

Fast and Robust Surface Normal Integration by a Discrete Eikonal Equation

Silvano Galliani
galliani@mia.uni-saarland.de

Michael Breuß
breuss@tu-cottbus.de

Yong Chul Ju
y.ju@mmci.uni-saarland.de

Mathematical Image Analysis Group
Saarland University, Germany

Institute for Applied Mathematics and Scientific
Computing, BTU Cottbus, Germany

Vision and Image Processing Group
Saarland University, Germany

Abstract

The integration of surface normals is a classic and fundamental task in computer vision. In this paper we deal with a highly efficient fast marching (FM) method to perform the integration. In doing this we build upon a previous work of Ho and his coauthors. Their FM scheme is based on an analytic model that incorporates the eikonal equation. Our method is also built upon this equation, but it makes use of a complete discrete formulation for constructing the FM integrator (DEFM). We not only provide a theoretical justification of the proposed method, but also illustrate at hand of a simple example that our approach is much better suited to the task. Several more sophisticated tests confirm the robustness and higher accuracy of the DEFM model. Moreover, we present an extension of DEFM that allows to integrate surface normals over non-trivial domains, *e.g.* featuring holes. Numerical results confirm desirable qualities of this method.

1 Introduction

The integration of normal vectors of a surface is a fundamental problem in computer vision. Important classic examples for techniques where this task arises are shape from shading and photometric stereo. There it is used to recover the three-dimensional (3-D) shape of objects from one or more input images, respectively [1, 2]. Several methods have been developed to solve the problem of surface normal integration over the last decades. As milestones let us mention here the method of Horn and Brooks based on the calculus of variations [3], the method of Frankot and Chellappa which employs the frequency domain [4] and the direct integration scheme of Wu and Li [5]. For an extensive review, we refer to [6].

While a variety of schemes has been proposed, there is, however, still a need for methods that combine accuracy, robustness and high efficiency. By the latter objective, an algorithmic approach that appears to be a natural candidate is the fast marching (FM) method, cf. [7, 8, 9, 10, 11]. Given that the task at hand can be formulated as a static eikonal-type equation, it can be solved with a complexity of $\mathcal{O}(N \log N)$, where N is the number of pixels of the computational domain. In the work of Ho *et al.* [12] this strategy has been adopted. Their approach is based on an analytic formulation of the integration task in terms of an eikonal

equation. However, while in [1] some promising results are presented, the authors also report significant problems with the robustness and accuracy of the scheme.

Let us emphasise that the discretisation of a hyperbolic partial differential equation (PDE) such as the eikonal equation is a delicate task, see *e.g.* [14] for a useful discussion in the context of this work. A mixture of discrete and analytic expressions as employed by Ho *et al.* [1] might induce numerical instability such as the amplification of discretisation errors, which would fit to observations in [7].

Our Contributions. In this paper we improve the scheme of Ho *et al.* by making it stable and accurate in the first step. We propose a complete discrete formulation of the model by Ho *et al.* in terms of a proper approximation of the underlying PDE. At hand of a simple example we explain why the resulting new algorithm DEFM is much better suited for the integration task than the method in [1]. We also give a theoretical analysis showing the difference between the schemes. By additional experiments we confirm the robustness and accuracy gain of DEFM compared to the method of Ho *et al.* Moreover, we present an extension of DEFM that enables to deal with more general, connected computational domains featuring *e.g.* holes. Our extended method relies on the use of a pre-computed geodesic distance as a metric on the computational domain.

Paper organisation Section 2 gives a brief review of the work by Ho *et al.* We give an exposition on the proposed approach by providing an example and detailed analysis in Section 3. Section 4 validates our approach by reporting experimental results. Afterwards, the paper is closed by concluding remarks.

2 The Previous Method

We briefly review the method by Ho *et al.* [1] upon which our approach is based. The basic idea relies on solving an eikonal equation using the FM method under certain assumptions.

