y})$ (see the orange point in Figure 2) and add 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 for z for each one by calculating p_1^*, p_2^*, α and β after determining what information is available as required by the forward schemes (30) or (32) derived in the next section;
3) for each newly visited pixel add its unvisited neighbors to the list of pixels to be visited;
4) in case of shadows we can change the wavefront propagation direction in order to *surround* the shadow sets (i.e. computing the boundary condition) as shown by the red arrows in Figure 2 and then solve the appropriate equation in (8) in the set of shadowed pixels when the wavefront expansion direction is in agreement with the shadowed pixel's characteristic direction;
5) the above steps are repeated until some stopping condition on convergence is fulfilled.

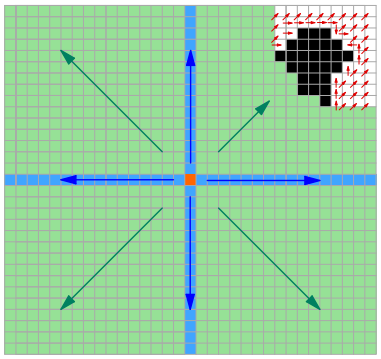

Figure 2. We fix an arbitrary point (in orange) and perform surface recovery along an expanding wavefront of visited pixels. When shadows occur, we orient the vector field in order to avoid them when their directions are inconsistent with the direction of expansion of the wavefront.

The main idea that we are proposing is that for all non shadow pixels we can orient the characteristic field so that its direction is convenient for use in our integration strategy. This is only possible using our new formulation in (13). In this way the advancing wavefront constantly makes new pixels available to be updated even when they are in the shadow set because at some point the direction of advance of the wavefront will be consistent with the direction of a pixel's characteristic which sits along the border of the shadow set. This pixel is then added to the set of pixels to be visited and in turn will act as a seed pixel that enables other points in the shadow set to be updated in the next incremental advance of the wavefront. In such a way the entire set of pixels can be visited and correctly updated.

IV. Numerical Schemes

We now describe the numerical methods that we will employ when we wish to determine the validity of the proposed model. The schemes originate from [11] where finite difference up-wind schemes and semi-Lagrangian schemes are used.

Let us start by considering a square domain Ω like the set $[a, b]^2$ and with a uniform discretization space step $\Delta = (b-a)/n$ where n is the number of intervals divided by the side of the square (that is $x_i = -1 + i\Delta$, $y_j = -1 + j\Delta$ with $i, j = 0, \ldots, n$). We will denote by $\overline{\Omega}_d$ all the points of the lattice belonging to $\overline{\Omega}$, by Ω_d all the internal points and by $\partial\Omega_d$ all the boundary points.

We recall the numerical schemes used for the forward approximation of (13) where the propagation of the information is considered according with the direction of the vector field that decide the direction of the derivatives. In order to simplify the notation we shall denote $b_3^w(x_i, y_j)$ by $b_{i,j} = (b_{i,j}^1, b_{i,j}^2)$ and $f_3^w(x_i, y_j)$ by $f_{i,j}$.

A. Forward up-wind scheme

Let us consider the following implicit up-wind scheme:

$$b^1_{i,j}\frac{Z_{i+1,j}-Z_{i-1,j}}{2\Delta}+b^2_{i,j}\frac{Z_{i,j+1}-Z_{i,j-1}}{2\Delta}=f_{i,j}+ |b^1_{i,j}|\frac{Z_{i+1,j}-2Z_{i,j}+Z_{i-1,j}}{2\Delta}+|b^2_{i,j}|\frac{Z_{i,j+1}-2Z_{i,j}+Z_{i,j-1}}{2\Delta} \tag{29}$$

for $i,j=1,\dots,n-1$. The artificial diffusion introduced in the right side of (29) allows to follow the vector field b by considering the most appropriate discretization for the fist derivative in order to track the characteristic lines. In particular it consists of a numerical scheme of consistency order equal to one with respect to both partial derivatives.

We can write (29) in the following iterative way:

$$Z^{t+1}_{i,j}=\frac{|b^1_{i,j}|Z^t_{i-\mathrm{sgn}\,(b^1_{i,j}),j}+|b^2_{i,j}|Z^t_{i,j-\mathrm{sgn}\,(b^2_{i,j})}+f_{i,j}\Delta}{|b^1_{i,j}|+|b^2_{i,j}|} \tag{30}$$

which is well-posed due to (**R1**) that guarantees that the division by $|b^1_{i,j}|+|b^2_{i,j}|$ does not involve any difficulty.

