AD-A089 907 DEFENCE RESEARCH ESTABLISHMENT VALCARTIER (QUEBEC) F/G 17/5
EVALUATION OF A CLASS OF SEGMENTERS FOR IR IMAGERY (EVALUATION --ETC(U)
MAY 80 L SEVIGNY
UNCLASSIFIED DREV-R-4172/80 NL

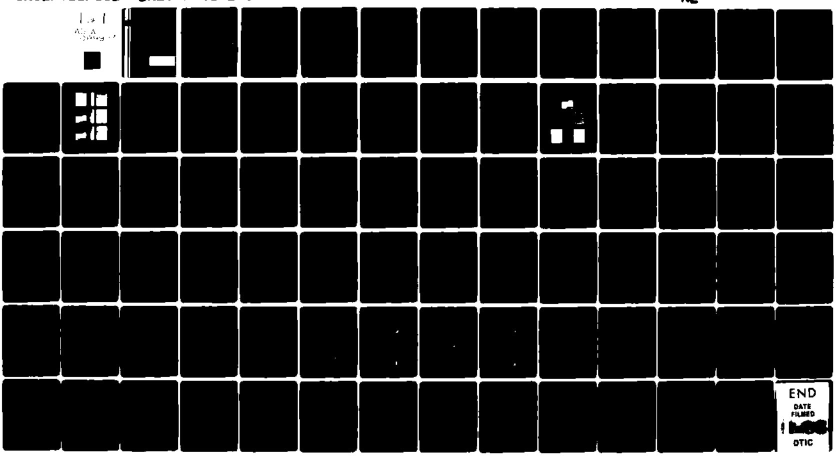

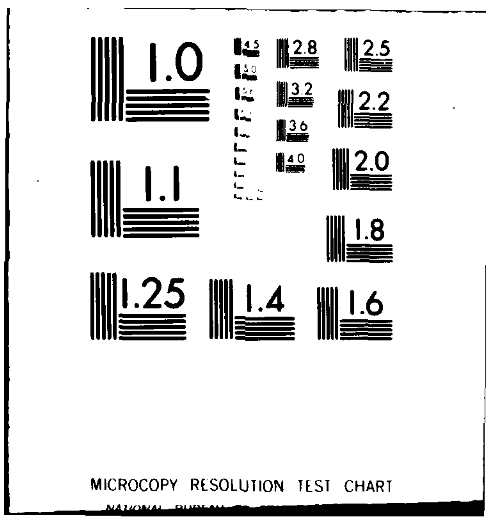
1.0
1.1
1.25
1.4
1.6
1.8
2.0
2.2
2.5
2.8
3.2
3.6
4.0
4.5
MICROCOPY RESOLUTION TEST CHART

UNLIMITED DISTRIBUTION ILLIMITÉE

UNCLASSIFIED
UNLIMITED DISTRIBUTION

CRDV RAPPORT 4172/80
DOSSIER: 3621J-011
MAI 1980

DREV REPORT 4172/80
FILE: 3621J-011
MAY 1980

3

AD A089907

LEVEL II

EVALUATION OF A CLASS OF SEGMENTERS FOR IR IMAGERY

L. Sevigny

Centre de Recherches pour la Défense
Defence Research Establishment
Valcartier, Québec

BUREAU - RECHERCHE ET DEVELOPPEMENT
MINISTERE DE LA DÉFENSE NATIONALE
CANADA

NON CLASSIFIÉ
DIFFUSION ILLIMITÉE

RESEARCH AND DEVELOPMENT BRANCH
DEPARTMENT OF NATIONAL DEFENCE
CANADA

80 9 29 025

CRDV R-4172/80
DOSSIER: 3621J-011

UNCLASSIFIED

14 DREV-R-4172/80
FILE: 3621J-011

11 May 80

1 **EVALUATION OF A CLASS OF SEGMENTERS FOR IR IMAGERY** (Evaluation d'une Classe de Segmenteurs pour Images IR),

by

10 L. Sevigny

12 81

CENTRE DE RECHERCHES POUR LA DEFENSE

DEFENCE RESEARCH ESTABLISHMENT

VALCARTIER

Tel. (418) 844-4291

Quebec, Canada

May/mai 1980

NON CLASSIFIE

404945

RESUME

Dans ce rapport nous évaluons l'efficacité de 6 algorithmes spécialisés dans la segmentation de cibles sur images IR. Ces algorithmes de segmentation, ou segmenteurs, reposent tous sur le principe voulant que la signature thermique d'une cible soit supérieure à celle de tout objet de l'arrière-plan. Les 3 premiers segmenteurs abordent l'image de front, en son entier, tandis que les 3 derniers incorporent la technique de redressement de l'arrière-plan (TRAP), visant à éliminer l'arrière-plan en tout ou en partie en le nivelant. Les divers algorithmes sont jugés d'après a) leur taux d'extraction, b) la fidélité du processus de segmentation en ce qui concerne les propriétés géométriques des cibles cibles et,finalement, d'après c) le degré d'individualisation imprimé aux cibles extraites par rapport aux pseudo-cibles. Les 3 segmenteurs centrés sur TRAP ont une meilleure probabilité d'extraction que les trois autres qui essayent d'éliminer l'arrière-plan simplement en morcelant l'image. Par ailleurs, la plupart des segmenteurs considérés dans ce rapport altèrent d'une façon ou d'une autre la forme des cibles. Les deux exceptions à cette règle sont les segmenteurs No. 1 (génerateur de silhouettes à seuil d'intensité unique) et No. 6 (le précédent segmenteur allié à une version particulière de TRAP). Les résultats expérimentaux montrent que l'intensité, le contraste et la variance sont les traits qui permettent le mieux de départager les cibles et les pseudo-cibles. Des expériences de classification réalisées à l'aide du segmenteur No. 6, lequel s'avère le meilleur, en fonction de ces traits caractéristiques indiquent que l'on peut espérer obtenir un taux de détection qui excède 90% avec un taux de fausses alarmes inférieur à 3%. (NC)

ABSTRACT

This report presents an evaluation of the performance of 6 algorithms dedicated to segmentation of targets in IR imagery. These segmentation algorithms or segmenters are based on the single assumption that the targets display a larger thermal signature than the background. The first 3 segmenters deal with an image in its entirety, whereas the last 3 incorporate the Background Elimination Technique (BET), which aims at eliminating wholly or partly the background by levelling it. The segmenters are judged according to a) their extraction rate; b) the fidelity of the segmentation with respect to the geometrical properties of the extracted targets; and c) the degree of distinctiveness imparted to the extracted targets as opposed to the nontargets. The 3 segmenters relying on BET have a better extraction rate than the other 3 that try to cope with the background simply by partitioning the image. Most segmenters here distort in one way or another the shape of the targets. The two exceptions are segmenter No. 1 (Single Intensity Threshold Silhouette Generator or SIT Generator) and No. 6 (SIT Generator in conjunction with a particular version of BET). The experimental results show that the intensity, contrast and variance features are the most effective in discriminating the targets from the nontargets. The classification results one can expect from these features together with the segmenter that proves to be the best (segmenter No. 6) amount to a detection rate in excess of 90% with a false alarm rate not greater than 3%. (U)

TABLE OF CONTENTS

1.0 INTRODUCTION 1

2.0 SEGMENTATION ALGORITHMS TO BE EVALUATED 5

3.0 IMAGERY USED FOR EVALUATION 18

4.0 FOUNDATIONS OF THE EVALUATION PROCESS 19

5.0 SEQUENTIAL FEATURE EXTRACTOR 30

6.0 EXPERIMENTAL RESULTS 36

6.1 Object Feature Histograms 38
6.1.1 SIT Generator 39
6.1.2 SIT Generator with BET (Mean) 39
6.2 Scatter Plots 40
6.3 Comments 59

7.0 CONCLUSION 72

8.0 REFERENCES 74

FIGURES 1 to 12

TABLES I and II

1.0 INTRODUCTION

A previous report (Ref. 1) describes various segmentation algorithms developed at DREV in relation to target acquisition in IR imagery. The present progress report evaluates the performance of these segmentation algorithms.

A total of 6 segmentation algorithms or segmenters are investigated. First, there is the Single Intensity Threshold Silhouette Generator. It is an early algorithm (Refs. 2 and 3) that has been successful in detecting targets in IR BOFORS imagery. In Ref. 1, we demonstrate that one can use a thresholding intensity function in lieu of a fixed and global threshold thus giving, among other possibilities, the Staircase Intensity Threshold Silhouette Generator and the Interpolated Staircase Intensity Threshold Silhouette Generator. These constitute the first 3 segmentation algorithms. The last 3 algorithms, unlike the aforementioned ones, do not deal with an image in its entirety. Instead, they try first to eliminate the background or, at the very least, to uniformize it. To this end, they incorporate the Background Elimination Technique (BET) expounded in Ref. 1, a technique which operates on a line-by-line basis and uses a narrow bandwidth low-pass filter to assess the general tendency of the background in order to subtract it from the signal corresponding to a line of the image. Because of its real-time implementation potential, we opted for a recursive filter and, more explicitly, for a 4-pole Butterworth filter (Ref. 1). Since BET can be applied either to the set of lines or columns of an image, it generates 2 images referred to as the Horizontal Fine Structure image and the Vertical Fine Structure image respectively.

