

RESYNTHESIS OF BIOLOGICAL IMAGES FROM TREE-STRUCTURED DECOMPOSITION DATA

RUSSELL A. KIRSCH

Applied Mathematics Division, National Bureau of Standards
Washington, D.C. 20234

Abstract: In this paper we are simultaneously concerned with methods for decomposing grey scale microscope images and with methods for verifying the correctness of these decompositions. One such method is resynthesis. Resynthesis is viewed as a procedure whereby an analyzed scene can be reconstituted and subjected to an analysis by human (informal) methods to determine the information preservation of the process. Several algorithms are presented for different ways of resynthesizing a decomposed image from its morphological decomposition analysis.

I. INTRODUCTION

In attempting to do pattern recognition with computers on continuous tone image sources of complex structure, one encounters the problem of decomposing the images for scene analysis. When one's goal in such pattern recognition is more than to assign the image to one of a number (usually small) of distinct classes, it becomes necessary to decompose the source image into subparts. These subparts usually reflect the structural organization of the underlying object which, through some transducer, produces the image for computer processing purposes. This structural organization can reflect not only the intrinsic parts of the image, but also the relations between the image and others of some class of images sharing similar structures. In this sense, the structural decomposition of an image reflects the way in which some subculture that views such images organizes them, usually for purposes not directly deducible from the images themselves. In this paper we treat one aspect of the image decomposition problem, namely that of attempting to verify the correctness of a decomposition of a computer-scanned image by resynthesis of the natural image or by creating a surrogate of the natural image from the decomposition data which characterizes its structural organization. The image decomposition method that we use is described by Kirsch (1971).¹ It is a connected regions type of analysis which is thus related to those of Brice and Fennema (1970),² Jarvis (1971),³ Rosenfeld and Thurston (1971),⁴ and most closely to that of Krakauer (1971).⁵

If we are to study a decomposition and verify its correctness, it is incumbent upon us to elucidate the criteria with respect to which we perform such verification. Several criteria which we wish to consider below have been implicitly suggested in the literature. Naturally, different decomposition procedures will appear to be differentially successful with respect to the distinct evaluation criteria.

II. EVALUATION CRITERIA

The several evaluation criteria which we list below are mostly informal because they require more or less subjective judgment. If these criteria appear appropriate, and if the decomposition-resynthesis experiments we report appear to be reasonably evaluated with them, then a problem worth future study is the formalization of these criteria. We do not attempt such a formalization here.

1. Criterion of Pragmatic Usefulness in Higher-Level Recognition

This criterion measures the success of an image decomposition, especially in such applications as scene analysis for complex objects (Guzman (1971)⁶) by determining the extent to which higher-level image processing routines (some of which are based on extra-image properties) can find low-level primitives needed for higher-level processing as outputs from the decomposition procedure. For example, in a heuristic pattern manipulation system such as Winograd (1971)⁷ a measure of the success of a scene analysis decomposition for images containing polyhedra would be the extent to which edges and regions, which natural language processing routines talk about, can be made available from the scene analysis programs.

2. Criterion of Succinctness for Characterizing Image Decomposition

If one's purpose in decomposing an image is encoding for such purposes as storage or image transmission, then a valuable decomposition procedure is one that reduces the information content of the image by redundancy elimination from that necessary to characterize the raw scan to a more compact one. In image analysis such a criterion finds particular utility for such applications as long-distance facsimile transmission, and for image storage in large data banks.

3. Criterion of Information Theoretic Invertibility

This criterion measures the success of a decomposition procedure by whether the original information can be completely resurrected from the decomposition data. The medial axis transform of Blum (for which see generalizations in Levi and Montanari (1970)⁸) has this property. In a sense, therefore, no information is lost during the decomposition and hence no pattern recognition is attempted, although transforms of the decomposition data may be simpler than transforms of the raw data in subsequent pattern recognition processes. This criterion is a somewhat more stringent version of the second criterion since the emphasis is on image storage and it forces a postponement of recognition to later phases of image processing.

