

A Non-Local Approach to Shape From Ambient Shading

Emmanuel Prados¹, Nitin Jindal¹, and Stefano Soatto²

¹ Perception Lab., INRIA Grenoble - Rhône-Alpes, FRANCE

² Computer Science Department, University of California, Los Angeles, USA

Abstract. We study the mathematical and numerical aspects of the estimation of the 3-D shape of a Lambertian scene seen under diffuse illumination. This problem is known as “shape from ambient shading” (SFAS), and its solution consists of integrating a strongly non-local and non-linear Integro-Partial Differential Equation (I-PDE). We provide a first analysis of this global I-PDE, whereas previous work had focused on a local version that ignored effects such as occlusion of the light field. We also design an original approximation scheme which, following Barles and Souganidis’ theory, ensures the correctness of the numerical approximations, and discuss about some numerical issues.

1 Introduction

Shape From Shading (SFS) refers to the problem of computing the three-dimensional shape of a surface, under certain assumptions on its reflectance and on the illumination, from a single grayscale image. By necessity, to render the problem tractable, these assumptions are rather coarse: Most restrict the illumination to a single point-light source at infinity [22, 4, 15, 8]. Only recently, [16] have shown that the problem actually simplifies when the attenuation of the light source at finite distance is taken into account. Nevertheless, due to inter-reflections and other complex phenomena, modeling illumination as a point source is very unrealistic even on a bright sunny day. Indeed, in most realistic conditions including indoors and outdoor overcast conditions, a uniform hemispherical illumination source is a more realistic model. The study of SFS under such illumination conditions has been pioneered by Langer et al. [11, 18, 10], and followed by others that we discuss shortly. In this work, we focus on the mathematical properties of the problem of “Shape From Ambient Shading” (SFAS), and seek for conditions that render the problem well-posed.

The paper is organized as follows. In section 2, we describe the modeling assumptions and introduce our notation. In section 3 we formulate the Shape From Ambient Shading problem as an integro-differential equation. We develop our theoretical results in section 4. We describe our numerical algorithm in section 5 and show some experimental results in section 6. Some discussions and future work are given in section 7.

1.1 Relation to Prior Work

Langer et al. [11, 18, 10] were the first to consider the case of *ambient lighting*, and to note that vignetting effects, far from being a nuisance, enable the inference of object shape similar to more traditional SFS, except for the added complication of the distributed source. In [19], Tian, Tsui and Yeung have proposed a *numerical* SFS algorithm for dealing with some non-punctual and multiple light sources (any combination of spherical, rectangular and cylindrical light sources). Following a more elaborate and physically motivated model of illumination, [14, 20, 10, 21] introduced methods to deal with interreflection. However, in none of these works [19, 11, 18, 10, 14, 20, 10] are the mathematical properties of the SFAS problem elucidated analytically. In particular, there are no results on the existence and uniqueness of solution for the ensuing global PDE.

At the opposite end of the spectrum, Lions, Rouy and Tourin [12] performed a *theoretical analysis* the SFS problem for multiple and continuous distributed light sources. As Tian, Tsui and Yeung [19], Lions, Rouy and Tourin neglect shadows (i.e. occlusions of the light sources by the surface itself); more specifically, they assume that for any fixed point x on the surface, all the light sources located on the hemisphere normal to the surface at x are visible from this point. This allows them to neglect the global nature of the equation, which in turn significantly simplifies the analysis.

As Langer et al. [11, 18, 10] we focus on *ambient lighting*. In their work, Langer et al. do not neglect the “shadows effect” and they model interreflections. They also underline the importance of ambient lighting in psychophysics. In this context, light comes from *all* directions and the assumption of Lions, Rouy and Tourin [12] is equivalent to assume that the solution is concave. Here, we do not want to limit ourself to concave objects. Therefore, Lions’ constraints are far too restrictive.³ The necessity to consider these phenomena takes us to mathematically uncharted territories. To the best of our knowledge, we are the first to provide theoretical results for the SFAS problem. Also, we introduce numerical algorithms verifying the properties of monotony, consistency and stability which typically ensure its convergence (see [1]).

2 Modeling Shape From Ambient Shading

Shape From Shading exploits assumptions on the illumination and reflectance properties of the scene (or of an object of interest within the scene) to relate its three-dimensional (3-D) shape to the measured grayscale image. The most typical assumptions are that the scene is Lambertian with constant diffuse albedo. This is akin to chalk and rough stone, and neglects specularities, translucency and other complex phenomena in the interaction of light with matter. While this assumption is clearly violated in most natural and man-made scenes, there are significant portions of scenes where the assumption is reasonable, and even objects that are far from Lambertian, such as human faces, have been successfully approximated as such for the purpose of analysis and inference (but not for synthesis, as humans are evolutionarily atuned to discriminate subtle features in human faces). Clearly, being SFS an ill-posed problem, there is no way to validate the assumptions on the data themselves, so applying SFS to a scene that is not Lambertian and that does not have constant diffuse albedo will results in gross errors even if the SFS algorithm used is provably correct and optimal. The second class of assumptions commonly made concern illumination. The most common assumption, that of a point light source, is made more for mathematical convenience than for realism. Under this model, anything hidden from direct line-of-sight to the sun would be invisible, clearly a far cry from reality. Modeling the entire sky as a constant-radiance hemisphere seems to be equally crude, but indeed it has been shown to be a better approximation than a single point-light source [9]. Clearly, both phenomena are important and we auspicate their eventual integration. In the next subsection we formalize these assumptions and introduce our notation.

