Local Shape from Specularity

by

Glenn Healey

Thomas O. Binford

Department of Computer Science

Stanford University

Stanford, CA 94305

LOCAL SHAPE FROM SPECULARITY

Glenn Healey and Thomas O. Binford

Artificial Intelligence Laboratory
Computer Science Department
Stanford University
Stanford, California 94305

Abstract

We show that highlights in images of objects with specularly reflecting surfaces provide significant information about the surfaces which generate them. A brief survey is given of specular reflectance models which have been used in computer vision and graphics. For our work, we adopt the Torrance-Sparrow specular model which, unlike most previous models, considers the underlying physics of specular reflection from rough surfaces. From this model we derive powerful relationships between the properties of a specular feature in an image and local properties of the corresponding surface. We show how this analysis can be used for both prediction and interpretation in a vision system A shape from specularity system has been implemented to test our approach. The performance of the system is demonstrated by careful experiments with specularly reflecting objects.

1. Introduction

When light is incident on a surface, some fraction of it is reflected. A perfectly smooth surface reflects light only in the direction such that the angle of incidence equals the angle of reflection. For rougher surfaces, e.g. the surface of a metal fork, specular effects are still observable. In this paper we analyze the properties of specular reflection from rough surfaces.

There are numerous reasons why the study of specular reflection deserves serious attention in computer vision. Specular features are almost always the brightest regions in an image. Contrast is often large across specularities; they are very prominent. In addition, the presence or absence of specular features provides immediate constraints on the positions of the viewer and light sources relative to the specular surface. Also, as we will show, the properties of a specularity constrain the local shape and orientation of the specular surface.

An ability to understand specular features is valuable for any vision system which must interpret images of glossy surfaces. This work, motivated by experience with ACRONYM [4], began in order to provide the SUCCESSOR system with the capability to reason about specular reflection from metal parts in the ITA project [6]. Images of these parts typically contain large specular regions (Figure 1).

Figure 1. Typical Image Containing Specularities

We examine what information can be inferred from an image of a rough surface by considering the physics of specular reflection. Particular emphasis is placed on finding symbolic quasi-invariant relationships which will hold in many different situations (e.g. different source, viewer configurations). In contrast to many intensity-based vision algorithms, we compute a small number of local surface statistics based on the properties of a relatively large number of pixels in an image. This allows us to observe predicted features and infer local surface shape in noisy intensity images and in cases where available specular models do not completely characterize the physics of specular reflection.

2. Review of Previous Work

Researchers in computer graphics have used increasingly realistic specular models. Several of these models will be discussed in the next section. In computer vision, however, relatively few attempts have been made to exploit the information encoded in specularities. Ikeuchi [16] employs the photometric stereo method [24] and uses distributed light sources to determine the orientation of patches on a surface. Grimson [ll] uses Phong's specular model [18] to examine specularities from two views in order to improve the performance of surface interpolation. Coleman and Jain [7] use four-source photometric stereo to identify and correct for specular reflection components. In more recent work, Blake [2] assumes smooth surfaces and single point specularities to derive equations to infer surface shape using specular stereo. He shows that the same equations can be used to predict the appearance of a specularity on a smooth surface when using a distributed light source.

Takai, Kimura, and Sata [22] describe a model-based vision system which recognizes objects by predicting **specular regions. As specular models and insights improve, we expect to see** more work which makes use **of the** properties **of specular reflection.**

3. Specular Reflectance Models

Given a viewer, a surface patch, and a light source, **a reflectance** model quantifies the intensity the viewer will perceive. General reflectance models represent the perceived intensity I as a sum of two reflection components

$$I = I_D + I_S \qquad (1).$$

I_D rcprcscnts the intensity of diffusely reflected light and I_S represents **the intensity** of specularly reflected light. In this paper we restrict our attention to **the** I_S reflection component.

We note that it is typically easy to separate the I_S reflection component from the I_D reflection component in an image. **There** arc several distinctive properties of specular reflection. Over most. of a surface I_S is zero, but in specular regions I_S is usually very large relative to I_D. In regions where the specular component is nonzero, I_S changes much more rapidly with surface geometry than I_D. Furthermore, **the** color of **the** I_S reflection component is almost always different from **the** color of the I_D reflection component.

Before discussing the various specular reflectance models, WC introduce the reflection geometry (Figure 2). We consider a viewer looking at a surface point P which is illuminated by a point light source. Define

V = unit vector from P in direction of viewer
$\vec{N}$ = unit surface normal at P
L = unit vector from P in direction of source
$\vec{H} = \frac{\vec{V}+\vec{L}}{\|\vec{V}+\vec{L}\|}$ (unit angular bisector of $\vec{V}$ and $\vec{L}$)
$\alpha = \cos^{-1}(\vec{N} \cdot \vec{H})$ (the angle between $\vec{N}$ and $\vec{H}$)


Figure 2. The Reflection Geometry

In describing specular models, we consider illumination from a single point light source. In principle, we lose no generality using this approach. In situations involving distributed light sources, we only need to integrate the effects of an equivalent array of point sources. A discussion of the geometry of extended sources is given in [14].

The simplest specular model assumes that specularities only occur where the angle of incidence equals the angle of reflection and $\vec{L}$, $\vec{N}$, and $\vec{V}$ all lie in the same plane. This corresponds to the situation $\alpha = 0$ in Figure 2. Unless the surface is locally flat, this model predicts that specularities will only be observed at isolated points on a surface. A few experiments, however, show that this model is inadequate for most real surfaces. Not only are observed 'specular features usually larger than single points, but highlights often occur in places which are not predicted by this model.