4. A Converse of the Third Criterion

A converse of the third criterion measures the degree of information reduction in a decomposition procedure. The highest form of such decomposition is the one which results in assigning a unique label to the pattern, as for example in character recognition. This may thus be viewed as an information reduction from n^2 bits into $\log_2 k$ bits where an n by n binary image is assigned to one of k classes.

5. Criterion of Completeness

The criterion of completeness is something of an elaboration on the third criterion. From the standpoint of this criterion an image decomposition is successful if it can result in the reproduction of the original image to within some (usually informal) fidelity criterion, and if in addition it can assign to the original image a structural description of the component parts. The decomposition must attempt to explicate how the components are represented as members of some larger class beyond the individual image being analyzed. There is a direct analogy between this criterion of completeness and the type of criterion used in natural language processing. There one wishes to do syntactic analysis of strings of symbols in such a way that the resulting analysis not only assigns to the original input string a unique code, but does this in a way that structural descriptions are assigned to related different strings which mirror the

relatedness of the symbol strings themselves. In order to do this it is, of course, necessary to ease up on both the restriction of information theoretic invertibility and that of information reduction because a structural description will necessarily be more elaborate than the minimal code for representing the information source.

6. Methodological Niceties

Methodological niceties provide a final criterion with which one can evaluate an image decomposition algorithm. Some of these niceties may appear extreme versions of the sine qua non criterion #1, but they are addressed to questions of how one attempts to generalize from particular solutions to the image decomposition problem to general solutions over larger image classes. One such nicety has to do with the clear separation of the morphological and syntactical aspects of image decomposition. In morphological decomposition one attempts to explicate the structure of the image based on properties which are measurable and intrinsic to the image, whereas in syntactic analysis one relates the image to others from the same class and thus explicates properties which are class properties rather than individual ones.

Another methodological nicety has to do with the generality of the decomposition over different classes of images. Needless to say, it is hopeless to attempt syntactic analysis over essentially distinct classes, but for the morphological aspect of decomposition it appears reasonable to attempt to use decomposition procedures which are neutral with respect to the intended interpretation of the images.

A third very important methodological criterion has to do with the ease of expansion of the decomposition algorithm to larger classes of images. If one views the problem of image analysis as essentially an open-ended one in which the class of images must necessarily increase in size with time, then a provision must be made whereby new information about image structure can be fed to a system without essentially disrupting the behavior of the analysis phase. Many complex information sources seem to require essentially open-ended descriptions because of the difficulty of anticipating a completely closed description of the information sources at any one particular time. Except in the most stylized problems, most natural image description has this property and thus has the need for allowing unlimited expansion of the information source description in dealing with open-ended classes of images.

In this paper we describe some experiments with an image decomposition procedure which seems to score highest in terms of the sixth of the above criteria (the methodological niceties), although if it should fail in terms of the first criterion (pragmatic) it would be of little value.

III. THE SYNTHESIS TEST

By a synthesis test for an image decomposition procedure we mean a test whereby an image is resurrected or resynthesized from decomposition information and then tested against the original image in terms of any one of a set of fidelity criteria. To understand how a synthesis test may be used we might compare the use of this test with the use of the third criterion mentioned above (information theoretic invertibility). For decomposition procedures which have the invertibility property the synthesis test is particularly simple to apply. One recomputes the original image and then identifies it or measures the degree of disparity between the resynthesized image and the original in terms of some common measure (cross-correlation, Hamming distance or others). In the case of invertible procedures the match between the resynthesized and original images is usually perfect. However, if there has been a loss of information during the

decomposition procedure the inverse operation is not unique and the results of the matching operation do not yield perfect agreement between the original image and the resynthesized one. This is where the notion of a human fidelity criterion is useful. By using a human judge to determine the degree of disparity between an original image and the one resynthesized from decomposition data, one may measure whether a decomposition procedure which demonstrably loses information nevertheless might preserve the information which is essential about some image. In general it is not possible to determine a priori whether a resynthesized image preserves the information that matters to some human observer, that property being very sensitively related both to the information source and to the degree of knowledge that the subject has about the information source.