2.1 Reflectance Assumptions

Let \mathcal{S} be a 2-D surface embedded in \mathbb{R}^3 that supports a bi-directional reflectance distribution function (BRDF) β with Lambertian reflection and constant diffuse albedo ρ . In other words, following [7], the BRDF at a point $p \in \mathcal{S}$ does not depend on the viewing direction ν_{px} , but only on the light source direction ν and on the position of the point itself $p \in \mathcal{S}$: $\beta(p; \nu_{px}, \nu) = \rho$. Because the intensity of the light source is not known, without loss of generality we can assume that the albedo to be equal to 1, and attribute the actual value to the light source.

³ For simplicity, however, we also neglect interreflections, as Lions et al. [12] did, and we lump their contribution into the ambient illumination term, up to additive errors.

2.2 Lighting Assumptions

Definition 1. [Ambient illumination] A power density distribution $R_L(\nu)d\nu$ defined for any ν in the unit sphere \mathbb{S}^2 of \mathbb{R}^3 is called ambient illumination.

We assume the dominating sky principle [11], so we neglect inter-reflections and, for any point of the surface, consider only radiant energy coming from the (spherical) sky. To simplify the problem, later we will also assume that the ambient illumination is homogeneous, that is to say, that the power density distribution is constant. As explained in section 2.3, this assumption is required if we want to get rid of other constraints while still keeping the problem manageable. For convenience, we assume $R_L(\nu) = 1$.

Now, unlike most previous work, we want to model the effect of self-occlusions, whereby the light source is only partly visible at each point. We then need to introduce the

Definition 2. [Light visibility] Let q be a point in \mathbb{R}^3 . We call visibility function and we denote $\chi_S(q; \nu)$ the indicator function of the directions $\nu \in \mathbb{S}^2$ from q that are not occluded by \mathcal{S} , i.e. the function $\chi_S(q; \nu) = 1$ if $\{q + \lambda\nu, \lambda \in \mathbb{R}_+\} \cap \mathcal{S} = \emptyset$, otherwise, $\chi_S(q; \nu) = 0$. The visibility cone at $q \in \mathbb{R}^3$ is

$$\mathbf{C}_{S,q} = \{\nu \in \mathbb{S}^2 : \chi_S(q; \nu) = 1.\}$$

The visibility function specifies if a point q is reached by the light ray of direction ν . If the 3-D point q sees the light ray, then $\chi_S(q; \nu) = 1$, else if this light ray is occluded by another part of the scene \mathcal{S} then $\chi_S(q; \nu) = 0$. The visibility cone assembles all the visible rays from a point q .

2.3 Resulting Radiance

Given the assumptions above, the radiance of the surface at a point p is given by

$$R_S(p) = \int_{\mathbb{S}^2} \chi_S(p; \nu) \langle \nu, \nu_p \rangle R_L(\nu) d\nu = \int_{\mathbf{C}_{S,p}} \langle \nu, \nu_p \rangle R_L(\nu) d\nu, \quad (1)$$

where ν_p is the unit normal vector to the surface S at p , see [7]. Here, the surface is implicitly assumed to be smooth. This ensures that all the light rays visible from a point come from above its tangent plane (the tangent plane would not be defined otherwise). So, for all points p on \mathcal{S} , all the light rays visible from that point are included in the hemisphere defined by the normal ν_p to the surface at that point; that is to say $\mathbf{C}_{S,p} \subset \text{Hemi}_{\nu_p}$. Therefore $\forall \nu \in \mathbf{C}_{S,p}$, $\langle \nu, \nu_p \rangle \geq 0$. To underscore the fact that our model would not make sense with discontinuous image intensities, but also for mathematical convenience, we will enforce this property in the radiance as:

$$R_S(p) = \int_{\mathbf{C}_{S,p}} \langle \nu, \nu_p \rangle^+ R_L(\nu) d\nu. \quad (2)$$

where for all a in \mathbb{R} , $a^+ = a$ if $a \geq 0$ and $a^+ = 0$ else.

