The Shape of Shadows

David A. Burton

Larry S. Davis

Department of Computer Sciences
University of Texas at Austin
Austin, Texas 78712

TR-147 July 1980

This research was supported in part by the National Science Foundation under Grant ENG-7904037 and the Air Force Office of Scientific Research under Grant AFOSR77-3190.

ABSTRACT

An important problem in image analysis is to identify useful scene features, such as shadows and occlusions, at a low level. This information is, hopefully, more directly useful to higher level scene analysis programs than the intensity data would be. This paper reviews some of this work, with special emphasis on techniques which are based upon examination of edges in the image, and contains an analysis of edges caused by shadows in the scene. It was found that the intensity profiles predicted for the boundaries of shadowed regions are quite distinctive. We describe our efforts to use this work to identify shadow edges in real images, which met with only limited success.

SECTION I

EDGES AND BOUNDARIES

Edges in digital images are local phenomena, corresponding to discontinuities in some image characteristic, usually observed intensity. They are commonly detected by applying an edge detection operator (generally a digital first derivative approximation) to the image and then, perhaps, thresholding the result. Each pixel classified as an edge point by this process has two numbers associated with it, an angular orientation and a strength. The process is quite noise sensitive, and tends to give multiple parallel weak edge responses for fuzzy or gradual edges. For these reasons, noise-cleaning, boundary-thinning, and gap-filling operations are usually then applied to the set of edge labelings. Relaxation techniques (Rosenfeld, et al [1] can sometimes be used to eliminate spurious labelings (due to noise) or mutually inconsistent labelings. Ideally, every image point which lies on the border between two regions of interest would then be labeled as an edge point with orientation along the border, and every interior point would be unlabeled.

Chains of adjacent edge points could then be grouped into boundaries.

In practice, a consistent edge-labeling is difficult to obtain by this method. Another approach to boundary detection is to find the regions directly by region-growing. Boundaries (edges) are then the borders between adjacent regions. It is also possible to combine the two techniques by using the edge-detecter responses to guide the region-growing process, or by first running a region-grower, and then using the edge detector responses to adjust the placement of borders between regions.

Whatever the technique used to detect them, it is useful to be able to classify edges within an image according to the physical phenomena which caused their appearance. This classification (labeling) can be very helpful in understanding the scene, since it conveys information about the positions and orientations of surfaces in the scene. Although this is difficult in the general case, it is often possible to arrive at probable labelings of some of the edges according to their causes, or to rule out some causes for some of the boundaries. Since the physics of the imaging system and the scene relates the possible labelings of different edges to one another, restricting the set of possible labels for one

edge implies constraints upon the plausible labelings of other edges. Thus, edge-label constraints can be propagated through the boundary network to further reduce the list of possible labelings for each boundary.

An observed boundary can be due to a discontinuity in any of several characteristics of the scene (e.g., reflectance, viewpoint/object distance, orientation, etc.). The process of labeling boundaries according to their types is equivalent to that of finding the discontinuities* in these "intrinsic images" (Barrow and Tenenbaum [2]). Unfortunately, if the surface characteristics of an object vary over the object (e.g., if the reflectance-intrinsic image can vary arbitrarily), discontinuities may be detected which will not correspond to actual boundaries in the scene. In an extreme case, if the object being examined were a uniformly illuminated photograph, it would be indistinguishable from the illustrated scene. The boundaries between regions in such a case could be considered to be correctly labeled either as reflectance discontinuities (in the photograph) or object boundaries, shadows, etc. (in the photographed scene). Fortunately, for most real applications, some constraining assumptions can be made about the nature and

*The term "discontinuity" is used loosely (since digital images are not continuous).

frequency of reflectance changes. For many applications, in fact, the reflectance is a single function of viewer and light source orientation over individual objects (that is, objects are a single solid color). Some researchers (Marr, [3]; Barrow and Tenenbaum [2]) to simplify analysis, make the assumption that object surfaces have lambertian reflectance functions. For a lambertian (matte) surface, the observed luminescence is proportional to the sine of the angle between the surface and incident light and independent of the view angle. Horn [4] notes that for a few applications (e.g., lunar landscapes), there is a single reflectance function for the entire image. Where valid, such assumptions can obviously be of great utility to a program which tries to classify boundaries according to their appearances.

Other boundary types are easier to discriminate between. Specularities (glare spots) on gradually varying glossy (non-lambertian) surfaces under highly directional illumination (e.g., sunlight) can sometimes be identified very simply because they are too bright to be anything else. Note that since the presence of glare at an image point implies that the incident angle of the light source is equal to the view angle, the three-dimensional orientation of the surface is easily determined if the light source direction is known. The

surface normal is simply the vector sum of the unit source direction vector and the unit vector from the viewpoint toward the specularity. For outdoor images where the sun is the light source, the source direction vector can be measured to within 1/2 the angular size of the source, approximately 1/2 degree. Hence, if the viewpoint-to-specularity direction vector is known precisely, the surface normal at the specularity can be determined (in theory) to within aproximately 1/4 degree. Thus, glare, which is commonly considered a nuisance, can give useful information about the scene, and the boundaries surrounding a glare spot can often be immediately classified. Horn [4] examines the phenomenon of specularities on non-lambertian surfaces in detail and gives an approximating reflectance function for some glossy white paints.

A graph of observed intensity versus displacement across an edge is called an intensity profile. Examination of the intensity profile can sometimes aid in classifying the boundary. One type of boundary for which this can be true occurs when a convex object with distinct adjoining faces is viewed from a perspective which allows both faces to be seen. If the corner between the faces is very sharp (as opposed to rounded), the intensity profile will be roughly step-shaped (a

"step-edge"). Horn [4], however, observed that most corners are, in reality, somewhat rounded. For this reason, a narrow specularity is sometimes observed along the boundary. This is, of course, only true for non-lambertian surfaces under directional lighting, and it implies that some linear combination of the two surface normals bisects the angle between the vector to the light source (from the surface) and the vector to the viewpoint. This can cause a narrow peak to be superimposed upon the step-shaped intensity profile. Horn found that this effect is less common (although possible) for concave object edges and object edges for which only one of the two faces is visible (occluding boundaries). Thus, he considers a step-edge with a narrow peak superimposed to be evidence of a convex object edge.

A concave boundary on a highly reflective (light colored) object, according to Horn [4], will sometimes exhibit a wide peak or roof superimposed upon the step, due to mutual illumination. He reports that this has proved to be an unreliable cue in practice, though, since his digitization system, an image dissector camera, can apparently produce a very similar distortion.

Yet another type of boundary is observed where one object partially occludes another in the image. A

step edge is likely to be observed in this case, although a specularity may be present at the boundary. Thus, this case can be difficult to discriminate from that of an object edge viewed from a perspective which allows both faces to be seen. Horn [4] notes that if the (hidden) surface adjoining the occluding surface is less directly illuminated or in shadow, a weak negative peak may be superimposed upon the step (Figure 1). Observation of such a peak is thus strongly indicative of an occluding boundary. If it is known that surface characteristics are constant over an object, and if color or textural information is available, then a boundary for which the two regions have different hue or texture corresponds to an occluding boundary. Some caution is necessary, though, since the same texture can appear very different for different orientations. In general, even if it is possible to classify a boundary as caused by an occlusion, it may be difficult to determine the sense of the occlusion.

Examination of the boundary ends (where several boundaries meet) can sometimes allow a classification to be made in spite of these difficulties. Waltz [5] classified trihedral boundary junctions into three types, arrow-junctions, T-junctions, and Y-junctions (Figure 2). Three common-sense notions allow the use of junction

