# EDGE RELAXATION AND BOUNDARY CONTINUITY*

Allen R. Hanson
Edward M. Riseman
Frank C. Glazer

COINS Technical Report 80-11

May 1980

This paper will appear in Consistent Labeling Problems in Pattern Recognition, Robert M. Haralick (Editor), Plenum Press, 1980.

*This research was supported by the Office of Naval Research under Grant N00014-75-C-0459 and by the National Science Foundation under Grant MCS79-18209.

## ABSTRACT

Many image analysis tasks require the construction of a boundary representation as a means of partitioning an image. This paper develops a parallel relaxation algorithm for updating initial heuristic estimates of the likelihood of edges so that continuous boundaries are formed. Bayesian probability theory is used to analyze the probability updating of a single edge based upon the joint probabilities of the edges in its local surrounding context. The relationships between edges, sometimes referred to as compatibility coefficients in relaxation algorithms, are embodied as conditional probabilities between the central edge and the context of edges. The set of conditional probabilities are theoretically derived from a model of desired line drawings that satisfy basic notions of boundary continuity. The local updating function attempts to drive the likelihood of each central edge into consistency with the surrounding context.

Experiments involving the iterative parallel application of this non-linear Bayesian updating function to all edge probabilities demonstrates serious problems in the formulation. A variety of heuristic modifications, guided by theoretical considerations, are examined empirically. The final formulation is an algorithm that performs in an effective manner on several very complex images.

TABLE OF CONTENTS

# I. INTRODUCTION

Many image processing tasks require construction of a two-dimensional boundary representation of an image. The sensory data is usually presented as an array of feature values. Each array element, referred to as a "pixel" (for "picture element"), is characterized by its spatial location within the array and by the values of the sensory features at that point. For our purposes it will be assumed that a pixel value represents a measure of the amount of light which lands on a small area of some imaging device. Typically, local edge operators [DAV75, ROS76, RIS77] are used to provide estimates of the local strength of edges throughout the image, but the results produced are almost always unreliable. The raw edge information requires far more structural organization in order to be useful.

In this paper we seek methods for partitioning the image by focussing on local edges between pixels and on the means by which they may be organized into continuous boundaries. The primitive element of a boundary is an edge element. A set of connected edge elements, located relative to the array of sensory data from which it was derived, forms a boundary. We sometimes will use the term "line" synonymously with "boundary" because the result of this form of segmentation is a "line drawing". The approach

described here addresses the problem of how to make use of local contextual information to reduce local edge ambiguity while producing a global organization of continuous boundaries. The initial array represents sensory data while the final array represents a boundary continuity interpretation of that data -- or controlled hallucination in imagining a line drawing consistent with the sensory data!

## I.1 Review of Related Literature

Methods vary for dealing with the uncertainty which remains after local edge operators have been applied. Marr [MARR76] uses a set of edge detectors of varying size to determine the appropriate width of an edge to be asserted in his primal sketch. Ehrich [EHR79a] has devised an edge detector whose output is not just a local edge assertion, but a data structure containing alternative interpretations of edge cross-sections.

Sequential edge tracking methods have benefited from current techniques for dealing with uncertain or noisy data. Montonari has embedded the properties of a curve in a figure of merit [MON71]. A dynamic programming technique is used to determine the optimal curve with respect to the given figure of merit.

Martelli [MART76] has reformulated this method as a graph search problem which finds a contour using an optimizing heuristic search. Cooper models the sequence of edge elements in a contour as a Markov process and uses maximum likelihood estimation to generate the contour [COO79]. Fischler et al. [FIS79] apply the optimization paradigm to the specific problem of detecting roads and linear structures in aerial imagery. They also introduce mechanisms for combining local evidence and constraints into the global optimization process.

Recently there has been a large amount of research in local iterative processes called relaxation labelling problems [ROS76, ZUC78]. The fundamental idea of relaxation labelling is to iteratively update the likelihood of local interpretations based upon expected local relationships between these interpretations. At each iteration the influence of a context of nodes will spread indirectly through the influence of the nodes which it affected. Ideally, the entire system will globally organize to be locally consistent, and possibly even display global characteristics which would have been very difficult to achieve as a global function specified directly in terms of all the data.

Current approaches to relaxation labelling have employed heuristic mechanisms for updating likelihoods. Much of this work is based on the formalization provided in [ROS76], referred to as the non-linear probabilistic rule. In these systems local contextual information is embodied in compatibility coefficients which parameterize the mutual effects of adjacent labellings. Some possible interpretations of these coefficients are presented in [PEL78] including measures of correlations and mutual information. Nevertheless, these parameters are still heuristic. Zucker has argued for an interpretation of these compatibility coefficients as measures of dependency [ZUC78]. He has reformulated the relaxation updating rule mentioned above into one involving conditional probabilities. While intended as an analysis of non-linear probabilistic updating, this formulation is similar to the most successful of the updating rules we present here. Peleg has also developed a probabilistic relaxation labelling scheme which uses conditional probabilities [PEL79]. This work shares our intent of providing an analytical basis for relaxation.

## I.2 Complexity of Segmentation

With relatively unconstrained natural images, any approach to segmentation which relies entirely on the sensory data, will be prone to error. The complexity of the data is a function of a variety of factors including the presence of highly textured objects, shadows and highlights on smooth or irregular surfaces, variation in surface reflectance, varied and uncontrolled lighting, and noise introduced in the digitization process. Few objects or surfaces in an image can be expected to exhibit uniform visual features. In the images considered here, such as the ones shown in Figure 1, it is very unclear exactly where some of the boundaries should be placed in relation to the local array of feature values.

Even human hand-drawn segmentations are inevitably prone to errors and tend to reflect implicit biases and explicit goals of the human perceiving the image. In many instances the boundaries would be conjectured based on prior expectations in the form of knowledge of objects and their shape, shadow effects, perspective and occlusion cues, etc. [HAN78b]. In short, it is generally accepted that a truly accurate segmentation is a function of the goals of the system and requires the application of "high-level" knowledge which is not directly available as measurable features

1(a)

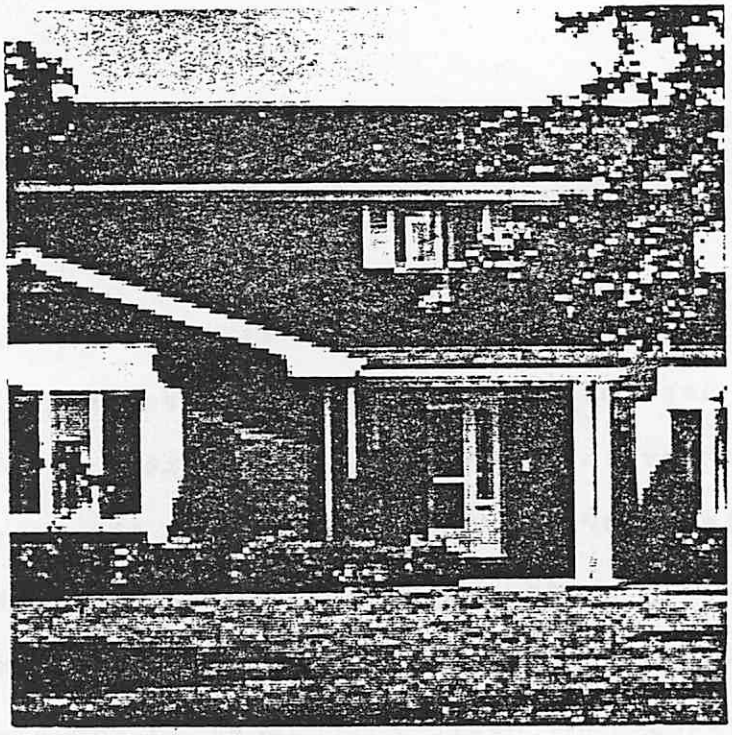

1(b)

Figure 1. Representative Images. (a,b) Intensity images of 128x128 portions of 256x256 color images. These images of natural outdoor scenes are representative of the input to the edge process developed in this paper. (c) 64x64 intensity subimage derived from (a). This data is used extensively throughout the paper as a test image. (d) 16x16 subimage showing detail in the bush area.

1(c)

1(d)

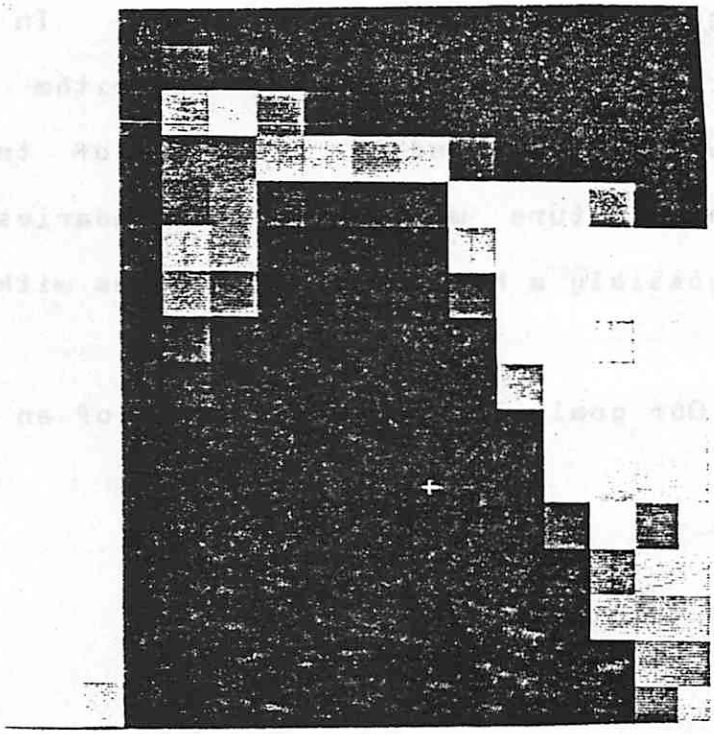

Figure 1 continued

of light. A discussion of the problems of and approaches to evaluation appears in [NAG79].

## I.3 Boundary Formation

The edge/boundary process presented in this paper cannot be expected to provide a perfect boundary segmentation of the image because no such segmentation exists. Rather we look for processes which can provide boundaries to a good first approximation, where the feature data exhibits relatively large and spatially consistent discontinuities. We will utilize the notion of variation in sensitivity of the boundary formation process so that it can be used to extract weaker and more ambiguous boundaries if so desired. In particular we do not expect this edge/boundary algorithm to distinguish between texture elements and the boundary of textured objects. Thus, strong texture will lead to boundaries around texture elements and possibly a high density of edges within a textured region.

Our goals in the segmentation of an image can be stated two ways:

1. to partition the image into (disjoint) sets of pixels, called regions, which have relatively invariant visual features.

2. to form continuous boundaries by placing edges between pixels which have relatively large differences in feature values.

When the visual data is very simple (e.g., a cartoon image with closed boundaries and without textural variation), these two approaches are equivalent. However, when the data is more complicated and has a high degree of ambiguity in certain places, then these approaches almost certainly will produce different segmentations and will use different analyses of the data to achieve the goal.

The edge/boundary process developed in this paper is a probabilistic edge labelling relaxation algorithm based on a simplified edge representation and a notion of boundary continuation in digital images. In this algorithm, information in the context of a given edge element is used in an attempt to drive the central edge into consistency with the surrounding context. The edge/boundary process is initially formulated in Bayesian probability theory. A model of desired line drawings is

used to determine theoretical estimates of the conditional probabilities relating an edge to the joint probability of edges in the context.

Some of the data used in this paper is shown in Figure 1. In all cases the feature used is intensity, which was derived from the original tri-stimulus color data. Our full images are actually subimages which are extracted from data whose original resolution was 256x256 pixels with at least 6 bits/color. In order to give a sense of our goals and methodology, the results produced by the edge/boundary process which is developed incrementally in the rest of the paper are previewed in Figures 2 and 3. The initial probability assigned to an edge is a function of both the local edge strength determined by application of a 1x2 edge mask and the strength of other edges in the immediate vicinity; this initial assignment is shown in Figure 2. Beginning with this initial data, the edge/boundary process updates every edge in the image simultaneously (in parallel) on the basis of data contained in a small neighborhood around each edge. This updating process is repeated until most of the edges do not change from one iteration to the next. The results shown in Figure 3(a)-(d) appear to be quite reasonable for most areas of the image.

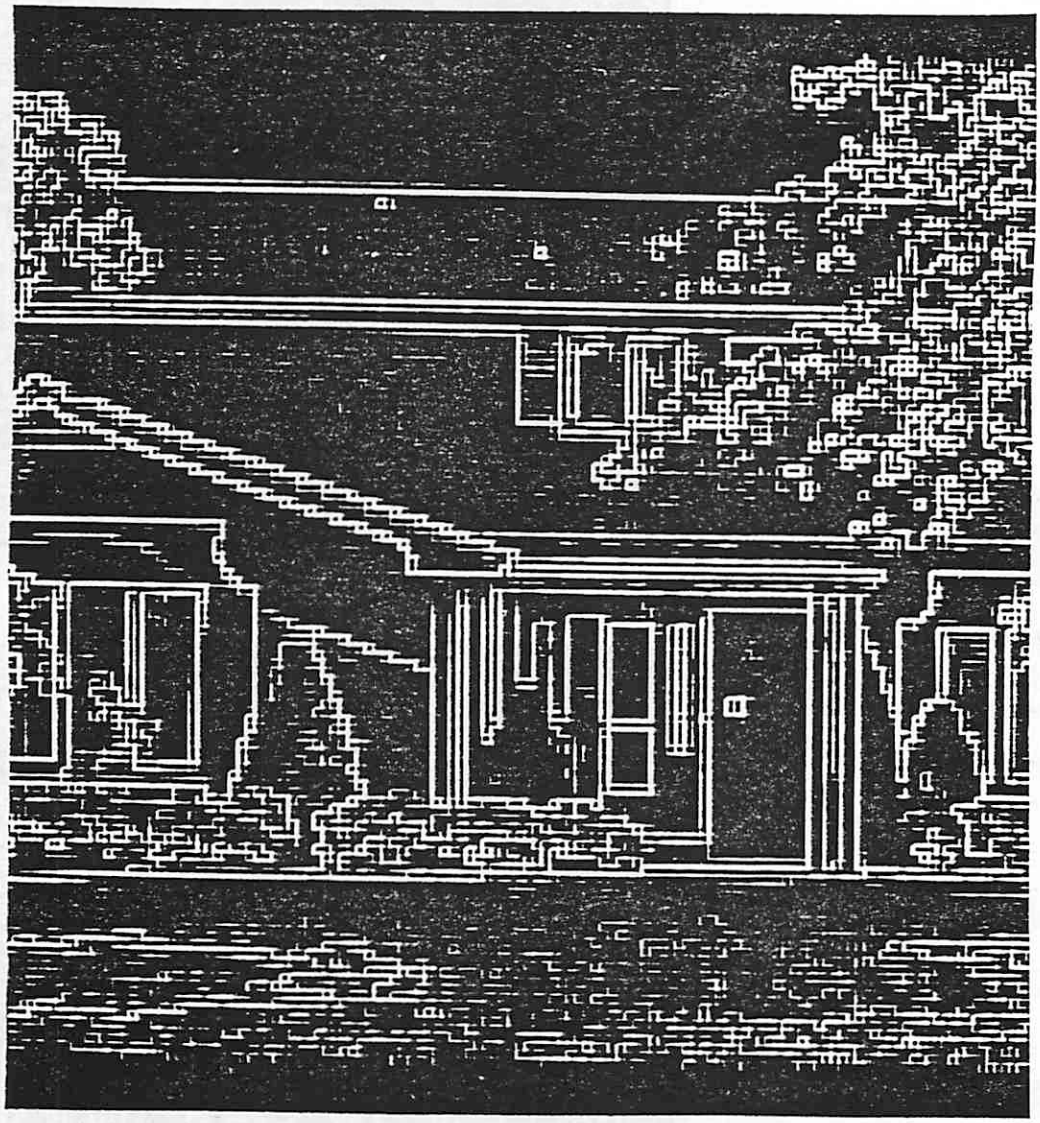

Figure 2. Transformation to Edge Probabilities. Edges derived from the intensity data are assigned a probability as a function of the total contrast and form of the gradient data. This forms the input data to the relaxation process.

Iteration 2

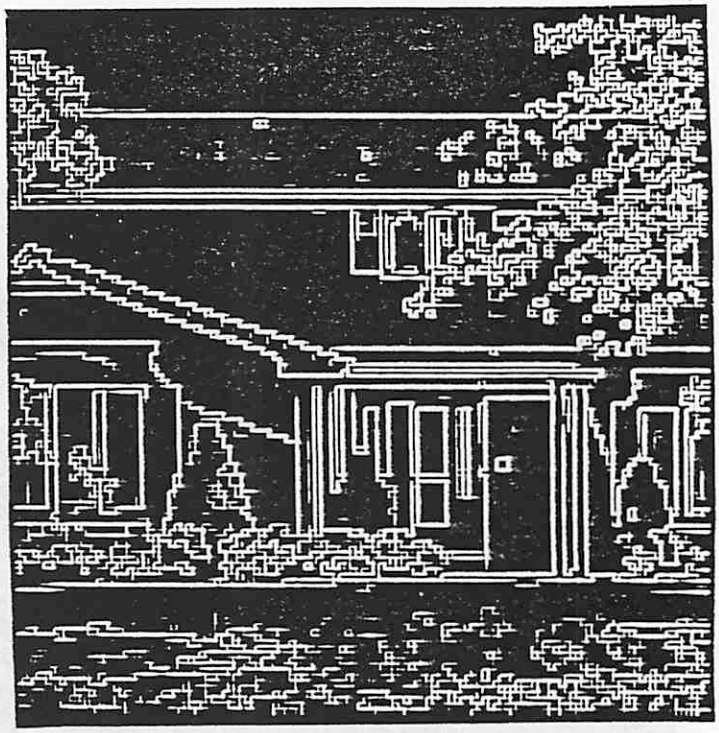

Iteration 5

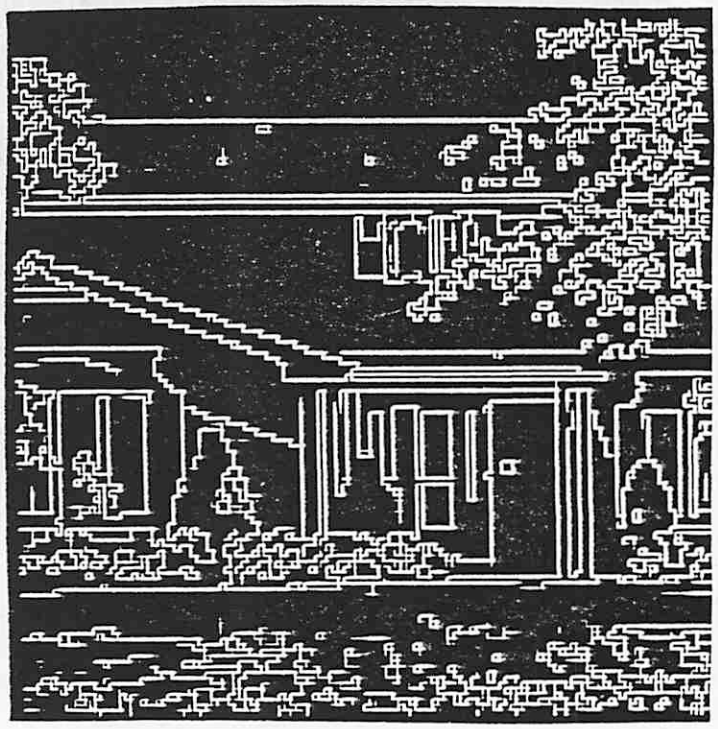

Figure 3. Results of Edge Relaxation for Boundary Continuity. Results from the final formulation of the edge relaxation process developed in Section IV. In most cases the boundaries are a very reasonable representation of the underlying data and with minor exceptions correspond quite well to visually apparent boundaries.

Iteration 10

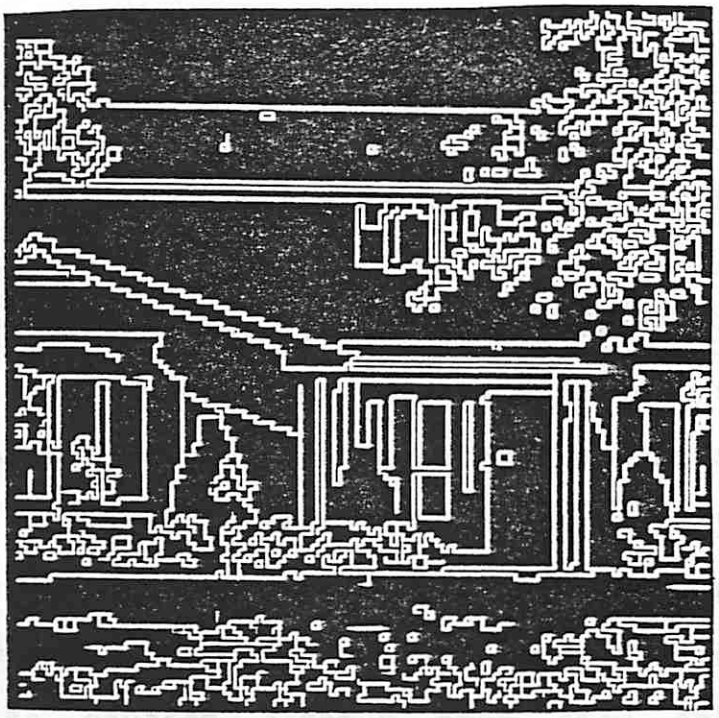

Iteration 20

Figure 3 continued

# II. A BAYESIAN VIEW OF EDGE CONTINUITY

## II.1 The Representation of Edges

Our representation of edge information is highly structured with each edge having an unambiguous location and orientation. In particular, a horizontal edge may be located between any pair of vertically adjacent pixels, a vertical edge may be located between any pair of horizontally adjacent pixels (refer to Figure 4). Edge presence at each horizontal and vertical edge location is encoded with some numeric measure of edge strength or likelihood. During the boundary formation process this representation avoids the ambiguities in precise location and orientation of edges that would occur if edge information relative to pixel locations were not encoded. The issues involved in this choice of this representation are discussed in [HAN78a, PRA80].