Already at this point one can immediately see the difficulty introduced by self-occlusions, for the integration domain of (2) is restricted to the visibility cone $\mathbf{C}_{S,p}$, which directly depends of the global geometry of the scene \mathcal{S} . This is unlike traditional SFS, where the radiance only depended on local properties of the scene, for instance the direction of the normal ν_p to the surface at a given point. This requires the deployment of a different arsenal of tools that traditionally considered in SFS.⁴

⁴ In order to simplify the problem and to remove this global dependency, Lions, Rouy and Tourin [12] assume that for all the points of the surface, all the light sources located on the normal hemisphere are

Unlike most prior work, we want to consider *full* ambient illumination $R_L(\nu)dL(\nu)$; i.e., we want to deal with R_L distribution such that $\text{supp}(R_L) = \mathbb{S}^2$. In such a case, the assumptions of [12] are equivalent to assuming that the surface is convex which is too restrictive an assumption. On the other hand, to simplify the notation and keep the problem tractable, we assume that the *ambient illumination is uniform*, i.e. $R_L(\nu)dL(\nu) = R_0 dL$, where $R_0 \in \mathbb{R}$. For convenience, we assume that $R_0 = 1$. In the next section we relate the measurements, i.e. the image greyscale, to the unknown – the 3-D shape of the scene – via the model above.

3 Shape from Ambient Shading

In this section we formalize the problem of SFAS as the solution of a global integro-partial differential equation, which we analyze in the next section.

3.1 Imaging Equation

We assume that we measure a greyscale image $I : \mathfrak{D} \subset \mathbb{R}^2 \rightarrow \mathbb{R}^+$; $x \mapsto I(x)$, on a closed domain \mathfrak{D} . Our goal is to characterize the surfaces \mathcal{S} which generate it. Note that in general there is no guarantee that the surface is unique. We now need to link the measurements (I) with the unknowns (\mathcal{S}). To do so we use the assumptions developed in the previous section, together with the so-called Radiance equation [7], which approximates the brightness of a pixel x of the image with the radiance of the point $\pi_S^{-1}(x)$ of the surface viewed in x : $I(x) = R_S(\pi_S^{-1}(x))$. Using the results from the previous section we have

$$I(x) = \int_{\mathcal{C}_{S,p}} \langle \nu, \nu_p \rangle^+ d\nu, \quad (3)$$

where ν_p is the outward-pointing normal vector to the surface \mathcal{S} at the point $p = \pi_S^{-1}(x)$. In what follows we are going to assume that the data I corresponds with an image of a scene verifying our modeling assumptions. In particular, for convenience, we rescale the range so as to have $0 \leq I(x) \leq \pi$. Also for simplicity, we assume that the camera performs an *orthographic projection* of the scene. This is a reasonable hypothesis provided that the domain of interest in the scene is small compared to its distance to the camera. Under these conditions, we can represent the surface as the graph of a function u , and write the outward unit normal vector explicitly:

$$\mathcal{S} = \{(x, u(x)); x \in \mathfrak{D}\}; \quad \nu_{(x,u(x))} = \frac{1}{\sqrt{1 + |\nabla u(x)|^2}} (-\nabla u(x), 1).$$

Finally, following [15], we could assume that the camera is a pinhole. This assumption could be forgone at the cost of a more complicated notation, but the core of the analysis in this paper would hold nevertheless.

visible. More specifically, they assume that $\text{supp}(R_L) \cap \text{Hemi}_{\nu_p} \subset \mathcal{C}_{S,p}$ where for a vector \mathbf{q} in \mathbb{R}^3 , the hemisphere $\text{Hemi}_{\mathbf{q}}$ is the set $\{\nu \in \mathbb{S}^2 \mid \langle \mathbf{q}, \nu \rangle \geq 0\}$, and where the support $\text{supp}(\cdot)$ of a function is the closure of the set on which this function does not vanish. In other words, Lions et al. assume that there are no self-shadows. Also, such an assumption simplifies strongly the problem because we have then $R_L(\nu) = 0$ outside of $\mathcal{C}_{S,p}$ and so $\int_{\mathcal{C}_{S,p}} \langle \nu, \nu_p \rangle R_L(\nu) d\nu = \int_{\mathbb{S}^2} \langle \nu, \nu_p \rangle R_L(\nu) d\nu$ which completely removes the global dependency of the radiance with respects to the whole shape.

3.2 Formulation as a Partial Differential Equation

With the orthographic camera model, the image formation model above can be interpreted as a Partial Differential Equation (PDE) in the unknown function u :

$$I(x) = \int_{\mathbf{C}_{u,(x,u(x))}} \left\langle \frac{1}{\sqrt{1+|\nabla u(x)|^2}} (-\nabla u(x), 1), \nu \right\rangle^+ d\nu, \quad (4)$$

where $\mathbf{C}_{u,p}$ denote $\mathbf{C}_{S,p}$ (the surface \mathcal{S} is represented by the function u). Solving the SFAS problem then amounts to integrating the PDE (4) given an image I . Clearly the result would be meaningful only if a solution exists, and if it is unique, or at least if one can characterize the set of functions u that are indistinguishable in the sense of all solving (4) for a given measured image I . Note that this equation is a first-order stationary **global integro-partial differential** equation of the general form: $H(x, u(x), \nabla u(x), u(\cdot)) = 0$, $\forall x \in \text{Int}(\mathfrak{D})$. The numerical and theoretical study of the solutions of these kind of equation is done via the *Hamiltonian*

$$H(x, t, p, u) = \int_{\mathbf{C}_{u,(x,t)}} \left\langle \frac{1}{\sqrt{1+|p|^2}} (-p, 1), \nu \right\rangle^+ d\nu - I(x).$$

