Shape Recovery Methods for Visual Inspection

Shree K. Nayar

Department of Computer Science, Columbia University, New York, N.Y. 10027

Abstract

The advancement of three-dimensional machine vision is closely related to the development of robust and efficient shape recovery methods. We address the recovery problem associated with three different classes of surfaces: (a) specular surfaces; (b) surfaces with varying reflectance; and (c) rough and textured surfaces. For specular surfaces, we have developed the structured highlight technique that uses a large array of point sources to recover surface orientation from image highlights. The sources are efficiently scanned using a binary coding scheme. For surfaces with spatially varying reflectance properties, the photometric sampling method is presented. This method uses a small array of extended sources to compute local estimates of surface orientation and reflectance, simultaneously. Photometric sampling can adapt to reflectance variation from specular to diffuse over the object surface. Finally, we address the recovery problem associated with rough and textured surface. We show that focus analysis provides an effective means of extracting the shapes of surfaces with complex roughness and reflectance properties. Three real-time machine vision systems have been developed based on these results. Experimental results demonstrate that the proposed methods and systems are applicable to a variety of visual inspection problems.

1 Introduction

Automatic inspection systems are vital to the advancement of manufacturing technology and quality control. Today, most real-time inspection systems use binary image processing techniques and pattern matching methods to compare inspected objects with their models. There are however several industrial problems that require the inspection of the three-dimensional shapes of manufactured parts. The advancement of three-dimensional machine vision is largely dependent on the development of robust and efficient shape recovery methods.

Machine vision research has produced several shape recovery techniques such as binocular stereo, structured light, photometric stereo, and shape from shading. Though these techniques are novel in their approaches, they each have limitations that often preclude them from being viable solutions to real-world problems. The current status of vision research suggests that there is no single method that can be applied to all types of surfaces. Surfaces are characterized by not only their shapes but also their reflectance and roughness properties. These properties may vary from one inspection task to another. Existing recovery methods are often incapable of adapting to substantial variations in surface reflectance. Further, there are types of surfaces that cannot be handled by any of these vision techniques.

On the assembly line, several difficult inspection tasks are still performed by human inspectors. We know that the human vision system is powerful and versatile; it can perceive the shapes and reflectance properties of a wide range of surfaces. However, there are surfaces, such as microscopic surfaces, that are difficult to perceive even for the human eye. For some of these problems, machine vision can provide effective solutions. By focusing on a particular problem, it is possible to develop machine vision systems that out-perform humans in both accuracy and speed.

In this paper, we present three shape recovery methods. These recovery methods are based on physical models of surface reflection and image formation. Each method has been developed to recover the shapes of a particular class of surfaces. The surface classes examined here are (a) specular surfaces; (b) surfaces with varying reflectance, and (c) rough and textured surfaces. In the following sections, we describe the problems posed by these surfaces and the vision techniques we have developed to solve them. These results indicate that the proposed shape recovery methods can be used to solve several challenging visual inspection tasks.

2 Specular Surfaces

Surfaces of machined or polished metal parts and solder joints are examples of specular surfaces. For a purely specular surface, light is reflected such that the angle of incidence equals the angle of reflection. Illumination of a specular surface using a point source of light does not produce smooth shading on the surface. Fig. 1 shows a camera image of several specular objects. Images of such objects are difficult to interpret as they are characterized by bright points or *highlights*, and inspection of surface shape is a challenging task.

The structured highlight technique uses a large array of point sources to illuminate the inspected object [Sanderson, Weiss, & Nayar 88] [Nayar et al. 90]. The point sources are uniformly distributed around the object, and images of the object are obtained using a camera. The

0-8186-2840-5/92 $03.00

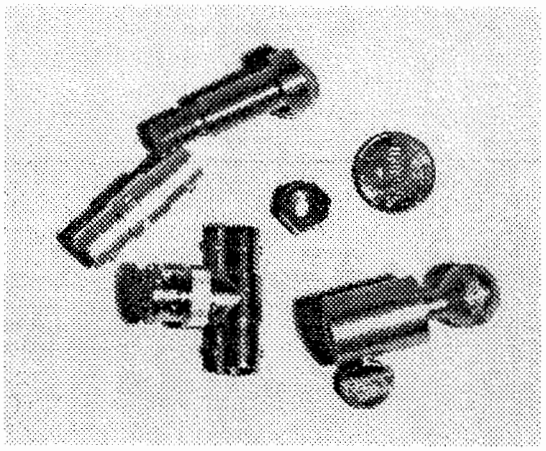

Figure 1: Image of several specular objects.

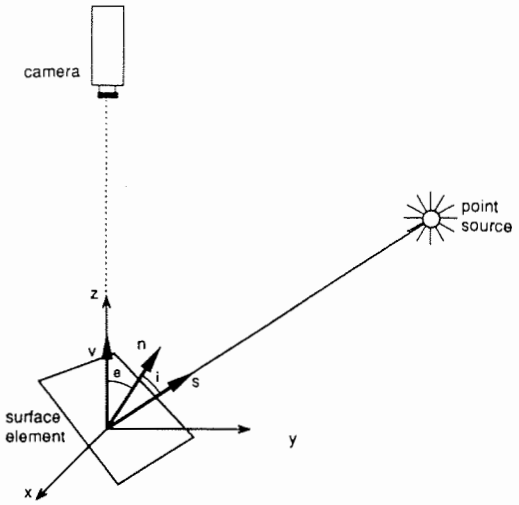

Figure 2: Surface element at the origin of a viewer-oriented frame.

point sources are activated and highlights on the object surface are used to compute local surface orientations. The sequential illumination of the sources is time consuming and would make it impossible to use the method for real-time inspection tasks. To overcome this problem, we have developed a binary coding scheme that activates selected sets of sources to obtain a small set of images. A prototype structured highlight inspection system, called SHINY, has been implemented and used to extract the shapes of specular objects including solder joints.

