

SHAPE FROM SHADING USING THE FACET MODEL

TING-CHUEN PONG,* ROBERT M. HARALICK,† and LINDA G. SHAPIRO†

*Computer Science Department, University of Minnesota, Minneapolis, MN 55455, U.S.A.;
and †Electrical Engineering Department, University of Washington, Seattle, WA 98195, U.S.A.

(Received 9 March 1988; in revised form 21 June 1988; received for publication 28 November 1988)

Abstract—In this paper, the facet model is used in the recovery of surface orientation from single and multiple images. Two methods for determining surface orientation from image shading are presented. The first method works in the single image domain and is formulated as a non-linear optimization problem. Since three-dimensional scene information available in a single image is usually ambiguous, the optimization procedure can result in multiple solutions. The possibility of adding boundary constraints to the optimization process is also investigated. In the second method, additional images are obtained from the same viewing position, but with changing illumination direction. With these additional images, local surface orientation is determined uniquely by a linear optimization procedure. Experimental results of the shape-from-shading methods are also described.

Computer vision Shape-from-shading Three-dimensional shape Facet model
Surface orientation

1. INTRODUCTION

A trend of the recent mid-level computer vision research is to study the recovery of intrinsic scene characteristics such as surface orientations, reflectance properties and distance information from images.⁽¹⁾ This recovery process is important because knowing these intrinsic scene characteristics will be a definite aid to higher-level scene analysis problems. Apparently, humans have the ability to determine these characteristics even for unfamiliar scenes. Since information about the scene being viewed is captured as observed gray level intensities in the image, it is reasonable to believe that some of these characteristics are recoverable from image intensity variations.

Many shape-from-*x* approaches have been proposed for determining intrinsic scene characteristics. Among these approaches, stereopsis^(5,18) and motion^(12,28) are currently the most studied subjects. Four other important sources of shape information are shading, texture, shadow and contour.

Smooth intensity variation (or shading) is an important clue for determining surface orientation. The shape-from-shading idea was first formulated by Horn.⁽¹¹⁾ Since then, a great deal of work has been done in this area.^(14,22,30) If we assume a uniformly textured surface, surface orientations may be inferred from the way the coarseness of the image texture changes across the image. Shape-from-texture is another area of recent research.^(16,29) When shadows are located in an image, the shapes of the shadows can be used to determine three-dimensional information about the objects in the scene. Shadow analysis can be referred to as the process of locating shadow

regions, finding correspondences between shadow casting objects and shadow regions, and deducing three-dimensional information about the objects involved in the shadow formation process. Theoretical work on shadow analysis can be found in Ref. (24). Shape-from-shadow methods have been found to be useful for estimating heights of objects in aerial images.⁽¹³⁾ Three-dimensional surface shape can also be inferred from the two-dimensional shapes of edges or curves in an image. Shape-from-contour methods^(3,19,26) have been found to be effective in determining the shape of a visible surface.

In the present work, the intrinsic property which we are primarily interested in is surface orientation. In this paper, the facet model^(7,8) is used in the recovery of surface orientations from single or multiple images. The facet model assumes that image intensity values are noisy sampled observations of an underlying intensity surface. Thus, any interpretation made on the basis of a neighborhood of pixel values should be understood through the analysis of its underlying intensity surface.

In past facet work,⁽⁹⁾ the underlying intensity surface was estimated by a least squares fit with a functional form consisting of a linear combination of the tensor products of discrete orthogonal polynomials. In the new facet shape-from-shading formulation, the three-dimensional object surface is assumed to be locally fit by a quadratic surface. The underlying gray tone intensity surface is then modeled using a quadratic object surface having Lambertian reflectance. Having estimated the free parameters of the Lambertian intensity surface, the surface orientation around each pixel is readily determined.

1.1. Overview

Two methods for determining surface orientation from image shading information are presented in this paper. The first method works in the single image domain and the second method works for multiple images. Before discussing these techniques in detail, the image geometry and the illumination model to be used in our presentation are defined in Section 1.2 and 1.3, respectively. The most pertinent related literature is briefly discussed in Section 2. The basic theory of the facet based shape-from-shading method is given in Section 3.1. In Section 3.1.1, the selection of starting value, which is required in our method, is discussed. A modification of the original technique is also suggested. The second method which works for multiple images is described in Section 3.2. A possible extension of these methods to handle specular reflection is given in Section 3.3. In Section 4, results of some preliminary experiments are discussed.

1.2. Imaging geometry

The relationship between scene coordinates and image coordinates is illustrated in Fig. 1. We assume that the camera lens is at the origin and that the z -axis is directed towards the image plane which is in front of the lens. The image plane is placed at a distance f , the focal length of the lens, in front of the origin so that the image is oriented in the same way as the scene. As seen from Fig. 1, the following relations hold for perspective projection:

$$u = \frac{fx}{z} \quad \text{and} \quad v = \frac{fy}{z}.$$

In our discussion, the perspective projection is approximated by an orthographic projection. This




Fig. 1. Relationship between scene coordinates and image coordinates.

approximation is good when the size of the objects being imaged is small compared to the viewing distance. In this case, appropriate coordinate systems can be chosen such that the following relations hold:

$$u = x \quad \text{and} \quad v = y.$$

1.3. Illumination model

In the following discussion, we will use a simple illumination model that assumes a distant point light source and a Lambertian reflectance model. A Lambertian surface looks equally bright in all directions. The brightness of a Lambertian surface illuminated by a distant point light source is given by:

$$I = I_0 N \cdot L \quad (1)$$

where I_0 is a constant depending on the surface albedo and the intensity of the light source, N is the unit surface normal vector, and L is the unit vector of the illumination direction.

The unit vector which points in the direction of the light source can be specified by the two angles shown in Fig. 1. The first is the azimuth (θ) which is the angle between the x -axis and the projection of the vector onto the x - y plane, while the second is the angle of elevation (ϕ) of the light source. If we represent this unit vector by $[a, b, c]$, then

$$a = \cos\theta \cos\phi,$$

$$b = \sin\theta \cos\phi$$

and

(2)

$$c = -\sin\phi.$$

In our discussion, we will consider only positive values of ϕ . Therefore, c is always less than zero.

