

Understanding Image Intensities¹

Berthold K. P. Horn

Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139, U.S.A.

Recommended by Max Clowes

ABSTRACT

Traditionally, image intensities have been processed to segment an image into regions or to find edge-fragments. Image intensities carry a great deal more information about three-dimensional shape, however. To exploit this information, it is necessary to understand how images are formed and what determines the observed intensity in the image. The gradient space, popularized by Huffman and Mackworth in a slightly different context, is a helpful tool in the development of new methods.

0. Introduction and Motivation

The purpose of this paper is to explore some of the puzzling phenomena observed by researchers in computer vision. They range from the effects of mutual illumination to the characteristic appearance of metallic surfaces—subjects which at first glance may seem to take us away from the central issues of artificial intelligence. But surely if artificial intelligence research is to claim victory over the vision problem, then it has to embrace the whole domain, understanding not only the problem solving aspects, but also the physical laws that underlie image formation and the corresponding symbolic constraints that enable the problem solving.

One reason for previous neglect of the image itself was the supposition that the work must surely already have been done by researchers in image processing, pattern recognition, signal processing and allied fields. There are several reasons why this attitude was misadvised:

Image processing deals with the conversion of images into new images, usually for human viewing. Computer image understanding systems, on the other hand, must work toward symbolic descriptions, not new images.

¹ This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's research is provided in part by the Advanced Research Projects Agency of the Department of Defence under Office of Naval Research contract N00014-75-C-0643.

Artificial Intelligence 8 (1977), 201-231

Copyright © 1977 by North-Holland Publishing Company

Pattern recognition, when concerned with images, has concentrated on the classification of images of characters and other two-dimensional input, often of a binary nature. Yet the world we want to understand is three-dimensional and the images we obtain have many grey-levels.

Signal processing studies the characteristics of transformations which are amenable to mathematical analysis, not the characteristics imposed on images by nature. Yet in the end, the choice of what to do with an image must depend on it alone, not the character of an established technical discipline.

Although we can borrow some of the techniques of each of these approaches we must still understand how the world forms an image if we are to make machines see. Yet I do not mean to suggest analysis by synthesis. Nothing of the sort! I propose only that if we are to solve the problem of creating computer performance in this domain, we must first thoroughly understand that domain.

This is, of course, not without precedent. The line of research beginning with Guzman and continuing through Clowes, Huffman, Waltz, Mackworth and others, was a study of how the physical world dictates constraints on what we see—constraints that once understood can be turned around and used to analyze what is seen with great speed and accuracy relative to older techniques which stressed problem solving expertise at the expense of domain understanding.

1. Developing the Tools for Image Understanding

An understanding of the visual effects of edge imperfections and mutual illumination will be used to suggest interpretations of image intensity profiles across edges, including those that puzzled researchers working in the blocks world. We shall see that a "sharp peak" or edge-effect implies that the edge is convex, a "roof" or triangular profile suggests a concave edge, while a step-transition or discontinuity accompanied by neither a sharp peak nor a roof component is probably an obscuring edge. This last hypothesis is strengthened significantly if an "inverse peak" or negative edge-effect is also present. (See Section 3.)

Next, it will be shown that the image intensities of regions meeting at a joint corresponding to an object's corner determine fairly accurately the orientation of each of the planes meeting at the corner. Thus we can establish the three-dimensional structure of a polyhedral scene without using information about the size or, support or nature of the objects being viewed. (See Section 3.4.)

Finally, we will turn to curved objects and show that their shape often can be determined from the intensities recorded in the image. The approach given here is supported by geometric arguments and does not depend on methods for solving first-order non-linear partial differential equations. It combines my previous shape-from-shading method [4, 2] with geometric arguments in gradient-space (Huffman and Mackworth [1, 2, 3, 9]). This approach to the image analysis problem enables us to establish whether or not certain features can be extracted from images. (See Sections 4 and 5.)

1.1. Image formation

The visual world consists of opaque bodies immersed in a transparent medium. The dimensionality of the two domains match: since only the object's surfaces are important for recognition and description purposes. On one hand we have two-dimensional surfaces plus depth and on the other, a two-dimensional image plus intensity. There are two parts to the problem of exploiting this observation to understand what is being imaged: one deals with the geometry of projection, the other with the intensity of light recorded in the image.

The relation between object coordinates and image coordinates is given by the well-known perspective projection equations derived from a diagram such as Fig. 1, where f if the focal length and,

$$x' = (x/z)f \text{ and } y' = (y/z)f$$

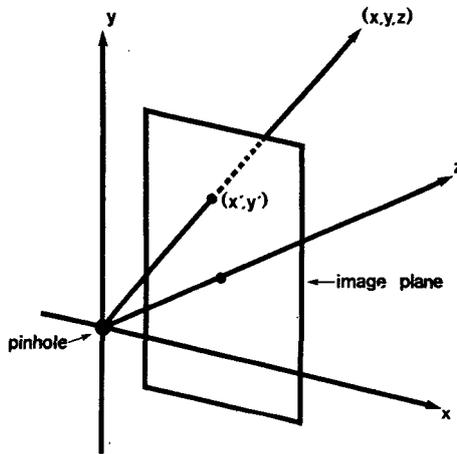


FIG. 1. Image projection geometry.

For the development presented here it will be convenient to concentrate on the case where the viewer is very far from the objects relative to their size. The resultant scene occupies a small visual angle as if viewed by a telephoto lens. This corresponds to orthographic projection, where z is considered constant in the equations above.

1.2. Surface orientation

We must understand the geometry of the rays connecting the lightsource(s), the object and the viewer in order to determine the light flux reflected to the viewer from a particular element of the object. The surface orientation in particular, plays a major role. There are, of course, various ways of specifying the surface orientation of a plane. We can give, for example, the equation defining the plane or the direction of a vector perpendicular to the surface. If the equation for the plane is $ax + by + cz = d$, then a suitable surface normal is (a, b, c) .

We extend this method to curved surfaces simply by applying it to tangent planes. A local normal to a smooth surface is $(z_x, z_y, -1)$, where z_x and z_y are the first

partial derivatives of z with respect to x and y . It is convenient to use the abbreviations p and q for these quantities. The local normal then becomes $(p, q, -1)$. It is clear then that the surface orientation thus defined has but two degrees to freedom. The quantity (p, q) is called the *gradient* (Fig. 2).

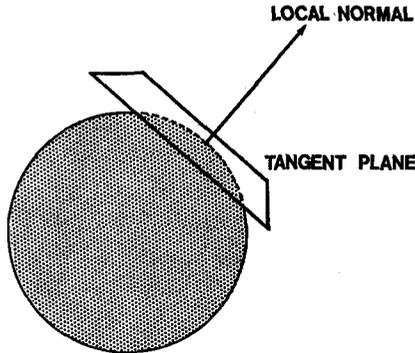


FIG. 2. Definition of surface orientation for curved objects.

1.3. Image intensity

The amount of light reflected by a surface element depends on its micro-structure and the distribution of incident light. Constructing a tangent plane to the object's surface at the point under consideration, we see that light may be arriving from directions distributed over a hemisphere. We first consider the contributions separately, from each of these directions and then superimpose the results.

For most surfaces there is a *unique value of reflectance and hence image intensity for a given surface orientation*. No matter how complex the distribution of light sources. We shall spend some time exploring this and develop the reflectance map in the process.

The simplest case is that of a single point-source where the geometry of reflection is governed by three angles: the incident, the emittance and the phase angles (Fig. 3). The incident angle, i , is the angle between the incident ray and the local normal,

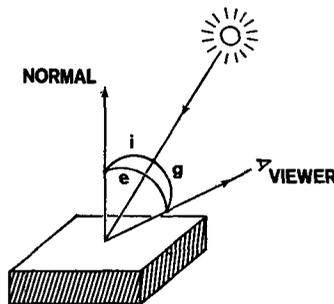


FIG. 3. The reflectivity is a function of the incident, emittance, and phase angles.

the emittance angle, e , is the angle between the ray to the viewer and the local normal, while the phase angle, g , is the angle between the incident and the emitted ray [4]. The reflectivity function is a measure of how much of the incident light is reflected in a particular direction. Superficially, it is the fraction of the incident light reflected per unit surface area, per unit solid angle in the direction of the viewer. More precisely:

Let the illumination be E (flux/area) and the resulting surface luminance in the direction of the viewer by B (flux/steradian/projected area). (The projected area is the equivalent foreshortened area as seen by the viewer.) *The reflectivity is then defined to be B/E .* It may be written $\phi(i, e, g)$, to indicate its dependence on the three angles involved.

Note that an infinitesimal surface element, dA , captures a flux $E \cos(i) dA$, since its surface normal is inclined i relative to the incident ray. Similarly, the intensity I (flux/steradian) equals $B \cos(e) dA$, since the projected area is foreshortened by the inclination of the surface normal relative to the emitted ray.

1.4 Reflectivity function

For some surfaces, mathematical models have been constructed that allow an analytical determination of the reflectivity function. Since such techniques have rarely proved successful, reflectivity functions are usually determined empirically. Often there will be more than one source illuminating the object. In this case one has to integrate the product of reflectivity and the incident light intensity per unit solid angle over the hemisphere visible from the point under consideration. This determines the total light flux reflected in the direction of the viewer.

The normal to the surface relates object geometry to image intensity because it is defined in terms of the surface geometry, yet it also appears in the equation for the reflected light intensity. Indeed two of the three angles on which the reflectivity function depends are angles between the normal and other rays. Although we could now proceed to develop a partial differential equation based on this observation, it is more fruitful to introduce first another tool—gradient space.

