IEEE COMPUTER SOCIETY REPRINT

THE RELATIONSHIP BETWEEN IMAGE IRRADIANCE AND SURFACE ORIENTATION

Grahame B. Smith

Reprinted from PROCEEDINGS OF IEEE 1983
CVPR — COMPUTER VISION AND PATTERN RECOGNITION

IEEE COMPUTER SOCIETY
1109 Spring Street, Suite 300
Silver Spring, MD 20910

IEEE COMPUTER SOCIETY PRESS

THE INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS, INC

THE RELATIONSHIP BETWEEN IMAGE IRRADIANCE AND SURFACE ORIENTATION

Grahame B. Smith

Artificial Intelligence Center, SRI International
Menlo Park, California 94025

ABSTRACT

A formulation of shape from shading is presented in which surface orientation is related to image irradiance without requiring detailed knowledge of either the scene illumination or the albedo of the surface material. The case for uniformly diffuse reflection and perspective projection is discussed in detail. Experiments aimed at using the formulation to recover surface orientation are presented and the difficulty of nonlocal computation discussed. We present an algorithm for reconstructing the 3-D surface shape once surface orientations are known.

1 INTRODUCTION

When the human visual system processes a single image, e.g., Figure 1, it returns a perceived 3-D model of the world, even when that image has limited contour and texture information. This 3-D model is underdetermined by the information in the 2-D image; the visual system has used the image data and its model of visual processing to reconstruct the 3-D world. While there are many information sources within the image, shading is an important source. Facial make-up or a cartoonist's shading, is an everyday example of the way shape, as perceived by our human visual system, is manipulated by shading information.

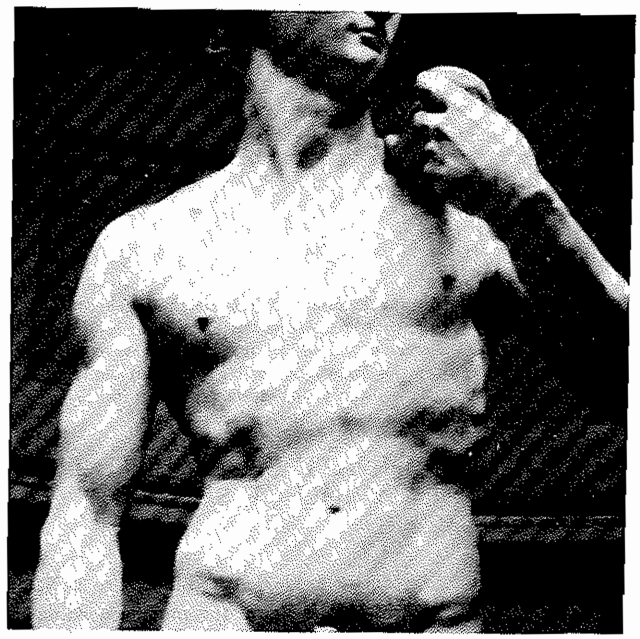

Figure 1 Shape from Shading.

A primary goal of computer vision is to understand this process of reconstructing the 3-D world from 2-D image data, to discover the model, or models that allow 2-D data to infer 3-D structure. The focus of this work is the recovery of the 3-D orientation of surfaces from image shading.

We present a formulation of the shape-from-shading problem, i.e., recovering 3-D surface shape from image shading, that is derived under assumptions of perspective projection, uniformly diffuse reflection,[1] and constant reflectance. This formulation differs from previous approaches to the problem in that we neither make assumptions about the surface shape [2], nor use direct knowledge of the illumination conditions and the surface albedo [3]. The cost we incur for dispensing with these restrictions is the introduction of higher-order differentials into the equations relating surface orientation and image irradiance. The benefits we gain allow us to investigate the strength of the constraint imposed by shading upon shape. Past attempts to solve the shape-from-shading problem, as well as our own efforts, have been aimed at recovering surface shape from image patches for which the reflectance (albedo) can be considered constant.

