Pattern Recognition, Vol 20, No 5, pp 487-496, 1987
Printed in Great Britain

0031-3203/87 $3 00 + .00
Pergamon Journals Ltd

SEGMENTATION OF NATURAL SCENES

BIR BHANU*† and BAHRAM A. PARVIN‡
*Department of Computer Science, University of Utah, Salt Lake City, UT 84112 and
‡Intelligent Systems Group, University of Southern California, Los Angeles, CA 90007, U.S.A.

(*Received* 26 *November* 1985; *in revised form* 5 *February* 1987)

Abstract—A new simple and computationally efficient approach to image segmentation via recursive region splitting and merging is presented. Unlike other techniques the criterion for splitting is based on a generalization of a two-class gradient relaxation method and merging uses a test for mean gray level equivalency for adjacent regions. The technique is illustrated by providing results for both synthetic and natural scenes.

Image understanding   Multiclass images   Outdoor natural scenes
Relaxation   Image segmentation   Split and merge   Thresholding

1. INTRODUCTION

Image segmentation is one of the crucial steps in image analysis. Segmentation algorithms can be categorized into three different types: edge detection, clustering and region splitting and merging. The goal of a segmentation algorithm is to partition an image into regions, each having a homogeneous property such as intensity, color, texture, etc. Ultimately, segmentation provides a set of symbols and interrelationships between the symbols that are necessary for machine perception. Recursive segmentation based on the analysis of the distribution of features is one of the most popular and commonly used techniques for image segmentation.[8, 9, 12, 14] Many of these techniques make use of an elaborate peak location and selection procedure,[8, 9] which provides threshold values for the purpose of image segmentation. As an example, Ohlander *et al.*[9] select a peak from the histograms of nine features of a color image using a heuristic peak precedence measure. The computation of maxima and minima is complicated since minor changes must be distinguished from major ones. One of the shortcomings of these techniques is that small regions in a large image may not show a distinct peak in the histogram, even if these regions are distinct from their immediate neighborhood. Therefore, in the application of these techniques, normally the image is partitioned artificially into a set of subimages and each subimage is segmented independently.[6, 8, 9] As a result, a remerging measure may be required[8] to merge regions that are arbitrarily broken at the subimage boundaries. Very often this merging step leads to some regions which remain unmerged or overmerged.

In natural scenes with many different regions, the gray level histogram may have only one peak because the range of intensities for each region will probably overlap with the ranges of other regions. As a result of this overlap, the histogram is usually almost unimodal. In Refs. (2, 3) we presented a two-class gradient relaxation algorithm for the segmentation of images having unimodal gray level distributions. This basic algorithm has been compared with the Rosenfeld, Hummel and Zucker algorithm,[13] and has been successfully used in the extraction of objects from biomedical images, aerial images and tactical images.[4, 5] The tactical images were both in the visible and infrared spectrum. The algorithm allows control over the segmentation results and the rate of convergence of the relaxation process. It provides automatic selection of the threshold for binarization. In this paper we generalize the use of this algorithm for the segmentation of complex natural scenes.[1, 10, 11] Region splitting is accomplished via recursive application of the two-class gradient relaxation algorithm.[2, 3] It is followed by merging to resolve fragmentation. Two adjacent regions are merged if their mean intensities are statistically equal at a prescribed level of confidence.

Since the technique presented here allows the splitting of the image to depend on the outcome of the two-class gradient relaxation technique,[1, 3] it avoids any heuristic and arbitrary measures for partitioning the image. Thus, it does not suffer from the fact that some of the regions may be artificially broken at the subimage boundaries, and does not require a robust merging step in order to remove the boundary artifacts.[8] It does not require a detailed peak location

† Present address of the corresponding author: Bir Bhanu, Honeywell Systems and Research Center, 3660 Technology Drive, Minneapolis, MN 55418, U.S.A.

and selection procedure. Further, the approach does not directly depend on the peaks and valleys of the gray level histogram to initiate the relaxation process (initial probability assignment) and uses image statistics. Thus, the initial step of the relaxation process is a simple one which is used over and over again at the various levels of the region splitting hierarchy.

2. TWO-CLASS GRADIENT RELAXATION ALGORITHM

Bhanu and Faugeras[1,3] proposed a gradient relaxation algorithm for the segmentation of images having unimodal distributions. These distributions are typically obtained when the image consists mostly of a large background area with other small but significant regions. In such cases, the selection and location of the valley in the histogram is a nontrivial problem. Since this method is used as the basis for the segmentation of multiclass images, a brief description of the method is given below.

Suppose a set of N pixels $i = 1, 2, \ldots, N$ fall into two classes λ_1 and λ_2 corresponding to the white (gray value = 255) and black (gray value = 0) classes. Reduced inconsistency and ambiguity of pixels with respect to their neighbors are achieved by maximizing the global criterion,

$$C(\mathbf{p}_1, \ldots, \mathbf{p}_N) = \sum_{i=1}^{N} \mathbf{p}_i \cdot \mathbf{q}_i \tag{1}$$

subject to the constraint that $\mathbf{p}_i$'s are probability vectors. $\mathbf{p}_i$ is the probability that the ith pixel belongs to class λ_1 and λ_2. $\mathbf{q}$, the compatibility vector, is a function of $\mathbf{p}_i$'s. It is defined as,

$$\mathbf{q}_i(\lambda_k) = \frac{1}{V_i} \sum_{j \in V_i} \sum_{l=1}^{2} c(i, \lambda_i, j, \lambda_l)\, p_j(\lambda_l);$$

$$k = 1, 2; \quad i = 1, \ldots, N \tag{2}$$

where compatibility,

$$c(i, \lambda_k, j, \lambda_l) = \begin{cases} 0 & \text{if } k \neq l;\ k = 1, 2; j \in V_i; \text{ for } all\ i \\ 1 & \text{if } k = l;\ k = 1, 2; j \in V_i; \text{ for } all\ i \end{cases} \tag{3}$$

and V_i is the number of elements in the set of the nearest neighbors of the ith pixel. $\mathbf{q}_i(\lambda_k)$ is in effect the average of $\mathbf{p}_i(\lambda_k)$ of the eight nearest neighbors ($V_i = 8$).

