**Surface Curvature from Photometric Stereo**

by

R.J. Woodham[1]

Technical Report 90-29

October 1990

Laboratory for Computational Vision
Department of Computer Science
University of British Columbia
Vancouver, B.C., CANADA

email: woodham@cs.ubc.ca

**Abstract**

A method is described to compute the curvature at each point on a visible surface. The idea is to use the intensity values recorded from multiple images obtained from the same viewpoint but under different conditions of illumination. This is the idea of photometric stereo. Previously, photometric stereo has been used to obtain local estimates of surface orientation. Here, an extension to photometric stereo is described in which the spatial derivatives of the intensity values are used to determine the principal curvatures, and associated directions, at each point on a visible surface. The result shows that it is possible to obtain reliable local estimates of both surface orientation and surface curvature without making global smoothness assumptions or requiring prior image segmentation.

The method is demonstrated using images of several pottery vases. No prior assumption is made about the reflectance characteristics of the objects to be analyzed. Instead, one object of known shape, a solid of revolution, is used for calibration purposes.

1. Fellow of the Canadian Institute for Advanced Research.

# 1 Introduction

The purpose of computational vision is to produce descriptions of a 3D world from 2D images of that world, sufficient to carry out a specified task. Computer vision systems are required to perform at least three generic tasks: 1) recognition 2) localization and 3) inspection. Research emphasis typically is placed on recognition tasks. But, increasingly applications require careful attention to localization and to inspection. Localization is the determination of the 3D position and attitude of a known object or surface. Inspection is the detailed monitoring of known surfaces for defects or changes. Robustness on any task is improved when computer vision systems make use of all the information available in an image, not just that obtained from a sparse set of features.

This paper demonstrates that reliable local estimates of the curvature at each point on a visible surface can be determined using photometric stereo. These estimates are obtained prior to image segmentation and are not based on any global measure of surface smoothness. Photometric stereo itself is not new. It was first described by Woodham in his thesis [1] and in subsequent journal articles [2, 3]. Photometric stereo determines a dense representation of the orientation at each visible point on a surface. First implemented by Silver [4], photometric stereo has since been used by Horn, Ikeuchi and colleagues both for recognition tasks [5] and for localization tasks [6, 7, 8, 9, 10].

The overall approach is based on principles of physical optics. An image irradiance equation is developed to determine image irradiance as a function of surface orientation. This equation cannot be inverted locally since image brightness provides only one measurement while surface orientation has two degrees of freedom. Photometric stereo allows for the local determination of surface orientation by using multiple images, obtained with the identical geometry but under different conditions of illumination. Three (or more) images overconstrain the solution. Computations of local surface orientation can be fast, accurate and robust. Once surface orientation is known, surface curvature also is determined locally from the partial derivatives of image irradiance. Local curvature estimates also are robust since, once again, the solution is overconstrained.

Physical optics determines that an image irradiance equation exists but says very little about the particular form that image irradiance equation must take. Of course, one can exploit situations where the reflectance properties of a material are known to be of a particular functional form. Formal analysis of these situations helps to establish the existence, uniqueness and robustness of solution methods under varying degrees of uncertainty and approximation. Implementation also is facilitated because the resulting computations typically involve equations of known form with unknown coefficients that can be determined as a problem of parameter estimation.

Here, no prior assumption is made about the reflectance characteristics of the objects to be analyzed. Instead, reflectance properties are measured using a calibration object of known shape, in this case a solid of revolution. Measurements from a calibration object of known shape are directly applicable to the analysis of other objects of unknown shape but made of the same material and illuminated and viewed under the same conditions. In principle, materials with any reflectance properties can be handled.

Section 2 provides the background and theory. Section 3 discusses a particular implementation and reports on the experiments performed. Finally, Section 4 provides a brief discussion and summary of the conclusions following from the work reported.

## 2 Background and Theory

A given spatial arrangement of objects, made of a given set of materials, illuminated in a given way, and viewed from a given vantage point, determine an image according to the laws of physical optics. Geometric equations determine where each point on a visible surface appears in the image and corresponding radiometric equations determine its brightness and color.

The standard geometry of shape from shading is assumed. That is, let the object surface be given explicitly by $z = f(x, y)$ in a left-handed Euclidean coordinate system, where the viewer is looking in the positive Z direction, image projection is orthographic, and the image XY axes coincide with the object XY axes. The *gradient*, $(p, q)$, is defined by

$$p = \frac{\partial f(x,y)}{\partial x} \quad \text{and} \quad q = \frac{\partial f(x,y)}{\partial y} \tag{1}$$

so that a surface normal vector is $[p, q, -1]$. An image irradiance equation can be written as

$$E(x,y) = R(p,q) \tag{2}$$

where $E(x, y)$ is the image irradiance and $R(p, q)$ is called the *reflectance map*. A reflectance map combines information about surface material, scene illumination and viewing geometry into a single representation that determines image brightness as a function of surface orientation.

Given an image, $E(x, y)$, and the corresponding reflectance map, $R(p, q)$, the problem of shape from shading is to determine a smooth surface, $z = f(x, y)$, that satisfies the image irradiance equation over some domain $\Omega$, including any initial conditions that may be specified on the boundary $\partial\Omega$ or elsewhere. There is a substantial, but scattered, literature on shape from shading. Two essential references are Horn's text [11] and the collection of papers edited by Horn and Brooks [12]. With a single image, shape from shading problems typically are solved by exploiting *a priori* constraints on the reflectance map, $R(p, q)$, *a priori* constraints on surface curvature, or global smoothness constraints. Photometric stereo, on the other hand, makes use of additional images.

### 2.1 Photometric Stereo

Photometric stereo uses multiple images obtained under the identical geometry but under different conditions of illumination. For example, two image irradiance equations

$$\begin{aligned} E_1(x,y) &= R_1(p,q) \\ E_2(x,y) &= R_2(p,q) \end{aligned} \tag{3}$$

provide two equations in the two unknowns, $p$ and $q$. But, these equations are, in general, nonlinear so that the solution typically is not unique. Three image irradiance equations

$$\begin{aligned} E_1(x,y) &= R_1(p,q) \\ E_2(x,y) &= R_2(p,q) \\ E_3(x,y) &= R_3(p,q) \end{aligned} \tag{4}$$

in general overconstrain the solution at each point, $(x, y)$, because three intensity measurements, $E_1(x, y)$, $E_2(x, y)$, and $E_3(x, y)$, are used to estimate two unknowns, $p$ and $q$.

Conceptually, the idea of photometric stereo is straightforward. Using a calibration object of known shape, one can build a lookup table mapping triples of measured brightness values, $[E_1, E_2, E_3]$, to the corresponding gradient, $(p, q)$. If each image is accurate to $2^8 = 256$ gray values then the full table would have $2^8 \times 2^8 \times 2^8 = 2^{24}$ entries. Literally building a table of this size may be prohibitive, in terms of memory. Various table compression techniques are possible. Indeed, one would expect the resulting table to be sparse since Equation (4) defines the parametric equations, in parameters $p$ and $q$, of a surface in $E_1, E_2, E_3$ space.

In any event, computation of the solution can be fast, accurate and robust, as various implementations have demonstrated. But, even though the computation of surface orientation can be accurate and robust, there can still be problems in locally differentiating surface orientation to obtain surface curvature. This would certainly be the case, for example, if the gradients represented were quantized into too small a set of possibilities.

Therefore, it is useful to examine what more information about surface curvature can be extracted from the image irradiance equation. Before doing this, it first is necessary to define a representation for surface curvature.

## 2.2 Surface Curvature

There are three degrees of freedom to the curvature at a point on a smooth surface. Consequently, three parameters are required to specify curvature. One representation is in terms of the $2 \times 2$ matrix of second partial derivatives of the surface $z = f(x, y)$. Let $\mathbf{H}$ be the matrix,

$$\mathbf{H} = \begin{bmatrix} \frac{\partial^2 f(x,y)}{\partial x^2} & \frac{\partial^2 f(x,y)}{\partial x \partial y} \\ \frac{\partial^2 f(x,y)}{\partial y \partial x} & \frac{\partial^2 f(x,y)}{\partial^2 y} \end{bmatrix} \tag{5}$$

$\mathbf{H}$ is called the *Hessian matrix* of $z = f(x, y)$. For notational convenience, let

$$p_x = \frac{\partial^2 f(x,y)}{\partial x^2}, \quad p_y = \frac{\partial^2 f(x,y)}{\partial x \partial y}, \quad q_x = \frac{\partial^2 f(x,y)}{\partial y \partial x} \quad \text{and} \quad q_y = \frac{\partial^2 f(x,y)}{\partial y^2} \tag{6}$$

It may appear that four parameters are required to specify **H**. But, for smooth surfaces, **H** is symmetric. That is, $p_y = q_x$. Therefore, only three parameters are required after all. **H** is a viewer-centered representation of surface curvature because its definition depends on the explicit form of the surface function, $z = f(x,y)$, and on the fact that the viewer is looking in the positive Z direction.

