Model-Based Segmentation of FLIR Images

BIR BHANU, Senior Member, IEEE
University of Utah

RICHARD D. HOLBEN
Ford Aerospace & Communications Corp.

We present the use of gray scale intensities together with the edge information present in a forward-looking infrared (FLIR) image to obtain a precise and accurate segmentation of a target. A model of FLIR images based on gray scale and edge information is incorporated in a gradient relaxation technique which explicitly maximizes a criterion function based on the inconsistency and ambiguity of classification of pixels with respect to their neighbors. Four variations of the basic technique are considered which provide automatic selection of thresholds to segment FLIR images. A comparison of these methods is discussed and several examples of segmentation of ship images are given.

Manuscript received June 5, 1987.

IEEE Log No. 32636.

Authors' current addresses: B. Bhanu, Honeywell Systems and Research Center, 3660 Technology Drive, Minneapolis, MN 55418; R. D. Holben, Odetics, Inc., Anaheim, CA 92802.

I. INTRODUCTION

FLIR (forward-looking Infrared) sensors image the thermal radiations emitted by a target. Since the target is in direct contact with its environment, there is an exchange of heat between the target boundary surface and the immediate surroundings (background). The loss of contrast between the target and background with range to the observer is mainly caused by atmospheric scatter and absorption [5]. As a result, the boundary of the target in the image is not very sharp and many times there is an intensity gradient across the boundary. In the automatic recognition of tactical targets in FLIR images, an accurate and precise representation of the boundary of the targets is desired [3, 7]. It is very important since the features used in the classification of the target are normally based on the shape and gray scale of the segmented target and therefore, the performance of a statistical or a structural classifier critically depends on the results of segmentation [2]. Generally, only the gray scale of the image is used to extract the target from the background. Thus, the resulting segmentation depends upon several parameters of the technique used. The problem addressed here is, Can a better segmentation be obtained more efficiently if additional information is used beyond that based on the original unprocessed gray scale image? It is possible to obtain better segmentation by using different levels of information present in the image, for example by extracting information based on low-level primitives like edges or high-level information such as contextual cues or temporal cues based on a sequence of images.

We specifically consider the use of gray scale together with the edge information (gradient of the gray scale) present in the image to obtain a more precise segmentation of a target than a segmentation obtained by using either gray scale or edge information alone. A model of FLIR images based on gray scale and edge information is incorporated in a gradient relaxation technique [4] which explicitly maximizes a criterion function based on the inconsistency and ambiguity of pixel classifications with respect to their neighbors. Four variations of the basic relaxation technique are considered which provide automatic selection of thresholds to segment FLIR images. A comparison of these methods is discussed and several examples of segmentation of ship images at various ranges are presented.

II. SEGMENTATION APPROACHES

There are three general approaches for image segmentation: edge based, region based, and clustering based. Edge-based approaches are suitable for detecting linear features in the image [12, 15, 20]. They have the disadvantage of producing very low-level primitives (segments) even after considerable

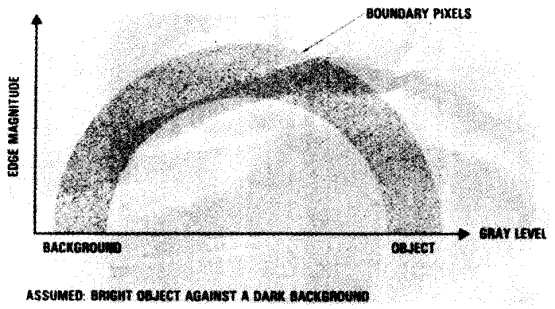

Fig. 1. Scatter plot of edge magnitude and gray value for typical FLIR image.

processing. A general approach to this problem of detecting significant edges is to treat a digitized gray scale image as a surface and attempt to associate large slope changes with edges [11]. However, modeling such a complicated surface near natural discontinuities is difficult and may require a simulated annealing approach to detecting boundaries [28]. Though techniques such as simulated annealing have reported success on degraded synthetic images [10], the large number of iterations preclude their consideration for real time applications. Region-based approaches use a region-growing or a region-splitting technique to achieve segmentation [6, 18, 23]. They have the advantage of producing higher level primitives, but the region so extracted may not correspond to actual physical objects unless the regions differ everywhere in intensity from the background [23]. Clustering-based methods use the cluster of local features to segment images [8, 18]. Since they do not use spatial relationship among pixels, they utilize region merging to obtain a better segmentation [18]. Perkins [22] uses an expansion-contraction approach to bridge gaps in the edge image of industrial parts, while Milgram [17] uses edge/border coincidence to segment targets in FLIR images. The "superslice" technique of Milgram is a nonprobabilistic technique that makes independent decisions based on gray scale and edge strength to determine if the edge border coincidence is sufficient to segment the target. Medioni [16] combines the advantages of edge-based and region-based techniques to segment aerial images, and Harlick and Shapiro [11] show that for nonreflective images better results are obtained by using binary edges to refine cluster-based segmentation.

