

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
PROJECT MAC

Artificial Intelligence
Memo No. 183

July 1970

ON BOUNDARY DETECTION

Annette Herskovits and Thomas O. Binford

Work reported herein was supported by the Warren McCulloch Laboratory, an M.I.T. research program sponsored by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract number N00014-70-A-0362-0002.

Reproduction of this document, in whole or in part, is permitted for any purpose of the United States Government.

ABSTRACT

A description is given of how edges of prismatic objects appear through a television camera serving as visual input to a computer. Two types of edge-finding predicates are proposed and compared, one linear in intensity, the other non-linear. A statistical analysis of both is carried out, assuming input data distorted by a Gaussian noise. Both predicates have been implemented as edge-verifying procedures, i.e. procedures aiming at high sensitivity and limited to looking for edges when approximate location and direction are given. Both procedures have been tried on actual scenes. Of the two procedures the non-linear one emerged as a satisfactory solution to line-verification because it performs well in spite of surface irregularities.

TABLE OF CONTENTS

TITLE	PAGE
Introduction	5
History of the problem	7
Preliminary studies: What the world looks like through our input devices	10
Some general considerations on detection procedures	22
Linear detection procedures	26
A non-linear detection procedure	35
Conclusion	55
Appendix I. Gradient at image for constant intensity at object	57
Appendix II Computation of detection characteristic curves	61
Bibliography	63

INTRODUCTION

The task of computer vision is defined as follows:

A television camera, the vidisector, scans a visual scene consisting of three-dimensional objects. The camera is connected to a computer, and an array of numbers representing the light intensity distribution is entered into the memory of the machine. Our goal is, using this information, to separate and identify the objects in the scene.

That visual perception can be accomplished by reducing the originally multidimensional information to the set of discontinuity lines, is amply confirmed by our ease at recognizing objects in line-drawings. There are many soft-curved objects in our world; but even then, the outline together with a few lines along the maximal curvatures, constitute a good representation of the object, for recognition purposes. It is thus natural that computer vision should start with line detection.

The difficulty of boundary detection lies in part in the nature of our detector, which provides with a set of set of digitized, noisy light intensities at discrete locations.

But the difficulty lies above all in discriminating surface structure from actual edges. In addition we must proceed in a reasonable amount of computing time.

The work in this thesis consists of three main parts:

- 1) an investigation of the nature of the world as seen through our input devices.
- 2) the definition of a boundary detection procedure proficient in the following situation: presence of a Gaussian detector noise and surfaces of uniform texture except for the presence of rare defects of small diameter.
- 3) Finally the detection method has been concretely worked out in a "line verifying program"; i.e. we assume that we know approximately the location and direction of the edge and we aim at a very high sensitivity, the price of which is reflected in a computing time too large for general purpose use. This defines a partial problem of line detection, justified by the overall conception of the vision system at Project MAC; i.e. our program is to be used only locally, in critical cases of "elusive" edges, under call of a higher level procedure, which looks for evidence that in some area, a less sensitive but faster procedure has missed a line.

HISTORY OF THE PROBLEM

In the beginnings of computer vision, the following local predicate was used:

A point is likely to belong to a boundary if the following condition is fulfilled: any one or both of the differences between the light intensity at it and at its horizontal and vertical neighbors are greater than a given threshold.

(see Roberts, 1963; his predicate is a variant of the above.)

However this simple idea soon proved itself insufficient. It is successful in the very limited case of surfaces of excellent homogeneity and with a high contrast between the regions limiting the edge. Otherwise any choice of value for the threshold will either give as boundary points any surface irregularity, or report as uniform areas where an edge lies.

'Coincidences Predicates' (T. Binford) constitute a fair improvement. Instead of a single large intensity difference one requires the occurrence of a pair of them, in parallel; i.e. if $I_{i,j}$ denotes the light intensity at the lattice point (i,j) , the differences

$(I_{i,j} - I_{i+1,j})$ and $(I_{i,j+1} - I_{i+1,j+1})$ must be both large and of the same sign (or alternatively $(I_{i,j} - I_{i,j+1})$ and $(I_{i+1,j} - I_{i+1,j+1})$). The actual predicate is a somewhat more complex elaboration, and improvement of the same principle.

This criterion will clearly be harsher on random perturbations than the preceding one. But true boundary points will pass through because an actual line 'extends along its own direction'.

With the additional noise filtration gained through the kind of line-fitting method which attempts to find an accumulation of boundary points along a line, this constitutes a reasonably good technique for general purpose use; it is economical in time and discriminates somewhat against surface defects. But it is limited in sensitivity (it uses for its decision the evidence of a small set of points, thus preventing the statistical law of large numbers to play very much in our favor by making us converge toward the underlying noiseless distribution). At least on some small subparts of the image, we need to do better.

In 'An operator which locates edges in digitized pictures' (Hueckel. 1969), an 'edge operator' is proposed, which accepts all brightness values within a small disc-shaped subarea of the picture; it then applies to them a low-pass filter which singles out 8 low-order Fourier components (wavelengths of order the diameter of the disc); some of the white noise, whose spatial frequency is higher is thus eliminated, but also some useful information. It will then attempt to fit to the transformed distribution an ideal noiseless brightness step. An elegant mathematical solution yields the location and amplitude of the best fitting step. If the amplitude is large enough, it is inferred that an edge effectively passes over the disc.

The operator appears to be reasonably fast and to have a fine resolution; although the paper does not include any discussion of the sensitivity, our evaluation of its lower limit showed that it was not sensitive enough for our needs. It does not provide for edges yielding a light intensity surface other than a step separating two horizontal planes; we will see that often this is not sufficient.

Griffith (P.H.D. thesis: in preparation) is working toward a "Theory of the optimal use of intensity information in the detection of lines in the visual field". Optimality is defined under the following set of conditions:

(i) The 'scanning pattern', namely the set of points at which brightness measurements are taken is fixed; i.e. no additional information about the picture can be obtained in the course of the procedure.

(ii) The noise is Gaussian. Surface defects are not taken in account.

(iii) The procedure is uniform with respect to position and orientation.

(iv) The a-priori probabilities of the presence of lines are known.

We do a somewhat more realistic treatment but a theoretical approach with these simplifying assumptions is of great interest.

PRELIMINARY STUDIES

WHAT THE WORLD LOOKS LIKE THROUGH OUR INPUT DEVICES

1 The light intensity units

The interface between vidisector and computer yields numbers which are logarithms of the light intensity. Calibration of the measure apparatus has given the following result:

Adding 96 to the logarithmic measure corresponds to division by a factor of 2 of the actual light intensity.

We will in the rest of this thesis call these units vidlog units or v.u.

2 The noise

a) Its sources can be traced to 4 phenomena.

i) Photon noise

Measuring the light intensity is done by measuring the time taken for a fixed number N of photons to hit a photosensitive device. The light intensity, I , is inversely proportional to that time. The number of photons arriving during successive equal intervals of time t , show a Poisson distribution, which for large numbers, can be approximated by a normal distribution of mean, say N , and standard deviation $\delta N = \sqrt{N}$. With a good approximation, if for a given intensity N photons arrive in a time t on average, then if δt is the deviation on time:

$$\frac{\delta t}{t} = \frac{\delta N}{N} = \frac{1}{\sqrt{N}}$$

As N is kept constant through all our measurements, $\frac{\delta t}{t}$ remains constant whatever the intensity and so is $\frac{\delta I}{I} = \frac{\delta t}{t}$. Thus, if $L = \text{Log } I$, δL is a constant; i.e. expressed in logarithmic units, the photon noise remains the same for any intensity level.

ii) Noise in the electronics

Very little is known about it. In any case it cannot be separated from the above source of noise.

iii) Ripple of the lamp

The sun gun lamp which we used had a large amount of 120 cycle noise, of which we became aware only recently, so that all the experiments described in the thesis were carried out without any protection against it. We now use photocathode reference for the measurements i.e. normalize by the average photocathode illumination. Small lamps are particularly noisy because they have small filaments. We are about to try D.C. supplies for the lamps.

iv) Inhomogeneity of the photocathode

The sensitivity of the measurements by the vidisector varies from one point of its field to another.

b) A measurement of the effective standard deviation

We want a measure of the noise which accounts for all the above factors. As we will see, our detection functions always involve a comparison of nearby values of the intensity. Thus a first approximation to the thresholds to use for testing the presence of edges, is based on a value of the standard deviation accounting for the spatial inhomogeneity of the sensitivity of the vidisector.

After having defocussed in order to have a surface with good homogeneity, we scanned a sheet of white paper, first horizontally

and then vertically, and at each point we took the following measurement: the difference between the average of 15 brightness values on a segment centered around the point and the value at the point.

The sum of the squares of these was averaged over the whole screen and the square root of the result is the required standard deviation. We found a value

$$\underline{\int L = 1.2 \text{ v.u.}}$$

As a guide for fixing thresholds, we worked according to the approximation that the overall noise was Gaussian. All the errors can be expected to be proportional to the intensity and therefore the logarithmic error is constant.