Let us begin with looking at an example of an integration problem in the one-dimensional (1D) case. In this instance, by the fundamental theorem of calculus an integration of v' gives an antiderivative v up to an additive constant c . Therefore, we attain

$$\int v'(x_1) dx_1 = v(x_1) + c. \quad (1)$$

In analogy to the 1D case the solution structure of (1) can be respected in 2D as well. Therefore in order to obtain v in 2D, Ho *et al.* consider the *auxiliary formulation*

$$w(x_1, x_2) := v(x_1, x_2) + \lambda f(x_1, x_2), \quad (2)$$

where $\lambda > 0$ is a constant parameter and f denotes a function. Here, one can note that the role of the constant c in (1) corresponds to the term λf in (2). Moreover, from a numerical viewpoint, in order to apply FM the information of the critical points of w should be available, *e.g.* that w admits only a single minimum at some initial point \mathbf{x}_0 . In this regard, the function f is designed in such a way that it should not modify any important structures of v , yet give us some useful information at the same time. In [1], the point \mathbf{x}_0 is set at the origin and as a function fulfilling the discussed requirements

$$f_{\text{ho}} := x_1^2 + x_2^2 \quad (3)$$

is chosen. Since the FM method is originally devised to solve the eikonal equation [4], for the deployment of FM as a numerical solver the expression in (2) should be put into an eikonal-type format. This leads to

$$|\nabla w| = |\nabla v + \lambda \nabla f_{\text{Ho}}| = \sqrt{(v_{x_1} + \lambda 2x_1)^2 + (v_{x_2} + \lambda 2x_2)^2} \quad (4)$$

with $v_{x_1} := \frac{\partial v}{\partial x_1}$ and $v_{x_2} := \frac{\partial v}{\partial x_2}$. Let us emphasise that Ho *et al.* make use of an analytic computation of ∇f_{Ho} in the above equation. Since all elements on the right hands side of (4) are known, the FM method allows to compute w from the PDE $|\nabla w| = |\nabla v + \lambda \nabla f_{\text{Ho}}|$. In order to obtain v from w , one has to subtract the known function f from w after that computation.

3 Our Approach

3.1 Upwind Discretisation

Our first advancement stems from the deployment of a proper discretisation method for the eikonal equation (4), namely an *upwind scheme* [5]. One important reason for this choice is directly related to the nature of an eikonal equation. It is a special case of a Hamilton-Jacobi PDE, and for this type of equations it is well-known that in general there does not exist a solution in the classic sense [6]. In fact, a solution we are looking for is defined in the viscosity sense [6]. Another reason comes from the nature of the FM method. Since FM is originally designed for solving an eikonal equation and the method itself already embodies an upwind strategy [4], a solution computed by FM relies also upon non-smooth solutions in the viscosity sense [5]. In the latter context, there is no theoretical justification for employing an analytic derivative such as ∇f_{Ho} in (4) for FM.

Thus, instead of the analytic expression ∇f_{Ho} in (4) we employ a complete discrete formulation using as discrete derivative the upwind difference from [5]. In 1D, this upwind discretisation reads as

$$\hat{f}_x := \max(D^- f, -D^+ f, 0) \quad (5)$$

with

$$D^- f = \frac{f_i - f_{i-1}}{\Delta x} > 0 \quad \text{and} \quad D^+ f = \frac{f_{i+1} - f_i}{\Delta x} < 0 \quad (6)$$

where Δx is the mesh width and f_j denotes the discrete function value at $j\Delta x$ with j being the index of the grid point. Each inequality in (6) holds for consistency since the upwind scheme chooses only one direction for the propagation of the information. This upwind formula can also be applied in a straightforward way in 2D.

In order to make the advantage of the proposed method over the method of Ho *et al.* clearly visible, in what follows we provide a simple example which shows that upwind differencing outperforms the analytic derivative.

3.1.1 Example

We consider the 1D version of the eikonal equation (4) solving it by FM with different discretisation methods. To this end, let us assume that we are given a three-pixel complete black input image, i.e. whose grey values are zeros. Furthermore, for simplicity suppose that the middle pixel is positioned at the origin and is a seed point for the FM method, and we let the solution value be 0 there.

