

Specular Surface Inspection Using Structured Highlight and Gaussian Images

SHREE K. NAYAR, MEMBER, IEEE, ARTHUR C. SANDERSON, SENIOR MEMBER, IEEE,
LEE E. WEISS, MEMBER IEEE, AND DAVID A. SIMON, MEMBER, IEEE

Abstract—Specular surfaces pose a challenging shape extraction problem; camera images of these surfaces are difficult to interpret because they are often characterized by highlights. The structured highlight inspection method uses an array of point sources to illuminate a specular object surface. The point sources are scanned, and highlights on the object surface resulting from each source are used to derive local surface orientation information. The extended Gaussian image (EGI) is obtained by placing at each point on a Gaussian sphere a mass proportional to the area of elements on the object surface that have a specific orientation. The EGI summarizes shape properties of the object surface and may be efficiently calculated from structured highlight data without surface reconstruction. Features of the estimated EGI including areas, moments, principal axes, homogeneity measures, and polygonality may be used as the basis for classification and inspection. The structured highlight inspection system (SHINY) has been implemented using a hemisphere of 127 point sources. The SHINY system uses a binary coding scheme to make the scanning of point sources efficient. Experiments have used the SHINY system and EGI features for the inspection and classification of surface-mounted solder joints. These experiments show excellent consistency with visual inspection and demonstrate the feasibility of the approach for production-line inspection systems.

I. INTRODUCTION

COMPUTER vision plays an important role in the automation of a large number of industrial processes. Surface inspection of machined parts, printed circuit boards, solder joint surfaces, and plastic sheets are a few examples of industrial tasks that are desirable to have automated because they are laborious for humans to perform. In recent years, considerable effort have been directed towards the development of automated inspection systems. A typical inspection system first attempts to extract surface features of objects placed in its field of view. The system then recognizes the objects on the basis of the extracted features and classifies them according to the application requirements.

A key step in the process of inspection is the extraction of the three-dimensional shape of the inspected surface. It is difficult to determine the shape of surfaces that have a dominant *specular* component of reflection [11]. Surfaces of machined or polished metal parts and solder joints are examples of spec-

ular surfaces. For a purely specular surface, light is reflected such that the angle of incidence equals the angle of reflection. Therefore, illumination of a specular surface using a point source of light does not produce smooth shading on the surface. Camera images of such surfaces are difficult to interpret because they are characterized by bright points or *highlights*, and inspection of surface shape is a challenging task.

The structured highlight technique uses a large number of point sources to illuminate the inspected object. The point sources are uniformly distributed around the object, and images of the object are obtained by using a video camera. The point sources are activated, and highlights on the object surface are used to compute local surface orientations. A prototype structured highlight inspection system, (SHINY) has been implemented and used to demonstrate the extraction of surface shape for several test objects including solder joints. Surface orientations measured by the SHINY system may be used to reconstruct the three-dimensional shape of a surface. Shape reconstructions are useful for the purpose of visual display. However, the process of reconstruction does not increase the amount of shape information that is available in the form of surface orientations. Therefore, in cases where surface orientation is measured rather than depth, surface inspection may be based on orientation values.

Extended Gaussian images (EGI) are used to represent the three-dimensional shape of an object surface [13]. The orientation histogram is used as a discrete estimate of the EGI, and shape features derived from the orientation histogram represent distribution properties of surface element orientations. This approach permits efficient shape classification without tedious registration and matching of local position features.

The structured highlight method coupled with EGI feature extraction is particularly well suited to solder-inspection applications. Solder surfaces have a dominant specular component of reflection, and thus, the structured highlight approach may be used to compute local surface orientations. The quality of solder joints is strongly related to the shape of solder flow, and the EGI features are effective in capturing this shape information. This approach is potentially more effective than other optical techniques for solder inspection [21]–[24], which use direct imaging of highlight patterns. In addition, the EGI features are very efficient to compute and are robust with respect to position and orientation of objects in the images.

This paper presents structured highlight as an efficient shape extraction technique for specular surfaces. The EGI representation of shape information and the derivation of features are

Manuscript received May 11, 1988; revised December 22, 1989. This work was supported by the Robotics Institute of Carnegie Mellon University and by the Westinghouse Electric Corporation.

S. K. Nayar, L. E. Weiss, and D. A. Simon are with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213.

A. C. Sanderson was with the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213. He is now with the Department of Electrical, Computer, and System Engineering, Rensselaer Polytechnic Institute, Troy, NY 12180-3590.

IEEE Log Number 9034764.

discussed. We then describe our implementation of the SHINY system and its application to solder inspection.

The Westinghouse Electric Corporation is currently making efforts to commercialize the SHINY system. Westinghouse has recently announced the completion of their first implementation of the SHINY system. With the use of special-purpose hardware, this system demonstrates the ability to perform high-speed inspection of specular surfaces. We will conclude this paper with a brief description of the Westinghouse SHINY system.

II. STRUCTURED HIGHLIGHT

Consider the illumination and imaging geometry shown in Fig. 1. A surface element is placed at the origin of a viewer-oriented frame and is illuminated by a point source of light. Light rays reflected by the surface in the viewing direction v are detected by the camera. For specular reflectance, the surface has to be properly oriented to reflect light from the point source into the camera. The source direction s is determined by the position of the source in the frame, and the viewing direction v is determined by the location and image projection model of 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 position of the source is known at the point of highlight detection, the orientation of the surface element may be determined by using the viewing direction, source direction, and the specular constraint that requires that the angle of reflection be equal to the angle of incidence.

The structured highlight technique [11] for determining local orientations of a specular surface uses an array of point sources to illuminate the surface. The estimation of orientations is based on the *distant source assumption*; the distances of the point sources from the origin of the viewer-oriented frame are assumed to be large compared with the distance of points on the object surface from the origin. Under this assumption, the direction of a particular point 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. Further, we assume an orthographic projection model for the camera; the viewing direction v is constant over the entire field of view. Thus, each point source produces highlights in the image at surface points that have the same orientation. The surface-orientation space is spanned by scanning an entire array of point sources.

