MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

AIM 422

May 1977

LIGHT SOURCE EFFECTS

by

K. Forbus

ABSTRACT:The perception of surface luster in achromatic single view images seems to depend on the existence of regions with source-like properties. These regions are due to the interaction of the specular component of the surface's reflectance and the illumination. Light source effects are broken down into three catagories according to gross aspects of the physical situation in which they occur, and criteria for detecting the regions they cause are suggested.

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

Introduction

From the gleam in a lover's eye to the dazzling brilliance of breaking waves, we are surrounded by examples of non-matte surface luster. The luster of a surface is a quality of its appearrance. Specific terms we use for describing surface luster include glossy, shiny, sheen, metallic, and matte. Since surface luster is often important for identifying materials, and is a conspicuous feature of human visual perception, it is important to understand how to compute these qualities.

Traditional theories dealing with the perception of suface luster have focused on factors such as disparity, color constancy, and lightness(see for example [Evans, 1974] [Lavin, 1973]), but none of them sufficiently accounts for human performance in describing surface luster. This conclusion comes from observing that people recognize non-matte lusters in both achromatic television images and in photographs, where neither color nor disparity are present. Our reasons for dismissing lightness will become clear later on. On what properties then do these descriptions rest?

For non-matte lusters, the interaction of the illumination with the specular component of the object's reflectance results in the formation of small regions or sub-regions which perceptually seem to have source-like properties. If we can detect which regions in an image have source-like properties, we may use their existence to provide a symbolic modifier to the surfaces that contain them. This modifier, we suggest, is the surface luster. A secondary consequence is that we will have simplified object recognition by "factoring out" some effects of the illumination.

This idea was also put forward by Beck[Beck, 1974] in similar terms. Beck also exhibited a strong piece of evidence for this view, in the form of a picture of a vase with and without highlights. (reproduced in figure 1) The vase with highlights looks shiny, the vase without looks matte. Other images with non-matte lusters are in figure 2. A light source effect is the interaction between the illumination and the specular component of a surface's reflectance that gives rise to source-like regions. This work classifies light source effects according to how they arise, and suggests methods for detecting the regions they produce.

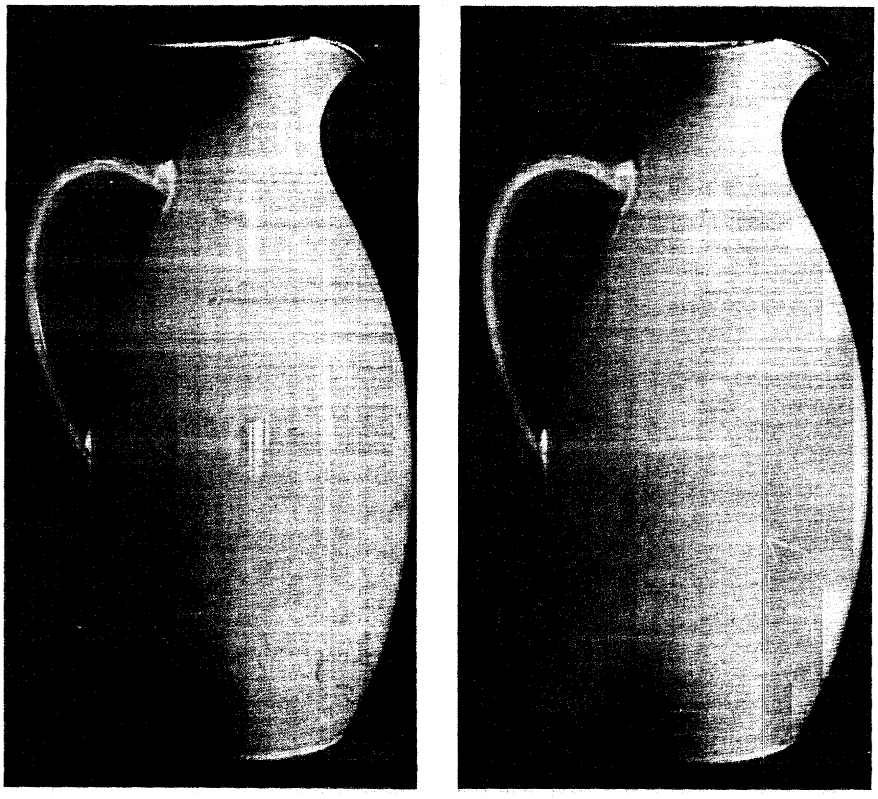

Figure 1

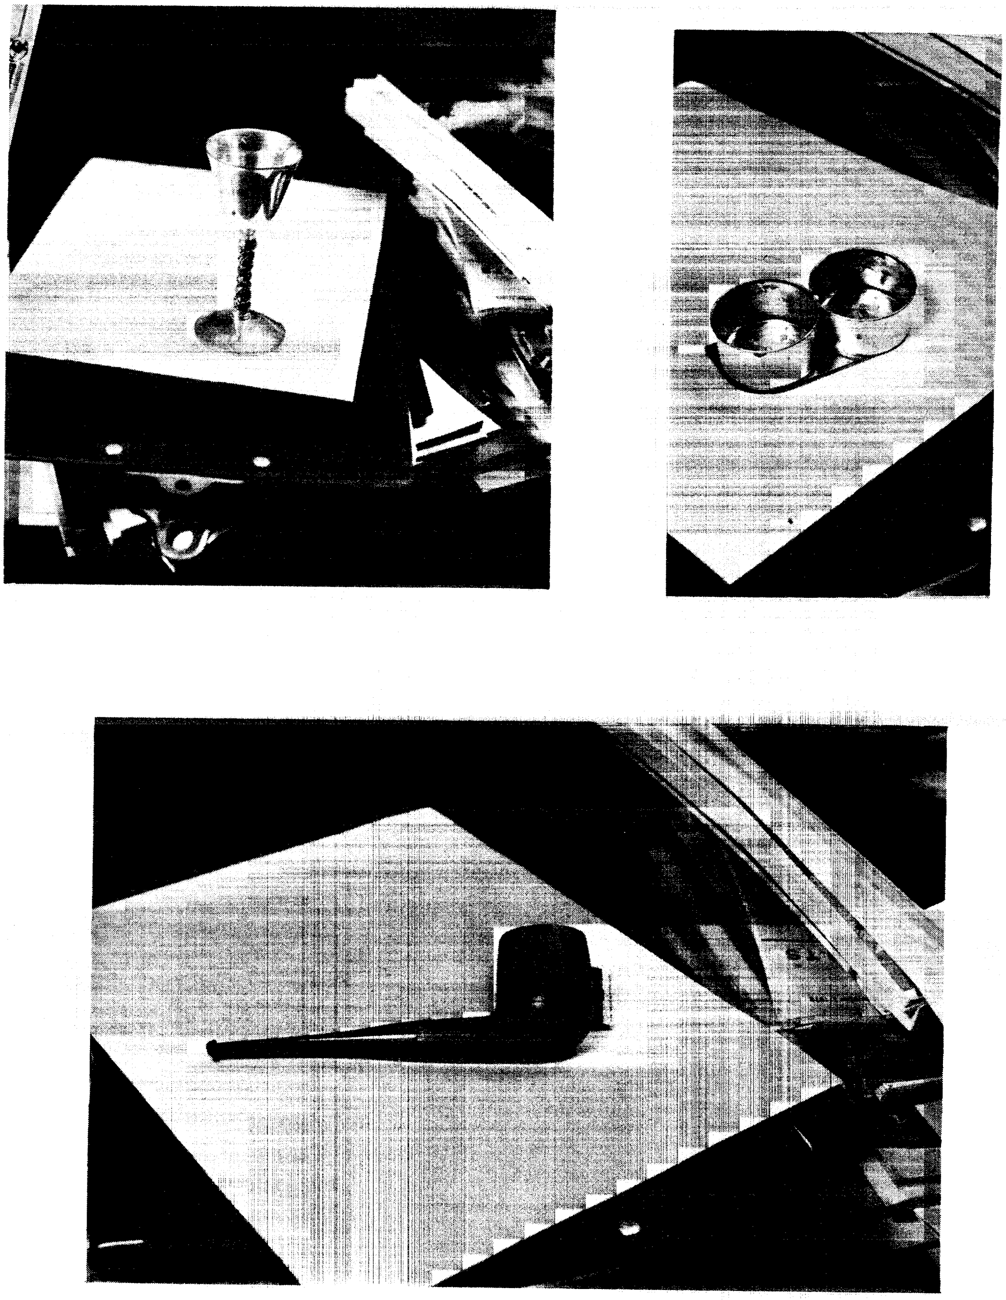

Figure 2

Some Conventions

Let **n** be the unit normal to an element of surface. We then define the angles between this vector and the source and view vectors as follows-