An empirical model for specular reflection has been developed by Phong [18] for computer graphics. This model represents the specular component of reflection by powers of the cosine of the angle between the perfect specular direction and the line of sight. Thus, Phong's model is capable of predicting specularities which extend beyond a single point. While Phong's model gives a reasonable approximation which is useful in some contexts, the parameters of this model have no physical meaning. It is possible to develop more accurate models by examining the physics underlying specular reflection.

The Torrance-Sparrow model [23], developed by physicists, is a more refined model of specular reflection. This model assumes that a surface is composed of small, randomly oriented, mirror-like facets. Only facets with a normal in the direction of $\vec{H}$ contribute to I_S. The model also quantifies the shadowing and masking of facets by adjacent facets using a geometrical attenuation factor. The resulting specular model is

$$I_S = FDA \qquad (2)$$

where

F = Fresnel coefficient
D = facet orientation distribution function
A = geometrical attenuation factor adjusted for foreshortening

We will analyze the effects of each factor in the model in the next few paragraphs. The results we present in this paper are derived from equation (2).

The Fresnel coefficient F models the amount of light which is reflected from individual facets. In general, F depends on the incidence angle and the complex index of refraction of the reflecting material. Cook and Torrance [8] have shown that to synthesize realistic images, F must characterize the color of the specularity. The Fresnel equations predict that F is a nearly constant function of incidence angle for the class of materials with a large extinction coefficient [21]. This class of materials includes all metals and many other materials with a significant specular reflection component.

. The distribution function D describes the orientation of the micro facets relative to the average surface normal $\vec{N}$. Blinn [3] and Cook and Torrance [8] discuss various distribution functions. All of these functions are very similar in shape. In agreement with Torrance and Sparrow we use the Gaussian distribution function given by

$$D = Ke^{-(\alpha/m)^2} \quad (3)$$

where K is a normalization **constant.** Thus, for a **given** α, D is proportional to the **fraction** of facets oriented in the direction $\vec{H}$. The constant m indicates surface **roughness and is** proportional to the standard deviation of the Gaussian. Small values of m describe smooth surfaces for which most of the specular reflection is concentrated in a single direction. Large values of m are used to describe rougher surfaces with larger **differences** in orientation between nearby facets. These rough surfaces produce **specularitics** which appear spread out on the reflecting surface. Figure 3 shows the effect of different values of m.

incident ray   reflected rays

Figure 3a. Specular Distribution for Small m

incident ray   reflected rays

Figure **3b.** Specular Distribution for Large m

The factor A quantifies the effects of **a** geometrical attenuation factor G corrected for foreshortening by dividing by $(\vec{N} \cdot \vec{V})$.

$$A = \frac{G}{\vec{N} \cdot \vec{V}} \qquad (4)$$

G is derived by Torrance and Sparrow in [23]. They assume that each specular facet makes up one side of a symmetric v-groove cavity. From this assumption, they examine the various possible facet configurations which correspond to shadowing or masking. The expression is

$$G = min\left\{1, \frac{2(\vec{N} \cdot \vec{H})(\vec{N} \cdot \vec{V})}{(\vec{V} \cdot \vec{H})}, \frac{2(\vec{N} \cdot \vec{H})(\vec{N} \cdot \vec{L})}{(\vec{V} \cdot \vec{H})}\right\} \quad (5).$$

We will show that in applications+ it is often possible to use a simpler expression for G.

Let μ be the angle between N and $\vec{V}$. As μ increases from 0 to $\frac{\pi}{2}$, the viewer gradually sees a larger part of the reflecting surface in a unit area in the view plane. Therefore, as μ gets larger, there are correspondingly more surface facets which contribute to the intensity perceived by the viewer. We take this phenomenon into account in *(4)* by dividing by $\vec{N} \cdot \vec{V}$.

4. Shape from Specularity

In this section, we demonstrate how we can use (2) to determine local surface properties from specularities. In almost all situations we do not require the full generality of (2) to infer these local properties. Our first assumption is that F is a constant with respect to viewing geometry. This is a very good approximation for metals and for many other materials. We can further simplify (2) byobserving that the exponential factor in (3) changes much faster than any of the terms of A. Therefore, except for a small range of angles near grazing incidence, A can be considered constant across the specularity. We will discuss the consequences of this assumption later. Hence, the form of (2) used to determine local surface properties is

$$I_S = K'e^{-(\alpha/m)^2} \qquad (6),$$

where K' is a constant.

Referring again to the geometry of Figure 2, we assume that the viewer and light source are distant relative to the dimensions of the surface. Therefore $\vec{V}$ and $\vec{L}$ may be regarded as constant; hence their angular bisector $\vec{H}$ is also constant. We assume that the positions of the viewer and light source are known. Finally, since the distance from the viewer to the surface is large, we can approximate the perspective projection of the imaging device with an orthographic projection.

4.1. Inferring Local Surface Shape

For a surface M on which the Gaussian curvature is locally nonzero, we will be able to locate a single point P_0 of maximum intensity in the image of the specularity. From (6) we

see that this point corresponds to the local surface orientation $\vec{N} = \vec{H}$ (i.e. $\alpha = 0$). Given such a surface where $\vec{H}$ is known, we can immediately determine the surface orientation at P_0.

