IEEE CVPR-94, Seattle, Washington, June 20-24, pages 377-384.

Analysis of Shape from Shading Techniques *

Ruo Zhang, Ping-Sing Tsai, James Edwin Cryer, Mubarak Shah
Computer Science Department
University of Central Florida
Orlando, FL 32816
(shah@cs.ucf.edu)

Abstract

Since the first shape-from-shading technique was developed by Horn in the early 1970s, different approaches have been continuously emerging in the past two decades. Some of them improve existing techniques, while others are completely new approaches. However, there is no literature on the comparison and performance analysis of these techniques. This is exactly what is addressed in this paper.

1 Introduction

Shape-from-shading (SFS) deals with the recovery of shape from a gradual variation of shading in the image. Artists have long exploited lighting and shading to convey vivid illusions of depth in paintings. In SFS, it is important to study how the images are formed. A simple model of image formation is the Lambertian model. According to the Lambertian model, the gray level at a pixel in the image depends on the light source location, and the surface normal. In SFS, given a gray level image, the aim is to recover the light source and a surface normal at each pixel in the image.

In this paper, a total of eight well-known SFS algorithms are implemented and compared in terms of timing and accuracy, in order to analyze the advantages and disadvantages of these approaches. The experiments were performed on various images with different light sources. The performance of the algorithms was analyzed using depth error, surface gradient error and CPU timing. The comparison showed that all of them have some limitations. None of the algorithms has consistent performance for all images, since they work well for certain images, but perform poorly for others. In general, global approaches are more robust, while local approaches are faster.

2 Shape from Shading

SFS techniques can be divided into two groups: Global approaches and local approaches. Global approaches can be further divided into global minimization approaches and global propagation approaches. Global minimization approaches obtain the solution by minimizing an energy function. Global propagation approaches propagate the shape information from known surface points (e.g., singular points) to the whole image. Local approaches derive shape only from the intensity information of the surface points in a small neighborhood.

One of the earlier global minimization approaches was by Ikeuchi and Horn [10]. Since each surface point has two unknowns for the surface normal, and each pixel in the image provides one gray value, therefore image gray levels alone are not enough to recover the shape. To overcome this, Ikeuchi and Horn introduced two constraints: The brightness constraint and the smoothness constraint. The brightness constraint requires that the reconstructed shape produce the same brightness as the input image at each surface point, while the smoothness constraint forces the gradient of the surface to change smoothly. The shape was computed by minimizing an energy function which consists of the above two constraints. Also using these same constraints, Brooks and Horn (B&H) [2] minimized the same energy function, in terms of surface normal instead of surface gradient. Frankot and Chellappa [5] enforced the integrability in B&H's algorithm in order to recover integrable surfaces (surfaces for which $z_{xy} = z_{yx}$). Surface slope estimates from the iterative scheme were expressed in terms of a linear combination of a finite set of orthogonal Fourier basis functions. The enforcement of integrability was done by projecting the nonintegrable surface slope estimates onto the nearest (in terms of distance) integrable sur-

*This work was supported by NSF grants CDA-9122006, CDA-9222798 and CDA-9200369. The extended version of this paper, the implementation of algorithms in C, and images used in this paper, are available by anonymous ftp under the *pub/tech_paper/survey* directory at *eustis.cs.ucf.edu* (132.170.108.42).

face slopes. This projection was fulfilled by finding the closest set of coefficients which satisfy integrability in the linear combination. Their results showed improvements in both accuracy and efficiency. Later, Horn [8] also replaced the smoothness constraint in his approach with an integrability constraint. The major problem with Horn's method is its slow convergence. Szeliski [20] sped it up using a hierarchical basis pre-conditioned conjugate gradient descent algorithm. Based on the geometrical interpretation of B&H's algorithm, Vega and Yang [22] applied heuristics to the variational approach so that the stability of B&H's algorithm was improved.

