

Shape From Shading: a well-posed problem?

E. Prados
UCLA Vision Lab.,
Los Angeles, USA
eprados@cs.ucla.edu

O. Faugeras
Odyssée Lab., INRIA
Sophia Antipolis, France
faugeras@sophia.inria.fr

Abstract

Shape From Shading is known to be an ill-posed problem. We show in this paper that if we model the problem in a different way than it is usually done, more precisely by taking into account the $1/r^2$ attenuation term of the illumination, Shape From Shading becomes completely well-posed. Thus the shading allows to recover (almost) any surface from only one image (of this surface) without any additional data (in particular, without the knowledge of the heights of the solution at the local intensity “minima”, contrary to [6, 23, 8, 25, 12]) and without regularity assumptions (contrary to [17, 10], for example). More precisely, we formulate the problem as that of solving a new Partial Differential Equation (PDE), we develop a complete mathematical study of this equation and we design a new provably convergent numerical method. Finally, we present results of our new Shape From Shading method on various synthetic and real images.

1. Introduction and related work

Shape from shading (SFS) has been a central problem in the field of computer vision from the early days on. The problem is to compute the three-dimensional shape of a surface from one image of that surface. It is well known that this problem is ill-posed. In particular a number of articles show that the solution is not unique [19, 16, 25, 3, 23]. This problem has often been illustrated by such concave/convex ambiguities as the one displayed in Figure 1. In this figure, the ambiguity is due to a change in the estimation of the parameters of the illumination. In fact this kind of ambiguity can be generalized: in [3], Belhumeur and colleagues prove that when the illumination direction¹ and the Lambertian reflectance (albedo) of the surface are unknown, then the same image can be obtained by a continuous family of surfaces. In other words, they show that neither shading nor

¹In the case of a distant light source.



Figure 1. Example of Shape From Shading ambiguities: the crater illusion [20]. We see two craters, a small and a big one. We can turn these craters into volcanoes (although upside down) if we imagine the light source to be at the bottom of the picture rather than at the top. This image is actually that of a pair of ash cones in the Hawaiian Island, not that of a pair of craters.

shadowing of an object, seen from a single viewpoint reveals its exact 3D structure. This is the “Bas-relief Ambiguity”. Being aware of these difficulties, we therefore assume in this paper that all the parameters of the light source, the surface reflectance and the camera are known. Nevertheless this knowledge is not sufficient to get rid of some concave/convex ambiguities. In effect, even with completely controlled experimental conditions, the classical SFS methods are hampered by this kind of difficulties, e.g. see [16, 25, 23]. For example, let us focus on the “Eikonal” framework used in [25]. Rouy and Tourin assume that the camera performs an orthographic projection of the scene, that the surface is Lambertian and that the light source is at infinity in the direction of the optical axis of the camera. In this situation, a concave/convex duality clearly appears. For example, the surfaces represented in Figures 2-a), 2-b) and 2-c) yield the same images. The surfaces in b) and c) have been obtained from the surfaces a) by applying horizontal symmetries. Today, this concave/convex ambiguity is well understood. It is due to the existence of singular points², see [25]. The presence of the singular points is due to the modeling. In particular it is due to the fact that the authors

²The singular points are the pixels of the image corresponding to points of the surface such that the surface normal coincides with the light direction. These points have maximal brightness.

assume that the scene is illuminated by a light source located at infinity (or to the fact that they neglect the $1/r^2$ attenuation term in the brightness equation, see the remark of section 2). This ambiguity holds even when the light source direction does not coincide with the camera axis. It happens in the orthographic SFS [25, 6] as well as in the perspective SFS [23] problem. For a complete and generic study, the reader can refer to [23]. Nevertheless, *in contradiction with all these results, we prove in this paper that the SFS problem can be well-posed* (existence and uniqueness of the solution). Not surprisingly, this result is obtained by considering a more realistic image formation model.




Figure 2. The concave/convex duality in the Eikonal framework: The surfaces a), b) and c) all yield the same image.

The results shown in some recent SFS surveys, e.g., [28] in the traditional framework, i.e. Lambertian surface + light source at infinity + orthographic projection, are quite unsatisfactory, even with very simple synthetic images verifying the modeling hypotheses. Some authors have attempted to improve the applicability of the SFS methods by modeling the physics of the problem in a more realistic manner. For example, Lee and Kuo [13] consider some non-Lambertian cases and several authors [13, 27, 23, 26, 4] amongst others, take into account the perspective effect. Nevertheless, we feel that this work is in a sense premature since they expand upon an ill-posed problem without changing its status. We show in this paper how a simple change in the modeling of the problem can make the SFS problem well-posed (modulo some weak prior on the behaviour of the solution near the image boundary).

In detail, we assume that the camera is a pinhole and that the light source is located at the optical center. Nevertheless, contrary to [23, 22], we do not neglect the $1/r^2$ attenuation term (see section 2). As proved in section 4, this “new” term makes the problem better posed. In particular, the notion of singular points disappears as well as the concave/convex duality.

