



Numerical Shape-From-Shading for Discontinuous Photographic Images

JOSEPH KAIN AND DANIEL N. OSTROV

Received August 30, 1999; Revised May 7, 2001; Accepted May 23, 2001

Abstract. The height, $u(x, y)$, of a continuous, Lambertian surface of known albedo (i.e., grayness) is related to $n(x, y)$, information recoverable from a black and white flash photograph of the surface, by the partial differential equation

$$\sqrt{u_x^2 + u_y^2} - n = 0.$$

We review the notion of a unique viscosity solution for this equation when n is continuous and a recent unique extension of the viscosity solution when n is discontinuous. We prove convergence to this extension for a wide class of the numerical algorithms that converge when n is continuous. After discussing the properties of the extension and the order of error in the algorithms simulating the extension, we point out warning signs which, when observed in the numerical solution, usually indicate that the surface is not continuous or that the viscosity solution or its extension does not correspond to the actual surface. Finally, we discuss a method that, in some of these cases, allows us to correct the simulation and recover the actual surface again.

Keywords: shape-from-shading, viscosity solutions, discontinuous Hamilton-Jacobi equations, convergent numerical methods

1. Introduction

The field of Shape-From-Shading studies how to recover a 3-D surface using a 2-D image of the surface. Consider the case where the 2-D image is a black and white photograph taken by a camera located far from the surface of interest using a flash as its only light source. We define the direction the camera points to be the negative z direction of a Cartesian (x, y, z) coordinate system and assume that the z coordinate of the surface can be described as a continuous function $u(x, y)$ (i.e., there are no shadows on the surface from the flash). If the albedo—i.e., the grayness—of the surface is known, then the brightness throughout the photographic image can be transformed into $I(x, y)$, the brightness of the photographic image that would occur if the surface were completely white. If the surface is Lambertian—i.e., the surface reflects light equally in all directions—then $I(x, y)$ is just the cosine of θ , the

angle between the positive z direction (due to the flash being a vertical light source) and the outward surface normal, which can be expressed in terms of the partial derivatives of u :

$$I(x, y) = \cos(\theta) = \frac{1}{\sqrt{1 + u_x^2 + u_y^2}}.$$

This Shape-From-Shading equation is easily transformed into the eikonal form

$$\sqrt{u_x^2 + u_y^2} - n(x, y) = 0 \quad (1)$$

where $n(x, y) \equiv \sqrt{\frac{1}{I^2} - 1}$ is known from the photographic information.

This paper is concerned with issues that arise in computing the solutions to (1) when n is a discontinuous function. In Section 2 of this paper, we review previous

work which has established that there is a unique solution to (1) for the following three cases: (1) when u is smooth, in which case there is a unique classical (i.e., smooth) solution, (2) when u is not necessarily smooth but n is continuous, in which case there is a unique viscosity solution, and (3) when n is not continuous, in which case there is a recently determined unique stable solution based on extending viscosity solutions. Section 2 will also review computational work with finite difference schemes that converge to the unique solution in these first two cases. In Section 3, we will quickly prove that, in the third case, most of these finite difference schemes must also converge to the unique solution. In Section 4, we will illustrate the use of this scheme and also discuss some potential problems that may arise. Specifically, while the viscosity solution is unique for case 2, it may not correspond to the actual surface. Because of this, the extension of viscosity solutions for case 3 may suffer from the same problem. For both cases 2 and 3, we will discuss how to look for warning signs in the computed solution that typically indicate that the unique solution differs from the actual surface. Finally, when these warnings signs are present, we suggest a method for finding a new solution which will sometimes (although not always) be able to resolve these concerns.

2. Shape-From-Shading History

Numerous papers over the past 25 years have studied the question of uniqueness and computation of the solution to (1) (e.g., Horn, 1975; Horn and Brooks, 1989; Frankot and Chellappa, 1988; Bruss, 1982; Saxberg, 1992; Eastwood, 1995). For example, if the surface is known to be smooth and contains a finite, non-zero number of *critical points* (i.e., locations where $n = 0$) then, using only the value and concavity of n from a single critical point, Dupuis and Oliensis (1993) have established a control theory based algorithm that is proven to converge to the actual surface (except possibly near the edges of the photograph) as its mesh size decreases.

If there is no a priori knowledge of the surface's smoothness, but n is continuous and the domain of the photo is the closure of a bounded, open region, Ω , in the (x, y) plane, then Rouy, Tourin, Lions, and Ishii (Rouy and Tourin, 1992; Lions et al., 1993; Ishii, 1987) have shown that (1) has a unique *viscosity* solution, provided we know the height of the surface on $\partial\Omega$ (i.e., the boundary of Ω) and at all critical points

in Ω . Using

$$H(x, y, Du) \equiv \sqrt{u_x^2 + u_y^2} - n(x, y), \quad (2)$$

the viscosity solution for (1) is defined (see Crandall and Lions, 1983; Crandall et al., 1984) as the bounded, uniformly continuous (not necessarily smooth) function u where, for any smooth function ϕ ,

$$H(x_0, y_0, D\phi(x_0, y_0)) \leq 0$$

at any local maximum, (x_0, y_0) , of $u(x, y) - \phi(x, y)$ and

$$H(x_0, y_0, D\phi(x_0, y_0)) \geq 0$$

at any local minimum, (x_0, y_0) , of $u(x, y) - \phi(x, y)$. (We will refer to this definition as the “original” viscosity solution definition through the remainder of this section.) This definition implies that the surface corresponding to the viscosity solution may contain kinks (that is, ridges) that point upwards, but the surface cannot contain kinks pointing downwards.