2.1 Structured Highlight

Consider the illumination and imaging geometry shown in Fig. 2. A surface element is located at the origin of a viewer-oriented coordinate frame and is illuminated by a point source of light. The source vector s is determined by the position of the source and the sensor vector v is determined by the location and image projection model of the camera. In the case of specular reflectance, the surface element has to be properly oriented to reflect light from the point source into the camera. If the point source is moved around the surface element, a highlight will be detected in the camera image at some source location. If the source vector s is known at the point of highlight detection, the normal (orientation) vector n of the surface element can be determined using the specular constraint that requires the source, sensor, and normal vectors be coplanar and the angle of reflection be equal to the angle of incidence:

$$\mathbf{n} = \frac{\mathbf{v} + \mathbf{s}}{\| \mathbf{v} + \mathbf{s} \|} \qquad (1)$$

Structured highlight uses a large array of point sources to determine local surface orientations of specular surfaces. The estimation of orientations is based on the *distant source assumption*; the distance of each source from the origin of the viewer-oriented frame is assumed to be large compared to the distance of surface points from the origin. Under this assumption, the direction of any particular source is the same for all points on the object surface, and is determined only by the position of the source in the viewer-oriented frame.

Consider a specular surface illuminated by two point sources from two different directions. The resulting image of the object may contain several highlights. Since the shape of the surface is not known, it is impossible to determine which of the two sources produced any particular highlight. This information, however, is necessary for the estimation of surface orientation. The simplest way to overcome this problem is by activating the point sources one at a time, obtaining an image of the object for each source, and computing orientations from the highlights. This approach is referred to as *sequential scanning*. In our implementation, we have used an alternative approach where the point sources are scanned using coded patterns to improve the efficiency of the scanning process. For a large array of sources, *coded scanning* is far more efficient than sequential scanning.

In Fig. 3, the coded scanning technique is illustratated for the case of 7 sources. The same approach can be used for a larger number of sources. First, the source numbers are converted into their corresponding binary codes. The numbers from 1 through 7 can each be uniquely expressed in binary by using 3 bits, namely, *Bit* 1, *Bit* 2, and *Bit* 3. For the first scan, namely, *Scan* 1, all point sources that have a high *Bit* 1 are activated and the remaining sources are turned off. An image of the surface is obtained and converted to a binary image using an appropriate threshold level. A "1" in the binary image corresponds to a highlight while "0" implies no highlight. The binary image corresponding to *Scan* 1 is shown as *Binary Image* 1. In a similar manner, *Binary Image* 2 and *Binary Image* 3 are obtained for *Scan* 2 and *Scan* 3, respectively. By reading the image values at the same pixel location (i,j) in all three binary images, we obtain a 3-bit pattern; (1,0,1) for the example considered. We assume that, due to high surface specularity, only a single source can generate a highlight at any particular point on the surface. Therefore, the bit pattern

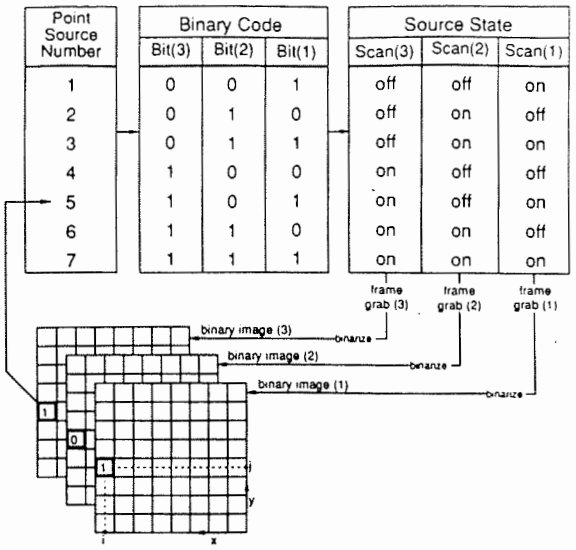

Figure 3: An example of point source scanning using the binary coding scheme. Seven point sources are scanned using three illumination patterns.

(1,0,1) could result only if the surface point corresponding to image point (i,j) reflected light from source 5 into the camera. The surface orientation at point (i,j) is computed by using the known direction of source 5, the viewing direction of the camera, and the specular reflectance model. This coding scheme can be generalized such that $2^N - 1$ point sources can be scanned using N illumination patterns. In our implementation of the structured highlight method, we have used a total of 127 sources and therefore 7 images of the object are needed to compute a complete orientation map.

2.2 SHINY System

The Structured Highlight Inspection System (SHINY) was implemented [Nayar et al. 90] to demonstrate the practical feasibility of the structured highlight method. A photograph of the orientation measuring device used by the SHINY system is shown in Fig. 4. The device consists of a hemispherical array of 127 point sources. Each point source is generated using a 1000 milli-candela light emitting diode coupled to a fiber-optic cable. The object is placed at the center of the hemisphere. The set-up provides a good distant-source approximation for objects that are 0.12 inches or less in diameter. Four CCD cameras are mounted on the hemisphere and all cameras are positioned and focused on a small region around the center of the hemisphere. Objects are brought under the view of the cameras using an X-Y-θ servo-controlled table. The object images are processed using a Sun work-station.

Experiments were conducted on a variety of test objects and the surface orientations measured by the system were found to be within 3% of the actual orientations. Fig. 5 shows the surface normal vectors of a polished sphere that is 1/16 inch in diameter, measured by the system.

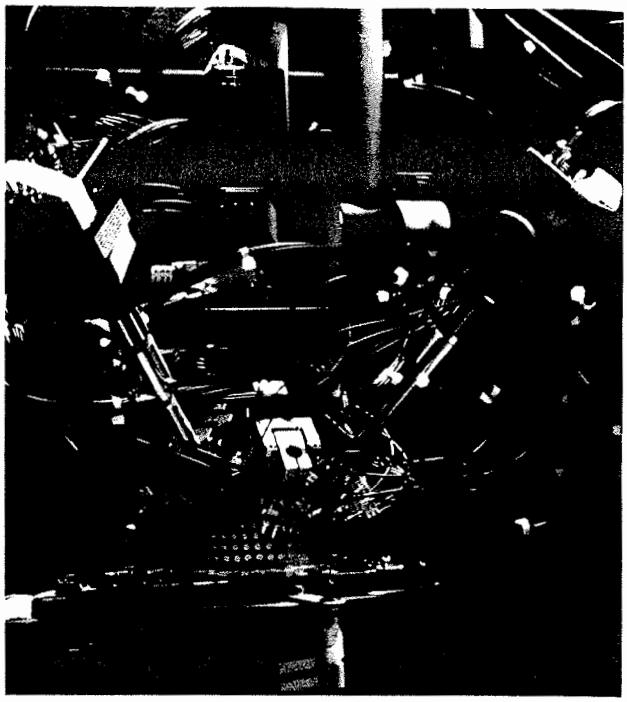

Figure 4: The structured highlight inspection system (SHINY) developed to recover the shapes of specular surfaces.