Instead of the smoothness constraint, Zheng and Chellappa (Z&C) [24] introduced an intensity gradient constraint, which specifies that the intensity gradients of the reconstructed image and the input image are close to each other in both the x and y directions. Leclerc and Bobick (L&B) [11] solved directly for depth by using a discrete formulation and employing a conjugate gradient technique. The brightness constraint and smoothness constraint were applied to ensure convergence, and a stereo depth map was used as an initial estimate. Recently, Lee and Kuo (L&K) [13] proposed an approach to recover depth using the brightness and the smoothness constraint. They approximated surfaces by a union of triangular patches. Unlike L&B's method, this approach did not require the depth from stereo as an initial value.

All of the above approaches deal with a single smooth surface. Malik and Maydan [14] developed the first solution for piecewise smooth surfaces. They combined the line drawing and shading constraints in an energy function, and recovered both surface normal and line labeling through the minimization of the energy function.

The first global propagation approach was the characteristic strip technique by Horn [7]. A characteristic strip is a line in the image along which the surface depth and orientation can be computed if these quantities are known at the starting point of the line. Horn's method constructs initial surface curves around the neighborhoods of singular points (singular points are the points with maximum intensity) using a spherical approximation. The shape information is propagated simultaneously along the characteristic strips outwards, assuming no crossover of adjacent strips. The direction of characteristic strips are identified as the direction of intensity gradients. In order to get a dense shape map, new strips have to be interpolated when neighboring strips separate too much.

Oliensis [15] observed that the smoothness constraint is only needed at the boundaries if we have initial values at the singular points. Therefore, the surface should be reconstructed from the interior of the image outward, instead of from the boundary inward. Based on this idea, Dupuis and Oliensis (D&O) [4, 16] formulated SFS as an optimal control problem, and solved it using numerical methods. The proof of equivalence between the optimal control representation and SFS was illustrated. Their initial algorithm [4] required priori depth information for all the singular points. A later extension [16] could determine this information automatically by assuming twice differentiable depth, isolated singular points and nonzero curvature at singular points. Following the main idea of D&O's approach, Bichsel and Pentland (B&P) [1] proposed a minimum downhill approach for SFS which converged in less than ten iterations.

Among the local approaches, two are by Pentland, one by Lee and Rosenfeld (L&R), and one by Tsai and Shah(T&S). Pentland [17] recovered shape information from the intensity, and its first and second derivatives. He used the assumption that the surface is locally spherical at each point. Under the same spherical assumption, L&R [12] computed the slant and tilt of the surface in the light source coordinate system through the first derivative of the intensity. A later approach by Pentland [18] used the linear approximation of the reflectance function in terms of the surface gradient, and applied a Fourier transform to the linear function to get a closed form solution for the depth at each point. Similar to Pentland's method, T&S [21] applied the discrete approximation of the gradient first, then employed the linear approximation of the reflectance function in terms of the depth directly. Their algorithm iteratively recovered the depth at each point without using any global information.

3 Experimental Images

It is very difficult to choose good test images for SFS algorithms. A good test image must match the assumptions of the algorithms, e.g. Lambertian reflectance model, constant albedo value, and infinite point source illumination. In this section, we describe the images chosen to test the SFS algorithms.

The synthetic images were generated using true depth maps, or range data obtained from a laser range finder. We simply computed the surface gradient ($p = \frac{\partial Z}{\partial x}$, $q = \frac{\partial Z}{\partial y}$) using the forward discrete approximation of the depth, Z, and generated shaded images using the Lambertian reflectance model. There are at least two advantages of using synthetic images. First, we can generate shaded images with different light source directions for the same surface. Second, with the true depth information, we can compute the error and compare the performance.

We used five synthetic surfaces: **Sphere**, **Synthetic Vase**, **Mozart**, **Penny**, and **Sombrero**. The true depth maps of the first two surfaces, Sphere and Vase, were generated mathematically.

Five real images were also used. Their light source directions, S, given below, either were estimated by the L&R method or provided with the images:

- **Lenna:** $S = (1.5, 0.866, 1)$.
- **Mannequin:** $S = (-0.345, 0.345, 0.875)$.
- **Pepper:** $S = (0.766, 0.642, 1)$.
- **David:** $S = (-0.707, 0.707, 1)$.
- **Vase:** $S = (-0.939, 1.867, 1.0)$.