A synthesis test for a decomposition procedure is particularly stringent in another way. The ability to synthesize members of what purports to be an original class of images can make very explicit the extent to which the resynthesized class either falls short of or exceeds the intended class of images to be interpreted. In natural language processing the synthesis test is a powerful one for demonstrating the extent to which a grammar exceeds or falls short of capturing an intended class of sentences. A similar benefit may be obtained by using synthesis methods for studying image decomposition procedures.

IV. RESYNTHESIS EXPERIMENTS

We wish now to consider some experiments in applying the notion of resynthesis to biological images which have been scanned with a scanning microscope and then decomposed by a particular decomposition algorithm. The images to be analyzed are two scans obtained with a scanning microscope (described by Stein, Lipkin and Shapiro (1969)²) of a human leukemic marrow blast cell in an autoradiograph preparation showing both the image structure and the structure of silver grains in a photographic emulsion overlying the image. The cell was first scanned with white light illumination and then with monochromatic blue light at a wavelength where the cell nucleus is substantially clear in the preparation used. For both the white and blue scans the blood cell image is represented as an array of 256 by 256 optical brightness values, each value having eight bits of information.

Figure 1 is a representation of the cell scanned in white light. Sixteen distinct printing characters are used to suggest the 256 optical brightness values discriminated by the scanner. Figure 2 is a corresponding scan of the same cell with the blue illumination.

The decomposition algorithm used is described by Kirsch (1971).¹ We give here a brief description of this algorithm. For the original image we first compute some brightness contrast function which is a locally computable point function relating to the gradient of the image brightness function. Larger values for this contrast function correspond to higher values of the brightness gradient. Next, connected regions are identified which have a contrast value equal to or below some low threshold. These are thus regions of maximum homogeneity of brightness. In an iterative procedure the contrast threshold is raised, allowing the identification of connected regions of successively less and less homogeneity. Each of the connected regions previously identified is properly or improperly included in the regions corresponding to higher threshold. When two disjoint regions coalesce these two regions and their coalescence are noted in a tree structure which is preserved during the successive iterations of the algorithm. The process terminates when all regions have coalesced into one maximal size region of minimal homogeneity, usually the whole image. The essential properties of this algorithm are, first of all, that it identifies regions by thresholding in the contrast transformation, and secondly, that the threshold used to identify regions is based neither on a single uniform threshold criterion nor a local one, but rather one which relates to the connectedness of regions under varying thresholds.




FIGURE 1.

SCAN OF HUMAN ADULT MARROW BLAST CELL FROM ACUTE MYELOGENOUS LEUKEMIA WITH WHITE LIGHT ILLUMINATION. IMAGE HAS 256 X 256 ELEMENTS WITH 16 PRINT CHARACTERS USED TO APPROXIMATE THE 256 BRIGHTNESS RESOLUTION VALUES MEASURED.




FIGURE 2.
BLUE LIGHT SCAN OF SAME CELL PICTURED IN FIGURE 1.

This (morphological) decomposition algorithm was applied to the image shown in Figure 1. In this experiment the contrast function was partitioned into equal intervals by choosing five equally spaced thresholds. In addition, because contrast functions are particularly sensitive to noise and local variations in an image structure, all regions having an area of less than 16 picture elements were rejected during the analysis phase. The consequent decomposition produced a single tree structure having 260 nodes, the top node representing the single region with $256 \times 256 = 65536$ points in it, and with the number of distinct regions at each threshold level and the total area subtended by them summarized in Table 1.

TABLE 1.

Summary of region decomposition for the cell shown in Figure 1. Only regions of area ≥ 16 were accepted. The contrast function range was divided by five equally spaced thresholds denoted level 0, ..., level 4.