3 Surface structure of the objects.

We have not dealt with texture but only with sparse surface irregularities of small diameter i.e. covering around 1% of the surfaces. In fact we were successful with surfaces considerably worse than we were expecting to use; several scenes (see figures 27,28 and 29 .) contain objects which were chosen to be very bad; they showed a linear wood texture and a pebbly texture of end-grain.

4 Light intensity distribution

Intensity profiles taken along straight lines on planes showed much larger gradients than we expected. The phenomenon was investigated, both theoretically and experimentally.

The intensity varies essentially because the light source induces a non-uniform intensity on the object. As Appendix I shows, with a logarithmic measure of the intensity, the gradient of the reflected light at the object and the gradient at the image are the same, to a second order approximation. Thus we will now consider the

(i) Variation of the light intensity on a plane in presence of a point light source.




figure 1

$$I_M = \frac{E_0}{4\pi} \frac{\cos \theta}{r^2}$$

$$\frac{\partial I}{\partial x} = \frac{E_0}{4\pi} \frac{\cos \theta}{r^2} \left[-\frac{2}{r} \frac{\partial r}{\partial x} - \operatorname{tg} \theta \frac{\partial \theta}{\partial x} \right]$$

LN is normal to the plane which intersects the plane of figure according to NMx. We want to compute $\frac{\partial I_M}{\partial x}$, I_M being the light intensity at M. (It is along the direction Nx that the gradient is maximum).

E_0 being the energy emitted the whole 4π solid angle.

$$\text{As } dx = r \frac{d\theta}{\cos \theta} \text{ and } \frac{\partial r}{\partial x} = \sin \theta \quad \frac{\partial I}{\partial x} = \frac{E_0}{4\pi} \frac{\cos \theta}{r^2} \frac{-3 \sin \theta}{r}$$

$$dI = \frac{\partial I}{\partial x} dx = -\frac{3}{2} \frac{\sin 2\theta}{d} dx$$

The gradient shows a maximum for an angle of 45° between the light rays and the plane of the object.

The experimental curves currently show gradients 4 times as large as the gradient computed according to the above formula, for the same relative position of the source and the object plane. Thus some other phenomena must interfere.

The two main causes seem to be that the lamps we use are far from being point light sources, and that the reflectivity of the surfaces varies with the angle of incidence of the light rays.

(ii) Experimental exploration of the light intensity on a plane surface

We used sheets of white and uniform paper and took intensity cross-sections in various locations and directions. We found the following results:

1) The most important factor is, the angular dependence of the lamp radiation. Instead of a continuous decrease of the intensity from N away along NM (figure 2), we get a maximum at C where the axis of the lamp intersects the plane of the sheet.




figure 2

2) Extreme values of the gradient are at points where the $\frac{\cos \theta}{r^2}$ effect is maximum and contribute to the gradient with the same sign as the angular dependence of the lamp. They are of the order of:

$$\underline{10 \text{ vidlog-units / cm}}$$

Defocussing

Because of defocussing the intensity profiles across a perfect edge (figure 3a) will look as in 3b.

b is the convolution of a step and a circle (see figure 5 too).

We will evaluate the extent of ab for usual deviations from perfect

focus.




figure 3





figure 4

Let us suppose an edge normal to the sheet of paper in O. Light rays which enter the aperture and issue from points on the circle of diameter Oβ (circle of confusion) converge toward the image of B (respectively Oα and A).

DEFOCUSED EDGE



Light rays apparently issued from any point P in the plane of best focus, converge after crossing the optical system, to a point image of P. If P is between A and B, the light rays will come from both sides of the edge. If P is outside AB, its image will be the point of convergence of light rays coming from the same side of the edge.

Thus the width of transition is the width of the image of $AB = \delta$ through the system. We have:

$$\delta = a \frac{d}{s-d} = a \eta$$

a being the diameter of the aperture and η the relative error in position with respect to perfect focus.

If s' is the distance aperture-image for a distance s aperture-object;

$$\frac{1}{s} + \frac{1}{s'} = \frac{1}{f} \quad f: \text{focal distance}$$

and the corresponding ratio image-size/object-size is $\frac{s'}{s}$. Thus δ' image of δ is:

$$\delta' = \frac{s'}{s} a \eta$$

As s is relatively large, the image of δ is practically in the focal plane and $f = s'$. Thus:

$$\delta' \approx \frac{f}{s} a \eta$$

We can evaluate an upper limit of δ' . The vidisector characteristics are:

$$f = 10''$$

$$a = \text{from } 2'' \text{ down.}$$

We will focus at worst to 2'' at 3 feet. Thus:

$$\delta' = \left(\frac{10}{36} \cdot 2 \cdot \frac{2}{34} \right)'' = 8 \text{ mm}$$

The diameter of the photocathode is 3". The ratio of the maximal width of a line to the photocathode diameter is:

$$\rho = \frac{1}{90}$$

The resolution of the photocathode being about 1000 points, δ' will be at most 11 grid-points.

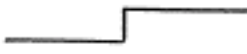
With B. Horn focussing program (B. Horn 1968), focus can be done to 1cm at 3 feet. Thus ρ becomes:


$$\rho = \frac{1}{28}$$


and $\delta' = 3 \text{ mm}$ or 4 grid-points.

6 Results. Characteristic aspects of edges.

We found mainly 3 different shapes of cross-sections normal to edges:

(i) Steps (figure 6a). 

(ii) Roofs (figure 6b). On each side of the edge the profiles have different constant slopes. Usually a small effect. 

(iii) Edge-effects (figure 7). On the edge is a small bump of the intensity surface. 

These three characters may be combined as in figure 8.