4 Experimental Results

We implemented eight of the twelve algorithms discussed in this paper. Szeliski's algorithm was not implemented since it is a faster version of B&H's algorithm. We did not implement Ikeuchi and Horn's algorithm because it resembles B&H's. Since B&P's algorithm is a simplification of Dupuis and Oliensis's, we implemented B&P's algorithm only.

Below we discuss some important points about the implementation of each algorithm, and analyze the results using 3-D plots of the depth maps. Due to space limitation, the 3-D plots are not included here.

4.1 Brooks and Horn (B&H)

B&H's approach requires occluding boundary information for the input image. Their algorithm computes the shape in terms of surface normals. In order to reconstruct the surface of the object, an integration step must be applied to compute the depth. However, neither finding the occluding boundary from an image nor integrating the surface normals are easy tasks. Since the primary objective of this survey paper is to study SFS algorithms, we only tested two synthetic images, sphere and synthetic vase with light source $(0, 0, 1)$, for which the occluding boundary information was available. The results of the reconstructed gray level images from the computed shape closely resemble the input images. A well-known problem of B&H's method is the slow convergence rate. In our implementation, we forced the algorithm to terminate when the error in the energy function starts to increase, or the error is less than some threshold. It took at least 400 iterations to achieve convergence for our test images.

4.2 Zheng and Chellappa (Z&C)

The implementation of Z&C's method is very straightforward. We used the forward difference approximation to compute the partial derivatives. For the border points, where the forward approximation could not be applied, we switched to the backward difference approximation for the first order partial derivatives and set the second order partial derivatives to zero. This method is very robust, since no parameters have to be tuned.

Our implementation works well for most of the real images, except for Mannequin and Vase. This is due to the dark background in these two images, which violates the uniform albedo assumption used in their algorithm. The basic shape for Lenna and Pepper are recovered with enough details. However, some errors are observed around the mouth and on the cheeks in Lenna. Their method also has a problem with light source $(0, 0, 1)$, which will zero out most of the terms in the approximation equation of the iterative method. In order to get reasonable results for the images with light source $(0, 0, 1)$, we used $(0.01, 0.01, 1)$ instead as the light source direction. Their results also show some error along the light source direction. We think this is due to the use of the intensity gradient constraint instead of the smoothness constraint used in their energy function, and the discrete approximation used for computing the partial derivatives.

4.3 Leclerc and Bobick (L&B)

L&B's approach was implemented without hierarchical structure, using the conjugate gradient routine. On the occluding boundary, the discrete approximations for the first order partial derivatives were changed from the central difference approximation to either the forward approximation or the backward approximation, and the second order partial derivatives were set to zero. Their approach requires the output from stereo as the initial estimate for the conjugate gradient method. Since we do not have stereo pairs for the test images, we used the true depth with $\pm 5\%$ uniform random noise as the initial estimate, and tested the algorithm on synthetic images only.

Their results depend heavily on the initial estimate and the initial weight of the smoothness term, λ. We observed that the algorithm works well on two sets of images, Sphere and Mozart, even when the light source is from the side. The initial value of the smoothness term was 0.25 for both sets of images, which was gradually reduced to 0.01 by a factor of 0.7. The maximum number of iterations for the conjugate gradient routine was set to 200. This method is basically a combination of stereo and shape from shading. It is hard to compare the performance of this method with other SFS methods. Because it heavily relies on the initial estimate from stereo, one has to take into account the computation and accuracy involved in stereo as well.