1.5. Gradient space

Gradient-space can be derived as a projection of dual-space or of the Gaussian sphere but it is easier here to relate it directly to surface orientation [2]. We will concern ourselves with orthographic projection only, although some of the methods can be extended to deal with perspective projection as well.

The mapping from surface orientation to gradient-space is made by constructing a normal $(p, q, -1)$ at a point on an object and mapping it into the point (p, q) in gradient-space. Equivalently, one can imagine the normal placed at the origin and find its intersection with a plane at unit distance from the origin.)

We should look at some examples in order to gain a feel for gradient-space. Parallel planes map into a single point in gradient-space. Planes perpendicular to

the view-vector map into the point at the origin of gradient-space. Moving away from the origin in gradient-space, we find that the distance from the origin equals the tangent of the emittance angle e , between the surface normal and the view-vector.

If we rotate the object-space about the view-vector, we induce an equal rotation of gradient-space about the origin. This allows us to line up points with the axes and simplify analysis. Using this technique, it is easy to show that the angular position of a point in gradient-space corresponds to the direction of steepest descent on the original surface.

Let us call the orthogonal projection of the original space the image-plane. Usually this is all that is directly accessible. Now consider two planes and their intersection. Let us call the projection of the line of intersection the image-line. The two planes, of course, also correspond to two points in gradient space. Let us call the line connecting these two points the gradient-line. Thus, a line maps into a line. The perpendicular distance of the gradient-space line from the origin equals the tangent of the inclination of the original line of intersection with respect to the image plane. We show by superimposing gradient-space on the image-space [2, 11] that the gradient-space line and the image-line are *mutually perpendicular*. Mackworth's scheme for scene analysis of line drawings of polyhedra depends on this observation [2].

1.6. Trihedral corners

The points in gradient-space corresponding to the three planes meeting at a trihedral corner must satisfy certain constraints. The lines connecting these points must be perpendicular to the corresponding lines in the image-plane (Fig. 4). This provides

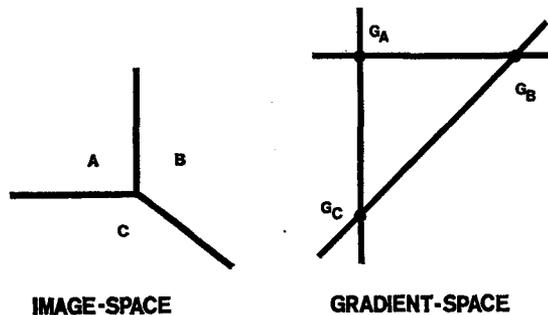


FIG. 4. Constraints on the gradient-space points corresponding to the planes meeting at a trihedral corner. The gradient-lines must be perpendicular to the image-lines.

us with three constraints but that is not enough to fix the position of three points in gradient-space. Three degrees of freedom—the *position* and *scale* of the triangle—remain undetermined. We see later that measurement of image intensities for the three planes provides enough information to specify their orientations, thus allowing a determination of the three-dimensional structure of a polyhedral scene.

2. The Reflectance Map

The amount of light reflected by a given surface element depends on the orientation of the surface and the distribution of light sources around it, as well as on the nature of the surface. For a given type of surface and distribution of light sources, there is a fixed value of reflectance for every orientation of the surface normal and hence for every point in *gradient-space*. Thus image intensity is a single-valued function of p and q . We can think of this as a map in gradient-space. This is *not* a transform of the image seen by the viewer. It is, in fact, independent of the scene and instead, only a function of the surface properties and the light source distribution (but see Section 2.6).

What can we do with this strange "image" of the world surrounding the object? If we measure a certain intensity at a given point on the object, we know that the orientation of the surface at that point is restricted to a subset of all possible orientations; we cannot, however, uniquely determine the orientation. The one constraint is that it be one of the points in gradient-space where we find this same value of intensity.

The use of the gradient-space diagram is analogous to the use of the hodogram or velocity-space diagram. The latter provides insight into the motion of particles in force field that is hard to obtain by algebraic reasoning alone. Similarly, the gradient-space allows geometric reasoning about surface orientation and image intensities.

2.1. Matte surfaces and a point-source near the viewer

A perfect lambertian surface or diffuser looks equally bright from all directions; the amount of light reflected depends only on the cosine of the incident angle. In order to postpone the calculation of incident, emittance and phase angles from p and q , we first place a single light source near the viewer. The incident angle then equals the emittance angle, the angle between the surface normal and the view-vector. The cosine of the incident angle is the dot product of the corresponding unit vectors:

$$\cos(i) = \frac{(p, q, -1) \cdot (0, 0, -1)}{|(p, q, -1)| |(0, 0, -1)|} = \frac{1}{\sqrt{1+p^2+q^2}}.$$

We obtain the same result by remembering that the distance from the origin in gradient-space is the tangent of the angle between the surface normal and the view-vector:

$$\begin{aligned} \sqrt{p^2+q^2} &= \tan(e), \\ \cos(e) &= \frac{1}{\sqrt{1+\tan^2(e)}}, \end{aligned}$$

and

$$e = i.$$

If we plot reflectance as a function of p and q , we obtain a central maximum of 1 at the origin and a circularly symmetric function that falls smoothly to 0 as we approach infinity in gradient-space. This is a nice, smooth reflectance map, typical of matte surfaces (Fig. 5).

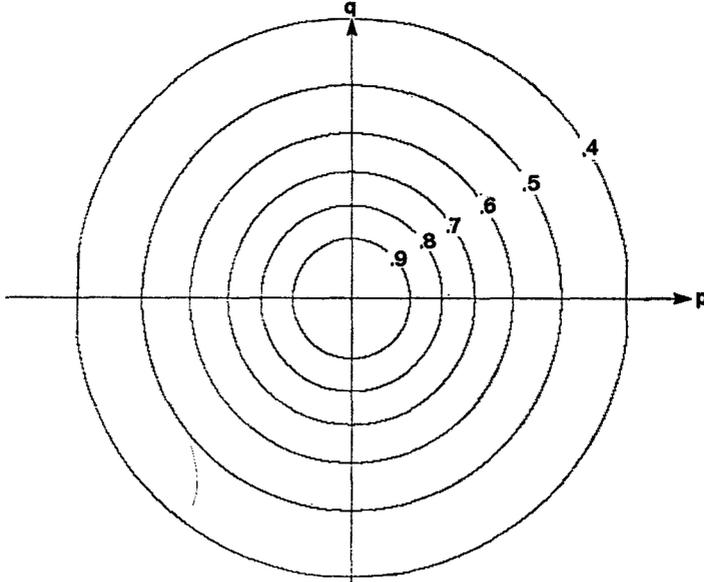


FIG. 5. Contours of $E = \cos(e)$. This is the reflectance map for objects with lambertian surfaces when there is a single light source and the light source is near the viewer.

A given image intensity corresponds to a simple locus in gradient-space, in this case a circle centered on the origin. A measurement of image intensity tells us that the surface gradient falls on a certain circle in gradient-space.

Since the light-source is not always likely to be near the viewer we now explore the more complicated geometry of incident and emitted rays for arbitrary directions of incident light at the object.

2.2. Incident, emittance, and phase angles

For many surfaces the reflectance is a smooth function of the angles of incidence, emittance and phase. It is convenient to work with the cosines of these angles, $I = \cos(i)$, $E = \cos(e)$, and $G = \cos(g)$. (These can be obtained easily from the dot products of the unit vectors.) If we have a single distant light source whose direction is given by a vector $(p_s, q_s, -1)$, and note that the view-vector is $(0, 0, -1)$, then,

$$G = \frac{1}{\sqrt{1+p_s^2+q_s^2}}, \quad E = \frac{1}{\sqrt{1+p^2+q^2}}$$

and

$$I = \frac{(1+p_s p + q_s q)}{\sqrt{1+p^2+q^2} \sqrt{1+p_s^2+q_s^2}} = (1+p_s p + q_s q)EG.$$

It is simple to calculate I , E , and G for any point in gradient-space. In fact G is constant given our assumption of orthogonal projection and distant light source. We saw earlier that the contours of constant E are circles in gradient-space centered on the origin. Setting I constant gives us a second-order polynomial in p and q and suggests that loci of constant I may be conic sections. The terminator—the line separating lighted from shadowed regions, for example, is a straight line obtained by setting $i = \pi/2$. That is, $I = 0$; or $1 + p_s p + q_s q = 0$. Similarly, the locus of $I = 1$ is the single point $p = p_s$ and $q = q_s$.

A geometrical way of constructing the loci of constant I is to develop the cone generated by all directions that have the same incident angle. The axis of the cone is the direction to the light-source $(p_s, q_s, -1)$. We find the corresponding points in gradient-space by intersecting this cone with a plane at unit distance from the origin. Varying values of I produce cones with varying angles. The cones form a nested sheaf. The intersection of this nested sheaf with the unit plane is a nested set of conic sections (Fig. 6). Note that our previous example (Fig. 5) is merely a special case in which the axis of the sheaf of nested cones points directly at the viewer.