Contrast Threshold	Number of Regions	Total Area
0	1	65536
1	16	37922
2	40	28888
3	76	17286
4	127	4732

Since it is not practical to exhibit the tree structure with all 260 nodes we will attempt to characterize some of the properties of the morphological decomposition tree of the image of Figure 1. Suppose that we take the decomposition tree and extract just those nodes corresponding to the first level below the root of the tree. As we see in Table 1 there are 16 different disjoint regions at that level in the tree subtending a total area of 37922 picture elements. In Table 2 we present a summary of the analysis of each of the individual regions in contrast threshold level 1 of the morphological decomposition tree.

TABLE 2.

Summary of Properties of the 16 Contrast Threshold 1 Regions

Arbitrary Sequence Number	Area	Upper Left Corner:		Brightness (Arbitrary Units)		Cut Membership S B
		Row	Col	White Scan	Blue Scan	
1	61	88	139	83	92	x x
2	1647	1	167	198	145	. .
3	442	1	215	75	117	x .
4	79	12	117	109	121	x x
5	202	161	251	92	110	. .
6	1369	223	178	95	108	. .
7	1667	1	1	216	152	. .
8	254	9	127	130	134	. .
9	36	9	141	156	143	x x
10	31695	16	71	154	125	. .
11	16	17	178	175	158	x x
12	346	86	246	207	145	x .
13	24	184	247	100	117	x x
14	21	225	216	128	115	x x
15	26	231	225	126	108	x x
16	37	245	245	186	119	x x

In the first column of Table 2 is an arbitrary sequence number assigned to each of the 16 regions. The second column gives the area of that region. We notice that the regions are widely distributed in size, ranging from a smallest one having area 16 to a largest one having area 31695. The next two columns in Table 2 give for reference purposes the coordinates of the upper lefthand corner of each region. The row numbers range from 0 to 255 from top to bottom and the column numbers are similarly ranged from left to right. In the fifth column of Table 2 we show the brightness of the picture elements constituting the region averaged over the whole region. Notice that insofar as this single number represents the mean brightness over the whole region, all internal structural information is lost. The scale of values for the fifth column ranges from that of a darkest element with value 0 to a brightest one with value 255. The next column gives the mean brightness of the corresponding region from the blue scan of Figure 2. As we would expect the different relative brightness among distinct regions in the white and blue scans relates to the different staining properties of the parts of the underlying cell from which the image derives. The last column of Table 2 designates membership in two subsets or "cuts" (discussed below) for each of the regions given in the table.

Another way to visualize this fragment of the morphological decomposition of the original image is to do a resynthesis of the 16 regions preserving shape information but replacing all internal structure by mean brightness values. Figure 3 exhibits such a resynthesis. Several properties of the resynthesized image in Figure 3 are noteworthy. First we note that the analysis of Figure 3 given in Table 2 shows that there are only 16 disjoint regions. These regions have different mean brightness values, thus the region in the upper righthand corner (sequence number 3, with area 442 at row 1, column 215) is the darkest of the regions, whereas another region (sequence number 7) is the brightest region and indeed corresponds to a part of the image background as can be seen from Figure 3.

Selection of Decompositions

We have suggested implicitly that within the complete morphological decomposition tree for an image like that of Figure 1 there appears sufficient information to characterize the morphology of the image insofar as that morphology is evident in intrinsic data, rather than external syntactic clues. To explore this matter more systematically we wish now to investigate various synthesis procedures. All of the procedures are based upon the morphological decomposition tree. We have already seen in Figure 3 a resynthesized image containing 16 regions chosen from the 260 regions in the whole decomposition tree. We also notice that in Figure 3 most of the gross morphology relating to such things as figure-ground separation and identification of large regions seems to have been made explicit.

In order further to investigate resynthesis procedures we introduce the notion of a "cut" for a tree. By a cut we mean a subset of nodes satisfying the following recursive definition:

- a. The set consisting of the root of a tree is a cut.
- b. If A is a cut of a tree and if B is obtained from A by replacing a member of A by all its immediate descendants in the tree, then B is a cut.
- c. All and only such sets of nodes are cuts.

