[4] R. E. Woods and R. C. Gonzalez, "Real time digital image enhancement," *Proc. IEEE*, to be published.

Patrenahalli M. Narendra (S'73-M'76) was born in Jog Falls, India, on March 16, 1951. He received the Bachelor of Engineering degree in electronics from Bangalore University in 1971 and the M.S. and Ph.D. degrees in electrical engineering from Purdue University, West Lafayette, IN, in 1973 and 1975, respectively.

He has been with Honeywell Systems and Research Center, Minneapolis, MN, since 1976, where he is presently Staff Research Scientist in the Signal and Image Processing Section. His current research interests span dynamic scene analysis, symbolic image processing, and VLSI device architectures for real-time implementation of image processing systems. He is the author of several publications in his areas of interest. In 1979, he won the H. W. Sweatt Award, the highest technical achievement award in Honeywell.

Dr. Narendra is a member of Phi Kappa Phi and the IEEE Computer, PAMI, and ASSP Societies.

Robert C. Fitch (S'73-M'77) was born on June 23, 1953. He received the B.S. degree in electrical engineering from Montana State University, Bozeman, in 1975 and the M.S. degree in electrical engineering from Colorado State University, Fort Collins, in 1977.

He is currently a Senior Research Scientist in the Signal and Image Processing Section at Honeywell Systems and Research Center, Minneapolis, MN, where he has worked since 1977. His [illegible] are in VLSI implementations for real-time image processing, automatic target recognition, and robotics.

Mr. Fitch is a member of Tau Beta Pi and Phi Kappa Phi.

Determining Surface Orientations of Specular Surfaces by Using the Photometric Stereo Method

KATSUSHI IKEUCHI, MEMBER, IEEE

Abstract—**The orientation of patches on the surface of an object can be determined from multiple images taken with different illumination, but from the same viewing position. The method, referred to as photometric stereo, can be implemented using table lookup based on numerical inversion of reflectance maps. Here we concentrate on objects with specularly reflecting surfaces, since these are of importance in industrial applications. Previous methods, intended for diffusely reflecting surfaces, employed point source illumination, which is quite unsuitable in this case. Instead, we use a distributed light source obtained by uneven illumination of a diffusely reflecting planar surface. Experimental results are shown to verify analytic expressions obtained for a method employing three light source distributions.**

Index Terms—**Bin of bolts and nuts, distributed light source, glossy object, reflectance map, shape from shading, surface inspection.**

Manuscript received November 29, 1979; revised February 26, 1981.

The author is with the Computer Vision Section, Electrotechnical Laboratory, Ministry of International Trade and Industry, Ibaraki 305.

I. Introduction

THIS paper addresses the problem of determining local surface orientation of specular materials from the intensity information under different illumination, but from the same viewing position. This method is referred to as photometric stereo and was first formulated by Woodham [2]. Here, we concentrate on objects with a specularly reflecting surface, since these are of importance in industrial applications. Previous methods [2], [11] intended for diffusely reflecting surfaces employed point source illumination, which is quite unsuitable in this case.

Historically, Horn solved the image intensity equations [11] in order to obtain an object shape from shading information. Horn used a method of characteristic strip expansion for solving the image intensity equation which is a nonlinear first-order partial differential equation. Horn then introduced the

equation. This map represents the relations between surface orientation and image intensity in the gradient space. Woodham developed a novel technique called photometric stereo [2] using the reflectance map. This photometric stereo is based on the following fact: if a pair of images of the same object are obtained by varying the direction of incident illumination but from the same viewing direction, we can draw a pair of different reflectance maps, because the reflectance map depends on the direction of the light source. So far, surface orientation is determined locally by the intensity pairs recorded at each image point as an intersection of constant brightness lines on the reflectance maps. This photometric stereo is very rapid and is free from noise compared with Horn's method, since surface orientation is determined locally. Horn and Sjoberg [1] showed how to calculate the reflectance map from NBS's BRDF (B_i-directional reflectance distribution function) analytically.

II. Basic Tools and Relationships

A. The Reflectance Map [1], [3] and Photometric Stereo [2]

The reflectance map represents the relationship between surface orientation and image intensity in gradient space [1]. We can express the geometric dependence of the reflectance characteristics of a surface in terms of the slope components p and q, used axes in gradient space [1],

$$p = \partial z/\partial x, \quad q = \partial z/\partial y \tag{1}$$

where z is the elevation of the surface and x, y are the spatial coordinates. If we take the direction from the surface to the viewer as the direction of the z-axis, then the reflectance properties of a surface patch depend on (p, q), the direction of the surface normal, and (ps, qs), the direction of the source [1]. Each point in the gradient space corresponds to a particular surface orientation based on the direction of the viewer. If we know the reflectance characteristics of an object, we can calculate how bright a surface element with that orientation will appear. It is convenient to use contour lines to connect those points in gradient space which correspond to surface orientations which give rise to the same apparent brightness. It is because of these contour lines that the resulting diagram is referred to as the "reflectance map" [1], [3]. The reflectance map is denoted by $R(p, q)$.

Using the reflectance map, the basic imaging equation is

$$E_i(x, y) = R_i(p, q, x, y) \tag{2}$$

where $E_i(x, y)$ is the brightness (image irradiance) in the image-forming system at the point (x, y) in the image plane. This equation contains two unknown variables p, q and one quantity E_i, which can be measured in the image.

In the above equation, the subscript i is used to denote different illumination conditions. For each value of the subscript, a different image is obtained, and a different reflectance map applies. If two images are taken, two such equations provide constraints on the possible values of p and q. This permits us to solve for the gradient. Because of the nonlinearity of the equations, however, a number of solutions may be found at times. In this case, a third image may be used to disambiguate the remaining possibilities. This is the principle of photometric stereo [2].

Orthographic projection can simplify the calculation considerably. If we can assume that the object is small compared with the distance to the source and the image-forming system, then the viewer direction can be approximated as the axis of the image-forming system and we can treat the system as orthographic. There are two merits to this approximation. One is that we can neglect the effect of position. The right-hand side of (2) depends only on (p, q). Namely, we can apply the same reflectance map on all points in the image. Another benefit is that we can calculate $R(p, q)$ more easily because the approximation fixes the viewing direction. So we can, for example, rotate the source keeping the phase angle constant (the phase angle is the angle between the source and the viewer, measured at the object). This means that we can obtain a new reflectance map just by rotating the old one [3].