Because H , the Hamiltonian defined in (2), is convex in Du , the viscosity solution also has the following control theory formulation (Rouy and Tourin, 1992; Lions, 1982; Ostrov, 2000):

$$u(x, y) = \inf_{\alpha(\cdot) \in A} \left[\int_0^S n(X(s)) ds + g(X(S)) \right], \quad (3)$$

where A is the set of all measurable functions with domain $[0, S]$ and range S^1 (that is, the set of all unit vectors in \mathbf{R}^2), $X(s)$ is a continuous path in Ω which is subject to the dynamics

$$\dot{X}(s) = \alpha(s), \quad X(0) = (x, y),$$

$$\text{and } X(S) \in \{\partial\Omega \cup \text{critical points of } \Omega\},$$

and g is the known height on $\{\partial\Omega \cup \text{critical points of } \Omega\}$. Note that since $\alpha(s)$ is always a unit vector, $X(s)$ moves with unit speed. Therefore, s , which denotes the arclength of $X(s)$ can also be thought of as time, so S denotes the first time that the path $X(s)$ crosses $\{\partial\Omega \cup \text{critical points of } \Omega\}$, the domain of g .

Both the original and the control theory formulations of the viscosity solution have been used to establish the convergence of computational schemes to the viscosity solution of (1). In Souganidis (1985) and Barles and Souganidis (1991), the original formulation

was exploited by Souganidis and Barles to establish the convergence of a wide class of stable, consistent, monotone finite difference schemes to the solution of many equations of the form $H(X, u, Du, D^2u) = 0$ where $X \in \mathbf{R}^n$. Rouy and Tourin (1992) applied this work to the specific Hamiltonian defined in (2) allowing them to establish a convergent scheme for the solution to (1). Whereas Rouy and Tourin approximated the values for their scheme by an iterative method, Sethian used a fast marching level set approach to show that the values in the scheme can be determined directly, allowing (1) to be solved in $O(m \ln(m))$ time where m is the number of grid points (Sethian, 1996). Related level set and front propagation approaches have also been obtained by Kimmel, Bruckstein, and others (Bruckstein, 1988; Kimmel, 1992; Kimmel and Bruckstein, 1995a, 1995b; Kimmel et al., 1995). At the same time, independent of Sethian, Tsitsiklis (1995) also found the exact same method of directly determining the values in the scheme, exploiting discretized versions of the control theory formulation of the solution which were proven convergent by Dupuis et al. (Dupuis and Oliensis, 1993; Kushner and Dupuis, 1992; Dupuis and Boué, 1999). Further, Tsitsiklis determined a method to compute the solution in $O(m)$ time (the theoretical minimum), but in general recommended using the $O(m \ln(m))$ method since the constant hidden by order notation in the $O(m)$ method is so large.

The viscosity solution requirement that n be continuous is a severe restriction since it is quite rare for n to be continuous when the surface is not smooth; i.e., it is almost always the case that discontinuities in the surface gradient will cause discontinuities in the photo brightness and therefore in n . Ishii (1985) defined an extension of the original viscosity solution definition which yields a unique solution for certain types of discontinuous Hamiltonians. Tourin (1992) generalized Ishii's approach with the following definition: a bounded, uniformly continuous function u is an (Ishii/Tourin) viscosity solution if, for any smooth ϕ ,

$$\limsup_{(x,y) \rightarrow (x_0,y_0)} H(x, y, D\phi(x_0, y_0)) \leq 0$$

at any local maximum, (x_0, y_0) , of $u(x, y) - \phi(x, y)$ and

$$\liminf_{(x,y) \rightarrow (x_0,y_0)} H(x, y, D\phi(x_0, y_0)) \geq 0$$

at any local minimum, (x_0, y_0) , of $u(x, y) - \phi(x, y)$.

Tourin then applied this definition to (1), but ran into obstacles in guaranteeing both uniqueness and existence. She did establish uniqueness and existence of an Ishii/Tourin viscosity solution under the following four conditions: (1) there is only one curve of discontinuity and it divides Ω into two regions, R_1 and R_2 , (2) the limit of n as any specific point on the curve is approached from R_1 is always non-zero and less than or equal to the limit of n as the point is approached from R_2 , (3) Ω contains no critical points and n is Lipschitz continuous except on the curve of discontinuity, and (4) $g \equiv 0$ (i.e., $u \equiv 0$ on $\partial\Omega$). Under these restrictions, Tourin applied the same algorithm used with Rouy for continuous n . The fourth restriction can be relaxed some, but if it is relaxed too much, existence problems may occur since the Ishii/Tourin viscosity solution, like the original viscosity solution, may only contain kinks that point upwards.