From the Hessian, **H**, and the gradient, $(p,q)$, one can determine a viewpoint invariant representation of surface curvature. Let **C** be the matrix,

$$\mathbf{C} = (1+p^2+q^2)^{-\frac{3}{2}} \begin{bmatrix} q^2+1 & -pq \\ -pq & p^2+1 \end{bmatrix} \mathbf{H} \tag{7}$$

Further, let $k_1$ and $k_2$ be the two eigenvalues of **C**, with associated eigenvectors $\omega_1$ and $\omega_2$. Then, $k_1$ and $k_2$ are the principal curvatures, with directions $\omega_1$ and $\omega_2$, at $z = f(x,y)$. The principal curvatures, $k_1$ and $k_2$, are viewpoint invariant surface properties since they do not depend on the viewer-centered XYZ coordinate system. In differential geometry, there are a variety of surface representations from which viewpoint invariant principal curvatures can, in principle, be determined. Equation (7) determines principal curvatures from the Hessian matrix, **H**, and the gradient, $(p,q)$. The terms in Equation (7) involving the gradient, $(p,q)$, can be interpreted as the corrections required to account for the geometric foreshortening associated with viewing a surface element obliquely.

The directions $\omega_1$ and $\omega_2$ are viewpoint dependent. Although the directions of principal curvature are orthogonal in the object-centered coordinate system defined by the local surface normal and tangent plane, they are not, in general, orthogonal when projected onto the image plane. Thus, $k_1$, $k_2$, $\omega_1$ and $\omega_2$ together constitute four independent parameters that can be exploited. (Because they are viewpoint dependent, the directions $\omega_1$ and $\omega_2$ are not often used in surface representations proposed for object recognition. Note, however, that Brady *et. al.* [13] argue that, in many cases, the lines of curvature form a natural parameterization of a surface).

The *Gaussian curvature*, K, also called the *total curvature*, is the product, $\mathrm{K} = k_1 k_2$, of the principal curvatures. The *mean curvature*, H, is the average, $\mathrm{H} = (k_1 + k_2)/2$, of the principal curvatures. It follows from elementary matrix theory that

$$\mathrm{K} = \det(\mathbf{C}) \;\text{ and }\; \mathrm{H} = \frac{1}{2}\,\mathrm{trace}(\mathbf{C}) \tag{8}$$

The expression for K further simplifies to

$$\mathrm{K} = \frac{1}{(1+p^2+q^2)^2}\det(\mathbf{H}) \tag{9}$$

Thus, the sign of det (**H**) is the sign of the Gaussian curvature. Besl and Jain [14, 15] classify sections of surface into one of eight basic types based on the sign and zeros of Gaussian and mean curvature.

Clearly, if one could locally determine the Hessian, **H**, then one could locally compute the curvature matrix, **C**, using the gradient, $(p, q)$, obtained from photometric stereo and Equation (7). Given **C**, one could examine its eigenvalue/eigenvector structure to determine any local curvature representation involving the principal curvatures, $k_1$ and $k_2$, and their associated directions, $\omega_1$ and $\omega_2$, including both the Gaussian curvature, K, and the mean curvature, H.

## 2.3 Determining the Hessian

By taking partial derivatives of the image irradiance Equation (2) with respect to $x$ and $y$, two equations are obtained which can be written as the single matrix equation

$$\begin{bmatrix} E_x \\ E_y \end{bmatrix} = \mathbf{H} \begin{bmatrix} R_p \\ R_q \end{bmatrix} \tag{10}$$

(Here, and in what follows, subscripts $x$, $y$, $p$ and $q$ denote partial differentiation and the dependence of $E$ on $(x, y)$ and of $R$ on $(p, q)$ often is omitted for clarity). The vector $[E_x, E_y]$ is normal to the contour of constant brightness in the image, $E(x, y)$, at the given point $(x, y)$. The vector $[R_p, R_q]$ is normal to the contour of constant brightness in the reflectance map, $R(p, q)$, at the given gradient $(p, q)$. Equation (10) alone is not enough to determine the Hessian, **H**. But, with photometric stereo, one such equation is obtained for each image. With two light source photometric stereo

$$\mathbf{H} = \begin{bmatrix} E_{1x} & E_{2x} \\ E_{1y} & E_{2y} \end{bmatrix} \begin{bmatrix} R_{1p} & R_{2p} \\ R_{1q} & R_{2q} \end{bmatrix}^{-1} \tag{11}$$

With three light source photometric stereo

$$\mathbf{H} = \begin{bmatrix} E_{1x} & E_{2x} & E_{3x} \\ E_{1y} & E_{2y} & E_{3y} \end{bmatrix} \mathbf{M} \left( (\mathbf{M}^T \mathbf{M})^{-1} \right)^T \tag{12}$$

where

$$\mathbf{M} = \begin{bmatrix} R_{1p} & R_{1q} \\ R_{2p} & R_{2q} \\ R_{3p} & R_{3q} \end{bmatrix} \tag{13}$$

($^T$ denotes matrix transpose). Equation (12) is the standard least squares estimate of the solution to an overdetermined set of linear equations. It can be extended, in the obvious way, to situations in which more than three light sources are used.

It is useful to consider whether the matrix inverses required in Equations (11) and (12) can be guaranteed to exist. The essential observation is that the matrices whose inverses are required are defined entirely in terms of the reflectance maps, $R_i(p,q)$, independent of the particular images, $E_i(x,y)$. Thus, for a particular surface material, the key factor to control the existence and robustness of the computation is the nature and the distribution of the light sources. No useful local information is obtained when $[R_p, R_q]$ is zero. This occurs at local extrema of $R(p,q)$ and at gradients, $(p,q)$, shadowed from the light source. There also may be gradients, $(p,q)$, where two of the $[R_p, R_q]$ vectors are nearly parallel. Local degeneracies can be eliminated and the effects of shadows minimized when three, rather than two, light source photometric stereo is used.

It also is important to note that in Equation (12) the magnitude of $[R_p, R_q]$ plays the role of a "weight" that pulls the three source solution towards an image irradiance equation for which the magnitude of $[R_p, R_q]$ is large (and consequently away from an image irradiance equation for which the magnitude of $[R_p, R_q]$ is small). This has a desirable effect because locations in an image at which the magnitude of $[R_p, R_q]$ is small will contribute minimal information, and thus it is good that they are discounted. Because of this, points that are shadowed with respect to one of the light sources need not be considered as a special case. Indeed, when one of the $[R_p, R_q]$ vectors is zero, the three light source solution, given by Equation (12), reduces to the two light source solution, given by Equation (11).

## 2.4 Integrability

As we saw above, if a surface is locally smooth then the Hessian, $\mathbf{H}$, is symmetric. Given two images, $E_i(x,y)$ and $E_j(x,y)$, $i \neq j$, with corresponding reflectance maps, $R_i(p,q)$ and $R_j(p,q)$, the Hessian, $\mathbf{H}$, is given by Equation (11). When one expands Equation (11) and checks the resulting matrix for symmetry, one finds that $\mathbf{H}$ is symmetric iff

$$[E_{ix}, E_{iy}] \cdot [R_{jp}, R_{jq}] \;=\; [E_{jx}, E_{jy}] \cdot [R_{ip}, R_{iq}] \tag{14}$$

($\cdot$ denotes vector inner product). Therefore, in a three light source configuration, one obtains three equations, as in Equation (14) above, involving the vectors $[E_{ix}, E_{iy}]$ and $[R_{jp}, R_{jq}]$, $i = 1,2,3$ $j = 1,2,3$ $i \neq j$. These equations can be used to check the consistency of the local estimates of the Hessian matrix, $\mathbf{H}$. Alternatively, it is possible that lack of consistency could be used to detect discontinuities in surface orientation.

In any event, when one estimates the Hessian, $\mathbf{H}$, using either Equation (11) or (12), it is unlikely that the estimated Hessian, $\hat{\mathbf{H}}$, will be exactly symmetric, owing to numerical error. Symmetry can be forced by projecting $\hat{\mathbf{H}}$ onto the closest symmetric matrix. An obvious candidate for the closest symmetric matrix is the so-called symmetric part of $\hat{\mathbf{H}}$, given by

$$\frac{\hat{\mathbf{H}} + \hat{\mathbf{H}}^T}{2} \tag{15}$$

Frankot and Chellappa [16] observed that their projection of a non-integrable set of gradients, $(p, q)$, onto a (closest) integrable set caused additional smoothing, leading to more rapid convergence in their shape from shading algorithm. Using Equation (15), one can project the estimated Hessian, $\hat{\mathbf{H}}$, onto a (closest) symmetric matrix. This can be done after estimating the Hessian, $\mathbf{H}$, with Equation (11) or (12) and prior to estimating the curvature matrix, $\mathbf{C}$, with Equation (7).