In order to demonstrate the accuracy of the system, the measured vectors are extended to intersect at the center of the sphere. A more recent version of the SHINY system is being used by Westinghouse Electic Corporation to inspect solder joints on surface-mount circuit boards [Nayar et al. 90]. Using customized hardware, the Westinghouse system inspects solder joints at a rate of 2 joints per seconds. A typical surface-mount board (with about 2000 joints) can be inspected by the SHINY system in approximately 18 minutes. A human inspector could take anywhere from 60 to 90 minutes to inspect the same board.

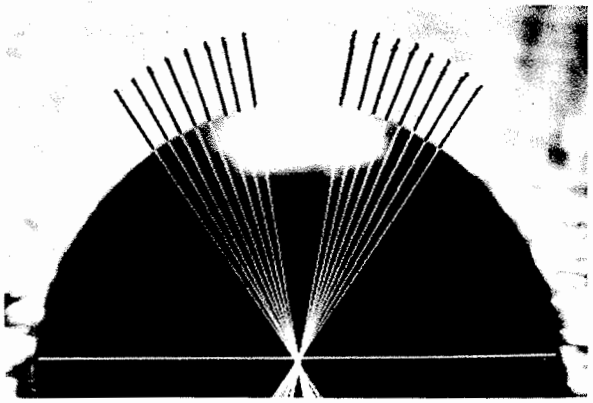

Figure 5: Surface normal vectors of a polished sphere computed using the SHINY system.

3 Surfaces with Varying Reflectance

The structured highlight method is effective and efficient only for the class of specular surfaces. Surfaces that have a significant diffuse component of reflection produce shading or brightness variation, and image intensity values must be analyzed to recover their shapes. Shape from shading [Horn 70][Pentland 84] and photometric stereo [Woodham 78][Coleman & Jain 82] are examples of techniques that extract three-dimensional shape information from image intensities. All of these techniques are applicable to diffuse surfaces but they rely on prior knowledge of the reflectance model of the surface. Further, these methods assume that all points on the object have the same reflectance properties. The method by Coleman and Jain [Coleman & Jain 82] can handle surfaces with diffuse and specular components but it relies on the existence of a significant diffuse component; it cannot be used to recover surface regions that are primarily specular. Many objects encountered in practice consist of surface regions that have very different reflectance properties. Fig. 6 shows several objects with regions of different reflectance characteristics; diffuse and specular. The above shape recovery methods are not applicable to objects such as these.

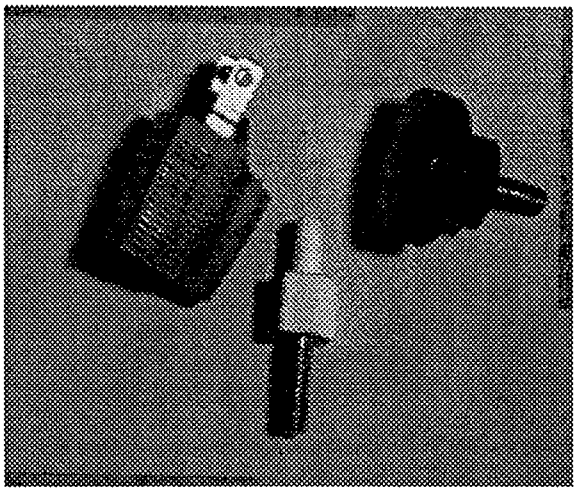

Figure 6: Objects with surface regions of different reflectance properties.

In this section, we present *photometric sampling* [Nayar, Ikeuchi, & Kanade 90], a shape-from-intensity method that can adapt to variations in surface reflectance. Reflectance invariance is achieved by computing both shape and reflectance, simultaneously. We use a hybrid reflectance model that is a linear combination of diffuse and specular components. The relative strengths of the diffuse and specular components are unknown and can vary from one surface point to the next. Ikeuchi [Ikeuchi 81] demonstrated that extended light sources are needed to obtain measurements from a continuous range of orientations on a specular surface. Based on this observation, an illumination technique is proposed that uses an array of extended sources. The source directions are selected by sampling the illumination space in a systematic manner. The sources are activated one by one and for each source an image of the object is obtained. An extraction algorithm uses the set of object images and the hybrid reflectance model to estimate both orientation as well as reflectance parameters at each surface point.

3.1 Photometric Sampling

Consider the illumination of an object by a point source of light, as shown in Fig. 7. Light energy reflected by the surface in the direction of the camera causes an image of the surface to be formed. We assume the surface to be hybrid

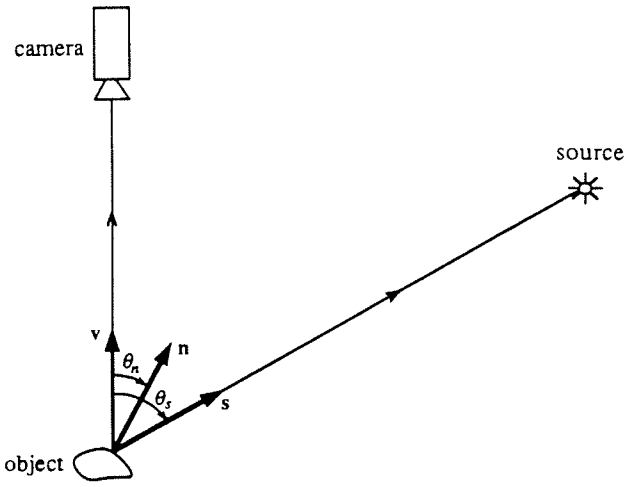

Figure 7: Two-dimensional illumination and imaging geometry. A surface element with orientation θ_n reflects light from the point source direction θ_s into the camera.

in reflectance; it has a diffuse and a specular component. The diffuse component is represented using the Lambertian model. The surfaces are assumed to be smooth. Hence, the specular intensity component is very sharp function of the source direction and may be approximated by the delta function. The *basic photometric function* relates image intensity to surface orientation, surface reflectance, and point source direction:

$$I = A\,cos(\theta_s - \theta_n) + B\,\delta(\theta_s - 2\theta_n) \qquad (2)$$

The constants A and B represent the relative strengths of the Lambertian and specular components of reflection, respectively. We call A and B the *reflectance parameters.*

3.1.1 Extended Light Sources

The photometric sampling method uses a set of extended light sources to illuminate the surface. Extended sources are chosen, rather than point sources, for the following reason. To detect specular reflections from surface points of all orientations, an infinite number of point sources need to

be positioned around the surface. Unlike a point source, an extended source emits light from an area rather than a single point. Therefore, a small number of extended sources may be used to ensure the detection of specular reflections.

