Light source direction from a single image: a performance analysis

Danny Gibbins*, Michael J. Brooks, Wojciech Chojnacki****

*Discipline of Computer Science, School of Information Science and Technology, Flinders University of South Australia, GPO Box 2100, Adelaide, SA 5001.
**Department of Computer Science, University of Adelaide, GPO Box 498, Adelaide, SA 5001.

A smooth object depicted in a photograph will often exhibit brightness variation, or shading. Of interest in computer vision is the problem of how object shape may be determined from image shading. Various computational techniques have emerged that perform reasonably well in solving this problem. However, these methods typically require substantial scene information prior to commencement. This requirement can be reduced by using pre-processing techniques to estimate, for a given image, the direction of the principal light source, or "sun". In this paper, we conduct a comparative performance analysis of these techniques, revealing their poor versatility. An informal study is undertaken of the ability of people to perceive the principal light source direction from synthesised images. Some implications of this work for models of congition are discussed.

Keywords and phrases: Light source direction, shape from shading, Lambertian surface, shape perception, models of cognition.

Manuscript received 5 August 1991.

1 INTRODUCTION

A black and white photograph of a smooth object will often exhibit brightness variation, or *shading*. Of interest in computer vision and perception psychology has been the inverse problem of how object shape may be extracted from image shading (eg see Horn, 1975; Horn and Brooks, 1989). Many computational techniques have been developed to solve the shading problem, and some of these perform reasonably well in favourable circumstances. However, a difficulty with most of the computational schemes is the need for considerable prerequisite information relating to the conditions under which a given image was formed. This is not surprising given that the shape from shading problem cannot be mathematically formulated, much less solved, without such information.

Contrasting with this is the apparent ease with which people perceive shape from shading, without prior knowledge of the scene conditions. Since (in a formal sense) this would appear to be impossible, it must be that people make certain assumptions, either explicitly or implicitly, relating to the scene conditions, and that these assumptions lead to fairly robust perceptions of shape in the presence of a variety of actual conditions. (Not that we should hold an exaggerated opinion of human capability in this regard.) It might be, for example, that people undertake a pre-processing step in order to determine the direction of the principal light source, or sun. It is of interest to note that psychological tests show that people have a decided predilection for perceiving shapes in an image that are consistent with a single sun in a specific direction (Ramachandran, 1988).

In this article, we analyse the performance of several computational techniques for automated recovery of light source direction. We also report on our informal studies of the human perception of source direction from single imagery. We conclude with some remarks concerning the likelihood of there being a human pre-processing phase to determine source direction, commenting on some relevant studies due to Mingolla and Todd (1986).

2 TESTING OF AUTOMATED TECHNIQUES

2.1 Methods Considered

All of the methods tested seek the slant and tilt of the source direction, as defined in Figure 1. Details of their derivations and formulations may be found in Gibbins, Brooks, Chojnacki (1991). The methods considered are:

Tilt Estimation

— Pentland Method: Under the assumption of a uniform distribution of surface normals, an estimate for light source tilt is expressed in terms of the mean intensity changes in several directions (Pentland, 1982).
— Lee and Rosenfeld Method: Assuming a uniform distribution of surface normals, an estimate for light source tilt is derived from the mean values of the first deriva-

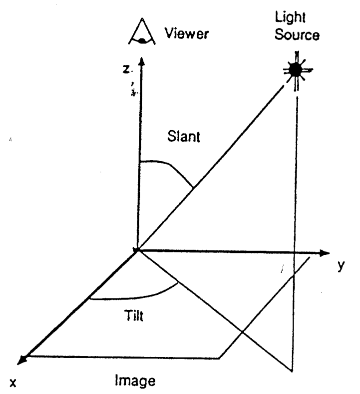

Figure 1. Slant and tilt definitions.

tives of image intensity (Lee and Rosenfeld, 1985).

— Contour Method: Assuming either the presence of an occluding boundary, or a boundary where the surface slant is constant and the surface tilt is perpendicular to the boundary, an estimate for light source tilt is expressed in terms of intensity around a closed boundary (Zheng and Chellappa, 1990; see also Zheng and Chellappa, 1991).

— Voting Method: This is a variation on the method of Pentland, where local changes of intensity are used to give a local estimate of tilt. By computing the mean of the sine and cosine of the local estimates for tilt, an overall value for light source tilt is obtained which is not biased towards larger intensity changes (Zheng and Chellappa, 1990).

Slant Estimation

— Pentland Method: By assuming that the surface approximates a sphere, an estimate for slant is derived from the calculated tilt and the variance of image intensity changes in several directions (Pentland, 1982).

— Lee and Rosenfeld Method: Assuming again that the surface approximates a sphere, an estimate for slant is derived from the mean intensity and the mean squared intensity over the illuminated portion of the surface (Lee and Rosenfeld, 1985).

— Zheng and Chellappa Method: This method is a modification of the slant estimator of Lee and Rosenfeld which additionally takes into account the portion of the image in shadow (Zheng and Chellappa, 1990).

— Disc Method: This method originates from an attempt to rationalise some of the arguments of Pentland. An estimate for slant is derived from the variance of intensity changes in a chosen direction, and the mean intensity changes along orthogonal axes within a circular subregion of the image which is free of shadow (Gibbins et al, 1991).

— Shadow Method: Assuming that the object approximates a sphere, a simple slant estimator is derived from the proportion of the surface in shadow as compared to that which is illuminated. This method is used in our analysis as a lower-bound technique that other methods ought to better (Gibbins et al, 1991).

2.2 Implementation Aspects

Each of the methods was tested on a number of images of simple smooth surfaces. Images were generated synthetically under the assumption that a point light source illuminates a Lambertian surface. Noise and ambient lighting were absent. The chief aim was to investigate the performance of the various methods when confronted with images of surfaces whose shapes were at variance with assumptions made in the derivations of the methods. It was felt that consideration of noise, ambient lighting, non-Lambertian reflectance and other factors would, at this stage, obfuscate our investigations. Most images were synthetised via orthographic projection onto a 32 x 32 grid, with intensity values in the continuous range [0, 1].

The shapes employed were:

— Sphere: this being an ideal surface for all of the methods.

— Ellipsoids of various eccentricities: these constituting a relatively mild departure from the sphere.

— Ellipsoid with an added lump: this being highly asymmetric.

— Stretched Gaussian: a Gaussian-like solid of revolution stretched along a horizontal axis, departing from the aforementioned surfaces in having a flattish periphery.

— Crater: this exhibiting a large concavity, and also being surrounded by a flattish periphery.

Figure 2 displays the various shapes, with the exception of the sphere.

