

REGION GROWING IN TEXTURED OUTDOOR SCENES*

By

A. R. Hanson*
E. M. Riseman†
P. Nagin†

COINS Technical Report 75C-3
February 1975

This research was supported by the Office of Naval Research under
Grant ONR N00014-67-A-0230-0007

* School of Language and Communication
Hampshire College
Amherst, MA 01002

† Department of Computer and Information Science
University of Massachusetts
Amherst, MA 01002

*This paper is the Proceedings of the 3rd Milwaukee Symposium on Automatic
Computation and Control, pages 407-417.

REGION GROWING IN TEXTURED OUTDOOR SCENES

Allen R. Hanson
School of Language & Communication
Hampshire College
Amherst, Mass. 01002

Edward M. Riseman
Computer and Information Science
University of Massachusetts
Amherst, Mass. 01002

Paul Nagin
Computer and Information Science
University of Massachusetts
Amherst, Mass. 01002

Abstract

Some progress has been made in scene analysis in finding boundaries between textured objects. However, standard line-finding and region growing algorithms have not been able to deal with strong local variations of intensity and color that repeat themselves more globally in images of natural scenes. These problems are significant with respect to both micro-texture, say the leaves of a tree, as well as macro-texture, the light green of leafy branches vs. dark green or black shadows between branches in a tree. How can local variations be globally bound together into a cohesive region for perceptual identification.

The approach that is being examined involves the reduction in size of the image while extracting features such as intensity, color, and color variation in local windows across the image. Two-dimensional histograms are being used to determine clusters of points with similar characteristics. Major clusters are labelled and the corresponding points in the reduced image are also labelled. This can directly detect regions of homogeneous micro-texture. Now an adjacency matrix of $N \times N$ clusters will denote the types of image points adjacent to each other. Peaks in the adjacency matrix can determine macro-texture.

All of this information can be used to grow regions in the reduced image and then in the detailed image. Experimental results on color images of natural outdoor scenes will demonstrate our results.

I. INTRODUCTION

One of the primary problems in scene analysis of natural two-dimensional images is the segmentation of the image into meaningful regions. Usually, the goal of such low level processing is to define regions whose identity can be used to form conceptual models of the world being viewed. Current region growing techniques work well in ideal environments, i.e., images in which regions are basically homogeneous with small variations in intensity (little texture). For this reason, images with highly textured regions containing variations in texture coarseness, di-

rectionality, and homogeneity are usually not examined. A notable exception to the constrained environment in which region growing takes place is the work of Yakimovsky and Feidman [1]. Extensive semantic direction for merging small regions was successfully applied to outdoor road scenes. However, we believe that the decision-theoretic techniques cannot deal with the wide variety of textures found in most natural scenes. Few researchers have examined the texture of a tree: in summer both the micro-texture of leaves and the macro-texture of shadows between leafy branches; in winter the micro-texture of webbed branches and the macro-texture of

blue sky showing through the branches.

Both region growers and line finders are subject to similar difficulties due to the variety of textures which may be present. Texture may involve a simple statistical distribution (in 2 dimensions) of intensity; it might involve a mixture of two or more colors (hues); there might be a structural relationship between the size and adjacency of several micro-textures in forming a macro-texture; in fact, there could be a variation in saturation as exhibited by clouds in a wispy sort of white and blue sky. We are therefore using the term "texture" in a wide sense; it may involve variation in any measurable perceptual attribute. When a local boundary is found is this part of some global boundary that separates distinct regions or is this local texture that is a defining characteristic across this region? The line finder does not know whether to track the local boundary; the region grower does not know whether to jump the boundary. Although algorithms have been suggested, none have proven to be sufficiently reliable.

The problems we wish to address ourselves to are:

- a) what features of texture should be extracted in order to deal with natural scenes.
- b) how can local characteristics of texture elements of varying color, size, and shape be globally bound together into a perceptual region. And
- c) how can an algorithm deal flexibly with micro- and macro-texture.

Previous research has examined a number of characteristics of texture. Rosenfeld [2, 3] has examined texture in order to find boundaries between regions of distinct texture by looking for variations in either average gray level or edge/unit area (i.e. average amount of variation). This is one of the few attempts to deal with micro- and macro-texture simultaneously. He has also worked on measures of texture orientation and coarseness, but the techniques have not been applied to the full complexity of natural scene analysis.