B. Relationship Between Source Radiance and Image Radiance

One of the main points of our discussion here is that we consider only specular components of reflectance when we calculate the reflectance map, since many industrial materials are made of metal and have strong specularity and little diffuse reflection. Experiments show that only 1 or 2 percent of the incident light is reflected diffusely from some metallic surfaces, with most of the rest reflected specularly. We cannot treat this kind of material using the usual Lambertian model for reflection of light from a surface. It is also clearly inappropriate to use point sources to illuminate such a surface, since very few surface patches will be oriented correctly to reflect any light and we will only see virtual images of the point sources.

Three relationships exist between a light source and the image plane. The first one is that between source radiance and incident radiance on a sample surface. Next, and most important, is that between incident radiance and emitting radiance from the surface (this is captured in the reflectance map). The last one is that between emitting radiance and image irradiance. (See Fig. 1.)

Since we use an extendent light source, incident radiance in one direction is the same as source radiance in the same direction in an extended light source. Namely,

$$L_i = L_s. \tag{3}$$

For a specular surface and an extended light source [1], [4]

$$L_e(\theta_e, \phi_e) = L_i(\theta_e, \phi_e + \pi). \tag{4}$$

Thus, even though the source distribution may be complicated, only the contribution from a single direction $(\theta_e, \phi_e + \pi)$ need be considered. The relationship between reflected radiance and image irradiance is

$$E_p = \{(\pi/4)(d/f_p)^2 \cos^4\alpha\} L_e \tag{5}$$

where F_p, d, α are effective focal length of lens, the diameter of the entrance aperture, and the off-axis angle, respectively [2].

Finally, using (3)-(5) we see that image irradiance at a particular point is proportional to the source radiance in a

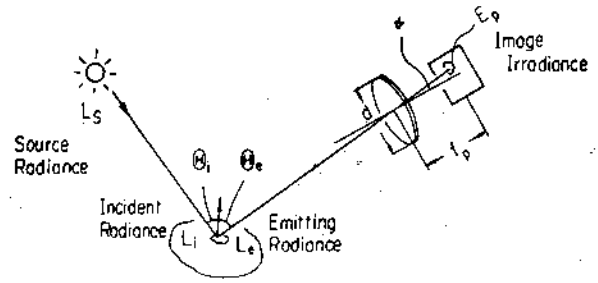

Fig. 1. Relationships between light source, surface, and image formation system are depicted.

direction which depends on the orientation of the corresponding surface patch [1]. That is, the brightness of a particular surface patch is simply equal to the brightness of the part of the extended source which it happens to reflect. Thus, even though the source distribution may be complicated, only the contribution from a single direction (p_s, q_s) need be considered at any time.

It is convenient to change from the local coordinate system to the viewer-centered coordinate system [1] because the reflectance map is defined in a viewer-centered coordinate system, and the light source distribution is given based on the direction of the viewer (see Fig. 2). We can rewrite image irradiance using the viewer-centered coordinate system:

$$E_p(p_n, q_n) = \gamma L_e(p_n, q_n) = \gamma L_s(p_s, q_s) \tag{6}$$

where γ is a constant. Namely, we observe this E_p as image irradiance under an extended light source on specular materials.

We can finally express brightness distribution in the gradient space:

$$R(p_n, q_n) = I_s(p_s, q_s) \tag{7}$$

where

$$\begin{aligned} p_n &= \partial z/\partial x = -\cos\phi_n \tan\theta_n \\ q_n &= \partial z/\partial y = -\sin\phi_n \tan\theta_n \\ p_s &= 2p_n/(1 - p_n^2 - q_n^2) \\ q_s &= 2q_n/(1 - p_n^2 - q_n^2). \end{aligned} \tag{8}$$

On the other hand, the inverse transformation is

$$\begin{aligned} p_n &= p_s(\sqrt{1 + p_s^2 + q_s^2} - 1)/(p_s^2 + q_s^2) \\ q_n &= q_s(\sqrt{1 + p_s^2 + q_s^2} - 1)/(p_s^2 + q_s^2). \end{aligned} \tag{9}$$

We will use a differentiable single-valued function as the reflectance function. Roughly speaking, since we determine surface orientation from brightness, if the function is not single-valued, the inverse function theorem does not apply, and this makes the situation very difficult.

III. Application Problems

A. Total Schema of the System

The technique requires two kinds of tasks; one is the off-line (precomputing) job and the other is the on-line (real-time) job which is rather simple compared with the off-line job. The simplicity implies rapid calculation, as desired in any hand-eye system. The off-line job consists of making the reflectance map and constructing a lookup table. The on-line job consists of reading the image brightness and determining orientations of a surface patch based on the lookup table. Fig. 3 shows the information flow between the on-line job and the off-line job.

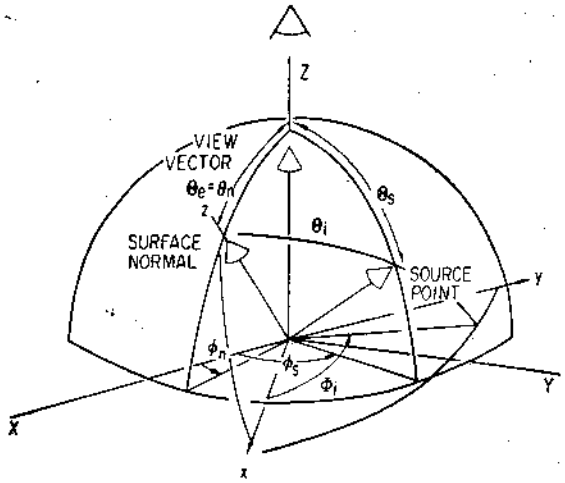

Fig. 2. Relationships between a local coordinate system and a viewer coordinate system.