We see thus that the set of 16 nodes from which Figure 3 was resynthesized constitutes a cut of the tree. In a certain sense a cut constitutes a maximal set of nodes for describing an image structure, there being many such maximal sets. To see how other cuts might be generated, consider the distribution of the terminal and non-terminal nodes of the decomposition tree as shown in Figure 4. We see the nodes of the tree at each level partitioned into those which are




FIGURE 3.
IMAGE RESYNTHESIS CONTAINING 16 REGIONS WITH EACH REGION EXHIBITED WITH UNIFORM BRIGHTNESS. THESE REGIONS FORM THE CUT IMMEDIATELY BELOW THE ROOT OF THE DECOMPOSITION TREE.

LEVEL	TERMINAL NODES	NON-TERMINAL NODES
0	0	1
1	8	8
2	31	9
3	66	10
4	127	0
TOTAL	<u>232</u>	<u>28</u>




FIGURE 4.
DISTRIBUTION OF TERMINAL AND NON-TERMINAL NODES IN THE DECOMPOSITION TREE
OF FIGURE 1.

terminal and have no constituent structure of their own, and those which are non-terminal and hence contain other subnodes corresponding to subregions. Another cut of the tree suggested by Figure 4 consists of all the terminal nodes (a total of 232). Figure 5 exhibits a resynthesis of the image consisting only of these 232 terminal nodes, again with the mean brightness for the nodes replacing the internal image structure in the resynthesized image. The resynthesized image of Figure 5 quite evidently consists of the fine structure of the original image just as the level 1 cut in Figure 3 exhibits the gross structure of the image. The area of this terminal node cut is 10308 picture elements. Those members of this cut which appear among the 16 in Table 2 are indicated by an "x" in column seven under "B".

As we view the morphological decomposition procedure it serves to provide a large set of alternatives to higher-level syntactical analysis routines, among which alternatives a choice may be made with respect to which set (usually a cut) is to be used for higher-level processing purposes.

An interesting example of how a higher-level procedure would invoke a morphological analysis procedure is illustrated by the next example. Here we make use of the additional information present in the scan of Figure 2 made in blue light. Because of the histochemical nature of the biological preparation we know that certain regions in the image should be dark in both the white light and in the blue light scan, whereas other regions might be distinguished by the two different sources of illumination. In order to select those nodes which are maximally similar in the white and the blue light scan we can choose a cut for the morphological tree in the following way. We start with the root of the tree and then successively invoke step 2 in the definition above of a cut, or not invoke it in such a way as to make the disparity between the average brightness of regions in the white and blue scan minimal. Thus the selection of subnodes in the tree, which is an option provided by the definition of a cut, is made according to what is effectively an external criterion, namely similar brightness in two different scans. By starting from the root of the tree and successively replacing nodes with subnodes when the mean two-color disparity in brightness of the subnodes is less than that of the parent node, we generate another cut from the tree. This cut consists of a mixture of terminal and non-terminal nodes having a total of 100 nodes and an area of 20719 picture elements. We may then resynthesize the image from mean brightness data using the cut thus generated and obtain a resynthesized image as shown in Figure 6. Those of the 100 nodes appearing in Table 2 are indicated by an "x" in the last column under "S".

We have thus seen a few of the many alternative procedures for decomposing an image using a morphological decomposition algorithm followed by a selection of a cut set of nodes to partition the image into disjoint regions. The test of the adequacy of such a partitioning with respect to some intended interpretation with a human fidelity criterion is the test of resynthesizing the image preserving the decomposition data and then inspecting the image for fidelity with respect to the original.

Other Contrast Functions

The nature of the set of decompositions available from an image is of course related to the choice of contrast function used for generating the morphologic tree. A contrast function related to the gradient of the image brightness seems natural for discriminating objects which are relatively uniform in their interior and bounded by comparatively sharp boundaries. In biological images obtained from the light microscope this largely constitutes the class of objects which are bounded by membranes. There are, however, many types of objects which we have elsewhere called constructive objects. They are characterized by variation in brightness and appear to be either not sharply bounded or else merely textural within their interior.