Bullock [4] has recently carried out an excellent comparison of several edge detection operators. He presents extensive experimental results showing the confusion between detection of "micro-structure edges" of texture elements, and "macro edges" of major surface boundaries. This work seems to have goals similar to those reported here. However, we believe the extraction of multiple features of a region may prove more powerful than the edge characteristics.

Haralick [5] has examined micro-texture using spatial gray level adjacency matrices. Here an adjacency matrix of gray level i vs. gray level j for all gray levels can characterize various types of texture.

Bajcsy [6] has applied Fourier analysis to extract characteristics of texture orientation and texture gradients in natural scenes. However, these descriptors can only be obtained when operating on high resolution data and are primarily concerned with micro-texture. It is our contention that in fact it is often macro-texture that is most useful in perceiving the world. In most cases, it is not the size or shape of an individual blade of grass or leaf that permits us to see a lawn or the crown of a tree, but rather the global interaction of many of these micro-structures. Thus, we seek more flexible measures on coarser data, both for the computational practicality as well as utility. Consequently, we have chosen measures similar to those of Rosenfeld but generalized to color space. We have embedded these in a unique computational structure for adapting to any distinctly perceivable type of macro- and micro-texture.

II. Perceptual Motivation for our Approach

A region that is perceived to be cohesive is bound by some features which are relatively consistent across the region. One of these features might itself be a measure of the amount of local variation. Thus, a smooth blue sky has a smaller value for the attribute "average variation" than a roughly textured region. Let us assume for the moment that we have a number of features which can be applied to a local window. Further, we assume that this window is sufficiently large so that in a region with homogeneous texture (not necessarily color or intensity) adjacent local windows would have approximately the same amounts of variation within them. Thus, we hypothesize that perception of a region as some cohesive entity requires the repetition of similar micro-texture across many local windows, although these windows need not be contiguous. This is a sort of statistical property requiring a certain frequency of distribution of a large number of these windows with similar properties.

The recognition of a macro-texture requires a distinctly different kind of perspective, one that analyzes certain structural relationships between types of micro-texture;

* Note that variance of these values would become greater as the window is reduced in size.

in particular, we will confine ourselves to the adjacencies between these types although other structural relationships could be examined. There is the trivial case of a large homogeneous region of micro-textured windows (call them type I) adjacent to each other; for example, a large expanse of blue sky has many type I windows of texture which are adjacent to each other.

On the other hand, an area of blue sky highly interspersed with green foliage should have many local windows of two types: smooth blue (type I) and roughly varied green (type II). Since there are many of each, by our first premise both will be noticeable. By our structural mechanism, many of these pairs are noticed to be adjacent and if these pairs occur contiguously across some large region, this will be detected as a macro-texture. We conjecture that at first glance people examining an image do not focus attention upon small textured areas or perceive such a macro-texture unless there are other prominent features of these regions. Thus, if only a small part of a sparse tree with blue and green is in the picture, the macro-texture may not be picked out and immediately recognized.

III. Computational Structure: Preprocessing Cones

This work is part of an attempt to develop a scene analysis system, called VISIONS (Visual Integration by Semantic Interpretation of Natural Scenes), which will build conceptual models of 2D color images of natural outdoor scenes [7-9].

In this paper, we will refer only to the front end of the system, the processing cone. It consists of a parallel computational structure which transforms and reduces large amounts of visual data in a layered fashion as shown in Figure 1. Information flows up, down, and laterally within the cone by defining local parallel functions which are duplicated to operate across the entire array. Parallel line finders, region growers, texture analyses, and color mappings, among others, operate on a 256^2 grid of raw image data and reduce it layer by layer to a 16^2 grid, and in some cases even to the 1×1 level, which contains information extracted from the entire scene.

There are three basic forms of processing available in the cone. The first is reduction as shown in Figure 1 which allows a function of a local window to pass a value to the next layer; thus, at each level a cell could have a value stored that is a function of the subcone below it; in this way, spatial relationships of extracted fea-