For all the computations, we employ $\lambda = 1$, grid size $\Delta x = 1$ and Neumann boundary conditions in case if boundary values are needed. This setting is equivalent to

$$x = [-1, 0, 1], \quad v_x = [0, 0, 0] \quad \text{and} \quad v = [0, 0, 0] \quad \text{and} \quad w = [w_{-1}, w_0 = 0, w_1], \quad (7)$$

where w_{-1} and w_1 denote the evaluation of w at $x = -1$ and $x = 1$, respectively. In addition, they are the unknowns that we want to find.

Let us stress that we expect from a proper solution method that it gives the true solution in such a simple toy example. For the comparison of methods let us construct it. Since all input grey values are identical in our toy example (and because we use Neumann boundary conditions), the derivatives v_x are all zero, so that by the given value $v = 0$ at the central point we have that v is identical to the zero function $v = 0$ at all grid points. By $\Delta x = 1$ and $\lambda = 1$, we evaluate $f := f(x) = x^2$ (centred upon the middle pixel) at the two grid points of interest as $f(\pm 1) = 1$, so that we expect by $w = v + \lambda f$ that the exact solution of our toy problem is given by the vector $w = [1, 0, 1]$.

Analytic method by Ho *et al.* First, we perform the procedure for w_{-1} in (7). In consideration of the 1D version of (3), we receive

$$f_{\text{ID Ho}} = x^2 = [1, 0, 1] \quad (8)$$

and

$$[f_{\text{ID Ho}}]_x = 2x = [-2, 0, 2]. \quad (9)$$

Therefore, by (4) and (9) we attain

$$|w_x|_{(x=-1)} \stackrel{(4)}{=} [f_{\text{ID Ho}}]_x|_{(x=-1)} \stackrel{(9)}{=} |-2| \Rightarrow |w_x|_{(x=-1)} = 2, \quad (10)$$

where $|w_x|_{(x=-1)}$ denotes the absolute value of the evaluation for the x -derivative of w at $x = -1$. The next step is to apply FM [14] at (10). This leads to

$$\sqrt{\max\left(\frac{w_{-1} - w_{-2}}{\Delta x}, -\frac{w_0 - w_{-1}}{\Delta x}, 0\right)^2} = \sqrt{(w_{-1})^2} = 2. \quad (11)$$

As a result, we have $w_{-1} = 2$, the case $w_{-1} = -2$ is automatically ruled out. For w_1 the same procedure is conducted. Finally, we obtain

$$w_{\text{ID Ho}} = [2, 0, 2]. \quad (12)$$

Proposed upwind discretisation. Now, we perform the same computation making use of upwind differencing instead of the analytic derivative. For the position at $x = -1$, by (5) we can obtain

$$\hat{f}_x|_{(x=-1)} = \max\left(\frac{f_{-1} - f_{-2}}{\Delta x}, -\frac{f_0 - f_{-1}}{\Delta x}, 0\right) = \max\left(\frac{1-1}{1}, -\frac{0-1}{1}, 0\right) = 1. \quad (13)$$

Analogously, for $x = 1$ we have $\hat{f}_x|_{x=1} = 1$. Thus, we have the upwind difference

$$\hat{f}_x = [1, 0, 1]. \quad (14)$$

Utilising FM as in the analytic case, we attain

$$w_{\text{ID upwind}} = [1, 0, 1]. \quad (15)$$

Evaluation. In summary, it can be recognised that our method achieves the desired result, while there is even in this simple toy example a significant discrepancy between exact solution and computed result when using the method of Ho *et al.*, see (12) and (15).

3.1.2 Scheme Analysis

One of the most desirable properties for an algorithmic approach is the robustness in that the computed result should not be highly sensitive to changes of input data, *e.g.* due to noise. A fundamental stability property implying robustness is the *monotonicity* of numerical schemes which we now investigate, cf. [□]. For shortness of presentation, we restrict our attention again to the 1D situation. Our investigation extends in a straightforward way to 2D.