4.4 Lee and Kuo (L&K)

L&K's algorithm was implemented using the V-cycle multigrid scheme to solve the linear system, as reported in their paper. We used Gauss-Seidel relaxation as the smoothing operator, and as the exact solver for the finest grid. Full-weighting restriction was applied to transfer the residual from finer grids to coarser grids, and bi-linear interpolation was applied to make the prolongation from the coarser grid to finer grids. The same stencil was used for the smoothness term as given in their paper. The nodal points in the finest grids were chosen to be the image pixels. Successive linearizations were done through a maximum of 10 successive iterations, and the number of V-cycles was set to 10 for the first iteration, 2 for the second, 1 for the rest. The initial values for depths of the finest grids, and corrections for the coarser grids, were all zeros. Since the algorithm does not work for light source direction $(0, 0, 1)$, we used $(0.0001, 0.0001, 1.0)$ as the input light source direction instead. This light source approximation is different from the $(0.01, 0.01, 1)$ used in Z&C's algorithm. The implementation of Z&C's algorithm does not work with $(0.0001, 0.0001, 1.0)$.

For most of the images, the smoothing factor was 2000, and the level of grids was computed by $L = \log(M) - 1$, where M is the size of the image. However, to eliminate the effect of over-smoothing, we used 200 as the smoothing factor for Sombrero and David, and we ran only 1 iteration for David, Mannequin and Penny. The depth maps, after the first iteration, contain more detail but have a smaller range. After 10 iterations, details are smoothed out, but the depth range is wider. This means that more iterations will provide more low frequency information, which overtakes the high frequency information from the initial iterations.

The algorithm works well, even when the light source is from the side, except in the cases of Sphere and Vase which create the most self-shadows. The recovered surfaces are well outlined, but lack details and have a tendency to be over-smoothed. Although different smoothing factors can be used for different images in order to get the best results, small changes in the smoothing factor will not affect the results very much.

4.5 Bichsel and Pentland (B&P)

In the implementation of B&P's algorithm the initial depth values for the singular points were assigned a fixed positive value, and the depth values for the other points were initialized a large negative value. Instead of computing the distance to the light source, only the local surface height was computed and maximized, in order to select the minimum downhill direction. This was based on the fact that the distance to the light source is a monotonically increasing function of the height when the angle between the light source direction and the optical axis (z-axis here) is less than 90 degrees. Height values were updated with a Gauss-Seidel iterative scheme and the convergence was accelerated by altering the direction of the pass at each iteration.

The algorithm provides the best results for the cases when the light source is on the side; even the sphere can be recovered very well when the light source comes from the side. However, the algorithm does not give good results for real images except for Pepper. This, we think, is due to the inaccuracy of the initial singular points, and noise in the real images. The algorithm is very fast; usually only 5 iterations are required to provide reasonable results.

4.6 Lee and Rosenfeld (L&R)

Their method estimates the depth of an image using local spherical assumption and intensity derivatives. This makes the algorithm unsuitable for non-spherical surfaces, and very sensitive to noise, which is observed in the depth maps obtained for the real images and some synthetic images, such as Penny or Mozart. The intensity of the real images varies slightly, causing the depth estimation to falter, while the synthetic images yield good depth maps, due to the smooth surfaces.

4.7 Pentland (P)

Pentland's algorithm [18] produces good results on most surfaces that change linearly, even if the surface has a naturally varying surface such as a person's face. However, this algorithm falls apart when the surface changes in a non-linear manner; this is clearly observed from the results of Sphere. For real images, the algorithm produces the best results except for Vase. The details of Mannequin are not recovered and inaccuracy is high around the eye regions in David and Lenna.

4.8 Tsai and Shah (T&S)

Their method works very well on smooth objects with the light source close to the viewing direction. However, it is sensitive to the intensity noise, such as the black hole on the nose of Mozart image or the shadow areas. The problem of the convex/concave ambiguity is clearly shown in the result for the Sombrero image with light source $(0, 0, 1)$. The results for real images are good for Mannequin, David and Vase, but noisy for Lenna and Pepper, especially in the top and bottom regions of Pepper, and the nose, eyes, and hat regions of Lenna. These are regions where there are sudden intensity changes, which cause roughness in the depth estimate due to the relationship between depth and intensity.

