

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

A.I. Memo No. 697

August, 1982

BINOCULAR SHADING AND
VISUAL SURFACE RECONSTRUCTION

W.E.L. Grimson

ABSTRACT: Zero-crossing or feature-point based stereo algorithms can, by definition, determine explicit depth information only at particular points in the image. To compute a complete surface description, this sparse depth map must be interpolated. A computational theory of this interpolation or reconstruction process, based on a *surface consistency constraint*, has previously been proposed. In order to provide stronger boundary conditions for the interpolation process, other visual cues to surface shape are examined in this paper. In particular, it is shown that, in principle, shading information from the two views can be used to determine the orientation of the surface normal along the feature-point contours, as well as the parameters of the reflective properties of the surface material. The numerical stability of the resulting equations is also examined.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-80-C-0505 and in part by National Science Foundation Grant 79-23110MCS.

1. Introduction

1.1. The Global Problem

One of the goals of a vision system is to compute the three-dimensional shape of the visible surfaces in a scene. The human visual system uses many cues to compute surface shape, with different modules of the system using varying sources of information in the images to infer information about surface shape. Examples include motion analysis, stereopsis, shading, and texture. How do these different visual modalities contribute to the computation of surface shape?

One computational approach to understanding the human visual system, pioneered by Marr and Poggio (see, for example, [Marr, 1976, 1982; Marr and Poggio, 1977]), views the computation of surface shape from images, in part, as a collection of transformations between two main representations. The first representation is the *primal sketch*, which makes explicit loci of changes in image irradiance at particular scales of resolution; the second is the $2\frac{1}{2}$ -D *sketch*, which makes explicit information about surface shape and reflective properties of the surface material. The modules that compute information feeding the $2\frac{1}{2}$ -D sketch from the primal sketch have generally been considered to a first approximation to be independent of one another. It is clear, however, that within the $2\frac{1}{2}$ -D sketch the different sources of information should interact, both to maintain consistency among the data provided by different modules, and to provide feedback to the modules in order to enhance the acquired data (for example, texture contours can facilitate stereopsis by driving vergence eye movements [Kidd, et al., 1979]). In this paper, we are interested in examining interactions at the level of the $2\frac{1}{2}$ -D sketch between modules of the early visual system. In particular, we will investigate some of the ways in which shading information can augment stereo data.

1.2. The Motivating Problem

The goal of the $2\frac{1}{2}$ -D sketch [see, for example, Marr, 1978, 1982] is to compute surface parameters, in particular, the distance to and orientation of small patches of the visible surfaces, the discontinuities in those surfaces (for example, the edges of objects), and possibly the properties of the surface material (for example, the amount of specularity, the colour and the albedo of the surface material). Representations similar to the $2\frac{1}{2}$ -D sketch have also been suggested by Horn [1979] and Barrow and Tenenbaum [1979]. As mentioned, the input to this representation is provided by a number of roughly independent modules, two important ones of which are stereopsis and motion perception. The input provided by these particular modules, which is characteristic of many early visual modules, consists of explicit information about surface shape or disposition only at particular points in the image.

This follows from the form of the primal sketch. Both psychophysical investigations [for example, Attneave, 1954; Barlow, 1961] and computational studies [for example, Roberts, 1965; Herskovitz and Binford, 1970; Horn, 1973; Marr, 1976; Marr and Hildreth, 1980, and reviews by Rosenfeld and Kak, 1976; and Pratt, 1978] of early visual processing suggest that most of the information in an image is carried by the intensity changes, and hence, that the first stage

of analysis of images is the detection of such changes. These changes in intensity are used as input by the modules that feed the $2\frac{1}{2}$ -D sketch and hence explicit information is obtained only at those locations. For example, in stereo computations [Marr and Poggio, 1979; Grimson, 1981; Baker and Binford, 1981; Mayhew and Frisby, 1981] the zero-crossing contours of the primal sketches [Marr and Hildreth, 1980] for the left and right eye are placed in correspondence, and the difference in projection of the corresponding contours in the two eyes is used to determine the depth of the surface along that contour. To create a complete surface representation it is necessary to interpolate between the known values of the raw $2\frac{1}{2}$ -D sketch. An initial theory of this process that implicitly takes into account some of the shading information available in the primal sketch has been proposed [Grimson, 1982a, 1982b; see also refinements by Terzopoulos, 1982].

While this interpolation or reconstruction theory constructs the "best" surface to fit through the known depth points provided by stereopsis, it is clear that, in principle, supplying additional constraints to the reconstruction process would result in an improved accuracy in the computed surface shape.

An extreme example of this can be seen in Figure 1, which shows one possible set of depth constraints obtained by stereoscopically viewing a sinusoidal depth grating of a uniform material from a particular viewpoint. The most conservative reconstruction, without explicitly taking the shading information into account, would be the plane shown in Figure 1b. On the other hand, if surface orientation information were also available along the depth contours, the more correct reconstruction of Figure 1c would be obtained. Thus, the motivation is to find algorithms for computing surface orientation at the known depth points provided by stereopsis.

1.3. The Specific Problem

Given that we want to augment the boundary conditions supplied to the reconstruction process, the specific question to be investigated here is whether shading information along the zero-crossing contours of the primal sketch can be used to provide information about the surface orientation along those contours.

The traditional approach to the shape-from-shading problem, pioneered by Horn [1970, 1975], has been monocular. The goal has been to extract explicit surface orientation values at each point in the image, from a single view of the scene. As stated, the problem is considerably underconstrained and thus additional information is required in order to obtain the surface orientation information (see for example [Ikeuchi and Horn, 1981; Bruss, 1981; Brooks, 1982]). For example, it is usually assumed that the direction and strength of the illuminant and the reflective properties of the surface material (which are assumed constant over large sections of the image) are known. One method for obtaining the additional information necessary to solve for the surface orientation, under these assumptions, is to use the technique of *photometric stereo* [Woodham, 1978, 1980, 1981; Ikeuchi and Horn, 1979; Silver, 1980]. In this case, multiple images obtained from a single viewpoint, but with different illuminant directions, are used to determine the surface orientation at points in the image.

The intention in this paper is to solve (at least partially) the shading problem with fewer






Figure 1. Part (a) shows an extreme example of a set of depth constraints that could be obtained by stereoscopically viewing a sinusoidal depth grating. If shading information is not explicitly accounted for, the best surface reconstruction is the plane shown in part (b). On the other hand, if the shading information is used to obtain surface orientation information along the depth contours, the more accurate reconstruction of part (c) is obtained.

requirements or assumptions about the parameters of the imaging process. To do so will require either additional sources of information about shading or weaker expectations on the delivered data, or both. Indeed, since the motivating problem is to obtain additional boundary conditions on the reconstruction process, we will only expect surface orientation information at the zero-crossings, rather than everywhere in the image. We will further only expect a coarse estimate of the surface normal, since the values will serve as boundary conditions for a surface approximation process, rather than constituting an explicit surface reconstruction in itself. Moreover, we will introduce an additional source of information to the process. In particular, we will investigate the added information provided by the differing shading information observed at nearby viewpoints. Hence, we are concerned with *binocular shading*, especially as it applies to the surface reconstruction problem. The main questions to be addressed here are: (1) Using binocular shading, under what conditions can we solve for the surface orientation at the zero-crossing contours that have been matched by stereopsis, and (2) How numerically stable is the resulting process? We will see that the method used to solve for the surface orientation is essentially an alternate form of photometric stereo, using images obtained from different viewpoints, rather than with different illuminants.