Because no single technique may work best even over one image, several researchers have taken an expert system approach which incorporates several different techniques [19, 27]. This has the advantage of using control rules which incorporate strategies that vary depending upon local conditions during processing. However, for tactical applications, the expert system approach has significant real time implementation problems.

Schachter [25] has given a survey of FLIR target segmentation algorithms. A survey paper on automatic target recognition (ATR) by Bhanu [3] provides a comprehensive review of various techniques and their evaluation for the segmentation of FLIR images. In the following, we provide a model-based segmentation approach for the segmentation of FLIR images. We have found this approach to be very effective and efficient for the segmentation of low contrast images and for a large class of natural scenes for which the gray level histogram is almost unimodal.

III. MODEL-BASED SEGMENTATION OF FLIR IMAGES

The simplest model of a FLIR image is a two-class gray value model in which the target is brighter (class 1) than the background (class 2). The next higher level of modeling involves the interaction of gray and edge values in the image. The ideal gray value and edge magnitude model of a FLIR image is shown in the scatter plot of Fig. 1. The background and target are characterized by low and high gray values, respectively. The edge magnitude is low inside a target or background since it is assumed that they are uniformly distributed in intensity. At the pixels where the edge magnitude is high, the gray value is intermediate to that of the target and background. The pixels, where the edge magnitude is in the intermediate range, may be inside the target or background or along the border of the target. As an example Fig. 2 shows FLIR images of ship targets and their scatter plots.

Panda and Rosenfeld [21] use such a model to segment targets in air-to-ground FLIR images. Their procedure is not automated and the valley selection to threshold the image is done manually. Weszka and Rosenfeld [29] use a similar model by considering the histogram of pixels having low or high edge values to select the threshold to segment FLIR images. Danker and Rosenfeld [9] use edge and gray value separately and together in a probabilistic relaxation algorithm to segment the targets. In Bhanu and Faugeras [4], the probabilistic relaxation technique of Rosenfeld, Hummel, and Zucker (RHZ) [24] is compared with the gradient relaxation technique. There, it is shown that unlike the RHZ method, the gradient relaxation method provides control over the relaxation process by selecting three parameters which can be tuned to obtain the desired segmentation at a faster rate. It also allows some smoothing at every iteration. Smoothing is required for the segmentation of FLIR images since some parts of a target are generally hot, other parts are cold, and the background may contain some noisy elements. Normally, in a gray scale FLIR image the hot and cold parts of a target appear as bright and dark areas, respectively. Fig. 3 shows an example of such a ship target.

The technique presented by Bhanu and Faugeras [4] is modified here to incorporate the FLIR image

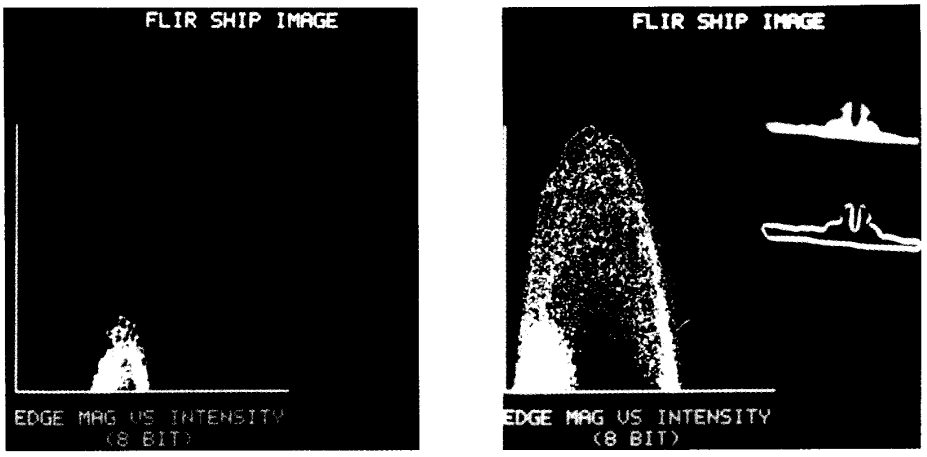

(a) (b)

Fig. 2. FLIR ship images and their scatter plots. (a) Low-contrast image. (b) High-contrast image.

Fig. 3. Ship image for which gray value distribution is not uniform.

model discussed above in the selection of compatibility coefficients and the assignment of initial probability to pixels so that the edge information and gray values interact in a controlled manner. Four variations of the technique are considered which provide automatic selection of thresholds to segment FLIR images.

A. Gray-Value-Based Segmentation—Basic Relaxation Technique

First we briefly describe the gray-value-based two-class segmentation technique [4] and then present several modifications which provide alternative ways of accomplishing model-based segmentation.