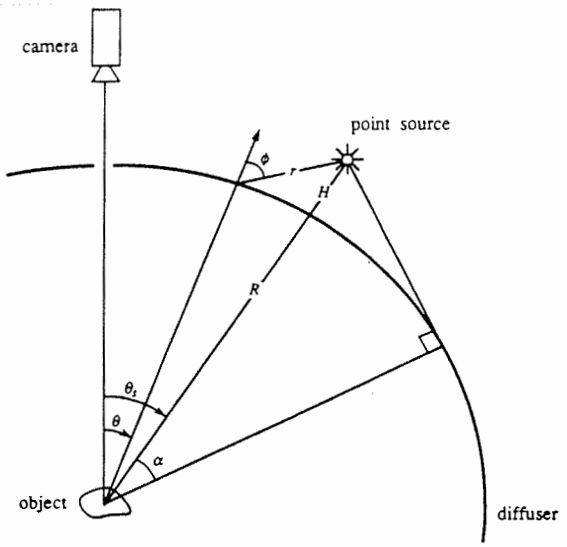

Figure 8: An extended light source.

Fig. 8 shows the technique we have used to generate extended sources. We assume that the diffuser is "ideal", i.e. any point on its surface appears equally bright from all directions. Then, the radiance of the inner surface of the diffuser is proportional to the irradiance of the outer surface. The direction of the extended source is denoted by the angle θ_s, i.e. the direction of the point source used to generate the extended source. The radiance function $L(\theta, \theta_s)$ is symmetric, or even, with respect to the source direction ($\theta = \theta_s$), and its magnitude decreases as θ deviates from θ_s. The expression for the extended source radiance is given in [Nayar, Ikeuchi, & Kanade 90] and is derived using the geometry shown in Fig. 8.

The photometric function for point source illumination (equation 15) needs to be modified for extended source illumination. An extended source may be thought of as a collection of point sources in which each point source has a radiant intensity that is dependent on its position on the extended source. Hence, the modified photometric function is determined by convolving the basic photometric function with the extended source radiance function $L(\theta, \theta_s)$. The result is [Nayar 90]:

$$I' = A' \cos(\theta_s - \theta_n) + B' L(2\theta_n, \theta_s) \quad (3)$$

Note that the Lambertian component remains a cosine function of the angle of incidence while the specular component is now represented by the extended source radiance function. The parameters A' and B' are proportional to the reflectance parameters A and B, respectively. In Fig. 9 image intensity I' is plotted as a function of extended source direction θ_s.

3.1.2 Sampling the Photometric Function

The process of measuring image intensities for different source directions is equivalent to sampling the modified photometric function $I'(\theta_s)$ (Fig. 9). Samples of the photometric function may be obtained by distributing an array of extended sources around the object. The entire array is sequentially scanned such that, for each scan, a single source is active and an image of the object surface is obtained. The sequential scanning of sources in the directions $\{\theta_i$: i=1,2,.....M$\}$ results in a set of image intensities $\{I'_i$: i=1,2,.....M$\}$ measured at each point on the object surface (Fig. 9). The number of intensities measured is determined by the frequency at which $I'(\theta_s)$ is sampled. In our implementation of the photometric sampling method we have chosen a total of 8 light sources. The reasons for this choice are given in [Nayar 90].

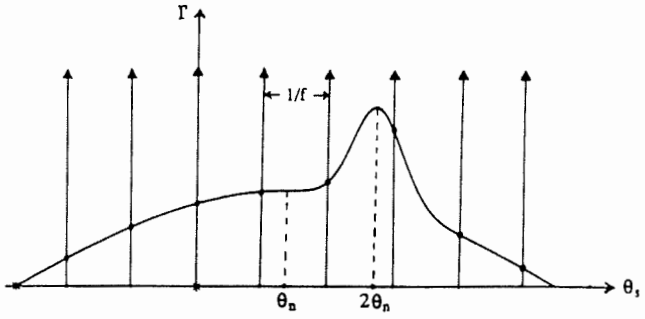

Figure 9: Sampling the modified photometric function.

3.1.3 Extracting Shape and Reflectance

The extraction problem is stated at follows: at each surface point, compute the orientation θ_n and the reflectance parameters A' and B' from the set of image intensities $\{I'_i$: i=1,2,.....M$\}$ and the source directions $\{\theta_i$: i=1,2,.....M$\}$. Equation 3 is non-linear in the unknown parameters A', B', and θ_n. Hence a closed-form solution for the parameters is difficult to obtain. We have developed an algorithm that first separates the Lambertian and specular components of each measured intensity value. It then computes orientation and reflectance parameters using the Lambertian and specular models.

A detailed description of the extraction algorithm is given in [Nayar, Ikeuchi, & Kanade 90]. It is not only applicable to hybrid surfaces that have both Lambertian and specular components, but also ideal Lambertian and ideal specular surfaces. Also noteworthy is the fact that the algorithm is local; orientation and reflectance parameters of a surface point are computed only from image intensities measured at that point. As a result, no assumptions need be made regarding the smoothness or continuity of the surface. Further, the reflectance parameters A' and B' can vary from one point on the object surface to the next.

3.2 Photo-Sampler System

A photometric sampling system was developed to recover the shape and reflectance of surfaces (Fig. 10). A 14-inch diameter lamp shade is used as a spherical diffuser and extended light sources are generated on the diffuser's surface by illuminating it using incandescent light bulbs. A total of 8 extended sources are generated in this manner. Each extended source has an angular width of 64 degrees, and adjacent sources are separated by an angle of 32 degrees. The object is placed at the center of the diffuser and is viewed by a CCD camera through a 1-inch diameter hole at the top of the diffuser. In this implementation, the light bulbs, camera, and object are all placed in the same plane. This two-dimensional set-up is capable of measuring only surface normal vectors that lie on a single plane in orientation space. For each extended source, an image of the object is digitized and stored in memory. The sequence of object images generated by scanning the array of extended sources is processed on a Sun work-station.

