March 1986 Report No. STAN-CS-86-1130

On Detecting Edges

by

Vishvjit S. Nalwa and Thomas O. Binford

Department of Computer Science

Stanford University
Stanford, CA 94305

ON DETECTING EDGES

Vishvjit S. Nalwa

Thomas O. Binford

ABSTRACT

An edge in an image corresponds to a discontinuity in the intensity surface of the underlying scene. It can be approximated by a piecewise straight curve composed of **edgels**, i.e., short, linear edge-elements, each characterized by a direction and a position. The approach to edgel-detection here, is to fit a series of ***one-dimensional*** surfaces to each window (kernel of the operator) and accept the surface-description which is adequate in the least squares sense and has the fewest parameters. (A *one-dimensional* surface is one which is constant along some direction.) The tanh is an adequate basis for the step-edge and its combinations are adequate for the roof-edge and the line-edge.

The proposed method of step-edge1 detection is robust with respect to noise; for (step-size / σ_{noise}) $\geq$ **2.5,** it has **subpixel** position localization ($\sigma_{position} < 1/3$) and an angular localization better than 10^0; further, it is designed to be insensitive to smooth shading. These results are demonstrated by some simple analysis, statistical data and edgel-images. Also included is a comparison, of performance on a *real* image, with a typical operator **(Difference-of-**Gaussians). The results indicate that the proposed operator is **superior** with respect to detection, localization and resolution.

I. Introduction

An edge in an image corresponds to an intensity discontinuity in the scene. Although it may correspond to an edge of an object in space, it need not. It might well be the image of a shadow (illumination discontinuity) or a surface mark (reflectance discontinuity).

It is hard to over-emphasize the importance of edge-detection in image understanding. Most modules in a conceivable vision system depend, directly or indirectly, on the performance of the edge-detector. Consequently, there has been a substantial effort in this direction. Despite this effort, many in the community believe that the problem is largely unsolved. In fact, it may be claimed with some justification, that research and motivation on other fronts (e.g., stereo and line-drawing interpretation) has been dampened by the ineffectiveness of existing edge-detectors.

Blicher [4] provides an insightful review of previous work on finding edges in image data [see also 7, 1]. Much of this work has been based on discrete approximation to differential operators [see 5]. Although step-edges do contain large first derivatives and zero-crossings of the second, the mapping is neither one-one, nor onto. It is well known that derivatives emphasize high-frequency noise. In fact, the higher the order of the derivative, the more pronounced is the effect. Further, operators that threshold on the first derivative respond to smooth shading. For example, the Nevatia-Babu Operator [14] and Canny's Operator [6] return false edges on smoothly shaded surfaces. Lateral inhibition has been proposed as a solution by Marr-Hildreth [12] and Binford [3]. However, this may involve taking 3^{rd} order derivatives.

The noise-characteristics of an operator depend on its size. The larger the operator, the more it averages out random noise. However, it is also more likely to overlap several edges or corners simultaneously and thus degrade its resolution capability. The detectability and localization of high-curvature edges also suffer. Further, as the operator size is increased, the assumptions invoked in its design

may breakdown, introducing large and unknown biases. The operator size in the foregoing discussion refers to the extent of the support used to make decisions. For example, if a (3 x 3) window is used to estimate the gradient at a point and then the decision is based on a local gradient maximum whose detection requires considering at least three adjacent estimates, then the operator width is 5 pixels and not 3.

Directional operators, like those of Nevatia-Babu [14] and Binford [3], introduce implicit averaging which is largely along the edge rather than across it. Isotropic operators, like Marr-Hildreth [12] and Shanmugam-Dickey-Green [18], on the other hand, offer simplicity and uniformity at the expense of smoothing across edges. Gaussian smoothing has been employed by Marr-Hildreth [12] and Canny [6] to reduce noise. This can be decomposed into two orthogonal 1-D Gaussian smoothing operations : one along the edge and the other across it. Let us consider the component along the edge. It is our claim that for a given support along a locally straight edge, Gaussian smoothing is less effective than simple averaging. The argument runs as follows. Given N equal intensities, each with independently, identically distributed additive Gaussian noise : the standard deviation of the weighted average of the intensities is minimized when the weights are all identical (the standard deviation in this case is reduced by a factor of $\sqrt{N}$). This argument can be equivalently carried over to the Fourier Domain. Gibb's phenomenon, although present, is not of any significance along the edge (it might, however, play a role at terminations).

Surface-fitting is among the other methods used to detect edges. It has been employed, both as a means to estimate derivatives, as by Prewitt [17] and Haralick [9], and as a classification technique, as by Hueckel [10]. The chief problem has been the choice of an adequate basis, i.e., a basis which can accurately represent the feature sought to be detected. Further, these attempts have largely failed to exploit the directional character of edges.

In this paper, a variant of the surface-fitting approach is used to detect

step-edges; however, there are significant differences from most previous approaches. **1)** An oriented *one-dimensional* surface, i.e., a surface constrained to be constant along some direction, is used. This results in effective noise-reduction without blurring the edges as severely as with circularly symmetric smoothing operators. 2) We do not seek to mark pixels as belonging to an edge, but to detect **edgels**, i.e., short, linear edge-elements, each characterized by a direction and position. Few [e.g., **10**] have treated edges thus. The directional information, we claim, is not only an essential descriptor of an edge at any point along its length, but is also valuable for linking [see **19**] and curve-fitting. **3)** The "blurring" effect of the imaging system, which can be approximately modeled by Gaussian convolution, is taken into account. "Blurring" averts undersampling and thus, as will be shown later, facilitates sub-pixel localization of the edge. Sub-pixel localization could also be achieved by deconvolution followed by the localization of discontinuities. Deconvolution, however, is an ill-conditioned problem [see **2**]. 4) An adequate basis has been found not only for most step-edges, but also for roof-edges and lincedges. These are various combinations of the tanh function with a constant. However, in this paper, we concern ourselves only with the detection of step-edges, which are by far the dominant type. 5) A **step-edge1** is declared to be present in a window if the basis constrained to have a step-shape has a better fit (in the least-squares sense) to the data than a **relatively** unconstrained basis with the same number of parameters. **Binford** suggested [3] that it is desirable to do away with thresholds altogether. Any method which selects a subset of candidate **edgels** has implicit thresholds (in our case this corresponds to the selection of the best-fitting surface). However, the choice of an explicit threshold (on the edge-contrast) does not **play a pivotal role in our** scheme as in most others. This will be illustrated in Section VII.

We make no claim to "optimality." Time and again, such claims have been made. Among others, the claimants include Hueckel [10], **Shanmugam-Dickey-Green** [18] and Canny [6]. However, often the analysis is in the continuous domain, the assumptions and criteria questionable, and the extensions to 2-D ad

hoc. An "optimal" solution is only as good as the optimality condition used.

We begin, in section II, by giving a definition of an edge in terms of the intensity profile of the viewed scene. Then, in section III, some of the problems associated with edge-detection based on zero-crossings of the second derivative are discussed. Much of the work to date has used a variant of this criterion. Section IV contains the details of our approach and Section V outlines the algorithm step-by-step as it has been implemented for step-edge1 detection. This is followed, in Section VI, by a detailed example.

The proposed approach to edgel-detection is robust with respect to noise. For $(step\text{-}size\ /\ \sigma_{noise}) \geq 2.5$, it has sub-pixel position localization ($\sigma_{position} < 1/3$) and an angular localization better than 10^0. Statistical supporting evidence in Section VII is accompanied by some simple analysis in Appendix III. Further, our operator is designed to be insensitive to high intensity gradients which do not correspond to edges. These claims are substantiated in section VIII with pictures of edgels estimated from several images. This section also includes a comparison between the performance of an implementation of the Marr-Hildreth Operator [12] and our operator. The pictures presented indicate that our operator is superior with respect to detection, localization and resolution. We conclude with section IX.

It should be pointed out that the problems of multiple scale and of linking edgels into extended edges are not considered here. Reliable edgel-detection should be expected to make these problems more manageable. Results on linked edges will be forthcoming in a sequel.

II. Definition of an Edge

Any extended edge in an image can be approximated by short linear segments called edgels, each characterized by a position and an angle. Edgels correspond to local discontinuities of various order in the intensity surface of a scene. A discontinuity of the n^{th} order is defined to be one whose n^{th} derivative contains a delta function. Hence, a line-edge is a 0^{th} order discontinuity, a step-

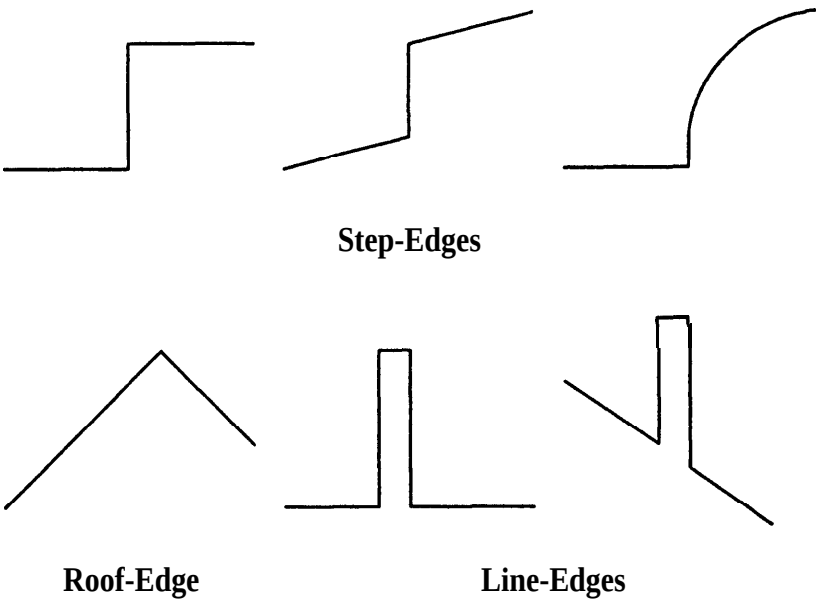

Fig. 1. ***Examples*** *oj* ***edge-profiles, as*** *they appear before being "blurred" by the imaging eystem.*

edge is a **1"** order discontinuity and **a roof-edge** is one of 2^{nd} order. Some examples are shown in Fig. **1.** However, the scene intensity-surface always undergoes "blurring" when registered by any imaging device.' This "blurring" is often assumed, for simplicity, to be linear-spaceinvariant and well-approximated by Gaussian convolution. "Blurring" is not completely undesirable, even though it limits the resolution, because it also bandlimits the signal before it is sampled. Its absence would result in severe **aliasing.**[2] As a consequence of "blurring," there are no intensity-discontinuities in the image. The importance of this will be illustrated in the following sections.

[1]**There are two fundamental sources for the "blur" associated with any imaging system. One is the diffraction limit [8] and the other is the finite sensor aperture. Optical aberrations and defocus may also contribute [see 8].**

2A manifestation of aliasing in a picture would be the "staircase" appearance of edges which are neither horizontal nor vertical. Beware of mistaking scanner-line jitter [see 2] for aliasing!

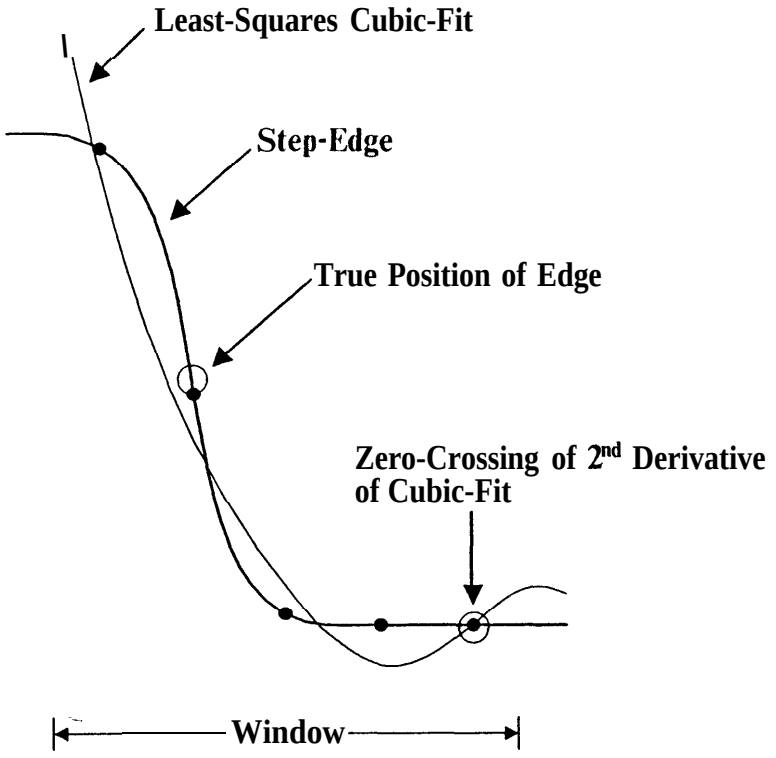

Fig. *2. Inadequacy oj the cubic-fit for a etep-edge* **cross-section** *which is positioned near the edge of a window.*

III. Zero-Crossings of the Second Derivative

Much of the work to date has used zero-crossings **of** the second **derivative to** detect and/or localize step-edges [12, 9, 6 etc.]. There are some problems associated with this approach. As indicated in the introduction, derivatives amplify high-frequency noise. In fact, the higher the order of the derivative, the more pronounced is the effect (taking the n^{th} derivative of a function is equivalent to multiplying its Fourier Transform by f^{n}). Further, if surface-fitting is used to estimate derivatives and the basis is inadequate, then the zero-crossing can result in extremely bad localization, e.g., consider the case of a cubic-fit for a step-edge cross-section which is located near the boundary of an image- window (see Fig. 2).

It is not hard to see that we can have **zero-crossings** in the absence of **a** step-edge. Zero-crossings of the second-derivative are essentially points of inflection and these need not correspond to edges, as in the case of **a** corrugated

intensity surface. It is our claim that zero-crossing operators do not adequately exploit the local intensity-profile of step-edges.

The intensity surface on the two sides of a step-edge will in general be sloped, as indicated in Fig. **1.** We will henceforth refer to such an edge as a generalized step-edge. In contrast, an ideal step-edge is constant on both sides and is a subset of the former. Note that whenever we refer to an ideal or generalized step-edge in an image, the imaging-system "blur" will be implicit. **A** simple analysis (Appendix I) of a generalized step convolved with a Gaussian shows that, in the continuous case, the localization based on zero-crossings would be biased by ($\Delta_{slope}\ \sigma_{blur}^2$ / step-size), where Δ_{slope} is the difference between the slopes on the two sides of the step and σ_{blur} is the standard-deviation of the effective blurring Gaussian mentioned in the previous section. On more than one occasion, authors have suggested Gaussian preconvolution as a method of noise reduction [12, 6]. It can be shown that this would effectively amount to having a blurring function with a variance equal to the sum of the two variances and hence, it would further degrade localization of generalized step-edges.