4 Analysis of the Shape From Ambient Shading Equation

In this section we consider the problem of uniqueness of solution of (4). While we show that the solution is, in general, not unique, we give an analytical characterization of all the different scenes that – under the given assumptions – yield the same measured image. This analysis is important both for the purpose of implementing viable numerical integration scheme, and also to make SFAS a useful tool in Computer Vision. This is akin to what is done in Structure From Motion [5], where the 3-D structure of a scene is in general not unique, but one can easily characterize the solutions as being equivalence classes under the similarity, affine or projective groups depending on knowledge on the camera calibration [6, 5, 13].

4.1 An Intrinsic Ambiguity

First, recall that $0 \leq R_S(p) \leq \pi$, $p \in \mathcal{S}$ and $\mathbf{C}_{S,p} \subset \text{Hemi}_{\nu_p}$, so one can easily show that $R_S(p) = \pi$ iff $\mathbf{C}_{S,p} = \text{Hemi}_{\nu_p}$. Now, let us consider a completely white image with a maximal intensity: $I(x) = \pi \forall x \in \mathfrak{D}$. With such an image, the solutions of equation (3) satisfy $\mathbf{C}_{S,p} = \text{Hemi}_{\nu_p}$ for all the points p on the surface. Therefore, if we represent the surface as the graph of the function u , it is easy to see that the surface lies below the tangent plane to the surface at the point $(x, u(x))$. So, the solutions u of (4) are concave, and so is the surface \mathcal{S} . Since inversely all concave functions generate such a white image then we can conclude that the set of solutions is comprised of all concave functions. In this case, the problem is clearly ill-posed because the image can be generated by a number of different surfaces, and therefore the solution cannot be unique. This problem does not arise only in this pathological case: It is patent as soon as the image contains a subset of pixels having the maximal intensity, as we illustrate in Figure 1. Pixels with maximal intensity are shown in red, and the green curve corresponds with a maximal solution when the blue gives the minimal one. Any curve between these two, which is concave on the set of points with maximal intensity, generates the same image as the one generated by the black curve. In the following sections, we will show that this condition is minimal, in the sense that the solution is unique if and only if there are no subsets of pixels having the maximal intensity. Also, when there are multiple solutions, they are characterized by in terms of their value on these subsets.




Fig. 1. Example of multiple solutions in dimension 2 when the image contains a subset of pixels having the maximal intensity. Any curve between the blue and the green curves, and which is concave on the set of points with maximal intensity, generates the same image as the one generated by the initial black curve.

4.2 Uniqueness Result and Characterization of the solutions

In this section we show that the solutions of the SFAS problem are characterized by their value on the subset $\{x \mid I(x) = \pi\} \subset \mathcal{D}$. To the end, let us define $\Omega = \{x \mid I(x) < \pi\}$ and let us complete the equation

$$H(x, u(x), \nabla u(x), u) = 0, \forall x \in \mathcal{D} \quad (5)$$

by some Dirichlet boundary conditions on $\mathcal{C}\Omega = \mathcal{D} - \Omega = \{x \in \mathcal{D} \mid I(x) = \pi\}$. In other words, we assume that we know the height of the solution on this subset. The equation then becomes

$$\begin{cases} H(x, u(x), \nabla u(x), u) = 0, & \forall x \in \Omega, \\ u(x) = \varphi(x) & \forall x \in \mathcal{C}\Omega. \end{cases} \quad (6)$$

For mathematical convenience, we also assume that the brightness image I is continuous (then Ω is an open subset of \mathcal{D}) and that the intensity is maximal on the boundary of the image (in other words, we assume that $\bar{\Omega} \subset \text{Int } \mathcal{D}$). We can now state the uniqueness theorem:

Theorem 1. *If u and v are two C^1 solutions to equation (6) then $u = v$ on \mathcal{D} .*

This theorem ensures that there exists at most a unique C^1 solution to equation (6). Also, it provides a characterization of the set of the solutions of equation (5), characterized by its values on the subset $\mathcal{C}\Omega$ (the region where $I(x) = \pi$). If the image never saturates ($\mathcal{C}\Omega$ is empty), then the solution is unique when complemented by a Dirichlet boundary condition. Equivalently, all solutions are parameterized by their boundary conditions. Because of space constraints, we cannot report the complete proof of theorem 1 here, and we refer the reader to our technical report [17] for details.