FIGURE 5.
IMAGE RESYNTHESIS CONTAINING 232 REGIONS. THESE REGIONS FORM THE TERMINAL
NODE CUT OF THE DECOMPOSITION TREE.




FIGURE 6.
IMAGE RESYNTHESIS CONTAINING 100 REGIONS OF MAXIMUM RELATIVE BLUE-WHITE
SIMILARITY. THESE REGIONS CONSTITUTE A CUT INTERMEDIATE BETWEEN THOSE
OF FIGURE 3 AND FIGURE 5.

In the scans of Figure 1 and 2 we have examples of such objects. These are the silver particle tracks produced by radioactive decomposition in a photographic emulsion overlying the blood cells of the image. Since these silver grains are substantially ultrastructural with respect to the light microscope and hence not resolvable, they appear as regions of heterogeneous brightness usually with a low mean value of brightness because they are opaque. This suggests that for such objects a different type of contrast function might be appropriate, namely one which is monotone-decreasing as a function of magnitude of the point value of the gradient of the brightness function. Such a function which is monotone-decreasing as opposed to the monotone-increasing function used in the above examples, was used to study the image of Figure 1. The set of terminal nodes obtained from such a decomposition (the analog of those used to produce Figure 5) was used to resynthesize an image, again using mean brightness values, this time for the different contrast function with the result shown in Figure 7. We notice that Figure 7 consists of two classes of objects, those which are evidently boundaries between real objects in the original image, and secondly those which are whole regions of substantial heterogeneity of brightness such as the dark silver grains.

V. REMAINING PROBLEMS

As with any interesting problem the use of morphological decomposition methods suggests more questions than it answers. We wish briefly to touch upon these questions as suggestions for further work.

1. Use of Formalized Syntactic Models

We have given an example of the use of extramorphological information corresponding to scans with different illumination wavelength as a method for selection of the proper set of regions from a morphological decomposition. However, there are many other ways in which extramorphological information can be used. Such syntactic methods have been discussed previously in Lipkin, Watt, and Kirsch (1966)¹⁰ where particular note is made of linguistic or descriptive models and their relation to the morphological data in images. Other types of syntactic information that can be used are taxonomic data relating to the taxonomy of a particular class of images, or developmental information. Image organization can be viewed as a reflection of image structure which in turn is a reflection of the morphological development (in such images as cells) which has produced the image structure as a consequence of underlying processes. A final example of the use of syntactic methods is in the work of Harlow (1971)¹¹ where a priori information about the anatomical organization of, in his case, X-rays of the chest can be used as a selection criterion for morphological discrimination.

2. Incremental Decomposition

The procedures for morphological decomposition described above are very expensive in terms of the computer time involved. In the example illustrated above an entire decomposition tree is produced and then 260 separate regions must be isolated and various features measured for each of those regions. On a large general-purpose computer this can take on the order of an hour of computer time for an image the size given in our example. If one has a particular set of syntactically generated selection criteria for morphological selection, it appears possible to avoid much of the computation necessary in morphological decomposition by syntax-driven techniques. Thus in the example where a cut was chosen based on color discrepancy information, the systematic nature of the process descending from the root of the tree to subsequent levels could be terminated whenever the syntactical decision criterion had been met. This would avoid much computation and could conceivably make the process acceptably efficient.




FIGURE 7.

IMAGE RESYNTHESIS OF TERMINAL NODE CUT FROM DECOMPOSITION USING MODIFIED CONTRAST FUNCTION WHICH EMPHASIZES REGIONS OF HIGH BRIGHTNESS HETEROGENEITY.

3. Selection of Decompositions by Feature Analysis Techniques

Since feature analysis and property filter methods (Rosenfeld (1969) in chapter 7)¹² are so widespread in pattern recognition applications it appears worthwhile to consider the set of regions in a morphological decomposition as candidates for feature analysis and subsequent pattern recognition by these relatively simple methods. These methods work well for pattern recognition on isolated patterns with no multiple region problems. But decomposition methods such as the one considered here provide just such isolated patterns. Thus, it appears worthwhile to pursue the question of selection of nodes from a morphological decomposition by feature analysis tests and the simple types of property filtration that can result. These methods may be intermixed with methods that exploit the structural information in the morphology tree to yield fairly elaborate decomposition criteria which can then be subjected to the same type of synthesis test that we used above.