B. Forward semi-Lagrangian scheme

A second type of numerical scheme is based on the semi-Lagrangian approximation by considering the following equivalent equation obtained by dividing the two sides of the equation in (13) by the norm of $b^w_3(x,y)$:

$$\nabla_\gamma z(x,y)=\frac{f^w_3(x,y)}{|b^w_3(x,y)|}\quad\forall(x,y)\in\Omega \tag{31}$$

with $\gamma(x,y)=\frac{b^w_3(x,y)}{|b^w_3(x,y)|}$.

We observe that here again the division by $|b^w_3(x,y)|$ does not involve any difficulties for the numerical scheme due to (**R1**). Now, considering the definition of the directional derivative (in the opposite direction of γ, i.e. of the characteristic field), we can write the semi-Lagrangian scheme as follows:

$$\text{(sL)}\quad Z^{t+1}_{i,j}=Z^t(x_i-h\gamma^1_{i,j},y_j-h\gamma^2_{i,j})+\frac{f_{i,j}}{|b_{i,j}|}h \tag{32}$$

where the parameter h is greater than zero and is assumed equal to the size of the grid Δ in order to reach the highest order of convergence equal to one ([11]).

V. Numerical Tests

In the first part of the numerical tests we show a numerical analysis of convergence of the numerical schemes introduced in the previous section. In this synthetic case we take into account a Lipschitz surface that, besides points of discontinuity, presents high slope (i.e. a large Lipschitz constant). The initial images shown in Figure 3 are synthesized using an analytical function of a surface z in the domain $\overline{\Omega}_d=[-1,1]^2$.

The images include artificial shadow regions together with a non constant albedo mask. We corrupt the image data with 5% Gaussian noise to further improve the realism of the simulation. The iterations for both schemes are stopped when the L^∞ discrete norm

$$E_{L^\infty}=\max_{(x_i,y_j)\in\Omega_d}|Z^t(x_i,y_j)-z(x_i,y_j)|, \tag{33}$$

is less than a parameter $\varepsilon=10^{-7}$. Table I displays the error together with the measured run time at convergence. The numerical schemes were all implemented in C++ MEX files and executed in Matlab using a single core of a 2012 Intel Core i7 2.3.GHz MacBook Pro. For this work the parallelism available to an implementation of the integration strategy has not been utilized. Despite this limitation the algorithm runs at 6 *fps* for images of size 500×500. It should be noted though that this is a synthetic case and clearly the presence of noise influences the convergence of the numerical schemes and so they do not preserve the same rate of convergence with respect to the images without noise and thus have longer runtimes. Figure 4 shows the surfaces recovered from noisy images with four megapixels. Even if the reconstruction accurately preserves the shape of the surface, it produces some artifacts due to the noise and the rectangular occlusions. In this algorithm no regularisation procedure has been adopted in order to make the surface smoother.

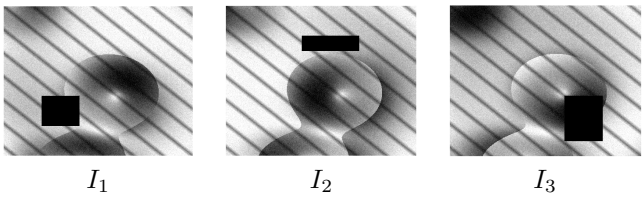

I_1 I_2 I_3

Figure 3. Set of synthetic images used in which the albedo is composed of diagonal stripes. 5% Gaussian noise and exaggerated black patches representing shadows are used to corrupt the images.

		up-wind		semi-Lagrangian	
	Δ	L^∞	time (sec)	L^∞	time (sec)
5 %	500	3.539×10^{-2}	0.162	2.332×10^{-2}	0.171
	1000	2.185×10^{-2}	0.561	1.166×10^{-2}	0.623
	2000	1.368×10^{-2}	2.783	6.248×10^{-3}	3.029
	500	6.635×10^{-2}	0.161	5.855×10^{-2}	0.274
	1000	3.578×10^{-2}	0.553	3.698×10^{-2}	1.201
	2000	3.917×10^{-2}	2.684	3.916×10^{-2}	4.354

Table I

IN ORDER TO SHOW THAT THE ORDER OF CONVERGENCE OF THE NUMERICAL SCHEMES IS ONE, WE COMPUTE THE L^∞ ERROR DOUBLING THE SIZE OF THE IMAGES (ADDING ALSO 5% OF GAUSSIAN NOISE) STARTING FROM DATA OF 500×500 TO 2000×2000 PIXELS.