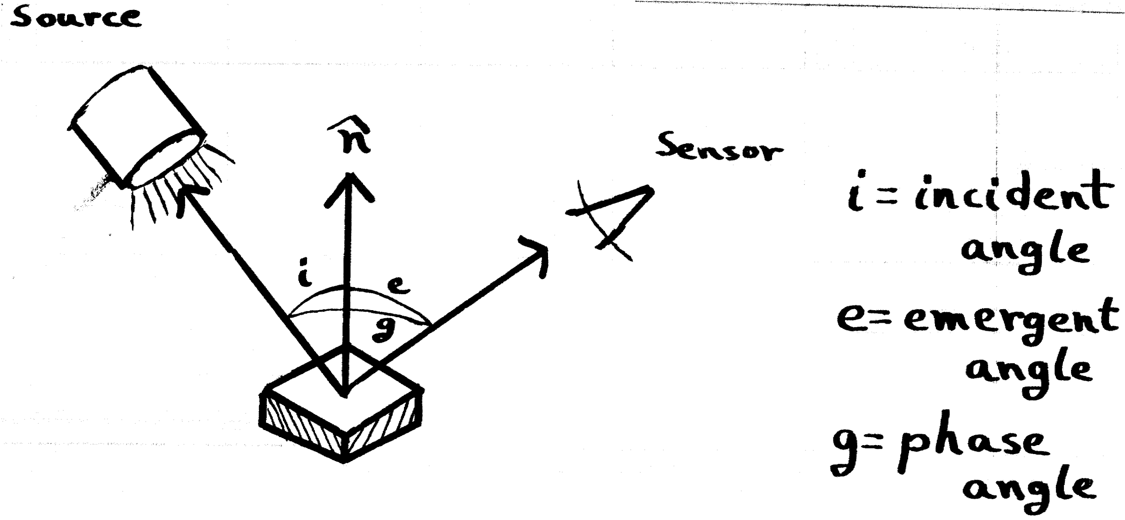

The model for surface reflectivity used here was taken from [Horn,1975], and is

$$R(I,E,G)=s(2IE-G)^{n}+mI$$

where

$I=\cos(i)$

$E=\cos(e)$

$G=\cos(g)$

s=relative strength of specular component
[between 0 and 1]

m=relative strength of matte component
= 1-s

n=degree of sharpness of specular component

An analysis of how well this model matches real surfaces may be found in [Forbus,in prep].

Analysis of the problem

There is a natural division of light source effects into three catagories, by the gross aspects of the physical situation in which they occur. Non-mnemonic names are given to these effects to avoid confusion with perceptual terms. Below is an overview.

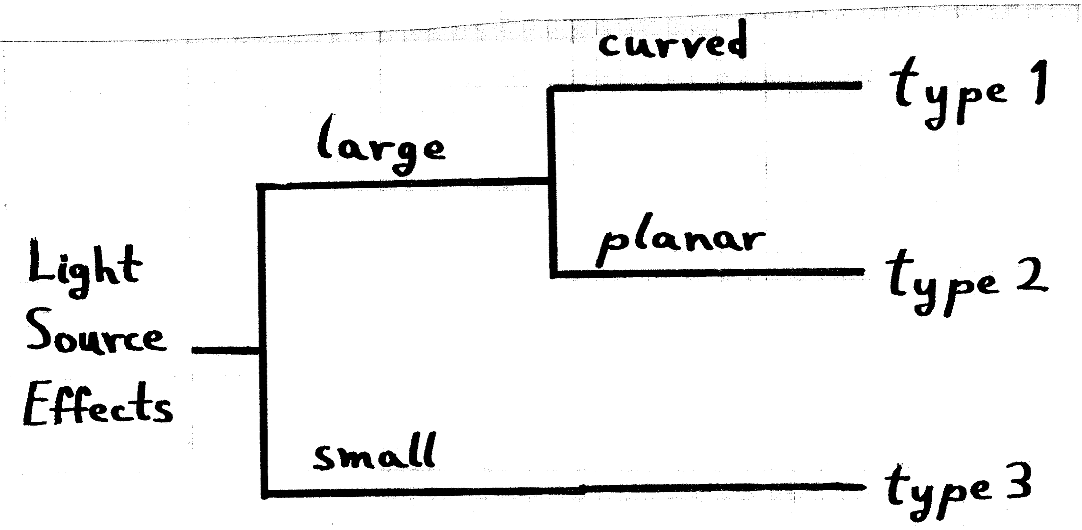

Type 1 effects are those that are due to an object's curvature. In perceptual terms they correspond to the highlights on curved objects. Some examples of regions they cause appear in figure 3.

Type 2 effects are those which are caused through the observer viewing a virtual image of a source by the specular component of a planar surface. Perceptually they correspond to shiny patches and other highlights. Examples of these regions may be found in figure 4.

Type 3 effects occur when the region that is specularly reflecting the source is small, or when the virtual image of the source is small. They often correspond perceptually to point specularities. An example is in figure 5.

Because of these physical differences, different detection methods are required for each type. Let us now analyze each in turn.

Figure 3

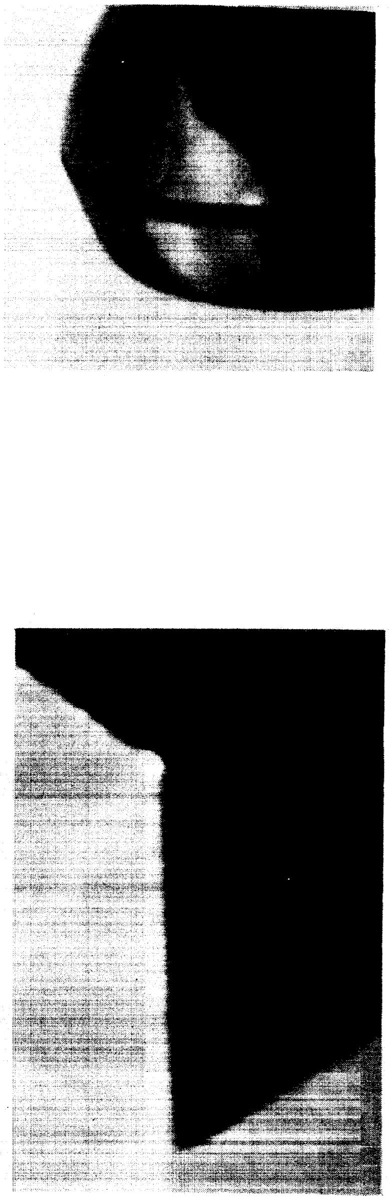

Figure 4

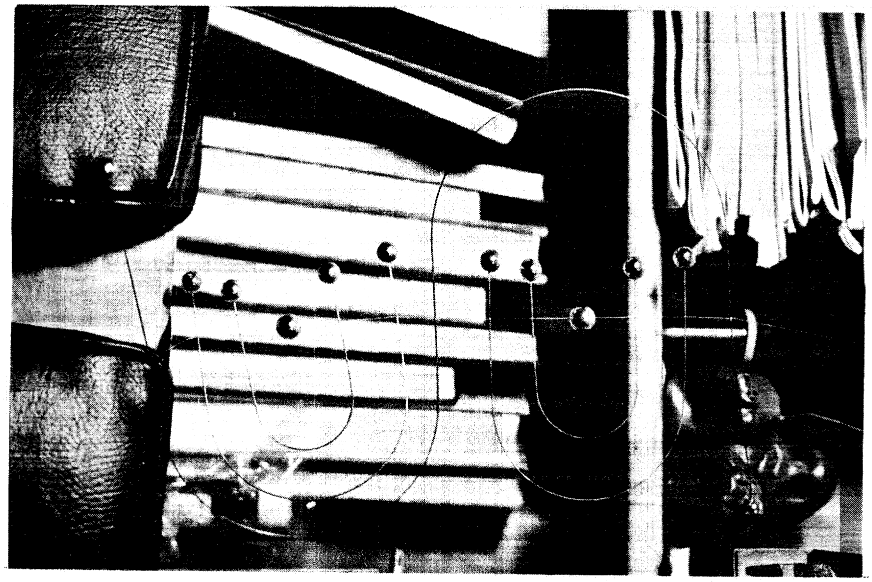

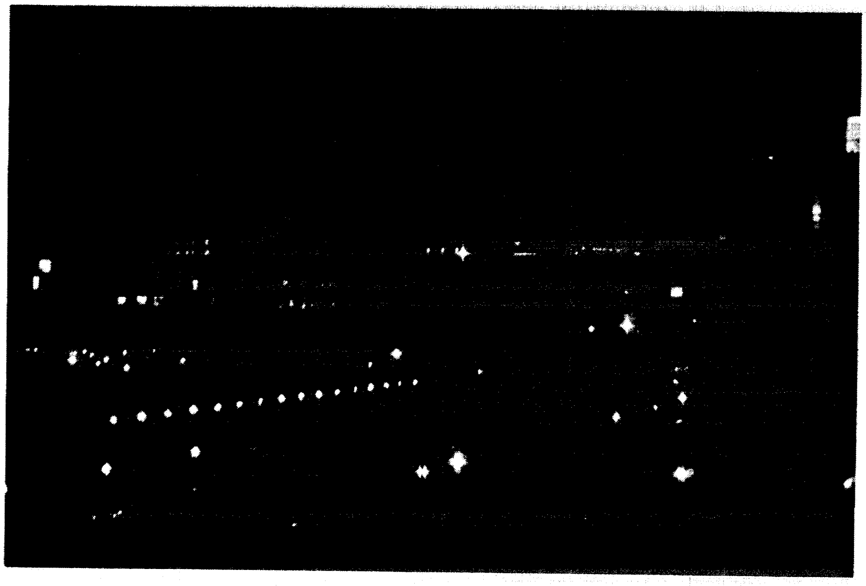

Figure 5

Type 1 Light Source Effects

A light source reflected in a curved specular surface gives rise to a highlight. This is caused by a type 1 light source effect.