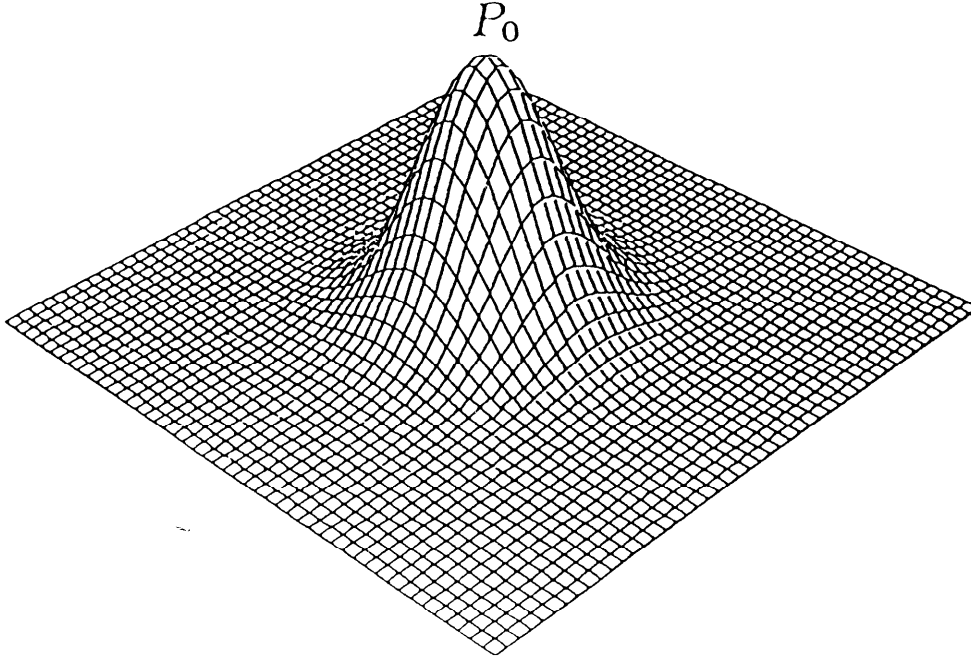Figure **4 shows a typical intensity surface for a specular image.**

$$P_0$$

Figure 4. Specular intensity Surface for a Curved Surface

The level sets are image curves of constant specular intensity. P_0 corresponds to $\alpha = 0$. As predicted by (6), specular intensity decreases as WC move away from P_0.

After locating P_0, we can transform the specular intensity image to the α angle image. Consider equation (6). If I' is the specular intensity corresponding to an arbitrary image point P' near P_0, then the angle α at the surface point imaging to P' is given by

$$|\alpha| = m\sqrt{-\ln\frac{I'}{K'}} \qquad (7)$$

We see that α is determined only up to sign. This will cause the sign of the normal curvature computed at P_0 to be ambiguous. In applications, this ambiguity can usually be resolved by considering other cues. From (7) we can compute the absolute value of α corresponding to each point P' in a neighborhood of P_0. The image of $|\alpha|$ values is called the α angle image.

From the α angle image, WC can compute local curvature properties of the surface. Let $T_{P_0}(M)$ be the tangent space to M at P_0. To compute curvatures, we take a finite number n of straight line samples of the α angle image intersecting P_0. To insure uniform angular resolution on the surface, these samples must be taken in equally spaced directions in

$T_{P_0}(M)$. In general, equally spaced directions in the image will not correspond to equally spaced directions in the tangent space. Thus, given a direction in the tangent space to the surface at P_0, we need to determine the corresponding direction in the image.

Consider a 3-D coordinate system such that the viewer is looking down the z-axis and P_0 is at the origin (Figure 5).


Figure 5. The Projection Geometry

. At P we have $\vec{N} = \vec{H}$ so that the vector $\vec{H}$ is normal to the tangent space to the surface at P. We **choose to define** angles in the tangent space in the counterclockwise sense from $y = 0$ **and in the half space $x \geq 0$.** Denote the normal to the surface at **P** by $\vec{N} = (\vec{N}_1, \vec{N}_2, \vec{N}_3)$. **The tangent space to** the surface at P is given by

$$\vec{N}_1 x + \vec{N}_2 y + \vec{N}_3 z = 0 \qquad (8)$$

Along y=0, the unit vector V_0 in the tangent space is

8

$$V_0 = (\frac{-\vec{N_3}}{\sqrt{\vec{N_3^2} + \vec{N_1^2}}}, 0, \frac{\vec{N_1}}{\sqrt{\vec{N_3^2} + \vec{N_1^2}}}) \quad (9)$$

To simplify the notation let

$$K_1 = \frac{-\vec{N_3}}{\sqrt{\vec{N_3^2} + \vec{N_1^2}}}, \quad K_2 = \frac{\vec{N_1}}{\sqrt{\vec{N_3^2} + \vec{N_1^2}}} \quad (10)$$

Let θ be the angle of interest in the tangent space. The goal is to find a unit vector $\vec{V_1} = (\vec{x_1}, \vec{y_1}, \vec{z_1})$ which lies in the tangent space and makes an angle θ with V_0. The angle θ provides the constraint

$$x_1 K_1 + z_1 K_2 = cos\theta \quad (11),$$

From (8) we must require

$$x_1 \vec{N_1} + y_1 \vec{N_2} + z_1 \vec{N_3} = 0 \quad (12),$$

and since V_1 is a unit vector we have

$$x_1^2 + y_1^2 + z_1^2 = 1 \quad (13).$$

The equations (11), (12), (13) may be solved uniquely for x_1, y_1, z_1 in the half space x $\geq$ 0. Briefly, the solution is

$$z_1 = \frac{-R_2 + \sqrt{R_2^2 - 4R_1 R_3}}{2R_1} \quad (14)$$

$$x_1 = C_4 - C_5 z_1 \quad (15)$$

$$y_1 = \frac{-x_1 \vec{N_1}}{\vec{N_2}} - \frac{z_1 \vec{N_3}}{\vec{N_2}} \quad (16)$$