B. Off-Line (Precomputing) Job

Consideration of Light Source: We use a Lambertian surface as a source plane. It is illuminated by a linear lamp as shown in Fig. 4. Although a spherical shaped source can easily cover directions of more than 90°, it is difficult to build such a device and difficult to control the distribution of the light on it, particularly if interflection is taken into account. On the other hand, if we use a planar source illuminated by a lamp, the brightness distribution is complicated, but we can calculate a reflectance map. It is possible to cover angles of more than 90° by making a box-like source. In that case, however, we have to treat each plane separately because the surface normal is not differentiable at the intersection of two planes. Thus, we consider one plane surface which is assumed to have the Lambertian characteristics in order to obtain some analytic results to help in the design of the system. This plane is illuminated by a line source as a representative case.

Brightness distribution on the surface is calculated using (11). Let f be the flux rate per unit source length [W/(m · sr)]. Then the total irradiance E[W/m^2] is

$$E = \int f \cos\theta_1 \cos\theta_2 / r^2 \, dt. \tag{10}$$

From $\cos\theta_1 = \sqrt{X^2 + l_0^2}/r$ and $\cos\theta_2 = l_0/r$, we can finally get

$$\begin{aligned} E = (f l_0/2a^2)\Bigg[\Bigg\{\tan^{-1}\left(y + \frac{L}{2}\right)\Big/a - \tan^{-1}\left(y - \frac{L}{2}\right)\Big/a\Bigg\} \\ + \Bigg\{a\left(y + \frac{L}{2}\right)\Big/\left(a^2 + \left(y + \frac{L}{2}\right)\Big/\left(a^2 + \left(y - \frac{L}{2}\right)^2\right)\right)\Bigg\}\Bigg] \end{aligned} \tag{11}$$

where $a = \sqrt{x^2 + l_0^2}$, L is the length of the line source, and l_0

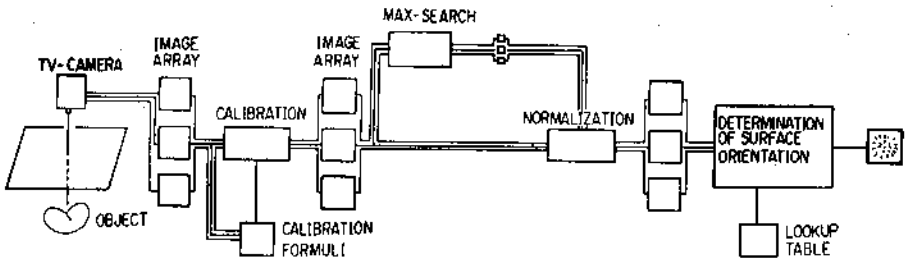

Fig. 3. The overall schema of the experiment. The technique consists of two types of jobs: off-line preparation of tables (represented by broken lines) and on-line processing. Image brightness is obtained by using a TV camera. We search for the maximum value of brightness in each array. Brightness arrays are calibrated and normalized. Surface orientations are obtained from the lookup table.

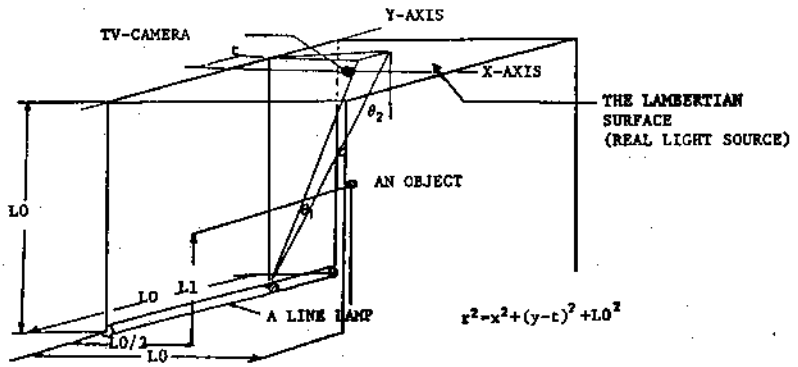

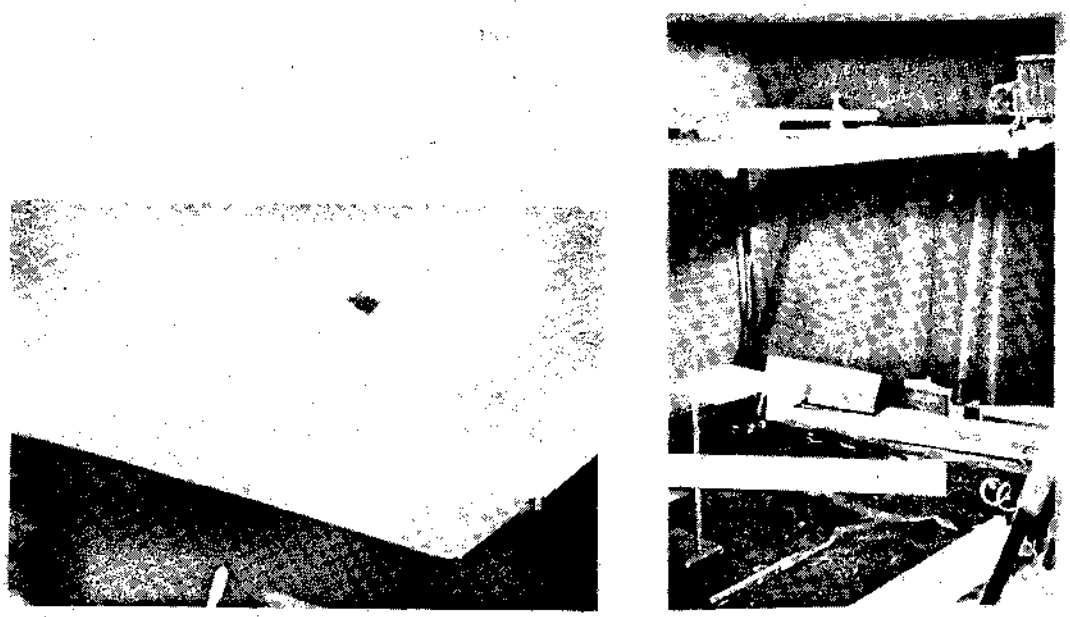

Fig. 4. The object is illuminated solely by the light reflected of an overhead Lambertian surface. This surface receives light from a linear lamp positioned below so as not to illuminate the object. The TV camera peers through a small hole in the overhead surface. Brightness distribution from a linear lamp on the Lambertian surface can be calculated in (12). We used three linear lamps, placed symmetrically, 120° apart. In the center of the lamps there is a hole. The object is observed through the hole.

is the distance between the line source and the Lambertian surface.