Initially, at every pixel, the assignment of probabilities is,

$$p_i(\lambda_1) = FACT\frac{I(i) - IBAR}{255} + 0.5 \tag{4}$$

where, $I(i)$ is the gray value at the ith pixel and $IBAR$ is the mean of the image. $FACT$ is a function of intensity which is taken to be equal to 1 if $I(i) > IBAR$, otherwise its value is between 0.5 and 1. $FACT$ is related with the expected number of white and black pixels in the image. It does not affect the rate of convergence very much, but it affects the segmentation results.[3]

A projection gradient technique is used to solve the problem as stated in (1). The gradients of the criterion C in (1) with respect to λ_1 and λ_2 are $2q_i(\lambda_1)$ and $2q_i(\lambda_2)$, respectively. By computing the projection of this gradient and simplifying the equations for quick convergence, the iterative equations for the relaxation process are obtained.[1,3] These are,

$$p_i^{n+1}(\lambda_1) = p_i^n(\lambda_1)[1 - \alpha_1] + \alpha_1; q_i(\lambda_1) > 0.5; 0 < \alpha_1 < 1.0 \tag{5}$$

$$p_i^{n+1}(\lambda_1) = p_i^n(\lambda_1)[1 - \alpha_2]; q_i(\lambda_1) < 0.5; 0 < \alpha_2 < 1.0. \tag{6}$$

The magnitudes of α_1 and α_2 control the degree of smoothing at each iteration and their ratio controls the bias towards a class. The magnitude of $FACT$ controls the initial assignment of probabilities. Since we use the projection gradient method to obtain iterative equations, we named this method the "gradient relaxation method" to distinguish it from the nonlinear relaxation method of Rosenfeld *et al.*[13]

Figure 1 shows the segmentation results at two iterations and the corresponding histograms, for the segmentation of an aerial image. Note that at the first iteration we get two peaks separated by a valley, which define the binarization threshold. Further iteration causes the two peaks to move farther apart. This leads to the increase in contrast and the probabilities converge, as expected. When the peaks are sufficiently far apart, thresholding can be done at the mean value.

2.1. *Evaluation of the two-class gradient relaxation algorithm*

Since the two-class gradient relaxation technique is used as the basis for the region splitting procedure, its performance needs to be understood and evaluated. Three factors are used to evaluate the performance: signal-to-noise ratio, region size and contrast. The following methodology analyses the first two parameters.

(1) A 100 × 100 synthetic image is created which consists of a square region against a constant background. The size of the square is 50 × 50 pixels, and the intensity is 130. The background has a magnitude of 100. The gray scale range is 0–255. This image is noise-free, and the objective is to extract the square from its immediate background (see Fig. 2(a), top left image).

(2) White noise is added to the original image and the signal-to-noise ratio is varied from 1 to 10. Signal-to-noise ratio is defined as the square of the step edge amplitude divided by the standard deviation of the Gaussian white noise.

(3) *A figure of merit* is established to assess the amount of deviation between the noise-free segmentation, and the segmentation obtained on the noisy images. It is defined as the sum of the number of the square region pixels which are segmented as background pixels, and the number of the background pixels which are segmented as the square region pixels.

(4) The size of the square region is gradually reduced from 50 × 50 to 5 × 5 pixels and steps (2) and (3) are

(a)

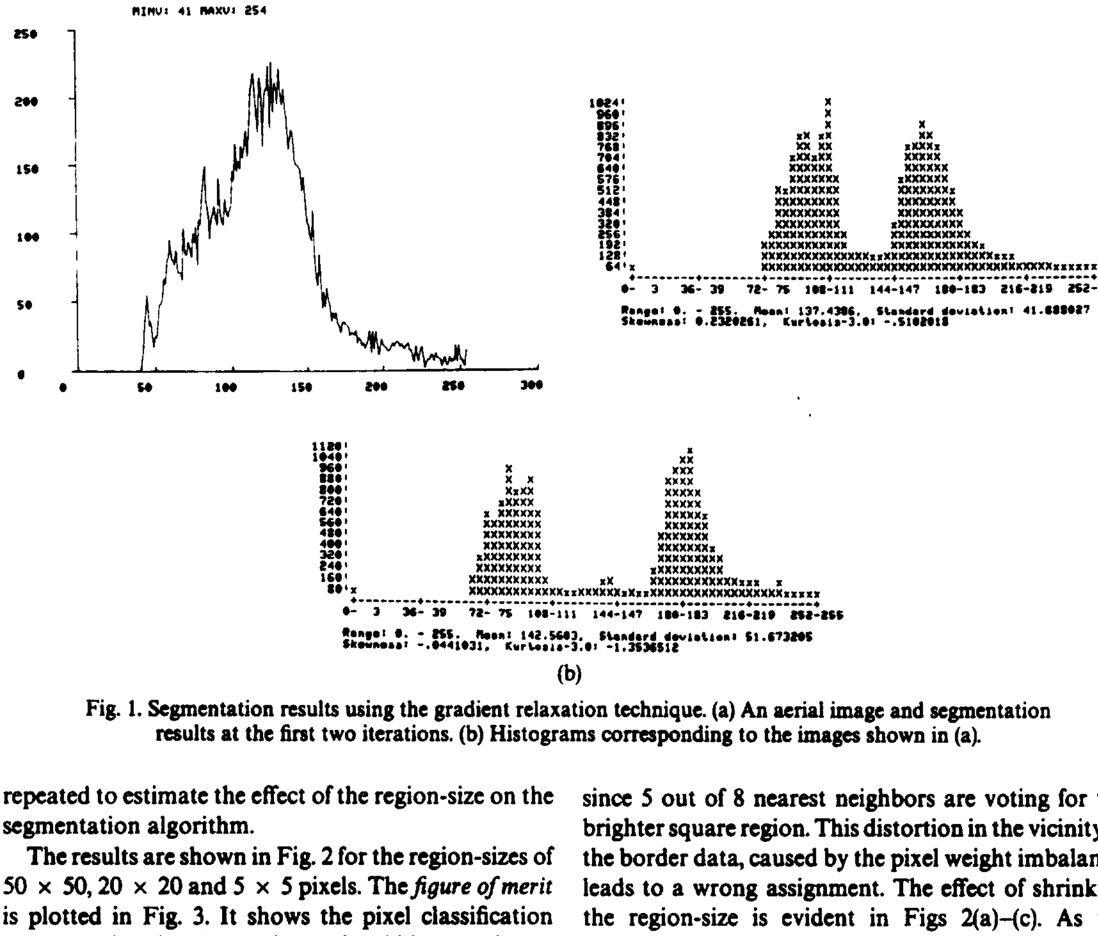

(b)

Fig. 1. Segmentation results using the gradient relaxation technique. (a) An aerial image and segmentation results at the first two iterations. (b) Histograms corresponding to the images shown in (a).

repeated to estimate the effect of the region-size on the segmentation algorithm.