Our method. Due to the fact that v_x is a pointwise given, precomputed quantity that does not influence the performance of the scheme, we deal without loss of generality with the following eikonal equation

$$\lambda |f_x| = 1. \quad (16)$$

Following [□] we consider the corresponding time-dependent formulation of (16) as

$$\lambda f_t + \lambda |f_x| = 0. \quad (17)$$

By applying Euler forward time discretisation and upwind spatial discretisation to (17), we have

$$\lambda \frac{f_i^{n+1} - f_i^n}{\Delta t} + \lambda \sqrt{\hat{f}_x^2} = 0 \Rightarrow \lambda f_i^{n+1} = \lambda f_i^n - \Delta t \lambda \sqrt{\hat{f}_x^2} =: \mathcal{H}, \quad (18)$$

where \hat{f}_x denotes the upwind discretisation given in (5).

Since monotonicity of a scheme means that the scheme update function \mathcal{H} in (18) is non-decreasing for all arguments [□, □], we need to show that it holds for all i

$$(i) \quad \frac{\partial \mathcal{H}}{\partial f_{i-1}} \geq 0, \quad (ii) \quad \frac{\partial \mathcal{H}}{\partial f_i} \geq 0, \quad (iii) \quad \frac{\partial \mathcal{H}}{\partial f_{i+1}} \geq 0. \quad (19)$$

For this purpose, we make case distinctions between the first two arguments in (5).

When $D^- f$ is chosen. Since $\frac{\partial \mathcal{H}}{\partial f_{i+1}} = 0$ in this case, we only have two cases. Then, we can receive

$$\frac{\partial \mathcal{H}}{\partial f_{i-1}} = \frac{\partial \mathcal{H}}{\partial \hat{f}_x} \frac{\partial \hat{f}_x}{\partial f_{i-1}} = \frac{\Delta t}{\Delta x} > 0 \quad (20)$$

and

$$\frac{\partial \mathcal{H}}{\partial f_i} = \lambda - \frac{\partial \mathcal{H}}{\partial \hat{f}_x} \frac{\partial \hat{f}_x}{\partial f_i} = \lambda - \frac{\Delta t}{\Delta x} \geq 0 \Rightarrow \lambda \geq \frac{\Delta t}{\Delta x} > 0. \quad (21)$$

The condition (21) states that the scheme is monotone when $\lambda \geq \frac{\Delta t}{\Delta x}$. Let us stress that Δt is an artificial parameter for which we may assume $0 < \Delta t \ll \Delta x$.

When $D^+ f$ is chosen. We apply the same procedure as previously. First, we have $\frac{\partial \mathcal{H}}{\partial f_{i-1}} = 0$. Then, we can obtain

$$\frac{\partial \mathcal{H}}{\partial f_{i+1}} = \frac{\partial \mathcal{H}}{\partial \hat{f}_x} \frac{\partial \hat{f}_x}{\partial f_{i+1}} = \frac{\Delta t}{\Delta x} > 0 \quad (22)$$

and

$$\frac{\partial \mathcal{H}}{\partial f_i} = \lambda - \frac{\partial \mathcal{H}}{\partial \hat{f}_x} \frac{\partial \hat{f}_x}{\partial f_i} = \lambda - \frac{\Delta t}{\Delta x} \geq 0 \Rightarrow \lambda \geq \frac{\Delta t}{\Delta x} > 0. \quad (23)$$

As a consequence, we can conclude that our proposed scheme is monotone in view of (19). Moreover, as is shown in (21) and (23) this investigation also tells us that our method leaves no restrictions for the choice of λ in contrast to the findings of Ho *et al.* in [10] because λ can be freely set as any predefined small positive values, see Figure 1.

Non-monotonicity of the analytic method. Here, we show the numerical instabilities of the method by Ho *et al.* [10] by presenting the non-monotonicity of the analytic method. The analytic expression corresponds to a central difference because $[f_{\text{IDHo}}]_x$ is identical to

$$Df = \frac{f_{i+1} - f_{i-1}}{2\Delta x} = 2x =: \hat{f}_x. \quad (24)$$

Plugging (24) into (18) and applying the criterion in (19) gives us

$$(a) \quad \frac{\partial \mathcal{H}}{\partial f_{i-1}} = 1 > 0, \quad (b) \quad \frac{\partial \mathcal{H}}{\partial \hat{f}_x} \frac{\partial \hat{f}_x}{\partial f_i} = -\frac{\Delta t}{2\Delta x} \frac{\hat{f}_x}{|\hat{f}_x|}, \quad (c) \quad \frac{\partial \mathcal{H}}{\partial \hat{f}_x} \frac{\partial \hat{f}_x}{\partial f_{i+1}} = \frac{\Delta t}{2\Delta x} \frac{\hat{f}_x}{|\hat{f}_x|}. \quad (25)$$

In the light of (b) and (c) in (25), it is clear that monotonicity does not hold in this case. Let us also mention that it is well known that central differences as in (24) yield unstable schemes when discretising hyperbolic PDEs, cf. [10].

