

Fast Shape from Shading for Phong-Type Surfaces

Oliver Vogel, Michael Breuß, Thomas Leichtweis, and Joachim Weickert

Mathematical Image Analysis Group,
Faculty of Mathematics and Computer Science, Building E1.1
Saarland University, 66041 Saarbrücken, Germany.
{vogel,breuss,leichtweis,weickert}@mia.uni-saarland.de

Abstract. Shape from Shading (SfS) is one of the oldest problems in image analysis that is modelled by partial differential equations (PDEs). The goal of SfS is to compute from a single 2-D image a reconstruction of the depicted 3-D scene. To this end, the brightness variation in the image and the knowledge of illumination conditions are used. While the quality of models has reached maturity, there is still the need for efficient numerical methods that enable to compute sophisticated SfS processes for large images in reasonable time. In this paper we address this problem. We consider a so-called Fast Marching (FM) scheme, which is one of the most efficient numerical approaches available. However, the FM scheme is not trivial to use for modern non-linear SfS models. We show how this is done for a recent SfS model incorporating the non-Lambertian reflectance model of Phong. Numerical experiments demonstrate that – without compromising quality – our FM scheme is two orders of magnitude faster than standard methods.

1 Introduction

Given a single 2-D image, the aim of *Shape from Shading (SfS)* is to infer the 3-D depth of the surface of depicted objects. For this, SfS uses the brightness variation in the image together with information on intensity and position of the light source. Much progress has been achieved in the last years in modelling SfS. As proper model components have been identified, SfS is now considered to be a well-posed problem. In recent model extensions, also non-Lambertian surfaces are taken into account within this well-posed framework. Thus, SfS has reached a reasonable level of maturity. However, these advances on the modelling side also lead to new challenges for numerical methods in this field. In order to obtain 3-D reconstructions of good quality it is recommended to use modern, highly non-linear SfS models together with large, high-resolution input images. Thus, a proper algorithm must be able to deal with the arising large non-linear problems in reasonable computing time. In this paper, we show how to use a *Fast Marching (FM)* scheme for this purpose. It turns out that this is not trivial because of the involved non-linearities.

Brief history of SfS models. The SfS-problem is a classic problem in computer vision. It was introduced in the works of Horn [1]. In particular, his model

assumptions of an orthographic camera and Lambertian surface reflectance became a standard for early SfS research, see the review article [2]. However, the authors of [2] also concluded that orthographic SfS models do not perform well on synthetic data, and even worse on real-world images.

In recent years, sophisticated models employing a more realistic perspective projection have been developed [3–5]. In [5], it was shown that the perspective camera model, together with a point light source at the optical centre of the camera and a non-linear light attenuation term, leads to the well-posedness of the SfS-task. Recently, this class of perspective SfS models has been extended to cover also non-Lambertian surface reflectance. In [6], the Lambertian diffuse reflection has been substituted by the model of Oren and Nayar [7] for the purpose of facial recognition. Another approach has been introduced in [8], where the reflectance model of Phong [9] well-known from computer graphics is used.

The Fast Marching method. The SfS models of interest infer the problem to solve boundary value problems for a class of non-linear hyperbolic *partial differential equations (PDEs)* called Hamilton-Jacobi equations. The fast marching (FM) method is an efficient technique for solving such problems. It was introduced by Tsitsiklis [10] and further developed by Sethian [11].

Our contribution. We show how to use the FM technique for the highly nonlinear, perspective SfS model given in [8] which especially incorporates light attenuation and the non-Lambertian reflectance model of Phong.

In particular, we address the following issues. We consider the problem to compute an initial guess of the depth in surface points with the minimal distance to the camera. The estimation is non-trivial for highly non-linear models such as the one we use. For this estimate, a suitable set of corresponding image points needs to be identified in advance. In order to realise the scheme, one also needs to perform in each discretisation point a fixed-point iteration, for which we give a well-working scheme here.