We put an object just under the point $(x, y) = (0.5l_0, 0.0)$. $E(x, y)$ is a symmetric function with respect to y. Thus, it is natural to put an object somewhere on the x axis. $E(x, 0)$ is an s-shape function and has an inflection point near $x = 0.5l_0$. At that point $E(0.5l_0, 0.0) = 0.715E(0.0, 0.0)$. So it is convenient to put an object just under this point. If we denote the distance from the surface to the object as l_1,

$$x = -l_1 p_s + l_0/2$$
$$y = -l_1 q_s. \tag{12}$$

Finally, we can get a reflectance map

$$R(p_n, q_n) = [fl_0/2\{(-l_1 p_s + l_0/2)^2 + l_0^2\}]$$
$$\cdot \tan^{-1}\{(-l_1 q_s + l_0/2)/\sqrt{(-l_1 p_s + l_0/2)^2 + l_0^2}\}$$
$$- \tan^{-1}\{(-l_1 q_s - l_0/2)/\sqrt{(-l_1 p_s + l_0/2)^2 + l_0^2}\}$$
$$+ \{\sqrt{(-l_1 p_s + l_0/2)^2 + l_0^2}\,(-l_1 q_s + l_0/2)\}/$$
$$\{(-l_1 p_s + l_0/2)^2 + l_0^2 + (-l_1 q_s + l_0/2)^2\}$$
$$- \{\sqrt{(-l_1 p_s + l_0/2)^2 + l_0^2}\,(-l_1 q_s - l_0/2)\}/$$
$$\{(-l_1 p_s + l_0/2)^2 + l_0^2 + (-l_1 q_s - l_0/2)^2\}] \quad (13)$$

where p_s and q_s are as defined above, and $L = l_0$, $2l_1 = l_0$ because these values are found to be optimal upon simulation.

We used three linear lamps, placed symmetrically 120° apart, about the hole through which the object is observed. One lamp is turned on at a time, giving rise to a reflectance map like the one shown in (13).

By using three reflectance maps, we can determine the surface orientation. Theoretically, it is possible to determine the orientation by using two maps only. However, there may be more than one solution because of the nonlinearity of the equations. If we use three maps, we can determine a unique solution easily as shown in Fig. 5.

The Lookup Table: The most convenient method for converting a triple of measured brightness to an orientation is by means of a lookup table made from the reflectance map. This table, indexed by quantized brightness measurements, contains surface orientations. Each dimension of the table corresponds to brightness measurements of surface patches under one of the three sources. Each entry of the table contains a surface orientation corresponding to a triple.

We, however, construct a two-dimensional lookup table. Although it is possible to make a three-dimensional lookup table, it is better to make a two-dimensional table. An observed triple often contains measurement error. This may cause an observed triple to give no solutions. In other words, it may occur that three contour lines on the reflectance maps do not intersect at a point under noisy situations. On the other hand, two contour lines always intersect giving approximate solutions. It means that a two-dimensional lookup table always gives two alternative solutions. The weakest brightness measurement often contains a relatively large amount of measurement error. We can use the weakest brightness only to choose a solution among two alternatives. Namely, we look up an entry using the two largest brightness values. The entry gives two alternative solutions. Then we use the weakest brightness for deciding the correct solution among the two alternatives. For that purpose, each entry of the lookup table should contain two alternatives, and each alternative should contain a surface orientation and the weakest brightness at that point in the gradient space. Thus, the lookup table can be two-dimensional (see Fig. 6).

We construct the lookup table by using Newton's method. An elegant method exists in the case of a Lambertian reflector and point sources [2]. In our case, however, we have to calculate it numerically. Specifically, we solve two expressions

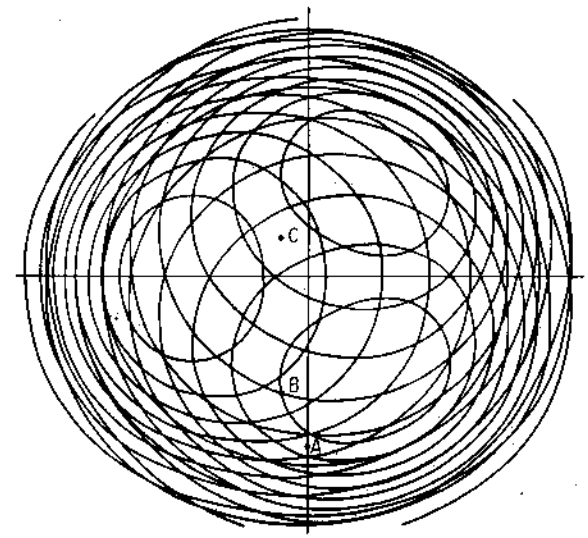

	CELL NUMBER		MEASURED BRIGHTNESS AT EACH CELL SOURCE 1	SOURCE 2	SOURCE 3	
A-----	108	258	0.59948421	0.47733079	0.83733455	
	108	259	0.69183096	0.49679881	0.85233261	
B-----	108	260	0.78222797	0.58294154	0.91505988	P= -0.1 Q= -0.4
	108	261	0.81491274	0.59471041	0.94813397	
	108	262	0.75885488	0.60671593	0.89395881	
C-----	108	268	0.84896324	0.65222758	0.70112918	
	108	269	0.78222797	0.53536265	0.51847315	
	108	270	0.51946413	0.56009784	0.26498424	

Fig. 5. How to use the reflectance map method. Three reflectance maps drawn in the same gradient space. Since the lamps are symmetrically configured, brightness distribution corresponding to a linear lamp is also symmetric having the same shape. Thus, we need to calculate only one distribution and then rotate it 120°. The resulting function is the desired one. We can determine (p, q) from three values of brightness, where each brightness value corresponds to each source condition. For example, switching on source 1 will yield a brightness value of 0.7822 at cell B. Similarly, source 2 and source 3 yield 0.5329 and 0.9150, respectively. From the above diagram, we note that the point (−0.1, −0.4) satisfies this triple. Hence, the surface orientation at the cell B is (−0.1, −0.4).

like (13) for p and q. The expressions differ in the brightness distribution of the source, obtained by turning on one of the three linear lamps.