In the first row of Figure 5 we consider the problem of the reconstruction of the well-known Beethoven bust after complicating the shape recovery by adding black occlusions. As shown in Figure 6 the computed shape does not have evident artifacts also in zones where the information has been removed.

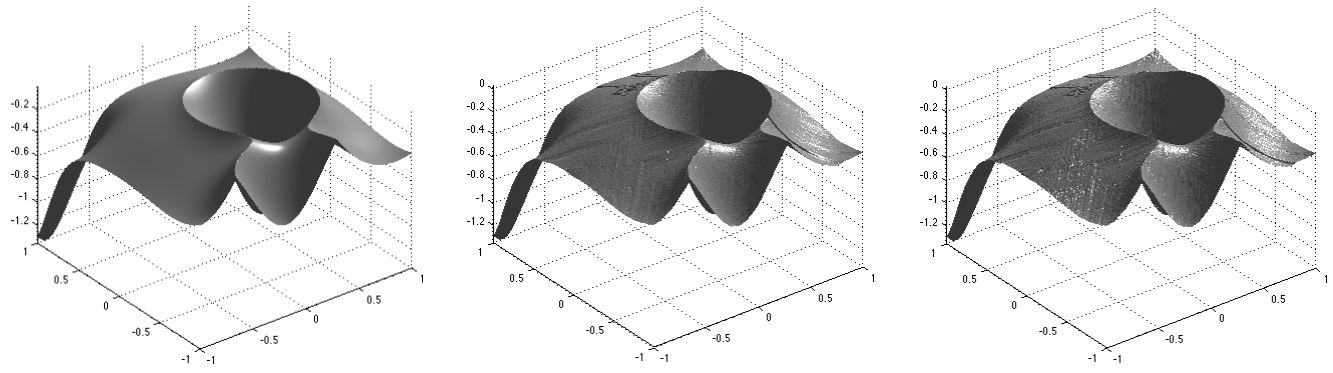

(original) (up-wind) (semi-Lagrangian)

Figure 4. All the surfaces shown here are related to the 2000×2000 pixels images with 5% of Gaussian noise.

I_1 I_2 I_3

Figure 5. Images of the Beethoven bust and coin with occlusions.

Figure 6. A novel view of the reconstruction of Beethoven using W-PS_3.

As a last example we consider the reconstruction of a coin using the images in the second row of Figure 5. These images have the shadows on different zones around the border of the human figure. The reconstruction in Figure 7 shows that the human shape is well computed and furthermore, that the flat region around this shape has been well preserved.

VI. Conclusions and Perspectives

In this work we have formulated a new approach to the multi-image shape reconstruction based on a direct PDE method. The advantage of this differential formulation stems from the direct computation of the height of the surface, which, along with proofs of the well-posedness of the PDE problem allows computation of the solution with a fast and stable algorithm even in the presence of noise. On a single core our non-optimized implementation converges in a few seconds even for sets of large images.

One direction of research we are currently investigating is improving the generality and richness of the shading model employed such as the use of near, non parallel-ray light sources and a perspective camera model. In particular the passage from the orthographic model to the perspective model using PS deserves attention, as this has been studied so far only by following the two-step shape recovery procedure and without taking into account shadows, as was recently done in [17].

The current work has mainly focused on the theoretical aspects of the suggested method but there are more practical elements which need to be pursued further to obtain a more coherent picture of the utility of the method. Using real digital photographs in a controlled physical setup is one such avenue. Of a similarly practical nature is the unused