2. The Basic Equations

The brightness recorded at a point in an image can generally be considered as a product of three factors, the total amount of incident light striking the surface, the percentage of such incident light which is reflected (as opposed to absorbed or transmitted), usually denoted by the *albedo*, ρ , of the surface, and the *reflectance* of the surface, the distribution of the reflected light as a function of direction. Here, we shall assume that the intensity of the incident light is normalized, by incorporating it into the definition of ρ . As a consequence, the parameter ρ is not restricted to the range 0 — 1, but rather can take on any positive value.

Using a right-hand coordinate system with origin at the left eye, x axis connecting the two eyes and the negative z axis pointing straight ahead, a surface may be represented by a Monge patch

$$\{x, y, -f(x, y)\}$$

and the surface normal at a point, by

$$\{p, q, 1\}$$

where

$$p = \frac{\partial f}{\partial x} \quad q = \frac{\partial f}{\partial y}.$$

The situation to be investigated here is one in which the illumination geometry consists of an arbitrary surface illuminated by a point source, sufficiently distant from the surface relative to the distance to the viewer. Under most circumstances (for example, the surface reflectance is isotropic), the reflectance properties of the surface can be combined with the illumination geometry and the viewer position into the convenient representation of a *reflectance map* [Horn, 1977].

For an arbitrary surface, the surface reflectance map will generally be a combination of two effects. The first is a diffuse or matte component, which results from the scattering of light that

penetrates some distance into the microstructure of the surface. The second is a specular or luster component, which results from the mirror-like reflection of light at a smooth interface between two materials of different refractive index. These two components can be combined into a single reflectance map in the following manner (see, for example, [Horn, 1981; Blinn, 1977; Phong, 1975]).

Let \mathbf{n} denote a unit vector in the direction of the surface normal, defined at a point on the surface, let \mathbf{s} denote a unit vector in the direction of the point light source, and let \mathbf{v} denote a unit view vector. (In general, we will consider objects to be distant relative to the viewer, so that \mathbf{v} is essentially constant, and will consider objects to also be distant relative to the light source, so that \mathbf{s} is essentially constant). Then the recorded brightness is given by

$$\mathcal{E} = \rho R(\mathbf{n})$$

and the reflectance map is given by

$$R(\mathbf{n}) = \left[(1 - \alpha)(\mathbf{n} \cdot \mathbf{s}) + \alpha \{ \mathbf{n} \cdot \mathbf{h} \}^k \right] \quad (1)$$

where \mathbf{h} is a unit vector in the direction of the off specular angle,

$$\mathbf{h} = \frac{(\mathbf{s} + \mathbf{v})}{\sqrt{2}\sqrt{1 + (\mathbf{s} \cdot \mathbf{v})}}$$

Here, ρ , α and k are parameters of the particular surface, with ρ and k non-negative and α ranging between 0 and 1.

Note that \mathbf{n} is a function of the surface f while \mathbf{s} is a function of the illumination geometry. For our particular choice of coordinate system, the view vector for the left eye is particularly simple, $\mathbf{v}_L = \{0, 0, 1\}$. Now suppose that we take a second view of the surface, with vergence angle β (see Figure 2). Here, the view vector is given by $\mathbf{v}_R = \{-\sin \beta, 0, \cos \beta\}$. Note that the vergence angle β and the illuminant direction are assumed to be known, so that the vectors \mathbf{v}_R and \mathbf{s} are constants.

The first term in expression (1) corresponds to a matte or Lambertian component of the surface. Since this component depends only on the angle between the light source and the surface normal (and hence is independent of \mathbf{v}), it is identical for both views. Thus, no additional information will be obtained from two views in the case of a perfectly matte surface. The second term corresponds to a specular component of the surface, and this term, in general, does change between the two views, as can be seen in expression (1).

Our intent is to use the two expressions for recorded brightnesses in the two eyes to determine the orientation of the viewed surface at some point. In order to compare irradiance values from the two views, we must be certain that the points in the images at which the irradiance measurements are taken are in correspondence (that is, we know which points in the two images correspond to the same point on the surface). In the ideal case of isolated intensity edges, matched zero-crossings from a stereo algorithm [for example, Grimson, 1981] are in correspondence, so by restricting our attention to such points, we are in principle guaranteed irradiance measurements that correspond to the same point on a physical surface. We should note, however, that in practice a certain amount




Figure 2. The geometry of two views of the surface, with the vergence angle β indicated.

of error will be associated with the matched zero-crossings, due to uncertainty in exactly localizing the underlying intensity changes. This error will reduce the effectiveness of the computation.

There are two stages to the computation of surface orientation. The first is to compute the surface reflectance parameters ρ , α and k (in some applications, these might be known *a priori*). Once the reflectance map is known, we must then solve for the surface orientation along the zero-crossing contours, assuming that ρ , α and k are constant along these contours. More precisely, we solve the following problem. Let the recorded brightnesses in the left and right eyes be

$$\begin{aligned} \mathcal{E}_L &= \rho[(1 - \alpha)m + \alpha u^k] \\ \mathcal{E}_R &= \rho[(1 - \alpha)m + \alpha v^k] \end{aligned}$$

respectively, where

$$\begin{aligned} m &= \mathbf{n} \cdot \mathbf{s} \\ u &= \frac{(\mathbf{s} + \mathbf{v}_L) \cdot \mathbf{n}}{\sqrt{2}\sqrt{1 + \mathbf{s} \cdot \mathbf{v}_L}} \\ v &= \frac{(\mathbf{s} + \mathbf{v}_R) \cdot \mathbf{n}}{\sqrt{2}\sqrt{1 + \mathbf{s} \cdot \mathbf{v}_R}} \end{aligned}$$

We will first solve for ρ , α and k over some portion of the image, and then use these values to solve for the surface orientation along zero-crossing contours in this region of the image.

3. Solving for the Reflectance Parameters

The task of determining the reflectance parameters would be considerably eased if the surface orientation (p, q) were already known. We thus restrict our attention to situations in which this is

true. Consider a point of high curvature along a zero-crossing contour, which in the limiting case might be a corner where two zero-crossing contours intersect. The orientation of the contour in the image plane, along either direction of the contour, is different (within the resolution of the image grid). A directional derivative taken along one direction of the contour yields a constraint on the relationship between p and q at this point. Since there are two different directions along the contour at such a point, we get two different constraints on the surface orientation, each one corresponding to a straight line in gradient space, defining a linear relationship between the values of p and q at the point. The intersection of these two lines uniquely defines the surface orientation at the corner.