First-difference approximations were used to compute derivatives of image intensity needed by most of the methods. Expectations were computed by summing values over an appropriate region and dividing by the number of points sampled. Four image directions were used in applying Pentland's tilt estimator, one in each of the directions of the x and y axes, and two others at 45° to these. It is

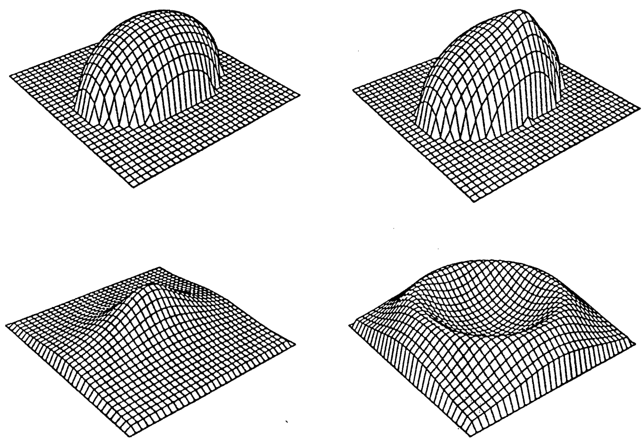

Figure 2. Shapes used in the performance analysis: Ellipsoid, Ellipsoid with lump, Stretched Gaussian and Crater.

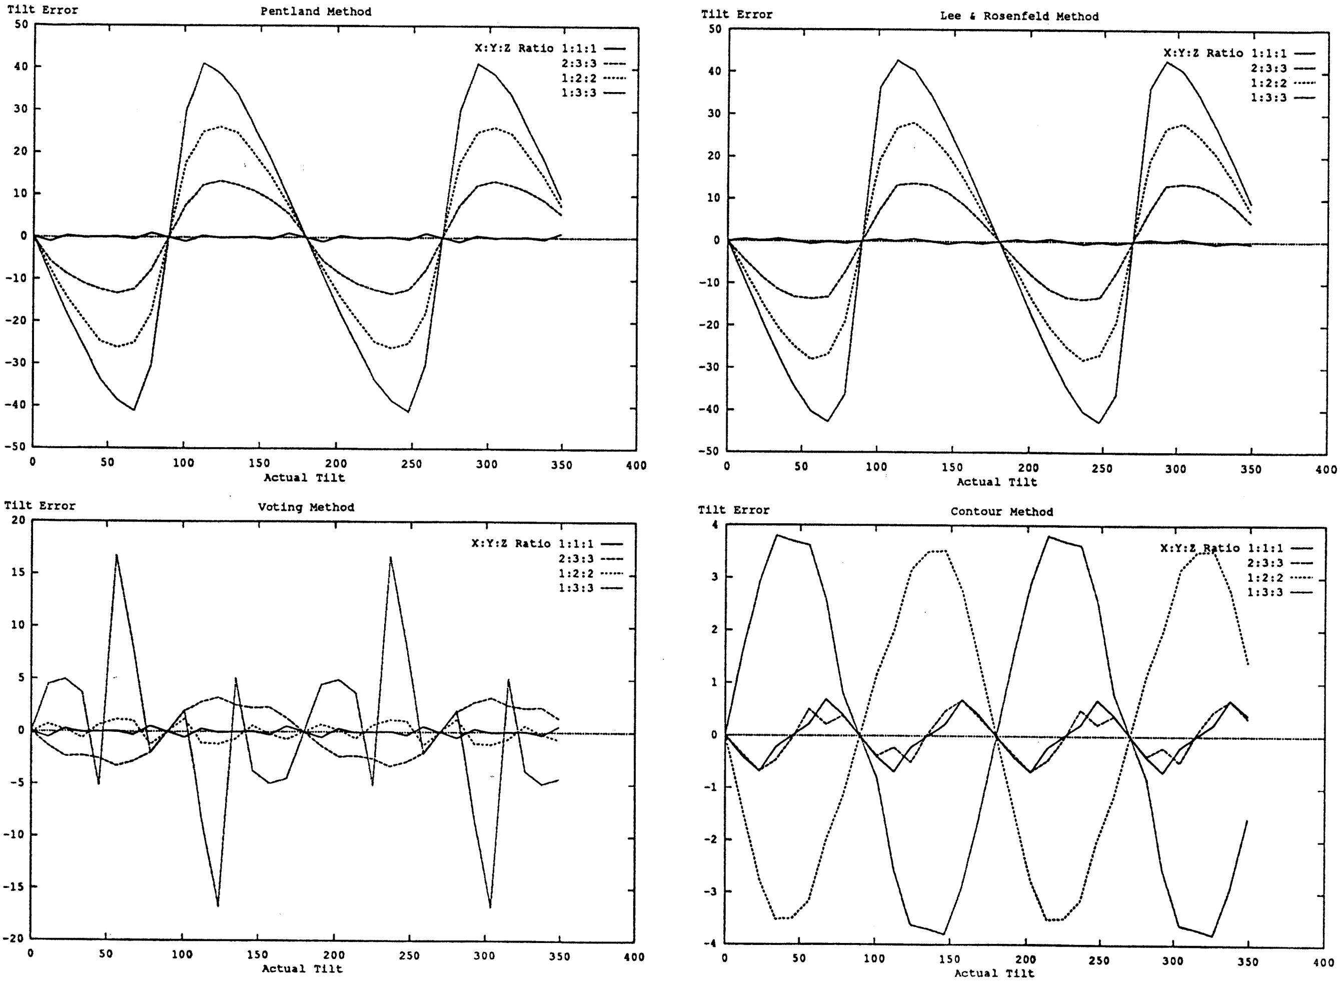

Figure 3. Tilt error against actual tilt for ellipsoids of varying eccentricity. Slant is fixed at 45 degrees, errors are shown in degrees.

worth noting that only two directions are actually required here, and if only two are used, then the method is equivalent to the tilt estimation method of Lee and Rosenfeld.

2.3 Estimating Tilt

We now examine results achieved by the Pentland, Lee and Rosenfeld, Contour, and Voting methods in determining the tilt component of the source direction.

Test 1: Response profiles for sphere and ellipsoids

In these tests, each method was applied to an image of a sphere, as well as to images of three ellipsoids. The principal *x, y, z* axes of the ellipsoids were in the ratios 2:3:3, 1:2:2, and 1:3:3. (Of course, the sphere may be regarded as having axis ratios of 1:1:1.) Source slant was fixed at 45°, and the angle of tilt was varied from 0° to 360°. Graphs are shown in Figure 3 of actual tilt against error in the estimate of tilt for each of the surface shapes and methods mentioned above. As is to be expected, all methods performed well on the image of the sphere. However, in all cases the estimate of tilt worsened as ellipsoid eccentricity increased. The estimators of Pentland, and Lee and Rosenfeld produced errors of up to 40°. Both the Voting and Contour methods worked relatively well, with the Contour method being least affected by variation in eccentricity, since it only uses points near the boundary.