Consider a specular surface illuminated by two point sources from two different directions. The resulting image of the object may contain many highlights. If the shape of the surface is not known, it is impossible to determine which of the two sources produced a 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*. We use an alternative approach where the point sources are scanned using encoded patterns to improve the time efficiency of the scanning process. For a large array of sources,




Fig. 1. Surface element is placed at the origin of a viewer-oriented frame. Light from the point source is reflected into the camera. For a specular surface, the angle of incidence i equals the angle of reflection e , and the vectors v , s , and n are coplanar.




Fig. 2. Reflectance map for a spherical array of distant point sources. Scanning of point sources corresponds to scanning of the reflectance map to determine local surface orientation.

coded scanning is far more efficient than sequential scanning. A detailed description of the source coding scheme is given in Section V-B. As a result of this scanning technique, an array of $(2^N - 1)$ sources can be scanned by taking only N images of the object.

Fig. 2 shows the *reflectance map* for a spherical array of point sources. Surface orientation is represented by the two orthogonal components p and q . A surface point with orientation (p, q) has the surface normal vector $(p, q, -1)$ in the viewer-oriented frame. Each ellipse-shaped region in the reflectance map corresponds to the range of surface orientations that produce reflections from a single point source. The shape of the region changes as a function of the radial distance from the origin of the map. This is because the (p, q) space is a nonlinear representation of surface orientation. For a purely specular surface, each one of the regions would shrink to a point. In practice, however, the size of each region increases as the sharpness of surface specularity decreases. The accuracy of the estimation of surface normals is determined by the sampling angle between adjacent point source directions and the sharpness of the surface specularity.




Fig. 3. Mapping of points from the surface of an object to the Gaussian sphere.

III. EXTENDED GAUSSIAN IMAGES

The unit normal vector at each point on a surface can be mapped to a unit sphere so that the tail of the normal vector is at the center of the sphere, and its head lies on the surface of the sphere. This sphere is called the *Gaussian sphere*, and each point on it corresponds to a unique surface orientation. The EGI of an object is obtained by placing at each point on the Gaussian sphere a mass proportional to the area of surface elements on the object that have the orientation at that point on the sphere [13]. Fig. 3 shows the mapping of a point from a surface to the Gaussian sphere. In the case of convex objects with a positive Gaussian curvature at all surface points, no two surface points have the same normal vector. Hence, the above mapping is unique, and there is only one convex object corresponding to a particular EGI.

The discrete representation of an EGI is an orientation histogram. EGI's may be computed from experimental data or a mathematical description of the object. The surface of the Gaussian sphere is tessellated, or divided, into cells, with each cell corresponding to a small range of orientations. The mass assigned to each cell in the EGI is proportional to the area of surface elements on the object with orientations that fall within the bounds of that cell.

Dividing the sphere into cells of equal area is a nontrivial problem that has been discussed by many authors [13]–[15]. Tessellations are found by projecting regular polyhedra onto the unit sphere after bringing their centers to the center of the sphere. In applications such as solder inspection, it is reasonable to assume that all rotations and translations of a solder joint are confined to the plane of the circuit board. For this reason, we selected the *longitudellatitude tessellation* shown in Fig. 4. The resulting cells do not have the same area and align with each other only for rotations about the axis through the poles P and P' , which in the case of solder joints, corresponds to rotations in the plane of the circuit board. The axis through the poles of the sphere is parallel to the viewing direction, as is shown in Fig. 4. All points in the cell c_{ij} of the EGI are assumed to have the same spherical angles (θ_j, ϕ_j) , where θ_j is the azimuth angle, and ϕ_j is the zenith angle. All cells with the same zenith angle have equal cell area, which is a property of the chosen tessellation that will be extensively used in the following sections.

In determining the mass that is to be assigned to each cell in the EGI, it may be noted that the projected area of a surface

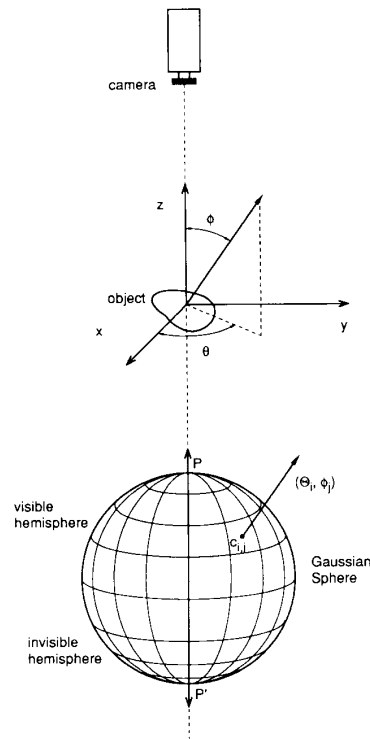


Fig. 4. Object positioned at the origin of a viewer-oriented frame. The Gaussian sphere is divided into cells along the meridians and lines of longitude, and its polar axis is oriented in the direction of the z axis of the frame.

element onto the image plane depends on the inclination of the surface with respect to the viewing direction. A surface that is steeply inclined away from viewing direction is foreshortened and appears smaller than it would if it were viewed head on. Fortunately, the angle of inclination at each surface point is exactly what has been measured as orientation, and hence, it is straightforward to correct for the above effect. Let each pixel in the image have unit area, and let n_{ij} be the number of pixels in the image with surface orientations that lie within the range corresponding to cell c_{ij} . Then, the mass m_{ij} assigned to cell c_{ij} is determined as [15]

$$m_{ij} = \frac{n_{ij}}{\cos \phi_j}. \quad (1)$$

For a given viewing direction, the EGI may be divided into the visible and invisible hemispheres. The visible hemisphere corresponds to the visible surface of the observed object. The invisible hemisphere has zero mass since it corresponds to the object surface that is occluded by the visible object surface. In Fig. 4, the portion of the Gaussian sphere above the equator is the visible hemisphere.

