Technical Report CS79009-R

CONTEXTUAL BOUNDARY FORMATION BY SCAN LINE MATCHING

by

R.W. Ehrich

F.H. Schroeder

Department of Computer Science

Virginia Polytechnic Institute and

State University

Blacksburg, Virginia 24061

September, 1979

The work reported in this paper was supported by the National Science Foundation under Grant ENG-7712105.

Abstract

In this paper an algorithm is given for generating linked edge boundaries between adjacent regions of different gray levels. In contrast with peak following algorithms, edges are treated as variable width regions, and the edge linking procedure is really a region grower. Edge linking is a parallel process on all the edges in pairs of adjacent scan lines, and contextual information in the direction of the scan lines is used to resolve ambiguous linking situations. The procedure relies heavily upon a one-dimensional edge detector that defers the formation of local edge interpretations until more informed decisions can be made by the edge linking procedure.

1. Introduction

The need for algorithms for producing line drawings from gray level images was recognized in the early years of computer vision research, and a number were constructed. Most, like the Binford-Horn algorithms [1], were designed to function in the blocks world, and they had built into them a substantial amount of knowledge of this restricted domain. Most algorithms are multi-step algorithms consisting of an edge detection phase followed by one or more tracking or linking phases. Because scene-dependent information could be used to correct the output of the edge detector whenever the interpretations it produced were erroneous, good results could be obtained without too much concern about the behavior of the edge detector. However, the analysis of natural scenes stimulated the development of numerous new edge detectors because the more difficult problem domain required the best possible edge information that could be obtained.

This paper describes the second phase of a system designed to produce linked edge boundaries between different gray level regions of natural scenes. The algorithm is a scan line matching procedure that links edges across an image in parallel. The matching procedure optimizes the linking in ambiguous situations while forming as many contiguous links as possible. The procedure depends upon an edge detector that produces edge assertions whose

interpretations are left open until they are fixed by the linking procedure. Thus the linker does not correct erroneous edge interpretations but rather forms the interpretations on the basis of contextual information. Most important, edges are considered variable width regions rather than just points of high slope or high edge confidence; consequently, the scan line matching procedure may be considered to be a region grower.

Line construction algorithms generally fall into four categories - those that track individual edges, those that track multiple edges in parallel through adjacent scan lines, those that generate all parts of all edges in parallel, and those that generate edges as a by-product of a region growing process. Typical of the first approach are tracking algorithms such as those discussed in Rosenfeld [2], heuristic search methods such as Martelli's [3], and adaptations of dynamic programming algorithms such as that described by Montanari [4]. The Binford-Horn linefinder [1] is typical of the one-dimensional scan line approach; their procedure is run once in both the vertical and horizontal directions, and then the results are merged. Examples of parallel procedures include that of O'Gorman and Clowes [5] who use the Hough transform to locate colinear points and that of Hanson and Riseman [6] who use relaxation to do simultaneous edge linking and enhancement.

Pavlidis has used a combination of the one-dimensional approach and region growing to generate boundaries by linking together portions of scan lines that have similar slopes [7,8]. His procedure is based upon the assumption that edges can be determined by approximating a scan line with straight line segments; edges would correspond to segments with particularly high slope. While this procedure yields good results on simple high contrast images using a very modest scan line matching algorithm, the procedure does not work well in complicated natural scenes. The difficulty is characteristic of the endemic problem of most edge construction algorithms: too little attention has been given to the extraction of information from the raw image data. As a consequence, the linking algorithms are frequently given incorrect data and have insufficient knowledge to recover from previous errors. Stated more directly, the data extraction phase that produces linker input frequently violates the principle of least commitment [9].

Marr's Principle of Least Commitment states that whenever insufficient information is available for making a reliable decision, the decision must be deferred until there is a sufficient basis for making the decision. In this particular problem there are interpretations of edge-like structures in an image that are completely impossible for an edge detector to make without more global contextual

information. An edge detector must not produce invariant interpretations of image data. Instead it must generate the most viable hypotheses about the nature of the image data and store it in a convenient representation. Then the edge constructor must evaluate the alternative interpretations based upon global context. In the case of the straight line approximation approach, edge interpretations are fixed at the time the approximation is computed, and there is no possibility of changing the interpretations without simply ignoring them.

It is inevitable that a system of the type being discussed here is going to have increased complexity, both in the edge detector which must form multiple hypotheses and in the linker which must evaluate numerous alternative hypotheses in context with another. One of our reasons for adopting a one-dimensional approach is to allow us to deal directly with the very complicated semantics of edges. Another reason is that meaningful interpretations of edges are determined principally by analyzing behavior along their normals, rather than in their principal directions. In fact, many popular edge detectors are essentially one-dimensional detectors whose outputs are combined so as to make them appear two-dimensional.

2. The Edge Detector

This paper is primarily concerned with the generation of gray level boundaries and not with the formation of texture boundaries. These are two very different problems. In the first problem, the detailed shape of an edge profile is critically important and is a consequence of tonal variation, illumination changes, reflectivity changes, and the curvatures of surfaces within the scene. Texture boundaries are places where one stochastic process changes into another and where the nature of the profile in the transition region has meaning only in a statistical sense. It is our view that these distinct problems require distinct techniques; it is a mistake to attempt the solution of both problems with a single technique.

Having stated a case for least commitment, we give a brief description of the edge detector we have proposed that is consistent with the requirements. A full discussion of the edge detector appears in [10]. Fundamental to the discussion of the edge detector is the notion that edgeness is not simply an attribute of a picture element. An edge is a region of variable width, having measurable attributes, in which an intensity change is taking place. In itself this view is not new. For example, Rosenfeld [2] and Haralick [11] discuss the representation of edge regions by fixed size sloped plane segments, and Pavlidis' straight line

approximations to intensity profiles require that edges to be treated as variable width regions. Also, Hanson and Riseman [6] give a particularly lucid exposition on the necessity for treating edges as regions and the failure of most detectors to interpret such regions properly.

The house scene in Figure 1 is a complicated combination of gray level boundaries involving the house and texture boundaries involving the trees, clouds, bushes, and grass. We are concerned here with the problem of obtaining the best possible line drawings of the gray level regions, and we are

Figure 1 - House scene.

less concerned with the texture boundaries. Figure 2a shows a part of a vertical intensity profile through the circled region at the top of the left window. Notice that the shadow on the window is a plateau between the unshaded window and the shadow above the window. There is no information in the profile itself that would give any clue as to whether the edge region should be interpreted as a single high contrast edge or as two lower contrast adjacent edges. The edge detector we are using will produce both alternative interpretations and rely on the edge linker to select the proper one on the basis of context.

The edge detector functions in the following way. First the first derivative is computed using a forward difference approximation as in Figure 2b. An imaginary threshold, T, is set at the maximum possible derivative value and reduced step by step to 0. As T is reduced, two peaks, P_1 and P_2, emerge that correspond to the two component slope regions of the edge. As T is lowered further, P_1 and P_2 merge into a single macropeak that corresponds to the entire edge region. The merger of peaks of the derivative can be represented as in Figure 2d by a relational tree [12]. Each vertex of the relational tree represents a different hypothesis about the interpretation of the edge, and each internal vertex of the tree is induced by the presence of a slope minimum. Each internal vertex is labeled by that peak among its descendants that is dominant or highest, and each such

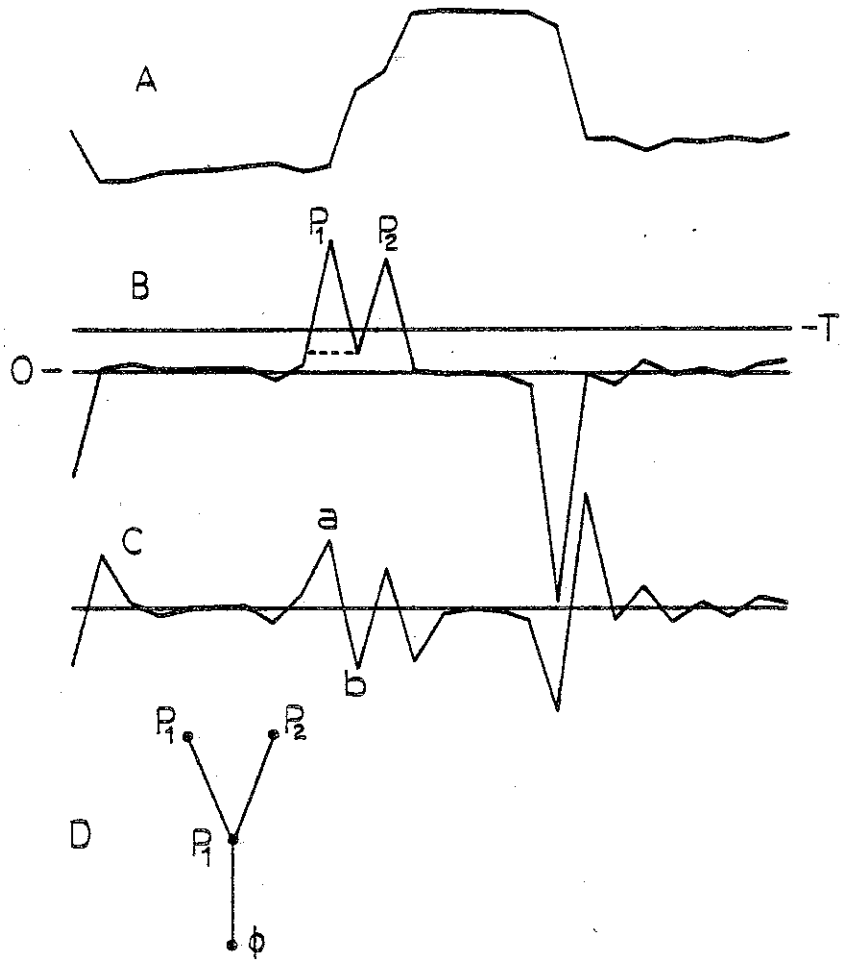

Figure 2 - Segment of vertical intensity profile from Figure 1 (a), First derivative (b), Second derivative (c), and Relational tree for edge (d).

vertex has a pointer to an attribute list that describes the edge hypothesis it represents.