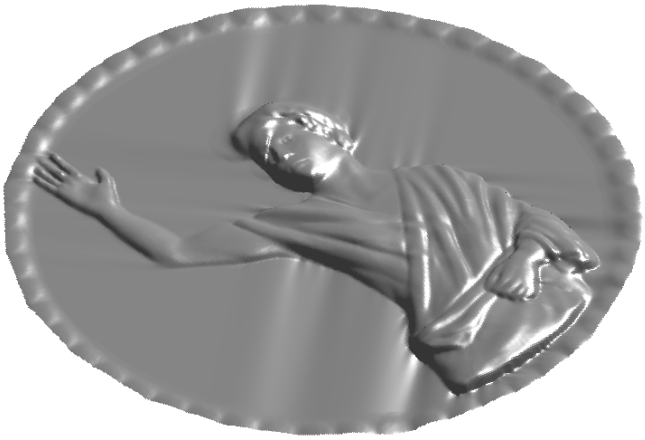

Figure 7. A novel view of the reconstruction of a coin model using W-PS$_3$.

parallelism inherently available in the proposed method. Our current implementation could potentially be accelerated by transferring it into a streaming programming environment such as CUDA to implement a fully parallelized version of the suggested integration strategy on a GPU.

REFERENCES

[1] S. Barsky and M. Petrou. The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 25(10):1239 – 1252, oct. 2003.

[2] P. Belhumeur and D. Kriegman. What is the set of images of an object under all possible lighting conditions? In *Computer Vision and Pattern Recognition, 1996. Proceedings CVPR '96, 1996 IEEE Computer Society Conference on*, pages 270 –277, jun 1996.

[3] A. M. Bruckstein. On shape from shading. *Computer Vision, Graphic, and Image Processing*, 44(2):139–154, 1988.

[4] K. O. Friedrichs. The identity of weak and strong extensions of differential operators. *Transactions of the American Mathematical Society*, 55(1):pp. 132–151, 1944.

[5] 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*, pages 290–303, Berlin, Heidelberg, 2008. Springer-Verlag.

[6] B. K. P. Horn. Height and gradient from shading. *The International Journal of Computer Vision, Winston*, 5:37–75, 1990.

[7] B. K. P. Horn and M. J. Brooks. The variational approach to shape from shading. *Computer Vision, Graphic, and Image Processing*, 33(2):174–208, 1986.

[8] B. K. P. Horn and M. J. Brooks. *Shape from Shading*. The MIT Press, 1989.

[9] R. Kimmel and A. M. Bruckstein. Tracking level sets by level sets: A method for solving the shape from shading problem. *Computer Vision and Image Understanding*, 62(1):47–58, 1995.

[10] R. Mecca and J. D. Durou. Unambiguous photometric stereo using two images. In *Image Analysis and Processing - ICIAP 2011 - 16th International Conference, Ravenna, Italy, September 14-16*, pages 286–295, 2011.

[11] R. Mecca and M. Falcone. Uniqueness and approximation of a photometric shape-from-shading model. *in printing on SIAM Journal on Imaging Sciences*, 2012.

[12] R. Mecca and S. Tozza. Shape reconstruction of symmetric surfaces using photometric stereo. In *International Dagstuhl Seminar on Innovations for Shape Analysis: Models and Algorithms, (to appear as a book in the Springer Verlag)*, page pp. 24, 2011.

[13] 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. *Computer Vision and Image Understanding*, 66(2):119 – 131, 1997.

[14] J. Oliensis. Shape from shading as a partially well-constrained problem. *Computer Vision, Graphic, and Image Processing: Image Understanding*, 54(2):163–183, 1991.

[15] R. Onn. and A. M. Bruckstein. Integrability Disambiguates Surface Recovery in Two-Image Photometric Stereo. *International Journal of Computer Vision*, 5(1):105–113, 1990.

[16] A. Shashua. Illumination and view position in 3d visual recognition. In *Advances in Neural Information Processing Systems 4*, pages 404–411. Morgan Kaufmann Publishers, 1992.

[17] A. Tankus, N. A. Sochen, and Y. Yeshurun. Shape-from-shading under perspective projection. *International Journal of Computer Vision*, 63(1):21–43, 2005.

[18] L. B. Wolff and E. Angelopoulou. 3-d stereo using photometric ratios. In *Proc. European Conference on Computer Vision*, pages 247–258. Springer, 1994.

[19] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):134–144, 1980.

[20] C. Wu, S. G. Narasimhan, and B. Jaramaz. A multi-image shape-from-shading framework for near-lighting perspective endoscopes. *International Journal of Computer Vision*, February 2009.

[21] R. Zhang, P. S. Tsai, J. E. Cryer, and M. Shah. Shape-from-shading: a survey. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, 21(8):690–706, 1999.