where

$$R_1 = C_1 C_5^2 + C_2 - C_3 C_5 \quad (17a)$$

$$R_2 = -2C_1 C_4 C_5 + C_3 C_4 \quad (17b)$$

$$R_3 = C_1 C_4^2 - 1 \qquad (17c)$$

$$C_1 = 1 + \frac{\vec{N}_1^2}{\vec{N}_2^2} \qquad (18a)$$

$$C_2 = 1 + \frac{\vec{N}_3^2}{\vec{N}_2^2} \qquad (18b)$$

$$C_3 = \frac{2\vec{N}_1 \vec{N}_3}{\vec{N}_2^2} \qquad (18c)$$

$$C_4 = \frac{cos\theta}{K_1} \qquad (18d)$$

$$C_5 = \frac{K_2}{K_1} \qquad (18e)$$

A special case occurs when $\vec{N}_2 = 0$. For this case we use (11) - (13) to arrive at

$$z_1 = \frac{-\vec{N}_1 cos\theta}{K_1 \vec{N}_3 - K_2 \vec{N}_2} \qquad (19)$$

$$x_1 = \frac{cos\theta - z_1 K_2}{K_1} \qquad (20)$$

$$y_1 = \sqrt{(1 - z_1^2 - x_1^2)} \qquad (21)$$

We use (14) - (21) to compute the components of V_1.

Assuming an orthographic projection and examining the geometry of Figure 5, we see that the x_1, y_1 components of V_1 give us the projected vector we are seeking in the image. Let θ' be the image angle corresponding to V_1. Then

$$\theta' = tan^{-1}\left(\frac{y_1}{x_1}\right) \qquad (22)$$

The next step is to use the α image to compute the normal curvature of the surface in the direction θ. The normal curvature is computed by taking a straight line L in the α image which intersects P_0 and is in the direction 8'. Under orthographic projection, L will project to a line L' in $T_{P_0}(M)$. The goa lis to compute the normal curvature of the curve C $\subset$ M where C is the orthogonal projection of L' onto M. Since C projects to a

line in $T_{P_0}(M)$, the magnitude of the geodesic curvature $|K_g|$ of C is 0 [17]. Thus local changes in α along C are due primarily to normal curvature along C. We compute the normal curvature κ_n in the direction θ by

$$\kappa_n = \left.\frac{da}{ds}\right|_{P_0} \quad (23)$$

where s is arc length in the direction θ. In other words, we are differentiating the a image along L with respect to arc length on the surface. From the local character of specularities, we see that, to a very good approximation, arc length on the surface is equal to length in $T_{P_0}(M)$. Therefore, length in the image and arc length on the surface are related by the scale factor $\sqrt{x_1^2 + y_1^2}$. Thus, computing normal curvature on the surface has been reduced to differentiation in the α image.

If we let θ vary in the range $0 \leq \theta \leq 2\pi$ we can compute κ_n in any number of directions at P_0. The principal curvatures of M at P_0 are defined to be the maximum and minimum values of κ_n; the corresponding directions are called the principal directions. Hence, using this technique it is possible to describe M locally to second order in terms of principal curvatures and principal directions. In the context of shape from shading [13], Bruss[5] and Deift and Sylvester [9] examine the assumptions required to generate higher order surface descriptions from an α image.

4.2. Special Cases

In this subsection we examine specular reflection from special classes of surfaces. In 4.2.1. and 4.2.2. we consider surfaces which are locally singly curved and planar respectively. For these surfaces, the Gaussian curvature is locally zero. In 4.2.3. we examine the case of corners and edges where surface normal is discontinuous but where specularities are frequently observed.

4.2.1. Singly Curved Surfaces

If one principal curvature of a surface is zero in a specular region (i.e. the surface is locally singly curved), we will not be able to infer immediately the local orientation as we did for a doubly curved surface. To understand why, consider Figure 6. Figure 6 shows a viewer looking at a tilted cylinder. To make the example concrete, assume that $\vec{L}$ is such that $\vec{H} = \vec{V}$. For this configuration there will be no point on the surface for which $\alpha = 0$ (recall that $\vec{H}$ is essentially constant), yet we will still observe a specularity in the image if at some point α is small enough to give a significant value for I_S in (6). Define ϕ to be the smallest value of a for a given surface-source-viewer configuration. Figure 7 shows a specularity generated by a cylinder which is oriented so that ϕ is 20^0. Note that a specular model which assumes a smooth surface would not predict a specularity for this case.

11

Figure 6. Viewer Observing a Singly Curved Surface

Figure 7. Specularity for a Tilted Cylinder

We observe that it is typically easy to detect that a surface is singly curved at a specularity. This is because we will observe *a* line of maximum intensity (along the line of zero curvature) instead **of the point maximum we observe for the doubly curved case.**

Figure **8 is a plot of I_S for a singly curved surface in a direction perpendicular to the** lines of zero curvature as we change ϕ. It is worth noting that both the magnitude and shape of I_S change as ϕ increases. Consequently, it is possible to recover significant local shape information for this class of surfaces.

Figure 8. I_S for different values of ϕ

4.2.2. Planes

For a planar surface, N is constant. Hence, recalling our basic assumptions, I_S is constant across a plane. If the plane is oriented such that α is small enough, then a viewer will perceive an I_S reflection component. As with the singly curved surface, the magnitude of the perceived intensity will depend on α. If α is not sufficiently small, then I_S will be zero at all points on the plane. These observations provide us with two useful pieces of information:

1. Glossy surfaces which don't generate specularities over a range of orientations are probably planar.
2. Surfaces which produce a specularity of constant intensity over a 2-D region in the image are locally planar.

13

4.2.3. Corners and Edges

Specularities are often observed at places of discontinuous surface normal on an object. Typical examples of these discontinuities are edges and corners on a polyhedron. For an ideal edge on a polyhedron, the surface normal is discontinuous across the edge (Figure 9).


