

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

AI Memo 479

June 1978

PHOTOMETRIC STEREO

Robert J. Woodham

ABSTRACT

Traditional stereo techniques determine range by relating two images of an object viewed from different directions. If the correspondence between picture elements is known, then distance to the object can be calculated by triangulation. Unfortunately, it is difficult to determine this correspondence.

This paper introduces a novel technique called photometric stereo. The idea of photometric stereo is to vary the direction of the incident illumination between successive views while holding the viewing direction constant. This provides enough information to determine surface orientation at each picture element. Since the imaging geometry does not change, the correspondence between picture elements is known a priori. This stereo technique is photometric because it uses the intensity values recorded in a single picture element, in successive views, rather than the relative positions of features.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-75-C-0643 and in part by the Office of Naval Research under Office of Naval Research contract N00014-77-C-0389.

1. Introduction

Work on image understanding has led to a need to model the imaging process. One aspect of this concerns the geometry of image projection. Less well understood is the radiometry of image formation. Relating the intensity values recorded in an image to object relief requires a model of the way surfaces reflect light.

A reflectance map is a convenient way to incorporate a fixed scene illumination, object photometry and imaging geometry into a single model that allows image intensity to be related directly to surface orientation. This relationship is not functional since surface orientation has two degrees of freedom and image intensity provides only one measurement. Local surface topography cannot, in general, be determined by the intensity value recorded at a single image point. In order to determine object shape, additional information must be provided.

This observation has led to a novel technique called photometric stereo in which surface orientation is determined from two or more images. Traditional stereo techniques determine range by relating two images of an object viewed from different directions. If the correspondence between picture elements is known, then distance to the object can be calculated by triangulation. Unfortunately, it is difficult to determine this correspondence. The idea of photometric stereo is to vary the direction of the incident illumination between successive views while holding the viewing direction constant. This provides enough information to determine surface orientation at each picture element. Since the imaging geometry does not change, the correspondence between picture elements is known a priori. This stereo technique is photometric because it uses the intensity values recorded at a single picture element, in successive views, rather than the relative positions of features.

2. The Reflectance Map

The fraction of light reflected by a surface in a given direction depends upon the optical properties of the object material, the surface microstructure and the spatial and spectral distribution and state of polarization of the incident light. A key photometric observation is:

No matter how complex the distribution of incident illumination, for most surfaces, the fraction of the incident light reflected in a particular direction depends only on the surface orientation.

The reflectance characteristics of an object material can be represented as a function $\phi(i, e, g)$ of the three angles i , e and g defined in figure 1. These are called, respectively, the *incident*, *emergent* and *phase* angles. The angles i , e and g are defined relative to the object surface. $\phi(i, e, g)$ determines the ratio of radiance to irradiance measured per unit surface area, per unit solid angle, in the direction of the viewer. The reflectance function defined here is related to the bi-directional reflectance-distribution function defined by the National Bureau of Standards [Nicodemus *et al* 77].

Image forming systems perform a perspective transformation as illustrated in figure 2(a). If the size of the objects in view is small compared to the viewing distance, then the perspective projection can be approximated as an orthographic projection as illustrated in figure 2(b). Consider an image forming system that performs an orthographic projection. To standardize the imaging geometry, it is convenient to align the viewing direction with the negative z -axis. Also, assume appropriate scaling of the image plane so that object point (x, y, z) maps onto image point (u, v) with $u = x$ and $v = y$. With these assumptions, image coordinates (x, y) and object coordinates (x, y) can be referred to interchangeably.

If the equation of a surface is given explicitly as:

$$z = f(x, y)$$

then a surface normal is given by the vector:

$$[\partial f(x,y)/\partial x, \partial f(x,y)/\partial y, -1].$$

If parameters p and q are defined by:

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

then the surface normal can be written as $[p, q, -1]$. The quantity (p, q) is called the *gradient*, and *gradient space* is the two-dimensional space of all such points (p, q) . Gradient space is a convenient way to represent surface orientation. It has been used in scene analysis [Mackworth 73]. In image analysis, it is used to relate the geometry of image projection to the radiometry of image formation [Horn 77].

One important simplification inherent in the assumption of an orthographic projection is that the viewing direction, and hence the phase angle g , is constant for all object points. Thus, for a standard light source and viewer geometry, the ratio of radiance to irradiance depends only on gradient coordinates p and q .