Test 2: Error surfaces for an ellipsoid

Attention was then turned to an image of an ellipsoid with axes in the ratios 5:8:7. Results for this test are given in Figure 4. Here, error surfaces are displayed, each having a circular (unit-disc) domain. A light source direction is associated with each point in the domain. This direction coincides with that of the surface normal of the unit hemisphere corresponding to the given point in the domain. Each height value of the error surface specifies the error in estimated tilt. The error surface thus arises out of tests on hundreds of images, each generated under a different light source direction. These directions were varied through the range of possible slant and tilt values. The ideal error surface is a flat disc in the zero height plane. Note that height is scaled so that an error surface corresponding to a unit hemisphere would have a maximum error of 1 radian (approximately 57°). Note also that when the slant is zero,

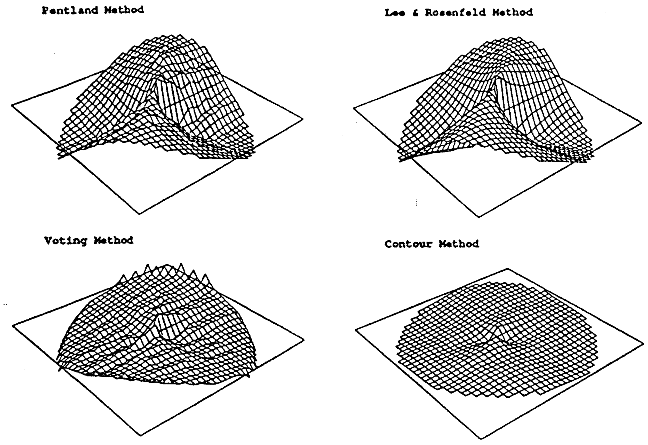

Figure 4. Tilt error surfaces describing the performance of various methods when confronted by images of an ellipsoid (see Section 2.3).

the tilt angle is undefined. An interpolated artificial value has been given at each surface centre. A spike nevertheless appears in some diagrams due to a tendency for the error to increase markedly as the centre is approached. As can be seen from the diagrams, the methods of Pentland, and Lee and Rosenfeld performed poorly at certain light source tilts, independently of slant; nevertheless, the error did not exceed 20°. Once again, the Voting and Contour methods performed reasonably well.

Test 3: Lumpy ellipsoid

Here the previous test was repeated with the only difference being the nature of the imaged surface shape. This time a smooth lump was added to the ellipsoid (refer to Figure 2). The lump was placed away from the centre so as to enhance the asymmetry of the surface. The aim here was to see how performance might be affected by additional surface undulation and asymmetry. The results illustrated in Figure 5 are not dissimilar to those of the previous

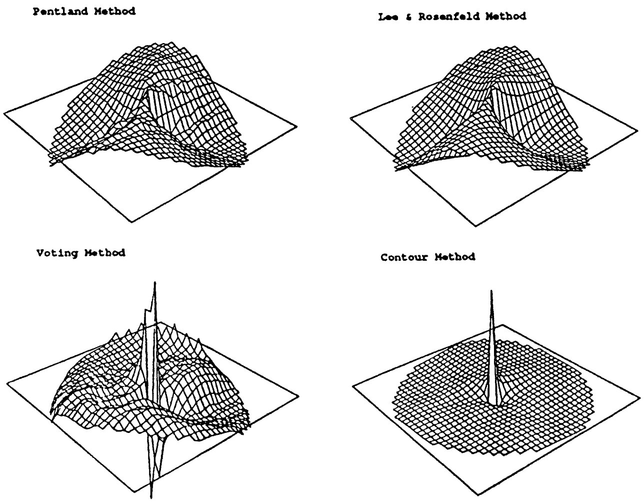

Figure 5. Tilt error surfaces describing the performance of various methods when confronted by images of an ellipsoid with lump.

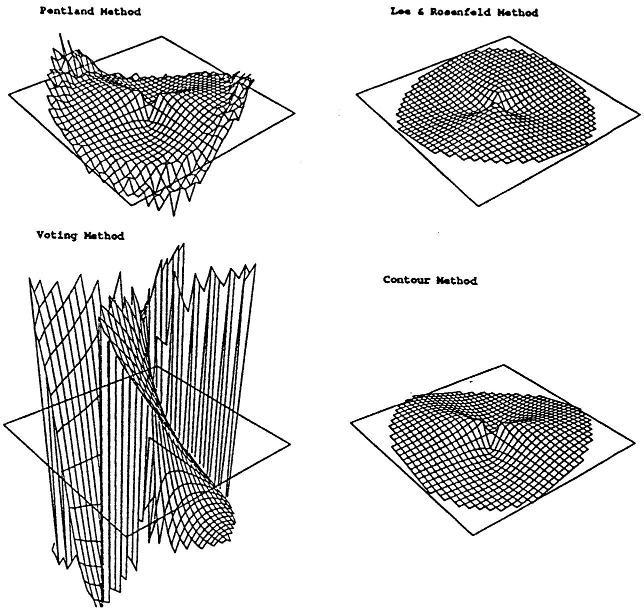

Figure 6. Tilt error surfaces describing the performance of various methods when confronted by images of a stretched Gaussian surface.

test, with the exception of the Voting method whose performance deteriorates.

Test 4: Stretched Gaussian and Crater

The methods gave almost ideal results for the Crater surface (not illustrated), and quite good results for the Stretched Gaussian (see Figure 6), with the notable exception of the Voting method, which performed very poorly.

2.4 Local Slant Estimation

The tests carried out on slant estimation were analogous to those for tilt estimation. One difference, however, was that the ellipsoids were produced by stretching a sphere in the direction of the z-axis instead of the x-axis. We compare the methods of Pentland, Lee and Rosenfeld, Zheng and Chellappa, along with the Disc method and the Shadow method. Recall that the last of these methods is very naive and is intended to give a lower bound on performance that other techniques ought to better.

Test 1: Response profiles for sphere and ellipsoids

Slant error against actual slant is displayed in Figure 7 for a variety of ellipsoids as well as the sphere. The various estimation techniques exhibited quite different responses, even when confronted with the simple sphere. Thus, for example, Pentland's estimator deteriorated (with an error of up to 40°) as light source slant was increased, while Lee and Rosenfeld's estimator deteriorated as slant decreased. A problem with the Pentland estimator is that as light source slant increases, so too the portion of the surface in shadow increases, with the consequence that the assumption of isotropic distribution of normals becomes less and less valid.