Our approach to boundary formation makes use of two basic assumptions:

1. the contrast assumption: the likelihood of a local edge is directly proportional to its contrast;

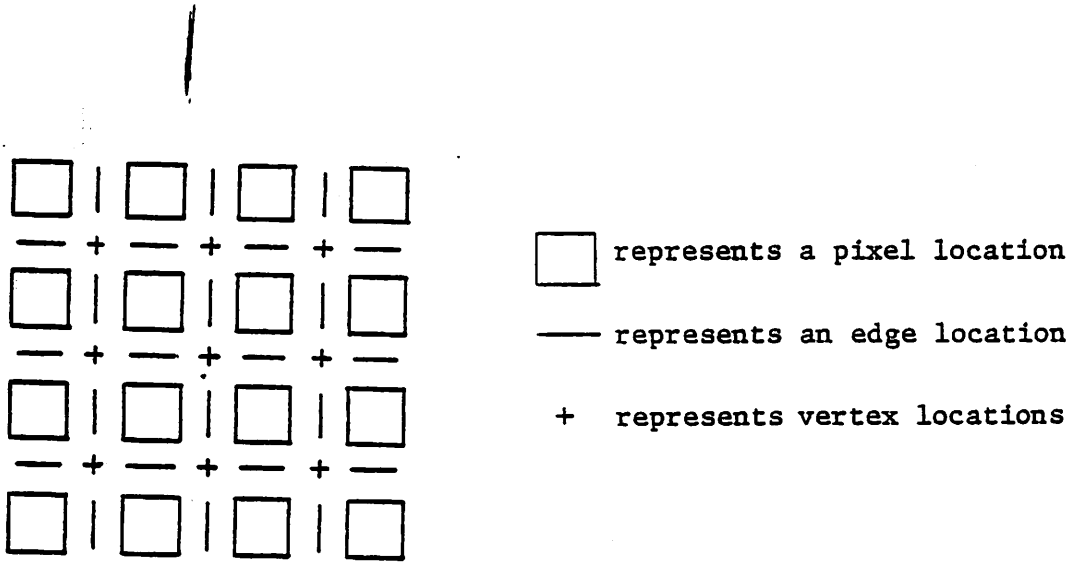


Figure 4. The Representation of Edges in VISIONS. This interpixel representation of horizontal and vertical edges between adjacent pixels allows the placement and orientation of edges to be precise. The location of possible boundary junctions and terminations follows from the edge positions.

and 2. the continuity assumption: boundaries in images are continuous almost everywhere; that is, every "real" edge element is joined to one or more adjacent edge elements almost everywhere.

The contrast assumption allows us to compute the likelihood of a local edge based solely on its contrast, while the continuity assumption provides a mechanism for propagating information from a "context" to the edge in question.

Our goal will be to define local, parallel mechanisms by which initial probabilities of edges can be updated so that uncertainty of edge presence is reduced while edges are aggregated into continuous boundaries. The means by which this will take place is based on the updating of an edge probability via a Bayesian view of the joint probability of edges in the local context surrounding that edge.

## II.2 The Probability of an Edge

As we have already noted, the image data to be analyzed is not very well-behaved. It is extremely difficult to make an absolute decision about the presence or absence of a particular edge element based purely on local information around that edge.

Consequently, it is desirable to avoid such absolute classification during the early stages of visual organization of information [MARR76, EHR79]; otherwise, one must face the computational burdens of undoing erroneous decisions. One plausible strategy would be to delay any decision at this point and to carry forward all possible decisions so that contextual data can be used in later stages to select the best choice. All possible local decisions -- or labellings -- at a single location are retained. The approach advocated here has some of this flavor mixed with the notion of spreading of contextual information via the parallel updating of edge likelihoods.

In our representation there are only two decisions -- or labellings -- at a single location: "edge" and "no-edge". A probability or degree of confidence will be associated with each label at each edge location. The term probability, rather than some heuristic notion of "confidence", will be used here because we will be guided by a Bayesian view of edge presence. Only a single value is required to store the probability of two labels at each edge location if the label set is assumed to be mutually exclusive and exhaustive.

Uncertainty in the existence of an edge can be expressed in various ways. We will denote the probability that some edge E is

present (true) or absent (false) with the random variable E, which can take on values T or F, respectively. Then, we can represent the probability of edge presence as P(E=T) and of edge absence as P(E=F) = 1 - P(E=T).

Via the contrast assumption, the probability of an edge is related to the difference in feature values of the pixels in the vicinity of that edge. Let us, leave aside, until the next section, the problems of gradual feature changes over a sequence of pixels, i.e., wide non-zero feature gradients. Thus, edge contrast varies over the same range as the feature values and can be scaled into a zero-one range to represent edge probability.

In this paper the conversion from edge contrast values to edge probabilities is performed by normalizing the maximum feature contrast over a large local neighborhood of an edge. The strongest edge in each neighborhood can have a probability of one and edges with zero feature difference will be mapped into probability zero. For each edge location in the image, we extract the maximum and minimum feature values, say Fmax and Fmin, within a local 11x11 window centered on the edge under consideration. The edge probability is then determined by

$$P(\text{edge}) = \text{MIN}\{1., E' \mid [K*(F_{MAX} - F_{MIN})]\} \qquad (1)$$

where E′ is the local edge obtained by differencing the adjacent pixel values (e.g., a 1x2 difference mask), and K is a local scaling factor which allows the contrast value which is considered to have a probability of one to be varied.

There are two parameters that have been heuristically set at intuitively reasonable, but arbitrary, constant values. Variation of the size of the neighborhood would change the locality of influence in setting initial probabilities, while variation of the scaling parameter K might serve as an edge sensitivity control. The effect of changing these parameters is not explored in this paper. The contrast neighborhood has been fixed at 11x11 and K has been set to .5 for all results presented in this paper. Figure 5 illustrates the conversion of edge contrast to edge probability for a subimage of Figure 1(a). Figure 5(a) represents the output of the 1x2 edge mask in which brightness encodes edge strength. Note that there are many weak edges which are not evident in the figure as it has been reproduced here. Figure 5(b) shows the edge probabilities obtained by applying equation 1 at every horizontal and vertical edge location. Note that the boundary on the left side of the bush is more complete and stronger. Because the normalization is performed locally over an 11x11 window, non-linear scaling may occur. A related effect, which can be seen in the wall area

(a)

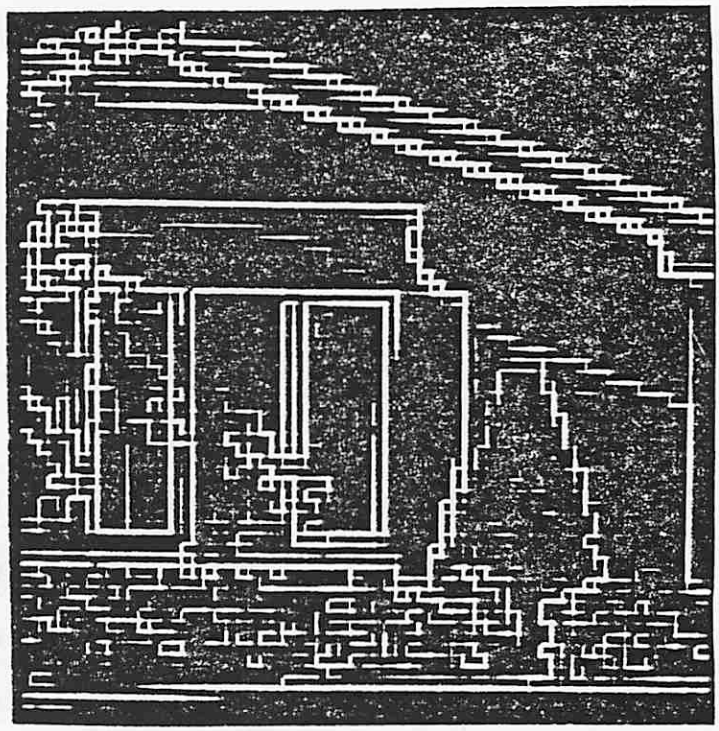

(b)

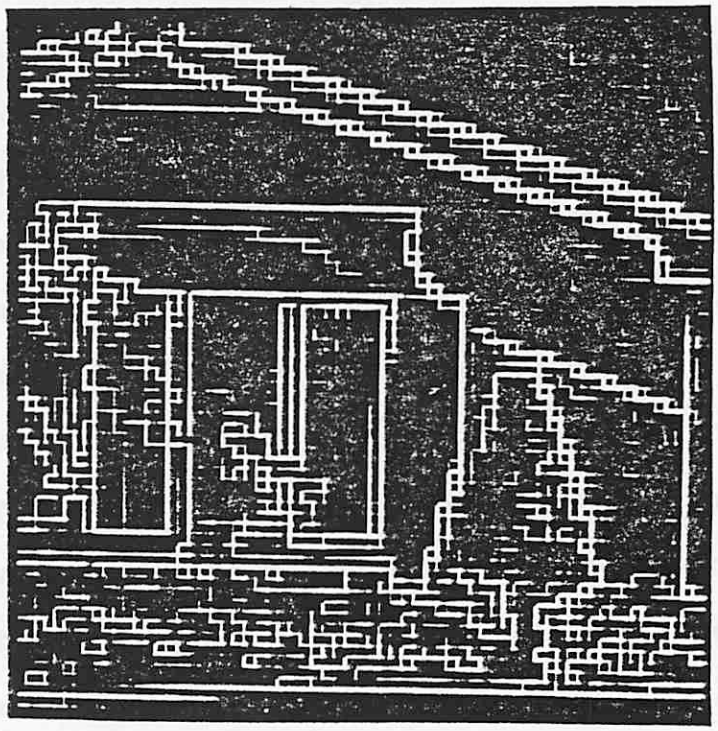

Figure 5. Conversion of Edge Strength to Edge Probability. (a) Result from 1x2 edge mask applied to Figure 1(c). Although it is not evident in the picture, there are a large number of very weak edges throughout the image. (b) Edge probabilities obtained by normalizing edge strength over an 11x11 local window centered on each edge. Compare the left side of the bush in the two images.

above the roof line, is the relative increase in the strength of weak edges in areas where there are only weak edges (and low maximum contrast).

## II.3 Gradient Collection

One objection which may be raised against the use of a 1x2 edge operator is related to those situations in which a boundary is actually represented by a non-zero feature gradient that is wide relative to the mask size. The highlight evident around the rim of the bush in Figure 6(a) is a function of the light source, the surface properties of the bush and background, and the relative orientations of the camera, the source, and the relevant surfaces. It might be argued that a sophisticated model of intensity changes is necessary, one perhaps in which the relationships between the geometry of the imaging process and the resulting image gradient are taken into account [HOR77, BAR78]. Although this approach may prove to be very effective, it has only been explored in artificially constructed domains. Extensions of the approach to an analysis of the physics and geometry of images containing foliage (for example) do not appear to be immediately feasible. We will view the total change in feature value across the width of a gradient as a single value to

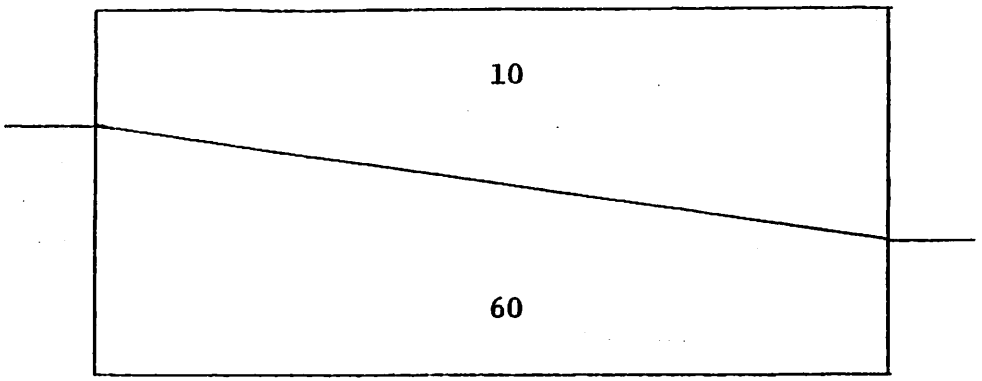


6(a)

| | | | | |
|---|---|---|---|---|
| 10 | 10 | 10 | 10 | 10 |
| 55 | 45 | 35 | 25 | 15 |
| 60 | 60 | 60 | 60 | 60 |

6(b)

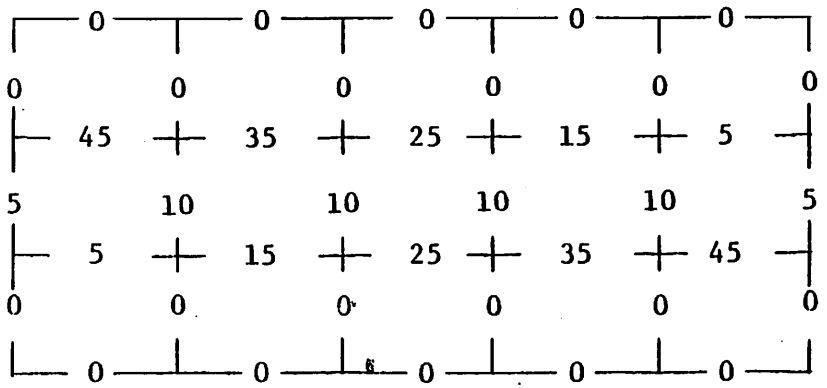


6(c)

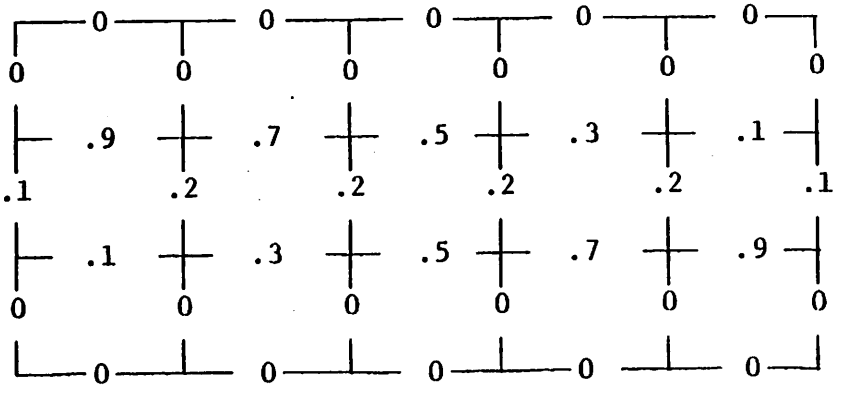


6(d)

Figure 6. The Gradient Collection Process. (a) The original data is a diagonal boundary separating two uniform regions. (b) Digitized data shows the effect of gradient formation due to finite sensor size and diagonal boundary. (c) Result of applying 1×2 edge mask (both horizontally and vertically) to digitized data. (d) Edge probabilities without any form of gradient collection, assuming a local maximum contrast of 100 in an 11×11 window centered on each edge location.

| Relative Position in Image | 0 | 1 | 2 | 3 | 4 | 5 | 0 | 1 |
|---|---|---|---|---|---|---|---|---|
| Edge Contrast | 5 | 10 | 10 | 10 | 10 | 5 | 35 | 15 |
| Center of Gravity | 2.5 | | | | | | .3 | |
| Location Factor | .375 | .625 | .875 | .875 | .625 | .375 | .925 | .825 |
| Relative Edge Strength | .5 | 1.0 | 1.0 | 1.0 | 1.0 | .5 | 1.0 | .43 |
| Collected Contrast Strength | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 | 1.0 |
| Gradient Collected Edge Probabilities | .188 | .625 | .875 | .875 | .625 | .188 | .925 | .354 |

6(e)

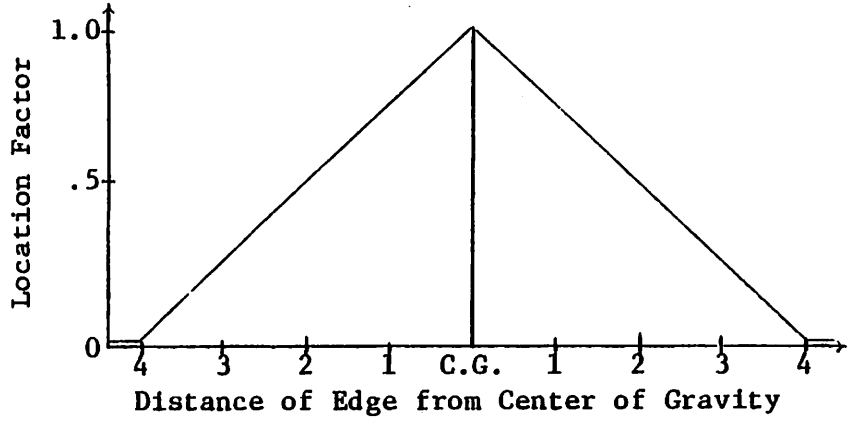


6(f)

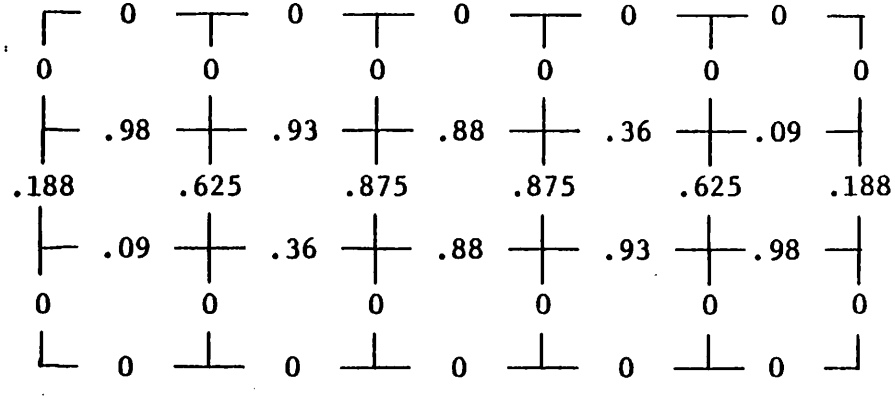


6(g)

Figure 6. The Gradient Collection Process. (e) Detailed computation of the gradient-collected edge probabilities for two of the gradients; the final probability is the product of the location factor, relative edge strength, and collected contrast strength. (f) The location factor is a linear interpolation around the center of gravity of the gradient. (g) Final values for the gradient collected edge probabilities; these values should be compared to (d).

be used in determining the likelihood of an edge somewhere within the gradient. This approach is described in more detail below and it seems to work reasonably well in our test images.

In many cases a gradient edge is produced whenever an environmental edge is not parallel to one of the axes of digitization (Figure 6(a,b)). The degree of "smearing" is a function of the slope of the edge and the digitizer characteristics (size of the active sensor, sampling rate, etc.). An example of this effect is clearly seen in the sloping roof edges of Figure 7(a) where the actual roof edge in the original photograph is a sharp edge.

In cases such as these, the mask will respond to only a portion of the full contrast which is distributed over several pixels. This has led to approaches in which a hierarchical set of increasingly larger masks is employed [ROS71, HAN74, MARR76, HAN80]. The approach suggested here -- referred to as "gradient collection" -- is to extract the total contrast of a boundary by collecting the local horizontal and vertical components of the gradient (as obtained from the 1x2 mask) and to place the resultant value at a representative edge location [HAN78a]. It is essentially an application of the contrast assumption to edges where the non-zero gradient is

spatially distributed. The probability of the resulting edge may then be determined using the method presented in the preceding section. This is discussed more fully in Section IV. Note that in effect this strategy is an avoidance of fixed mask sizes and geometries; the effective mask size is determined dynamically as a function of the local data.

Our gradient collection process involves selection of an edge location within the width of the gradient as the representative canonical location of the entire gradient edge, as well as the specification of the likelihood to be associated with the gradient edge as a function of the total contrast. In general, edges will have to be placed at several representative locations to minimize misalignment of adjacent edges. This misalignment can occur if adjacent views of the gradient differ due to noise or systematic changes in the extent of the gradient.

Gradient collection is performed during the assignment of initial edge probabilities. All horizontal runs of consecutive vertical edges and vertical runs of consecutive horizontal edges across a non-zero gradient and with the same direction of contrast are considered as gradient edges. As illustrated in Figure 6, three multiplicative factors determine the final probability assignment to each constituent edge in the gradient:

1. the location factor - the center of gravity (c.g.) of the edge strengths in a gradient is used as the point about which an inverse distance weighting is computed; the location factor will be assigned a value of 1.0 at the c.g., a value of 0 for edges at a distance of 4 pixels or greater from the c.g., and linear interpolation for all edges at a distance from the c.g. between 0 and 4;

2. the relative edge strength factor - all edges in a gradient are divided by the maximum edge so that the places of strongest contrast change will receive a greater weighting; and

3. the collected edge strength factor - the collected strength of the gradient (i.e., the sum of all constituent edges) is assigned a probability by normalizing by the maximum local contrast, which is scaled by a factor K and bound by a maximum value of 1, as in equation (1).

These three factors when multiplied give the final edge probabilities along a gradient. It is easy to verify that this

method has several desirable properties. The inverse distance weighting about the center of gravity focusses evidential weight on a few likely candidates. Normalization of individual edge strengths maintains some information about the relative edge strength in a gradient, while the collected edge strengths over the gradient allows the full contrast to influence the final probability of edge presence.