One can then attempt to extract targets by segmenting one of the fine structure images, or a combination of both, with the aid of, say, the Single Intensity Threshold Silhouette Generator. The options retained are explained in Sec. 2.0.

As the evaluation process is necessarily based on some sort of imagery its scope is somewhat limited. The imagery we used is known as the Alabama Data Base and consists of 43 thermoscopic images. These contain tanks, armoured personnel carriers, jeeps and, in one instance, a bus. Although they represent ground scenes, we would term their background as moderately cluttered. On the other hand, the images are relatively clean and, for all practical purposes, can probably be considered as noise free. Hence, this imagery constitutes a good test of the segmentation algorithms although it might not be representative of real-life battlefield situations.

The effectiveness of a particular segmenter is generally characterized:

a) first, by its extraction rate, that is, its ability to segment all the targets present in the imagery;

b) by the fidelity of the segmentation process as regards the geometrical properties of the targets, this aspect is important to further discriminate the targets into classes;

c) and, finally, by what we would call the degree of distinctiveness introduced among the segmented objects and, in particular, between targets and nontargets.

The first two points are quite easy to evaluate since we know beforehand the exact number of targets as well as their respective location. The last point is a bit more tricky for the segmented objects cannot be dissimilar in every way. So the problem is really twofold: determine the most discriminatory feature or set of features and measure how well it separates the segmented objects corresponding to targets from those corresponding to nontargets. Section 4 lists and defines all the features that were extracted. We limited ourselves to those that can be extracted sequentially in the space of a single pass over the image. The sequential extractor used is based on the Labelling-by-Tracking Algorithm a short description of which is given in Sec. 5. More information about this extractor can be found in Ref. 4. The ability of a given feature to discriminate between target and nontarget segmented objects is judged according to the histograms of that feature respectively for the targets as a whole and the nontargets as a whole. If both histograms peak at the same feature value, the feature in question is useless. On the contrary, if the 2 histograms do not overlap at all, that feature alone is sufficient to isolate the targets. Hence, the amount of overlapping is a measure of the discrimination power. We have gathered together in Sec. 6.1 all the histograms that were determined for 2 segmenters out of 6, and in Sec. 6.2 some scatter plots of the most useful features. Section 6.3 discusses the pros and cons of the various segmentation algorithms to finally conclude that the best segmenter here is the Single Intensity Threshold Silhouette Generator in conjunction with the arithmetic mean of the 2 fine structure images.

This work was performed at DREV between April and November 1979 under PCN 21J11 "Automatic Target Acquisition".

2.0 SEGMENTATION ALGORITHMS TO BE EVALUATED

The segmentation algorithms investigated are:

1) The Single Intensity Threshold Silhouette Generator, hereafter designated as SIT Generator.

2) The Staircase Intensity Threshold Silhouette Generator, hereafter designated as SCIT Generator.

3) The Interpolated Staircase Intensity Threshold Silhouette Generator, hereafter designated as ISCIT Generator.

4) The SIT Generator together with the Horizontal Fine Structure image, hereafter designated as SIT Generator with BET(HFS); BET stands for Background Elimination Technique.

5) The SIT Generator together with the Maximal Fine Structure image, hereafter designated as SIT Generator with BET(Max).

6) The SIT Generator together with the Mean Fine Structure image, hereafter designated as SIT Generator with BET(Mean).

We will also sometimes refer to these segmentation algorithms as segmenter No. followed by the appropriate number.

The SIT Generator is an early algorithm that was used to detect targets in IR BOFORS imagery (Refs. 2 and 3). The defining procedure of this segmenter as applied to the imagery used for evaluation is:

a) Divide the image into 16 sub-images by quartering both axes.

b) Determine the histogram of each sub-image.

c) Find out the cutoff gray level of each partial histogram. Scanning the histogram from the highest bin down, the cutoff gray level is defined as the gray level of the first bin occupied by at least 3 pixels.

d) Discard the cutoff gray levels less than the 80th percentile of the histogram of the whole image and then choose as a global intensity threshold the smallest of the remaining cutoff gray levels.

e) In extremis, if it ever happens that all the cutoff gray levels are equal, use the 80th percentile as a threshold.

The SCIT and ISCIT generators are variants of the SIT Generator. Formally, the defining procedure of the SCIT generator is:

a) Partition the image horizontally into 4 independant sections.

b) Divide each section into 4 sub-images.

c) Determine the histogram of each sub-image within each section.

d) For each section, find out the cutoff gray level of all the histograms.

e) For each section, discard the cutoff gray levels less than the 80th percentile of the sectional histogram and choose as a global intensity threshold for that section the smallest of the remaining cutoff gray levels.

This procedure generates 4 discrete intensity thresholds or, if we plot the threshold for each line of the image against the line number, a staircase-like discontinuous thresholding intensity function. As explained in Ref. 1 and evidenced in Fig. 1, the SCIT Generator is bound to create artifacts whenever the thresholds of two adjacent sections differ widely. A manifest way to eliminate these artifacts consists in smoothing the transition between two sections by linearly interpolating the relevant thresholds. The continuous thresholding intensity function that results thereof defines the ISCIT Generator. This generator as well as the SIT and SCIT Generators are depicted in Fig. 1. Although we did not implement it, it might be worthwhile to add to the SCIT and ISCIT Generators a last-resort alternative, similar to e) above, for the case where all the cutoff gray levels of a particular section are equal.

The 3 segmenters we will now outline, unlike the previous ones, do not deal with the image in its entirety. In fact, they all include a common technique which aims to suppress all or part of the background. This technique, referred to as BET and described in detail in Ref. 1, operates on a one-dimensional signal (a given line or column of an image) and uses a narrow bandwidth low-pass filter to assess the general

FIGURE 1 - Segmenters No. 1, 2 and 3

A) Image ALA 6 3 from the Alabama Data Base
(1: raw; 2: histogram equalized;
3: sub-images delineated)
B) Thresholding Intensity Functions
C) Segmented Images
1) Single Intensity Threshold
2) Staircase Intensity Threshold
3) Interpolated Staircase Intensity Threshold

