Estimation of the Light Source Distribution and its Use in Integrated Shape Recovery from Stereo and Shading *

Darrell R. Hougen and Narendra Ahuja

Beckman Institute and Coordinated Sciences Laboratory
University of Illinois, Urbana, IL 61801

Abstract

This paper is concerned with estimation of the light source distribution and its use in surface shape estimation by the methods of stereo and shading. There are many advantages to integrating stereo and shading information. However, shape from shading algorithms are limited in their applicability by the assumption of overly simplistic models of the light source distribution. A more complete representation is described in which the lighting model makes use of multiple fixed point sources located at infinity. Methods of estimating the model parameters are developed and a method of estimating the surface shape given the albedo and source distribution is presented. The shading algorithm is combined with a stereo algorithm in an integrated approach that is designed to handle albedo variations through the use of color images.

1 Introduction

This paper is concerned with estimation of the light source distribution and its use in surface shape estimation by the methods of stereo and shading. There are many advantages to combining stereo and shading information in an integrated approach [9,7]. However, shape from shading algorithms are limited in their applicability by the assumption of overly simplistic models of the light source distribution.

Several methods have been developed for recovering the direction of a single point light source [18,2,13]. A method has also been developed for recovering the directions of a small number of point light sources [17]. At this date, however, no method exists for estimating more complex light source distributions.

This paper describes a more complete representation of the light source distribution that makes use

of multiple point sources located at infinity. Methods of estimating the parameters are developed and described and a method of estimating the surface shape given the source distribution is presented. The shape estimation algorithm is based on a new, single source algorithm that is an improvement over existing shape from shading algorithms.

The new shape from shading algorithm is combined with a stereo depth estimation algorithm to form an integrated approach. The competing requirements of shading and stereo lead to a combined approach that involves image segmentation and albedo estimation. Color images are used to facilitate the albedo estimation process. This approach has recently been used by Hougen and Ahuja [7] and Funt et al [4].

2 Source distribution from shading

The method described in this section is designed to address the problem of complex lighting. However, it is still necessary to make several assumptions. It is assumed that all sources are located at infinity. Surface patches are assumed not to shadow each other. The problem of interreflections is not handled. This limitation can be partially overcome by using the method of Nayar, Ikeuchi and Kanade [11]. Finally, this method works better if the reflectance function is not highly specular.

Beyond the assumptions stated above, the validity of the derivation depends only upon the fact that the luminance appears as a multiplicative constant in the reflectance function. It is a consequence of basic electromagnetic theory [10] that this assumption holds for all reflectance functions corresponding to real objects or scenes. Examples are provided by the models of Beckmann and Spizzichino, Torrance and Sparrow, Healey [5], Wolff [16], and Nayar et al [12].

For a surface region illuminated by a continuous light source located at infinity, the total image irradi-

ance is

$$E(x,y) = \rho(x,y) \int_{\Theta} R(n(x,y), l(\theta)) d\Omega$$

where ρ is the albedo, R is the reflectance map written as a function of the normal, n, and the incident light source distribution function, l, the angular position parameters are represented by θ and the integral is taken over the sphere of possible incident directions.

The light source distribution function may be approximated by a set of m distinct light source vectors, $l_1, l_2, \ldots, l_m$ where l_s is an average value of $l(\theta)$ over a neighborhood of θ_s. Each source vector may be written $l_s = \lambda_s \hat{l}_s$ where $\hat{l}_s$ is a unit vector pointing in the same direction as l_s and λ_s is the intensity of the sth source. Now define the intensity independent reflectance function, $\tilde{R}$, by, $R(n,l) = \lambda \tilde{R}(n, \hat{l})$. This definition is possible because λ is guaranteed, to appear as a multiplicative constant in R. Let E be measured at k image points. If a term representing the contribution of ambient light is included, the resulting image irradiance equation is, for $j = 1, \ldots, k$,

$$E_j = \rho_j(\lambda_0 \tilde{R}_0 + \sum_{s=1}^{m} \lambda_s \tilde{R}(n_j, \hat{l}_s))$$

This is a set of linear equations for $\lambda_0, \ldots, \lambda_m$. If $k \geq m + 1$ and the equations are linearly independent, a unique solution exists to the corresponding least squares problem. In order to use this set of equations, the albedo and surface orientation must be given for at least m points. Here, an initial estimate of the surface is given by stereo and the albedo is estimated by the method explained in section 8. In addition, the unit source vectors, $\hat{l}_1, \ldots, \hat{l}_m$, must be known. Fortunately, it is not necessary to solve for the positions of the unit source vectors. In fact, the positions of the unit source vectors may be chosen by any suitable method before the computation begins.