figure 6a




figure 6b

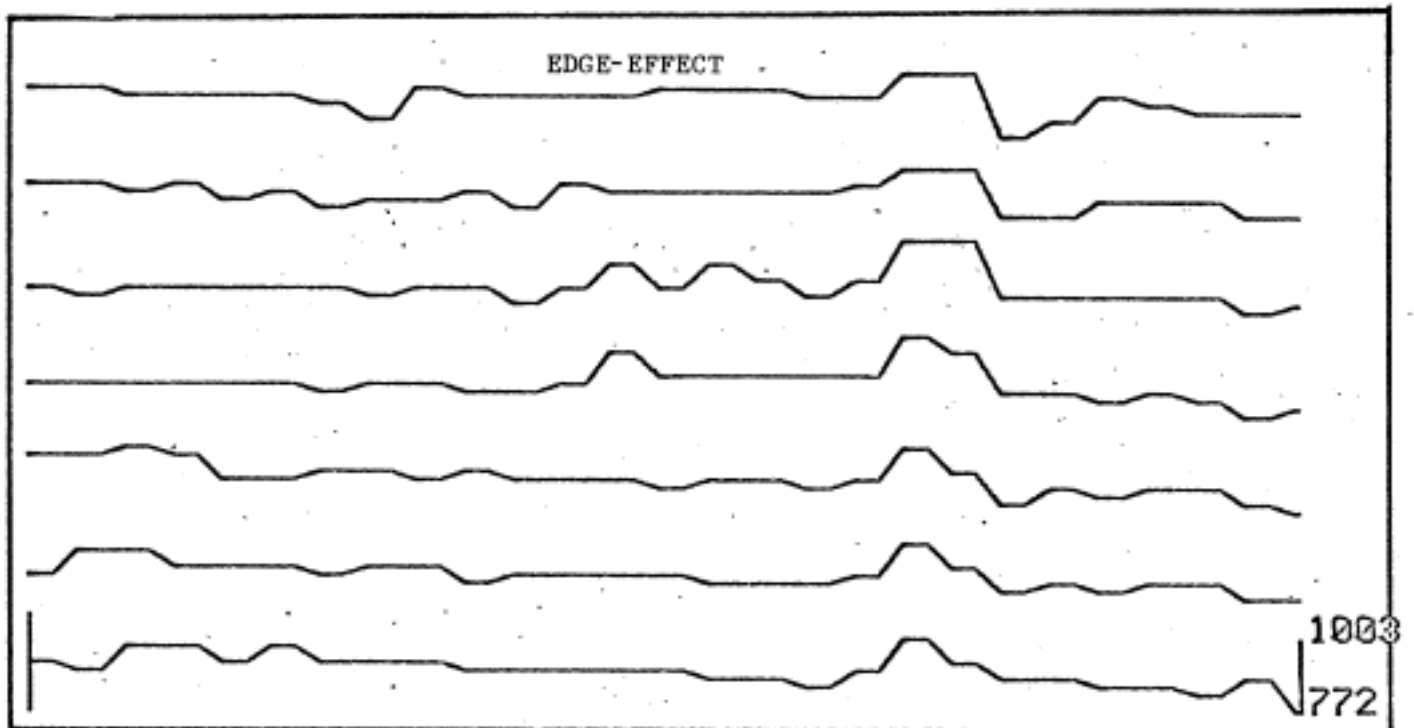


figure 7

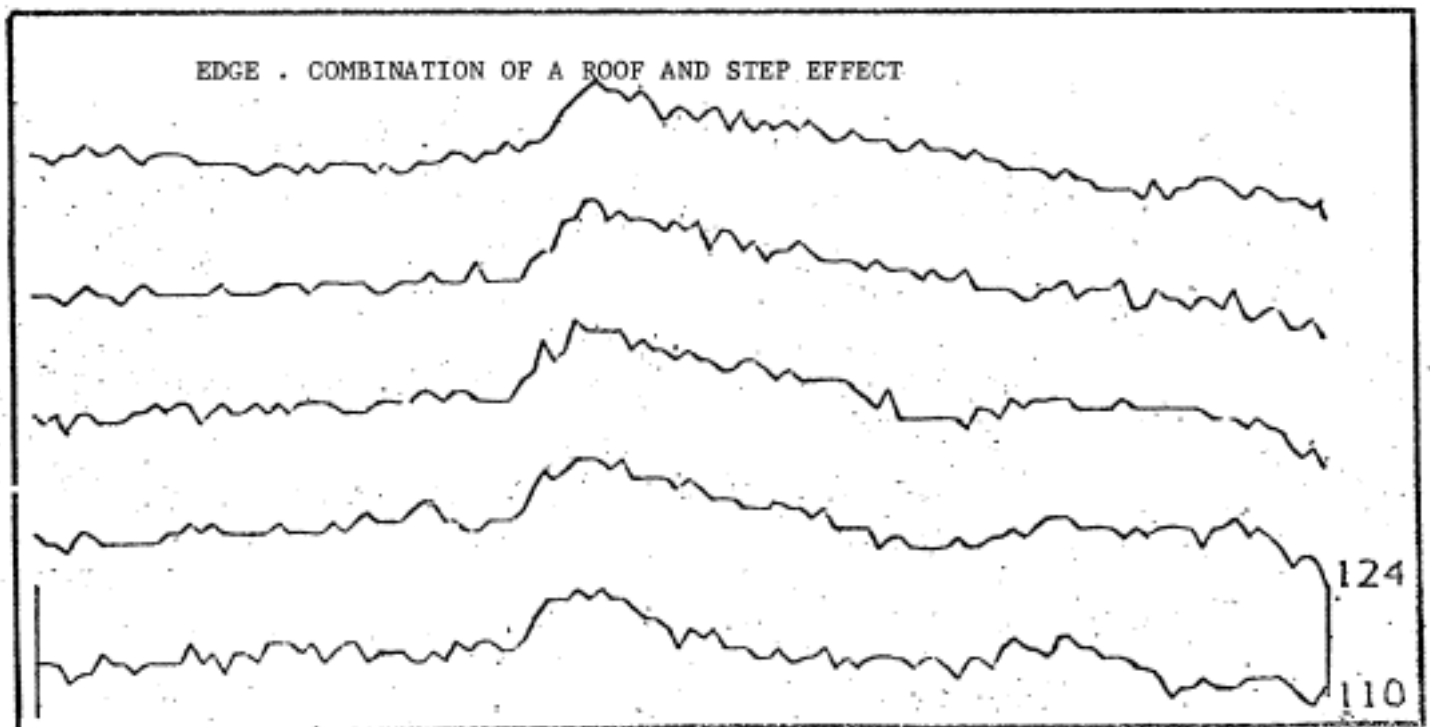


figure 8

SOME GENERAL CONSIDERATIONS ON DETECTION PROCEDURES

The number of possible schemes for boundary detection is very large. However they all have in common that noise discrimination is obtained, in part, by integrating some function of the light intensities over some subset of the plane, (a one-dimensional (segment) or two-dimensional (rectangle or circle) neighborhood).

Around each point of the plane is centered such a neighborhood \mathcal{A} ; the intensity at any point P of \mathcal{A} can be written as the sum of a pure signal S(P) and a noise N(P). If we then apply a transform T to the intensity surface, noise discrimination will occur under the assumption that:

$$\int_{\mathcal{A}} [T S(P) + N(P)] \quad \text{tends toward} \quad \int_{\mathcal{A}} T [S(P)]$$

when the area (or length) of \mathcal{A} tends to infinity; i.e. the integral of the transform of the combined signal tends toward the integral of the transform of the pure signal.

This is true for the transformations we have considered under the conditions that the noise at every point is Gaussian, has mean zero and a constant standard deviation; and provided that noise and signal are uncorrelated.

The transform must be such that the integral reflects the presence of a discontinuity of the pure signal within the domain of integration \mathcal{A} , a discontinuity characteristic of rectilinear edges.

The transform can be linear or non-linear in intensity. In the rest of the thesis we will examine two methods; one linear and one strongly non-linear. We adopted the latter. It uses cutoffs on the result of processing of the intensity surface by some differential operator; consequently the original brightness surface is transformed into a 3-level surface, a drastic reduction of the original information. As a result, it cannot be optimally sensitive and linear predicates behave better in this respect. But it has one important advantage, namely:

Contribution of surface defects to the integral over a basic neighborhood is according to their extent not their amplitude. This gives to the method a good discrimination against structural noise.

In linear methods, surface defects bias the computation proportionally to their extent and amplitude, which is often very large. This is particularly unfit for a line-verifier system where high sensitivity must be obtained by integration over a relatively large neighborhood; thus the number of points in the plane whose surrounding neighborhood includes the defect is large and the error is propagated far away from the defect itself. We have kept here the analysis of a linear method, because it sheds light on the properties of real edges and gives some idea of the theoretical sensitivity limits.

Up to now we have been speaking of the local part of the procedure. Detection methods will then differ by:

- (1) What the local stage returns for information.

(ii) How this local information is put together to reach a global answer to the detection question, i.e. the global set of lines.

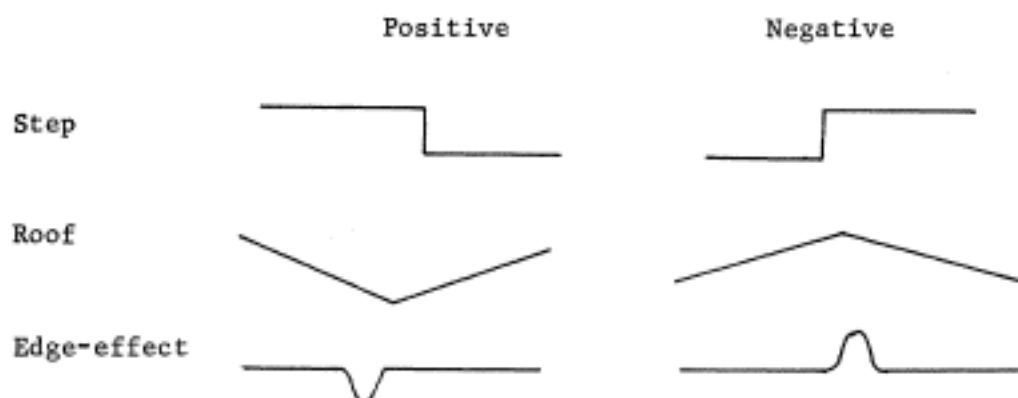
The local processing may return likely "edge-points". Then, at the next stage, we will look for an accumulation of edge-points along some direction.

Or it may return "edge-vectors", i.e. vectors tangent to the edges. Then the vectors will string according to collinearity.

Or the information carried along to the global process may be still more detailed; for example the edge-vectors may be accompanied by information about the type and amplitude of the discontinuity they reveal. We will then string together only vectors of the same type and amplitude.

We have listed these methods in an order corresponding to an increased noise filtration by the local-to-global process. The amount of noise which can then be allowed to pass through the first noise filtering, (the local predicate) is correspondingly greater, i.e. fainter edges can pass through and sensitivity improves. But the necessary amount of work increases.

The method we have opted for in the line-verifier returns from the local stage with an edge-vector accompanied by the type and the "sign" of the discontinuity, the sign being defined according to the following:



It did not seem meaningful here to introduce several levels of amplitude because we are looking only for very small signals anyway.

With the following chapter we will start examining some linear transforms; we will first see how they act on one dimension, i.e. along a profile normal to the edge. Then we will briefly describe a global method which we tried and how it failed. The last chapter will be devoted to the non-linear method.

LINEAR DETECTION PROCEDURES

1 One-dimensional linear transforms

1.1 Function F1. First differences integrated

The most usual intensity profile across an edge being roughly a step, let us consider the following function, F1:

At each point of the profile, we compute the difference between the average intensity after and before the point; both averages are taken over a given length 2Δ (2Δ points).

Let $L(X)$ be the intensity profile;

$$F1(X) = \frac{1}{2\Delta} \sum_{i=1}^{2\Delta} [L(X+i) - L(X-i)]$$

This amounts to taking the first differences with an interval 2Δ :

$$D1(X) = L(X+\Delta) - L(X-\Delta)$$

and averaging them over 2Δ ;

$$\begin{aligned} \frac{1}{2\Delta} \sum_{i=-\Delta}^{+\Delta} D1(X+i) &= \frac{1}{2\Delta} \sum_{i=-\Delta}^{+\Delta} [L(X+\Delta+i) - L(X-\Delta-i)] \\ &= F1(X) \end{aligned}$$

Figure 9 shows L , $D1$ and $F1$ for an ideal noiseless step.




figure 9

We will now discuss the sensitivity of detection on a single intensity profile.