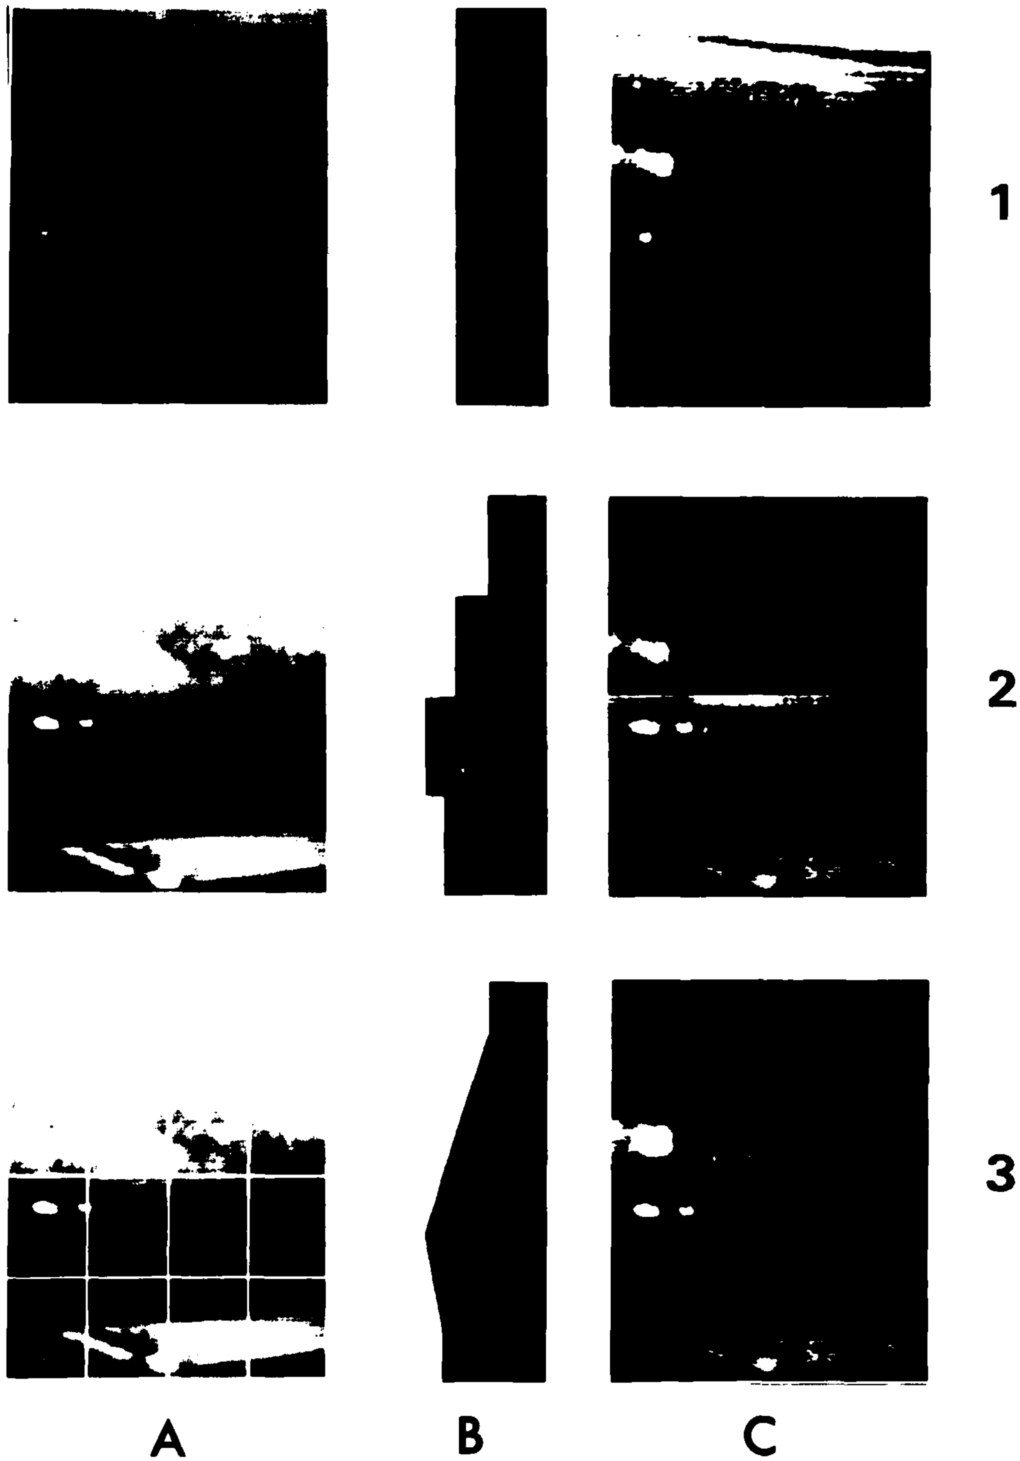
1
2
3
A
B
C

tendency of the background and then subtract it from the signal itself. Because of its real-time implementation potential, we opted for a recursive infinite impulse response filter and, to be more specific, for a 4-pole Butterworth filter (FPBF). Such a filter can be realized (Ref. 1) as a cascade of 2 second-order systems. The resulting set of linear difference equations is:

$$\begin{aligned}
f_1(nT) &= x\,[(n-2)T]\ , \\
f_2(nT) &= f_1(nT) - b_1 f_2[(n-1)T] - b_2 f_2[(n-2)T]\ , \\
f_3(nT) &= f_2(nT) - b_3 f_3[(n-1)T] - b_4 f_3[(n-2)T]\ , \\
y\,(nT) &= b_0 f_3(nT)
\end{aligned} \qquad [1]$$

with

$$\begin{aligned}
b_0 &= (1 + b_1 + b_2)(1 + b_3 + b_4)\ , \\
b_1 &= -\,(Z_1 + Z_1^*)\ , \quad b_2 = Z_1\, Z_1^* \\
b_3 &= -\,(Z_2 + Z_2^*)\ , \quad b_4 = Z_2\, Z_2^*
\end{aligned} \qquad [2]$$

where

$$\begin{aligned}
Z_1 &= \exp[-2\pi f_c(\cos 67.5^\circ - j \sin 67.5^\circ)/f_s] \\
Z_2 &= \exp[-2\pi f_c(\cos 22.5^\circ - j \sin 22.5^\circ)/f_s]
\end{aligned} \qquad [3]$$

In these equations, x designates the input signal, y the filtered output signal, T the sampling interval, f_c the 3-dB cutoff frequency of the filter and f_s the sampling frequency of the signal. The asterisk in [2] denotes the complex conjugate. frequency of the signal. The asterisk in [2] denotes the complex conjugate.

To illustrate BET we will use the signal of Fig. 2, which corresponds to line 175 of image 6 from the Alabama Data Base, and assume that the 3-dB normalized cutoff frequency (f_c/f_s) of the low-pass FPBF digital filter is equal to 0.01 (to process the evaluation imagery we used a cutoff frequency of 0.05; see Ref. 1). The filtered signal generated by such a filter is shown in Fig. 2a along with the input signal. Two points are worth mentioning about the filtered signal:

a) There is a droop in the filtered signal at its origin.

b) The filtered signal is shifted to the right.

The first anomaly can be easily corrected by selecting the initial conditions so that there is no transient at the origin. It can be shown (Ref. 1) that the required initial conditions are:

$$f_3(nT) = H / b_0$$

and

$$f_2(nT) = H / (1 + b_1 + b_2) \qquad [4]$$

FIGURE 2 - Background Elimination Technique (BET) The illustrative signal is line 175 of image 6 from the Alabama Data Base. The 3 peaks correspond respectively (from left to right) to a tank, an APC and a jeep. The cutoff frequency of the filter is 0.01.

A) FPBF filter initially at rest
B) FPBF filter with nonzero initial conditions; the solid line is the left filtered signal while the dashed line is the right filtered signal.
C) Arithmetic mean of the 2 filtered signals
D) Fine structure or fluctuating component of the input signal

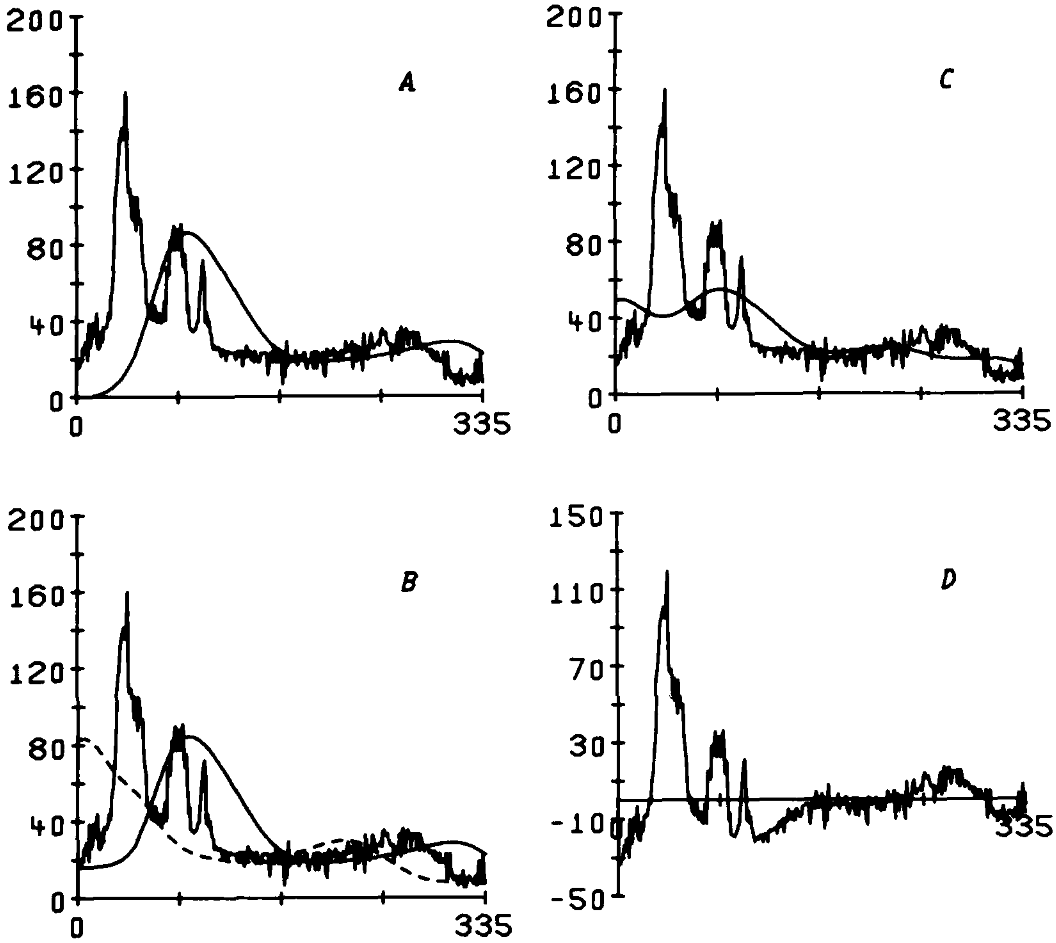
A
B
C
D
200
160
120
80
40
0
0
335
150
110
70
30
-10
-50

for $n<0$; H is the value of the input signal at $t=0^+$. Fig. 2b shows the filtered signal (solid line) that results when we use these new initial conditions. The second anomaly can be as easily corrected by shifting the filtered signal to the left. However, rather than rectifying this anomaly we will take advantage of it to clip the peaks. Let us consider Fig. 2b. The signal is fed to the filter from left to right. Normally, we would expect the filtered signal to peak at, or close to, the position of the main spike in the input signal. Instead, it overshoots to the right. Therefore, had the signal been fed from right to left, the overshoot would have occurred to the left (dashed line in Fig. 2b). By combining both filtered signals in some fashion, we can expect to end up with a curve that will bypass entirely the peaks to follow only the broad characteristics of the input signal. Various combinations were tried (Ref. 1). All things considered, the arithmetic mean (Fig. 2c) was judged most satisfactory. Fig. 2d exhibits the fine structure (fluctuating component) of the illustrative signal, that is, what is left of the signal once the estimated trend of the background is removed.