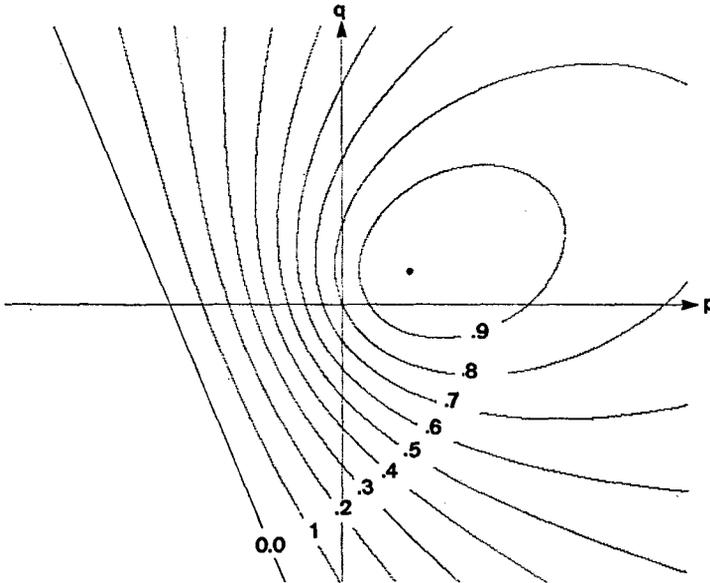


FIG. 6. Contours of $I = \cos(i)$. The direction to the single source is $(p_s, q_s) = (0.7, 0.3)$. This is the reflectance map for objects with lambertian surfaces when the light source is not near the viewer.

If we measure a particular image intensity, we know that the gradient of the corresponding surface element must fall on a particular conic section. The possible normals are confined to a cone—in this case, merely a circular cone. In the case of more general reflectivity functions, the locus of possible normals constitutes a more general figure called the *Monge cone*.

2.3. Specularity and glossy surfaces

Many surfaces have some specular or mirror-like reflection from the outermost layers of their surface, and thus are not completely matte. This is particularly true of surfaces that are smooth on a microscopic scale. For specular reflection $i = e$, and the incident, emitted, and normal vectors are all in the same plane. Alternatively, we can say that $i + e = g$. In any case, only one surface orientation is correct for reflecting the light source towards the viewer. That is, a perfect specular reflection contributes an impulse to the gradient-space image at a particular point.

In practice, few surfaces are perfectly specular. Glossy surfaces reflect some light in directions slightly away from the geometrically correct direction [8]. It can be shown that the cosine of the angle between the direction for perfectly specular reflection and any other direction is $(2IE - G)$ [11]. This clearly equals 1 in the ideal direction and falls off towards 0 as the angle increases to a right-angle. By taking various functions of $(2IE - G)$, such as high powers, one can construct more or less compact specular contributions.

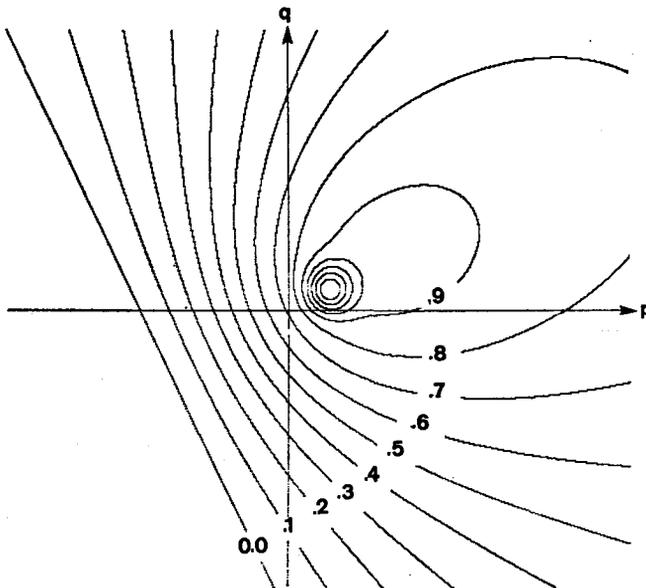


FIG. 7. Contours for $\phi(I, E, G) = \frac{1}{2}s(n+1)(2IE - G)^n + (1-s)I$. This is the reflectance map for a surface with both a matte and a specular component of reflectivity illuminated by a single point-source.

For example, a good approximation for some glossy white paint can be obtained by combining the usual matte component with a specular component defined in this way— $\phi(I, E, G) = \frac{1}{2}s(n+1)(2IE - G)^n + (1-s)I$. Here s lies between 0 and 1 and determines the fraction of incident light reflected specularly before penetrating the surface, while n determines the sharpness of the specularity peak in the gradient-space image (Fig. 7).

2.4. Finding the gradient from the angles

In order to further explore the relation between the specification of surface orientation in gradient-space and the angles involved, we solve for p and q , given I , E , and G . We have already shown that the opposite operation is simple to perform. One approach to this problem is to solve the polynomial equations in p and q derived from the equations for I , E , and G . It can be shown that [11]:

$$p = p' \cos(\theta) - q' \sin(\theta),$$

$$q = p' \sin(\theta) + q' \cos(\theta),$$

where

$$p' = \frac{(I/E - G)}{\sqrt{1 - G^2}} \quad \text{and} \quad q' = \frac{(\Delta/E)}{\sqrt{1 - G^2}},$$

$$\Delta^2 = 1 + 2IEG - (I^2 + E^2 + G^2),$$

$$\cos(\theta) = \frac{p_s}{\sqrt{p_s^2 + q_s^2}} \quad \text{and} \quad \sin(\theta) = \frac{q_s}{\sqrt{p_s^2 + q_s^2}}.$$

It is immediately apparent that there are two solution points in gradient space for most values of I , E , and G . Notice that θ is the direction of the light source in gradient-space, that is, the line connecting (p_s, q_s) to the origin makes an angle θ with the p -axis. So p' and q' are coordinates in a new gradient-space obtained after

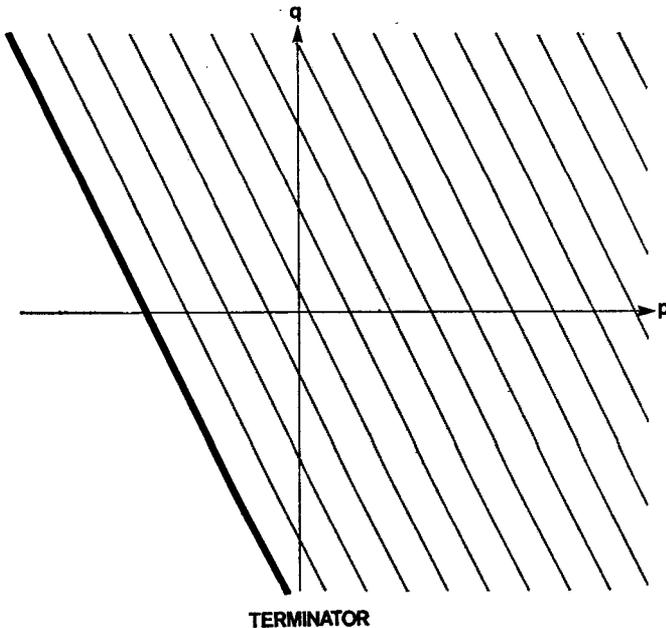


FIG. 8. Contours of $\phi(I, E, G) = I/E$. Contour intervals are 0.2 units wide. The reflectivity function for the material in the maria of the moon is constant for constant I/E (constant luminance longitude).

simplifying matters by rotating the axes until $q_s = 0$. The light source is then in the direction of the x' -axis. Notice that p' is constant if I/E is constant (remembering that G is constant anyway.) Hence the loci of constant I/E are straight lines. These lines are all parallel to the terminator, for which $I = 0$. This turns out to be important since some surfaces have constant reflectance for constant I/E (Fig. 8).

2.5. Smooth metallic surfaces

Consider a smooth metallic surface: a surface with a purely specular or mirror-like reflectance. Each point in gradient-space corresponds to a particular direction of the surface normal and defines a direction from which incident light has to approach

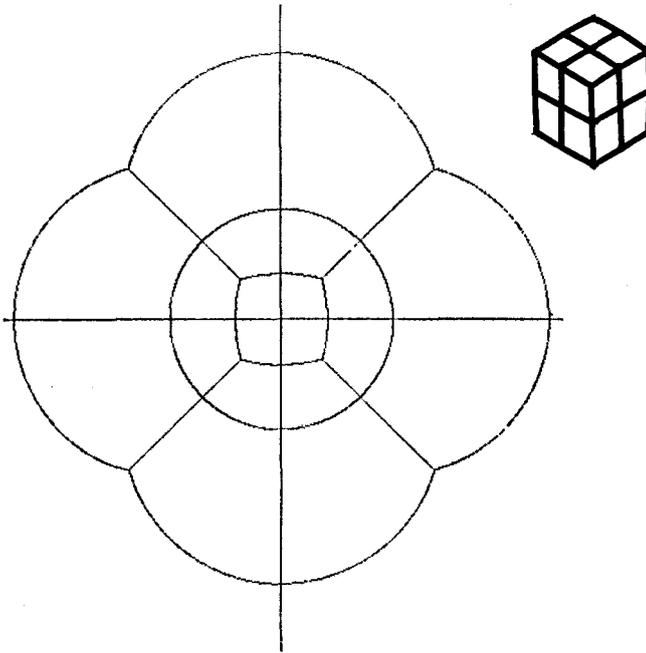


FIG. 9. Reflectance map for a metallic object in the center of a very large wire cube. Equivalently one can think of it as the reflection of the wire cube in a paraboloid with a specularly-reflecting surface.

the object in order to be reflected towards the viewer. In fact, we can produce a complete map of the sphere of possible directions as seen from the object. At the origin, for example, we have the direction towards the viewer. Now for each incident direction there is a certain light-intensity depending on what objects lie that way. Consider recording these intensities at the corresponding points in gradient-space. Clearly one obtains some kind of image of the world surrounding the given metallic object. In fact, one develops a stereographic projection, a plane projection of a sphere with one of the poles as the center of projection. Another way of looking at it is that the image we construct in this fashion is like one we

Artificial Intelligence 8 (1977) 201-231

obtain by looking into a convex mirror—a metallic paraboloid to be precise (Fig. 9).