## 2.5 Validating the Estimate of the Hessian

One would like to exploit the redundancy inherent in an overdetermined problem in order evaluate the validity of the estimated solution. The validity of estimate of the Hessian, $\hat{\mathbf{H}}$, can be expressed quantitatively using an extension of Equation (10). Define a matrix of residuals, $\mathbf{R} = [r_{ij}]$, by

$$\mathbf{R} = \begin{bmatrix} E_{1x} & E_{2x} & E_{3x} \\ E_{1y} & E_{2y} & E_{3y} \end{bmatrix} - \hat{\mathbf{H}} \begin{bmatrix} R_{1p} & R_{2p} & R_{3p} \\ R_{1q} & R_{2q} & R_{3q} \end{bmatrix} \tag{16}$$

Using the Frobenius matrix norm, a relative error term is given by the ratio

$$\left( \frac{\sum r_{ij}^2}{{E_{1x}}^2 + {E_{1y}}^2 + {E_{2x}}^2 + {E_{2y}}^2 + {E_{3x}}^2 + {E_{3y}}^2} \right)^{\frac{1}{2}} \tag{17}$$

This error term combines components due to measurement uncertainty in $[E_{ix}, E_{iy}]$ and to systematic modeling error either in the image irradiance equations (4) or in the determination of $[R_{ip}, R_{iq}]$, $i = 1, 2, 3$.

# 3 Implementation and Experimental Results

## 3.1 Experimental Setup

The objects used in this study were pottery vases. Pottery, in bisque form, is a reasonably diffuse reflector although no particular assumption is made (or required) concerning the underlying surface reflectance function. The vases were imaged from a distance of about 3m using a CCD camera equipped with a C-mount 50mm telephoto lens. In this configuration, the variation in surface relief is small compared to the distance to the camera so that image projection is well modeled as orthographic. Each vase was imaged three times from the identical viewpoint but under different conditions of illumination. The different illuminations were achieved using three different light sources, each of which was a standard 35mm slide projector. Each projector was focussed, to the extent possible with standard lenses, to produce a collimated beam. The light sources were sufficiently distant that any variation in surface relief was negligible compared to the distance to the light source. Therefore, it is reasonable to assume that scene irradiance was independent of depth. The CCD camera was operated with its automatic gain control (AGC) suppressed. This is required so that the brightness values obtained for images of different scenes, but under identical conditions of illumination, can be directly compared. The effect of inter-reflection was minimized by housing the vases in a custom "studio" with matte black walls and ceiling. All other lights in the room were turned off prior to each image acquisition.

## 3.2 Calibration

One way to obtain a reflectance map is to measure it using a calibration object of known shape. Measurements from a calibration object support the analysis of other objects provided the other objects are made of the same material and are illuminated and viewed under the same imaging conditions. Calibration by measurement has the added benefit of automatically compensating for the transfer characteristics of the sensor. Ideally, the calibration object would have visible surface points spanning the full range of gradients, $(p, q)$. A sphere, for example, is a good choice. Solids of revolution are another choice. For a solid of revolution, it is possible to determine the gradient, $(p, q)$, and the Hessian matrix, $\mathbf{H}$, at each visible surface point by geometric analysis of the object's bounding contour.

To simplify calibration, we make the axis of the solid of revolution parallel to the image plane. Without loss of generality, we can assume a coordinate system in which this axis is aligned with the image Y-axis. Under these assumptions, a solid of revolution is the volume swept out by moving a circular cross section along the image Y-axis while magnifying or contracting it in a smoothly varying way. Let $r = h(y)$ be the radius of the circular cross section and let $h'(y)$ and $h''(y)$ be its first and second derivatives with respect to $y$.

If $(x, y)$ is any visible surface point, the corresponding gradient, $(p, q)$, is given by

$$p = \frac{x}{\sqrt{r^2 - x^2}} \quad \text{and} \quad q = \frac{-r\, h'(y)}{\sqrt{r^2 - x^2}} \tag{18}$$

and the elements of the corresponding Hessian matrix, $\mathbf{H}$, are given by

$$\begin{aligned} p_x &= \frac{r^2}{(r^2 - x^2)^{\frac{3}{2}}} \\ q_y &= \frac{x^2\, h'(y)^2 - r\,(r^2 - x^2)\, h''(y)}{(r^2 - x^2)^{\frac{3}{2}}} \\ p_y &= q_x = \frac{-r\, x\, h'(y)}{(r^2 - x^2)^{\frac{3}{2}}} \end{aligned} \tag{19}$$

Figure 1 shows the three images of the calibration vase obtained from the three different light source positions. The object's bounding contour is easily determined. The particular method used consists of three steps. First, the three images are summed together. Second, the intensity histogram of the sum is computed. Third, a single threshold is selected to separate object points from background points. Simple thresholding of the summed image is sufficient because, by design, the histogram always is distinctly bimodal. Once the object silhouette is determined, values for $h'(y)$ and $h''(y)$ can be estimated. Of course, some smoothing of the bounding contour is required to overcome local quantization effects. Here, smoothing and derivative estimation is combined by filtering the boundary curve with the first and second derivatives of the Gaussian using the particular method described by Lowe [17].

Thus, a solid of revolution is a useful calibration object because it is possible, by geometric analysis alone, to determine the gradient, $(p, q)$, the Hessian matrix, $\mathbf{H}$, and, using Equation (7), the principal curvatures and associated directions at each visible surface point. At each $(x, y)$, the five parameters required, $p$, $q$, $p_x$, $q_y$ and $p_y = q_x$, depend only on $r$, $h'(y)$ and $h''(y)$. In practice, the measurement of $r$, $h'(y)$ and $h''(y)$ can be made both accurate and robust.

## 3.3 Lookup Tables for Photometric Stereo

Photometric stereo does not, in fact, require that reflectance maps be determined explicitly. It is sufficient that the required information be represented implicitly in the lookup tables used to determine the gradient, $(p, q)$, as a function of the measured triple of brightness values, $[E_1, E_2, E_3]$.

Here, the lookup table constructed was of dimension $2^5 \times 2^5 \times 2^5 = 2^{15}$. For the given table size, the goal is to use the available entries to provide the best resolution possible in the estimate of the gradient, $(p, q)$. The technique of histogram equalization was used to achieve table compression. Table entries are disproportionately allocated to brightness values that occur frequently on the calibration object. This provides good discrimination between nearby brightness values that occur frequently at the expense of poorer discrimination in ranges of brightness values that occur less frequently or not at all.

Each image point on the calibration object was sampled. The brightness triple, $[E_1, E_2, E_3]$, determined the table location. The corresponding gradient was calculated using Equation (18). The gradient was then converted into a unit surface normal and entered into the table. Multiple triples, $[E_1, E_2, E_3]$, can map to the same table location. The resulting surface normals were averaged and converted back to gradients as a post-processing step. Some table interpolation is done, at the user's discretion, also as a post-processing step to fill in missing table entries.

By construction, this lookup table achieves good discrimination between brightness triples, $[E_1, E_2, E_3]$, that correspond to possible measurements from points on the calibration object. Space in the table is not effectively utilized for this purpose when entries are allocated to impossible brightness triples, $[E_1, E_2, E_3]$. Thus, the particular lookup table design is not best for segmentation, defined in this context to be the separation of object points from non-objects points.

Once the lookup table is established, images from the other objects are quickly analyzed. For each additional object, the input is the triple of images, $E_1(x, y)$, $E_2(x, y)$, and $E_3(x, y)$, and the output is a file giving the gradient, $(p, q)$, at each point. A bitmap file also is produced to indicate those points at which the gradient, $(p, q)$, was estimated.

## 3.4 Determining the Reflectance Maps

The reflectance maps, $R_1(p, q)$, $R_2(p, q)$, and $R_3(p, q)$, are required for curvature estimation since their partial derivatives, with respect to $p$ and $q$, are needed in Equations (11) and (12). To determine the reflectance map, it is convenient to solve the inverse of Equation (18). That is, given a gradient, $(p, q)$, we want to determine a visible surface point, $(x, y)$, at which to measure brightness. This can be done in a two step process. First, given $(p, q)$, we find a value of $y$ such that

$$h'(y) = \frac{-q}{\sqrt{1 + p^2}} \tag{20}$$

In general, this value of $y$ will not be unique. In practice, it is convenient to choose a range for $y$ over which the function $h'(y)$ is monotonic. This guarantees that a unique value of $y$ can be obtained. Given $y$, we also know the radius, $r = h(y)$, and we obtain the required $x$ as

$$x = \frac{r\,p}{\sqrt{1+p^2}} \tag{21}$$

Figure 2(a) shows the analysis of the occluding boundary of the calibration object of Figure 1. Superimposed upon the boundary, as a thicker curve, is the portion chosen over which $h'(y)$ is monotonic. Normals to the boundary curve also are plotted at selected intervals. Unfortunately, not all values of the gradient, $(p, q)$, occur in Figure 2(a). One way to obtain additional coverage is to repeat the analysis using three additional images (not shown) of the calibration object simply placed upside down. This analysis is shown in Figure 2(b). Figure 3 shows the three reflectance maps obtained by combining the results of Figures 2(a) and 2(b). In Figure 3, the axis tick marks are one unit apart so that the range covered is $-2.5 \leq p \leq 2.5$ and $-2.5 \leq q \leq 2.5$. Some areas still are missing (i.e., black) in all three reflectance maps, $R_1(p,q)$, $R_2(p,q)$ and $R_3(p,q)$, because the results of Figures 2(a) and 2(b) do not yet span the full range of possible gradients. One remedy would be to combine the above results with measurements from three additional images of the same calibration vase re-positioned so that the axis of revolution was aligned with the image X-axis. This was not done here.