The Background Elimination Technique can be applied either to the set of lines or columns of an image thus producing 2 distinct images (Fig. 3) referred to as the Horizontal Fine Structure (HFS) image and the Vertical Fine Structure (VFS) image respectively. Although these images turn out to be highly textured, they do not exhibit, unlike the parent image, large-scale fluctuations. This is important for large-scale fluctuations may easily fool a segmenter like the SIT Generator based on the single assumption that the targets present a larger thermal signature than the background. The SCIT and ISCIT generators do try to circumvent the problem by slicing the image into sections whose background may be considered as "uniform", but then the crux centers on the manner in which the sections are defined. With the 2 fine structure images, this crucial question does not arise because their background, on a large-scale basis, is inherently uniform. In consequence, the SIT Generator should be well suited for thresholding the fine structure images. The veracity of this affirmation is confirmed by the results of Figs. 3d and 3e. The procedure leading to Fig. 3d defines segmenter No. 4: SIT Generator with BET (HFS).

The segmentation of both the HFS image (Fig. 3d) and the VFS image (Fig. 3e) results in targets whose shape is slightly distorted. However, since the distortion is more outstanding in one direction than in the other, and since the dimension affected is different whether HFS or VFS is involved, it should be possible to maintain intact the shape of the targets by thresholding a joined image resulting from some combination of HFS and VFS. The following sensible combinations were formed:

a) Maximal Fine Structure image, where the value at any given location corresponds to the maximum of HFS and VFS for that location.

b) Mean Fine Structure image, where the value at any given location corresponds to the arithmetic mean of HFS and VFS for that location.

FIGURE 3 - Thresholding of the fine structure images derived from image 6 3 of the Alabama Data Base:

a) Original histogram - equalized image
b) Horizontal Fine Structure (HFS) image
c) Vertical Fine Structure (VFS) image
d) Segmented image generated by thresholding HFS with the SIT Generator; this defines segmenter No. 4: SIT Generator with BET (HFS).
e) Segmented image generated by thresholding VFS with the SIT Generator

The images b and c were postprocessed, for display purpose, first by adding a constant bias, so as to remove negative gray levels, and then by stretching the gray levels bounded by the 5th and 95th percentiles linearly over the display range.

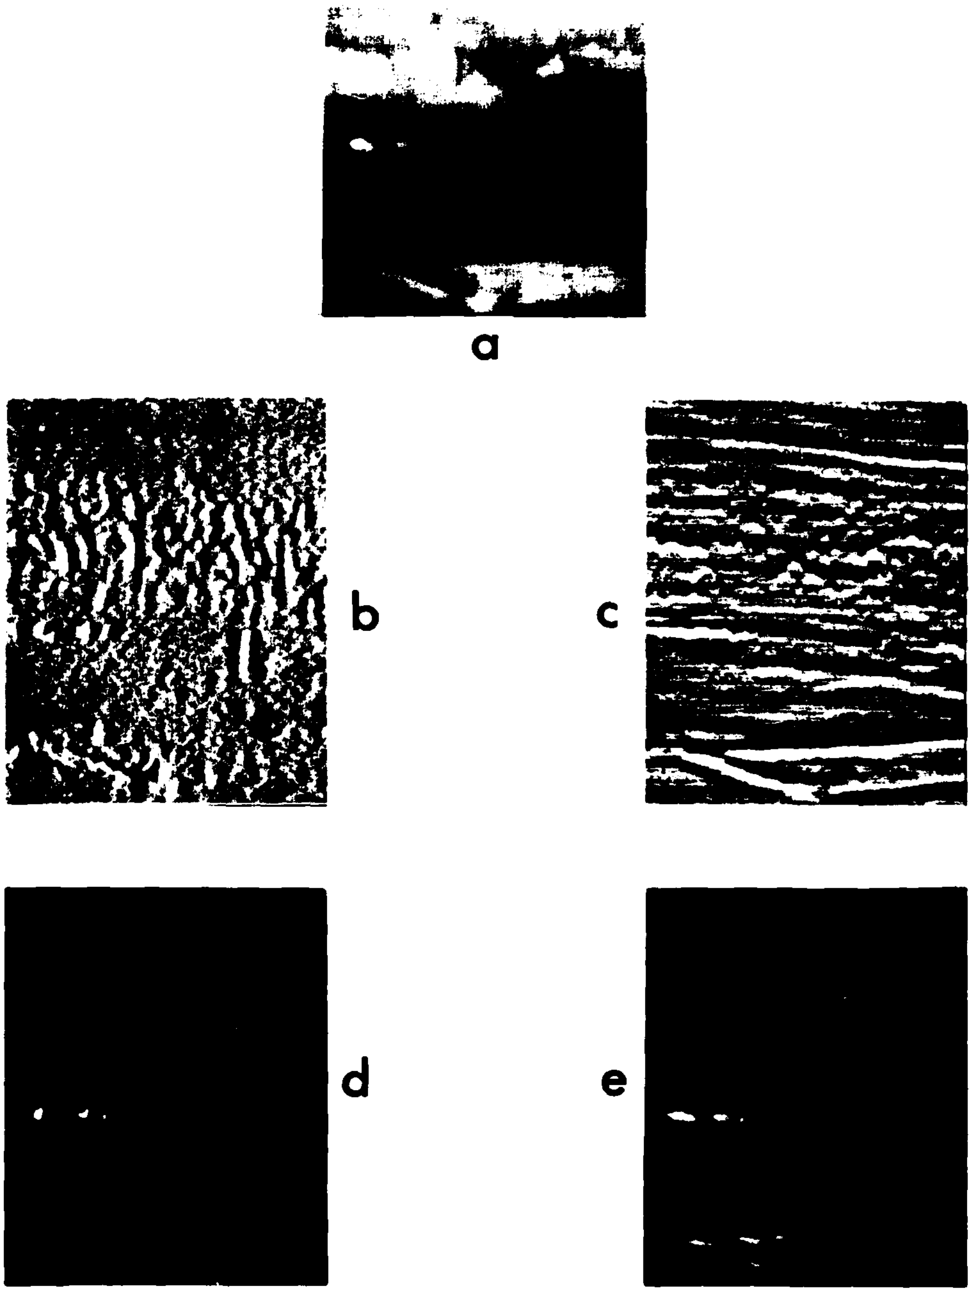
a
b
c
d
e

The first combination defines segmenter No. 5 (SIT Generator with BET(Max)) and the second one segmenter No. 6 (SIT Generator with BET(Mean)).

3.0 IMAGERY USED FOR EVALUATION

To reliably evaluate the performance of a particular segmenter, we need some sort of imagery to start with. However, this very fact somewhat limits the scope of the evaluation to a certain type of background, noise, image quality etc. The imagery used here for evaluation is known as the Alabama Data Base and consists of 43 thermoscopic images. The spectral region of the majority of them (30 out of 43) corresponds to the 8-14μm band, and that of the remaining ones to the 3-5μm band. Altogether the images contain 85 targets, some of them so close to each other as to form a distinct entity, distributed as follows: 40 tanks, 29 armoured personnel carriers (APC), 15 jeeps and, finally, a bus. The number of targets in a single image never exceeds 3 and no image contains 2 targets of the same type. Although the images represent ground scenes, we would term their background as moderately cluttered. On the other hand, the images are relatively clean and, for all practical purposes, can probably be considered as noise free. The size of the images is 420 x 335 pixels and they are digitized according to a 256-level grayscale. The images were in no way preprocessed prior to segmentation but, for display purpose (e.g. Figs. 1 and 3), they were postprocessed by histogram equalization, which almost consistently yields "good-looking" images.

4.0 FOUNDATIONS OF THE EVALUATION PROCESS

An overview of the scientific literature devoted to automatic target acquisition would reveal that the effectiveness of a particular segmenter is generally characterized

a) first, by its extraction rate, that is, its ability to segment all the targets present in the imagery;

b) by the fidelity of the segmentation as regards the geometrical properties of the targets; this aspect is important to further discriminate the targets into classes;

c) and, finally, by what we would call the degree of distinctiveness introduced among the segmented objects and, in particular, between targets and nontargets.

The first two points are quite easy to evaluate since we know beforehand the exact number of targets as well as their respective location. The last point is a bit more tricky for the segmented objects cannot be dissimilar in every way. So the problem is really twofold: determine the most discriminatory feature or set of features and measure how well it separates the segmented objects corresponding to targets from those corresponding to nontargets. In this section, we list and define all the candidate features. We limited ourselves to those that can be extracted sequentially in the space of a single pass over the image.

We will first define a quantity that appears in many expressions below. Given an arbitrary segmented object S, the (p,q) th moment $m_{p,q}$ of S is defined as

$$m_{p,q} = \frac{1}{N} \sum\sum x^p y^q \; n(x,y) \quad p,q = 0,1,2 \;\ldots$$

where N denotes the number of points in summation, x and y the spatial coordinates, and where the value of n at any point (x,y) is proportional to the brightness (or gray level) of S at that point. We can now proceed with the definition of the features involved in the evaluation process.