In order to construct reflectance maps for various surfaces and distributions of light sources, we superimpose the results in gradient-space for each light source in turn. We now examine a flaw in this approach and attempt a partial analysis of mutual illumination.

2.6. Mutual illumination

The reflectance map is based on the assumption that the viewer and all light sources are distant from the object. Only under these assumptions can we associate a unique value of image intensity with every surface orientation. If the scene consists of a single convex object these assumptions will be satisfied, but when there are several highly reflective objects placed near one another, mutual illumination may become important. That is, the distribution of incident light no longer depends only on direction, but is a function of position as well. The general case is very difficult, and we shall only study some idealized situations applicable to scenes consisting of polyhedra.

Two important effects of mutual illumination are a reduction in *contrast* between faces, and the appearance of *shading* or gradation of light on images of plane surfaces. In the absence of this effect, we would expect plane surfaces to give rise to images of uniform intensity since all points on a plane surface have the same orientation.

2.7. Two semi-infinite planes

First, let us consider a highly idealized situation of two semi-infinite planes joined at right angles and a distant light source. Let the incident rays make an angle α with respect to one of the planes (Fig. 10). Assume further that the surfaces reflect

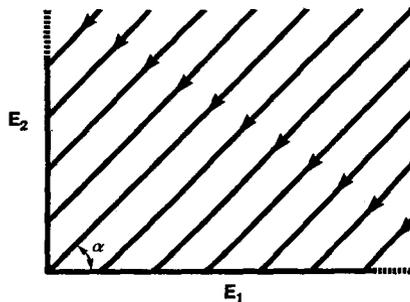


FIG. 10. Mutual illumination of two infinite lambertian planes of reflectance r , when illuminated by a single distant source.

a fraction r of the light falling on them, and that the illumination provided by the source is E (light flux/unit area). Picking any point on one of the half-planes, we find that one-half of its hemisphere of directions is occupied by the other plane,

Artificial Intelligence 8 (1977), 201–231

one-half of the light radiated from this point hits the other plane, while one-half is lost. Since both planes are semi-infinite, the geometry of this does not depend on the distance from the corner. The light incident at any point is made up of two components: that received directly from the source and that reflected from the other plane. The intensity on one plane does not vary with distance from the corner—a point receives light from one-half of its hemisphere of directions no matter what its distance from the corner. Put another way, there is no natural scale factor for a fluctuation in intensity. Let the illumination of the plane be E_1 and E_2 (light flux/unit area); then,

$$\begin{aligned} E_1 &= \frac{1}{2}rE_2 + E \cos(\alpha), \\ E_2 &= \frac{1}{2}rE_1 + E \sin(\alpha). \end{aligned}$$

Solving for E_1 and E_2 , we get:

$$\begin{aligned} E_1 &= E[\cos(\alpha) + \frac{1}{2}r \sin(\alpha)] \frac{1}{[1 - (\frac{1}{2}r)^2]}, \\ E_2 &= E[\sin(\alpha) + \frac{1}{2}r \cos(\alpha)] \frac{1}{[1 - (\frac{1}{2}r)^2]}. \end{aligned}$$

Had we ignored the effects of mutual illumination we would have found $E_1 = E \cos(\alpha)$ and $E_2 = E \sin(\alpha)$. Clearly the effect increases as reflectance r increases (it is not significant for dark surfaces). When the planes are illuminated equally, for $\alpha = \pi/4$, we find

$$E_1 = E_2 = (E/\sqrt{2})/(1 - \frac{1}{2}r).$$

When $r = 1$, we obtain *twice* the illumination and hence twice the brightness than that obtained in the absence of mutual illumination. If the angle between the two planes varies, we find that the effect becomes larger and larger as the angle becomes more and more acute. By choosing the angle small enough, we can obtain arbitrary "amplification". Conversely, for angles larger than $\pi/2$, the effect is less pronounced.

In the derivation above, we did not make very specific assumptions about the angular distribution of reflected light, just that it is symmetrical about the normal and that it does not depend on the direction from which the incident ray comes. Hence, a lambertian surface is included, while a highly specular one is not. The effect is indeed less pronounced for surfaces with a high specular component of reflection, since most of the light is bounced back to the source after two reflections. Another important thing to note is that even if the planes are not infinite, the above calculations are approximately valid close to the corner. For finite planes we expect a variation of intensity as a function of distance from the corner; the results derived here apply asymptotically as one approaches the corner.

2.8. Two truncated planes

The geometry becomes quite complex if the planes are of finite extent, but we can develop an integral equation if we allow the planes to be infinite along their line of intersection and truncate them only in the direction perpendicular to this. Suppose *Artificial Intelligence* 8 (1977), 201–231

two perpendicular planes extend a distance L from the corner, and that $\alpha = \pi/4$. This produces a particularly simple integral equation [11]; nevertheless I have been unable to solve it analytically. Numerical methods show that the resultant illumination falls off monotonically from the corner (Fig. 11), that the value at the corner is

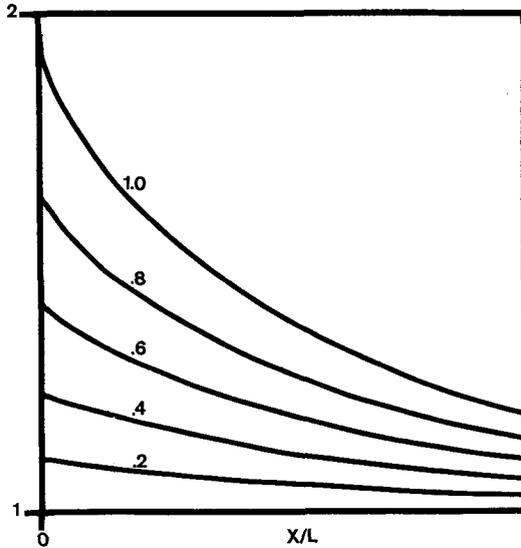


FIG. 11. Surface luminance plotted versus fractional distance from a right-angle corner. The curves are for reflectances of 0.2, 0.4, 0.6, 0.8, and 1.0.

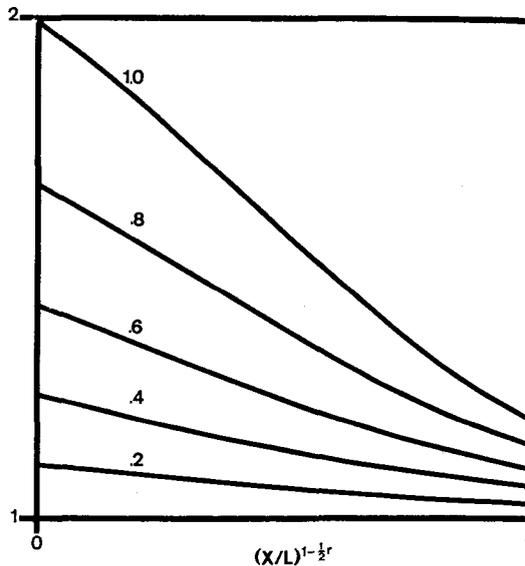


FIG. 12. Surface luminance plotted versus $(x/L)^{1-\frac{1}{2}r}$ to illustrate asymptotic behavior near the corner. The curves correspond to reflectances of 0.2, 0.4, 0.6, 0.8, and 1.0.

Artificial Intelligence 8 (1977), 201-231

edge may be convex and not give rise to a peak) if the line connecting the two points

Artificial Intelligence 8 (1977), 201-231

indeed what was predicted in the previous section, and the fall-off near the corner is governed by a term in $-(x/L)^{(1-\frac{1}{r})}$ (Fig. 12). A good approximation appears to be $1 + E_{2/r}(x/L)$, where $E_n(x)$ is the elliptical integral and x is the distance along the plane measured from the edge where the planes meet.

3. The Semantics of Edge-Profiles

We are now ready to apply the tools developed so far. First let us consider the interpretation of intensity profiles taken across edges. If polyhedral objects and image sensors were perfect, if there were no mutual illumination, and if light sources were distant from the scene, images of polyhedral objects would be divided into polygonal areas, each of uniform intensity. It is well known that in real images, image intensity varies within the polygonal areas and that an intensity profile taken across an edge separating two such polygonal regions does not have a simple step-shaped intensity transition. Herskovitz and Binford determined experimentally that the most common edge transitions are step-, peak-, and roof-shaped [7] (Fig. 13). So

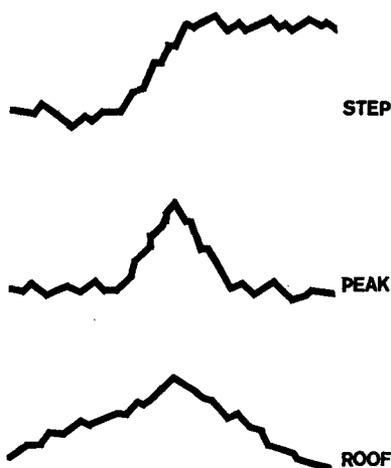


FIG. 13. Most common intensity profiles across images of polyhedral edges.

far this has been considered no more than a nuisance, because it complicates the process of finding edges. We now discuss the interpretation of these profiles in terms of the three-dimensional aspects of the scene.