If the height of the object surface above the x - y plane is expressed as a function of x and y , that is, $z = S(x, y)$, then the surface normal is given by the vector:

$$N = \frac{[S_x, S_y, -1]}{(1 + S_x^2 + S_y^2)^{1/2}}$$

where S_x and S_y denote first partials of S with respect to x and y , respectively. By carrying out the dot product in equation (1), it follows that

$$I = I_0 \frac{aS_x + bS_y - c}{(1 + S_x^2 + S_y^2)^{1/2}}.$$

2. RELATED LITERATURE

A great deal of work has been done in determining surface orientation from image shading. The shape from shading idea is first formulated in Horn's paper.^(1,1) Horn approaches the shape from shading problem by solving the image irradiance equation which is expressed in the form of a first-order partial

differential equation. He suggests this equation be solved using a modified characteristic strip-expansion method. While this method works well in the absence of noise, its behavior is uncertain when applied to a noisy image.

While Horn estimates surface orientation by solving the image irradiance equation by direct integration, Ikeuchi and Horn⁽¹⁴⁾ estimate surface normal by an iterative method. The iterative method attempts to minimize the difference between the observed image intensities and the predicted values which are obtained iteratively from two constraints. The first constraint is based on the image irradiance equation. The second constraint is provided by the assumption that the surface is smooth. As an additional constraint for both Horn's and Ikeuchi's methods, it is required that both the reflectance property of the surface and the illumination condition of the light source be known. Ikeuchi's method requires that occluding boundary information be specified as boundary conditions.

Stereographic projection is used in Ikeuchi's method because the gradients of an occluding boundary map to infinity in gradient space. In stereographic projection, orientations of surface normals are represented uniquely by points on the Gaussian sphere. A point on the Gaussian sphere is then projected onto a plane tangent at the north pole with the center of projection at the south pole. Therefore, an occluding boundary, which is represented by points on the equator of the Gaussian sphere, is projected into a circle of radius two. A modification of Ikeuchi's method can be found in Ref. (25). Other results in shape from shading from a single image include the work of Woodham⁽³¹⁾ and Pentland.⁽²²⁾ These studies determine surface orientations of visible surfaces through local analysis of image shading.

Woodham⁽³⁰⁾ introduces a technique to determine local surface orientation from images which are taken by varying the illumination direction, while keeping the viewing direction fixed. This technique is referred to as photometric stereo. In this method, since images are taken without altering the camera position, there is no problem in finding corresponding points in the images. This technique determines surface orientation at each pixel by finding the intersection of the reflectance maps of the images in the gradient space. In general, at least three images are needed to determine a unique surface orientation.

While Woodham's method is designed for matte surfaces, Ikeuchi⁽¹⁵⁾ applies the photometric stereo technique to determine surface orientations of specular surfaces. The point light source used in Woodham's method is unsuitable for specularly reflecting surfaces, since the specular component in the image due to point light source only occupies a very small area. Instead, Ikeuchi employs a distributed planar light source with spatially varying irradiation.

Although both Woodham's and Ikeuchi's methods work well for noise free images, these methods obtain surface orientation in an image on a pixel-by-pixel

basis. Therefore, we would expect these methods to be very sensitive to noise. The problem is yet to be solved if the input images are known to be noisy.

3. THE FACET MODEL APPROACH

In the following sections, two methods which determine surface orientation from image shading information are presented. The first method works in the single image domain and is formulated as a non-linear optimization problem. Since three-dimensional scene information available in a single image is usually ambiguous, the optimization procedure can result in multiple solutions. In the second method, additional images are obtained from the same viewing position, but with changing illumination direction. With these additional images, the non-linear system of image irradiance equations are transformed into a set of linear equations. The least squares procedure described in Section 3.2 enables a unique solution to be obtained from this linear system of equations.

3.1. Shape from shading from a single image

The first method determines surface orientation from the shading information of a single image. It is formulated as a non-linear optimization problem. According to the facet model, image intensity values are the results of noisy sampling of an underlying intensity surface. Therefore, the image intensity of a Lambertian surface illuminated by a distant point light source under orthographic projection can be expressed as:

$$I(x, y) = I_0 \frac{aS_x(x, y) + bS_y(x, y) - c}{(1 + S_x^2 + S_y^2)^{1/2}} + n(x, y) \quad (3)$$

where n denotes the independent and identically distributed noise. Furthermore, a smoothness constraint on the object surface is imposed in the following way. Let R be an arbitrary neighborhood over which equation (3) holds. We assume that

$$S(x, y) = \sum_{i=1}^n k_i h_i(x, y) \quad \text{for } (x, y) \in R,$$

where the h_i 's are some basis functions and the k_i 's are constants. Different neighborhoods have different coefficients. The form for S that we currently use is a bivariate quadratic polynomial:

$$S(x, y) = k_0 + k_1x + k_2y + k_3x^2 + k_4xy + k_5y^2.$$

By using the intensity values $\{I(x, y) | (x, y) \in R\}$, we can determine a least squares estimator \hat{k} of k (note that k denotes the vector $[k_1, k_2, \dots, k_n]$).

The least squares approximation problem is to determine coefficients k_i such that

$$e^2 = \sum_{(x, y) \in R} \left\| I(x, y) - I_0 \frac{aS_x(x, y) + bS_y(x, y) - c}{(1 + S_x^2 + S_y^2)^{1/2}} \right\|^2 \quad (4)$$

is minimized (note that $\|\cdot\|$ denotes the l_2 norm). In the case of a quadratic object surface,

$$S_x(x, y) = k_1 + 2k_3x + k_4y,$$

and

$$S_y(x, y) = k_2 + k_4x + 2k_5y.$$

where $S_x(x, y)$ and $S_y(x, y)$ denote the partial derivatives of S with respect to x and y . Equation (4) can be rewritten as:

$$e^2 = \sum_{(x,y) \in R} \|I(x, y) - F(x, y, \underline{k})\|^2 \quad (5)$$

where

$$F(x, y, \underline{k}) = I_0 \frac{a(k_1 + 2k_3x + k_4y) + b(k_2 + k_4x + 2k_5y) - c}{[1 + (k_1 + 2k_3x + k_4y)^2 + (k_2 + k_4x + 2k_5y)^2]^{1/2}}.$$

Equation (5) is non-linear in the k_i 's. One of the methods for solving this non-linear system is the Marquardt Method.⁽⁴⁾ This method is basically an improvement of the Taylor series linearization method and the steepest descent method. The ideas of the linearization method and the steepest descent method are explained briefly here.