The results are shown in Fig. 2 for the region-sizes of 50×50, 20×20 and 5×5 pixels. The *figure of merit* is plotted in Fig. 3. It shows the pixel classification error vs the signal-to-noise ratio. This experiment reveals that (a) the segmentation algorithm has a noise cleaning effect, (b) most of the erroneous labelings occur at the boundary pixels, and (c) as the region-size decreases, the number of mislabeled pixels in the background increases.

The noise cleaning effect is an important attribute of the iterative relaxation technique. The relaxation process allows the labeling at any pixel location to depend on the results of the previous iteration. Thus, the process becomes better informed as the analysis proceeds. The erroneous labeling of the border pixels is due to their inherent instability. For example, all the border pixels in the inner square (with the exception of the corner pixels) are biased toward the inner square, since 5 out of 8 nearest neighbors are voting for the brighter square region. This distortion in the vicinity of the border data, caused by the pixel weight imbalance, leads to a wrong assignment. The effect of shrinking the region-size is evident in Figs 2(a)–(c). As the region-size is reduced, the contribution of the smaller region to the global mean is decreased. This increases the initial probability and compatibility measure at each pixel location, which contributes to the increased number of mislabeled pixels.

The contrast can be defined as the difference between the square region's mean intensity and the background's mean intensity. It was found through experiments that the relaxation algorithm, as defined by equations (4)–(6), can partition the above synthetic images (Figs 2(a)–(c)), provided that the signal mean is approx. 25% above the background mean. Figure 4 shows the effects of contrast and noise on the region/target extraction. To enhance the performance of the segmentation algorithm, the estimation of the

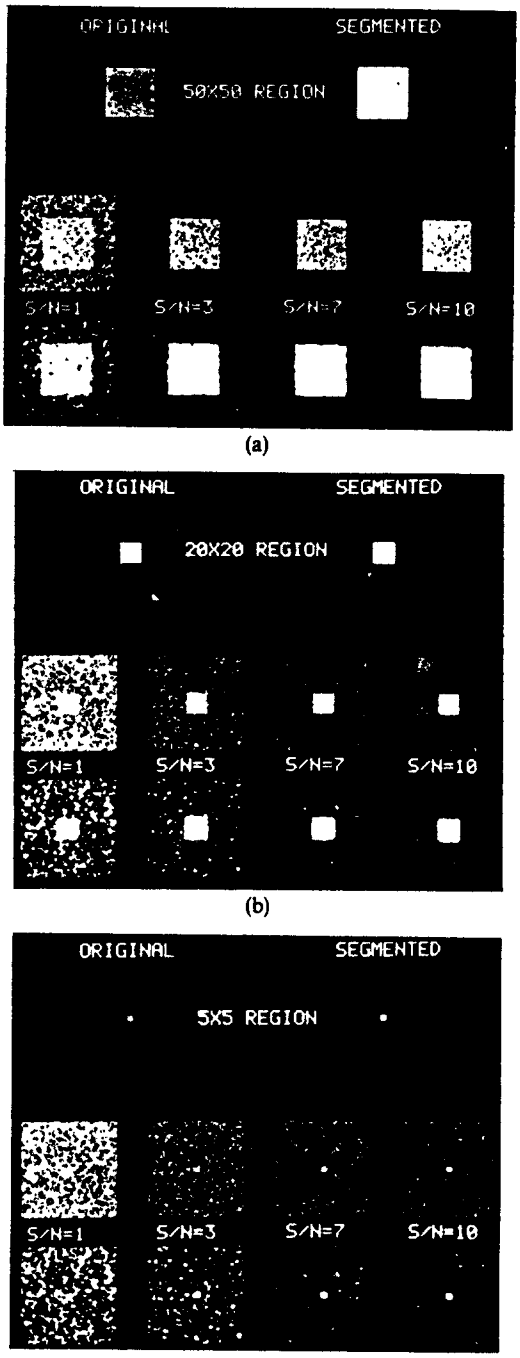

Fig. 2. Effect of noise on the two-class gradient relaxation technique for regions of various sizes with varying amount of signal-to-noise ratio. (a) Region size 50 × 50. (b) Region size 20 × 20. (c) Region size 5 × 5.

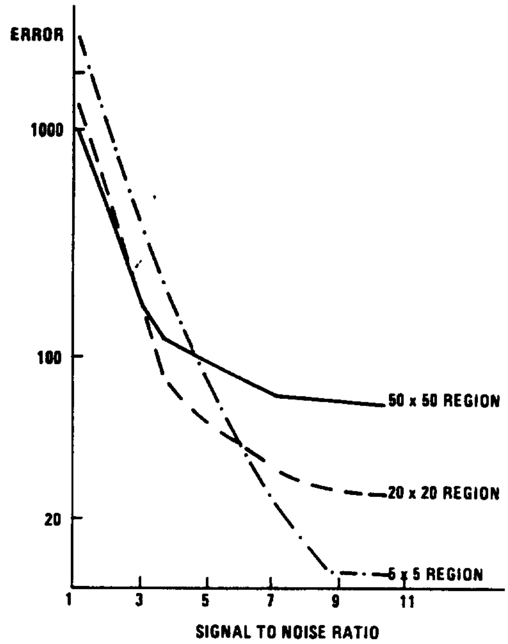

Fig. 3. Pixel classification error vs signal-to-noise ratio for the images in Fig. 2.