The relevance of this result from the standpoint of Computer Vision is that if we know the depth of the scene on the subset where the image is saturated, then there exists a unique solution to the Shape From Ambient Shading problem. This means that, elsewhere on the image, ambient shading is sufficient to recover the original surface which generated the image. In the next section we develop an approximation scheme for numerically integrating (6).

5 Approximation Scheme And Numerical Algorithm

In section 3 we have formalized the SFAS problem as the solution of a partial differential equation of the form $H(x, u(x), \nabla u(x), u) = 0$. We have then added Dirichlet boundary conditions on $\mathcal{C}\Omega = \mathcal{D} - \Omega$ to arrive at a unique solution when the image is not saturated. In order to compute a reliable numerical solution to this equation, we use machinery available for Hamilton-Jacobi equations. The key point consists then in designing approximation schemes which are monotone [2, 1].

5.1 A Monotonic Scheme

An approximation scheme is a functional equation of the form

$$T(h, x, u^\rho) = 0 \quad \forall x \in \bar{\Omega};$$

where T is a real function defined on $\mathbb{R}^+ \times \bar{\Omega} \times B(\bar{\Omega})$; $h \in \mathbb{R}^+$ defines the size of the grid that is used in the corresponding numerical algorithms (a 2D Cartesian grid); $B(\bar{\Omega})$ is the space of bounded functions defined on the set $\bar{\Omega}$. u^ρ is the unknown (u^ρ is a function). Also, we are interested in the solution u^ρ of the scheme T . Generally, we say that a scheme is monotone if for all $h \in \mathbb{R}^+$ and $x \in \bar{\Omega}$, the function $T(h, x, \cdot) : B(\bar{\Omega}) \rightarrow \mathbb{R}$ is monotone. That is, for all $y \in \bar{\Omega}$, $u(y) \geq v(y)$, then $T(h, x, u) \geq T(h, x, v)$. Following [1], we introduce the representation S of a scheme T as

$$S(h, x, u^\rho(x), u^\rho) = 0 \quad \forall x \in \bar{\Omega}, \quad (7)$$

where $S : \mathbb{R}^+ \times \bar{\Omega} \times \mathbb{R} \times B(\bar{\Omega}) \rightarrow \mathbb{R} : (h, x, t, u) \mapsto S(h, x, t, u)$. Note that a representation of a scheme is also a scheme. This last mathematical object allows us to take advantage of the tools developed by Barles and Souganidis [1] which require that the scheme be monotonic with respect to all the values of u up to the value at one point, generally $u(x)$. We then isolate $u(x)$ from the other values of u . Also, let us stress that the result demonstrated by Barles and Souganidis [1] is optimal for a large class of (static as well as evolutive) Hamilton-Jacobi equations.

The introduction of representations is also a way to simplify the computations. In effect, the representation of a scheme $T(h, x, u^\rho) = 0$ by a scheme of the form $S(h, x, u^\rho(x), u^\rho) = 0$ suggests an iterative algorithm for computing a numerical approximation of the solution of the scheme. Given u^n (the approximation of u^ρ at step n), and a point x of $\bar{\Omega}$, the associated algorithm consists in solving the equation

$$S(h, x, t, u^n) = 0 \quad (8)$$

with respect to t . A solution of (8) is the updated value of u^n at x . Here, we are then going to use the definition of monotonicity given by Barles and Souganidis in [1]:

Definition 3 (monotonicity). *The scheme $S(h, x, u^\rho(x), u^\rho) = 0$ defined in $\bar{\Omega}$, is monotone if $\forall h \in \mathbb{R}^+, \forall x \in \bar{\Omega}, \forall t \in \mathbb{R}$ and $\forall u, v \in B(\bar{\Omega})$,*

$$u \leq v \implies S(h, x, t, u) \geq S(h, x, t, v)$$

(the scheme is non-increasing with respect to u).

The interest of the monotonicity is twofold. (i) With other basic assumptions (monotonicity with respect to t , existence of a subsolution, bound for the subsolutions), this property is the key to ensure that the scheme is stable (existence of the solution and of an upper bound), that the computed approximations converge towards the solution of the scheme, see [15]. (ii) Combined with some stability and consistency properties, the monotonicity ensures that the solutions of the scheme converge towards the continuous solution of the considered PDE when the grid vanishes see [1]. In what follows, we are going to design a monotonic approximation scheme for the SFAS problem in order to take advantage of all these benefits.

5.2 Monotonic Scheme for the SFAS problem

For readability, we denote $H_{u,t}(x, p) = H(x, t, p, u)$. Let us recall that the Hamiltonian of interest in SFAS is

$$H_{u,t}(x, p) = \int_{\mathbf{C}_{u,(x,t)}} \left\langle \frac{1}{\sqrt{1+|p|^2}}(-p, 1), \nu \right\rangle^+ d\nu - I(x).$$

One can verify easily that $\mathbf{C}_{u,(x,t)}$ is decreasing (in the sense of inclusion) with respect to u and increasing with respect to t . Also, it follows that $H_{u,t}$ verifies exactly the same monotonic properties.