The attribute list for each vertex contains all the information necessary to make a decision about the validity of the hypothesis it represents. First the wide edge region

is defined by the base of the peak in the first derivative. For the frontier peak P_1 this is the dashed line in Figure 2b, and for the internal P_1 vertex this is the entire region where the combined peaks P_1 P_2 are non-zero. Next, the maximum slope point within the wide edge region is stored because it is the best single representative point of the edge. As edge boundaries are formed, links will be made by connecting these points. Next, the edge contrast is computed by subtracting the minimum and maximum values of the intensity profile over the wide edge region, and it is stored in the attribute list for the edge hypothesis. This gives the degree of importance of the edge, and since it is a slope-independent measure, it is independent of the width of the edge region. Consequently, low slope shading edges are not discriminated against because they happen to have low slope. Finally the narrow edge region is determined by locating the extrema of the second derivative of the profile within the wide edge region and on either side of the maximum slope point. In Figure 2, either vertex labeled P has a narrow edge region delineated by extrema a and b in Figure 2c.

As a consequence of the way in which the edge detector is designed it has no parameters that need to be set before it is used. Instead it produces a set of alternative hypotheses for each edge, and these are stored in the relational tree for that edge. Normally, the relational

tree is trivial, but the structure must be provided for edges with ambiguous interpretation. It should be added that while any type of differentiator may be used in the edge detector, any one other than the one used here will produce complex interpretations of image data which violate least commitment. For example, no other differentiator will properly resolve the two distinct edges of Figure 2.

3. Scan Line Matching

In this section a scan line matching algorithm is introduced that approaches the problems associated with boundary formation in a manner very different from earlier approaches such as those of Pavlidis. In its most general form, the matching algorithm would be a parallel tree matching algorithm that matches the relational trees of the edges in adjacent scan lines. The matching algorithm described here is more modest in the sense that while alternative edge hypotheses are evaluated, they are formed by the linker itself rather than taken from the relational trees. The edge detector that produces the simple edge hypotheses for the linker is a simplified version of the detector described above. All it does is detect those edge regions where the slope is greater than some minimum slope, S_0, and the wide edge regions for these simple hypotheses are simply the coordinates between which the slope exceeds

S_o. Linking is achieved through a search procedure that determines the optimal linking arrangement based upon the strengths of the matches between the various edge hypotheses that are generated from the simple hypotheses put forth by the edge detector.

As discussed in the previous section, it is important for the linking procedure to have the ability to decide the interpretations of the edge regions in an intensity profile on the basis of context. This can be done by generating additional edge hypotheses called compound hypotheses from the simple hypotheses formed by the edge detector whenever an edge region is sufficiently ambiguous. An example of where it is important to form a compound hypothesis is shown in Figure 3. Here, region A has two possible interpretations. In the first, the minor peak in the center of region A is considered to be a true peak, and the region is composed of the three simple hypotheses ab, bc, and cd. The second hypothesis is that the minor peak in the center of region A is caused by a very local event, and the region is correctly represented by the compound hypothesis ad. All of these edge hypotheses must be considered, and the three simple hypotheses and the compound hypothesis mutually constrain each other.

The linking procedure is based upon several observations about adjacent profiles.

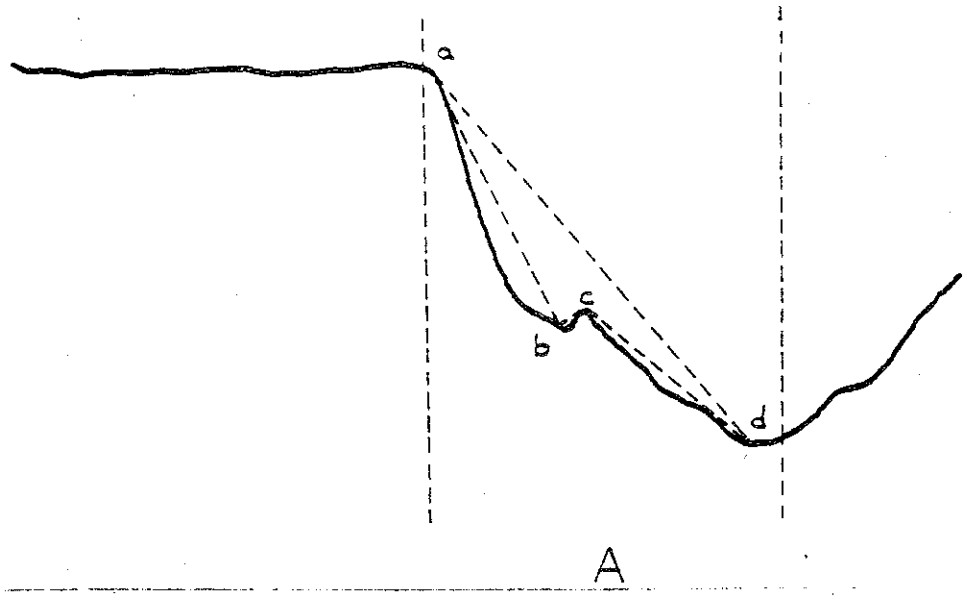

Figure 3 - Simple and compound hypotheses.

a) Smooth edge boundaries in the image that are within 45 degrees of vertical have corresponding edge hypotheses in adjacent intensity profiles that are spatially aligned in the horizontal direction.

b) Homogeneous edge boundaries in an image have properties such as contrast, slope and spatial extent that vary slowly along their lengths.

c) An edge hypothesis in one intensity profile may have no match in an adjacent intensity profile for three possible reasons: other events have altered the continuity of the edge boundary, the end of an edge has been located, or the edge turns and becomes horizontal. The linking procedure searches through the competing edge hypotheses and generates

an optimal linking arrangement between adjacent intensity profiles by considering the profiles in overlapping pairs.

After the edge detector produces the simple edge hypotheses, the linking algorithm has the following tasks:

1) Generate compound edge hypotheses where needed.
2) For each pair of adjacent intensity profiles, generate an edge pair table that gives the cost associated with each feasible edge link.
3) For each pair of adjacent intensity profiles, search for lowest cost consistent linking arrangement.
4) Postprocess the linking arrangement to increase the number of contiguous edge segments.

The most significant aspects of the edge formation process are:

a) Linking is a one-dimensional process.

b) The procedure is highly parallel, and the context across the entire image is used to link edge regions in the direction normal to the context.

c) Linking is achieved through a search that involves what is called the "strongest first" paradigm, and the procedure closely adheres to the principle of least commitment.

d) The output consists of linked directed edge segments rather than independent edge elements.

e) No a priori knowledge of the image is needed except that edges are smooth and continuous.

f) Links are not allowed to cross one another.

g) Once an edge region has been determined it may be involved in only one link to each of its adjacent intensity profiles.

3.1 - Generation of Edge Hypotheses

The generation of edge hypotheses is a very important part of processing edge regions. The edge hypotheses include the simple hypotheses that were produced by the edge detector and the compound edge hypotheses that are formed by the linker. The generation of compound edge hypotheses is based upon the spatial relationships between adjacent edge hypotheses. If the relative positioning of adjacent simple hypotheses implies that they could have originated from the same edge region, a compound edge hypothesis is formed that represents the combination of those two simple hypotheses. This merging procedure is based upon the differences of the end point values for the edge hypotheses as shown in Figures 4a and 4b. Let the two sets of end points $(X_{a1}, Y_{a1}), (X_{a2}, Y_{a2})$ and $(X_{b1}, Y_{b1}), (X_{b2}, Y_{b2})$ represent the edge hypotheses in both cases.

Then,

$$d_1 = |Y_{b1} - Y_{a2}|$$

$$d_2 = |X_{b1} - X_{a2}|$$

$$d_3 = |Y_{a2} - Y_{b2}|$$

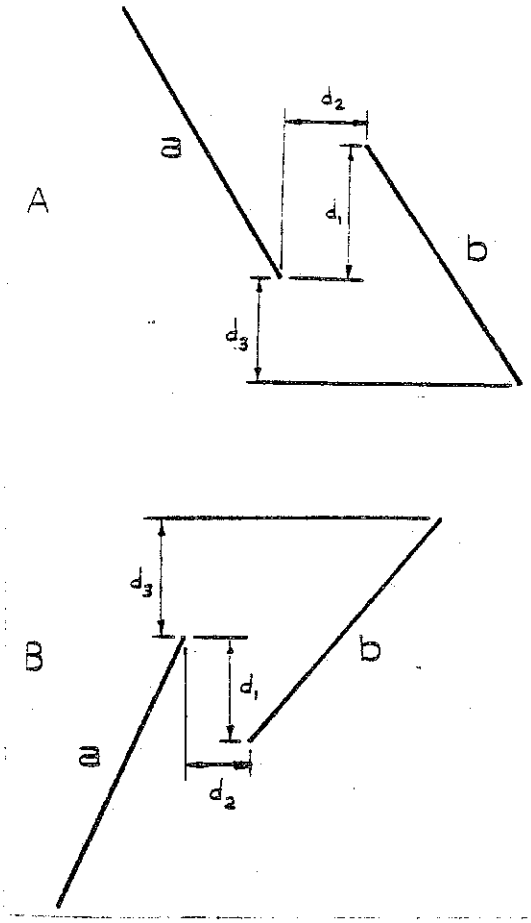

Figure 4 - Formation of compound edge hypotheses.

The two edge hypotheses are joined together if and only if:

$$d_1 < D_1$$

$$d_2 < D_2$$

$$d_3 > D_3$$

for some threshold values, D_1, D_2, and D_3. If the two hypotheses are joined, the resulting edge hypothesis is defined by the end points $\{(X_{a1}, Y_{a1}), (X_{b2}, Y_{b2})\}$. In Figure

5, two adjacent intensity profiles from an image are shown together with the compound hypotheses generated using $D_1 = 25$, $D_2 = 3$, $D_3 = 2$.