While generating the EGI, all information regarding the spatial distribution of surface orientations is discarded. For certain applications, this characteristic may prove to be a limitation of the EGI. On the other hand, matching object EGI's with those of stored models is, in general, computationally ef-

ficient as compared with matching shapes using depth maps. In addition, in shape-extraction methods such as structured highlight, the measurement of surface orientation, rather than depth, renders the EGI directly applicable. However, the usefulness of the EGI largely depends on the geometry of the object of interest. As stated earlier, the EGI works best for convex objects and is least suited for the representation of complex surfaces that are characterized by several concavities.

In addition to its computational advantages, the EGI simplifies the problem associated with using *multiple views* of the inspected object. In the structured highlight approach, point sources are positioned on a hemisphere around the object. Point sources that lie on the x - y plane for the viewer-oriented frame have the largest zenith angle of $\pi/2$. Since, for purely specular surfaces, the angle of reflection equals the angle of incidence, the camera in Fig. 4 is only capable of measuring surface orientations with zenith angles that lie in the range $0 < \phi < \pi/4$. For many applications, such a small detectable range of orientations may prove to be insufficient. To enhance the detectable range of the system, additional cameras may be distributed around the object. Multiple viewing directions, however, lead to the well-known correspondence problem: the task of determining which point in one camera corresponds to a particular point in another camera. Cooperative algorithms exist that perform heuristic matching of points based on features such as brightness and surface orientation. Once correspondence has been established, depth of surface elements may be obtained by triangulation. One such algorithm for stereo reconstruction based on the structured highlight approach is described in [11]. However, in using the EGI as a basis for shape representation and feature extraction, the spatial matching of image points in different cameras is unnecessary, and the fusion of information from several cameras is simplified. Each camera is used to measure an exclusive range of orientations; no two cameras are used to compute the orientation of the same surface point. Given the attitude of each camera, orientations measured in all cameras can be transformed to a common frame of reference.

Considerable attention has been given to recognizing three-dimensional objects by *matching* their EGI's [14], [15]. A measure of success of a match may be determined by finding the square of the difference in mass between each cell on the object EGI and the corresponding cell on the model EGI. The best match is one that minimizes the sum of squares for all pairs of corresponding cells. The orientation of the EGI that results in the best match determines the orientation of the object in space. However, in applications such as solder inspection, it is not possible to define a single model that represents a good solder joint. Two solder joints of different shapes may both be classified as joints of good quality. One approach to this problem is to store all possible models of acceptable solder joints. In the case of a defective solder joint, the matching process would be repeated for all models. When large numbers of shapes belong to the same class, this approach is computationally intensive. An alternative approach is presented in this paper where an object is recognized by extracting a set of features from its EGI that relate well to its physical shape.

IV. EGI SHAPE FEATURES

Global and local features may be extracted from the EGI of an object. Global features represent some of the more general characteristics of the object's physical shape. Some of the global features are determined by using the properties of the EGI [13], and these features are computed over the entire visible hemisphere of the EGI. In contrast, local features are used to describe subtleties in the shape of the object. Local features are computed over a ring on the EGI where a ring is constituted by cells that have the same zenith angle ϕ (Fig. 4). We have chosen rings to represent local characteristics since shapes such as those of solder joints show some degree of symmetry about the z axis of the viewer-oriented frame shown in Fig. 4.

A. Global Features

1) *Surface Area*: Orientation histograms plotted on a sphere have some useful properties. Since the mass m_{ij} assigned to each cell c_{ij} corresponds to the area of elements on the surface with the orientation (θ_j, ϕ_j) , the total mass distributed over all the cells is equal to the total area A_s of the observed surface. If there are k longitudinal sections and l latitudinal sections in the visible hemisphere of the EGI, we can express the surface area as

$$A_s = \sum_{i=1}^k \sum_{j=1}^l m_{ij}. \quad (2)$$

Similarly, areas over bounded regions on the EGI may be evaluated by changing the limits in the summations in (2).

2) *Center of Mass*: Let x_{ij} , y_{ij} , and z_{ij} be the coordinates of the cells c_{ij} . The coordinates of the EGI's center of mass are given by

$$X_c = \frac{\sum_{i=1}^k \sum_{j=1}^l m_{ij} x_{ij}}{\sum_{i=1}^k \sum_{j=1}^l m_{ij}} \quad (3)$$

$$Y_c = \frac{\sum_{i=1}^k \sum_{j=1}^l m_{ij} y_{ij}}{\sum_{i=1}^k \sum_{j=1}^l m_{ij}} \quad (4)$$

$$Z_c = \frac{\sum_{i=1}^k \sum_{j=1}^l m_{ij} z_{ij}}{\sum_{i=1}^k \sum_{j=1}^l m_{ij}}. \quad (5)$$

3) *Area Ratio*: Since the EGI is generated on a unit sphere, we see that z_{ij} equals $\cos \phi_j$ in (5). In addition, the product $m_{ij} \cos \phi_j$ is the foreshortened area of all surface elements with orientation (θ_i, ϕ_j) . Hence, the z component of the EGI's center of mass corresponds to the ratio of the

projected area A_{cs} of the object surface to the actual area A_s of the object surface [15]:

$$Z_c = \frac{A_{cs}}{A_s}. \quad (6)$$

4) *Mean and Variance of Zenith Angle:* The mass distribution over the set of rings that constitute the visible hemisphere of the EGI may be described by the sample mean ϕ_m and the sample variance ϕ_v of the zenith component of surface orientation:

$$\phi_m = \frac{\sum_{i=1}^k \sum_{j=1}^l m_{ij} \phi_j}{A_s} \quad (7)$$

$$\phi_v = \frac{\sum_{i=1}^k \sum_{j=1}^l m_{ij} (\phi_j - \phi_m)^2}{A_s - 1}. \quad (8)$$