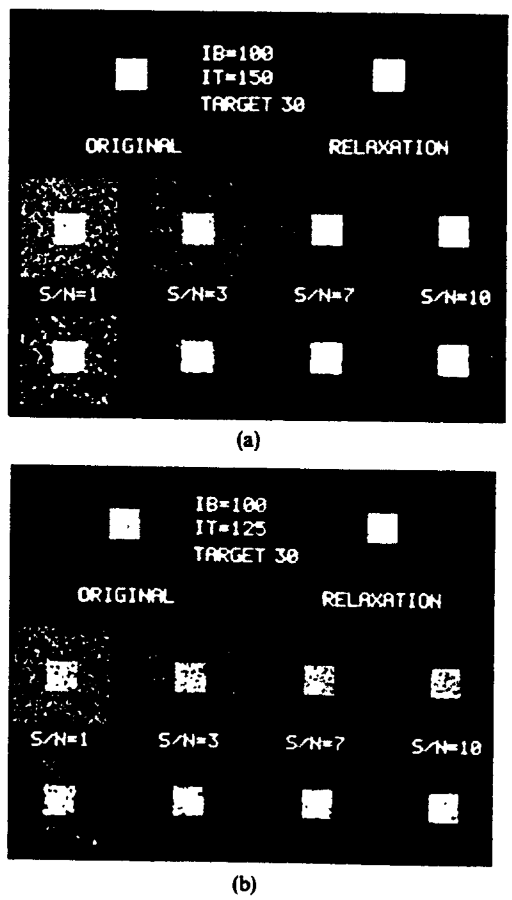

Fig. 4. Effect of contrast and noise on target extraction. (a) Target size 30 × 30, target intensity 150, background intensity 100. (b) Target size 30 × 30, target intensity 125, background intensity 100.

initial probability (equation (4)) was revised to include both mean and variance of the image. This is done by replacing *IBAR* with (*IBAR–bias*), where bias is defined in Fig. 5. Its value is obtained in such a way that the segmentation error across a wide range of contrast is minimized. With this change, the small square region was successfully extracted from the lower contrast image (with signal intensity = 110, and the background intensity = 100). Consider what has

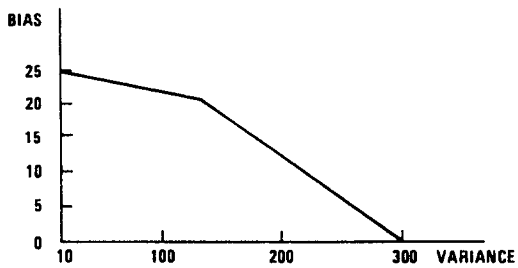

Fig. 5. Relation between bias and variance used in the initial assignment of probabilities.

occurred: as the variance of the image is decreased, the value of the bias is increased, thus increasing the initial probability assignment and local compatibility measure. Consequently, region extraction for low contrast images is possible.

2.2. *A real-time architecture for the two-class gradient relaxation algorithm*

To emphasize the simplicity of the two-class segmentation algorithm, its real-time implementation is also examined and evaluated. The implementation of a pipelined architecture shown in Fig. 6 consists of three components:

(1) computation of the initial probability,
(2) computation of the compatibility vector, and
(3) updating the probability vector.

Since the computation of the initial probability depends on the global mean and variance of the image, it is assumed that these two variables are computed from the previous frame. In other words, the global mean and variance do not change significantly from one frame to another. In many image understanding problems such as biomedical data analysis, the repeated computation of mean and variance is not necessary, since either the mean and variance are known *a priori*, or they have to be computed only once. The second component of the relaxation algorithm is the computation of a compatability vector, which is defined as the mean probability of the neighborhood. To align the eight nearest neighbors, two shift-registers and nine registers are employed (Fig. 6), and a tree of carry-save-adders is used to compute the local mean. Next, the most significant bit of the compatibility measure is checked to either increase or decrease the current probability assignment.

This architecture makes the application of the relaxation algorithm a viable approach to image segmentation in a real-time environment. Currently, the VLSI implementation of this algorithm is in progress.