5 Error Analysis

In this section, we will quantitatively analyse the results for the synthetic images, for which the true depth maps are available, by using the following error measures:

- **Mean and standard deviation of depth error** (Tables 1-2). For each algorithm, we compared the recovered depth with the true depth from the range image. The output depth from each algorithm was first normalized according to the range data, then compared with the range data for mean and standard deviation of depth error.

- **Mean gradient error** (Table 3). This indicates the error in the surface orientation. The standard deviation is not used here, since it does not have any physical meaning. The forward discrete approximation was used to compute the gradient from the recovered depth.

- **Difference images of the absolute depth error** which provide the depth error distribution over the images to show the dependence of error on the underlying surface structure and image location. The depth error images were obtained by first calculating the absolute depth error at each point, then rescaling it using the minimum and the maximum value over the whole image. The regions which have the least depth error and the regions which have the most depth error can be easily identified from these images.

- **The histograms of the percentage depth error** which show percentages of depth errors with respect to true depths and distributions of these percentages. The Y-axis of each plot represents the number of pixels. The X-axis of each plot represents the percentage depth error which was computed pixel by pixel using the following formula:

  $$\frac{|true\ depth\ -\ estimated\ depth|}{true\ depth} \times 100\%.$$

  There are some pixels with more than 100% error. This may happen at the points in the shadow areas, the points with convex/concave ambiguity, or at the object boundaries. All pixels which have more than 100% error were plotted as 101% error.

For those algorithms which compute the surface gradient together with the depth, we still use the discrete approximation of the depth to calculate the surface gradient in the gradient error table, in order to be consistent with the other algorithms. Due to space limitation, the difference images and histograms are not included here.

From Table 1 (depth error), L&R's method gives the best results for the sphere; this is due to the spherical assumption used in the algorithm. However, it provides poor results for the real images. T&S's approach produces very good results for Sphere with light source $(0, 0, 1)$, but not for images with the light source from the side. This is due to the linearization of the reflectance function in terms of depth. When the intensity cannot globally reflect the depth information, the algorithm falls apart. Pentland's approach also has this problem when the surface shape changes nonlinearly, as with spherical surfaces. L&B's conjugate gradient approach produces the least depth error since the initial depth, used in our tests, is close to the ground truth. On the average, B&P's minimum downhill approach gives good results even when the light source is from the side, and the results of L&K's approach are close to the results of Z&C's approach.

The standard deviation of depth (Table 2) agrees with the average depth error in the sense that the one with smaller average error would have smaller standard deviation in most cases.

From the gradient error (table 3), we find that Pentland's approach gives the best results for most test images, except for Sphere and Sombrero images with light source direction $(0, 0, 1)$. This suggests that local intensity information is sufficient for a good shape estimation.

From these three error tables, we conclude that there is no strict ordering for the accuracy of the algorithms, however, overall L&B's is the best, since it uses good initial estimates. L&K's places second, especially in terms of the gradient error. Z&C's algorithm takes third place, followed by B&P's, Pentland's, T&S's, and finally L&R's.

From the difference images, we observe that L&B's algorithm is still the best, most of errors occur at the boundaries of the objects. Among the remaining six algorithms which do not require accurate initial values, L&R's algorithm only has errors along half of the object boundaries for spheres with light sources $(5, 5, 7)$ and $(1, 0, 1)$. B&P's algorithm also has low error distribution for these two images. The error is also very small in the center of all three vase images for B&P's, and L&R's algorithms, and in the center of vase image with light source $(0, 0, 1)$ for T&S's algorithm. For the Mozart images with light source $(5, 5, 7)$ and $(1, 0, 1)$, L&K's, and B&P's methods have the lowest error at the face areas. For Penny and Sombrero images, the errors for all six algorithms are equally distributed over the whole image.

The histograms show that L&B's algorithm gives the best results, since we used the true depth with $\pm5\%$ uniform random noise as the initial estimate. None of the other algorithms give good results for Penny and Sombrero images, since there are large number of pixels with more than 100% error in the histograms. For Mozart image, L&K's algorithm gives the best results even for the image with light source from the side. L&R's, and B&P's algorithms give better results for Sphere and Vase images.