3.2 Scheme Extension for Handling Non-Trivial Domains

It can be clearly noticed that the metric function f in (2) plays a major role, especially when λ is large. However, the usage of the Euclidean norm (3) leads to difficulties when the topology of the domain for an integration is qualitatively different from that of a disc, *e.g.* as in the case of a mask shape with two holes inside, see Figure 3(e) and 4(e). By these figures, one can realise that the Euclidean norm fails in measuring a meaningful distance from the image centre to places behind the two holes. In such a case, we inevitably should detour around the holes and at the same time still have to find a suitable distance from the centre. In order to take care of this matter, instead of the Euclidean distance we incorporate as indicated a more general *geodesic distance* which can handle this situation.

4 Experimental Results

We evaluate the performance of the proposed method by carrying out numerical tests on both synthetic and real data.


Experiment I: Monkey saddle function. In order to test the stability of the proposed scheme, our first numerical experiment is carried out with the monkey saddle function,

$$Z = x(x^2 - 3y^2) + 20, \quad (26)$$


see Figure 1(a). As in the work by Ho *et al.*, for the comparison purpose we also use the same relative error measure

$$\varepsilon = \frac{|Z^* - Z|}{|Z|}, \quad (27)$$


where Z^* denotes the computed solution. For different λ values, the computation results of mean, median and standard deviation of the relative error measure (27) are reported in Figure 1(b), 1(c) and 1(d), respectively. From the plots, one can clearly notice that for the




(a) The input monkey saddle function (26).



(b) Mean of the relative error (27).



(c) Median of the relative error (27).



(d) Standard deviation of the relative error (27).


Figure 1: The input monkey saddle function described in (26) and plots of mean, median and standard deviation of the relative error measurements (27) depending on the value of λ . The value of λ varies in the range $0 < \lambda \leq 100$ and is represented by the horizontal axes.

whole range of λ the proposed scheme is highly stable comparing to the method by Ho *et al.*, in all error measures proposed in [4]. Furthermore, this also confirms the results of our theoretical investigation presented in Section 3.1.2.

Experiment II: Lena. For a visual comparison of schemes, we perform the method of Ho *et al.* and our integration method on x - and y -derivatives of the test image “Lena”, see Figure 2(a), 2(b) and 2(c). The corresponding reconstruction results by each method are presented in Figure 2(d) and 2(e), respectively. The quality difference between Figure 2(d) and 2(e) is conspicuous. Notice also the big difference of λ between two experiments, $\lambda_{\text{Ho}} = 0.2$ versus $\lambda_{\text{Proposed}} = 1000000$. Relying on our analysis in Section 3.1.2, we chose $\lambda_{\text{Proposed}}$ in a generic way and in a range that may be disadvantageous for the scheme accuracy. Since the method by Ho *et al.* does not give a reasonable result for large λ values, we only apply very small values of λ and present the best result. Nevertheless, the proposed method outperforms the method of Ho *et al.* in all error measures, see Table 1.

Table 1: Error measurements for Lena experiment given in Figure 2.

	Mean	Median	Standard deviation
Ho <i>et al.</i> ($\lambda = 0.2$)	0.3060	0.2079	0.3604
Our method ($\lambda = 1000000$)	0.0785	0.0364	0.1325



(a) The original Lena image.

(b) x -derivative of Figure 2(a).(c) y -derivative of Figure 2(a).(d) Optimal result by the scheme of Ho *et al.*

(e) Generic result by our method.

Figure 2: **Top row.** The original Lena image and its x - and y -derivatives used as input images. **Bottom row.** Reconstruction results by each method.