Having solved these problems, we compare the method with other schemes in the Lambertian case, confirming that our FM scheme is two orders of magnitude faster without a trade-off in accuracy. Then we apply the FM scheme directly for Phong-type non-Lambertian SfS of objects in real-world images taken with a standard digital camera. We show that our FM scheme delivers high-quality results in just a few seconds of computing time, while the method from [8] we compare with takes hours for computing comparable results.

Relation to previous work. It is quite well-known that FM schemes may outperform other discretisation methods for the class of problems we are interested in, provided it is possible to construct such a scheme. The potential usefulness of FM schemes has also been noticed by other authors in the field of SfS. The first one who applied FM to the SfS problem was Sethian [11]. The model he considered was the classic orthographic Lambertian model with a single far light source. Later, Kimmel and Sethian used FM for the same set-up but with an oblique light source [12]. In [13], Yuen et al. apply FM at a Lambertian model incorporating a perspective projection. Let us note that this model is formulated in terms of unknown surface normals – in contrast to the unknown

depth as in [5, 3] – and it does not include a light attenuation term. Tankus et al. perform FM at a perspective Lambertian model also not incorporating light attenuation [14]. In [15], Prados and Soatto develop a FM approach based on ideas from optimal control theory. However, while they claim that their approach holds for perspective SfS with Lambertian reflectance, they only show computational results of their scheme for the classic orthographic model also considered in [11]. The paper [16] is an extension of the work [13], addressing problems with strong gradients of the authors’ previous method arising by occluded regions.

As it is of importance in the context of this paper, let us stress that up to now the light attenuation has not been taken consequently into account in FM, and that non-Lambertian reflectance models have not been considered at all within an FM scheme. Note that exactly terms corresponding to these model assumptions yield strong non-linear contributions.

Paper organisation. After briefly introducing the Phong-type SfS model in Section 2, we describe in Section 3 in detail its discretisation of the SfS model, making use of the FM method. We then proceed elaborating on the choice of points featuring the initial guess of the 3-D depth in Section 4. Sections 5 and 6 are devoted to the experimental evaluation and a conclusion, respectively.

2 The Perspective SfS Model with Phong-type Reflectance

The SfS model we deal with in this paper is given in [8]. We briefly review here the developments in that work.

The Phong reflection model. It is adequate to begin the presentation of the SfS model with the modeling ansatz given by the *brightness equation* due to Phong [9]. Assuming thereby the presence of only one light source, it reads as

$$I = k_a I_a + \frac{1}{r^2} (k_d I_d \cos \phi + k_s I_s (\cos \theta)^\alpha) \quad (1)$$

where $I := I(x)$ is the normalised grey value of the image pixel located at $x = (x_1, x_2)^T \in \mathbb{R}^2$, and $r = uf$ is the distance of the surface point from the light source. In (1), the intensities of ambient, diffuse, and specular components of light are denoted by I_a , I_d and I_s , respectively. In analogous notation, the constants k_a , k_d , and k_s with $k_a + k_d + k_s \leq 1$ denote the ratio of ambient, diffuse, and specular reflection.

Discussing the light reflection contributions, the ambient light models a base intensity in the depicted scene, i.e., a basic illumination present everywhere. The diffusely reflected light in each direction is proportional to the cosine of the angle ϕ between surface normal and light source direction. In our scenario, the latter is identical to the direction of the optical centre. The amount of specular light also reflected in this direction is proportional to $(\cos \theta)^\alpha$, where θ is the angle between the ideal (mirror) reflection direction of the incoming light and the optical centre. The number α is a constant depending on the roughness of the material. An ideal mirror reflection can be described via $\alpha \rightarrow \infty$. Note also

that the cosine in the specular term is to be set to zero if it yields negative values.