### 3.5 Determining Surface Curvature

To determine surface curvature at each surface point, we need to know the gradient, $(p, q)$, and the twelve partial derivatives, $E_{ix}$, $E_{iy}$, $R_{ip}$, $R_{iq}$, $i = 1, 2, 3$. The gradient, $(p, q)$, is obtained as the output from photometric stereo, as described in Section 3.3. The reflectance maps, $R_i(p, q)$, $i = 1, 2, 3$, are obtained as described in Section 3.4. The images, $E_i(x, y)$, and reflectance maps are further processed to estimate partial derivatives. One could compute the twelve partial derivatives explicitly by combining the appropriate directional derivative with some degree of local Gaussian smoothing. In the current implementation, there was insufficient memory available to store all twelve partial derivative results. Therefore, each image and reflectance map was smoothed with a 2D Gaussian and the required partial derivatives were estimated using simple local differencing, as required. This reduced the on-line storage requirement to the six files, $E_i(x, y)$, $R_i(p, q)$, $i = 1, 2, 3$, and the gradients, $(p, q)$.

For each object point for which the gradient is known and for which $R_i(p, q)$ is defined, Equation (12) is used to estimate the Hessian matrix, $\mathbf{H}$. The resulting Hessian is made symmetric using Equation (15). The curvature matrix, **C**, is determined using Equation (7). From the matrix **C**, the principal curvatures, $k_1$ and $k_2$, their associated directions, $\omega_1$ and $\omega_2$, the Gaussian curvature, K, and the mean curvature, H, are derived, as described in Section 2.2. The relative error of the estimate of **H** also is determined using Equation (17).

### 3.6 Results

This paper is not about surface reconstruction directly. Nevertheless, it is convenient to perform some surface reconstruction in order to examine the results of photometric stereo.

Reconstructing a surface, $z = f(x, y)$, given the gradient, $(p, q)$, is a problem of integration. Given a point, $(x, y)$, and the corresponding gradient, $(p, q)$, the change in height, dz, corresponding to a small movement [dx, dy] is given by

$$\mathrm{dz} = p\,\mathrm{dx} + q\,\mathrm{dy} \tag{22}$$

This suggests a very simple surface reconstruction scheme. Given an initial point, $z_0 = f(x_0, y_0)$, one can use Equation (22) to trace a path along which depth is reconstructed. Figure 4 shows a plot of a surface reconstructed from the results of photometric stereo applied to the calibration object of Figure 1. Given an initial value, $z_0 = f(x_0, y_0)$, Equation (22) was used to reconstruct a single depth profile in the vertical (i.e., column) direction. Using the values so obtained as initial conditions, Equation (22) was then used to reconstruct depth profiles in the horizontal direction, one profile for each image row.

The images shown in Figure 1 are $256 \times 256$. To avoid clutter, Figure 4 plots every fourth row and column. No smoothing has been performed. The quality of the reconstruction is due to the local accuracy achieved with photometric stereo. Of course, one would expect this example to be accurate since, after all, it is the example of the calibration object itself.

Figure 5 shows curvature results for the calibration object. In this and in the examples to follow, the figure consists of three parts, (a), (b) and (c). The mean curvature, H, is shown in Figure 5(a). The Gaussian curvature, K, is shown in Figure 5(b). The results are scaled to $2^8 = 256$ values and offset so that middle gray represents zero, lighter points represent positive values and darker points represent negative values. The relative error of the estimate of the Hessian, computed using Equation (17), is shown in Figure 5(c). Darker points represent larger relative error. As one would expect, curvature remains relatively constant over the body of the calibration vase. At the base and at the neck; negative curvatures are seen.

Figure 6 shows three images of the calibration vase now with its top in place. Each image is $320 \times 256$. These images were obtained with the identical conditions of illumination as in

Figure 1. Figure 7 shows a plot of a surface reconstructed from the results of photometric stereo. The surface reconstruction and subsequent plotting were performed exactly as described above for Figure 4. Again, no smoothing has been performed. The quality of the reconstruction is due to the local accuracy achieved with photometric stereo. Detail at the neck and at the top is well maintained.

Figure 8 shows curvature results for the topped vase. As expected, the results for the body of the vase are very close to those obtained in Figure 5. Interesting local detail emerges at the top. High positive mean and Gaussian curvature is seen at the top tip and at the lip where the top sits on the base. Negative mean and Gaussian curvature separates the lip of the top from the main body of the top.

Figure 9 shows the three images obtained for the most complex object analyzed, a long necked vase with a handle. Each image is $320 \times 192$. Again, these images were obtained with the identical conditions of illumination as in Figure 1. Figure 10(a) shows a plot of a surface reconstructed from the results of photometric stereo. The surface reconstruction and subsequent plotting were performed exactly as described above for Figure 4. In this case, while the simple surface reconstruction algorithm is sufficient to give some idea of the local accuracy achieved with photometric stereo, it is not sufficient to obtain a reconstruction over the entire surface. Any scheme based on Equation (22) will be suspect since it will be difficult to make errors that propagate along the neck compatible with errors that propagate along the handle.

Figure 10(b) shows a plot of a surface reconstructed from the results of photometric stereo using an implementation of Harris' method of surface reconstruction [18, 19]. This method finds the surface, $z = f(x, y)$, that minimizes

$$\int\int\left(\left((z_x - p)^2 \; + \; (z_y - q)^2\right) \; + \; \lambda\left(p_x^2 + p_y^2 + q_x^2 + q_y^2\right)\right) \mathrm{dx}\ \mathrm{dy} \qquad (23)$$

Imposing a global smoothness term, $\left(p_x^2 + p_y^2 + q_x^2 + q_y^2\right)$, helps to combine the handle and the neck in the reconstruction. At the same time, in the absence of prior boundary conditions, the global smoothness term causes the depth discontinuities at the top and at the bottom of the vase to be lost. The result in Figure 10(a) preserves local detail but is incomplete. The result in Figure 10(b) is complete but fails to preserve local detail.

Figure 11 shows curvature results for the long necked vase with a handle. This example provides the greatest variety of curvature values of any of the examples considered. The long neck has moderately high positive mean curvature and negative Gaussian curvature, except for its middle section where it is essentially cylindrical and the Gaussian curvature is effectively zero. The handle has very high positive mean curvature throughout. Its Gaussian curvature is highly positive at the top where the handle curves sharply. Elsewhere along the handle, the Gaussian curvature is nearly zero since the handle is nearly cylindrical. Where the handle joins the main body of the vase, both the mean curvature and the Gaussian curvature become sharply negative. The relative error measure is quite low on the neck, on the handle and especially where the handle joins the main body of the vase. This indicates that the local curvature estimates are reliable.

Finally, it is important to point out that all the curvature results shown in Figures 5, 8 and 11 are dense, local estimates. Any apparent global coherence of the result is a consequence of the computation, not a consequence of any global smoothness assumption embedded therein. In particular, the object boundary plays no role in the computation. The local results would be identical even if the object boundary was obscured.

# 4 Discussion and Conclusions

Multiple images acquired under the identical viewing geometry but different conditions of illumination provide additional local constraint to determine shape from shading. One way to exploit this fact is to obtain better estimates of local surface orientation. Photometric stereo does this accurately and robustly with minimal computational requirements compared to the iterative schemes typically required for shape from shading from a single image.

Here, it has been demonstrated that one can obtain, in addition, estimates of local second-order surface curvature. This represents a new capability for local shading analysis. The computational requirements remain quite modest. The method described does require knowledge of $[R_p, R_q]$ for each image. But, this also is the case for methods for determining shape from shading from a single image. The method also requires estimates of $[E_x, E_y]$ for each image. Interestingly, Horn's original method for shape from shading, based on characteristic strip expansion, also required estimates of $[E_x, E_y]$. Modern variational formulations of shape from shading, however, typically do not require $[E_x, E_y]$. Of course, estimates of $[E_x, E_y]$ are required for several other vision problems, including edge detection and optical flow. Thus, multiple light source curvature estimation does not place undue demands on what is to be measured, compared to alternatives. The resulting computation is direct without requiring any iteration steps.