Experiment III: Buddha. We now apply the proposed method at the reconstruction of a 3D surface by making use of the photometric stereo (PS) algorithm based on the work by Tankus and Kiryati [16]. To this end, we have used three input images of the same object with different illumination conditions, see Figure 3(a), 3(b) and 3(c). In order to validate if the proposed geodesic metric works properly for integration of surface normals delivered by PS, we have done two experiments i.e. for the integration domain with a mask and without a mask. The mask has two holes as often used with basic PS methods such as the one we employ here to avoid effects of the complex eye regions. The rendering results of the 3D reconstruction depending on the integration domain are shown in Figure 3(d) and 3(e), respectively. It can be clearly recognised that the proposed geodesic metric can handle the topological changes in the integration domain accordingly.

Experiment IV: Beethoven. Finally, we test our algorithm on the real-world data set “Beethoven” shown in Figure 4(a), 4(b) and 4(c). As the case with “Buddha”, we perform the 3D reconstruction from three different input images by using first the aforementioned PS technique for computing a field of surface normals on the integration domain with a mask

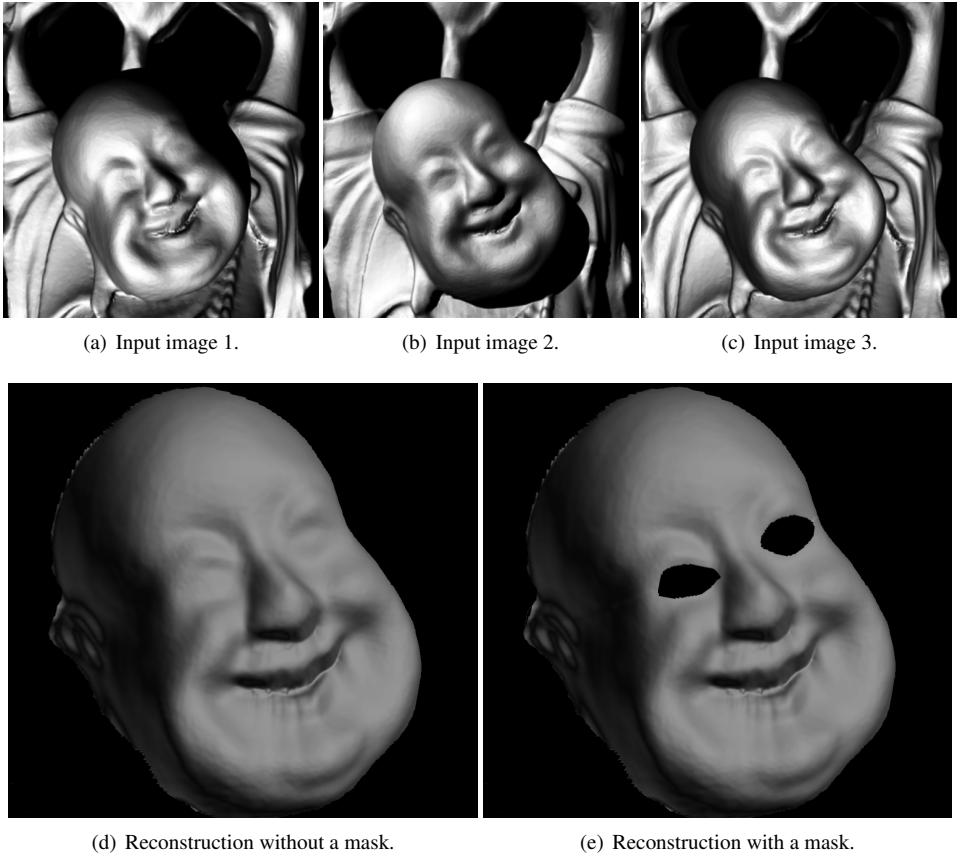



Figure 3: **Top row.** The three input images of Buddha with different illumination conditions. **Bottom row.** Renderings of the Buddha face after 3D surface integration.

and without a mask, respectively. The outcomes after surface normal integration with the extended DEFM method are exhibited in Figure 4(d) and 4(e). As can be observed manifestly, the proposed method works appropriately with realistic data as well.