1) Area (A) - The area of S is just the number of points in S:

$$A = N \qquad [6]$$

2) Perimeter (P) - Rosenfeld and Kak (Ref. 5) give 4 possible definitions of the perimeter:

a) The number of pairs of points (u,v) with u in S and v not in S.

b) The number of steps taken by a border-following algorithm in following all the borders of S.

c) The same, but with diagonal steps counting $\sqrt{2}$ each, while horizontal and vertical steps count only 1 each.

d) The number of border points of S.

We have adopted the first definition for it can be easily computed sequentially and also because it yields the right answer for a square. It should be emphasized here that not only the perimeter of the outside border but that of all the borders of the segmented object are computed.

3) Thinness Ratio(T) - The thinness ratio of a segmented object of area A and perimeter P is defined (Ref. 6) by

$$T = 4\pi\left(\frac{A}{P^2}\right) . \qquad [7]$$

It can be shown that T has a maximum value of 1, which it achieves if the segmented object in question is circular. Loosely speaking, the fatter a segmented object is the greater will be the associated thinness ratio; conversely, line-like or largely perforated objects will have a thinness ratio close to zero. Moreover, the thinness ratio is dimensionless and hence depends only on the shape (but not the scale) of the segmented object.

4) Average Intensity(B) - The average intensity or brightness is a function of the average temperature of the underlying object and is defined as

$$B = m_{0,0} \qquad [8]$$

5) Average Contrast (C) - Contrast has many definitions. The one used here is

$$C = (B_t - B_b)/B_b \qquad [9]$$

where B_t refers to the brightness of the segmented object and B_b to that of the background immediately surrounding the object. We already know how to compute B_t. We can compute B_b in the same way but, then, we have to specify what the background $\bar{S}$ of S is. Obviously, $\bar{S}$ should not include S or another segmented object. Moreover, it should not stretch out too far away from S in order that C be a local measure of contrast. These requirements are easily met by extending (Fig. 4) in both directions independently all the runs of segmented pixels present on any given scan line. The extension of one end of a run proceeds until another run of segmented pixels is hit, or until the run has been extended by half its length. In this way, the area of S is about the same as that of S. Although this scheme may distort the measurement of C by introducing a certain degree of directionality, this is not a major drawback for the segmented objects more often resemble objects 1 and 2 in Fig. 4 than objects 3, 4 or 5 where the effect of directionality is most damaging. On the other hand, an important asset of this scheme is its ease of implementation.

6) Relative Intensity (B') - The relative intensity is obtained by mapping the average intensity of a segmented object within a given frame into a scale from 0 to 1, that is, a scale independent of the

```
     0110
   00111100
  00111100                                        0330
 00111100                                         0330
00111100                                          0330
  0110




                    02220
                  002222200
                0002222222000
                 02220 02220
                 02220 02220
                 02220 02220
                0002222222000
                  002222200
                    02220
                     020


                                            0044440550
                                            0044440550
```

DIRECTION OF SCAN

FIGURE 4 - The pixels used to compute the brightness of B_b of the background of the 5 numbered objects are marked with 0's. The pixels marked with 0's enter into the computation of 2 different B_b's.

overall characteristics of the imagery. We used to define the relative intensity as

$$B'_i = (B_i - B_{min})/(B_{max} - B_{min}) \quad i = 1,2 \ldots M \qquad [10]$$

where M is the number of segmented objects within the frame in question and B_{min} and B_{max} are respectively the largest and the smallest of all the B_i's. However, this expression is not flawless. For example, if M = 2, the relative intensity of one segmented object will be 1 and that of the other 0 regardless of their respective brightness. From the standpoint of discrimination, this is not desirable for it might well occur that both segmented objects are targets. In fact, the same situation is bound to happen each time the number of segmented objects is equal to or less than the number of expected targets. So, when the number of segmented objects is small, in our case small means less than 6 (this number is greater than the maximum number of targets one can find in any image of the Alabama Data Base in order to account for multiply segmented targets), we use the following expression instead of [9]:

$$B'_i = B_i/B_{max} \quad i = 1,2,\ldots M<6. \qquad [11]$$

One may wonder why [11] alone is not used. It is simply because experimentation shows that in most cases [10] better discriminates the targets from the nontargets.

7) Relative Contrast(C') - The relative contrast is defined as the relative intensity.

8) Centroid(x,y) - The centroid of a blob (segmented object) is the point (x,y) whose coordinates are given by

$$\bar{x} = m_{1,0} / m_{0,0} ,$$

and

$$\bar{y} = m_{0,1} / m_{0,0} . \qquad [12]$$

9) Principal Axis(Θ) - The principal axis Θ is the angle for which the moment of inertia of a blob about a line through its centroid is as small as possible. It can be shown that

$$\Theta = \frac{\arctan\,(2M_{1,1} / (M_{2,0} - M_{0,2})}{2} , \qquad [13]$$

where $M_{p,q} = m_{p,q} / m_{0,0} - \bar{x}^{p}\,\bar{y}^{q}$.

The $M_{p,q}$'s are nothing but central moments, that is, the above moments evaluated around (x,y) as the origin.

10) Overall Width(ℓ) - The overall width of a blob is defined as

$$\ell = x_R - x_L \qquad [14]$$

where x_R is the abscissa of the lower right corner of the smallest rectangle circumscribed around the blob and x_L that of the upper left corner of the same rectangle.

11) Overall Height(h) - The overall height is simply the difference between the ordinates corresponding to x_R and x_L:

$$h = y_R - y_L . \qquad [15]$$

12) ℓ/h Ratio

13) Bulkiness(e) - The bulkiness is the proportion of the circumscribed rectangle occupied by the blob:

$$e = A/(\ell x h) . \qquad [16]$$

14) Major and Minor Diameters(d_1,d_2) - The eigenvalues of the matrix

$$\begin{pmatrix} M_{2,0} & M_{1,1} \\ M_{1,1} & M_{0,2} \end{pmatrix}$$

of second central moments are

$$r_1^2 = M_{0,2} + M_{1,1} / \tan\Theta$$
$$[17]$$
$$r_2^2 = M_{2,0} + M_{1,1} / \tan\Theta .$$

These eigenvalues are the principal moments of inertia of the blob and it can be shown that the larger eigenvalue corresponds to the principal axis. The major and minor diameters are then defined as

$$d_1 = 1 + 2r_1 \text{ and } d_2 = 1 + 2r_2 . \qquad [18]$$

15) Aspect Ratio(a) - The aspect ratio is given by

$$a = d_1 / d_2 \qquad [19]$$

where d_1 denotes the major diameter. This quantity is then always greater than (or possibly equal to) 1.

The last 8 features aim at characterizing the statistics of S (segmented object) and $\bar{S}$ (background of S). A complete specification of the statistics of say S is possible if one knows its moments m_k (there should be no confusion between $m_{p,q}$ and m_k since the latter has only one index) defined by

$$m_k = E\,\{n(x,y)^k\} = \frac{1}{N} \sum_S n(x,y)^k \ . \qquad [20]$$

Clearly,

$$m_1 = B$$

the brightness or the average intensity of S. The corresponding central moments are given by

$$M_k = E\,\{(n(x,y) - B)^k\} = \frac{1}{N} \sum_S (n(x,y)-B)^k \ . \qquad [21]$$

These can be expressed in terms of m_k:

$$M_k = \sum_{r=0}^{k} \frac{k!}{r!(k-r)!} (-1)^r B^r m_{k-r} \ . \qquad [22]$$

In particular

$$m_2 = m_2 - B^2 , \qquad [23]$$

$$M_3 = m_3 - 3\ Bm_2 + 2\ B^3 , \qquad [24]$$

$$M_4 = m_4 - 4\ Bm_3 + 6\ B^2\ m_2 - 3\ B^4 . \qquad [25]$$

Given these expressions we readily obtain:

$$\sigma^2 = M_2 . \qquad [26]$$

16) Blob Variance (σ^2)

17) Blob Relative Variance - This quantity is defined as the relative intensity.

18) Blob Skewness

$$M_3 / \sigma^3 \qquad [27]$$

19) Blob Kurtosis

$$M_4 / \sigma^4 - 3 \qquad [28]$$

and their counterpart as regards $\overline{S}$.

20) <u>Background Variance</u>

21) <u>Background Relative Variance</u>

22) <u>Background Skewness</u>

23) <u>Background Kurtosis</u>

5.0 SEQUENTIAL FEATURE EXTRACTOR

In this section we briefly describe (more information can be found in Ref. 4) the sequential algorithm used to extract from a segmented image the various features listed in the preceding section. Such an algorithm is called a feature extractor and it can be regarded as an emulation of a hypothetical hardware unit of the same name. This dual meaning is reflected in the terminology used here for its description. The basic procedure of the extractor in question is sometimes referred to in the scientific literature as the Labelling-by-Tracking (LT) Algorithm (Ref. 5). References 4 and 8 describe a more complex extractor, the Boundary Continuation Algorithm, that can also fulfill the same task. Given a thresholding intensity function of the kind defined in Sec. 2, both extractors can segment the image, identify the objects generated, and extract the relevant features in a single image scan.