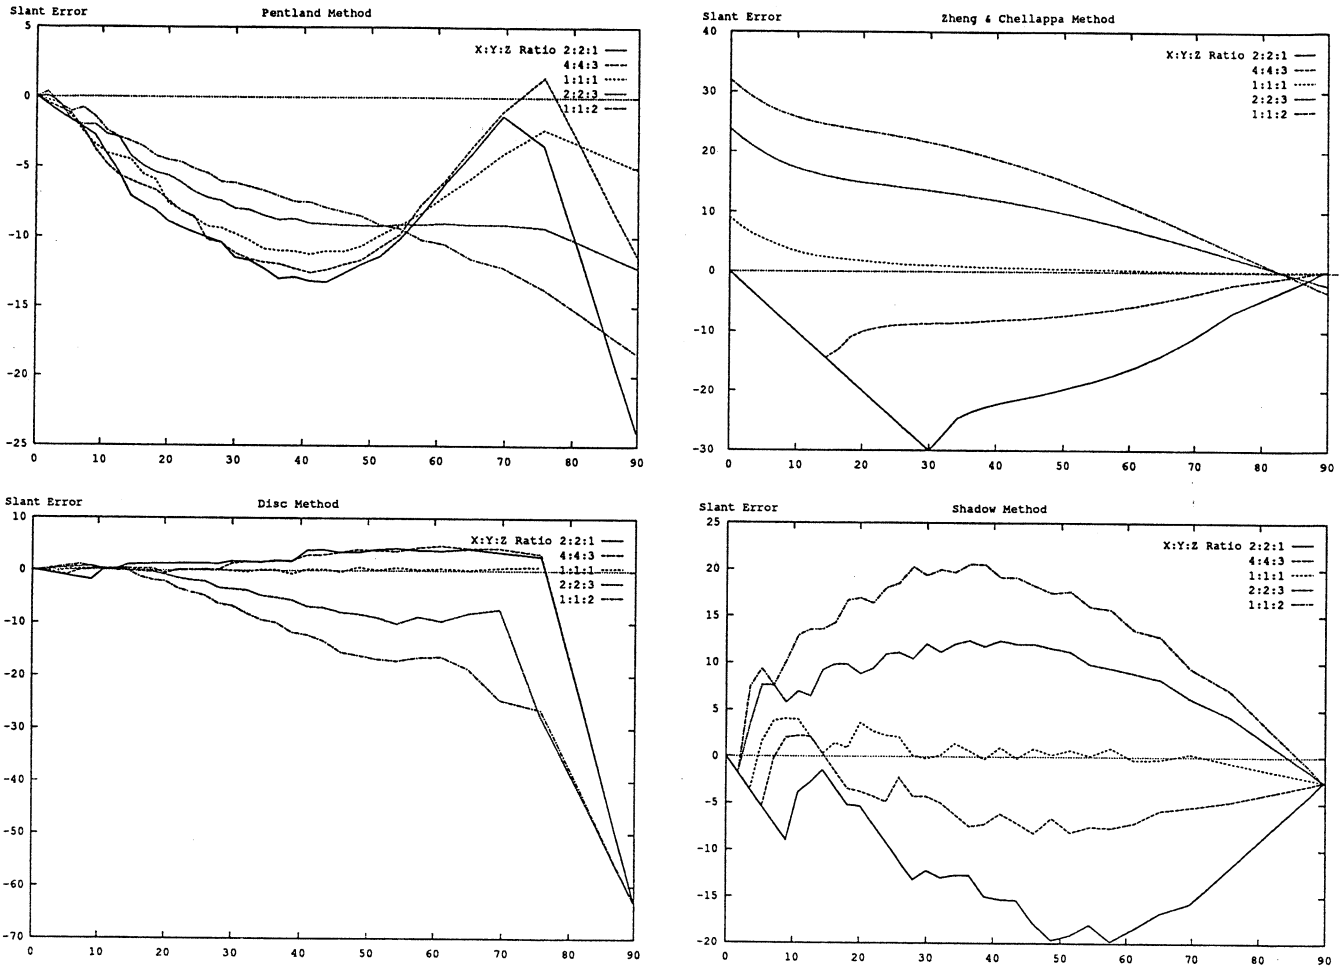

Figure 7. Slant error against actual slant for ellipsoids of varying eccentricity. Tilt is fixed at 0 degrees, errors are shown in degrees.

The Disc estimator performed best of all, except in the case of large slant when there were insufficient illuminated points in the disc for the estimates of image derivatives to be reliable. This estimator is designed precisely to overcome the aforementioned problem with the Pentland method. The Shadow estimator performed well, but with some instabilities when few shadow points existed in the image, this occurring with small slant.

Figure 7 also shows responses to ellipsoids of various kinds. Pentlant's method remained relatively stable over the range of eccentricities, with errors not worsening much over those recorded for the sphere. On the other hand, the methods of Lee and Rosenfeld, and Zheng and Chellappa were very sensitive to the changes in eccentricity. Note that only the better performing Zheng and Chellappa method is illustrated, this yielding small errors for high values of slant. However, with an ellipsoid of eccentricity 2, the error exceeded 25° for any slant of less than 15°. The Disc method showed deterioration as slant was increased.

The Shadow estimator did not perform as well over midrange slant values (20°-60°) as at the extremities of slant. Interestingly, however, the errors exhibited were comparable with those obtained for the Zheng and Chellappa method; indeed, the Shadow method gave less significant errors in the case of ellipsoids with larger eccentricity.

Test 2: Error surfaces for an ellipsoid

Here we again examine the performance of the methods when applied to images of an ellipsoid generated under a range of source directions. (Again, these directions were varied through the range of possible slant and tilt values.) Figure 8 clearly illustrates the poor performance of the Lee and Rosenfeld estimator as compared to Zheng and Chellappa's improved version. Performance of the other estimators was quite good.