After having derived a *new explicit equation* for this modeling (sections 2 and 3), we perform a *complete theoretical study* of this equation (section 4). We then design a *new provably convergent algorithm* allowing to compute numerical solutions of the problem (section 5). Finally, we demonstrate the practical relevance of our new SFS method by displaying some experimental results (section 6).

2. Modeling of the SFS problem

In this section, we recall the image formation process. To do this, we will follow a light ray from the optical system to the light source. For more details we refer to [9].

The image brightness (assumed to be equal to the image irradiance E_i) is substantially proportional to the surface radiance³ L_s :

$$E_i = \mu L_s, \quad (1)$$

where the parameter μ depends on some intrinsic parameters of the camera (such as the diameter of the lens, the focal length, etc...). Next we assume that the scene is illuminated by a single point light source and that there are not interreflections. In this case, the relationship between the radiance L_s of a point of the surface³ with the surface normal (at this point) and the light source direction is described by the Bi-directional Reflectance Distribution Function (BRDF), see [9], which is constant for a *Lambertian* surface:

$$L_s = \frac{\alpha}{\pi} E_s, \quad (2)$$

where α is the albedo and E_s is the irradiance of the surface. Finally, the irradiance E_s of the surface point is given by (see [9]):

$$E_s = I_0 \frac{\cos \theta_i}{r^2} \quad (3)$$

where I_0 is the intensity of the light source, r is the distance between the light source and the surface point, and θ_i is the angle between the object surface normal and the direction to the light source.

Combining (1), (2) and (3), the brightness image is given by:

$$E_i = \sigma \frac{\cos \theta_i}{r^2}, \quad (4)$$

where σ is a constant coefficient related to the parameters of the imaging system, the intensity of the light source and the albedo of the surface.

Remark: If the light source is located far enough from the surface, the variations of the brightness of the image is essentially due to those of $\cos \theta_i$ and we can assume that r is constant. In this case, the brightness image is: (σ_2 is a constant coefficient)

$$E_i = \sigma_2 \cos \theta_i. \quad (5)$$

Contrary to most of the other classical SFS methods (which model the problem with a *far* point light source and a *orthographic projection*), in this paper we model the camera as a *pinhole* and we assume that the scene is illuminated by a point light source *located at the optical center*. Note that, this modeling is quite relevant for many applications. In effect, it approximately corresponds to the situation encountered in some medical protocols like endoscopy

³i.e. the radiance of the surface in the direction of the viewer.

[14, 22], and to the situation encountered when we use a simple camera equipped with a flash [22].

3. Mathematical formulation of the SFS problem

In this section, we formulate the SFS problem as that of solving a Partial Differential Equation (PDE). We describe an explicit equation arising from equation (4) and we detail the associated Hamiltonians.

Let Ω be an open subset of \mathbb{R}^2 . Ω represents the image domain, e.g., the rectangle $]0, X[\times]0, Y[$. As in [22], we represent the scene by a surface $\mathcal{S} = \{S(x); x \in \overline{\Omega}\}$, which can be explicitly parameterized by using the function $S : \overline{\Omega} \rightarrow \mathbb{R}^3$ defined by $S(x) = \frac{f \cdot u(x)}{\sqrt{|x|^2 + f^2}} (x, -f)$. $f > 0$ denotes the focal length. For such a surface \mathcal{S} , a normal vector $\mathbf{n}(x)$ at the point $S(x)$ is given by⁴:

$$\mathbf{n}(x) = \frac{f \nabla u(x) - \frac{f u(x)}{|x|^2 + f^2} x, \nabla u(x) \cdot x + \frac{f u(x)}{|x|^2 + f^2} f}{\sqrt{f^2 |\nabla u(x)|^2 + (\nabla u(x) \cdot x)^2 + Q(x)^2}} \cdot$$

For $y \in \mathbb{R}^3$, we denote $\mathbf{L}(y)$ the unit vector representing the light source direction at the point y . Here we assume that the *light source is located at the optical center*, so the vector $\mathbf{L}(S(x))$ is equal to: $\mathbf{L}(S(x)) = 1/\sqrt{|x|^2 + f^2} (-x, f)$.

Next we assume that the surface is *Lambertian*. If we denote $I(x) = \frac{E_i(x)}{\sigma}$, the brightness equation (4) becomes⁵:

$$I(x) = \frac{\cos \theta_i}{r^2}. \quad (6)$$

Since $\cos \theta_i$ is the dot product $\mathbf{L}(S(x)) \cdot \frac{\mathbf{n}(x)}{|\mathbf{n}(x)|}$ and $r = f u(x)$, we obtain from (6) the following PDE⁶:

$$I(x) f^2 \frac{\sqrt{[f^2 |\nabla u|^2 + (\nabla u \cdot x)^2]/Q(x)^2 + u^2}}{u} - u^{-2} = 0 \quad (7)$$