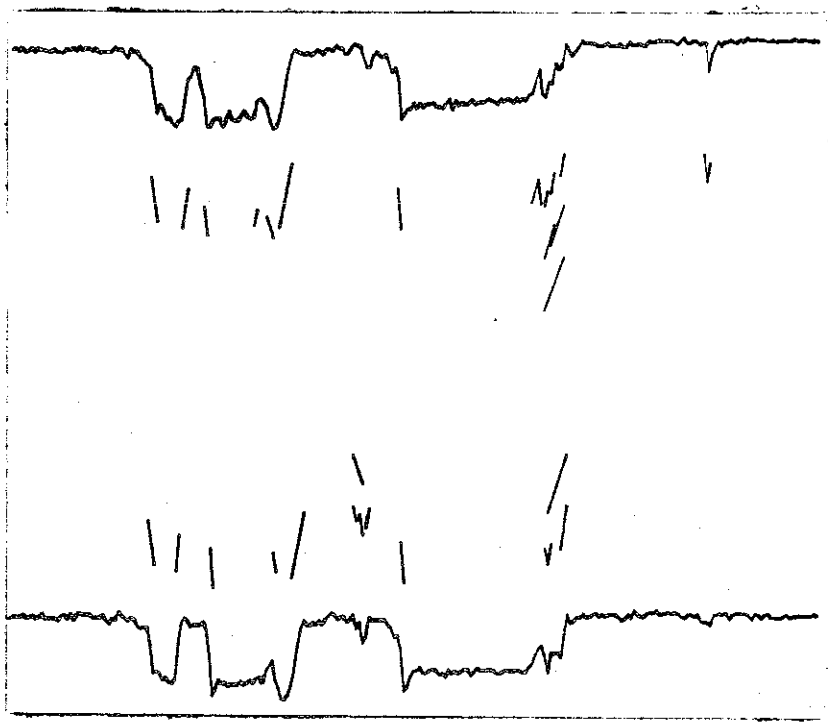

Figure 5 - Two adjacent profiles from an image and all simple and compound edge hypotheses.

It can be seen that the parsing of intensity profiles into hypotheses that correspond to the most significant interpretations of the edge regions can be achieved by a relatively simple procedure. Unlike earlier methods, the edge construction algorithm presented here is not based upon

approximation methods and is more sensitive to the structure of the intensity profile.

3.2 - Generation of the Edge Pair Table

As we have seen, the edges of an image that are nearly vertical have edge regions in adjacent intensity profiles that have similar slopes and spatial translation. There should then exist a corresponding pair of edge hypotheses that are similar. If there is more than one edge hypothesis for an edge region, a decision must be made as to which one should be involved in a link to an adjacent intensity profile.

A simple method for making this decision is to link the most similar pair of edge hypotheses. As seen in Figure 6, each edge region in two adjacent intensity profiles is represented by three hypotheses a,b,c and d,e,f, respectively, and there are eight possible links ad, ae, af, be, bf, cd, ce, and cf. On the basis of similarity alone the link cf would be formed. Unfortunately, this simple linking procedure is not always correct. Figure 7 shows two adjacent intensity profiles across an edge boundary oriented at 45 degrees. It is obvious that the correct set of links is ad, be, and cf, but because of the minor differences in the intensities, the first link to be formed on the basis of similarity would be cd. This link would then constrain the others since links are not allowed to cross. Therefore,

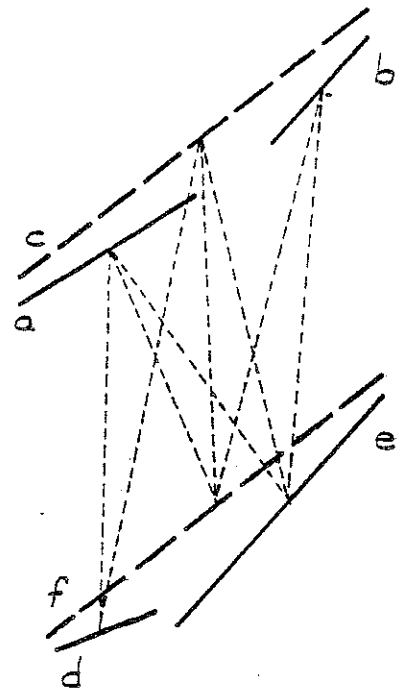

Figure 6 - Edge hypotheses in adjacent intensity profiles.

the linking process cannot be based solely upon the similarity of edge hypotheses. It must also consider the global effect that each link will have on the entire linking arrangement. Because of this, the linking procedure must involve a search of some kind that includes the "strongest first" paradigm.

The "strongest first" paradigm states that the strongest links should be made first unless they constrain links with similar strength. The strongest first paradigm is closely related to least commitment because it is undesirable to match weak hypotheses first and allow them to constrain

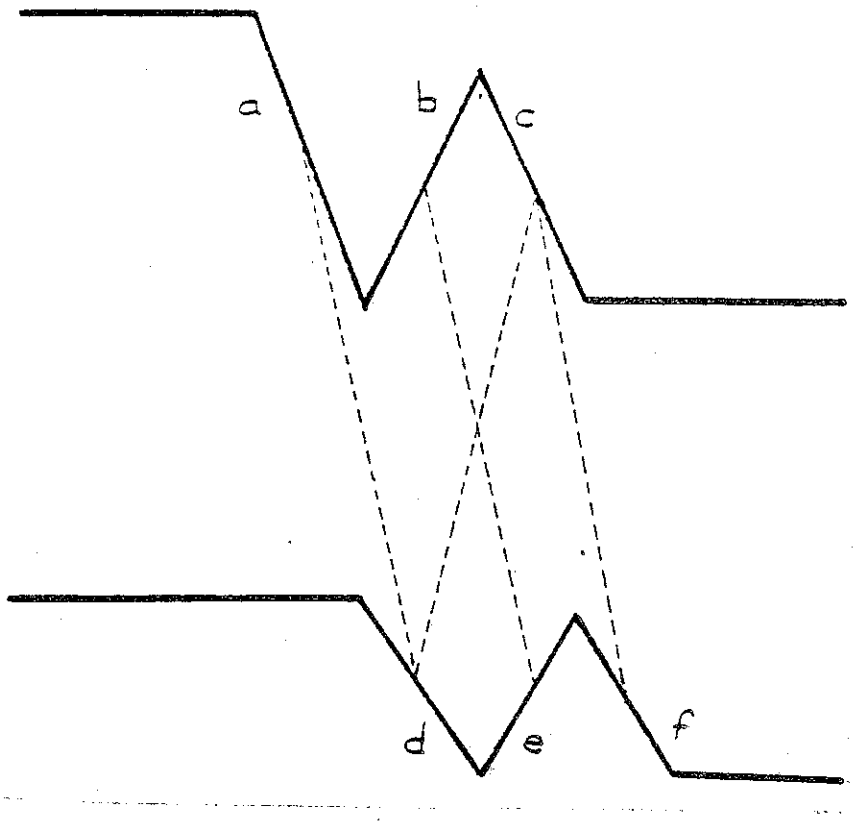

Figure 7 - Possible edge links between adjacent intensity profiles.

other links about which one is more confident. The strength of a link is based upon the similarity between the edge hypotheses to be linked. Since we wish to define a cost measure that decreases in proportion to the similarity between the candidate hypotheses, similarity is defined on the basis of endpoint proximities. In cases where two or more links with similar strengths constrain one another as in Figure 7, the decision for linking must be based upon other criterea. Before linking begins, an edge pair table

is formed that contains the costs for each feasible link between a pair of edge hypotheses in adjacent intensity profiles.

The cost of forming an edge link is a combination of the match between the two hypotheses and the strengths of the individual hypotheses. Roughly, the cost is the ratio

$$\text{cost} = \frac{\text{Differences between edge hypotheses}}{\text{Strength of individual edge hypotheses}}$$

The cost value is determined by the differences in the end points of the edge hypotheses as shown in Figure 8. Let the two sets of end points $\{(X_{a1}, Y_{a1}), (X_{a2}, Y_{a2})\}$ and $\{(X_{b1}, Y_{b1}), (X_{b2}, Y_{b2})\}$ represent the end points of pairs of edge hypotheses from adjacent intensity profiles. Then

$DV_1 = |Y_{a1} - Y_{b1}|$

$DV_2 = |Y_{b2} - Y_{a2}|$

$DH_1 = |X_{a1} - X_{b1}|$

$DH_2 = |X_{b2} - X_{a2}|$

S_1 = contrast of edge hypothesis a

S_2 = contrast of edge hypothesis b

The edge contrasts are determined by the edge detector and stored in the attribute lists. Given these difference values, the cost of the link is given by

$$\text{cost} = \frac{DV_1 + DV_2 + DH_1^3 + DH_2^3}{S_1 + S_2}$$

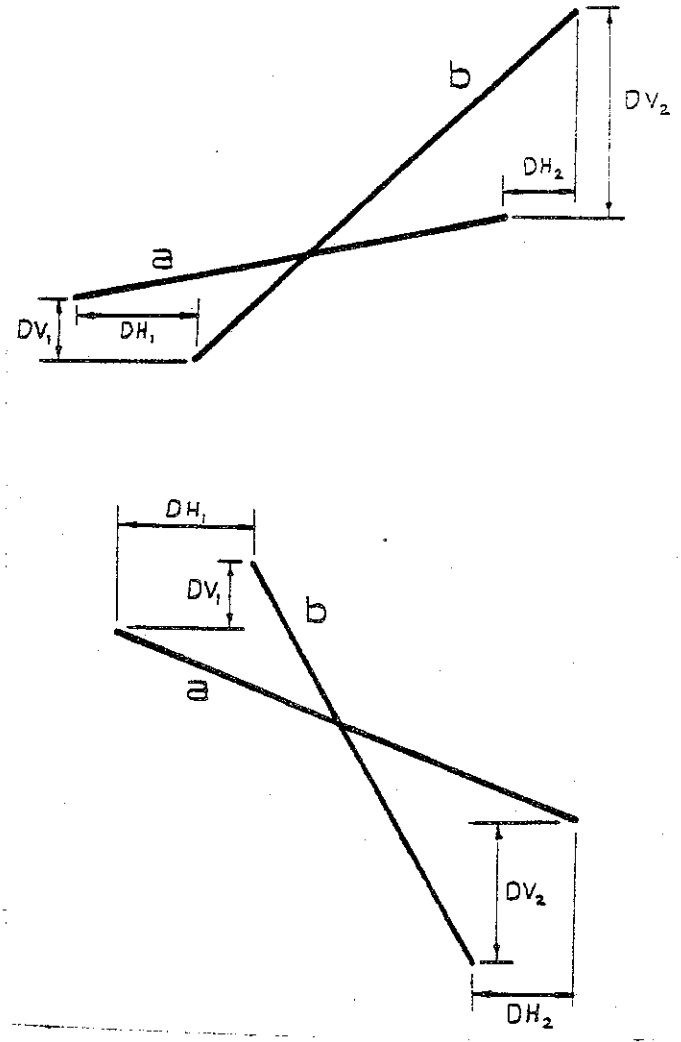

Figure 8 - Pairs of superimposed edge hypotheses from adjacent intensity profiles.