Test 3: Lumpy ellipsoid

When confronted with an image of an ellipsoid with protuberance, all of the methods suffered some degradation in performance (see Figure 9). Pentland's method gave poor estimates of steeper light source slants. The Zheng and Chellappa method performed somewhat better. The Disc method struggled severely in this test, producing

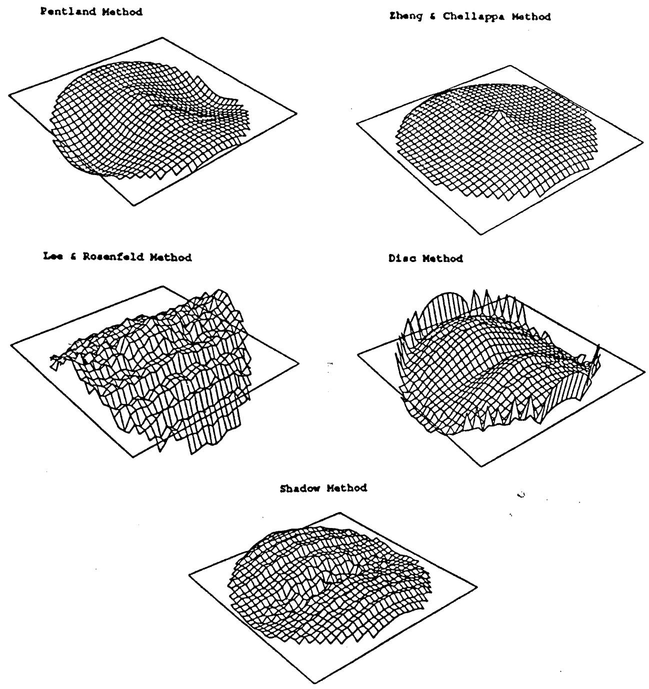

Figure 8. Slant error surfaces describing the performance of various methods when confronted by images of an ellipsoid.

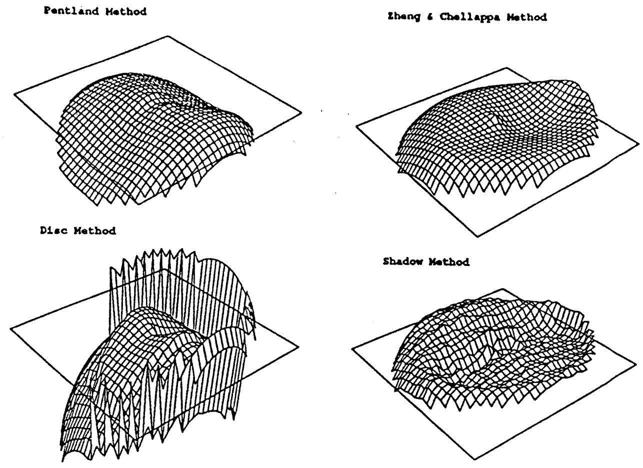

Figure 9. Slant error surfaces describing the performance of various methods when confronted by images of an ellipsoid with lump.

highly unreliable estimates for steeper slants. Perhaps surprisingly, the Shadow method coped well with images of the test surface, again performing much like the Zheng and Chellappa method.

Test 4: Stretched Gaussian and Crater

Here, the Gaussian and Crater images were used, and all of the methods failed to produce remotely tolerable results (consequently no figures are presented). The assumptions implicit in the derivations of all methods are violated in the case of these surface shapes.

2.5 Comments

It is evident from the testing that each method has domains in which it works well, and domains in which it performs badly. All of the methods are essentially based on the assumption that the depicted surface is spherical. As mentioned previously, this strong assumption is made in order to render tractable a highly ill-posed problem. When the assumption is violated by a given surface, then estimates of source direction tend to degrade. (Moreover, estimates are similarly unreliable when an imaged surface satisfies the assumption, but is partially occluded from view by other objects, or is only partly visible due to image clipping.) Also, determination of slant appears to be a more difficult problem than determination of tilt, which is not surprising given that the former is not confined to the image plane.

Comparison of Lee and Rosenfeld's method with the Zheng and Chellappa improvement shows clearly that the latter is more stable and accurate. The two methods differ only in that the latter incorporates knowledge of the portion of the region in shadow, the improved performance therefore being due to the adoption of this additional information.

It is interesting to reflect on the performance of the Shadow slant estimator. Although it is much simpler than methods such as those of Pentland, and Lee and Rosenfeld, it nevertheless produces comparable results. This is despite the fact that it uses considerably less data than its counterparts, making no recourse to image intensity values or derivatives, and being simply based on the portion of the object's image which is in shadow. That a naive method performs comparably to the more sophisticated techniques raises obvious concerns.

It is possible that further improvement in estimates might be obtained by judicious combination of various techniques. Nevertheless, it is hard to envisage an automated technique having true versatility, giving good responses for a wide range of surface types. For example, none of the techniques considered is tuned to the situation in which an essentially planar landscape exhibits undulations. This is commonly found in aerial photographs of terrain. Horn (1990) has suggested that a useful slant estimator here might be the inverse cosine of intensity averaged over the image.

3 TESTING HUMAN PERFORMANCE

3.1 Previous Work

First we consider some psychological studies relevant to shading analysis.

Ramachandran (1988) suggests that shading is perhaps one of the most primitive cues for depth, having been developed early in the evolution of animal life. Thus it may be that some fish exhibit "counter-shading" in order to neutralise depth-perception and, presumably, to reduce the risk of falling prey to predators. It also appears that shape information obtained from shading can be used as