We illustrate these difficulties and motivate a new notion of solution presented in Ostrov (2000) when n is discontinuous by considering the 1-D analogue of (1):

$$\left| \frac{du}{dx} \right| - n(x) = 0$$

where $u(0) = a$ and $u(3) = b$ (so $\Omega = (0, 3)$). If $a = b = 0$ and $n \equiv 1$ (a continuous function), the viscosity solution is

$$u = \begin{cases} x & \text{if } x \in [0, 1.5] \\ 3 - x & \text{if } x \in [1.5, 3] \end{cases}$$

and contains a single upward kink at $x = 1.5$. If $a = b = 0$ and $n \equiv 2$ for $x \in (0, 1]$ and $n \equiv 1$ for $x \in (1, 3)$ (a discontinuous function) then an Ishii/Tourin viscosity solution exists:

$$u = \begin{cases} 2x & \text{if } x \in [0, 1] \\ 3 - x & \text{if } x \in [1, 3] \end{cases}$$

which also contains one upward kink. However, if we consider $a = 4, b = 0$, and $n \equiv 2$ for $x \in (0, 1]$ and $n \equiv 1$ for $x \in (1, 3)$, the obvious solution

$$u = \begin{cases} 4 - 2x & \text{if } x \in [0, 1] \\ 3 - x & \text{if } x \in [1, 3] \end{cases}$$

contains a downward kink and is therefore not an Ishii/Tourin solution, and, in fact, no Ishii/Tourin solution satisfying the boundary conditions exists.

In Ostrov (2000), a new extension of the viscosity solution for (1) with discontinuous n was introduced which defines a unique, stable solution that may contain downward kinks as in the example above and is subject to almost none of Tourin's restrictions. Specifically, the discontinuities of n within Ω can be comprised of any finite number of curves (or points) C_i that conform to the following (numerous, but very weak) restrictions: (1) the arclength of each C_i is finite, (2) the number of locations where each C_i fails to be C^2 is finite, (3) the number of intersections of any curve with itself or the other curves is finite and the angle between curves when they meet is non-zero, (4) the limit of n as any C_i is approached from either side is non-zero and as one progresses down each C_i , the side of the curve where n 's limit is higher cannot switch an infinite number of times, and (5) $\partial\Omega$ is smooth on either side of a point of intersection between $\partial\Omega$ and any C_i , (although $\partial\Omega$ may fail to be smooth at the point of intersection, so C_i is allowed to intersect $\partial\Omega$ at a corner of $\partial\Omega$.) We also assume that n has an upper bound, B , and that n is Lipschitz continuous except, of course, on the C_i 's. Throughout the remainder of this paper we will assume all discontinuous n conform to these restrictions.

The new extension is motivated by stability concerns. Specifically, let Ω' denote the set of all locations in Ω where n is continuous and define for each $(x, y) \in \Omega$,

$$\begin{aligned} n_\varepsilon(x, y) &\equiv \inf_{(\xi, \eta) \in \Omega'} \left(n(\xi, \eta) + \frac{B}{\varepsilon} \sqrt{(x - \xi)^2 + (y - \eta)^2} \right), \\ n^\varepsilon(x, y) &\equiv \sup_{(\xi, \eta) \in \Omega'} \left(n(\xi, \eta) - \frac{B}{\varepsilon} \sqrt{(x - \xi)^2 + (y - \eta)^2} \right). \end{aligned}$$

The n_ε are continuous functions that monotonically increase to n on Ω' as $\varepsilon \rightarrow 0$; the n^ε are continuous functions that monotonically decrease to n on Ω' as $\varepsilon \rightarrow 0$. Since the n_ε and n^ε are continuous, they correspond to unique viscosity solutions u_{n_ε} and u_{n^ε} . In Ostrov (2000), it is shown that for any $(x, y) \in \Omega$, the limit as $\varepsilon \rightarrow 0$ of u_{n_ε} and u_{n^ε} coincide, allowing the solution u to be defined as

$$u(x, y) \equiv \lim_{\varepsilon \rightarrow 0} u_{n_\varepsilon}(x, y) = \lim_{\varepsilon \rightarrow 0} u_{n^\varepsilon}(x, y). \quad (4)$$

Because this solution, u , is approached by the viscosity solutions corresponding to n_ε and n^ε (and many

other continuous approximating n_i , including convolutions of n), we will refer to u as an *extended viscosity solution*, even though it is often not itself a viscosity solution. Also, in Ostrov (2000), it is established that the extended viscosity solution, u , has a control theory representation similar to the representation of the viscosity solution when n is continuous; specifically,

$$u(x, y) = \inf_{\alpha(\cdot) \in A} \left[\int_0^S n^*(X(s)) ds + g(X(S)) \right] \quad (5)$$

where A is the set of all measurable functions from $[0, S]$ to S^1 , $X(s) = \alpha(s)$, $X(0) = (x, y)$, $X(S) \in \{\partial\Omega \cup \text{critical points of } \Omega\}$, and n^* is defined to be equal to n at every (x, y) where n is continuous and, for every (x, y) where n is discontinuous, n^* can take any value in the interval

$$\left[\liminf_{(\xi, \eta) \rightarrow (x, y)} n(\xi, \eta), \limsup_{(\xi, \eta) \rightarrow (x, y)} n(\xi, \eta) \right].$$