(i) Sensitivity limit imposed by random noise

Let us consider a flat intensity distribution with a random Gaussian noise superimposed, of standard deviation δL .

F_1 , linear function of 4Δ independent variables is normally distributed with average 0 and standard deviation:

$$\delta F_1 = \sqrt{4\Delta} \frac{\delta L}{2\Delta} = \frac{\delta L}{\sqrt{\Delta}}$$

A step h which we can detect must be such that:

the value of F_1 on an edge with step $h \gg \delta F_1$, the noise.

Thus:

$$(1) \quad h \gg \frac{\delta L}{\sqrt{\Delta}}$$

(ii) Sensitivity imposed by systematic effectsa) Constant slow gradient

In order to discriminate an edge from a constant slow slope, the value of F_1 on an edge must be larger than its value for the maximal slow slopes encountered.

Let α be the slope. Then:

$$F_1 = 2\Delta\alpha$$

A step h can be safely detected only if:

$$(2) \quad h \gg 2\Delta\alpha$$



figure 10

b) Consequent optimal value of Δ

The maximal values of α encountered are of the order of 10 v.u./cm.

To improve the sensitivity Δ should be increased according to (1)

and decreased according to (2). Thus the optimal value occurs for:

$$2\Delta\alpha = \frac{\delta L}{\sqrt{\Delta}}$$

$$\Delta^{3/2} = \frac{\delta L}{2\alpha}$$

At a distance of 3 feet, 1 cm is imaged into 40 grid-points, and with a noise $\delta L = 1.2$ v.u. ;

$$2\Delta_{\text{optimal}} = 4 \text{ points}$$

With the normal distribution assumption, if we then put a cutoff on $F1$ at 2.6 times $\delta F1$, i.e. if we predicate:

$$\text{there is an edge} \Leftrightarrow F1 > 2.6 \delta F1$$

we will reduce the probability of error in the absence of an edge to 1%. Detection will be successful at 50% for an edge h yielding a value of $F1$ equal to the cutoff, i.e.:

$$h = 2.6 \frac{\delta L}{\sqrt{\Delta}}$$

$$\underline{h = 1.6 \text{ v.u.}}$$

c) Defocussing

Let us approximate the smeared intensity profile by straight lines as shown in figure 11. Let δ be the width of the edge.

As we have seen δ may currently be of

the order of 8 points. For $2\Delta = 4$,

$F1(X)$ will have a peak value of $h/2$

instead of h . If we use the same

cutoff on $F1$ as above, 50% detection

will occur only for a step:

$$\underline{h = 3.2 \text{ v.u.}}$$




figure 11

The sensitivity is clearly unsatisfactory, and whatever the global decision procedure, it cannot be conspicuously improved because of the presence of constant gradients on planes.

One is therefore lead to think of criteria such as: there must be a local maximum of large amplitude of the function F1 on the considered profile. This amounts to considering differential operators of higher order, using second and third differences; these are the transforms which we will now consider.

1.2 Function F3

Let $L(X)$ be an intensity profile and:

$$D2(X) = -2L(X) + L(X+\Delta) + L(X-\Delta)$$

$$F3(X) = \frac{1}{\Delta} \left[\int_{i=1}^{\Delta} D2(X+i) - \int_{i=1}^{\Delta} D2(X-i) \right]$$

$D2$ is the second differences with an interval Δ . $F3$ is the difference between the mean of $D2(X)$ over Δ points after X and over Δ points before X .

Figure 12 shows $D2(X)$ and $F3(X)$ for an idealized step-type edge.




figure 12

Our reason for introducing $F3$ and not simply taking a function which averages $D2$ over an interval is that we prefer to deal with a signal which is large rather than zero on the edge.

(i) Sensitivity limit imposed by random noise.

With random noise, D2 has mean zero and variance:

$$(\delta D2)^2 = (2\delta L)^2 + \delta L^2 + \delta L^2$$

$$\delta D2 = \sqrt{6} \delta L$$

Thus:

$$\delta F3 = \frac{\sqrt{2\Delta} \times \sqrt{6} \delta L}{\Delta} = 2 \sqrt{\frac{3}{\Delta}} \delta L$$

F3 on the edge equals 2h.

Therefore a reliably detectable step h must verify:

$$2h \gg 2 \sqrt{\frac{3}{\Delta}} \delta L$$

$$h \gg \sqrt{\frac{3}{\Delta}} \delta L$$

(ii) Systematic effects

a) F3 is obviously zero for constant gradients. F3 is zero too for any intensity curve which may be approximated by a parabola. So that the angular dependance of the light source which is approximately parabolic, gives no contribution.

b) The effect of defocussing depends on the relative values of the width δ of the line and of Δ . The limitation on Δ being much less tight because we don't suffer from constant gradients, we can assume that $2\Delta \gg \delta$ and that defocussing does not affect us.

(iii) Choice of Δ

The greater Δ the better the sensitivity. Our limitation will then

come from:

- a) The physical size of the object
- b) The light intensity distribution on a plane has to be

possibly approximated by a second degree polynomial.

Now to compute one value of F3 at a point X, we need 4Δ points, 2Δ on each side of X. Thus 4Δ must not exceed, say 1/5 of the side of a current cube. The side of a current cube being about 1/5 of the scene we must have:

$$4\Delta < \frac{1000}{5 \cdot 5} = 40 \text{ points}$$

$$\underline{\Delta_{\text{optimal}} = 10 \text{ points}}$$

Then, if we put a cutoff on F3 corresponding to $2.6\sqrt{3}F3$, we guarantee an error $< 1\%$ on a flat distribution, and the corresponding step h, for which detection will occur with a probability of 50% is:

$$h = 2.6 \sqrt{\frac{3}{10}} \delta L$$

$$\underline{h = 0.53 \text{ v.u.}}$$

1.3 Function F2

When the edge does not appear as a step, but as an edge-effect or a roof, the function F3 is not effective. We must instead use something like the integral of the second differences, which expresses the radius of curvature. Let:

$$D2(X) = -2L(X) + L(X+\Delta) + L(X-\Delta)$$




figure 13

$$D2(X) = \frac{1}{2\Delta} \sum_{i=X-\Delta}^{X+\Delta} D2(X+1)$$

F2 is the average of the second differences over 2Δ . Figure shows the resulting curve for a symmetric roof with slopes α and $-\alpha$.

Expectation of F2 for random noise

With pure random noise, F2 has mean zero and standard deviation:

$$\begin{aligned} \delta F2 &= \frac{\sqrt{2\Delta} \delta D2}{2\Delta} = \sqrt{6} \frac{\delta L}{\sqrt{2\Delta}} \\ \delta F2 &= \sqrt{\frac{3}{\Delta}} \delta L \end{aligned}$$

When dealing with an edge-effect, Δ should be of the order of 3 points.

With a roof, something around 10 points is a reasonable value for Δ .

In the case of an edge-effect, the width of the bump is of the order of 3 to 10 points, under current conditions; by a one-dimensional processing of a cross-section, we will discriminate against irregularities less than 3 points wide; therefore we will keep most structure defects, which are not distinguishable from an edge-effect. Our global decision procedure must then rely essentially on the comparison of parallel cross-sections.

2 A simple linear global process

We tried these one-dimensional linear transforms within a completely linear decision scheme.

To be concrete we will assume that we are dealing with an edge of the step type, and that we apply to the intensity cross-sections the function $F3$.

We want to know whether an edge lies between $P1 \pm dx$ and $P2 \pm dx$. We consider 15 equidistant profiles between $P1$ and $P2$ and sum the values of $F3$ along a number




figure 14

of lines joining a point of $P1 \pm dx$ to a point of $P2 \pm dx$.

We look for the largest of those sums and decide in favor of an edge when it is greater than some prefixed threshold.

An evaluation of the sensitivity limit:

The standard deviation on a sum S of 15 independent values of $F3$ is:

$$\delta S = \sqrt{15} \delta F3 = 2.0 \text{ v.u.}$$

A cutoff at $2.6 \delta S$ as we usually do, gives for the step detectable at

$$50\%: \quad 15 \times 2 h = 2.6 \delta S$$

$$h = 0.15 \text{ v.u.}$$

This was programmed and tried. It was very good at finding very faint edges but the rate of errors on faces was very large, due to surface defects.

A NON-LINEAR DETECTION PROCEDURE

As usual, we will describe the method by distinguishing two stages of processing, the local and the global.

1) The local predicate

We make a local decision based not on a single profile, but on a small number of parallel adjacent profiles. In other words, we consider a narrow strip perpendicular to the expected direction of the edge, and the local decision process returns possible edge-segments which cross it almost vertically (for brevity we will always say vertical for the expected direction of the edge although it can be any direction along which we choose to look).