The two values DH_1 and DH_2 have the largest influence in the cost function since they represent the differences in horizontal placement which is very important in the matching procedure. The cost function was made inversely proportional to the contrasts of the edge hypotheses since a link between strong edges of given similarity should have lower cost than a link between weaker edges with the same similarity.

Given two sets of edge hypotheses, not every pairing is entered into the cost table. Since edge region pairs in the intensity profiles are spatially aligned in the horizontal direction and have similar slopes, it is unreasonable to enter pairings that are drastically different. Therefore, the following contraints must be met before a pair is entered. First, edge hypotheses must overlap each other horizontally or have a common X coordinate in their end points. Second, their slopes must have the same sign. Third, given that the first two contraints have been met, the cost of the link must be below a preset value. This value is the same for all images and was determined by examining the costs of linking edge hypotheses with varying similarities and contrasts. The value chosen was 1.5.

3.3 - The Linking Procedure

The first step in the development of the linking procedure was the determination of the search cost function. The cost function must allow as many links to be made as possible so long as the total cost for the linking arrangement remains reasonable. If this was not done, the optimal linking arrangement would be the trivial one where no links are made. The search is to find that linking arrangement that maximizes the total number of links subject to the condition that no link cost exceeds a preset threshold t.

Let P_i and P_{i+1} be two adjacent intensity profiles with respective sets of edge hypotheses H_i and H_{i+1} where $H_i = \{e_1, e_2, \ldots, e_n\}$ and $H_{i+1} = \{f_1, f_2, \ldots, f_m\}$. Then $H_i X H_{i+1}$ is a set that contains all possible links between P_i and P_{i+1}.

Let $c(e_j f_k)$, $e_j f_k \in (H_i X H_{i+1})$ be the cost of linking e_j to f_k.

Then the *minimal* cost *arrangement* is given by

$$\min_h C = \sum_{j,k \in K_h} c(e_j f_k) + \frac{t}{2}(N - 2K_h)$$

where h = total number of feasible linking arrangements

K_h = set of links in a given linking arrangement h

$N = |H_i| + |H_{i+1}|$

t = maximum cost link to be considered

Paths are followed in the search tree until the cost of the lowest cost unused link exceeds a threshold t. The longest path in the search tree then represents the linking arrangement to be used. If the longest path is not unique, the one with the lowest cost is used.

Since the cost depends upon the entire search path and, in fact, decreases as a goal node is approached in the search tree, it appears that a full depth first search is required to determine the optimal path. This does not cause

as many difficulties as it appears. Consider Figure 9 which shows two separate edge regions A and B in adjacent

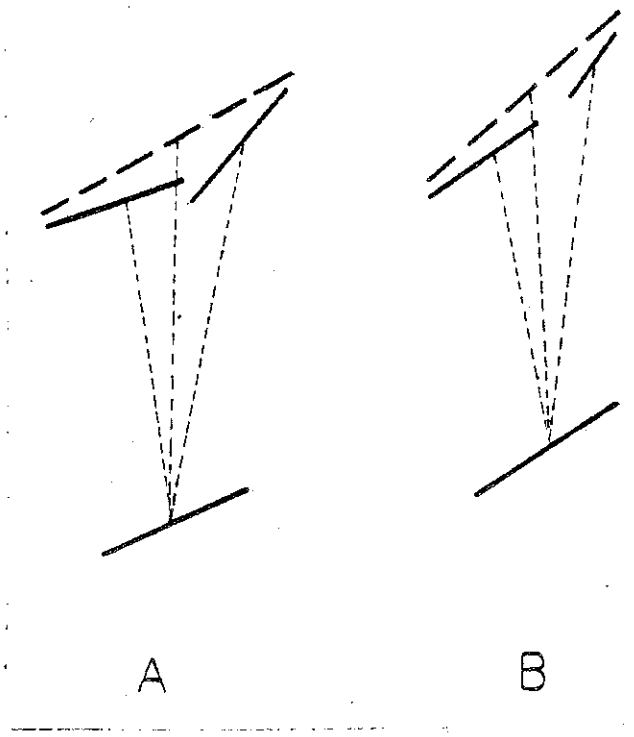

Figure 9 - Two separate edge regions in adjacent intensity profiles.

intensity profiles. It is clear that the links in one region do not interfere with the links of the other. Therefore, the search should consider these regions separately. In general, the search is then a series of local searches which drastically reduces the size of the search tree.

The size of the search tree can be reduced again by using the "strongest first" paradigm. Let $l_1, l_2, \ldots, l_n$ be a set of links over which a local search is to take place. If

link l_i is the lowest cost link in this set and its cost is much less than those of the links it constrains, then link l_i may be included in every linking arrangement. If l_i interferes with a link that has a similar cost then both possibilities must be kept open.

3.4 - Edge Postprocessing

Next, we have implemented a postprocessing phase that generates additional links in cases where there are no ambiguities. If no compound edge hypotheses were involved in the linking procedure, then the edge linking process for a pair of adjacent intensity profiles would be complete. If compound edge hypotheses are involved, then postprocessing can be done to increase the number of links.

As seen in Figures 10a and 10b, the same edge regions are present in three adjacent profiles although the edge boundary is not continuous because linking is done between pairs of profiles. Cases such as this can be located and corrected very easily at this stage of the processing. Figures 11a and 11b show the results after forcing a link through the middle edge regions.

This type of processing is not always possible. As seen in Figures 12a, 12b, 12c, and 12d there exist ambiguous cases. At this level of processing no further linking decisions can be made. Further linking must be left to a

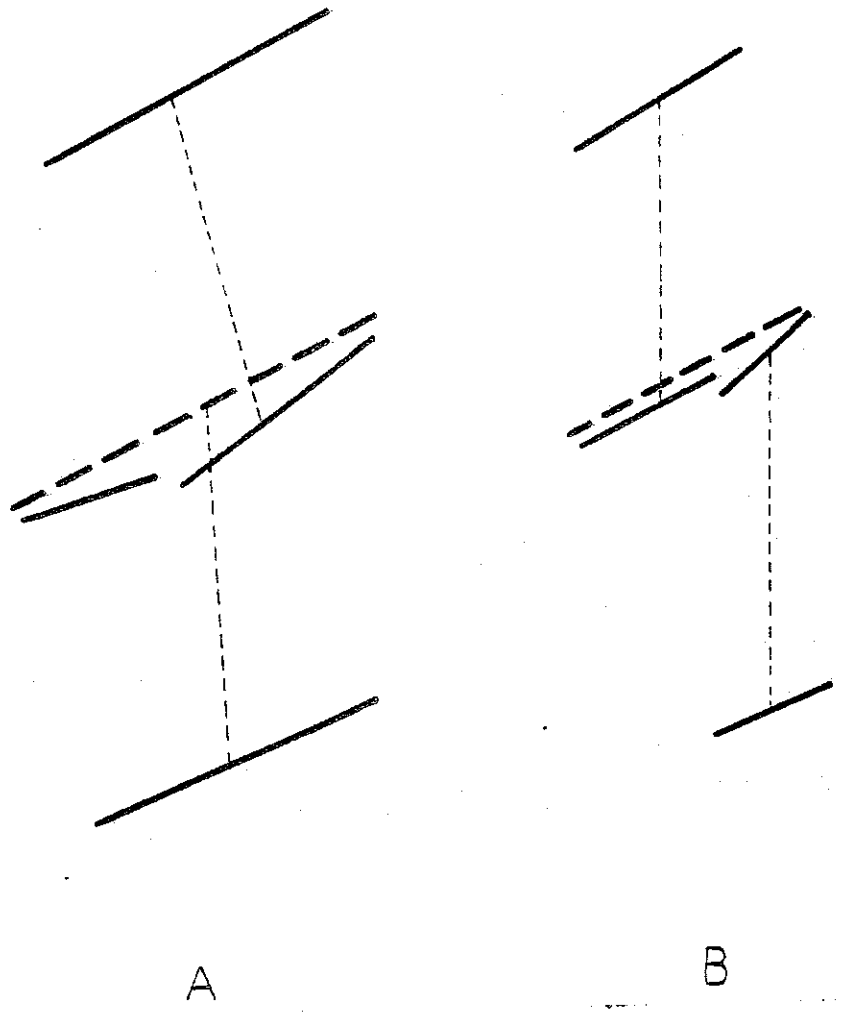

Figure 10 - Two edge regions in three adjacent intensity profiles.

process that can examine more global information or to a process that has a priori knowledge of the image.

4. Results

In this section, the results of the edge linking procedure are presented and are compared to the unlinked output of the Roberts Cross and Sobel operators.