Figure 7 illustrates the application of the gradient collection process to actual data. For illustrative purposes, a small portion of the roof trim (Figure 7a,b) was chosen because of the horizontal and vertical gradients both above and below the trim area. Figure 7(c) shows the result of applying a 1x2 edge mask to the intensity data; in this figure edge intensity encodes edge strength. The picture clearly shows the multiplicity of edges caused by the gradients. Figure 7(d) is the results obtained from the gradient collection process described above; here, brightness encodes edge probability. Without changes in the display scaling the effects of the gradient collection process are hard to see, but there is more consistency in the presence of one or two strong edge locations down the entire length of the boundary. Otherwise any weak spot, particularly in the vertical position, cause the boundary to break up during relaxation. It is difficult to evaluate the

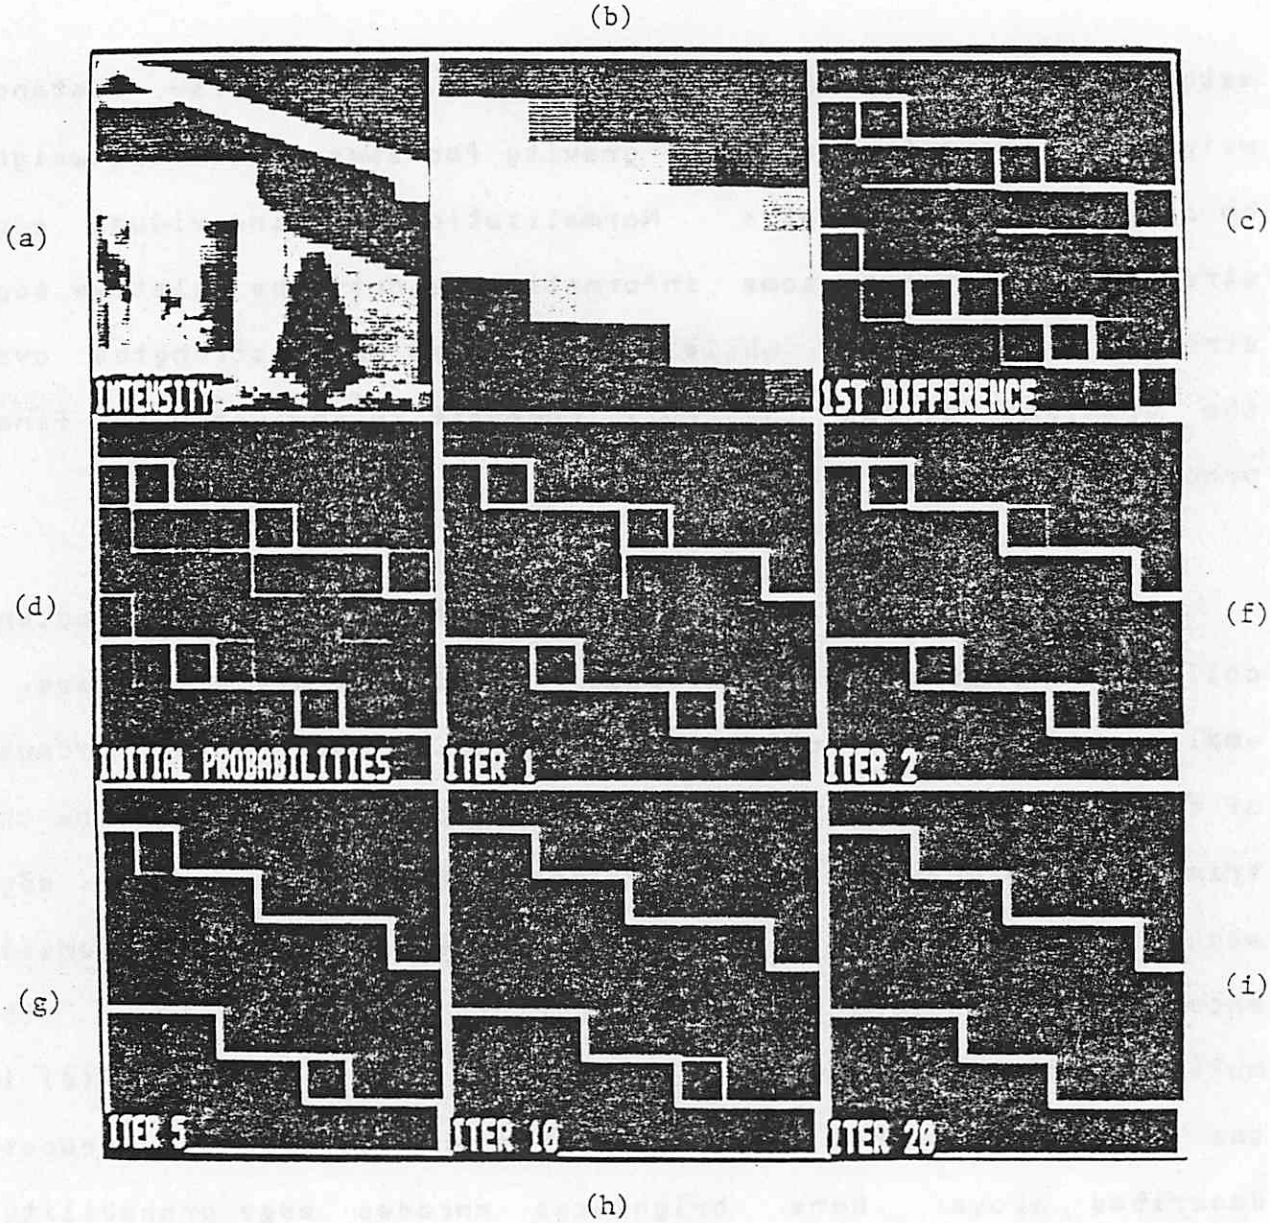


Figure 7. Application of the Gradient Collection Process. (a) Intensity image of house showing location of roof trim subimage. (b) 8x8 section of roof; note wide non-zero gradients both above and below the roof trim. (c) Results of 1x2 edge mask. (d) Gradient-collected edge probabilities; edge brightness is a function of probability (0 = black, 1 = white). (e-i) Iterations 1, 2, 5, 10, and 20 (respectively) of the edge relaxation process developed in the paper based on the initial gradient-collected edge probabilities; the final result is a reasonable boundary representation of the roof trim subimage.

effectiveness of the collection in isolation from the relaxation process which will make use of the data. Figure 7(e-i) shows the application of the edge relaxation process developed in the remainder of the paper. It is apparent that by 20 iterations (Figure 7i) the boundary obtained is a reasonable representation of the visual boundary in 7(b). Note that the boundary in 7(i) is composed only of horizontal and vertical edge elements. As a result, there is no information locally available about global characteristics of the boundary, such as the slope of the roof trim. In our system, the final boundary results shown here are processed further and relevant features extracted [YORBO].

While we have presented a totally heuristic strategy for assigning initial probabilities; the factors are all relevant to edge placement and provide intuitively reasonable values. Better models of gradient collection, perhaps guided by psychological considerations in boundary perception, are called for.

## II.4 A Bayesian View of the Local Context of an Edge

The simplest view of the boundary formation process involves the use of the information contained in the context of a single edge to affect the degree of belief in the existence of that edge

as part of a continuous boundary. Given the horizontal/vertical edge representation that we employ, the smallest neighborhood which allows contextual information to propagate by means of boundary continuity is shown in Figure 8. Such a neighborhood must include the three edge locations to either side of a given edge where continuation of that edge can occur. Shortly, we will expand the local context to include the adjacent parallel edges so that competing alternate locations for the placement of wider gradient changes can be handled in a local parallel organization of the information. However, let us first examine the simpler context of only the six potential continuation edges around any given horizontal or vertical edge.

Given the conditions of uncertainty in the presence of most edges, the goal is to use the neighborhood context to improve our estimate of the probability of the central edge. It is natural to turn to Bayesian probability theory for the definition of the relationship of the central edge to the neighborhood context. Letting random variable H be associated with the central edge, the probability that central edge H is present will be P(H=T) and can be determined as the sum of the joint probability of each of the mutually exclusive 2**6 = 64 ways that H can occur within the context of the surrounding edges:

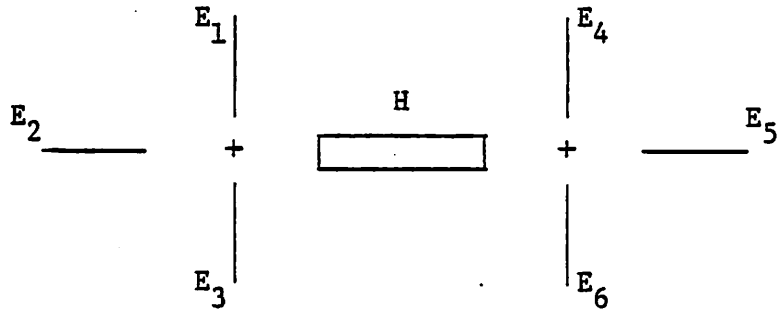


Figure 8. Context of Edge H. The smallest neighborhood which allows contextual information to propagate to H must include the possible boundary continuations to the left (edges $E_1$, $E_2$, and $E_3$) and to the right (edges $E_4$, $E_5$, and $E_6$). Note that + represents vertex locations where horizontal and vertical edges can meet.

$$P(H=T) = P(H=T,\ E_1=T,\ E_2=T,\ldots,E_6=T)$$
$$+ P(H=T,\ E_1=F,\ E_2=T,\ldots,E_6=T)$$
$$+ \ldots$$
$$+ P(H=T,\ E_1=F,\ E_2=F,\ldots,E_6=F)$$

This can be shortened to:

$$P(H=T) = \sum_{E_1} \ldots \sum_{E_6} P(H=T,\ E_1,\ E_2,\ldots,E_6) \qquad (2)$$

where the summation involving each Ei, i = 1,...,6, is enumerated over Ei = {T,F}. In effect the likelihood of the central edge given a context of six local edges, where each edge Ei in the context may be uncertain, is a linear interpolation of the likelihood of the 64 six-tuples for which there is certainty about the joint presence and absence of each edge.

Using the standard definition of conditional probability the joint probability of the central edge and context can be decomposed:

$$P(H=T) = \sum_{E_1} \ldots \sum_{E_6} P(H=T|E_1,\ldots,E_6)\ P(E_1,\ldots,E_6) \qquad (3)$$

The conditional probability of the central edge, given a state of certainty in the context of six edges, will embody the semantics of boundary continuity.

## II.5 Equivalence Classes of Edge Contexts

By taking into account natural symmetries in the local neighborhood, the 64 cases which must be explicitly considered can be significantly reduced [HAN78a, PRA80]. To accomplish this, a labelling convention for the edge patterns that can appear in the left and right contexts of the central edge is required. With the generality of our continuity assumption, it is not necessary to take into account the orientation of the extension of the central edge into the context. Note that this would be important if straight boundaries (or boundaries with some other parameterized 2D curvature) were to be extracted, but in this paper no constraints on boundary curvature other than continuity are considered.

Extensions of the central edge into the context can best be categorized by labelling the two vertices of any given edge. A vertex label will be a function of the number of edges present in the three edge locations which can serve as the continuation of a boundary which passes through the central edge. Thus, we will represent the four possible equivalence classes of vertices as V0, V1, V2, and V3, where Vi represents the case where there are exactly i edges present at a vertex in the context. Thus, the three possible single edge continuations at a vertex of the

central edge are treated as an equivalence class, as shown in Figure 9.

The notation naturally extends to the full representation of our context of six edges by considering jointly the vertices to the right and left. Thus, Vij will be used to represent a context with a vertex labelled Vi to the left and a vertex labelled Vj to the right. For vertical central edges Vij will represent top and bottom vertices. Now, any of our 64 six-tuples of edges which are present or absent can be represented by one of the 16 equivalence classes Vij, $0 \leq i, j \leq 3$. Thus, the equivalence classes produced by symmetry of the number of edges reduced the 64 contexts to 16 contexts. By the additional left-right symmetry of vertices (which does not, of course, affect continuity considerations), we can further reduce the set of equivalence classes to the ten Vij, $0 \leq i \leq j \leq 3$. Figure 10 is an enumeration of these ten classes in which an unfilled bar represents the central edge H, a solid line represents an edge in the context which is present, and a dotted line represents an edge location where an edge is absent.

In summary, there has been a significant reduction in the number of local contexts which must be considered for boundary continuation. It was achieved by assuming left-right symmetry

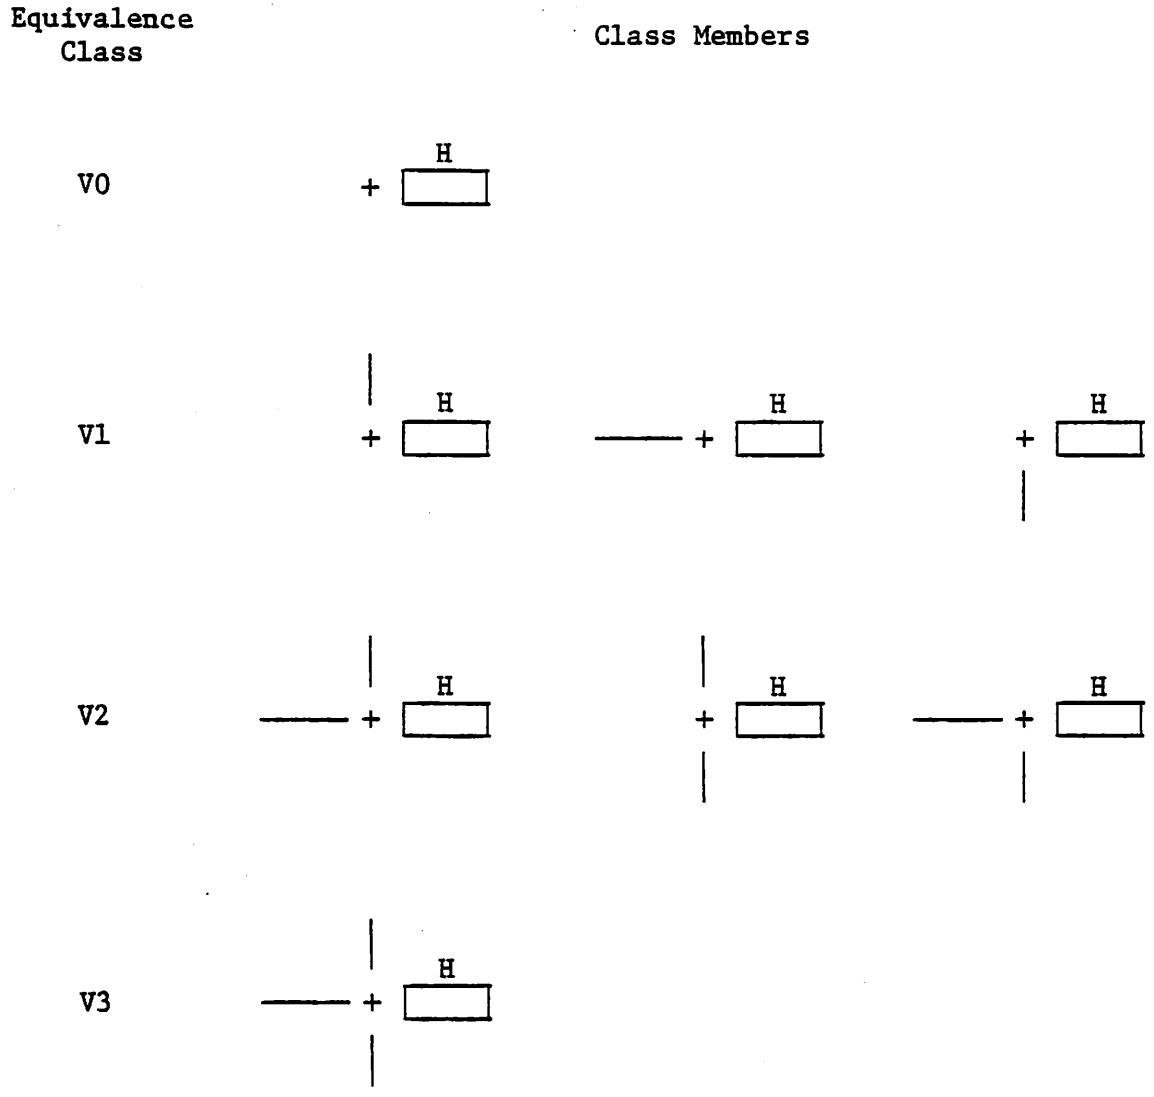


Figure 9. Equivalence Classes of Vertex Types. Each type represents an equivalence class of line continuations in which the number of possible continuations is fixed. Note that by symmetry a similar labelling convention holds when the central edge is to the left of the vertex.

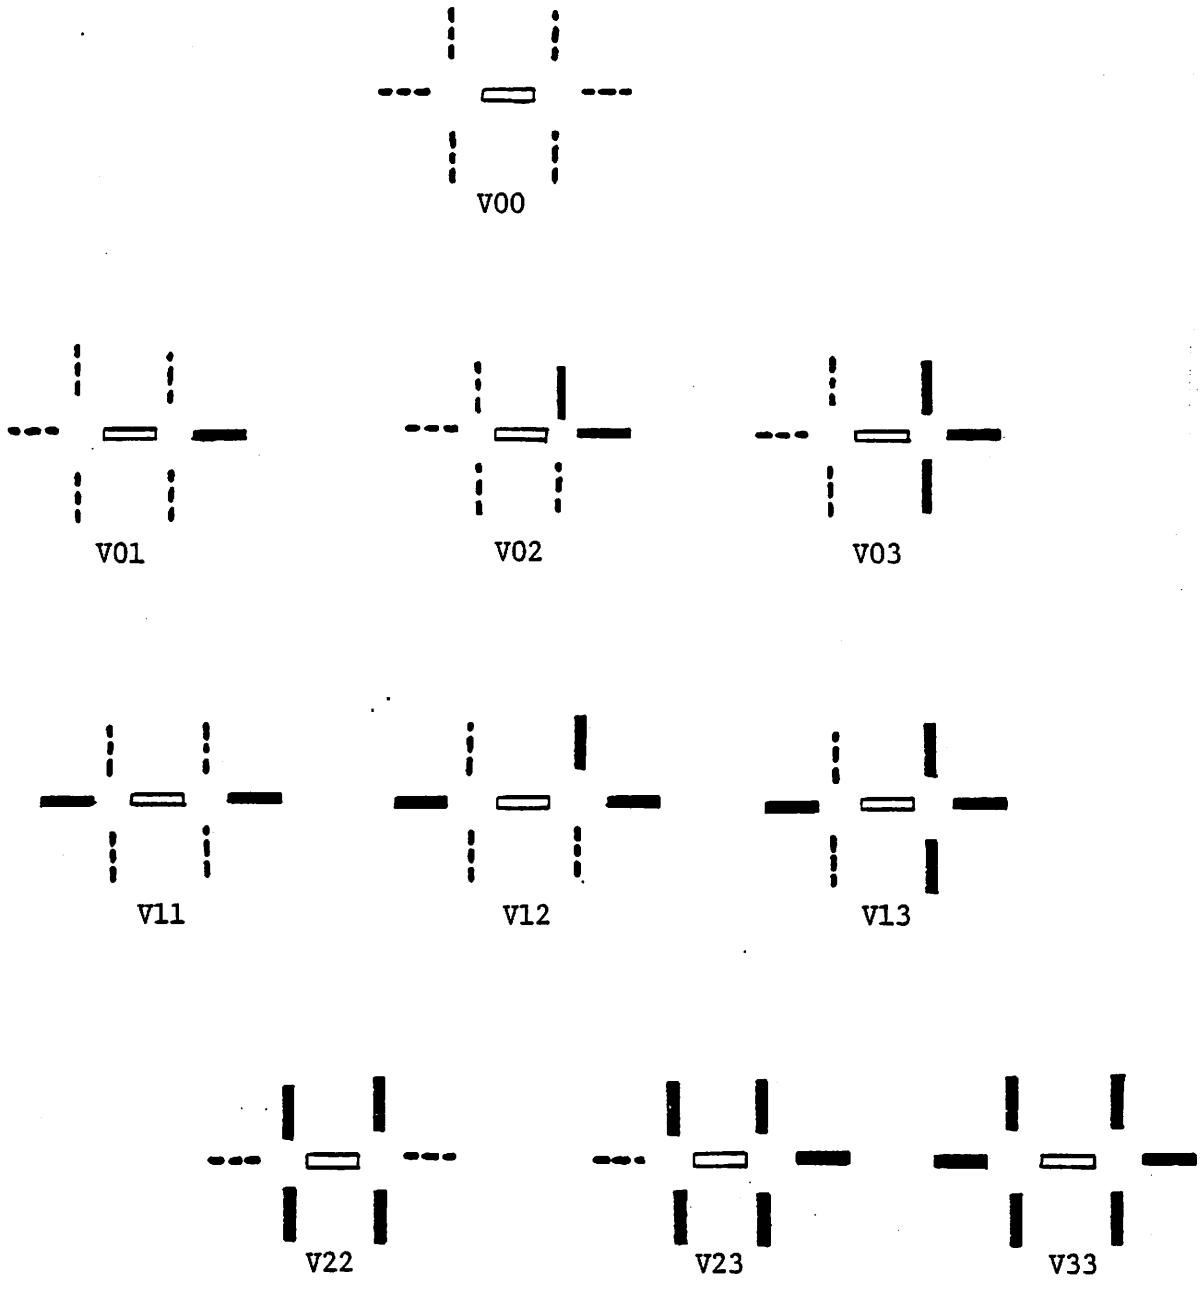


Figure 10. Ten equivalence classes of edge contexts based on the number of edges in each vertex and subject to left-right symmetry.

and counting only the number of edges within each half of the six edge context while disregarding the actual edge orientation and edge placement. As we have pointed out, orientation can be very useful, and it is employed in our system at a later stage of processing and at a different level of representation [HAN78a, YOR80].

## II.6 Determination of the Conditionals

Our primary purpose in updating the likelihoods of edges is to reduce uncertainty in the presence of elemental edges by relating them to their local contexts. The organization of edges into boundaries can be viewed as the derivation of a likely "line drawing" that is "close" to the initial edge estimates, and at the same time is "consistent" with desired properties of the line drawings.

The conditional probabilities in equation (3) provide the theoretical basis for embodying the contextual information of boundary continuity in the updating of edge probabilities. The conditional probabilities relate the joint probability distribution of the central edge with its surrounding context to the distribution of the surrounding context; by definition

$$P(H|E_1,\ldots,E_6) = P(H, E_1,\ldots,E_6)/P(E_1,\ldots,E_6) \qquad (4)$$

These conditionals will be estimated from the class of all line drawings that are acceptable or desired; or in other words the class of "likely" line drawings that are expected to be derived by the edge process. Figure 11 is an example line drawing demonstrating the characteristics of the boundaries that are expected to appear in our desired line drawings -- lines are continuous and either terminate or meet at junctions with other lines. One should notice that a primitive edge element can appear in only a few types of local patterns in our class of edge images. Isolated edges and lines with small gaps do not appear. Any edge which appears as part of a boundary will have a boundary continuation through at least one, and probably both, vertices of that edge. Furthermore, an edge which participates in a boundary continuation might take the form of either a single line continuation, or the edge may complete a line junction of two or more lines.