The memory of the LT-extractor consists of a scan line array plus a feature array. The first array contains the current scan line of the image being processed as well as the immediately preceding the scan line. It is initially set to zero and afterwards updated by replacing the precedent scan line by the current one and reading in the subsequent one. The scan line array is equivalent to viewing the image through a downward moving slit whose width matches that of the image, but has only 2 pixels in height. The feature array has an arbitrary number of lines whereas the number of columns is a function of the number of features to be extracted. The number of columns is not exactly equal to the number of features because some of these are nothing but a combination of other features (e.g. the aspect ratio), while, on the contrary, it is necessary to accumulate more than one quantity to determine other features (e.g. the principal axis). On the other hand, there are no fixed rules as regards the number of lines. We can say, for sure, that it should at least be equal to the maximum number of expected targets and nontargets in a single frame but, in practice, it should be chosen larger than that for at some point in the extraction process an object may temporarily be split up into several components. These components will most probably be ultimately merged but in the meantime the feature array should offer sufficient room for them. If the number of lines of the feature array is not large enough, the extraction process will not necessarily be impeded but it might be slowed down appreciably because of frequent updates (Ref. 4). One column of the feature array is set apart for a substitution table that keeps track of all the components of the various objects. This table is the key for updating the feature array (Ref. 4).

We detail hereafter the procedure used by the LT-extractor to identify the objects generated by the segmentation process. The identification is done by labelling the various objects, that is, assigning a specific number to each of them. It should be obvious from this procedure that the pixels belonging to an object are assumed to be 8-connected (Refs. 4 and 5). Let $n_i(j)$ be the label assigned to the jth pixel (Fig. 5) of scan line i (even though both use the same letter it is unlikely to confuse the gray level n(x,y) with the label $n_i(j)$) and n_o the last number utilized to label a new object. Furthermore, let

	J-1	J	J+1	J+2
SCAN LINE I-1:	X	X	X	X
SCAN LINE I :	X	*	X	X

FIGURE 5 - a) Fraction of the scan line array with the current pixel (i,j) marked with an asterik

us assume that all the pixels belonging to the background are set to zero. Then,

1) if $$n_{i-1}(j-1) \times n_{i-1}(j+1) \neq 0$$

and

$$n_{i-1}(j-1) \neq n_{i-1}(j+1)$$

we conclude that the relevant pixels are connected through the diagonal, and consequently that they belong to the same object. We determine which label has the greatest value, say $n_{i-1}(j-1)$, and replace in the substitution table all the $n_{i-1}(j-1)$'s by $n_{i-1}(j+1)$

2) Next, we put

$n_i(j) = n_{i-1}(j-1)$ if $n_{i-1}(j-1) \neq 0$; otherwise
$n_i(j) = n_{i-1}(j)$ if $n_{i-1}(j) \neq 0$; otherwise
$n_i(j) = n_{i-1}(j+1)$ if $n_{i-1}(j+1) \neq 0$; otherwise
$n_i(j) = n_i(j-1)$ if $n_i(j-1) \neq 0$; otherwise

we conclude that the pixel (i,j) is not the continuation of an existing object but the beginning of a new one. We can then set

3) $$n_i(j) = n_o + 1$$

but in this way all the objects, whatever their size, will be labelled, that is, even 1-pixel and 2-pixel objects. However, these objects are obviously (here) nontargets and labelling them unnecessarily overloads the feature array and, in some instances, can possibly saturate it. It is then preferable to eliminate them at once. To this end, it suffices to replace the preceding step 3) by

3')
$n_i(j) = 0$ if $n_i(j+1) = 0$; otherwise
$n_i(j) = n_i(j+1) = n_{i-1}(j+2)$ if $n_{i-1}(j+2) \neq 0$; otherwise
$n_i(j) = n_i(j+1) = 0$ if $n_i(j+2) = 0$; otherwise

we conclude that the pixels (i,j) and (i,j+1) belong to a new object and we set

$$n_i(j) = n_i(j+1) = n_o+1.$$

It is worth noting that this step also eliminates slanted lines (lines parallel to the scan direction remain but they are eliminated later when the feature array is updated) whose width is less than 3 pixels as well as line-like object protuberences jutting out counter to the scan direction. This might be a source of distortion of the object shape, but probably not a serious one considering that the boundary of the targets is in general relatively smooth (this is true mainlv because the targets are small and line-like features such as a tank's gun are unresolved).

4) If $n_i(j-1) \neq 0$
and $n_i(j-1) \neq n_i(j)$

we are faced with 2 object segments connected by one end. This piece of information is entered into the substitution table as described in 1).

Given that the number of lines of the feature array is fixed, it might well happen that the extraction process will have to be halted despite repeated updating operations because the feature array is full. To reduce the number of updates and to prevent the feature array from saturating, the above procedure can be modified so as to utilize the smallest number of labels to identify the objects. Let $J = j-1$ and L_i be the number of pixels labelled $n_i(j)$. Then 4) is replaced by:

4') If $n_i(J) \neq 0$ and $n_i(J) \neq n_i(j)$ set

$$L(j) = L(j) + 1,$$
$$L(J) = L(J) - 1,$$
$$n_i(J) = n_i(j),$$

and we repeat for $J = J-1$ if $n_i(J-1) \neq 0$; otherwise we put

$$n_o = n_o - 1 \text{ if } L(J) = 0; \text{ otherwise}$$

we conclude that there exists on scan line i-1 an object segment labelled $n_i(J)$ that belongs to the same object as the segment $n_i(J)$. We then modify the substitution table accordingly. This simple step frequently allows us to save a label (Ref. 4).

5) Repeat from 1) with the next pixel different from 0.

6.0 EXPERIMENTAL RESULTS

We mentioned in Sec. 4 that one valuable attribute of any segmenter is the degree of distinctiveness it introduces among the segmented objects and, in particular, between targets and nontargets. We also pointed out that we cannot expect the objects to be dissimilar in every way and, consequently, that the problem is really twofold:

a) Determine the most discriminatory feature or set of features.

b) Measure how well this feature or set of features separates the objects corresponding to targets from those corresponding to nontargets.

We have defined in Sec. 4 the complete set of features we intend to consider for this purpose. The means to be used to assess the discriminatory power of a given feature will consist in a comparison of the histograms of that feature both for targets and nontargets, and this for each one of the 6 segmenters described in Sec. 2. The total system of operations, illustrated in Fig. 6, is:

1) The 43 raw images of the Alabama Data Base are processed in turn by all 6 segmenters.

2) The segmented images are passed on to the *LT-extractor* and the object features extracted according to the procedure outlined in Sec. 5.

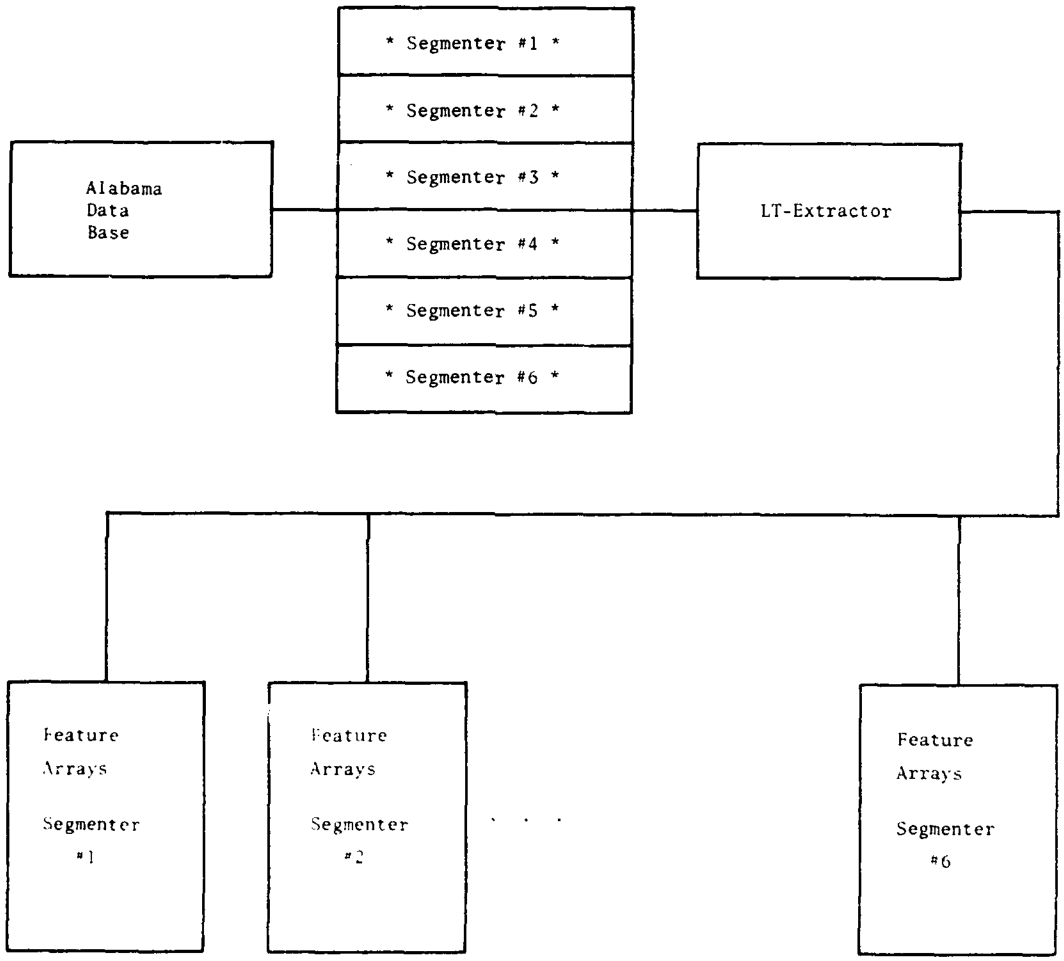

FIGURE 6 - System of operations leading to the determination of the target and nontarget feature histograms

3) The resulting feature arrays (each segmenter gives rise to 43 feature arrays), duly updated, are stored in APL files (the number of lines of whatever feature array is equal to the number of objects in the corresponding segmented image and the line numbers are the labels associated to those objects as a result of the last update).