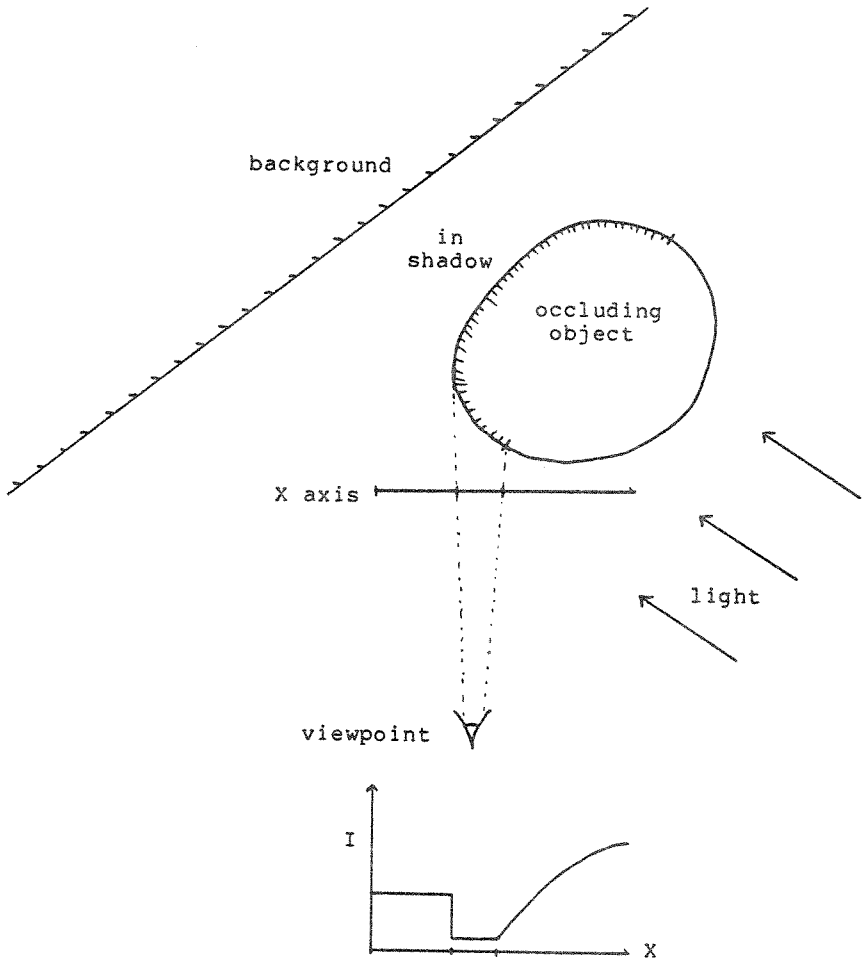

Figure 1) A situation giving rise to a negative peak in the intensity profile for an occluding edge. A portion of the shadowed surface is visible at the boundary.

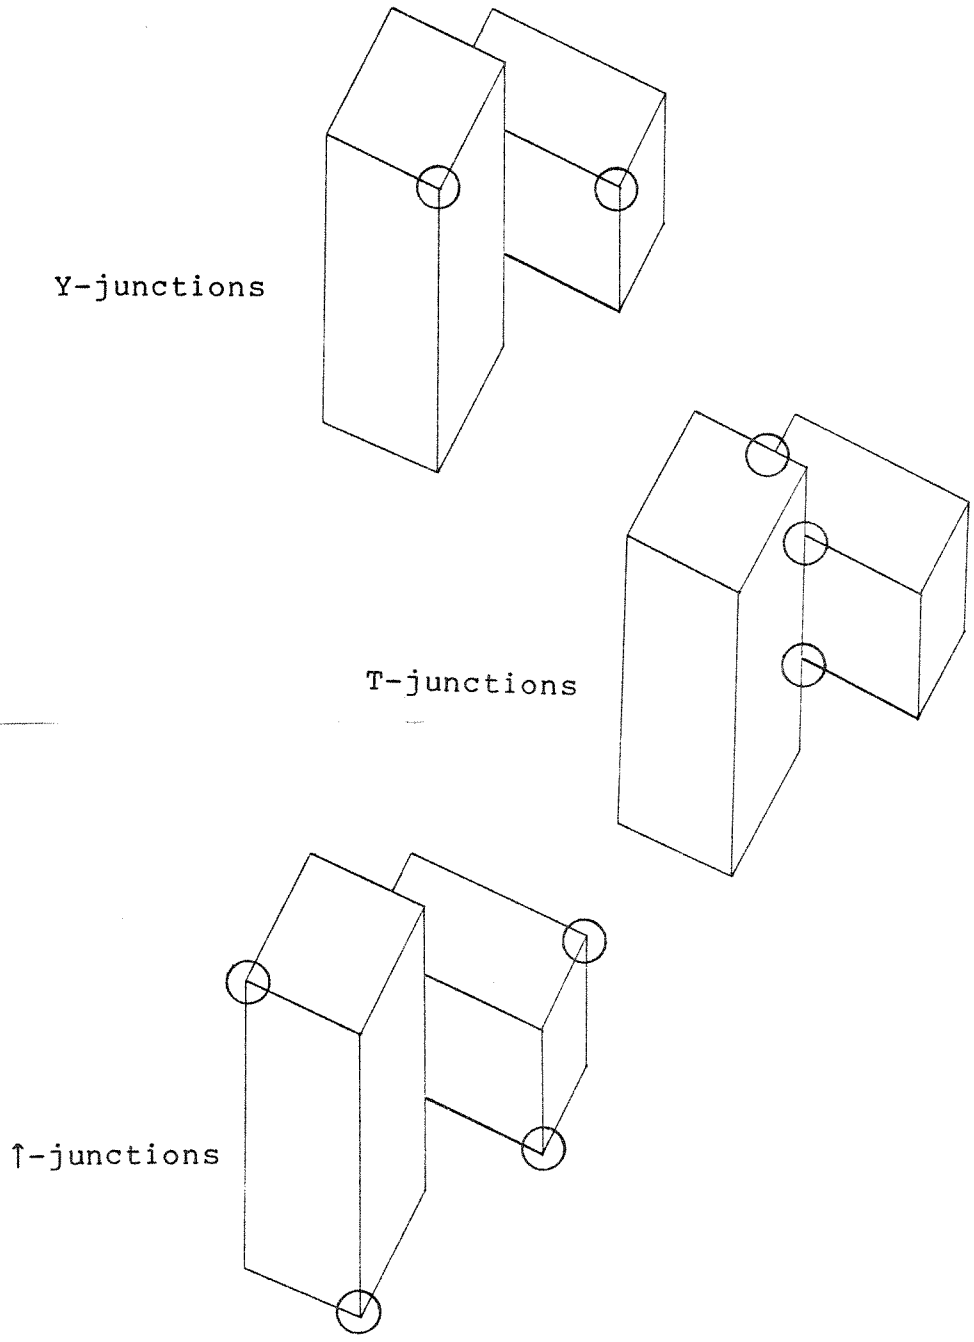

Figure 2) Trihedral junctions in a simple scene.

types to classify boundaries. The first is that boundaries, even curved boundaries, tend to be locally straight because surfaces tend to be locally flat except at well-defined surface boundaries. The second is that singular views (coincidences) are uncommon. Singular views are views for which a minute change in view position could result in a change in the topology of the image. The last notion is that object boundaries are most frequently convex.

If we make these assumptions, we can infer the following: i) that T-junctions are generally caused by an object occluding another boundary; ii) that arrow-junctions are usually due to an object vertex for which one of the three faces is turned away from the viewpoint so that a background object is visible; and iii) that Y-junctions are generally due to object vertices for which all three surfaces are visible. Thus, if a boundary terminates in a Y-junction or the "shaft" of an arrow-junction, the boundary is probably due to an object edge (surface junction) with both surfaces visible. If the boundary ends at one side of an arrow head, the boundary is probably due to an occlusion with the arrow pointing to the occluded surface. Similarly, if the boundary ends at the top of a T-junction, it is likely that it was caused by an occlusion with the

occluding object at the top of the T. If the base of the T is caused by a shadow, however, the sense of the occlusion cannot be directly determined.

The last two types of boundaries commonly found in an image are both caused by shadows. When one object casts a shadow upon another, a boundary will be seen. There are several cues which can sometimes be used to discriminate shadow boundaries from other types. If the boundary is fuzzy, it is likely to be due to a shadow, although blurring in the imaging process can have a similar effect. Also, if the lighting conditions are known, inferences can be made about the ratio of the distances from the surface to the viewpoint and from the surface to the object casting the shadow. If the distance from the shadowed surface to the object is not constant, but varies along its extent, then the width of the observed boundary at the shadow edge will vary as well. If this is observed, the boundary can immediately be classified with high probability as a shadow. For most real surfaces, the color is the same regardless of the direction from which it is illuminated or viewed. Thus, if color information is available and there is some secondary lighting so that the shadowed regions do not appear black, the observed hue will be approximately the same on both sides of the boundary for a shadow.

Examination of the boundary ends can also allow identification of cast-shadow boundaries. If the scene does not contain transparent objects, coincidences, or object vertices of degree greater than three, then any degree-4 junctions observed must be due to a shadow falling across another type of boundary. They fall into two distinguishable classes, illustrated in figure 3. If the shadow is cast across a reflectance discontinuity, the junction will appear X-shaped, i.e., as the intersection of two (locally) straight lines. If the shadow is cast across an orientation discontinuity (object edge, where two adjoining visible surfaces meet), one of the two lines will be bent at the junction.* The bent boundary is the shadow edge.

The other type of shadow boundary is illustrated in Figure 4. It is due to self-shadowing in a convex object surface. This type of edge can be due either to a sharp object edge or the gradual curvature of the object surface. The apparent width of the boundary will be proportional to the radius of curvature of the object surface as well as to the angular size of the directional light source. As with a cast shadow, the hue of the two surfaces will be the same unless the boundary also

*It will be bent unless the viewpoint is in the plane determined by the edge casting the shadow and the line from the source to the junction, a coincidence.