I have observed that humans see non-linear intensity gradients as caused by either surface curvature or non-uniform illumination. To recognize highlights, one must be able to distinguish them from object curvature. Let us examine a simple situation for insight.

In figure 6, we have the equations derived for the intensity across a one dimensional profile, assuming both the source and the viewer are very far away from a uniformly curved surface. We see that the object will give rise to a "hill" in intensity, with a seperate peak superimposed on top. The position of the peak and its extent correspond with the places on the surface where the source is specularly reflected. Figure 7 shows the same equation graphed for various n and s. Notice that a lower n yields a wider peak, thus "smearing" the mirror effect of the specular component. By sweeping the surface normal, the curvature reveals to us the specular component.

To determine what parameters are relevant to perceiving highlights, a number of different intensity profiles were used to generate images, either on a tv screen or transparencies for viewing with a light table. I will call the contribution from the specular component the peak, and the rest of the intensity the surround. A set of pictures was produced with the equation derived in figure 6, and the results looked like reasonable cylinders(see figure 8).

Several other analytic approximations for the peak were tried on a constant surround, with different combinations of surround intensity and peak height. Since they all appeared as highlights, the exact shape of the peak cannot matter, so long as it is not a discontinuous step in intensity. The lowest contrast of the peaks was around 0.08, so contrast is not very important, although the highlight looks "better" with higher contrast.

Next a surround corresponding to the matte component of intensity of a cylinder was used, but the width of the peak as a percentage of the width of the surround was varied. All but the broadest (peak width 75% of the total width)appeared as highlights. This suggests that the width is not critical, so long as the surround is visible. A similar set of peak widths was then superimposed over a constant surround. On peaks that were wider than about 20% of the surround, there was no clear interpartation as a highlight.

These results indicate that context is an important factor in classifying the results of this type of light source effect. If the local context is right(second derivative of intensity non-zero), a peak will be interperted as a highlight, subject to the limits above. Otherwise the peak must be reasonably narrow and not a step change in intensity to be interperted in this manner.

both viewer
and source
at infinity

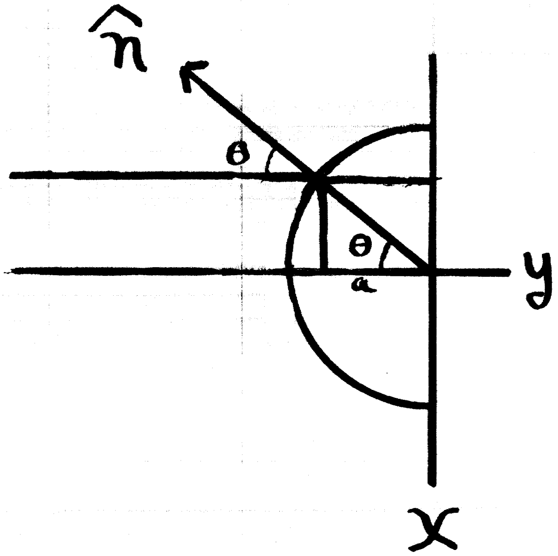

we can let $r=1$ with no loss of generality

then

$$\cos(\theta)=a \quad \text{, and} \quad a=\sqrt{1-x^2}$$

$$i=e=\theta,\ g=0 \quad \text{so } \cos(g)=1$$

so

$$\Phi(i,e,g)=s(1-2x^2)^n+m\sqrt{1-x^2}$$

Figure 6

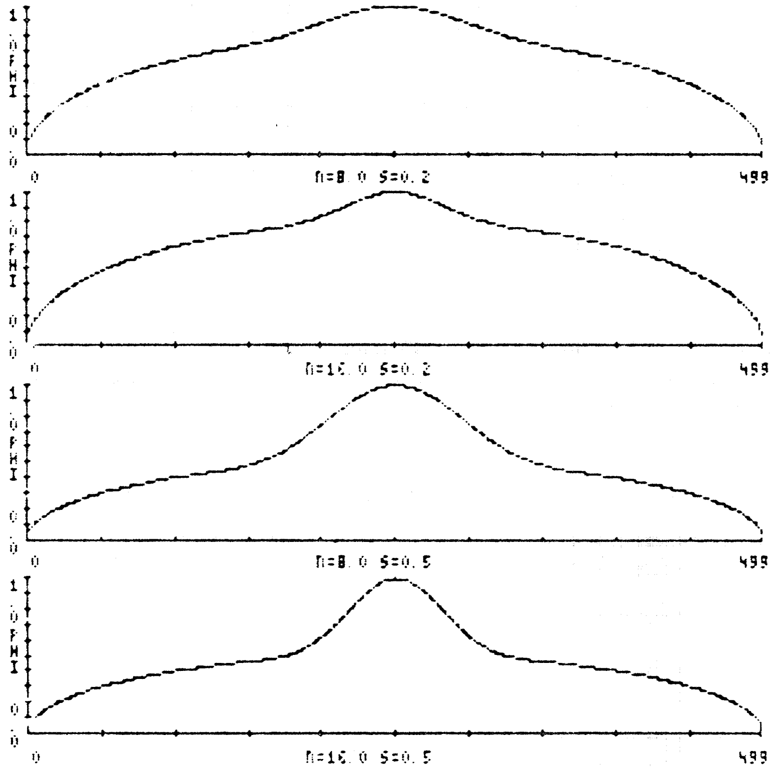
1.0
FHI
0.0
0
n=8.0 S=0.2
499
1.0
FHI
0.0
0
n=16.0 S=0.2
499
1.0
FHI
0.0
0
n=8.0 S=0.5
499
1.0
FHI
0.0
0
n=16.0 S=0.5
499

Figure 7

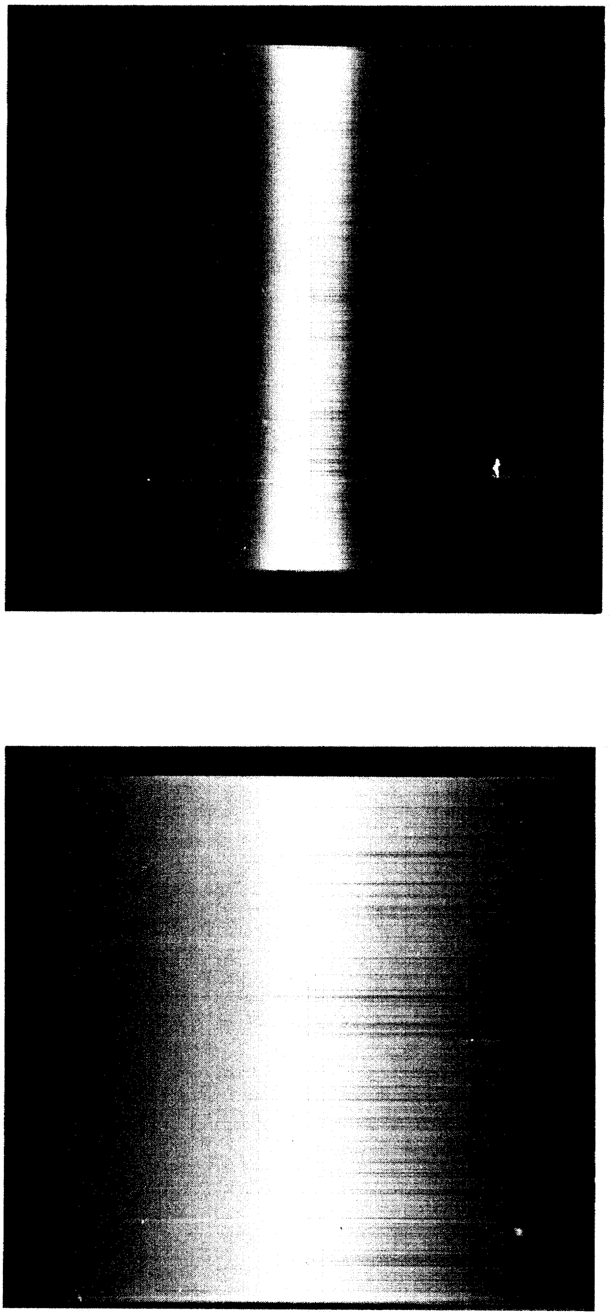

Figure 8

Type 2 Light Source Effects

When the surface being viewed is flat, a specular component will create a virtual image of the source at some set of viewing angles. This virtual image of the source is called a type 2 light source effect.

In the limiting case of a mirror ($s=1$), one can resolve all the details of the field of view for that angle. However most surfaces also contain a considerable matte component, so only very rough characteristics of the region, such as average local intensity and average local gradients are preserved.

A virtual image of a source should have the same properties as a real source. Hence, if we can detect light sources, we can detect type 2 light source effects. It has been claimed that source detection may be a side effect of a lightness computation[Lavin, 1973]. However, [Ullman,1975] demonstrated this is unfeasible, and defined an autonomous operator, the S-operator, which detects sources in an achromatic mondrian image.