The accuracy of the above method of approximating the source distribution is limited by the number of sources used in the representation. If the reflectance function is highly sensitive to the incident direction the estimate will only be accurate if the space of possible source directions is densely sampled. However, in the presence of noisy surface estimates, the number of parameters in the model must be kept small to avoid overfitting problems. Therefore, this approach is better suited to nearly diffuse reflectors.

2.1 Representation of the distribution

The choice of unit source vectors serves the same purpose in the present case as the choice of knot points

serves in the definition of the cubic spline. Just as it is convenient in the case of the cubic spline to establish equally spaced knot points, in the present case it is convenient to establish uniformly distributed unit source vectors. The important thing to note is that the particular choice of unit source vectors is independent of the data and may be fixed in advance.

The unit source vectors may be treated as points on a unit sphere. Unfortunately, it is not possible in general to uniformly distribute an arbitrary number of points on a sphere [3]. In the present implementation, points are placed on the hemisphere of directions on the camera side of the image plane by an ad hoc method that is designed to approximate uniformity. A more optimal distribution might be obtained by estimating the probability density of incident radiant energy as a function of direction for real scenes and distributing the unit source vectors accordingly.

2.2 Solving for the distribution

For the solution to the multisource equations to be physically meaningful, the intensity of every light source including the ambient source must be nonnegative. More precisely, the following relations must hold, $0 \leq \lambda_s, s = 0, \ldots, m$. This section explores the effects of these constraints on the solution.

If the only concern is the recovery of surface shape, it is possible to ignore the nonnegativity constraints. In this case the sources are called virtual sources and can be either positive or negative. Clearly, negative sources are not realizable. However, the use of virtual sources has the advantage that the solution can be found by ordinary methods of linear least squares.

If the nonnegativity constraints are enforced, the problem becomes an instance of the problem of nonnegative linear least squares [8]. Because the error surface contains no local minima, the unique global minimum is always found. This fact may be established by observing the convexity of the objective function and the convexity of the constraints. In this case the solution is constrained to lie in the positive hypercone.

Other methods of solving the multiple source equations involve letting some, but not all, of the coefficients be negative. In particular, it is sometimes convenient to let the ambient term be unconstrained. The ambient term acts to compensate real ambient light and also to adjust for a voltage gain cutoff in the digitization process. The effect of a voltage cutoff is to reduce the brightness of every pixel by a fixed amount. Allowing the ambient term to be negative allows this effect to be modelled.

To solve the source equations when some but not all of the variables in the model are constrained it is pos-

149

sible to use techniques for least squares with inequality constraints or least squares with bounded variables [8]. In the present case the ambient term is bounded by the maximum brightness so the problem can be solved by adding a negative constant to the brightness.

3 Shape from shading algorithm

This section describes a shape from shading algorithm that is based on the direct height from shading method of Leclerc and Bobick [9] but which incorporates the derivative calculations based on offset grids given in Horn [6]. The resulting hybrid method has been found to perform better than either foregoing method and can be easily converted into a multiple light source method as is done in the next section.

In the direct height from shading method of Leclerc and Bobick the surface height (or depth) is represented as a discrete array $\{z_{ij}\}$ and the objective function to be minimized is given by

$$Q = \sum_{ij}(1 - \gamma)(E_{ij} - R(p_{ij}, q_{ij}))^2 + \gamma(u_{ij}^2 + v_{ij}^2)$$

where E is the image irradiance, R is the reflectance function, p, q, u and v are the symmetric first and second finite difference approximations to the corresponding derivatives, z_x, z_y, z_{xx} and z_{yy}, and γ is a weight that is gradually reduced to zero as if it were a continuation parameter.