Figure 9. An Ideal Edge

For an ideal edge joining two planes, **we should** not expect to **observe a specularity** unless either $\vec{N}_1$ or $\vec{N}_2$ is oriented in a direction which is sufficiently close to the perfect specular direction $\vec{H}$. But for this case, as discussed in 4.2.2., we would expect to observe a spread out specular feature on one of the two planes joined by the edge. So why do we frequently see specular refl ections along edges? On real polyhedra, **surface normals are** usually continuous across edges. Instead of the normal vector **changing discontinuously,** the normal usually changes smoothly from $\vec{N}_1$ to $\vec{N}_2$ by taking **values which are linear** combinations of $\vec{N}_1$ and $\vec{N}_2$. As WC cross **an** edge, the surface normal moves rapidly through a large range of angles. If any of these normals is oriented in a direction **sufficiently near** the direction $\vec{H}$, WC will observe a specularity. Therefore, WC often observe a specularity on an edge. Figure 10 shows an image of an edge specularity.

Figure 10. Image of an Edge Specularity

The situation is similar for trihedral vertices. As with edges, the normal vector is usually continuous at a corner. For trihedral vertices, the normal vector typically takes on values which are linear combinations of the three normals corresponding to the three planes defining the vertex.

From experiments with polyhedra, we have developed a useful model for the behavior of surface normal across edges. Define r to be the edge sharpness parameter and assume the coordinate system of Figure 11. P_1 and P_2 denote two planes intersecting to form an edge.

Figure 11. The Geometry of a Rounded Edge

15

The y axis is aligned in the direction of $\vec{N_1}$ and the **origin is** a distance r from P_1 such that the normal to the surface P_1 begins to **turn away** from $\vec{N_1}$ when x becomes positive. The model for the normal as it turns from $\vec{N_1}$ to $\vec{N_2}$ is

$$\vec{N} = \frac{\sqrt{r^2 - x^2}}{r}\vec{N_1} + \frac{x}{r}\vec{N_2} \text{ for } 0 \leq x \leq r|\vec{N_1} \times \vec{N_2}| \quad (24)$$

In other words, the normal is assumed to turn through a curve of constant curvature $\frac{1}{r}$. Here the parameter r is used to specify the sharpness of the edge. Small values of **r** indicate sharp edges, while larger values of r indicate more rounded edges. Figure **12** shows the profile of a specularity across a sharp **90⁰** edge which is similar to the profile predicted by the continuous normal variation of (24).
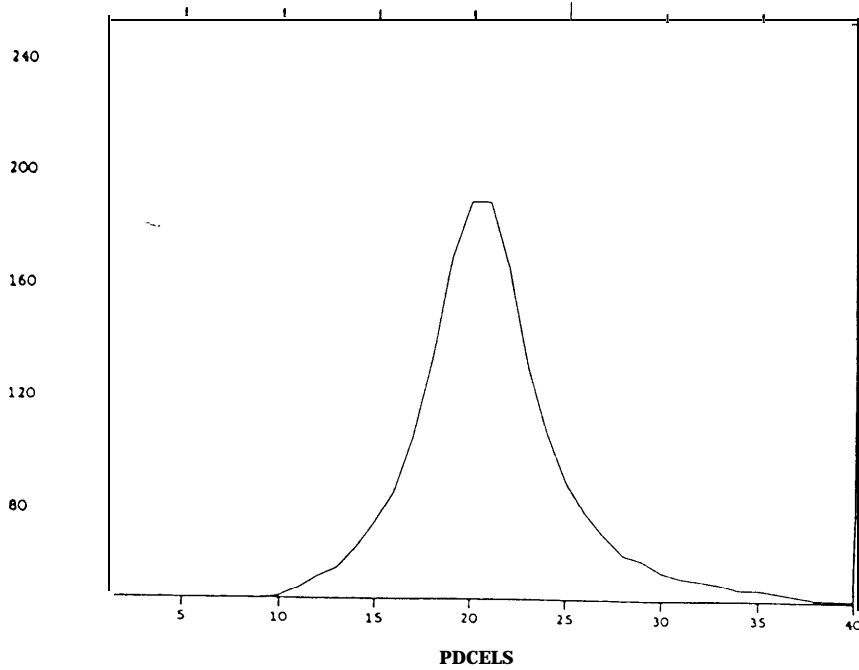

Figure 12. Specular Intensity Profile Across an Edge

4.3: Predicting A

In the previous analysis we have assumed that over most configurations of viewer, source, and surface the adjusted geometrical attenuation factor A of (4) will have a small constant value across the specularity. For large angles of incidence, however, the character of A changes remarkably (Figure 13). In particular, for large angles of incidence (glancing incidence) we see that

1. $\dot{A}$ becomes large relative to its value for other incidence angles (Figure 13).
2. A causes a shift in the peak of the specular profile toward larger angles of incidence.
3. A causes the specular profile to be unsymmetric as a function of α.

16

It **is** not surprising that when these effects are present in an image, they are rather easy to detect. For this reason, it is useful to make qualitative predictions about A in applications where large angles of incidence are possible.
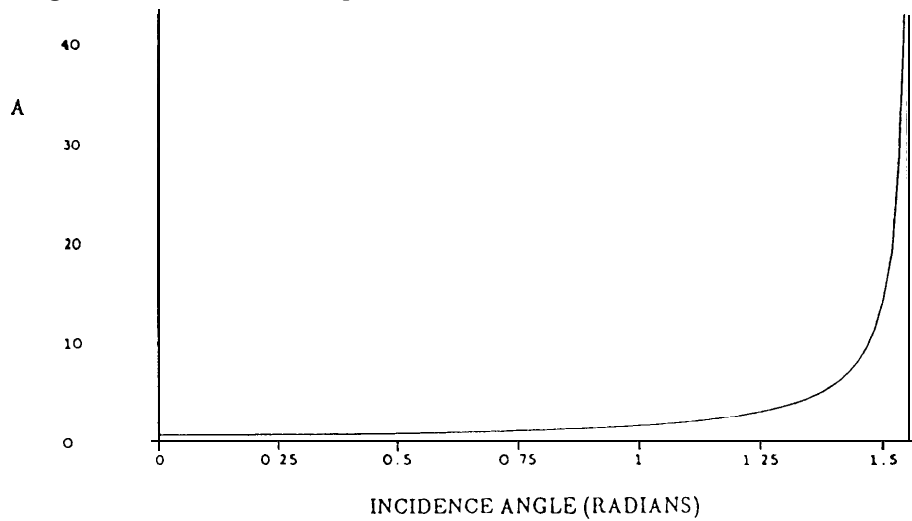