The SfS model. Plugging in appropriate expressions, the brightness equation (1) yields a nonlinear Hamilton-Jacobi equation. For details of the derivation see [8]. One (usual) important model assumption not mentioned up to now is the visibility of the surface. This means that it is in the front of the optical centre, so that the unknown 3-D depth u is strictly positive. Employing then the change of variables $v := v(x) = \ln(u(x))$, the resulting model is given by

$$JM - k_d I_a \exp(-2v) - \frac{M k_s I_s}{Q} \exp(-2v) \left(\frac{2Q^2}{M^2} - 1 \right)^\alpha = 0, \quad (2)$$

where $J(x) = (I - k_a I_a) f^2 / Q$, $M(x) = \sqrt{f^2 |\nabla v|^2 + (\nabla v \cdot x)^2 + Q^2}$, and $Q(x) = f / \sqrt{|x|^2 + f^2}$. In this description, $|\cdot|$ is the Euclidean vector norm and f is the focal length relating the optical centre of the camera and the retinal plane.

The terms occurring in (2) can be distinguished by their ordering corresponding to their appearance within the brightness equation (1).

Note that $\nabla v = \left(\frac{\partial}{\partial x_1} v, \frac{\partial}{\partial x_2} v \right)^T =: (v_{x_1}, v_{x_2})^T$ contains first-order spatial derivatives, and thus the given model is a first-order PDE. It needs to be supplemented by boundary conditions: for details see the section concerned with experiments. The expressions in (2) are also the basis for our numerical implementation of the FM scheme.

3 Discretisation and Fast Marching Implementation

It is of importance to discretise the occurring spatial derivatives in the correct fashion, as in the case of hyperbolic PDEs like the currently given Hamilton-Jacobi equation it is well-known that simply using central differences leads to a blow-up of numerical solutions. In order to ensure the stability of our algorithm as well as the validity of reasonable theoretical properties, we thus employ an upwind method as in [5, 4, 8].

Spatial Discretisation. We use the following conventions:

- $v_{i,j}$ denotes the approximation of $v(ih_1, jh_2)$, where
- i and j are the coordinates of the pixel (i, j) in x_1 - and x_2 -direction, respectively, and
- h_1 and h_2 are the corresponding mesh widths in our pixel grid.

Then the spatial discretisation of derivatives reads as

$$v_{x_1}(ih_1, jh_2) \approx h_1^{-1} \min(0, v_{i+1,j} - v_{i,j}, v_{i-1,j} - v_{i,j}), \quad (3)$$

$$v_{x_2}(ih_1, jh_2) \approx h_2^{-1} \min(0, v_{i,j+1} - v_{i,j}, v_{i,j-1} - v_{i,j}). \quad (4)$$

Terms like Q , I and $\exp(-2v)$ can be evaluated pointwise at (i, j) , so that we have completely defined the spatial discretisation of (2). We refrain from writing

down the complete discrete expression of the scheme, as this is quite cumbersome and does not give more insight.

Fast Marching. Let us now turn to the FM method. We only sketch here the idea behind it, as there are many extensive descriptions available in the literature, see especially [11].

The basic principle behind the FM scheme applied in the SfS setting is to advance monotonically a front from the foreground of the depicted object to the background. Thereby, the pixels are distinguished by the labels 'known', 'trial' and 'far', respectively, referring thereby via 'known' and 'trial' to the corresponding 3-D depth.

In the beginning, all pixels are labelled as 'far' with their depth values set to infinity. However, since the FM method propagates information from the foreground to the background, it relies on correct depth values being supplied in the pixel which is most in the foreground, i.e. the pixel with minimum depth. In the case of complex images which consist of multiple segments, for each of these segments the correct depth in the point with minimum depth must be supplied. These points are called *singular points*. These singular points are then marked as 'trial', which concludes the initialisation of the method.

For FM methods on SfS it is common to just require this data to be provided. Other methods like [5], however, do not require the knowledge of given initial depth data. We therefore aim at estimating very precisely the locations of singular points and obtain a SfS method using the FM scheme that does not rely on any depth information to be provided. The task of estimating this data will be the subject of the next section.