The N-dimensional Newton method involves solving the simultaneous linear equations

$$f_j + \sum_{k=1}^{n} \partial f_i/\partial x_k|x_k = x_k^{[j]}(x_k^{[j+1]} - x_k^{[j]}) = 0 \quad (14)$$

where f_i is the ith element of a vector function of n dimensions and $x_i^{[j]}$ is the jth element of a solution vector of the jth iteration.

We can rewrite (13) as (17) in our case of a planar source illuminated by a linear lamp. The two mapping functions are

$$f(p_s, q_s) = R(p_s, q_s) - E_n$$
$$g(p_s, q_s) = R(p_s \cos\alpha - q_s \sin\alpha, p_s \sin\alpha + q_s \cos\alpha) - E_m \quad (15)$$

where α is an angle between two light sources and E_n, E_m are image brightness corresponding to the (n, m) element of the lookup table. We can finally get

$$X^{[j+1]} = X^{[j]} - (F')^{-1}F \quad (16)$$

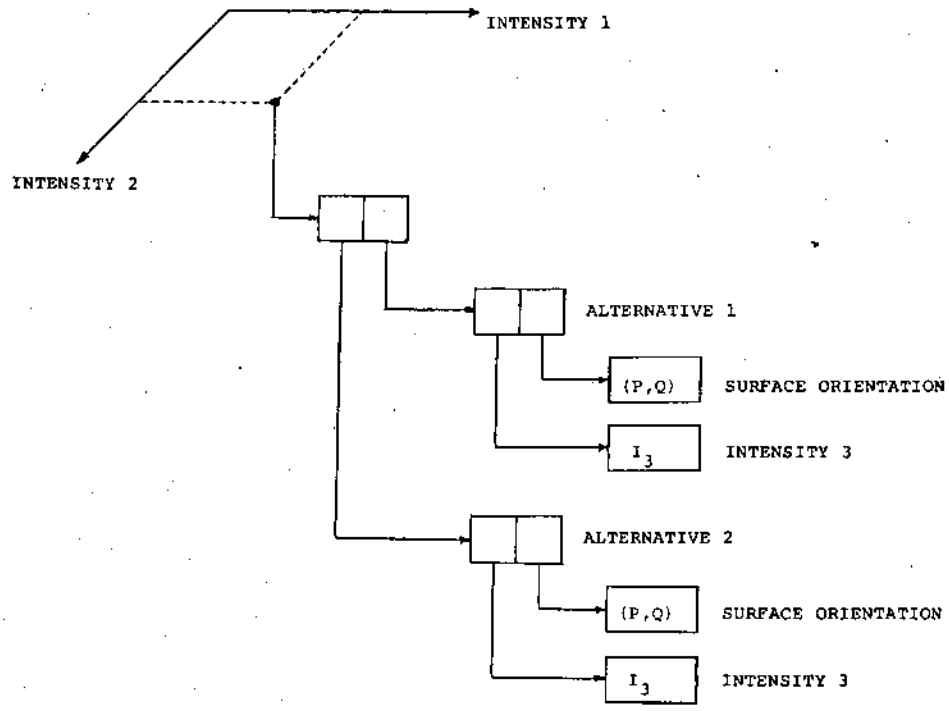

Fig. 6. The structure of the lookup table. Each entry of the lookup table contains two alternatives. Each alternative contains a surface orientation and the weakest brightness at that point in the gradient space. Thus, we can look up an entry using the two largest brightness values. The entry gives two alternative solutions. Then we can use the weakest brightness obtained for deciding the correct solution among the two alternatives on comparing the value with two values contained in the entry.

where $X = {}^{I}(p_s, q_s)$, $F = {}^{I}(f, g)$ and $(F')^{-1}$ is an inverse of the Jacobian matrix. (See Fig. 7.)

C. On-Line (Real-Time) Job

Image brightness is obtained from the TV camera. To reduce the noise typical of these devices, we took more than one picture per light source, and arrays corresponding to the same light source were averaged. The resulting three brightness arrays, one for each light source, were the input to the photometric stereo system.

Before applying the reflectance map (the lookup table) to image arrays obtained, we have to calibrate and normalize the image arrays. The output of the TV camera is not linearly proportional to original brightness. The algorithm has to convert the output into some values linearly proportional to the original brightness values. We call this process calibration. In the meantime, original brightness always contains an effect of albedo. Normalization process cancels the effect of albedo. In other words, the algorithm has to identify a brightness value, say 100 units, in an image array with some value, say 0.5, in the reflectance map. The algorithm normalizes the brightness using the maximum brightness in an image array as a base value.

Brightness calibration is done by using a six-step Kodak gray scale. Since the output of the TV camera is not linearly proportional to original brightness values, the algorithm has to convert the output to certain values proportional to image brightness. We always put a Kodak scale besides the object so that all pictures taken by the TV will contain a gray scale. By using obtained brightness corresponding to the scale, the algorithm makes a formula for linear interpolations. The algorithm can convert the output to certain values proportional to image brightness based on this formula.

The algorithm has to normalize image brightness. Although the scale tells the algorithm exact information about the relationship between real brightness and the output of the TV camera, the algorithm cannot apply the relationship to the object because the albedo ratio depends on materials of objects. In other words, the gray scale has a property of a Lambertian surface, and metals reflect light as a specular surface. Also, the amount of light reflected depends on materials of objects, even though objects reflect light specularly. Thus, the algorithm needs to cancel the effect of albedo.

The maximum brightness in an image array comes from 1.0 in the reflectance map. If an object is convex, you can find any direction among directions of the surface patches. In other words, a convex object always has at least one surface patch corresponding to any point in the gradient space. Thus, an image array always contains a brightness value corresponding to 1.0 in the reflectance map, provided that the object is convex and the visual angle is wide enough.