On the other hand, in order to get a consistent approximation scheme, we have to replace ∇u (represented by the variable p in the above Hamiltonian) in the PDE by one of its numerical approximations (finite differences). The difficulty is then to find such a discretization while maintaining monotonicity. In order to get a monotonic scheme, we take inspiration from *Lax-Friedrichs scheme* for conservation laws [3, 2]. We chose:

$$S(h, x, t, u) = H_{u,t}(x, \mathbf{D}u(x)) - \theta \mathbf{L}u^t(x), \quad (9)$$

where $\mathbf{D}u(x)$ is the vector obtained by a centered discretization of $\nabla u(x)$, more precisely, the i^{th} component of $\mathbf{D}u(x)$ is

$$[\mathbf{D}u(x)]_i = \frac{u(x + h\vec{e}_i) - u(x - h\vec{e}_i)}{2h}$$

and where $\mathbf{L}u^t(x)$ is the classical discretization of the Laplacian $\Delta u(x)$ (in which one replaces $u(x)$ by t), i.e.

$$\mathbf{L}u^t(x) = \sum_{i=1..N} \frac{u(x + h\vec{e}_i) + u(x - h\vec{e}_i) - 2t}{h^2}.$$

This scheme, however, is still not necessarily monotonic. To satisfy this property, we need to find an adequate value for θ . By differential calculus, one can verify that a sufficient condition to ensure this property is $\max_{i=1..N} h |\partial_{p_i} H_{u,t}(x, Dz)| \leq 2\theta$; see [17] for a detailed proof. By the same tools, one can also easily prove that $|\partial_{p_i} H_{u,t}(x, p)| \leq 2\sqrt{2}\pi$. The scheme(9) is then monotonic as soon as

$$\theta \geq \sqrt{2}\pi h.$$

Also, to limit the smoothing due to the Laplacian term introduced in the scheme (term which can be interpreted as a regularization), θ must be as small as possible.

On the other hand, under the assumptions of section 4.2, one can verify that any deep enough function is a subsolution of the scheme (9) (because the visibility cone becomes arbitrarily small). Moreover, the subsolutions are necessarily bounded by the function corresponding to convex hull defined by the Dirichlet boundary constraints. Since the scheme is also increasing with respect to t and verifies $\lim_{t \rightarrow +\infty} S(h, x, t, u) \geq 0$ then theorems 3.1 and 3.5 of [15] ensure that the scheme (9) is stable and that the iterative approximations converge towards the solution of the scheme.

In practice, we can start from any subsolution and we have just to update the surface with scheme (9) until convergence. Finally, our scheme being also consistent with the SFAS I-PDE, relying on Barle and Souganidis theorem [1], we can conjecture that the computed approximations converge towards the continuous solution of the I-PDE. This guarantees the reliability of our numerical approximations toward the theoretical solution of our problem.

6 Numerical Experiments

We focus here on the numerical results obtained by the algorithm associated to the scheme (9). As described in section 5.1, approximation schemes of the form (7) suggest an iterative numerical algorithm, whose updating step (at point x) consists in solving equation $S(h, x, t, u) = 0$ (equation in t), where u is the approximation of the whole solution at the previous step. Here, to solve

equation $H_{u,t}(x, \mathbf{D}u(x)) - \theta \mathbf{L}u^t(x) = 0$, we rewrite this equation as a fixed point equation $t = g(t)$, where

$$g(t) = \frac{1}{4} \left(\sum_{i=1,2} (u(x + h\vec{e}_i) + u(x - h\vec{e}_i)) - \frac{h^2}{\theta} H_{u,t}(x, \mathbf{D}u(x)) \right)$$

and then process the iterations $t_{n+1} = g(t_n)$. In practice this process systematically converges after less than 5 iterations (we assign t_0 to the previous value of $u(x)$). The numerical algorithm starts with a subsolution as a very steep valley such that visibility is closed to 0 for all points in the domain of the image. We refer the reader to [17] for further implementation details.

To test our algorithm, we consider some scenarios for which the problem is well-posed. In other words, we limit the computation domain to a subset of $\Omega = \{x \mid I(x) < \pi\}$. This computation domain is delimited by the red box in the corresponding figures. On the other part of the image domain, we enforce Dirichlet boundary conditions.