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. The probability that the ith pixel belongs to class λ_1 and λ_2 is represented by $\mathbf{p}_i$, while $\mathbf{q}_i$, the compatibility vector, is a function of $\mathbf{p}_i$s. It is defined as

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

where compatibility,

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

and V_i is the set of 9 pixels consisting of the ith pixel and its 8 nearest neighbors. The term $\mathbf{q}_i(\lambda_k)$ is in effect the average of $\mathbf{p}_i(\lambda_k)$ of the pixel and its eight nearest neighbors. In the rest of this paper, for the sake of simplicity, we denote $c(i,\lambda_k,j,\lambda_l)$ as $c(i,k,j,l)$.

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

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

where, $I(i)$ is the gray value at the ith pixel and IBAR is related to the mean and variance of the image [4, 6]. 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 [4].

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

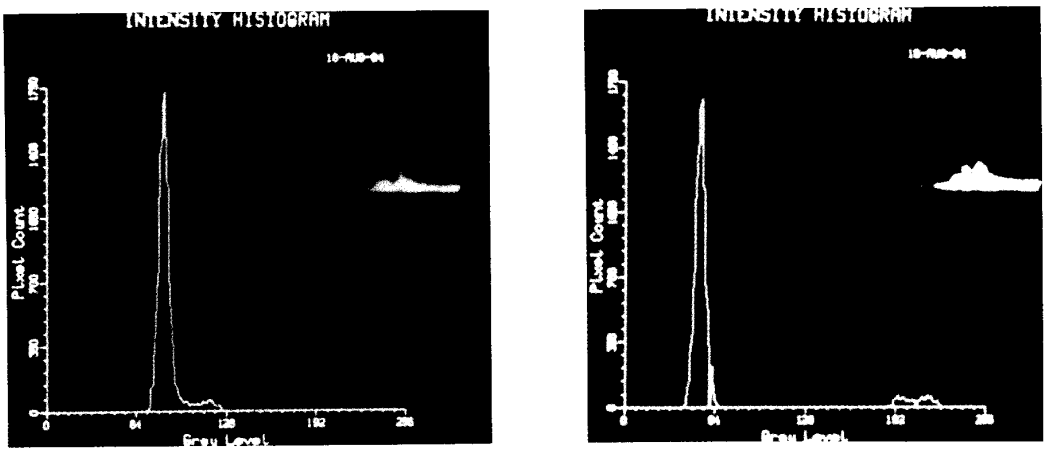

(a) (b)

Fig. 4. Segmentation results of basic relaxation technique. (a) Ship image and its histogram. (b) Segmentation of image shown in Fig. 4(a) and resulting histogram.

this gradient and simplifying the equations for quick convergence, the iterative equations for the relaxation process are obtained [1, 4]. These are,

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

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

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, Hummel and Zucker [24].

A few iterations of (5) and (6) result in the segmentation of images by allowing the automatic selection of thresholds. Fig. 4(a) shows the image of a ship target with its gray level histogram. Fig. 4(b) shows the segmentation results and the resulting histogram. Note that the selection of the threshold to segment the image is trivial. Fig. 5 gives a similar example of the segmentation of a ship target. As another example, Fig. 6(a) shows the image of a ship target and Figs. 6(b) and 6(c) show the results after two and five iterations. Note that a good definition of the boundary is obtained. However, one problem with the technique is that as the number of iterations increases, the target boundary extends more and more into the background. Here the use of edge information can be made to obtain a precise and accurate boundary and to automatically determine the number of iterations to stop the relaxation process.

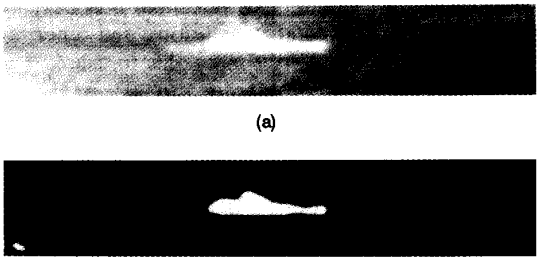

(a)

(b)

Fig. 5. Segmentation results of basic relaxation technique. (a) Ship image. (b) Segmentation of image shown in Fig. 5(a).

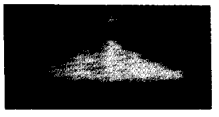

(a)

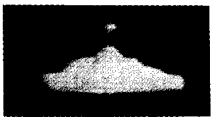

(b)

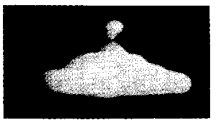

(c)

Fig. 6. Segmentation results of basic relaxation technique. (a) Ship image. (b) Segmentation results after 2 iterations. (c) Segmentation results after 4 iterations.

B. Modification of Gray-Value-Based Segmentation—Technique 1

In this approach, it is assumed that the edge magnitude and direction are available at every pixel. They can be computed by using an edge operator such as a Sobel operator which can be efficiently implemented or by a series of direction masks such as 6 masks in steps of 30° which assign each pixel the L_∞ norm of the 6 edge magnitude outputs and the corresponding direction. The choice of an edge operator depends on the quality of the edges present in the image, their location, response, and susceptibility to noise. Ideally, an operator is desired which is not very sensitive to noise. Thus, the use of edge operators such as Robert's is avoided.