B. Ring Features

1) *Strength:* The longitude/latitude tessellation of Fig. 4 is characterized by unequal cell areas. However, all cells within the same ring have the same area and shape. The ring whose cells have the zenith angle ϕ_j is referred to as ring j . The strength of ring j is the total mass contained within the ring and is determined as

$$S_j = \sum_{i=1}^k m_{ij}. \quad (9)$$

Before the remaining ring features are computed, each ring is normalized such that the total ring mass equals unity. The normalized cell masses p_{ij} are determined as

$$p_{ij} = \frac{m_{ij}}{S_j} \text{ for } 1 < i < k. \quad (10)$$

2) *Center of Mass:* The ring center of mass C lies in the plane of the ring, and its coordinates are given by

$$x_{cj} = \sum_{i=1}^k p_{ij} \cos \theta_i \sin \phi_j \quad (11)$$

$$y_{cj} = \sum_{i=1}^k p_{ij} \sin \theta_i \sin \phi_j. \quad (12)$$

3) *Inertia:* For the purpose of discussion, assume that ring j is displaced such that it lies on the x - y plane of the viewer-oriented frame, and its center coincides with the origin of the frame. The moment of inertia of the ring about any axis that passes through the origin and lies in the plane of the ring is determined by the inertia tensor

$$I_j = \begin{bmatrix} I_{xx} & I_{xy} \\ I_{xy} & I_{yy} \end{bmatrix} \quad (13)$$

where I_{xx} and I_{yy} are the moments of inertia about the x and

y axes, respectively, and I_{xy} is the product of inertia. These quantities are computed as follows [25]:

$$I_{xx} = \sum_{i=1}^k p_{ij} (\sin \theta_i \sin \phi_j)^2 \quad (14)$$

$$I_{yy} = \sum_{i=1}^k p_{ij} (\cos \theta_i \sin \phi_j)^2 \quad (15)$$

$$I_{xy} = \sum_{i=1}^k p_{ij} \sin \theta_i \cos \theta_i (\sin \phi_j)^2. \quad (16)$$

4) *Principal Axes:* Since the tensor I_j is symmetric, we can always find a set of axes for which the tensor is diagonal. These are the principal axes of inertia. Since the inertia tensor is computed with respect to the center of the ring, the principal axes pass through the center and lie on the x - y plane. The principal axes are determined by the equation [25]

$$\tan 2\psi_j = \frac{-2I_{xy}}{I_{xx} - I_{yy}} \quad (17)$$

where the two solutions to ψ_j correspond to the orthogonal directions of the principal axes. If there is symmetry in the distribution of mass over the ring, the axis of maximum symmetry would lie in one of the two principal directions.

5) *Local Homogeneity:* EGI's of smooth surfaces generally have homogeneous mass distributions. However, a global measure of homogeneity does not reveal much information, except in the case of a small domain of continuous surfaces. In addition, such a measure would be sensitive to orientation discontinuities resulting from occlusion of the viewed object by itself or other objects. On the other hand, local homogeneity is a useful feature because it represents the continuity of the surface for a small range of orientations. Ring homogeneity is defined as

$$H_j = \frac{1}{k^2} \sum_{a=1}^k \sum_{b=1}^k \frac{1}{1 + \left((p_{aj} - p_{bj}) \cos \left(\frac{\theta_a - \theta_b}{2} \right) \right)^2}. \quad (18)$$

The mass in each cell is compared with masses in all other cells of the ring. Each comparison is weighted by the square of a cosine function so that immediate neighbors of the cell have most influence on the measure. Homogeneity has the maximum value of 1 when all the cells in the ring have equal mass.

6) *Polygonality:* As mentioned in Section III, the EGI of a convex object is unique, and its surface may be unambiguously recovered from the EGI. The EGI of a convex polyhedron would have masses concentrated in a few cells; the number of cells equals the number of faces of the polyhedron. We define ring polygonality as

$$P_j = \sum_{i=1}^k p_{ij}^2. \quad (19)$$

Since p_{ij} is the normalized mass in each cell such that the total ring mass equals unity, we see that $0 < P_j < 1$. The minimum possible value of P_j depends on the number of cells in the ring, and $P_j = 1$ when only one cell in the ring has a nonzero mass. For a convex object, as the number of discrete masses on the EGI decreases, the object takes the shape of a polyhedron. Therefore, a ring with a large value of P_j is said to have a high degree of polygonality.

A *classification algorithm* uses a set of extracted features to assign an inspected object to one of many classes. The different classes are defined by the requirements of the inspection task to be automated. If the features are effective in reflecting the characteristics of the various classes, objects that belong to the same class tend to *cluster* together in feature space. By using a "large" number of objects of known classes and by observing the clustering of their feature values in feature space, it is possible to assign a region in feature space to each class. An inspected object is classified by finding the feature space region in which its feature values lie.

The features discussed in this section are not mutually independent. Some of the features supply redundant information and require careful attention in the construction of *decision criteria*. There is a tradeoff between work in the feature space and the reliability of the classification. The optimum number of features for a specific application is always difficult to determine and is generally deduced from a large number of experiments that test the discriminating power and consistency of each feature.

V. EXPERIMENTS

A. SHINY System

The SHINY prototype system was implemented to demonstrate the practical feasibility of the structured highlight technique. A photograph of the surface orientation measuring device used by the SHINY system is shown in Fig. 5. The experimental apparatus consists of a hemispherical array of 127 point sources. Each point source is generated using a 1000-mcd light emitting diode coupled to a fiberoptic cable. The object is placed at the center of the hemisphere. The setup provides a good distant source approximation for objects that are 0.12 in or less in diameter. Four cameras are mounted on the hemisphere, and all cameras are positioned and focused to observe a small region around the center of the hemisphere. Each camera is a model WV-CD22 Panasonic camera with a 512x480 pixel resolution. The physical resolution of the camera and optical system used for solder inspection is approximately 0.1 mils per image pixel.