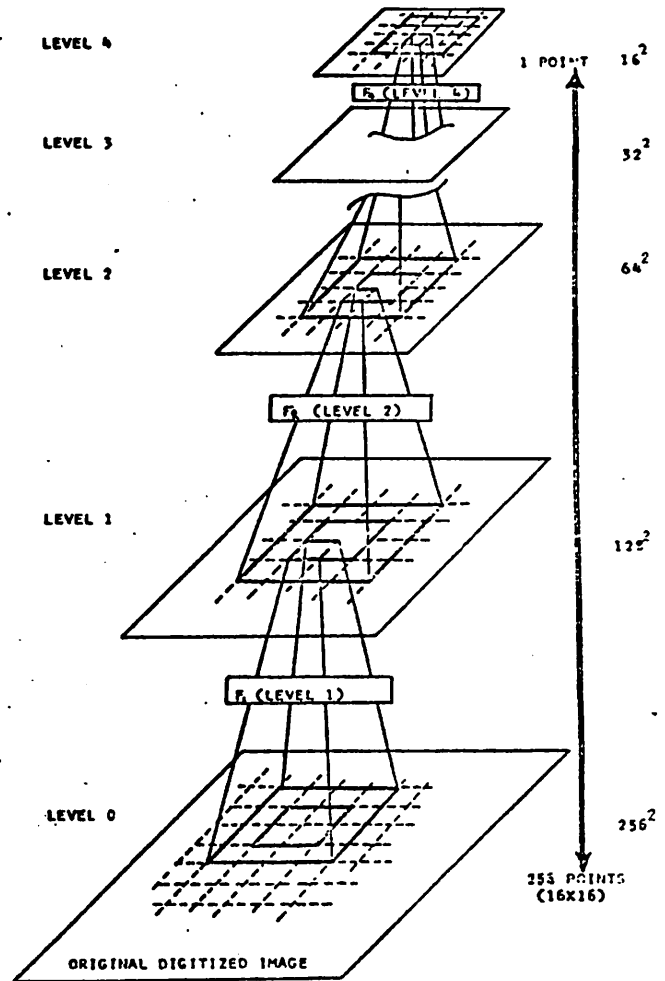



Figure 1. PROCESSING CONE

tures are preserved. An iteration process allows the function of a local 5×5 window to be mapped into the center cell of the window. The last type is called projection and allows information to be passed down to a cell from the unique ancestral cell at each level above it. At any particular time, only a single layer of functions are being computed, but it is being applied uniformly across a layer with each local function applied to its local window. It should also be noted that each cell has a memory of several words so that a number of values can be saved for one or more processes.

Algorithms for extracting features in the cone consist of a sequence of parallel reductions, iterations, and projections, possibly interleaved with sequential programs which may analyze and change information in a layer. In order to make this process a bit clearer, we will give a few simple examples of features that we will employ. We start with digitized color data: gray level intensities measured through blue, green, and red filters. This raw data can be mapped upward in the following manner.

An average image in any color combination can be computed by letting each local function average its 2 x 2 center cells. Average intensity (a measure of the brightness) can be computed by summing the three components and averaging upward



Hue and saturation [10] can be extracted by first normalizing each color component; two of these normalized components can then be used as coordinates of a digitized chromaticity diagram to map hue (the kind of color). Of course, this computation can be performed by our iteration type of function at any level, on the normalized raw data at level 0 or on the normalized averaged data at a higher level. The latter computation would extract average hue. A simple computation of saturation is a function of the normalized components



$$S = 1 - 3 * \min(b_n, g_n, r_n).$$

All of the above functions measure some type of average value across a local area of the original image (whose size is dependent upon the level). Several measures of variation can also be extracted. The parameter that has been most valuable to us so far is average color variation in 3-space. This is just a differentiation and averaging algorithm.

There are many possible differencing mechanisms; for computational simplicity, we choose $|b_1 - b_2| + |g_1 - g_2| + |r_1 - r_2|$.

Within the 2 x 2 center of a local window, we choose the max in the horizontal and vertical directions; i.e., the max of the four directions shown. Thus, at level 3 we have computed the average color variation in an

8 x 8 window of the original data. Although we will not discuss them further, variations of intensity, saturation, etc., can be computed in a similar manner.



IV. DETECTION OF MICRO TEXTURE

We have shown how the cone structure allows us to compute values for various parameters (features) which characterize certain aspects of the image. The ultimate goal is to extract from this information the types of textures in various parts (re-

gions) of the image without attempting to quantify them by arbitrary predefined values of texture attributes. These parameters may be used individually to separate textures. For example, a histogram plot of hue from a region may indicate the contribution of the hue attribute to a perceived texture. One might expect, however that single parameters are not sufficient to adequately define the points comprising a textured region. Single parameter separations, such as average intensity across a region, fail when the region is composed of two types of textures, both of which exhibit the same or nearly the same average intensity. Clearly these remarks apply to any arbitrarily chosen attribute. For example, suppose we have a region composed of four texture types; we define a texture type as characterized by distinct values of the parameters chosen to describe the texture. Let us assume we choose average intensity and maximum intensity variation over a window as our parameters. Figure 2 illustrates a histogram derived from an image composed of these four textures.