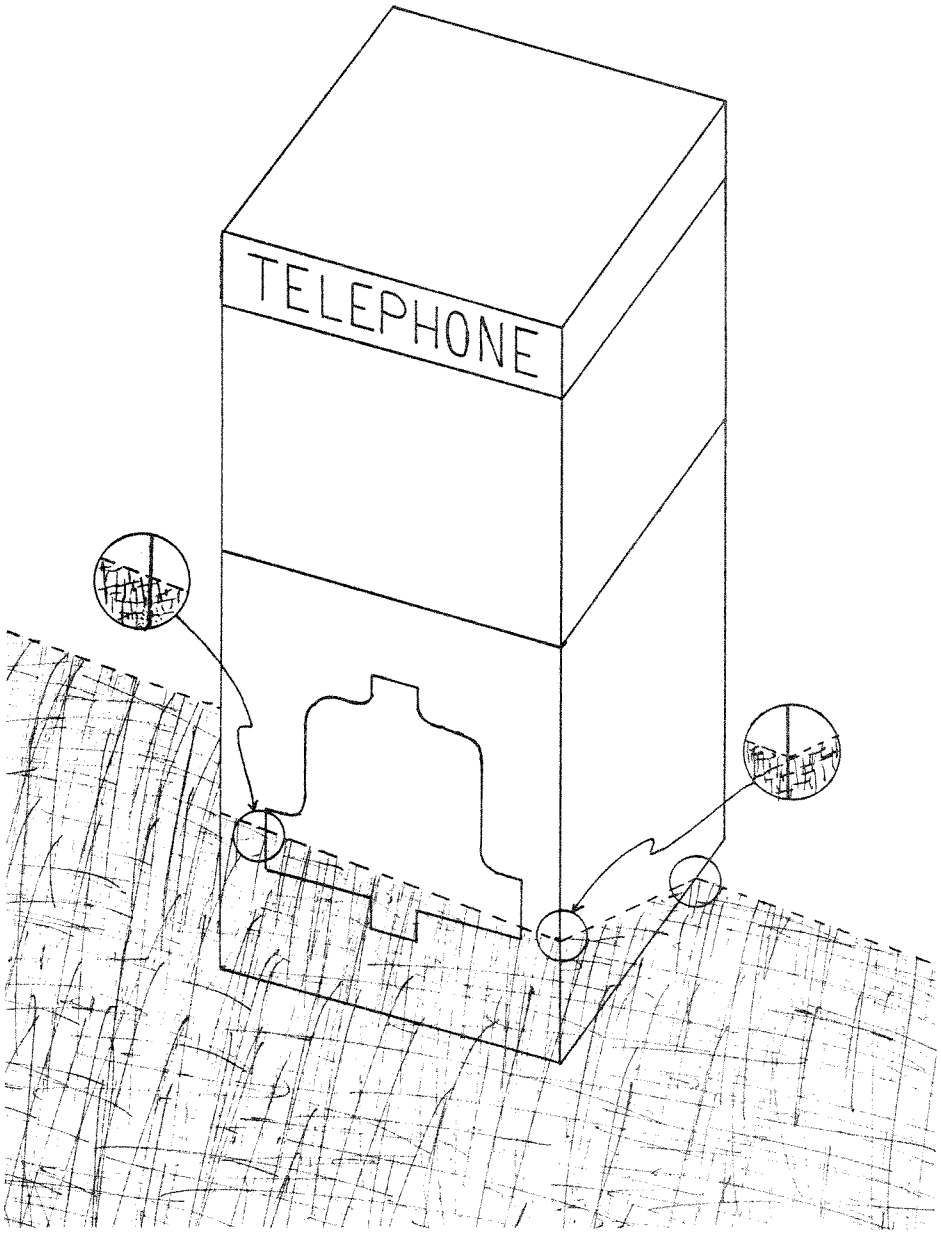

Figure 3) Degree-4 junctions due to shadows cast across reflectance discontinuities (left) and orientation discontinuities (right).

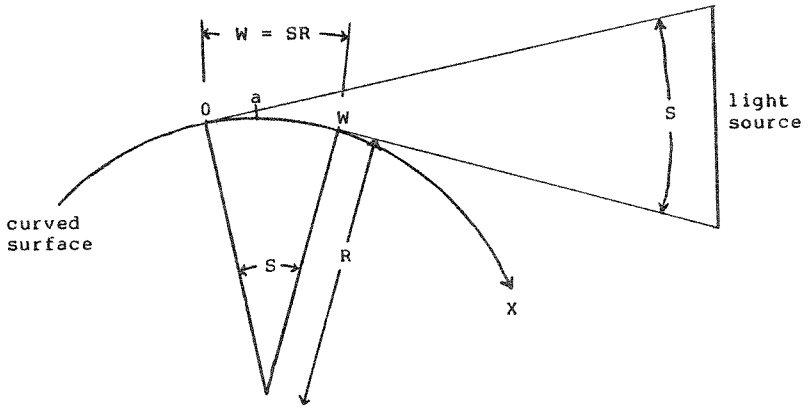

(4c)

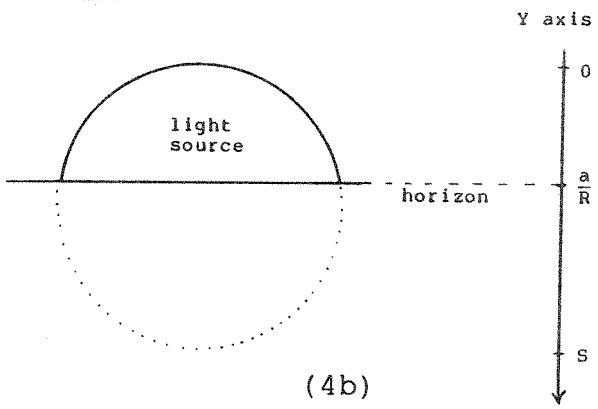

(4b)

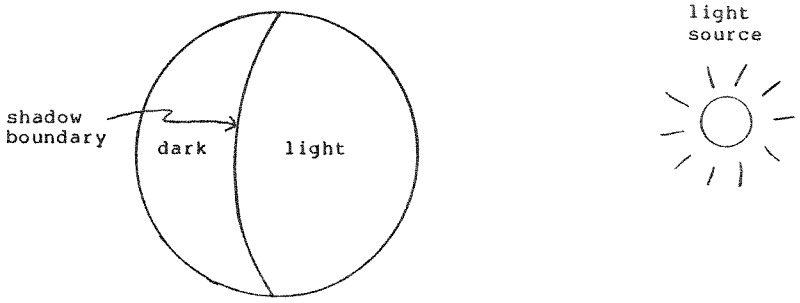

(4a)

Figure 4) Self-shadowing on a curved surface. 4b shows what the source would look like if viewed from the point marked 'a' in 4c.

corresponds to a reflectance discontinuity, as when different object faces are different colors. Unlike a cast shadow edge, however, this type of boundary will not end in a quadrahedral junction (barring a coincidence).

Barrow and Tenenbaum [2] have taken a different approach to the identification of shadow boundaries. They studied a restricted set of images for which flat surfaces and sharp object edges were forbidden. They assumed that all reflectances were lambertian and did not vary over the extent of an object. They modeled the lighting conditions by a point source plus a uniform secondary illumination, corresponding to the sun and the light sky, respectively. Mutual illumination and partial shadowing of the surfaces from the secondary light source were not accounted for. Because of the point source assumption, shadow boundaries were perfectly sharp, so edge width could not be used to identify them. Instead, they identified shadows directly. Since the observed grey-levels within a region depended only upon surface orientation relative to the point light source, shadowed regions always had constant grey-level and illuminated regions never did. Unfortunately, this is not always true for real images.

Marr [6] attempts to classify boundaries in real images according to their widths, using the symbolic

fuzziness classifications sharp, slightly fuzzy, and fuzzy. His method for measuring the widths is interesting. When doing the initial edge detection, he convolves* step edge masks at several orientations with the image (Figure 5). This process is similar to applying a digital first derivative operator to the image. It is equivalent to computing the local correlation between a perfect step edge and the image at every point. He was also able to detect edges using a second derivative approximation achieved by convolving the image with a bar mask.

Marr observed that wide (fuzzy) edges give weak responses for narrow edge masks, but that sharp boundaries give equally strong responses for narrow and wide masks. Therefore, he assesses the spread of the boundary by comparing the relative sizes of response peaks for several mask widths. Some care is necessary to avoid the effects of cross-talk between adjacent boundaries.

A more recent theory of edge detection and analysis based on the zero-crossings in a Laplacian filtered image is given in Marr and Hildreth [7].

*He suggests that, to save computation, this could be done by multiplying in the Fourier domain.