where $Q(x) = \sqrt{f^2/(|x|^2 + f^2)}$. Let us assume that the surface \mathcal{S} is visible (i.e. in front of the optical center); consequently u is nonnegative (see [22]). We can therefore simplify equation (7) by using the change of variables $v = \ln(u)$:

$$-e^{-2v(x)} + J(x) \sqrt{f^2 |\nabla v(x)|^2 + (\nabla v(x) \cdot x)^2 + Q(x)^2} = 0, \quad (8)$$

where $J(x) = \frac{I(x) f^2}{Q(x)}$ is a positive function. To this equation, we associate the *Hamiltonian*

$$H_F(x, u, p) = -e^{-2u} + J(x) \sqrt{f^2 |p|^2 + (p \cdot x)^2 + Q(x)^2}.$$

⁴The two columns of the Jacobian $DS(x)$ are tangent vectors to \mathcal{S} at the point $S(x)$. Their cross-product is a normal vector.

⁵We assume that all the parameters of the camera, of the light source and of the surface are known. Therefore σ and so $I(x)$ are known.

⁶Note: in equation (6): ∇u and u depend on x

4. Shape from Shading can be a completely well-posed problem!

4.1. Related work

To our knowledge, only Okatani and Deguchi [14], and Prados and Faugeras [23, 22] deal with the model considered here (pinhole camera and light source at the optical center).

In [14], Okatani and Deguchi do not formalize the problem with explicit PDEs (in particular equation (8)), and the associated Hamiltonians H_F . Let us emphasize that stating the problem as that of solving PDEs is a fundamental preliminary step for a theoretical study, for example for proving the uniqueness of the solution. Also, Okatani and Deguchi do not address at all the theoretical question of existence and uniqueness of a solution. They only propose a numerical method based on the propagation of the iso-distance contours, turning the static equation (4) into an evolution equation.

In [23, 22], Prados and Faugeras deal with equation (5) but not with equation (4). They neglect the $1/r^2$ term. By simplifying the modeling, they hope to simplify the problem. On the contrary, they make it more complex. In effect, as explained in [23] in this context the problem becomes ill-posed. Due to the existence of the singular points, the uniqueness of the solution does not hold. To get around this difficulty, Prados and colleagues [23] characterize a maximal solution by using sophisticated mathematical tools. They propose a viscosity method which does not necessarily require boundary data, in the sense that they are able to compute a solution⁷ just by fixing the height of the solution at one singular point. Nevertheless, even if they obtain excellent numerical results with real images containing a single singular point [22], they are unable to recover a surface containing several local “minima” such as the ones in figure 3-I without additional data (they need to know the values of the solution at all the local “minima”).

Aware of the major role played by the singular points, Oliensis and Dupuis [17] and Kimmel and Bruckstein [10] propose some “global” methods (based on the nature of these particular points) allowing to recover some *smooth* and constrained surfaces⁸ (in the case where these surfaces exist). In [15], Okatani and Deguchi use the isophotes of the image for classifying part of the singular points. They also suggest how to use the informations they obtain for improving the global methods of Dupuis and Oliensis [17] and of Kimmel and Bruckstein [10]. Note that although the papers [17, 10, 15] are based on an orthographic camera

⁷Which can be different from the original surface.

⁸They consider C^2 surfaces with second order derivatives satisfying some properties: for example, Kimmel and Bruckstein [10] consider surfaces represented by Morse functions (with non-degenerate Hessians).

and a single distant light source, the work therein can easily be extended to the more realistic modeling of [22] by using the tools developed in [23]. Nevertheless in practice (because of noise, of errors on the parameters and of incorrect modeling) the SFS equations such as those considered in [17, 10, 15, 22] (and more generally such as those described in [23]) do not have smooth solutions! Also, these global methods are quite disappointing when applied to real images.

Opposite to all the previous work [25, 23, 22, 17, 10, 15] which are hampered by the existence of the singular points, let us stress that *the notion of singular points does not make any sense as soon as we do not neglect the $1/r^2$ attenuation term*. As a consequence the difficulties described above completely disappear.

4.2. Well-posedness of the SFS problem

The results presented in this section are based on the notion of (discontinuous) viscosity solutions of the Hamilton-Jacobi PDEs. Let us recall that the viscosity solutions are solutions in a weak sense and that the classical (differentiable) solutions are particular viscosity solutions. For more details about this notion of weak solutions, we refer the reader to [1]. For an intuitive approach connected to computer vision, see for example [23] and references therein.