IV. The Details

A variant of the surface-fitting approach is used here to detect step-edges. However, unlike previous work, our basis is constrained to be directional. This results in effective noise-reduction without blurring the edges as severely as with circularly symmetric operators. Further, we take into account the fact that the image consists of samples of the true intensity profile "blurred" by the imaging system. The standard deviation of the effective Gaussian blurring function can be determined by an examination of the image of a point or a step-edge. As a result of this "blur," we have an image with no underlying discontinuities. The spectrum is bandlimited, avoiding aliasing and making sub-pixel localization possible?

[3] In the absence of "blurring," a step-edge in the image can only be localized to lie within the interval bounded by the samples on its either side. This implies a fundamental uncertainity in its position and angle. Rephrasing the above argument, it is clear that

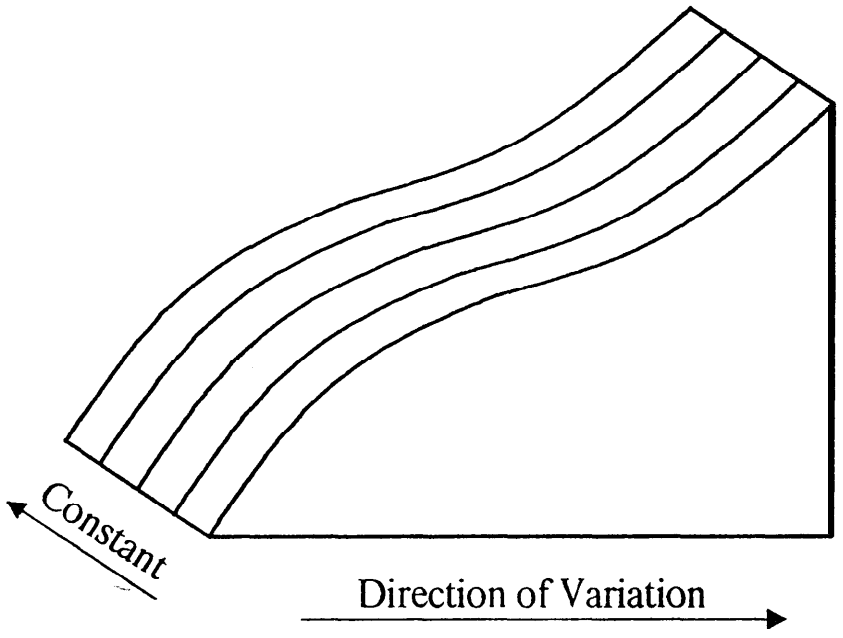

Fig. 3. One-Dimensional *Surface*

The noise in an image could be reduced by averaging data in a direction parallel to the edge if the direction of the edge could be reliably found. This, of course, relies on the fact that the window, i.e., the kernel of the operator, is small enough for the edgesegment in it to be modeled as an **edgel**. We achieve the above mentioned smoothing by fitting to each window a *one-dimensional* oriented surface, i.e., a surface which is constant in one direction, as shown in Fig. **3** (the direction of invariance would be parallel to the **edgel**). Fitting this *1-D* surface is equivalent to treating the data as strictly one dimensional by projecting it along the direction of invariance onto a plane.

Now we address the question of a reliable direction-finder for windows hypothesized to contain **edgels**. A first-approximation for the direction of **variation** can be obtained from the gradient of a least-square-error planar-fit to the

any effective blurring function which allows a many-to-one mapping of step-edges leads to an irretrievable loss of information about their position and angle.

window. However, this leads to a substantial systematic bias for rectangular windows [15], which is what we have used. A more general surface can be used to refine the first estimate and reduce the associated bias. We fit a least-squares *one-dimensional* cubic surface to the nearest 5^0. To clarify, a *1-D* cubic surface is constant in one direction and is described by a cubic polynomial in the orthogonal direction. Starting with the initial estimate of the direction, which is obtained from the planar-fit, the search for the orientation of the cubic-fit is generally not more than a few steps. It should also be pointed out that for a window with an edgel, the plot of the square-error vs angle for a *1-D* cubic-fit is bowl-shaped and centered around the true angle. Hence, once within the bowl, standard techniques like Newton's Method can be used to find the minimum. Appendix II contains. all the relevant equations for the various least-square-fits performed.

It should be emphasized that there cannot be any one unique basis which is appropriate to describe the image data in all windows. If we attempt to do this, we will obtain incorrect results when the basis is inadequate and noise-sensitive results if the basis is not minimal. Perhaps, a simple illustration of this important observation is called for. Consider a hypothetical situation where we are given some noisy samples of "y," which is a polynomial in "x," and are told to determine a description of the underlying curve. For the sake of argument, assume that "y" is a quadratic function of "x." Obviously, fitting a straight line ($y = a_0 + a_1 x$) to the data is not going to give us an adequate description. More importantly, fitting a polynomial in "x" which is of higher order than a quadratic, is going to have non-zero coefficients for x^i, $i > 2$, owing to the noisy nature of the data. Hence, even though polynomials of order greater than 2, give a smaller least-squares-error than a polynomial of order 2, it is desirable to fit a quadratic. The reader will probably recognize this to be a restatement of the underlying principle of linear regression analysis in statistics. Considerations similar to the ones just detailed have been investigated for one dimensional steps by Leclerc and Zucker [11].

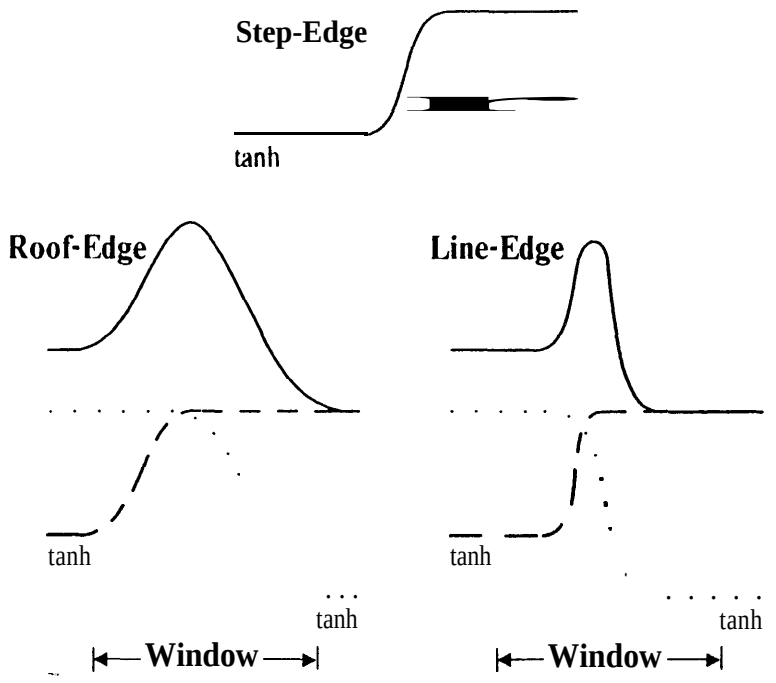

Fig. 4. Adequate *bases jot* ***edge-profiles*** *in the image are combinations* ***of*** *the tanh /unction with a constant.*

Now, consider the choice of an adequate basis. For most step-edges the tanh function with a constant, i.e., $s.\tanh(f[z+p]) + k$ where s ,p and k are the parameters and f is a constant determined by the "blur" of the imaging-system, will be adequate. As can be seen from Fig. **5,** the maximum error in **approximat-**ing an ideal step-edge by the tanh is less than **1%** of the step-size. One important by-product of employing the tanh is a reliable estimate of the contrast of the edge. From our case studies it seems that the contrast is helpful not only in linking, but also in interpretation. For roof-edges and line-edges, combinations of the tanh function, as depicted in Fig. **4,** seem to be adequate bases.

It is obvious that the above choice of basis is non-linear in its parameters. It is important for the reader to distinguish between a non-linear basis and a basis which is non-linear in its parameters: to illustrate, $(a_0 + a_1x + a_2x^2 + a_3x^3)$ is linear in its parameters while $(b + b^2x)$ is not. Whereas any least-squares surface-fitting method whose basis is linear in its parameters can equivalently be

formulated as a convolution, surface-fitting with a basis which is non-linear in its parameters cannot be thus formulated. Fitting a non-linear basis is necessarily an iterative process and is therefore computationally more expensive than fitting a linear basis.

For edges which have large deviations from an ideal step-edge, the tanh basis is inadequate and a cubic or a tanh with a cubic might be adequate. The latter has some problems because the tanh and cubic are not completely independent. It should also be noted that the cubic is inadequate for most step edges and that derivative estimates based on a cubic-fit can be quite unreliable due to the wiggles which are characteristic of polynomials. It may be desirable to employ splines when the tanh and the cubic are inadequate bases. We have used a cubic, with a check for consistency in the position estimate with the tanh-fit, in one version our detector. Our window is too small *(5* x *5)* for finding the parameters of a tanh with a cubic or of splines, in the case of horizontal and vertical edges. The position estimate based on the zero-crossing of the second derivative for a cubic-fit is biased for reasons similar to those highlighted in Appendix I. Hence, for large values of Δ_{slope}, refinement of the initial estimate may be desirable. If one uses a general basis like the cubic, it is also desirable to confirm that a dominant component of the cubic-fit is indeed a step-edge. In our implementation we accomplish this by basing our estimate of the step-size on a tanh-fit even when the basis used for detection is the cubic. We do not consider our handling of non-ideal steps to be completely satisfactory.

We compare the least-square-error of a quadratic-fit with that of a tanh-fit and choose the one with the smaller error to determine the existence or absence of a step-edgel.[4] This discriminates against smooth shading and reduces the significance of subsequent thresholding. In the initial stages, we had used the

[4] **It should be noted that both the fits have the same number of unknown parameters. This justifies our comparison of the two least-square-errors to determine which basis describes the data more accurately. The formulation of the F-Statistic corresponding to the tanh and quadratic fits is not possible, even if one ignores the non-linearity, because the bases are not nested.**

χ^2-Statistic to determine the adequacy of the basis. It was found that this was unnecessary and perhaps undesirable because of inadequate modeling of the error. A procedure similar to the one just described can be used to detect steps with large deviations from an ideal step-edge. For example, if a cubic basis is being used, then the 1, 20 F-Statistic corresponding to the quadratic and cubic fits could be employed to verify the appropriateness of the cubic-fit.

At this juncture, we would like to bring to the reader's notice some of the reasons to expect an improved performance from the use of a directional tanh-surface. First, our basis requires the specification of only four parameters which determine the orientation, the position and the upper and lower intensities of the step-edge. It is immediately seen that this is the minimum number required to describe a step-edgel with predetermined "blur." Contrast this with eight required by Hueckel's Method [10] and ten by Haralick [9]. As a result, we can use smaller windows than most previous equally sophisticated approaches. This implies better resolution capabilities and improved performance on high-curvature edges. Second, the highly constrained nature of our basis (which is borne out by the presence of only four unknown parameters) should be expected to offer noise-robustness analogous to matched-filtering classification wherein noisy patterns are categorized based on their closest "match" to noiseless representatives of the different classes. Our approach distinguishes between two classes: step-edgels and non-step-edgels. Step-edgels are characterized by a step-component of variable intensity, orientation and position. Non-step-edgels can be better described by quadratic surfaces. Of course, these assumptions may break down as the window size is increased.

As mentioned previously, we have carried out our initial investigation only for step-edges, which are by far the dominant type. Numerically, it was determined that for $\sigma_{blur} = 1$ and an ideal step-edge, the optimum scaling factor for the argument of the tanh function was 0.85. This factor was determined by minimizing the square-error. This is not surprising, as equating the slopes of the two functions at the origin would give us a value of 0.8. Hence, a rule of thumb

for the scaling factor is $(0.85/\sigma_{blur})$, where σ_{blur} determines the scale of detection. The normalized error-profile, using this factor, is shown in Fig. 5. The detection scheme is not particularly sensitive to this factor, as will be illustrated in Section VIII.

The window size is determined by the standard deviation of the blurring Gaussian. It is not hard to see that the minimum window size, irrespective of the "blur," has to be larger than **(3** x **3)** because, as illustrated in Fig. 6, there is no way to distinguish a horizontal or vertical step-edge from smooth shading if we take three symmetric samples of the edge. We chose (5 x 5) square windows. Not surprisingly, detection of zero-crossings of the 2^{nd} derivative requires a minimum lateral support of **5** pixels in the symmetric case. As the window size is increased for a' fixed "blur," we trade-off resolution for improved detection and localization of locally straight resolvable edges. However, the detection and localization of high-curvature edges will deteriorate because of the invalidity of our implicit edgel-model. Resolution refers to the ability to discriminate between and detect adjacent edges. It should be noted that we have not investigated the trade-offs accompanying different window shapes. As an illustration, notice that decreasing the window size along the edge will improve the response of the operator along high-curvature edges at the expense of deterioration in performance along straight edges.

Before we outline the algorithm, a point about intensity quantization deserves mention. Intensity quantization can cause planar (and other) surfaces to have a "staircase" appearance. Hence, local examination of the image may indicate step-edges where there are none. To see this, notice that we can generate a "step" of **1** gray level by simply quantizing the samples of a linear intensity profile with slope less than **1/2**, e.g., **(0.1x** + 0.55) can be sampled and quantized to give **{0 0 0 1 1 1).** In fact, quantization of the samples of a linear profile, with slope in the interval **(1/2, 2/3),** can produce a "step" of 2 gray levels in a 5 pixel window, e.g., **(0.6x** + **1)** can be sampled and quantized to give **{0 0 1 2 2).** The immediate implication of this phenomenon is that local edge-detectors cannot

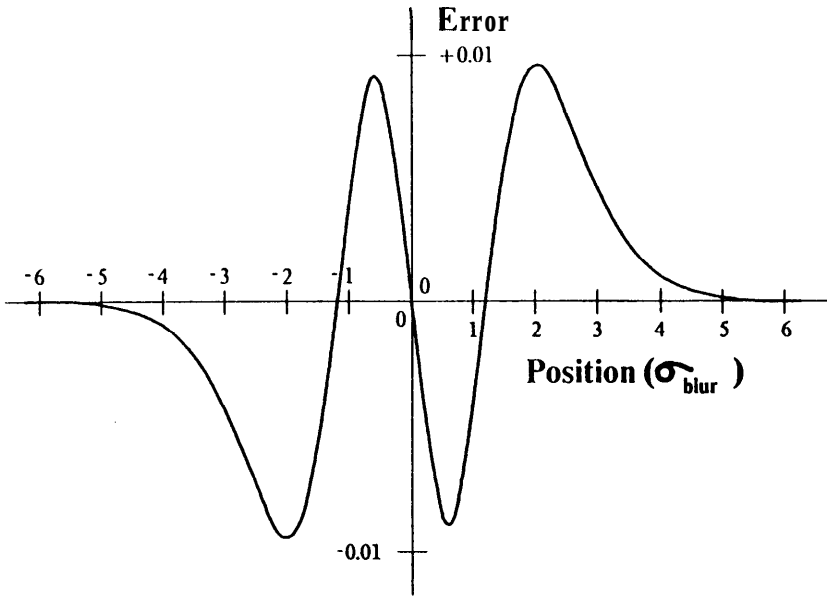

Fig. *5. Error profile resulting* ***from*** *the* ***approximation of*** *an ideal unit step-edge* ***by*** $[0.5 + 0.5 \tanh(\frac{0.85\,x}{\sigma_{blur}})]$.