Previously we examined the influence exerted by the assumption of uniformly diffuse reflection [1], and indicated that the equations relating surface orientation to image irradiance could be expected to yield useful results even in cases in which the reflection is not uniformly diffuse. In that examination we assumed orthographic rather than perspective projection. A comparison of our previous work with this paper, however, shows that the structure of the formulation is not dependent upon the projection used.

If we add additional assumptions, e.g., constraints on the surface type, we can simplify the relationship between surface orientation and image irradiance. While it is not our goal to add constraints upon surface type, the assumption that the surface is locally spherical allows the approximate surface orientation to be recovered by local computation.

The research reported herein was supported by the Defense Advanced Research Projects Agency under Contract MDA903-83-C-0027 and by the National Aeronautics and Space Administration under Contract NASA 9-16664. These contracts are monitored by the U.S. Army Engineer Topographic Laboratory and by the Texas A&M Research Foundation for the Lyndon B. Johnson Space Center.

[1] We prefer the expression *isotropic scattering* to either *uniformly diffuse reflection*, or *Lambertian reflection*, as it emphasis that scene radiance is isotropic. However, uniformly diffuse reflection, and Lambertian reflection are the terms commonly used to indicate that the scene radiance is isotropic.

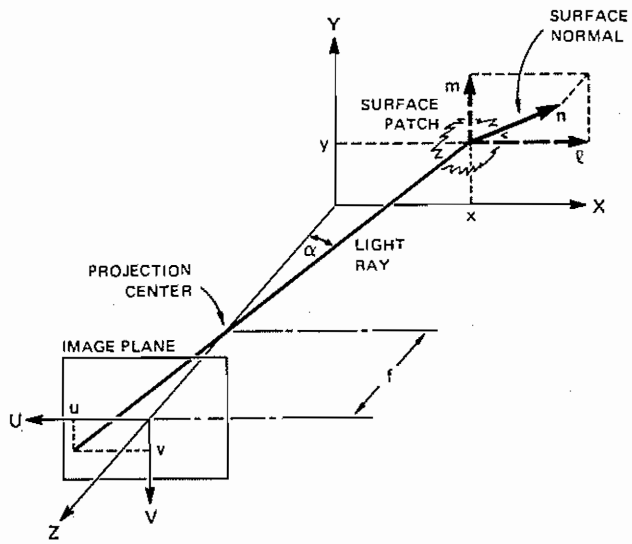

Figure 2 Coordinate Frame. X,Y,Z are the scene coordinates, U,V the image coordinates, and the image plane is located a distance f from the scene coordinate's origin - the projection center. α is the angle between the Z axis (the viewing direction) and the ray of light from the scene point (x, y, z) to the image point (u, v). l and m are the X and Y components of the surface normal **n**.

2 THE COORDINATE FRAME AND REPRESENTATION OF SURFACE ORIENTATION

The coordinate system we use is depicted in Figure 2. X,Y,Z are the scene coordinates and U,V are the image coordinates. The image and scene coordinates are aligned so that X and U axes are parallel, as are the Y and V axes. The U and V axes are inverted with respect to the X and Y axes, so that positive X and Y coordinates will correspond to positive U and V coordinates. The image plane is located at a distance f from the (perspective) projection center, the origin of the scene coordinates. A ray of light from the point (x, y, z) in the scene to the image point (u, v) makes an angle α with the viewing direction (i.e., the Z axis).

There are many parameterizations of the surface orientation: we choose to use (l, m), which are the X and Y components of the unit surface normal. In Figure 2, **n** is the unit normal of the surface patch located at (x, y, z); l and m are the components of this surface normal in the X and Y directions. From our viewing position we can see at most half the surfaces in the scene (i.e., those that face the viewer). The Z component of the surface normal has the magnitude $\sqrt{1 - l^2 - m^2}$, the sign determining whether the surface is forward-facing (has a positive Z component), or backward-facing (has a negative Z component). For large off-axis angle α, we see backward-facing surfaces near the edges of objects. The two components of the surface normal, l and m, do not provide an adequate parameterization of the surface in this case. Additionally, we need to know the sign of the Z component. Here we restrict ourselves to forward-facing surfaces. This minor restriction amounts to assuming that α is not too large and that we are not adjacent to an object's edge. Consequently, in this discussion we assume that the Z component of the surface normal is positive and that l and m constitute an adequate parameterization of scene surfaces.