3.1. Imperfections of polyhedral edges

A perfect polyhedron has a discontinuity in surface normal at an edge. In practice, edges are somewhat rounded off. A cross-section through the object's edge show that the surface normal varies smoothly from one value to the other and takes on values that are linear combinations of the surface normals of the two adjoining planes (Fig. 14). What does this mean in terms of reflected light intensity? Intensity varies smoothly at the edge, instead of jumping from one surface normal value to

Artificial Intelligence 8 (1977), 201-231

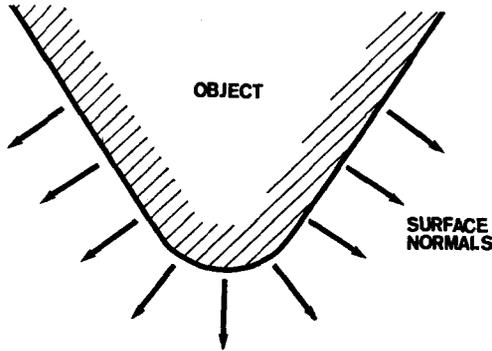


FIG. 14. The normals at an imperfect, rounded off edge are positive linear combinations of the normals of the two adjoining faces.

another. The important point is that it may take on values *outside* the range of those defined by the two planes. The best way to see this is to consider the situation in gradient-space. The two planes defined two points in gradient-space. Tangent planes on the corner correspond to points on the line connecting these two points. If the image intensity is higher for a point somewhere on this line, we will see a peak in the intensity profile across the edge. (Fig. 15.)

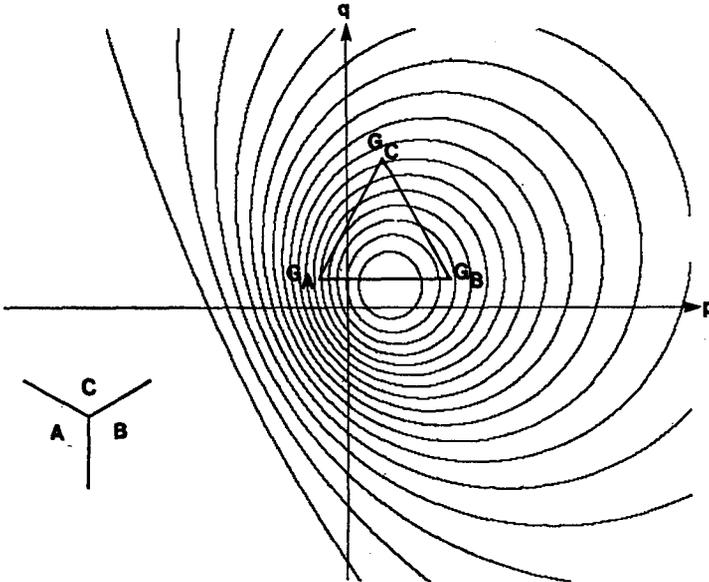


FIG. 15. Image of tri-hedral corner and corresponding gradient-space diagram. The image intensity profile across the edge between face A and face B will have a peak or highlight. The others will not.

If we find an edge profile with a peak shape or step with a peak superimposed, it is most likely that the corresponding edge is convex. The converse is not true (an edge may be convex and not give rise to a peak) if the line connecting the two points

in gradient space has intensity varying monotonically along its length. The identification is also not completely certain since under peculiar lighting conditions with objects that have acute angles between adjacent faces, a peak may appear at an obscuring edge. Notice that the peak is quite compact, since it extends only as far as the rounded-off edge.

At a corner, where the planes meet, we find that surface imperfections provide surface normals that are linear combinations of the three normals corresponding to the three planes. In gradient space, this corresponds to points in the triangle connecting the three points corresponding to the planes. If this triangle contains a maximum in image intensity we expect to see a highlight right on the corner (Fig. 15).

3.2. Mutual illumination

We have seen that mutual illumination gives rise to intensity variations on planar surfaces—intensity roughly decreases linearly away from the corner. Notice that this affects the intensity profile over a large distance from the edge, quite unlike the sharp peak found due to edge imperfections. Clearly, if we find a roof-shaped profile or step with a roof-shaped superimposed we should consider labelling the edge concave.

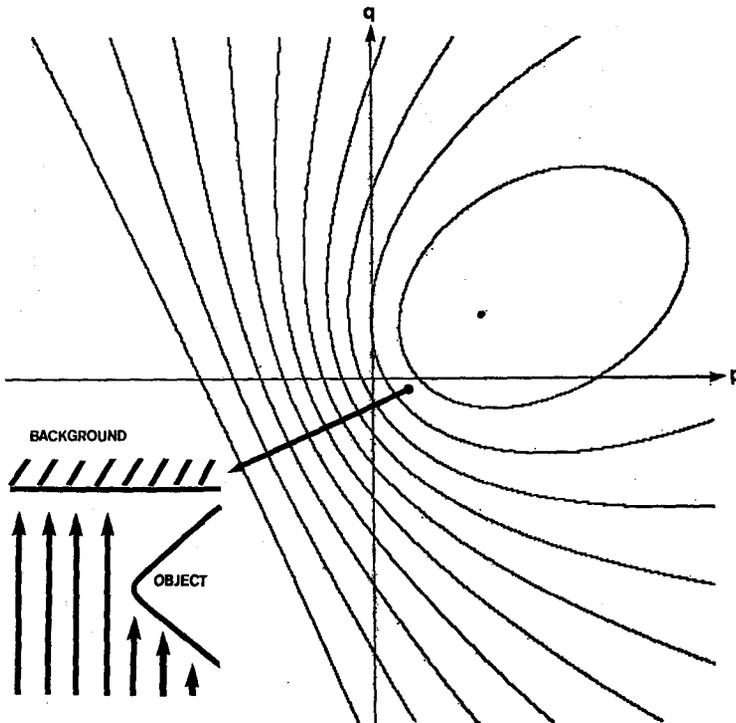


FIG. 16. Generation of a negative peak at an obscuring edge facing away from the light source. *Artificial Intelligence* 8 (1977), 201-231

The identification is, however, partly unreliable since some imaging device defects can produce a similar effect. Image dissectors, for example suffer from a great deal of scattering—areas further from a dark background appear brighter. So we may see a smoothed version of a roof-shape in the middle of a bright scene against a dark background. Experimentation with high quality image input devices such as a PIN-diode mirror-deflection system has confirmed that this is an artifact introduced by the image dissector. When the light source is close to the scene, significant gradients can appear on planar surfaces as pointed out by Herskovitz and Binford [7]. Lastly, the roof-shaped profiles on the two surfaces may be due to mutual illumination with other surfaces, not each other. Nevertheless, a roof-shaped profile usually suggests a concave edge.

3.3. Obscuring edges

Although they can be found with both convex and concave edges, step-shaped intensity profiles occur most often when objects obscure one another. If the obscuring surface adjoins a self-shadowed surface however, edge imperfections will produce a negative peak on the profile, since the line connecting the points corresponding to the two surfaces of the object passes through the terminator in gradient-space. Hence a negative peak or a step with a superimposed negative peak strongly suggests obscuration. Unfortunately it is difficult to tell on which side the obscuring plane lies (Fig. 16).

3.4. Determining the three-dimensional structure of polyhedral scenes

Mackworth's approach to understanding line drawings of polyhedra allows us to take into account some of the quantitative aspects of the three-dimensional geometry of scenes [2]. It does not, however, allow us to fully determine the orientation of all the planes. The scale and position of the gradient space diagram is undetermined by his technique. To illustrate this, consider a single trihedral corner. Here we know that the three points in gradient space representing the three planes meeting at the corner have to satisfy certain constraints: they must lie on three lines perpendicular to the image lines (see fig. 4).

It takes six parameters to specify the position of three points on a plane, leaving three degrees of freedom after the introduction of these constraints. Measuring the three image intensities of the planes supplies another three. The constraints are due to the fact that the points in gradient-space have to lie on the correct contours of image intensity. The triangle can be stretched and moved until the points correspond to the correct image intensities as measured for the three planes.

Since this process corresponds to solving three non-linear equations for three unknowns, we can expect a finite number of solutions. Often there are but one or two—prior knowledge often eliminates some. (For a numerical example, see Appendix).

When more than three planes meet at a corner, the equations are over-determined,

Artificial Intelligence 8 (1977), 201–231

and the situation is even more constraint. Conversely, we cannot determine much with just two planes meeting at an edge—there are too few equations and an infinite number of solutions. The possible ambiguity at a trihedral corner is not very serious when we consider that in a typical scene there are many “connect” edges, either convex or concave as determined by Mackworth’s program. In such a case the overall constraints may allow only one consistent interpretation. A practical difficulty is that it is unclear which search strategy leads most efficiently to this interpretation.

Measurements of image intensity are not very precise and surfaces have properties that vary from point to point as well as with handling. We cannot expect this method to be extremely accurate in pinning down surface orientation. In general, the equations for a typical scene are over-determined; a least-squares approach may improve matters slightly. The idea of stretching and shifting can be generalized to smooth surfaces. We know that the image of a paraboloid is the same as the reflectance map. If we can stretch and shift the reflectance map to fit the image of some object, then the same deformations turn the paraboloid into the object.