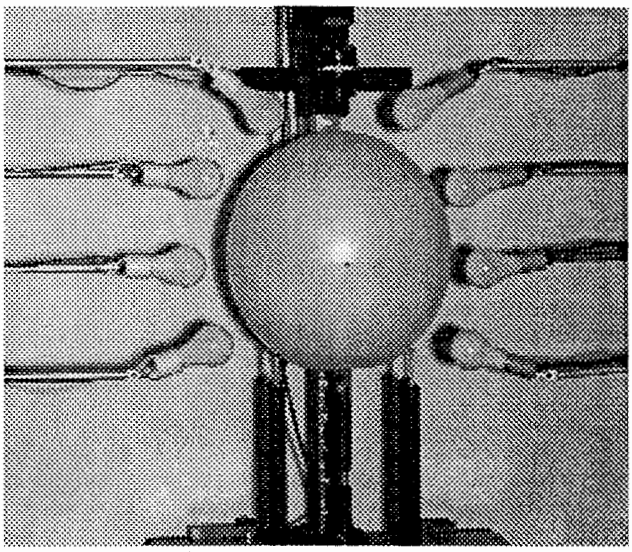

Figure 10: The photo-sampler system.

The photo-sampler and the extraction algorithm were used to compute the shape and reflectance of a variety of objects, including, matte painted objects, smooth metal objects, and plastic objects [Nayar, Ikeuchi, & Kanade 90]. Fig. 11 shows a metal bolt with two regions of different reflectance properties. The painted surface of the head is diffuse in reflection, while the threaded section of the bolt is specular. Figs. 11b and 11c show the reflectance images (Lambertian and specular) computed by the extraction algorithm. The intensities in the Lambertian and specular reflectance images, are proportional to the computed Lambertian albedo (A') and specular albedo (B'), respectively. Fig. 11d shows the needle map computed by the extraction algorithm. The needle map is a representation of surface orientations. At each point on a needle map, the length of

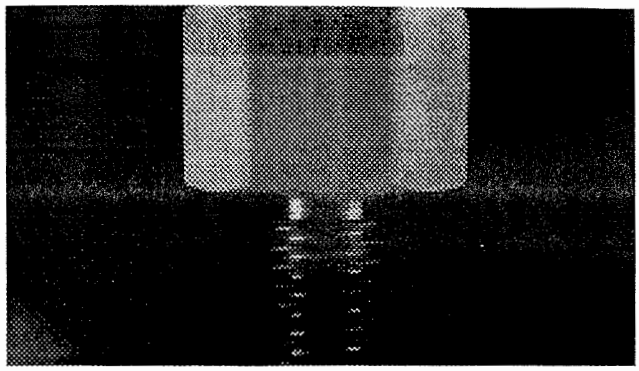

OBJECT

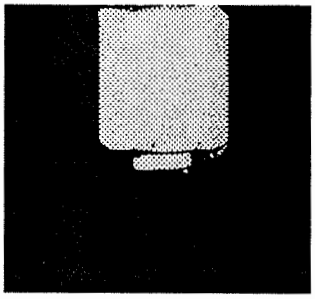

LAMBERTIAN STRENGTH

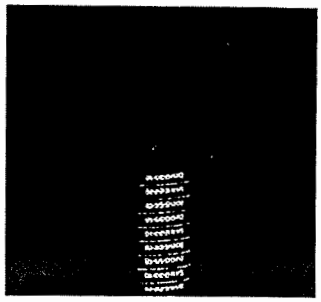

SPECULAR STRENGTH

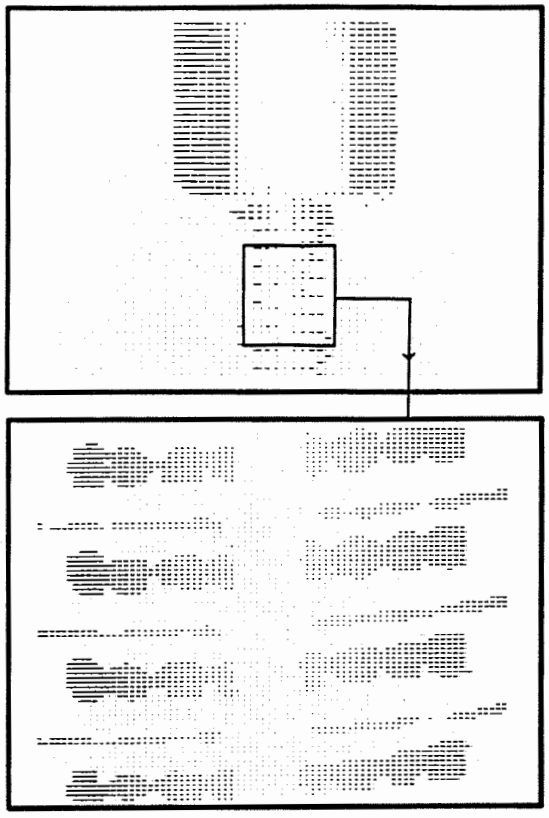

NEEDLE MAP

Figure 11: Shape and reflectance of a metal bolt recovered using the photo-sampler system. The head of the bolt is painted and has a diffuse surface, while the threaded section is specular.

the needle is proportional to the slant of the surface away from the viewing direction of the camera. An error analysis was conducted to estimate the accuracy of the computed normals. The mean of the absolute error in computed surface orientations was found to be less than 2 degrees.

In the photo-sampler implementation described above, all extended sources are positioned in the same plane. To recover the shape and reflectance of general three-dimensional surfaces, a three-dimensional version of the photo-sampler has been developed and described in [Sato et al. 90].

4 Rough and Textured Surfaces

All surfaces encountered in practice are rough at some level of detail. In many vision applications, surface variations are comparable in dimensions to the resolution of the imaging system. Fig. 12 shows a microscope image of a via-hole filling on a ceramic circuit board. Image intensities produced by such surfaces vary in an unpredictable manner from one sensor element (pixel) to the next. Hence, it is difficult to obtain dense and accurate surface shape information by using existing passive or active sensing techniques, including, the structured highlight and photometric sampling methods described in the previous sections. A practical and reliable solution to this rather difficult extraction problem is desirable.