The S-operator solves for the strength of a source by checking the intensities and gradients of the region and a neighbor, which are both assumed to have the same orientation and the same illumination. It is based on the observation that a source should have a higher intensity than its gradients would suggest. To make this notion more precise, I shall re-derive the S-operator here.

Let L=incident illumination

$r_{1,2}$=reflectance of surface 1,2

$I_{1,2}$=Intensity on surface 1,2

$G_{1,2}$=Intensity gradients across surface 1,2

where the intensity and gradients are measured at points close to the boundary between the two regions.

Then

$$I_1=(l_1r_1)+S \qquad I_2=L_2r_2$$

where S=source strength

$$G_1=G_0r_1 \qquad G_2=G_0r_2$$

$$L_2=L_1+G_0d$$

where d is the distance between points of measurement, and

G_0=Illumination gradient

we want to solve for S,

$$S=I_1-L_1r_1$$
$$S=I_1-r_1(L_2-G_0d)$$

then

$$S=I_1-((I_2G_1)/G_2)+G_1d$$

Notice that we have the source strength in terms of quantities which we can measure in the intensity. Now let us remove the restriction that the orientation of surfaces 1 and 2 remain the same. To do this we must create a new expression for the gradient across a surface.

There is a foreshortening of distance due to the viewing angle

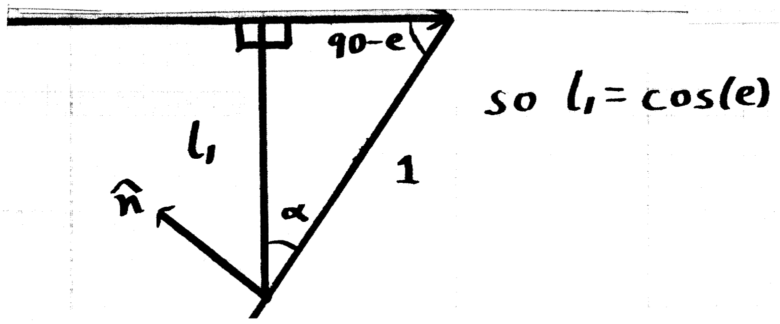

Thus the incident gradient should be divided by cos(e). We must also consider the reduction in intensity due to absorbtion and beam spread. For small distances it is resonable to use a linear approximation.

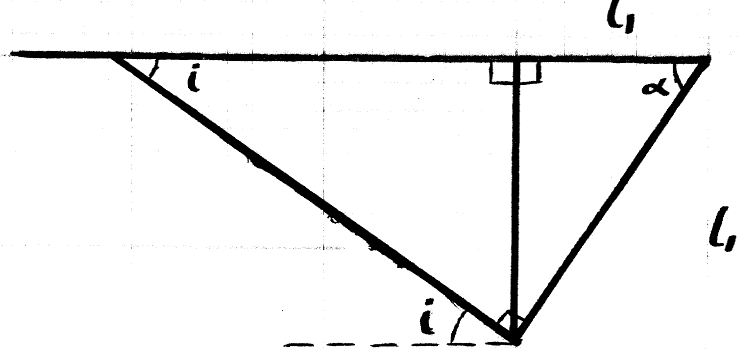

$\alpha = 90 - i$

$l_1 = \cos(\alpha) = \sin(i)$

So we now have the gradient G as proportional to sin(i)/cos(e). There are several other effects which are important, but since they are hard to put in an analytic form, and are less general, I shall postpone them until later. For now let us see the result of using our new expression for the gradient.

Let G_0 be a constant of proportionality which is specific to the illumination. Then

$$r_1G_0=G_1\cos(e)/\sin(i) \qquad r_2G_0=G_2\cos(e+\alpha)/\sin(i+\alpha)$$

where α=the angle between the planes. Then

$$S=I_1-(2G_1/G_2)[\cos(e)\sin(i+\alpha)/\cos(e+\alpha)\sin(i)]+G_1d(\cos(e)/\sin(i))$$

Notice that we no longer have S only in terms of quantities directly measured from the image. The introduction of a non-zero angle α between the surfaces forces us to consider their relative orientation. Therefore adding more regions will not help the analysis, for each region will introduce another angle.

We can estimate the effect of this new term by considering the maximum value of S we can get if there is no source involved.

$$S=I_1[1-(\cos(e)\sin(i+\alpha)/\cos(e+\alpha)\sin(i))+G_0d\cos(e)/\sin(i)]$$

which reduces to

$$S=I_1[(\sin(\alpha)\cos(i+e)+G_0d\cos(e)\cos(e+\alpha))/\sin(i)\cos(e+\alpha)]$$

We can of course take the partial derivatives, but they are not too revealing. When one considers the rational function which multiplies I_l, we discover first of all that d should be kept as small as possible for good accuracy, and that the only time it gets very large is when $i=0$ or $e+\alpha=pi/2$, which correspond to uniform lighting and to total obscuration, and thus is no surprise. In fact, the value rarely gets larger in magnitude than 1. Therefore we should be able to get a threshold which varies according to the intensities we measure, and use the s-value so computed as a discriminant.

Now let us re-consider the other factor affecting the intensity gradient. If the source spreads out spatially, the angle of incidence will be slightly different. An analysis of the effect this has on gradients may be found in [Herskovits & Binford, 1970].

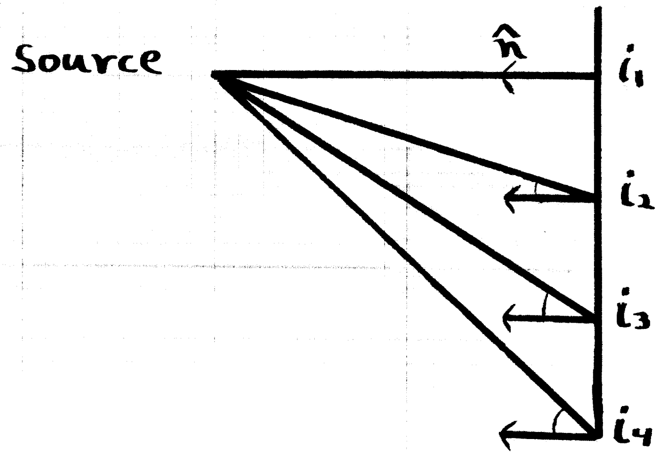

Even more important are the effects caused by the lack of spatial uniformity in actual light sources and the effects of mutual illumination. To handle these in any analytic fashion is very difficult. For an example of a simple case see [Horn,1975]. For this reason a better analytic model for gradients was not undertaken. There are some predictions which may be made however using the rough approximation.

Suppose we shine an additional source of light perpendicular to one surface, and not the other. Since this adds to the region's intensity without changing the gradient, the regions should then appear sourcelike. This should not be considered a drawback, for people perceive a source under the same conditions! Figure 9 is an example.

This effect was also noted by [Wallach, 1963].

A set of programs was implemented to run the S-operator at points in an image. It is evalutated at varying distances from an edge(for details of the parser, see [Marr, 1974]), ignoring values which are within an edge mask width of the center. A suitably conservative criterion is to consider only those s-values which exceed twice the sum of the intensities at the point of measurement. Some examples of the operator's results may be found in figure 10.

An interesting sidelight concerns the values the S-operator delivers when run across an edge which is caused by a shadow. High source values occur, which suggest that shadows might be aptly characterized as regions where the intensity is smaller than the gradients would suggest. Thus knowing the s-values might allow the classification of shadow regions. This conjecture has not been fully explored.