Since the CCD sensors have finite size, we assume that Ω is *bounded*. In this case, it is well known that the Hamilton-Jacobi equations of the form $H(x, u(x), \nabla u(x)) = 0, \forall x \in \Omega$, (and therefore our new SFS equation (8)) do not have a unique viscosity solution [1]. Nevertheless, for ensuring the uniqueness, it is sufficient to add “*state constraints*” on the boundary of the image. Yet, in [23] (but also implicitly, in [6]) it is shown that the idea of state constraints provides a more convenient notion of boundary condition than Dirichlet’s or Neumann’s. The “*state constraint*” is a boundary condition which is reduced to

$$H(x, u(x), \nabla u(x)) \geq 0 \quad \text{on } \partial\Omega,$$

in the viscosity sense (see for example [1]). This constraint corresponds to the Dirichlet’s conditions

$$\forall x \in \partial\Omega, \quad u(x) = \varphi(x) \quad \text{with } \varphi(x) = +\infty$$

in the viscosity sense. One may see adding state constraints to a PDE as a way of choosing the largest viscosity solution. The interest of the notion of state constraints is twofolds: 1) in contrast with the Dirichlet and Neumann boundary conditions, the state constraints do not require any data⁹. 2) for

⁹The Dirichlet (respectively, Neumann) boundary conditions require the knowledge of the values of the solution (respectively, of its normal derivative) on the boundary of the image. In the SFS problem such data is rarely available.

our SFS equation, the notion of state constraints is roughly equivalent to “ $u(x)$ increases when x tends to $\partial\Omega$ ”¹⁰. Let us emphasize that this constraint is in fact not a strong one since for example, the condition is satisfied as soon as the image to be processed contains an object of interest in front of a background.

In order to prove the uniqueness of the solution with state constraints, we have the theorem¹¹

Theorem 1 *Let Ω be bounded and smooth enough¹¹ and $H : \Omega \times \mathbb{R} \times \mathbb{R}^2 \mapsto \mathbb{R}$ continuous. If the hypotheses (H1)-(H2) and the boundary hypotheses (H3)-(H4) (described below) hold, then equation $H(x, u(x), \nabla u(x)) = 0, \forall x \in \Omega$, with state constraints on $\partial\Omega$ has a unique (discontinuous) viscosity solution on Ω .*

The hypotheses (H1), (H2), (H3) and (H4) are: $\forall 0 < R < +\infty$, there exists $m_R(t) \rightarrow 0$ when $t \rightarrow 0$ s.t.

- (H1) $H(x, u, p) - H(x, v, p) \geq \gamma_R(u - v), \forall x \in \Omega (\gamma_R > 0)$;
- (H2) $|H(x, u, p) - H(x, u, q)| \leq m_R(|x - y|(1 + |p|)), \forall x, y \in \Omega;$
 $\Gamma \subset \mathbb{R}^2$ being a neighborhood of $\partial\Omega$,
- (H3) $|H(x, u, p) - H(x, u, q)| \leq m_R(|p - q|), \text{ for all } x \in \Gamma;$
- (H4) $H(x, u, p) \xrightarrow[p \rightarrow +\infty]{} +\infty$ uniformly with respect to $x \in \Gamma$ and p in all compact subset of \mathbb{R}^2 ;

for all $-R \leq v \leq u \leq R$ and $p, q \in \mathbb{R}^2$.

We can prove¹² that Theorem 1 applies to the SFS Hamiltonian H_F as soon as the brightness image I is differentiable and verifies: there exist $\delta > 0$ and M s.t. $\delta \leq I(x) \leq M$ and $|\nabla I(x)| \leq M$. These assumptions are reasonable because with our modeling the brightness image is bounded and there are no black shadows. Hence the SFS equation (8) with state constraints on the boundary of the image has a unique viscosity solution. Therefore, if the actual surface producing the image roughly increases when x gets close to the boundary of the image, then we are able to characterize and recover it without any boundary data! Let us emphasize that this uniqueness result and the correspondence of the solution with the actual surface hold even when the actual surface has several local minima on Ω . In particular, we do not need to impose constraints at the singular points¹³. In other words the concave/convex ambiguities linked to the presence of the singularities completely disappear.

5. A provably convergent numerical method

In [14], Okatani and Deguchi describe a numerical method based on the propagation of the equal-distance contours. As Bruckstein and Kimmel [11] for the Eikonal equation, they design an evolution equation and propose to solve

¹⁰A more rigorous description of this constraint is detailed in our technical report [24]. Because of space, we cannot develop it here.

¹¹Theorem 4.6 of [1]. Ω “smooth enough” means $\Omega \in W^{2,\infty}$.

¹²We detail completely the proof in our technical report [24].

¹³Let us recall that the notion of singular points does not have sense in this setup!

it by using a level-set method [18]. Although the method of Okatani and Deguchi is rigorous, it suffers from an important drawback because it requires an *initial equal-distance contour*. The consequences are twofolds. First, it decreases the applicability of the method since such data (initial equal-distance contour and height of this contour) are usually not available. Second because these data, when available, are noisy, they in fact may perturb the reconstruction! This comes in contrast with our approach where the characterization of the solution (and therefore its computation) does not require any additional data! This also shows the interest of the theoretical analyses such as those presented in section 4.