Further, suppose each surface element receives the same incident illumination. Then, the amount of the incident light reflected in a particular direction depends only on the surface orientation. The assumption of an orthographic projection also allows one to relate the amount of light reflected per unit surface area per unit solid angle in the direction of the viewer directly to image intensity. Thus, for the given imaging geometry, for a given distribution of incident illumination and a given object material, the image intensity corresponding to a surface point with gradient (p, q) is unique.

The *reflectance map* $R(p, q)$ determines image intensity as a function of p and q . A reflectance map captures the surface photometry of an object material for a particular light source, object surface and viewer geometry. It explicitly incorporates both the geometry and radiometry of image formation into a single model.

If the viewing direction and the direction of incident illumination are known, then expressions for $\cos(i)$, $\cos(e)$ and $\cos(g)$ can be derived in terms of gradient space coordinates p and q . Suppose vector $[p_s, q_s, -1]$ defines the direction to the light source. Then:

$$\cos(i) = \frac{1 + pp_s + qq_s}{\sqrt{1 + p_s^2 + q_s^2} \sqrt{1 + p^2 + q^2}}$$

$$\cos(e) = \frac{1}{\sqrt{1 + p^2 + q^2}}$$

$$\cos(g) = \frac{1}{\sqrt{1 + p_s^2 + q_s^2}}$$

These expressions allow one to transform an arbitrary surface photometric function $\phi(i, e, g)$ into a reflectance map function $R(p, q)$.

Reflectance maps can be determined empirically, derived from phenomenological models of surface reflectivity or derived from analytic models of surface microstructure. One simple, idealized model of surface reflectance is given by:

$$\phi_a(i, e, g) = \rho \cos(i)$$

This reflectance function corresponds to the phenomenological model of a perfectly diffuse (Lambertian) surface which appears equally bright from all viewing directions. Here, ρ is a reflectance factor and the cosine of the incident angle accounts for the foreshortening of the surface as seen from the source. The corresponding reflectance map is given by:

$$R_a(p, q) = \frac{\rho (1 + pp_s + qq_s)}{\sqrt{1 + p_s^2 + q_s^2} \sqrt{1 + p^2 + q^2}}$$

A second reflectance function, similar to that of materials in the maria of the moon and rocky planets, is given by:

$$\phi_b(i, e, g) = \rho \cos(i) / \cos(e)$$

This reflectance function corresponds to a surface which reflects equal amounts of light in all

directions. The cosine of the emergent angle accounts for the foreshortening of the surface as seen from the viewer. The corresponding reflectance map is given by:

$$R_D(p, q) = \frac{\rho (1 + pp_s + qq_s)}{\sqrt{1 + p_s^2 + q_s^2}}$$

Reflectance maps are independent of the shape of the objects being viewed. To emphasize that a reflectance map is not an image, it is convenient to present $R(p, q)$ as a series of "iso-brightness" contours in gradient space. Figure 3 and figure 4 illustrate the two simple reflectance maps $R_a(p, q)$ and $R_b(p, q)$, defined above, for the case $p_s = 0.7$ $q_s = 0.3$ and $\rho = 1$.

3. Reflectance Map Techniques

Using the reflectance map, the basic equation describing the image-forming process can be written as:

$$I(x, y) = R(p, q) \quad (1)$$

This equation has been used in image analysis to explore the relationship between image intensity and object relief. Direct solution is tedious. One can think of (1) as one equation in the two unknowns p and q . Determining object shape from image intensity is difficult because (1) is underdetermined. In order to calculate shape, additional assumptions must be invoked.

Recent work has helped to make these assumptions explicit. For certain materials, such as the material of the maria of the moon, special properties of the reflectance map simplify the solution [Rindfleisch 66] [Horn 75] [Horn 77]. Other methods for determining object shape from image intensity embody assumptions about surface curvature [Horn 77] [Woodham 77]. Simple surfaces have been proposed for use in computer aided design [Huffman 75]. When properties of surface curvature are known a priori, these can be exploited in image analysis [Woodham 78]. This is useful, for example, in industrial inspection since there are often constraints on surface curvature

imposed by the drafting techniques available for part design and by the fabrication processes available for part manufacture. One purpose of these studies is to deepen our understanding of what can and can not be computed directly from image intensity.

Other reflectance map techniques use (1) directly to generate shaded images of surfaces. This has obvious utility in graphic applications including hill-shading for automated cartography [Horn 76] and video input for a flight simulator [Strat 78]. Synthesized imagery can be registered to real imagery to align images with surface models. This technique has been used to achieve precise alignment of LANDSAT imagery with digital terrain models [Horn & Bachman 77].