3 IMAGE IRRADIANCE

The image irradiance equation we use is [4]

$$I(u, v) = R(l, m) \cos^4 \alpha \quad ,$$

where $I(u, v)$ is the image irradiance as a function of the image coordinates u and v, and $R(l, m)$ is the surface radiance as a function of l and m, the components of the surface normal.[2] The term $\cos^4 \alpha$ represents the off-axis effect of perspective projection. When α is small, $\cos^4 \alpha$ is approximately unity; we then have the more familiar form of the image irradiance equation. From Figure 2 we see that

$$\cos \alpha = \frac{f}{\sqrt{u^2 + v^2 + f^2}} \quad .$$

Differentiating the image irradiance equation with respect to the image coordinates u and v, we obtain

$$I'_u = R_l l_u + R_m m_u \quad ,$$
$$I'_v = R_l l_v + R_m m_v \quad ,$$

$$I'_{uu} = R_{ll} {l_u}^2 + R_{mm} {m_u}^2 + 2R_{lm} l_u m_u + R_l l_{uu} + R_m m_{uu} \quad ,$$
$$I'_{vv} = R_{ll} {l_v}^2 + R_{mm} {m_v}^2 + 2R_{lm} l_v m_v + R_l l_{vv} + R_m m_{vv} \quad ,$$
$$I'_{uv} = R_{ll} l_u l_v + R_{mm} m_u m_v + R_{lm}(l_u m_v + l_v m_u) + R_l l_{uv} + R_m m_{uv} \quad ,$$

where subscripted variables denote partial differentation with respect to the subscript(s), and

$$I'_u = \left(\frac{1}{\cos^4 \alpha}\right)\left(I_u + \frac{4uI}{u^2 + v^2 + f^2}\right) \quad ,$$

$$I'_v = \left(\frac{1}{\cos^4 \alpha}\right)\left(I_v + \frac{4vI}{u^2 + v^2 + f^2}\right) \quad ,$$

$$I'_{uu} = \left(\frac{1}{\cos^4 \alpha}\right)\left(I_{uu} + \frac{8uI_u}{u^2 + v^2 + f^2} + \frac{8u^2 I}{(u^2 + v^2 + f^2)^2} + \frac{4I}{u^2 + v^2 + f^2}\right) \quad ,$$

$$I'_{vv} = \left(\frac{1}{\cos^4 \alpha}\right)\left(I_{vv} + \frac{8vI_v}{u^2 + v^2 + f^2} + \frac{8v^2 I}{(u^2 + v^2 + f^2)^2} + \frac{4I}{u^2 + v^2 + f^2}\right) \quad ,$$

$$I'_{uv} = \left(\frac{1}{\cos^4 \alpha}\right)\left(I_{uv} + \frac{4vI_u}{u^2 + v^2 + f^2} + \frac{4uI_v}{u^2 + v^2 + f^2} + \frac{8uvI}{(u^2 + v^2 + f^2)^2}\right) \quad .$$

[2]Image irradiance is the light flux per unit area falling on the image, i.e., incident flux density. Scene radiance is the light flux per unit projected area per unit solid angle emitted from the scene, i.e., emitted flux density per unit solid angle.

If we are to use these expression to relate image measurements, e.g., I'_{uu}, to surface parameters l and m, then we must remove the derivatives of R.

4 UNIFORMLY DIFFUSE REFLECTION

To provide the additional constraints we need for relating surface orientation to image irradiance, we introduce constraints that relate properties of $R(l,m)$, — that is, constraints that specify the relationship between surface radiance and surface orientation. Such constraints are

$$(1-l^2)R_{ll} = (1-m^2)R_{mm} \quad ,$$
$$(R_{ll} - R_{mm})lm = (l^2 - m^2)R_{lm} \quad ,$$

where R_{ll} is the second partial derivative of R with respect to l, R_{mm} is the second partial derivative of R with respect to m, and R_{lm} is the second partial cross-derivative of R with respect to l and m.