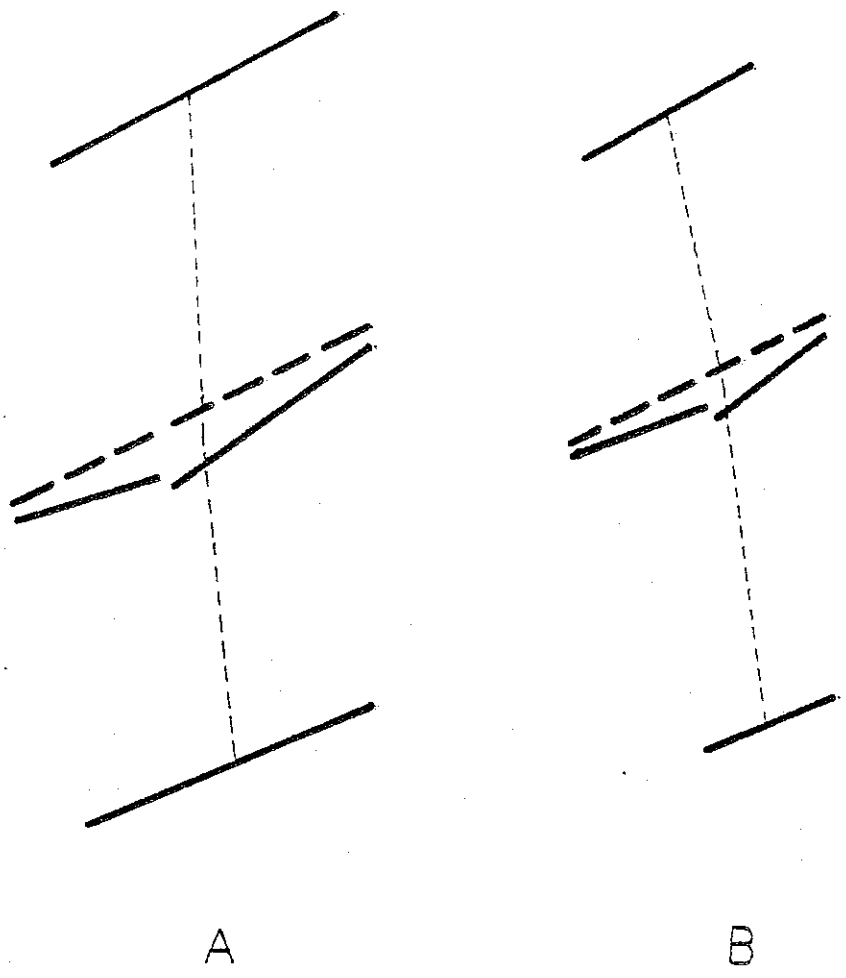

Figure 11 - Postprocessed links in three adjacent intensity profiles.

Figure 1 shows the natural scene containing a house, trees, grass, and sky from which the first set of results were produced. Edges are very difficult to extract from an image of this type because of the large amount of texture present. Figure 13 shows the thresholded results of the Roberts Cross and Sobel operators on Figure 1. Each operator produces a slightly different edge output, and each edge image has both good and bad aspects. The Sobel

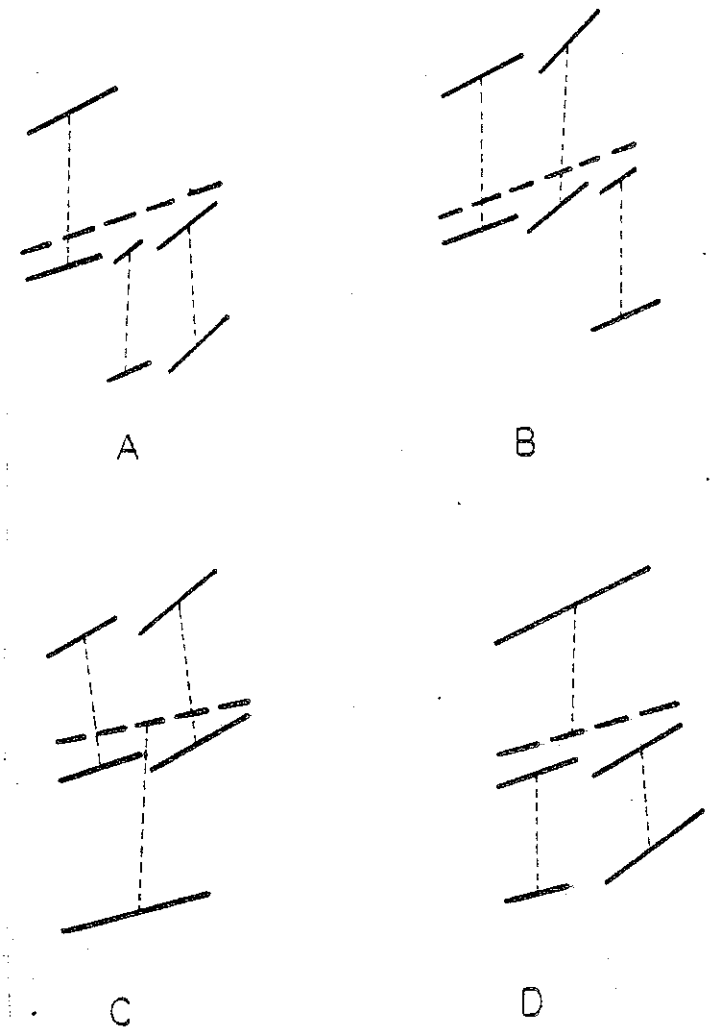

Figure 12 - Ambiguous links in three adjacent intensity profiles.

operator appears to produce better results for areas A and B of Figure 13 while the Roberts Cross operator appears to be superior for areas C and D. However, neither does very well on the obscured corner of the left set of windows in area E. The differences between the operators are due, of course, to the fact that the Sobel operator is a larger operator. Therefore it will do a better job on texture boundaries but a much poorer job on fine detail.

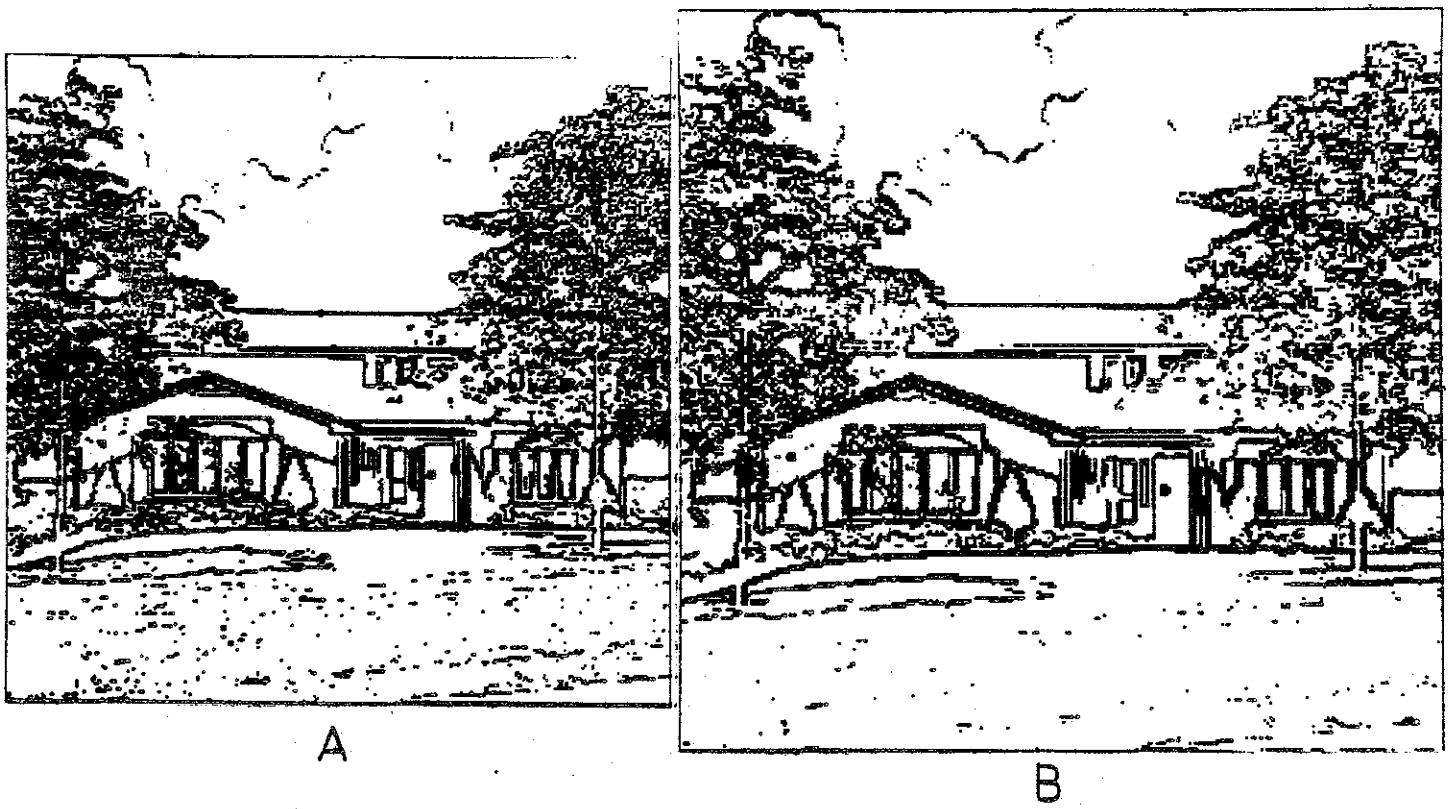

Figure 13 - Roberts Cross (a) and Sobel operator (b) output from Figure 1, thresholded at 45.

Figure 14 is an overlay of vertical and horizontal linked edges produced by the linker from Figure 1. In all there are over 14,000 edge elements and over 3,700 lines. However, the line drawing can be simplified dramatically by discarding the short edges. Figure 15 shows the results of eliminating lines containing fewer than 3 or 10 edge elements. Notice that in Figure 15b a significant portion of the house remains even though only 293 lines are required. Notice also the larger branches of the trees that are difficult to observe when shorter lines are included.