The conditional probabilities can be estimated from a large sample of line drawings that are representative of the class of desired line drawings. Mapping these drawings onto the rectilinear edge representation used in this paper provides a large number of local contexts for estimating the joint

probabilities of equation (4). Rather than actually collect such a set of samples and exhaustively tabulate the statistics of the ten equivalence classes, we can construct a general model of the class of drawings and theoretically specify values of the conditional probabilities from the model.

It is straightforward to tabulate how many times an edge or no-edge occurs with each context Vij and then:

$$P(H|V_{ij}) = P(H, V_{ij})/P(V_{ij})$$
$$= P(H, V_{ij})/[P(H=T, V_{ij}) + P(H=F, V_{ij})]$$

Figure 12 summarizes the joint occurrence of H with Vij for each of the 10 equivalence classes based upon our model of desired line drawings.

The equivalence class of V00 is fairly obvious in that we do not expect (or want) to see "isolated" edges — i.e., a single edge with no continuation of edges in its context. Thus, P(H=T|V00) is close to zero. The V11 case is also clear in that the central edge should be present and, therefore, P(H=T|V11) is close to one.

The case of V01 is somewhat more subtle. If we consider a boundary which terminates as a "line ending" (i.e., with no

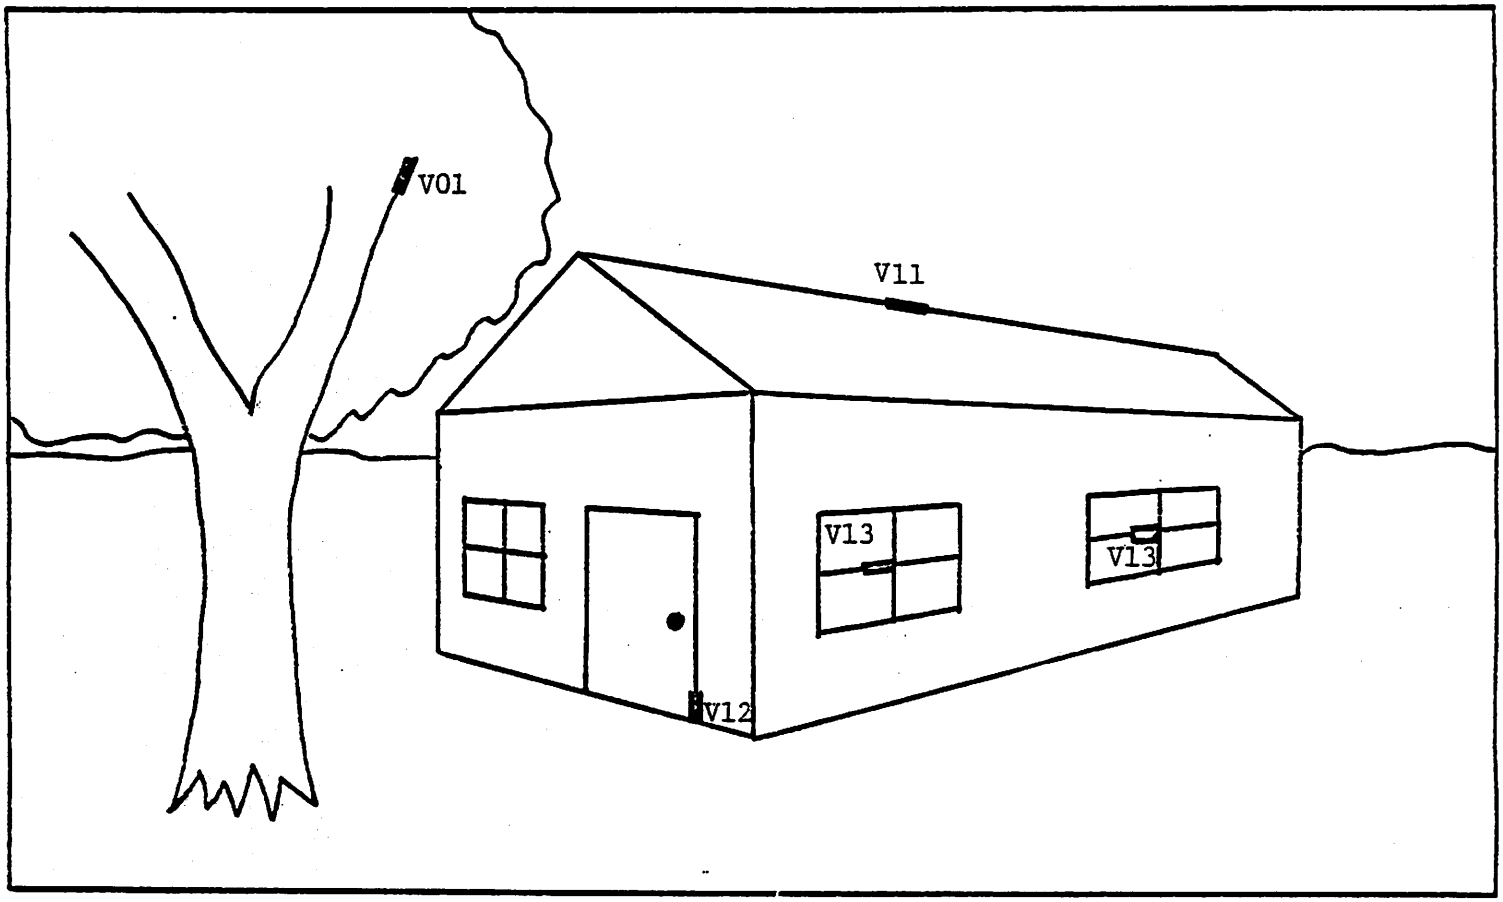


11(a)

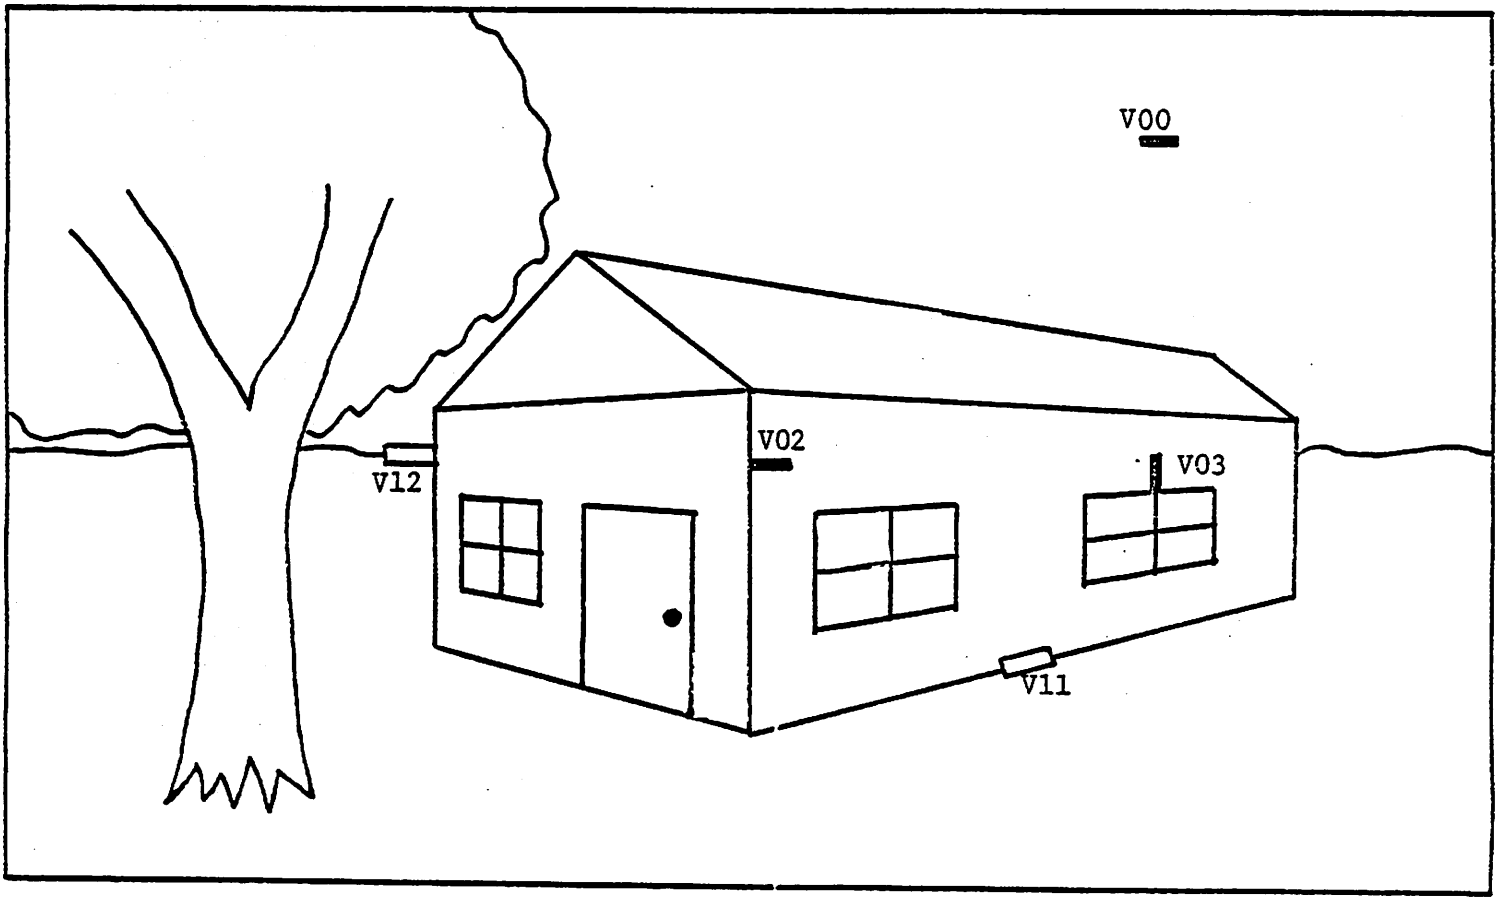


11(b)

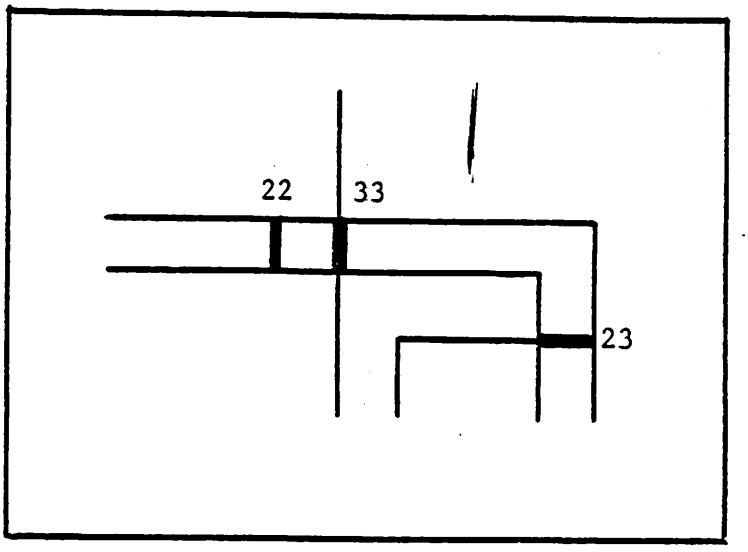


11(c)

Figure 11. Boundary Characteristics Expected in Experimental Domain. Example line drawing showing characteristics of boundaries and vertices expected in the general class of line drawings. (a) Edges expected to appear. The local edge contexts shown around the V01, V11, V12, V13 contexts are the only ways, subject to symmetry, that an edge can participate in a boundary in our class of line drawings. (b) The edges marked V00, V02, and V03 clearly do not participate in closed boundaries and should not appear in the line drawing. The absence of edges, as shown by open rectangles, in the V11, V12, and V13 context would result in a break in boundary continuity and these cases also are not expected to appear. (c) Contexts with inherent ambiguity. There is some degree of ambiguity concerning the need for the edges in the V22, V23, and V33 contexts. None of them are required for boundary continuity, although they may appear because they accurately reflect the underlying data in the pixel perpendicular to the edge. One would also expect these cases to occur in dense edge areas, such as heavily textured portions of an image, or, as we have seen, in gradients (cf. Figures 6 and 7).

continuity to one side), then the last edge appearing in that boundary will have a V01 context. However, for each such case there are three edge locations forming the V0 vertex of that last edge, which, when considered as the central edge location themselves, will contribute to the count of H=F and V01; thus, there are three times as many joint H=F and V01 occurrences as joint H=T and V01 occurrences.

The cases of V12 and V13 involve the edge locations leading into the junction of boundaries, and they derive from the fact that a boundary is not expected to terminate just before meeting another boundary. Finally, the V22, V23, and V33 all involve cases which are somewhat ambiguous. The absence of the central edge does not really affect edge continuity in the neighborhood since there are independent continutions to each side of the edge. However, if the underlying data implies the presence of the central edge, its survival should not be discouraged. Since it is much more likely that two nearby boundaries are not connected by "bridging" edges (as in the case of the boundaries associated with narrow horizontal or vertical stripes), a low value for $P(H=T|V22)$ is arbitrarily chosen, say .1. The V23 and V33 contexts have similar analyses. Note that the V22, V23, and V33 contexts might be expected to be more frequent in textured areas or in places of non-zero gradients.

| Vij | Joint Occurrence of H and Vij | Joint Occurrence of $\overline{H}$ and Vij | Approximate P(H\|Vij) |
|---|---|---|---|
| V00 | low | high | 0 |
| V01 | number of line endings | three times number of line endings | .25 |
| V02 | low | high | 0 |
| V03 | low | high | 0 |
| V11 | high | low | 1 |
| V12 | high | low | 1 |
| V13 | high | low | 1 |
| V22<br>V23<br>V33 | number of times two or more nearby lines and/or line junctions have a bridge | number of times two or more nearby lines and/or line junctions do not bridge | << .5 ? |

(a)

| Vij | P(H\|Vij) |
|---|---|
| V00 | 0 |
| V01 | .25 |
| V02 | 0 |
| V11 | 1 |
| V12 | 1 |
| V22 | .1 |

(b)

Figure 12. Estimation of the Conditional Probabilities from our Model of the Class of Line Drawings. (a) The ten equivalence classes. (b) Reduction to six equivalence classes based upon identical semantics of the V2 and V3 vertices.

The set of ten equivalence classes of local edge contexts can be further reduced. An examination of Figure 12(a) shows that the conditionals associated with the V2 and V3 contexts are the same. This reduces the set of equivalence classes to six, namely Vij, $0 \leq i \leq j \leq 2$, as shown in Figure 12(b), where the V2 vertex class now includes V3 also.

## II.7 The Patterns of Boundary Continuity

It may be helpful to briefly examine the semantics of these six equivalence classes over the 64 patterns of edge contexts:

- $P(H=T|V00) = 0$ – an "isolated edge" which does not appear to participate in a continuous boundary;
- $P(H=T|V01) = .25$ – an "ambiguous boundary continuation" where the line may terminate or else continue in one or more of the three possible directions;
- $P(H=T|V02) = 0$ – a "spur" which is an edge whose absence will not affect the boundary continuity of the other edges which are present in the context;
- $P(H=T|V11) = 1$ – local "continuity" for a single boundary is ensured through the central edge;
- $P(H=T|V12) = 1$ – boundary continuation via "junction completion" is ensured through the central edge;
- $P(H=T|V22) = .1$ – "ambiguity of bridge" linking nearby boundaries; the edge context provides no information concerning the presence of the central edge.

Hopefully, these descriptions provide some insight into the semantics of local patterns of continuity. Four of the patterns -- V00, V02, V11, V12 -- appear to be quite reasonable. However, uncertainty in the V01 and V22 cases will give us some degree of difficulty. In the V01 case the problem is immediately obvious when one considers that the abrupt termination of a boundary composed of elemental edges with probability one will update the probability of all three adjacent edges to a value of .25. This appears to be a diffusion of information rather than a coherent organizing process.

There is also a potential problem in the use of the V22 conditionals in the case of two lines which are close together. For example, consider two strong parallel lines one pixel apart, as in the case of a black stripe on a white background; the probability of an edge linking the right and left boundary of the stripe at each possible "bridging" location will be set to a constant (say K), no matter what is implied by the underlying data of the pixels in the stripe. Thus, any value of K, high or low, will be wrong in certain cases! These issues will be set aside for the moment but will be addressed again later.

## II.8 Relaxation Updating as Bayesian Consistency

The development, thus far, has focused on determining the likelihood of a single central edge given its local context of neighboring edges, and was based to a reasonable degree on a Bayesian view of this relationship. Now we wish to embed this relationship within an iterative updating scheme which is to be applied simultaneously to all edge locations in the image, realizing full well that this step takes us outside the contraints imposed by Bayesian theory. This parallel updating will be viewed as a relaxation labelling process whereby an inconsistent set of likelihoods are driven towards consistency with their neighborhoods. The conditional probabilities serve as compatibility coefficients in relating the neighborhood influences of edges in the context.

Since the parallel updating scheme is to be repeatedly applied, it will be convenient to extend the notation to reflect the iterative progression of the computation by including the notion of time in the updating function:

$$P^{t+1}(H) = \sum_i \sum_j P^t(H|V_{ij}) P^t(V_{ij}) \qquad (5)$$

Let us make the following important simplifying assumption: the conditional probability distribution of H upon its context is

given and is stationary. Consequently the conditionals in equation (5) will not require a superscript t. We now observe that the unary marginal probabilities of edges P(Ei) are determined a priori on the basis of external physical evidence, but the joint probability of edges P(E1, E2, E3, E4, E5, E6) are not available. This implies that the joint probability cannot be computed without additional information or assumptions, the most obvious being independence of the edges in the context. By using the assumption of independence of left and right vertices to approximate the joint probability of edges, yielding

$$P^{t+1}(H) = \sum_i \sum_j P(H|V_{ij})P^t(V_i)P^t(V_j) \qquad (6)$$

We can further approximate the probability of vertex types using independence of edges in each vertex

$$P(V_0) = P(E_1=F,\ E_2=F,\ E_3=F) = P(E_1=F)\ P(E_2=F)\ P(E_3=F)$$

$$P(V_1) = P(E_1=T,\ E_2=F,\ E_3=F) + P(E_1=F,\ E_2=T,\ E_3=F) + P(E_1=F,\ E_2=F,\ E_3=T)$$

$$P(V_2) = P(E_1=T)\ P(E_2=T)\ P(E_3=F) + P(E_1=T)\ P(E_2=F)\ P(E_3=T) + P(E_1=F)\ P(E_2=T)\ P(E_3=T)$$

$$P(V_3) = P(E_1=T)\ P(E_2=T)\ P(E_3=T)$$

The assumption of edge independence may cause the careful reader

some concern. We are using the dependencies between every edge and its neighboring context to update edge probabilities, yet portions of the context are assumed to be independent! In particular when one of the edges, say E1, is considered as the central edge, then the updating function will be based on the dependencies of E1 upon E2 and E3 and its other surrounding edges in the context -- yet we have already assumed E1, E2, and E3 are independent. Thus, we have violated theory here and have entered the realm of heuristics. Even without the assumption of independence, our mechanism for parallel updating is still not theoretically sound from a Bayesian viewpoint. While any single context could be used to update an edge probability, that central edge should not also be used to update the value of the other edges which were in its context. Otherwise, this allows the probability of an edge in equation (5) to have an effect upon itself as closed loops of probability updates. An examination of the theoretical issues in this approach to relaxation labelling and the propagation of information in inference networks [DUD76] is treated in detail in [LOW80].

Despite some of the theoretical difficulties, we will examine the use of such an updating function. In general, initial values of individual edge probabilities are consistent with the underlying feature data which gave rise to them, but are

not consistent with their surrounding contexts in terms of the conditional probabilities derived from the class of ideal line drawings. If the conditionals remain constant, then one can view the updating process as a means by which the likelihood of a central edge becomes consistent with its surround.

If the array of edge probabilities ever reaches a state where the updating function leaves every edge probability at its current value, i.e., $P^{t+1}(H) = P^{t}(H)$ for all H in the image, then the system is at a fixed point. Further iterative updating will leave all values unchanged. Note also that at a fixed point all edges are consistent with their local surrounding contexts because the updated values implied by the contexts are the same as their current values. Thus, fixed points are consistent in this pseudo-Bayesian sense.

## II.9 Examination of Specific Cases

In this section we examine two cases of initial probabilities which demonstrate certain difficulties in our formulation of edge relaxation, and then present some results which confirm these difficulties. One case demonstrates the decay of edge patterns which should remain stable, while the

other case demonstrates the growth of 'noise' which should disappear.

Case 1: A closed loop of edges with probability q on a background of zero probability edges.

Note that q = 0 and q = 1 are fixed points. However, if $0 < q < 1$, a fixed point is approached in the updating process, but there is total loss of the starting information that represented the boundary (refer to Figure 14). Let us analyze why this takes place by representing $P(H=T|V01)$ with variable W. Then the updating equation for central edges in the closed loop (each with probability q) is given in Figure 13(a), while the update for zero probability edges that are touching the boundary are given in Figure 13(b).

It is interesting to note here that a value of W = .25, as defined by our estimates in Figure 12, will cause the updated value of edges H in the loop to decrease in probability if $q < 1$, even if q is arbitrarily close to 1; thus

$$P^{t+1}(H=T) < P^{t}(H=T) \quad \text{when} \quad P^{t}(H=T) < 1.$$

At the same time, edges adjacent to the boundary of probability q will change to a non-zero edge probability. As the updating

process continues, the information is diffused, as shown in Figure 14a, until it approaches a fixed point. All values of q other than 0 and 1 will not allow the initial pattern to remain stable.

Simple examination of the updating equation shows that a value of W = .5 will allow the edges in the loop to remain constant. A value of W > .5 will allow edges in the loop to increase, that is, $P^{t+1}(H=T) > P^{t}(H=T)$. In either case, however, the zero probability edges hanging off the boundary will still increase and eventually the information is diffused as shown in Figure 14(b) and (c).

Case 2: A uniform field of low equiprobability edges q.

In this case all edge locations have an identical context in the updating process. If q has a low value, one might consider it desirable for all edge probabilities to converge to a fixed point of 0. Unfortunately, our formulation, which has been guided by theoretical considerations, is not so well-behaved. Given the analysis of Case 1, let us set the value of P(H=T|V01) to .5 and determine the updated value at all edges as a function of q.