A modification to the compatibility coefficients in (3) is made while keeping the initial assignment of probabilities in (4) the same. By thresholding the edge magnitude image, it is classified into two classes, high and low edge values. This threshold may be a function of range, weather conditions, and other environmental factors. Very simply, this threshold can be determined by considering 5 to 20% of the pixels where the edge magnitude is large to belong to the high edge value category. The actual percentage depends upon the complexity of the scene and range involved. Let the subset of pixels in the set V_i, where the edge magnitude is low or high, be denoted by $V_{i\mathrm{L}}$ and $V_{i\mathrm{H}}$, respectively. If the edge magnitude at the ith pixel is low, it is likely to be an interior point of the object or background. In this case, the compatibility coefficients can be obtained as given by (3) or only those pixels in the set V_i can be considered where the edge magnitude is not high. In the latter case,

$$c(i,k,j,l) = \begin{cases} 1, & \text{if } k = l \text{ and } j \in V_{i\mathrm{L}} \\ 0, & \text{if } k \neq l \end{cases} \tag{7}$$

and

$$q_i(\lambda_k) = \frac{1}{V_{i\mathrm{L}}} \sum_{j \in V_{i\mathrm{L}}} p_j(\lambda_k). \tag{8}$$

If the edge magnitude at the ith pixel is high, the pixel could be a noisy point of the object/boundary or a boundary point. To determine if it is an interior point of the object/background, the edge magnitude of its eight nearest neighbors is checked. If the edge magnitude of the majority of the neighboring pixels is low, it is a noisy point and the compatibility is determined as in (3). Otherwise, it is a boundary point and for this situation, the compatibilities are computed as

$$c(i,k,j,l) = \begin{cases} 1, & \text{if } k = l \text{ and } j \in V_D \\ 0, & \text{if } k \neq l \end{cases} \tag{9}$$

where V_D is the subset of pixels in $V_{i\mathrm{H}}$ such that the directions of the edges at pixel i and j are within ±30°.

Fig. 7. Segmentation of ship image shown in Fig. 3 by using technique 1.

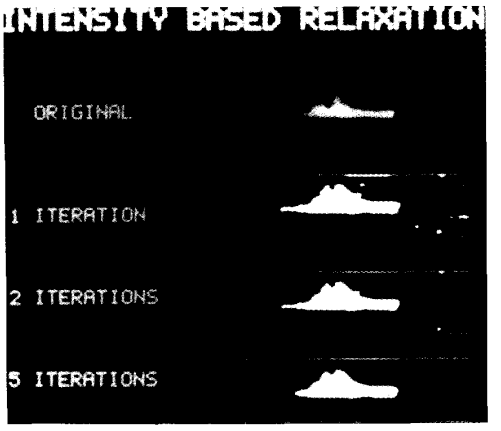

Fig. 8. Segmentation results at various iterations using technique 1.

Then,

$$q_i(\lambda_k) = \frac{1}{V_D} \sum_{j \in V_D} p_j(\lambda_k). \tag{10}$$

When (5) and (6) are iterated by incorporating the assignment of the compatibilities as described above, a bimodal histogram results where one peak corresponds to the background and the other peak to the boundary and target pixels. These peaks are separated sufficiently to provide an automatic selection of the threshold. Fig. 7 shows the segmentation results on the image shown in Fig. 3. Note that in spite of the significant gray value distribution in the target, a good definition of the boundary is obtained. Fig. 8 provides another example of segmentation with the results displayed at various iterations. In comparison with Figs. 6(b) and 6(c) from Fig. 8, note that as the number of iterations increase, boundary refinement takes place.

C. Edge Relaxation—Technique 2

In this technique, the model consists of two classes (edge/no edge) and an orientation parameter. The edge response at a pixel is reinforced by considering edge responses at neighboring pixels. Initially, the edge magnitude and direction are computed at every pixel in the image. The initial probability assignment at the

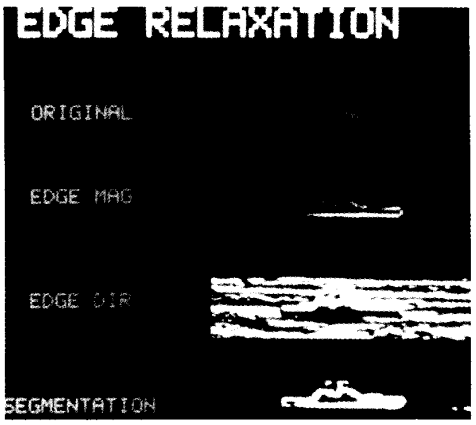

Fig. 9. Segmentation results using edge relaxation, technique 2.