4) A target identification array is assigned to each APL file. This 43-line array contains the labels of the objects corresponding to targets. It is defined according to target location data collected beforehand.

5) An APL program automatically determines and plots the various target and nontarget feature histograms.

Being stored in APL files, the feature arrays can be analyzed interactively, and thus the above systems of operations offers a great flexibility.

6.1 Object Feature Histograms

All the histograms presented in this report have the same number of bins, namely, 20. The target and nontarget histograms of a particular feature are plotted side by side and both the horizontal and the vertical scales are the same for ease of comparison. The range (horizontal scale) of a histogram is generally that of the feature itself with the exception of the area (limited to 400) and the perimeter (limited to 200). The number of elements in a bin (vertical axis) is expressed as a percentage of the total number of targets or nontargets as the case may be. In all instances where the numerical value of a feature (e.g. thinness ratio), as defined in Sec. 4, is always less than 1, this value is multiplied by 100 in order to get rid of fractional numbers.

6.1.1 SIT Generator

The SIT Generator manages to extract (segment from the background) 73 targets out of 85. Most missed targets are APC's. The number of nontargets generated by this segmenter, on the other hand, is rather high, that is 1011. It was included in the present study mostly for historical reasons (Refs. 2 and 3) but also as a standard by which the results of more sophisticated segmenters are evaluated. The histograms arising from this segmenter made up Fig. 7. There are 13 histogram pairs corresponding to as many features. Certain features listed in Sec. 4 were excluded whether because they turn out to be useless (e.g. skewness) or because they are not distinctive characteristics in themselves (e.g. centroid). These results (as well as those of the next 2 sections) will be commented further in a subsequent section.

6.1.2 SIT Generator with BET (Mean)

This segmenter is better suited for the task at hand since it extracts 83 targets out of 85 while producing only half as many nontargets (584) as the precedent segmenter. The 13 histogram pairs that sum up the results obtained with this segmenter are given in Fig. 8, which is the counterpart of Fig. 7.

6.2 Scatter Plots

Figures 9 and 10 show scatter plots of the following features:

a) Relative intensity

b) Relative contrast

c) Relative blob variance

that happen to be the most useful as far as discrimination between targets and nontargets is concerned. These scatter plots might be somewhat misleading, however, for a plotted point often corresponds to more than one datum. In other words, 2 pairs of features from 2 segmented objects may well match each other and, consequently, the relevant objects may be represented by a single point in the scatter plots (for example, there are 1011 nontargets associated with segmenter No. 1 but only 232 plotted points in Fig. 10a). So one should not attempt to draw conclusions based on the density of the points plotted in these figures. It is also worth mentioning that it is not the variance (σ^2) which is actually plotted in Figs. 9 and 10 but σ, the positive square root (standard deviation) of the variance. This is equally true of Figs. 7 and 8.

FIGURE 7 - The 13 histogram pairs derived from segmenter No. 1. The associated features are:

a) Area
b) Perimeter
c) Thinness Ratio
d) Relative Intensity
e) Relative Contrast
f) Overall Width
g) Overall Height
h) Width/Height Ratio
i) Bulkiness
j) Minor Diameter
k) Major Diameter
l) Aspect Ratio
m) Blob Relative Variance