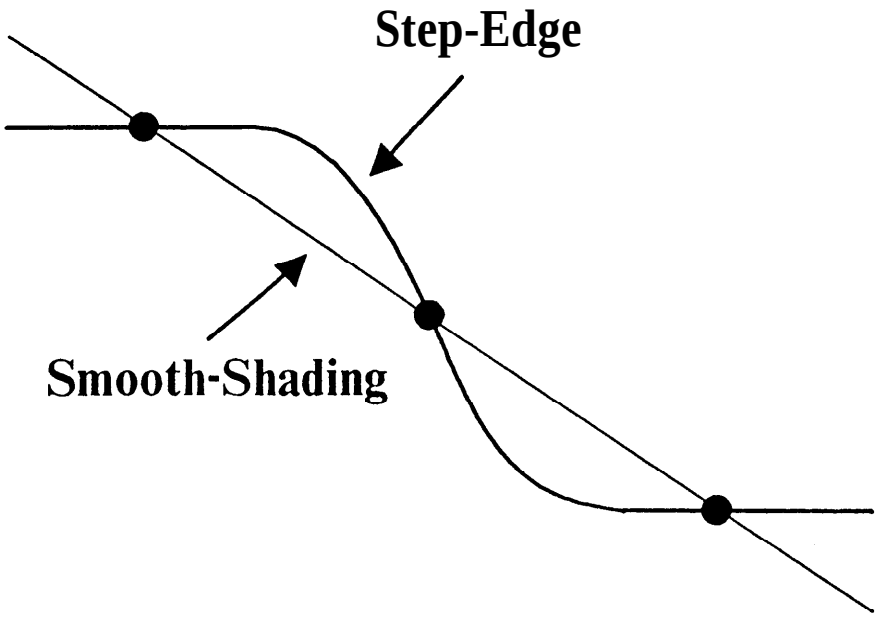

Fig. **6.** *Ambiguity in profile, given* **3** *symmetric samples* ***of*** *a* ***step-edge*** *cross-section.*

cope with fundamental ambiguities resulting from intensity quantization. Keeping this discussion in view, it may be advisable to have a lower bound of 2 gray levels for the edge-contrast threshold in a (5 x 5) window.

V. Outline of Algorithm for Step-Edge1 Detection

The following is an outline of the procedure used to detect the presence of an edge1 in an image-window. This procedure is to be repeated over the whole image by shifting the window in l-pixel steps in the x and y directions.

(All the relevant equations and statistics are listed in Appendix II.)

(i) Perform a least-squares planar-fit to the window and use the gradient of this fit to obtain an initial estimate for the direction of variation in the window, assuming that the underlying intensity surface is *1-D.*

(ii) Refine the above estimate of the direction of variation by fitting a *1-D* cubic surface with the least-squares-error criterion. The resulting equations are non-linear in the angle. However, owing to the reliable initial estimate, the search is typically a couple of steps. We find the angle to the nearest 5^0.

(ii)' *{Optional)* Calculate the 2, 20 F-Statistic for the planar and cubic fits obtained in (i) and (ii). If it is less than the 75% threshold, then declare the absence of an edgel. This thresholding serves the purpose of reducing computation by considering only those windows which exhibit a statistically significant reduction in the least-square-error by employing a cubic basis rather than a planar one.

(iii) Find the least-squares *1-D* tanh surface oriented in the direction found in (ii). The tanh-fit is localized to the nearest **0.1** pixel. As will be seen in Section VII, for low and moderate S.N.R.'s the position accuracy is not determined by the quantization error associated with the search steps.

(iv) Find the least-squares *1-D* quadratic surface oriented in the direction found in (ii). If the least-square-error in this case is less than that for the tanh-fit, then declare the absence of an edgel.

(v) The least-squares tanh-fit performed in (iii) determines the intensities on the two sides of the step and its position in the window. The position of the step-edge is given by the displacement of the tanh term and its intensities are determined by the sum and difference of the constant term in the basis and the coeffkient of the tanh term.

(vi) Threshold on the step-size determined from (v). To improve the reliability of the detection process, it may also be desirable to require the edge1 to be localized within some central sub-window, e.g., 2-pixel x 2-pixel.

N.B.If one wants to detect step-edgels which have large deviations from an ideal step-edgel, steps similar to (iii) and (iv), but with a basis different from the tanh, will have to be added. Of course, the appropriate statistical formulation will have to be used.

VI. An Example

We now proceed to illustrate the algorithm outlined in the previous section with an example. Consider the image-window in Fig. 7-b which is a noisy version of that in Fig. 7-a. The underlying intensity step-edge shown in Fig. 7-a has grey-levels 64 and 128 on its two sides and $\sigma_{blur} = 0.6$. The edge is located at a distance of 0.2363 pixel from the center of the window and at an angle of 34.4' to the x-axis. The noise in Fig. 7-b is additive white zero-mean Gaussian with $\sigma_{noise} = 8$. The detected edge is located at a distance of 0.1679 pixel from the center of the window and at an angle of 30^0 to the x-axis. The error in position is 0.0684 pixel and the error in angle is 4.4'. Recall that the position quantization error is $\pm$ 0.05 pixel and that the angle quantization error is $\pm 2.5^0$. As mentioned in the previous section, the relevant equations are listed in Appendix II. The z-axis shown in Fig. 7-c is the estimated direction of variation in the window and is orthogonal to the estimated orientation of the edge.

Underlying Edge

128	128	128	124	108
128	127	118	97	75
124	110	86	69	64
99	76	66	64	64
70	64	64	64	64

Fig. 7-a.

Detected Edge

135	132	132	127	118
121	133	110	101	88
111	110	72	75	61
108	78	61	69	71
71	63	53	61	66

Fig. 7-b.

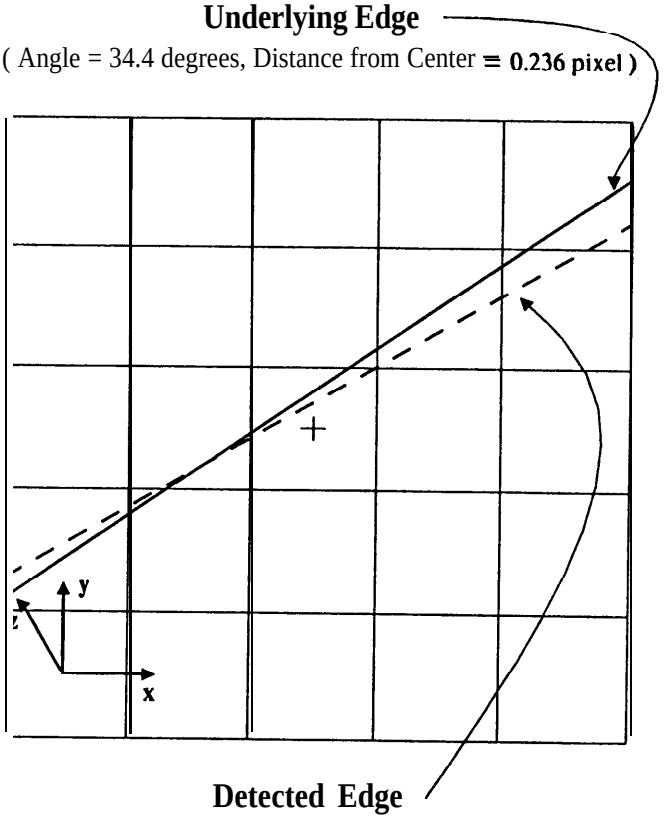

Fig. 7-c.

Fig. *7. Example : (a) Original image-window with step-size* $= 64$*; (b) Noisy image-window with (step-size* $/ \sigma_{noise}$ *)* = 8; (c) ***Underlying and detected edges.***

(i) Least-Squares Planar-Fit

$I[x, y] = 74.72 - 7.342 + 16.52y$

Least -Squares -Error = 2683, $\theta_0 = \tan^{-1}\frac{16.52}{-7.341} = 115^0$ *(to nearest* 5^0*)*

θ_0 is the direction of the gradient of the planar-fit and is used as a first estimate for the direction of variation in the window.

(ii) Least-Squares **1-D** Cubic-Fit

$I[x, y] = 71.74 + 21.83z + 6.19z^2 - 2.16z^3$

$z = x\cos(\theta) + y\sin(\theta)$, $\theta = 120^0$ *(to nearest* 5^0*)*

Least -Squares -Error = **1295**

θ, a refined estimate of θ_0, is the final estimate of the direction of variation in the window and is orthogonal to the direction estimate for the **edgel,** if any.

(ii)' *(Optional)*

The 2, **20** F-Statistic for the planar and cubic fits is **10.7** and it does exceed the 75% threshold which is **1.47.** Hence, we continue with the rest of the algorithm.

(iii) Least-Squares **1-D** Tanh-Fit along θ

$$I[x, Y1 = 95.57 + 32.52\tanh\left(\left(\frac{0.85}{0.6}\right)(z - p)\right), \quad p = 0.9 \text{ (to nearest } 0.1 \text{ pixel)}$$

$z = x\cos(120^0) + y\sin(120^0)$

Least -Squares -Error = **1203**

p is the estimate of the position of the edge1 along the z-axis.

(iv) Least-Squares **1-D** Quadratic-Fit along θ

$I[x, y] = 77.80 + 15.86z + 1.45z^2$

$z = x\cos(120^0) + y\sin(120^0)$

Least -Squares -Error = **2615**

The quadratic-fit least-squares-error is more than the tanh-fit **least-squares-**error. Hence, an edge1 has been detected.

(v) Edge Parameters

The intensities on the two sides of the step are estimated from (iii) to be **63.1** and **128.1** (i.e., **95.57 ± 32.52).** The orientation of the edge1 is determined from (ii) to be 30^0, i.e., orthogonal to θ, the direction of variation. Its position is determined from (iii) to be 0.9 pixel from the origin along the **z**-axis or equivalently **0.1679** pixel from the center of the window.

VII. Statistical Data

We now present some statistical results obtained for our edge-detector. The algorithm outlined in steps (i) through (vi), excluding (ii)', of Section V, was implemented. Let us begin by clarifying our notation. **Signal-to-Noise-Ratio (S.N.R.)** is defined as *(step -size* / σ_{noise}) where σ_{noise} is the standard **deviation** of the noise. The noise is assumed to be additive white zero-mean Gaussian. A false-positive occurs when no edge is present in the window and an edge with contrast greater than the threshold is declared. A true positive occurs when an edge is present in the window and it is identified as such, with its contrast greater than the threshold, the error in position (perpendicular distance from the center of the window) less than 0.7 pixels (half the diagonal of a pixel-support) and the error in angle less than 15^0. $\sigma_{position}$ is the root-mean-square of the error in the position and σ_{angle} is the root-mean-square of the error in the angle. We use σ to denote the r.m.s. values because they closely approximate the standard deviation of the errors. This is a consequence of the bias in the position and angle estimates being relatively small. The threshold is on the edge-contrast and is always stated in units of σ_{noise} .

Fig. 8 shows a plot of the false positives vs the threshold. Windows of size (5 x 5) with a constant intensity surface, $\sigma_{blur} = 0.6$ and additive white zero-mean Gaussian noise were used for this simulation. The value of σ_{blur} was chosen to be **0.6** because this was found to be a typical estimate in the real images considered in the next section. It can't be much smaller than **0.5** as then we should expect aliasing and if it's much larger, the edge is at a larger scale and we need a

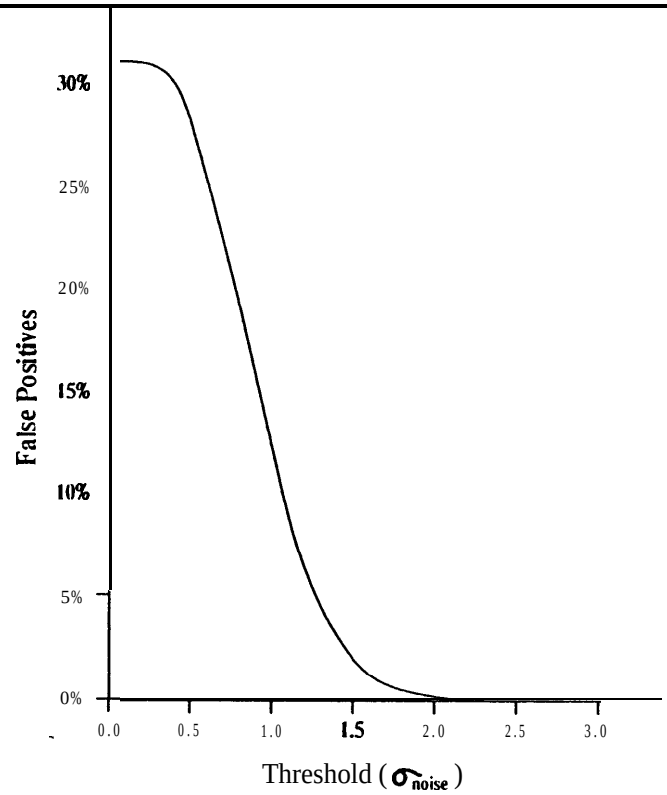

Fig. *8. Plot* ***of*** *false positives detected in windows with constant intensity and additive white Gaussian* **noise,** *as a junction* ***of*** *the threshold.*

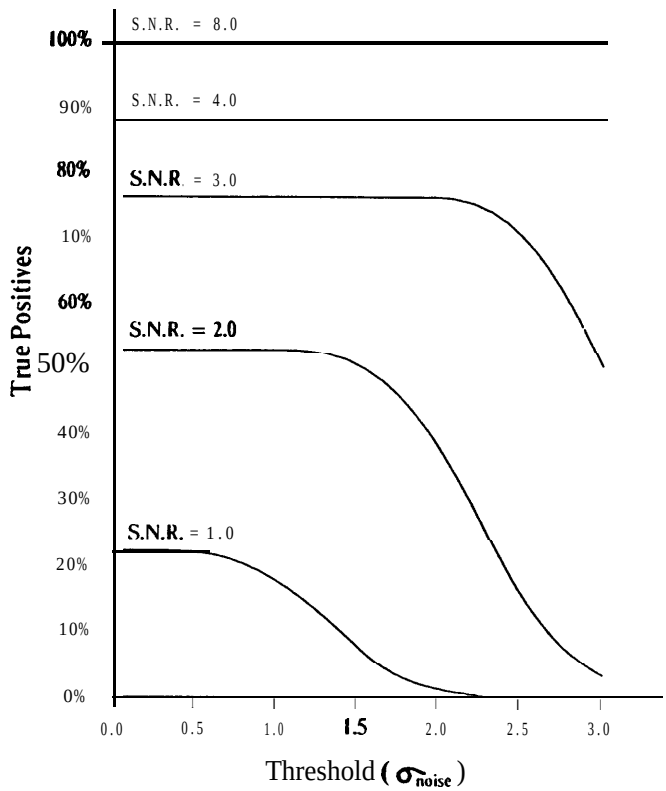

Fig. 9. *Plot* ***of*** *true positives detected in windows with synthesized step-edges and additive white Gaussian noise, as a junction oj the threshold* ***for*** *different S. N.R.* s.

correspondingly larger support. Notice, that even for a zero threshold, false positives are declared in only **31%** of the cases. This is in sharp contrast with gradient thresholding schemes which would give 100% false positives. The disparity is a result of our detection scheme requiring a certain step-like "correlation" among the samples for an edge to be declared. This requirement stems from our choice of the tanh as a basis. We have F.P.< 2.5% for a threshold of $1.5\sigma_{noise}$; F.P.< 0.2% for $2\sigma_{noise}$ and F.P.< **0.01%** for $2.5\sigma_{noise}$.