We do this by computing a certain function F of the intensities within a basic neighborhood which is a small elongated rectangle which we slide along the strip. Large local maxima of this function F indicate possible locations of almost vertical edge-segments.

For a step, the function F is defined as follows:




figure 15

we compute the second differences along every horizontal profile in the strip with an interval Δ . Figure 16 recalls the shape of the curves. Thus, if we put a two-sided cutoff on the second differences (and assuming the step high enough), the distribution of cutoff points shows two Δ -points wide adjacent rectangles, one with all values below the negative cutoff, and the other with all values above the positive cutoff (figure 16).

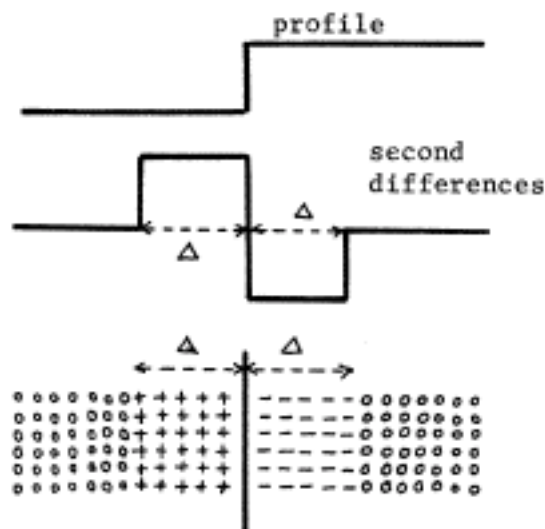



figure 16

The function F_{step} is then computed over a neighborhood 2Δ , and equals the difference between:

$R = (\# \text{'positive' points} - \# \text{'negative' points})$ in the right half of the neighborhood and:

$L = (\# \text{'positive' points} - \# \text{'negative' points})$ in the left half.

This function shows a sharp extremum on the edge (figure 17).



The function F_{step} has the following advantages:

(1) The use of a two-dimensional basis neighborhood increases the possible number of points entering in its computation, thus improving the noise discrimination.




figure 17

(ii) It accounts for the directional property of the edge, i.e. the repetition of some accident in the 'vertical' direction.

(iii) The weight of a surface defect is equal to the number of points it covers, i.e. taxes us minimally.

One difficulty is the "wings" shape of F_{step} ; the relevant extremum is framed by two 'parasite' extrema of amplitude half that at the edge. Thus some precaution must be taken to eliminate the 'satellites'. Our first idea was to compute the function S which is the sum of the two terms involved in computing F_{step} , i.e. S is:

(# positive points - # negative points) in the whole basis neighborhood.

$$S = L + R \quad (\text{notation as in the definition of } F_{\text{step}})$$

S is plotted below F_{step} ; as it is zero on the edge but large on the satellites, it offers a possibility of discrimination.

As we will see, in the case of the line-verifier, this proved an unsatisfactory solution and had to be completed by some measure at the global level.

Remarks: 1) The height of the strip should be big enough to allow the 'directional property' of the edge to play its role; but small enough, so that for the given angular uncertainty of the edge which we allow, the horizontal coordinate of the edge varies negligibly while

crossing the strip (i.e. at most 1 grid-point).

2) In actuality, when computing F_{step} , we separate left and right half of the neighborhood by a narrow vertical corridor, 3 or 4 points wide, to allow for imperfect focus.

3) The second differences are in fact the following function:

$$D2'(X) = L(X) - L(X+\Delta) - L(X+\Delta+1) + L(X+2\Delta+1)$$

i.e. instead of taking twice the center value as in the usual second differences $D2$, we add the values at two nearby points. Then;

$$(\int D2')^2 = 4(\int L)^2$$

$$\int D2' = 2\int L \quad \text{instead of } \sqrt{6}\int L \text{ for } \int D2.$$

$D2'$ is less noisy.

1) Statistical properties of the local function F_{step}

We define the probability of local detection as the probability that the value of F_{step} at a given edge be greater than some prefixed threshold $\theta_{F_{\text{step}}}$. As we will see, the additional condition which we have introduced to eliminate the satellite extrema of F_{step} keeps the actual probability of local detection equal to the one in this definition.

We will compute "detection characteristic curves" which represent the probability of local detection as a function of the signal, considering Gaussian noise only. This approximation will give us a starting point for fixing the parameters.

The parameters which bear on the probabilities of detection are as follows: the size of the basic rectangular neighborhood, the cutoff value for the second differences, and the threshold for the function F_{step} . The computations constitute a simple combinatorial problem (see Appendix II) whose results follow:

1) Figure 18 (next page) shows the detection characteristics for various sizes of the neighborhood, expressed as the number of grid-points it contains.

To allow comparison, the same cutoff on the second differences has been used (i.e. zero) for all the curves, and the threshold on F_{step} has been fixed to a value giving a probability of error of 1% in the absence of an edge (signal = 0). (When made possible by the available parameters, this is the quantity we will keep constant through all characteristic curves; as it is more relevant for performance comparisons). The signal is measured in number of light intensity standard deviations.

The sensitivity of course improves with the number N of grid-points. A measure of the sensitivity is the signal for which detection is 50% successful. Figure 19 plots this quantity as a function of N ; sensitivity improves approximately as \sqrt{N} .

Geometric considerations limit the size of the neighborhood. For the line-verifier, 70 grid-points, a rectangle of length 14 and height 5, seemed an adequate choice.

DETECTION CHARACTERISTIC CURVES FOR VARIOUS SIZES




figure 18

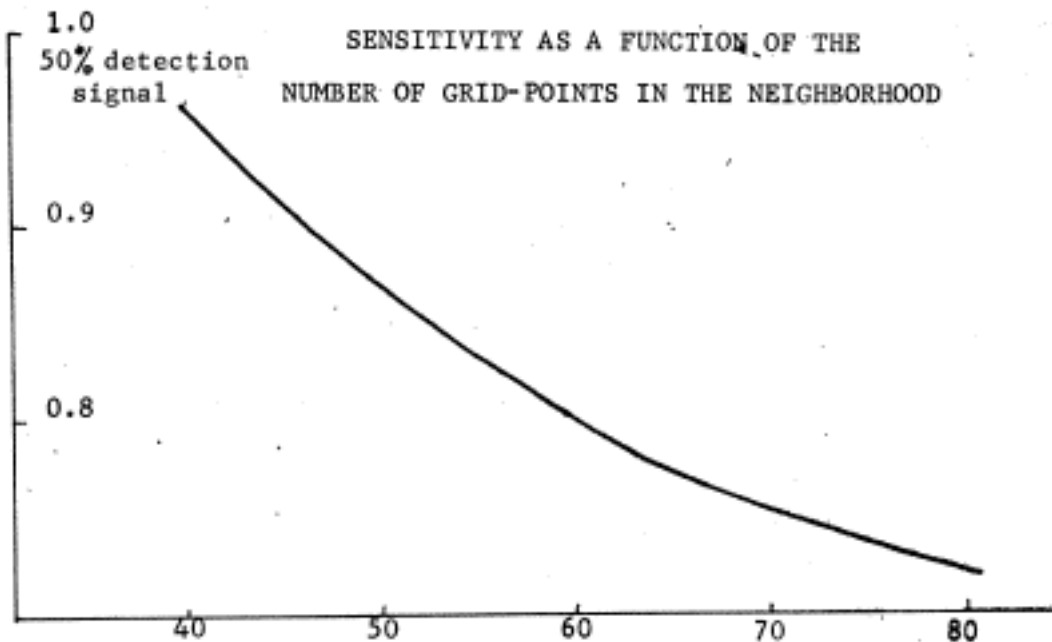


figure 19




figure 20

All numerical results, in the remaining of the thesis, are based on a basic neighborhood containing 70 grid-points.

2) Next, we vary the cutoff on the second differences. The sensitivity, defined in exactly the same manner as above, shows a flat extremum for a cutoff around 1.2 standard deviation (figure 20). With $\sqrt{L} = 1.2$ v.u., the best we can approach this cutoff, with the digitization of the intensities, is by using a threshold of 1 v.u..

The resulting detection characteristic for the optimal choice of parameters is plotted in figure ; it has the usual 1% error for a zero-signal and 50% detection occurs for a step:

$$\underline{h = 0.56 \text{ v.u.}}$$

This fixes the threshold on F_{step} to $\theta_{F_{\text{step}}} = 16$.

II) Effects of the introduction of a condition on the function S

The function S defined on page 37 can be used to discriminate the center extremum of F_{step} from its satellites.

Both at the edge and at the satellites, we computed the probabilities that F_{step} be greater than $\theta_{F_{\text{step}}}$ and S be smaller than some threshold θ_S .