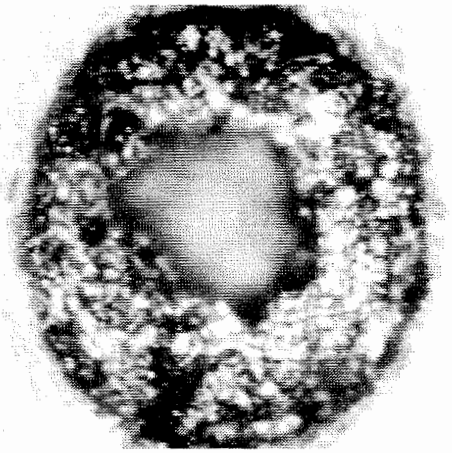

Figure 12: Image of a visibly rough surface.

Focus analysis provides a powerful means of recovering the shapes of visibly rough surfaces. Previously, focus analysis has been used to recover sparse depth maps of general scenes (see [Pentland 87], [Subbarao 89], [Krotkov 87], and [Darell & Wohn 88]). Here, we present a shape-from-focus method [Nayar & Nakagawa 90] for the accurate estimation of visibly rough surfaces. The method uses different focus levels to obtain a sequence of object images. The sum-modified-Laplacian (SML) operator is developed to provide local measures of the quality of image focus. The operator is applied to the image sequence of the object to determine a set of focus measures at each image point. A Gaussian model is used to describe the variation of the SML focus measure due to defocusing. This model is used by a depth estimation algorithm to interpolate a small number of focus measure values and obtain accurate depth estimates. An automated shape from focus system has been implemented using an optical microscope and tested on a variety of industrial samples.

4.1 Shape from Focus

Fig.13 shows an object of unknown shape placed on a translational stage. The reference plane shown corresponds to the initial position of the stage. The configuration of the optics and sensor defines a single plane, the "focused plane," that is perfectly focused onto the sensor plane. The

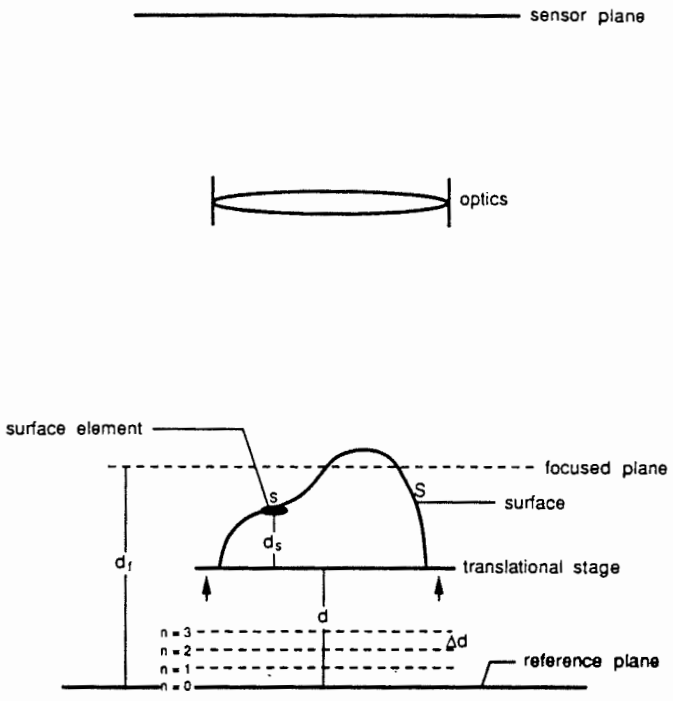

Figure 13: Shape from focus.

distance d_f between the focused and reference planes, and the displacement d of the stage with respect to the reference plane, are always known by measurement. Consider the surface element, s, that lies on the object surface, S. If the stage is moved towards the focused plane, the image of s will gradually increase in its degree of focus (high frequency content) and will be perfectly focused when s lies on the focused plane. Further movement of the element s will again increase the defocusing of its image. If we observe the image area corresponding to s and record the stage displacement $d = \overline{d}$ at the instant of maximum focus, we can compute the height d_s of s with respect to the stage as $d_s = d_f - \overline{d}$. This procedure may be applied independently to all surface elements to obtain the shape

of the entire surface S.

To automatically detect the instant of "best" focus, we will develop an image focus measure. In the above discussion, the stage motion and image acquisition were assumed to be continuous processes. In practice, however, it is not feasible to acquire and process such a large number of images in a reasonable amount of time. Therefore, we obtain only a finite number of images; the stage is moved in increments of Δd, and an image is obtained at each stage position. An interpolation algorithm computes accurate depth estimates from a small number of focus measures.

4.1.1 A Focus Measure Operator

In [Nayar & Nakagawa 90] the effects of defocusing are analyzed in frequency domain. Defocusing attenuates the high frequencies in the focused image and hence plays the role of a *low-pass* filter where the bandwidth decreases with increase in defocusing. Therefore, the effects of defocusing are more pronounced and detectable if the image has strong high-frequency content. An effective focus measure operator, therefore, must high-pass filter the image. One way to high-pass filter an image is to determine its second derivative. For two-dimensional images, the *Laplacian* may be used:

$$\nabla^2 I = \frac{\partial^2 I}{\partial x^2} + \frac{\partial^2 I}{\partial y^2} \tag{4}$$

where I is the image intensity at the point (x,y). In [Nayar & Nakagawa 90] we have shown that the result of applying the Laplacian operator to an image point may be assumed to be a Gaussian function of the defocusing parameter σ_h. This general behavior is expected irrespective of the focus measure operator used.

In the context of textured images, the Laplacian poses a problem as a focus measure operator. Note that in the case of the Laplacian the second derivatives in the x and y directions can have opposite signs and tend to cancel each other. In the case of textured images, such instances may occur frequently and the Laplacian may at times behave in a unstable manner. We overcome this problem by defining the *modified Laplacian* as:

$$\nabla^2_M I = \left|\frac{\partial^2 I}{\partial x^2}\right| + \left|\frac{\partial^2 I}{\partial y^2}\right| \tag{5}$$

Note that the modified Laplacian is always greater or equal in magnitude to the Laplacian.