Fig. 10. Threshold and thinned image corresponding to image in Fig. 5(a), technique 3.

ith pixel for edge/no edge is done by,

$$p_i(\lambda_1) = (E(i)/E_{\max}) \tag{11}$$

where λ_1 is the edge class, $E_{\max}$ is the maximum edge magnitude over the entire image and $p_i(\lambda_1) + p_i(\lambda_2) = 1$. λ_2 is the no-edge class. Initially, the orientation parameter θ_i at the ith pixel is taken as the direction of the edge at that pixel. Compatibility of a pixel with its neighbors is determined by,

$$c(i,k,j,l) = \begin{cases} 1, & \text{if } k = l \text{ and } j \in V_E \\ 0, & \text{if } k \neq l \end{cases} \tag{12}$$

where V_E is the subset of pixels in V_i such that the directions of the edges at pixel i and j are within $\pm 30°$. The compatibility vector is given by,

$$q_i(\lambda_k) = \frac{1}{V_E} \sum_{j \in V_E} p_j(\lambda_k). \tag{13}$$

Now the process is iterated as before in the two-class gray value case. The new orientation parameter at the end of an iteration is taken as the average of the orientation at the neighboring pixels which are within $\pm 30°$ i.e.,

$$\theta^{n+1}(i) = \frac{1}{V_E} \sum_{j \in V_E} \theta^n(j). \tag{14}$$

Note that in (12), we have defined the compatibility in a very simple way. A more accurate method is to use the cosine of the difference of the edge directions as a compatibility measure. Fig. 9 shows the edge relaxation results. A Sobel operator is used for edge detection. The segmentation result obtained is not as good as was achieved using technique 1.

D. Gray Value and Edge Relaxation in Parallel—Technique 3

In this method, a gray value relaxation is carried out as described in technique 1. Also, an edge value relaxation (technique 2) is carried out in the same two-class formulation. This provides an advantage in that we use the same technique once with gray values and the other time with edge values. At the end of each iteration, the coincidence of edge and border values is determined. This computation is performed after thresholding and thinning the edge magnitude image and thresholding the gray value image to obtain the boundaries of the objects. Thinning and thresholding the edge magnitude image is done using the following set of rules:

1) edge magnitude of the central pixel > TH1, a threshold determined from the histogram of edge magnitude;
2) edge magnitude of the central pixel > edge magnitude at its 2 neighbors in a direction normal to the direction of this edge;
3) edge direction of the two neighboring pixels (as above) are within $\pm 30°$ of the central pixel.

When the coincidence factor is greater than 50%, the process stops and the gray value image provides the segmentation of target. Fig. 10 shows the results obtained after thresholding and thinning the edge relaxation results on the image shown in Fig. 5(a). Other than automatically stopping the relaxation process, the boundary definition is not substantially better than the first two techniques. Also, it is more expensive computationally than any of the other techniques described here.

E. Interaction of Gray and Edge Values in a Joint Relaxation Process—Technique 4

In technique 3 as presented above, gray value and edge information do not directly interact. The joint relaxation process described below provides edge and gray value interactions in the initial assignment of probabilities and interprocess compatibilities. The initial assignment at a pixel for class λ_1 (white, gray value = 255) and λ_2 (black, gray value = 0) is done by

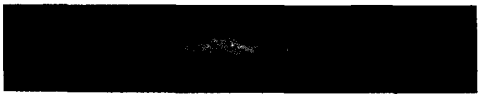

Fig. 11. Segmentation results using joint relaxation, technique 4.

Fig. 12. Segmentation of image shown in Fig. 5(a), technique 4.

using,

$$p_i(\lambda_1) = B_1\left(\frac{I(i) - \text{IBAR}}{255}\right) + B_2\left(\frac{E(i) - \text{EBAR}}{E_{\max}}\right) + 0.5 \quad (15)$$

$$p_i(\lambda_2) = 1 - p_i(\lambda_1) \quad (16)$$

where B_1 and B_2 are taken as 0.5 multiplied by the inverse of the standard deviation of gray value and edge magnitude, respectively, and EBAR is the average of the edge magnitude. IBAR is the same as defined in (4). The compatibility coefficients are obtained as described in the modification of the basic relaxation technique. They are based on gray scale and edge direction. Edge direction at every iteration is modified by (14). Note that in this method, edge magnitude is used only once in the initial assignment of probabilities and edge direction is used in the compatibility coefficients. Since the edge direction is updated at every iteration, it is effectively used in going from one iteration to the next. The result obtained using this technique for the image shown in Fig. 5(a) is given in Fig. 11. Observe that although the top of the target is little smoothed out, a better overall boundary representation of the ship is obtained. Another example is shown in Fig. 12 where the initial probability assignment is also included. Figs. 13 and 14 present several examples of segmentation of ship targets at various ranges using technique 4. In these figures the left-hand column shows the original images and the right-hand column shows the segmentation results. In Fig. 14, we have shown the boundaries of the segmented ship targets. Note that the technique is quite effective and produces a faithful representation of the boundary. This is specially true in those difficult situations when the targets are quite far away and contrast is very poor (see Fig. 13).