4. Space-Filling Requirements

It is immediately obvious from the synthesized images exhibited above that the regions constituting the resynthesized image are disjoint since they are separated by white space in our figures. The resynthesized images are thus non-committal with respect to their handling of the background for any particular resynthesis. Several ways suggest themselves for handling background. One would be based on growth and propagation techniques. One could propagate regions maintaining their mean brightness values until intersection takes place between disjoint regions. This would result in a space-filling type of resynthesis and in images that might satisfy other types of more stringent synthesis tests.

VI. CONCLUSION

In our discussion we have employed an image decomposition method which we hope can be of fairly broad use at least within the class of biological microscope images and hopefully, in a still broader class of images. We have attempted to evaluate the kind of decomposition produced by invoking the synthesis test of resynthesizing images from decomposition data. The real reason for using a synthesis test, despite the many arguments we offer above in support of it, is that it represents one of the best ways of trying to understand how it is that a pattern recognition procedure works when it does. Since so much effort is necessarily devoted to getting pattern recognition procedures for complex images merely to exhibit satisfactory behavior, it may seem premature to be concerned with issues of why such procedures work. It does seem, however, that ultimately one must encounter such questions and answer them. The present paper is an attempt to suggest how to obtain such answers.

ACKNOWLEDGMENTS

We wish to thank Dr. Lewis E. Lipkin for preparing the scan data used in the above experiment and Mrs. Anne Meininger, Mr. Don Orser, and Mrs. Ida Rhodes who wrote the various programs which achieved the degree of efficiency that made the algorithms practical to run. The work reported here was done under an agreement between the National Cancer Institute and the National Bureau of Standards.

This report is a publication of the National Bureau of Standards and is not subject to copyright.

REFERENCES

- Brice, C.R., and C.L. Fennema (1970), Scene Analysis Using Regions. Artificial Intelligence 1, pp. 205-226.
- Guzman, A. (1971), Analysis of Curved Line Drawings Using Context and Global Information. In Machine Intelligence 6, edited by Bernard Meltzer and Donald Michie. New York: American-Elsevier Press.
- Harlow, C.A., and S.A. Eisenbeis (1971), Analysis of Radiographic Images. Proceedings of the Conference on Two-Dimensional Signal Processing. University of Missouri, Columbia, Missouri.
- Jarvis, C.L. (1971), Picture Element Identification by Labelling Transformations. Computers and Biomedical Research 4, pp. 167-180.
- Kirsch, R.A. (1970), Computer Determination of the Constituent Structure of Biological Images. Computers and Biomedical Research 4;3. pp. 315-328.
- Krakauer, L.J. (1971), Computer Analysis of Visual Properties of Curved Objects. Massachusetts Institute of Technology Project MAC Report #MAC-TR-82. (PhD. Thesis)
- Levi, G. and U. Montanari (1970), A Grey-Weighted Skeleton. Information and Control 17, pp. 62-91.
- Lipkin, L.E., W.C. Watt, and R.A. Kirsch (1966), The Analysis, Synthesis and Description of Biological Images. Annals of the New York Academy of Science 128;3, pp. 984-1012.
- Rosenfeld, A. (1969), Picture Processing by Computer. New York: Academic Press.
- Rosenfeld, A. and M. Thurston (1971), Edge and Curve Detection of Visual Scene Analysis. IEEE Transactions on Computers C-20;5, pp. 562-569. (See especially comments on "grouping.")
- Stein, P.G., L.E. Lipkin, and H.M. Shapiro (1969), Spectre II: General Purpose Microscope Input for a Computer. Science 166, pp. 328-333.
- Winograd, T. (1971), Procedures as a Representation for Data in a Computer Program for Understanding Natural Language. Massachusetts Institute of Technology Project MAC Report #MAC-TR-84.