The discrete approximation to the Laplacian is usually a 3×3 operator. In order to accommodate for possible variations in the size of texture elements, we compute the partial derivatives by using a variable spacing (*step*) between the pixels used to compute the derivatives. Hence, the discrete approximation to the modified Laplacian is computed as:

$$\begin{aligned} ML(x,y) &= |\,2I(x,y) - I(x-step,y) - I(x+step,y)\,| \\ &+ |\,2I(x,y) - I(x,y-step) - I(x,y+step)\,| \end{aligned} \tag{6}$$

Finally, the focus measure at a point (i,j) is computed as the sum of modified Laplacian values, in a small window around (i,j), that are greater than a threshold value:

$$F(i,j) = \sum_{x=i-N}^{i+N} \sum_{y=j-N}^{j+N} ML(x,y) \quad for\ ML(x,y) \geq T_1 \tag{7}$$

where the parameter N determines the window size used to compute the focus measure. In contrast to auto-focusing methods [Krotkov 87], we typically use a small window of size 3×3 or 5×5, i.e. $N = 1$ or $N = 2$. We shall refer to the above focus measure as the *sum-modified-Laplacian* (SML).

4.1.2 Depth Estimation

The focus measure function at any given image point may be denoted by $F(d)$; the SML focus measure as a function of the distance of the corresponding surface point from the focused plane (Fig.13). As stated in the previous section, the peak of $F(d)$ can be modeled using a Gaussian distribution with mean value $\overline{d}$ and standard deviation σ_F (Fig.14). Using the Gaussian model, the focus measure

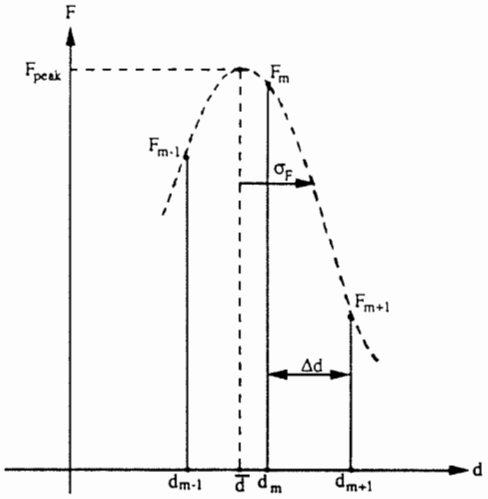

Figure 14: Gaussian interpolation of focus measures.

function may be expressed as:

$$F = F_p\ exp\left\{-\frac{1}{2}\left(\frac{d-\overline{d}}{\sigma_F}\right)^2\right\} \tag{8}$$

Using natural logarithm, we obtain:

$$ln\,F = ln\,F_p - \frac{1}{2}\left(\frac{d-\overline{d}}{\sigma_F}\right)^2 \tag{9}$$

The above Gaussian model enables us to compute depth $\overline{d}$ from just three focus measures. The fringes of the focus measure function may deviate from the Gaussian distribution since the magnification of the imaging system can vary substantially from one fringe to the other. We use three focus measures that lie on the largest mode of $F(d)$ (Fig.14).

4.2 Automated Shape from Focus System

We have implemented a fully automated shape from focus system for the recovery of microscopic objects. A photograph of the system is shown in Fig.15. A Nikon Alphaphot-2 model microscope is used to image the objects. Objects can be magnified using objective lenses with $\times 10$, $\times 40$, and $\times 100$ magnification. The object is illuminated using bright field illumination where light energy is focused on the object by the same lenses that are used to magnify the object. A CCD camera with 512×480 pixels is mounted on the microscope to obtain digital images of the object. The z-axis of the microscope stage is driven by a stepper motor and the position of the stage can be computer controlled with a resolution and accuracy of 0.02 μm. The shape from focus algorithm is programmed and executed on a Sun workstation.

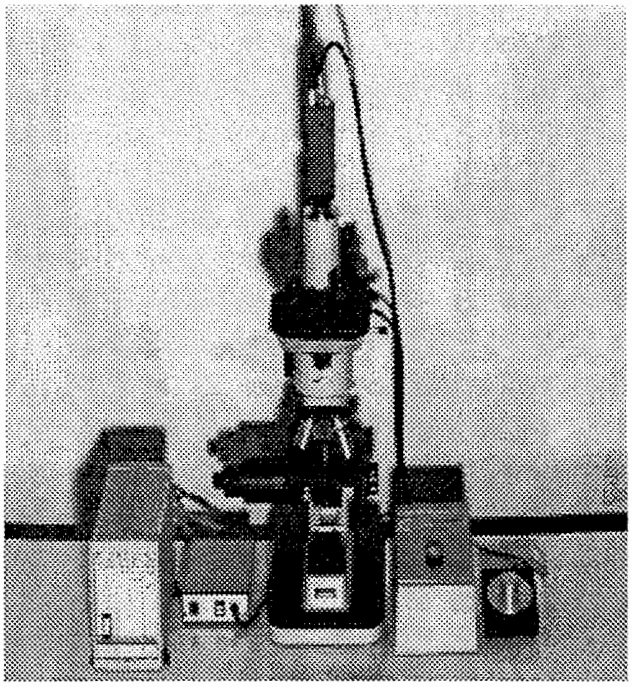

Figure 15: Automated shape from focus system.

The object is placed on the microscope stage and the appropriate objective lens is used to magnify it. The focus measure parameters (T_1 and *step*) and the stage displacement (Δd) are provided by the user. The program then automatically increments the stage position, digitizes and stores an image for each new position, and uses the image sequence to compute a depth map of the object. Fig.16 shows results obtained for a tungsten paste filling in a via-hole on a ceramic substrate. The via-hole is approximately 70 μm in diameter. A total of 18 images of the via-hole were obtained using stage position increments of $4\mu m$. Several other results are reported in [Nayar & Nakagawa 90].