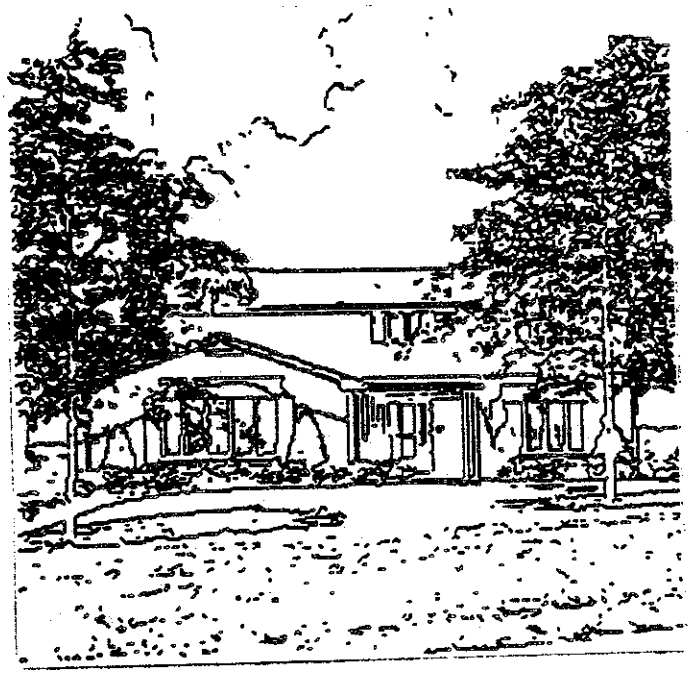

Figure 14 - Overlayed linked edges from linking procedure applied in vertical and horizontal directions.

Figure 16 shows enlargements of four areas of special interest in the house scene of Figure 1. All lines of length greater than 2 have been shown. In Figure 16a, the procedure generates very accurate and complete edges even in the area where the window is obscured by the tree. It also locates the boundaries caused by shadows at the top of the window and near the peak of the roof. In comparison, the Roberts Cross generated ambiguous patches of strong edge points between the windows, and the Sobel operator generated unaligned points for the shadow near the peak of the roof.

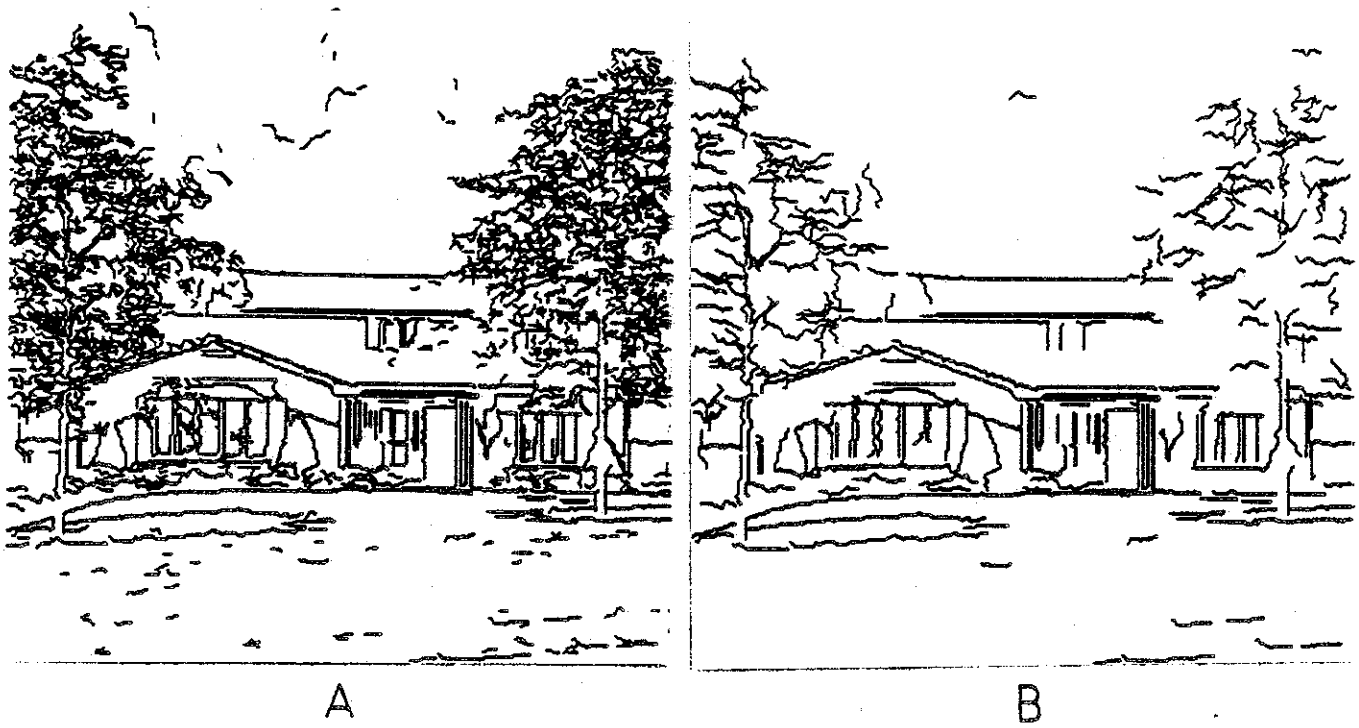

Figure 15 - Same as Figure 14 except lines shorter than 3 edgels removed (a) and lines shorter than 10 edgels removed (b).

As shown in Figure 16b, the edge linking procedure locates most of the detail near the door of the house scene. The Sobel operator appears to generate a more complete representation for the first window to left of the door, but it fails to separate the individual pillars to its right and the windows to the far left of the door. Figure 16c shows a window that has been obscured by tree branches. Both the Sobel and Roberts Cross operators generate ambiguous connections between the edges of the windows and the tree while the edge linking procedure tends to segment the window from the background. Finally, Figure 16d shows that the