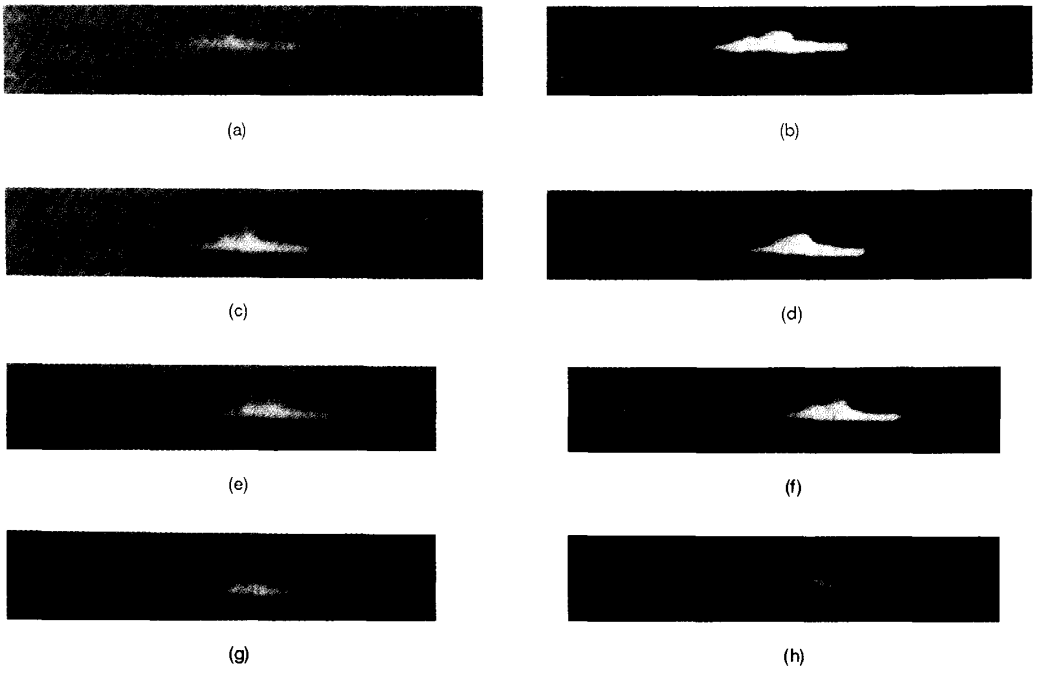

Fig. 13. Segmentation results for images at far ranges, technique 4.

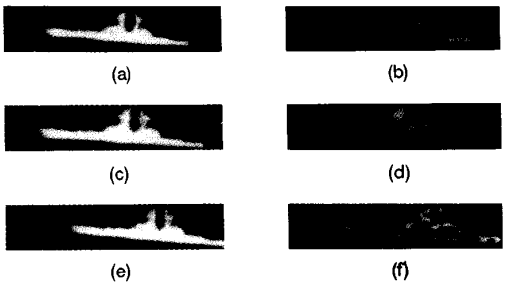

Fig. 14. Segmentation results for images at relatively close ranges, technique 4.

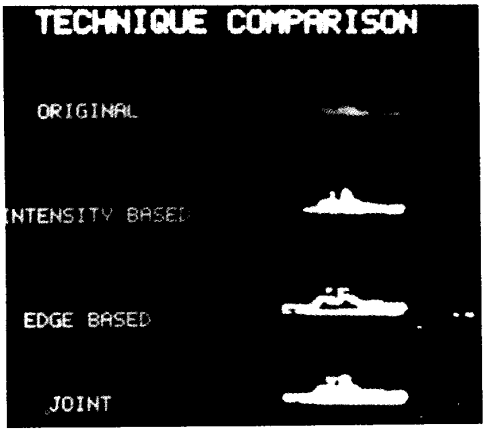

Fig. 15. Comparison of techniques 1, 2, 4.

IV. CONCLUSIONS

We presented four techniques for the segmentation of targets in images. Fig. 15 shows the comparison of techniques 1, 2, and 4. Note that technique 4, where we make use of gray scale together with the edge information, produces a more precise segmentation of the target than a segmentation obtained by using gray scale or edge information alone. It is also very computationally efficient. Technique 3 is the worst in the amount of computational effort required. In practice, for a given class of images such as ship images, we stop after a few predetermined number of iterations (normally 2 to 5). The technique has been tested on a large data base of about 7500 images. Although a number of performance criteria, such as the number of pixels misclassified, correlation coefficient, mean square error, etc., exist to evaluate the segmentation results quantitatively [3, 6], we have done a qualitative evaluation of results on our data base. Our qualitative evaluation has found that the technique 4 is very effective for the segmentation of low-contrast images. This technique is in the spirit of the basic two-class gradient relaxation algorithm for which we have done a comprehensive evaluation with respect to signal-to-noise ratio, region size, and contrast [6]. We find that the technique can extract target even when the contrast of the target is as low as 10% with respect to its immediate background. The technique also has a noise-cleaning effect. This is an important attribute of the technique where the relaxation process allows the labeling at any pixel location to depend upon the results of the previous iteration. Thus, the process becomes better informed as the analysis proceeds and allows smoothing within the image. This smoothing leads to the extraction of targets which have significant intensity gradient across them. Recently the technique has been successfully implemented in a 3 micron CMOS VLSI technology, using the path programmable design methodology (PPL), to work at TV frame rates [13]. We have also used the technique for the segmentation of low-contrast cell images, aerial images, ground targets, and the segmentation of natural scenes [3, 4, 6].