Fig. **9** shows a plot of the true positives vs. the threshold. Square (5 x 5) windows with ideal step-edges, $\sigma_{blur} = 0.6$ and additive white zero-mean Gaussian noise were used for the simulation. Each step-edge passed through the **1**-pixel square in the center of the window and its position and angle were independently uniformly distributed. Constraining the edge to pass through the central pixel-support is justified because each segment of an edge, which is not near the picture border, will pass through the center-pixel of one window or another. To reduce the contribution of gray scale quantization effects, the edge contrast was chosen to be 64 levels on a scale of **0..255.** Notice, that even for zero threshold, we do not get **100%** detection for low **S.N.R.'s.** In contrast, gradient **thresholding** schemes would give **100%** true positives. But then, they would declare any distribution to be an edge! Thus, they would have **100%** false positives too. Also notice the relatively flat profile of the plots when the S.N.R. is less than the corresponding threshold (the "knee" of the plot for a particular S.N.R. occurs when the threshold is equal to the step-size). If we synthesized images rather than windows, we should expect somewhat higher detection since each non-border edge-segment in a pixel-support is "scanned" 25 times and as pointed out earlier, we detect **edgels** and not **edge-points.**

Fig. 10 shows the plot of σ_{angle} vs S.N.R. for the true positives which would be detected in Fig. 9 if the threshold were zero and the constraints in the position-error and angle-error were removed. This curve decays to $\sigma_{angle} \approx 1.9^0$ for large **S.N.R.'s** (the diagram is cut-off at **S.N.R.=8).** This is about 30% more than what we should expect from the quantization error for a uniformly

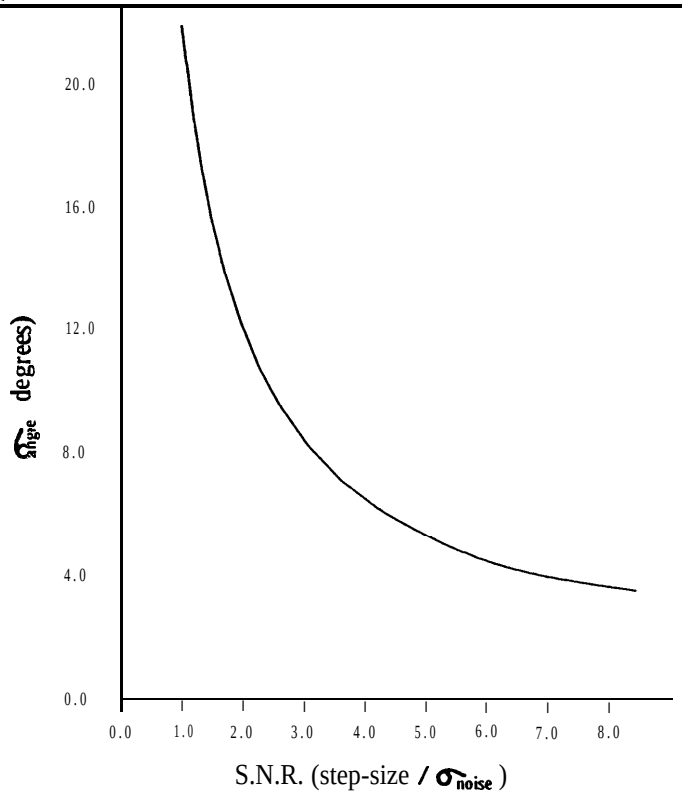

Fig. 10. *Plot of the standard deviation of the angular error in the true positives detected with threshold = 0, vs the S.N.R..*

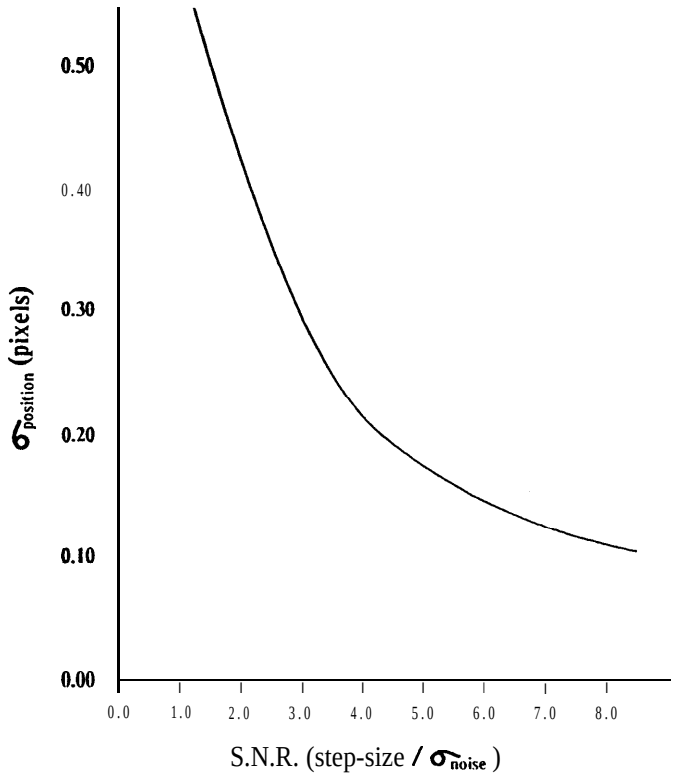

Fig. 11. *Plot of the standard deviation of the positional error in the true positives detected with threshold = 0, vs the S.N.R..*

distributed random variable [see 16]. It suggests that the bias associated with cubic-fit angle estimates, from the (5 x 5) windows, is small in comparison to the quantization interval, i.e., 5^0.

Fig. 11 shows the plot of $\sigma_{position}$ vs S.N.R. under the same conditions as in Fig. 10. This curve decays to $\sigma_{position} \approx 0.032$ for large S.N.R.'s (the diagram is cut-off at S.N.R.=8). This differs by about 10% from what we expect from the quantization error for a uniformly distributed random variable. This suggests that the position estimates from the (5 x 5) windows have a negligible bias in comparison to the quantization interval. In Appendix III, we derive an expression for $\sigma_{position}$ in the 1-D high-S.N.R. case under the assumption of known contrast. It can be shown that this is equivalent to a vertical or horizontal edge in 2-D, with the effective S.N.R. being $\sqrt{5}$ times the actual S.N.R.. It turns out that the values of $\sigma_{position}$ we would expect for vertical or horizontal edges using the expression derived in the Appendix are are within 25% of those shown in Fig. 11. This is despite the fact that the contrast of the edge is not available to the detector and the errors in the angle estimates propagate to introduce errors in the position estimate. The asymptotic value is within 10%.

The reader may also wish to know the effect of the inclusion of step (ii)' on the statistics. Although the shapes of the false-positives and true-positives plots remain more or less the same, their sizes get scaled. The false-positives plot now starts out at 11% for a zero threshold and decays to F.P.< 1% for a threshold of $1.5\sigma_{noise}$; F.P.< 0.1% for $2\sigma_{noise}$ and F.P.< 0.01% for $2.5\sigma_{noise}$. The plot of true-positives for S.N.R.=8 remains unchanged, the plot for S.N.R.=4 now starts out at 78% instead of 88%, S.N.R.=3 at 58% instead of 76%, S.N.R.=2 at 33% instead of 53% and S.N.R.=1 at 11% instead of 22%. The plots of σ_{angle} and $\sigma_{position}$ remain approximately the same.

Comparisons of the statistics of various operators are valid only if the size of the support used to make decisions is the same. Increasing the support size, which in our case is (5 x 5), would increase the fraction of true positives and

decrease the fraction of false positives, for any given S.N.R.. Also, $\sigma_{position}$ and σ_{angle} would decrease. This, however, would be at the expense of resolution between adjacent edges and the detection and localization of high-curvature edges.

We end this section with a word of caution. The analysis in the Appendix and the statistical data of this section are for ideal step-edges. They can, at best, only be indicative of the performance on real images owing to the numerous simplifications and assumptions invoked. For example, non-constant intensity surfaces have a higher likelihood of false positives than constant surfaces like those used for the statistics. The results are of no value if our edge-model is seriously flawed. Hence, although theoretical and statistical support is desirable, in the final analysis the edge-detector must also "work."

VIII. Three Case-Studies and a Comparison

It is essential to point out some details concerning the photographs displayed. a) Only the step-edge1 detector outlined in Section V (including (ii)') has been implemented. (The exclusion of step (ii)' marginally improves the detection of low-contrast edges.) b) The edgel-images have the intensity of their edges proportional to the contrast. Owing to characteristics of the display system, the high-intensity edges seem thicker than they actually are. (This artifact also occurs in the highlights of the original images.) This can easily be confirmed by an examination of the superimposed images. c) σ_{blur} is not constant over most images. This is particularly true of images with substantial "depth," e.g., Fig. 14-a. In such cases, it is desirable to use multiple values of σ_{blur} and then integrate the resulting edge1 images. However, we have restricted ourselves here to the use of a single value of σ_{blur} for each processed image. The detection process is not particularly sensitive to the value of σ_{blur} used, although the localization of the edgels may be affected. d) All edges displayed are composed of *raw* edgels with no post-processing, like linking, thinning, cleaning etc.. e) Some of the low contrast edgels have inevitably been thresholded out. A significant

improvement in their detection is likely if we were to instead use two thresholds at the linking stage as suggested by Canny [6], i.e., a high threshold to create a new linked edge and a low threshold to extend it. This procedure exploits the fact that isolated **edgels** are more likely to be false positives than those that are not. f) Degradation resulting from the various reproduction processes would make it difficult to confirm some of the edges present in the original image. This is particularly true in the high intensity regions which saturate the display well below the highest gray level. g) The pictures with the **edgels** and the superimposed **edgels** are displayed on a grid with twice the linear resolution of the original image because of our sub-pixel localization. Further, pixels in the vicinity of **edgels** in the superimposed images have been reduced to the lowest gray level, for clarity. h) It is important to bear the size of the original image in mind when scrutinizing the pictures.

(i) Industrial Setting : Bin of Parts **(*Size* : 128 x 128; $\sigma_{blur} \approx 0.6$)**

Refer to Figs. 12-a (the original image), **12-b** (the edge1 image) and **12-c** (the superimposed image). This picture was chosen to demonstrate the resolution capability of the detector and its performance on high-curvature edges. The pins of the various parts have false negatives. This is because they are bounded by dark lines, and our edge detector has currently been implemented only for step-edgels. The outer edges of the lines have been detected although not well-localized, but the inner edges exceed the resolution capabilities of our detector. Notice that some of the circular regions detected have a diameter of just a few pixels.

(ii) Aerial View : San Francisco Bay **(*Size* : 256 x 256; $\sigma_{blur} \approx 0.4$)**

Refer to Figs. 13-a (the original image), **13-b** (the edge1 image) and **13-c** (the superimposed image). This picture was chosen because of its complexity. On first glance, it may seem that there are a host of false positives. However, a closer examination of the superimposed image reveals this to be

untrue. The long lines in the sea correspond to silt lines. It may not be possible to confirm them in the photographs you will see. In any case, notice the continuity in most edges. Long continuous false positives are statistically unlikely. Also notice the detection of the small island in the mid-right of the image. In the superimposed image, the edgels are seen to impose a structure based on local intensity changes.

(iii) Indoor Scene : Telephone, Cup and Pencil (*Size* : 256 x 256; $\sigma_{blur} \approx 1.1$)
Refer to Figs. 14-a (the original image), 14-b (the edge1 image) and 14-c (the superimposed image). This image has σ_{blur} varying in the range [0.5, 1.5] owing to the "depth," and the resulting defocus, in the image. The central portion of the flower on the cup is not well detected because the resolution capabilities of our detector are exceeded.