The above formulation lends itself to efficient numerical solutions using the conjugate gradient algorithm as has been confirmed in our laboratory. In addition, it has been found to be very stable in comparison with algorithm of Horn which seems sensitive to image noise and errors in the light source direction. Unfortunately, the particular choice of derivative approximations leads to unnecessarily large errors in the final surface estimate. These errors result from the use of the symmetric finite difference approximation to the first derivative.

The symmetric first finite difference does not contain information about the relative heights of adjacent grid elements. In one dimension, this results in two independent chains of height estimates that do not communicate with each other. In two dimensions, there are four independent sets of estimates arranged in disconnected grids. The smoothing term does contain information about the relationships between adjacent pixels so that when the smoothing parameter is large, the method works reasonably well. However, as the smoothing term approaches zero a pattern of alternating dark and light colored pixels, a "checkerboard" pattern, emerges in the reconstructed image. Even

for relatively large values of the smoothing parameter, ridges may develop in the depth map.

To avoid the difficulties associated with symmetric differences mentioned above, the derivative approximations of Horn [6], based on staggered grids, are used here. The quantities p_{ij}, q_{ij} and E_{ij} are defined on a grid that is aligned with the input image. The quantities z_{ij}, u_{ij} and v_{ij} are defined on a grid that is one row and one column larger than the input image and offset by one half pixel. In this case, the definitions of the derivatives are,

$$p_{ij} = \frac{1}{2}(z_{i+1,j+1} + z_{i+1,j} - z_{i,j+1} - z_{i,j})$$

$$q_{ij} = \frac{1}{2}(z_{i+1,j+1} - z_{i+1,j} + z_{i,j+1} - z_{i,j})$$

$$u_{ij} = \frac{1}{2}(p_{i,j} + p_{i,j-1} - p_{i-1,j} - p_{i-1,j-1})$$

$$v_{ij} = \frac{1}{2}(q_{i,j} - q_{i,j-1} + q_{i-1,j} - q_{i-1,j-1})$$

Note that because the smoothing term was never a problem, it is still possible to use the definitions of u and v given by Leclerc and Bobick. The gradient of the objective function is given by

$$\frac{\partial Q}{\partial z_{i,j}} = -2(1 - \gamma)\sum_{k,l}(E_{k,l} - R_{k,l})\frac{\partial R_{k,l}}{\partial z_{i,j}} + 2\gamma\sum_{k,l}u_{k,l}\frac{\partial u_{k,l}}{\partial z_{i,j}} + v_{k,l}\frac{\partial v_{k,l}}{\partial z_{i,j}}$$

where R_{ij} is an abbreviation for $R(p_{ij}, q_{ij})$ and the sum is from $(i - 1, j - 1)$ to (i, j). Evaluation of the partial derivatives of the reflectance function is straightforward.

The shape from shading method described in this section has been found to be faster and more reliable than the method of Horn and more accurate than the method of Leclerc and Bobick, requiring only two or three iterations to reduce γ to its minimum value.

4 Multisource shape from shading

It is a relatively simple matter to extend the single source shape from shading algorithm to the multiple source case. To simplify the following expressions, the multisource reflectance function R^M is defined by,

$$R_{i,j}^M = \lambda_0\tilde{R}_0 + \sum_{s=1}^{m}\lambda_s\tilde{R}_{i,j,s}$$

where $\tilde{R}_{i,j,s}$ is shortened notation for $\tilde{R}(p_{ij}, q_{ij}, \hat{l}_s)$. Combining the multisource image irradiance equation

150

with the smoothing term given in the previous section leads to the following multisource objective function,

$$Q^M = \sum_{ij}(1 - \gamma)(E_{ij} - R_{ij}^M)^2 + \gamma(u_{ij}^2 + v_{ij}^2)$$

Computing the gradient of this new objective function yields the following expression,

$$\frac{\partial Q^M}{\partial z_{i,j}} = -\frac{2}{n}(1 - \gamma)\sum_{k,l}(E_{k,l} - R_{k,l}^M)\frac{\partial R_{k,l}^M}{\partial z_{i,j}}$$
$$+ \frac{2}{n}\gamma\sum_{k,l}u_{k,l}\frac{\partial u_{k,l}}{\partial z_{i,j}} + v_{k,l}\frac{\partial v_{k,l}}{\partial z_{i,j}}$$

where the sum is from $(i - 1, j - 1)$ to (i, j) and the derivative of the multisource reflectance function is given by

$$\frac{\partial R_{i+k,j+l}^M}{\partial z_{i,j}} = \sum_{s=1}^{m}\lambda_s\frac{\partial \tilde{R}_{i+k,j+l,s}}{\partial z_{i,j}}$$

The derivatives of the reflectance function and the smoothing term are easily computed.