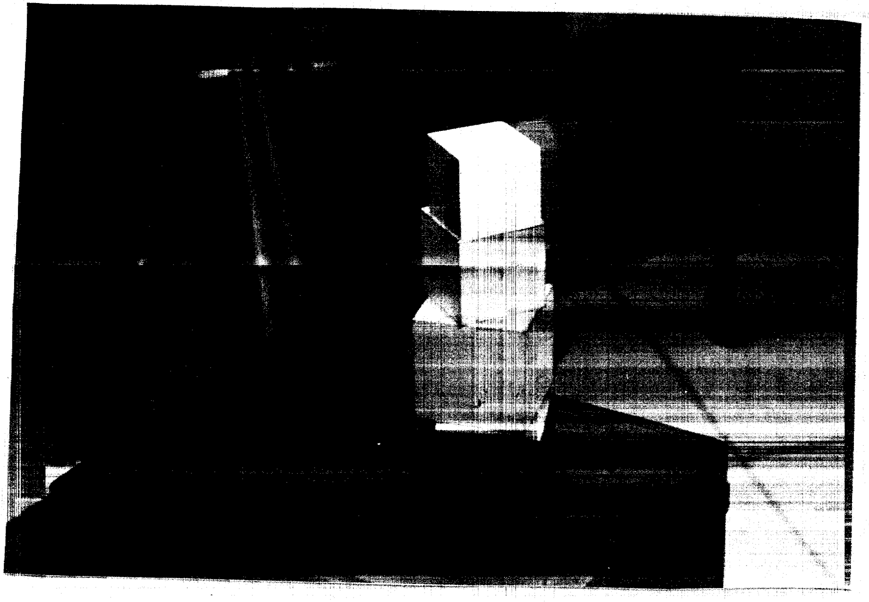

Figure 9

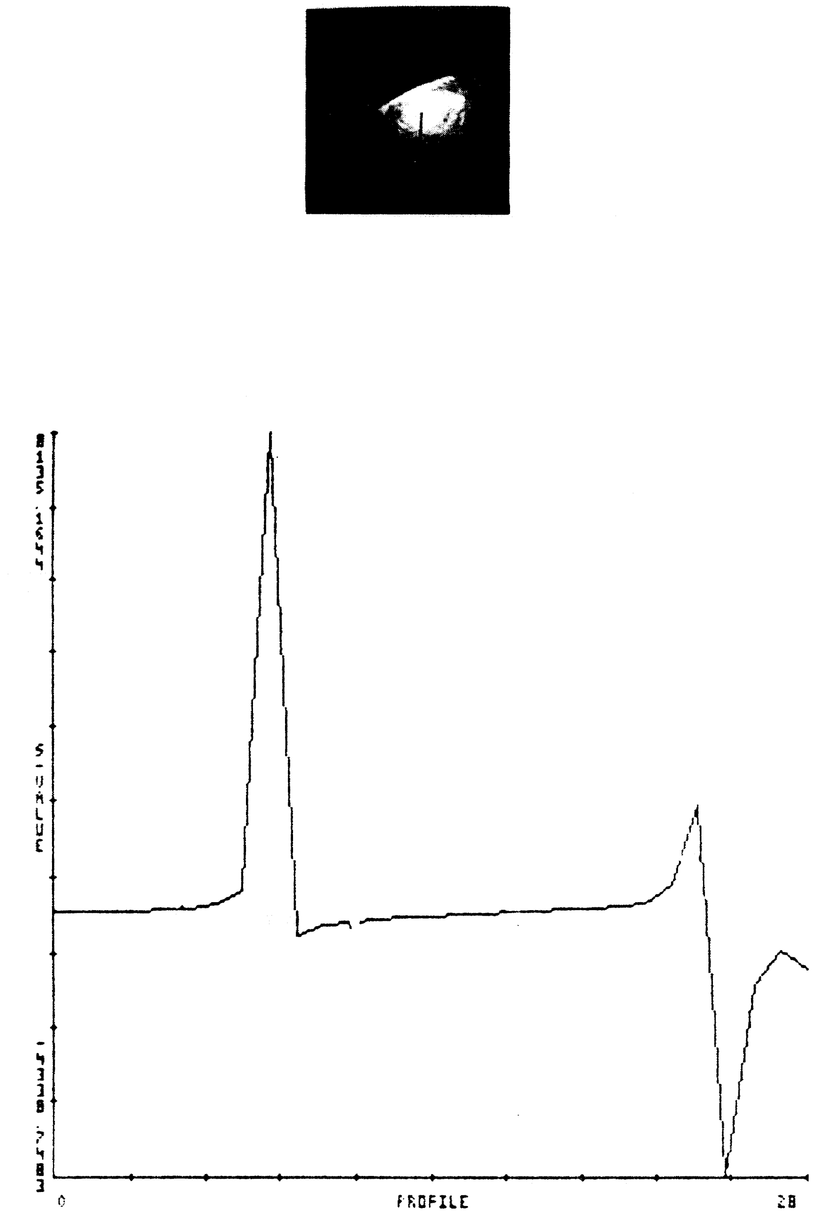

Figure 10(a)

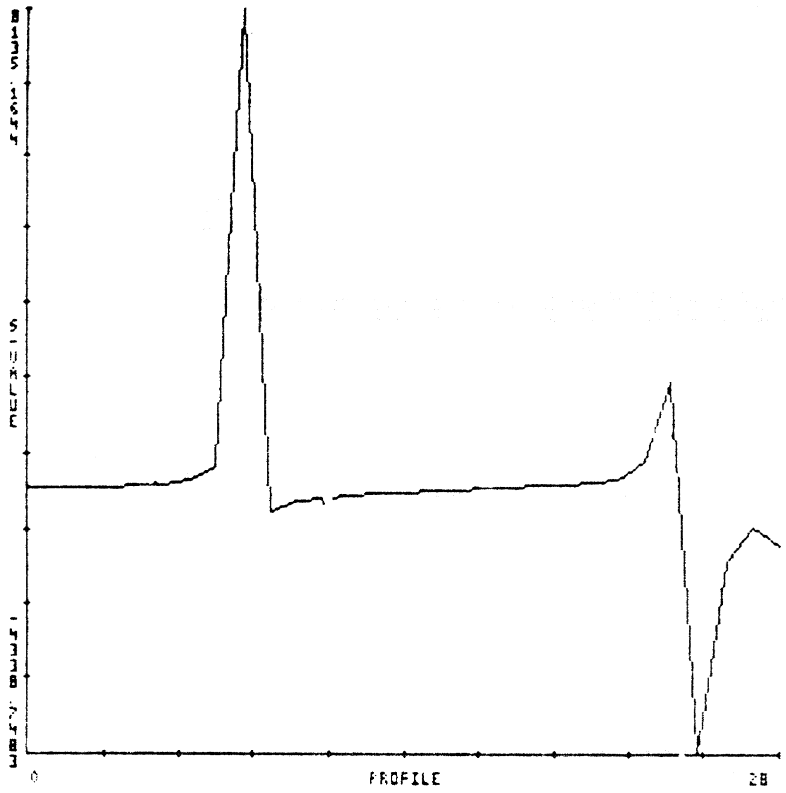

Figure 10(b)

Type 3 Light Source Effects

When the virtual image of the source is small, either because the source is far away or the surface reflecting it is small, a small bright blob occurs in the visual field. This is the type 3 light source effect.

Usually if one were able to look closer, the effect would be distinguishable as either a type 1 or type 2 effect. Since it is too small for adequately measuring gradients, one must use other techniques.

The only relevant property at first inspection is its contrast. To look for a contrast threshold, a surface with a small hole was painted with matte paint, and placed in front of a mirror. By using a variable-strength source and measuring the light coming off the mirror and the surface seperately the contrast of the blob thus formed was computed. The blob begins to look like a source when the contrast reached around 0.86(intensity ratio of about 14 to 1). Since this is rather high, we can conclude that high contrast is sufficient for a blob to be considered as the result of a type 3 light source effect.

One additional factor which must be considered is flare. Flare is the beam-like extensions from the image of a source, and is caused by diffraction in the imaging system. This makes it a very strong clue to a blob being caused by a light source effect, because below a certain diameter of spot size, the diameter is a function of incident intensity rather than angular size[Graham,1965]. Therefore flare is also a sufficient condition for ascribing a blob to a type 3 light source effect. One may find flare also in cases where Type 1 or Type 2 methods are applicable, especially if the reflected intensity is quite high, and should be regarded as a sufficient condition for those cases also.

Conclusions

Light source effects form valuable clues to the perception of surface luster. They cause source dependent regions in one of three ways, each of which is detectable in the context of single view achromatic images.

Suggestions for Future Work

1) People seem to distinguish between calling some region a source or a light source effect according to the context of the regions which contain it. For example, inducing a point specularity in the middle of a planar surface does not cause the surface to become shiny, yet in other cases such as an automobile bumper, it does. There may also be some ranges of contrast and/or s-values which correspond to "shine" rather than "source"(Cox, in preperation).

2) As mentioned earlier, the s-operator reacts strongly at shadow boundaries. This may enable a semi-local process to isolate which regions in an image are shadows, a result that would be nicely complementary to light source effects in deciding how a vision system should arrive at illumination assumptions, as well as simplifying recognition even more.

3) If one assumes that a well defined gradient and average intensity exist on a region, the ratio of the region's gradient and intensity is approximately proportional to sin(i)/cos(e), where the constant of proportionality is related to the illumination. This is perhaps the simplest quantity which can be measured locally in an image which directly links the illumination and the orientation of a surface, and therefore might be useful in giving a rough orientation estimate given an assumption about the illumination. A piece of evidence which supports the linking of illumination assumptions with orientation estimates is the Mach illusion [Beck, 1974].

Acknowledgements

I wish to thank David Marr for advice, encouragement, and many long talks. I would also like to thank Berthold K.P. Horn for our discussions of vision, and Patrick Winston for honing my writing and moral support.

This work has also profited from discussion with other workers in vision at the MIT Artificial Intelligence Lab, particularly Shimon Ullman, Mark Lavin, Bob Woodham, and Tomas Lozano-Perez. Karen Prendergast provided valuable assistence in darkroom work.