The algorithm can cancel the effect of the varying albedo using the obtained maximum value:

$$\rho = I(x, y)/R(p, q) \tag{17}$$

where ρ is the albedo. At the brightest part, $R = 1.0$. The

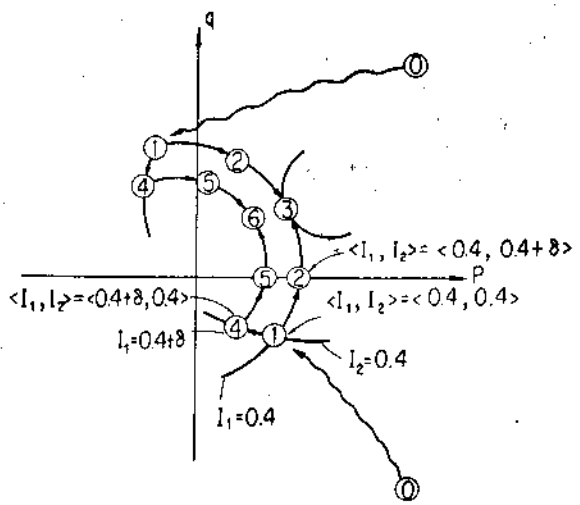

Fig. 7. How to get entire solutions using Newton's method.

algorithm has I_{max} (the maximum value in an image array) and R (1.0 in the reflectance map) at the brightest part. Thus, ρ is obtained as $I_{max}/1.0$. Entire values in the image array will be normalized using this albedo ratio ρ. Finally, each point on the image has a triple of quantized brightness applicable to the reflectance map (the lookup table). It is interesting that the maximum point always contains important information and that this search process resembles Land's lightness search model [7].

From the lookup table we can get the surface orientation. The lookup table entry is accessed using the two largest brightness values. To increase the accuracy of computation and economy of memory use, the first dimension represents the largest image brightness. The second dimension represents the second largest. Each element of the matrix contains the corresponding surface orientation and the smallest image brightness.

Since the nonlinearity gives rise to two solutions for each intensity pair, we have to decide which one is the real one. We choose a surface orientation by comparing the distance between the actual third image brightness and the element of the matrix. The third brightness is always weak and likely to contain errors. It is only used to decide which of the two alternatives obtained from the other brightness is the correct one.

IV. Experiment and Discussion

Our experimental results are shown in Fig. 8. Fig. 8(a) shows three brightness arrays. These arrays are input information to the photometric system. No surface normals are shown in areas where insufficient information was available in the three images to determine them accurately [see Fig. 8(b)]. The choice of light source distribution affects the extent of these regions as well as the accuracy with which surface normals can be found.

We tried two kinds of relaxation methods. The first one is suitable to on-line systems. This method reads from the lookup table twice, exchanging the second and third brightness values, and determines the solution by averaging the two previous solutions. The simple schema can give good results as shown in Fig. 8(c).

The other method is to find $(p_{i,j}, q_{i,j})$ at each mesh point which minimizes

$$e_{ij} = (p_{i,j} - p_{i,j})^2 + (q_{i,j} - q_{i,j})^2 + \lambda[\{I_1 - R_1(p_{i,j}, q_{i,j})\}^2 + \{I_2 - R_2(p_{i,j}, q_{i,j})\}^2 + \{I_3 - R_3(p_{i,j}, q_{i,j})\}^2] \quad (18)$$

where

$$P^*_{i,j} = \tfrac{1}{4}(p_{i,j-1} + p_{i,j+1} + p_{i-1,j} + p_{i+1,j})$$
$$q^*_{i,j} = \tfrac{1}{4}(q_{i,j-1} + q_{1,j+1} + q_{i-1,j} + q_{i+1,j})$$

and I_1, I_2, I_3 are measured image irradiance under the light source 1, 2, 3, respectively, and R_1, R_2, R_3 are their reflectance maps. The former part of (18) is called the surface smoothness constraints [8], and the latter is called the image irradiance constraints. (See [8] for more detail.) Namely, the solution is pulled towards each constraint line of the reflectance map and towards the local constraints from the surface smoothness. It ends up in some compromise position which minimizes overall "strain" [8]. This method is implemented iteratively. Fig. 8(d) shows the result obtained using this iterative method. The algorithm not only smoothed the output, but also extended the area in which a solution could be found. This method is appropriate for off-line systems. Fig. 8(e) is a generated surface from their surface orientations.

A direct application of our technique is an industrial hand-eye system that picks up an object out of a jumble of material [see Fig. 9(a)]. Although our technique cannot correctly determine the surface orientation when there is mutual illumination, it does provide the means for detecting this condition, since the three measurements will be inconsistent. In this fashion erroneous results are avoided. This is important, since the manipulator might otherwise be sent to a position where it would collide with other parts.

Another application of this technique is the inspection of the surface condition of metals. If a surface has a crack, stain, or finger print, the image brightness triple yields inconsistent values in the area of the blemish (see Fig. 10).

This technique may be combined with the Marr-Poggio-Grimson stereo technique now being developed [9], [10]. Their technique detects depth cues and works well when the image contains many discontinuities. On the other hand, our method detects surface orientation directly and works well when the object is smooth. The combination can be established if the output of one of the two stereo cameras is fed to our system while the two outputs from two cameras are fed to their system. We feel that the composite system will produce an excellent representation of the object, just as people are believed to use both stereo and shading information to construct a symbolic image of the visual world.