Figure 13. Plot of A as function of incidence angle

5. The Laboratory Setup

A laboratory arrangement has been set up to test the derived relationships (Figure 14). This section of the paper describes the laboratory **setup.** Section 6 examines factors which must be considered to successfully interpret real images. In Section 7, we describe an implemented system which has been used to infer local surface properties from specularities. Section 8 presents experimental results.


Figure 14. The Laboratory Setup

To insure accurate measurements, the experiments are conducted on a 4x6 foot optical table. High precision rotation and translation stages are used to position the objects being viewed. A halogen light source with a 5 mm wide filament is placed 20 feet from the object surface to approximate a point source. Monochromatic image data is obtained using a video camera and an image digitizer. A 210 mm lens is used with the video camera to obtain high resolution across the specularity. The resulting images are in the form of 256x256 arrays of pixels. Each pixel has eight bits of gray level resolution. A precise positioning device has been built to position the camera relative to the surface. Camera-object distances of at least 24 inches are enforced to insure that the assumed distant object condition is met. Using this setup, it is possible to obtain more than 40 pixels across a specular feature which is less than a centimeter wide on the surface. Metal cylinders and spheres of varying curvature are used to test the predicted relationships (Figure 15).


Figure 15. Some Experimental Specular Surfaces

6. Practical Considerations

. This section examines factors which must be considered to enable a shape from specularity system to successfully interpret real images.

6.1. Gaussian Blur

Unfortunately, the formation of an image by an optical system introduces some amount of degradation. We can model this degradation as a convolution with a spatially invariant Gaussian point spread function [I]. The standard deviation of this blur is typically less than one pixel. For small specular features, taking into account the effects of this blur allows a more accurate determination of surface shape.

Our system uses a module called BLURINVERT to deblur the input specular image. For general 'L-dimensional functions, inverting Gaussian blur is an unstable process. However, an explicit deblurring convolution kernel has been derived under certain assumptions in [15]. The 1-D continuous version of the kernel is given by

$$M_N(x) = \frac{e^{-x^2}}{\sqrt{\pi}k!2^k} \sum_{k=0}^{(N-1)/2} (-1)^k H_{2k}(x) \quad (25)$$

where N is an odd integer denoting the order of the kernel and H_{2k} is the Hermite polynomial of degree 2k. Larger values of N give better deblurring filters (i.e. they recover exactly a larger space of blurred input functions), but are more costly to compute. The value of N that is chosen in applications depends on the intensity characteristics of the images that will be processed by the system. Using a 2-dimensional discrete version of (25), the BLURINVERT module allows our system to produce accurate shape descriptions from small specular features in images.

6.2. Quantization Effects

On a surface of high curvature, it is unlikely that we will measure the correct maximum specular intensity K' in (6). The problem is that for highly curved surfaces we are unable to shrink a pixel down to where the surface area it images is approximately planar. Even within the single pixel of maximum intensity, α is changing and cannot be considered constant. Hence the intensity value at the maximum pixel will be an average specular intensity over a small range of α and will not give the true K' of (6). This must be corrected for in applications. An artifact of this phenomenon is that measured K' seems to increase as surface curvature decreases. It follows that if we wish to measure K' for a material, we should use a surface of small curvature, ideally a plane.

Since specularities are usually the brightest features in images, specular intensities are often too large to be represented in the number of bits per pixel allowed by the digitizing hardware or within the dynamic range of the camera. If this is the case, the specularity is truncated. Figure 16 shows I for a truncated specularity. The obvious way to deal with this situation is to avoid it. One avoidance technique is to take multiple images in which differing amounts of light are allowed to reach the camera. This can be achieved either by adjusting the lens aperture or by using filters. Another possible solution is to control the illumination to eliminate the possibility of truncation.

Figure 16. A Truncated Specularity

If inferences must be made from a single image, then it is arguably better **to allow** truncation to occur. In the case where input images have eight bits per pixel, intensities will range from 0 to 255. In many applications it is possible to weaken the incident illumination so that no truncation occurs. In doing this, however, we cause pixels on the I_S curve which previously had significant specular intensities (on the truncated specular feature) to have negligible specular intensities. **The** net effect of eliminating truncation is to **decrease the** width of the specular feature and make measurements more susceptible **to small** errors.

7. **A Shape from Specularity System**

A system has been implemented which computes local surface properties from images of specular surfaces [12]. The system currently stands alone, but will be used in the more general context of the SUCCESSOR vision system. **The** shape from specularity system is primarily designed to perform the computations described **in Section 4**. This section describes the implementation of the system.

7.1. **Overview of System Structure**

` At a high level of abstraction, the problem is best solved in two steps. The first step is **to deblur** the **input** specular intensity image. The second step is to compute local surface properties from the image resulting from step 1.