Several other variations of our basic technique are possible. For example, initial probability assignment can be based on other statistical models of the image such as those discussed by Jordan and Flachs [14] and Shazeer [26] for gray scale images. Furthermore, the assignment of compatibilities can be done so that we can weight the interior and exterior of the target differently. Also a more general formulation of technique 1 can be obtained by considering white, black, edge, no edge, and orientation as separate classes.

ACKNOWLEDGMENT

The authors are grateful to Ford Aerospace Corporation, Newport Beach, CA, where this work was performed.

REFERENCES

[1] Bhanu, B. (1981)
Shape matching and image segmentation using stochastic labeling.
Technical Report USCIPI 1030, Image Processing Institute, University of Southern California, Los Angeles, Aug. 1981.

[2] Bhanu, B. (1983)
Evaluation of automatic target recognition algorithms.
Proceedings of Society of Photo-Optical Instrumentation Engineers, **435** (Aug. 1983), 18–27.

[3] Bhanu, B. (1986)
Automatic target recognition: state of the art survey.
IEEE Transactions on Aerospace and Electronic Systems, **AES-22**, 4 (July 1986), 364–379.

[4] Bhanu, B., and Faugeras, O. D. (1982)
Segmentation of images having unimodal distributions.
IEEE Transactions on Pattern Analysis and Machine Intelligence, **PAMI-4**, 4 (July 1982), 408–419.

[5] Bhanu, B., and Holben, R. D. (1984)
Model based segmentation of FLIR images.
Proceedings of the Society of Photo-Optical Instrumentation Engineers, **504** (Aug. 1984), 10–18.

[6] Bhanu, B., and Parvin, B. A. (1987)
Segmentation of natural scenes.
Pattern Recognition, **20**, 5 (1987), 487–496.

[7] Bhanu, B., Politopoulos, A. S., and Parvin, B. (1983)
Intelligent autocueing of tactical targets.
In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, June 1983, 502–503.

[8] Coleman, G. B. (1977)
Image segmentation by clustering.
Technical Report USCIPI 750, Image Processing Institute, University of Southern California, Los Angeles, 1977.

[9] Danker, A. J., and Rosenfeld, A. (1981)
Blob detection by relaxation.
IEEE Transactions on Pattern Analysis and Machine Intelligence, **PAMI-3**, 1 (Jan. 1981), 79–92.

[10] Geman, S., and Geman, D. (1984)
Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images.
IEEE Transactions on Pattern Analysis and Machine Intelligence, **PAMI-6**, 6 (Nov. 1984), 721–741.

[11] Harlick, R. M., and Shapiro, L. G. (1985)
Image segmentation techniques.
Computer Vision, Graphics and Image Processing, **29** (1985), 100–132.

[12] Healy, D. J. (1987)
Improved edge-based image segmentation.
Proceedings of the Society of Photo-Optical Engineers, **845** (1987), 71–77.

[13] Bhanu, B., Hutchings, B. L., and Smith, K. F. (1990)
VLSI design and implementation of a real-time segmentation processor.
International Journal of Machine Vision and Applications, January 1990.

[14] Jordan, J. B., and Flachs, G. M. (1987)
Statistical segmentation of FLIR images.
Proceedings of the Society of Photo-Optical Engineers, **754** (1987), 220–228.

[15] Lineberry, M. (1982)
Image segmentation by edge tracing.
Proceedings of the Society of Photo-Optical Engineers, **359** (1982), 361–368.

[16] Medioni, G. R. (1983)
Segmentation of images into regions using edge information.
Technical Report 101, Intelligent Systems Group, University of Southern California, Los Angeles, Mar. 1983.

[17] Milgram, D. L. (1979)
Region extraction using convergent series.
Computer Graphics and Image Processing, **11** (1979), 1–12.

[18] Nagin, P., Hanson, A., and Riseman, E. (1982)
Studies in global and local histogram guided relaxation algorithms.
IEEE Transactions on Pattern Analysis and Machine Intelligence, **PAMI-4**, 3 (May 1982), 263–277.

[19] Nazif, A. M., and Levine, M. D. (1984)
Low level image segmentation: an expert system.
IEEE Transactions on Pattern Analysis and Machine Intelligence, **PAMI-6**, 5 (Sept. 1984), 555–577.

[20] Nevatia, R., and Babu, K. R. (1980)
Linear feature extraction and description.
Computer Graphics and Image Processing, **13** (1980), 257–269.