The 'trial' candidate with the smallest computed depth is then marked as 'known', taking the computed 3-D depth in this point as the estimate. The pixels adjacent in terms of the stencil to the new set of known points are updated with respect to their label, marking them as 'trial'. The described process is then repeated until all image pixels are marked 'known'.

Fixed-Point Iteration. Updating the depth at 'trial' points consists of solving the discrete form of (2) for v in this point. In contrast to other SfS techniques using FM, we need to solve a nonlinear equation. This is not trivial in our case, since near the solution, the derivative of (2) is very low, making standard solvers like the Newton method diverge in most cases. To avoid this, we employ the *Regula Falsi*: Starting with two values v_1 and v_2 such that $v_1 < v_2$ and the left-hand side L of (2) is negative in v_1 and positive in v_2 , one chooses

$$v_3 := \frac{L(v_2)v_2 - L(v_1)v_1}{L(v_2) - L(v_1)}, \quad (5)$$

which is between v_1 and v_2 . If L at v_3 is negative, set $v_1 := v_3$, otherwise set $v_2 := v_3$. Repeating this until v_1 and v_2 are very close together yields an estimate for the solution of (2) in this pixel. Note that computing the derivatives involves computing a minimum. Depending on v_1, v_2 and v_3 , these minima might change within the estimation process. Thus, it is necessary to update the values of v_1, v_2, v_3 during the process.

4 Estimating the Initial Depth

The FM methods for SFS rely on the knowledge of ground truth data at singular points, i.e. at points with locally minimal depth. However, in general this kind of data is not given. Thus, these depth values need to be estimated. In the experimental section, we will show that a good estimate is crucial for the reconstruction quality.

In most other works, this issue is neglected. In [4], the problem is solved by obtaining an initial estimate for the depth using an orthographic SFS method. Their perspective method, however, is not comparable with the one used in this paper, since they neglect the light attenuation term. By doing this, their solution is invariant to multiplicative scalings of the depth. This is not true in our case. To obtain a working method, we either need to know the correct depth at singular points or estimate both the singular points and their depth.

In this section, we will introduce ways to estimate the locations of singular points and estimate their depths as correctly as possible.

Lambertian Case. For simplicity, we first focus on the Lambertian case, i.e. $k_a = k_s = 0, k_d = 1$. In this simplified model, the brightness of a pixel is determined by two main factors: (i) The angle between surface normal and light source direction ϕ and (ii) the light attenuation because of the distance of the surface point to the light source. Directly from the model (1) we obtain the simple equation

$$I = I_d \frac{\cos \phi}{u^2 f^2}. \quad (6)$$

Assuming the surface to be continuously differentiable, the points of minimal depth are the points where the derivatives of the depth vanish, which means the surface normal points directly to the viewer. This results in $\phi = 0$, which leads by use of $\cos 0 = 1$ and re-arranging (6) to

$$u = \sqrt{I_d \frac{1}{I f^2}}. \quad (7)$$

Knowing the coordinates of singular points, we can compute the depth. It remains to determine the coordinates of singular points. Singular points are local minima in depth. Since minima in depth mean both less attenuation and a maximum Lambertian reflectance, this suggests that local maxima in image brightness are the singular points. At the image boundary, it might happen that we have brightness maxima that do not satisfy $\phi = 0$. In this case, there can be errors. In most cases, this does not affect the reconstruction quality significantly.

Due to sampling and quantisation artifacts, it is possible that this estimate might be slightly off, both in the location of singular points and in the estimated depth. This effect is usually rather small.

In conclusion, we propose to search local maxima in the image and estimate their depth according to equation (7). Boundary pixels should not be considered, since the estimate might be incorrect due to ϕ not being zero. The points

obtained in this way should be marked as 'trial' points for the subsequent FM method. In the Phong case which follows we use the same approach.