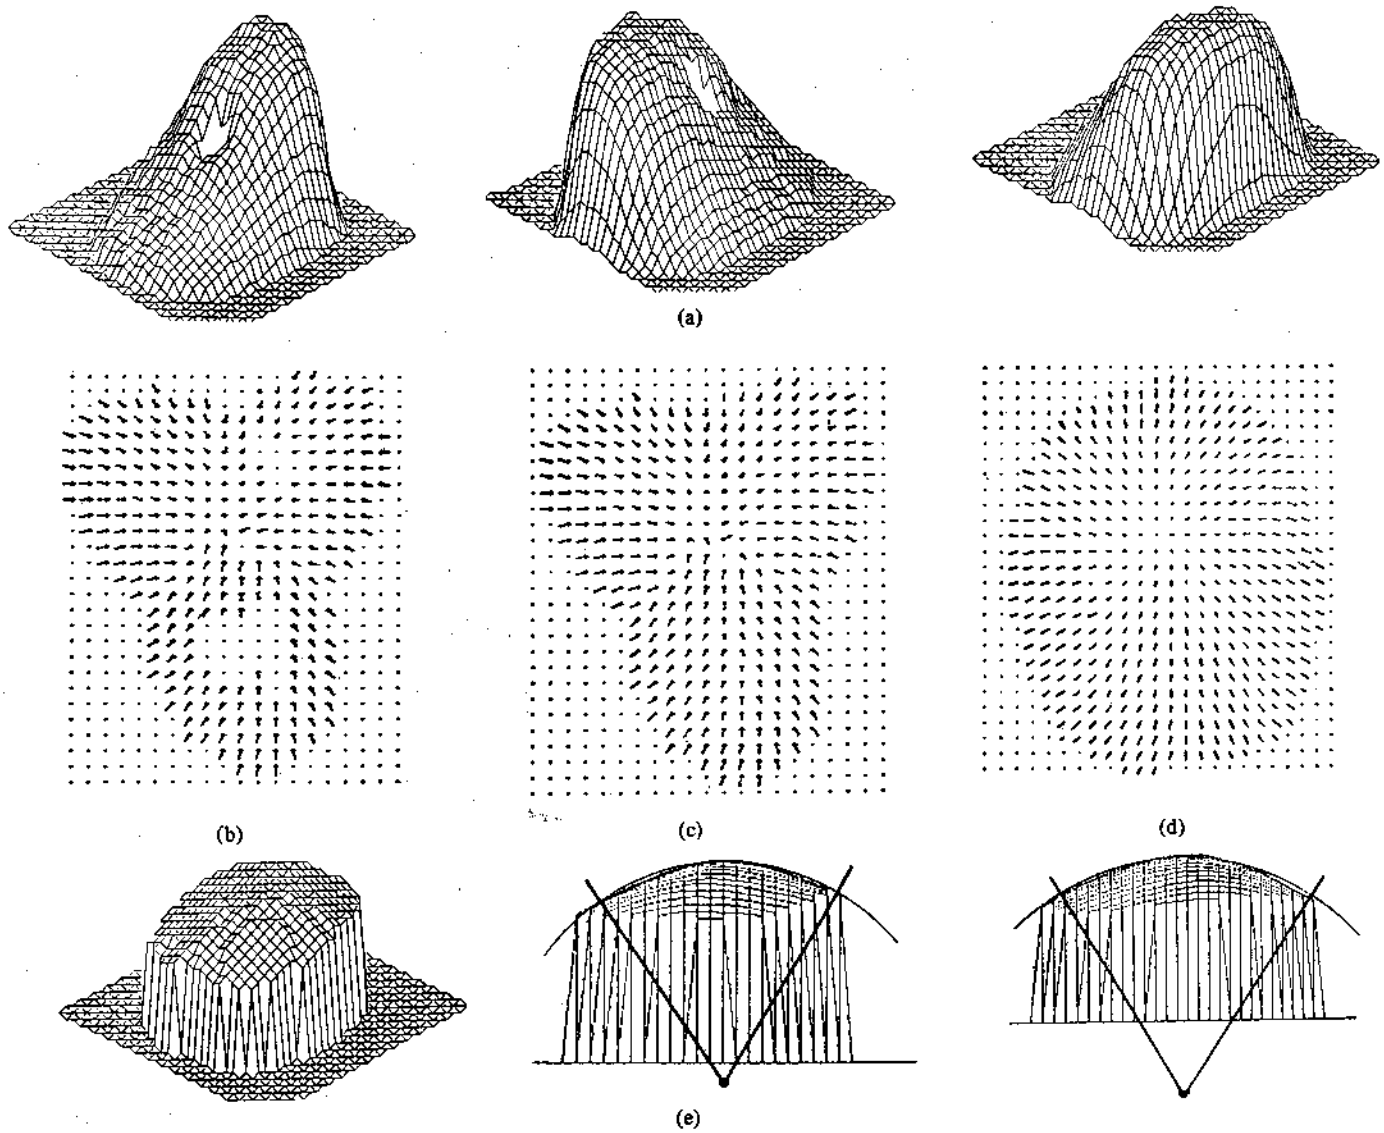

Fig. 8. (a) Three brightness arrays. Each array corresponds to one of the three lamps. These arrays are input information to the photometric stereo system. (b) Direct output from the system. No surface normals are shown in areas where there was insufficient information. (c) Output from the first relaxation method. This method would be suitable for a hand-eye system when real-time response is vital. The system reads from the lookup table twice, exchanging the second and third brightness values. The solution is determined by averaging the two previous values. This simple schema can give good results. (d) Output from the iterative relaxation method. This method finds solutions which minimize the difference of the actual brightness and the theoretical brightness using a constraint derived from surface continuity. (e) i) Generated surfaces from the needle diagrams. ii) A surface along the x axis. The bold line represents the actual surface. The diagram is obtained from the photometric stereo system. iii) A surface along the y axis.

ACKNOWLEDGMENT

The author would like to extend his sincere appreciation to Prof. B. K. P. Horn and Prof. P. H. Winston of M.I.T. Discussions with Dr. Y. Shirai of ETL, H. Murota of the University of Tokyo, and R. Sjoberg of M.I.T. were helpful. Thanks go to B. Roberts of M.I.T. and M. Brooks of the University of Essex for proofreadings.