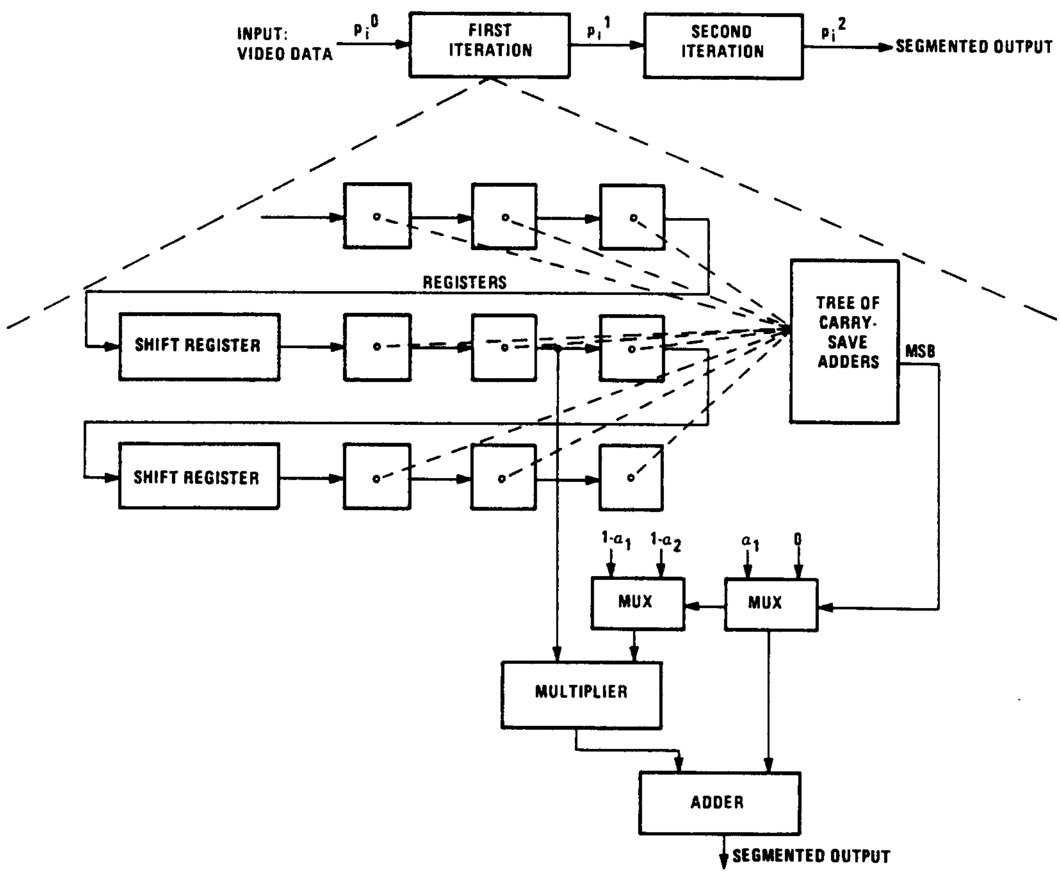

Fig. 6. Real-time implementation of the two-class gradient relaxation algorithm.

3. EXTENSION TO MULTICLASS IMAGES

3.1. *Recursive region splitting*

The multiclass segmentation algorithm is fundamentally a region splitting technique. Historically,[9, 14] these techniques rely on the analysis of the gray level (or color) histogram. In this context, the image is split into separate parts on the basis of region inhomogeneity. The location of sharp boundaries is determined either by locating the valley in the histogram, which often requires histogram smoothing to eliminate small peaks, or by partitioning the histogram into intervals where each interval satisfies an elaborate set of rules[14] such as,

(i) the interval must contain a marked peak,
(ii) the area of the interval should be above a prespecified threshold,
(iii) the highest peak in an interval should be above a certain amplitude, and
(iv) the valleys on each side of the interval should be sufficiently low.

The splitting is done by finding a well separated peak. If splitting can not be done, then either the region is considered to be segmented or in cases where there is a lack of spectral features (such as monochromatic images) arbitrary partitioning of the image is done and the splitting is done recursively for each partition.

The proposed multiclass segmentation algorithm depends on the recursive application of the two-class segmentation algorithm discussed earlier. Consequently, splitting of the image on the basis of the global histogram is avoided, and the segmentation at each level of the hierarchy depends on the local neighborhood.

Another important issue in region splitting is the ambiguity of the border pixels. An approach to this problem is "conservative thresholding"[14] where the original threshold is actually replaced by two thresholds, one to the left and the other one to the right of the original threshold. The pixels to the right of the rightmost threshold and left of the leftmost threshold are labeled normally, whereas the pixels that lie between the two thresholds are not labeled and region growing is used to fill in the labels for these pixels. In the proposed multiclass segmentation algorithm the border pixels are masked off and are not labeled initially. However, the border pixels are used to compute the compatibility coefficients for the pixels directly adjacent to them. After the segmentation process, region growing is used to label the border pixels.

In summary, the two-class segmentation algorithm is applied to the entire image, splitting the image based on its global and local properties into two distinct classes. Next, all the connected-components of the segmented image are isolated and labeled. Then each connected-component is used as a binary mask on the original image to partition it further. A region is said to be "partitionable" if it has a valley in its histogram after the application of the two-class relaxation procedure. As shown in Fig. 1, if such a valley exists, the selection of the threshold is trivial. This process continues recursively until a region is either nonpartitionable (its histogram remains unimodal), or its area is small enough such that further partitioning is of no interest. Currently, any region that has an area of less than one percent of the total image area is not partitioned further. In addition, any region that is less than 0.05% of the total image area is not labeled at all. These holes are either due to noise or sharp local intensity variations, and are filled in during the final analysis. The steps used in the multiclass segmentation algorithm are:

(1) Set a two-dimensional binary mask array true on the entire image. Initially, the entire image is one label or one segment.

(2) Turn off the mask pixels which touch the boundary of the label.

(3) Apply the two-class relaxation algorithm to the intensity image, wherever the mask array is set. The output of this step is a binary image.

(4) Perform region growing and assign a label (black or white) to the border pixels that were not processed by the previous step.

(5) Find all the connected components of the segmented image. For each component whose area exceeds the area threshold, (a) set mask pixels true in the region specified by the label and false everywhere else, (b) repeat steps 2–5 recursively.

Note that the recursive segmentation algorithm as described above does not assume that the number of classes is *a-priori* known. It is unsupervised and allows classes to be generated as required. This is accomplished by first extracting the dominant regions based on the global feature activity and then using these regions to structure finer segmentation in a hierarchical manner, so that the local feature activity is revealed.

3.2. *Region merging*

A side effect of any pure recursive region splitting algorithm is fragmentation which may lead to poor segmentation results for some images.[7] Therefore, we use a merging step which removes fragmentation and improves the performance of the relaxation algorithm for a wide class of images.

After the splitting algorithm has been applied, the image consists of a set of regions each of which is uniquely labeled. The goal of merging is to test whether the histogram of two adjacent regions is unimodal. If the test is successful then the two regions are merged; otherwise they remain separated. Such a test is usually achieved by evaluating a measure of the homogeneity of the gray levels of the larger region formed by merging the two regions. Assuming that adjacent regions X1 and X2 are of size $n1$ and $n2$ pixels respectively, there are $(n1 + n2)!/(n1!)*(n2!)$ ways to merge the two regions. Since $n!$ grows as n^n, several criteria have been used for merging.[8, 15] In this paper, adjacent regions are merged if their mean values are statistically equal at a certain level of confidence. Initially, we used a one way analysis of variance for merging adjacent regions, since it allows multiple segments to be tested simultaneously. Contextual and shape features were used to select a set of adjacent regions for analysis. However, this technique proved to be computationally intensive, and eventually only two neighbors were tested for region merging. An efficient technique for testing if two region's mean are equal is the Student's t test. This test assumes that the pixels intensity values in a region are independent. However, the intensity information in a given image is spatially correlated. For this reason, only a fraction of the pixels are used to compute the statistics of the adjacent regions. This fraction is 5% of pixels in each region. Further, if regions X1 and X2 are merged, then the statistic of the new region is updated and all of the new neighbors are recursively examined for further merging.