Each circle in Figure 2 represents a distribution (or cluster) of image points having similar values of the parameters. In this case, removal of either coordinate results in a loss of discrimination; III and IV become one cluster or I and II become a single cluster.

If there are n parameters, one could map each image point into n-space, forming an n dimensional histogram in an attempt to discriminate texture types. We have chosen 2-space (two parameters) because a) the computation described below becomes unwieldy in n-space even for reasonably small values of n; and b) the cone structure is eminently suitable for 2-space manipulations in parallel as will become evident in the following discussion.

As a further simplification, let us assume we have only two texture types making up a region R as shown in Figure 3. If we




Figure 2. Histogram of Image Points Representing Texture Types




Figure 3. Distribution of two Textures in R

form a histogram of the values of two parameters, average intensity and maximum intensity variation, over each of the small windows shown in the figure, the clusters shown in Figure 4 could be formed. In the ideal case, in which the clusters are well separated, a simple region growing algorithm may be applied to the histogram in order to determine which histogram points comprise each cluster. In effect, we have formed a clustering algorithm in the cone that will produce results similar to standard clustering algorithms applied to two dimensional data. The clusters are labeled I and II as shown in Figure 4. Each window in region R may then be labelled type I or type II, since each of these image points map into one cluster or the other.

A running example of this algorithm applied to real data will now be introduced. Figure 5 is a B & W version of a digitized



Figure 5. B & W Version of color image used in examples




Figure 4. Histogram derived from Figure 3

color image kindly supplied to us by R. Reddy at Carnegie-Mellon. Interesting results have been obtained by applying the algorithm to the following pairs of features:

- intensity vs. color variation
- hue vs. color variation
- hue vs. saturation
- and hue vs. intensity.

This example will be restricted to intensity vs. color variation. The average intensity and average color variation from level 3 (32 x 32) in the cone is computed as we described in the previous section. This level of the reduced image is chosen because the local windows are reasonably large (8 x 8) and the 32² image still preserves some detail. These results are shown in Figures 6 and 7. The reduced 32 x 32 image is now mapped into a histogram of intensity vs. color variation, part of which is shown in Figure 8. All values greater than 61 are mapped into 61 so that it is a 61 x 61 matrix.

The real data shown in Figure 8 seldom conforms to our ideal examples for the following reasons. Two clusters may be linked by "bridges" of very low values which will cause the region growing algorithm to grow them into one region when in fact at a "gestalt" level, they should be distinct clusters. Furthermore, clusters which are distributed over an area may contain minor breaks which would halt the region grower before completely connecting the points while at a "gestalt" level, the points comprise one cluster. This implies that intermediate processing must be performed on the histogram in order to obtain the clusters.

One characteristic of clusters is that they are composed of a set of points spatially related according to some distance property which provides a gestalt grouping. Inherent in the cone structure is a similar distance property. Points that lie within a certain distance of each other (e.g., no