The linearization method begins by expanding $F(x, y, \underline{k})$ into a Taylor Series, that is,

$$F(x, y, \underline{k}) = F(x, y, \underline{k}_0) + \sum_{i=1}^n \left[\frac{dF(x, y, \underline{k})}{dk_i} \Big|_{\underline{k}=\underline{k}_0} (k_i - k_0(i)) \right] \quad (6)$$

where \underline{k}_0 is a starting value of \underline{k} and $k_0(i)$ denotes the i^{th} component of \underline{k}_0 . If we let $g_i = k_i - k_0(i)$, then equation (6) becomes linear in the g_i 's. Therefore, the original problem becomes a linear regression problem and can be solved by the least squares procedure for linear models. The least squares estimator thus obtained is used as a new starting value for the Taylor series expansion and the process is repeated until convergence (i.e. until e^2 is less than some predetermined tolerance).

The idea of the steepest descent method is to move the estimate, from a starting point \underline{k}_0 , along a direction of steepest descent as long as e^2 decreases. The direction of steepest descent is estimated as the direction obtained by differentiating (5) with respect to \underline{k} . In general, the steepest descent method works well when the starting value is far from the desired solution and the linearization method does a better job when a good initial value is provided. However, Marquardt's method is an iterative methods that combines the good features of the two methods and works well for most practical problems. There are many possible implementations of the Marquardt algorithm.⁽¹⁷⁾ The algorithm used here is based on the idea of More.⁽²¹⁾

The problem with using the above method for the

shape from shading task is that it is very starting-value sensitive. A bad choice of starting value may lead to only a local minimum instead of a global minimum. Furthermore, multiple solutions may exist if the information available from the input image is ambiguous.

3.1.1. *Initial estimate and boundary constraints.* The nonlinear optimization procedure described above is an iterative method. An important starting point is to obtain a good initial estimate. Better starting values not only increase the rate of convergence, they also reduce the chances of converging to local minima. We also believe that supplying additional constraints to the shape recovery process would result in an improved accuracy in the reconstructed surface. We want to explore in this section the possibility of using two kinds of boundary constraints, depth and surface orientation constraints, to bootstrap the shape-from-shading process. Sparse depth values can be obtained from stereo computations. A detailed discussion of stereo computation can be found in Ref. (18). The second constraint can be obtained from local analysis of image shading. Local analysis of image shading is given in Refs (22) and (23).

3.1.2. *Depth constraints.* In binocular stereopsis, surface depth can be determined by matching image features which are usually extracted as zero-crossing edges.⁽²⁰⁾ Given a pair of stereo images, depth of object surface can be determined by simple trigonometric transformation once the correspondence problem is solved. In stereo computations, corresponding points in the two images are determined only at locations with high information content. Since stereo matching can produce only a sparse depth map, some kind of interpolation technique is required to compute a complete surface description.

A computational theory of visual surface interpolation is proposed by Grimson⁽⁶⁾ and refined by Terzopoulos.⁽²⁷⁾ While this interpolation method constructs surface based on information available along zero-crossing contours, its reconstructed surface fails to account fully for the intensity changes which occur away from the zero-crossing contours. Therefore, the reconstructed surface may not be consistent with the observed intensity variations. Our intention here is to compute a complete surface reconstruction by augmenting shading information with depth constraints.

Let R be an arbitrary connected region over which the smoothness constraint $S(x, y)$ is satisfied. Let B be a set of pixels for which the depths are known. We have for $(x_i, y_i) \in B$,

$$k_0 + k_1x_i + k_2y_i + k_3x_i^2 + k_4x_iy_i + k_5y_i^2 \doteq s_i$$

for $i = 1, \dots, n$

where n is the number of pixels in B , and s_i gives the depth of the surface at (x_i, y_i) . In matrix notation, we have

$$A \underline{x} \doteq \underline{y}$$

where

$$\underline{x} = (k_0, k_1, k_2, k_3, k_4, k_5)^t,$$

$$y_i = s_i,$$

and

$$a_i = (1, x_i, y_i, x_i^2, x_i y_i, y_i^2), \quad i = 1, \dots, n.$$

Given enough points in B , an initial estimate for the surface in R can be determined by a least-squares method which minimizes

$$e^2 = \sum_{i=1}^n \|(k_0 + k_1 x_i + k_2 y_i + k_3 x_i^2 + k_4 x_i y_i + k_5 y_i^2) - s_i\|^2$$

or equivalently,

$$e^2 = \|A\underline{x} - \underline{y}\|^2.$$

The solution to this least-squares problem is discussed below.

3.1.3. *Singular value decomposition.* There are many algorithms for solving the least-squares problems. The most reliable method is the singular value decomposition method. A singular value decomposition of a rectangular $m \times n$ matrix A with $m \geq n$ is a factorization of the form

$$A = USV^t$$

where U is an $m \times m$ orthogonal matrix, V is an $n \times n$ orthogonal matrix, and S is an $m \times n$ diagonal matrix with nonnegative diagonal entries. The diagonal elements of S are called singular values of A .

Using the singular decomposition of A , the linear system $A\underline{x} = \underline{y}$ becomes

$$USV^t \underline{x} = \underline{y}.$$

The error,

$$e^2 = \|A\underline{x} - \underline{y}\|^2,$$

to be minimized in the least-squares problem becomes

$$e^2 = \|USV^t \underline{x} - \underline{y}\|^2.$$

Since orthogonal matrices preserve norm, we have

$$e = \|U^t(USV^t \underline{x} - \underline{y})\|$$

$$= \|SV^t \underline{x} - U^t \underline{y}\|.$$

If we let $\underline{z} = V^t \underline{x}$ and $\underline{v} = U^t \underline{y}$, we then have

$$e = \|\underline{S}\underline{z} - \underline{v}\|.$$

Since S is a diagonal matrix, it is easily seen that the vector \underline{z} that produces the minimum e is given by

$$z_i = v_i/s_i \quad \text{if } s_i \neq 0, \text{ and}$$

$$z_i = 0 \quad \text{if } s_i = 0.$$

The transformation $\underline{x} = V\underline{z}$ then determines the least-squares solution.