(iv) A Comparison : Bin of Parts (Size : 256 x 256; $\sigma_{blur} \approx 0.6$)
Fig. 15-a displays the original image. Fig. 15-b is the corresponding edge1 image with σ_{blur} = 0.6 used in the argument of the tanh function. Notice that the lower edge of the specularity on the head of the plunger (located in the top center) has not been detected despite its very high contrast. This is not surprising because the edge resulting from the "saturated" specularity does not resemble an ideal edge. What is surprising is that the various other edges of specularities have been detected despite their large deviations from an ideal step-edge. In fact, the edges of the cylindrical portions, like all lambertian folds, also deviate from ideal step-edges; the extent of deviation depending on the curvature of the cross-section. A careful examination of the shaft of the plunger should reveal that two edges on either side of the specularity have also been missed. While the upper missed edge is a generalized step, the lower one is a roof edge. These conclusions were drawn by an examination of the individual pixel gray levels. While, tuning σ_{blur} to 1 allows us to detect the missed edge on the plunger head as shown in Fig. 15-c, we are still unable to detect the generalized step-edge along the

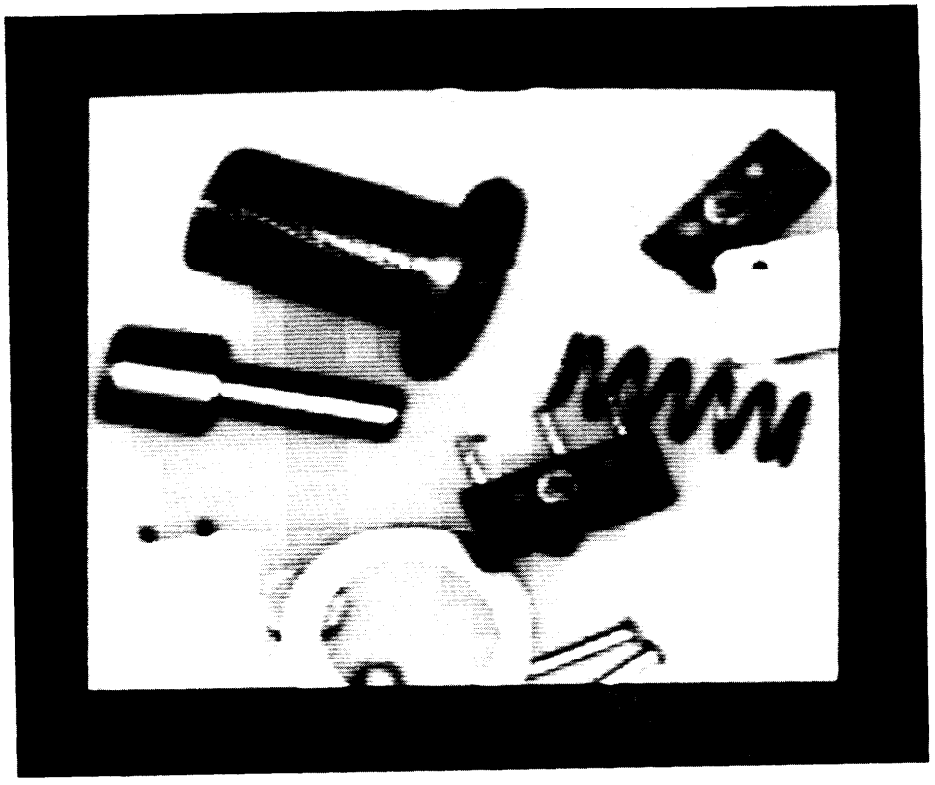

Fig. 12-a. *Bin of Parts : Original Image* **(128 x 128)**

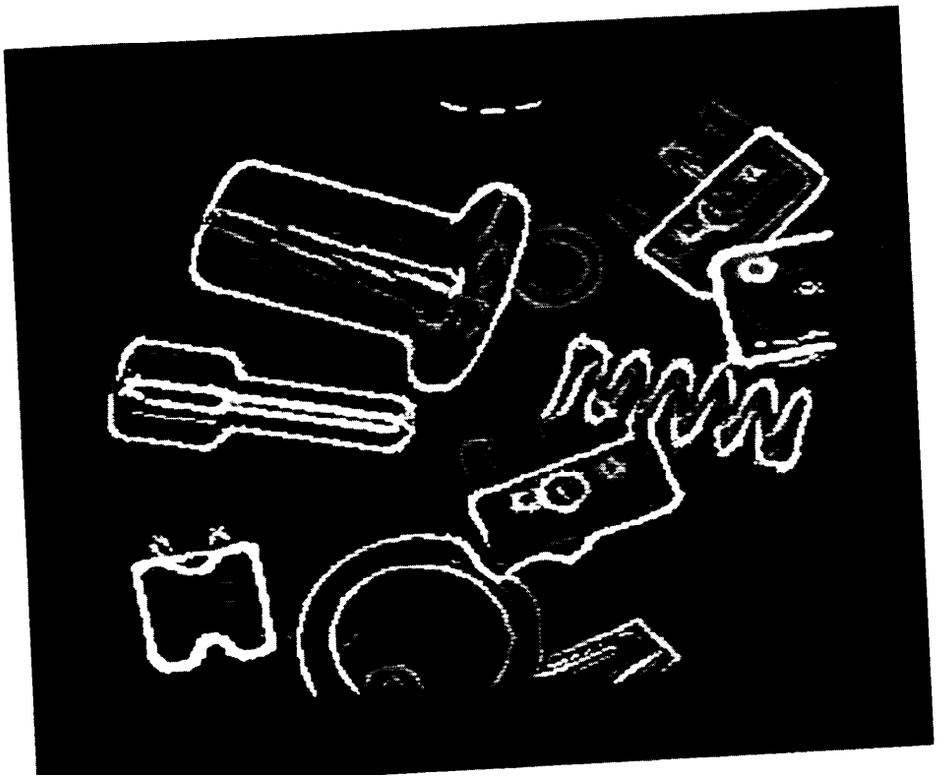

Fig. 12-b. *Bin of* **Parts** *: Edge1 Image (tanh)*

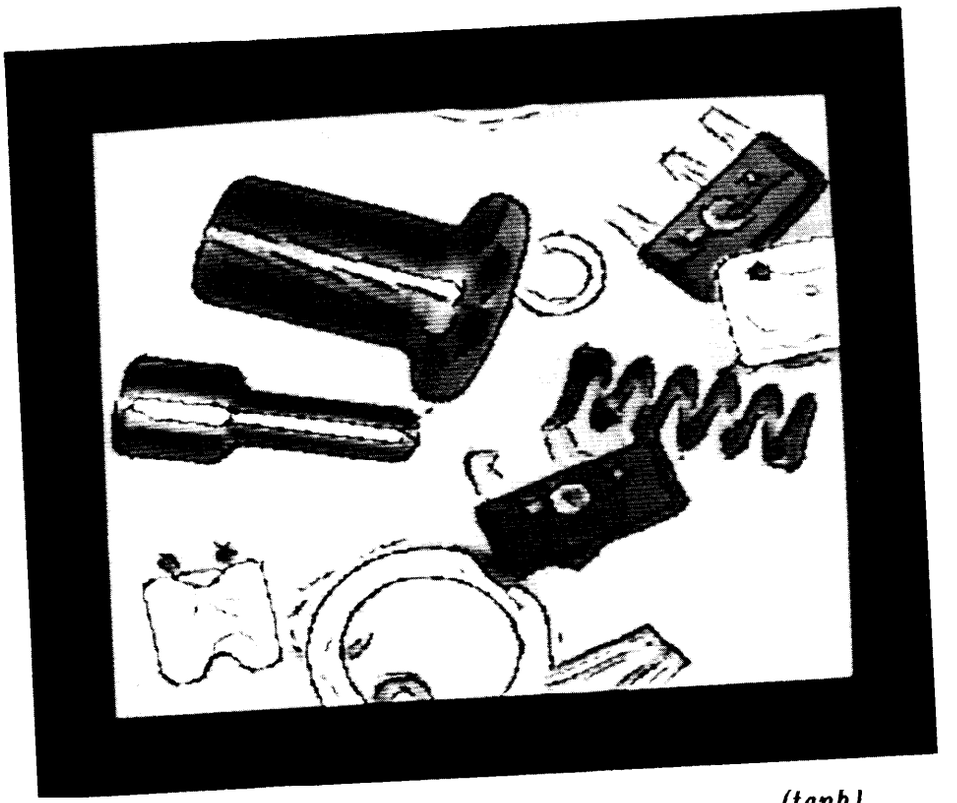

Fig. *12-c. Bin of Parts : Superimposed Image (tanh)*

Fig. 13-a. *San Francisco Bay : Original Image* (256 x 256)

Fig. 13-b. *San Francisco Bay : Edge1 Image (tanh)*

Fig. 13-c. *San Francisco Bay : Superimposed Image (tanh)*

Fig. 14-a. *Indoor Scene : Original Image* (256 x 256)

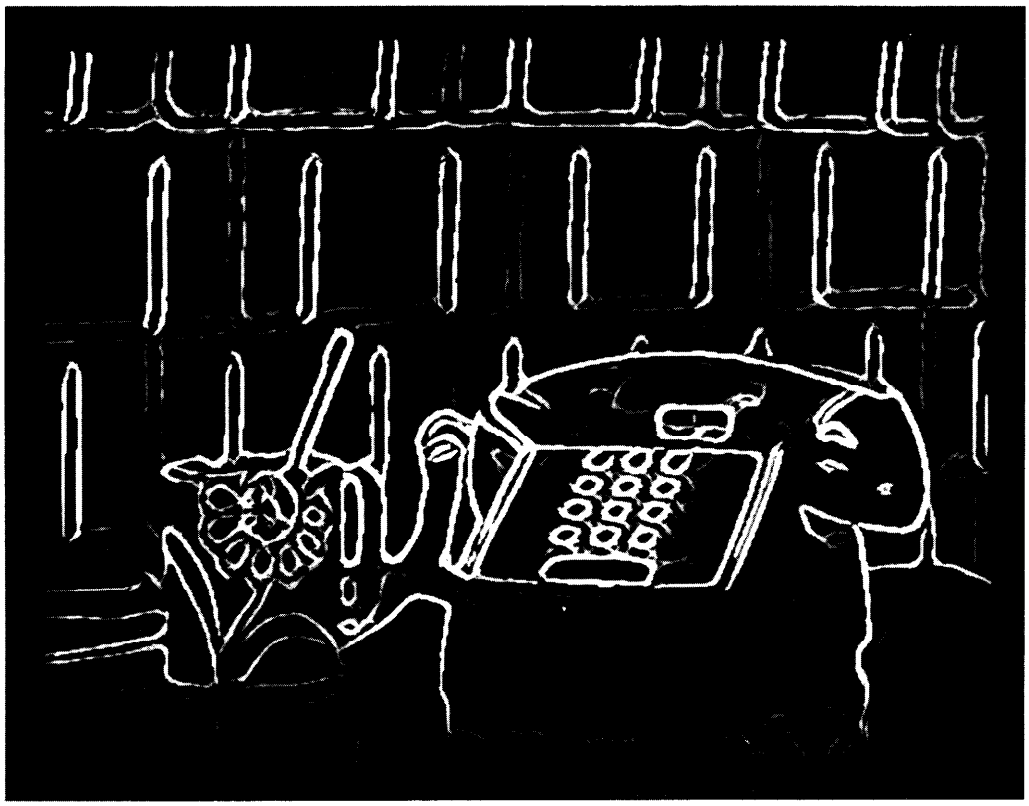

Fig. 14-b. *Indoor Scene : Edge1 Image (tanh)*

Fig. 14-c. *Indoor Scene : Superimposed Image (tanh)*

Fig. 15-a. *Original Image* (256 x 256)

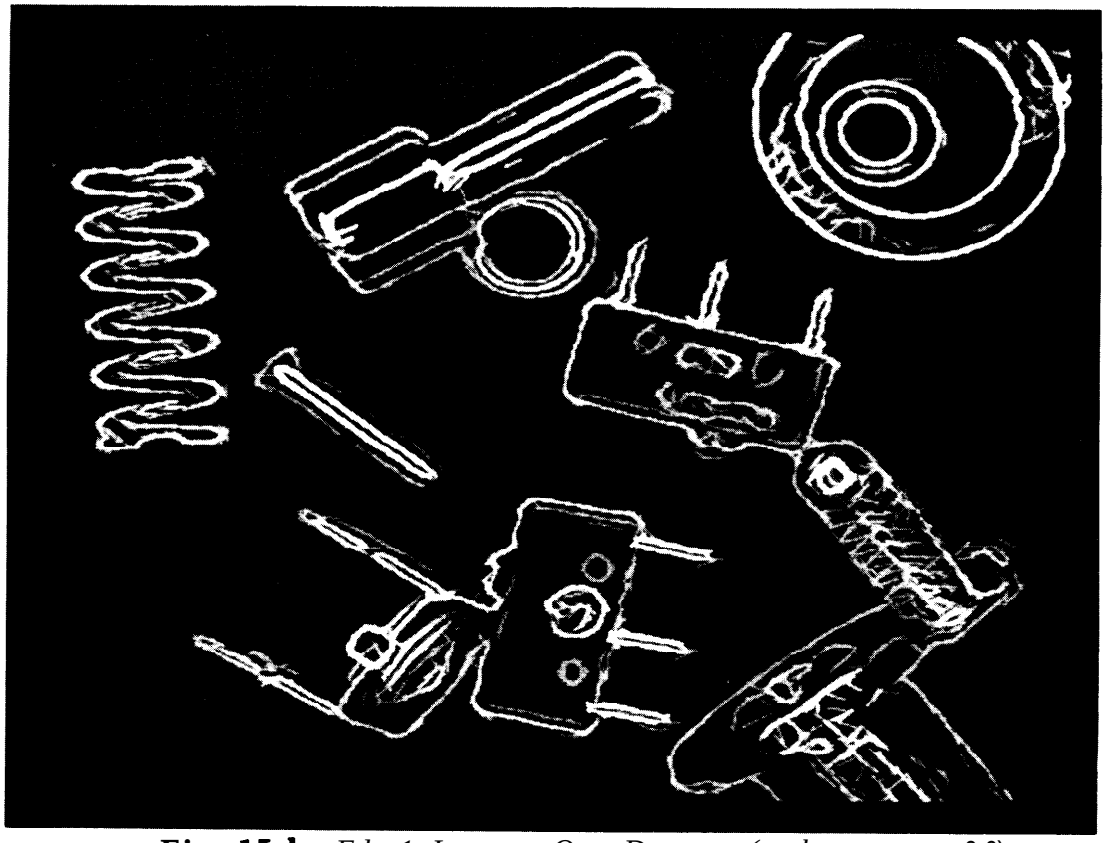

Fig. 15-b. *Edge1 Image - Our Detector (tanh,* σ_{blur} = 0.6)