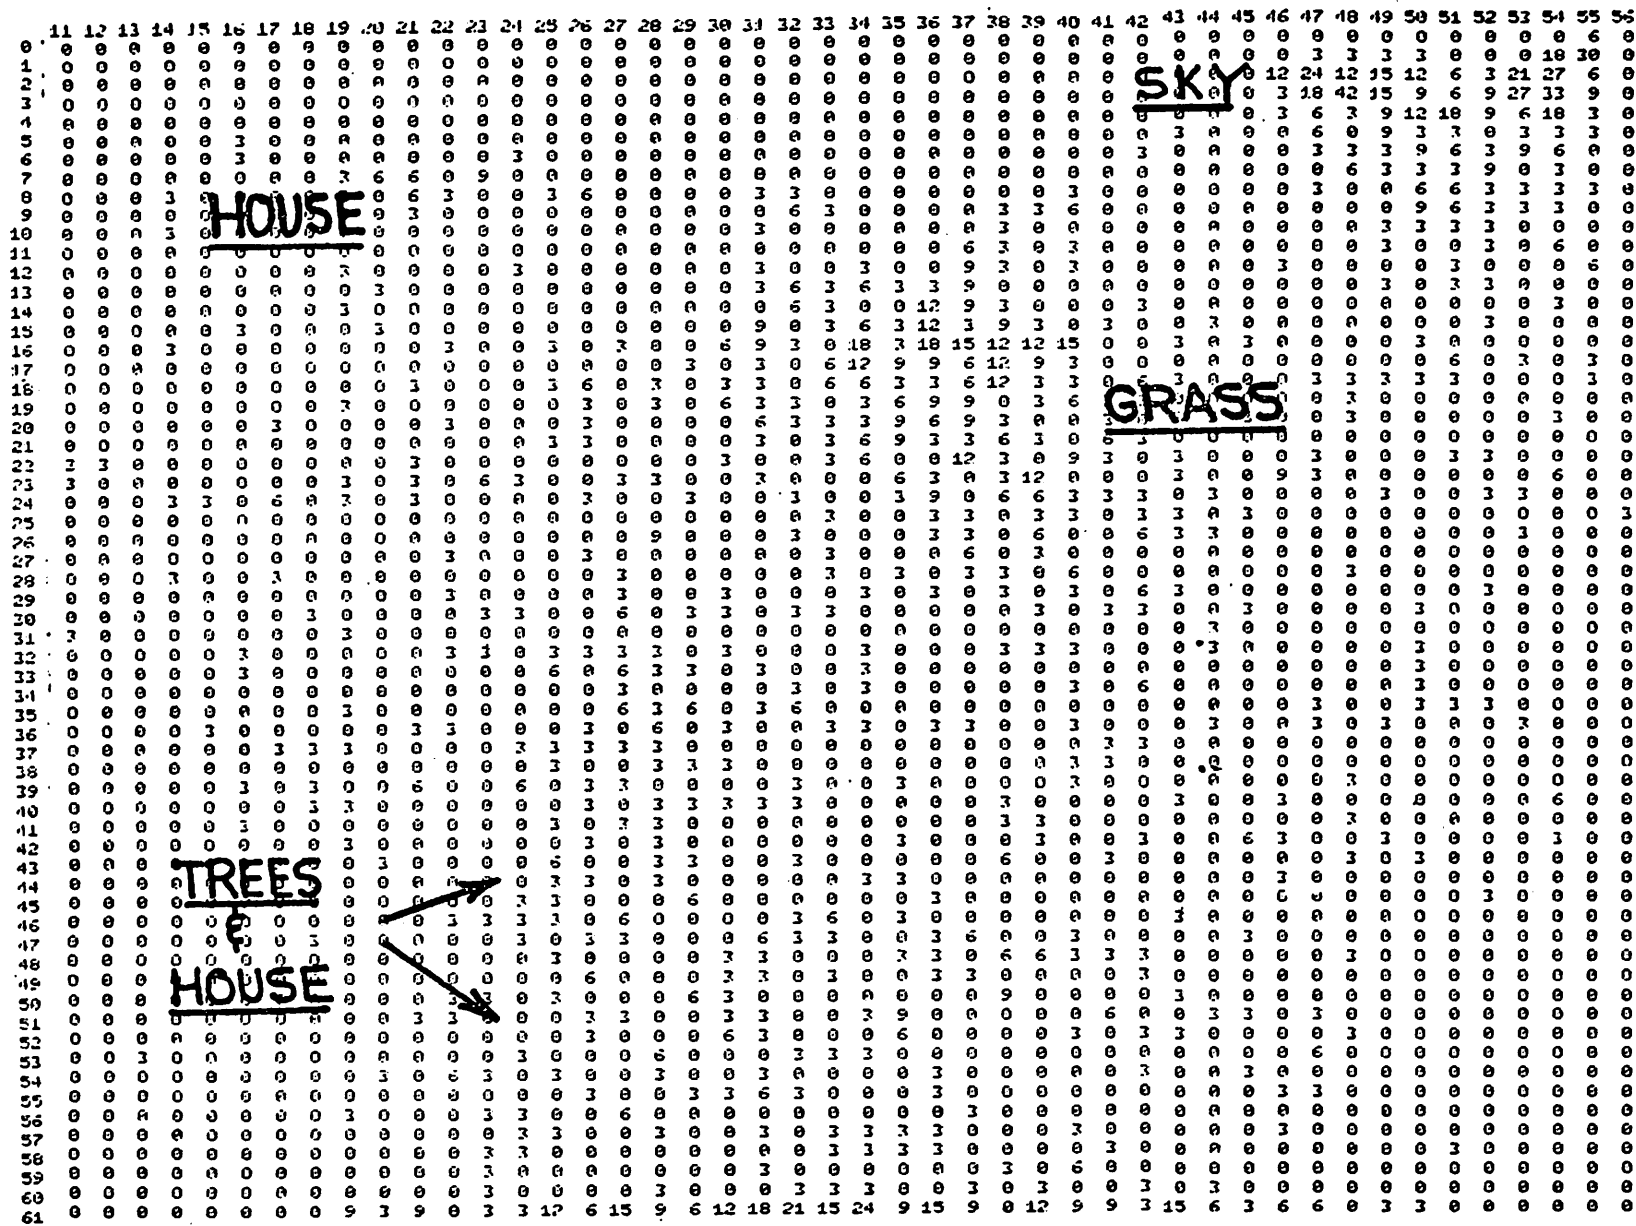


Figure 8. Histogram of 3-Color Differences (y coordinate) versus Intensity (x coordinate)

more than one window width apart) must end up in the same cell or in adjacent cells at the next level in the cone; at intervals of more than one level, the window width is the size of one side of the subcone below a cell. We may therefore treat the histogram as a pseudo-image and insert it into the cone at the appropriate level; in the analyses to follow, the histogram is inserted at the 64^2 level since the values of the parameters are in the range 0 to 63. The reason for treating the histogram as a pseudo-image is that we can now apply, in parallel, the complete range of operators available in the cone; e.g., reduction, scaling, region growing, projection, etc.

Once the histogram has been inserted in the cone, it is reduced to the 16^2 level via a scaling and averaging operation. By only using the 2×2 window below each cell during reduction from level 2 to level 4, each cell at level 4 is an average of the 4×4 window in the original histogram. This operation has the effect of eliminating small values bridging clusters and tightening the clusters themselves; e.g., significant clusters remain intact while small diffuse clusters and stray points are averaged out. The results of this operation are shown in Figure 9 below. The simple region grower may now be applied to this highly reduced histogram at the 16^2 level. Once the clusters have been labelled, as shown in Figure 10, the 16^2 level may be projected down to the 64^2 level and the region growing algorithm applied once again using the already labelled points as the seed regions to

0	0	0	0	0	0	0	0	0	0	3	7	11	0
0	0	0	0	1	0	0	0	0	0	0	5	4	0
0	0	0	0	0	0	0	0	0	0	0	2	1	0
0	0	0	0	0	0	0	2	4	0	0	0	0	0
0	0	0	0	0	0	2	4	0	2	0	0	0	0
0	0	0	0	0	0	2	3	1	1	0	0	0	0
0	0	0	0	0	0	0	2	1	0	0	0	0	0
0	0	0	0	0	0	0	0	0	1	0	0	0	0
0	0	0	0	0	1	1	0	0	0	0	0	0	0
0	0	0	0	0	1	0	0	0	0	0	0	0	0
0	0	0	0	0	1	0	1	0	0	0	0	0	0
0	0	0	0	0	1	0	1	0	0	0	0	0	0