The particular formulation for curvature, given in Equation (7), was proposed in [1, 3]. As argued above, this expression nicely decouples a viewer-centered representation for curvature, given by the Hessian, $\mathbf{H}$, from viewpoint dependent foreshortening, given in terms of the gradient, $(p, q)$. The Hessian matrix is key because it relates the intensity gradient, $[E_x, E_y]$, and the reflectance map gradient, $[R_p, R_q]$. The notion that intensity gradients can be used to estimate the Hessian, $\mathbf{H}$, was noted in [1] and further developed, in the context of multiple light sources, by Wolff [20, 21]. Penna and Chen [22] also considered shape from shading with multiple light sources. They derived analytic expressions for surface curvature based on the assumption of Lambertian reflectance and single distant point light sources.

The current work, including an earlier conference paper [23], represents the first experimental demonstration that surface curvatures can, in fact, be estimated locally in shape from shading. The approach makes no prior assumption about the reflectance properties of the objects in view. In particular, it does not assume Lambertian (or any other) reflectance function. The robustness of the approach does depend on the accuracy with which the reflectance maps, $R(p, q)$, and their gradients, $[R_p, R_q]$, are known. Careful calibration can lead to accuracy in measurement. Careful choice and positioning of the light sources can be used to make the computation of the Hessian, $\mathbf{H}$, well-conditioned.

As shapes analyzed become more complex, segmentation based on surface curvature be-comes more important. To specify the local properties of a surface up to curvature requires six parameters since there is one degree of freedom for range, two for surface orientation and three for curvature. If only a single measurement, say range, is available locally, then the problem is underconstrained. The usual approach is to interpolate a high-order surface with desired properties using measurements otained over extended regions. But, this begins to beg the question since one no longer has reliable local estimates of curvature upon which segmentation can be based. The approach demonstrated here is to seek additional local information. Multiple images acquired from the same viewpoint but under different conditions of illumination provide useful additional information. In principle, each image provides three independent pieces of local information, one for brightness and two for the two partial spatial derivatives of brightness. (To be truly independent, one would need an image sensor that measured partial derivatives directly). With three images one obtains nine local measurements which is sufficient to overconstrain the local solution. Thus, accuracy in local curvature estimation relates directly to the quality of imaging. Robustness can be achieved by overdetermining the computation locally, rather than by imposing global smoothness constraints.

Segmentation has always been a "chicken-and-egg" problem in computational vision. The contribution of this work is to demonstrate that surface curvature can be computed locally and reliably prior to segmentation. This, in turn, should allow future segmentation schemes to be more robust.

## Acknowledgement

The work described in this paper benefited from discussions with A.K. Mackworth, D.G. Lowe, J.J. Little and J.M. Varah. L.B. Wolff pointed out the connection to his work. D.G. Lowe supplied software for curve smoothing. Y. Li implemented a version of Harris' surface reconstruction method and produced the surface shown in Figure 10(b). Y. Li and S. Kingdon assisted with image acquisition in the laboratory. Pottery used in the experiments was obtained with the help of the staff at Steveston Ceramics, Steveston, BC. Support for the work described was provided, in part, by the Natural Science and Engineering Research Council of Canada (NSERC) and, in part, by the Canadian Institute for Advanced Research (CIAR).

# References


[1] R. J. Woodham, "Reflectance map techniques for analyzing surface defects in metal castings," AI-TR-457, MIT AI Laboratory, Cambridge, MA, 1978.

[2] R. J. Woodham, "Photometric method for determining surface orientation from multiple images," *Optical Engineering*, vol. 19, pp. 139–144, 1980. Reprinted in [12].

[3] R. J. Woodham, "Analysing images of curved surfaces," *Artificial Intelligence*, vol. 17, pp. 117–140, 1981.

[4] W. M. Silver, "Determining shape and reflectance using multiple images," SM thesis, MIT Dept. Electrical Engineering and Computer Science, Cambridge, MA, 1980.

[5] K. Ikeuchi, "Recognition of 3-D objects using the extended Gaussian image," in *Proc. 7th Int. Joint Conf. on Artificial Intelligence*, pp. 595–600, 1981.

[6] P. Brou, "Using the Gaussian image to find the orientation of objects," *Int. J. Robotics Research*, vol. 3, no. 4, pp. 89–125, 1984.

[7] B. K. P. Horn and K. Ikeuchi, "The mechanical manipulation of randomly oriented parts," *Scientific American*, vol. 251, pp. 100–111, August 1984.

[8] K. Ikeuchi, "Determining attitude of object from needle map using extended Gaussian image," AI-Memo-714, MIT AI Laboratory, Cambridge, MA, 1983.

[9] K. Ikeuchi, B. K. P. Horn, S. Nagata, T. Callahan, and O. Feingold, "Picking up an object from a pile of objects," in *Robotics Research: The First International Symposium* (M. Brady and R. Paul, eds.), pp. 139–162, Cambridge, MA: MIT Press, 1984.

[10] K. Ikeuchi, H. K. Nishihara, B. K. P. Horn, P. Sobalvarro, and S. Nagata, "Determining grasp configurations using photometric stereo and the prism binocular stereo system," *Int. J. Robotics Research*, vol. 5, no. 1, pp. 46–65, 1986.

[11] B. K. P. Horn, *Robot Vision.* Cambridge, MA: MIT Press, 1986.

[12] B. K. P. Horn and M. J. Brooks, eds., *Shape from Shading.* Cambridge, MA: MIT Press, 1989.

[13] M. Brady, J. Ponce, A. Yuille, and H. Asada, "Describing surfaces," *Computer Vision Graphics and Image Processing*, vol. 32, pp. 1–28, 1985.

[14] P. J. Besl and R. C. Jain, "Three-dimensional object recognition," *ACM Computing Surveys*, vol. 17, pp. 75–145, 1985.

[15] P. J. Besl and R. C. Jain, "Invariant surface characteristics for 3D object recognition in range images," *Computer Vision Graphics and Image Processing*, vol. 33, pp. 33–80, 1986.

[16] R. T. Frankot and R. Chellappa, "A method for enforcing integrability in shape from shading algorithms," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 10, pp. 439–451, 1988. Reprinted in [12].

[17] D. G. Lowe, "Organization of smooth image curves at multiple scales," *Int. J. Computer Vision*, vol. 3, pp. 119–130, 1989.

[18] J. G. Harris, "The coupled depth/slope approach to surface reconstruction," AI-TR-908, MIT AI Laboratory, Cambridge, MA, 1986.

[19] J. G. Harris, "A new approach to surface reconstruction: the coupled depth/slope model," in *Proc. 1st Int. Conf. on Computer Vision*, pp. 277–283, 1987.

[20] L. B. Wolff, "Surface curvature and contour from photometric stereo," in *Proc. DARPA Image Understanding Workshop, 1987*, pp. 821–824, 1987.

[21] L. B. Wolff, "Accurate measurement of second order variations of a smooth object surface using reflectance maps," in *Proc. SPIE Vol. 1003 Sensory Fusion: Spatial Reasoning and Scene Interpretation (1988)*, pp. 59–62, 1988.

[22] M. A. Penna and S. S. Chen, "Shape-from-shading using multiple light sources," *Int. J. Intelligent Systems*, vol. 1, pp. 263–291, 1986.

[23] R. J. Woodham, "Determining surface curvature with photometric stereo," in *Proc. IEEE Conf. Robotics and Automation, 1989*, pp. 36–42, 1989.

## Figure Captions

**Figure 1** Three images of the vase used as the calibration object. Light source 1 comes from a direction to the left and slightly above the viewing direction. Light source 2 comes from a direction to the right and slightly above the viewing direction. Light source 3 comes almost exactly from the viewing direction. $E_1(x,y)$, $E_2(x,y)$ and $E_3(x,y)$ are shown respectively as (a), (b), and (c).

**Figure 2** Geometric analysis of the boundary of the calibration object is used to determine the gradient, $(p,q)$, at each image point, $(x,y)$. Superimposed upon the boundary, as the thicker curve, is the portion chosen over which $h'(y)$ is monotonic. Selected normals to the boundary curve also are shown. The analysis of the images of Figure 1 is shown in (a). In order to obtain additional coverage in the gradient space, the analysis was repeated using three additional images (not shown) of the calibration object placed upside down. This analysis is shown in (b).

**Figure 3** The three reflectance maps determined by measurements from the calibration object. $R_1(p,q)$, $R_2(p,q)$ and $R_3(p,q)$ are shown respectively as (a), (b), and (c). The axis tick marks are one unit apart so that the range covered is $-2.5 \leq p \leq 2.5$ and $-2.5 \leq q \leq 2.5$. The black areas in (c) correspond to gradients not obtained from the calibration object, either from Figure 2(a) or from Figure 2(b). Additional dark areas in (a) and (b) correspond to points shadowed from the light source.

**Figure 4** Plot of the surface of the calibration vase reconstructed from the gradients, $(p,q)$, obtained from photometric stereo.

**Figure 5** Curvature results for the calibration vase. The mean curvature, H, is shown in (a). The Gaussian curvature, K, is shown in (b). In each case the result has been scaled so that middle gray is zero, lighter points represent positive values and darker points represent negative values. The relative error of the estimate of the Hessian is shown in (c). Darker points represent larger relative error.