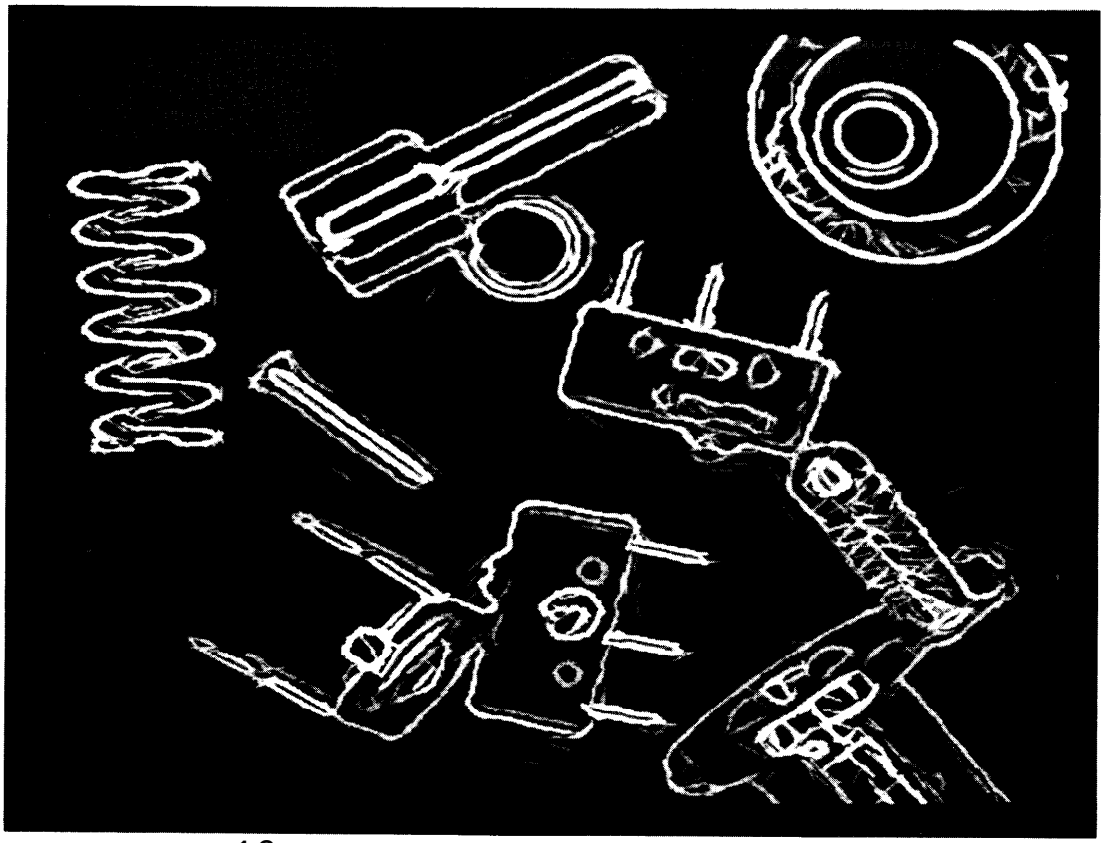

Fig. 16-c. *Edge1 Image - Our Detector (tanh,* σ_{blur} = 1.0)

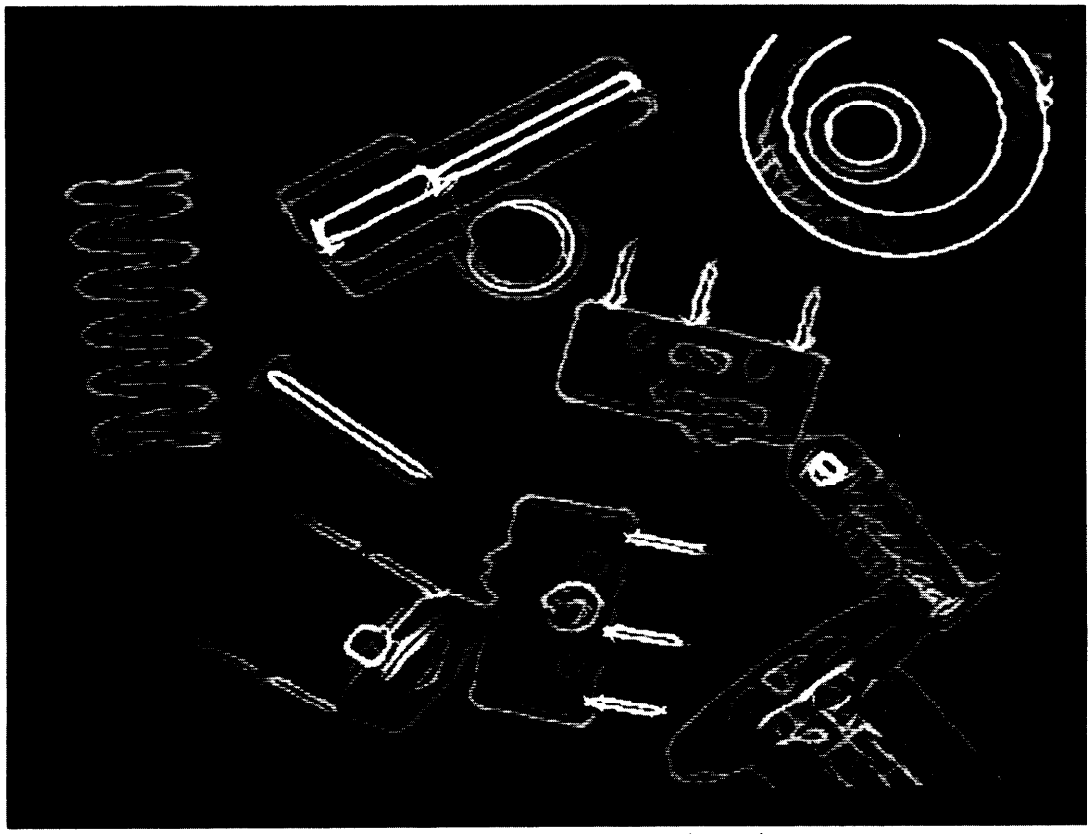

Fig. 16-d. *Edge1 Image - Our Detector* ($tanh/cubic$, σ_{blur} = 0.6)