3. Convergence of Finite Difference Schemes

From the representation of the extended viscosity solution, u , in (5), it is intuitive (and shown rigorously in Ostrov, 2000) that for continuous or discontinuous functions n_1 and n_2 where $n_1(x, y) \geq n_2(x, y)$ for all $(x, y) \in \Omega$, the corresponding solutions u_{n_1} and u_{n_2} also have the property $u_{n_1}(x, y) \geq u_{n_2}(x, y)$ for all $(x, y) \in \Omega$ —given that g is the same for both cases. We now show that any finite difference scheme that reflects this monotonicity in n and converges to the solution when n is continuous must also converge to the solution when n is discontinuous.

Specifically, let us define $u_n^\rho(x, y)$ to be the numerical approximation to the extended viscosity solution, u , corresponding to n . (Since g is fixed in our discussion, we suppress the dependence on g in the notation for the approximation.) The approximation is defined for all (x, y) on the mesh $M \equiv \Omega \cap \{\Delta x \mathbf{Z} \times \Delta y \mathbf{Z}\}$ and we use ρ to denote the specific mesh $(\Delta x, \Delta y)$ and the mesh size in the sense that $\rho \rightarrow 0$ means $\Delta x, \Delta y \rightarrow 0$. We will define a scheme to be *monotonic in n* if $n_1 \geq n_2$ for all $(x, y) \in M$ implies $u_{n_1}^\rho \geq u_{n_2}^\rho$ for all $(x, y) \in M$. We can now establish the convergence of finite difference schemes for discontinuous n :

Theorem 1. *Any finite difference scheme for (1) that is monotonic in n and converges to the viscosity*

solution when n is continuous must also converge to the extended viscosity solution, u , when n is discontinuous.

Proof: If the scheme is monotonic in n , the fact that $n_\varepsilon \leq n^* \leq n^\varepsilon$ implies that

$$u_{n_\varepsilon}^\rho \leq u_{n^*}^\rho \leq u_{n^\varepsilon}^\rho$$

for all $\rho, \varepsilon > 0$ and $(x, y) \in M$. Now let the mesh size $\rho \rightarrow 0$. Since the scheme converges for continuous n , $\lim_{\rho \rightarrow 0} u_{n_\varepsilon}^\rho = u_{n_\varepsilon}$ and $\lim_{\rho \rightarrow 0} u_{n^\varepsilon}^\rho = u_{n^\varepsilon}$, therefore

$$u_{n_\varepsilon} \leq \liminf_{\rho \rightarrow 0} u_{n^*}^\rho \leq \limsup_{\rho \rightarrow 0} u_{n^*}^\rho \leq u_{n^\varepsilon}$$

for all $\varepsilon > 0$. Letting $\varepsilon \rightarrow 0$ and applying (4), we have that

$$u \leq \liminf_{\rho \rightarrow 0} u_{n^*}^\rho \leq \limsup_{\rho \rightarrow 0} u_{n^*}^\rho \leq u,$$

which implies that $\lim_{\rho \rightarrow 0} u_{n^*}^\rho = u$,

where u is the extended viscosity solution corresponding to the discontinuous n^* . \square

It would be surprising to find a convergent scheme that failed to be monotonic in n given that the monotonicity must exist in the continuum case. For example, the class of schemes formed by discretizing the control theory representation of the solution and proven to be convergent by Dupuis et al. (Dupuis and Oliensis, 1993; Kushner and Dupuis, 1992; Dupuis and Boué, 1999) are easily seen to be monotonic in n for the same reason their continuum counterparts are monotonic. Though less intuitively obvious, Rouy and Tourin's scheme is also monotonic in n . This follows from the fact that Sethian's direct computation of the Rouy and Tourin solution in Sethian (1996) for $\Delta x = \Delta y$ is the same algorithm as Tsitsiklis's, and Tsitsiklis (1995) showed that his algorithm was in the class of discretized control theory schemes considered by Dupuis.

The Sethian/Tsitsiklis scheme is a modification of the Dijkstra algorithm in the sense that the scheme works by building the solution up from the lowest values of u to the highest values of u . Therefore, reflecting the control theory perspective in (5), it starts at the smallest value of g on the boundary/critical points and works upward, never utilizing the values of g that correspond to local maxima of u , which we will exploit in the next section.

4. Numerical Experiments

In this section we will use the term “(extended) viscosity solution” to simultaneously refer to the viscosity solution for continuous n and the extended viscosity solution for discontinuous n . Also, as in the introduction, u will represent the actual surface, as opposed to the previous two sections where u represented the (extended) viscosity solution. In both of the examples presented in this section, we define the domain of the photograph, Ω , to be $(0, 1) \times (0, 1)$ and, unless otherwise specified, we approximate Ω by a 100×100 grid (i.e., $\Delta x = \Delta y (= \rho) = 0.01$). Because of its speed and simplicity, we will employ the Sethian/Tsitsiklis algorithm to determine u^ρ , the numerical approximation on the grid to the (extended) viscosity solution. Often, the (extended) viscosity solution will equal u , the actual surface, but this is not guaranteed. In this section, we will discuss warning signs which typically indicate cases where the two are not the same, and we will also describe how to remedy this problem in some of these cases.