3.1.4. *Possible modification.* We have shown how

boundary constraints can be used to obtain initial estimates. As an extension of the algorithm described in Section 3.1, we may also consider a minimization of the form

$$w \sum_{(x,y) \in R} \|I(x,y) - F(x,y,\underline{k})\|^2$$

$$+ (1-w) \sum_{(x',y') \in B} \|Ax' - y'\|^2$$

where w is a constant between zero and one. The first part of the above expression is acquired from equation (5) which gives the error between the observed and the estimated gray level intensities. The second part designates the deviation between the estimates and the boundary conditions. The constant w is set close to one if the boundary conditions are not certain to be reliable. The same Marquardt method can be used to solve this minimization problem. Experiments show that this formulation is superior to the original formulation.

3.1.5. *Extracting regions.* So far, we have described a shape-from-shading method. One problem remaining to be solved is to find regions over which the smoothness constraints are satisfied. This region extraction process can be accomplished by analyzing the topographic primal sketch of the image intensity surface. While a detailed description of the region extraction process is given in Ref. (23), the overall process is summarized as follows.

1. Extract edge structures based on the facet model second directional derivative zero-crossing edge operator.

2. Extract maximally connected components of non-edge structures as an initial set of region segments.

3. Extract regions by assembling topographic structure within the initial set of region segments.

3.2. Shape from shading from multiple images

As mentioned in Section 3.1, three-dimensional scene information available in a single image is sometimes ambiguous. Woodham⁽³⁰⁾ suggests a method to resolve this ambiguity by using multiple images. In this method, multiple images are obtained by varying the illumination direction, while keeping the viewing direction fixed. In this section, we will describe a variation of Woodham's method. This method is also based on the facet model. We believe that our method is more robust than Woodham's method.

Suppose two images $I_1(x,y)$ and $I_2(x,y)$ are obtained by varying the position of a point light source. The irradiance equations corresponding to $I_1(x,y)$ and $I_2(x,y)$ are

$$I_1(x,y) = I_0 \frac{S_x(x,y)a_1 + S_y(x,y)b_1 - c_1}{(1 + S_x^2(x,y) + S_y^2(x,y))^{1/2}}$$

and

$$I_2(x,y) = I_0 \frac{S_x(x,y)a_2 + S_y(x,y)b_2 - c_2}{(1 + S_x^2(x,y) + S_y^2(x,y))^{1/2}}$$

where $[a_1, b_1, c_1]^t$ and $[a_2, b_2, c_2]^t$ are the unit vectors which point in the directions of the light sources.

In the case of a quadratic object surface,

$$I_1(x, y) = I_0 \frac{a_1(k_1 + 2k_3x + k_4y) + b_1(k_2 + k_4x + 2k_5y) - c_1}{[1 + (k_1 + 2k_3x + k_4y)^2 + (k_2 + k_4x + 2k_5y)^2]^{1/2}} \quad (7)$$

$$I_2(x, y) = I_0 \frac{a_2(k_1 + 2k_3x + k_4y) + b_2(k_2 + k_4x + 2k_5y) - c_2}{[1 + (k_1 + 2k_3x + k_4y)^2 + (k_2 + k_4x + 2k_5y)^2]^{1/2}}. \quad (8)$$

Note that $I_1(x, y)$ and $I_2(x, y)$ do not depend on k_0 . This is expected because a distant point light source is assumed; a translation of the object along the viewing direction does not affect the observed image intensity.

The above irradiance equations are nonlinear in the k_i 's. By taking the ratio of the two irradiance equations, we will show that local surface orientation can be determined by a linear optimization procedure.

3.2.1. *Estimation of surface orientation.* Although (7) and (8) are nonlinear in the k_i 's, we can linearize the problem by taking their ratio. That is

$$\frac{I_1(x, y)}{I_2(x, y)} = \frac{a_1(k_1 + 2k_3x + k_4y) + b_1(k_2 + k_4x + 2k_5y) - c_1}{a_2(k_1 + 2k_3x + k_4y) + b_2(k_2 + k_4x + 2k_5y) - c_2}. \quad (9)$$

If we let $I(x, y) = I_1(x, y)/I_2(x, y)$, equation (9) can be rewritten as:

$$\begin{aligned} (I(x, y)a_2 - a_1)k_1 + (I(x, y)b_2 - b_1)k_2 \\ + 2x(I(x, y)a_2 - a_1)k_3 + [x(I(x, y)b_2 - b_1) \\ + y(I(x, y)a_2 - a_1)]k_4 + 2y(I(x, y)b_2 - b_1)k_5 \\ - (I(x, y)c_2 - c_1) = 0. \end{aligned} \quad (10)$$

Equation (10) is linear in the k_i 's. We can estimate the k_i 's by fitting (10) to the observed $I(x, y)$ s in the neighborhoods of I_1 and I_2 . For simplicity, we will consider only odd sized rectangular neighborhoods and the coordinates of the center of the neighborhood $R \times C$ will be set to (0, 0). It is important to note that the surface normal at the center of the neighborhood is given by $[k_1, k_2, -1]^t$. For $(x, y) \in R \times C$, the system of equations can be represented in matrix notation by

$$B\hat{a} = y, \quad (11)$$

where, if $\underline{b}(x, y)$ is the row vector

$$\begin{aligned} [I(x, y)a_2 - a_1, I(x, y)b_2 - b_1, 2x(I(x, y)a_2 - a_1), \\ x(I(x, y)b_2 - b_1) + y(I(x, y)a_2 - a_1), \\ 2y(I(x, y)b_2 - b_1)] \end{aligned}$$

then

$$B = \begin{bmatrix} \underline{b}(-R, -C) \\ \underline{b}(-R, -C + 1) \\ \vdots \\ \underline{b}(0, 0) \\ \vdots \\ \underline{b}(R - 1, C - 1) \\ \underline{b}(R, C) \end{bmatrix}, \quad y = \begin{bmatrix} I(-R, -C)c_2 - c_1 \\ I(-R, -C + 1)c_2 - c_1 \\ \vdots \\ I(0, 0)c_2 - c_1 \\ \vdots \\ I(R - 1, C - 1)c_2 - c_1 \\ I(R, C)c_2 - c_1 \end{bmatrix}$$

and $\hat{a} = [k_1, k_2, k_3, k_4, k_5]^t$.

A least squares procedure which minimizes

$$e^2 = \|y - B\hat{a}\|^2$$

is used to determine the estimator \hat{a} . The normal equation estimator for a is given by

$$\hat{a} = (B'B)^{-1}B'y. \quad (12)$$

If there is exact linear dependency between the columns of B , the rank of B is less than 5 and $B'B$ is singular. In general, the columns of B are near linearly dependent. In this case, $B'B$ is ill-conditioned and the estimate of \hat{a} suffers from large error. A more reliable method for solving this least-squares problem is based on a matrix factorization known as the singular value decomposition which is described in Section 3.1.3. However, if the problem is rank deficient (i.e. at least one of the singular values of A is zero), the solution which minimizes e^2 is not unique. Two approaches are suggested to avoid this problem.