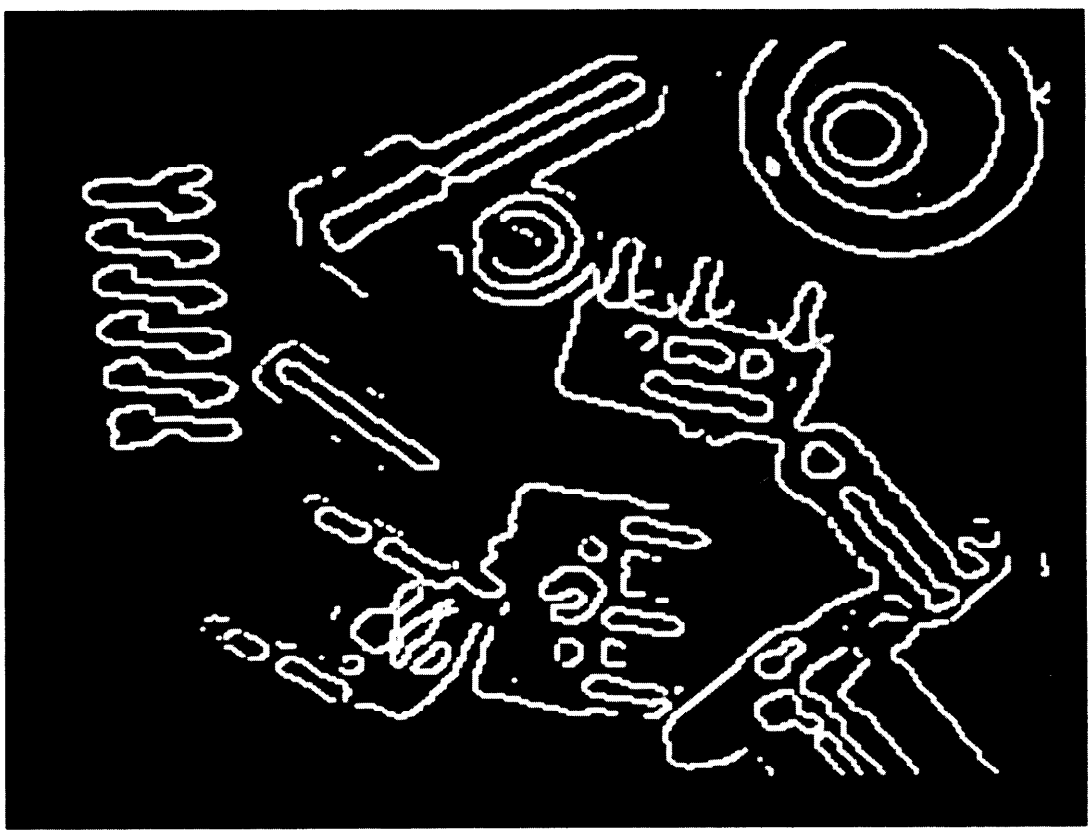

Fig. 15-e. *Edge Image - Marr-Hildreth Operator*

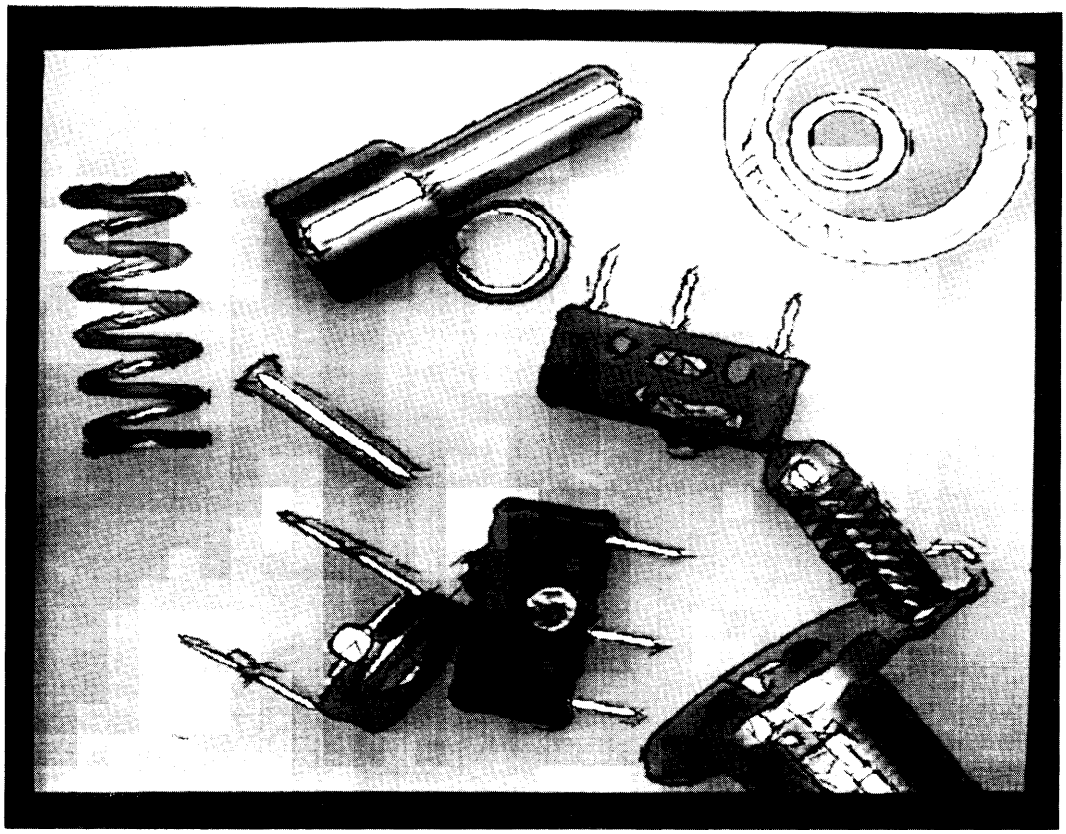

Fig. 15-f. *Superimposed Image - Our Detector (tanh/cubic)*

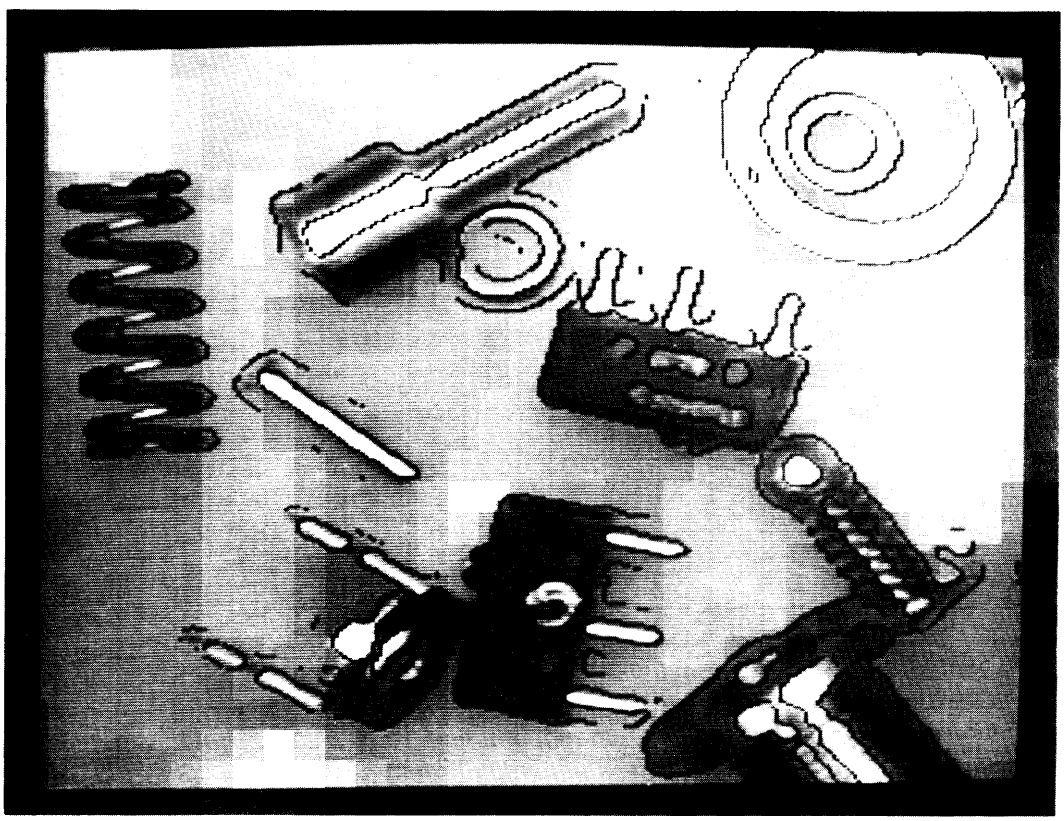

Fig. *16-g. Superimposed Image - Marr-Hildreth Operator*

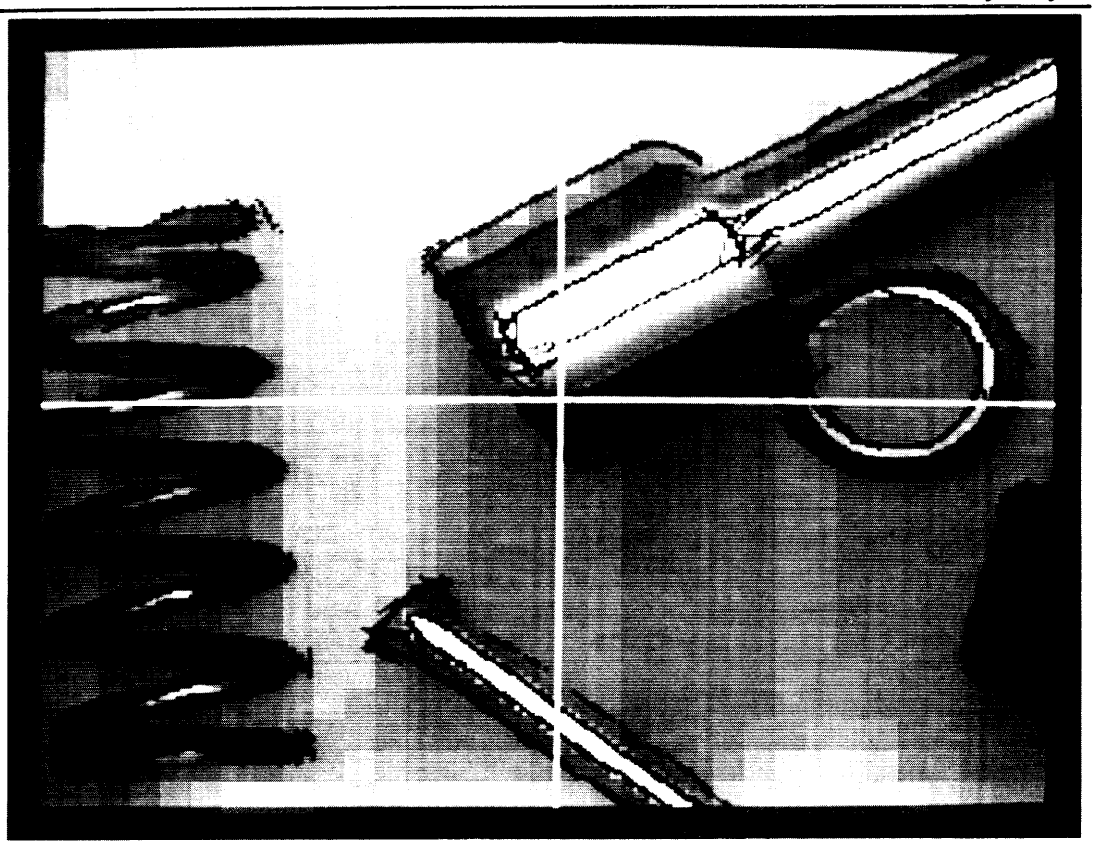

Fig. 15-h. *Close-Up of Superimposed Image - Our Detector* ***(tanh/cubic)***

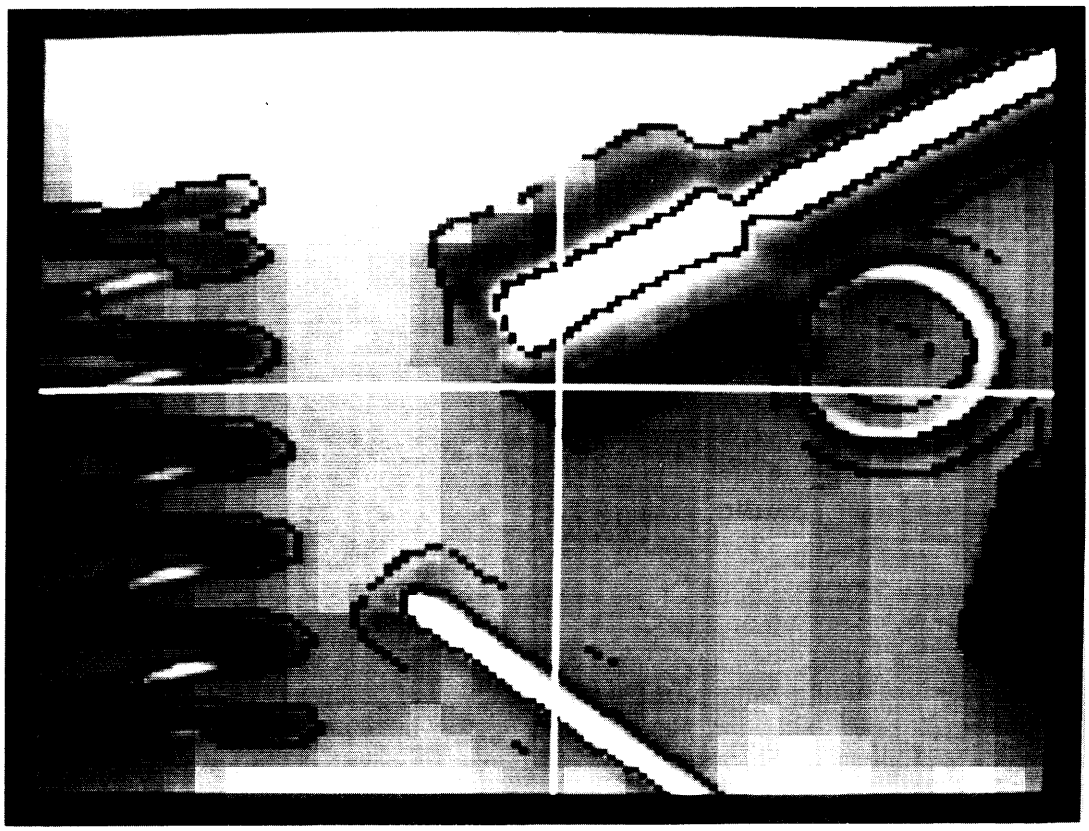

Fig. 16-i. *Close-Up of Superimposed Image - Marr-Hildreth Operator*

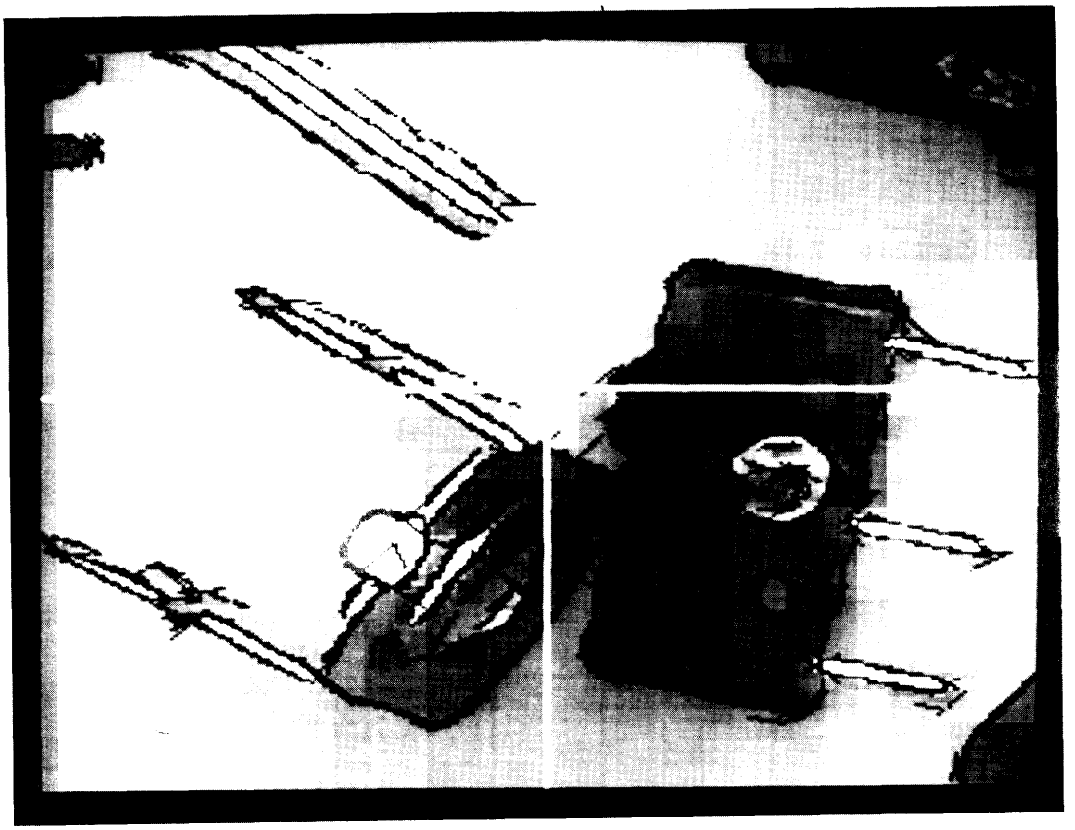

Fig. 15-j. *Close-Up of Superimposed Image* - *Our Detector* ***(tanh/cubic)***

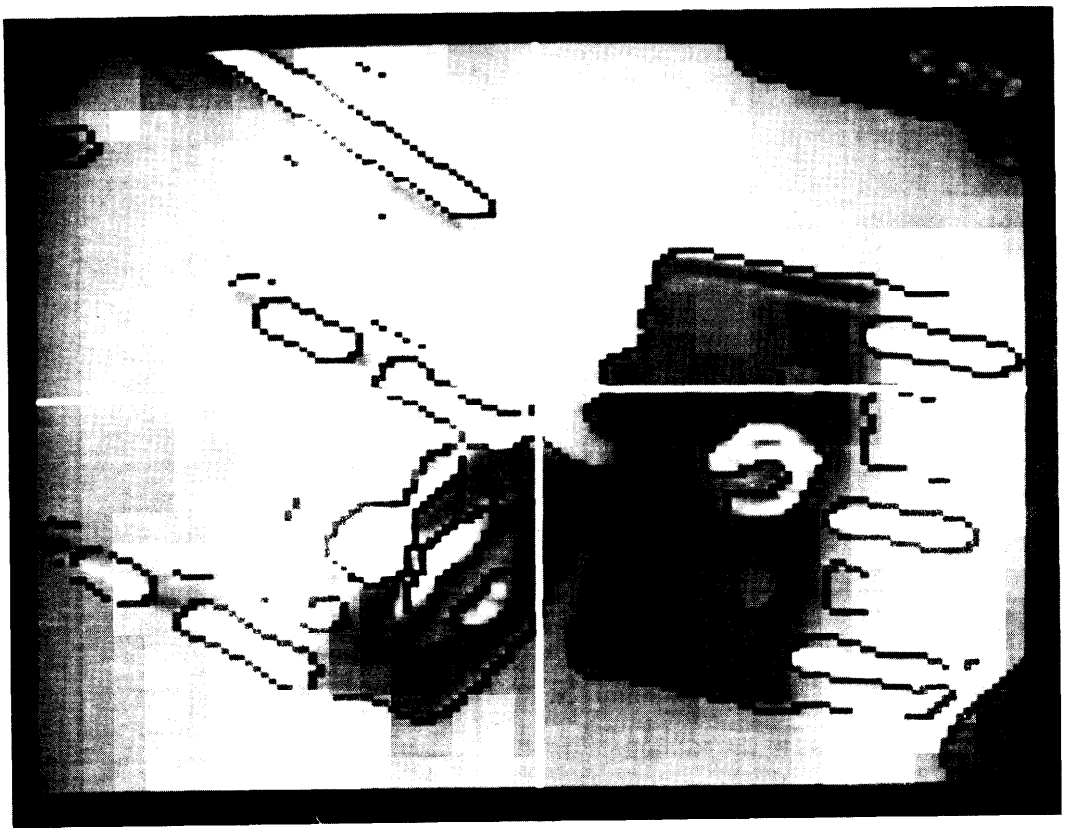

Fig. 15-k. *Close-Up of Superimposed Image* - ***Marr-Hildreth Operator***

plunger shaft. Fig. 15-d is the edge1 image using the **tanh/cubic** fit to detect generalized step edges, as explained in Section IV. The value of σ_{blur} used in this image is 0.6 and not **1.** Both the generalized step-edges which were absent from Fig. 15-b have now been detected. We would also like to draw the reader's attention to the false positives between the individual coils of the plunger/spring assembly located in the mid-right. These false positives suggest that it might be advisable to include in the algorithm a check for the validity of the assumption that the underlying intensity surface within each window is *1-D.*

Fig. 15-e is the edge image for a version of the Marr-Hildreth Operator [5] [12]. Figs. 15-f and 15-g are the superimposed images corresponding to Figs. 15-d and 15-e respectively. In order to facilitate a comparison between the two superimposed images, we *zoom-in* on **(128 x 128)** subsections in Figs. 15-h, **-i, -j** and **-k**. For reasons mentioned in the beginning of this section, it might not be possible to confirm all the detected edges. In any case, a careful examination is instructive to discover the differences in performance between the two operators with respect to detection, resolution and localization (especially of high curvature edges).

IX. Conclusion

We dealt with the problem of edge-detection by using directional *one-dimensional* surfaces. Edges were defined in terms of short, linear segments called edgels. Detection of edgels was claimed to be more appropriate than that of edge-pixels. Some shortcomings of derivative operators were then presented.

[5] **The choice of the Marr-Hildreth Operator was based solely on convenience. It was used by S.R.I. International for the I.T.A. Project in which Stanford was also a participant. As the displayed image was among those used in the Project, we expect that the operator has been tuned for optimum performance on it. The implementation used the Difference of Gaussians (D.O.G.) with $\sigma_i = 1.6$, with a (11 x 11) support, and $\sigma_e = 1$, with a (7 x 7) support. The choice of $\sigma_i / \sigma_e = 1.6$ results in a close ap proximation to the Laplacian of a Gaussian [12]. The zero-crossings were thresholded on their slope. It is conceivable that a different implementation of the operator will produce better results, but it seems unlikely that the improvement will be dramatic.**

An adequate basis for most step-edgels was shown to be the tanh. It is likely that other adequate bases exist and, in fact, if one were going to use a table look-up to perform surface fitting, the exact profile of the ideal step-edge can be stored. This is a functional of the Point-Spread-Function of the imaging system. A detailed discussion on the design of the operator was followed by an outline of the algorithm and an example. Robustness to noise, sub-pixel position localization ($\sigma_{position} < 1/3$) and better than 10^0 angular localization were statistically established for S.N.R. $\geq$ *2.5.* This was accompanied by some simple analysis and a variety of images demonstrating the performance of our operator. It was indicated that it might be advisable to include in the algorithm a check for the validity of the assumption that the underlying intensity surface within each window is *l-D,* in order to reduce false positives. It was also mentioned that we do not consider our handling of non-ideal step-edges to be completely satisfactory.

The statistical data and theoretical analysis presented in the course of this paper were intended to provide the reader with insights into the behavior of the proposed algorithm rather than numerical results applicable to "real" images. We reiterate that the numerous assumptions invoked in Section VII and Appendix III are only first-order approximations. For example, ideal step-edges were assumed. Non-zero slopes on the two sides of the step would contribute to bias the position estimates. The imaging system was implicitly assumed to be linear-space-invariant. It was also assumed that the "blur" of the imaging system could be approximated by Gaussian convolution. Although the main lobe of the diffraction pattern for a circular aperture closely resembles a Gaussian [see 8], some modifications may be warranted by the particular sensor response and the degree of defocus [see 8]. It was further assumed that the noise was independently, identically distributed (i.i.d.) additive zero-mean Gaussian. This is an accurate model for the thermal noise in electrical circuits [see 2]. However, the noise associated with photoelectronic emission is intensity dependent." At high

[6]The photoelectronic noise is often approximated to be Poisson distributed with its standard deviation equal to the square-root of the mean signal [see 2]. For high light levels, the Poisson distribution converges to a Gaussian distribution with the same standard

intensities photoelectronic noise may dominate thermal noise, rendering our noise-model invalid. Under such circumstances it may be advisable to have a signal-dependent threshold on the estimated step-size. Of course, if the thermal and photoelectronic noise are both relatively low, then the intensity quantization noise may become significant. Under general conditions, this noise has a uniform distribution in the interval [-0.5, 0.5] and its standard deviation is $1/\sqrt{12}$ [see 16].