For our first example, the extended viscosity solution will equal the surface, u . The surface we consider is

$$u(x, y) = \max(81[(x - 0.1)^2 + (y - 0.5)^2], [(x + 0.9)^2 + (y - 0.5)^2]),$$

which describes the intersection of two paraboloids: one centered at $(0.1, 0.5)$, the other at $(-0.9, 0.5)$. In Fig. 1(a) we show the n function determined from the photograph of u ; in Fig. 1(b) we show $u^{0.01}$, the approximation of the extended viscosity solution determined from n and the value of u along the square perimeter of the domain (note that there are no critical points). In the figure we present only the subsection of the domain where $(x, y) \in [0, 0.3] \times [0.25, 0.75]$ so that the intersection of the two paraboloids can be seen clearly.

Two types of characteristic curve patterns occur on the circle where the two paraboloids intersect, $(x - \frac{9}{80})^2 + (y - \frac{1}{2})^2 = (\frac{9}{80})^2$. For Eq. (1), the characteristic curves always flow in the direction of the gradient, Du . Therefore, outside the circle of the intersection, where the surface behaves like the paraboloid centered at $(0.1, 0.5)$, the characteristics emanate in straight lines from $(0.1, 0.5)$; inside the circle where the surface behaves like the paraboloid centered at $(-0.9, 0.5)$, the characteristics emanate in straight lines from