These two partial differential equations embody the assumption of uniformly diffuse reflection. For uniformly diffuse reflection, $R(l,m)$ has the form

$$R(l,m) = al + bm + c\sqrt{1-l^2-m^2} + d \quad ,$$

where a, b, c, and d are constants, their values depending on illumination conditions and surface albedo. Note that l, m, and $\sqrt{1-l^2-m^2}$ are the components of the unit surface normal in the directions X, Y, and Z. $R(l,m)$ can be viewed as the dot product of the surface normal vector $(l, m, \sqrt{1-l^2-m^2})$ and a vector (a,b,c) denoting illumination conditions. As the value of a dot product is rotationally independent of the coordinate system, the scene radiance is independent of the viewing direction — which is the definition of uniformly diffuse reflection.

It is clearly evident that $R(l,m) = al + bm + c\sqrt{1-l^2-m^2} + d$ satisfies the pair of partial differential equations given above. In [1] we showed that $R(l,m) = al + bm + c\sqrt{1-l^2-m^2} + d$ is the solution of the pair of partial differential equations. These partial differential equations are an alternative definition of uniformly diffuse reflection.

It is worthy of note that $R(l,m) = al+bm+c\sqrt{1-l^2-m^2}+d$ includes radiance functions for multiple and extended illumination sources, including that for a hemispherical uniform source such as the sky. Of course, at a self-shadow edge R is not differentiable, so that the surfaces on each side of the self-shadow boundary have to be treated separately. The assumption of uniformly diffuse reflection restricts the class of material surfaces being considered, not the illumination conditions.

From the constraints for uniformly diffuse reflection, we derive the relationships

$$R_{ll} = \frac{1-m^2}{lm} R_{lm} \quad ,$$
$$R_{mm} = \frac{1-l^2}{lm} R_{lm} \quad .$$

Substituting these relationships for R_{ll} and R_{mm} in the expressions for I'_{uu}, I'_{vv}, and I'_{uv}, we obtain

$$[l_u{}^2(\frac{1-m^2}{lm}) + m_u{}^2(\frac{1-l^2}{lm}) + 2l_u m_u]R_{lm} = I'_{uu} - R_l l_{uu} - R_m m_{uu} \quad ,$$
$$[l_v{}^2(\frac{1-m^2}{lm}) + m_v{}^2(\frac{1-l^2}{lm}) + 2l_v m_v]R_{lm} = I'_{vv} - R_l l_{vv} - R_m m_{vv} \quad ,$$
$$[l_u l_v(\frac{1-m^2}{lm}) + m_u m_v(\frac{1-l^2}{lm}) + l_u m_v + l_v m_u]R_{lm} = I'_{uv} - R_l l_{uv} - R_m m_{uv} \quad .$$

By removing R_{lm} and substituting the expressions for R_l and R_m, defined by the expressions for I'_u and I'_v, we produce two partial differential equations relating surface orientation to image irradiance:

$$\alpha\theta l_{uu} + \beta\theta m_{uu} - \alpha\gamma l_{uv} - \beta\gamma m_{uv} = \chi\theta I'_{uu} - \chi\gamma I'_{uv} \quad ,$$
$$\alpha\theta l_{vv} + \beta\theta m_{vv} - \alpha\delta l_{uv} - \beta\delta m_{uv} = \chi\theta I'_{vv} - \chi\delta I'_{uv} \quad ,$$

where

$$\alpha = I'_u m_v - I'_v m_u \quad ,$$
$$\beta = I'_v l_u - I'_u l_v \quad ,$$
$$\gamma = l_u{}^2(1-m^2) + m_u{}^2(1-l^2) + 2l_u m_u lm \quad ,$$
$$\delta = l_v{}^2(1-m^2) + m_v{}^2(1-l^2) + 2l_v m_v lm \quad ,$$
$$\theta = l_u l_v(1-m^2) + m_u m_v(1-l^2) + (l_u m_v + l_v m_u)lm \quad ,$$
$$\chi = l_u m_v - l_v m_u \quad .$$

These equations relate surface orientation to image irradiance by parameter-free expressions. We make no assumptions about surface shape, nor do we need to know the parameters specifying illuminant direction, illuminant strength, and surface albedo. Our assumptions are about the properties of reflection in the world; these alone are sufficient to relate surface orientation to image irradiance. The above equations have been derived for the case of perspective projection; for orthographic projection, the primed ($'$) quantities are replaced by their unprimed counterparts, e.g., I'_u is replaced by I_u. The form of the equations is not a function of the projection used.