A schematic diagram of the complete inspection system is shown in Fig. 6. Objects are brought under the view of the cameras using an X-Y-θ table. A servo controller is used to position the table within 0.1 mils of the desired position. The light source array is scanned, and images of the object are digitized and processed using a SUN 2/170 workstation. 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. 7 shows the surface normal vectors of a polished sphere that is 1/16 in in diameter, measured by the system. The normal vectors are overlaid




Fig. 5. Photo of the orientation-measuring device used by the SHINY system.

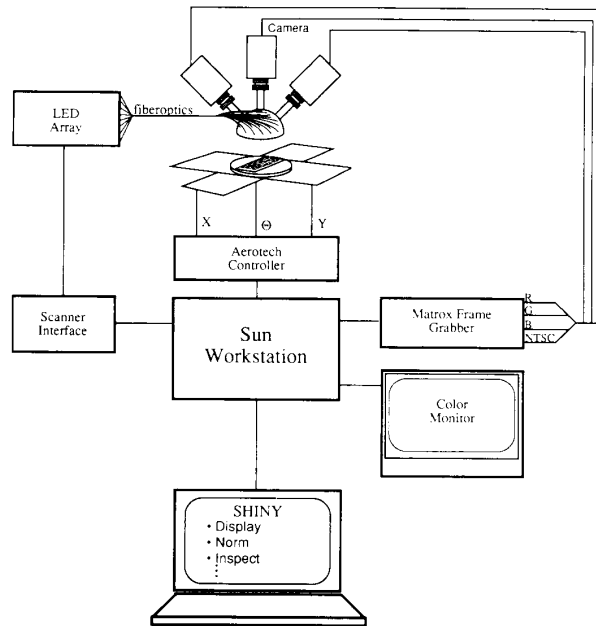


Fig. 6. Schematic diagram of the SHINY system.

on a profile image of the sphere. In order to demonstrate the accuracy of the system, the measured vectors are extended to intersect at the center of the sphere.

B. Binary Coding of Sources

The SHINY system uses a binary coding scheme to speed up the scanning of point sources. For a large number of point




Fig. 7. Photograph of surface normal vectors of a polished sphere determined by the SHINY system. The vectors are overlaid on a profile image of the sphere and extrapolated to intersect at the center of the sphere. The intersection of the normals demonstrates the accuracy of the system.

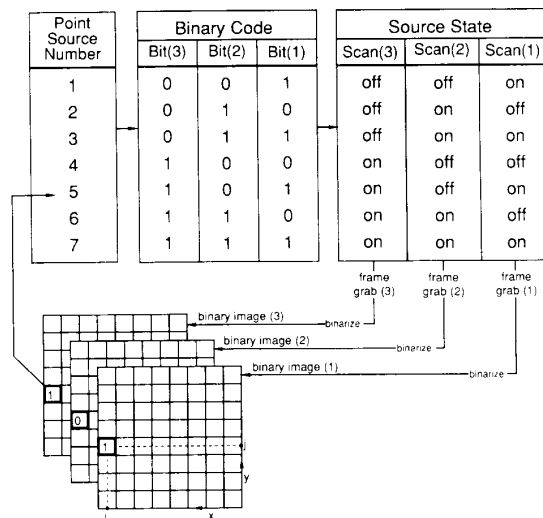


Fig. 8. Example of point source scanning using the binary coding scheme. Seven point sources are scanned in three frame grabs.

sources, coded scanning is far more efficient than sequential scanning. For example, if each frame grab takes 60 ms, then sequential scanning of 127 point sources would take a total of 7.62 s. This is unacceptable for many high-speed applications such as solder inspection. An example of point source scanning using binary coding is illustrated in Fig. 8. In the example, only seven point sources are shown, although the coding may be extended to a larger number of sources. 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 three 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 turned on, and the remaining sources are turned off. An image of the surface is grabbed into the frame buffer of the computer and converted to a binary image using a threshold. A "1" in the binary image corresponds to a highlight, whereas a "0" means no highlight. The binary image corresponding to Scan(1) is shown as Binary Image(1). In a similar manner, the Binary Image(2)




Fig. 9. Solder joint of a surface-mounted component.

and Binary Image(3) are obtained for Scan(2) and Scan(3), respectively. By reading the contents of the same pixel location (i, j) in all three binary images, we obtain a three-bit pattern, namely, (1, 0, 1) in the example. We make the assumption that, due to high surface specularity, only a single point source can generate a highlight at any particular point on the surface. Therefore, the bit pattern (1, 0, 1) in the binary images could result only if the surface point corresponding to image point (i, j) reflected light from point 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. On generalizing the above example, we see that $(2^N - 1)$ point sources can be scanned in N frame grabs. In the current system, 127 point sources are scanned in seven frame grabs. An odd or even parity scan is included to eliminate some of the errors resulting from more than one point source generating highlights at the same surface point, which is a likely event for surfaces with relatively low sharpness of specularity.

C. Solder Inspection

The reflectance model of a typical solder surface has a dominant specular component. This renders solder inspection a suitable candidate for structured highlight inspection. A large number of solder joints were inspected using the SHINY system. The orientations measured on the solder surfaces were used to generate the EGI's. Features were extracted from the EGI's, and joints were classified based on the feature values.