We can also perform this computation in a somewhat non-standard representation, by using unit normals. The representation we choose to utilize is obtained by orthonormally projecting the Gaussian hemisphere of all visible unit normals onto the plane spanned by the first two components of the unit normal. (Note that this is a one-to-one, onto projection, since we are only interested in the half of the Gaussian sphere composed of the unit normals corresponding to visible surfaces). Thus, rather than performing computations in gradient space, we instead use the components of the unit normals lying in the unit disk. In this case, each directional derivative constrains the first two components of the unit normal to lie on a half ellipse in the unit disk. The intersection of two such half ellipses uniquely defines the components of the unit surface normal at the point.

3.1. Using Irradiance Derivatives

We now consider methods for determining the surface reflectance parameters at such a point. Initially we have three unknowns and only two equations. To uniquely solve for the unknowns, we will need additional information. One possibility is to consider using derivatives of the image irradiances to solve for the parameters. Taking first directional derivatives of the irradiance equations introduces three new variables, $p_x, p_y (= q_x)$ and q_y . (Note that q_x is assumed to be equal to p_y , since we shall assume that the surfaces are at least twice continuously differentiable.) This gives us a total of 8 unknown variables, the three listed, as well as p, q, ρ, α and k . At a corner, we can compute the values of p and q . Furthermore, second directional derivatives along the depth contour yield two additional linear constraints on p_x, p_y and q_y , so there are only 4 independent variables. We also have four equations, namely:

$$\begin{aligned}\frac{\partial \mathcal{E}_L}{\partial x} &= \rho(1 - \alpha)\{m_p p_x + m_q p_y\} + \rho \alpha k u^{k-1} \{u_p p_x + u_q p_y\} \\ \frac{\partial \mathcal{E}_L}{\partial y} &= \rho(1 - \alpha)\{m_p p_y + m_q q_y\} + \rho \alpha k u^{k-1} \{u_p p_y + u_q q_y\} \\ \frac{\partial \mathcal{E}_R}{\partial x} &= \rho(1 - \alpha)\{m_p p_x + m_q p_y\} + \rho \alpha k v^{k-1} \{v_p p_x + v_q p_y\} \\ \frac{\partial \mathcal{E}_R}{\partial y} &= \rho(1 - \alpha)\{m_p p_y + m_q q_y\} + \rho \alpha k v^{k-1} \{v_p p_y + v_q q_y\}\end{aligned}$$

where the subscripts L and R denote the left and right images, and the other subscripts denote partial differentiation.

3.2. Solving The Equations

We now seek a closed form solution to this set of equations; that is, we seek expressions for determining ρ , α and k as functions of the image irradiance derivatives, the illuminant direction and the vergence angle. While it is possible to obtain such a closed form solution (for completeness, the solution is given in the appendix), it is not very useful for computational purposes, since it is numerically very sensitive, as discussed below.

3.3. Numerical Stability

While the appendix describes a solution to the problem of determining the parameters of the surface reflectivity, numerical simulations indicate that the solution is not numerically stable. This is illustrated by the following example, used to determine whether the image values $\frac{\partial}{\partial x}(\mathcal{E}_L - \mathcal{E}_R)$ and $\frac{\partial}{\partial y}(\mathcal{E}_L - \mathcal{E}_R)$ can be reliably measured from an image.

Since a visual system is not an infinite resolution device, it is important to take this into account when simulating the solution to the above equations. In particular, one cannot use analytic expressions for terms such as $\frac{\partial}{\partial x}(\mathcal{E}_L - \mathcal{E}_R)$ but rather, one must use finite difference approximations to such differential expressions, in order to reflect the discrete computational structure of the vision system. Moreover, the range of values for \mathcal{E}_L and \mathcal{E}_R should also be discrete, rather than continuous, since any sensor will have a finite resolution of brightness values. These considerations turn out to be critical in evaluating the numerical stability of the computation. In particular, the closed form solution for the reflectance parameters given in the appendix in principle provides exact values for those parameters, if infinite precision data is available. When discrete approximations are used, however, this exact solution is no longer possible, as is illustrated by the following example.

A synthetically shaded sphere of radius 160 pixels with a total irradiance range of 8 bits (or 256 grey levels) was constructed using equation (1) with the parameters $\rho = 256$, $\alpha = 0.8$ and with k varying from 1 to 10. Note that this implies that recorded brightness was discretized to a finite precision, thereby restricting the resolution of measurements of spatial change in the difference between the image brightnesses.

The ratio of interocular separation to fixation distance was set at 6:100. The total number of points visible to both eyes was counted, and the percentage of such points for which the finite differences corresponding to the above partial derivatives exceeded a threshold of 1 grey level, or 1 part in 256, was measured. This percentage was observed to be small, no more than 4%, even for large values of k . This held true over a range of illuminant directions. Some examples are shown in Table 1. (Similar results hold for other values of α .)

k	vis	%	k	vis	%
1	99.9	2.27	1	89.7	2.64
2	99.9	0.00	2	89.7	0.08
3	99.9	0.00	3	89.7	0.00
4	99.9	0.00	4	89.7	0.24
5	99.9	1.34	5	89.7	1.77
6	99.9	1.99	6	89.7	2.43
7	99.9	2.33	7	89.7	3.03
8	99.9	2.56	8	89.7	3.42
9	99.9	3.19	9	89.7	3.94
10	99.9	3.67	10	89.7	4.36

Table 1: Sensitivity in computing derivatives of image differences

The table lists the percentage of points on the sphere, visible in both eyes, for which the finite difference approximation of the partial derivative of the change in image intensity exceeded a threshold of one part in 256. Here, k denotes the exponent of the specular component, and vis is the percentage of all possible points on the sphere which are visible to both eyes. The first two components of the unit source normal, p_s and q_s , are given by $p_s = 0, q_s = 0$ in the left columns, and $p_s = 0.557, q_s = 0.239$ in the right, and the vergence angle is $\beta = \tan^{-1} \frac{6}{100}$.

This observation suggests that the system of equations listed in the appendix, while providing a solution in the ideal case, will not provide a numerically stable solution. One explanation for this can be seen by noting that any partial derivatives of the irradiance equations will involve partial derivatives of the components of the surface normal. Unless the surface has a significant curvature at these points, the partial derivatives of the image irradiances will be small and hence the computation of the surface parameters will be numerically unstable. For example, if a smaller sphere had been used in the above example, the numerical stability of the solution to the reflectance parameters would be improved, since the effective curvature of the sphere would be increased.