References

Beck,Jacob 1974 "*Surface color perception*" Cornell University Press

Evans, R.M. 1974 "*The perception of color*" N.Y. London, Sydney, Toronto:John Wiley & Sons

Forbus,K An analysis of reflectance models *in preparation*

Graham, C.H., editor, 1965 "*Vision and visual perception*" New York:John Wiley & Sons,

Herskovits & Binford 1970 On boundary detection *MIT A.I. Lab Memo No. 183*

Horn, B.K.P. 1975 Image intensity understanding *MIT A.I. Lab Memo No. 335*

Krakauer, Lawrence J. 1971 Computer analysis of visual properties of curved objects *MIT A.I. Lab Technical Report AI-TR-234*

Lavin, Mark 1973 The gloss of glossy things *MIT A.I. Lab Working Paper No. 41*

Marr, David 1974 The low-level symbolic representation of intensity changes *MIT A.I. Lab Memo No. 325*

Ullman, S. 1975 On visual detection of light sources *MIT A.I. Lab Memo No. 333*

Wallach, Hans 1963 The perception of neutral colors *Scientific American, January 1963*

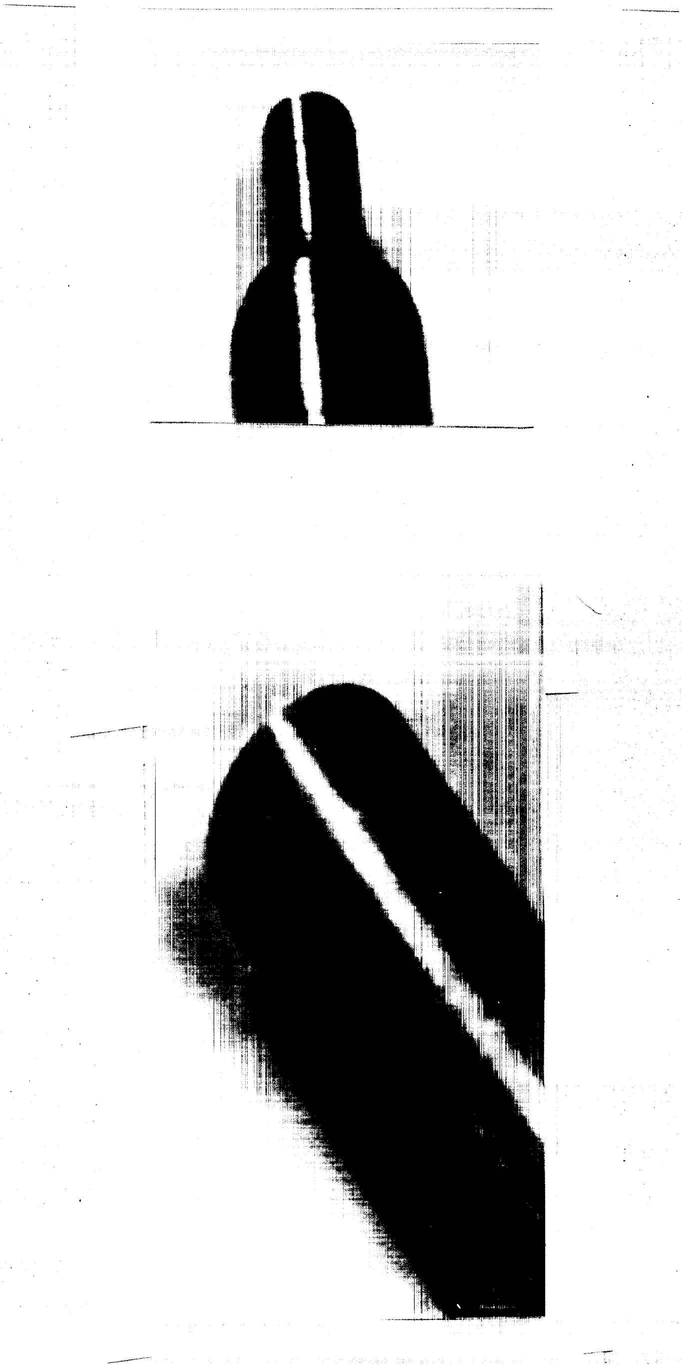

Figure 3

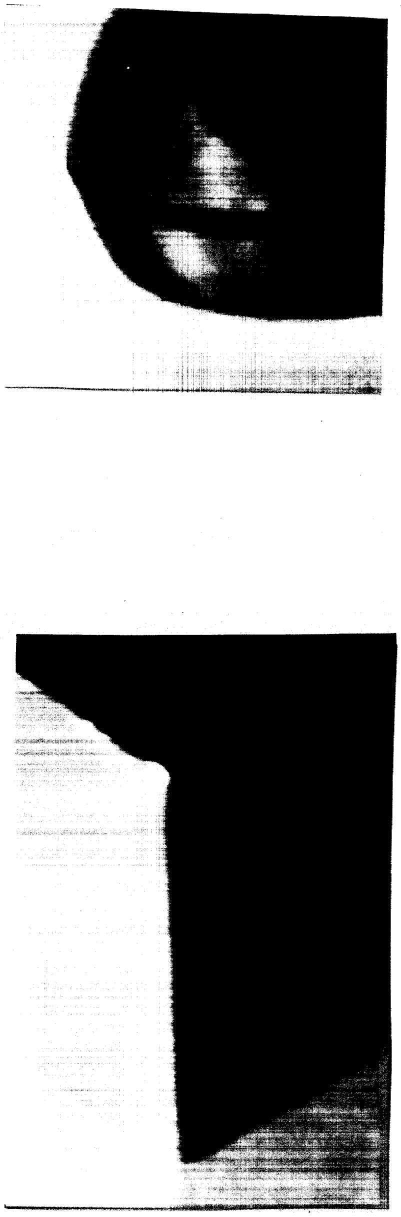

Figure 4

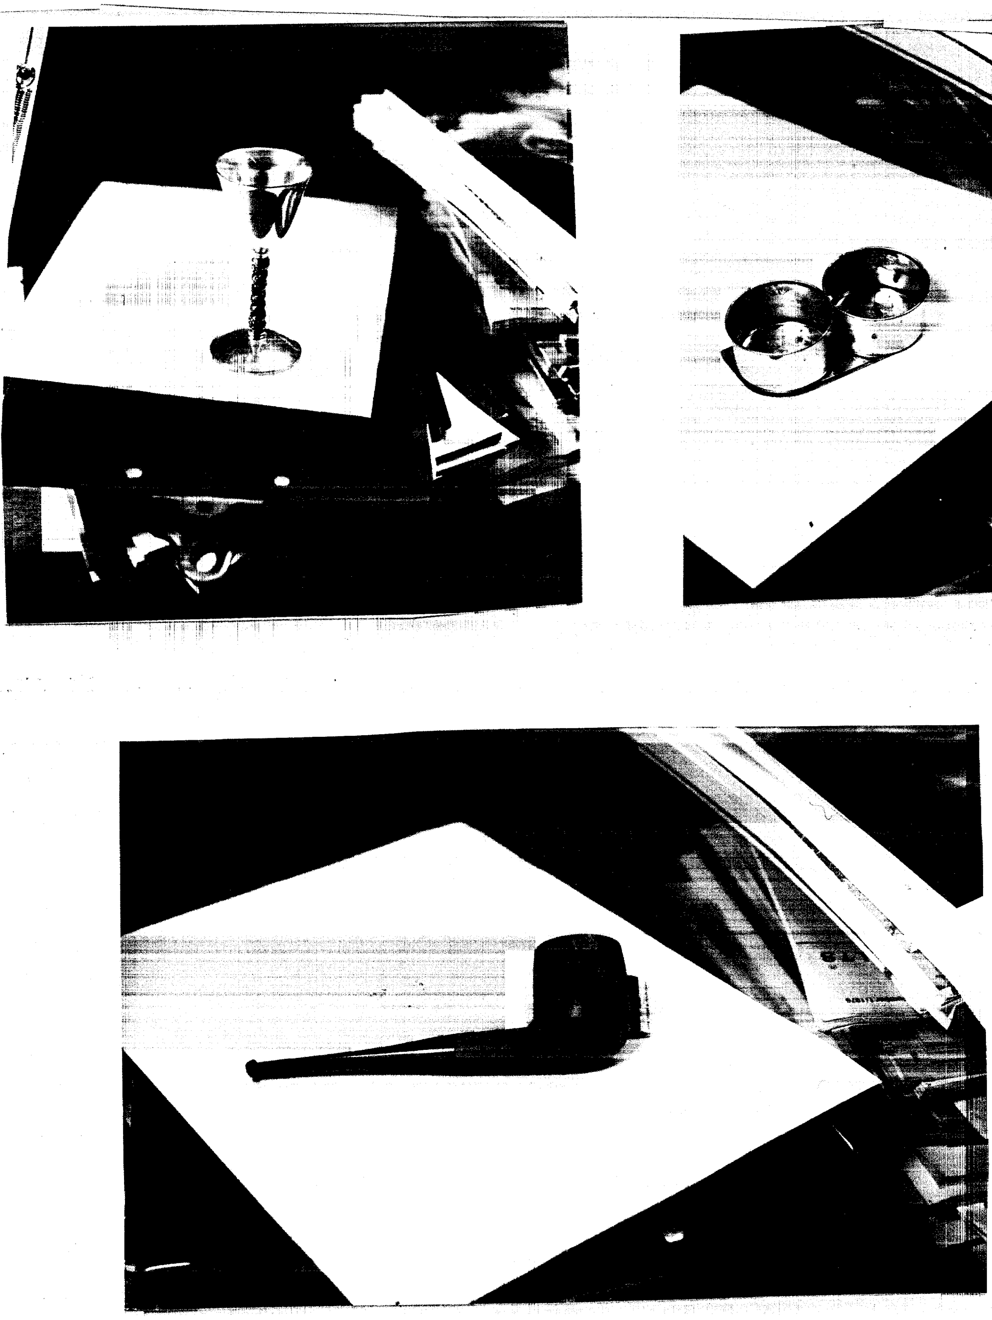

Figure 2

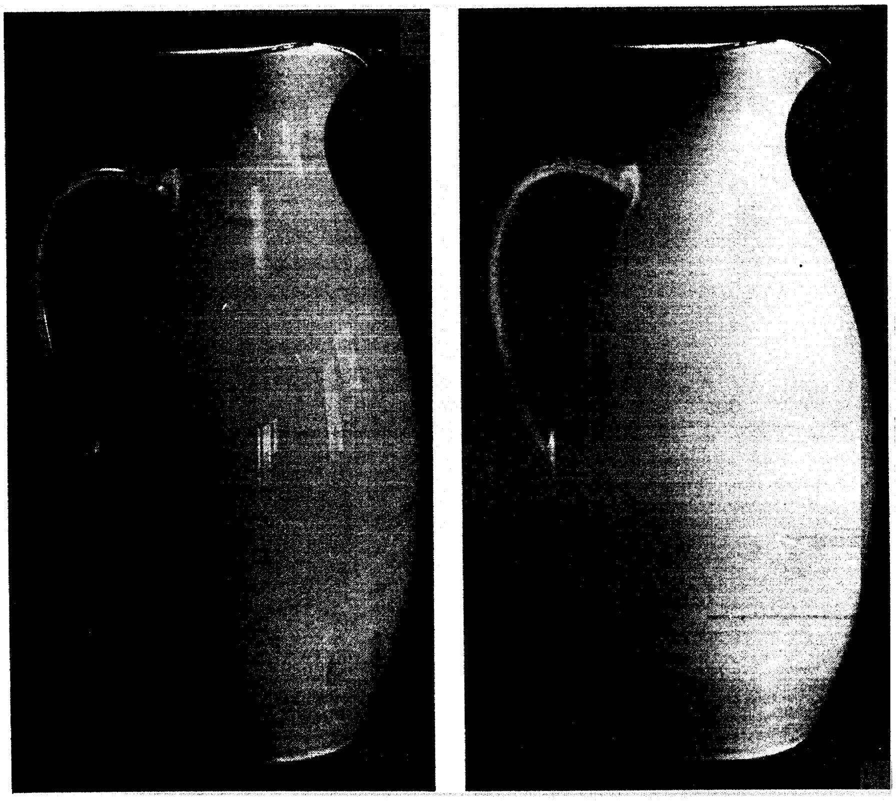

Figure 1

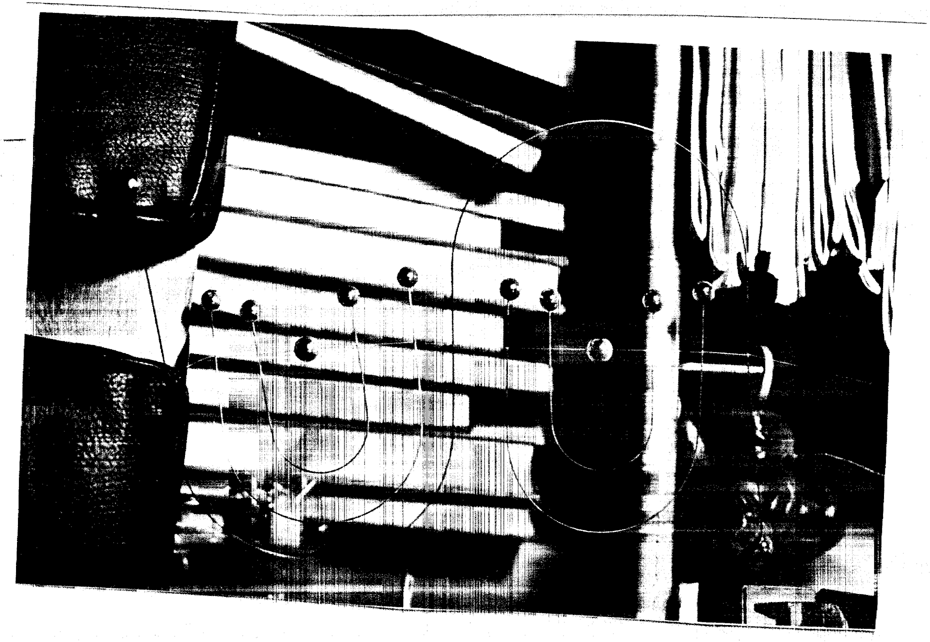

Figure 5

Type 1 Light Source Effects

A light source reflected in a curved specular surface gives rise to a highlight. This is caused by a type 1 light source effect.

I have observed that humans see non-linear intensity gradients as caused by either surface curvature or non-uniform illumination. To recognize highlights, one must be able to distinguish them from object curvature. Let us examine a simple situation for insight.