6 Timing

CPU timing was computed on a SUN SPARC 4. The disk I/O time was not included, and only the computational time was considered. Due to space limitation, the CPU timing table is not included either. The results show that the three local approaches are significantly faster than the global approaches; their times depend only on the size of the input image. For the global approaches, time not only depends on the size of the input image, but also varies from scene to scene. Among the global approaches, B&P's algorithm is the most efficient. L&B's algorithm, without hierarchical structure, takes the most time. L&K's algorithm is also time consuming, since it involves multigrid iterations. Z&C's algorithm is reasonably fast with pyramid implementation. The order of the algorithms according to CPU time, from the slowest to the fastest, is L&B's algorithm, L&K's algorithm (in most cases, L&B's algorithm is slower than L&K's), then followed by Z&C's, B&P's, Pentland's, L&R's, and T&S.

7 Conclusions and Future Research

In this paper, we analysed a total of twelve existing algorithms and grouped them into three different categories: global minimization techniques, global propagation techniques, and local techniques, These groupings are based on the conceptual differences among the algorithms. Eight representatives out of the twelve were implemented in order to compare their performance in terms of accuracy and time.

Overall, the global minimization techniques are more robust to different scenes and noise. Among them, L&B's algorithm yields very good results due to the use of good initial estimates from stereo. L&K's algorithm produces the second best results. The global propagation techniques provide almost perfect results if the estimates for singular points are accurate. The local approaches tend to have more error for real, noisy images, especially for L&R's approach which is based on intensity derivatives and the spherical assumption. The conclusions from the timing is that the local approaches are faster than the global approaches, and the global propagation approaches are a lot faster than the global minimization approaches. The execution times for local appraoches depend only on the size of the input image. While for the global approaches, time not only depends on the size of the input image, but also varies from scene to scene.

There are several possible directions for future research. As we noted, reflectance models used in SFS are too simplistic; recently, more sophisticated models have been proposed. This not only includes more accurate models for Lambertian, specular, and hybrid reflectance, but also includes replacing the assumption of orthographic projection with perspective projection, which is a more realistic model of cameras. The traditional simplification of lighting conditions, assuming an infinite point light source, can also be eliminated by either assuming a non-infinite point source, or simulating lighting conditions using a set of point sources. This trend will continue. SFS methods employing more sophisticated models will be developed to provide more accurate, and realistic, results.

Another direction is the combination of shading with other cues. One can use results of stereo or range data to improve results of SFS (such as [11] and [19]), or use results of SFS or range data to improve results of stereo. A different approach is to directly combine results from shading and stereo (such as [3]).

Multiple images can also be employed by moving either the viewer (as in [6]) or the light source (as in [23]) in order to successively refine the shape. The successive refinement can improve the quality of estimates by combining estimates between image frames, and reduce the computation time since the estimates from the previous frame can be used as the initial values for the next frame, which may be closer to the correct solution. By using successive refinement, the process can be easily started at any frame, stopped at any frame, and restarted if new frames become available. The advantage of moving the light source over moving the viewer is the elimination of the mapping of the depth map (warping) between image frames.

One problem with SFS is that the shape information in the shadow areas is not recovered, since shadow areas do not provide enough intensity information. This can be solved if we make use of the information available from shape-from-shadow (shape-from-darkness) and combine it with the results from SFS. The depth values on the shadow boundaries from SFS can be used either as the initial values for shape-from-shadow, or as constraints for the shape-from-shadow algorithm. In the case of multiple image frames, the information recovered from shadow in the previous frame can also be used for SFS in the next frame.

Acknowledgements

Thanks to Dr. Ascher, Dr. Bichsel, Prof. Bobick, Prof. Carter, Prof. Dupuis, Dr. Lee, Dr. Leclerc, Prof. Healey, Prof. Kuo, Prof. Woodham and Dr. Zheng for helpful discussions and/or providing images.