[21] Panda, D. P., and Rosenfeld, A. (1978)
Image segmentation by pixel classification in (gray level, edge value) space.
IEEE Transactions on Computers, **C-27**, 9 (Sept. 1978), 875–879.

[22] Perkins, W. A. (1980)
Area segmentation of images using edge points.
IEEE Transactions on Pattern Analysis and Machine Intelligence, **PAMI-2**, 1 (Jan. 1980), 8–15.

[23] Price, K. E. (1976)
Change detection and analysis in multi-spectral images.
Ph.D. dissertation, Department of Computer Science, Carnegie Mellon University, Pittsburgh, PA, Dec. 1976.

[24] Rosenfeld, A., Hummel, R., and Zucker, S. (1976)
Scene labeling by relaxation operations.
IEEE Transactions on Systems, Man and Cybernetics, **SMC-6**, 6 (June 1976), 420–433.

[25] Schachter, B. J. (1982)
A survey and evaluation of FLIR target detection/segmentation algorithms.
In *Proceedings of the DARPA Image Understanding Workshop*, Sept. 1982, 49–57.

[26] Shazeer, D. J. (1982)
Performance measures for statistical segmentation.
Proceedings of the Society of Photo-Optical Engineers, **359** (1982), 348–360.

[27] Shu, J., and Freeman, H. (1987)
An expert system for image segmentation.
Proceedings of the Society of Photo-Optical Engineers, **829** (1987), 240–252.

[28] Terzopoulos, D. (1984)
Regularization of inverse visual problems involving discontinuities.
IEEE Transactions on Pattern Analysis and Machine Intelligence, **PAMI-8**, 4 (July 1984), 413–424.

[29] Weszka, J. S., and Rosenfeld, A. (1979)
Histogram modification for threshold selection.
IEEE Transactions on Systems, Man, and Cybernetics, **SMC-9**, 1 (Jan. 1979), 38–52.

Bir Bhanu (S'72—M'82—SM'87) received the B.S. degree (with Honors) in electronics engineering from the Institute of Technology, B.H.U., Varanasi, India, the M.E. degree (with Distinction) in electronics engineering from Birla Institute of Technology and Science, Pilani, India, 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 Image Processing Institute, University of Southern California, and the M.B.A. degree from the University of California, Irvine. He also received a Diploma in German from B.H.U., Varanasi, India.

Dr. Bhanu is a Research Fellow at Honeywell Systems & Research Center, where he serves as the Principal Investigator of the Scene Dynamics Program from DARPA, Obstacle Detection Program from NASA, and the Machine Learning Program from a government agency. Additionally, he is conducting sponsored and IR&D research efforts on robotic vehicle navigation, machine learning, object modeling, contextual analysis, multisensor integration, parallel algorithms, and photointerpretation and surveillance. He has also worked with IBM on image processing, INRIA-France on 3-D object recognition and Ford Aerospace & Communications Corporation on automatic target recognition. While on the faculty of the Department of Computer Science, University of Utah, he was the Principal Investigator on several NSF and industry-funded research projects in machine intelligence. Presently, he is an adjunct faculty member in Computer Science at the University of Utah. He has given national short courses on intelligent automatic target recognition and he is a reviewer to over two dozen technical publications and government agencies. His current interests are computer vision, robotics, machine learning, neural networks and neurocomputers, object modeling, multisensor integration, distributed sensing and control, parallel computer architectures, and applications of artificial intelligence.

Dr. Bhanu has over 100 reviewed technical publications and 5 patents in the areas of his interests. He is the guest editor of a special issue of *IEEE Computer* on "CAD-Based Robot Vision" published in August 1987. He is listed in the *American Men and Women of Science*, *Who's Who in the West*, *Personalities of Americas* and *5000 Personalities of the World*. He is a member of the Association for Computing Machinery, the American Association for Artificial Intelligence, Sigma Xi, Pattern Recognition Society, Society for Photo-Optical Instrumentation Engineers, Institute of Electrical and Electronics Engineers (Senior Member), the IEEE Computer Society.

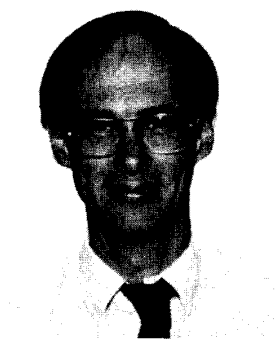

Richard D. Holben received a B.A. degree in mathematics and a B.S. degree in physics from Washington State University, Pullman, and the M.A. degree in physics from the University of Rochester, New York, in 1969.

Mr. Holben worked on active and passive sensor systems for object acquisition, tracking, and identification at Ford Aerospace, Newport Beach, Calif., from 1969 to 1988 where he was head of an algorithm development group for computer vision. Currently, he is a principal engineer in the Artificial Intelligence group at Odetics, Inc. working on vision systems. His current areas of interest include motion analysis and sensor fusion.

Mr. Holben is a member of the IEEE Computer Society, Society of Photo-Optical Instrumentation Engineers, and the American Association for Artificial Intelligence.