Figure 21 plots those probabilities as a function of the signal, for

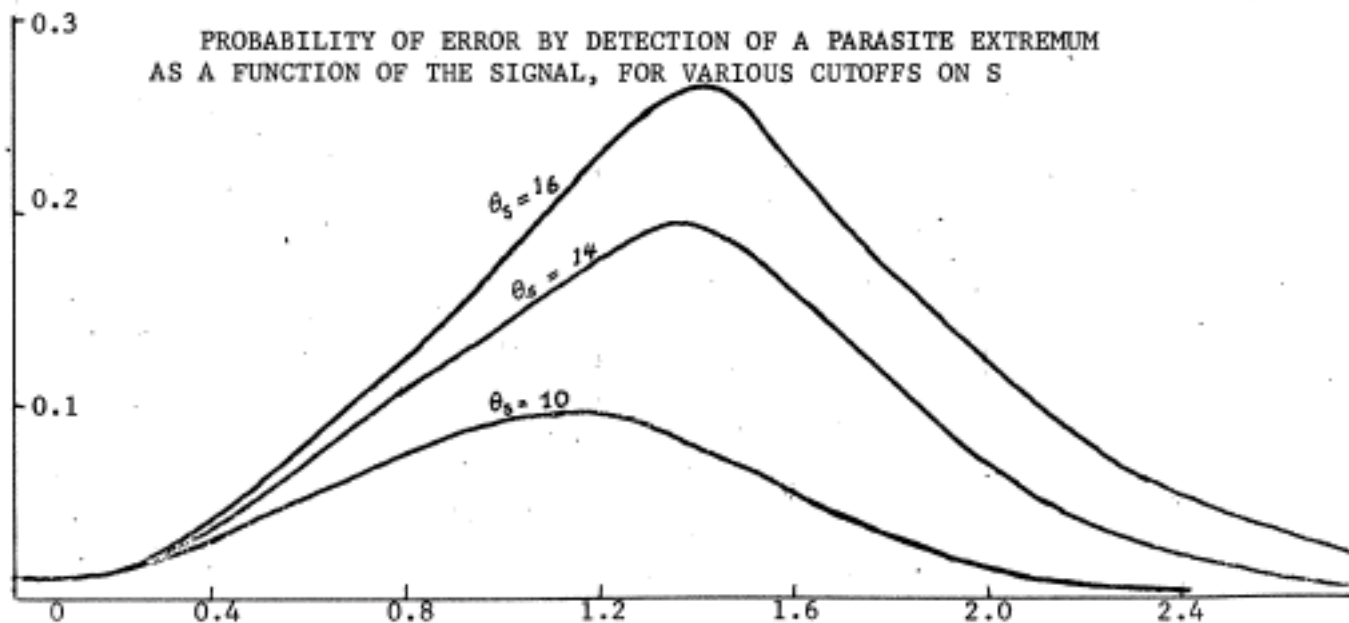



figure 21

various values of θ_S . (θ_F has been fitted in each case to get 1% of error for a zero-signal).

For $\theta_S = 10$, the decrease in sensitivity becomes too great for a line-verifier. Down to $\theta_S = 14$, this decrease is negligible. But then, the probability of error by detection of the satellites has an upper bound of 0.2, which is large.

Nonetheless we used a condition on S, as it takes no time to compute and it does perform discrimination for a whole range of signals, i.e. those ≥ 2.8 v.u. .

Remark: more thorough investigation of this condition on S were carried out. It seemed to show that no detection system, even with different sensitivity requirements, could use it alone. Reducing the error enough involves reducing the probability of detection too much. Another alternative which we would consider, if we were to build a general purpose detection system, is : look for the zero-crossings of the function S where L and R are both greater than some threshold.

2) The global procedure for the line-verifier

Between the two points given as possible extremities of an edge, 25 intensity profiles are taken, grouped by bands of 5 (figure 22). The functions F_{step} and S are computed on each band and the cutoff points are stored with the sign of F recorded.






figure 22

The next step is to find a good vertical alignment of cutoff points with the same sign of F_{step} . Note that the cutoff points are extrema of F_{step} along a band, so that we are left with very few points at this stage - from 1 to 4 per band depending on the state of the surface- ; "vertical linking" can then be very fast; usually the number of possible lines is just one, and the localization is precise (see figure 23 which shows a typical distribution of cutoff points).

The search for an approximate vertical alignment is done as follows: we consider in turn various directions within the angular uncertainty of the edge. The cutoff points are projected along each direction, until at least 4 of them, each on a different band, yield approximately the same direction.



SECOND DIFFERENCES
CUTOFF POINTS



CUTOFF POINTS AFTER PROCESSING
BY F_{STEP}
(LOCAL EXTREMA OF F_{STEP})




figure 23

Let us evaluate an upper bound of the probability that an edge following the satellite extrema be detected. The probability for a parasite extremal point on a band to appear is, as we saw, at most 0.2 (with $\Theta_S = 14$). Assume that about 3 consecutive points on a band have that same probability of passing the cutoffs, Then the global probability of detection of the 'wrong' line is:

$$2 \times \binom{5}{4} \times 3 \times 0.2^4 = 0.024$$

At the global stage, and with the surface imperfections, this is too large. So we do the following; note that lines along parasite extrema follow points where F_{step} has a sign opposite to the sign on the edge. At the global stage, we process separately the points where F_{step} is positive and the points where it is negative; when each processing yields us lines, we compute, for each line, the sum of the values of F_{step} along it. For the real edge, this sum is largest (actually it is about twice the others). This is a very reliable statistic because it is global. This additional work however is necessary only when discrimination by a condition on S has failed.

In order to evaluate the performance of the global procedure - with Gaussian noise only- we used simulation. We generated steps of various amplitude with a Gaussian noise superimposed. For each given amplitude, 1000 such edges were generated and processed by the line-verifier equipped with the set of optimal parameters, and the number of successes recorded. We thus obtained an approximation to the global

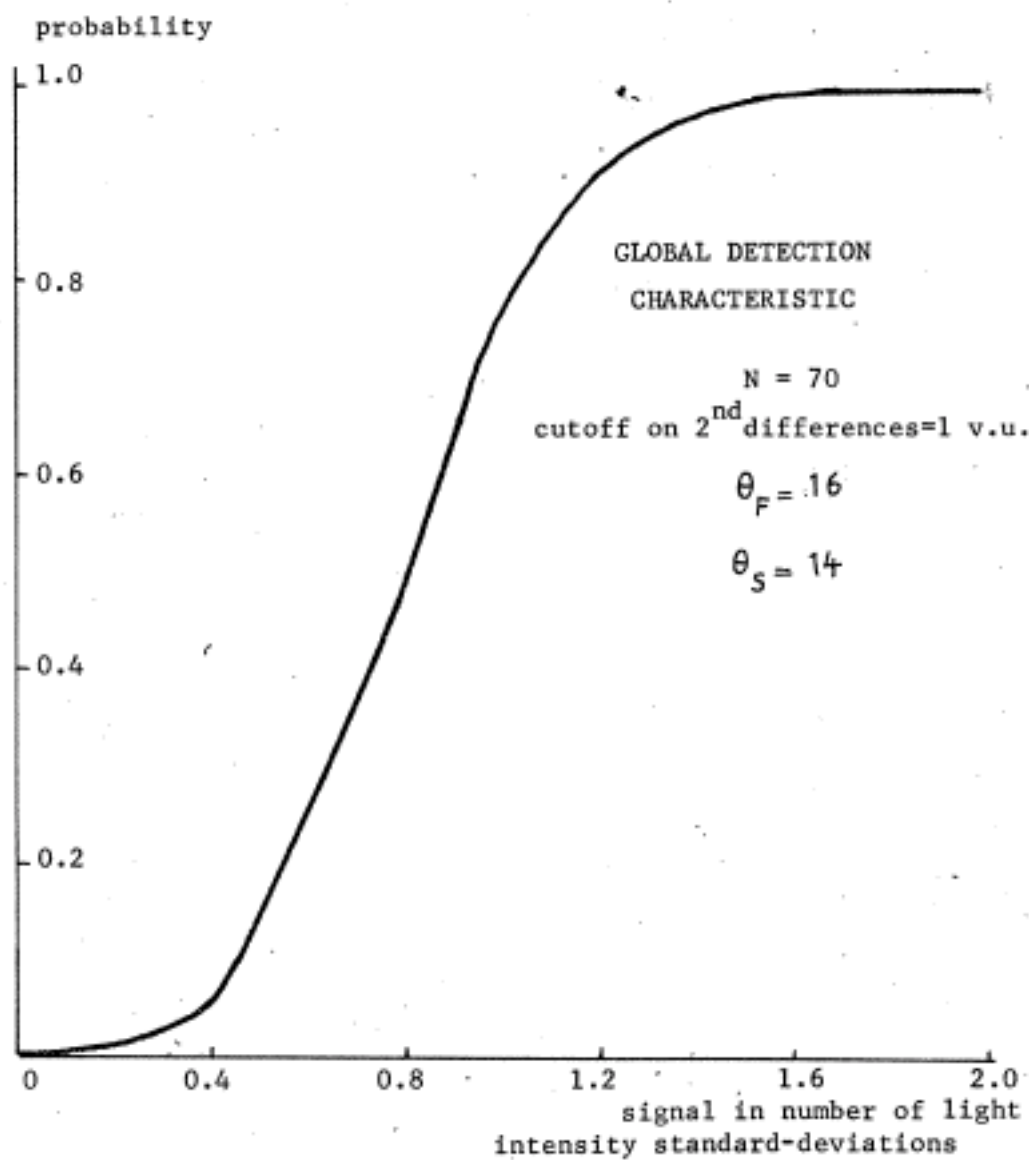


figure 24

detection characteristic (figure 24), and the following result:

With a probability of error of 0.3% on a flat distribution, 50% detection occurs for a step amplitude of

$$\underline{1 \text{ v.u.}}$$

Remark: to improve localization, whenever we find 4 points approximately vertically aligned we do not stop, using them to define the line, but we pick up all points near enough to this first approximation of the edge, and try to fit a line to them; we then throw out every point too remote from this second definition of the line and check whether we are left with enough points i.e. at least 4, all on different bands.