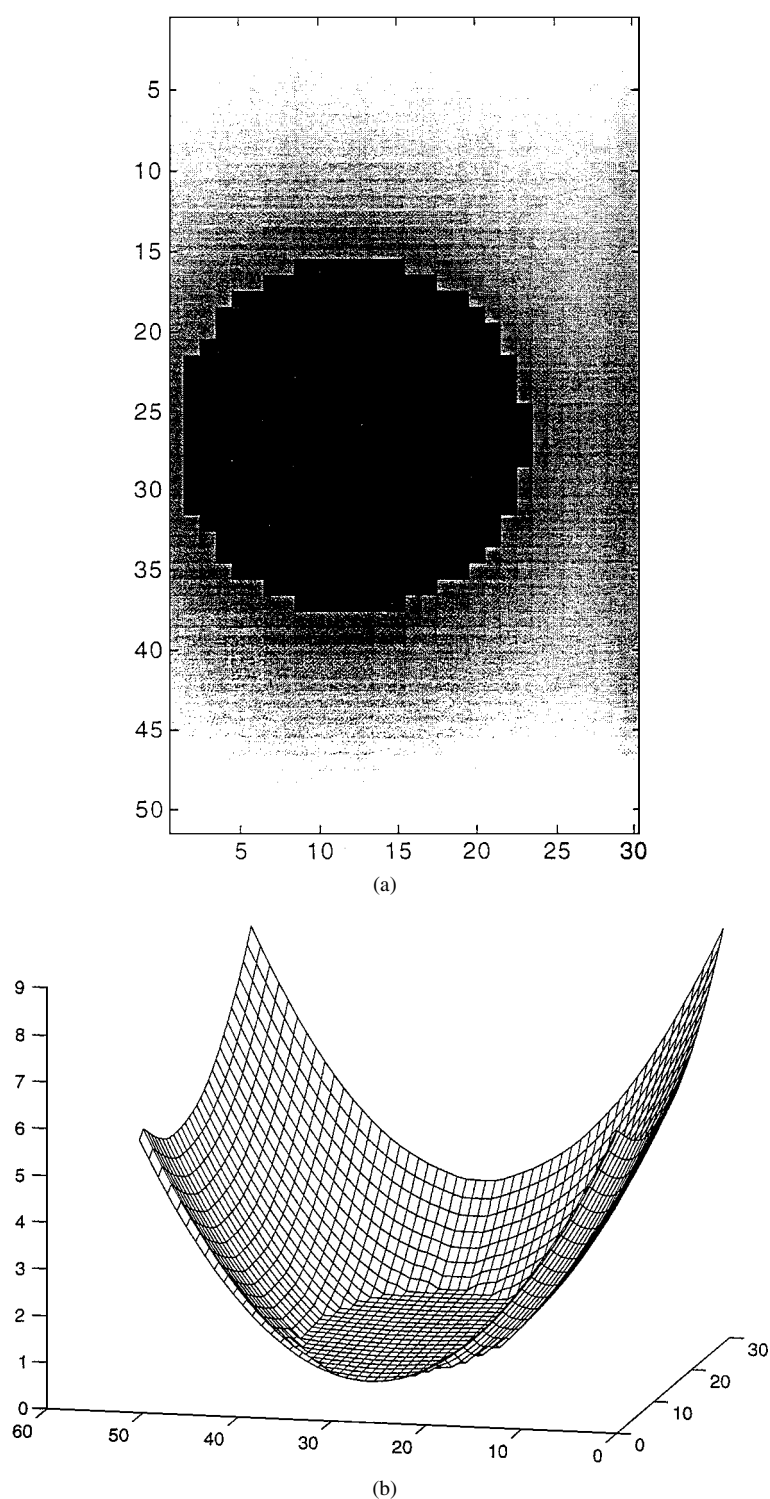


Figure 1. Reconstruction of two intersecting convex paraboloids where $\Delta x = \Delta y = 0.01$. In Fig. 1(a), we show the $n(x, y)$ function determined from the photograph. The horizontal axis represents the values of x between 0 and 0.3; the vertical axis represents the values of y between 0.25 (at the top) and 0.75 (at the bottom). In Fig. 1(b), we show the reconstruction of the paraboloids over the same region. The x -axis goes into the plane of the page, the y -axis is horizontal, and the value of $u(x, y)$ is on the vertical axis.

$(-0.9, 0.5)$. This means that characteristics propagate through the section of the circle where $x > \frac{18}{205}$, but, more importantly, on the section of the circle where $x < \frac{18}{205}$, characteristics propagate out of both sides of the circle (see Fig. 2), which can never happen in a viscosity solution where n is continuous (due to the parallels with entropy solutions in conservation laws). In an Ishii/Tourin viscosity solution, not only is it impossible for characteristics to emanate from both sides of a discontinuity in n , but it is also impossible to have characteristics propagate through discontinuities in the manner seen for $x > \frac{18}{205}$, since this also leads to downward kinks in the solution.

Interesting behavior also occurs in the error between u and its approximation u^ρ . Souganidis (1985) established that—when n is continuous—the maximum (L^∞) error between u and u^ρ is, at worst, of order $\rho^{\frac{1}{2}}$; experimentally, order ρ behavior is observed (e.g., remark 3 in Rouy and Tourin (1992)). For the discon-




Figure 2. Characteristic behavior in the solution. Note that characteristics are created on both sides of the circle when $x < \frac{18}{205}$, which cannot happen at shocks in viscosity solutions when n is continuous. The characteristics pass through the circle when $x \geq \frac{18}{205}$.

tinuous n here, however, we observe different convergence behavior. For larger meshes ($\rho > 0.00116$), the maximum error occurs near the points $(1, 0)$ or $(1, 1)$, where the value of the reconstructed solution is affected by characteristics that propagate *through* the discontinuity and the error appears to be first order (see Fig. 3(a)). For smaller meshes ($\rho < 0.00116$), however, the maximum error occurs near the points $(0, 0)$ or $(0, 1)$, where the value of the reconstructed solution is affected by characteristics *created* at the discontinuity (i.e., coming from the $x < \frac{18}{205}$ region of the circle) and the order of the error appears to be lower than $\frac{1}{2}$ (see Fig. 3(b)).

We next consider cases where there are potential problems in using the (extended) viscosity solution—and therefore its corresponding numerical reconstruction—to describe the actual surface. One possible problem is that the (extended) viscosity solution may not be continuous at some critical or boundary points. Because the solution is built upward from $g(x, y)$, the given boundary and critical point data, if the reconstructed (extended) viscosity surface is discontinuous at a point (x_0, y_0) , then $g(x_0, y_0) > \lim_{(x,y) \rightarrow (x_0,y_0)} \lim_{\rho \rightarrow 0} u^\rho(x, y)$, i.e., g must always be *above* the (extended) viscosity solution when there is a discontinuity at boundary or critical points. However, one of the advantages of using (extended) viscosity solutions to construct our solution is that the (extended) viscosity solution is always the greatest possible weak solution (i.e., $u_{\text{visc.}} \geq u_{\text{weak}}$ for all (x, y)), which means that if the computed solution is not continuous at some boundary/critical point, then no weak solution can be, and therefore the given boundary and critical point data cannot correspond to *any* continuous surface. Therefore, we can conclude in this case that the actual surface is discontinuous (or not Lambertian) and hence is beyond the scope of surfaces considered in this paper.

Sometimes the (extended) viscosity solution is continuous but still does not correspond to the actual surface, even if the actual surface is also continuous. For instance, if our surface in the first example had been

$$u(x, y) = \max(81[(x - 0.1)^2 + (y - 0.5)^2], [(x + 0.9)^2 + (y - 0.5)^2 - 0.01]) \quad (6)$$

—so the second paraboloid is now translated 0.01 units downward—we would have locations in the interior of the domain where the surface is lower than all the