20

The system designed to solve the problem preserves this two step structure. Figure 17 is a diagram of the modules in our system with arcs indicating module interactions. From this diagram we see that there is a clean separation between the deblurring task and the task of computing local surface properties. First the main program invokes a function called **BLURINVERT** to deblur the input image. After the deblurring task is completed, the function CURVATURES is **called** to compute local surface properties. **The** next two subsections give overviews of the BLURINVERT and CURVATURES functions.
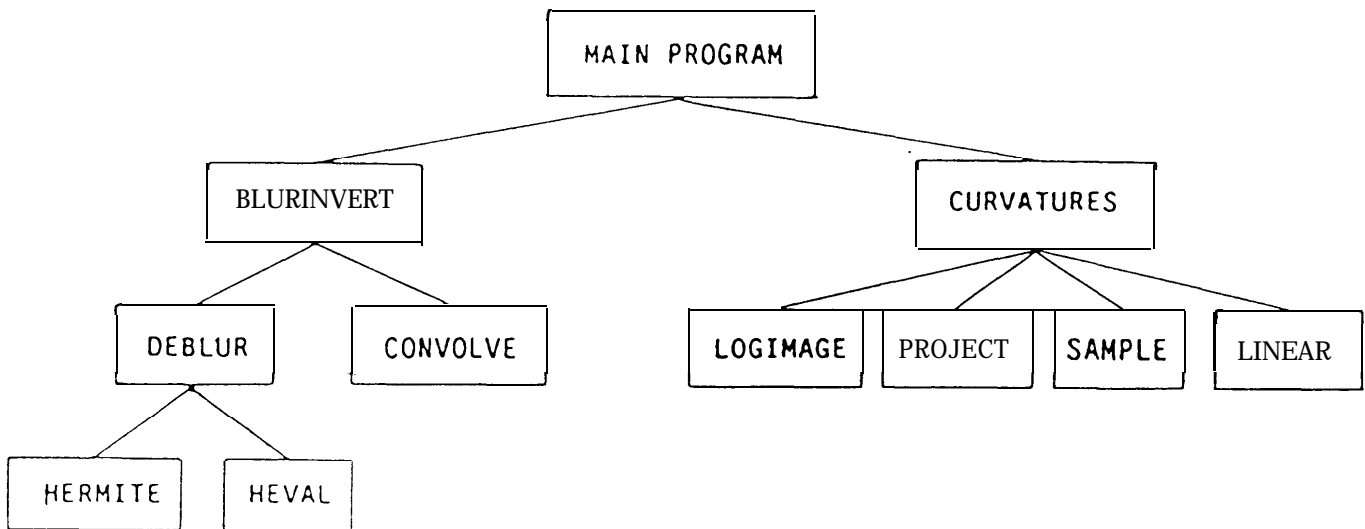
Figure 17. System Structure

7.2. Overview of BLURINVERT

The BLURINVERT function is used to deblur the Gaussian blurred input image. This is accomplished in two stages. First, the deblurring convolution kernel is generated by DEBLUR. Then the CONVOLVE function is called to perform the convolution of the blurred input image with the constructed deblurring kernel. Two functions are used by DEBLUR to manipulate the Hermite polynomials required to generate the deblurring filter. The function HERMITE uses a dynamic programming scheme to compute coefficients of all Hermite polynomials up to some specified degree. The function HEVAL is used to evaluate Hermite polynomials at fixed values of the polynomial parameter.

7.3. Overview of **CURVATURES**

Given a deblurred specular intensity image, the CURVATURES **function** computes the principal curvatures and directions of the surface at specular points. CURVATURES first uses the function LOGIMAGE to transform the intensity image into the α angle image of Section 4. The CURVATURES function then systematically computes 1-D curvature at different directions in the tangent space to the surface. The function PROJECT is used

to compute the metric transform between the tangent space to the surface and the image. This is necessary to insure that the system samples the angle image at equidistant angles on the surface. The function SAMPLE is used to sample the α angle image in a specified direction. Finally, the function LINEAR is used to compute the least squares curvature given the data generated by PROJECT and SAMPLE.

8. Experimental Results

The system described in Section 7 has consistently generated accurate surface descriptions from images of specular surfaces. In this section we give examples of our system's performance on real images of metal objects illuminated by a point source. Figures 18(a), 19(a), and 20(a) are images of circular cylinders of varying radii. Each cylinder is oriented such that its axis is perpendicular to the axis of the imaging device. Figure 21(a) is the image of a sphere. The actual statistics of the surfaces are given in Table 1. The dotted lines in the images indicate the direction of maximum curvature as determined by the system. Figures 18(b), 19(b), 20(b), and 21(b) are plots of intensity along the dotted lines in 18(a), 19(a), 20(a), and 21(a). Note that in Figure 21 the specularity is truncated, but we are still able to compute accurate surface statistics. Table 2 gives the second order surface statistics computed by the system. Error represents the percent of error in the computation of the largest curvature of the surface. The small errors can be attributed to quantization effects, noise introduced during the measurement process; and the various simplifications made to the specular model.

9. Summary and Implications

Understanding specular reflection is important for any vision system which must interpret a world containing glossy objects. Using a model developed by optics researchers, we have shown that we can accurately predict the appearance of specular features in an image. In addition we have shown how to compute the local orientation and principal curvatures and principal directions of a specular surface by examining image intensities on a specularity. These statistics give a complete local characterization of the surface up to second order. Unlike previous work, our derivations have included the effects of surface roughness and microstructure on the appearance of specular features.

A system has been implemented which computes local surface properties from images of specular objects. A laboratory setup has been arranged which allows us to capture images to test our system. The system has consistently produced accurate surface descriptions despite the fact that the high intensity and small spatial extent of specularities makes measurements difficult. Significant aspects of the implementation are discussed in Section 7. Examples of experimental results are given in Section 8.