Note that since we are not arbitrarily partitioning the image, it is possible to avoid this merging step altogether if a smoothing operation (expansion and contraction) is carried out on the binary mask to eliminate small regions, holes and thin connections between regions. At present we have not implemented this idea.

3.3. *Results of the multiclass technique on synthetic and natural scenes*

Synthetic images: Three synthetic images are used to evaluate the performance of the multiclass segmentation algorithm under a controlled environment. These images and their grey values have also been used by Nagin *et al.*[8] and Price.[12] We have chosen these images so that our results can be directly compared with their results.

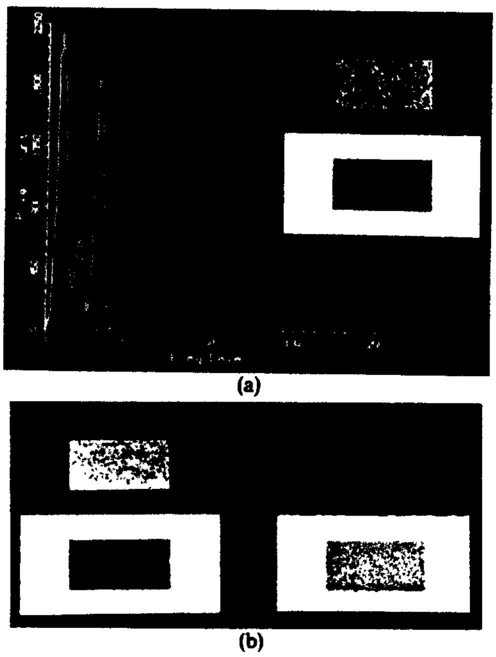

(a)

(b)

Fig. 7. (a) An artificial image and its histogram. The means are 10 and 25 for the top outer and inner rectangles and 40 and 14 for the bottom ones (with standard deviation = 3 for all cases). (b) Segmentation results using the recursive region splitting technique.

Example 1: Figure 7(a) shows a synthetic image with its gray level histogram. There are four regions in this image. The means are R1 = 10 and R2 = 25 for the top outer and inner rectangles, and R3 = 40 and R4 = 14 for the bottom ones with standard deviation = 3 for all cases. The histogram reveals three distinct peaks, as the distribution of the regions R2 and R4 is represented by the same peak. However, since R2 and R4 are not spatially adjacent, they are labeled as two separate regions. Note that R1 and R4 would be merged if they were spatially adjacent. In Fig. 7(b) the left image is the original image and the right image shows the segmentation results after region splitting. The multiclass technique produces results which are similar to those in Nagin *et al.*[8] and Price.[12]

This experiment was repeated by changing the intensity of the region R4 from 14 to 17, and the multiclass segmentation algorithm produced an identical result. The importance of this test is that as the intensity of R4 is changed from 14 to 17, its distribution becomes completely ambiguous in the histogram.[8, 12] As a result, classification techniques that depend solely on the global histogram peaks are subject to failure.[8]

Example 2: Figure 8(a) shows the previous image with R4 replaced with a ramp function which varies from 11 to 18 from the top to the bottom. The result is shown in Fig. 8(b). Evidently, the relaxation scheme did not correctly label a very small number of pixels. This is due to the very low contrast variation in R4.

Natural images: Two natural images of approximately 200 × 200 pixels size are examined. The first one is a green component of the color house image. It is shown in Fig. 9 together with its gray level histogram. The same image has been used by many researchers. Figure 10 shows the result of the two-class gradient relaxation algorithm on the original house image. It is evident from Fig. 10(b) that the selection of the binarization threshold is trivial. Figure 11(a) shows the segmentation result following region splitting. Figure 11(b) shows the outline of the regions obtained after the split-and-merge process. The results indicate that the sky, roof, front wall, bushes and windows are well segmented and most of the details in the image are maintained. Segmentation problems have occurred, however, at three locations: (i) the left corner of the roof where it is split into three parts; (ii) the left tree where it is merged with the left window; and (iii) the roof which is merged with the right tree.

The second natural image is shown in Fig. 12 together with its histogram. It is an aerial image consisting of several roads and a storage tank complex. The segmentation results after recursive region splitting are shown in Fig. 13(a), and the boundaries of the extracted regions after the split and merge process are shown in Fig. 13(b). Note that the results indicate

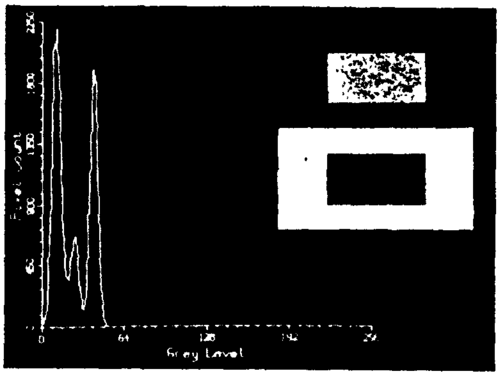

(a)

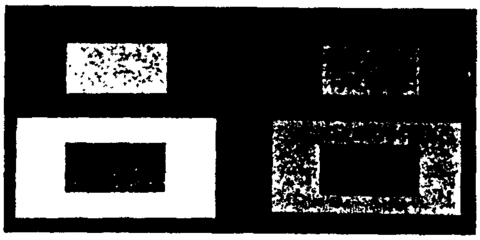

(b)

Fig. 8. (a) An artificial image and its histogram. The image is the same as in Fig. 7(a) except that the mean for the lower inner rectangle increases from 11 to 18 from the top to the bottom. (b) Segmentation results using the recursive region splitting technique.

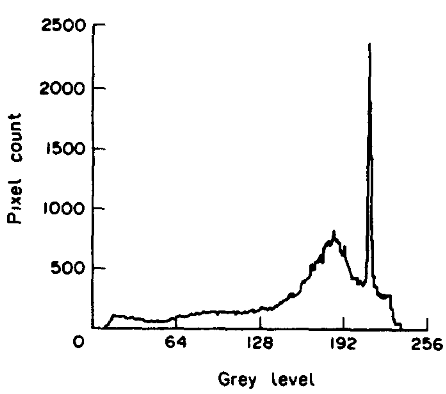

Fig. 9. A house image and its gray level histogram.