DISCUSSION

Rosenfeld:

I am a little unclear as to where the background comes from. Are the points left white on your output simply the sum total of those areas that were too small to get put into the tree?

Kirsch:

No, the area that is white in those images is accounted for in two ways: first, the way that you mentioned, those regions which are too small to be accepted in the filtering process; secondly, the regions which are above the chosen threshold, which do not have to be thin lines because these are continuous images. You see, since we are not dealing with line images the regions that form boundaries in such images might themselves have substantial area. In the paper I give the total area of these objects and you see how much is in fact accounted for as boundary. This suggests immediately that some type of thinning algorithm might be appropriately used to fill in pictures of this sort.

Rosenfeld:

I am still a little bit lost. When you have a criterion for the homogeneity of a region and you then take connected components with respect to that criterion, I can see how very inhomogeneous regions break up into tiny components and thereby get left out, but if by definition a region means a connected component with respect to a homogeneity criterion, I don't see how there can be a region which is homogeneous. Can you enlighten me a little further, or maybe we will have to do it privately?

Kirsch:

Well, let me try to characterize the thing that's worrying you. You are concerned, as I understand it, with regions which are heterogeneous, that is which are not homogeneous, and you ask how it is that heterogeneous regions show up as single coherent connected objects. The reason is that the heterogeneity is a variable parameter and as one raises the value of this parameter, raises the threshold, regions that are heterogeneous now are, with respect to that threshold, no longer sufficiently heterogeneous to break up into separate parts, hence they become single connected objects. As you move in this type of tree diagram you are making more and more heterogeneous regions into single connected regions until finally the very top node represents one region, the whole image, which has a heterogeneity which is less than or equal to the maximum heterogeneity in the whole image.

Rosenfeld:

Your Table 1 shows the number of regions as a function of the contrast threshold; I want to mention a curious fourth power law discovered by Mott-Smith et al, namely: if you look at connected components of constant grey level, the number of such components is inversely proportional to the fourth power of the area. It would be of some interest empirically to ask what sort of power law is revealed by the sorts of components that you are obtaining.

Kirsch:

That's a very interesting point. Intuitively I would say that our experiments confirm that, although we haven't in fact done any formal experiment of that sort.

McCormick:

There is a great deal of similarity between what you are doing and the classical paper of Rogers on Numerical Taxonomy. As far as I can see, you first form homogeneous regions by introducing a contrast function and then map these regions into an association graph. Then you break up the graph by setting a threshold

criterion on the strength of the coefficients of association between neighbouring regions. This strategy, as far as I can determine, decomposes the association graph into a tree structure by a mechanism identical to the one used by Rogers.

But Rogers goes on and does something which I didn't hear discussed here and on which I would like your comments. Many of the nodes in this process may be, in some sense, superfluous. That is, one can have neighbouring points of nodal break-off for small perturbations of the coefficient of association. It's like making a phylogenic tree: one wants to get more significant nodes, nodes which are stable so that a little change in the threshold will not suddenly make a new configuration. For that reason Rogers introduced the concept of moat, i.e. how much you have to be able to move the contrast function to get a new node, and introduced other selection processes to select out the more stable nodes - i.e., in your work, the preferred regions. Do you have a similar clustering strategy in your work, or not?

Kirsch:

No, I don't. I view the decomposition tree as a set of morphological choices provided to a later syntactical process which selects the particular decomposition for an image. In the tree diagram I'm using there are many different ways of characterizing the whole image by choosing different sets of nodes across the tree, the so-called cuts. Now, which one is better? I don't know. I'm suggesting that the choice of nodes should be done by essentially a syntactical process, a process that uses other information than what is intrinsic in the image itself. This I offer as a suggestion for how to continue the program that Narasimhan and I, and others, have been interested in for so long: attempting to impose linguistic constraints which are, of course, not in the image but in the culture for looking at the image.