The first approach is to use a third image. An additional set of equations can be obtained by taking the ratio of the irradiance equation of the third image and that of the first (or second). A new matrix B is obtained by combining these equations. It is clear that the effect of multicollinearity in the new B can be reduced, provided the vectors of the incident illuminations do not all lie in a plane. The second approach is to constrain the locations of the light sources.

3.2.2. *Locations of the light sources.* Recall from Section 1.3 that the direction of a distant point light source is specified by two angles: the azimuth (θ) and the elevation (ϕ). It follows from equation (2) that if θ is 90° , then a becomes zero, if θ is 0° then b becomes zero and if ϕ is 90° both a and b are zero. The angle of elevation of the light source is always made greater than zero, therefore, c is always non-zero.

We can simplify equation (10) by restricting the positions of the light sources. If both a_1 and a_2 are set to zero, equation (10) can be rewritten as

$$\begin{aligned} (I(x, y)b_2 - b_1)k_2 + x(I(x, y)b_2 - b_1)k_4 \\ + 2y(I(x, y)b_2 - b_1)k_5 - (I(x, y)c_2 - c_1) = 0 \end{aligned}$$

or

$$k_2 + xk_4 + 2yk_5 = J(x, y),$$

where $J(x, y) = (I(x, y)c_2 - c_1)/(I(x, y)b_2 - b_1)$. Therefore, the matrix B in (11) can be replaced by

$$B = \begin{bmatrix} 1 & -R & 2(-C) \\ 1 & -R & 2(-C + 1) \\ & \vdots & \\ 1 & 0 & 0 \\ & \vdots & \\ 1 & R & 2(C - 1) \\ 1 & R & 2C \end{bmatrix},$$

\underline{y} becomes a column vector of the $J(x, y)$'s and \underline{a} becomes $[k_2, k_4, k_5]^t$. Note that B is independent of the I 's and because the center of each neighborhood has (x, y) equal to $(0, 0)$, $B'B$ is the same for all the neighborhoods in the image. Furthermore, $B'B$ in (12) is always invertible and can be computed accurately. Therefore, the computation of $\hat{\underline{a}}$ from equation (12) is stable and efficient.

We can determine k_2, k_4 and k_5 accurately if a_1 and a_2 are both zero. Similarly, k_1, k_3 and k_4 can be obtained by making both b_1 and b_2 equal to zero. Hence, if we obtain a first image with $\theta = 90^\circ$ and $\phi = 90^\circ$, a second image with $\theta = 90^\circ$ and arbitrary ϕ , and a third image with $\theta = 0^\circ$ and arbitrary ϕ , then $\hat{\underline{a}}$ can be determined accurately and efficiently. If (k_1, k_3, k_4) and (k_2, k_4, k_5) are determined independently as described, it is possible that the k_4 's computed from the two processes disagree. In the case where the k_4 's differ by a significant amount, we can either take their average or use the three images together to solve for a unique \underline{a} .

If only two images are to be used, we can first compute k_2, k_4 , and k_5 by setting a_1 and a_2 to zero (or k_1, k_3 and k_4 by setting b_1 and b_2 to zero) and then obtain k_1 (or k_2) from the following derivations.

If $a_1 = 0$ and $a_2 = 0$, one obtains from (7)

$$I(0, 0) = \frac{I_0(k_2 b_1 - c_1)}{(1 + k_1^2 + k_2^2)^{1/2}} \quad (13)$$

or

$$1 + k_1^2 + k_2^2 = [I_0(k_2 b_1 - c_1)]^2 / I(0, 0)^2. \quad (14)$$

Similarly,

$$1 + (k_1^2 - 2k_1 k_4 + k_4^2) + (k_2^2 - 4k_2 k_5 + 4k_5^2) = [I_0((k_2 - 2k_5)b_1 - c_1)]^2 / I(0, -1)^2 \quad (15)$$

and

$$1 + (k_1^2 + 2k_1 k_4 + k_4^2) + (k_2^2 + 4k_2 k_5 + 4k_5^2) = [I_0((k_2 + 2k_5)b_1 - c_1)]^2 / I(0, 1)^2. \quad (16)$$

If I_0 is known, one can obtain k_1 directly from (13). Otherwise, let

$$h_0 = (k_2 b_1 - c_1)^2 / I(0, 0)^2, \\ h_{-1} = ((k_2 - 2k_5)b_1 - c_1)^2 / I(0, -1)^2$$

and

$$h_1 = ((k_2 + 2k_5)b_1 - c_1)^2 / I(0, 1)^2,$$

then

$$(15) - (16) \Rightarrow 4k_1 k_4 + 8k_2 k_5 = I_0^2 (h_1 - h_{-1}) \quad (17)$$

and

$$(14) - (15) \Rightarrow 2k_1 k_4 - k_4^2 + 4k_2 k_5 - 4k_5^2 = I_0^2 (h_0 - h_{-1}). \quad (18)$$

If $k_4 \neq 0$, by taking the ratio of (17) and (18), one finds

$$k_1 = \frac{h(4k_2 k_5 - k_4^2 - 4k_5^2) - 8k_2 k_5}{2k_4(2 - h)},$$

where $h = (h_1 - h_{-1}) / (h_0 - h_{-1})$. If $k_4 = 0$, by taking the ratio of (15) and (16), one finds

$$k_1^2 = \frac{h[1 + (k_2^2 + 4k_2 k_5 + k_5^2)] - [1 + (k_2^2 - 4k_2 k_5 + 4k_5^2)]}{1 - h}.$$

Where $h = h_{-1}/h_1$. In this case, the sign of k_1 is undetermined.

In this section, a photometric stereo method has been presented. By taking the ratio of the irradiance equations of a pair of images, we are able to determine local surface orientation uniquely by a linear optimization procedure. It was also shown that adding a third image and restricting the position of the light sources allows us to determine surface orientations accurately and efficiently. The photometric stereo technique is particularly useful in controlled environments where the intensity and position of light source can be adjusted and measured accurately. Some possible applications of this technique include robotic hand-eye systems and industrial inspection tasks.