0	0	0	0	0	0	1	0	1	0	0	0	0	0
0	0	0	0	0	0	0	1	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0	0	0	0	0	0	0
0	0	0	0	1	2	2	4	2	2	1	0	0	0

Figure 9. Histogram of Figure 8 reduced to 16^2 level.

Note: the leftmost and rightmost columns are zero and do not appear in Figure 9. In addition, the values shown in this figure are a result of a multiplicative scale factor of 3 (already reflected in Figure 8).


the algorithm. This has the effect of adding connected fringe points to the clusters identified at the higher level. The results of this algorithm at the 16^2 level are shown in Figure 11. Simple additional processing results in most histogram points being labelled according to the cluster (texture) type to which they belong; randomly scattered points, however, are not labelled.

Returning to the original scene data at the 32^2 level, a cluster label may now be associated with each image point since an ordered pair (intensity, color variation) of values uniquely identifies the cluster into which it maps. The results of labelling the original 32×32 reduced image are shown in Figure 12.

V. Detection of Macro Texture

At no point has the micro-texture analysis utilized any spatial information which could bind textured regions together; this information is not reflected in the clusters. Now we must try to detect the regularly occurring spatial distributions of texture types which we call macro-texture. By forming an adjacency matrix of N clusters by N clusters, some sense of the spatial relationships between texture types may be determined. A homogeneous region composed of points from a single cluster i will cause large values in the adjacency matrix on the $(i,i)^{th}$ diagonal element. Macro-texture is then determined by examining the large off-diagonal entries (if any). That is, a large number of adjacencies between points of two cluster types might signal a cohesive region that has a repetitive mixture of cluster types. The light green and dark green of a tree in bright sunlight is one example, as discussed above.