REFERENCES

[1] B. K. P. Horn and R. W. Sjoberg, "Calculating the reflectance map," *Appl. Opt.*, vol. 18, no. 11, 1979.
[2] R. J. Woodham, "Photometric stereo: A reflectance map technique for determining surface orientation from image intensity," *Proc. SPIE*, vol. 155, 1978.
[3] B. K. P. Horn, "Understanding image intensity," *Artificial Intell.*, vol. 8, no. 11, 1977.
[4] F. E. Nicodemus *et al.*, "Geometrical considerations and nomenclature for reflectance," Nat. Bur. Stand., U.S. Dep. Commerce, NBS Monograph 160, 1977.
[5] M. P. D. Carmo, *Differential Geometry of Curves and Surfaces.* Englewood Cliffs, NJ: Prentice-Hall, 1976.
[6] D. M. Erway, "Exact color reproduction," B.S. thesis, Dep. Elec. Eng. Comput. Sci., Massachusetts Inst. of Technol., Cambridge, 1978.
[7] E. Land and J. McCann, "Lightness and retinex theory," *J. Opt. Soc. Amer.*, vol. 61, no. 1, 1971.
[8] K. Ikeuchi, "Numerical shape from shading and occluding contours in a single view," Artificial Intell. Lab., Massachusetts Inst. of Technol., Cambridge, AI-Memo. 566, 1980.
[9] D. Marr and T. Poggio, "Cooperative computeration of stereo disparity," Artificial Intell. Lab., Massachusetts Inst. of Technol., Cambridge, AI-Memo. 364, 1976.

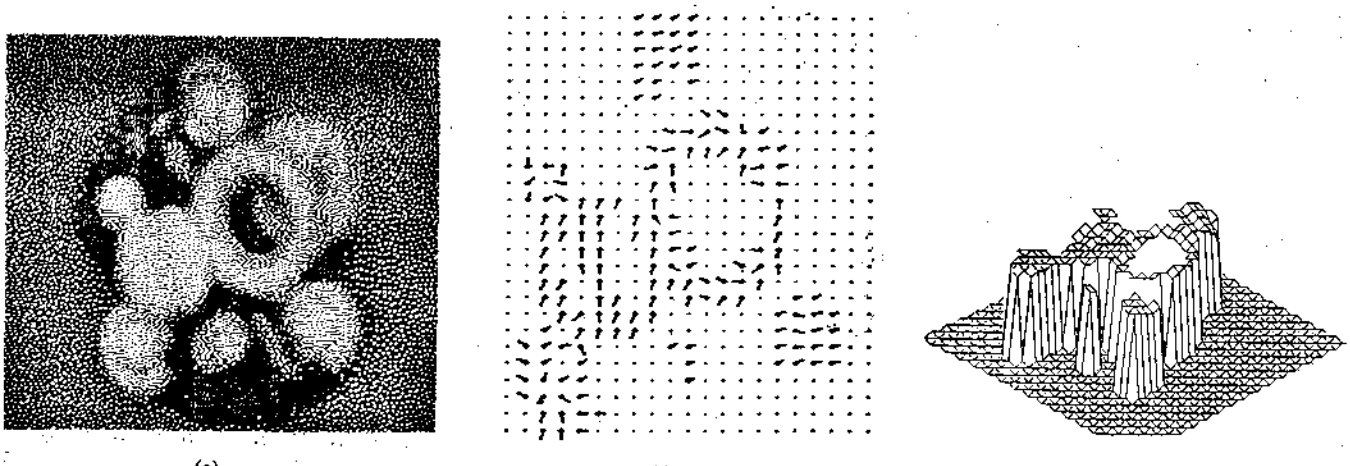

(a) (b) (c)

Fig. 9. (a) Binary image of a jumble of nuts and bolts. (b) Needle diagram obtained from image of nuts and bolts. (c) The generated surfaces are elevated in order to distinguish easily from the surrounding area.

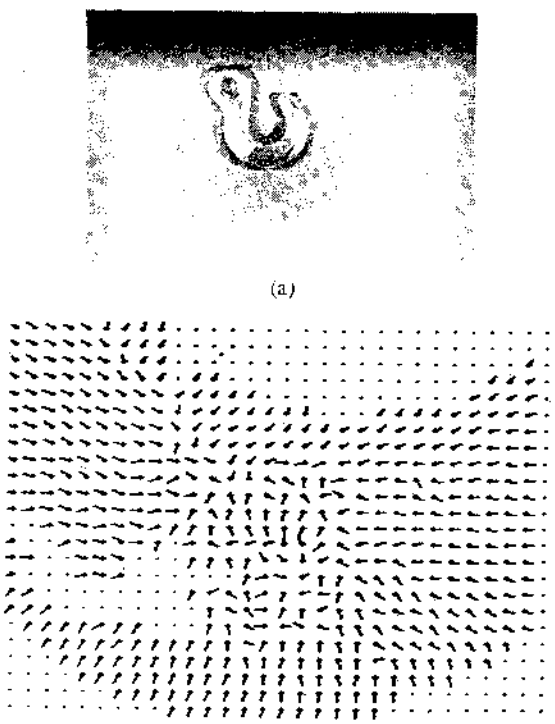

(a) (b)

Fig. 10. (a) A picture of a hook which contains a lot of cracks and stains. (b) Needle diagram of the hook. There exists a printed figure "3" in the central area of the object. This can be seen as randomness of needles in the area. At the low right side, the surface is rather dark and the system is unable to determine the surface orientations at that area. Note that a relaxation method would fail to identify this central region as different from its surroundings.

[10] W. E. L. Grimson and D. Marr, "A computer implementation of a theory of human stereo vision," in *Proc. Image Understanding Workshop*, Science Applications, Inc., Washington, DC, 1979.

[11] B. K. P. Horn, "Determining shape from shading," in *The Psychology of Computer Vision*, P. H. Winston, Ed. New York: McGraw-Hill, 1975.

Katsushi Ikeuchi (M'78) was born in Osaka, Japan, on May 29, 1949. He received the B.Eng. degree in mechanical engineering from Kyoto University, Kyoto, Japan, in 1973, and the M.Eng. and D.Eng. degrees in information engineering from the University of Tokyo, Tokyo, Japan, in 1975 and 1978, respectively.

From 1978 to 1980 he was a member of the Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge. Since 1980 he has been with the Electrotechnical Laboratory, Ministry of International Trade and Industry, Ibaraki, Japan, where he has been working on computer vision.

Dr. Ikeuchi is a member of the Institute of Electronics and Communication Engineers of Japan.