3.3. Specular reflectance model

A more general shading model consists of two components. The first component is Lambertian reflection. Lambertian surfaces (or matte surfaces) look equally bright in all directions. Specular or mirror-like reflection accounts for the second component. A surface having specular reflection reflects light unequally in different directions. Specular reflection is usually observed as in the image.

Let θ and ϕ be the angles defined in Fig. 2. The angle θ is the angle between the incident ray (L) and the surface normal (N). Note that the angle between the reflected ray (R) and the surface normal is also θ . The angle ϕ is the angle between N and the vector half way between the viewing direction (V) and L . According to the shading model described in Ref. (10), the intensity of a specular surface illuminated by a point light source can be expressed as:

$$I = I_1 \cos \theta + I_s \cos^n \phi \quad (19)$$

where the values of I_1 and I_s depend on the properties




Fig. 2. Relationship between θ and ϕ .

of the surface material, the distance of the light source to the surface element, and the intensity of the light source, and n is a constant. The value of n ranges roughly from 1 to 200, depending on the surface. A shiny metallic surface will have a very large n . A dull surface will have a small value. If we let \bar{L} , \bar{N} , and \bar{V} be, respectively, the unit vectors of L , N , and V , then equation (19) can be rewritten as:

$$I = I_1(\bar{L} \cdot \bar{N}) + I_s(\bar{N} \cdot \bar{H})^n.$$

where

$$\bar{H} = (L + V)/|L + V|.$$

It is clear that $I_1(\bar{L} \cdot \bar{N})$ is the component due to Lambertian reflection, and $I_s(\bar{N} \cdot \bar{H})^n$ is the component due to specular reflection. Note that the intensity of the specularly reflected light falls off rapidly as ϕ increases.

The shape-from-shading method described above can easily be extended to handle specular reflection. The observed intensity of a specular quadratic object surface can be approximated by

$$I(x, y) = I_1 \frac{a(k_1 + 2k_3x + k_4y) + b(k_2 + k_4x + 2k_5y) - c}{[1 + (k_1 + 2k_3x + k_4y)^2 + (k_2 + k_4 + 2k_5y)^2]^{1/2}} + I_s \left[\frac{d(k_1 + 2k_3x + k_4y) + e(k_2 + k_4x + 2k_5y) - f}{[1 + (k_1 + 2k_3x + k_4y)^2 + (k_2 + k_4 + 2k_5y)^2]^{1/2}} \right]^n$$

when $L = (a, b, c)$ and $H = (d, e, f)$. The same optimization procedure as described in Section 3.1 can be used to determine surface orientations for specular surfaces by minimizing the sum of squares of the differences between the observed intensity data and the estimated intensity values computed from the equation above.

3.4. Justification

The methods presented here have some advantages over the previous shape-from-shading approaches. Since a neighborhood of pixels are considered when the surface orientation at a particular location is

estimated, results show that this method is less sensitive to noise.

Furthermore, the above methods can produce more than just the surface normal at each pixel. For each neighborhood, it also produces the coefficients of the fitting polynomial. Using these coefficients permits us to determine whether or not neighboring pixels belong to the same surface patch.

4. RESULTS

The non-linear optimization procedure described in Section 3.1 has been tested on a 50×50 image of a simple surface. Figure 3 shows the surface generated by the polynomial: $40x + 40y - 2x^2 - xy + 2y^2$. The shaded image resulting from illuminating the surface of Fig. 3 from light direction $(0, 0, -1)$, which means directly above the center of the surface, is shown in Fig. 4. By inputting the starting value $\underline{k}_0 = (1, 1, 1, 1, 1)$ or $\underline{k}_0 = (40, 40, -2, -1, 2)$ to the optimization procedure, \underline{k} was found to be $(38.5, 43.8, -2.0, -1.3, 2.0)$, which is very close to the




Fig. 3. The polynomial surface: $40x + 40y - 2x^2 - xy + 2y^2$.




Fig. 4. The shaded image of the surface of Fig. 3.

original surface within the range of (x, y) used. By using $k_0 = (-40, -40, 2, 1, -2)$, the estimated surface is shown in Fig. 5. By using $k_0 = (0.86, -1.61, -0.03, 0.04, -0.01)$, which was obtained by fitting a second degree polynomial to the intensity surface of the shaded image, the estimated surface is shown in Fig. 6. By illuminating the three estimated surfaces from light direction $(0, 0, -1)$, we obtained essentially the same shaded image. This shows the ambiguity in image shading information.

Figure 7 shows a pair of synthetic stereo images. A sparse depth map was determined by the stereo matching procedure described in Ref. (23). The depth map is illustrated in Fig. 8. The depth constraints obtained from the stereo pair of images were used to determine an initial estimate of the object surface in




Fig. 5. The reconstructed surface for the shaded image of Fig. 4.




Fig. 6. Another reconstructed surface for the shaded image of Fig. 4.




Fig. 7. Synthetic stereo images of a spherical surface.




Fig. 8. The depth map resulting from matching the stereo images of Fig. 7.

the images. The singular value decomposition method was used to obtain such an initial estimate. By incorporating the depth constraints into the shape-from-shading algorithm, an accurate reconstructed surface was obtained (Fig. 9). It is worth mentioning




Fig. 9. The reconstructed surface for object in the stereo images.

that multiple surface reconstructions are possible if boundary constraints are unavailable.

To investigate the effect of boundary constraints on the optimization process, the shape-from-shading algorithms with and without boundary constraints have been tested on the image of Fig. 10. One-hundred sets of randomly generated initial estimates were input to the methods described in Section 3.1 and Section 3.1.4. A set of boundary constraints was obtained along the object boundary. Results show that the method converged to three different but correct solutions for eighty-four out of the one hundred initial estimates if no boundary constraint was used. The method converged to a unique solution ninety-four times if the set of boundary constraints was included in the optimization process. We also found that the average rate of convergence was improved by a factor of about two if the boundary constraints were used. Similar results were obtained when these methods were run on the image of Fig. 10.

The photometric stereo method described in Section 3.2 has been tested on both artificially generated images and real images. An object surface corresponding to a portion of a spherical surface is shown in Fig. 11. Three shaded images were obtained by illuminating the surface of Fig. 11 from three different positions: (1) the light direction is $(0, 0, -1)$ which means directly above the center of the surface; (2) the light direction is $(0, 0.259, -0.966)$ which translates to azimuth 90° and elevation 75° ; (3) the light direction is $(0.259, 0, -0.966)$ which translates to azimuth 0° and elevation 75° . The shaded images are shown in the first row of Fig. 12.