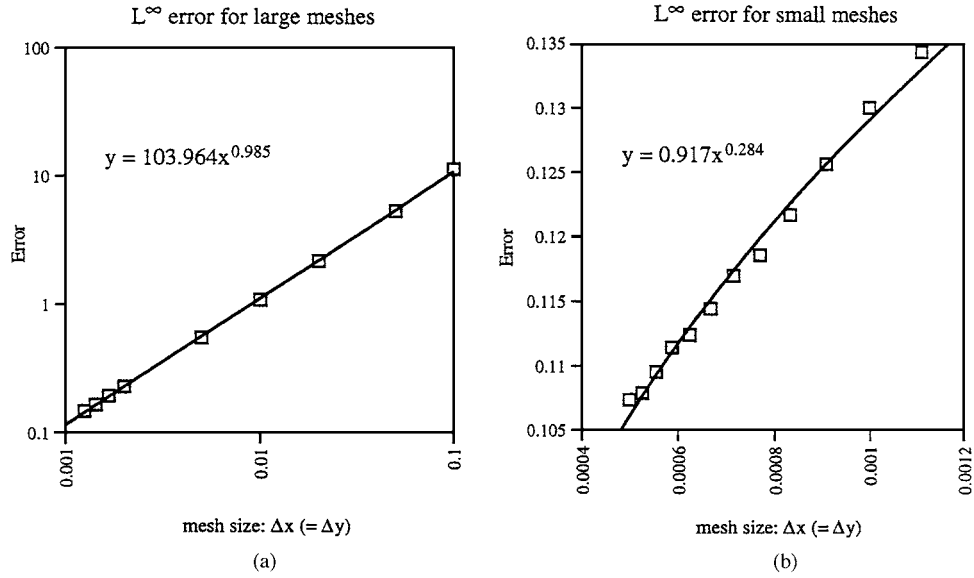


Figure 3. Order of error in the numerical simulation. In Fig. 3(a), we show the L^∞ error of the numerical solution for large meshes, which conforms with the order 1 error observed when n is continuous. In Fig. 3(b), we show the L^∞ error for smaller meshes, and the observed order of the error, 0.284, is below the theoretical bound of 0.5 for continuous n .

boundary and critical points. Since, by (5), the solution is built upwards from g , the boundary and critical point data, the (extended) viscosity solution and its corresponding numerical reconstruction cannot possibly correspond to the actual surface. (This phenomenon is easier to visualize for the following 1-D case: if $u = |x|$ and $\Omega = (-1, 1)$ then $n = 1$ and g only tells us that $u(-1) = u(1) = 1$, so, from the 1-D analogue of (3), the viscosity solution must be $2 - |x|$, which is the actual surface flipped over.)

How can we determine when this has happened? We recommend the following rule: *reject (or at least be extremely suspicious of) any (extended) viscosity solution containing gradient discontinuities that do not correspond to discontinuities in n .* Real surfaces that exhibit this type of behavior are *extremely* unlikely since they would require that at *every* point where the gradient is discontinuous and n is continuous, the angle between the unit normal to the surface on one side of the discontinuity and the z axis would have to be *exactly* the same as the angle for the other side of the discontinuity. When the (extended) viscosity solution exhibits this behavior, as it will for our example in (6), it is almost always indicative that the (extended) viscosity solution is not the same as the actual surface. (Note that this recommendation

implies that when n is continuous, only *smooth* viscosity solutions, that is, classical solutions, should be accepted.)

When (extended) viscosity solutions exhibit this behavior, it is sometimes possible to remedy the problem. Consider the n function in Fig. 4(a), which corresponds to our second example, the surface

$$u(x, y) = \max \begin{pmatrix} 1 - 6 \left(x - \frac{1}{4}\right)^2 - 3 \left(y - \frac{3}{4}\right)^2, \\ \frac{3}{2} - 6 \left(x - \frac{3}{4}\right)^2 - 6 \left(y - \frac{1}{4}\right)^2, \\ \frac{1}{2}x + \frac{1}{2}y \end{pmatrix}.$$

The reconstruction of the extended viscosity solution using g —which specifies u on the domain's perimeter and at the two critical points, $(0.25, 0.75)$ and $(0.75, 0.25)$ —and n leads to the picture in Fig. 4(b). The extended viscosity solution represented in this picture does not correspond to u , the original surface, and—as we expected from our rule—it also contains a gradient discontinuity where n is continuous between the two critical points. However, by using $-u^\rho$ where u^ρ is the solution corresponding to $-g$ and n , we get a reconstruction that *does* correspond to the original





Figure 4. Reconstruction of two concave paraboloids intersecting with a plane where $\Delta x = \Delta y = 0.01$. In Fig. 1(a), we show the $n(x, y)$ function determined from the photograph. The horizontal axis represents the values of x between 0 and 1; the vertical axis represents the values of y between 0 (at the top) and 1 (at the bottom). In Fig. 4(b), we show the reconstruction from Eq. (1), which contains a false crease between the tops of the two paraboloids. Note that n is continuous where the crease occurs. The x -axis appears to the right of the y -axis in the figure, and the value of $u(x, y)$ is on the vertical axis. In Fig. 4(c), we show the reconstruction from Eq. (6), the negation of Eq. (1). Here, the crease is “removed” and the reconstruction corresponds to the original surface.

(Continued on next page.)




Figure 4. (Continued).

surface (see Fig. 4(c)). This same trick also produces the original surface given in (6).