5 RECOVERY OF SURFACE ORIENTATION

It is difficult to solve the equations relating surface orientation to image irradiance, and thus to recover surface shape from observed image irradiance. We have used numerous integration schemes that characterize two distinct approaches. The two differential equations can be directly integrated in a step-by-step manner or, given some initial solution, a relaxation procedure may be employed. The difficulties that arise are twofold: numerical errors and multiple solutions.

Solutions of the equation $\chi = 0$ (the developable surfaces, e.g., a cylinder) are also solutions of the equations relating surface orientation to image irradiance. If the image intensities were known in analytic form, the analytic approach to solving the equations could then employ boundary conditions to select the appropriate solution. However, since the analytic form for the image intensities is unknown, numerical procedures must be employed. The use of such procedures to directly integrate the equations inevitably introduces small errors. Such errors 'mix in' multiple solutions even when those solutions are incompatible with the boundary conditions. Instability of the numerical scheme seems responsible for the fact that such errors eventually dominate the recovered solution. A scheme that is representative of our various trials at direct integration is outlined.

We transform our equations into finite-difference equations by using a three-point formula for the differentials of l and m. If $l(i,j)$ and $m(i,j)$ are the values of l and m at the (i,j)th pixel in the image, then at this pixel we use the finite-difference formulas,

$$l_u = \frac{l(i+1,j) - l(i-1,j)}{2} \quad ,$$

$$l_{uu} = l(i+1,j) + l(i-1,j) - 2l(i,j) \quad ,$$

$$l_{uv} = \frac{l(i+1,j+1) + l(i-1,j-1)}{4} - \frac{l(i+1,j-1) + l(i-1,j+1)}{4} \quad ,$$

and similar formulas for the other differentials. If we consider the 3 x 3 image patch centered on the (i,j)th pixel,

	j-1	j	j+1
i+1	O	O	&
i	O	O	O
i-1	O	O	O

we could hope that the two finite difference equations, relating the eighteen values of l and m on the patch, could be solved explicitly for $l(i+1,j+1)$ and $m(i+1,j+1)$, (the (&) cell). Such a solution would allow l and m at the (&) cell to be calculated from the l's and m's at the (o) cells. Starting at some boundary at which we know l and m at the (o) cells, we can move along the image's row and then along the successive rows, calculating l and m at the (&) cell. However, examination of the surface-orientation-to-image-irradiance equations shows that we cannot solve these equations explicitly for l_{uv} and m_{uv} and that, consequently, we cannot obtain finite-difference equations that are explicit in the l and m of the (&) cell.

We avoid this difficulty by combining the two surface-orientation-to-image-irradiance equations into one and using surface continuity to provide the additional equation. Removing l_{uv} and m_{uv} from the differential equations, we have

$$\alpha(\delta l_{uu} - \gamma l_{vv}) + \beta(\delta m_{uu} - \gamma m_{vv}) = \chi(\delta I'_{uu} - \gamma I'_{vv}) \quad .$$

Surface continuity requires that $\frac{\partial^2 z}{\partial x \partial y} = \frac{\partial^2 z}{\partial y \partial x}$, from which it follows that

$$l_y(1-m^2) + m_y lm = m_x(1-l^2) + l_x lm \quad .$$

Provided that u and v are small compared with z (e.g., in the eye or in a standard-format camera), then

$$l_v(1-m^2) + m_v lm = m_u(1-l^2) + l_u lm \quad .$$

These two equations, which do not involve l_{uv} or m_{uv}, form a basis for finite difference equations that calculate l and m at the (-) cell from values of l and m at (+) cells.