In order to obtain the surface shape, the multi-source objective function is minimized using a conjugate gradient algorithm as was done in the single source case.

5 Integration of stereo and shading

The methods of stereo and shape from shading naturally complement each other and there are several advantages to combining them [1,9,7]. First, an integrated system can operate under a wider range of conditions than either method used alone. Stereo can operate in highly textured regions but not in feature-less regions. Shape from shading can operate in featureless regions but not highly textured regions. In addition, the stereo depth estimate can be used to calculate the light source distribution, the initial conditions and boundary conditions needed by the shape from shading algorithm. Also, because the accuracy of stereo methods decrease with the square of depth [14] and shape from shading is locally more accurate than stereo, the integrated approach provides more detailed results than could be provided by stereo alone.

In order to use existing shape from shading methods, an image of a scene with albedo variations must first be segmented into regions of uniform reflectivity. Color images are used to facilitate the segmentation process because abrupt changes in color are more likely to correspond to albedo changes than are abrupt
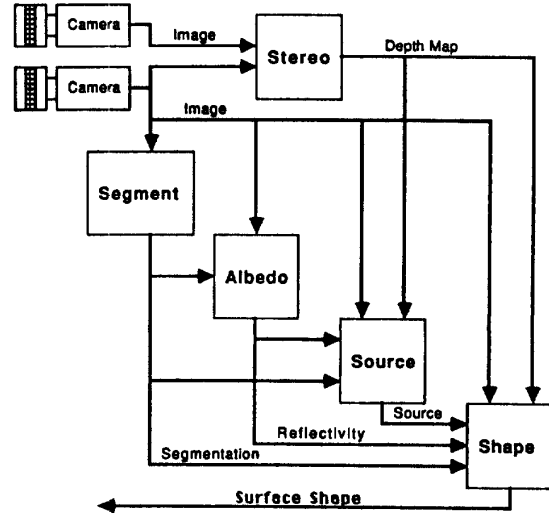

Figure 1: Data flow diagram for the integrated system. The stereo module computes depth directly from the input images and the segmentation module segments one of them. The segmentation is used in the computation of the reflectivity coefficients which are used along with the stereo depth estimate by the source estimation module. The output of every module is used by the shape from shading module in the final shape estimation.

changes in gray level. It is important to avoid detecting edges that correspond to changes in surface orientation because the edges will later be used in the estimation of the reflectivity coefficients.

The integrated system is composed of five main parts (as shown in figure 1): (1) A stereo module that provides an initial depth estimate. (2) A segmentation module that segments one of the images into regions of uniform color. (3) An albedo module that estimates the reflectivity coefficients of the various regions. (4) A source module that uses the stereo depth estimate and the reflectivity map in the estimation of the source distribution. And (5) a shape from shading module that uses the initial depth estimate provided by the stereo module, the reflectivity map, and the source distribution in the estimation of the final surface shape.

6 Stereo module

In an integrated system, other modules must be considered in the choice of a stereo algorithm. In the present approach, the source estimation module requires the surface orientation for a wide range of values. If a feature matching algorithm is used, inter-

polation of the isolated depth points must be done in order to estimate the surface orientation. For simplicity, an area-comparison based algorithm is used here. Keeping each module of the integrated system simple has the virtue of making modifications to the system easier. However, since depth map accuracy affects the quality of the source estimate, it should also be considered

7 Segmentation module

As in the case of the stereo module, the results of the segmentation module will be used throughout the integrated system. In this work, a segmentation algorithm based on minimal spanning trees, is used. Minimal spanning trees were previously used for segmentation by Suk and Cho[15]. Here the interpixel distance is defined to be the length of the vector in color space from one normalized color vector to the other. More precisely, let c_i and c_j be the color vectors associated with two pixels. Then the edge length, e, is defined to be $e = |c_i/|c_i| - c_j/|c_j||$. This distance function is intensity independent and works well in practice. Additional algorithm details are given in [7].