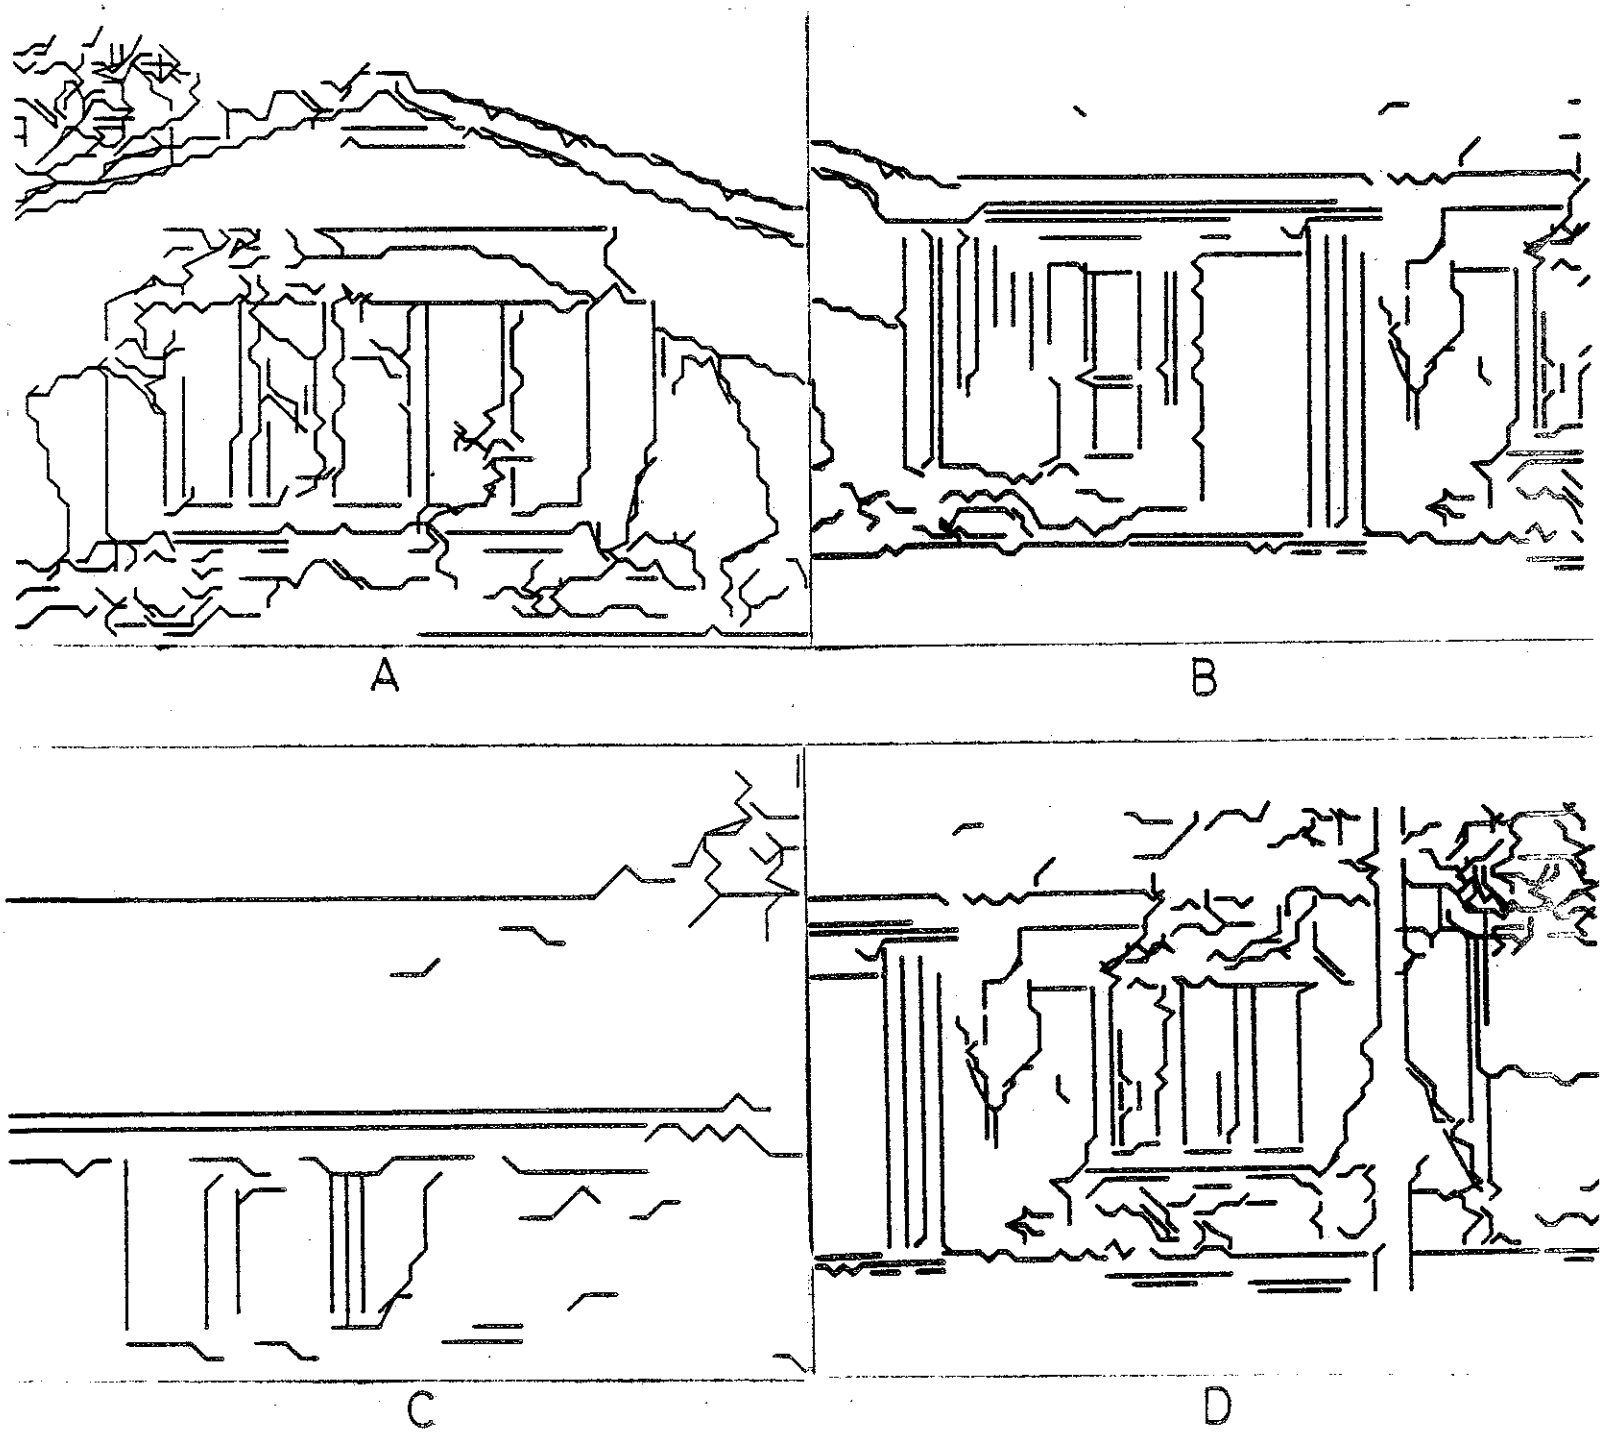

Figure 16 - Enlargements of left window (a), door (b), right window (c), upper window (d) showing edges of length 3 or more.

detail of the upstairs window is much better in the line drawing generated by the edge linking procedure, particularly for the middle window. The Roberts Cross

operator tends to fragment the portion obscured by the tree into ambiguous patches while the edge linker generates a smooth boundary.

Continuing with some other examples, Figure 17 shows a high contrast industrial scene with a noisy background. Figure 18a shows the Roberts Cross thresholded at 35, and Figure 18b shows the lines produced by the edge linker having length at least 10 edgels. Notice the smoothness of the edges. The double edges are caused by shadows, and these edges can be distinguished since they have opposite polarities.

Finally, Figure 19 shows another very complicated industrial scene with occluded objects and complicated surface texture. Figure 20a shows the Roberts Cross with threshold 50, and Figure 20b shows the lines produced by the linker with length at least 5 edgels. Notice how clearly regions A, B, and C are defined. The parallel edges of the capacitor leads can not be confused, since they have opposite polarities.

In general, the edge linking procedure generates very good results for most images. In regions of low contrast the edges tend to become fragmented, but this is a problem encountered by all other edge detectors as well. The linking procedure does extremely well in regions where ambiguous edge points are generated by local operators.

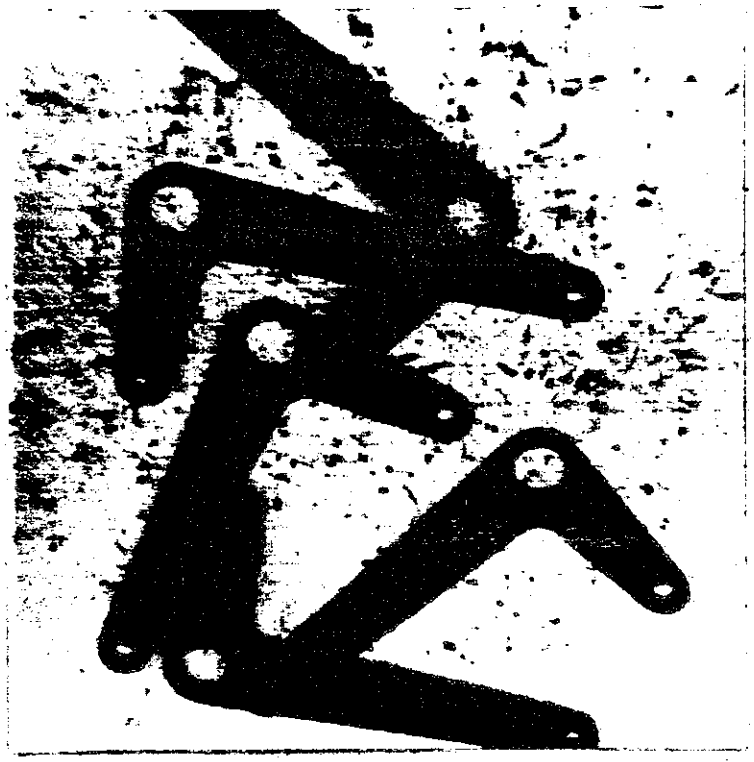

Figure 17 - High contrast industrial scene.

In all the experimental results, it must be emphasized again that even though the linker output may appear dense, the edge elements are linked together, and the edge polarities are known. Hence, vastly more information is available than appears in these images.

5. Conclusions

By treating edges as regions and by observing least commitment, good line drawings can be produced from gray level images. There are a number of cases where the

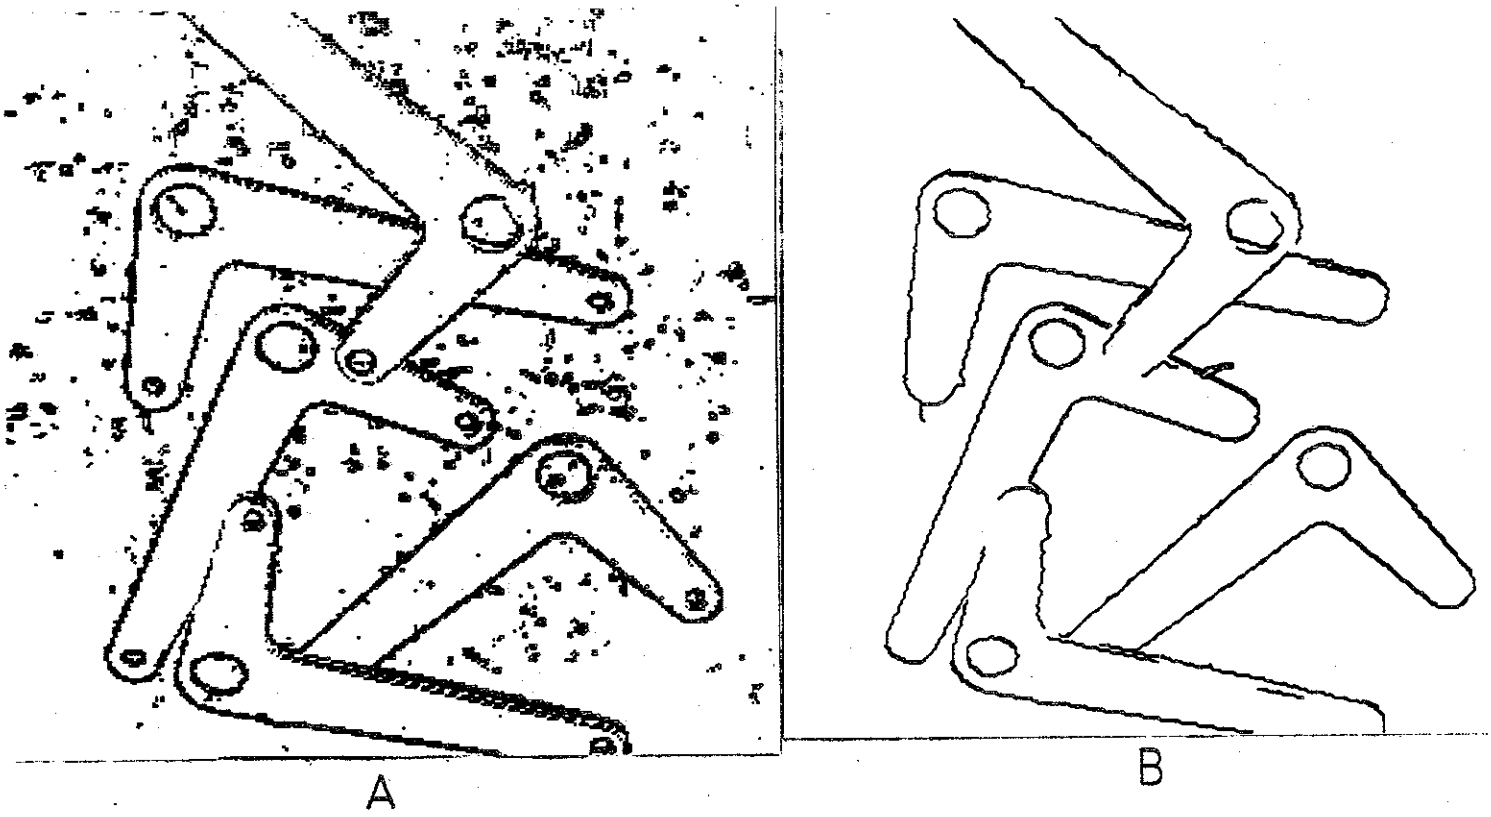

Figure 18 - Results of processing Figure 17. Roberts Cross thresholded at 35(a) and lines produced by linker with length at least 10 edgels (b).

procedure does not generate contiguous lines; however, these are ususally natural texture boundaries that do not meet the initial requirement of smooth and continuous edges. Texture boundaries such as these must be dealt with by algorithms designed explicitly for this purpose.