4. Determining Lunar Topography

When viewed from a great distance, the material in the maria of the moon has a particularly interesting reflectivity function. First, note that the lunar phase is the angle at the moon between the light source (sun) and the viewer (earth). This is clearly the angle we call g , and explains why we use the term phase angle. For constant phase angle, detailed measurements using surface elements whose projected area as seen from the source is a constant multiple of the projected area as seen by the viewer have shown that all such surface elements have the same reflectance. But the area appears foreshortened by $\cos(i)$ and $\cos(e)$ as seen by the source and the viewer respectively. Hence the reflectivity function is constant for constant $\cos(i)/\cos(e) = I/E$.

In this case each surface element scatters light uniformly into its hemisphere of directions, quite unlike the lambertian surface which favors directions normal to its surface. This is not an isolated incident. The surfaces of other rocky, dusty objects when viewed from great distances appear to have similar properties. For example, the surface of the planet Mercury and perhaps Mars as well as some asteroids and atmosphere-free satellites fit this pattern. Surfaces with reflectance a function of I/E thus form a third species we should add to matte surfaces where the reflectance is a function of I and glossy surfaces where the reflectance is a function of $(2IE - G)$.

4.1. Lunar reflectivity function

Returning to the lunar surface, we find an early formula due to Lommel Seelinger [6].

$$\phi(I, E, G) = \frac{\Gamma_0(I/E)}{(I/E) + \lambda(G)}$$

Here Γ_0 is a constant and the function $\lambda(G)$ is defined by an empirically determined table. A somewhat more satisfactory fit to the data is provided by a formula of Fesenkov's [6]:

$$\phi(I, E, G) = \frac{\Gamma_0(I/E)[1 + \cos^2(\alpha/2)]}{(I/E) + \lambda_0[1 + \tan^2(\lambda/2)]}$$

Where Γ_0 and λ_0 are constants and $\tan(\alpha) = -p' = -(I/E - G)/\sqrt{1 - G^2}$. This formula is also supported by a theoretical model of the surface due to Hapke. Note that given I , E , and G , it is straightforward to calculate the expected reflectance. We need to go in the reverse direction and solve for I/E given G and the reflectance as measured by the image intensity. While it may be hard to invert the above equation analytically, it should be clear that by some iterative, interpolation, or hill-climbing scheme one can solve for I/E . We shall ignore for now the ambiguities that arise if there is more than one solution.

4.2. Lunar reflectance map

Next, we ask what the reflectance map looks like for the lunar surface illuminated by a single point source. The contours of constant intensity in gradient-space will be lines of constant I/E . But the contours of constant I/E are straight lines! So the gradient-space image can be generated from a single curve by shifting it along a straight-line—the shadow-line, for example (see Fig. 8). The contour lines are perpendicular to the direction defined by the position of the source (that is, the line from the origin to p_s, q_s).

Now what information does a single measurement of image intensity provide? It tells us that the gradient has to be on a particular straight line. Again, we ignore for the moment the possible existence of more than one contour for a given intensity.

What we would like to know, of course, is the orientation of the surface element. We cannot completely determine the local orientation, but we can determine its component in the direction perpendicular to the contour lines in gradient-space.

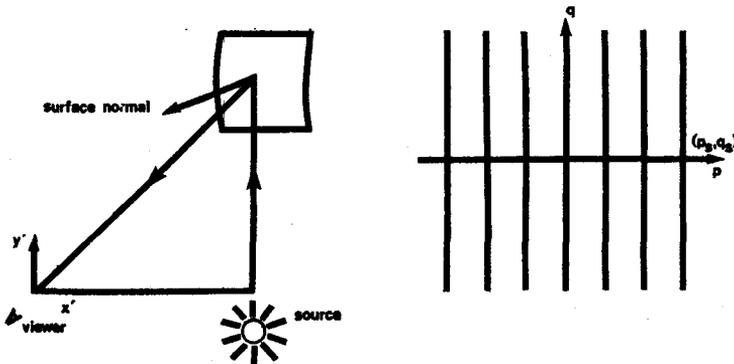


FIG. 17. Rotation of coordinate system to simplify gradient-space geometry.

We can tell nothing about it in the direction at right angles. That is, knowing I/E and G determines p' , as previously defined, and tells us nothing about q' .

This favored direction lies in the plane defined by the source, the viewer, and the surface element under consideration. If we wish, we can simplify matters by rotating the viewer's coordinate system until the x axis lies in this plane as well. Then $q_s = 0$, and the contours of constant intensity in gradient-space are all vertical lines. Evidently, an image intensity measurement determines the slope of the surface in the x' direction, without telling us anything about the slope in the y' direction (Fig. 17). We are now ready to develop the surface by advancing in the direction in which we can determine locally the surface slope.

4.3. Finding a surface profile by integration

We have:

$$p' = \frac{dz}{ds} = \frac{I/E - G}{\sqrt{1 - G^2}}.$$

The distance s from some starting point is measured in the object coordinate system and is related to the distance along the curve's projection in the image by $s' = s(f/z_0)$.

$$\frac{dz}{ds'} = \frac{f}{z_0} \frac{I/E - G}{\sqrt{1 - G^2}}.$$

Integrating, we get:

$$z(s') = z_0 + \frac{f}{z_0} \int_0^{s'} \frac{I/E - G}{\sqrt{1 - G^2}} ds',$$

where I/E is found from G and the image intensity $b(x', y')$. Starting anywhere in the image, we can integrate along a particular line and find the relative elevation of the corresponding points on the object.

The curves traced out on the object in this fashion are called *characteristics*, and their projections in the image-plane are called *base characteristics*. It is clear that here the base characteristics are parallel straight lines in the image, independent of the object's shape.

4.4. Finding the whole surface

We can explore the whole image by choosing sufficient starting points along a line at an angle to the favored direction. In this way we obtain the surface shape over the whole area recorded in the image (Fig. 18). Since we cannot determine the gradient at right angles to the direction of the characteristics, there is nothing to relate adjacent characteristics in the image. We have to know an initial curve, or use assumptions about reasonable smoothness. Alternatively, we can perform a second surface calculation from an image taken with a different source-surface-observer geometry. In this case, we obtain solutions along lines crossing the surface at a different angle, tying the two solutions together. This is not

Artificial Intelligence 8 (1977), 201-231

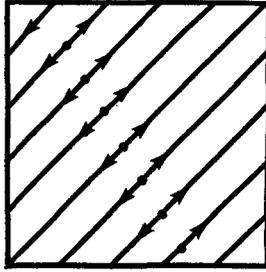


FIG. 18. Finding the shape of the surface by integrating along several base characteristics starting from different initial points.

quite as useful as one might first think, because it does not apply to pictures taken from earth. The plane of the sun, moon and earth varies little from the ecliptic plane. The lines of integration in the image vary little in inclination. This idea however *does* work for pictures taken from near the moon.

4.5. Ambiguity in local gradient

What if more than one contour in gradient-space corresponds to a given intensity? Then we cannot tell locally which gradient to apply. If we are integrating along some curve however, this is no problem, since we may assume that there is little change in gradient over small distances, and pick one close to the gradient last used. This assumption of smoothness leaves us with one remaining problem: what happens if we approach a maximum of intensity in gradient-space and then enter areas of lower intensity (Fig. 19). Which side of the local maximum do we slide down?

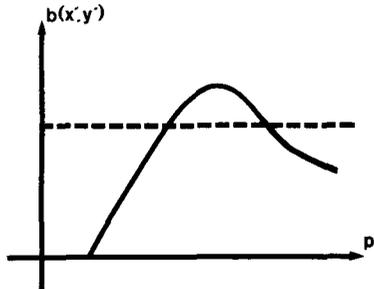


FIG. 19. Problem of ambiguity caused by non-uniqueness of slope for a particular observed image intensity.

This is an ambiguity which cannot be resolved locally, and the solution has to be terminated at this point. Under certain lighting conditions the image is divided into regions inside each of which we can find a solution. The regions are separated by ambiguity edges, which cannot be crossed without making an arbitrary choice [4].

4.6. Low sun angles

This problem can be avoided entirely if one deals only with pictures taken at low sun angles, where the gradient is a single-valued function of image intensity. This

is a good idea in any case, since the accuracy of the reconstruction depends on how accurately one can determine the gradient, which in turn depends on the spacing of the contour lines in gradient-space. If they are close together, this accuracy is high (near a maximum, on the other hand, it is low). It is easy to convince oneself that pictures taken at low sun angle have "better contrast," show the "relief in more detail," and are "easier to interpret." An additional reason for interest in low sun angle images is that the contours of constant intensity near the shadow-line in gradient-space are nearly straight lines even if we are *not* dealing with the special reflectivity function for the lunar material! An early solution to the problem of determining the shape of lunar hills makes use of this fact by integrating along lines perpendicular to the terminator [5].

Working at low sun angles introduces another problem, of course, since shadows are likely to appear. Fortunately, they are easy to deal with since we simply trace the line in the image until we see a lighted area again. Knowing the direction of the source's rays, we easily determine the position of the first lighted point. The integration is then continued from there (Fig. 20). In fact, no special attention to this



FIG. 20. Geometry of grazing ray needed to deal with shadow gaps in solution.

problem is needed since a surface element oriented for grazing incidence of light already has the correct slope. Thus, simply looking up the slope for zero intensity and integrating with this value will do. Some portion of the surface, of course, is not explored because of shadows. If we take one picture just after "sunrise" and one just before "sunset," most of this area is covered.