8 Albedo module

Two methods of estimating the reflectivity coefficients are summarized here. The first method involves the minimization of a functional defined on each region of the segmented image and is referred to as the region based method. The second method involves comparing neighboring points along the boundaries between different regions and is referred to as the boundary based method.

8.1 Region based method

Consider minimizing the following functional with respect to the color vectors associated with each region

$$\int\int (E(x,y) - c(\alpha(x,y)) \sum_{i=1}^{m} R^*(n(x,y),l_i))^2 dxdy$$

where $c(\alpha(x,y))$ is the color of region α and R^* is the albedo independent reflectance function. The above functional may be minimized separately for each region of constant color. Also, because the color vector, $c(\alpha)$ is independent of the image coordinates in each region, the equation may be rearranged and solved for each color vector. Let Ω_α be the αth region. The solution is,

$$c(\alpha) = \frac{\int\int_{\Omega_\alpha} E \sum_{i=1}^{m} R^*(n,l_i)dxdy}{\int\int_{\Omega_\alpha} (\sum_{i=1}^{m} R^*(n,l_i))^2 dxdy}$$

An estimate of the normal to the surface can be obtained from the initial surface estimate provided by stereo. The main drawback of this approach is that the light source distribution must be known to estimate the albedo.

8.2 Boundary based method

The albedo of each image point depends upon the reflectivity of the material and the surface orientation at that point. If it is assumed that the surface orientation varies slowly at every color region boundary, then the ratio of the albedo of one region to that of another is proportional to the ratio of gray levels of neighboring points on opposite sides of a boundary.

Given N regions, $r_l, l = 1,\ldots,N$, let ρ_j be the albedo of region r_j and ρ_k be the albedo of region r_k. Let (E_{jki}, E_{kji}) be the ith pair of pixel intensities drawn from regions r_j and r_k, respectively, and define the albedo ratio statistic, β_{jk}, by

$$\beta_{jk} = \frac{\sum_{i=1}^{N_{jk}} E_{jki}}{\sum_{i=1}^{N_{jk}} E_{kji}}$$

Define $\mathcal{E}_{jk} = \sum_{i=1}^{N_{jk}} E_{jki}$ to be the sum of the pixel intensities along one side of a boundary segment. Then the albedo ratios may be written $\beta_{jk} = \rho_j/\rho_k = \mathcal{E}_{jk}/\mathcal{E}_{kj}$. This is a set of linear equations for the albedos as is easy to see upon rearrangement, $\mathcal{E}_{kj}\rho_j = \mathcal{E}_{jk}\rho_k$. This system of equations can be simplified by noticing that they may be combined to yield equations of the following form $\sum_{j, j\neq k} \mathcal{E}_{kj}\rho_j = \left(\sum_{j, j\neq k} \mathcal{E}_{jk}\right)\rho_k$. In words, these equations state that the weighted sum of the irradiance along the interior boundary of a region where the weights are the albedos of the adjoining regions is equal to the sum of the irradiance along the exterior boundary multiplied by the albedo of the given region. Let $\tilde{\mathcal{E}}_{kk} = \sum_{j, j\neq k} \mathcal{E}_{jk}$. Then the matrix F may be defined

$$F = \begin{pmatrix} -\tilde{\mathcal{E}}_{11} & \mathcal{E}_{12} & \mathcal{E}_{13} & \cdots & \mathcal{E}_{1N} \\ \mathcal{E}_{21} & -\tilde{\mathcal{E}}_{22} & \mathcal{E}_{23} & \cdots & \mathcal{E}_{2N} \\ \vdots & & & & \vdots \\ \mathcal{E}_{N1} & \mathcal{E}_{N2} & \cdots & \mathcal{E}_{N,N-1} & -\tilde{\mathcal{E}}_{NN} \end{pmatrix}$$

Finally, the equations may be written, $F\vec{\rho} = 0$ which is a linear, homogeneous system. An additional constraint is required to obtain meaningful results. Presently, the equations are solved by setting the albedo of the brightest region equal to one and solving the remaining inhomogeneous system by a least squares method.