4. Photometric Stereo

Photometric stereo is a novel reflectance map technique that uses two or more images to solve (1) directly. The idea of photometric stereo is to vary the direction of incident illumination between successive views while holding the viewing direction constant. Suppose two images $I_1(x,y)$ and $I_2(x,y)$ are obtained by varying the direction of incident illumination. Since there has been no change in the imaging geometry, each picture element (x,y) in the two images corresponds to the same object point and hence to the same gradient (p,q) . The effect of varying the direction of incident illumination is to change the reflectance map $R(p,q)$ that characterizes the imaging situation.

Let the reflectance maps corresponding to $I_1(x,y)$ and $I_2(x,y)$ be $R_1(p,q)$ and $R_2(p,q)$ respectively. The two views are characterized by two independent equations:

$$I_1(x,y) = R_1(p,q) \quad (2)$$

$$I_2(x,y) = R_2(p,q) \quad (3)$$

Two reflectance maps $R_1(p,q)$ and $R_2(p,q)$ are required. But, if the phase angle g is the same in both views (i.e., the illumination is simply rotated about the viewing direction), then the two

reflectance maps are rotations of each other.

For reflectance characterized by $R_p(p, q)$ above, (2) and (3) are linear equations in p and q . If the reflectance factor ρ is known, two views are sufficient to uniquely determine surface orientation at each image point, provided the directions of incident illumination are not collinear. Here, (2) and (3) suggest that a 90° angle between the directions of incident illumination would be optimal for photometric stereo.

In general, equations (2) and (3) are nonlinear so that more than one solution is possible. One idea would be to obtain a third image:

$$I_3(x, y) = R_3(p, q) \quad (4)$$

to overdetermine the solution.

For reflectance characterized by $R_a(p, q)$ above, three views are sufficient to uniquely determine both the surface orientation and the reflectance factor ρ at each image point [Horn 78].

Let $I = [I_1, I_2, I_3]'$ be the column vector of intensity values recorded at a point (x, y) in each of the three views (' denotes vector transpose). Further, let

$$\begin{aligned} n_1 &= [n_{11}, n_{12}, n_{13}]' \\ n_2 &= [n_{21}, n_{22}, n_{23}]' \\ n_3 &= [n_{31}, n_{32}, n_{33}]' \end{aligned}$$

be unit column vectors which point in the direction of the three positions of the incident illumination. Construct the matrix N where:

$$N = \begin{bmatrix} n_{11} & n_{12} & n_{13} \\ n_{21} & n_{22} & n_{23} \\ n_{31} & n_{32} & n_{33} \end{bmatrix}$$

Let $n = [n_1, n_2, n_3]'$ be the column vector corresponding to a unit surface normal at (x, y) .

Then,

$$I = \rho N n$$

so that,

$$\rho n = N^{-1} I$$

provided the inverse N^{-1} exists. This inverse exists if and only if the three vectors n_1 , n_2 and n_3 do not lie in a plane. In this case, the reflectance factor and unit surface normal at (x, y) are given by:

$$\begin{aligned} \rho &= |N^{-1} I| && \text{and} \\ n &= (1/\rho)N^{-1} I && (5) \end{aligned}$$

Unfortunately, since the sun's path across the sky is planar, this simple solution does not apply to outdoor images taken at different times during the same day.

Equation (5) suggests that three mutually orthogonal directions of incident illumination would be optimal for lambertian reflectance. In any stereo technique, however, there is some trade-off to acknowledge. In photometric stereo, choosing a larger phase angle g leads to more accurate solutions. At the same time, a larger phase angle causes a larger portion of gradient space to lie in the shadow region of one or more of the sources. A practical compromise is achieved by using four light sources and a relatively large phase angle [Woodham 78]. Solutions are accurate and most of gradient space lies in regions illuminated by at least three of the sources. Three image intensity measurements overdetermine the set of equations and establish a unique solution.

The images required for photometric stereo can be obtained by explicitly moving a single light source, by using multiple light sources calibrated with respect to each other or by rotating the object surface and imaging hardware together to simulate the effect of moving a single light source. The equivalent of photometric stereo can also be achieved in a single view by using multiple illuminations which can be separated by color.

Photometric stereo is fast. It has been developed as a practical scheme for environments in which the nature and position of the incident illumination is known or can be controlled. Initial computation is required to determine the reflectance map for a particular experimental situation. Once calibrated, however, photometric stereo can be reduced to simple table lookup and/or search

operations.

5. Applications of Photometric Stereo

Photometric stereo can be used in two ways. First, photometric stereo is a general technique for determining surface orientation at each image point. For a given image point (x, y) , the equations characterizing each image can be combined to determine the corresponding gradient (p, q) .