What we have actually done is compute the (extended) viscosity solution corresponding to

$$-\sqrt{u_x^2 + u_y^2} + n(x, y) = 0, \quad (7)$$

the negation of (1). For continuous n , the viscosity solution for (7) can only contain kinks pointing *downwards* (as opposed to (1) where the kinks must point *upwards*), which can potentially lead to a different solution both when n is continuous or discontinuous. From the control theory (or computational) perspective, we are now starting at the boundary/critical points and building our solution downwards as opposed to upwards. We note that this only works sometimes: if, for example, both the maximum height of the surface is greater than the maximum of g and the minimum height of the surface is less than the minimum of g , then neither the unique (extended) viscosity solution to (1) nor the unique (extended) viscosity solution to (7) will correspond to the actual surface. In this case there is no clear recipe for finding the actual surface without additional information.

Acknowledgment

I would like to acknowledge the financial support of the National Science Foundation under Grant No. DMS-9704864.

References

- Barles, G. and Souganidis, P.E. 1991. Convergence of approximation schemes for fully nonlinear second order equations. *Asymptotic Anal.*, 4:271–283.
- Bruckstein, A.M. 1988. On shape from shading. *Comput. Vision Graphics Image Process.*, 44:139–154.
- Bruss, A.R. 1982. The Eikonal equation: Some results applicable to computer vision. *J. Math. Phys.*, 23:890–896.
- Crandall, M.G., Evans, L.C., and Lions, P.L. 1984. Some properties of viscosity solutions of Hamilton-Jacobi equations. *Trans. Amer. Math. Soc.*, 282:487–502.
- Crandall, M.G. and Lions, P.L. 1983. Viscosity solutions of Hamilton-Jacobi equations. *Trans. Amer. Math. Soc.*, 277:1–42.
- Dupuis, P.G. and Boué, M. 1999. Markov chain approximations for deterministic control problems with affine dynamics and quadratic cost in the control. *SIAM J. Num Anal.*, 36(3):667–695.
- Dupuis, P.G. and Oliensis, J. 1993. An optimal control formulation and related numerical methods for a problem in shape reconstruction. *Annals of Applied Probability*, 4:287–345.
- Eastwood, M. 1995. Some remarks on shape from shading. *Advances in Applied Mathematics*, 16:259–268.
- Frankot, R.T. and Chellappa, R. 1988. A method for enforcing integrability in shape-from-shading algorithms. *IEEE Trans. Pattern Anal. Machine Intell.*, 10:439–451.
- Horn, B.K.P. 1975. Obtaining shapes from shading information. *Psychology of Computer Vision*. McGraw-Hill: New York, pp. 115–155.
- Horn, B.K.P. and Brooks, M.J. (Eds.). 1989. *Shape From Shading*. MIT Press: Cambridge, MA.
- Ishii, H. 1985. Hamilton-Jacobi equations with discontinuous Hamiltonians on arbitrary open subsets. *Bull. Fac. Sci. Engrg. Chuo Univ.*, 28:33–77.
- Ishii, H. 1987. A simple, direct proof of uniqueness for solutions of the Hamilton-Jacobi equations of Eikonal type. *Proc. Amer. Math. Soc.*, 100(2):247–251.

- Kimmel, R. 1992. Shape from shading via level sets. M.Sc. Thesis, Technion-Israel Institute of Technology, June.
- Kimmel, R. and Bruckstein, A.M. 1995a. Tracking level sets by level sets: A method for solving the shape from shading problem. *Comput. Vision Image Understanding*, 62(2):47–58.
- Kimmel, R. and Bruckstein, A.M. 1995b. Global shape from shading. *Comput. Vision Image Understanding*, 62(3):360–369.
- Kimmel, R., Siddiqi, K., Kimia, B.B., and Bruckstein, A.M. 1995. Shape from shading: Level set propagation and viscosity solutions. *Int. J. Comput. Vision*, 16:107–133.
- Kushner, H.J. and Dupuis, P.G. 1992. *Numerical Methods for Stochastic Control Problems in Continuous Time*. Springer-Verlag: New York.
- Lions, P.L. 1982. *Generalized Solutions of Hamilton-Jacobi Equations*, Pitman Research Notes Series, Pitman: London.
- Lions, P.L., Rouy, E., and Tourin, A. 1993. Shape-from-shading, viscosity solutions, and edges. *Numer. Math.*, 64:323–353.
- Ostrov, D.N. 2000. Extending viscosity solutions to eikonal equations with discontinuous spatial dependence. *Nonlinear Anal., Ser. A: Theory Methods*, 42(4):709–736.
- Rouy, E. and Tourin, A. 1992. A viscosity solutions approach to shape-from-shading. *SIAM J. Numer. Anal.*, 29(3):867–884.
- Saxberg, B.V.H. 1992. *Int. J. Robot. Res.*, 11(3):202–224.
- Sethian, J.A. 1996. A fast marching level set method for monotonically advancing fronts. In *Proc. Natl. Acad. Sci.*, 93:1591–1595.
- Souganidis, P.E. 1985. Approximation schemes for viscosity solutions of Hamilton-Jacobi equations. *J. Differential Equations*, 59:1–43.
- Tourin, A. 1992. A comparison theorem for a piecewise Lipschitz continuous Hamiltonian and application to shape-from-shading problems. *Numer. Math.*, 62:75–85.
- Tsitsiklis, J.N. 1995. Efficient algorithms for globally optimal trajectories. *IEEE Trans. on Automatic Control*, 40(9): 1528–1538.