In figure 6, we have the equations derived for the intensity across a one dimensional profile, assuming both the source and the viewer are very far away from a uniformly curved surface. We see that the object will give rise to a "hill" in intensity, with a seperate peak superimposed on top. The position of the peak and its extent correspond with the places on the surface where the source is specularly reflected. Figure 7 shows the same equation graphed for various n and s. Notice that a lower n yields a wider peak, thus "smearing" the mirror effect of the specular component. By sweeping the surface normal, the curvature reveals to us the specular component.

To determine what parameters are relevant to perceiving highlights, a number of different intensity profiles were used to generate images, either on a tv screen or transparencies for viewing with a light table. I will call the contribution from the specular component the peak, and the rest of the intensity the surround. A set of pictures was produced with the equation derived in figure 6, and the results looked like reasonable cylinders(see figure 8).

Several other analytic approximations for the peak were tried on a constant surround, with different combinations of surround intensity and peak height. Since they all appeared as highlights, the exact shape of the peak cannot matter, so long as it is not a discontinuous step in intensity. The lowest contrast of the peaks was around 0.08, so contrast is not very important, although the highlight looks "better" with higher contrast.

Next a surround corresponding to the matte component of intensity of a cylinder was used, but the width of the peak as a percentage of the width of the surround was varied. All but the broadest (peak width 75% of the total width)appeared as highlights. This suggests that the width is not critical, so long as the surround is visible. A similar set of peak widths was then superimposed over a constant surround. On peaks that were wider

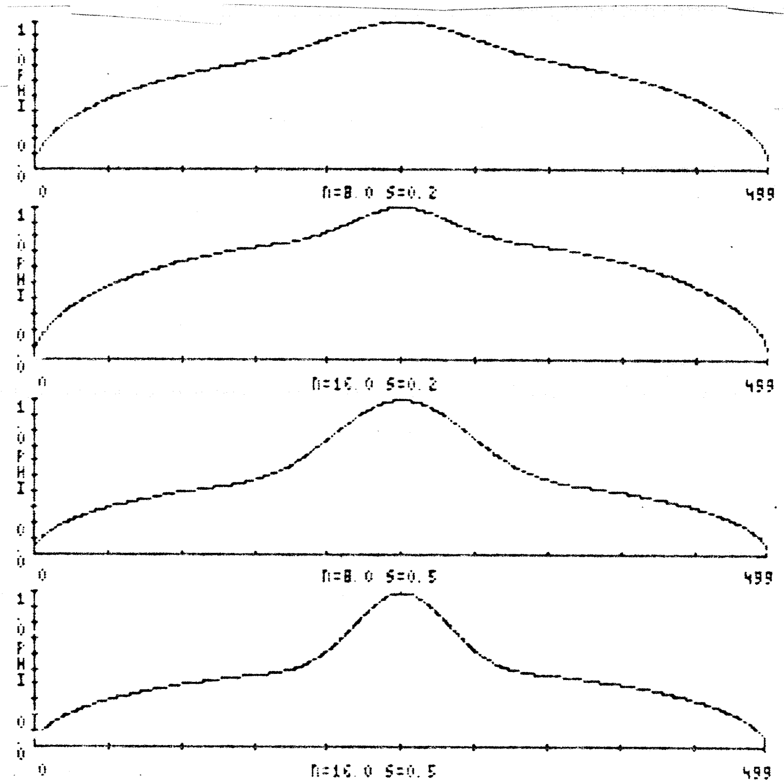
1.0
FHI
0.0
0
N=8.0 S=0.2
499
1.0
FHI
0.0
0
N=16.0 S=0.2
499
1.0
FHI
0.0
0
N=8.0 S=0.5
499
1.0
FHI
0.0
0
N=16.0 S=0.5
499