Phong Case. To obtain a good estimate for singular points in the general case, we review the model equation again. Essentially, we have

$$I = k_a I_a + \frac{k_d I_d \cos \phi + k_s I_s (\cos \theta)^\alpha}{u^2 f^2} . \quad (8)$$

At singular points, we have $\phi = \theta = 0$, which simplifies equation (8) to

$$I = k_a I_a + \frac{k_d I_d + k_s I_s}{u^2 f^2} . \quad (9)$$

Now, after shifting the grey values down by the ambient brightness to $I - k_a I_a$, we can separate diffuse and specular light and compute the diffuse brightness I' by

$$I' = \frac{k_d I_d}{k_d I_d + k_s I_s} (I - k_a I_a) . \quad (10)$$

Now, we can make use of the equation (7) using I' instead of I .

5 Experiments

In this section, we evaluate the presented method on both synthetic and real-world images. We discuss the accuracy and importance of the initial estimates at singular points. In comparison to other methods in the field, we evaluate the accuracy and the performance of our method.

Note that for none of the experiments, any a-priori depth information is used. In the cases where we need depth initialisation at singular points, we use the estimation method introduced in Section 4.

Lambertian case. First, we restrict the method to diffuse reflection only. We compare the reconstruction quality and performance with the methods of Prados et al. [5], Cristiani et al. [17], and Vogel et al. [18], which use all the same Lambertian model, but different schemes. Visually, the reconstructions of these methods are almost identical. Their performance, however, is different.

Figure 1 shows the vase surface [2], a classic test surface for SfS algorithms, and a rendered version of this surface using a Lambertian model. The rendering parameters are $f = 492$, $I_d = 100000$, 128×128 pixels.

When detecting the local maxima, we notice that around the maxima, we have more than just one point with the same maximal grey value. This is a result from the quantisation of the image. Since we set the depth estimates of these points to 'trial', only one of them will be used as an actual depth estimate. This might not be the actual position of the singular point, but it is close.

Figure 2 shows the reconstructions of the vase surface using both the presented method and the reference methods. The results are visually very similar. In Figure 2, also a reconstruction can be found where we manually chose a wrong depth at the singular points, i.e. we multiplied the estimates with a factor 0.9. We can see that this distorts the reconstruction.




Fig. 1. Vase surface and Lambertian rendered image.




Fig. 2. Reconstruction of the vase using a Lambertian model. Left: Reference methods. Middle: Our method. Right: Wrong depth estimate.

Table 1 supports our visual impression. It shows the relative average depth errors for the reconstructions. With the correct depth estimates, our method is about as good as the three reference methods. In fact, all reconstructions are nearly perfect. The quality of the reconstruction using the faulty estimation technique is a lot worse. This means the correctness of our initial guess is crucial for the reconstruction quality.

Phong case. Now we evaluate the method on a synthetic image rendered using the Phong reflectance model. We compare to the same methods as before, but since the reference methods only consider a Lambertian model, we additionally compare to the method of Vogel et al. [8] using the Phong model.

Figure 3 shows a rendered version and the ground truth of the Mozart face [2], a classic test image. This time, we rendered the image using the Phong reflectance model. Parameters for the rendering are $f = 500$, $I_d = I_s = 100000$, $k_d =$

Table 1. Results of the Lambertian vase experiment.

Method	Depth Error	Initialisation Time	Computation Time
Prados et al.	0.39%	$\approx 0s$	36.99s
Cristiani et al.	0.31%	$\approx 0s$	28.89s
Vogel et al.	0.32%	$\approx 0s$	2.96s
Our method	0.39%	0.02s	0.39s
Wrong initialisation	8.15%	0.02s	0.39s




Fig. 3. Mozart surface and rendered image using the Phong model.




Fig. 4. Reconstructions of the Mozart surface. Left: Lambertian methods. Middle: Vogel et al. using the Phong model. Right: Our method.

$0.7, k_s = 0.3, \alpha = 5$. Note that the Mozart face is a perfect test image for multiple sectors in an image, of which each has its own singular point.