(a)

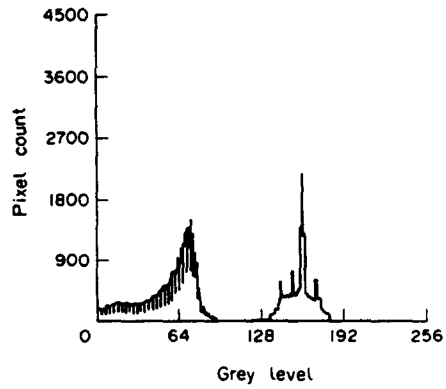

(b)

Fig. 10. (a) Effect of two-class gradient relaxation algorithm on the house image. (b) Histogram of the image corresponding to Fig. 10(a).

(a)

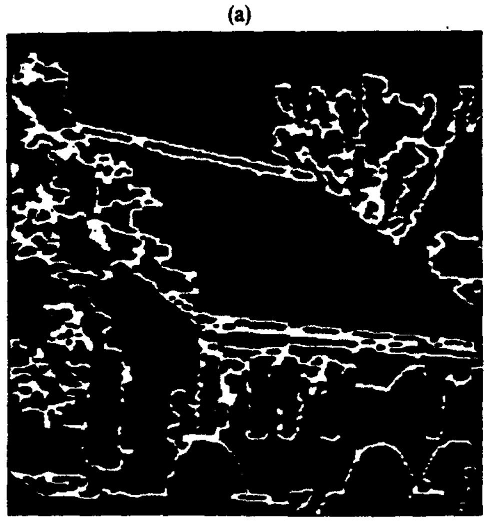

(b)

Fig. 11. (a) Results of recursive region splitting. (b) Edge results of the segmented image following the region merging step.

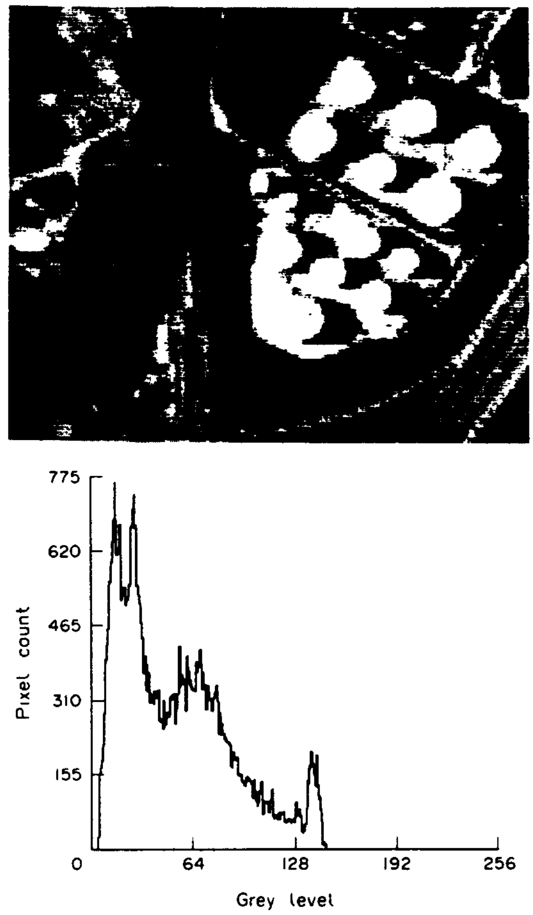

Fig. 12. An aerial image and its gray level histogram.

that the tank complex, main highway, curved road and the background are well segmented. However, some parts of the highway are merged with the background due to the small width of the road.

4. CONCLUSIONS

In this paper, the performance of the two-class gradient relaxation algorithm was evaluated. The evidence indicates successful segmentation in the presence of noise, low contrast and small region size. The simplicity of the algorithm was demonstrated by introducing a suitable real-time architecture. This implementation is possible because the algorithm requires only local processing for updating the initial probabilities in a subimage. Next, the two-class gradient relaxation algorithm was used as the basis for a recursive region splitting algorithm which was followed by a region merging algorithm to resolve fragmentation. One shortcoming of the proposed technique is its inability to extract line segments such as roads and highways. One possible solution is to make use of the intensity and edge information in an integrated manner.[4]

The multiclass segmentation technique presented here is a conceptually straightforward technique. It is a much simpler technique and provides results on natural scenes that are as good as the results of other techniques published in the literature.[8,9] It uses only the intensity feature as opposed to color or texture features, and does not require a detailed peak location and selection procedure. Furthermore, regions are not artificially broken at the subimage boundaries.

SUMMARY

Recursive segmentation based on the analysis of distribution of features is one of the most popular and commonly used techniques for image segmentation. Many of these techniques make use of an elaborate peak location and selection procedure which provides threshold values for the purpose of image segmentation. The computation of peak maxima and minima is complicated since minor changes must be distinguished from major ones. One of the shortcomings of these techniques is that small regions in a large image may not show a distinct peak in the histogram, even if they are distinct from their immediate neighborhood. Therefore, in the application of these techniques, normally the image is partitioned artificially into a set of subimages and each subimage is segmented and split further independently. As a result, a remerging measure may be required to merge the regions that are arbitrarily broken at the subimage boundaries. Very often this leads to some regions which remain unmerged or overmerged.

(a)

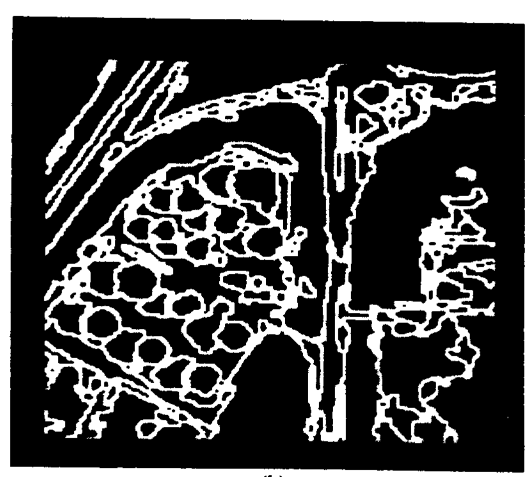

(b)

Fig. 13. (a) Results of recursive region splitting. (b) Edge results of the segmented image following the region merging step.