We next propose a new numerical method. Contrary to [14], *it does not require any data* (in particular, it does not require an initial equal-distance contour). Moreover, our method is *provably convergent*: We prove that our approximation scheme is *stable, consistent* and that *their solutions converge toward the unique viscosity solution* of the problem. We then prove that the numerical solutions computed by our iterative algorithm converge toward the solutions of our scheme. Let us note that Okatani and Deguchi do not even consider such questions whose practical importance should not be underestimated because, for example, they allow to certify algorithms, to guarantee their robustness and to describe their limitations. . .

5.1. A new approximation scheme

In this section, we propose a finite difference approximation scheme. The reader unfamiliar with the notion of approximation schemes can refer to [2]. Let us just recall that, following [2], an approximation scheme is a functional equation of the form

$$S(\rho, x, u(x), u) = 0 \quad \forall x \in \bar{\Omega},$$

which ‘‘approximates’’ the considered PDE. S is defined on $\mathcal{M} \times \bar{\Omega} \times \mathbb{R} \times B(\bar{\Omega})$ into \mathbb{R} , $\mathcal{M} = \mathbb{R}^+ \times \mathbb{R}^+$ and $\rho = (h_1, h_2) \in \mathcal{M}$ defines the size of the mesh that is used in the corresponding numerical algorithms. $B(D)$ is the space of bounded functions defined on a set D .

5.1.1 Design of the approximation scheme

Let us consider the function $H_C : \Omega \times \mathbb{R}^2 \mapsto \mathbb{R}^2$ given by:

$$H_C(x, p) = J(x) \sqrt{f^2 |p|^2 + (p \cdot x)^2 + Q(x)^2}.$$

We therefore have

$$H_F(x, u, p) = -e^{-2u} + H_C(x, p).$$

As Prados and Faugeras in [23], we can prove that

$$H_C(x, p) = \sup_{a \in A} \{-f_c(x, a) \cdot p - l_c(x, a)\}$$

where A is the closed unit ball of \mathbb{R}^2 ($f_c(x, a)$ and $l_c(x, a)$ are detailed in our technical report [24], because of space). Following [23], we thus approximate $H_C(x, \nabla u(x))$ by:

$$H_C(x, \nabla u(x)) \approx \sup_{a \in A} \left\{ \sum_{i=1}^2 (-f_i(x, a)) \frac{u(x) - u(x + s_i(x, a)h_i \vec{e}_i)}{-s_i(x, a)h_i} - l_c(x, a) \right\} \quad (9)$$

where $f_i(x, a)$ is the i^{th} component of $f_c(x, a)$ and $s_i(x, a)$ is its sign. Thus, we obtain the approximation:

$$H_F(x, u(x), \nabla u(x)) \approx -e^{-2u(x)} + \sup_{a \in A} \left\{ \sum_{i=1}^2 (-f_i(x, a)) \frac{u(x) - u(x + s_i(x, a)h_i \vec{e}_i)}{-s_i(x, a)h_i} - l_c(x, a) \right\}. \quad (10)$$

By multiplying (10) by a fictitious time increment $\Delta\tau(x) > 0$ and by adding $u(x)$ and $-u(x)$ we obtain the scheme $S(\rho, x, u(x), u) = 0$ with S defined by:

$$S(\rho, x, t, u) = t - \Delta\tau e^{-2t} + \sup_{a \in A} \left\{ -1 - \Delta\tau \sum_{i=1}^2 \frac{|f_i(x, a)|}{h_i} \right\} u(x) - \Delta\tau \sum_{i=1}^2 \frac{|f_i(x, a)|}{h_i} u(x + s_i(x, a)h_i \vec{e}_i) - \Delta\tau l_c(x, a) \quad (11)$$

5.1.2 Stability of our approximation schemes

Definition: We say that a scheme S is stable if for all fixed mesh size ρ it has solutions and if all the solutions are bounded independently of ρ (see [2]).

For proving the stability of our scheme, we require mainly to ensure the monotonicity of the scheme (i.e. the function $u \mapsto S(\rho, x, t, u)$ is nonincreasing) and that the function $t \mapsto S(\rho, x, t, u)$ is nondecreasing.

Here, the function $t \mapsto S(\rho, x, t, u)$ is obviously nondecreasing. But the function $u \mapsto S(\rho, x, t, u)$ is not monotonous for any $\Delta\tau(x)$. In order to satisfy this condition, we choose

$$\Delta\tau(x) = \left(\sum_{i=1}^2 |f_i(x, a_0)|/h_i \right)^{-1},$$

where a_0 is the optimal control¹⁴ of (9). We can then prove¹⁵ the stability of our scheme:

Proposition 1 *If there exist δ and M such that $\forall x \in \bar{\Omega}$, $0 < \delta \leq I(x) \leq M$, then the scheme S is stable.*

¹⁴i.e. the a in A for which the maximum of (9) is reached.

¹⁵ Details of proofs can be found in our technical report [24] or in [21].