Again, let us assume we have determined two micro-textures. Figures 13a and 13b



(a)

(b)

Figure 13. Example distributions of two texture-types and their adjacency matrices.




Figure 10. Labelled clusters at 16^2 level before including fringe points.




Figure 11. Final labelling of clusters at 16^2 level (after additional processing to retain fringe points).




Figure 12. Original 32^2 reduced image labelled by cluster membership.

illustrate two possible configurations of these texture types.

An adjacency matrix of these texture types can be formed in which a_{ij} represents the number of times texture-type i is adjacent to texture-type j (using four neighbor adjacency):

	I	II
I	a_{11}	a_{12}
II	a_{21}	a_{22}

In the case of Figure 13a, a_{11} and a_{22} are large because of the spatial distribution of textures; a_{21} is small and represents the boundary between the two regions. Since the adjacency matrix is symmetric, a_{21} may be ignored and set to 0. For Figure 13b, both a_{11} and a_{12} are large while a_{22} is small due to the absence of a homogeneous region of type II texture.

One of the problems we face is that some entries in the adjacency matrix are artificially inflated due to the interaction of texture types across the boundary of adjacent regions. In Figure 13b, it is not clear whether type I elements should be grouped with type II. If that is allowed, one big region will be formed as opposed to a homogeneous subregion of I's and a macro-textured subregion of I's and II's. The removal of the homogeneous region will make this distinction evident. In order to remove its contribution, a modified region growing algorithm may be applied to the texture type labels associated with the image points. This algorithm finds large connected homogeneous regions which are subsequently removed; construction of the adjacency matrix may then proceed without the contribution of boundary conditions. In the case of Figure 13b, this results in an adjacency matrix of the form

	I	II
I	x	○
II		x

where the circle represents a large value, the contribution due to macro texture, and the x's smaller values. Figure 14 shows the adjacency matrices for the 5 micro-textures appearing in the labelled image of Figure 12 before and after the few major homogeneous regions (and their boundary effects) are deleted. The macro-texture of types 4 and 5 in the trees on both sides of the image are detected, but do not stand out as strongly as we would like.

There are a number of algorithms that are currently being examined to improve the

	1	2	3	4	5
1	367				
2	1	21			
3	67	10	294		
4	9	43	150	300	
5	4	14	47	216	178

14(a)

	1	2	3	4	5
1	2				
2	1	21			
3	6	8	71		
4	1	35	72	119	
5	2	14	40	122	111

14(b)

Figure 14. Cluster-type adjacency matrix for Figure 12.

performance of this portion of the analysis but no conclusions are available at the time of this writing. One significant weakness that must be dealt with is the problem of improper placement of the local windows (8 x 8 in our examples). Some of the windows fall entirely across a single type of micro-texture; others will overlap different types. If the boundaries of the trees (types 4 & 5) and sky (type 1) are examined in figure 12, one will note a line of points of type 3 which are a result of this effect. A number of ways to suppress these points are being examined.

VI. Conclusions

The algorithms discussed in the paper represent a portion of an on-going research effort. The power of the system resides in the maintenance of both a local and global perspective. The local extraction of texture and color features is blended with a global labelling of cluster types in the histogram. The experiments are incomplete and the results should be considered as preliminary. While we feel the results conclusively reveal the merits of this approach, several avenues of research are still open.

The results shown here are based on intensity versus color variation. While we have mentioned several other possibilities they have not been examined in any detail. One of the more exciting aspects here is the expectation that other pairs of parameters will yield additional information. Redundancies will produce regions with greater confidence. However, different feature pairs may extract different regions with different texture characteristics. The integration of this different information should yield a more complete characterization of the image.

The flexibility of the processing cones allow projection of an entire region down to lower levels of the cone. Examinations are underway of techniques for quickly growing regions at a relatively high level, say 32^2 , and refining the regions by alternately projecting downward and region growing at lower levels.

All of the analyses to date have been derived from the original reduced image at the 32^2 level. While this level is sufficient for the detection of reasonably large regions, the quantification of both macro- and micro-texture clearly depends upon the resolution of the original image as well as the size and spacing of the texture elements. Regions containing small macro-texture elements or closely spaced regions (such as the house shown in Figure 5) may be analyzed at a lower level (64^2 or even 128^2) in precisely the same manner as described here. Furthermore, we can confine our processing on these lower levels to only those areas of the image which require further analysis. This type of low-level processing is being integrated into a full scene analysis system.