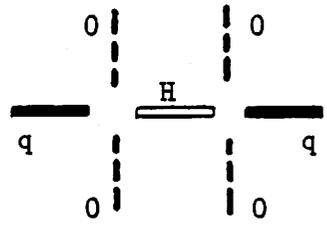


$P(V0) = 1-q$

$P(V1) = q$

$P(V2) = 0$

for left and right edges

| Vij | P(Vij) | P(H\|Vij) |
|---|---|---|
| V00 | $(1-q)^2$ | 0 |
| V01 | $2q(1-q)$ | w |
| V02 | 0 | 0 |
| V11 | $q^2$ | 1 |
| V12 | 0 | 1 |
| V22 | 0 | .5 |

$$P(H) = \sum_{i} \sum_{j \geq i} P(H|Vij)P(Vij) = q^2(1-2w) + q(2w)$$

$P(H) < q$ if $w < .5$

$P(H) = q$ if $w = .5$

$P(H) > q$ if $w > .5$

13(a)

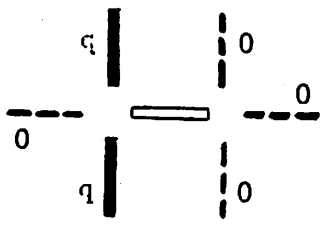


| left vertex | right vertex |
|---|---|
| $PL(V0) = (1-q)^2$ | $PR(V0) = 1$ |
| $PL(V1) = 2q(1-q)$ | $PR(V1) = 0$ |
| $PL(V2) = q^2$ | $PR(V2) = 0$ |

| Vij | P(Vij) | P(H\|Vij) |
|---|---|---|
| V00 | $(1-q)^2$ | 0 |
| V01 | $2q(1-q)$ | w |
| V02 | $q^2$ | 0 |
| V11 | 0 | 1 |
| V12 | 0 | 1 |
| V22 | 0 | .5 |

$P(H) = 2wq(1-q) > 0$ for $w > 0$ and $0 < q < 1$

13(b)

Figure 13. Case 1 Analysis. (a) Update equation for a central edge in a closed loop of edges with probability q. The vertex probabilities are used to compute the probability of V00,...,V22 edge contexts, from which the update equation can be obtained. (b) Analysis similar to (a) but for the edges adjacent to the closed loop.

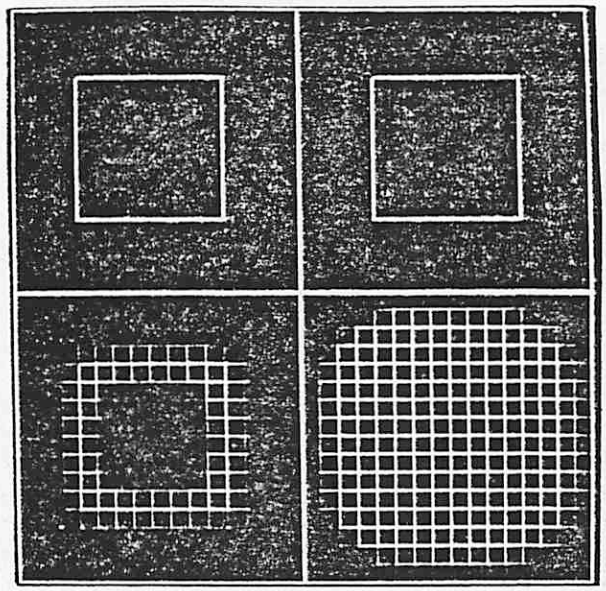

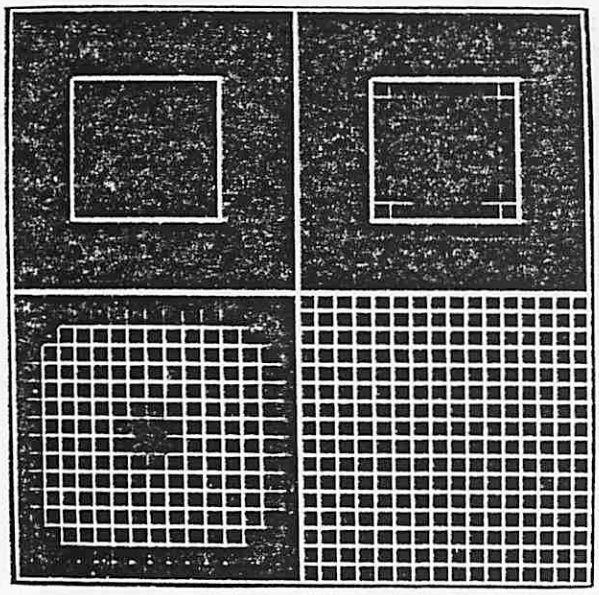

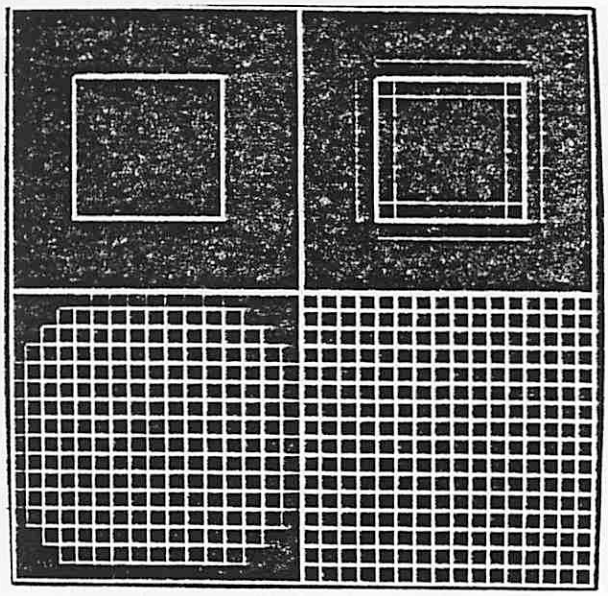

Figure 14. Case 1: Closed Loop of Edges. Each edge in the loop has an initial probability of .8 while background edges are of probability 0. Each image shows the original data (upper left corner) followed by the edge relaxation after iterations 2 (upper right), 5 (lower left), and 10 (lower right). The parameters used other than P(H/V01) were those given in Figure 12(b).

First, it is helpful to examine the probabilities of the four vertex types V0, V1, V2, and V3, given a set of three edges each of probability q. Figure 15(a) lists the vertex likelihoods for several low values of q. With a value of q = .2, the likelihood of V1 is relatively high: .384. Interestingly, this causes the updating of the central edge to be more heavily weighted by V01 than V00 because P(V01) > P(V00). Figure 15(b) depicts the relative contribution of each context Vij in updating the central edge to .424. Thus, in a single iteration a uniform field of edges with probability .2 would be raised to .424 and then even higher on the next iteration. This process converges to a fixed point about .65.

While this increase of low probability edges may disturb the intuition of some, the reason is that P(V00) = (1-q)**6 and for q = .1, .2, and .3, P(V00) equals .53, .26, and .12 respectively. Thus, a context of six low probability edges can lead to a relatively high likelihood of the V01 case, because it can become quite unlikely that all six edges are absent. The updating process increases the probability of the central edge in response to the many uncertain, but possible, ways of extending the boundary.

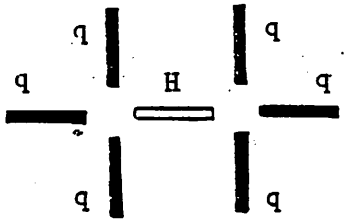


$$P(V0) = (1-q)^3$$

$$P(V1) = 3q(1-q)^2$$

$$P(V2) = 3q^2(1-q)$$

$$P(V3) = q^3$$

| | q = 0 | q = .1 | q = .2 | q = .3 |
|---|---|---|---|---|
| P(V0) | 1 | .729 | .512 | .343 |
| P(V1) | 0 | .243 | .384 | .441 |
| P(V2) | 0 | .027 | .096 | .189 |
| P(V3) | 0 | .001 | .008 | .027 |

(a)

For q = .2

| Vij | P(Vij) | P(H\|Vij) | P(H\|Vij)P(H\|Vij) |
|---|---|---|---|
| V00 | $(.512)^2 = .262$ | 0 | 0 |
| V01 | 2(.512)(.384) = .393 | .5 | .197 |
| V02 | 2(.512)(.104) = .106 | 0 | 0 |
| V11 | $(.384)^2 = .147$ | 1 | .147 |
| V12 | 2(.384)(.104) = .080 | 1 | .080 |
| V22 | $(.104)^2 = .001$ | .5 | 0 |

P(H) = .424

(b)

Figure 15. Analysis of Case 2: A uniform field of edges of low probability q. (a) Vertex probabilities as a function of q. (b) Contributions of each context Vij to the probability update of the central edge, assuming q = .2.

In order to examine the impact of these problems on an actual scene, the relaxation algorithm was applied to the bush subimage (Figure 1b). This subimage was chosen because it exhibits a range of characteristics including clear straight lines at several orientations, relatively weak edges in many places, texture edges, house edges in the textured areas, non-zero intensity gradients of several forms, etc.

The results of the edge update process after iterations 0, 1, 2, and 5, using the conditionals described in Figure 12b, are shown in Figure 16. The results are obviously poor, and improvements were not achieved by variation in the conditionals. Clear edge patterns immediately decay with diffusion into the surround, while there is little control in the growth of low values of edge probability -- which for convenience we shall call "noise" here even though the source of this data is varied. Attempted remedies for controlling the low probability edges via techniques involving fuzzy probabilities and thresholding also were not successful.

It is difficult to determine the source of the problems in the formulation. Clearly, the most obvious theoretical problems are the assumption of independence which allows the joint probability of the context to be decomposed, and the simultaneous

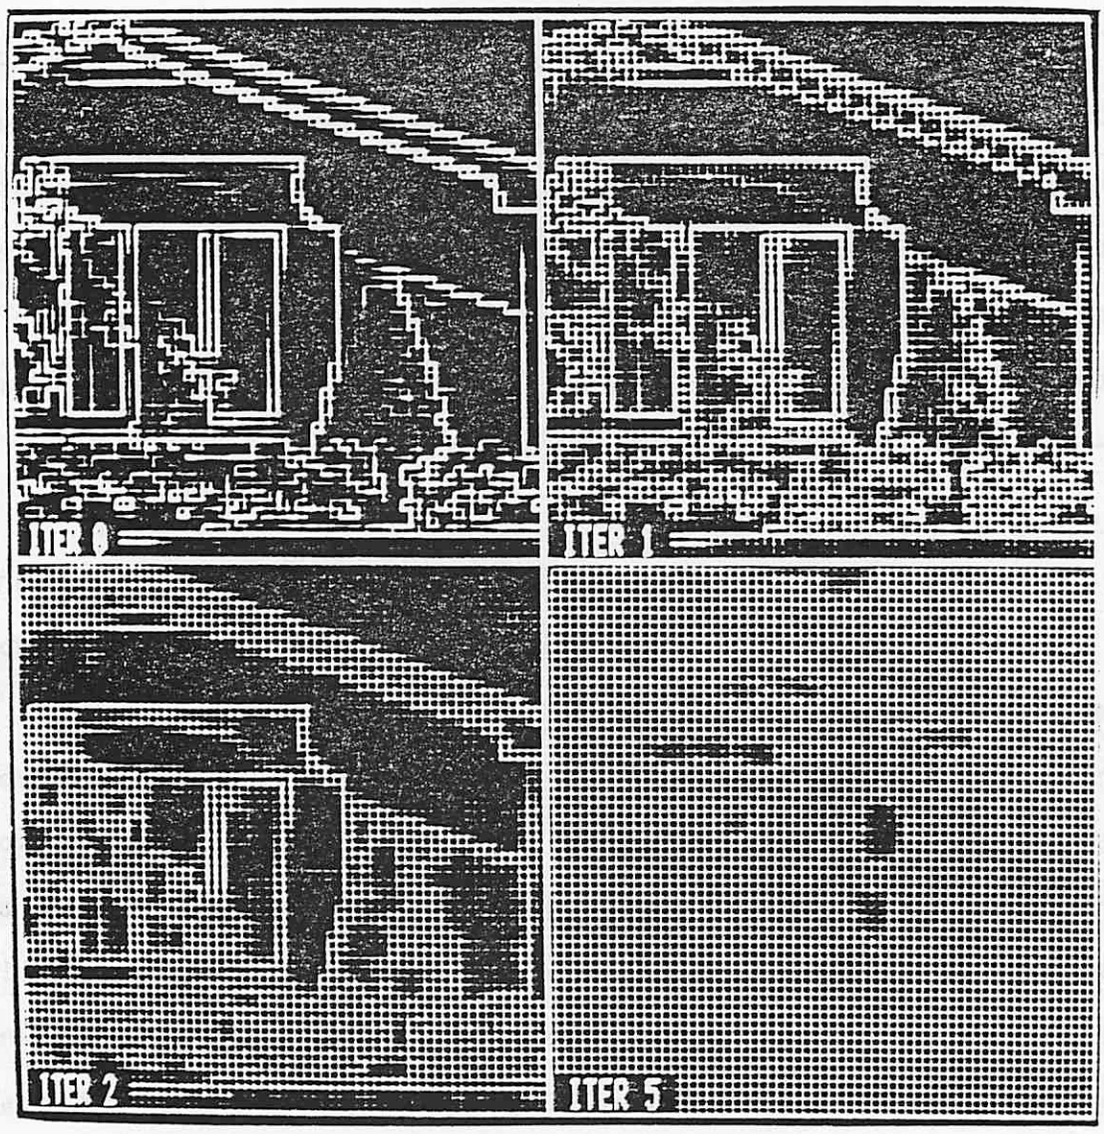


Figure 16. Results from Edge Update Process. (a) Initial probabilities. (b) After 1 iteration. (c) 2 iterations. (d) 5 iterations. By the second iteration, the diffusion of the original data is well underway. The uncontrolled growth continues in subsequent iterations until the original edge information is gone.

and iterative updating of all edges. Rather than question these assumptions upon which our basic approach rests, in the next two sections we incorporate additional forms of contextual information and explore variations in the updating equation as we seek to develop a formulation which performs as expected.

# III. PARALLEL EDGE SUPPRESSION AND CONTRAST DIRECTIONALITY

Up to this point, the definition of the local context to be used in updating was based primarily on the requirements imposed by the continuity assumption. The local context defined thus far is the smallest meaningful neighborhood for an edge which permits the translation of the continuity assumption into a set of constraints defined over the neighborhood. However, the presence of edges defined by non-zero feature gradients means that an extension of the local context to include the two adjacent edges parallel to the central edge (edges E7 and E8 in Figure 17a) will provide highly relevant information. These locations represent alternative placements for the collected gradient edges (Section II.3). Their absence makes it impossible to control the selection of that edge location (within the set of alternative choices of the gradient) which best supports continuity in the local context.

Without the expanded local context, in the current system all parallel edges could, and often do, find support in their context and grow to their maximum strength. The overall effect varies, depending on the situation. In the worst case, most non-zero edges in the image grow, resulting in a dense mesh of edges over much of the image, as shown in Figure 16. Some form

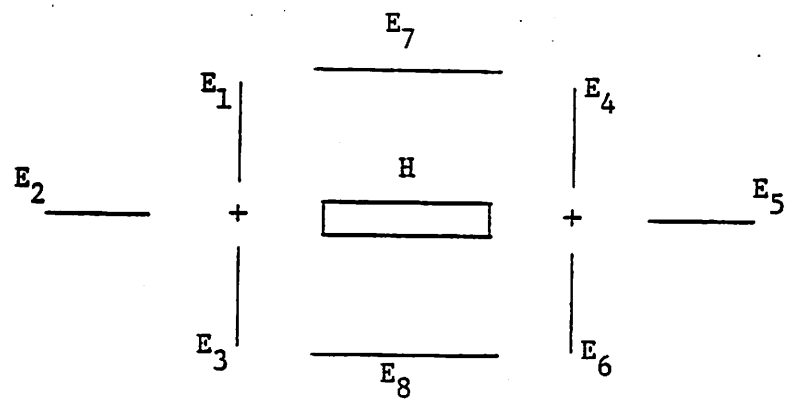


(a)

| | G-ON | | | | | G-OFF | | | |
|---|---|---|---|---|---|---|---|---|---|
| $E_7$ | S | F | S | S | 0 | F | F | 0 | 0 |
| $E_8$ | F | S | S | 0 | S | F | 0 | F | 0 |

F ≡ no-edge
S ≡ edge of same sign as central edge
0 ≡ edge of opposite sign as central edge

(b)

Figure 17. Extension of the Local Neighborhood. (a) The local neighborhood is extended to include the two edges E7 and E8 parallel to the central edge. This extension permits gradient information to propagate into the context. (b) By maintaining contrast directionality of the edges relative to the central edge, the effects of parallel edges on the central edge are easily decomposed into two classes G-ON and G-OFF. For case G-ON, the effect is inhibitory and suppression is on; for case G-OFF, the edges are parts of different boundary gradients and no suppression takes place.

of parallel edge suppression (or gradient suppression) via lateral inhibition is necessary to remedy this. Adjacent parallel edges can be structured to compete for survival, with the outcome decided by the support from their respective local contexts.

The process of parallel edge suppression must take into account the direction of contrast of parallel adjacent edges. Otherwise a one-pixel wide black stripe on a white background, for example, would cause the two distinct boundaries of the stripe to compete until only a single boundary remained, a result that is clearly undesirable. Consequently, parallel (and adjacent) edges should compete to become the representative of a given gradient edge, but this process should be deactivated if the parallel edges represent different boundaries.

The contrast directionality of edges in the initial data provides sufficient information for the decision of activation of parallel edge suppression for every pair of parallel edges. In fact it is precisely this feature of edges that is used in the gradient collection algorithm to group contiguous parallel edges (or equivalent runs of consecutive pixels) into units to be considered as potential gradient edges. Thus, the direction of feature contrast change will be maintained as a sign, "+" or "−".

The required properties of the edge suppression process are:

1. the central edge will be inhibited when edge E7 or E8 possesses the same sign of contrast change as the central edge;
2. parallel suppression will be decoupled and there will be no effect from edges E7 or E8 (or both) if they have a different contrast sign; and
3. the degree of parallel suppression will be directly proportional to the likelihood of the parallel competing edge(s).

For simplicity the effect of parallel edge suppression will be separated from the contextual influence of boundary continuity by introducing a heuristic assumption of independence between the parallel edges E7, E8, and the six continuing edges in Vij. Thus,

$$P(H=T) = \sum_i \sum_j \sum_{E_7} \sum_{E_8} P(H=T|V_{ij}, E_7, E_8)\, P(V_{ij}, E_7, E_8)$$

may be rewritten as:

$$P(H=T) = \sum_i \sum_j \sum_{E_7} \sum_{E_8} P(H=T|V_{ij}, E_7, E_8)\, P(V_{ij})\, P(E_7)\, P(E_8)$$

The key problem remaining is the determination of the conditionals relating the central edge to the eight edges in the context. One approach is to enumerate the 256 cases of the joint presence and absence of these edges and estimate the associated conditionals from the model of desired line drawings. The analysis is further complicated by the need to include the contrast directionality for controlling suppression of parallel edges, thereby significantly increasing the number of cases to be considered. The approach employed here is simpler.

First, let us note that the edges E7 and E8 can each be in one of three possible states: 1) no edge, labelled "F", 2) an edge of the same contrast direction as the central edge, labelled "S", or 3) an edge of the opposite contrast as the central edge, labelled "O". There are nine possible joint occurrences of the states of E7 and E8. These are divided into two groups O-ON and O-OFF, corresponding to the cases of parallel edge suppression being activated or deactivated, respectively, as shown in Figure 17(b). The contrast signs of edges are fixed during initial gradient collection, and consequently only some of these nine cases are relevant; in particular only one of the pair of signs (+, +), (+, -), (-, +), and (-, -) for any pair of edges can have a non-zero probability, while P(E7=F) for example can be non-zero no matter what the sign of edge E7. The probability of the nine

cases will vary depending upon the signs and magnitudes of P(E7) and P(E8), and can be computed as a product of the presence or absence of the two edges.

Inclusion of the parallel edges E7 and E8 in the local context of the central edge can be viewed in terms of G-ON and G-OFF instead of E7 and E8:

$$P(H=T) = \sum_i \sum_j [P(H=T|V_{ij}, \text{G-OFF}) * P(V_{ij}, \text{G-OFF}) + P(H=T|V_{ij}, \text{G-ON}) * P(V_{ij}, \text{G-ON})]$$

and by the independence of Vij, and E7 and E8

$$P(H=T) = \sum_i \sum_j [P(H=T|V_{ij}, \text{G-OFF}) * P(V_{ij}) * P(\text{G-OFF}) + P(H=T|V_{ij}, \text{G-ON}) * P(V_{ij}) * P(\text{G-ON})]$$

The required effect of cases G-ON and G-OFF suggest a rather simple means of integrating their effect into the conditionals. Case G-OFF should have no effect on the updating; therefore

$$P(H|V_{ij}, \text{G-OFF}) = P(H|V_{ij})$$

For case G-ON an assumption of conditional independence of Vij and G-ON given H=T leads to a simple embodiment of parallel inhibition.

$$P(H=T|V_{ij},\ G\text{-}ON) = \frac{P(V_{ij},\ G\text{-}ON|H=T) * P(H=T)}{P(V_{ij},\ G\text{-}ON)}$$

$$= \frac{P(V_{ij}|H=T) * P(G\text{-}ON|H=T) * P(H=T)}{P(V_{ij}) * P(G\text{-}ON)}$$

$$= \frac{P(H=T|V_{ij}) * P(H=T|G\text{-}ON)}{P(H=T)}$$

$$= P(H=T|V_{ij}) * \frac{P(H=T|G\text{-}ON)}{P(H=T)}$$

If we let

$$C = \frac{P(H=T|G\text{-}ON)}{P(H=T)} \qquad \text{where } 0 \le C \le 1$$

then, a setting of constant C to a value less than one has the desired effect of parallel edge suppression using

$$P(H=T|V_{ij},\ G\text{-}ON) = C * P(H=T|V_{ij})$$

for all vertex classes Vij. This parameter will allow variable control in the overall effect of the suppression mechanism. Thus, we get

$$P(H=T) = \sum_i \sum_j [P(H=T|V_{ij})\ C\ P(V_{ij})\ P(G\text{-}ON) + P(H=T|V_{ij})\ P(V_{ij})\ P(G\text{-}OFF)]$$

finally yielding

$$P(H=T) = [C * P(G\text{-}ON) + P(G\text{-}OFF)] * [\sum_i \sum_j P(H=T|V_{ij})\ P(V_{ij})]$$

Note that the effect of parallel edge suppression and boundary continuity have been decomposed into product terms. The second term is exactly the update before parallel suppression was considered, while the effects of parallel suppression have been isolated in the first term. The required effects of gradient suppression can now be attained. As condition G-OFF becomes more certain, the value of the first term approaches 1.0 which is the case without parallel suppression. Alternatively as condition G-ON becomes dominant because of likely parallel competing edges, the value of the first term goes to C. If C = 1, parallel suppression is turned off; for C = 0, maximal suppression is achieved. In practice a value of C = .5 appears reasonable given the particular gradient collection algorithm in use.