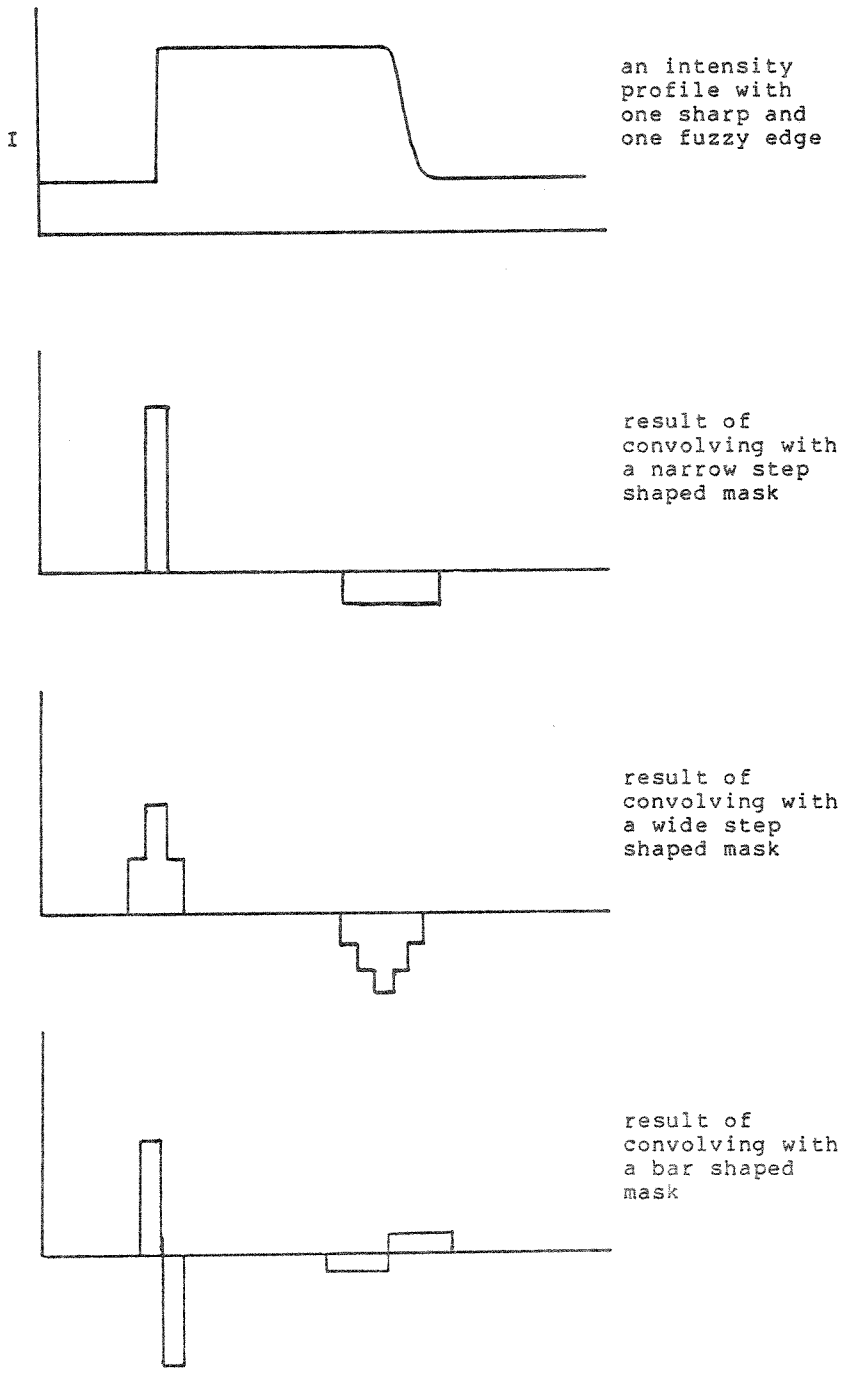
I
an intensity profile with one sharp and one fuzzy edge
result of convolving with a narrow step shaped mask
result of convolving with a wide step shaped mask
result of convolving with a bar shaped mask

Figure 5)

SECTION II

THE SHAPES OF SHADOWS

If shadow boundaries are to be discriminated according to their intensity profiles, it must be assertained what those profiles should be for the two different types of shadows described in the previous chapter. This is clearly impossible if the lighting distribution is arbitrary and unknown. It is straightforward, however, if some reasonable simplifying assumptions are made.

We will model the illumination by the combination of a single small circular source of uniform brightness and known angular size plus a locally constant secondary illumination. This is a reasonable approximation for most outdoor scenes on sunny days, and some indoor scenes. Note that this does not require that there be only one directional light source, merely that those light sources not causing the shadow contribute approximately uniform illumination in the vicinity of a

shadow edge, and that the light source causing the shadow be small and round.

The case of a shadow cast from one object onto a locally flat surface is shown in Figure 6. For outdoor scenes, S is the angle subtended by the sun, approximately 1/2 of one degree (the contribution of the corona to the total illumination is negligible (Teschan [8])). D is the average distance from the obstruction to the viewed surface. V is the average angle of incidence of the sunlight upon the viewed surface. If V is much larger than S, the sine of the angle of incidence will be nearly constant across the shadow boundary. Since the amount of light incident upon a unit surface area is proportional to the sine of the angle of incidence*, it is useful to make the simplifying assumption that this angle is constant across the width of the boundary and equal to V (it actually varies slightly across the width of the boundary, from (V - S/2) to (V + S/2)). The defraction effects are negligible for any reasonable size of source (Teschan [8]). The observed width of the shadow boundary, W, is then approximately

S D / sin V or D sin S / sin V (since S is small).

*Note that this is true regardless of the surface characteristics. If the observed intensity from a specific viewpoint is to be predicted, however, some knowledge of the surface characteristics is needed.

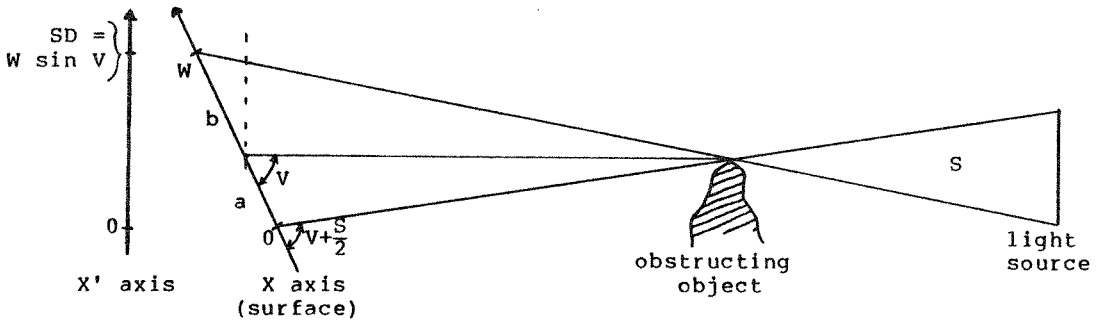

(6c)

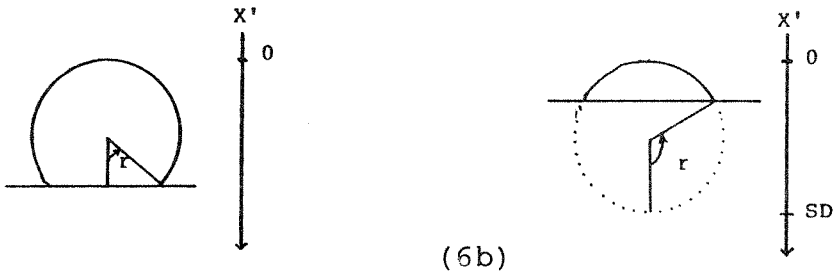

(6b)

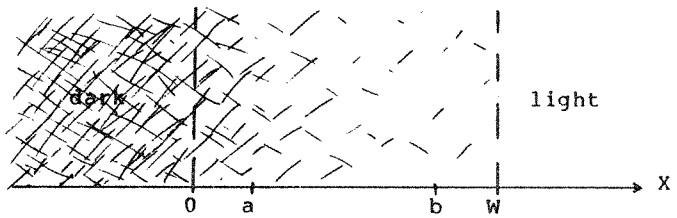

(6a)

Figure 6) A cast-shadow boundary. 6b shows what the source would look like if it were viewed from the point marked 'a' (right) or 'b' (left) in 6c. 'r' in 6b is arccos(1 - 2X'/SD).