In Figure 4, we show reconstructions of the Mozart face using our method, the method of Vogel et al., and the three Lambertian reference methods. The Phong reconstructions are clearly more accurate than the Lambertian ones. Table 2 shows the reconstruction errors and computation times of the Mozart experiment. We notice that the error of our method is about equal to the one of the method of Vogel et al., and it the Lambertian methods w.r.t. quality.

Again, our method is up to two orders of magnitude faster than any of the other methods. Another important observation is that the performance gain is much larger on the Mozart test image compared to the vase. The reason for this is the larger size of the Mozart image.

Table 2. Results of the Phong Mozart experiment. Methods marked with (L) use a Lambertian model for the reconstruction.

Method	Depth Error	Initialisation Time	Computation Time
Prados et al. (L)	12.58%	0.02s	158.62s
Cristiani et al. (L)	12.17%	0.02s	170.37s
Vogel et al. (L)	12.56%	0.02s	16.33s
Vogel et al.	5.39%	0.02s	68.76s
Our method	5.07%	0.03s	1.85s



Fig. 5. Real input image: Rook, knight, and pawn.




Fig. 6. Reconstructions of chess figures. Left: [8] with Phong. Right: Our method.

This suggests that on high-resolution images, our method might have a clear advantage over other methods in the field. This is particularly interesting for real-world images. Many authors apply their methods to relatively small test images, at most 256×256 pixels, usually even much less. We now apply our method to a full-size real-world image with 8 megapixels. On such images, the reference methods take very long to converge.

A Real-World Experiment. Figure 5 shows a photograph of three chess figures: a rook, a knight, and a pawn. The original image has size 3264×2448 and has been taken with a digital camera with flash in a darkened room.

For the reconstruction, we used the known square pixel sizes of $1.61\mu m$ and the known focal length of $70.2mm$. This gives for pixel size 1 a relative focal length of $f = 43478$, which we used for the reconstruction. Since scaling I_a , I_d and I_s only stretches the reconstructed surface by a factor that depends quadratically on the scaling factor, their magnitude is not important for the reconstruction process. For simplicity, we just chose them all equal to 100000. We manually estimated the other parameters to $k_d = 0.6$, $k_s = 0.4$ and $\alpha = 10$. We neglected ambient light, i.e. we set $k_a = 0$.

Table 3. Run times of the chess experiment. (S) marks experiments on a downsampled image of size 408×306 , (F) denotes the full, high-resolution image.

Method	Iterations	Initialisation Time	Computation Time
Vogel et al. (S)	296	0.03s	139.8s
Our method (S)	1	0.07s	2.8s
Vogel et al. (F)	1207	1.98s	38645s
Our method (F)	1	2.9s	263.2s

Figure 6 shows reconstructions of the high-resolution version of the image using our method and the method of Vogel et al., both with the same parameters. The reconstruction using our method looks much smoother than the one with the method of Vogel et al.. This can be explained by the different numerics of both methods. Our method starts at singular points and reconstructs the surface from near to far, while the other method treats all depths equally. For images like this, i.e. images with light objects in front of a dark background, our method has the clear advantage of recovering the object of interest first, such that this part is not distorted by artifacts caused by the background.

Table 3 shows the computation times compared to a test using a downsampled version of the image. While the computation times of our method are very low, the computation times of the iterative reference method are extremely high, especially for the large input image. This makes our method still applicable even on large images, outshining other methods with respect to computation time. It also shows that the performance gain of using FM for SfS increases with the size of the input image.

6 Conclusion

We have shown that the FM scheme is the method of choice for modern SfS models that incorporate light attenuation and non-Lambertian reflectance. Without compromising quality it is two orders of magnitude faster than other approaches. We demonstrated that it is possible to estimate initial depths to obtain a method that does not rely on the knowledge of initial data.

By combining state-of-the-art SfS models and proper numerical methods, it becomes possible to tackle real-world data with image sizes of many megapixels. This is far beyond the size of the model problems that are considered in many SfS papers.