The ability to predict intensity-based features such as specularities opens up interesting possibilities for model-based vision. Previous model-based vision systems have restricted their predictions to the structure of edges which will be observed for a given model. An ability to predict intensity-based features will significantly enhance the top-down capabilities of a model-based vision system. Clearly it is advantageous to be able to make stronger predictions about an image by using additional information about the imaging

Figure 18(a)

Figure 18(b)

Figure 19(a)

Figure 19(b)

Figure 20(a)


PIXELS

Figure 20(b)

Figure 21(a)

Figure 21(b)

26

Object	κ_1	θ_1	κ_2	θ_2
cylinder$_1$	0.286	0.0	0.0	1.571
cylinder$_2$	0.400	o-o	0.0	1.571
cylinder$_3$	1.333	o-o	0.0	1.571
sphere$_1$	0500	—	0.500	—

Table **1.** Actual Surface Statistics

Object	κ_1	θ_1	κ_2	θ_2	Error
cylinder$_1$	0.297	0.0	0.00 1	1.571	3.9%
cylinder$_2$	0.397	0.0	0.00 1	1.571	0.9%
cylinder$_3$	1.356	0.0	0.002	1.571	1.7%
sphere$_1$	0.514	0.0	0.534	1.571	2.8%

Table 2. Computed Surface Statistics

process. Another important advantage of predicting intensity-based features is that this prediction can provide strong guidance to low level intensity-based visual processes. By making predictions about the appearance of intensity patches in an irnage we can hope to further unify the goals of the low level and high level mechanisms of a model-based vision system.

More important than being able to predict the appearance of specularities from surface models is our system's ability to invert the process. We have shown how to inier local second order surface shape from specular images. This capability provides a vision system with strong generic information about a surface in a scene using strictly bottom-up proccssi ng. Inferring shape information from specularity is particularly important when viewing metal surfaces because other shape cues such as shading and texture are often not present. For other kinds of surfaces, shape information from specularity can be combined with shape information obtained using other cues to improve the 3-D surface descriptions generated by a vision system.

Acknowledgements

References

[1] Andrews, H. and Hunt, B., "Digital Image Restoration," Prentice-Hall Inc., Englewood Cliffs, 1977.

[2] Blake A., "Specular Stereo," *Proceedings of IJCAI-9* (Los Angeles: August 1985), 973-976.

[3] Blinn, J., "Models of Light Reflection for Computer Synthesized Pictures," *Computer Graphics,* 11(2) (1977), 192-198.

[4] Brooks, R., "Symbolic Reasoning Among 3-D Models and 2-D Images," *Artificial Intelligence,* 17 (1981), 285-348.

[5] Bruss, A., "The Image Irradiance Equation: Its Solution and Application," MIT AI Memo 623, June 1981.

[6] Chelberg, D. and Lim, H. and Cowan, C., "ACRONYM Model-based Vision in the Intelligent Task Automation Project ," *Proceedings of Image Understanding Workshop* (1984).

[7] Coleman, E.N. Jr. and Jain, R., "Obtaining 3-D Shape of Textured and Specular Surfaces Using Four-Source Photometry," *Computer Graphics and Image Processing,* 18 (1982)) 309-328.

[8] Cook, R. and Torrance, K., "A Reflectance Model for Computer Graphics", *Computer Graphics,* 15(3) (1981), 307-316.

[9] Deift, P., and Sylvester, J., "Some Remarks on the Shape-from-Shading Problem in Computer Vision," *Journal of Mathematical Analysis and Applications,* Vol. 84, No. 1, pp.235-248, November, 1981.

[10] Foley, J. and Van Dam, A., *Fundamentals of Interactive Computer Graphics,* Addison-Wesley, 1982.

[ll] Grimson, W.E.L., "Binocular Shading and Visual Surface Reconstruction," MIT AI Memo 697 (1982).

[12] Healey, G., "Solving the Inverse Specular Problem", Stanford Computer Science Department Programming Project, March 1986.

[13] Horn, B., "Obtaining Shape from Shading Information", in "The Psychology of Computer Vision" (P. Winston, Ed.), McGraw-Hill, New York, 1975.

[14] Horn, B., *Robot Vision,* MIT Press, 1986.

[15] Hummel, R. and Kimia, B. and Zucker, S., "Gaussian Blur and the Heat Equation: Forward and Inverse Solutions," *Proceedings of CVPR* (San Francisco: June, 1985).

[16] Ikeuchi, K., "Determining Surface Orientations of Specular Surfaces by Using the Photometric Stereo Method", *IEEE PAMI* 3(6) (1981), 661-669.

[17] Kreyszig, E., *Introduction to Differential Geometry and Riemannian Geometry,* University of Toronto Press, 1968.

[18] Phong, B., "Illumination for Computer Generated Pictures," **Communications** *of* **the ACM** 18 (1975), 311-317.

[19] Shafer, S., "Optical Phenomena in Computer Vision," Univ. of Rochester TR 135 (1984).

[20] Spivak, M., Differential Geometry - **Volume II,** Publish or Perish, Inc., 1979.

[21] Sparrow, E. and Cess, R., **Radiation Heat Transfer,** McGraw-Hill, New York, 1978.

[22] Takai, K. and Kimura, F. and Sata, T., "A Fast Visual Recognition System of Mechanical Parts by Use of Three Dimensional Model," source unknown. (First author is with CANON INC. in Tokyo.)

[23] Torrance, K. and Sparrow, E., "Theory for Off-Specular Reflection from Roughened Surfaces ," **Journal** *of* **the Optical Society** *of America,* 57 (1967), 1105-1114.

[24] Woodham, R., "Photometric Stereo: A Reflectance Map Technique for Determining Surface Orientation from Image Intensity," Proc. SPIE, vol. 155 (1978).