152

9 Experimental results

Three types of experiments were performed to determine the performance of the multiple source and shape from shading algorithms. The results are summarized below.

9.1 Source estimation tests

The algorithms for estimating the source distribution were tested on a large number of synthetic images and the results were compared with the results obtained from the single source estimation algorithm of Brooks and Horn [2].

The depth maps and images were constructed by a two step procedure: (1) Two depth map types were defined by the equations, $z = A\cos(\eta r)$, and $z = Br\cos(\pi\eta\alpha w/2)$ where $r = \sqrt{x^2 + y^2}$ and $\alpha = \arctan(y/x)$ and A, B, η and w are constants. These map types contain a wide range of surface slopes and avoid depth discontinuity problems. (2) The depth map was illuminated by a random collection of light source vectors including an ambient term. Ten 100×100 images of each depth map type were generated using $1, 3, 9, 27$ and 81 light sources and two values of the ambient term for a total of 200 test images.

Both the virtual and positive light source algorithms were tested using 5, 9, and 17 sources and sometimes an ambient term. The tests were conducted as follows: (1) For each test condition, the source algorithm was applied to the test depth maps and images. (2) The resulting source distribution was used to construct a new image. (3) The root mean squared (RMS) difference between the original and reconstructed images was computed. In total, 2600 tests were performed.

Figure 2 summarizes some of the results for the first type of surface using 17 model sources and no ambient term. The average RMS error is plotted versus the number of light sources in the test distribution. Results for the single light source program and various versions of the multiple source program are shown.

Several conclusions may be drawn from the complete battery of tests. The multiple source methods almost always perform better than the single source method except in the case of one test source. Increasing the number of model sources or including an ambient term increases performance. The virtual source method performs better than the positive source method. As the scene complexity increases, the advantage of using multiple source methods increases.

Evidently, the hardest distributions to fit are those with several distinct light sources. As the number of


Figure 2: Performance of source distribution algorithms on surface type 1 using 17 model sources.

test sources increases, each light source becomes less distinct and the distribution becomes more uniform. This explains the decrease in variance and improvement in performance of all methods as the number of test sources becomes large.

9.2 Shape estimation tests

The single source and multisource shape from shading algorithms were compared on several 100×100 synthetic images generated from the depth map $z = A\cos(\eta r)$ using $1, 8$ and 64 light sources.

The source was computed as in the previous section using the positive and virtual source methods with 9 and 17 sources and, in some cases, an ambient term. The method of Brooks and Horn [2] was also used. The shape from shading program for each source type was applied to each image to arrive at a new depth map. In each case, the resulting depth map was compared with the original and the mean error percent, $\epsilon_p = \sum_{ij}(z_{ij} - \tilde{z}_{ij})^2 / \sum_{ij} z_{ij}^2$ was computed. In total, 19 tests were performed.

Figure 3 summarizes the results for the first type of surface using 17 model sources. The average percentage error in height is plotted versus the number of light sources in the test distribution. Results for the single source, shape from shading program and various versions of the multiple source, shape from shading program are shown in each case. The data show several things. The multiple source methods generally performed better than the single source method on images that were generated with many sources but not

153

Figure 3: Performance of the shape from shading algorithms on surface type 1 using 17 model sources.

as well on images generated with a single source. The performance of the multiple source methods increased as the number of model sources increased. There was no consistent difference in the performance of the virtual and positive source methods.

9.3 Integrated estimation tests

The integrated shape recovery system was tested on a stereo pair of images acquired in the laboratory. One view of a rectangular solid viewed from the edge is shown at the top of figure 4. The surface of the solid is covered with a crenulated, colored material. The material has no surface texture other than that caused by shaded undulations. The two vertical bands appearing near the edges are blue and rest of the surface is white or cream colored. The scene is illuminated by ordinary ceiling lights which are an array of diffuse rectangular illumination sources.