An attempt was made to highlight some of the issues and concerns in edge-detection, as we see them. Analytical, statistical and empirical tools were employed to demonstrate the performance of the proposed algorithm. Computational effkiency was not a criterion in the design of our algorithm. We concerned ourselves solely with adequacy. The current implementation is in C on a VAX11/780. The processing time is typically 3 C.P.U. minutes for a (128 x 128) image. We expect that this can be reduced by a factor greater than 2. The algorithm is implementable as a strictly parallel process and has natural extensions for roof and line edges.

Appendix I : Zero-Crossing Bias

Let $E(x)$ be a generalized step of height S at the origin, and G (x) be a normalized Gaussian with "standard deviation" σ_{blur} .

$$E(x) = \begin{cases} k_1.x & \text{if } x < 0 \\ k_2.x + S & \text{if } x > 0 \end{cases}$$

$$G(x) = n.e^{-x^2/2.\sigma_{blur}^2}$$

Then $(E(x) * G(x))$ is the corresponding step-edge (where $*$ denotes convolution) and it can be shown, that $(E(x) * G(x))'' = E(x) * G''(x)$.

$$\begin{aligned} E(x) * G''(x) &= \int_x^{+\infty} k_1.(x-u).G''(u)\,du + \int_{-\infty}^{x} [k_2.(x-u) + S].G''(u)\,du \\ &= \int_{-\infty}^{+\infty} k_1.(x-u).G''(u)\,du + \int_{-\infty}^{x} [(k_2-k_1).(x-u) + S].G''(u)\,du \\ &= [S - (k_2-k_1).x].G'(x) + (k_2-k_1).[x.G'(x) - G(x)] \end{aligned}$$

deviation.

$$= S.G'(x) - (k_2 - k_1).G(x)$$

Equating this to zero we get, $x = \Delta_{slope} .\sigma^2_{blur}/S$ where $\Delta_{slope} = (k_1 - k_2)$ and S is the step-size. This is the biased zero-crossing of the second derivative.

Appendix II : Least-Squares Criteria and Statistics

Least-Squares Criterion for a Planar-fit :

$$\xi_P = \sum_{x,y=0}^{4} (Image\ [x,y] - (a_0 + a_x.x + a_y.y))^2$$

(minimize w.r.t. a_0, a_x and a_y)

Initial Estimate of θ (used below) : $\theta_0 = \tan^{-1}(a_y / a_x)$

Least-Squares Criterion for a Cubic-fit :

$$\xi_C = \sum_{x,y=0}^{4} (Image\ [x,y] - (a_0 + a_1.z + a_2.z^2 + a_3.z^3))^2$$

$z = x.cos(\theta) + y.sin(\theta)$

(minimize w.r.t. a_0, a_1, a_2, a_3 and θ)

θ is the angle by which the axes have to be rotated to align the x-axis with the edge1 cross-section. Its initial estimate, from the L.S.E. planar-fit, is refined here. The equations to be solved are non-linear in θ.

Least-Squares Criterion for a Tanh-fit :

$$\xi_T = \sum_{x,y=0}^{4} (Image\ [x,y] - (s.tanh(f.[z + p]) + k))^2$$

$z = x.cos(\theta) + y.sin(\theta)$

(minimize w.r.t. s , p and k)

The value of θ determined from the L.S.E. cubic-fit is used here. f is determined from the rule of thumb mentioned in section IV i.e., $(0.85 / \sigma_{blur})$. The edge contrast is 2s and p is the position. The equations to be solved are non-linear in p.

Least-Squares Criterion for a Quadratic-fit :

$$\xi_Q = \sum_{x,y=0}^{4} (Image\ [x,y] - (a_0 + a_1.z + a_2.z^2))^2$$

$z = x.\cos(\theta) + y.\sin(\theta)$

(minimize w.r.t. a_0, a_1 and a_2)

The value of θ determined from the L.S.E. cubic-fit is used here.

Statistics :

ξ_P / σ^2_{noise} follows the X^2-Statistic with 22 D.O.F..

ξ_C / σ^2_{noise} approximately follows the X^2-Statistic with **20** D.O.F..

ξ_T / σ^2_{noise} approximately follows the X^2-Statistic with **21** D.O.F..

ξ_Q / σ^2_{noise} approximately follows the X^2-Statistic with **21** D.O.F..

$\{10.(\xi_P - \xi_C)/\xi_C\}$ approximately follows the 2, **20** F-Statistic.

$\{20.(\xi_Q - \xi_C)/\xi_C\}$ approximately follows the 1, **20** F-Statistic.

The above -formulations are inexact because of the non-linearity of the cubic and tanh bases (in 8 and *p* respectively) and the fact that the value of θ used in the tanh and quadratic fits is predetermined.

Appendix III : Localization of Tanh-Fit

Let $E(x)$ be an ideal step of height S at the origin, $G(x)$ be a normalized Gaussian with "standard deviation" σ_{blur} and $\Sigma\, \delta(x - k)$ represent a discrete sampling function.

$$E(x) = \begin{cases} 0 & \text{if } x < 0 \\ S & \text{if } x > 0 \end{cases}$$

$$G(x) = n.e^{-x^2/2.\sigma^2_{blur}}$$

$$\delta(x) = \begin{cases} 1 & \text{if } x = 0 \\ 0 & \text{if } x <> 0 \end{cases}$$

Further, let $\eta(k)$ represent additive white zero-mean Gaussian noise with standard deviation σ_{noise}. Then, we can model a one-dimensional step-edge at position *p*, as below (where * denotes convolution).

$$f(k) = \left[E(x-p) * G(x)\right].\, \delta(x-k) + \eta(k) \qquad k = \cdots -1\ 0\ +1\ . **$$

Let the fitting function be $S.\{0.5 + 0.5\ \tanh(\ \frac{0.85}{\sigma_{blur}}[x - p - \epsilon]\ )\}$, where ϵ is the error in position-localization. The factor 0.85 was chosen to minimize the total-square-error, in the absence of noise. Then, the total-square-error ξ is given by

$$\xi = \sum_k \left[f(k) - S\left\{0.5 + 0.5 \tanh\left(\frac{0.85}{\sigma_{blur}}[x - p - \epsilon] \right) . \delta(x-k) \right\} \right]^2$$

Minimizing ξ w.r.t. ϵ by equating $\frac{\partial \xi}{\partial \epsilon}$ to zero, we get

$$\sum_k \left\{ S. \frac{0.85}{\sigma_{blur}} . sech^2\left(\frac{0.85}{\sigma_{blur}}[x - p - \epsilon] \right) . \delta(x-k) \right\} . [\ .\ .\ .\ .\ ] = 0$$

where $[\ \ldots\ .\ I$ *represents the error term from the preceding equation*

Now, let's assume the signal to have a high S.N.R.. Then, ϵ is small and we can substitute the $sech^2$ and $tanh$ terms by the first two terms of their Taylor series expansion w.r.t. ϵ. Dropping the ϵ^2 term and simplifying, we get

$$\left[\sum_k sech^2\left(\frac{0.85}{\sigma_{blur}}[x-p] \right) . \delta(x-k) . \left\{ \Omega(k) + \frac{\eta(k)}{S} \right\} \right] + \epsilon . \left[\sum_k \frac{0.425}{\sigma_{blur}} . sech^4\left(\frac{0.85}{\sigma_{blur}}[x-p] \right) . \delta(x-k) \right]$$

$$+ \epsilon . \left[\sum_k \frac{1.7}{\sigma_{blur}} . sech^2\ I\ \frac{0.85}{\sigma_{blur}}[z-p] \right) . \tanh\ _I\ \frac{0.85}{\sigma_{blur}}[x-p] \right) . \delta(x-k) . \left\{ \Omega(k) + \frac{\eta(k)}{S} \right\} \Bigg] = 0$$

where $\Omega(k) = \left[\frac{1}{S} . E(x-p) * G(x) - (0.5 + 0.5 \tanh\left(\frac{0.85}{\sigma_{blur}}[x-p] \right) \} \right] . \delta(x-k)$

Notice that $\Omega(k)$ are samples of the error profile shown in Fig. 5. Hence, the magnitude of $\Omega(k)$ is bounded by **0.01.** Invoking the high S.N.R. assumption once again, for typical values of σ_{blur} ($\approx$ **0.5**), we can drop the last term by comparison with the other ϵ term. Then,

$$\epsilon = -\frac{\sum_k sech^2\left(\frac{0.85}{\sigma_{blur}}[x-p] \right) . \delta(x-k) . \left\{ \Omega(k) + \frac{\eta(k)}{S} \right\}}{\sum_k \frac{0.425}{\sigma_{blur}} sech^4 \left| \frac{0.85}{\sigma_{blur}}[x-p] \right) \bullet \quad 6(z-k)}$$

where $\Omega(k)$ *is as defined above*

Taking the expected value w.r.t. η, the noise, it follows that ϵ, the error, is biased. The bias is a function of the position, p, of the original step and for a typical σ_{blur} (= 0.6), it closely resembles a sinusoid with a period of 1 pixel-width and amplitude 1.06E-2 pixel. In practice, we would be required to quantize the position we determine from the tanh-fit and in all likelihood the quantization error will be an order of magnitude more than the bias. Taking the expectation of ϵ^2 w.r.t. η, we get

$$E[\epsilon^2 \mid p] = \frac{\left[\sum_k sech^2\left(\frac{0.85}{\sigma_{blur}}[x-p]\right) \cdot \delta(x-k) \cdot \Omega(k)\right]^2}{\left[\sum_k \frac{0.425}{\sigma_{blur}} sech^4\left(\frac{0.85}{\sigma_{blur}}[x-p]\right) \delta(Z-K)\right]} + \frac{\left\{\frac{\sigma^2_{noise}}{S^2}\right\}}{\sum_k \left\{\frac{0.425}{\sigma_{blur}}\right\}^2 sech^4\left(\frac{0.85}{\sigma_{blur}}[x-p]\right) \cdot \delta(x-k)}$$

Now, let's obtain an expression for the root-mean-square of the total-error, taking position-quantization into account. Let the quantization interval, Δ_q, be 0.1 pixel and the quantization levels be centered around the origin. Further, consider p, the actual location of the edge, to be a uniformly distributed random variable in the interval (-0.5 +0.5) and let σ_{blur} = 0.6. Then, it can be shown that the quantization error is approximately uncorrelated to the bias. Hence, the mean-square of the total-error is the sum of the expectation of $E[\epsilon^2 \mid p]$ w.r.t. p and $\frac{\Delta_q^2}{12}$. The latter is the variance of the quantization error [see 16]. For the above choices of σ_{blur} and Δ_q, the root-mean-square error can be numerically evaluated to be $\sqrt{8.9E\text{-}4 + \frac{2.15}{S.N.R.^2}}$. Perhaps, it should be pointed out that the exact choice of the range of summation in the above expressions does not matter, as the first two terms on either side of the origin dominate the calculation. Under simulation with a 5-pixel-width window centered about the origin, this expression was found to be in error by less than 5% for $S.N.R. \geq 8$ and less than 10% for $S.N.R. \geq 4$. Note, that intensity quantization effects were neither accounted for in the analysis, nor present in the simulations. Also, it was assumed that a step-edge of known contrast was being localized.

Acknowledgement

V.S.N. is indebted to Ron Fearing, who was a constant source of encouragement, and to Brian Wandell for his excellent editorial commments. Jim Herson and Cregg Cowan, both at S.R.I. International, made it possible for us to offer a comparison of our results with those of an implementation of the Marr-Hildreth Operator. Sathya Narayanan made many useful suggestions to improve the running time of the implementation.

References

[1] I.E.Abdou, W.K.Pratt: "Quantitative Design and Evaluation of Enhancement/Thresholding Edge Detectors," Proc. IEEE, Vo1.67, No.5, May 1979, 753-763.

[2] H.C.Andrews, B.R.Hunt: "Digital Image Restoration," Prentice-Hall Inc., Englewood Cliffs, 1977.

[3] T.O.Binford: "Inferring Surfaces from Images," Artificial Intelligence, 17, August 1981, 205-244.

[4] P.Blicher: "Edge Detection and Geometric Methods in Computer Vision," Ph.D. Thesis, Math. Dept., U.C. Berkeley, October 1984.

[5] M.Brady: "Computational Approaches to Image Understanding," Computing Surveys, Vo1.14, No.1, March 1982, 3-71.

[6] F.J.Canny: "Finding Edges and Lines in Images" AI-TR 720, M.I.T. A.I. Lab., June 1983.

[7] L.S.Davis: "A Survey of Edge Detection Techniques," Computer Graphics and Image Processing, Vo1.4, No.3, Sep. 1975, 248-270.

[8] J.W.Goodman: "Introduction to Fourier Optics," McGraw-Hill, Inc., New York, 1968.

[9] R.M.Haralick: "Digital Step Edges from Zero Crossing of Second Directional Derivatives," IEEE Trans. PAMI-6, No.1, Jan. 1984, 58-68.

[10] M.H.Hueckel: "An Operator which Locates Edges in Digitized Pictures," Journal of the ACM, Vo1.18, No.1, Jan. 1971, 113-125.

[11] Y.Leclerc, S.W.Zucker: "The Local Structure of Image Discontinuities in One Dimension," TR-83-19R, Computer Vision and Robotics Lab., McGill Univ., May 1984.

[12] D.C.Marr, E.Hildreth: "Theory of Edge Detection," Proc. R. Soc. Lond., I3 207, 1980, 187-217.

[13] V.S.Nalwa: "On Detecting Edges," presented at the Image Understanding Workshop, New Orleans, Louisiana, Oct. 1984.

[14] R.Nevatia, K.R.Babu: "Linear Feature Extraction and Description," Computer Graphics and Image Processing, Vol. 13, 1980, 257-269.

[15] F.O'Gorman: "Edge Detection using Walsh Functions," Artificial Intelligence, 10, 1978, 215-233.

[16] A.V.Oppenheim, R.W. Schafer: "Digital Signal Processing," Prentice-Hall Inc., Englewood Cliffs, 1975.

[17] J.M.S.Prewitt: "Object Enhancement and Extraction," in Picture Processing and Psychopictorics, B.S.Lipkin and A.Rosenfeld, Eds., Academic Press, N.Y., 1970, 75-149.

[18] K.S.Shanmugam, F.M.Dickey, J.A.Green: "An Optimal Frequency Domain Filter for Edge Detection in Digital Images," IEEE Trans. PAMI-1, No.1, Jan. 1978, 37-49.

[19] K.Turner: "Computer Perception of Curved Objects using a Television Camera," Ph.D. Thesis, A.I. Lab., Univ. of Edinburgh, Nov. 1974.