While using finite differences in the change in image brightness between the two views did not provide detectable values, it was observed that the simple image differences frequently did. Examples are shown in Table 2. Here, we observe that even for strongly specular surfaces ($k = 10$) the percentage of orientations with a detectable (≥ 1 grey-level) difference in image brightnesses is over 55%. Allowing for image sensor noise (for example, image differences ≥ 3) still leaves an appreciable percentage of orientations with measurable differences.

$p_s = 0.557 \quad q_s = 0.239$				
k	% total	% ≥ 1	% ≥ 2	% ≥ 3
1	96	85	53	25
2	92	91	63	37
3	89	89	62	39
4	89	83	57	38
5	88	77	53	36
6	88	71	49	34
7	88	66	45	32
8	88	62	42	31
9	88	58	40	29
10	88	55	38	28

Table 2: Sensitivity in computing image differences.

The table lists the percentage of points on the sphere, visible in both eyes with a recorded intensity of at least 5 parts in 256, for which the difference in image intensity exceeded a given threshold. Here, k denotes the exponent of the specular component, $\beta = \tan^{-1} \frac{6}{100}$, and $p_s = 0.557$ and $q_s = 0.239$. The first column lists the percentage of points on the sphere which are visible in both eyes with sufficient brightness. The remaining three columns list the percentage of such points which have an irradiance difference in the two eyes of the amount listed.

This suggests that while we may not be able to reliably measure the spatial derivatives of the difference in image brightness, due to the discrete nature of the sensing of brightness values and the discrete approximations to the spatial derivatives, we may be able to rely on the simple differences in image brightness. We thus turn to the problem of finding a solution for the surface reflectivity parameters from the image brightnesses alone.

4. A Second Solution

We first observe that even recording the irradiances at a corner (defined here as a point of high curvature), where the surface orientation is known, is not sufficient to solve the problem, since we have two equations in three unknowns. We thus require an additional source of information and will assume that we know the surface orientation at two distinct points, (we will show later that these could be two nearby corners or simply one corner and a nearby point).

At each corner, we note that the following expressions hold (where i denotes the image point in question)

$$\begin{aligned}\alpha_i &= \frac{2m_i(\mathcal{E}_L - \mathcal{E}_R)_i}{(\mathcal{E}_L + \mathcal{E}_R)_i(u_i^k - v_i^k) - (\mathcal{E}_L - \mathcal{E}_R)_i(u_i^k + v_i^k - 2m_i)} \\ &= \frac{1}{1 + \frac{(\mathcal{E}_R)_i u_i^k - (\mathcal{E}_L)_i v_i^k}{m_i(\mathcal{E}_L - \mathcal{E}_R)_i}} \\ \rho_i &= \frac{(\mathcal{E}_L + \mathcal{E}_R)_i(u_i^k - v_i^k) - (\mathcal{E}_L - \mathcal{E}_R)_i(u_i^k + v_i^k - 2m_i)}{2m_i(u_i^k - v_i^k)} \\ &= \frac{(\mathcal{E}_L - \mathcal{E}_R)_i}{u_i^k - v_i^k} \frac{1}{\alpha_i}\end{aligned}$$

This gives us a solution for ρ and α , presuming we know the value of k . To determine k , we assume that the surface reflectivity parameters are constant over the region of the image spanned by the two known points. In this case, $\alpha_1 = \alpha_2$ and $\rho_1 = \rho_2$ and this leads to the following implicit equations for k :

$$\frac{(\mathcal{E}_L - \mathcal{E}_R)_1}{(\mathcal{E}_L - \mathcal{E}_R)_2} = \frac{m_2(u^k \mathcal{E}_L - v^k \mathcal{E}_R)_1}{m_1(u^k \mathcal{E}_L - v^k \mathcal{E}_R)_2}$$

and

$$\frac{(\mathcal{E}_L - \mathcal{E}_R)_1}{(\mathcal{E}_L - \mathcal{E}_R)_2} = \frac{(u^k - v^k)_1}{(u^k - v^k)_2}$$

Either equation can be solved numerically to determine the value of k and hence of the other reflectivity parameters α and ρ .

4.1. Numerical Stability

While the equations derived in section 3 were observed to be numerically unstable, the above solution is much more sound. This can be observed by the following simulation. Consider a sphere illuminated by a point source, and observed by two sensors whose angle of vergence is $\beta = \tan^{-1}(\frac{6}{100})$ (roughly equivalent to a human observer fixating on an object 1 meter removed). The surface of the sphere can be given any surface reflectance properties by specifying the parameters k , ρ and α . For some set of parameters, we compute the observed brightness value in each sensor to one part in 256. We then apply the equations derived above (assuming initially that the value of k is known) to compute the parameters α and ρ . The accuracy of the computed values is compared to the original value under a number of situations.

Since a certain amount of error will follow from the truncation of brightness accuracy, we measured the error in computing the parameters as a function of the difference in observed brightness values. It will be noted that we can trade off accuracy of computed parameters with the percentage of orientations for which such a computation can be made (see Table 3). In particular, as we increase the required lower bound on the observed differences in brightness, we decrease the total percentage of orientations for which a computation can be made, but increase the percentage of such points which give rise to computed surface parameters within a particular error range. Table 3 illustrates some examples.

k	vis	$\rho-1$	$\rho-5$	$\rho-10$	$\alpha-1$	$\alpha-5$	$\alpha-10$
1	25	21.9	52.6	62.0	5.1	17.8	31.0
2	37	69.7	95.8	98.4	6.8	28.5	54.6
3	39	66.3	95.2	99.4	11.9	39.8	66.6
4	38	60.1	93.6	99.0	13.9	44.2	72.0
5	36	57.7	92.2	98.6	13.9	46.1	76.0
6	34	56.4	91.8	98.4	15.2	49.3	78.6
7	32	56.4	91.0	98.2	16.3	52.4	80.3
8	31	55.3	90.6	98.1	15.0	52.7	82.7
9	29	55.4	90.9	98.3	18.3	55.5	84.1
10	28	55.2	90.1	98.0	17.7	56.8	85.0

k	vis	$\rho-1$	$\rho-5$	$\rho-10$	$\alpha-1$	$\alpha-5$	$\alpha-10$
1	0	0.0	0.0	0.0	0.0	0.0	0.0
2	0	0.0	0.0	0.0	0.0	0.0	0.0
3	2	51.3	98.0	100.0	0.0	0.0	13.3
4	7	73.2	99.0	100.0	20.0	53.7	81.2
5	11	72.5	99.1	100.0	17.1	52.9	86.5
6	13	70.6	98.5	100.0	18.6	57.6	87.9
7	14	71.0	98.8	100.0	19.4	61.8	90.1
8	14	69.6	98.6	100.0	16.7	61.3	92.7
9	14	69.9	98.4	100.0	22.6	64.7	92.7
10	14	70.0	98.5	100.0	21.2	66.2	94.7

Table 3: Errors in computing α and ρ .