4.7. Generalization to perspective projection

All along we have assumed orthographic projection—looking at the surface from a great distance with a telephoto lens. In practice, this is an unreasonable assumption for pictures taken by artificial satellites near the surface. The first thing that changes in the more general case of perspective projection is that the sun-surface-viewer plane is no longer the same for all portions of the surface images. Since it is this plane which determines the integration lines, we expect that these lines are no longer parallel. Instead they all converge on the anti-solar point in the image which corresponds to a direction directly opposite the direction towards the source (Fig. 21).

The next change is that z is no longer constant in the projection equation. So $s' = f(s/z)$. Hence,

$$p' = \frac{dz}{ds} = \frac{f dz}{z ds'} = \frac{I/E-G}{\sqrt{1-G^2}}$$

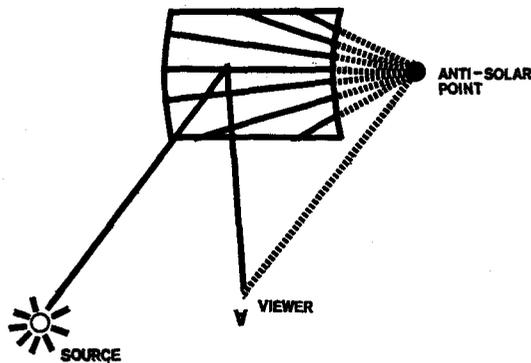


FIG. 21. The base-characteristics converge on the anti-solar point.

We can no longer simply integrate. It is easy, however, to solve the above differential equation for z by separating terms:

$$\log(z) = \frac{1}{f} \int \frac{I/E - G}{\sqrt{1 - G^2}} ds',$$

and so

$$z(s') = z_0 \exp\left(\frac{1}{f} \int_0^{s'} \frac{I/E - G}{\sqrt{1 - G^2}} ds'\right).$$

Finally, note that the phase angle g is no longer constant. This has to be taken into account when calculating I/E from the measured image intensity. On the whole, the process is still very simple. The paths of integration are predetermined straight lines in the image radiating from the anti-solar point. At each point we measure the image intensity, determining which value of I/E gives rise to this image intensity. Then we calculate the corresponding slope along the straight line and take a small step. Repeating the process for all lines crossing the image we obtain the surface elevation at all points in the image. The same result can be obtained by a very complex algebraic method [6].

4.8. A note on accuracy

Since image intensities are determined only with rather limited precision, we must expect the calculation of surface coordinates to suffer from errors that accumulate along characteristics. A "sharpening" method that relates adjacent characteristics reduces these errors somewhat [4].

It also appears that an object's shape is better described by the orientations of its surface normals than by the distances from the viewer to points on its surface. In part, this may be because distances to the surface undergo a more complicated transformation when the object is rotated than do surface normal directions. Note that the calculation of surface normals is *not* subject to the cumulative errors mentioned.

Finally, we should point out that the precise determination of the surface is not the main impetus for the development presented here. The understanding of how image intensities are determined by the object, the lighting, and the image forming system is of more importance and may lead to more interesting heuristic methods.

5. The Shape of Surfaces with Arbitrary Reflectance Maps

The simple method developed for lunar topography is inapplicable if the contours of constant intensity in gradient-space are not parallel straight lines. We will still be able to trace along the surface, but the direction we take at each point now depends on the image and changes along the profile. The base characteristics no longer are predetermined straight lines in the image. At each point on a characteristic curve we find that the solution can be continued only in a particular direction. It also appears that we need more information to start a solution and will have to carry along more information to proceed. Reasoning from the gradient-space diagram is augmented here by some algebraic manipulation.

Let $a(p, q)$ be the intensity corresponding to a surface element with a gradient (p, q) . Let $b(x, y)$ be the intensity recorded in the image at the point (x, y) . Then, for a particular surface element, we must have:

$$a(p, q) = b(x, y).$$

Now suppose we want to proceed in a manner analogous to the method developed earlier by taking a small step (dx, dy) in the image. It is clear that we can calculate the corresponding change in z as follows:

$$dz = z_x dx + z_y dy = p dx + q dy.$$

To do this we need the values of p and q . We have to keep track of the values of the gradient as we integrate along the curve. We can calculate the increments in p and q by:

$$dp = p_x dx + p_y dy \quad \text{and} \quad dq = q_x dx + q_y dy.$$

At first, we appear to be getting into more difficulty, since now we need to know $p_x, p_y = q_x$ and q_y . In order to determine these unknowns we will differentiate the basic equation $a(p, q) = b(x, y)$ with respect to x and y :

$$a_p p_x + a_q q_x = b_x \quad \text{and} \quad a_p p_y + a_q q_y = b_y.$$

While these equations contain the right unknowns, there are only two equations, not enough to solve for three unknowns. Note, however, that we do not really need the individual values! We are only after the linear combinations $(p_x dx + p_y dy)$ and $(q_x dx + q_y dy)$.

We have to properly choose the direction of the small step (dx, dy) to allow the determination of these quantities. There is only one such direction. Let $(dx, dy) = (a_p, a_q) ds$, then $(dp, dq) = (b_x, b_y) ds$. This is the solution we were after. Summarizing, we have five ordinary differential equations:

$$\dot{x} = a_p, \dot{y} = a_q, \dot{z} = a_p p + a_q q, \dot{p} = b_x \quad \text{and} \quad \dot{q} = b_y.$$

Here the dot denotes differentiation with respect to s , a parameter that varies along the solution curve.

5.1. Interpretation in terms of the gradient-space

As we solve along a particular characteristic curve on the object, we simultaneously trace out a base characteristic in the image and a curve in gradient-space. At each point in the solution we know to which point in the image and to which point in the gradient-space the surface element under consideration corresponds.

The intensity in the real image and in the gradient space image must, of course, be the same. The paths in the two spaces are related in a peculiar manner. The step we take in the image is perpendicular to the contour in *gradient-space* and the step we take in gradient-space is perpendicular to the intensity contour in the *image-plane*. (See Fig. 22.)

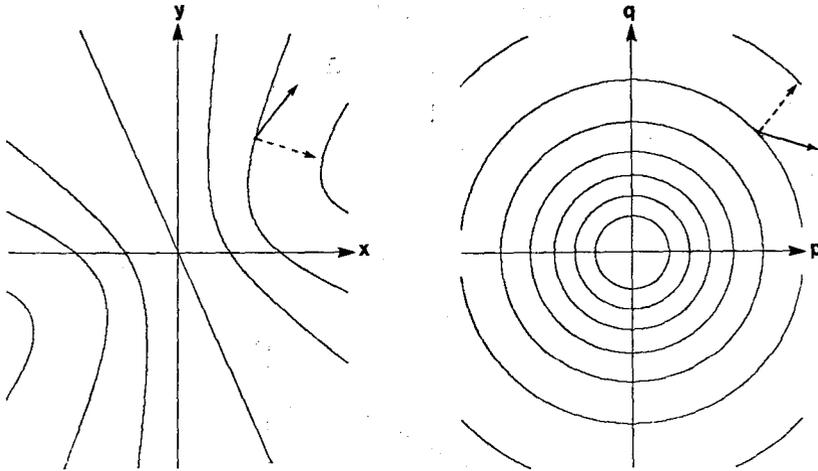


FIG. 22. The solution curve in image-space is along the line of steepest descent in gradient-space and the solution curve in gradient-space is along the line of steepest descent in image-space.

5.2. Generalization to source and viewer near the object

The last solution method, while correct for arbitrary reflectivity functions, still assumes orthographic projection and a distant source. This is a good approximation for many practical cases. In order to take into account the effects of the nearness of the source and the viewer, we discard the gradient-space diagram, since it is based on the assumption of constant phase angle. The problem can still be tackled by algebraic manipulation, much as the last solution. It turns out that we are really trying to solve a first order non-linear partial differential equation in two independent variables. The well-known solution involves converting this equation into five ordinary differential equations, quite like the ones we obtain in the last Section [4].

Appendix

Using the reflectance map to determine three-dimensional structure of polyhedral scenes

What follows is a simple numerical example to illustrate the idea that information about surface reflectance can augment the gradient-space diagram and lead to a solution for the orientation of three planes meeting at a vertex. We will assume a lambertian reflectance for the object and assume that the light-source and viewer are far removed from the object, but close to each other. Suppose now that we are given the partial line-drawing as in Fig. 23 showing edges separating three planes

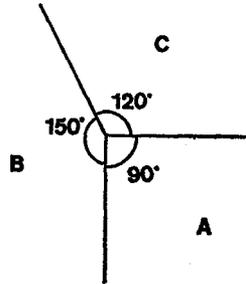


FIG. 23. Image-lines of a tri-hedral vertex. The orientations of the planes A , B and C are to be found.

A , B and C . The gradient-space diagram showing the three points G_A , G_B and G_C corresponding to these three planes will be as in Fig. 24. The scale and position of the indicated triangle are not yet determined. In fact the whole diagram could be flipped around if the scale is negative.

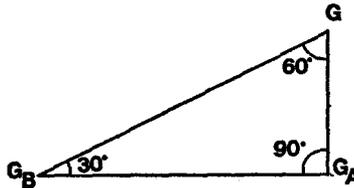


FIG. 24. The gradient-space diagram corresponding to the previous figure. The planes A , B and C map into the points G_A , G_B and G_C . The scale and position of this triangle is as yet undetermined.