# IV. REFORMULATION OF THE LOCAL CONTEXT TO INCLUDE THE CENTRAL EDGE

In Section II of this paper the relaxation process for updating the probability of edges was motivated and guided by considering the Bayesian relationships of edges in a local context. Empirical results, however, were quite disappointing. One apparent difficulty to be explored in this section is the lack of direct influence of the current central edge probability on the updated central edge probability. While it is clear that the current probability of a central edge will have an indirect effect upon itself over successive iterations by influencing the edges in its context which, in turn, influence it, the updating process fails to capture meaningful information in a single updating of a local context. For example in the case of a boundary termination, the three edges which are absent, as well as the last edge actually present in the boundary, will be updated to similar values because all four of the edges see a VO1 context. There is a lack of discrimination between the last edge in the boundary and the adjacent absent edges. Much of the relevant information that is being ignored is available in the current probability of the central edge.

We will present two formulations that include the central edge directly in the updating process. They are both heuristic

as are most of the approaches in this paper when viewed from strict theoretical grounds. The first is of some interest because it is theoretically motivated by a loose re-interpretation of the Bayesian formulation, but it still does not produce effective results. The second formulation involves updating of both edge and no-edge labels with normalization. It retains the same values of conditional probabilities that were originally suggested theoretically, while capturing the effect of competition between labels via normalization. Highly effective results are achieved in the second formulation.

## Formulation A: Edges as Distinct Events Over Time

It is difficult to justify a Bayesian formulation which includes the central edge in the context unless we modify our basic definition of the events whose probability distributions are being manipulated. Up to this point we have associated a probability with each edge location, treating each edge as a particular event on the spatial array of image points. The problem then becomes one of updating the probability of these events so that consistency with the set of fixed conditional probabilities is achieved. Thus, one can view the updating process as adjustment (over time or, equivalently, iteration

number) of the original estimates of the probability of edge presence.

A different interpretation of these events may be obtained by viewing as distinct events edges which are in the same spatial location, but are updated over time. Thus, edges at time t and t+1 in the same spatial location will be considered distinct entities. However, local neighborhoods and conditional probabilities will still be used to relate events in successive iterations. The updating process becomes one of deriving a probability distribution of another set of edges which has been derived from the probability distribution of the original edge array. In fact, it is natural to think of both the initial and final probability distribution of edges to be available to further processing; the initial array represents sensory data while the final array represents a boundary continuity interpretation of that data.

The superscript t is now related to the random variable of the edges themselves instead of the probability of the random variable. The relationship between the probability of the central edge at time t+1 and the seven edges in our local context at time t is then:

$$P(H^{t+1}=T) = \sum_i \sum_j [P(H^{t+1}=T,\ H^t=T,\ V^t_{ij}) + P(H^{t+1}=T,\ H^t=F,\ V^t_{ij})]$$

$$P(H^{t+1}=T) + \sum_i \sum_j [P(H^{t+1}=T|H^t=T,\ V^t_{ij})\ P(H^t=T,\ V^t_{ij}) + P(H^{t+1}=T|H^t=F,\ V^t_{ij})\ P(H^t=F,\ V_{ij})]$$

Note that the same theoretical difficulty in estimating the joint probabilities of edges still persists. Only the unary marginals are available, yet we know that the edges in the context are not independent. Nonetheless, for lack of a better solution, an assumption of independence gives us

$$P(H^{t+1}=T) = \sum_i \sum_j [P(H^{t+1}=T|H^t=T,\ V^t_{ij})\ P(V^t_{ij})\ P(H^t=T) + P(H^{t+1}=T|H^t=F,\ V^t_{ij})\ P(V^t_{ij})\ P(H^t=F)]$$

$$P(H^{t+1}=T) = \sum_i \sum_j P(V^t_{ij})\ [P(H^{t+1}=T|V^t_{ij},\ H^t=T) * P(H^t=T) + P(H^{t+1}=T|V^t_{ij},\ H^t=F) * P(H^t=F)]$$

The key difference in this updating formula is that the conditionals are based upon seven edges in the context; this introduces a multiplicative factor involving the probability of the central edge on the previous iteration.

Figure 18 summarizes the new set of conditionals based on estimates from the desired line drawings discussed earlier. The six equivalence classes Vij have now become 12 because the

| $V_{ij}^{t}$ | $P(H^{t+1} = T \mid V_{ij}^{t}, H^{t} = T)$ | $P(H^{t+1} = T \mid V_{ij}, H^{t} = F)$ |
|---|---|---|
| $V_{00}^{t}$ | 0 | 0 |
| $V_{01}^{t}$ | 1 | 0 |
| $V_{02}^{t}$ | 0 | 0 |
| $V_{11}^{t}$ | 1 | 1 |
| $V_{12}^{t}$ | 1 | 1 |
| $V_{22}^{t}$ | 1 | 0 |

Figure 18. Estimation of Conditional Probabilities with Central Edge in the Context (Formulation A). The conditionals are derived from the model of desired line drawings described in Figure 11.

central edge may be present or absent with each. It is quite interesting to note that all conditionals are 0 or 1 in the seven edge context. It appears that inclusion of the central edge has greatly clarified the desired states in the updating process.

In the case of the V00, V02, V11, and V12 contexts, the conditionals are the same whether the central edge is present or absent. Thus, in these cases there is effectively no change in the updating because the central edge will factor out of the new updating formula. However, in the V22 case the values of 1 and 0 for $P(Ht+1=T|Ht=T, Vij)$ and $P(Ht+1|Ht=F, Vij)$, respectively, cause a linear interpolation between the values for $P(Ht=T)$ and $P(Ht=F)$. When $P(Vt22) = 1$ then

$$P(H^{t+1}=T) = P(H^{t}=T)$$

which appears to be exactly what we want. The central edge will stabilize at the current probability that was the best estimate based on local evidence; if there was little evidence for a bridging link between lines, the edge probability can remain low.

In the previous formulation, the V01 context also caused some difficulty in defining the appropriate conditionals. This was due to the symmetry in the cases where the central edge was present and where it was absent; it was impossible to

distinguish these cases, yet it was clearly desirable to do so. Inclusion of the central edge in the conditional allows the last edge present in the boundary to be distinguished from the three adjacent edges that are absent. Thus, a boundary termination in terms of the four edge probabilities is stable in the updating process when the boundary has probability one. There will be no diffusion of the information in the case of a boundary with edges of probability one that abruptly terminates.

Once again experimental results demonstrate that intuition can be quite deceiving, particularly in the case of parallel local updating of complex information. Figure 19 depicts several iterations of the seven-edge context using the conditionals of Figure 18 and no gradient suppression. The end result of this process is the diffusion and growth of edges across much of the image. Figure 20 depicts runs with moderate and strong parallel edge suppression, and while this algorithm appears marginally better than the basic approach of Section II and Figure 16, again no effective results are achieved. The dependency of results on particular values of conditionals was explored by varying several values of conditionals, but this parameter diddling yielded no significant improvement.

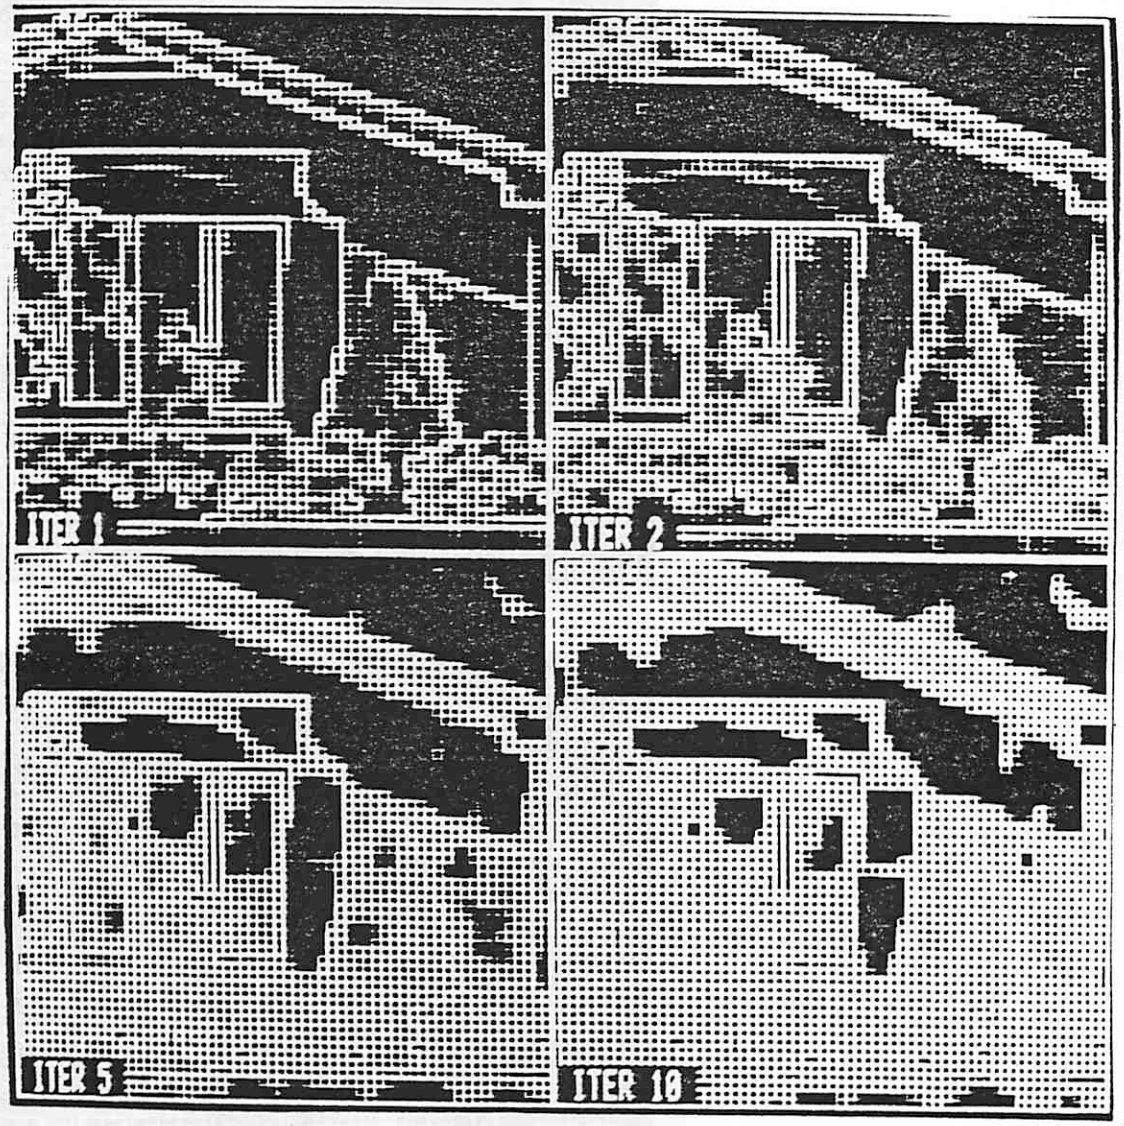


Figure 19. Edge Results for Formulation A: Inclusion of Central Edge. For these results the conditionals shown in Figure 18 were used and no gradient suppression was performed. The results are discouraging, particularly when compared to Figure 16.

(a)

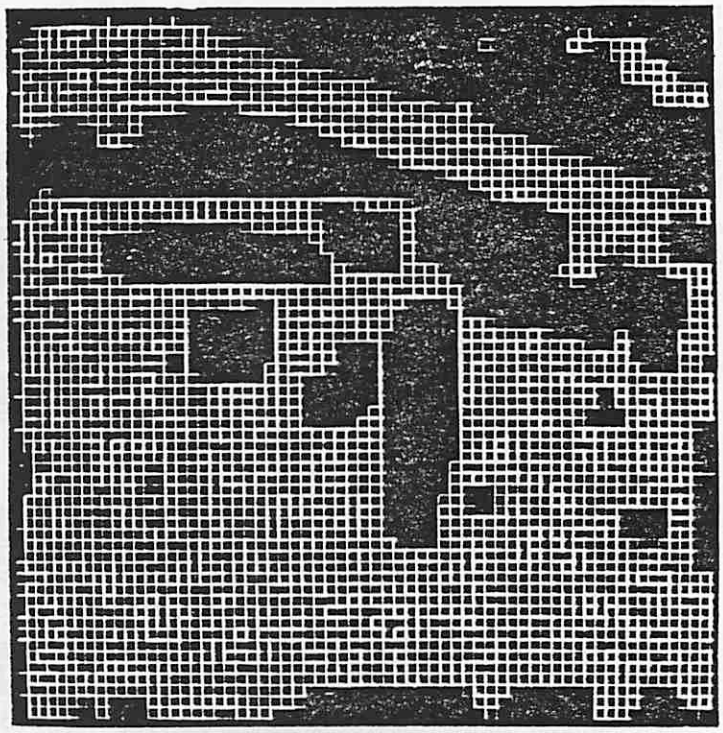

(b)

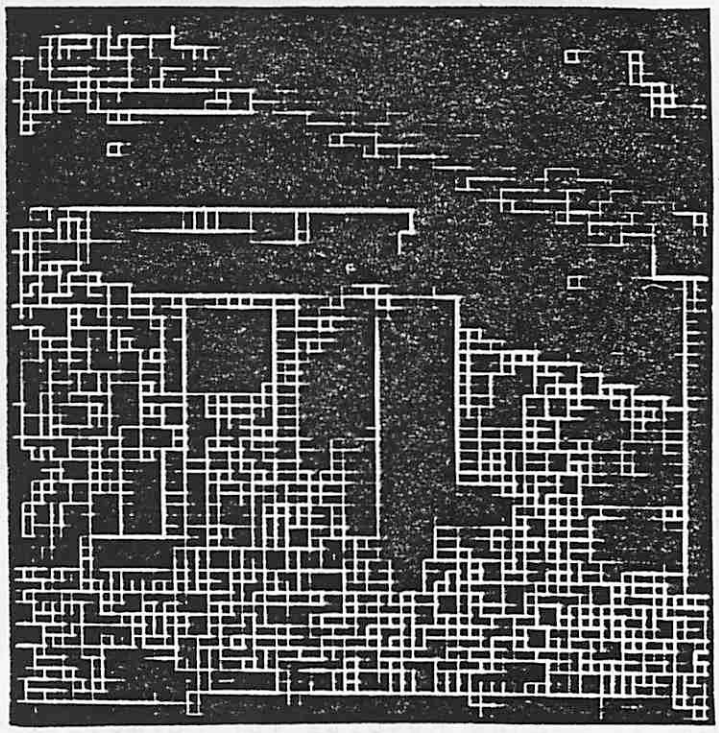

Figure 20. Edge Results for Formulation A: Inclusion of Central Edge. Results for (a) moderate and (b) strong parallel edge suppression after 10 iterations exhibit a similarity to those in Figure 19. Evidently the poor results are not sensitive to the degree of parallel edge suppression. Similar explorations of the dependency of the results on other parameter settings yielded no significant change in their quality.

Part of the problem of rapid growth of edges in Figure 19 suprisingly comes from the V12 case. Figure 21 is an analysis of the effect of updating given that one uncertain boundary at a probability of .5 is in the vicinity of a boundary whose probability is one. Analysis shows that the central edge will be updated to at least .5 no matter what the probability of the central edge is, and that the problem is due to the context of Ht=F and V12. The probability of the central edge is a factor in one term, but the probability of the absence of the central edge is the multiplicative factor in the other term, and therein lies the difficulty. Again we have a problem whereby the probability of the central edge is changing rapidly with little relationship to the value on the previous iteration.

## Formulation B: Updating Edge and No-Edge Labels with Normalization

The failure of Formulation A to produce acceptable results is discouraging. The probability of the central edge appears to change, almost independently of the actual data. Let us consider the second updating formulation obtained by treating the central edge and its context as two separate, independent sources of information. From this perspective, the information in the

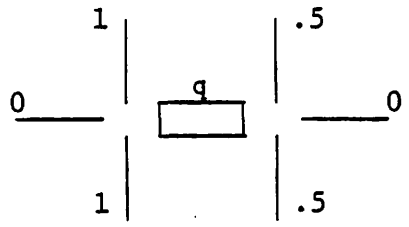


PL(V0) = 0 PR(V0) = .25

PL(V1) = 0 PR(V1) = .5

PL(V2) = 1 PR(V2) = .25

(a)

| $V_{ij}^t$ | $P(V_{ij}^t)$ | $P(H^{t+1}=T \mid V_{ij}^t, H^t=T) * P(H^t=T)$ | $P(H^{t+1}=T \mid V_{ij}^t, H^t=F) * P(H^t=F)$ | $P(H^{t+1}=T \mid V_{ij}^t) * P(V_{ij}^t)$ |
|---|---|---|---|---|
| V00 | 0 | 0 | 0 | 0 |
| V01 | 0 | q | 0 | 0 |
| V02 | .25 | 0 | 0 | 0 |
| V11 | 0 | q | 1-q | 0 |
| V12 | .5 | q | 1-q | .5 |
| V22 | .25 | q | 0 | .25q |

(b)

Figure 21. Formulation A: Analysis of the update to an edge of probability q between two parallel lines, one of probability 1 and the other of probability .5. (a) Neighborhood of central edge and vertex probabilities. (b) Contributions of each context Vij to the probability update of the central edge; note that the central edge is updated to at least .5, independently of its initial value.

surrounding context is combined in a multiplicative way with the information contained in the central edge itself. In this formulation it is difficult for either piece of information to completely override the other. The growth of weak edges is slower because the update increments are smaller; they receive only a portion of the support available from the context. Strong edges, on the other hand, for which there is evidence of support in the context, grow relatively faster. The conjecture is that the growth of edges will be more controlled and, where necessary, slow enough to allow evidence of support (or non-support) to propagate from "islands of reliability" [LES77] to local contexts.

Thus, Formulation B weights the contextual evidence by the current central edge probability to produce updated values P't+1(H):

$$P'^{t+1}(H=T) = P^t(H=T) * [\sum_i \sum_j P^t(H=T|V_{ij})\ P^t(V_{ij})]$$

and

$$P'^{t+1}(H=F) = P^t(H=F) * [\sum_i \sum_j P^t(H=F|V_{ij})\ P^t(V_{ij})$$

For clarity, we let

$$f = \sum_i \sum_j P^t(H=T|V_{ij})\ P^t(V_{ij})$$

and then rewrite the equations as

$$P'^{t+1}(H=T) = P^{t}(H=T) * f \qquad (7)$$

and $$P'^{t+1}(H=F) = [1-P^{t}(H=T)] * [1-f]$$

Clearly, these equations are not well-founded in probability theory and represent a different heuristic departure from the earlier formulation. In particular, since the updated central edge values P't+1(H) do not sum to 1, renormalization is necessary to obtain the final updated values:

$$P^{t+1}(H=T) = P'^{t+1}(H=T) \,/\, [P'^{t+1}(H=T) + P'^{t+1}(H=F)]$$

and

$$P^{t+1}(H=F) = 1 - P^{t+1}(H=T)$$

There is an interesting symmetry in the update equations (7). The current central edge probability and the contextual evidence really "weight" each other equally. Neither is favored by the formulation. This formulation is not subject to the problem noted in Figure 21 of Formulation A: edges between parallel edges are updated to a value which is a function of the probability of the central edge. Figure 22 summarizes the analysis of updating two parallel lines, one of probability 1. and the other of probability" .5, one pixel apart. Figure 22(d) shows that the central edge update in this case is very

close to initial probability and is much more reasonable than the results of Figure 21.

These results from Formulation B are encouraging and application to natural images shows the algorithm is effective. Figure 3 and Figures 23-26 summarize various aspects of these results. Figure 23 shows the necessity for gradient collection and suppression; 23(a) is the result of applying Formulation B in the absence of gradient collection/suppression. The failure of the edge process is most obvious in the roof trim and bush boundary. The gradients associated with the boundaries result in a series of "squares" or "rectangles" connected by single edges. Thus, the boundaries consist of many V22 vertices and as the iterations proceed, the edges surrounded by the V22 vertices eventually disappear, until by 20 iterations the boundaries are highly fragmented. Figure 23(b) shows the same edge process applied to the same image but in this case gradient suppression was included. The boundaries are quite good, although not all of the bush boundary and shadow boundary have been found.

The V01 conditional determines the amount of support the central edge gets from the context when continuation of the central edge occurs to one side only. Variation of this parameter has the effect of a "sensitivity" control over the

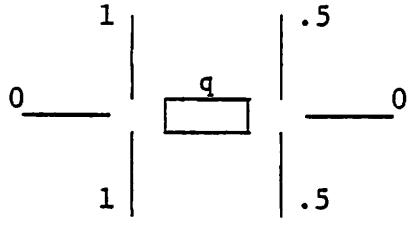


22(a)

| Vij | $P^t(Vij)$ | P(H=T\|Vij) | P(H=T\|Vij)*$P^t$(Vij) | P(H=F\|Vij) | P(H=F\|Vij)*$P^t$(Vij) |
|---|---|---|---|---|---|
| V00 | 0 | 0 | 0 | 1 | 0 |
| V01 | 0 | .25 | 0 | .75 | 0 |
| V02 | .25 | 0 | 0 | 1 | .25 |
| V11 | 0 | 1 | 0 | 0 | 0 |
| V12 | .5 | 1 | .5 | 0 | 0 |
| V22 | .25 | .1 | .025 | .9 | .225 |
| | | | u = .525 | | $\bar{u}$ = .475 |

22(b)

$$P'(H^{t+1}) = P(H^t) * u = .525q$$

$$P'(\bar{H}^{t+1}) = P(\bar{H}^t) * \bar{u} = .475(1-q)$$

$$P(H^{t+1}) = \frac{P'(H^{t+1})}{P'(H^{t+1}) + P'(\bar{H}^{t+1})} = \frac{.525q}{.475 + .05q}$$

22(c)

| q | $P(H^{t+1})$ |
|---|---|
| 0 | 0 |
| .03 | .033 |
| .25 | .269 |
| .5 | .525 |
| .75 | .768 |
| .97 | .973 |
| 1.0 | 1.0 |

22(d)

Figure 22. Formulation B: Analysis of the update to an edge of probability q between two parallel lines, one of probability 1. and the other of probability .5. (a) Neighborhood of edge of probability q. (b) Contributions of each context Vij to the probability update of the central edge. (c) Probability of central edge at next iteration expressed as a function of q. (d) Table of $P(H^{t+1})$ as a function of q; note that the update is much more reasonable than that shown in Figure 21.