If we superimpose an X axis upon the surface and perpendicular to the boundary (fig 6a), we can consider the intensity of the incident light at a point on the boundary as a function of X. This quantity is di, the integrated incident light flux on a differential surface area, dA. di/dA could be called the incident intensity. If we ignore other light sources, which will only contribute a constant additive illumination across the entire edge, and if the light source has intensity I, then the incident intensity when the entire light source is visible, i.e., with X=W, will be approximately

$$\frac{di}{dA} = (\pi S^2/4)\ I \sin V$$

and

$$di = (\pi S^2/4)\ I \sin V\ dA$$

where $(\pi S^2/4)$ is the area of the light source, and sin V dA is the area of the projection of the surface element upon the plane perpendicular to the direction of illumination. If we consider the projection of the X axis upon this plane to be a new axis, X', where X' = X sin V, subsequent analysis will be simplified. For a point X'=a, $0 < a < SD = W \sin V$, the incident intensity will be approximately

$$\frac{di}{dA} = \frac{1}{D} \int_0^a S \sin r\ dX'\ I \sin V, \text{ where}$$

$$r = \arccos\left(1 - \frac{2X'}{SD}\right) \text{ and the area of}$$

the light source to which a is exposed

$$\text{is } \frac{1}{D}\int_0^a S \sin r \, dX'$$

Letting $X' = \frac{S\,D}{2}(1 - \cos r)$, then

$$dX' = -\frac{S\,D}{2}\, d\cos r, \text{ and}$$

$$\frac{di}{I\,dA} = -\sin V \int_0^{\arccos(1-\frac{2a}{SD})} \frac{S}{2}^2 \sin r \, d\cos r$$

$$= \frac{S}{2}^2 \sin V \int_0^{\arccos(1-\frac{2a}{SD})} \sin^2 r \, dr$$

$$= \frac{S}{2}^2 \sin V \left[\frac{1}{4}\sin\left(2 \arccos\left(1 - \frac{2a}{SD}\right)\right) - \frac{1}{2}\arccos\left(1 - \frac{2a}{SD}\right)\right]$$

$$= \frac{S}{8}^2 \left[2 \arccos\left(1 - \frac{2a}{SD}\right) - \sin\left(2 \arccos\left(1 - \frac{2a}{SD}\right)\right)\right] \sin V$$

This function is plotted in figure 7 verses a/SD.

To analyze the case of the boundary observed on a self-shadowing object, we make the assumption that the cross section of the object perpendicular to the shadow edge and the object surface is locally circular (near the shadow edge), and has radius of curvature R. If the shadow boundary is due to a sharp object edge between two faces of the object, R will be small and the intensity

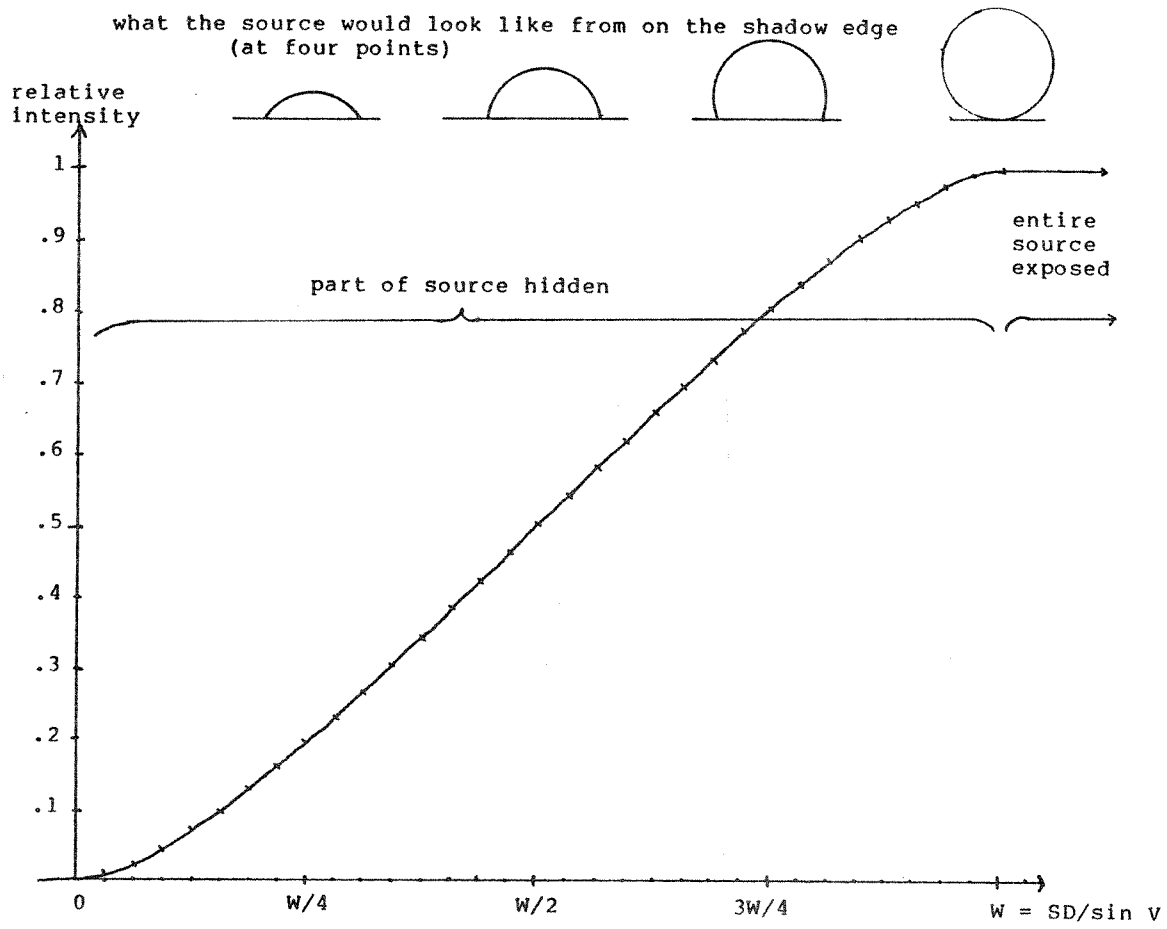

Figure 7) The intensity profile for a theoretical cast shadow boundary.

profile will be perceived as a step. Otherwise, the shadow boundary will have a "width" of SR, as shown in Figure 4c, although the gradation in illumination extends beyond the shadow edge, as will be shown.

Analysis of this case is more complicated because the angle of incidence of the light upon the surface is small compared to S. Thus, the sine of that angle cannot be approximated by the constant term, sin V. Also, note that since the surface is not flat, the displacement across the edge (in the X direction), a, will be measured as R multiplied by an angle (i.e., the X axis is slightly curved):

$$\frac{di}{dA} = \frac{I}{R} \int_0^a (\sin \frac{a-x}{R})\ S \sin r\ dx, \quad \text{where}$$

$r = \arccos(1 - \frac{2x}{SR})$, as before, except that the edge width is SR instead of SD, and $(\sin \frac{a-x}{r})$ is analagous to sin V. Then,

$$\frac{di}{I\ dA} = \frac{S}{R} \int_0^a (\sin \frac{a-x}{R})\ \sin \arccos(1 - \frac{2x}{SR})\ dx, \qquad 0 < a < W$$

Letting $y = \frac{x}{R}$ (see Figure 4b), we have:

$$\frac{di}{I\ dA} = S \int_0^{a/R} \sin\ (\frac{a}{R} - y)\ \sin \arccos\ (1 - \frac{2\ y}{S})\ dy, \qquad 0 < a < W$$

$$= S \int_0^{a/R} \sin\left(\frac{a}{R} - y\right) \sqrt{1 - \left(1 - \frac{2y}{S}\right)^2}\, dy,$$

$$0 < a < W$$

$$= 2 \int_0^{a/R} \sin\left(\frac{a}{R} - y\right) \sqrt{y\,(S - y)}\, dy,$$

$$0 < a < W$$

This function, integrated numerically*, is tabulated in table 1a and plotted in figure 8 verses a/SR with S=0.54 degrees (the angle subtended by the sun).

Because the surface is not flat, the effective illumination is not constant for X greater than W, but rather increases due to the sin ((a-X)/R) term. The form is nearly the same:

$$\frac{di}{I\, dA} = S \int_0^{S} \sin\left(\frac{a}{R} - y\right) \sin \arccos\left(1 - \frac{2y}{S}\right) dy,$$

$$w < a < \{\text{That point where the radius of curvature begins to differ significantly from R}\}$$

$$= 2 \int_0^{S} \sin\left(\frac{a}{R} - y\right) \sqrt{y\,(S - y)}\, dy$$

This function, integrated numerically* is tabulated in Table 1b and plotted in the right part of Figure 8. Note that it is

*For a << R, sin (a/R - y) can be approximated by (a/R - y), thus rendering the function integrable (make the substitution z = y - S/2).