**Figure 6** Three images of the calibration vase with its top in place. The three images were obtained with the identical conditions of illumination as in Figure 1. $E_1(x,y)$, $E_2(x,y)$ and $E_3(x,y)$ are shown respectively as (a), (b), and (c).

**Figure 7** Plot of the surface of the calibration vase with its top in place reconstructed from the gradients, $(p,q)$, obtained from photometric stereo.

**Figure 8** Curvature results for the calibration vase with its top in place. The mean curvature, H, is shown in (a). The Gaussian curvature, K, is shown in (b). The relative error of the estimate of the Hessian is shown in (c).

**Figure 9** Three images of the most complex object analyzed, a long necked vase with a handle. Again, the three images were obtained with the identical conditions of illumination as in Figure 1. $E_1(x,y)$, $E_2(x,y)$ and $E_3(x,y)$ are shown respectively as (a), (b), and (c).

**Figure 10** Plots of the surface of the long necked vase with a handle. The result with the algorithm used for Figures 4 and 7 is shown in (a). The result using an implementation of Harris' method of surface reconstruction is shown in (b).

**Figure 11** Curvature results for the long necked vase with a handle. The mean curvature, H, is shown in (a). The Gaussian curvature, K, is shown in (b). The relative error of the estimate of the Hessian is shown in (c).

**Figure 1**

(a)

(b)

(c)

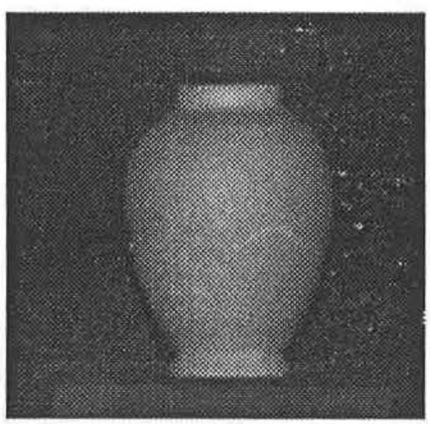

**Figure 2**

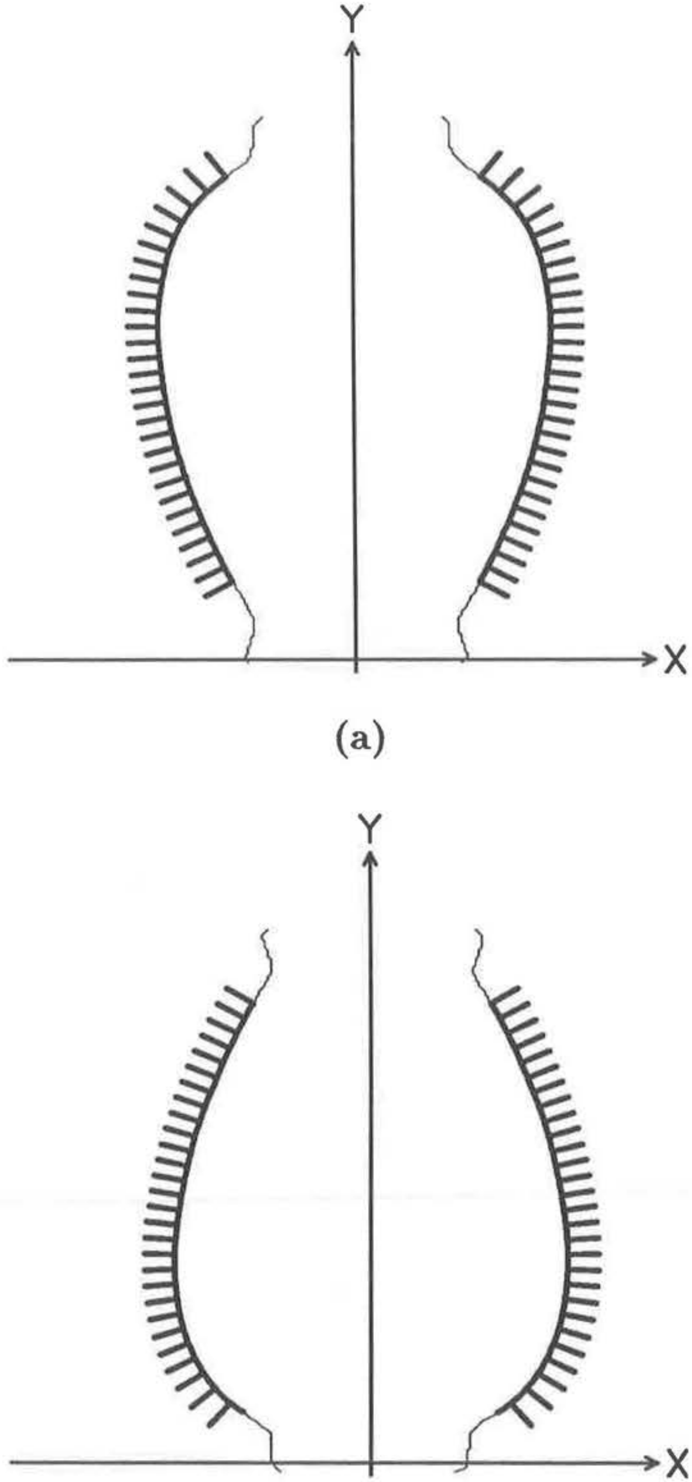


(a)

(b)

**Figure 3**

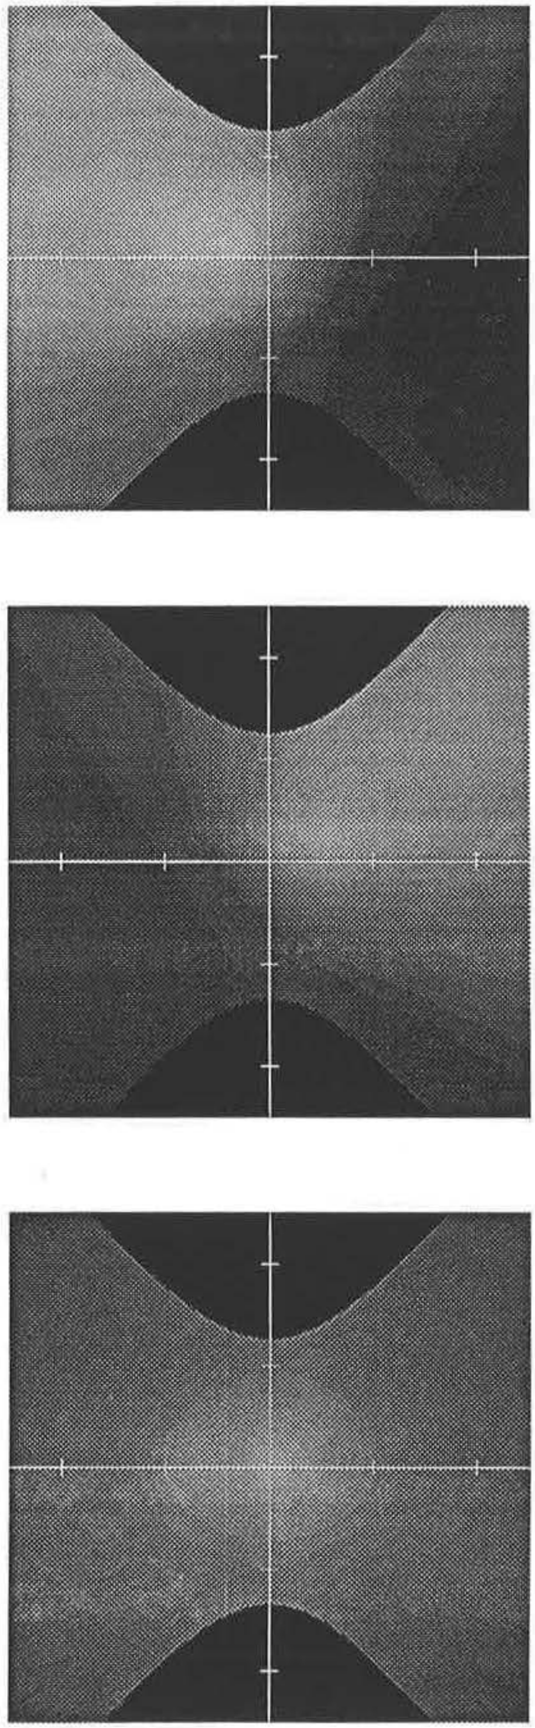

(a)

(b)

(c)