Nevertheless, we now have available three linear constraints on the coordinates (p_a, q_a) , (p_b, q_b) and (p_c, q_c) of the points G_A , G_B , G_C .

$$p_a = p_c, q_a = q_b, \text{ and } (q_c - q_a) = (p_a - p_b)t.$$

Where $t = \tan(30^\circ) = 1/\sqrt{3}$. We are now told that image measurements suggest reflectances of 0.707, 0.807 and 0.577 for the three faces respectively. From the information about the position of the source and viewer, we know that $G = 1$ and that $I = E$. Next, since we are dealing with a lambertian surface we calculate the reflectance from $\phi(I, E, G) = I$, which here equals $E = \cos(e) = 1/\sqrt{1+p^2+q^2}$. We immediately conclude that the surface normals of the three planes are inclined 45° , 36.2° and 54.8° respectively with respect to the view vector.

It also follows that the points G_A , G_B and G_C must lie on circles of radii 1.0, 0.732 and 1.415 in gradient-space, since distance from the origin in gradient space $\sqrt{p^2+q^2}$ equals $\tan(e)$. That is, the points have to lie on the appropriate contours of reflectance in the reflectance map as in Fig. 25.

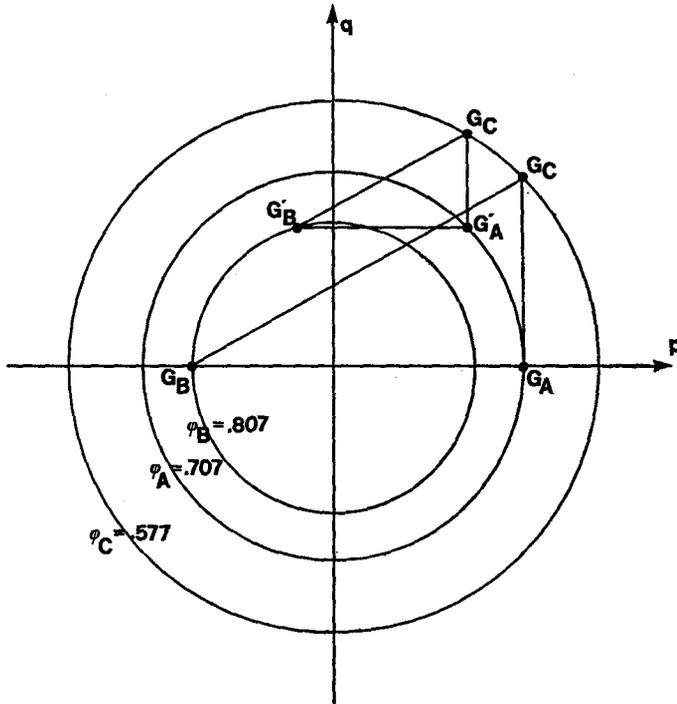


FIG. 25. Contours of constant reflectance corresponding to the reflectances of the three planes. Two solutions are superimposed on the reflectance map.

This gives us three further constraints, unfortunately non-linear ones. Let us call the radii r_a , r_b and r_c respectively, then

$$p_a^2 + q_a^2 = r_a^2, \quad p_b^2 + q_b^2 = r_b^2 \quad \text{and} \quad p_c^2 + q_c^2 = r_c^2.$$

If the source had *not* been near the viewer, these equations would have involved linear terms in p and q as well, since then the contours of equal reflectance would have been conic sections other than circles. In the general case these three equations could be more complicated and in fact it is possible that the reflectance map is only known numerically. Then one will have to proceed iteratively from here. In our simple example however it is possible to solve the three linear equations and the three non-linear ones we have developed directly. As usual one proceeds by judiciously eliminating variables.

Let us use the three linear equations to eliminate the unknowns p_c , q_b and q_c from the three non-linear equations, then,

$$p_a^2 + q_a^2 = r_a^2, \quad p_b^2 + q_a^2 = r_b^2 \quad \text{and} \quad p_a^2 + [q_a + i(p_a - p_b)]^2 = r_c^2.$$

We now have three non-linear equations in three unknowns. First note that

$$(p_a - p_b)(p_a + p_b) = p_a^2 - p_b^2 = r_a^2 - r_b^2.$$

Then expand the last of the three equations to get:

$$p_a^2 + q_a^2 + 2tq_a(p_a - p_b) + t^2(p_a - p_b)^2 = r_c^2.$$

Now using $p_a^2 + q_a^2 = r_a^2$ and the previous equation for $(p_a - p_b)$:

$$[2tq_a + t^2(p_a - p_b)](r_a^2 - r_b^2) = (p_a + p_b)(r_c^2 - r_a^2).$$

This last question is linear! It is of the form $ap_a + bp_b + cq_a = 0$, where a , b and c can be evaluated and found to be -0.845 , -1.155 and 0.536 .

The simplest next operation is the elimination of p_a and p_b , using the equations for r_a^2 and r_b^2 .

$$-a\sqrt{r_a^2 - q_a^2} - b\sqrt{r_b^2 - q_a^2} = cq_a.$$

A single equation in a single unknown, at last. Squaring both sides leads to:

$$2ab\sqrt{(r_a^2 - q_a^2)(r_b^2 - q_a^2)} = (a^2 + b^2 + c^2)q_a^2 - (a^2r_a^2 + b^2r_b^2).$$

Letting $e = (a^2 + b^2 + c^2)/(2ab)$ and $f = (a^2r_a^2 + b^2r_b^2)/(2ab)$ and squaring again to get rid of the square-root:

$$(q_a^2)^2(1 - d^2) + (q_a^2)[2de - (r_a^2 + r_b^2)] + (r_a^2r_b^2 - e^2) = 0.$$

This quadratic equation in q_a^2 can be further simplified by evaluating the terms to:

$$(q_a^2)^2 - 0.5(q_a^2) = 0.0.$$

The solutions are $q_a = 0.0$, -0.7071 and $+0.7071$. Not all of these may be solutions of the original equations, so we will have to check the results. Trying $q_a = 0$, leads to $q_b = 0$, $p_a = \pm 1$, $p_c = \pm 1$, $p_b = \pm 0.732$, and $q_c = 0.577(\pm 1 \pm 0.732)$. Two of these combinations satisfy the original equations.

$$(p_a, q_a) = (1, 0)$$

$$(p_b, q_b) = (-0.732, 0)$$

$$(p_c, q_c) = (1, 1).$$

The other solution is the mirror image of this one, with the same numerical values, but reversed signs. One of these solutions is seen superimposed on the contours of the reflectance map in Fig. 25.

If one tries the other possible set of values for q_a , ± 0.7071 one finds, $p_a = q_b = \pm 0.7071$, $p_b = \pm 0.1895$, $p_c = \pm 0.1895$, and $q_c = \pm 0.7071 + 0.577(\pm 0.7071 \pm 0.1895)$. As before only two combinations satisfy the original equations. One of these is the following:

$$(p_a, q_a) = (0.707, 0.707),$$

$$(p_b, q_b) = (-0.189, 0.707),$$

$$(p_c, q_c) = (0.707, 1.225).$$

The other solution again simply has the signs reversed. One of these solutions is also shown in figure 25. The symmetry of the problem here contributes to the plethora of solutions, more usually one finds but two.

Clearly it would be desirable to avoid this tedious solution of simultaneous non-linear equations. Graphical techniques work and iterative Newton-Raphson techniques are appropriate for computer implementations of this method. For a numerical example see [11].

ACKNOWLEDGMENTS

I wish to thank Blenda Horn and Eva Kampitts for help in the preparation of the paper, Karen Prendergast for the drawings and Kathy Van Sant for an early version of the numerical solution for the truncated-plane mutual-illumination problem. David Marr and Patrick Winston provided much appreciated stimulation and discussion.

REFERENCES

1. Huffman, D. A., Impossible objects as nonsense sentences, *Machine Intelligence* 6, Meltzer, R., and Michie, D. (Eds.), (Edinburgh University Press, 1971) 295-323.
2. Mackworth, A. K., Interpreting pictures of polyhedral scenes, *Artificial Intelligence* 4 (1973), 121-137.
3. Huffman, D. A., Curvature and creases: a primer on paper, *Proc. Conf. Computer Graphics, Pattern Recognition and Data Structures* (May 1975) 360-370.
4. Horn, B. K. P., Obtaining shape from shading information, in: Winston, P. H., (Ed.), *The Psychology of Machine Vision* (McGraw-Hill, NY, 1975) 115-155.
5. Van Diggelen, J., A photometric investigation of the slopes and heights of the ranges of hills in the maria of the moon, *Bull. Astron. Inst. Netherlands* 11 (1951) 283-289.
6. Rindfleisch, T., Photometric method for lunar topography, *Photogrammetric Eng.* 32 (1966) 262-276.
7. Herskovits, A. and Binford, T. O., On boundary detection, MIT Artificial Intelligence Memo 183 (July 1970) 19, 55, 56.
8. Phong, Bui Tuong, Illumination for Computer Generated Pictures, *CACM* 18 (1975) 311-317.
9. Hilbert, D. and Cohn-Vossen, S., *Geometry and the Imagination* (Chelsea Publishers, New York, 1952).
10. Hildebrand, F. D., *Methods of Applied Mathematics* (Prentice-Hall, New Jersey, 1952) 222-294.
11. Horn, B. K. P., Image intensity understanding, MIT Artificial Intelligence Memo 335 (August 1975).

Received January 1976; final version received August 1976