The following parameters are defined:

vis is the percentage of viewed orientations on a sphere for which the difference in observed brightnesses exceeds some threshold;

$\alpha - \epsilon$ is the percentage of vis which lie within a range $\pm \epsilon\%$ of the correct value of α ;

$\rho - \epsilon$ is the percentage of vis which lie within a range $\pm \epsilon\%$ of the correct value of ρ ;

ϵ is the percentage error bar allowed about the correct value.

In this case, the illuminant direction is given by $p_s = 0.557, q_s = 0.239$. In the first table, the threshold of image differences was set at 3, in the second table, the threshold was 5.

To compute the value of k , we must solve an implicit equation in k . This can be done numerically, for example, by Newton-Raphson. While computer simulations indicate that this computation is very robust in the case of perfect image brightnesses, the computation is much less stable when considering truncated brightness values (i.e. if infinite precision brightness values are used, the computation returns very accurate values for k , but when the finite resolution of the brightness sensor is taken into account, the accuracy of the computation drops off quickly). Table 4 indicates an example of the error associated with computing k numerically. This table was computed by choosing two surface orientations at random, such that the difference in brightness exceed a threshold of 3 grey levels, and solving the implicit equation for k by a Newton-Raphson method.

One possible method for improving the accuracy of computing not only k , but also ρ and α , is to consider the reflective properties of the surface to be constant over some extended region of the image. In this case, a number of points may be sampled, and the computed values for the parameters averaged. Table 4 illustrates the reduction in error associated with computing k .

	1	3	5	10
1%	2.7	2.7	2.8	4.3
5%	14.8	15.6	16.2	18.6
10%	27.8	30.6	29.7	35.3
20%	47.2	50.7	54.6	65.3
30%	62.6	65.9	73.6	84.5
40%	69.3	77.4	83.9	94.7
50%	75.0	85.4	92.0	98.5

Table 4: Errors in computing k .

Two different orientations were chosen at random, such that the difference in brightness values at each point exceeded a threshold of 3 grey levels. The percentage of such pairs of points with percentage error in computing k within the given range are indicated. In the case illustrated below, the illuminant direction was given by $p_s = 0.557, q_s = 0.239$, with $\rho = 256, \alpha = 0.5$ and $k = 5$. The percentages are shown for the case of sampling 1, 3, 5, and 10 pairs, and averaging the results.

5. Solving for the Surface Orientation

Given that the equations derived in Section 4 can be used to determine the surface reflectivity parameters k, ρ and α , we turn to the final stage, which is to compute the surface orientation along the matched zero-crossing contours of the stereo algorithm.

The algorithm described below is essentially a modification of the photometric stereo algorithm [Woodham, 1978, 1980, 1981; Ikeuchi and Horn, 1979; Silver, 1980]. It can be sketched as follows.

- (1) Compute the reflectance map [Horn, 1977, 1981; Horn and Sjoberg, 1979] associated with the computed values of ρ , α and k . Since R is a function of the unit surface normal, it can be calculated over the unit disk, with coordinate axis given by the first and second components of the unit normal.
- (2) Use the measured value of \mathcal{E}_L at a point to determine a contour of isobrightness in the unit disk. This defines the set of possible surface orientations, which are consistent with the value of \mathcal{E}_L .
- (3) Use the measured value of \mathcal{E}_R at the corresponding point to determine a second contour of isobrightness in the unit disk. The intersections of the two contours define possible surface orientations. Note that in general there will be at least two such points of intersection.
- (4) We can disambiguate the possible orientations by applying a third constraint. Take a directional derivative of depth along the contour through the point. This will yield a linear constraint between p and q . When translated to the unit disk (and a constraint between the components of the unit surface normal), this constraint becomes a half ellipse in the unit disk. The intersection of this curve with the two isobrightness contours defines the correct surface orientation for the point in question.

The key question still to be considered is whether this algorithm results in a unique solution. We first observe that the isobrightness contours for either a matte or a specular surface consist of a series of nested ellipses or half-ellipses. In particular, for a given value $R(\mathbf{n}) = c$, the ellipse has center in the unit disk of (cp_s, cq_s) where p_s and q_s are the first and second components of the unit source vector \mathbf{s} . The minor axis of the ellipse is oriented along the line from the origin to the projection of \mathbf{s} and the major axis is perpendicular to it.

While it is possible to derive this analytically, it more easily observed by noting that for a Lambertian coated sphere, with source located along the same direction as the viewer, the isobrightness contours are nested circles. For an off-axis light source, the corresponding isobrightness contours can be obtained by rotating the sphere relative to the viewer, and then projecting the previous isobrightness contours onto the unit disk, clearly resulting in a series of nested ellipses. A similar argument holds for the isobrightness contours of a purely specular surface.

When considering a convex combination of two such reflectance maps, the isobrightness contours become somewhat more complex. In general, however, the contours are slightly distorted ellipses (see Figure 3). As a consequence, the intersection of isobrightness contours from the two eyes will in general result in a pair of points in the unit disk. (This excludes the degenerate case of a pure matte surface, in which case the intersection will be a complete ellipse.) We must now determine the circumstances under which the third constraint will disambiguate this pair of points.

Suppose that the zero-crossing contour at the point of interest has a tangent in the image plane whose angle with respect to the x axis is γ , and whose derivative of depth along the contour




Figure 3. Examples of isobrightness contours on the unit disk, for different reflectance maps. In all four cases, $p_s = 0.7, q_s = 0.3$. For cases (a), (b), and (c), the exponent $k = 10$, and the parameter α is 0.1, 0.5 and 0.9 respectively. In case (d), the exponent is $k = 5$ and $\alpha = 0.5$.




Figure 4. Constraints on the unit surface normal, as a function of the component of the surface gradient along a contour. Each contour delineates the possible values of the unit surface normal, for a different value of the surface gradient along the zero-crossing contour. In each of the contours illustrated here, the parameter γ was negative. The straight line corresponds to the case in which the depth derivative along the zero-crossing contour was zero.

is given by the value d . This derivative determines the component of the surface normal along the contour. When translated into a constraint on the components of the unit surface normal, the set of possible unit normals which are consistent with such a component along the contour are given by the parametric equations (assuming γ is non-zero)

$$\begin{aligned} n_1(t) &= \frac{t \sin \gamma}{(\sin^2 \gamma + d^2 - 2dt \cos \gamma + t^2)^{\frac{1}{2}}} \\ n_2(t) &= \frac{d - t \cos \gamma}{(\sin^2 \gamma + d^2 - 2dt \cos \gamma + t^2)^{\frac{1}{2}}} \end{aligned}$$

These parametric equations define a half ellipse in the unit disk with center at the origin and with major axis the line through the origin with slope $-\cot \gamma$ of length 2. If γ is positive, the half ellipse is that which connects the two endpoints of the major axis and whose half of the minor axis lies in the positive n_2 half disk, if γ is negative, the half ellipse lies with its half of the minor axis in the negative n_2 half disk (see Figure 4).