Second, photometric stereo is a general technique for determining object points that have a particular surface orientation. This use of photometric stereo corresponds to interpreting the basic image-forming equation (1) as one equation in the unknowns x and y . For a given gradient (p, q) , the equations characterizing each image can be combined to determine corresponding object points (x, y) .

This latter use of photometric stereo is appropriate for the so called industrial bin-of-parts problem. The location in an image of "key" object points is often sufficient to determine the position and orientation of a known object tossed onto a table or conveyor belt.

A particularly useful special case concerns object points whose surface normal directly faces the viewer (i.e. object points with $p = 0$ and $q = 0$). Such points form a unique class of image points whose intensity value is invariant under rotation of the incident illumination about the viewing direction. Object points with surface normal directly facing the viewer can be located without explicitly determining the reflectance map $R(p, q)$. Whatever the nature of the function $R(p, q)$, the value of $R(0, 0)$ is not changed by varying the direction of incident illumination, provided only that the phase angle g is held constant.

These applications of photometric stereo are illustrated using a simple, synthesized example. Consider a sphere of radius r centered at the object space origin. The explicit representation of this object surface, corresponding to the viewing geometry of figure 2(b), is given by:

$$z = f(x,y) = -\sqrt{r^2 - x^2 - y^2} \quad (6)$$

The gradient coordinates p and q are determined by differentiating (6) with respect to x and y . One finds:

$$p = -x/z \quad \text{and} \quad q = -y/z$$

Suppose that the sphere is made of a perfectly diffusing object material and is illuminated by a single distant point source at gradient point (p_s, q_s) . Then, the reflectance map is given by $R_a(p, q)$ above so that the corresponding image is:

$$I(x,y) = \begin{cases} 0 & \text{if } x^2 + y^2 > r^2 \\ \max\{0, R_a(-x/z, -y/z)\} & \text{otherwise} \end{cases} \quad (7)$$

Equation (7) generates image intensities in the range 0 to ρ . In the example below, $r = 60$ and $\rho = 1$.

Multiple images are obtained by varying the position of the light source. Consider three different positions. Let the first be $p_s = 0.7$ and $q_s = 0.3$, as above. Let the second and third correspond to rotations of the light source about the viewing direction of -120° and $+120^\circ$ respectively (i.e., $p_s = -0.610$, $q_s = 0.456$ and $p_s = -0.090$, $q_s = -0.756$). Let the three reflectance maps be $R_1(p, q)$, $R_2(p, q)$ and $R_3(p, q)$. The phase angle g is constant in each case. Let the corresponding images generated by (6) be $I_1(x, y)$, $I_2(x, y)$ and $I_3(x, y)$.

First, consider image point $x = 15$, $y = 20$. Here, $I_1(x, y) = 0.942$, $I_2(x, y) = 0.723$ and $I_3(x, y) = 0.505$. Figure 5 illustrates the reflectance map contours $R_1(p, q) = 0.942$, $R_2(p, q) = 0.723$ and $R_3(p, q) = 0.505$. The point $p = 0.275$, $q = 0.367$ at which these three contours intersect determines the gradient corresponding to image point $x = 15$, $y = 20$.

Second, consider gradient point $p = 0.5$, $q = 0.5$. Here, $R_1(p, q) = 0.974$, $R_2(p, q) = 0.600$ and $R_3(p, q) = 0.375$. Figure 6 illustrates the image intensity contours $I_1(x, y) = 0.974$, $I_2(x, y) = 0.600$ and $I_3(x, y) = 0.375$. The point $x = 24.5$, $y = 24.5$ at which these three contours intersect determines an object point whose gradient is $p = 0.5$, $q = 0.5$.

Finally, Figure 7 repeats the example given in Figure 6 but for the case $p = 0$, $q = 0$. Here, $R_1(p, q) = R_2(p, q) = R_3(p, q) = 0.796$. Object points with surface normal directly facing the viewer form a unique class of points whose image intensity is invariant for rotations of the light source about the viewing direction. The point $x = 0$, $y = 0$ at which these three contours intersect determines an object point with surface normal directly facing the viewer. This result holds even if the form of $R(p, q)$ is unknown.

6. Conclusions

Surface orientation at a point on an object can be determined from the image intensities obtained under a fixed imaging geometry but with varying lighting conditions. Because the images used in photometric stereo arise from the same viewpoint, there is no difficulty determining points in one image which correspond to the same object point in another image. The calculations required in photometric stereo are local and can be implemented in a straightforward manner. Photometric stereo does not depend on the recognition of features in a scene but instead makes use of the intensity values recorded at each picture element.