The inspection experiments were conducted on solder joints of surface-mounted components. Fig. 9 illustrates the shape of a typical surface-mount solder joint. The solder joint is approximately 0.06-in long and 0.035-in wide. The flow of solder around the perimeter of the lead relates to the quality of the joint. Some common solder defects such as insufficient solder, excess solder, high lead, lead overhang, and locked lead are shown in Fig. 10. Insufficient solder on the perimeter of the joint suggests a weak joint. In some cases, as a result of excessive solder deposition, it is difficult to determine whether or not the lead is under the blob of solder. Improper bending of the lead can result in high leads that do not make contact with the solder pad. The lead overhang defect results when the lead is not placed within the confines of the solder pad, and a locked lead usually has an excess of solder at the joint heel. The challenge of solder inspection lies in the fact that




Fig. 10. Some common defects associated with solder joints of surface-mounted components.

within each of these defect classes, a large number of joint shapes are possible. For this reason, the traditional approach of matching EGI's or depth maps is not a practical one.

Fig. 11(a) shows the photograph of a good joint as seen by a camera that is looking down on the circuit board. Since it is difficult to obtain good gray-level images of specular surfaces, all 127 point sources were used to illuminate the joint surface. The bright spots in the photo are highlights generated by the point sources. In order to generate the EGI, surface orientations are measured only within the region in the image where the lead makes contact with the solder pad. The bend of the lead as it emerges from the component is not of relevance to solder inspection. We exclude this part of the lead by placing a window in the image around the soldered region. The window is oversized to accommodate for inaccuracies in the placement of the lead on the solder pad and the registration of the solder pad in the image.

The EGI for the solder joint in Fig. 11(a) is shown in Fig. 11(b). The figure shows the visible hemisphere of the EGI as seen from a point on the positive z axis in Fig. 4. The zenith angle ϕ varies as we move radially away from the center of the EGI, and the azimuth angle θ varies along the concentric circles. A resolution of 16 zenith and 16 azimuth angles was chosen. However, only the eight zenith angles that belong to the visible hemisphere are seen in the photo. All cells of the same shape and size have equal zenith angles and thus belong to the same ring. The brightness of each cell is proportional to the mass in that cell. The reference frame shown in Fig. 9 may be used to understand the mapping of surface points to the EGI. Flat surfaces result in high masses near the center of the EGI, whereas surfaces that are inclined away from the viewing direction produce masses in the outer rings of the EGI. A good joint has a smooth solder flow along the perimeter of the




Fig. 11. (a) Photograph of the top view of a good solder joint; (b) Gaussian image of the solder joint shown in (a).




Fig. 12. (a) Photograph of the top view of an insufficient solder joint; (b) Gaussian image of the solder joint shown in (a).

lead. The corresponding EGI is expected to have good mass distribution along the vertical and horizontal axes. For a good joint, the top surface of the joint lead is flat and not covered with solder. This results in nonzero masses in the innermost ring. Fig. 12 shows the photos of a joint with insufficient




Fig. 13. (a) Photograph of the top view of an excess solder joint; (b) Gaussian image of the solder joint shown in (a).

solder and its EGI. Due to the lack of solder, the joint has relatively more flat surface. Consequently, the EGI has large mass in the inner rings and less mass in the outer rings. On the other hand, a solder joint with excess solder, like the one shown in Fig. 13(a), has almost no flat area since a bump of solder covers the flat lead surface. Therefore, the corresponding EGI, shown in Fig. 13(b), has less mass in the inner ring but a uniform mass distribution in the outer rings.

Features are extracted from the EGI of each solder joint, and a classification algorithm uses the set of extracted features to assign the joint to one of many classes. Fig. 14 shows a two-dimensional example of a feature space divided into regions that are each assigned to a class. Two features, namely, the strength S_1 (9) of the innermost ring of the EGI and the mean ϕ_m (7) of the zenith angle, may be used to classify joints into three classes: good, insufficient, and excess. A sample size of 15 joints is used, and as seen in the figure, joints belonging to the same class tend to cluster in a particular region of the feature space. The lines P and Q may be used to classify joints into good, insufficient, and excess classes. The features S_1 and ϕ_m work well for the three classes for this data set. However, the number of features must be increased with the number of classes. The solder inspection system is also capable of detecting lead overhang, locked leads, unsoldered leads, and cold solder joints.

In addition to classifying, the system furnishes information pertaining to symmetry in solder distribution, excessive bends in the lead, and specularities of the solder surface. The idea behind inspection is not only to classify solder joints into broad defect classes but also to provide details of the physical shape of the solder joint. Such information is valuable from the per-




Fig. 14. Solder joints are classified into good, insufficient, and excess classes by using features S_1 and ϕ_m .

spective of process control. By studying joint shapes in detail, the parameters that effect the quality of the soldering process can be controlled. By using an inspection system to close the process control loop, the number of defective solder joints can be minimized.

Approximately 170 solder joints were inspected using the current system, and only 3% of the joints were incorrectly classified. This error rate is easily within the variance observed between solder inspection experts. The SHINY system inspects at a rate of 12 s per joint. The cycle time of 12 s includes the time required to move from one joint to another, measure orientations in four cameras (which are scanned separately), generate the Gaussian image, extract features, and classify the joint. In our implementation, all four cameras use the same frame grabber, and thus, the orientation measurement requires the repetition of point source scanning for each camera. In addition, EGI's were generated by accessing the orientation at each image pixel in software. Considerable reduction in cycle time can be achieved by using a separate frame grabber for each camera and customized hardware for the generation of EGI's from orientation maps.