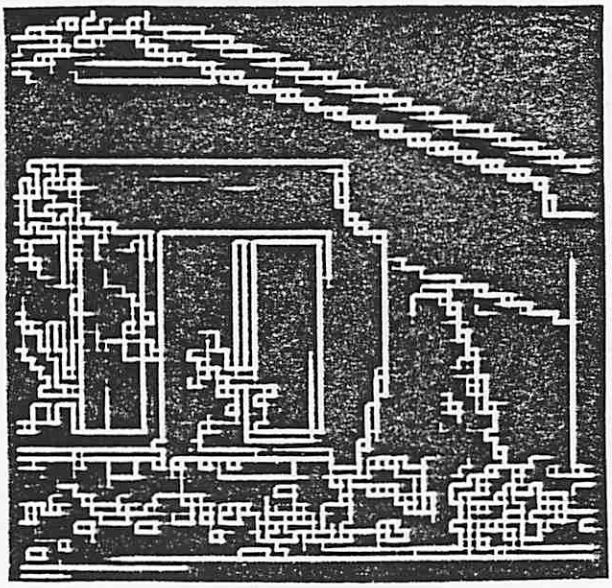

Iteration 2

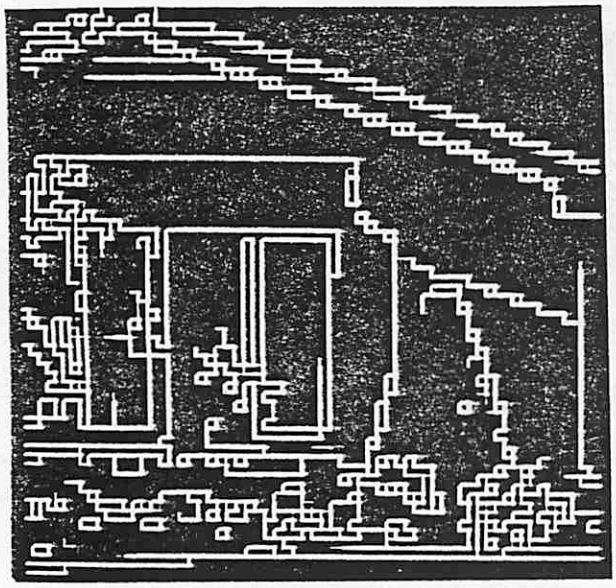

Iteration 5

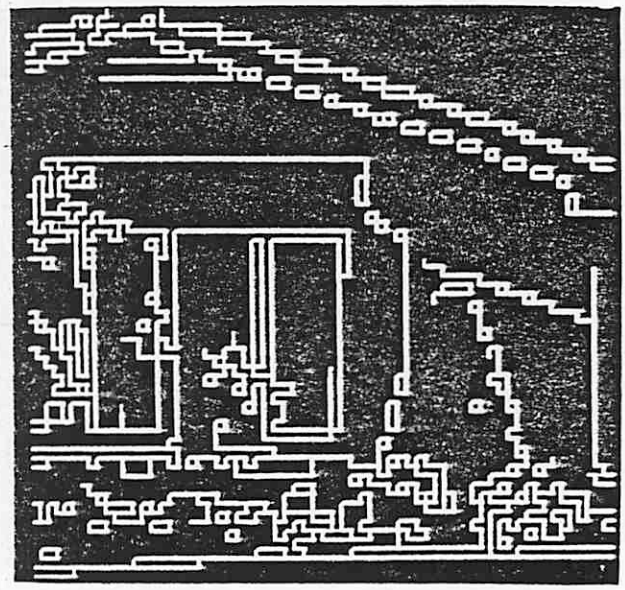

Iteration 10

Figure 23. Results from Formation B after 2, 5, and 10 Iterations. (a) No gradient suppression. (b) With gradient suppression. In each case, the conditionals used are those shown in Figure 12(b).

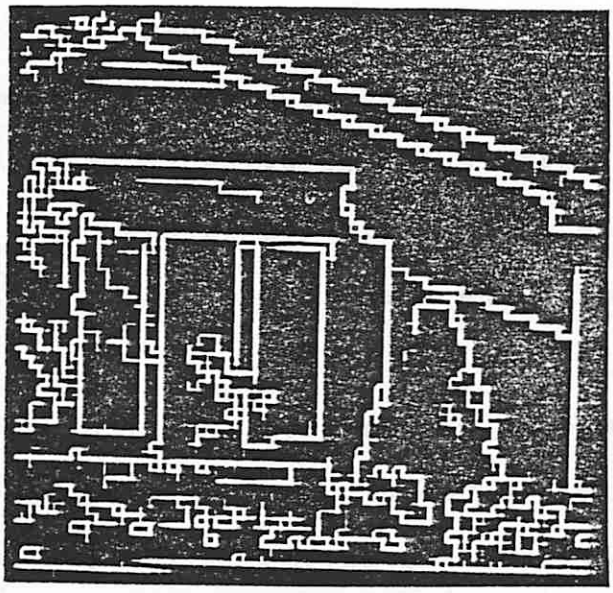

Iteration 2

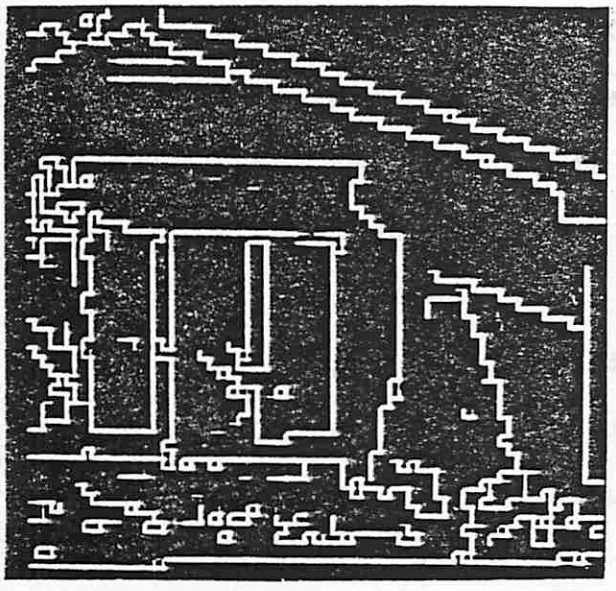

Iteration 5

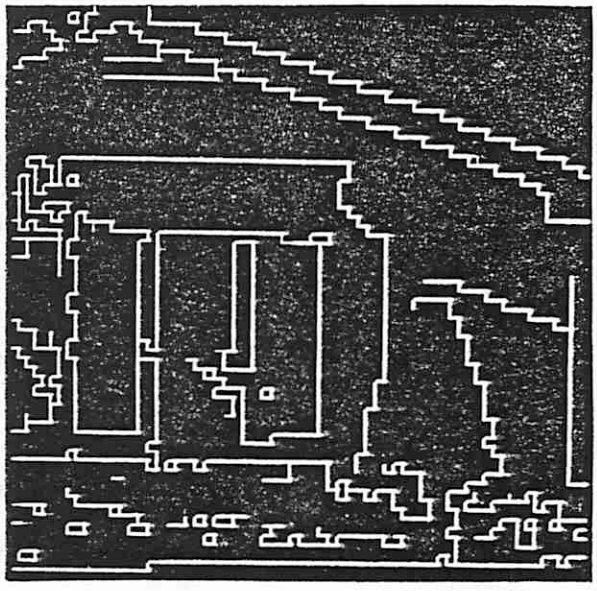

Iteration 10

Figure 23(b)

continuation edges which will eventually survive into the final boundary images. The effect is shown in Figure 24; note that for $P(H|VO1) = .67$, the bush boundary is completed but other "noisy" kinds of edges are also evident, notably in the tree area to the left of the windows, in the shrubs under the window, and in the right side of the bush.

The left boundary of the bush, particularly where it crosses the window trim boundary, is an interesting area and it highlights some lingering problems with the edge process, as shown in Figure 25. The intensity close-up in 25(b) clearly shows that the placement of the bush boundary is not at all obvious. The underlying problem with this area is that a gradient exists between the (white) trim and the interior of the bush (dark). Within this gradient are two real boundaries one pixel apart: the boundary between the trim and wall and the boundary between the wall and bush. Furthermore, the sign of the contrast of both boundaries is the same; both can be seen in Figure 25(d), which represents the initial edge probabilities. Figure 25(e-h) show the results after 1, 2, 5, and 20 iterations, respectively. The wall/bush boundary is suppressed by the wall/trim boundary because it is a much weaker boundary and is of the same direction. Once this possible connection is gone, the

(a)

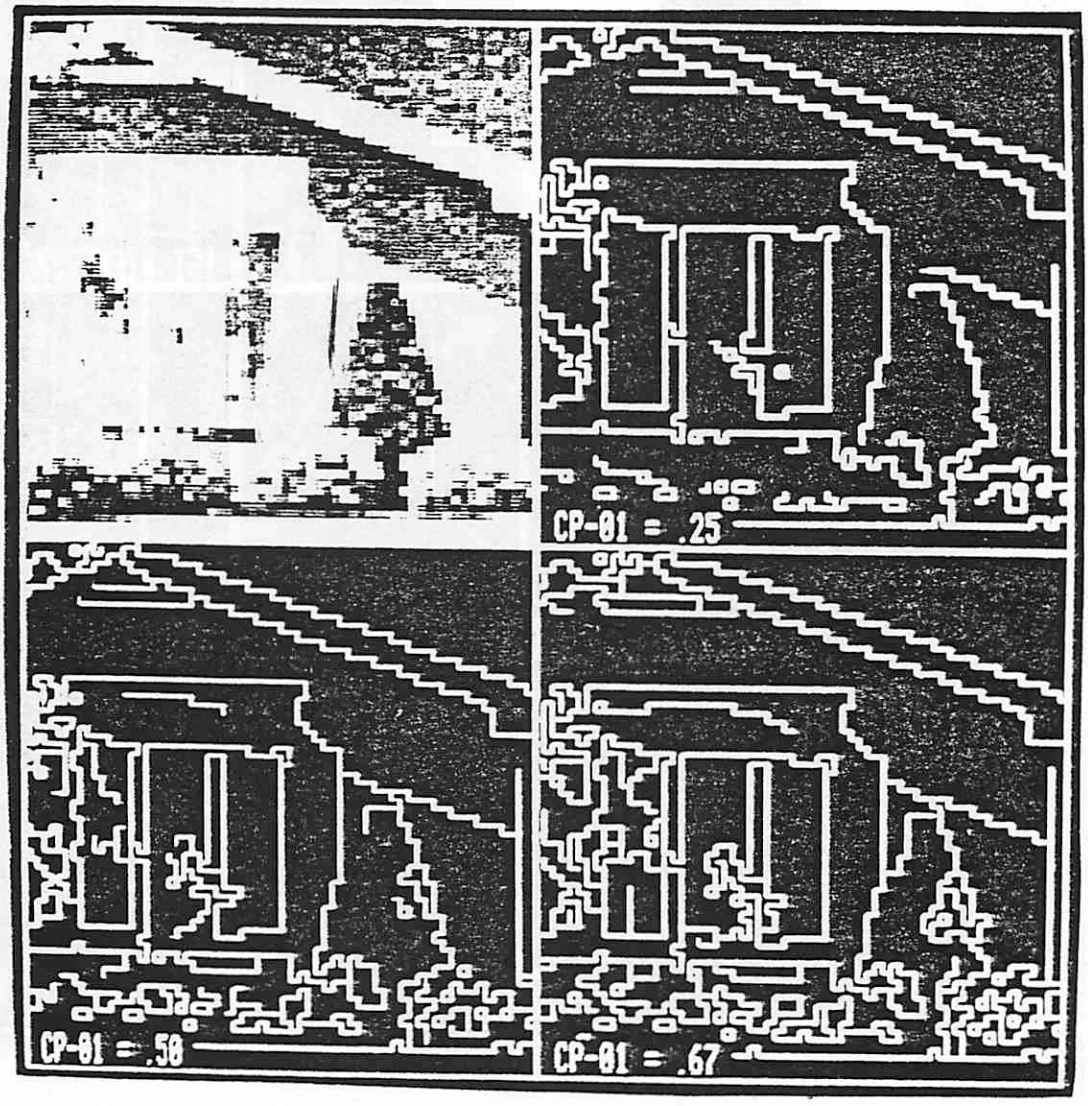


(b)

(c)

(d)

Figure 24. Results from Formation B of the Edge Process Using the V01 Conditional as a Boundary Sensitivity Control. (a) Original image. (b) P(H/V01) = .25. (c) P(H/V01) = .5. (d) P(H/V01) = .67. All results are obtained after 20 iterations of relaxation updating using the conditionals from Figure 12(b) (except for V01).

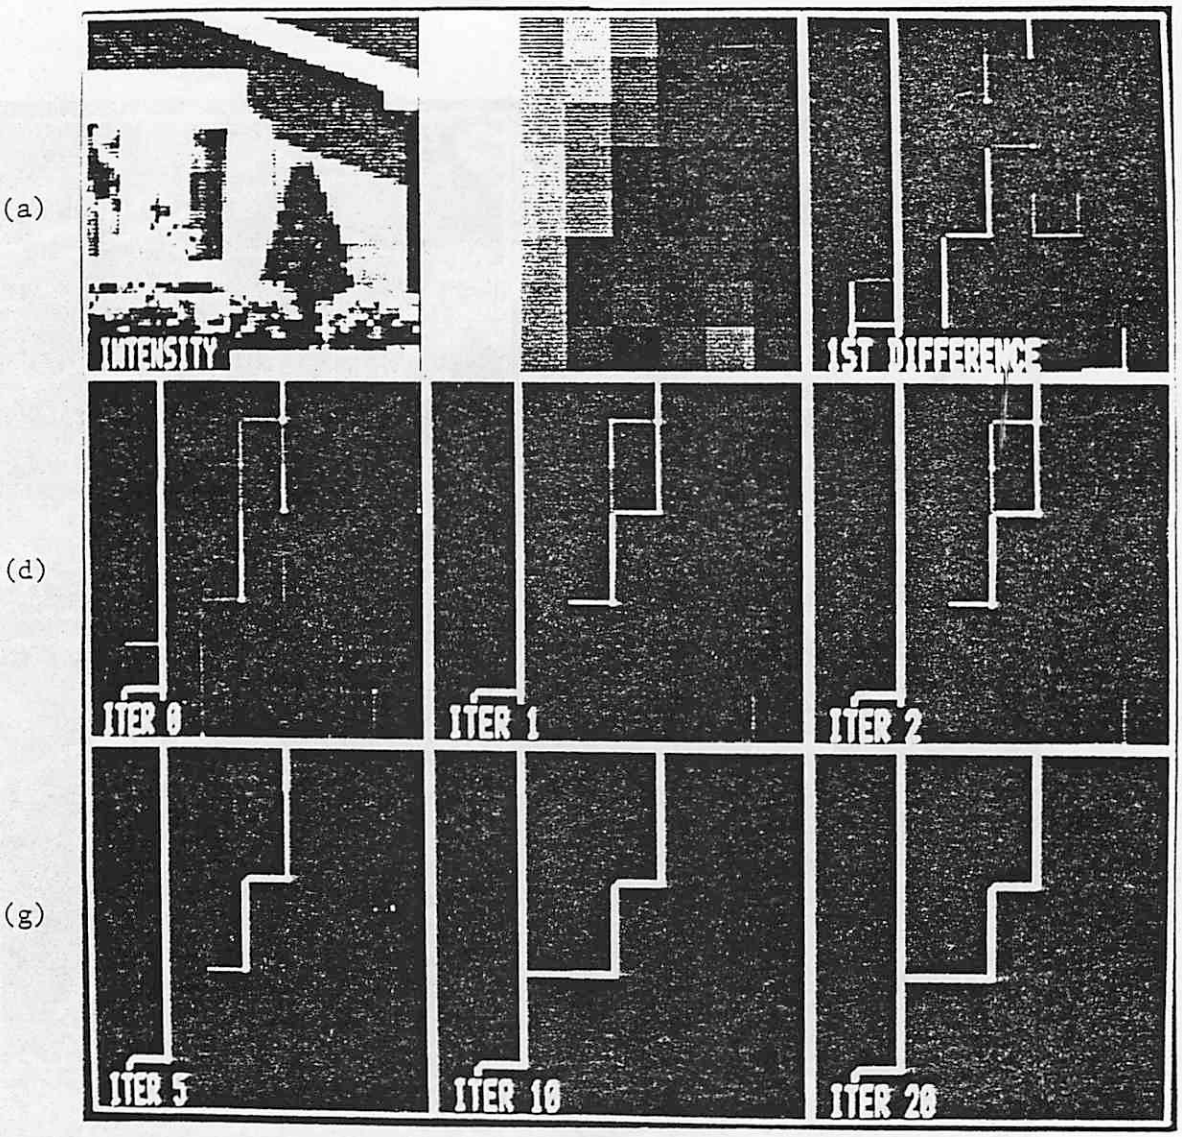


Figure 25. Results from Formulation B for Bush/wall/trim Closeup. (a) Intensity image of house showing location of bush/wall/trim subimage. (b) 8x8 blowup of bush area; note the gradient between trim/wall/bush in lower left quadrant. (c) Result of 1x2 mask; very weak edges, although not visible, are present. (d) Gradient collected edge probabilities. (e-h) Iterations 1, 2, 5 and 20 (respectively) of the edge relaxation process. The junction of the bush boundary with the trim/wall represents a situation in which two parallel boundaries of the same sign and one pixel apart should be maintained, but are not. The final results are still a reasonable approximation to the boundary.

context of the missing edge in Figure 25(g) switches primarily to V21 and the connection edge grows in (Figure 25h).

The loss of the wall/bush boundary is inevitable given the conditions in the image (parallel edges of same sign, one pixel apart) and the current version of the gradient collection/suppression mechanisms. These mechanisms assume that any non-zero gradient gives rise to at most one final boundary. The edge process then attempts to find the best placement of the individual edges making up the boundary. In this case it is doubtful whether any purely local process could do much better.

Figure 26 shows results from Formulation B on the second house image (Figure 1b). Two sequences are given, one for VO1 = .25 and one for VO1 = .67; within each sequence, results after 1, 5, and 10 iterations are shown. Again, the results are quite reasonable given the starting data. For the first time, an undesirable effect of local normalization over the 11x11 window is evident. There are many fairly weak edges within the left wall area to the right of the rightmost large window. Edges within the influence of the strong boundaries in this area (e.g., the right window edges) are assigned a low value because of the domination by the strong boundary. Edges in the central part of this area only see other weak edges and hence are assigned fairly

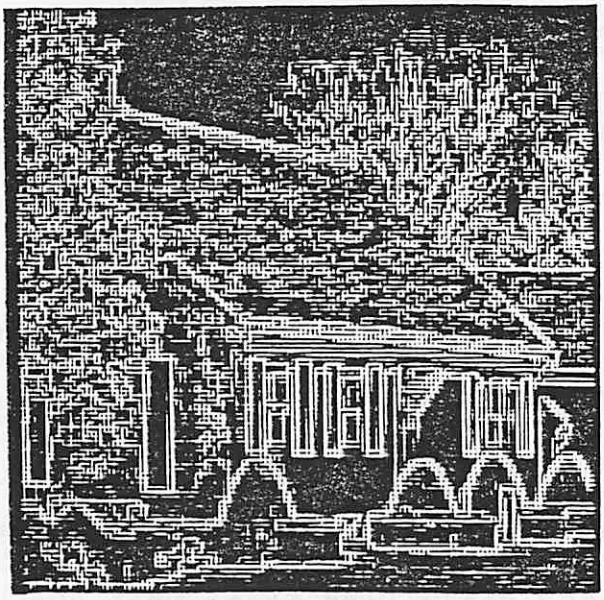

Iteration 1

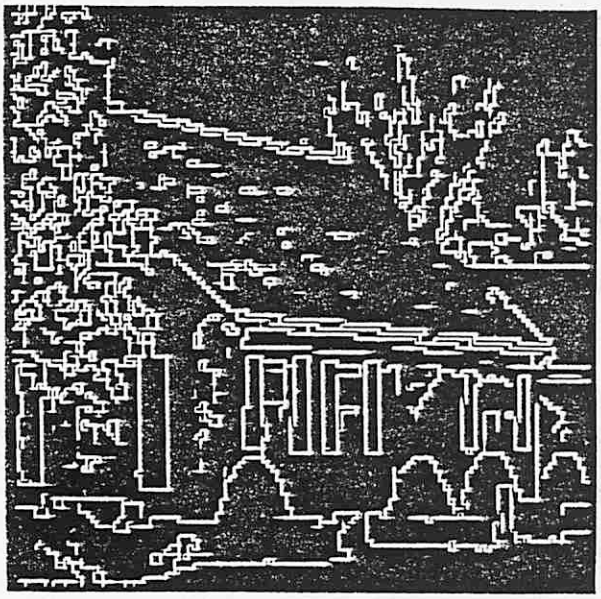

Iteration 5

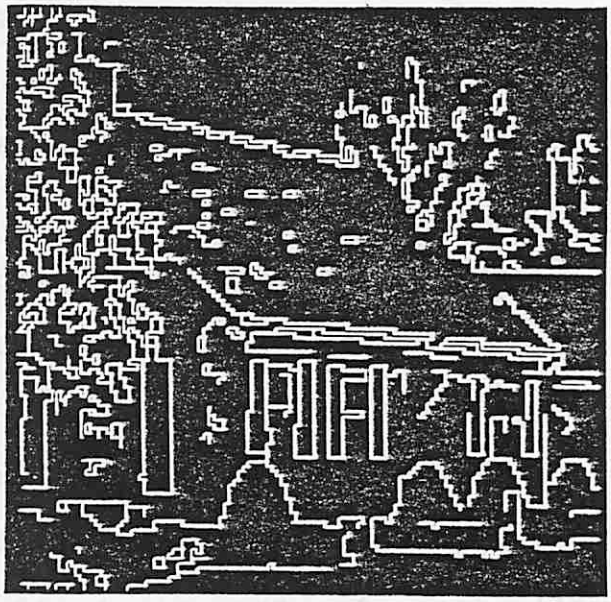

Iteration 10

Figure 26. Results from Formulation B for Second House Image. (a) Iterations 1, 5, and 10 for VOI = .25. (b) Iterations 1, 5, and 10 for VOI = .67.