3) Roofs

The overall procedure is the same except for a different function F . Indeed figure 25 shows a profile and the corresponding second differences. A cutoff on the second differences will give, in an area containing a roof, a single vertical strip of all positive or all negative points.

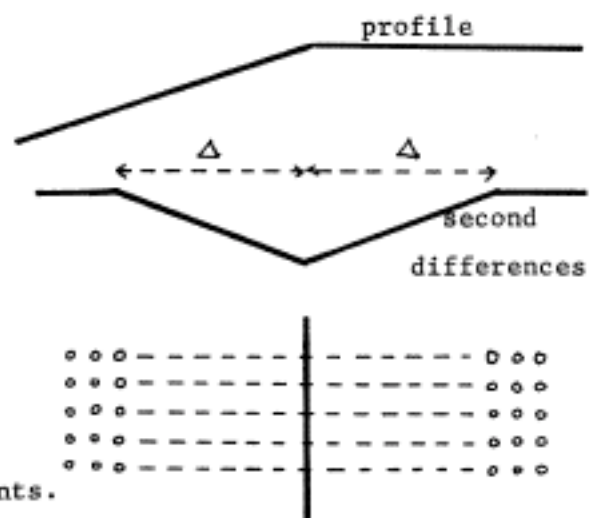


figure 25

So we replace F_{step} by F_{roof} , the difference between the number of positive points and the number of negative points within a rectangular neighborhood of properly chosen size. The local predicate will simply be based on large extrema of F_{roof} .

If the second differences are computed with an interval Δ , they are non-zero on a width 2Δ , for an idealized roof. We chose a rectangular neighborhood of height 5 and width $\frac{2}{3}(2\Delta) = \frac{4}{3}\Delta$.

The definition of a proper cutoff on F_{roof} can be done as follows: a flat noisy intensity distribution is generated, with a given standard deviation. Using the same cutoff of 1 v.u. on the second differences as we did for a step, we applied the whole line-verification process for roofs, with a given cutoff on F_{roof} . The operation is repeated a great number of times (~ 5000) and the number of wrong decisions (yes) recorded. This is repeated for several cutoffs on F_{roof} until one is found which gives us at most 1 error in a 1000.

In fact, this procedure has a limited range of efficiency. Indeed the height of the peak of the second differences varies, for a Δ of 8 (which for geometrical reasons is about maximal) from average values of 0.6 v.u. to maximals of 2 v.u. ; clearly there is a whole range of roofs which we would not see. In fact, detection for those small values of the curvature is very difficult, because there are some flat surfaces on which the curvature of the intensity surface is as large as on an edge,

(for example the translucent plastic cubes), so that there is no way of distinguishing the two. Detection then should be made by another type of analysis; i.e. the analysis of shape from shading of B.Horn (Horn 1970) will not now handle this but could reasonably be expected to extend to this problem.

4) Edge-effects


The procedure is the same as for roofs, with a different size of basic rectangular neighborhood. Figure 26 shows the distribution of the second differences cutoff points for a typical edge-effect; the cutoffs on the second differences are respectively 2, 3 and 4 v.u. for a,b,c. This suggests a cutoff of 3 and a neighborhood of width 4.

Definition of a cutoff on F_{effect} is determined as in the case of roofs; by global simulation, the criterion being to make low enough the probability of error on a flat distribution.


SECOND DIFFERENCES CUTOFF POINTS ON AN EDGE-EFFECT

FOR


VARIOUS CUTOFFS



a) cutoff at 2 v.u.



b) cutoff at 3 v.u.



c) cutoff at 4 v.u.

figure 26

5) Experiments with the line-verifier

To test the line-verifier , we started with cubes having relatively good surfaces (scene 1: figure 27), i.e. quite uniform and with little reflectivity.

First we looked for a proper value of the threshold on F by trying the verifier on faces, starting with the value of θ_F computed with Gaussian noise. As expected, we had to raise θ_F , but only from 16 - the computed value- to 18. Then, out of about 30 tries, no failure was recorded.

We then tried the verifier on the edges, modifying the lighting in order to get a wide range of different conditions. Failures occurred only by carefully adjusting the intensity on both sides of the edge, using the information returned by the verifier as to which side was darker; we could not tell otherwise as typically edges were detected such that, on their profiles as displayed on the scope, nothing was visible by eye. Actually the range of positions of a cube where failures occur is very small (translation of 1 mm in 50 mm) and we could only find a position where failures occurred only half of the time. The localization was good and it was insensitive to displacements of the edge from its expected position within the ranges which we expected; i.e. the extremities of the edge can be within ± 10 points from the expected.

Very dim reflections of an edge against another surface were detected

as edges.

To test the limits of functioning of the verifier, we set up two other scenes with very bad objects (scenes 2 and 3; figure 28 and 29). When keeping the same constants as in scene 1, we made errors most of the time on plane surfaces. Then, instead of raising θ_p , which would have diminished our sensitivity, we used 10 bands instead of 5, and required at least 8 of them to carry a cutoff point. Then we could record no failures on flat surfaces, although some of them practically showed texture; end-grain and wood-lines (note that the procedure ignores every line less than 3 points wide). The sensitivity on edges was quite similar to that described for scene 1.

Rounded edges as in scene 2 were successfully detected. Trying the verifier across an actual edge does not cause errors; i.e. the answer is no. Roofs and edge-effects are detected as such (roofs appear essentially with the translucent plastic cubes).

We think that the line-verifier should adapt its constants depending on data made available by the higher level procedures; like the space probably available for taking cross-sections without overlapping other edges, corners etc...; one could possibly use trapezoidal neighborhoods when one looks near corners. Data about the state of the surfaces also could be used to fix the best values of the parameters for each particular verification.

The computation time is about 1/2 second, (the program is in MIDAS).

This page is
missing from
the original
document.

This page is
missing from
the original
document.

CONCLUSION

We have proposed a boundary detection method which proved valuable for line verification. Its noise discrimination is the result of three principles; i.e. it uses summation of values of a non-linear function function of the intensities, and this function accounts for the geometry of edges.

It can detect an intensity step of 1 v.u. = 0.7% with 50% success. There are theoretically certain types of edges for which it fails (certain roofs and mixed-type edges) but their occurrence is rare and we were not able to find one in our experiments.

The most interesting property of the verifier is its very good filtering of surface irregularities, as shown by the success of experiments with practically textured surfaces. The speed of the verifier is satisfactory; most of the time is spent in scanning.

One defect is that it will have trouble distinguishing parallel edges less than 16 grid-points apart (twice the interval for the second differences) which in usual conditions is 1.6% of the scene. The trouble thus incurred is linked to the relatively low resolution of the vidisector and the fact that one cannot escape a trade-off between sensitivity and observed space.

The method could certainly be implemented as a general purpose line-finder, trading excellent sensitivity for time economy by using a looser grid and making summation over smaller neighborhoods. It would of course retain its insensitivity to surface defects which is so essential for a practical edge-detector.

APPENDIX I


Gradient at image for constant intensity at object

figure 30

The reflected intensity at the surface of the object is \mathcal{J}_o . Let M be a point on the object and $d\vec{A}_o$ a surface element on the object surrounding M ; then $\mathcal{J}_o |d\vec{A}_o|$ is the light reflected and if $d\vec{A}_i$ is the image of $d\vec{A}_o$;

$$|d\vec{A}_i| = \left(\frac{d_i}{d_o}\right)^2 |d\vec{A}_o| \quad \begin{array}{l} d_i: \text{distance aperture-image} \\ d_o: \text{distance aperture-object} \end{array}$$

If \vec{r}_a is the vector joining M to the center of the aperture, the light entering the aperture and issuing from the object element $d\vec{A}_o$ is:

$$dL = \mathcal{J}_o |d\vec{A}_o| \frac{1}{4\pi} \frac{\vec{A}_a \cdot \vec{r}_a}{|\vec{r}_a|^3}$$

solid angle defined by M and
the aperture \vec{A}_a .