DELTA= 9.4248E-06 S= 9.4248E-03

X/SA	VALUE	+/-	NORMALIZED	+/-
5.0000E-02	2.4632E-10	7.1430E-12	7.4929E-04	2.1728E-05
.1000	1.3801E-09	2.0040E-11	4.1981E-03	6.0958E-05
.1500	3.7613E-09	3.6508E-11	1.1442E-02	1.1105E-04
.2000	7.6323E-09	5.5733E-11	2.3217E-02	1.6953E-04
.2500	1.3174E-08	7.7230E-11	4.0075E-02	2.3492E-04
.3000	2.0526E-08	1.0066E-10	6.2438E-02	3.0620E-04
.3500	2.9794E-08	1.2577E-10	9.0629E-02	3.8257E-04
.4000	4.1056E-08	1.5235E-10	.1249	4.6342E-04
.4500	5.4367E-08	1.8024E-10	.1654	5.4826E-04
.5000	6.9758E-08	2.0929E-10	.2122	6.3664E-04
.5500	8.7241E-08	2.3940E-10	.2654	7.2822E-04
.6000	1.0680E-07	2.7044E-10	.3249	8.2266E-04
.6500	1.2842E-07	3.0234E-10	.3906	9.1970E-04
.7000	1.5202E-07	3.3502E-10	.4624	1.0191E-03
.7500	1.7755E-07	3.6839E-10	.5401	1.1206E-03
.8000	2.0488E-07	4.0240E-10	.6232	1.2241E-03
.8500	2.3388E-07	4.3699E-10	.7114	1.3293E-03
.9000	2.6437E-07	4.7211E-10	.8042	1.4361E-03
.9500	2.9611E-07	5.0771E-10	.9007	1.5444E-03

Table 1a

1.000	3.2874E-07	5.4375E-10	1.000	1.6540E-03
1.200	4.6024E-07	6.9163E-10	1.400	2.1039E-03
1.400	5.9173E-07	8.4410E-10	1.800	2.5677E-03
1.600	7.2322E-07	9.9987E-10	2.200	3.0415E-03
1.800	8.5471E-07	1.1580E-09	2.600	3.5226E-03
2.000	9.8620E-07	1.3180E-09	3.000	4.0092E-03
2.500	1.3149E-06	1.7231E-09	4.000	5.2414E-03
3.000	1.6436E-06	2.1327E-09	5.000	6.4874E-03
4.000	2.3009E-06	2.9587E-10	6.999	8.9998E-04
6.000	3.6147E-06	4.6212E-10	11.00	1.4057E-03
8.000	4.9272E-06	6.2874E-10	14.99	1.9125E-03
10.00	6.2380E-06	7.9534E-10	18.97	2.4193E-03
15.00	9.5042E-06	1.2108E-09	28.91	3.6831E-03
20.00	1.2749E-05	1.6238E-09	38.78	4.9394E-03
30.00	1.9148E-05	2.4383E-09	58.24	7.4168E-03
40.00	2.5376E-05	3.2312E-09	77.19	9.8286E-03
60.00	3.7103E-05	4.7243E-09	112.9	1.4370E-02
80.00	4.7516E-05	6.0501E-09	144.5	1.8403E-02
100.0	5.6246E-05	7.1616E-09	171.1	2.1784E-02
125.0	6.4327E-05	8.1904E-09	195.7	2.4914E-02
150.0	6.8853E-05	8.7666E-09	209.4	2.6666E-02

Table 1b

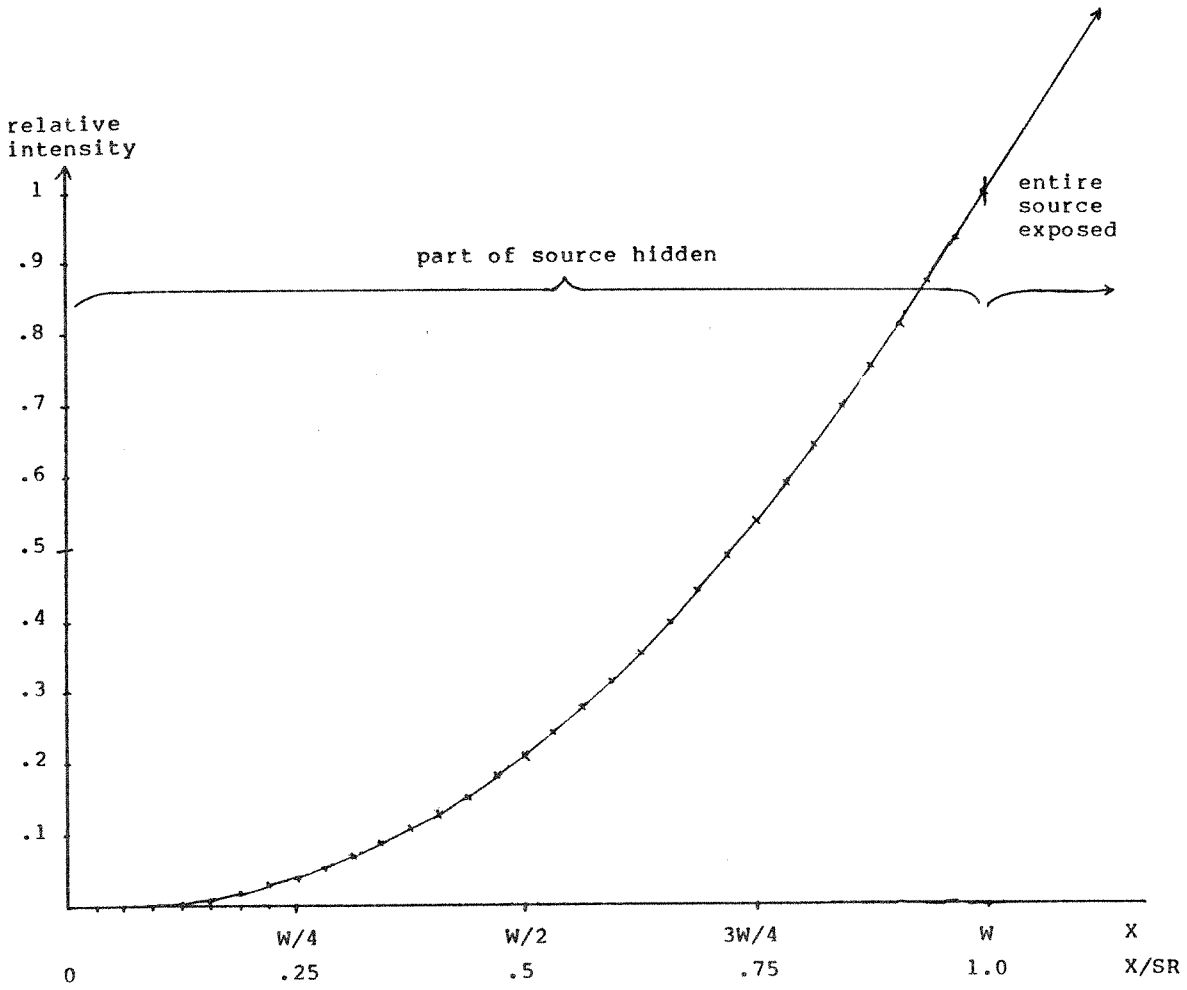

Figure 8) The intensity profile for a shadow edge on a self-shadowing object.

very different from the intensity profile predicted for cast-shadow edges. For cast-shadows, the incident intensity is approximately constant for points exposed to the entire source (i.e., for $a > W$). For shadows due to self-shadowing on a curved surface, however, the incident intensity increases with the sine of the angle of incidence, which is approximately linear near the boundary. Thus, in theory, at least, it is possible to identify the two types of shadow boundaries by examining the intensity profiles of their cross-sections.

SECTION III

APPLICATION

To be of much value, an analysis like that of Section II should enable us to identify features in real images. We sought to apply our observations about shadow boundaries by writing programs which could identify shadows by examining the intensity profiles across their edges. These routines were then to have been incorporated into a general purpose region-growing system which would be able to identify and merge shadowed and unshadowed regions of the same object.

We wrote a system of interactive programs for use on UT's Control Data time sharing system. They allow the display of small grey scale images on the screen of a Tektronix 4012 or 4010 graphics terminal. A user can identify boundaries by marking endpoints with the terminal's crosshairs. Graphs of the intensity profiles of the perpendicular cross-sections of the boundaries are

then displayed on the screen of the 4012. The cross-correlations (convolutions) of these profiles with arbitrary prototypical edges can also be graphed.

Some of our images were obtained by digitizing 35mm slides using a flying spot scanner (Underwood et al [9]) controlled by a PDP-11/34, the rest were digitized using a high resolution television camera and the same computer. The images were typically 128x160 pixels, either monochromatic or with three color bands of six significant bits each.

Although the programs can work with any or all of the three color bands, or with "normalized colors" (e.g., R/(R+G+B)), we have generally used the sum of the three colors (the intensity) when examining edge intensity profiles. Linear interpolation was used where it was necessary to obtain intensity measurements for arbitrary picture coordinates.