In our tests, we use the $\sin(x)\sin(y)$ surface. For the first test, we restrict the computation domain to a subset on which the surface is convex. As shown in Figure 2, the computed iterative solution converges accurately towards the original surface. In the second test, we want to extend




Fig. 2. Left: image generated by the $\sin x * \sin y$ surface with $h = 0.05$ and region of interest where we run the algorithm; middle: original surface (groundtruth) on the region interest; right: surface reconstructed by our algorithm (result).

the computation domain to both concave and convex areas. To remove the ambiguity due to points with maximal intensity, we reduce the intensity of the image by placing the $\sin(x)\sin(y)$ surface in a box, i.e. surrounded by four walls of a cube with the roof open. In this test, the algorithm converges towards the solution in both concave and convex regions. Nevertheless, as shown Figure 3, when the reconstruction is very accurate in the convex region, there is a significant error in the concave region.




Fig. 3. Left: image generated by the $\sin x * \sin y$ surface with $h = 0.05$ inside a cubical box and region of interest where we run the algorithm; middle: original surface (groundtruth) on the region of interest; right: surface reconstructed by our algorithm (result).

Table 1 shows the minimum and maximum values of the original surfaces in the regions of interest (where the algorithm is applied). It also shows the L_1 , L_2 and L_∞ errors. The top row shows the errors for the first test ($\sin(x) \sin(y)$ surface) illustrated in Figure 2. The second row shows the errors for $\sin(x) \sin(y)$ surface inside a box; it corresponds with the result of Figure 3. In our experiments, we have used the L_1 error to test for convergence.

In the second test, one can understand the error on the concave region as a result of the introduction of the regularization term (which was needed to make the scheme monotonic). To further analyze this effect, we focus on the concave part and we perform the following two experiments. 1) We run our algorithm with an input image containing the regularization term. More precisely, we use

$$\tilde{I}(x) = \int_{C_{u,(x,u(x))}} \left\langle \frac{1}{\sqrt{1+|\mathbf{D}u(x)|^2}} (-\mathbf{D}u(x), 1), \nu \right\rangle^+ d\nu - \theta \mathbf{L}u(x)$$

as input to our algorithm. So, in practice, the algorithm computes the solution of equation $\int_{C_{u,(x,u(x))}} \left\langle \frac{1}{\sqrt{1+|\mathbf{D}u(x)|^2}} (-\mathbf{D}u(x), 1), \nu \right\rangle^+ d\nu - \tilde{I}(x) - \theta \mathbf{L}u(x) = 0$ and the computed solution should then better coincide with the original surface. We then make this third test with the $\sin x \sin y$ surface inside the box (with a computation domain reduced to the concave part). As shown in table 2 and Figure 4, the algorithm is now able to recover accurately the surface. 2) Finally, since the regularization parameter θ is linearly dependent with the size of the grid h , then the regularization effect should reduce when the size of the grid vanishes. We then redo the second test ($\sin x \sin y$ surface inside a box, with the original image I , with the same reduced computation domain as previously) with smaller and smaller grid sizes: $h = 0.2, 0.1, 0.08, 0.05, 0.04$. Also, as we can see in Figure 5 and Table 3, the computed approximations actually converge towards the original surface when the grid size is reduced. In addition to confirm the above assertion, this also validates our methodology and our theory which ensures a well-posed algorithm whose the output converges towards the continuous solution when the grid vanishes.




Fig. 4. $\sin x * \sin y$ image with regularization and region of interest where we run the numerical scheme. Results of the numerical scheme with (right) and without (left) regularization in input image.

Table 1. Errors for the first two tests.

	min value	max value	L_1 errors	L_2 errors	L_∞ errors
$\sin x \sin y$, Fig. 2	-0.999707	0.066750	0.006191	0.009792	0.033867
$\sin x \sin y$ in box, Fig. 3	-0.999707	0.999568	0.188896	0.240712	0.372564

Table 2. Errors by adding the regularization term in the input image.

	Min Value	Max Value	L_1 Error	L_2 Error	L_∞ Error
without regularization	-0.999707	0.999568	0.186037	0.189434	0.207331
with regularization	-0.999707	0.999568	0.065627	0.067900	0.078941




Fig. 5. Reconstruction with different grid sizes h .

Table 3. Errors with respect to h .

grid sizes (h)	$h = 0.2$	$h = 0.1$	$h = 0.08$	$h = 0.05$	$h = 0.04$
L_1 error	0.504147	0.358676	0.270054	0.186037	0.151427
L_2 error	0.526644	0.371685	0.276862	0.189434	0.153691
L_∞ error	0.658852	0.424875	0.308127	0.207331	0.166671

7 Conclusion and Future Work

In 3-D reconstruction approaches to Computer Vision, illumination is rarely modeled explicitly. With few notable exceptions, most work in Structure From Motion assumes that illumination is constant and therefore it ascribes all photometric effects to the radiance of the scene, regardless of how it comes to be. In Shape From Shading, where the illumination is key, most existing work models it as an ideal point light source. In this paper we focus on the opposite abstraction, where the illumination is diffuse, and indeed it is constant. Outdoor scenes on a cloudy day, or indoor scenes in modern offices are reasonably well approximated by these conditions. Clearly one would like to account for arbitrary unknown radiant distributions, and possibly also illumination, but this would render the analysis prohibitive.