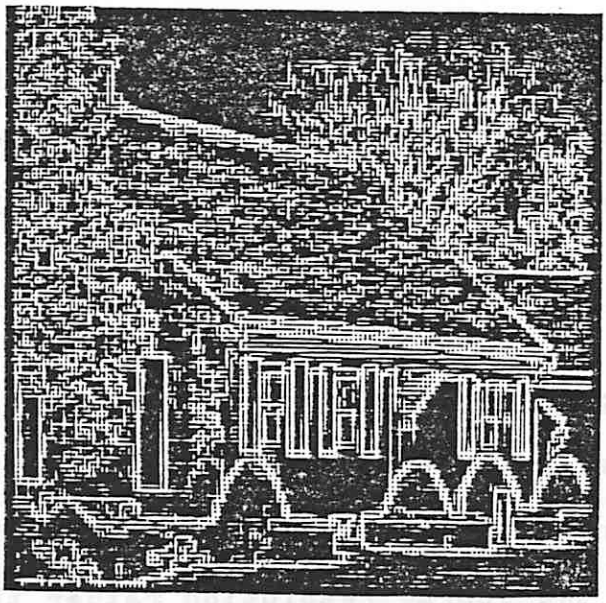

Iteration 1

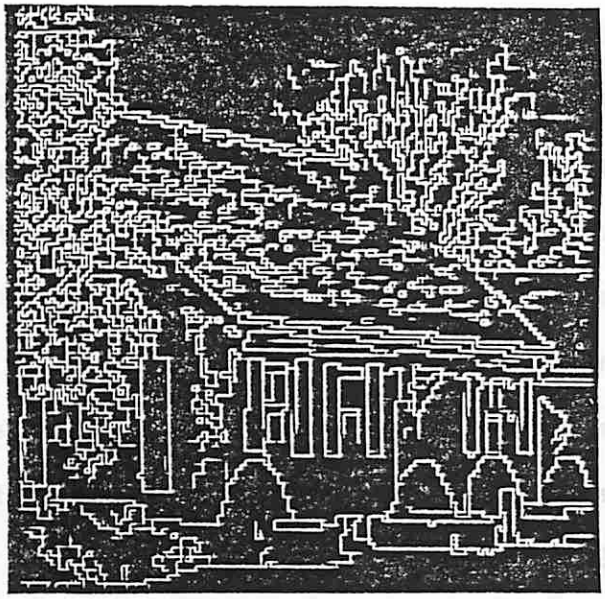

Iteration 5

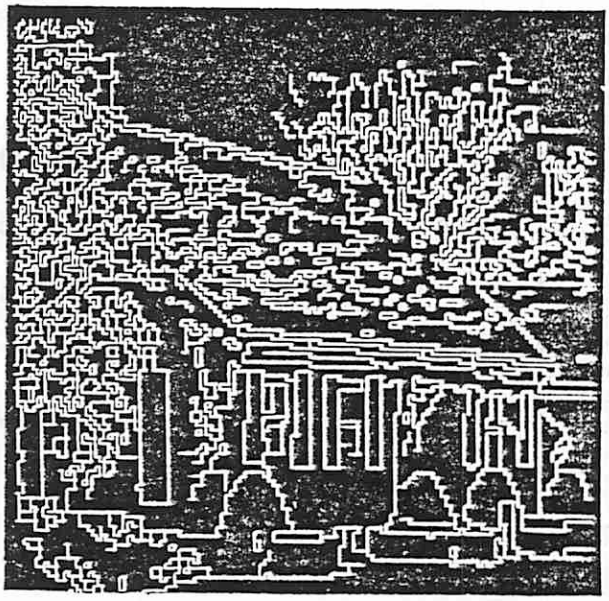

Iteration 10

Figure 26(b)

high values. This non-linear scaling gives rise to the vertical strip of edges in the center of the wall which can be seen in all the images. A similar, though less pronounced, effect can be seen in the roof.

There is a problem in extracting continuous boundaries for the dark stripe which forms the upper border of the roof. The stripe is fragmented into pieces, but this difficulty appears to be related to the coarse resolution rather than the quality of the algorithm. Any very narrow diagonal region can be expected to fragment.

The textural lines that form in the roof in 26(b) are reasonable in that they derive from the visual data. Other processes would be required to suppress this detail, although some of that effect takes place by running the algorithm to extract only the stronger boundaries as in Figure 26(a). However, that would not solve the problem of removing lines in a strong contrast brick wall. We leave that for higher levels of organization and processing [HAN78b].

Finally, if one examines results in close detail, there are strange effects which occur in some cases as boundaries approach each other. Other related problems might be the result of

attempts to extend incomplete boundaries after the initial weak contextual information has disappeared during relaxation. There are a number of plausible extensions to the current system which would provide improved edge fidelity; these are discussed in the conclusions.

# V. MEASURES OF PERFORMANCE

## V.1 Fixed Points, Entropy, and Consistency

The implication of our assumption that the conditionals remain constant is that any fixed point in the relaxation process is a state in which all edges are consistent with their contexts; i.e., all fixed points are consistent. It is of theoretical and practical interest to understand which states are fixed points and whether any given set of initial probabilities will converge to a fixed point. Rosenfeld et. al [ROS76] demonstrated that iterative linear updating functions will converge to a fixed point that is independent of the initial probabilities, a characteristic that obviously makes that form of processing useless. One solution is to replace the linear function with a non-linear one, although this requires a normalization process to restore the condition that $P^t(H) + P^t(H) = 1$.

The updating function, defined in equation 6, is linear in its combination of the conditional edge probability with its context. However, the classical relaxation process as described above treats each label in the context independently, that is each of the labels enters into a linear update of the central label. In the process described here, the labels in the context

are jointly used in the update function via the vertex probabilities. Viewed from the vantage of the individual labels at the various edge locations in the context, this is a highly non-linear process.

While we do not have a clear theoretical understanding of the properties of the relaxation process described here, it is obvious that there are a very large number of fixed points given the set of conditionals already defined. Consider any array of edge probabilities where all are zero except for closed loops of edges with probability one. The only edge contexts with non-zero probability are V00, V01 (for edge locations adjacent to the probability one boundary), and V11. The updating conditionals defined in Figure 12(b) will leave each edge unchanged; hence any array with only sets of closed loops of certain edges is a fixed point.

The array of probabilities has zero entropy when the probability of each edge in the array is 0 or 1. The closed-loop boundary examples just given are zero-entropy arrays which are fixed points in the updating process. However, not all zero-entropy arrays are fixed points, as in the case of a closed loop of probability one edges with one edge missing. Conversely, there are fixed points which do not have zero-entropy. It is

also easy to show that sets of conditionals exist which produce oscillation of zero-entropy arrays.

One can view the goal of the updating process to be the transformation of an array of edge likelihoods with non-zero entropy into a zero-entropy fixed point which is by some measure "closest" to the initial values. There are a variety of possible measures of closeness, such as the mean square difference between the initial and final arrays.

## V.2 Global Measures of Uncertainty, Drift, and Inconsistency

One of the key concerns of those interested in relaxation labelling algorithms is the lack of any criteria for global optimization. The problem of finding the minimum distance fixed point with zero entropy could be formulated as a linear programming problem on the set of edge values, but it does not appear computationally feasible when the number of such edges is of the order of 500,000. Instead we have chosen parallel mechanisms for organizing local contexts in a goal-oriented way in order to achieve a global organization of this information. This heuristic approach allows efficient local processing at the risk of producing a non-optimal global result. Therefore, it is

useful to have empirical measures of the global performance of the iterative updating process over time. The following measures defined across the W edges in our array will hopefully lend some empirical understanding of the processing.

The first measure, entropy, characterizes the degree of ambiguity or uncertainty in the edge array:

$$Q_1^t = -[\sum_{i=1}^{W} P^t(H_i)\log P^t(H_i)]$$

This measure is zero only when each edge in the array has probability zero or one. Although minimization of entropy is desirable in most cases, one must remember that most of the huge number of possible global zero-entropy cases will entirely violate the desired relationships embodied in the conditionals.

A more important measure is inconsistency:

$$Q_2^t = \sum_{i=1}^{W} [P^t(H_i) - P^{t-1}(H_i)]**2$$

This measure is the only one related to the conditionals which embodies the semantics of the domain. It is a function of the degree to which likelihoods of local edges are inconsistent with their context. Due to the nature of our updating formula, this measure follows directly from the change in the distribution of

likelihoods from one iteration to the next. Consequently, it is a measure of stability of the relaxation process, since it will reflect convergence of the distributions to a fixed point, as well as oscillatory behavior.

Finally we define "drift", a measure of the divergence from the initial distribution of edge probabilities:

$$Q_3^t = \sum_{i=1}^{W} [P^t(H_i) - P^0(H_i)]^{**2}$$

While one may choose the goal to be a set of likelihoods which is consistent throughout the network, the result should be a function of the initial data. It is desirable that the final result be close "almost everywhere" to the initial probabilities from which the updating has been driven.

Each of the three measures appears to focus upon a different aspect of the relaxation labelling process. Ideally, one would like the system to converge to a fixed point which simultaneously minimizes some function of entropy and drift. The manner in which this can be globally achieved via local updating processes is an open question, and in fact may only be possible in restricted situations.

# VI. CONCLUSION

The theoretical Bayesian formulation which is valid for a single edge and its context was expanded into a heuristic process for iteratively updating in parallel the array of edge probabilities. Unfortunately, experimental results demonstrated the ineffectiveness of this formulation. A variety of heuristic modifications, often suggested by theoretical considerations, were explored, finally resulting in an algorithm that performs in an effective manner on a number of very complex images. It is interesting that the process effectiveness was dependent upon a process of normalization which introduced competition between alternative labels.

## VI.1 Theoretical Specification of Conditionals

It is worthy of note that in all algorithms explored the values of the set of conditional probabilities that were derived from the model of desired line drawings has remained basically fixed. It has not been necessary to "tune" a set of weights for different images, or even for application of the algorithm to the first image. The theoretical estimates have worked very well once an effective form of the algorithm was developed, and at the

same time results were consistently bad for all variations in the parameters with the ineffective algorithms. Variation of one conditional probability in the final algorithm provides a mechanism for varying the contrast sensitivity of the boundaries extracted.

## VI.2 Model for Gradient Boundary

Another key aspect of our edge-boundary algorithm is the use of a model for the gradient of boundaries. The representative strength and location of an edge must be derived from a wide non-zero gradient across the feature values of pixels. Such gradient boundaries are typical in images and in many cases results which match human perception and intuition are dependent upon the extraction of the total contrast and placement of that edge strength within the gradient width. Further studies have been initiated to improve this model by using factors of boundary contrast, boundary width, and values of the first and second derivatives of the changes in the feature values.

## VI.3 Use of Boundary Width

The algorithm presented in this paper can be extended by employing additional local contextual information beyond the joint probabilities which reflect the likelihood of edge presence in the context. Local support via boundary continuity can be dependent upon the continuing edges having similar characteristics of edge width, with support decoupled to the degree that edge width of a pair of edges differs. This avoids the problem of labelling edge types, yet still takes advantage of edge widths without violation of Marr's principle of least commitment; it just requires the parameterization of width as a feature instead of classification into a small number of types of edge widths [MARR76].

It may be possible to take full advantage of all the information in a gradient, by integrating the boundary gradient model more fully into the updating process. Ehrich [EHR79b] and Quam [QUA78] have both employed measures of similarity of the intensity profiles of adjacent scan lines in the grouping of edges into boundaries.

## VI.4 Utility of Initial Sensory Data

There are other potentially useful mechanisms by which the local context can serve to further guide the boundary formation process. The characteristics of adjacent pixels to either side of a pair of edges can be compared so that similar region properties will lend further support. In this way the initial sensory data maintains an influence throughout the updating process.

Another problem is that by the time that contextual information propagates in from surrounding and less ambiguous areas, it is quite possible that ambiguous data will have disappeared during relaxation updating, leaving no trace of the original information. The initial sensory data can be used to alleviate this problem by restarting the iterative updating process after N iterations with a value at each edge H of MAX(PO(H), PN(H)). The effect is to allow boundaries, which have been organized by N iterations of relaxation updating, to influence the further organization of ambiguous edges without the problem of properly timing the propagation of this contextual information. Some initial experiments in this direction have shown this technique to be promising.

## VI.5 Variation of Contrast Sensitivity

An important aspect of the algorithm is the ability to vary the sensitivity of the edge contrast in forming boundaries. For example, the boundary of an object in the image may lie partially in shade causing reduced contrast between object and background, yet the continuity and consistency of that boundary make the boundary obvious. The algorithm must vary its sensitivity to edge strength. The conditional probability P(H=T/VO1) has been used to extend boundary terminations in an effort to extract less certain boundaries, and has yielded encouraging results. An alternative which ought to be explored is variation of the initial probabilities accomplished by means of changes in the value of K in equation 1. Reducing K has the effect of increasing the initial edge probabilities.

## VI.6 Larger Local Context

Finally, the use of a local larger context than the eight-edge context developed here would allow more effective patterns of boundary continuity to be employed. There would be more extensive overlap of adjacent local contexts, thereby giving a better perspective of the relationship between contiguous edges

in the consistent organization of continuous boundaries. The computational overhead, however, increases exponentially with the number of edges in the context and may not be feasible even for a context of edges of length two to either side. This would involve conditionals based upon 22 edges around the central edge, which is already beyond consideration via the approach developed in this paper.

# VII. BIBLIOGRAPHY

[BAR78] H. G. Barrow and J. M. Tenenbaum, "Recovering Intrinsic Scene Characteristics from Images," in Computer Vision Systems (A. R. Hanson and E. M. Riseman, Eds.), Academic Press, pp. 3-26, 1978.

[COO79] D. B. Cooper, "Maximum Likelihood Estimation of Markov-Process Blob Boundaries in Noisy Images," PAMI-1, No. 4, October 1979.

[DAV75] L. S. Davis, "A Survey of Edge Detection Techniques," Computer Graphics and Image Processing, 4, pp. 248-270, 1975.

[DUD76] R. O. Duda, P. E. Hart and N. J. Nilsson, "Subjective Bayesian Methods for Rule-Based Inference Systems," Proc. of the National Computer Conference, 1976.

[EHR79a] R. W. Ehrich, "One-Dimensional Edge Detection and Representation," Tech. Report CS79007-R, Dept. of Computer Science, Virginia Polytechnic Institute and State University, September 1979.

[EHR79b] R. W. Ehrich and F. H. Schroeder, "Contextual Boundary Formation by Scan Line Matching," Tech. Report CS79009-R, Dept. of Computer Science, Virginia Polytechnic Institute and State University, September 1979.

[FIS79] M. A. Fischler, J. M. Tenenbaum and H. C. Wolf, "Detection of Roads and Linear Structures in Aerial Imagery by Computer," AI Technical Note 200, SRI International, Menlo Park, California, October 1979.

[HAN74] A. Hanson and E. Riseman, "Preprocessing Cones: A Computational Structure for Scene Analysis," COINS Technical Report 74C-7, University of Massachusetts, Amherst, September 1974.

[HAN78a] A. R. Hanson and E. M. Riseman, "Segmentation of Natural Scenes," in Computer Vision Systems (A. R. Hanson and E. M. Riseman, Eds.), Academic Press, pp. 129-163, 1978.

[HAN78b] A.R. Hanson and E.M. Riseman, "VISIONS: A Computer System for Interpreting Scenes," in Computer Vision Systems (A.R. Hanson and E.M. Riseman, Eds.), Academic Press, pp. 303-333, 1978.

[HAN80] A.R. Hanson and E.M. Riseman, "Processing Cones: A Computational Structure for Image Analysis," in Structured Computer Vision (S. Tanimoto, Ed.), Academic Press, 1980.

[HOR77] B.K.P. Horn, "Understanding Image Intensities," Artificial Intelligence, 8, pp. 201-231, 1977.

[LES77] V.R. Lesser and L.D. Erman, "A Retrospective View of Hearsay-II Architecture," Proc. IJCAI-5, Cambridge, MA, August 1977.

[LOW80] J.D. Lowrance, "Dependency-Graph Models of Evidential Support," Ph.D. Thesis, COINS Dept., University of Massachusetts, Amherst, expected June 1980.

[MARR76] D. Marr, "Early Processing of Visual Information," Phil. Trans. Royal Soc. London 275 on Biological Sciences, October 1976, pp. 483-524.

[MART76] A. Martelli, "An Application of Heuristic Search Methods to Edge and Contour Detection," Communications of the ACM, V19, No. 2, February 1976.

[MON71] V. Montanari, "On the Optimal Detection of Curves in Noisy Pictures," Communications of the ACM, V14, No. 5, May 1971.

[NAG79] P.A. Nagin, "Studies in Image Segmentation Algorithms Based on Histogram Clustering and Relaxation," COINS Technical Report 79-15 and Ph.D. Thesis, University of Massachusetts, Amherst, September 1979.

[PEL78] S. Peleg and A. Rosenfeld, "Determining Compatibility Coefficients for Curve Enhancement Relaxation Processes," IEEESMC-8, No. 7, July 1978.

[PEL79] S. Peleg, "Ambiguity Reduction in Probabilistic Networks," Ph.D. Thesis, Computer Science Dept., University of Maryland, College Park, 1979.

[PRA80] J.M. Prager, "Extracting and Labeling Boundary Segments in Natural Scenes," IEEE Trans. Pattern Analysis and Machine Intelligence, Vol. PAMI-2, January 1980, pp. 16-27.

[QUA78] L. Quam, "Road Tracking and Anomaly Detection," Proc. of Image Understanding Workshop, May 1978, pp. 51-55.

[RIS77] E.M. Riseman and M.A. Arbib, "Computational Techniques in the Visual Segmentation of Static Scenes," Computer Graphics and Image Processing, 6, 1977, pp. 221-276.

[ROS71] A. Rosenfeld and M. Thurston, "Edge and Curve Detection for Visual Scene Analysis," IEEETC, 1971, pp. 562-569.

[ROS76] A. Rosenfeld, R.A. Hummel and S.W. Zucker, "Scene Labelling by Relaxation Operations," IEEESMC-6, No. 6, June 1976.

[YOR80] B.W. York, "Shape Representation in Computer Vision," Ph.D. Thesis, COINS Dept., University of Massachusetts, Amherst, expected June 1980.

[ZUC78] S.W. Zucker and J.L. Mohammed, "Analysis of Probabilistic Relaxation Labelling Processes," Tech. Report TR-78-3R, Dept. of EE, McGill University, January 1978.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE *(When Data Entered)*

| REPORT DOCUMENTATION PAGE | | READ INSTRUCTIONS BEFORE COMPLETING FORM |
|---|---|---|
| 1. REPORT NUMBER<br>COINS TR 80-11 | 2. GOVT ACCESSION NO. | 3. RECIPIENT'S CATALOG NUMBER |
| 4. TITLE *(and Subtitle)*<br>EDGE RELAXATION AND BOUNDARY CONTINUITY | | 5. TYPE OF REPORT & PERIOD COVERED<br>INTERIM<br>6. PERFORMING ORG. REPORT NUMBER |
| 7. AUTHOR(s)<br>Allen R. Hanson<br>Edward M. Riseman<br>Frank C. Glazer | | 8. CONTRACT OR GRANT NUMBER(s)<br>ONR N00014-75-C-0459 |
| 9. PERFORMING ORGANIZATION NAME AND ADDRESS<br>Computer and Information Science Department<br>University of Massachusetts<br>Amherst, Massachusetts 01003 | | 10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS |
| 11. CONTROLLING OFFICE NAME AND ADDRESS<br>Office of Naval Research<br>Arlington, Virginia 22217 | | 12. REPORT DATE<br>5/80<br>13. NUMBER OF PAGES<br>110 |
| 14. MONITORING AGENCY NAME & ADDRESS*(if different from Controlling Office)* | | 15. SECURITY CLASS. *(of this report)*<br>UNCLASSIFIED<br>15a. DECLASSIFICATION/DOWNGRADING SCHEDULE |

16. DISTRIBUTION STATEMENT *(of this Report)*

Distribution of this document is unlimited.

17. DISTRIBUTION STATEMENT *(of the abstract entered in Block 20, if different from Report)*

18. SUPPLEMENTARY NOTES

19. KEY WORDS *(Continue on reverse side if necessary and identify by block number)*

segmentation
relaxation
edges
boundaries
scene analysis
image analysis

20. ABSTRACT *(Continue on reverse side if necessary and identify by block number)*

Many image analysis tasks require the construction of a boundary representation as a means of partitioning an image. This paper develops a parallel relaxation algorithm for updating initial heuristic estimates of the likelihood of edges so that continuous boundaries are formed. Bayesian probability theory is used to analyze the probability updating of a single edge based upon the joint probabilities of the edges in its local surrounding

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-014-6601

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE *(When Data Entered)*

20.

context. The relationships between edges, sometimes referred to as compatibility coefficients in relaxation algorithms, are embodied as conditional probabilities between the central edge and the context of edges. The set of conditional probabilities are theoretically derived from a model of desired line drawings that satisfy basic notions of boundary continuity. The local updating function attempts to drive the likelihood of each central edge into consistency with the surrounding context.

Experiments involving the iterative parallel application of this non-linear Bayesian updating function to all edge probabilities demonstrates serious problems in the formulation. A variety of heuristic modifications, guided by theoretical considerations, are examined empirically. The final formulation is an algorithm that performs in an effective manner on several very complex images.