In this paper we present a simple and computationally efficient technique which does not have the above disadvantages and provides results that are as good as others published in the literature. It is based on the generalization of a two-class gradient relaxation algorithm for the segmentation of natural scenes. The two-class algorithm provides automatic selection of the threshold for binarization. First, we evaluate the performance of the two-class gradient relaxation technique with respect to signal-to-noise ratio, size of the objects and contrast present in the image and provide a real-time architecture. It is then extended for the segmentation of multiclass images. Region splitting is accomplished via recursive application of the two-class gradient relaxation algorithm. It is followed by merging which attempts to resolve fragmentation. Two adjacent regions are merged if their mean values are close to each other. The technique is illustrated by providing results for both synthetic and natural scenes.

REFERENCES

1. B. Bhanu, Shape matching and image segmentation using stochastic labeling, Technical Report USCIPI 1030, USC Image Processing Institute, Los Angeles, California (August 1981).
2. B. Bhanu, Segmentation of images using a gradient relaxation technique, *Proc. 6th Int. Conf. on Pattern Recognition*, Munich, West Germany, pp. 915–917 (1982).
3. B. Bhanu and O. D. Faugeras, Segmentation of images having unimodal distributions, *IEEE Trans. Pattern Anal. Mach. Intell.* **PAMI-4**(4), 408–419 (1982).
4. B. Bhanu and R. D. Holben, Model based segmentation of FLIR images, in *Proc. SPIE, Applications of Digital Image Processing VII*, No. 504, pp. 10–18. A. G. Tescher, Ed. (1984).
5. B. Bhanu, A. S. Politopoulos and B. Parvin, Intelligent autocueing of tactical targets, *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, pp. 502–503 (1983).
6. C. K. Chow .and T. Kaneko, Automatic boundary detection of the left ventricle from cineangiograms, *Comput. Graphics Biomed. Res.* **5**, 388–410 (1972).
7. A. R. Hanson, E. M. Riseman and P. A. Nagin, Author's reply, *IEEE Trans. Pattern Anal. Mach. Intell.* **PAMI-6**(2), 249 (1984).
8. P. Nagin, A. Hanson and E. Riseman, Studies in global and local histogram guided relaxation algorithms, *IEEE Trans. Pattern Anal. Mach. Intell.* **PAMI-4**(3), 263–277 (1982).
9. R. Ohlander, K. Price and D. R. Reddy, Picture segmentation using a recursive region splitting method, *Comput. Graphics Image Processing* **8**, 313–333 (1978).
10. B. A. Parvin, A split and merge algorithm for segmentation of natural scenes, *Proc. Seventh Int. Conf. Pattern Recognition*, pp. 294–296 (1984).
11. B. A. Parvin and B. Bhanu, Segmentation of images using a relaxation technique, *Proc. IEEE Conf. Computer Vision Pattern Recognition*, pp. 151–153 (1983).
12. K. Price, Image Segmentation: A comment on 'Studies in Global and Local Histogram Guided Relaxation Algorithms', *IEEE Trans. Pattern Anal. Mach. Intell.* **PAMI-6**(2), 247–249 (1984).
13. A. Rosenfeld, R. Hummel and S. Zucker, Scene labeling by relaxation operations, *IEEE Trans. Systems, Man and Cybernetics* **SMC-6**(6), 420–433 (1976).
14. S. Shafer and T. Kanade, Recursive region segmentation by analysis of the histogram, *Proc. Int. Conf. Acoustics, Speech Signal Processing*, pp. 1166–1171 (1982).
15. Y. Yakimovsky, Boundary and object detection in real world images, *J. Ass. Comp. Mach.* **23**(4), 599–618 (1976).

About the Author—BIR BHANU received the S.M. and E.E. degrees in Electrical Engineering and Computer Science from the Massachusetts Institute of Technology, the Ph.D. degree in Electrical Engineering from the University of Southern California and the M.B.A. degree from the University of California, Irvine.

At MIT Research Laboratory of Electronics he worked on the problem of computation of complex cepstrum. At USC he was at the Image Processing Institute where he worked on the problems of recognition of 2-D and 3-D objects, segmentation and motion analysis. In 1978 he worked with the computer science department of IBM research laboratory at San Jose in the image processing area. During the Winter of 1980 he was a research fellow at INRIA, the French National Research Institute in Computer Science where he worked on 3-D shape representation. From July 1981 to March 1984 he was an Engineering Specialist with the Aeronutronic Division of Ford Aerospace and Communications Corporation, Newport Beach, California. At Ford he worked on various aspects of the automatic target recognition problem and the industrial and tactical applications of image processing, pattern recognition and artificial intelligence. He has also been on the computer science faculty at the University of Utah. Currently he is working as a Research Staff Scientist at Honeywell Systems and Research Center, Minneapolis. He is also an Adjunct Associate Professor of Computer Science at the University of Utah. He is engaged in research in the areas of computer vision, robotics, pattern recognition and artificial intelligence.

He is a member of IEEE, IEEE Computer Society, ACM, AAAI, Sigma XI, SPIE and Pattern Recognition Society.

About the Author—BAHRAM PARVIN received his MSEE from Purdue University, Lafayette, Indiana. From 1978 to 1981, he was a Research Assistant at the Advanced Automation Research Laboratory, Purdue University. From 1981 to 1984 he was a Research Engineer at the Ford Aerospace & Communications Corporation, Newport Beach, California, where he was responsible for developing algorithms for tactical applications. During 1984–85 he was with Photonic Automation, Santa Ana, California, where he developed systems for machine vision applications. Currently he is with the Intelligent Systems Group of the University of Southern California. His areas of interest are pattern recognition, computer vision and artificial intelligence.