Figure 4

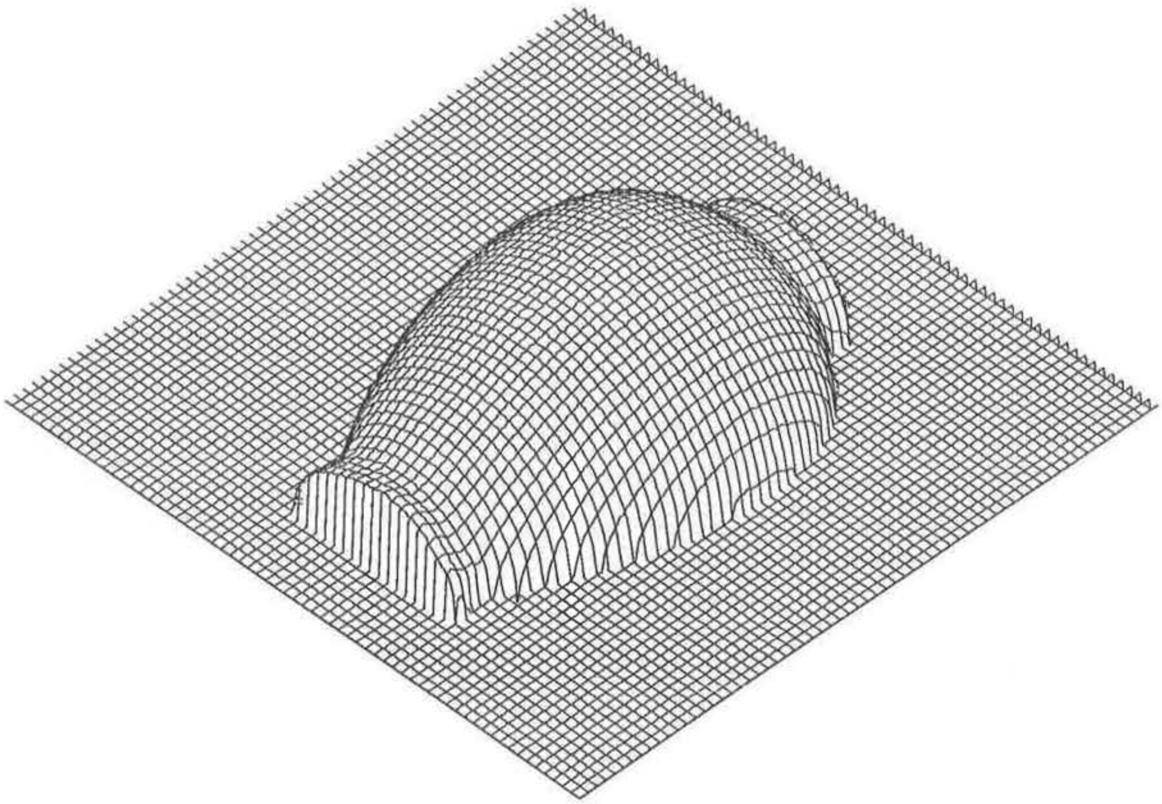

**Surface plot of calibration vase**

Figure 5

(a)

(b)

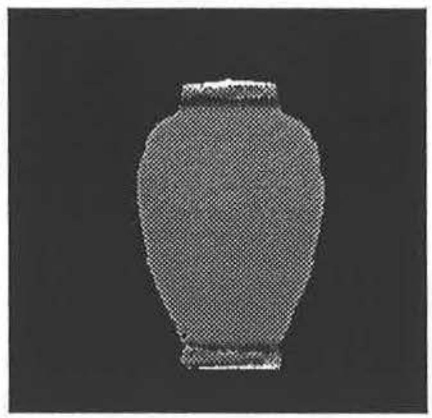

(c)

**Figure 6**

(a)

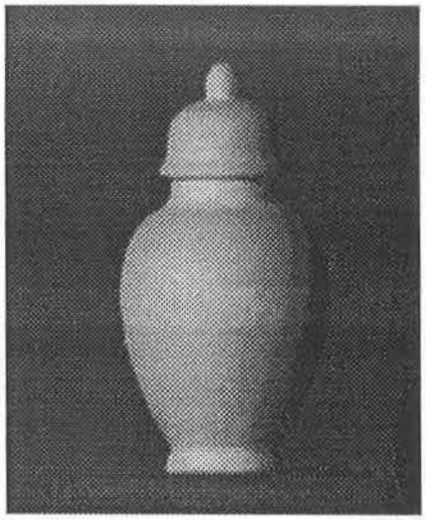

(b)

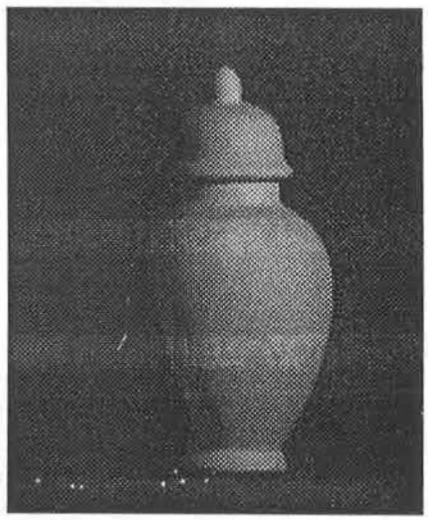

(c)

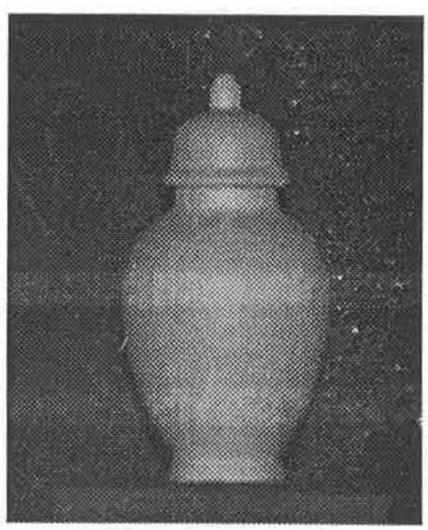

**Figure 7**

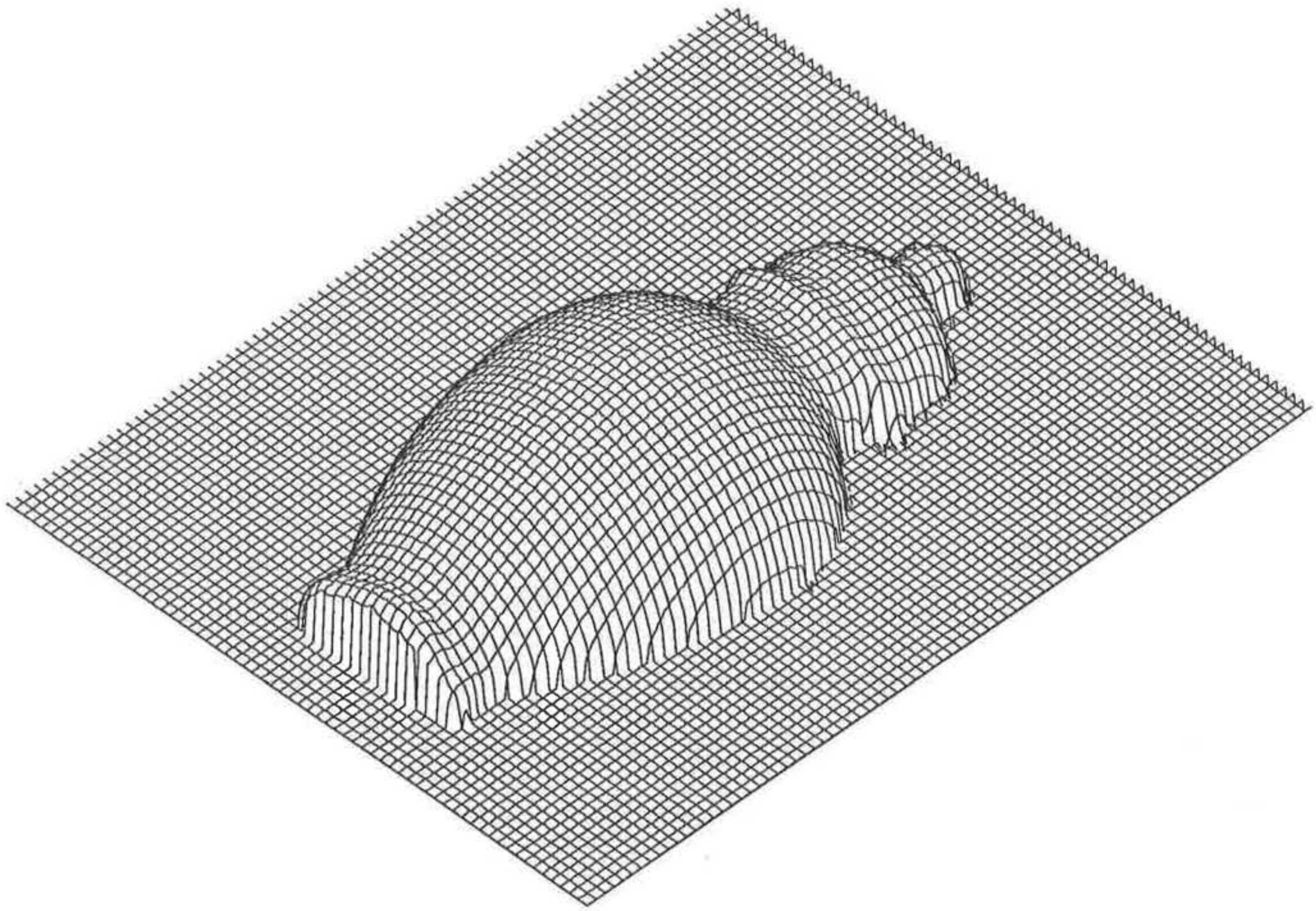

## Surface plot of topped vase

Figure 8

(a)