D. Westinghouse SHINY System

The Westinghouse Electric Corporation is currently in the process of developing the next generation of SHINY systems. The CMU prototype of the SHINY system was implemented to demonstrate the basic concepts underlying the structured highlight approach. It is also being used to aid in the design and development of the Westinghouse system. A photograph of the first implementation of the Westinghouse SHINY system is shown in Fig. 15. Similar to the CMU system, this system




Fig. 15. SHINY system implemented by Westinghouse Electric Corporation.

uses three cameras and 127 point sources to generate orientation maps of the inspected objects. The new system, however, uses three frame buffers to facilitate the parallel measurement of surface orientations by the three cameras. The EGI generation, feature extraction, and classification are executed on a Motorola 68010 computer. The computer is supplied with a map of the circuit board layout, and joints on the board are automatically positioned over the sensing device by means of an X - Y - θ servo controlled table. A high-resolution color monitor is used to display the circuit board layout. As the joints are inspected, they are displayed in colors that represent different defect classes. Further, a history database enables the user to return to any inspected joint and obtain details of its shape. Such information would be valuable for controlling the soldering process.

The additional hardware mentioned above enables the Westinghouse SHINY system to inspect joints at a rate of 2 joints/s. This is approximately 25 times faster than the CMU system. With this speed improvement, the system is able to inspect a typical surface-mount board (with about 2000 joints) in approximately 18 min. A human inspector would take anywhere from 60 to 90 min to inspect the same board.

VI. SUMMARY

We have described in this paper the development of a vision system for the automation of a challenging industrial task: the inspection of specular surfaces. The structured highlight technique uses the active illumination of a point source array to extract local surface orientations. A prototype implementation uses a binary coding scheme to expedite the process of shape extraction. Experiments conducted on a large number of test objects have produced accurate estimates of surface shape. The Gaussian image is generated directly from the measured surface orientations and serves as a powerful tool for shape representation. We have developed Gaussian image features that are effective in describing surface shape and provide a basis for the classification of objects into types that are relevant to the application.

The system described in this paper was implemented to demonstrate and evaluate the basic principles. Although the feature selection and classification have proven reliable for the inspection of sample surface-mounted circuit boards, the

versatility of the system needs to be enhanced by incorporating a self-training classifier. Effort in this direction is essential to make the system directly applicable to production lines. The results obtained so far are promising and provide a basis for the development of automatic inspection systems for specular surfaces.

ACKNOWLEDGMENT

We would like to thank Dr. R. J. Stewart of the Westinghouse Electric Corporation for his encouragement and support. We also wish to thank M. Zaremsky and J. Christ for their contributions towards the implementation of the SHINY system and Dr. J. L. Swedlow for his valuable comments. We are grateful to K. Mohnkern and R. Levine for their assistance in preparing the manuscript.

REFERENCES

- [1] B. K. P. Horn and R. W. Sjöberg, "Calculating the reflectance map," *Appl. Optics*, vol. 18, no. 1, pp. 1770-1779.
- [2] B. K. P. Horn, "Determining shape from shading," in *The Psychology of Computer Vision* (P. H. Winston, Ed.), New York: McGraw-Hill, 1975.
- [3] B. K. P. Horn, "Image intensity understanding," *Artificial Intell.*, vol. 8, no. 2, 1977.
- [4] R. J. Woodham, "Photometric stereo: A reflectance map technique for determining surface orientation from image intensity," *Proc. SPIE*, vol. 155, pp. 136-143, 1978.
- [5] R. J. Woodham, "Photometric method for determining surface orientation from multiple images," *Opt. Eng.*, vol. 19, no. 1, pp. 139-144, 1980.
- [6] M. D. R. Babu, C-H. Lee, and A. Rosenfeld, "Determining plane orientation from specular reflectance," *Pattern Recognition*, vol. 18, no. 1, pp. 53-62, 1985.
- [7] K. Ikeuchi, "Numerical shape from shading and occluding contours in a single view," Artificial Intell. Lab., MIT, Cambridge, MA, AI-Memo 566, 1980.
- [8] K. Ikeuchi, "Determining surface orientations of specular surfaces by using the photometric stereo method," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. PAMI-3, no. 6, pp. 661-669, Nov. 1981.
- [9] E. N. Coleman and R. Jain, "Obtaining 3-dimensional shape of textured and specular surface using four-source photometry," *Comput. Graphics Image Processing*, vol. 18, no. 4, pp. 309-328, Apr. 1982.
- [10] S. K. Nayar and A. C. Sanderson, "Determining surface orientations of specular surfaces by intensity encoded illumination," *Proc. SPIE*, vol. 850, pp. 122-127, Nov. 1987.
- [11] A. C. Sanderson, L. E. Weiss, and S. K. Nayar, "Structured highlight inspection of specular surfaces," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 10, no. 1, pp. 44-55, Jan. 1988.
- [12] S. K. Nayar, L. E. Weiss, D. A. Simon, and A. C. Sanderson, "Structured highlight inspection of specular surfaces using extended Gaussian images," *Proc. SPIE*, vol. 1005, no. 32, Nov. 1988.
- [13] B. K. P. Horn, "Extended Gaussian images," *Proc. IEEE*, vol. 72, no. 12, Dec. 1984.
- [14] K. Ikeuchi, "Recognition of objects using the extended Gaussian image," in *Proc. IJCAI-81*, 1981, pp. 595-600.
- [15] B. K. P. Horn and K. Ikeuchi, "Mechanically manipulating randomly oriented parts," *Sci. Amer.*, vol. 251, no. 2, pp. 100-113, Aug. 1984.
- [16] R. M. Haralick, K. Shanmugam, and I. Dinstein, "Textural features for image classification," *IEEE Trans. Syst., Man, Cybern.*, vol. SMC-3, no. 6, Nov. 1973.
- [17] R. P. Kruger, W. B. Thompson, and A. F. Turner, "Computer diagnosis of pneumoconiosis," *IEEE Trans. Syst., Man, Cybern.*, vol. SMC-4, no. 1, Jan. 1974.
- [18] J. G. Davy, "A comprehensive list of wave solder defects and their probable causes," *Brazing Soldering*, no. 9, pp. 50-59, 1985.
- [19] R. Vanzetti, A. C. Traub, and J. S. Ele, "Hidden solder defects detected by laser infrared system," in *Proc. IPC 24th Ann. Mtg.*, Apr. 1981, pp. 1-15.
- [20] R. Vanzetti and G. Fikiet, "New laser soldering with vision," in *Proc. 13th IEEE Ind. Electron. Soc. Conf.*, Nov. 1987.
- [21] S. L. Bartlett *et al.*, "Automatic solder inspection," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 10, no. 1, pp. 31-43, Jan. 1988.