7. References

- Horn, B. K. P. (1975), "Obtaining Shape from Shading Information", in *The Psychology of Computer Vision*, P. H. Winston (ed.), McGraw-Hill, pp 115-155, 1975.
- Horn, B. K. P. (1976), "Automatic Hill-Shading Using the Reflectance Map", (unpublished), 1976.
- Horn, B. K. P. (1977), "Understanding Image Intensities", in *Artificial Intelligence*, Vol 8, pp 201-231, 1977.
- Horn, B. K. P. & Bachman, B. L. (1977), "Using Synthetic Images to Register Real Images with Surface Models", AI Memo 437, M.I.T. AI Laboratory, August 1977.
- Horn, B. K. P. (1978), "Three Source Photometry", (personal communication), 1978.
- Huffman, D. A. (1975), "Curvature and Creases: A Primer on Paper", in *Proc. of Conf. on Computer Graphics, Pattern Recognition and Data Structures*, pp 360-370, 1975.
- Mackworth, A. K. (1973), "Interpreting Pictures of Polyhedral Scenes", in *Artificial Intelligence*, Vol 4, pp 121-137, 1973.
- Nicodemus, F. E., Richmond, J. C. & Hsia, J. J. (1977), "Geometrical Considerations and Nomenclature for Reflectance", NBS Monograph 160, National Bureau of Standards, Washington, D.C., October 1977.
- Rindfleisch, T. (1966), "Photometric Method for Lunar Topography", *Photogrammetric Engineering*, Vol 32, pp 262-276, March 1966.
- Strat, T. M. (1978), "Shaded Perspective Images of Terrain", AI Memo 463, M.I.T. AI Laboratory, March 1978.
- Woodham, R. J. (1977), "A Cooperative Algorithm for Determining Surface Orientation from a Single View", in *Proceedings of IJCAI-77*, pp 635-641, August 1977.
- Woodham, R. J. (1978), "Reflectance Map Techniques for Analyzing Surface Defects in Metal Castings", TR-457, M.I.T. AI Laboratory, (in press).




Figure 1 Defining the three photometric angles i , e and g . The incident angle i is the angle between the incident ray and the surface normal. The view angle e is the angle between the emergent ray and the surface normal. The phase angle g is the angle between the incident and emergent rays.




Figure 2 Characterizing image projections. Figure 2(a) illustrates the well-known perspective projection. [Note: to avoid image inversion, it is convenient to assume that the image plane lies in front of the lens rather than behind it.] For objects that are small relative to the viewing distance, the image projection can be modeled as the orthographic projection illustrated in figure 2(b). In an orthographic projection, the focal length f is infinite so that all rays from object to image are parallel.




Figure 3 The reflectance map $R_a(p, q)$ for a Lambertian surface illuminated from gradient point $p_s = 0.7$ and $q_s = 0.3$ (with $\rho = 1.0$). The reflectance map is plotted as a series of contours (spaced 0.1 units apart).




Figure 4 The reflectance map $R_b(p, q)$ for a perfect diffusing surface illuminated from gradient point $p_s = 0.7$ and $q_s = 0.3$ (with $\rho = 1.0$). The reflectance map is plotted as a series of contours (spaced 0.2 units apart).




Figure 5 Determining the surface orientation (p, q) at a given image point (x, y) . Three (superimposed) reflectance map contours are intersected where each contour corresponds to the intensity value at (x, y) obtained from three separate images (taken under the same imaging geometry but with different light source position). $I_1(x, y) = 0.942$, $I_2(x, y) = 0.723$ and $I_3(x, y) = 0.505$.




Figure 6 Determining image points (x,y) whose surface orientation is a given gradient (p,q) . Three (superimposed) image intensity contours are intersected where each contour corresponds to the value at (p,q) obtained from three separate reflectance maps. (Each reflectance map characterizes the same imaging geometry but corresponds to a different light source position.) $R_1(p,q) = 0.974$, $R_2(p,q) = 0.600$ and $R_3(p,q) = 0.375$.




Figure 7 Determining image points whose surface normal directly faces the viewer. Three (superimposed) image intensity contours are intersected where each contour corresponds to the value at $(0,0)$ obtained from three separate reflectance maps. (Each reflectance map characterizes the same imaging geometry but corresponds to a different light source position.) Note that the reflectance map value at $(0,0)$ does not change with light source position (provided the phase angle g is held constant).