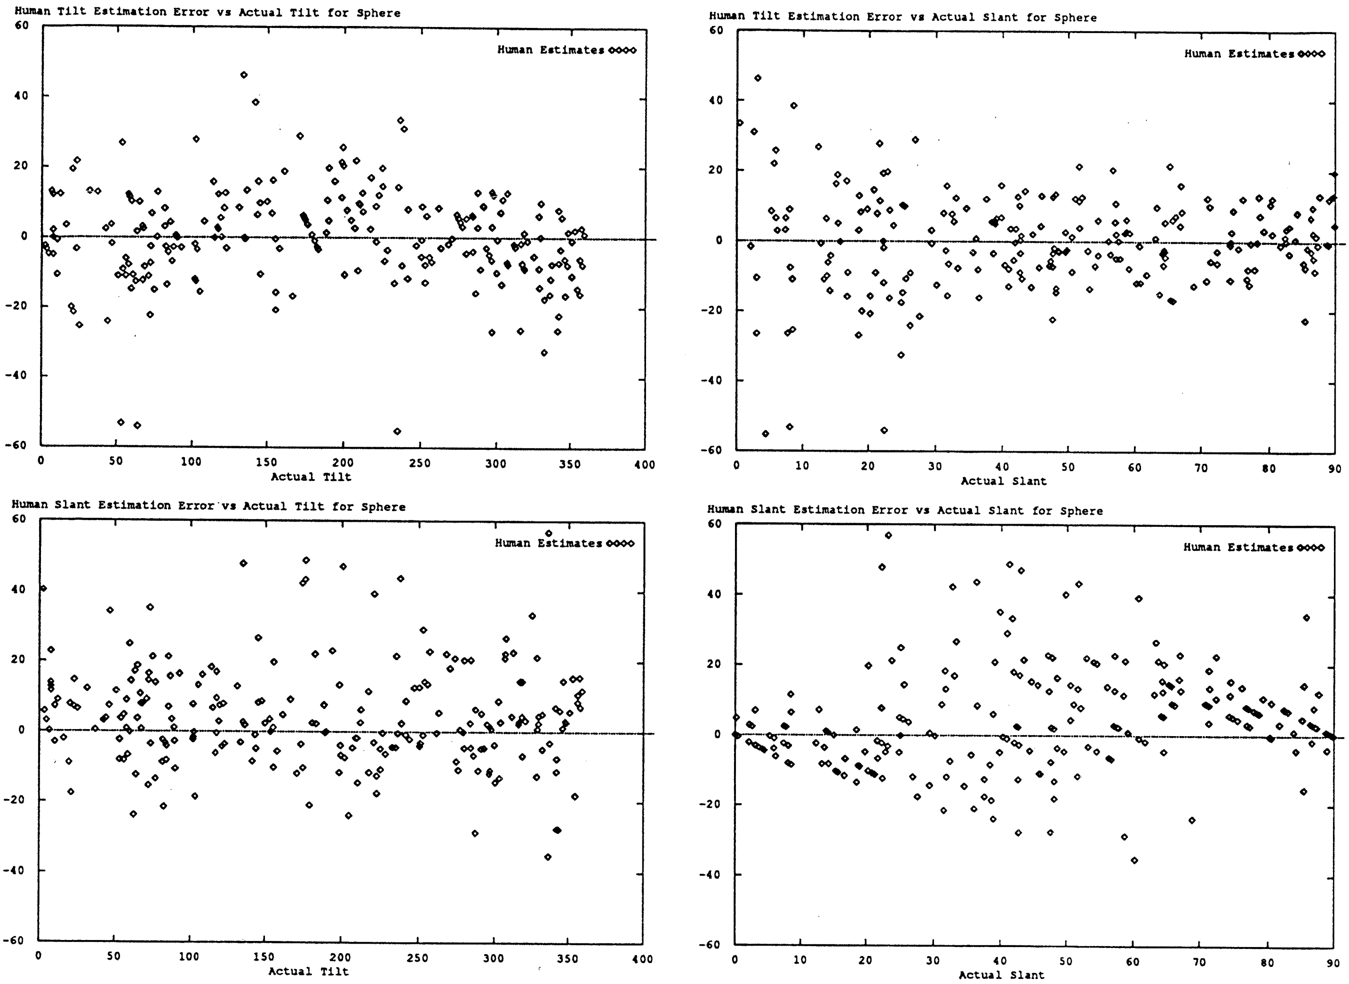

Figure 10. Human performance analysis for images of a sphere. Results show: Tilt error vs tilt, tilt error vs slant, slant error vs tilt, and slant error vs slant.

input to the processing of motion data, and the determination of figure-ground relations. Contours are shown to be an important influence on the perception of shape from shading. Also, as is well known, the human visual system appears to have a predilection for a single overhead source, in the sense that shape interpretations of multiple objects in a scene are generally kept consistent with an assumption of a single sun in a specific direction.

Berbaum et al. (1983) considered the influence of an image of a well-known object (such as a human hand) on the perception of a possibly ambiguous image of an unknown object. Again, it appears that humans generally achieve consistency by having the perception of the familiar surface affect the interpretation of the unknown object. Note that the image of the human hand might be considered similar in information content to the reflectance map of Horn and Sjoberg (1979). Further studies are conducted on the memorising of source direction and its conditioning effect on the subsequent interpretation of a sequence of images.

It is in the work of Mingolla and Todd (1986) that contributions most relevant to our concerns may be found. Here, the authors observe that subjects tend to show low correlation between the quality of their source-direction estimates, and the quality of their shape estimates. That is, on a given test, a good estimate may be obtained for source direction, yet shape might be estimated poorly. Alternatively, shape may be well estimated, but the source directon might be substantially in error. This, the authors claim, argues against automated techniques requiring prior knowledge of the light source direction. The claim is also made that Lambertian surface reflectance appears not to be a default assumption in human perception. This is because the perception of shape seems almost unaffected by the addition of an arbitrarily significant specular component to an underlying Lambertian model. From these observations, Mingolla and Todd speculate that the human visual system is not based on an inverting of the image-forming process. The (global) analysis of contours of constant intensity due to Koenderik and van Doorn (1980) is cited as an example of an approach that is perhaps more consistent with human performance.