We wish to determine the conditions in which there is *not* a unique solution, which corresponds to determining the conditions under which the half ellipse constraint passes through both points arising from the intersection of the isobrightness contours. Note that the family of all possible depth constraints can be represented by the pairs (γ, b) where γ ranges from $-\pi$ to π and b is the length of the minor axis of the half ellipse. Consider a point (p_0, q_0) lying in the unit disk. Let

$$\xi = \text{sgn}(q_0) \tan^{-1} \frac{q_0}{p_0}.$$

Then the valid range of possible γ is from $\xi - \frac{\pi}{2}$ to $\xi + \frac{\pi}{2}$, modulo 2π . Thus, the set of possible depth constraints which pass through this point can be specified by the graph given in Figure 5, which has a range of π in γ and a maximum in b of $\sqrt{p_0^2 + q_0^2}$. For a second, distinct point in the unit disk, a similar mapping is obtained, which intersects the first at exactly one point. This implies that there is exactly one ellipse, with orientation and eccentricity given by this intersection, which passes through both points in the unit disk. This in turn implies that the depth constraint will almost always uniquely specify the surface orientation of the zero-crossing point, since only a particular alignment of the light source and the depth contour will result in an ambiguity, and in general, the physical processes which give rise to the depth contour constraint are independent of the light source. In the case of continuous value ranges for γ and d , the set of conditions in which a unique answer is not possible will, in fact, have measure zero. In the more practical case of a discrete range of values for γ and d , assume that the orientation γ can take on one of n different values, while the component of the depth gradient along the contour d can take on one of m different values. Then assuming that the possible values for γ and d are uniformly distributed, and assuming that the constraints due to the reflectance map isobrightness contours are independent of the constraint due to the depth derivative implies that the probability of an ambiguous computation of surface orientation is $\frac{1}{mn}$. Clearly, if we have a reasonable resolution for γ and d , this probability will be very small.

6. Discussion

We note that it is possible to solve for the reflectivity parameters without requiring two corners. Suppose instead that we have one corner, and that the surface is at least twice continuously differentiable. Then by noting that integration of surface orientation around a closed loop must yield zero (note that a similar constraint has been used by Ikeuchi and Horn [1981]) we can determine the surface orientation at a second point, and the analysis of Section 4 still holds.

We also note that the numerical properties of the algorithm sketched here probably preclude its use by the human visual system. In particular, to have reasonable numerical stability over a large range of surface orientations, a large vergence angle is required (on the order of 3°). For the human system, this is roughly equivalent to viewing objects from a distance of 1 meter or less. In cases of machine vision, where control of the positioning of the sensors is possible (i.e. we are not restricted to interocular separations of 6 cm) large vergence angles are more feasible.

7. Error Analysis

As well as performing simulations of the computation of surface parameters, we can also make an estimate of the error associated with computing the parameters. This analysis will also indicate a means for determining those situations in which the error is small.

We again consider the situation in which the parameter k is known exactly, as well as the values of p and q and hence of m , u and v . We let




Figure 5. Determining the uniqueness of the solution. Parts (a) and (b) represent constraints between the parameters b and γ , for different values of ξ . When a second point is added (c), the intersection of the two contours defines the only possible half ellipse in the unit disk which passes through both points in the unit disk and has appropriate major and minor axes. Thus, given two isobrightness contours in the unit disk, from the two views of a point on a zero-crossing contour, there is exactly one depth constraint which will not disambiguate these points.

α_m = measured value for α
 ρ_m = measured value for ρ
 α_0 = actual value for α
 ρ_0 = actual value for ρ
 δ_L = error in \mathcal{E}_L
 δ_R = error in \mathcal{E}_R .

The goal is to determine estimates for the ratio of the measured values to the actual values, that is, to derive expressions for

$$\epsilon_\alpha = \frac{\alpha_m}{\alpha_0}$$

$$\epsilon_\rho = \frac{\rho_m}{\rho_0}.$$

By expansion, we obtain the following expression

$$\rho_m = \left(\frac{\mathcal{E}_L - \mathcal{E}_R + \delta_L - \delta_R}{u^k - v^k} \right) \frac{1}{s_m}$$

$$= \left(1 + \frac{\delta_L - \delta_R}{\mathcal{E}_L - \mathcal{E}_R} \right) \frac{\rho_0}{\epsilon_s}.$$

Thus,

$$\epsilon_\rho = \left(1 + \frac{\delta_L - \delta_R}{\mathcal{E}_L - \mathcal{E}_R} \right) \frac{1}{\epsilon_\alpha}.$$

Similarly,

$$\alpha_m = \frac{m(\mathcal{E}_L - \mathcal{E}_R + \delta_L - \delta_R)}{(u^k - m)(\mathcal{E}_R + \delta_R) - (v^k - m)(\mathcal{E}_L + \delta_L)}$$

$$= \alpha_0 \left(1 + \frac{\delta_L - \delta_R}{\mathcal{E}_L - \mathcal{E}_R} \right) \left(1 - \frac{(u^k - m)\delta_R - (v^k - m)\delta_L}{(u^k - m)(\mathcal{E}_R + \delta_R) - (v^k - m)(\mathcal{E}_L + \delta_L)} \right)$$

so that

$$\epsilon_\alpha = \left(1 + \frac{\delta_L - \delta_R}{\mathcal{E}_L - \mathcal{E}_R} \right) \left(1 - \frac{(u^k - m)\delta_R - (v^k - m)\delta_L}{(u^k - m)(\mathcal{E}_R + \delta_R) - (v^k - m)(\mathcal{E}_L + \delta_L)} \right).$$

The goal is to determine what circumstances will result in small errors in the computed parameters (i.e. under what conditions will $\epsilon_\rho \approx 1$ and $\epsilon_\alpha \approx 1$).

Algebraic substitution yields the following alternative forms:

$$\epsilon_\rho = 1 + \frac{\delta_L - \delta_R}{\mathcal{E}_L - \mathcal{E}_R} + \frac{1}{\rho} \frac{\delta_R \mathcal{E}_L - \delta_L \mathcal{E}_R}{m(\mathcal{E}_L - \mathcal{E}_R)}$$

$$\epsilon_\alpha = \left(1 + \frac{\delta_L - \delta_R}{\mathcal{E}_L - \mathcal{E}_R} \right) \left(1 - \frac{\delta_L - \delta_R + \frac{1}{\rho m}(\mathcal{E}_L \delta_R - \mathcal{E}_R \delta_L)}{\mathcal{E}_L - \mathcal{E}_R + \delta_L - \delta_R + \frac{1}{\rho m}(\mathcal{E}_L \delta_R - \mathcal{E}_R \delta_L)} \right).$$

We can see that the only circumstances in which the error will be large are when $|\mathcal{E}_L - \mathcal{E}_R|$ is not significantly greater than $|\delta_L - \delta_R|$. Since in general, $|\delta_L - \delta_R|$ will be no more than 1, by restricting our attention to portions of the image in which $|\mathcal{E}_L - \mathcal{E}_R|$ is larger than some threshold, we can reduce the probable error in our estimation of the surface reflectivity parameters.