References

[1] M. Bichsel and A. P. Pentland. A simple algorithm for shape from shading. In *CVPR*, pages 459–465, 1992.

[2] M. J. Brooks and B. K. P. Horn. Shape and source from shading. In *Proceedings of International Joint Conference on Artificial Intelligence*, pages 932–936, 1985.

[3] J. Cryer, P. Tsai, and M. Shah. Integration of shape from x modules: Combining stereo and shading. In *CVPR*, pages 720–721, June 1993.

[4] P. Dupuis and J. Oliensis. Direct method for reconstructing shape from shading. In *CVPR*, pages 453–458, 1992.

[5] R. T. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. *PAMI*, 10:439–451, 1988.

[6] J. Heel. Dynamic motion vision. *Robotics and Autonomous Systems*, 6(1), 1990.

[7] B. K. P. Horn. *Shape from Shading: A Method for Obtaining the Shape of a Smooth Opaque Object from One View.* PhD thesis, MIT, 1970.

[8] B. K. P. Horn. Height and gradient from shading. *IJCV*, pages 37–75, 1989.

[9] B. K. P. Horn, R. S. Szeliski, and A. L. Yuille. Impossible shaded images. *PAMI*, 15(2):166–170, 1993.

[10] K. Ikeuchi and B.K.P. Horn. Numerical shape from shading and occluding boundaries. *Artificial Intelligence*, 17(1-3):141–184, 1981.

[11] Y. G. Leclerc and A. F. Bobick. The direct computation of height from shading. In *CVPR*, pages 552–558, 1991.

[12] C.H. Lee and A. Rosenfeld. Improved methods of estimating shape from shading using the light source coordinate system. *Artificial Intelligence*, 26:125–143, 1985.

[13] K. M. Lee and C. C. J. Kuo. Shape from shading with a linear triangular element surface model. *Signal and Image Processing Institute, University of Southern California, Technical Report 172*, 1991.

[14] J. Malik and D. Maydan. Recovering three dimensional shape from a single image of curved objects. *PAMI*, 11, 1989.

[15] J. Oliensis. Shape from shading as a partially well-constrained problem. *CVGIP: IU*, 54:163–183, 1991.

[16] J. Oliensis and P. Dupuis. A global algorithm for Shape from shading. In *ICCV*, pages 692-701, 1993.

[17] A. P. Pentland. Local shading analysis. *PAMI*, 6:170–187, 1984.

[18] A. Pentland. Shape information from shading: a theory about human perception. In *ICCV*, pages 404–413, 1988.

[19] H. Pien and J. Gauch. A variational approach to sensor fusion using registered range and intensity data. In *SPIE OE/Aerospace Sensing*, April 1993.

[20] R. Szeliski. Fast shape from shading. *CVGIP: IU*, 53:129–153, 1991.

[21] P. S. Tsai and M. Shah. A simple shape from shading algorithm. In *CVPR*, pages 734–736, 1992.

[22] O. E. Vega and Y. H. Yang. Shading logic: a heuristic approach to recover shape from shading. *PAMI*, 15(6):592–597, 1993.

[23] R. Zhang, P. T. Tsai, and M. Shah. Shape from photomotion. In *CVPR*, pages 740–741, 1993.

[24] Q. Zheng and R. Chellappa. Estimation of illuminant direction, albedo, and shape from shading. *PAMI*, 13(7):680–702, 1991.

Table 1: Average Z error for synthetic images.