In spite of the fact that a spherical surface cannot be fit exactly by a second degree polynomial, the reconstructed surface by the photometric stereo method is extremely close to the original one. The root mean square difference between the elevation data (range from 100 to 255) of the original surface and the reconstructed surface was found to be 1.20.

To examine the performance of our method on noisy images, random noise of mean = 0 and standard deviation = 10 was added to the shaded images of




Fig. 10. Shaded images of a curved surface.




Fig. 11. Portion of a spherical surface.




Fig. 12. Shaded images and noisy images of the surface of Fig. 11.

Fig. 12. The noisy images are shown in the second row of Fig. 12. The reconstructed surface by our method is shown in Fig. 13. The RMS difference between the original surface and the reconstructed surface is 2.67. As a comparison, Woodham's method was implemented and was run on the same noisy images. The reconstructed surface by Woodham's method is shown in Fig. 14. The RMS difference between the original surface and the reconstructed surface is 5.82. These results show that our method is less sensitive to noise.

Another experiment has been performed on the shaded images of an egg. The images were taken with the camera positioned directly above the egg and the light sources positioned at $(\theta, \phi) = (90^\circ, 90^\circ)$, $(\theta, \phi) = (90^\circ, 45^\circ)$, and $(\theta, \phi) = (0^\circ, 45^\circ)$. The corresponding images are shown in Fig. 15. The recovered surface is shown in Fig. 16. The result is considered satisfactory because the light source used was not a point light and the illumination angles measured may differ from the true angles by as much as $\pm 10^\circ$.

In addition to the above experiments, images of a




Fig. 13. The reconstructed surface by the photometric stereo method.




Fig. 16. The reconstructed surface from the images of Fig. 15.




Fig. 14. The reconstructed surface by Woodham's method.

more complex object were also used in testing the photometric stereo method. Three shaded images were obtained by illuminating the object from $(\theta, \phi) = (90^\circ, 90^\circ)$, $(\theta, \phi) = (90^\circ, 60^\circ)$, and $(\theta, \phi) = (0^\circ, 60^\circ)$. The shaded images are shown in Fig. 17. Since the object is composed of different types of surfaces, the smoothness constraint was not expected to hold at locations where these surfaces met. Thus, an edge operator was first used to detect these discontinuities, and the shape-from-shading method was only applied to locations where no discontinuity was detected. The reconstructed surface obtained by using the photometric stereo method with a 5×5 window is shown in Fig. 18. The CPU time for running our algorithm on a VAX/780 for these images took a little over 1.5 min.

5. SUMMARY

Two methods to recover 3-D surface orientation from image shading information have been presented.




Fig. 15. Shaded images of an egg.




Fig. 17. Shaded images of a complex object.




Fig. 18. The reconstructed surface from the images of Fig. 17.

The first method works in the single image domain and is posted as a non-linear optimization procedure. Experiments have been performed on synthetic and real images and the results are found to be satisfactory. Since inherent ambiguity exists in the image, the optimization procedure may lead to undesirable or multiple solutions. By using additional boundary constraints, good initial estimates are obtained for the optimization process and the performance of the method is improved.

By using additional images as constraints, the second method linearizes the optimization process. Unique and accurate solutions have been obtained by experimenting with the second method on simple images. So far, the proposed methods have been tested only on images which are obtained from orthographic projection. In order for these methods to handle perspective projection and more complex images, more work needs to be done in the future.

REFERENCES

1. H. Barrow and J. M. Tenenbaum, Recovery of intrinsic scene characteristics from images, *Computer Vision Systems*. A. Hanson and E. Riseman, Eds, pp. 3–26. Academic Press, New York (1978).
2. H. Barrow and J. M. Tenenbaum, Interpreting line drawings as three-dimensional surfaces, *Artificially intelligence*. Vol. 17, pp. 75–116 (1981).
3. M. Brady and A. Yuille, An extremum principle for shape from contour, *IEEE Trans. Pattern Analysis Mach. Intell.* PAMI-6, 288–301 (1984).
4. N. R. Draper and H. Smith, *Applied Regression Analysis*, 2nd Edn. Wiley, New York (1981).
5. W. Grimson, *From Images to Surfaces*. MIT Press, Cambridge, MA (1981).
6. W. E. L. Grimson, An implementation of a computational theory of visual surface interpolation, *Comput. Vision Graphics Image Process.* 22, 39–69 (1983).
7. R. M. Haralick, Edge and region analysis for digital image data, *Comput. Graphics Image Process.* 12, 60–73 (1980).
8. R. M. Haralick and L. T. Watson, Facet model for image data, *Comput. Graphics Image Process.* 15, pp. 113–129 (1981).
9. R. M. Haralick, Digital step edges from zero crossing of second directional derivatives, *IEEE Trans. Pattern Analysis Mach. Intell.* PAMI-6, 58–68 (1984).
10. S. Harrington, *Computer Graphics: a Programming Approach*. McGraw-Hill, New York (1983).
11. B. K. P. Horn, Obtaining shape from shading information, *The Psychology of Computer Vision*, P. H. Winston, Ed., pp. 115–155. McGraw-Hill, New York (1975).
12. B. K. P. Horn, Determining optical flow, *Artif. Intell.* 17, 185–203 (1981).
13. A. Huertas, Using shadows in the interpretation of aerial images, University of Southern California, ISG Report 104, R. Nevatia, Ed., pp. 55–84 (1983).
14. K. Ikeuchi and B. K. P. Horn, Numerical shape from shape from shading and occluding Boundaries, *Artif. Intell.* 17, 141–184 (1981).
15. K. Ikeuchi, Determining surface orientation of specular surfaces by using the photometric stereo method, *IEEE Trans. PAMI* 3, 661–669 (1981).
16. J. R. Kender, A computational paradigm for deriving local surface orientation from local textural properties, *Proc. IEEE Workshop Comput. Vision Rindge, NH*, pp. 143–152 (1982).
17. D. W. Marquardt, An algorithm for least-squares estimation of nonlinear parameters, *SIAM J. appl. Math.* 11, 431–441 (1963).
18. D. Marr and T. Poggio, A theory of human stereo disparity, *Science* 194, 283–287 (1976).
19. D. Marr, Analysis of occluding contour, *Proc. R. Soc. Lond. B* 197, 441–475 (1977).
20. D. Marr and E. C. Hildreth, Theory of edge detection, *Proc. R. Soc. Lond. B* 207, 187–217 (1980).
21. J. J. More, The Levenberg–Marquardt algorithm implementation and theory, *Numerical Analysis*, G. A. Watson, Ed., *Lecture Notes in Mathematics* 630 (1977).
22. A. Pentland, Local shading analysis, SRI International, Technical Note 272 (1982).
23. T. C. Pong, Determining intrinsic scene characteristics from images, Virginia Polytechnic Institute and State University, Ph.D. Dissertation (1984).
24. S. A. Shafer and T. Kanade, Using shadows in finding surface orientations, *Comput. Vision Graphics Image Process.* 22, 145–176 (1983).
25. G. B. Smith, The relationship between image irradiance and surface orientation, *Proc. IEEE Conf. Comput. Vision Pattern Recognition*, pp. 14–19 (1983).
26. K. A. Stevens, The visual interpretation of surface contours, *Artif. Intell.* 17, 47–73 (1981).
27. D. Terzopoulos, Multilevel computational process for visual surface reconstruction, *Comput. Vision Graphics Image Process.* 24, 52–96 (1983).
28. S. Ullman, *The Interpretation of Visual Motion*. MIT Press, Cambridge, MA (1979).
29. A. P. Witkin, Recovering surface shape and orientation from texture, *Artif. Intell.* 17, 17–45 (1981).
30. R. J. Woodham, Photometric method for determining surface orientation from multiple images, *Opt. Engng* 19, 139–144 (1980).
31. R. J. Woodham, Analyzing images of curved surfaces, *Artif. Intell.* 17, 117–140 (1981).