References

1. Horn, B.K.P., Brooks, M.J.: Shape from Shading. Artificial Intelligence Series. MIT Press (1989)
2. Zhang, R., Tsai, P.S., Cryer, J.E., Shah, M.: Shape from shading: A survey. IEEE Transactions on Pattern Analysis and Machine Intelligence **21**(8) (1999) 690–706

3. Tankus, A., Sochen, N., Yeshurun, Y.: A new perspective [on] shape-from-shading. In: Proc. Ninth International Conference on Computer Vision. Volume 2. IEEE Computer Society Press, Nice, France (October 2003) 862–869
4. Tankus, A., Sochen, N., Yeshurun, Y.: Perspective shape-from-shading by fast marching. In: Proc. 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. Volume 1. IEEE Computer Society Press, Washington, DC (June–July 2004) 43–49
5. Prados, E., Faugeras, O.: Shape from shading: A well-posed problem? In: Proc. 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. Volume 2. IEEE Computer Society Press, San Diego, CA (June 2005) 870–877
6. Ahmed, A., Farag, A.: A new formulation for shape from shading for non-Lambertian surfaces. In: Proc. 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. Volume 2. IEEE Computer Society Press, New York, NY (June 2006) 17–22
7. Oren, M., Nayar, S.: Generalization of the Lambertian model and implications for machine vision. *International Journal of Computer Vision* **14**(3) (1995) 227–251
8. Vogel, O., Breuß, M., Weickert, J.: Perspective shape from shading with non-Lambertian reflectance. In Rigoll, G., ed.: *Pattern Recognition*. Volume 5096 of *Lecture Notes in Computer Science*. Springer, Berlin (June 2008) 517–526
9. Phong, B.T.: Illumination for computer-generated pictures. *Communications of the ACM* **18**(6) (1975) 311–317
10. Tsitsiklis, J.N.: Efficient algorithms for globally optimal trajectories. *IEEE Transactions on Automatic Control* **40**(9) (September 1995) 1528–1538
11. Sethian, J.A.: *Level Set Methods and Fast Marching Methods*. Second edn. Cambridge University Press, Cambridge, UK (1999)
12. Kimmel, R., Sethian, J.A.: Optimal algorithm for shape from shading and path planning. *Journal of Mathematical Imaging and Vision* **14**(3) (2001) 237–244
13. Yuen, S.Y., Tsui, Y.Y., Chow, C.K.: Fast marching method for shape from shading under perspective projection. In: *Proceedings of the 2nd IASTED International Conference on Visualization, Imaging and Image Processing*, Malaga, Spain (September 2002) 584–589
14. Tankus, A., Sochen, N., Yeshurun, Y.: Shape-from-shading under perspective projection. *International Journal of Computer Vision* **63**(1) (June 2005) 21–43
15. Prados, E., Soatto, S.: Fast marching method for generic shape from shading. In Paragios, N., Faugeras, O., Chan, T., Schnörr, C., eds.: *Variational, Geometric, and Level Set Methods in Computer Vision*. Volume 3752 of *Lecture Notes in Computer Science*. Springer (2005) 320–331
16. Yuen, S.Y., Tsui, Y.Y., Chow, C.K.: A fast marching formulation of perspective shape from shading under frontal illumination. *Pattern Recognition Letters* **28** (2007) 806–824
17. Cristiani, E., Falcone, M., Seghini, A.: Some remarks on perspective shape-from-shading models. In Sgallari, F., Murli, F., Paragios, N., eds.: *Scale Space and Variational Methods in Computer Vision*. Volume 4485 of *Lecture Notes in Computer Science*. Springer, Berlin (2007) 276–287
18. Vogel, O., Breuß, M., Weickert, J.: A direct numerical approach to perspective shape-from-shading. In Lensch, H., Rosenhahn, B., Seidel, H.P., Slusallek, P., Weickert, J., eds.: *Vision, Modeling, and Visualization*. AKA, Berlin (November 2007) 91–100