5 Conclusion

We have constructed and validated a very fast, highly robust and reasonably accurate numerical scheme for surface integration by employing a discrete eikonal equation, based on a previous method of Ho *et al.* We extended our proposed method by employing a geodesic metric in a way that the new algorithm can handle topological changes in integration domains. Since surface normal integration methods are of considerable importance and have many potential applications in computer vision, we believe that our new method can be an attractive tool for applications in the field.

Acknowledgements. This work has been partly funded by the Fraunhofer Institute for Industrial Mathematics (ITWM). Moreover, Yong Chul Ju gratefully acknowledges the Cluster of Excellence “Multimodal Computing and Interaction” within the Excellence Initiative of the German Federal Government.



(a) Input image 1.

(b) Input image 2.

(c) Input image 3.

(d) Reconstruction without a mask.

(e) Reconstruction with a mask.

Figure 4: **Top row.** The three input images of Beethoven bust with different illumination conditions. **Bottom row.** Renderings of 3D surface integration results.

References

- [1] G. Aubert and P. Kornprobst. *Mathematical Problems in Image Processing: Partial Differential Equations and the Calculus of Variations*, volume 147. Springer, 2nd edition, 2006.
- [2] M.G. Crandall and P.-L. Lions. Viscosity solutions of Hamilton-Jacobi equations. *Transactions of the American Mathematical Society*, 277(1):1–42, 1983.
- [3] M.G. Crandall and A. Majda. Monotone difference approximations for scalar conservation laws. *Mathematics of Computation*, 34(149):1–21, 1980.
- [4] J.-D. Durou, J.-F. Aujol, and F. Courteille. Integrating the normal field of a surface in the presence of discontinuities. In *Energy Minimization Methods in Computer Vision and Pattern Recognition (EMMCVPR)*, volume 5681, pages 261–273, 2009.
- [5] R.T. Frankot and R. Chellappa. A method for enforcing integrability in shape from

- shading algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(4):439–451, 1988.
- [6] J.J. Helmsen, E.G. Puckett, P. Colella, and M. Dorr. Two new methods for simulating photolithography development in 3d. In *Optical Microlithography IX*, volume 2726, pages 253–261. SPIE, 1996.
- [7] J. Ho, J. Lim, M.-H. Yang, and D.J. Kriegman. Integrating surface normal vectors using fast marching method. In *Proc. European Conference on Computer Vision (ECCV)*, pages 239–250, 2006.
- [8] B.K.P. Horn and M.J. Brooks. The variational approach to shape from shading. *Computer Vision, Graphics, and Image Processing*, 33(2):174–208, 1986.
- [9] B.K.P. Horn and M.J. Brooks. *Shape from Shading*. The MIT Press, 1989.
- [10] R.J. LeVeque. *Numerical Methods for Conservation Laws*. Birkhäuser, 1992.
- [11] E. Rouy and A. Tourin. A viscosity solutions approach to shape-from-shading. *SIAM Journal on Numerical Analysis*, 29(3):867–884, June 1992.
- [12] J.A. Sethian. A fast marching level set method for monotonically advancing fronts. In *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, pages 1591–1595, 1995.
- [13] J.A. Sethian. Fast-marching level-set methods for three-dimensional photolithography development. volume 2726, pages 262–272. SPIE, 1996.
- [14] J.A. Sethian. *Level Set Methods and Fast Marching Methods*. Cambridge University Press, 2nd edition, 1999.
- [15] J.A. Sethian. Evolution, implementation, and application of level set and fast marching methods for advancing fronts. *Journal of Computational Physics*, 169(2):503 – 555, 2001.
- [16] A. Tankus and N. Kiryati. Photometric stereo under perspective projection. In *Proc. IEEE International Conference on Computer Vision (ICCV)*, volume 1, pages 611–616, 2005.
- [17] J.N. Tsitsiklis. Efficient algorithms for globally optimal trajectories. *IEEE Transactions on Automatic Control*, 40(9):1528–1538, 1995.
- [18] R.J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980.
- [19] Z. Wu and L. Li. A line-integration based method for depth recovery from surface normals. *Computer Vision, Graphics, and Image Processing*, 43(1):53–66, 1988.