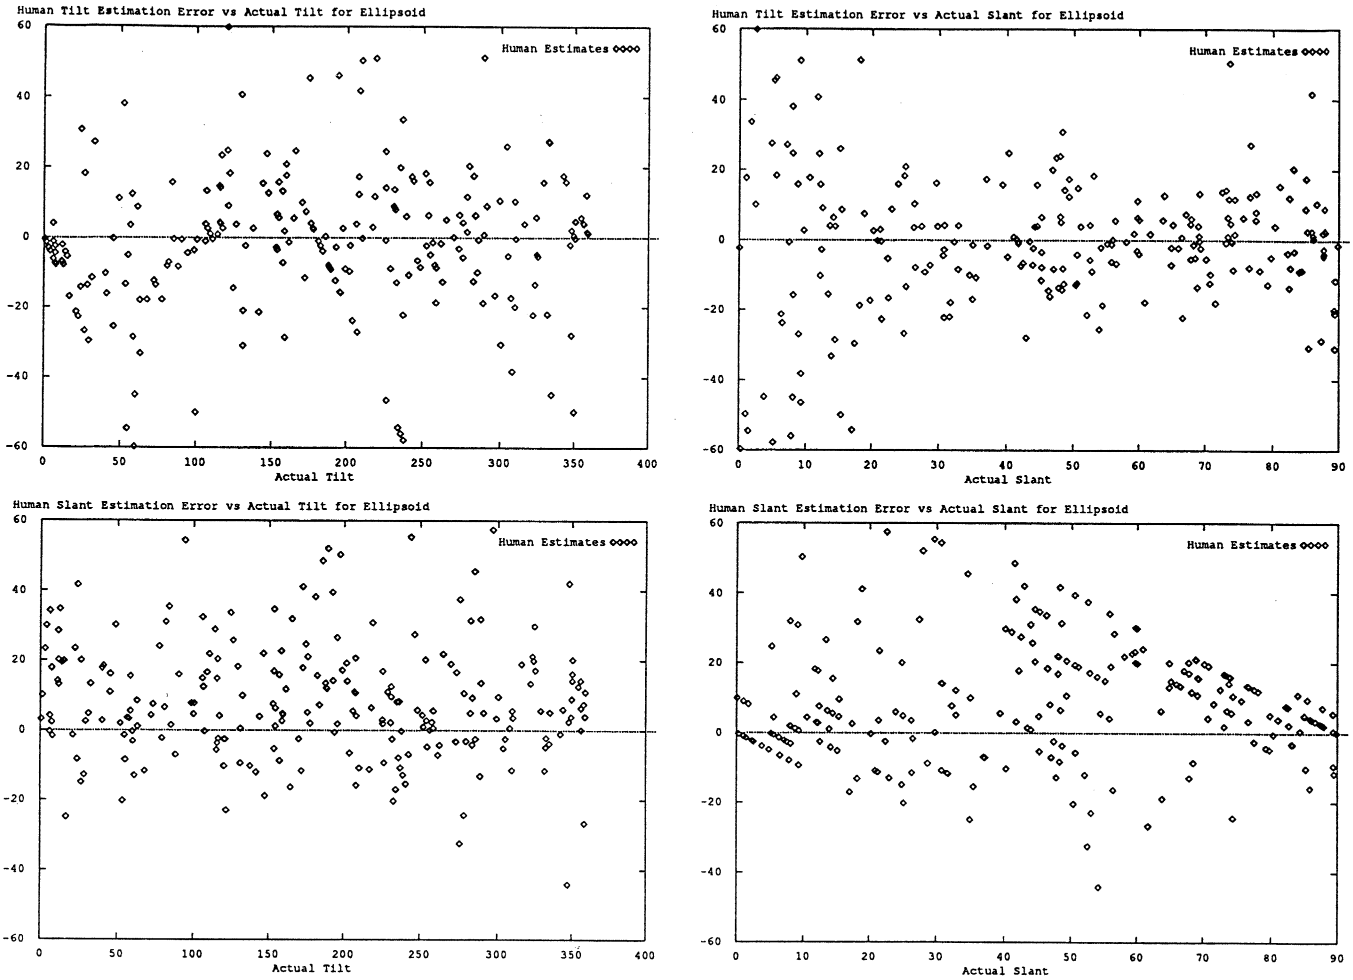

Figure 11. Human performance analysis for images of an ellipsoid. Results show: Tilt error vs tilt, tilt error vs slant, slant error vs tilt, and slant error vs slant.

3.2 The Experimental Environment

Informal tests were carried out on the ability of people to estimate light source direction from various images of a sphere and an ellipsoid. Results were obtained from sessions conducted with 24 individuals. Subjects were asked to examine a given synthetic image, and to gauge the direction of the sun. They were trained to express their results in terms of slant and tilt, and were provided with diagrams useful for estimating angles. Each subject examined twenty 512 x 512 images of Lambertian surfaces presented in random order (10 spheres and 10 ellipsoids). In all, therefore, 480 human estimates of source direction were obtained. The ellipsoid used in the tests was the same as that used previously to test the automated methods.

3.3 Performance Analysis

Various results were obtained in response to an image of a sphere, and these are shown in Figure 10. The following informal observations are made in relation to the various plots:

— Tilt estimation error against correct tilt: the graph suggests that estimation of tilt is unaffected by rotation of the image.

— Tilt error against correct slant: as the "sun goes down" tilt error seems to decrease. Of course, it should be recalled that tilt is undefined at zero slant.

— Slant error against tilt: here it appears that slant estimation is invariant of image rotation.

— Slant error against correct slant: in this test, the variance of the estimates appears to decrease for small or large slants. There would also seem to be a tendency to underestimate small slant values, and to overestimate large slant values.

Tests were also carried out on images of the ellipsoid, with similar results (Figure 11).

Finally, histogram-type graphs were plotted of estimation error against the frequency of making that error (5° ranges were used). One diagram deals with slant error, the other with tilt error (see Figure 12). Each diagram plots a graph for the sphere and the ellipsoid. A mean estimation error of close to zero is indicated in each case. The variance remains reasonably small, increasing somewhat in

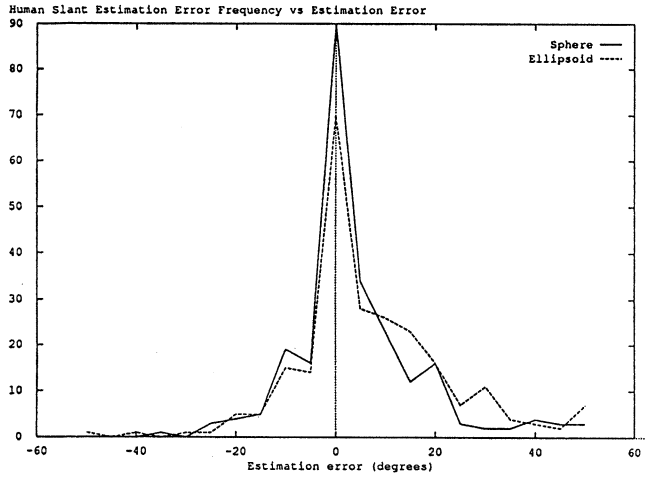

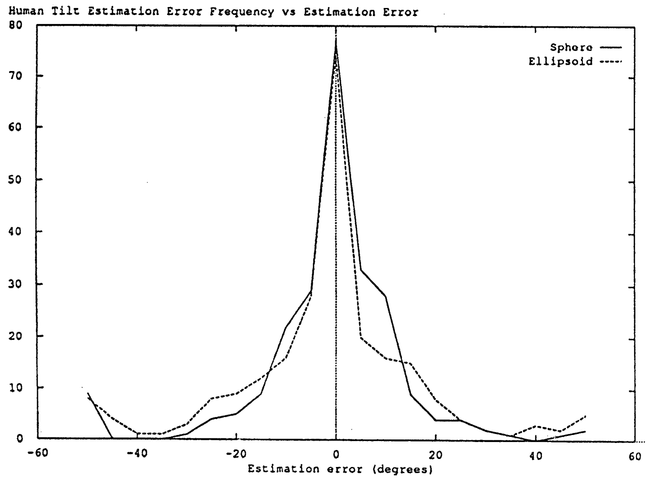

Figure 12. Histogram-like plots of error frequency against estimation errors in slant and tilt, respectively.

the case of the ellipsoid. Other informal tests not recorded here suggest that people respond reasonably well in a wide variety of situations, including, for example, aerially-viewed terrain.

4 CONCLUSIONS

We now return to the main question posed earlier: Might prior determination of source direction be a useful approach to reducing the prerequisite needs of computational schemes, and might this also be the strategy employed by the human visual system?

An obvious caveat needs to be registered from the outset in relation to psychological studies in which the subject is asked to report on such as perceived light source direction. Perhaps it is the case that humans determine source direction prior to shape. Another possibility is that recovery of source direction and shape is a coupled process, with an intermingling of activity. Whichever strategy is employed, it may well be that the subject has no conscious access to the processes involved. If this is correct, a subject confronted with questions on source direction or shape will be basing answers upon the outcome of the complete process of shape and source recovery, not the process itself. Worse still, it is even possible that the method used by humans to introspectively compute light source direction from an image is different from the method that humans use unconsciously during recovery of shape from shading. It might even be the case that introspective computation of light source direction uses the already determined shape of the object. Clearly, it is essential that results of psychological tests are interpreted with this in mind.