Methods	Images														
	Sphere			**Vase**			**Mozart**			**Penny**			**Sombrero**		
	S1	S2	S3	S1	S2	S3	S1	S2	S3	S1	S2	S3	S1	S2	S3
Zheng & Chellappa	10.3	10.9	26.6	8.5	12.3	8.5	15.1	16.6	10.5	7.2	4.8	4.8	7.0	6.6	5.6
Leclerc & Bobick	2.1	3.6	3.7	1.8	3.0	*.*	1.7	4.7	5.5	2.2	*.*	*.*	1.2	*.*	*.*
Lee & Kuo	16.0	10.3	10.9	10.0	7.5	7.9	16.0	8.8	11.4	7.6	4.7	4.4	6.9	5.4	7.7
Bichsel & Pentland	0.7	9.4	5.2	10.0	8.8	7.9	20.5	17.8	7.7	12.1	8.0	8.4	13.7	11.0	6.4
Lee & Rosenfeld	0.8	3.8	4.3	8.1	8.4	11.0	18.3	17.8	17.6	11.3	8.2	7.9	11.7	8.8	8.7
Pentland	17.3	20.1	14.0	11.2	13.6	9.0	15.7	22.5	19.7	7.4	6.4	6.6	7.3	7.6	7.3
Tsai & Shah	0.1	16.4	16.4	8.3	11.8	12.7	18.5	20.1	20.0	11.0	8.5	8.6	12.6	10.2	10.1

Table 2: Standard deviation of Z error for synthetic images.

Methods	Images														
	Sphere			**Vase**			**Mozart**			**Penny**			**Sombrero**		
	S1	S2	S3	S1	S2	S3	S1	S2	S3	S1	S2	S3	S1	S2	S3
Zheng & Chellappa	13.1	15.9	17.4	11.1	17.0	13.9	18.4	17.2	15.9	13.5	7.4	5.5	11.9	10.5	10.4
Leclerc & Bobick	2.4	5.0	5.1	2.9	4.1	*.*	2.1	7.1	7.7	3.1	*.*	*.*	2.4	*.*	*.*
Lee & Kuo	15.9	18.9	20.6	13.2	12.9	15.39	19.2	15.9	22.1	14.0	7.3	6.6	12.2	9.2	15.0
Bichsel & Pentland	1.2	13.4	9.0	13.8	13.6	16.9	37.4	21.9	14.6	23.4	11.7	16.6	26.6	20.3	12.5
Lee & Rosenfeld	0.4	5.8	6.6	14.6	16.4	22.3	33.0	29.8	30.3	21.2	15.2	14.8	22.5	17.0	16.9
Pentland	17.5	18.3	19.3	12.6	18.9	11.1	18.2	24.2	20.5	12.2	10.6	11.1	12.2	12.6	12.6
Tsai & Shah	0.1	20.9	21.0	15.0	16.9	19.7	33.3	30.7	30.5	20.6	15.4	15.6	24.3	18.4	18.4

Table 3: Average p-q error for synthetic images.

Methods	Images														
	Sphere			**Vase**			**Mozart**			**Penny**			**Sombrero**		
	S1	S2	S3	S1	S2	S3	S1	S2	S3	S1	S2	S3	S1	S2	S3
Zheng & Chellappa	2.8	1.6	1.7	2.2	1.5	1.3	2.3	1.1	1.1	1.3	1.1	1.0	1.3	1.0	0.7
Leclerc & Bobick	0.8	4.5	4.4	1.2	3.1	*.*	0.5	8.4	9.5	1.2	*.*	*.*	0.6	*.*	*.*
Lee & Kuo	2.3	1.4	1.4	1.6	0.9	0.9	1.7	0.7	0.6	1.3	1.1	1.0	0.8	0.7	0.6
Bichsel & Pentland	0.3	5.8	2.5	2.7	4.9	1.9	3.1	8.1	1.9	1.7	4.4	1.1	1.2	3.3	0.5
Lee & Rosenfeld	0.1	6.5	6.7	1.3	3.3	2.2	6.8	13.7	12.8	4.3	8.4	7.0	1.3	2.5	2.3
Pentland	2.2	2.9	4.7	1.8	1.3	1.2	1.3	1.3	1.3	1.3	1.3	1.2	1.1	1.1	1.0
Tsai & Shah	0.1	0.9	0.9	1.4	1.4	2.6	6.7	5.5	5.6	4.2	5.2	4.8	1.2	1.5	1.5

S1, S2, and S3 stand for three different light sources, $(0, 0, 1)$, $(5, 5, 7)$, and $(1, 0, 1)$, and '*.*' stands for unavailable data.