8. Summary

We have demonstrated that shading information can be combined with stereo data along feature point contours in order to determine the surface orientation along those contours, under the following conditions.

- (1) The surface is illuminated by a distant point source, whose direction relative to the viewer is known.
- (2) The surface is viewed from two different positions, and the angle of vergence of the two views is known. For numerical stability, this angle should be on the order of 3 degrees or larger.
- (3) The surface is at least twice continuously differentiable.
- (4) The reflectance map of the surface can be represented in the form

$$R(\mathbf{n}) = \left[(1 - \alpha)(\mathbf{n} \cdot \mathbf{s}) + \alpha \{ \mathbf{n} \cdot \mathbf{h} \}^k \right]$$

where \mathbf{s} is a unit vector in the direction of the light source, \mathbf{h} is a unit vector in the direction of the off specular angle,

$$\mathbf{h} = \frac{(\mathbf{s} + \mathbf{v})}{\sqrt{2}\sqrt{1 + (\mathbf{s} \cdot \mathbf{v})}}$$

and \mathbf{n} is a unit surface normal. The variables α , ρ and k are unknown, and are to be computed from the irradiance data. It is assumed that α is non-zero, so that the surface is not perfectly matte.

Under these conditions, we have shown that using the brightness values recorded at corresponding points in the two images (determined by the stereo correspondences) we can compute the reflectivity parameters for the surface, and using in addition the component of the surface gradient along the stereo contour, we can apply an algorithm similar to photometric stereo to compute the surface orientation along the stereo contours.

The numerical stability of the computation was also explored and methods for reducing the possible error in the computation were suggested. The key observation concerning the computation was that while accurate solutions were possible given high resolution input (e.g. the image brightnesses were known to extremely high accuracy), when the discrete nature of the brightness sensitivity and spatial sensitivity of the imaging devices was taken into account, the numerical stability of the computations degraded rapidly. This suggests that in its current form, computations such as those outlined here may be of limited practical use, until methods for improving the numerical stability are derived.

9. Acknowledgements

The author wishes to express his gratitude to Tomaso Poggio, Tomàs Lozano-Pérez, Berthold Horn and Ellen Hildreth for many useful comments and discussions.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-80-C-0505 and in part by National Science Foundation Grant 79-23110MCS.

10. Appendix

The solution for the reflectivity parameters, using derivatives of the image irradiances is summarized as follows. The solution to k is given by

$$k = 1 + \frac{\ln \{(a_2 + m_p a_3)u_q - (a_1 + m_q a_3)u_p\} - \ln \{(a_2 - m_p a_3)v_q - (a_1 - m_q a_3)v_p\}}{\ln v - \ln u}$$

the albedo ρ is given by

$$2\rho = \left(\frac{\Phi}{\Theta}\right) \frac{\Theta(\Sigma_x + \Sigma_y) - [\mu C + \omega D]\Phi - [(1 - \nu)C + (1 - \xi)D]\Psi}{(\mu m_p + \omega m_q)\Phi + [(1 - \nu)m_p + (1 - \xi)m_q]\Psi}$$

and the specular scalar α is given by

$$\alpha = \frac{1}{\rho} \frac{\Phi}{\Theta}$$

where

$$\begin{aligned} a_1 &= \mu m_p (\Delta_x \Sigma_y - \Delta_y \Sigma_x) + m_q (\mu \Delta_y \Sigma_y - \omega \Delta_x \Sigma_x) + \Delta_y \Sigma_x m_q (\mu \xi - \omega \nu) \\ a_2 &= \omega m_q (\Delta_x \Sigma_y - \Delta_y \Sigma_x) + m_p (\mu \Delta_y \Sigma_y - \omega \Delta_x \Sigma_x) + \Delta_x \Sigma_y m_p (\mu \xi - \omega \nu) \\ a_3 &= \omega \Delta_x^2 + (\omega \nu - \mu \xi) \Delta_x \Delta_y - \mu \Delta_y^2 \end{aligned}$$

where

$$\begin{aligned} \Phi &= (\Delta_x + \nu \Delta_y)A - (\Delta_y + \xi \Delta_x)B \\ \Psi &= \mu A \Delta_y - \omega B \Delta_x \\ \Theta &= \mu A^2 + (\omega \nu - \mu \xi)AB - \omega B^2 \end{aligned}$$

and where

$$\begin{aligned} \Delta_x &= \frac{\partial \mathcal{E}_L}{\partial x} - \frac{\partial \mathcal{E}_R}{\partial x} \\ \Delta_y &= \frac{\partial \mathcal{E}_L}{\partial y} - \frac{\partial \mathcal{E}_R}{\partial y} \\ \Sigma_x &= \frac{\partial \mathcal{E}_L}{\partial x} + \frac{\partial \mathcal{E}_R}{\partial x} \\ \Sigma_y &= \frac{\partial \mathcal{E}_L}{\partial y} + \frac{\partial \mathcal{E}_R}{\partial y} \\ A &= k(u_p u^{k-1} - v_p v^{k-1}) \\ B &= k(u_q u^{k-1} - v_q v^{k-1}) \\ C &= k(u_p u^{k-1} + v_p v^{k-1}) - 2m_p \\ D &= k(u_q u^{k-1} + v_q v^{k-1}) - 2m_q \end{aligned}$$

and where the second directional derivatives along the two directions of the zero-crossing contour, whose values are d_1 and d_2 with associated directions γ_1 and γ_2 , determine the constants

$$\begin{aligned}\mu &= \frac{d_1 \sin^2 \gamma_2 - d_2 \sin^2 \gamma_1}{\sin(\gamma_1 + \gamma_2) \sin(\gamma_2 - \gamma_1)} \\ \omega &= \frac{d_s \cos^2 \gamma_1 - d_1 \cos^2 \gamma_2}{\sin(\gamma_1 + \gamma_2) \sin(\gamma_2 - \gamma_1)} \\ \nu &= \frac{2 \sin \gamma_1 \sin \gamma_2}{\sin(\gamma_1 + \gamma_2)} \\ \xi &= \frac{2 \cos \gamma_1 \cos \gamma_2}{\sin(\gamma_1 + \gamma_2)}.\end{aligned}$$

Finally, the surface orientation at this point is given by the three parameters

$$\begin{aligned}f_{xy} &= \frac{\Delta_y \mu A - \Delta_x \omega B}{(\Delta_x + \nu \Delta_y)A - (\Delta_y + \xi \Delta_x)B} \\ f_{xx} &= \mu - \nu f_{xy} \\ f_{yy} &= \omega - \xi f_{xy}.\end{aligned}$$

10.1. Discussion of the Equations

Several comments on the form and derivation of the equations are appropriate. First, we note that in deriving the above expression for α , the form obtained is valid only in the case of both $\Delta_x + \Delta_y$ and $\Delta_x - \Delta_y$ non-zero. This provides an interesting side effect, since these expressions both vanish if and only if the surface material has only a matte component, and no (measurable) specular component, or the surface is a plane. If the surface material has no specular component, $\alpha = 0$ and this is exactly the expression obtained by substitution of $\Delta_x = 0$ and $\Delta_y = 0$ into the expression obtained for α .