Most of our attention was devoted to examination of cast-shadow edges and object edges in a scene

containing blocks and other small objects (Figures 9 through 11). We also examined an aerial photograph of an airplane tail on a runway and several other outdoor scenes containing shadows, as well as a number of close-ups of the two types of shadow boundaries. We used the convolution technique of Marr [6] to compute first and second derivative approximations of the intensity profiles. The results for a fairly typical edge are shown in Figures 12 through 14. A wide step, for the first derivative, or bar convolution profile, for the second derivative, was used to improve noise immunity by including more points in the computation.

We were readily able to find the approximate midpoints of the shadow boundaries in both cases; they correspond to peaks in the first derivative and zero-crossings in the second derivative. Measurement of the widths by measuring the distance between second derivative peaks in the natural scenes, though, proved unreliable. The difficulty was apparently due to the presence of noise and the limited resolution of our images; a mask width of at least 20 pixels was required to make the peaks clear, and this is as wide as most of the shadow edges.

In hopes of obtaining improved edge width measurements for degraded pictures, and of discriminating between the two types of fuzzy shadow edges, we attempted

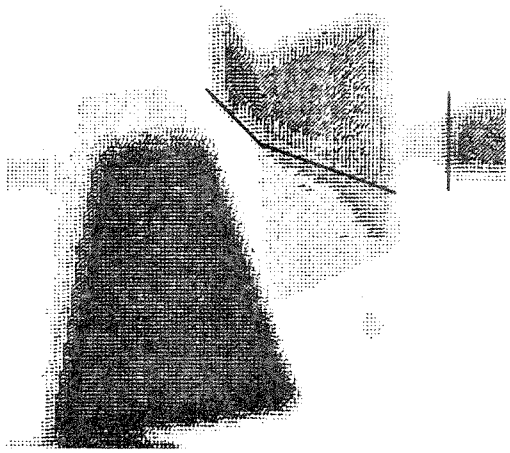

Figure 9) The (negative) blocks scene with three shadow edges identified. Two blocks and a hammer handle are visible. The hammer handle has apparently shrunk in the shadowed region (due to blurring). This image was digitized with our flying spot scanner system [see Underwood].

Figure 10) A (negative) closeup of the three shadow edges in the blocks scene.

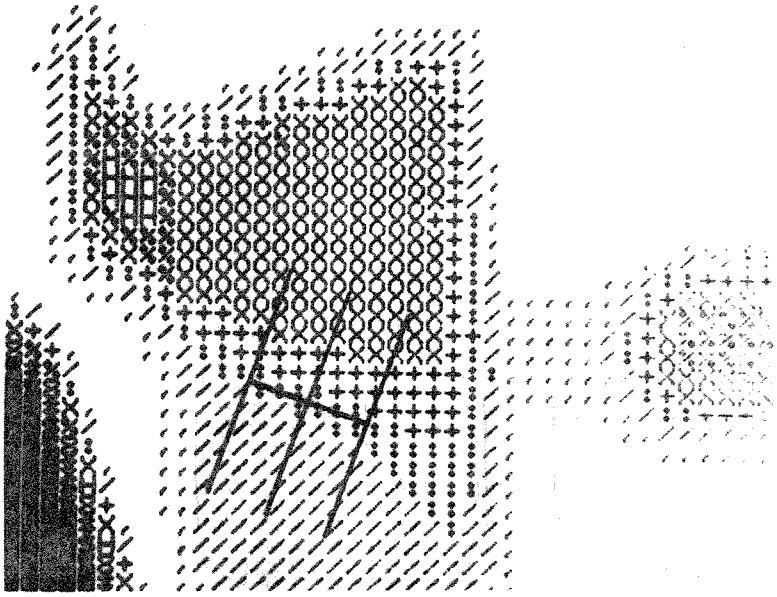

Figure 11) This is a Tektronix 4012 psuedo grey-scale reproduction of the (negative) blocks scene with one of the shadow edges and three cross-sections marked. The bottom side of the edge is in shadow. The image is displayed at 1/4 resolution because of the limited screen size.

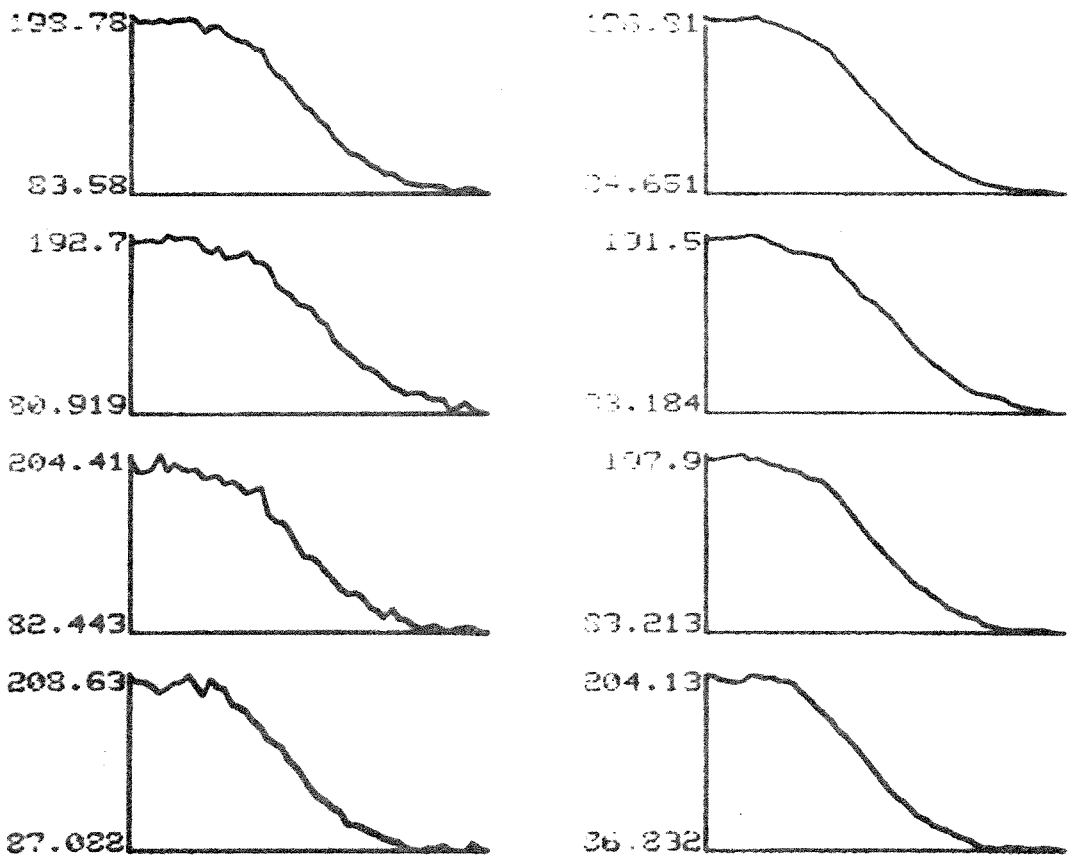

Figure 12) The intensity profiles for the cast-shadow cross-sections in figure 11. The left column shows the actual intensities; the data in the right column has been smoothed by averaging over groups of four points. The cross sections in figure 11 are profiled in the bottom three pairs of graphs; the top two graphs are averages over the three cross-sections graphed beneath them. Thus, the top right graph is the result of averaging in two directions. The bottom intensity profiles correspond to the left-hand cross-section in figure 11. The X axis is 50 points long.