Already under the restrictive assumptions we have chosen to operate under, the problem of recovering the 3-D shape of the scene translates to a global integro-differential equation that, to the best of our knowledge, has never been analyzed. Although algorithms have been explored in the past to exploit diffuse shading for recovering properties of the scene, a thorough theoretical study of the mathematical properties of this problem has been lacking.

We believe we are the first to study the uniqueness of SFAS, to show that – in general – it is not unique, and to characterize the set of scenes that are indistinguishable, in the sense of satisfying the assumptions of SFAS and generating the same image.

While we believe that the main contribution of this paper is analytical, we do validate our results empirically in simulation. To that end, we propose a monotonic scheme for numerically integrating the SFAS equation, and show experimental results that highlight the features, and challenges, of this method.

Moving forward it would be desirable to develop both numerical schemes that are robust to noise, and – most importantly – algorithms that can provide competing explanations for the

image so that different assumptions, among which those for SFAS, can be applied to different portions of the image. For instance, it would not make sense to apply SFAS in regions of the image with albedo discontinuities (where stereo works well), and vice-versa one would not want to use stereo or Structure From Motion where albedo is constant. It would also be desirable to integrate SFAS with multiple-view reconstruction by providing constraints on multiple images of the same scene seen from multiple viewpoints. Finally, it would be desirable to relax the assumptions of pinhole orthographic camera to a perspective camera with a finite aperture.

Acknowledgement: Research was supported by the Flamenco project ANR-06-MDCA-007 and ONR N00014-08-1-0414.

References

1. G. Barles and P.E. Souganidis. Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic Analysis*, 4:271–283, 1991.
2. M. G. Crandall and P. L. Lions. Two approximations of solutions of Hamilton-Jacobi equations. *Mathematics of Computation*, 43(167):1–19, 1984.
3. M. G. Crandall and A. Majda. Monotone difference approximations for scalar conservation laws. *Mathematics of Computation*, 34(149):1–21, January 1980.
4. J.-D. Durou, M. Falcone, and M. Sagona. Numerical methods for shape-from-shading: A new survey with benchmarks. *Computer Vision and Image Understanding*, 109(1):22–43, jan 2008.
5. O.D. Faugeras. Three-dimensional computer vision: A geometric viewpoint. *MIT Press*, 1993.
6. R. Hartley and A. Zisserman. *Multiple view geometry in computer vision*. Cambridge University Press, 2000.
7. B.K. Horn. *Robot Vision*. mit-press, 1986.
8. B.K. Horn and M.J. Brooks, editors. *Shape from Shading*. The MIT Press, 1989.
9. J.J. Koenderink, S. C. Pont, A. J. van Doorn, A. M. L. Kappers, and J. T. Todd. The visual light field. *Perception*, 36:1595–1610, 2007.
10. M. S. Langer and H. H. Bulthoff. Depth discrimination from shading under diffuse lighting. *Perception*, 29(6):649–660, 2000.
11. M. S. Langer and S. W. Zucker. Shape from shading on a cloudy day. *Journal of Optical Society of America*, 11:467–478, 1994.
12. P.-L. Lions, E. Rouy, and A. Tourin. Shape-from-shading, viscosity solutions and edges. *Numer. Math.*, 64:323–353, 1993.
13. Y. Ma, S. Soatto, J. Kosecka, and S. Sastry. *An invitation to 3D vision, from images to models*. Springer Verlag, 2003.
14. S. Nayar, K. Ikeuchi, and T. Kanade. Shape from interreflections. *IJCV*, 6(3):173–195, 1991.
15. E. Prados. *Application of the theory of the viscosity solutions to the Shape From Shading problem*. PhD thesis, Univ. of Nice-Sophia Antipolis, 2004.
16. E. Prados and O. Faugeras. Shape from shading: a well-posed problem ? In *Proceedings of CVPR'05*, volume II, pages 870–877. IEEE, jun 2005.
17. E. Prados, N. Jindal, and S. Soatto. A non-local approach to shape from ambient shading. Technical report, INRIA, 2009.
18. A.J. Stewart and M. S. Langer. Towards accurate recovery of shape from shading under diffuse lighting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(9):1020–1025, 1997.
19. Y.L. Tian, H.T. Tsui, S.Y. Yeung, and S. Ma. Shape from shading for multiple light sources. *Journal of the Optical Society of America*, 16(1):36–52, January 1999.
20. T. Wada, H. Ukida, and T. Matsuyama. Shape from shading with interreflections under proximal light source-3D shape reconstruction of unfolded book surface from a scanner image. In *ICCV*, 1995.
21. J. Yang, D. Zhang, N. Ohnishi, and N. Sugie. Determining a polyhedral shape using interreflections. In *CVPR'97*, page 110, 1997.
22. R. Zhang, P.-S. Tsai, J.-E. Cryer, and M. Shah. Shape from Shading: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(8):690–706, August 1999.