Second, we note that, subject to possible numerical errors, the solution for the surface reflectance parameters is unique.

Third, we note that it is important to consider the measurability of the image knowns. In other words, can we extract the required measurements for the images in a reliable manner? We first note that the zero-crossing contours, along which the stereo depth information is known, will frequently outline discontinuities in the surface or in the surface material, for example, occluding edges of objects or changes in surface albedo. Since we cannot measure infinitesimal derivatives in the image, does this preclude our ability to measure derivatives of the images? The answer is no, since we can again rely on directional derivatives. For example, we need to measure $\frac{\partial}{\partial x}(\mathcal{E}_L - \mathcal{E}_R)$ and $\frac{\partial}{\partial y}(\mathcal{E}_L - \mathcal{E}_R)$. This can be done without crossing a discontinuity in the image irradiances by considering the two directional derivatives along the contour obtained at a corner in the zero-crossing contour. At this point, we have two different directional derivatives, and simple algebra allows us to then determine the parametric partial derivatives.

11. References

- Attneave, F. "Informational aspects of visual perception," *Psychol. Rev.* 61 (1954), 183-193.
- Baker, H.H. and Binford, T.O. "Depth from edge and intensity based stereo," *Seventh International Joint Conference on Artificial Intelligence*, 1981, 631-636.
- Barlow, H.B. "Three points about lateral inhibition," *Sensory Communication*, W. A. Rosenblith, (ed) MIT Press (1961), 217-234; 782-786.

- Barrow, H.G. and Tenenbaum, J.M. "Recovering intrinsic scene characteristics from images," in *Computer Vision Systems*, E. Riseman and A. Hanson (eds), Academic Press (1979).
- Blinn, J.F. "Models of light reflection for computer synthesized pictures," *Computer Graphics (SIGGRAPH)* 11, 2 (1977), 192-198.
- Brooks, M.J. Shape from shading discretely, Ph. D. Thesis, University of Essex, 1982.
- Bruss, A.R. The image irradiance equation: its solution and application, MIT Artificial Intelligence Laboratory Tech. Rep. 623, 1981.
- Grimson, W.E.L. "A computer implementation of a theory of human stereo vision," *Phil. Trans. Roy. Soc. Lond.B* 292 (1981), 217-253.
- Grimson, W.E.L. "A computational theory of visual surface interpolation," *Phil. Trans. Roy. Soc. Lond.B* 298 (1982a), 395-427.
- Grimson, W.E.L. "An implementation of a computational theory of visual surface interpolation," *Computer Graphics and Image Processing* (1982b), to appear.
- Herskovitz, A. and Binford, T.O. On boundary detection, MIT Artificial Intelligence Laboratory Memo 183, 1970.
- Horn, B.K.P. Shape from shading: a method for obtaining the shape of a smooth opaque object from one view, MIT Project MAC Technical Report, MAC TR-79, 1970.
- Horn, B.K.P. The Binford-Horn Line Finder, MIT Artificial Intelligence Memo 285, 1973.
- Horn, B.K.P. "Obtaining shape from shading information," in *The Psychology of Computer Vision*, P.H. Winston (ed) McGraw-Hill, (1975) 115-155.
- Horn, B.K.P. "Understanding image intensities," *Artificial Intelligence* 8, 11 (1977), 201-231.
- Horn, B.K.P. Sequins and quills - representations for surface topography, MIT Artificial Intelligence Memo 536, 1979.
- Horn, B.K.P. "Hill shading and the reflectance map," *Proceedings of the IEEE* 69, 1 (1981), 14-47.
- Horn, B.K.P. and Sjoberg, R.W. "Calculating the reflectance map," *Appl. Opt.* 18, 11 (1979), 1770-1779.
- Ikeuchi, K. and Horn, B.K.P. "An application of photometric stereo," *Sixth International Joint Conference on Artificial Intelligence*, 1979, 413-415.
- Ikeuchi, K. and Horn, B.K.P. "Numerical shape from shading and occluding boundaries," *Artificial Intelligence* 17, (1981), 141-184.
- Kidd, A.L., Frisby, J.P., and Mayhew, J.E.W. "Texture contours can facilitate stereopsis by initiating appropriate vergence eye movements," *Nature* 280 (1979), 829-832.
- Marr, D. "Early processing of visual information," *Phil. Trans. R. Soc. Lond.* 275, 942 (1976), 483-534.
- Marr, D. "Representing visual information," *AAS 143rd Annual Meeting, Symposium on Some Mathematical Questions in Biology, February, 1977*. Published in *Lectures in the Life Sciences* 10 (1978), 101-180.
- Marr, D. *Vision: A computational investigation in the human representation and processing of visual information*, W.H. Freeman, San Francisco, 1982.
- Marr, D. and Hildreth, E.C. "Theory of edge detection," *Proc. R. Soc. Lond.B* 207 (1980), 187-217.
- Marr, D. and Poggio, T. "From understanding computation to understanding neural circuitry," *Neurosciences Research Program Bulletin* 15 (1977), 470-488.
- Marr, D. and Poggio, T. "A theory of human stereo vision," *Proc. Roy. Soc. Lond.B* 204 (1979), 301-328.
- Mayhew, J.E.W. and Frisby, J.P. "Psychophysical and computational studies towards a theory of human stereopsis," *Artificial Intelligence* 17 (1981), 349-385.

- Phong, B.T. "Illumination for computer generated pictures," *Communications of the ACM* **18**, 6 (1975), 311-317.
- Pratt, W. *Digital Image Processing*, John Wiley and Sons, New York, 1978.
- Roberts, L.G. "Machine perception of three-dimensional solids," in *Optical and Electro-Optical Information Processing*, Tippett, et al. (eds), MIT Press, Cambridge (1965), 159-197.
- Rosenfeld, A. and Kak, A. *Digital Picture Processing*, Academic Press, New York, 1976.
- Silver, W. Determining shape and reflectance using multiple images, M.S. Thesis, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, 1980.
- Terzopoulos, D. Multi-level reconstruction of visual surfaces: variational principles and finite element representations, MIT Artificial Intelligence Memo 671, 1982.
- Woodham, R.J. "Photometric stereo: A reflectance map technique for determining surface orientation from image intensity," *Image Understanding Systems and Industrial Applications, Proc. S.P.I.E.* **155** (1978).
- Woodham, R.J. "Photometric method for determining surface orientation from multiple images," *Optical Engineering* **19**, 1 (1980), 139-144.
- Woodham, R.J. "Analysing images of curved surfaces," *Artificial Intelligence* **17** (1981), 117-140.