	−	
+	+	+
	+	

The results obtained with the above integration scheme, together with many variations of it, are poor. Accurate values for l and m are obtained only within approximately five to ten rows of the known boundary. This is the case for noise-free image data. These results can be understood by examination of the finite-difference equations. The explicit expressions for l and m at the (-) cell are functions of the differences of l and m at the (+) cells. Such schemes are usually numerically unstable, making step-by-step integration impossible. While the failure to find a stable numerical scheme does not imply that one does not exist, our difficulty highlights the problem of finding numerical schemes, based on differential models, to propagate information from known boundaries. (One wonders whether nature experienced the same difficulties when designing the human vision system.)

Although the alternative to direct integration, a relaxation procedure to solve the equations, seems to offer relief from the numerical instability of direct integration, it nevertheless poses its own problems. The approach we used parallels the one in [3] for solving the image irradiance equation when the surface albedo and illumination conditions are known. For each image pixel we form three error terms: the residuals associated with the two surface-orientation-to-image-irradiance equations, and with the one surface continutiy equation. Minimizing the sum of the errors over the whole image with respect to l and m at each pixel produces an updating rule for l and m at each pixel. Given an initial solution, i.e., assignment of values for l and m at each pixel, a relaxtion scheme, like the one described, is useful only if it converges. While the constraint imposed by the underlying model is most important in ensuring convergence, the importance of a good initial solution for a relaxation method cannot be overemphasized. Simplifying the two partial differential equations (by using additional assumptions) provides a method for obtaining an good initial solution.

The spherical approximation assumes that we are viewing a spherical surface. This implies $l_y = 0, m_x = 0$, and $l_x = m_y$, — namely, constant curvature that is independent of direction. Provided that u and v are small compared with z, then $l_v = 0, m_u = 0$ and $l_u = m_v$. For this case, the partial differential equations become relationships between image irradiance and its derivatives, on the one hand, and the components of the surface normal, on the other:

$$\frac{1-m^2}{lm} = \frac{I'_{uu}}{I'_{uv}} \quad ,$$
$$\frac{1-l^2}{lm} = \frac{I'_{vv}}{I'_{uv}} \quad .$$

The spherical-approximation results for perspective projection are similar to those Pentland was able to obtain [2] for orthographic projection through local analysis of the surface. Besides providing a mechanism for obtaining an initial solution for a relaxation-style algorithm, they allow surface orientation to be estimated by purely local computation. Such an estimate will be exact when the surface is locally spherical.

The results of our experiments with relaxation procedures are easily summarized: the relaxation procedures were not convergent. While such nonconvergence is hardly unusual, the reasons for failure, however, are instructive. The residuals associated with both the surface-orientation-to-image-irradiance equations, and the surface continuity equations remain small during the relaxation, even when the solution is starting to diverge. Of course the residuals are not as small as they are when on the verge of solution, but they are small enough to make one believe that a solution has been obtained, particularly when the image is not noise-free. Apparently the equations are insensitive to particular values of l and m, being more concerned with the values of l_u,l_v,m_u, and m_v. As with direct integration, relaxation models need boundary conditions to select a particular solution. We used various boundary conditions in our relaxation experiments, but it is difficult to believe that a model, apparently insensitive to surface orientations, could be overly influenced by the surface orientations at a boundary.

Our two approaches, direct integration and relaxation, have not yielded a computational solution to the problem of recovering surface orientation from shading. The attractiveness of local computation is clear; it has neither numerical instability nor divergent behavior, but the cost it imposes is that assumptions must be made about surface shape. A compromise between some local computation and some information propagation may offer an approach that is not overly restrictive in its assumptions about surface shape. However, the question needs to be considered: Is the model underconstrained? Is shape recovery dependent on information other than shading? What other information (that is obtainable from the image), is necessary to enable the construction of effective shape-recovery algorithms?

6 RECONSTRUCTION OF THE SURFACE SHAPE

Surface orientation is not the same as surface shape. However, once we have obtained the surface orientation as a function of image coordinates, i.e., $l(u,v)$ and $m(u,v)$, we can use these to reconstruct the surface shape in the scene coordinates X,Y,Z. We derive a suitable formula.