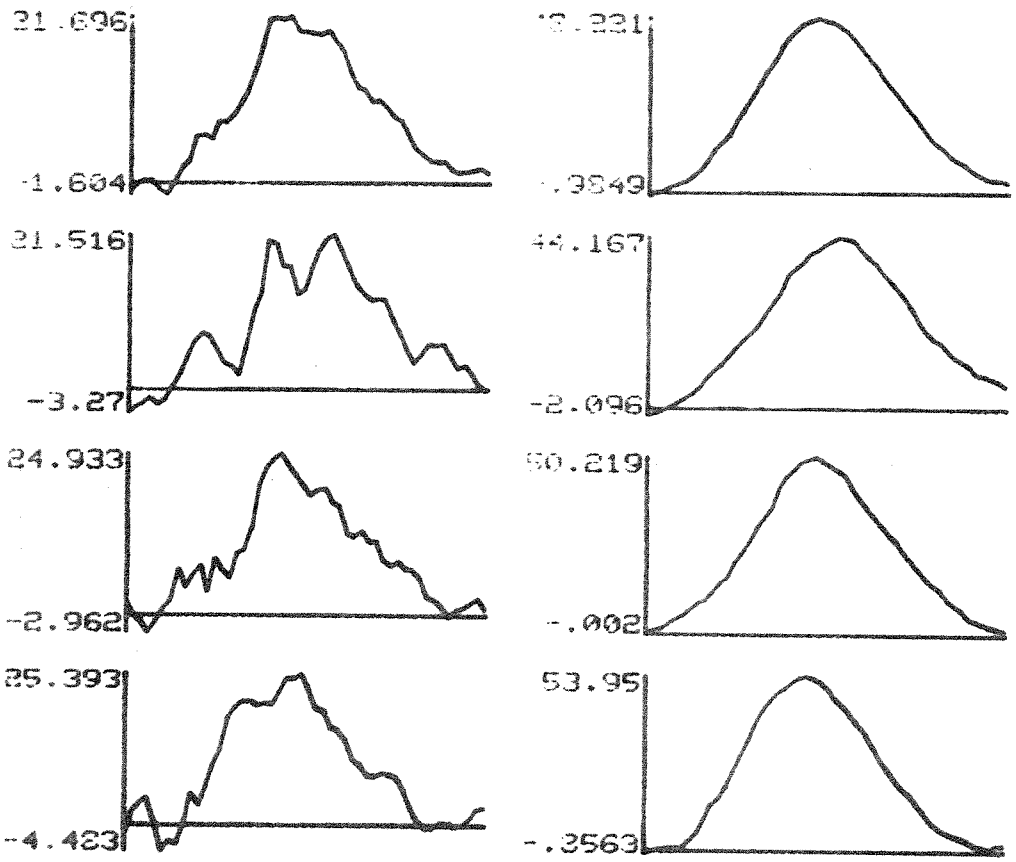

Figure 13) This is the result of cross-correlating step-edge masks of width 8 (left) and 20 pixels (right) with the intensity profiles in the left column of figure 12.

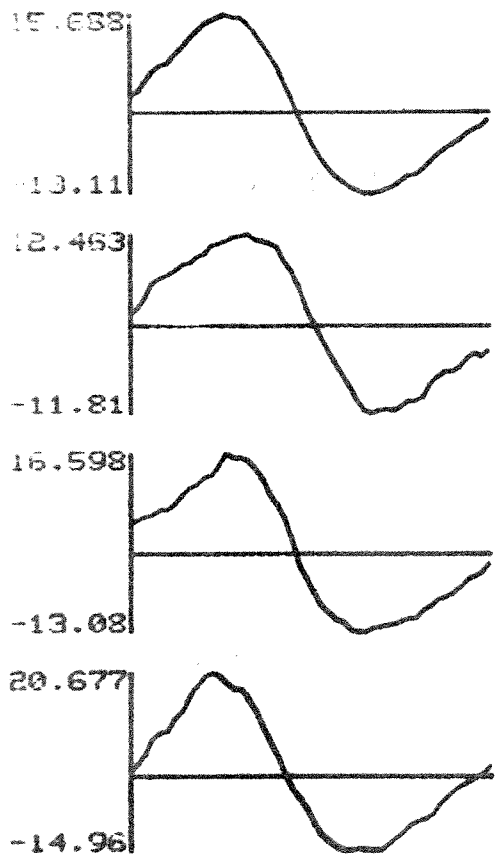

Figure 14) This is the result of cross-correlating a bar-mask of width 30 with each of the profiles in the left column of figure 12. This process is a kind of second-derivative (Laplacian) operator. The distance between the positive and negative peaks is a metric of the edge width. The "width" of the bottom graph (which corresponds to the sharpest part of the edge) is not appreciably less than that of the other two. A gradation in width would be evidence of a shadow edge.

to directly match theoretical shadow edge profiles to the observed profiles. A high resolution cast-shadow profile was synthesized according to our model and stored in an array for this purpose. Scaling and position parameters were adjusted to best fit the theoretical profile to the observed profile. The summed squared error was used as the measurement of the quality of fit. A steepest descent ("hill climbing") non-linear optimization algorithm using finite difference approximations to the first derivatives (and modified to prevent jamming) (Luenberger [10]) was used to adjust the parameters.

Unfortunately, although the algorithm converged to a solution after an acceptable amount of computation, the parameters found were frequently unreasonable. As with the second derivative approximation measurements, the problem was apparently caused by image degradation.

Since the profile shape and, hence, the degree of fit to an ideal shadow cross-section will be affected by image degradation, it is appropriate at this point to consider the effects which various types of image degredation can be expected to have on attempts to discriminate and parameterize shadow boundaries by this technique. Clearly, a combination of noise and low resolution can be fatal to the success of the method.

Blurring, however, need not be. Obviously, blurring the image will alter the observed intensity profiles along both shadow and object boundaries. It is well known (Rosenfeld and Kak [11]) however, that some types of blurring (due to camera movement or defocussing, for instance) are readily reversible. The techniques used to deblur whole pictures, though, can be complex and time consuming. Fortunately, they are not needed for this method. It is much simpler to deblur the one-dimensional intensity profile perpendicular to the edge than to deblur the entire image. To deblur the profile, we need a model of the blurring process. For most types of blurring, a reasonable model is that each blurred picture point is a weighted average of nearby points in the unblurred image. If we examine only cross-sections where the boundary is locally straight, the blurred picture point is the weighted average of unblurred points along the cross-section. For example, if B_n is the grey-level of the blurry cross-section at pixel n and U_n is the corresponding grey-level for the unblurred cross-section, a width 5 model for the blurring operation might be:

$$B_n = .1U_{n-2}+.2U_{n-1}+.4U_n+.2U_{n+1}+.1U_{n+2}$$

Determination of the actual coefficients requires either prior knowledge of the nature of the blurring operation or examination of a blurred boundary for which

the deblurred profile is known. This would normally be an object boundary (which should appear as a step-profile). Usually, it may be assumed that the sharpest edge in the picture is on such a boundary. If it is known that the entire picture has been subjected to the same blurring, the coefficients are easily determined. If, however, the blurring is due to camera motion, the sharp object boundary chosen should be approximately parallel to the boundary to be deblurred, or, if the direction of motion is known, the cross-sectional profile can be scaled appropriately.

Once the coefficients are known, reversing the process (and thus deblurring the blurry cross-section) is done by solving a large system of simultaneous linear equations. This is easily done (without a matrix inversion) if either end of the cross-section extends into an area of constant grey-level, so that the first K-1 deblurred grey-levels are known (and equal to the blurred values), where K is the width of the blurring operation model. For our width 5 example, we start with U_1 through U_5 known. Then,

$$U_5 = (B_3 - .1U_1 - .2U_2 - .4U_3 - .2U_4)/.1$$

$$\text{and, } U_n = (B_{n-2} - .1U_{n-4} - .2U_{n-3} - .4U_{n-2} - .2U_{n-1})/.1$$

In practice, this solution technique would be very sensitive to noise, and so would have to be modified somewhat for robustness.

SECTION IV

CONCLUSION

We concluded that, for most applications, edge classification on the basis of cross-sectional profiles does not work very well. The limited resolution generally available for real applications implies that the width of a typical fuzzy edge is only a few pixels, and noise distorts the profile beyond recognition, making reliable width measurements impossible. Where greater resolution and improved signal to noise ratios are available, the techniques described here may be useful for discriminating shadows. Also, in such cases, it may be possible to eliminate the effects of blurring in the imaging process quite easily. In any case, the method shows promise mostly as a component in a comprehensive image processing system, providing one of many types of clues to aid in understanding the image.

REFERENCES

1. Rosenfeld, Hummel and Zucker, "Scene Labeling Using Relaxation Operations", IEEE-SMC, Vol 6, pp. 420-433.

2. Barrow, H. and Tenenbaum, J., "Recovering Intrinsic Scene Characteristics from Images", Computer Vision Systems, A. Hanson and E. Riseman, eds. (Academic Press, N.Y., 1978)

3. Marr, David, 1975, "Early Processing of Visual Information", Philosophical Transactions of the Royal Society of London, Vol. 275B, pp. 483-524.

4. Horn, B.K.P., 1977, "Understanding Image Intensities", Artificial Intelligence, Vol. 8, pp. 201-230, April 1977.

5. Waltz, D.L., "Generating Semantic Descriptions from Drawings of Scenes with Shadows", MIT Artificial Intelligence Lab., Technical Report AI-TR-271, 1972.

6. Marr, David, 1974, "The Low Level Symbolic Representation of Intensity Changes in an Image", MIT Artificial Intelligence Memo 325.

7. Marr, D. and Hildreth, E., "Theory of Edge Detection", MIT Artificial Intelligence Memo, April, 1979.

8. Teschan, Paul E., Jr., Department of Physics, University of Texas at Austin, Personal communication, 1980.

9. Underwood, Williams, Costa and Aggarwal, "A Flying Spot Scanner System", Instrument Society of America Trans., Vol. 14, No. 4, pp. 340-357, 1975.

10. Luenberger, David, Chapter 7 of Introduction to Linear and Nonlinear Programming, Addison-Wesley, Reading, MA., 1973.

11. Rosenfeld and Kak, Digital Picture Processing, Academic Press, NY, 1976.