References

1. Y. Yakimovsky and J. A. Feldman, "A Semantics-Based Decision Theory Region Analyzer," Proceedings of Third Joint Conference on Artificial Intelligence, August 1973, pp. 580-588.
2. A. Rosenfeld and M. Thurston, "Edge and Curve Detection for Visual Scene Analysis," IEEETC, pp. 562-569, May 1971.
3. K. C. Hayes, Jr., A. N. Shah, and A. Rosenfeld, "Texture Coarseness: Further Experiments," IEEETSMC, pp. 467-472, September 1974.
4. B. Bullock, "The Performance of Edge Operators on Images with Texture," Hughes Research Laboratories Technical Report, Hughes Aircraft Company, Malibu, California, October 1974.
5. R. Haralick, "Textured Features for Image Classification," IEEETSMC, pp. 610-621, November 1973.
6. R. Bajcsy, "Computer Description of Textured Surfaces," Proc. of 3rd IJCAI August 1973, pp. 572-579.
7. E. M. Riseman and A. R. Hanson, "Design of a Semantically Directed Vision Processor (Rev. & Updated)," Tech. Report 74C-1, Computer and Information Science, University of Massachusetts, January 1974.
8. A. R. Hanson and E. M. Riseman, "Pre-Processing Cones: A Computational Structure for Scene Analysis," Tech. Report 74C-7, Computer and Information Science, University of Massachusetts, September 1974.
9. D. Fishman, A. Hanson, E. Riseman, "Considerations in Model-Building for Scene Analysis," Tech. Report 75C-, Computer and Information Science, University of Massachusetts, January 1975.
10. R. O. Duda and P. E. Hart, Pattern Classification and Scene Analysis, John Wiley & Sons, 1973.

Biographies

Dr. Allen R. Hanson received his Ph.D. degree from Cornell University in 1969. He taught in the Department of Computer Information and Control Sciences at the University of Minnesota from 1969-1973. He is currently an Assistant Professor at Hampshire College. His research interests have included a variety of topics in pattern recognition and artificial intelligence, particularly scene analysis.

Dr. Edward M. Riseman received his Ph.D. degree from Cornell University in 1969. Since then, he has been a member of the Computer and Information Science Department at the University of Massachusetts and is currently an Associate Professor. His research interests span the use of context in character recognition, AI systems for mechanized theory formation, and scene analysis.

Paul Nagin is a graduate student in the Computer and Information Science Department at the University of Massachusetts. He received his M.S. degree in 1974 and currently is working towards his doctorate. His research is in scene analysis systems, particularly the processing of texture and the segmentation of regions in complex scenes.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER COINS Technical Report 75C-3	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) REGION GROWING IN TEXTURED OUTDOOR SCENES		5. TYPE OF REPORT & PERIOD COVERED INTERIM
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) A. R. Hanson, E. M. Riseman, P. Nagin		8. CONTRACT OR GRANT NUMBER(s) ONR N00014-67-A-0230-0007
9. PERFORMING ORGANIZATION NAME AND ADDRESS Computer and Information Science Dept. University of Massachusetts Amherst, MA 01002		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS Office of Naval Research Code 437 Washington, D.C.		12. REPORT DATE February 1975
		13. NUMBER OF PAGES 13
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		15. SECURITY CLASS. (of this report) UNCLASSIFIED
		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 This paper was presented at the 3rd Milwaukee Symposium on Automatic Computation and Control, April 18 and 19, 1975, Milwaukee, Wisconsin.		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) Some progress has been made in scene analysis in finding boundaries between textured objects. However, standard line-finding and region growing algorithms have not been able to deal with strong local variations of intensity and color that repeat themselves more globally in images of natural scenes. These prob- lems are significant with respect to both micro-texture, say the leaves of a tree, as well as macro-texture, the light green of leafy branches vs. dark green or black shadows between branches in a tree. How can local variations		

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

be globally bound together into a cohesive region for perceptual identification.

The approach that is being examined involves the reduction in size of the image while extracting features such as intensity, color, and color variation in local windows across the image. Two-dimensional histograms are being used to determine clusters of points with similar characteristics. Major clusters are labelled and the corresponding points in the reduced image are also labelled. This can directly detect regions of homogeneous micro-texture. Now an adjacency matrix of $N \times N$ clusters will denote the types of image points adjacent to each other. Peaks in the adjacency matrix can determine macro-texture.

All of this information can be used to grow regions in the reduced image and then in the detailed image. Experimental results on color images of natural outdoor scenes will demonstrate our results.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)