- [22] P. J. Besl, E. J. Delp, and R. Jain, "Automatic visual solder joint inspection," *IEEE J. Robotics Automat.*, vol. RA-1, pp. 42-56, Mar. 1985.
- [23] P. A. Merrill and M. D. Levine, "Visual inspection of solder joints by computer," McGill Univ. Rep. TR-85-1R, Jan. 1985.
- [24] Y. Nakagawa, "Automatic visual inspection of solder joints on printed circuit boards," *SPIE Robot Vision*, vol. 336, pp. 121-127, 1982.
- [25] S. Timoshenko and J. M. Gere, *Mechanics of Materials*. New York: Van Nostrand Reinhold Co., 1972.



Shree K. Nayar (M'90) received the B.S. degree in electrical engineering from Birla Institute of Technology, Ranchi, India, in 1984, and the M.S. degree in electrical and computer engineering from the North Carolina State University, Raleigh, in 1986.

He is currently a Research Assistant at the Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, and is working towards the Ph.D. degree in electrical and computer engineering. Since 1988, he has been supported for his dissertation work by a Westinghouse research fellowship.

In the summer of 1989, he was a Visiting Research Scientist at the Production Engineering Research Laboratory, Hitachi Ltd., Yokohama, Japan. He has been awarded several patents for his inventions related to machine vision. His research interests include machine vision, robotics, and artificial intelligence.



Arthur C. Sanderson (M'74-SM'86) received the B.S. degree from Brown University, Providence, RI, in 1968, and the M.S. and Ph.D. degrees from Carnegie-Mellon University, Pittsburgh, PA, in 1970 and 1972, respectively.

From 1968 to 1970, he was a Research Engineer at Westinghouse Research Laboratories and worked on design and simulation of solid-state electronic devices. From 1972 to 1973, he was a Visiting Research Fellow at Delft University of Technology, Delft, The Netherlands, and conducted research in

the areas of signal processing and pattern recognition. From 1973 to 1987, he was a faculty member at Carnegie-Mellon University in the Department of Electrical and Computer Engineering. From 1980 to 1987, he was Professor of Electrical and Computer Engineering. He participated in the founding and development of the Robotics Institute at CMU and was Associate Director of the Robotics Institute from 1980 to 1987. As Associate Director, he coordinated many of the research initiatives in the Institute. From 1985 to 1987, he was on leave from CMU and held the position of Director of Information Sciences Research at Philips Laboratories, Briarcliff Manor, NY. As Research Director, he developed new technical programs in artificial intelligence, computer architecture, computer-aided design, and robotics and flexible automation. In 1987, he joined Rensselaer Polytechnic Institute, Troy, NY, as Professor and Department Head of the Electrical, Computer, and Systems Engineering Department. At RPI, he is Co-Director of the New York State Center for Advanced Technology in Automation and Robotics and a Co-Principal Investigator in the RPI/NASA Center for Intelligent Robotics Systems in Space Exploration. He is the author of over 120 technical publica-

tions and proceedings. His current research interests include planning systems for robots and automation systems, sensor-based control, computer vision, and applications of knowledge-based systems.


Dr. Sanderson is President of the IEEE Robotics and Automation Society. He was Associate Editor of the IEEE TRANSACTIONS ON ROBOTICS AND AUTOMATION from 1984-1989, Program Chairman for the 1987 IEEE International Conference on Robotics and Automation, and Chairman of the 1989 IEEE International Symposium on Intelligent Control. He is a member of AAAI, SME, and AAAS.



Lee E. Weiss (S'83-M'83) received the B.S. degree in electrical engineering in 1972 from the University of Pittsburgh, PA, and the M.S. degree in bioengineering and the Ph.D. degree in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, PA, in 1974 and 1984, respectively.

From 1974 to 1976, he was a Research Assistant associated with numerous Pittsburgh hospitals. During this period, he was involved with the development of several prosthetic and implantable devices

and with the signal processing aspects of computerized fetal monitoring systems. From 1976 to 1979, he was a bioengineer at ARCO (Atlantic Richfield) Medical Products Company, where he was responsible for the design and manufacture of cardiac pacemaker heartleads. In 1979, he became the first graduate student to be supported by the Robotics Institute. There, he has designed and implemented several systems, including a flexible assembly cell for electronic PC boards, a series of intelligent robot grippers, an automated solder joint inspection station, an adaptively controlled direct drive robot arm, a robot leg for a Mars Rover autonomous walking machine, and a rapid tool manufacturing system. He is currently a Research Scientist in the Robotics Institute and in the Engineering Design Research Center of Carnegie Mellon University. He is also the Staff Scientist for The Department of Radiology of Shadyside Hospital in Pittsburgh. His current research includes sprayed metal tooling and biomedical applications of micromechanisms. He holds several patents in both the robotics and bioengineering fields and has authored numerous articles on dynamic sensory robot control, learning, automated visual inspection, and tool manufacturing.



David A. Simon (S'85-M'87) received the B.S.E. degree in electrical engineering and computer science from Princeton University, Princeton, NJ, in 1983, the M.S. degree in electrical and computer engineering from Carnegie Mellon University, Pittsburgh, PA, in 1987, and is currently a candidate in the Robotics Ph.D. program at Carnegie Mellon.

From 1983 to 1985, he worked at Siemens Research and Technology Labs in Princeton, NJ, as a Research Associate. From 1987 to 1989, he was a

Research Engineer in the Robotics Institute at Carnegie Mellon. His research interests include robotics and man-machine interfaces.