Further work in this area could involve not only matching pairs of adjacent profiles but also larger sets of adjacent profiles. This would provide much more information for the linking procedure, and it would eliminate many of

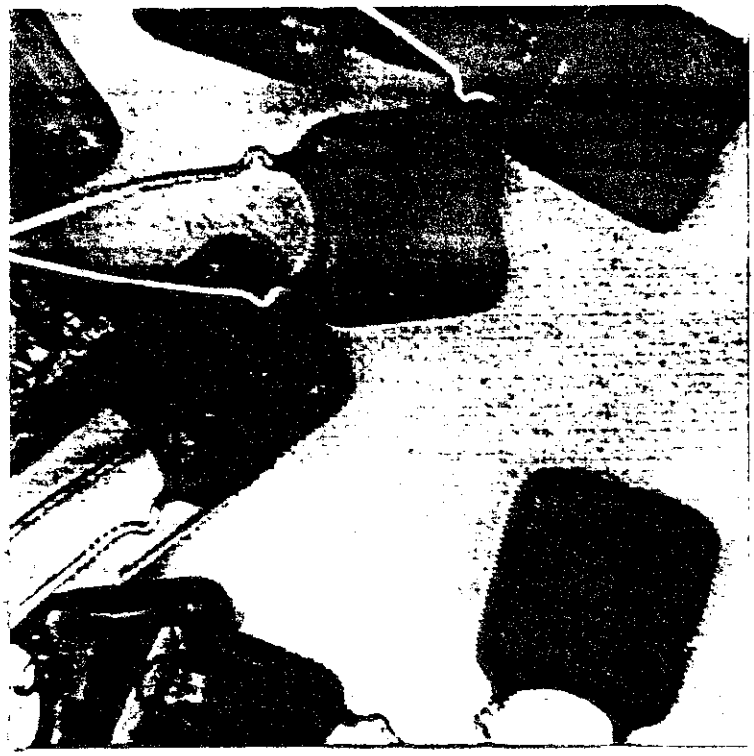

Figure 19 - Complex industrial scene.

the ambiguous cases encountered by the postprocessing procedure.

The edge formation procedure presented in this paper does not produce adequate results for textural boundaries since it was not designed for this purpose. Also, the problem of matching the full relational trees for edges has not been solved. This does not imply that the concepts are inappropriate for textural boundary formation; rather, the current implementation is inadequate. We believe that

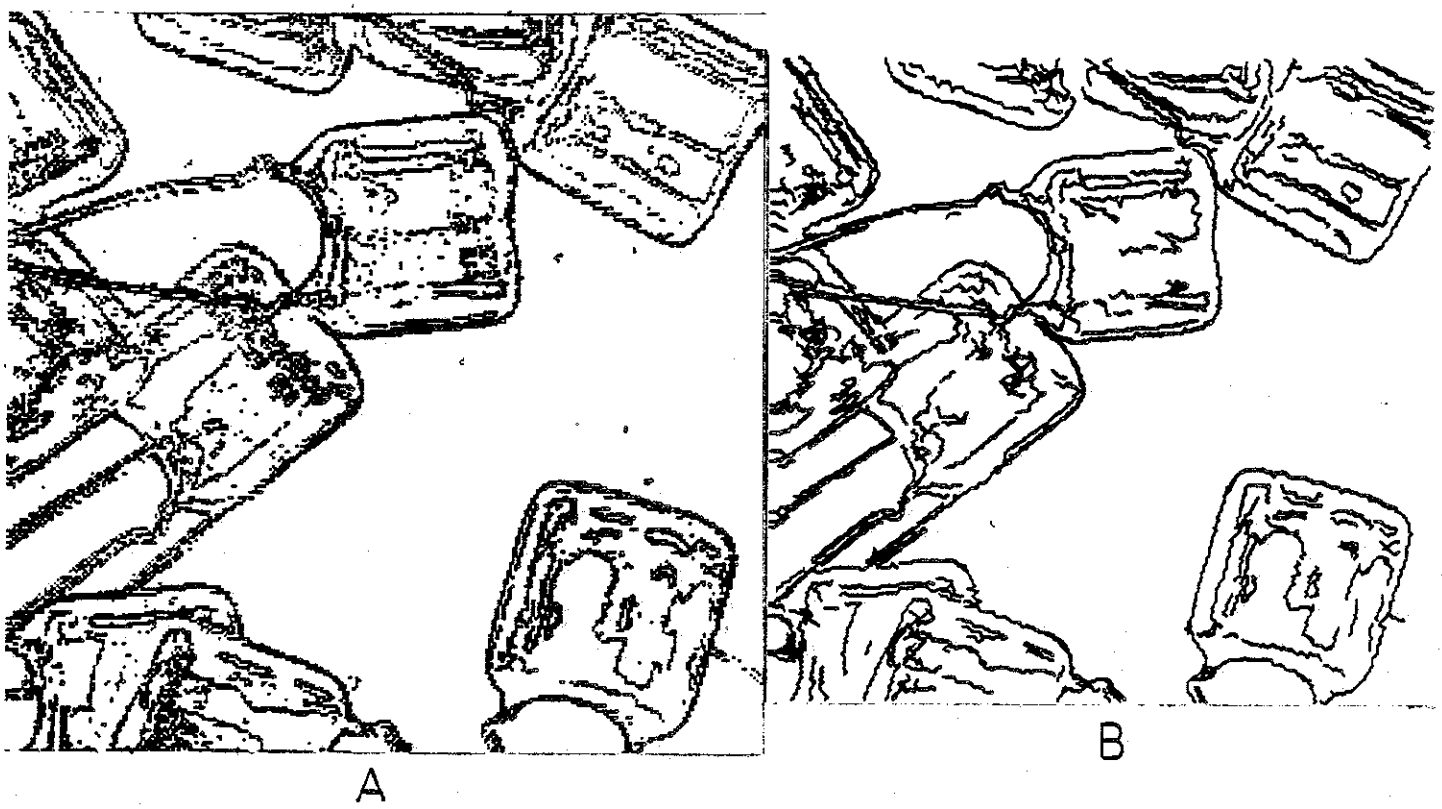

Figure 20 - Results of processing Figure 19. Roberts Cross thresholded at 50 (a) and lines produced by linker with length at least 5 edgels (b).

generalized region matching algorithms will be very powerful tools for region growing as well as for boundary formation.

The next step in the processing of these raw line drawings is a phase in which regions are formed, subjective contours are completed, and unwanted objects are removed. For this purpose we are using a special data base system for manipulating line drawings. In this system information such as polarity, length, and contrast of a line is retained so that subsequent processing steps will have the necessary

information. The processing of the raw line drawings is a sophisticated process that will be the subject of a later paper.

References

1. Horn, B.K.P., "The Binford-Horn linefinder," Artificial Intelligence Memo 285, MIT AI Laboratory, Cambridge, March 1973.

2. Rosenfeld, A. and A.C. Kak, Digital Picture Processing, Academic Press, New York, 1976.

3. Martelli, A., "Edge detection using heuristic search methods," Computer Graphics and Image Processing 1, August 1972, pp. 169-182.

4. Montanari, V., "On the optimum detections of curves in noisy pictures," Comm. ACM 14, 1971, pp. 335-345.

5. O'Gorman, F. and M. Clowes, "Finding picture edges through collinearity of feature points," Proc. 3rd JCAI, August 1973, pp. 543-555.

6. Hanson, A.R. and E.M. Riseman, "Segmentation of natural scenes,"in Computer Vision Systems, A.R. Hanson and E.M. Riseman, eds., Academic press, 1978.

7. Pavlidis, T., "A Minimum storage boundary tracing algorithm and its application to automatic inspection," IEEE Transactions on Systems, man, and Cybernetics 8, January 1978, pp. 66-69.

8. Feng, H.Y. and T. Pavlidis, "The generation of polygonal outlines of objects from gray level pictures," IEEE Transaction on Circuits and Systems 22, May 1975, pp. 427-439.

9. Marr, D., "Early processing of visual information," Phil. Trans of the Royal Society of London 275, Biological Sciences, 19 October 1976, pp. 483-519.

10. Ehrich, R.W., "One dimensional edge detection and representation," Technical Report CS79009-R, CS Department, Virginia Polytechnic Institute and State University, September 1979.

11. Haralick, R.M. and L. Watson, "A facet model for image data," Proc. Conf. on Pattern Recognition and Image Processing, August 1979, pp. 489-497.

12. Ehrich, R.W. and J.P. Foith, "Representation of random waveforms by relational trees," IEEE Transactions on Computers 25, July 1976, pp. 725-736.

Table of Captions

Figure 19 - Complex industrial scene.

Figure 20 - Results of processing Figure 19. Roberts Cross thresholded at 50 (a) and lines produced by linker with length at least 5 edgels (b).