The image intensity is thus:

$$\mathcal{J}_i = \frac{dL}{|d\vec{A}_i|} = \mathcal{J}_o \frac{|d\vec{A}_o|}{|d\vec{A}_i|} \frac{1}{4\pi} \frac{\vec{A}_a \cdot \vec{r}_a}{|\vec{r}_a|^3} = \mathcal{J}_o \left(\frac{d_o}{d_i}\right)^2 \frac{\vec{A}_a \cdot \vec{r}_a}{|\vec{r}_a|^3}$$

Using the thin lens relation;

$$\frac{1}{d_i} + \frac{1}{d_o} = \frac{1}{f}$$

we get:

$$\frac{d_i}{d_o} = \frac{f}{d_o - f}$$

Thus:

$$J_i = \frac{J_o}{4\pi} \frac{d_o - f}{f}{}^2 \frac{\vec{A}_a \cdot \vec{r}_a}{|\vec{r}_a|^3}$$

Let \vec{dx} be an infinitesimal vector on the object surface; to a second order approximation, we will compute the logarithmic intensity variation at the image, $\frac{\delta J_i}{J_i}$, for a displacement \vec{dx} of the object point.

$$\frac{\delta J_i}{J_i} = 2 \frac{\delta(d_o - f)}{d_o - f} + \frac{\delta(\vec{A}_a \cdot \vec{r}_a)}{\vec{A}_a \cdot \vec{r}_a} - 3 \frac{\delta|\vec{r}_a|}{|\vec{r}_a|}$$

We will calculate the separate terms:

$$\frac{\delta(\vec{A}_a \cdot \vec{r}_a)}{\vec{A}_a \cdot \vec{r}_a} = - \frac{\vec{A}_a \cdot \vec{dx}}{\vec{A}_a \cdot \vec{r}_a} = - \frac{\vec{n}_a \cdot \vec{dx}}{\vec{n}_a \cdot \vec{r}_a} \quad \vec{n}_a \text{ being a unit-}$$

vector normal to the aperture.

$$\begin{aligned} \delta|\vec{r}_a| &= \delta \sqrt{\vec{r}_a \cdot \vec{r}_a} = \sqrt{(\vec{r}_a - \vec{dx})^2} - \sqrt{(\vec{r}_a)^2} \\ &= (\vec{r}_a^2 - 2 \vec{r}_a \cdot \vec{dx} + \vec{dx}^2)^{1/2} - (\vec{r}_a^2)^{1/2} \\ &= - \frac{\vec{r}_a \cdot \vec{dx}}{|\vec{r}_a|} - \frac{|\vec{dx}|^2}{2|\vec{r}_a|} \end{aligned}$$

Given the orientation of \vec{n}_a and \vec{r}_a (See figure), we have:

$$d_o = -\vec{n}_a \cdot \vec{r}_a \quad \text{and,}$$

$$\delta(d_o - f) = \delta d_o = \vec{n}_a \cdot d\vec{x}$$

Thus;
$$\frac{\delta J_i}{J_i} \approx 2 \frac{\vec{n}_a \cdot d\vec{x}}{d_o - f} - \frac{\vec{n}_a \cdot d\vec{x}}{\vec{n}_a \cdot \vec{r}_a} + 3 \left(\frac{\vec{r}_a \cdot d\vec{x}}{|\vec{r}_a|^2} + \frac{|d\vec{x}|^2}{2|\vec{r}_a|^2} \right)$$

Now:
$$\frac{1}{d_o - f} \approx \frac{1}{d_o} \left(1 + \frac{d_o}{f} \right)$$

and if \vec{x} is the distance from the optical axis to M;

$$\begin{aligned} \frac{\vec{r}_a \cdot d\vec{x}}{|\vec{r}_a|^2} &= \frac{\vec{r}_a \cdot d\vec{x}}{|d_o \vec{n}_a + \vec{x}|^2} \approx \frac{\vec{r}_a \cdot d\vec{x}}{d_o^2 + 2d_o \vec{n}_a \cdot \vec{x}} && \text{And as } \vec{n}_a \cdot \vec{x} = 0 \\ &\approx - \frac{(d_o \vec{n}_a + \vec{x}) \cdot d\vec{x}}{d_o^2} \end{aligned}$$

Thus:

$$\begin{aligned} \frac{\delta J_i}{J_i} &\approx \frac{2 \vec{n}_a \cdot d\vec{x}}{d_o} \left(1 + \frac{d_o}{f} \right) + \frac{\vec{n}_a \cdot d\vec{x}}{d_o} + 3 \left(- \frac{(d_o \vec{n}_a + \vec{x}) \cdot d\vec{x}}{d_o^2} + \frac{|d\vec{x}|^2}{2|\vec{r}_a|^2} \right) \\ &\approx 2 \vec{n}_a \cdot d\vec{x} \frac{f}{d_o^2} - 3 \frac{\vec{x} \cdot d\vec{x}}{d_o^2} + \frac{3}{2} \frac{|d\vec{x}|^2}{d_o^2} \end{aligned}$$

$$\boxed{\frac{\delta J_i}{J_i} \approx \frac{d\vec{x}}{d_o^2} \cdot \left(2f \vec{n}_a - 3 \vec{x} + \frac{3}{2} d\vec{x} \right)}$$

The first order term in $\frac{\delta J_i}{J_i}$ cancels. If we write; $J_i = J_o S$

$$\frac{\delta J_i}{J_i} = \frac{\delta J_o}{J_o} + \frac{\delta S}{S}$$

To a first order approximation, $\frac{\delta S}{S} = 0$.

Thus $\frac{\delta J_i}{J_i} = \frac{\delta J_o}{J_o}$ and gradients at the image and at the object are the same.

We obtain a typical magnitude of the gradient at the image for a constant intensity at the object, by taking:

$$|\vec{dx}| = 1'' \quad f = 10''$$

$$|\vec{x}| = 3'' \quad d_o = 50''$$

Roughly;

$$\frac{\delta J_i}{J_i} = \frac{1}{2500} \times 25 = 1\%$$

APPENDIX II

Computation of detection characteristic curves

1) Consider the signal of figure 31.

We sample K points where the amplitude is τ ;
at each such point X , $d(X)$ is normally
distributed with mean τ and variance 1.




figure 31

We put a two-sided cutoff at θ_d on those
 K values; let

n_+ be the number of points where $d(X) > \theta_d$

n_- the number of points where $d(X) < -\theta_d$

We will evaluate

$$\pi_{\tau}(k) = \text{probability that } (n_+ - n_-) = k$$

The tables of the normal distribu-
tion give us the probabilities p_{τ}^+ (resp. p_{τ}^-)
that at a given point X where the mean is τ ,
 $d(X)$ be $> \theta_d$ (resp. $< -\theta_d$). (figure 32).

normal curve of error




figure 32

Let

$$p_{\tau}^0 = \text{probability that } -\theta_d \leq d(x) \leq +\theta_d$$

$$= 1 - (p_{\tau}^+ + p_{\tau}^-)$$

We will have $(K - n_+ - n_-)$ points where

The probability $\pi_{\tau}(k)$ will be a sum of trinomial terms.

$$\pi_{\tau}(k) = \sum_{n_+ = k}^K \binom{K}{n_+} (p_{\tau}^+)^{n_+} \binom{K-n_+}{n_-} (p_{\tau}^-)^{n_-} (p_{\tau}^0)^{K-(n_++n_-)}$$

with $n_- = n_+ - k$

2) Now consider $D2(X)$ for a step.

We sample K points right of C where

the signal is τ , and K points left of C where it is $-\tau$.

$$F = (n_+ - n_-)_{\text{right}} - (n_+ - n_-)_{\text{left}}$$

$$= k - k'$$

and

$$S = k + k'$$




figure 33

Then the probability that

$$|F| > \theta_F \quad \text{and} \quad |S| < \theta_S \quad \text{at } C \text{ is:}$$

$$P_C = \sum_{\substack{k, k' \\ |k-k'| > \theta_F \\ |k+k'| < \theta_S}} \pi_{\tau}(k) \pi_{-\tau}(k')$$

3) The computation of the similar probability at A (parasite extremum)

is given by the same formula except that in the right half the proba-

bilities are π_0 instead of π_{τ} .

$$P_A = \sum_{\substack{k, k' \\ |k-k'| > \theta_F \\ |k+k'| < \theta_S}} \pi_0(k) \pi_{\tau}(k')$$

This was programmed with a little different organization in order to

keep the computation time from getting too large.

BIBLIOGRAPHY

- Binford, T. O. 1970. A vision system. MIT Project MAC Artificial Intelligence Memo. In progress.
- Griffith, A.K. 1970. Computer recognition of prismatic solids
MIT PHD thesis.
- Hueckel, M. 1968. An operator which locates edges in digitized pictures. Stanford Artificial Intelligence Project Memo AIM 105.
- Horn, B.K.P. 1968. Focussing. MIT Project MAC Artificial Intelligence Memo 160.
- Horn B.K.P. 1970. Shape from shading. A method for finding the shape of a smooth opaque object from one view.
MIT PHD thesis.
- Roberts, L.G. 1963. Machine perception of three-dimensional solids
MIT PHD thesis.