5.1.3 Convergence toward the viscosity solutions

By construction, our scheme (11) is consistent (following [2]) with the SFS equation (8) as soon as the brightness image I is Lipschitz continuous. So, using the stability and the monotonicity of our scheme and the uniqueness of the solution of the SFS equation (8), it follows directly from [2] that the solutions of our approximation scheme (11) converge towards the unique viscosity solution of equation (8) when the mesh size vanishes¹⁵. This result also proves¹⁵ the existence of the viscosity solution of equation (8).

5.2. A numerical algorithm

We now describe an *iterative algorithm* that computes an approximation of the solutions of the scheme (11) for all fixed $\rho = (h_1, h_2)$. We denote, for $k \in \mathbb{Z}^2$, $x_k = (k_1 h_1, k_2 h_2)$, and $Q := \{k \in \mathbb{Z}^2 \text{ s.t. } x_k \in \overline{\Omega}\}$. We call ‘‘pixel’’ a point x_k in $\overline{\Omega}$. Since $\overline{\Omega}$ is bounded; therefore the number of pixels is finite. The following algorithm computes for all $k \in Q$ a sequence of approximations U_k^n of $u(x_k)$:

Algorithm:

1. *Initialisation* ($n = 0$): $\forall k \in Q$, $U_k^0 = u_0(x_k)$;
2. *Choice of a pixel x_k and modification (step $n + 1$) of U_k^n* : we choose U_k^{n+1} such that

$$\begin{cases} U_l^{n+1} = U_l^n & \text{if } l \neq k, \\ S(\rho, x_k, U_k^{n+1}, U^n) = 0; \end{cases}$$

3. *Choose the next pixel x_k (using alternating raster scans [5]) and go back to 2.*

If $u_0(x) = -\frac{1}{2} \ln(I(x)f^2)$ then we can prove¹⁵ that step 2 of the algorithm has always a unique solution and that the computed numerical solutions converge (when $n \rightarrow +\infty$) toward the solutions of the scheme. *Details about the implementation of the algorithm can be found in our technical report [24].*

6. Experimental results

In all the experiments that follow, we have not used any additional data and have not fixed the value of the solution at any point.

Our algorithm being iterative, it requires a *stopping criterion*. We have chosen to stop the iterations when $\frac{1}{\text{card}(Q)} \sum_{k \in Q} |U_k^{n+1} - U_k^n| \leq s$, where $\text{card}(Q)$ is the number of pixels (in practice, we have fixed $s = 10^{-10}$). In the results displayed in the sequel, n_s corresponds to the number of iterations required for reaching this threshold.

6.1. Experiments with synthetic images

In Figure 3, we show the original object (groundtruth) in a), the input image obtained from the original object in b) and the reconstructed surface in c).

Let us recall that one of the most significant improvements of our method is that it can recover surfaces containing *several local minima without any additional data* contrary to the other propagation/PDE’s methods (e.g., [8, 25, 6, 11, 12, 23]) which require the knowledge of the values of the solution at the points of the local minima or which require stronger boundary conditions. We have first tested our algorithm with images synthesized from surfaces containing *several local minima*, e.g., the field of bumps displayed in the first row of Figure 3. As can be seen, the result is completely satisfying. The second row of Figure 3 displays the result of our algorithm from an image of the classical Mozart’s face [28]. In [23], Prados and Faugeras need to fix the height of the singular point on the nose for obtaining a relatively satisfying solution. Here without any information, we obtain a quite satisfying result.


In order to test the ability of our algorithm to deal with discontinuous images and nonsmooth surfaces, we have applied it to synthetic images generated by shapes containing (decreasing and increasing) edges. Experimentally, we have noticed that the reconstructions are visually perfect when the local maxima¹⁶ are smooth (see our technical report [24] for more numerical examples...). The third row of Figure 3 shows an example of reconstruction of a surface containing edges and crisp maxima. In this case, the reconstruction is satisfying (but not perfect). Let us note that the global methods of Dupuis and Oliensis [17] and of Kimmel and Bruckstein [10] are unable to deal with such nonsmooth surfaces.

In order to demonstrate the applicability of our method, we have shown the stability of our algorithm with respect to three types of errors: 1) image intensity errors due to uniformly distributed white noise; 2) errors on the gamma factor; 3) incorrect estimation of the focal length f of the camera. Let us note that we can prove¹⁵ that the errors on the σ parameter¹⁷ simply involve changes of scale. Figure 4 shows the results obtained from images Fig.3-II-b) of Mozart’s face distorted by *pixel noise*, *gamma corrections*, *albedo errors* and with *wrong focal length*. In addition to the accumulation effect, let us note that the errors imposed are quite large. Surprisingly, the algorithm produces relatively satisfying results given the large difference between the original image and the input images. In other respects, in practice with real images, the *light source is never located exactly at the optical center*. For a camera equipped

¹⁶ The maxima of the visualized surface correspond approximately to the minima of depth modulation. The surface is below the optical center.

¹⁷ Errors made on the parameters of the imaging system and the intensity of the light source.