(b)

(c)

**Figure 9**

(a)

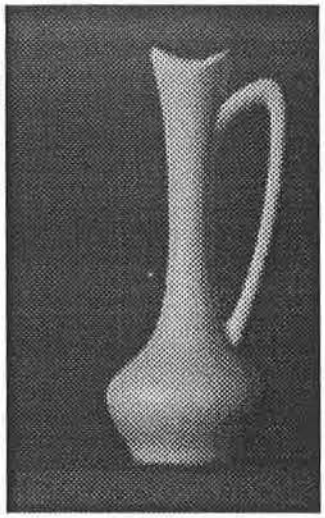

(b)

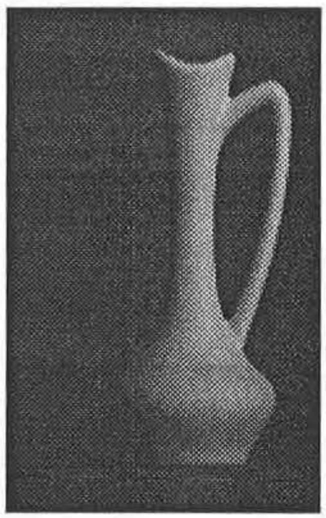

(c)

Figure 10

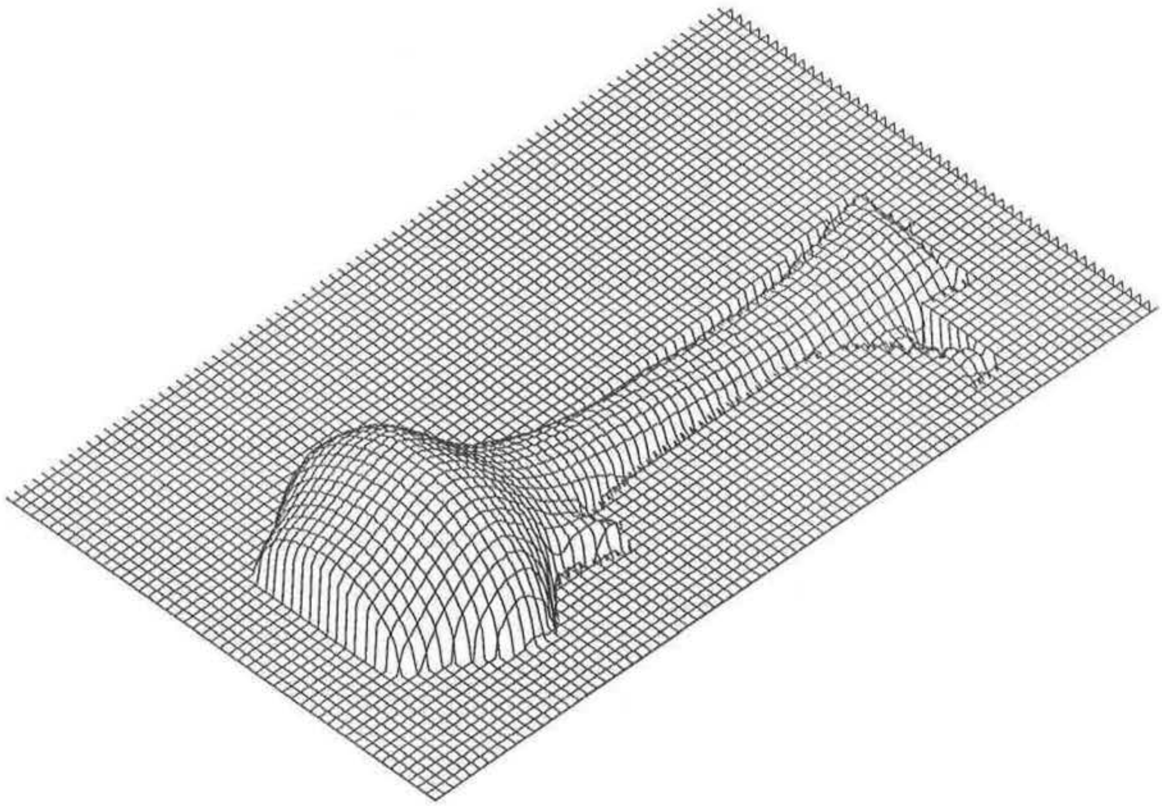

(a)

**Surface plot of handled vase**

Figure 10

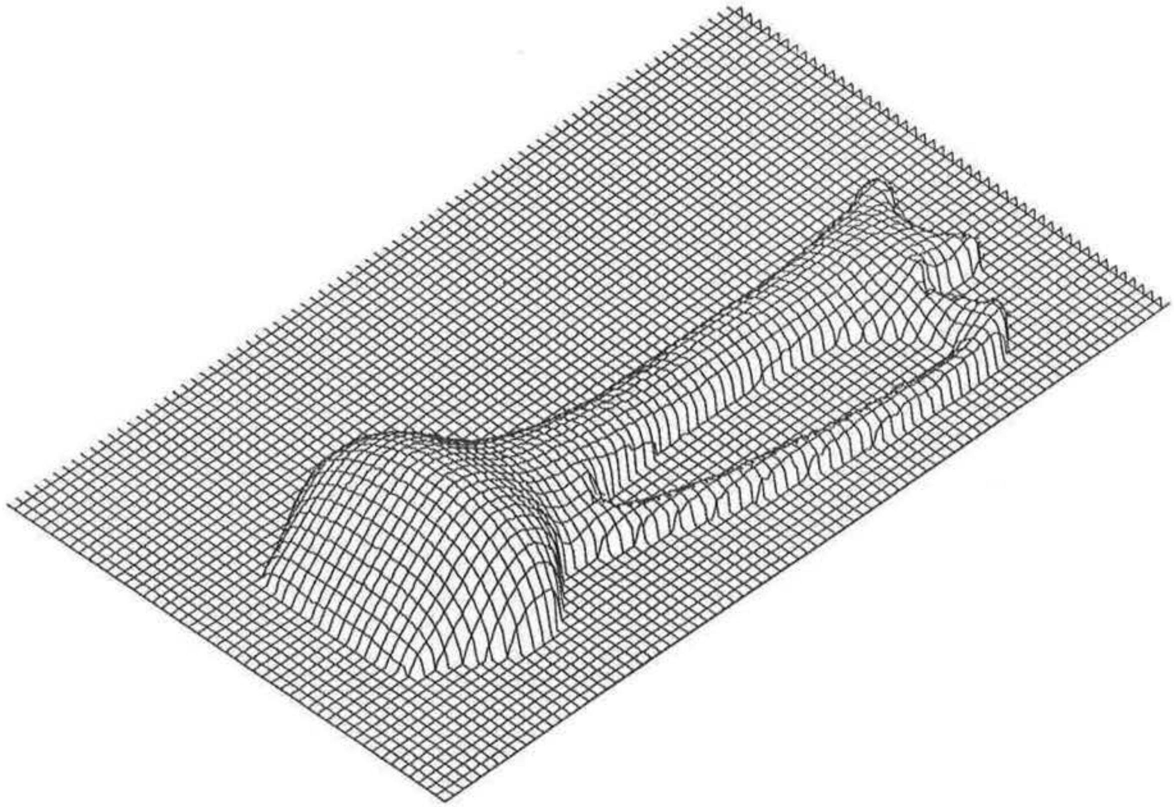

(b)

**Surface plot of handled vase (Harris reconstruction)**

Figure 11

(a)

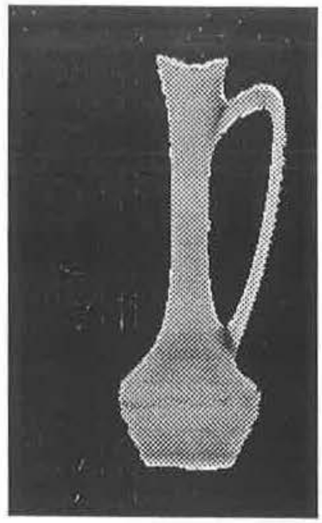

(b)

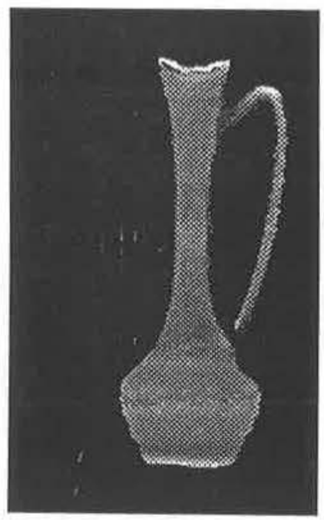

(c)