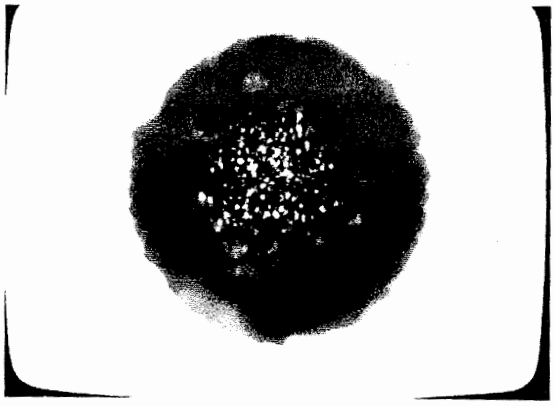

Camera image

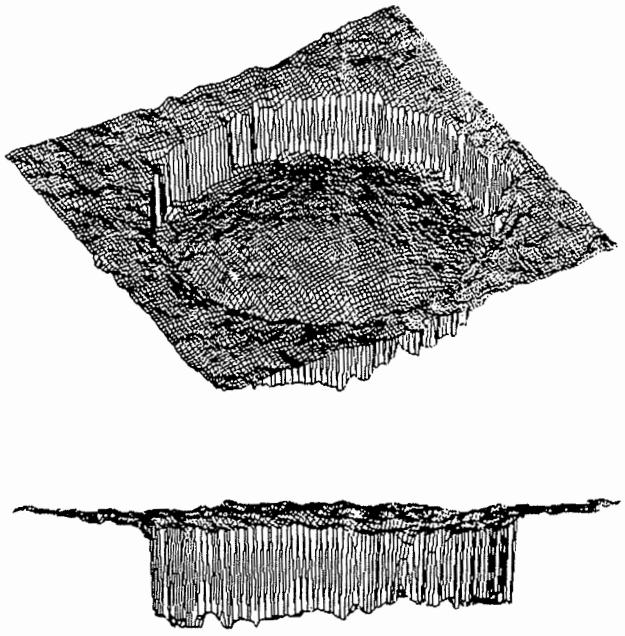

Depth maps.

Figure 16: Depth maps of a via-hole filling computed using the shape from focus system.

5 Conclusion

In this paper, we have presented new shape recovery methods for three-dimensional machine vision. We conclude with the following remarks:

- The results presented here demonstrate that it is possible to develop effective shape recovery techniques but focusing on a particular class of surfaces. The three classes we have examined here are: specular surfaces; surfaces with varying reflectance; and rough and textured surfaces. Each class encompasses a wide range of real-world objects and inspection problems.
- The shape recovery methods developed here are: structured highlight; photometric sampling; and shape from focus. These methods are based on physical models of surface reflectance and image formation. In general, the estimation of three-dimensional shape from a single two-dimensional image is an under-constrained problem. We have used active illumination and sensing techniques to constrain the problem and recover accurate shape estimates.
- Three automated vision systems have been developed that are based on the above recovery techniques: the SHINY system for specular surfaces; the photo-sampler system for surfaces with varying reflectance; and the shape from focus system for rough and textured surfaces. Each system has been tested on industrial samples. The results demonstrate that these systems can provide practical solutions to a variety of inspection problems.

Acknowledgements

The systems presented in this paper were developed in collaboration with Arthur Sanderson, Lee Weiss, and David Simon (SHINY system); Katsushi Ikeuchi and Takeo Kanade (Photo-sampler system); and Yasuo Nakagawa and Ushir Shah (shape from focus system). This research was conducted at the Robotics Institute, Carnegie Mellon University; the Production Engineering Research Lab., Hitachi Ltd.; and the Center for Research on Intelligent Systems, Columbia University.

References

[Coleman & Jain 82] E. N. Coleman and R. Jain, "Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry," *Computer Graphics and Image Processing*, Vol. 18, No. 4, pp. 309-328, April, 1982.

[Darell & Wohn 88] T. Darrell and K. Wohn, "Pyramid Based Depth from Focus," *Proc. Computer Vision and Pattern Recognition Conf.*, pp. 504-509, 1988.

[Horn 70] B. K. P. Horn, "Shape from Shading: A Method for Obtaining the Shape of a Smooth Opaque Object from One View," MIT Project MAC Internal Report TR-79 and MIT AI Laboratory Technical Report 232, November, 1970.

[Ikeuchi 81] K. Ikeuchi, "Determining surface orientations of specular surfaces by using the photometric stereo method," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 3, No. 6, pp. 661-669, November, 1981.

[Krotkov 87] E. Krotkov, "Focusing," *International Journal of Computer Vision*, Vol. 1, pp. 223-237, 1987.

[Nayar et al. 90] S. K. Nayar, A. C. Sanderson, L. E. Weiss, D. D. Simon, "Specular Surface Inspection Using Structured Highlight and Gaussian Images," *IEEE Tran. on Robotics and Automation*, Vol. 6, No. 2, pp. 208-218, April, 1990.

[Nayar & Nakagawa 90] S. K. Nayar and Y. Nakagawa, "Shape from Focus: An Effective Approach for Rough Surfaces," *Proc. IEEE Robotics and Automation Conf.*, pp. 218-225, May 1990.

[Nayar, Ikeuchi, & Kanade 90] S. K. Nayar, K. Ikeuchi, T. Kanade, "Determining Shape and Reflectance of Hybrid Surfaces by Photometric Sampling," *IEEE Trans. on Robotics and Automation*, Vol. 6, No. 4, pp. 418-431, August, 1990.

[Nayar 90] S. K. Nayar, *Shape Recovery using Physical Models of Reflection and Interreflection*, Ph.D. Dissertation, Department of Electrical and Computer Engineering, Carnegie Mellon University, December, 1990.

[Nayar, Ikeuchi, & Kanade 91] S. K. Nayar, K. Ikeuchi, T. Kanade, "Surface Reflection: Physical and Geometrical Perspectives," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 13, No. 7, pp. 611-634, July 1991.

[Pentland 84] A. P. Pentland, "Local Shading Analysis," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 6, No. 2, pp. 170-187, March, 1984.

[Pentland 87] A. Pentland, "A New Sense for Depth of Field," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 9, No. 4, pp. 523-531, July 1987.

[Sanderson, Weiss, & Nayar 88] A. C. Sanderson, L. E. Weiss, and S. K. Nayar, "Structured Highlight Inspection of Specular Surfaces," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 10, No. 1, pp. 44-55, January, 1988.

[Sato et al. 90] H. Sato, S. K. Nayar, and K. Ikeuchi, "Implementation and Evaluation of a Three-Dimensional Photometric Sampler," Technical Report, Carnegie Mellon University, November 1990.

[Subbarao 89] M. Subbarao, "Efficient Depth Recovery through Inverse Optics," *Machine Vision for Inspection and Measurement*, edited by H. Freeman, Academic Press, 1989.

[Woodham 78] R. J. Woodham, "Photometric stereo: A reflectance map technique for determining surface orientation from image intensity," *Proc. SPIE*, Vol. 155, pp. 136-143, 1978.