**Results for an image
of a surface with several local minima:**



I-a) I-b) I-c)


“field of bumps”: $f = 23mm$, size= 400×400 pixels; $n_s \simeq 70$;



II-a) II-b) II-c)

Mozart's face: $f = 25mm$, size= 250×250 pixels; $n_s \simeq 50$;

Results for an image containing discontinuities




III-a) III-b) III-c)

$f = 23mm$, size= 250×250 pixels; $n_s \simeq 85$;

a) original surface (groundtruth); b) image obtained from the original surface a); c) surface reconstructed from the image b).

Figure 3. Examples of results with synthetic images.



a) b) c) d)

a) image of Fig.3-II-b) distorted by noise ($SNR = 4.04$), gamma distortion ($\gamma = 2$) and albedo distortion ($\sigma_\epsilon = 1.2\sigma$); b) surface reconstructed from image a) with a wrong focal length parameter $f_\epsilon = 15mm$ (the correct one is $f = 25mm$).

c) distorted image with $SNR = 3.36$, $\gamma = 0.5$, $\sigma_\epsilon = 0.8\sigma$;

d) surface reconstructed from c) with $f_\epsilon = 45mm$.

Figure 4. Results for the image of Mozart's face distorted by pixel noise, by gamma distortions, by albedo errors, and with a wrong focal length parameter.




Figure 5. Surfaces reconstructed from the images synthesized with light source located at $10cm$ to the left and $20cm$ above of the optical center (height of face $\simeq 25cm$, distance of the face to the optical center $\simeq 90cm$).

with a flash, the light source is located at several centimeters from the optical center. Figure 5 shows the robustness of our method to this kind of errors: it shows the reconstruction obtained from images of Mozart's face synthesized with the light source located at $10cm$ to the left and $20cm$ above of the optical center. For more experiments see [24] or [21].

6.2. Experimental results on real images

We have tested our new algorithm on a small database of real images of faces verifying approximately our modeling assumptions. This database is available online and partly displayed in Figure 6. For all these images, the size of the pixels is estimated to $0.018 \times 0.018mm^2$ (the images size is 512×384); the focal length is $5.8mm$. In these tests we have fixed¹⁸ $\sigma = 1000$ and we do not have made gamma correction. Our modeling of the problem assumes that the scene is lambertian and the albedo is constant. Because of eyes, for example, this hypothesis does not hold. To reduce the effects of this fault, we have painted them manually¹⁹. On the right of each image, the Figure 6 shows the surface reconstructed from it (after having inpainted the eyes). Let us mention that the number of iterations required for recovering the surfaces of Figure 6 is less than 80. For more details, see [24, 21]. Let us emphasize again that contrary to all²⁰ the other PDEs methods (which are based on the classical modeling of SFS, for example [8, 25, 6, 23]; and contrary to [14]), in the experiments presented here, besides the intensity image, we do not use additional data except the intrinsic parameters of the camera (focal length and size of pixels). For example, contrary to [22] in which also deals with real images of a faces, we do not need to fix the height of the singular point on the nose anymore.

7. Conclusion

The main difficulties encountered by the previous SFS work are due to the ill-posedness of the problem. Here,

¹⁸The errors on the σ parameter just involve changes of scale.

¹⁹This step can done automatically for example by matching the image to a model image already segmented; see for example [7].

²⁰With the exceptions of the global methods [17, 10] which require important regularity assumptions which are generally not available in practice...

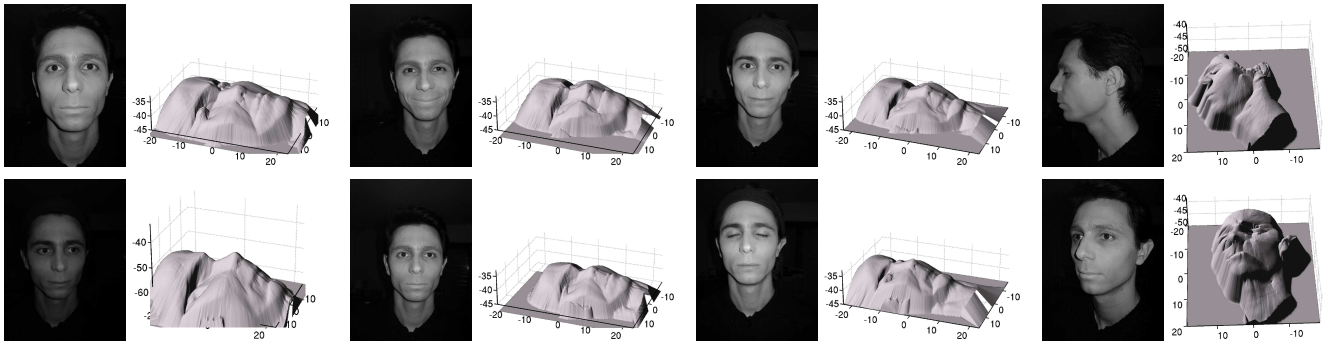


Figure 6. Examples of reconstructions from real images of faces.

we have proved that these difficulties disappear completely when we model the SFS problem in a realistic enough way in the case of a proximal light source. In other words, we have shown that in the previous SFS work, the problem is ill-posed because of the choices made in the modeling step. Also, we have proved that the SFS problem can be completely well posed. More precisely, after having detailed a new (explicit) SFS equation, we have developed a complete mathematical study of its solutions, we have proposed a provably convergent numerical method which requires neither additional data (contrary f.e. to [8, 25, 6, 12, 23, 14]) nor regularity assumptions (contrary f.e. to the global methods [17, 10]) We have tested it with success on synthetic and real images.