Suppose we know the depth z_0 at scene coordinates (x_0, y_0, z_0), corresponding to (u_0, v_0) in the image. For the point $(x_0 + \Delta x, y_0 + \Delta y)$ we use the approximation

$$z(x_0 + \Delta x, y_0 + \Delta y) = z(x_0, y_0) + \Delta x \left.\frac{\partial z}{\partial x}\right|_{x_0,y_0} + \Delta y \left.\frac{\partial z}{\partial y}\right|_{x_0,y_0} \quad .$$

Similarly,

$$z(x_1 - \Delta x, y_1 - \Delta y) = z(x_1, y_1) - \Delta x \left.\frac{\partial z}{\partial x}\right|_{x_1,y_1} - \Delta y \left.\frac{\partial z}{\partial y}\right|_{x_1,y_1} \quad .$$

If $x_1 = x_0 + \Delta x$ and $y_1 = y_0 + \Delta y$, then

$$z(x_1, y_1) = z(x_0, y_0) + \frac{x_1 - x_0}{2}\left(\left.\frac{\partial z}{\partial x}\right|_{x_0,y_0} + \left.\frac{\partial z}{\partial x}\right|_{x_1,y_1}\right) + \frac{y_1 - y_0}{2}\left(\left.\frac{\partial z}{\partial y}\right|_{x_0,y_0} + \left.\frac{\partial z}{\partial y}\right|_{x_1,y_1}\right) \quad .$$

Using the perspective transformation $u = -f\frac{x}{z}$ and $v = -f\frac{y}{z}$ to remove x and y, we obtain

$$z(u_1, v_1) = z(u_0, v_0) \times \frac{2f + u_0\left(\left.\frac{\partial z}{\partial x}\right|_{u_0,v_0} + \left.\frac{\partial z}{\partial x}\right|_{u_1,v_1}\right) + v_0\left(\left.\frac{\partial z}{\partial y}\right|_{u_0,v_0} + \left.\frac{\partial z}{\partial y}\right|_{u_1,v_1}\right)}{2f + u_1\left(\left.\frac{\partial z}{\partial x}\right|_{u_0,v_0} + \left.\frac{\partial z}{\partial x}\right|_{u_1,v_1}\right) + v_1\left(\left.\frac{\partial z}{\partial y}\right|_{u_0,v_0} + \left.\frac{\partial z}{\partial y}\right|_{u_1,v_1}\right)} \, .$$

As $\frac{\partial z}{\partial x} = \frac{-l}{\sqrt{1-l^2-m^2}}$ and $\frac{\partial z}{\partial y} = \frac{-m}{\sqrt{1-l^2-m^2}}$, we have the means of reconstructing the surface in scene coordinates from the values of surface orientation in image coordinates.

7 CONCLUSION

In this formulation of the shape-from-shading task, we have eliminated the need to know the explicit form of the scene radiance function by introducing higher-order derivatives into our model. This model is applicable to natural scenery without any additional assumptions about illumination conditions or the albedo of the surface material. However, without a computational scheme to reconstruct surface shape from image irradiance we may wonder if we have surrendered too much. The difficulties of finding a computational scheme must induce one to ask whether the model is underconstrained. Have we applied too few restrictions, thereby making shape recovery impossible? Notwithstanding the general concern about underconstraint of the model, the numerical difficulties encounted makes local computation of scene parameters attractive. Information propagation methods must always cope with the problem of accumulated errors. In our model, however, to achieve local computation we must make assumptions with regard to surface shape. What other information, besides shading, do we need to know if we are to recover surface shape? Can we find moderate restrictions that allow mostly local computation of the surface shape parameters? We are actively engaged in the pursuit of such procedures.

REFERENCES

1. Smith, G.B., From image irradiance to surface orientation, (submitted for publication). Available as Technical Note 273, Artificial Intelligence Center, SRI International (1982).
2. Pentland, A.P., The visual inference of shape: computation from local features, Ph.D. Thesis, Department of Psychology, Massachusetts Institute of Technology (1982).
3. Ikeuchi, K. and Horn, B.K.P., Numerical shape from shading and occluding boundaries, *Artificial Intelligence* **17** (1981) 141-184.
4. Horn, B.K.P. and Sjoberg R.W., Calculating the reflectance map, A.I. Memo 498, Artificial Intelligence Laboratory, Massachusetts Institute of Technology (1978).