Returning to the low correlation observed by Mingolla and Todd between the quality of estimates of source direction and shape, it is useful to consider the Brooks and Horn (1985) method, which is the only approach to date that seeks to determine, in a coupled manner, both shape and source direction from shading information. In this technique, a point light source is assumed to illuminate a smooth Lambertian surface. A dual scheme is employed in which, repeatedly, shape is computed using an estimate of source direction, and then a new source direction is computed from the revised shape. The search space for this scheme is not entirely convex, and so the technique is susceptible to falling into local minima. It is quite possible with this method that there will be a difference in the quality of estimates for source direction and shape. Of course, it is also possible that both estimates will be poor, or both estimates will be good, these presumably being possibilities that also arise in human performance.

Given that the Brooks and Horn method can deliver estimates for shape and source-direction of differing quality, we believe that the Todd and Mingolla conclusion that present shape from shading techniques are wholly inappropriate models of cognition is stronger than the evidence will bear. (What is clear, however, is that the performance of these techniques is poor in comparison with human perception, but this is a separate point. We are concerned here with the possible validity of the methodology.) Our view is that the question of whether such an approach could be refined and developed into an adequate model remains open. A further complicating factor is whether the Koenderink and van Doorn (1986) approach favoured by Mingolla and Todd is indeed more consistent with the observed human differences in quality of estimates, especially given that there is no shape from shading implementation based on this analysis. We suggest that similar conjectures of Mingolla and Todd concerning the inadequacy of both the assumption of Lambertian-ness, and the use of "local" units of shading data (intensities or intensity gradients, as opposed to, say, global contours of constant intensity) are also contentious.

On the basis of data before us, early indications are that it will be extremely difficult (and perhaps impossible) to generate a robust automated technique for determining source direction from images exhibiting a wide range of surface types. This is because the strong assumptions needed in the derivation of such a technique are unlikely to be appropriate to a wide class of shapes. If a robust preprocess for determination of source direction turns out to be infeasible, it might then be that the relatively versatile skills of the human subject are due to a coupled approach to recovery, in which shape and source are sought in parallel. Further psychological work is clearly needed to offer more clues.

ACKNOWLEDGEMENTS

The authors are grateful for the valuable advice given by the referees. This research was in part funded by the Australian Research Council.

REFERENCES

BERBAUM, K., BEVER, T. and CHUNG, C.S. (1983): Light Source Position in the Perception of Object Shape. *Perception* 12, pp. 411-416.

BLAKE, A., ZISSERMAN, A. and KNOWLES, G. (1985): Surface Descriptions from Stereo and Shading. *Image and Vision Computing* (4) 3, pp. 183-189. Also in Horn and Brooks, 1989, pp. 29-52.

BRADY, M. (1979): Inferring the Direction of the Sun From Intensity Values on a Generalised Cone. In *Proceedings of the International Joint Conference on Artificial Intelligence,* Los Angeles, CA, August, pp. 88-91.

BROOKS, M.J. and HORN, B.K.P. (1985): Shape and Source from Shading. In *Proceedings of the International Joint Conference on Artificial Intelligence,* Los Angeles, CA, August, pp. 932-936. Also in Horn and Brooks, 1989, pp. 53-68.

GIBBINS, D., BROOKS, M.J. and CHOJNACKI, W. (1991): Determining light source direction from a single image. Technical Report TR91-26, School of Information Science and Technology, Flinders University.

HORN, B.K.P. (1975): Obtaining Shape from Shading Information. In *The Psychology of Computer Vision.* P.H. Winston (ed). McGraw-Hill, New York, pp. 115-155. Also in Horn and Brooks, 1989, pp. 123-172.

HORN, B.K.P. (1990): Private communication.

HORN, B.K.P. and BROOKS, M.J. (1986): The Variational Approach to Shape from Shading, *Computer Vision, Graphics and Image Processing* (2) 33, pp. 174-208.

HORN, B.K.P. and BROOKS, M.J. (eds) (1989): *Shape from Shading.* MIT Press, Cambridge, Mass.

HORN, B.K.P. and SJOBERG, R.W. (1979): Calculating the Reflectance Map. *Applied Optics* (11) 18, pp. 1770-1779. Also in Horn and Brooks, 1989, pp. 215-244.

KOENDERINK, J.J. and VAN DOORN, A.J. (1980): Photometric Invariants Related to Solid Shape. *Optica Acta* (7) 27, pp. 981-996.

LEE, C.H. and ROSENFELD, A. (1985): Improved Methods of Estimating Shape from Shading Using the Light Source Coordinate System. *Artificial Intelligence* (2) 26, pp. 125-143. Also revised version in Horn and Brooks, 1989, pp. 323-347.

MINGOLLA, E. and TODD, J.T. (1986): Perception of Solid Shape from Shading. *Biological Cybernetics* 53, pp. 137-151. Also in Horn and Brooks, 1989, pp. 409-441.

PENTLAND, A.P. (1982): Finding the Illuminant Direction. *Journal of the Optical Society of America* (4) 27, pp. 448-455.

RAMACHANDRAN, V.S. (1988): Perceiving Shape from Shading. *Scientific American* (2) 259, pp. 58-63.

TODD, J.T. and MINGOLLA, E. (1983): Perception of Surface Curvature and Directions of Illumination from Patterns of Shading. *Journal of Experimental Psychology: Human Perception and Performance* (4) 9, pp. 583-595.

ZHENG, Q. and CHELLAPPA, R. (1990): A Robust Algorithm for Inferring Shape from Shading. USC-SIPI Report No. 159, Signal and Image Processing Institute, Department of Electrical Engineering Systems, University of Southern California.

ZHENG, Q. and CHELLAPPA, R. (1991): Estimation of Illuminant Direction, Albedo, and Shape from Shading. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition,* Lahaina, Maui, Hawaii, June, pp. 540-545.

BIOGRAPHICAL NOTES

Danny Gibbins is a Tutor and PhD student in the Discipline of Computer Science at Flinders University. His research is concerned with reducing prerequisite requirements of vision algorithms.

Michael Brooks is an Associate Professor in the Department of Computer Science at Adelaide University, leading a research group investigating a range of vision problems.

Wojciech Chojnacki is a Professor of Mathematics in the Institute of Applied Mathematics and Mechanics at Warsaw University. He is currently a Senior Research Fellow in the Department of Computer Science at Adelaide University, working on the shape-from-shading problem in computer vision.