References

- [1] G. Barles. *Solutions de Viscosité des Equations de Hamilton–Jacobi*. Springer–Verlag, 1994.
- [2] G. Barles and P. Souganidis. Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic Analysis*, 4:271–283, 1991.
- [3] P. N. Belhumeur, D. J. Kriegman, and A. L. Yuille. The bas-relief ambiguity. *IJCV*, 35(1):33–44, Nov. 1999.
- [4] F. Courteille, A. Crouzil, J.-D. Durou, and P. Gurdjos. Towards shape from shading under realistic photographic conditions. In *ICPR*, 2004.
- [5] P.-E. Danielsson. Euclidean Distance Mapping. *Computer Graphics and Image Processing*, 14(3):227–248, Nov. 1980.
- [6] P. Dupuis and J. Oliensis. An optimal control formulation and related numerical methods for a problem in shape reconstruction. *The Annals of Applied Probability*, 4(2):287–346, 1994.
- [7] G. Hermosillo and O. Faugeras. Dense image matching with global and local statistical criteria: a variational approach. In *CVPR*, 2001.
- [8] B. Horn. Obtaining shape from shading information. In *The Psychology of Computer Vision*. McGraw-Hill, 1975.
- [9] B. Horn. *Robot Vision*. MIT Press, 1986.
- [10] R. Kimmel and A. Bruckstein. Global shape-from-shading. *CVGIP: Image Understanding*, 1995.
- [11] R. Kimmel and A. Bruckstein. Tracking level sets by level sets : A method for solving the shape from shading problem. *CVIU*, 62(2):47–58, July 1995.
- [12] R. Kimmel and J. Sethian. Optimal algorithm for shape from shading and path planning. *JMIV*, 14(2):237–244, May 2001.
- [13] K. Lee and C. Kuo. Shape from shading with a generalized reflectance map model. *CVIU*, 67(2):143–160, Aug. 1997.
- [14] T. Okatani and K. Deguchi. Shape reconstruction from an endoscope image by sfs technique for a point light source at the projection center. *CVIU*, 66(2):119–131, 1997.
- [15] T. Okatani and K. Deguchi. On classification of singular points for global shape from shading. In *ACCV’98*, pages 48–55, 1998.
- [16] J. Oliensis. Shape from shading as a partially well-constrained problem. *CVGIP: Image Understanding*, 54(2):163–183, 1991.
- [17] J. Oliensis and P. Dupuis. A global algorithm for shape from shading. In *Proceedings of ICCV’93*, pages 692–701, 1993.
- [18] S. Osher and J. Sethian. Fronts propagating with curvature dependent speed: algorithms based on the Hamilton–Jacobi formulation. *J. of Computational Physics*, 79:12–49, 1988.
- [19] A. Pentland. Local shading analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:170–187, 1984.
- [20] A. Pentland. Local shading analysis. *Trans. PAMI*, 6:170–187, 1984.
- [21] E. Prados. *Application of the theory of the viscosity solutions to the Shape From Shading problem*. PhD thesis, Univ. of Nice-Sophia Antipolis, 2004.
- [22] E. Prados and O. Faugeras. A rigorous and realistic shape from shading method and some of its applications. Technical Report 5133, INRIA, Mar. 2004.
- [23] E. Prados and O. Faugeras. Unifying approaches and removing unrealistic assumptions in Shape from Shading: Mathematics can help. In *Proceedings of ECCV’04*, 2004.
- [24] E. Prados, O. Faugeras, and F. Camilli. Shape from shading: a well-posed problem? Technical Report RR-5297, INRIA, Aug. 2004.
- [25] E. Rouy and A. Tourin. A Viscosity Solutions Approach to Shape-from-Shading. *SIAM Num. An.*, 29(3):867–884, June 1992.
- [26] A. Tankus, N. Sochen, and Y. Yeshurun. Perspective Shape-from-Shading by Fast Marching. In *Proceedings of CVPR’04*, 2004.
- [27] S. Yuen, Y. Tsui, Y. Leung, and R. Chen. Fast marching method for shape from shading under perspective projection. In *Proceedings of VIIP’02*, pages 584–589, Sept. 2002.
- [28] R. Zhang, P.-S. Tsai, J.-E. Cryer, and M. Shah. Shape from Shading: A survey. *Trans. on PAMI*, 21(8):690–706, 1999.