The source distribution was modelled by 17 sources plus an ambient term. The solution was found by the positive source method. The depth map produced by stereo is shown in the upper graph of figure 5. Although this result is globally quite accurate, the surface detail is not accurate. The lower picture in figure 4 shows the shaded depth map after applying the shape estimation program to the output of the stereo program. The lower graph in figure 5 is the corresponding depth map showing improvement in the sharpness of the corner.


Figure 4: Performance of the integrated stereo and shape from shading system on an image of a box. The original image and final result are shown.

10 Conclusions

The primary result of this paper is a method of approximating and using complex light source distributions in an integrated approach that combines stereo and shading information. The combined approach is applicable to a wider range of scenes and is more accurate than either method used alone. A method of estimating the reflectivity parameters using color images has also been incorporated.

11 Acknowledgements

154

Figure 5: Performance of the integrated stereo and shape from shading system on an image of a box. The depth map produced by stereo and the final result are shown.

References

[1] H. Bulthoff and H. Mallot, "Interaction of different modules in depth perception," *Proc. First Intl. Conf. Comput. Vis.*, pp. 295–305, 1987.

[2] M. J. Brooks and B. K. P. Horn, "Shape and source from shading," *Proc. Intl. Joint Conf. Artif. Intell.*, pp. 932–936, 1985.

[3] H. T. Croft, K. J. Falconer, R. K. Guy, *Unsolved Problems in Geometry*, Springer-Verlag, 1991.

[4] B. V. Funt, M. S. Drew, and M. Brockington, "Recovering shading from color images," *Proc. Sec. European Conf. Comput. Vis.*, pp. 124–132, 1992.

[5] G. Healey, "Using color for geometry-insensitive segmentation," *J. Opt. Soc. Am. A*, Vol. 6, No. 6, pp. 920–937, 1989.

[6] B. K. P. Horn, "Height and gradient from shading," *Intl. J. of Comput. Vis.*, Vol. 5, No. 1, pp. 37–75, 1990.

[7] D. R. Hougen and N. Ahuja, "Integration of stereo and shape from shading using color," *Proc. Second Intl. Conf. Auto. Robotics and Comput. Vis.*, Vol. 1, pp. cv-6.6.1–6.6.5, 1992.

[8] C. L. Lawson and R. J. Hanson, *Solving Least Squares Problems*, Englewood Cliffs, New Jersey: Prentice Hall, Inc., 1974.

[9] Y. G. Leclerc and A. F. Bobick, "The direct computation of height from shading," *Proc. Comput. Vis. Patt. Recog.*, pp. 552–558, 1991.

[10] P. Lorrain and D. Corson, *Electromagnetic Fields and Waves, second edition*, New York: W. H. Freeman and Company, 1970.

[11] S. K. Nayar, K. Ikeuchi, and T. Kanade, "Shape from Interreflections," *Proc. Intl. Conf. Comput. Vis.*, pp. 2–11, 1990.

[12] S. K. Nayar, K. Ikeuchi, and T. Kanade, "Surface reflection: physical and geometrical perspectives," *Patt. Recog. Machine Intell.*, Vol. 13, No. 5, pp. 185–212, 1991.

[13] A. P. Pentland, "Finding the illuminant direction," *J. Opt. Soc. Am.*, Vol. 72, No. 4, pp. 448–455, 1982.

[14] J.J Rodriquez and J.K. Aggarwal, "Quantization error in stereo imaging," *Comput. Vis. Patt. Recog.*, pp. 153–158, 1988.

[15] M. Suk and T.H. Cho, "Segmentation of images using minimum spanning trees," *Intl. Soc. Opt. Engr., SPIE*, Vol. 397, pp. 180–185, 1983.

[16] L. B. Wolff, "Diffuse Reflection," *Proc. Comput. Vis. Patt. Recog.*, pp. 472–478, 1992.

[17] Y. Yang and A. Yuille, "Sources from shading," *Proc. Comput. Vis. Patt. Recog.*, pp. 534–539, 1991.

[18] Q. Zheng and R. Chellappa, "Estimation of illuminant direction, albedo, and shape from shading," *Proc. Comput. Vis. Patt. Recog.*, pp. 540–545, 1991.