About the Author—TING-CHUEN PONG was born in Hong Kong in 1957. He received the B.S. degree in Physics and Mathematics from the University of Wisconsin, Eau Claire in 1978, and the M.S. and Ph.D. degrees from Virginia Polytechnic Institute and State University in 1981 and 1984, respectively, both in Computer Science.

He is presently an Assistant Professor of Computer Science at the University of Minnesota, Minneapolis. His research interests include computer vision, pattern recognition, and artificial intelligence.

Dr Pong is a member of the IEEE Computer Society the Association for Computing Machinery, and the American Association for Artificial Intelligence.

About the Author—ROBERT M. HARALICK was born in Brooklyn, New York on 30 September 1943. He received a B.A. degree in Mathematics from the University of Kansas in 1964, a B.S. degree in Electrical Engineering in 1966 and a M.S. degree in Electrical Engineering in 1967. In 1969, after completing his Ph.D. at the University of Kansas, he joined the faculty of the Electrical Engineering Department there where he last served as Professor from 1975 to 1978. In 1979 Dr Haralick joined the Electrical Engineering Department at Virginia Polytechnic Institute and State University where he was a Professor and Director of the Spatial Data Analysis Laboratory. From 1984 to 1986 Dr Haralick served as Vice President of Research at Machine Vision International, Ann Arbor, MI. Dr Haralick now occupies the Boeing Clairmont Egtvedt Professorship in the Department of Electrical Engineering at the University of Washington.

Professor Haralick has made a series of contributions in computer vision. In the high-level vision area, he has identified a variety of vision problems which are special cases of the consistent labeling problem. His papers on consistent labeling, arrangements, relation homomorphism, matching, and tree search translate some specific computer vision problems to the more general combinatorial consistent labeling problem and then discuss the theory of the look-ahead operators that speed up the tree search. This gives a framework for the control structure required in high-level vision problems. Recently, he has extended the forward-checking tree search technique to propositional logic.

In the low-level and mid-level areas, Professor Haralick has worked in image texture analysis using spatial gray tone co-occurrence texture features. These features have been used with success on biological cell images, X-ray images, satellite images, aerial images and many other kinds of images taken at small and large scales. In the feature detection area Professor Haralick has developed the facet model for image processing. The facet model states that many low-level image processing operations can be interpreted relative to what the processing does to the estimated underlying gray tone intensity surface of which the given image is a sampled noisy version. The facet papers develop techniques for edge detection, line detection, noise removal, peak and pit detection, as well as a variety of other topographic gray tone surface features.

Professor Haralick's recent work is in shape analysis and extraction using the techniques of mathematical morphology. He has developed the morphological sampling theorem which establishes a sound shape/size basis for the focus of attention mechanisms which can process image data in a multiresolution mode, thereby making some of the image feature extraction processes execute more efficiently.

Professor Haralick is a Fellow of IEEE for his contributions in computer vision and image processing. He serves on the Editorial Board of *IEEE Transactions on Pattern Analysis and Machine Intelligence*. He is the computer vision area editor for *Communications of the ACM*. He also serves as associate editor for *Computer Vision, Graphics, and Image Processing* and *Pattern Recognition*.

About the Author—LINDA G. SHAPIRO was born in Chicago, Illinois in 1949. She received the B.S. degree in mathematics from the University of Illinois, Urbana in 1970, and the M.S. and Ph.D. degrees in computer science from the University of Iowa, Iowa City in 1972 and 1974, respectively.

She was an Assistant Professor of Computer Science at Kansas State University, Manhattan, from 1974 to 1978 and was an Assistant Professor of Computer Science from 1979 to 1981 and Associate Professor of Computer Science from 1981 to 1984 at Virginia Polytechnic Institute and State University, Blacksburg. She was Director of Intelligent Systems at Machine Vision International in Ann Arbor from 1984 to 1986. She is currently an Associate Professor of Electrical Engineering at the University of Washington. Her research interests include computer vision, artificial intelligence, pattern recognition, robotics, and spatial database systems. She has co-authored a textbook on data structures with R. Baron.

Dr Shapiro is a senior member of the IEEE Computer Society and a member of the Association for Computing Machinery, the Pattern Recognition Society, and the American Association for Artificial Intelligence. She is Editor of *Computer Vision, Graphics, and Image Processing* and an editorial board member of *Pattern Recognition*. She was General Chairman of the IEEE Conference on Computer Vision and Pattern Recognition in 1986, General Chairman of the IEEE Computer Vision Workshop in 1985, and Co-Program Chairman of the IEEE Computer Vision Workshop in 1982; and she has served on the program committees of a number of vision and AI workshops and conferences.