Figure 7

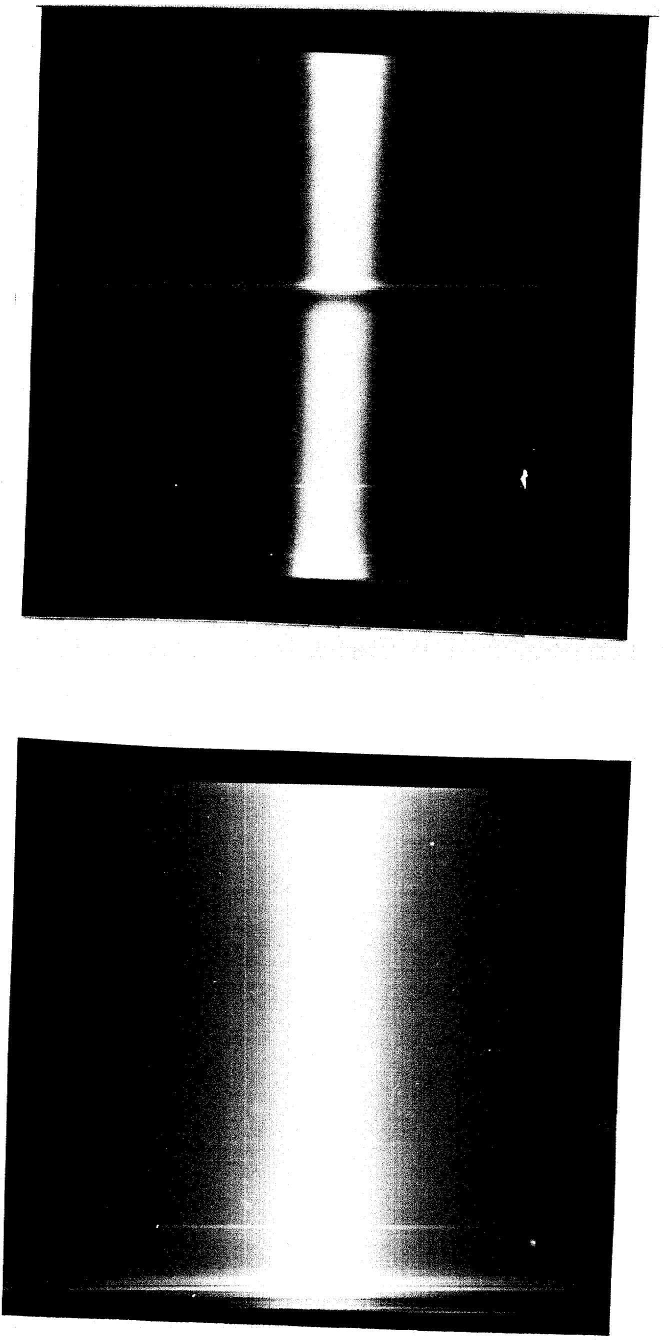

Figure 8

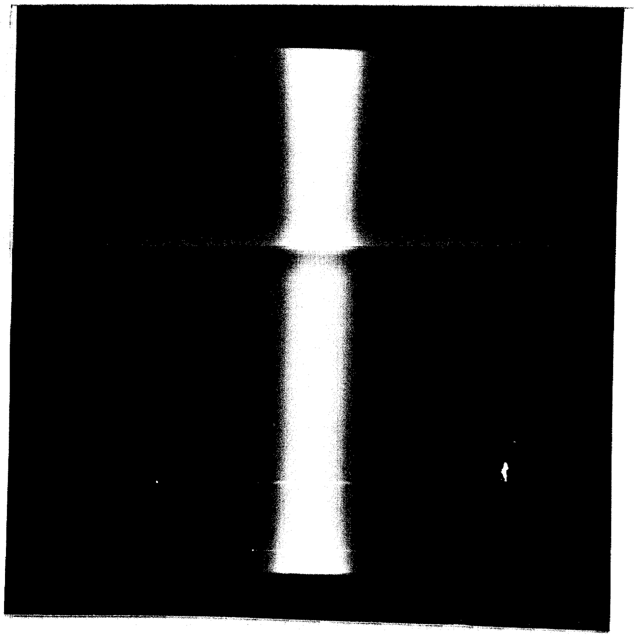

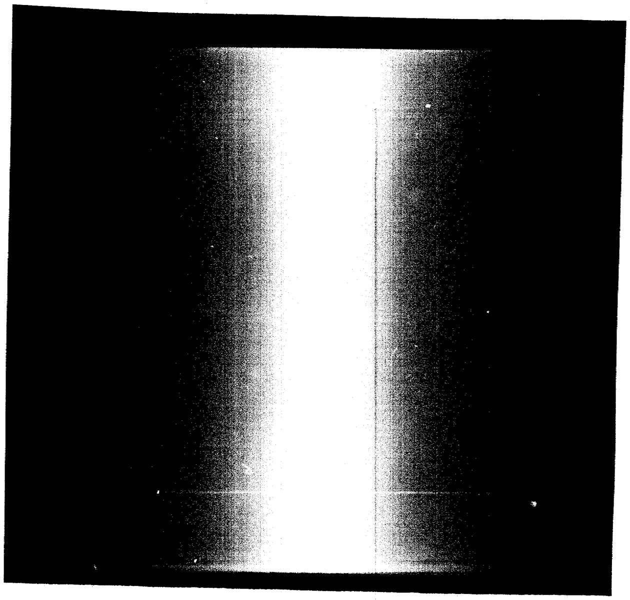

Figure 8

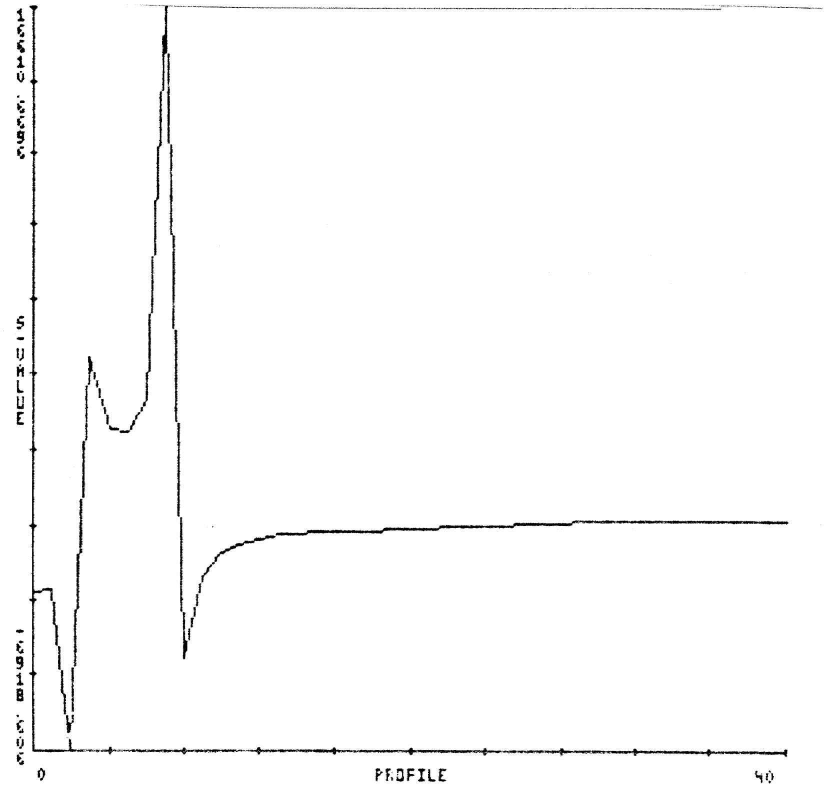

Figure 10(a)

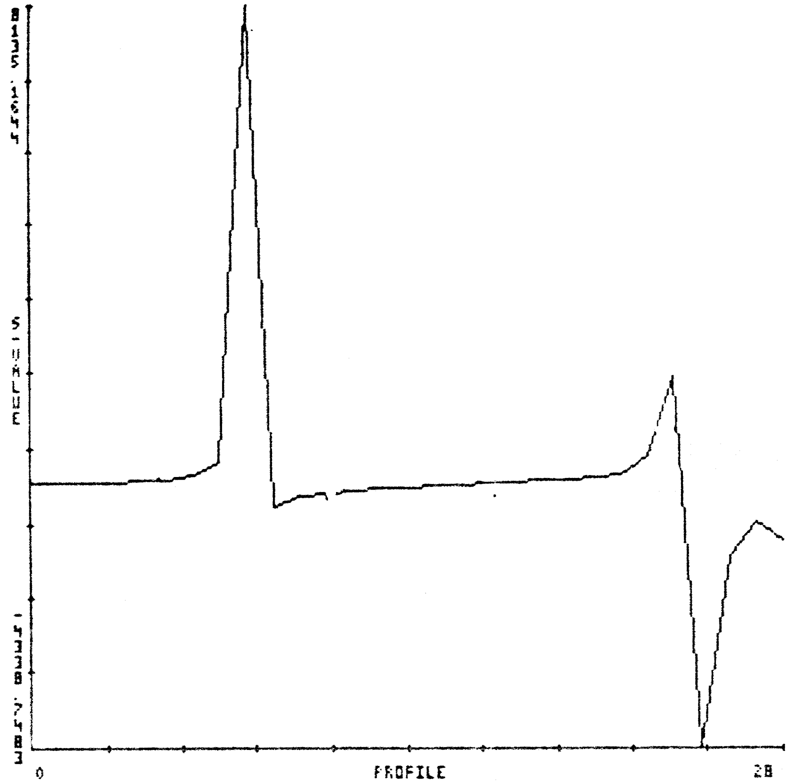

Figure 10(b)