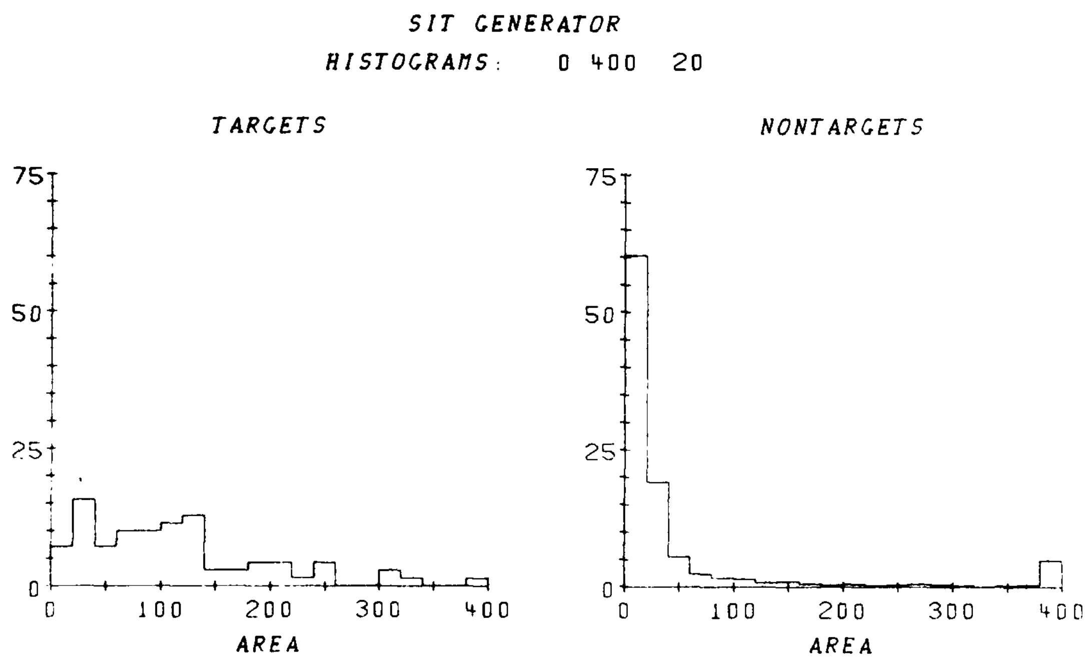

SIT GENERATOR

HISTOGRAMS: 0 200 20

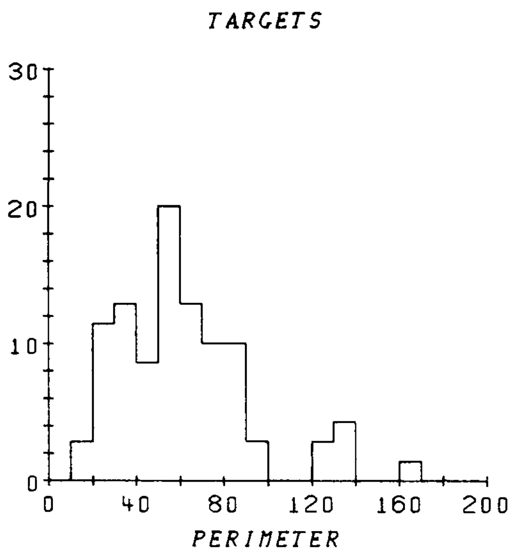

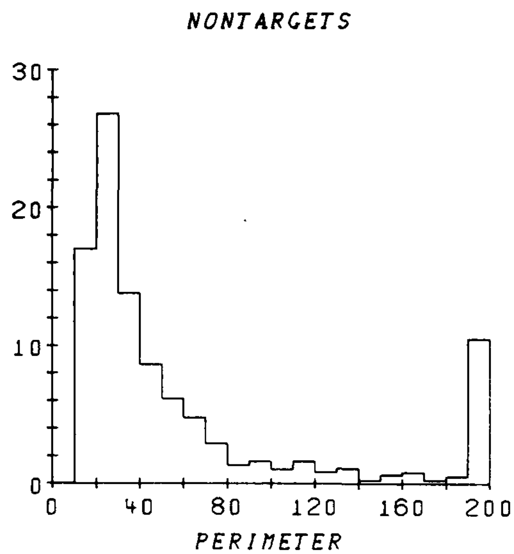

SIT GENERATOR

HISTOGRAMS: 0 100 20

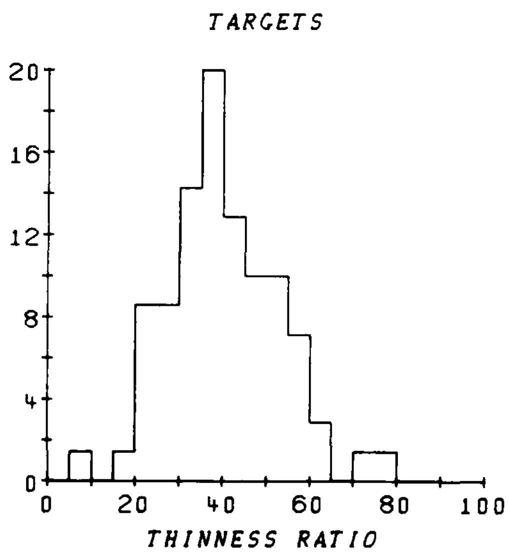

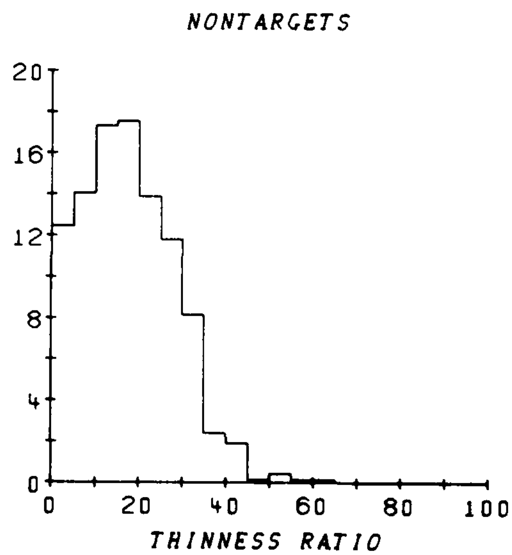

SIT GENERATOR

HISTOGRAMS: 0 100 20

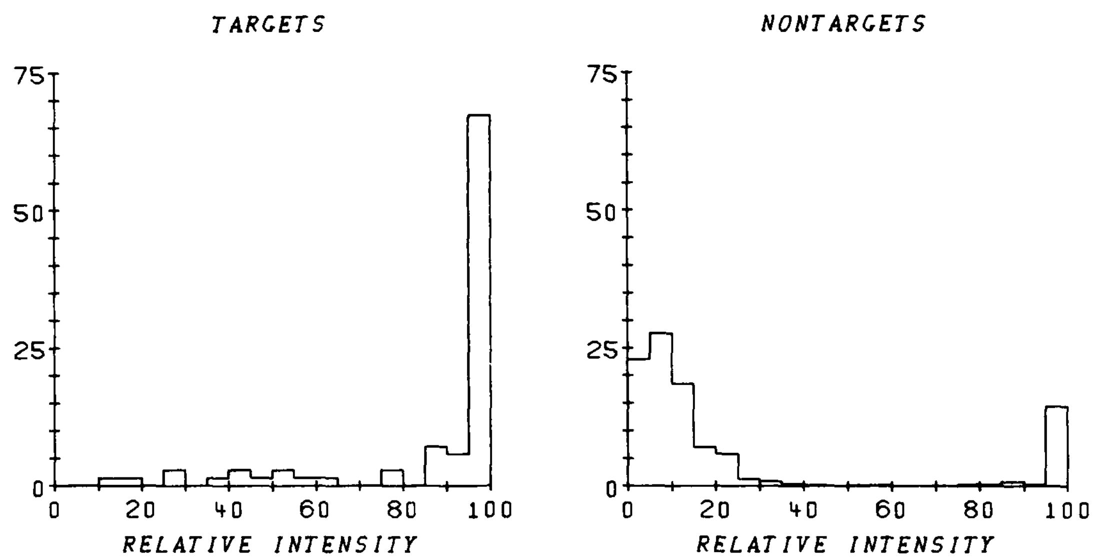

SIT GENERATOR

HISTOGRAMS: 0 100 20

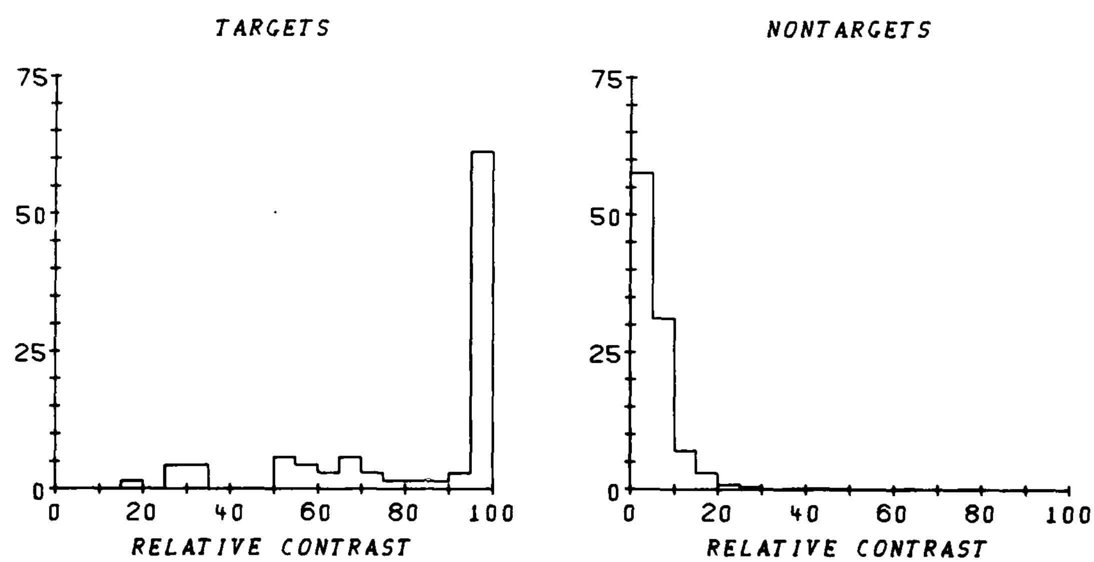

SIT GENERATOR
HISTOGRAMS: 0 100 20

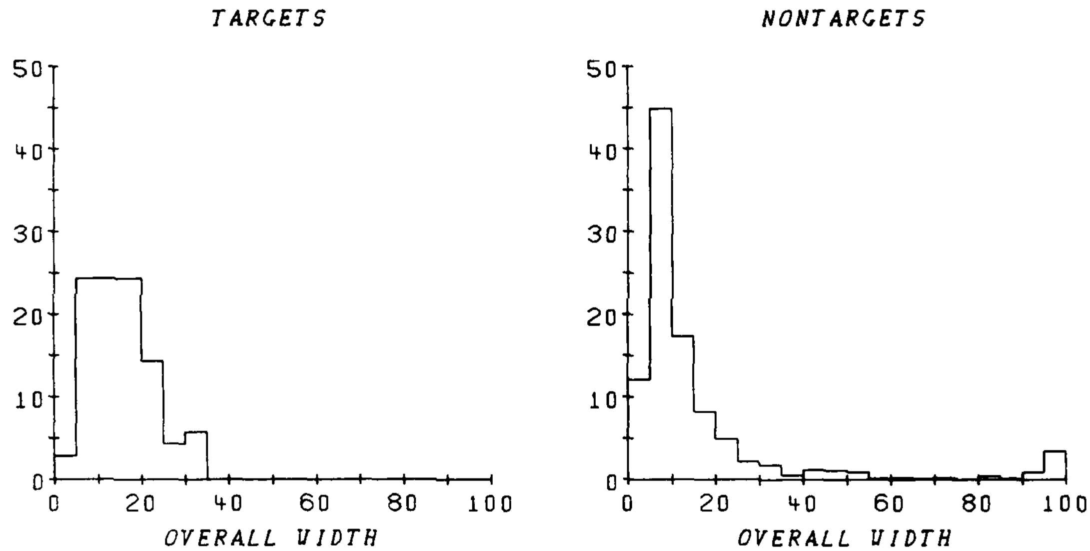

SIT GENERATOR
HISTOGRAMS: 0 100 20

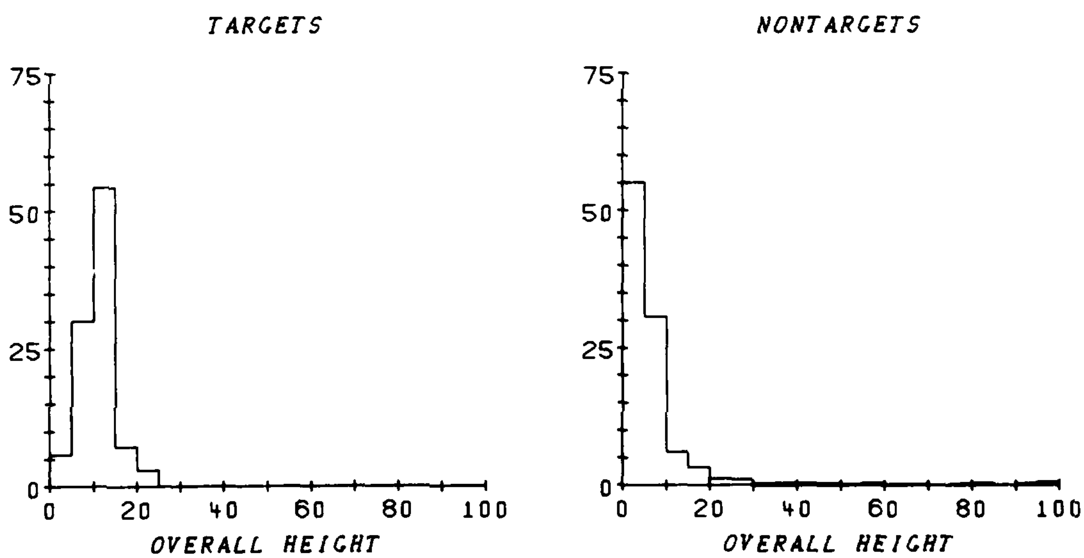

SIT GENERATOR

HISTOGRAMS: 0 10 20

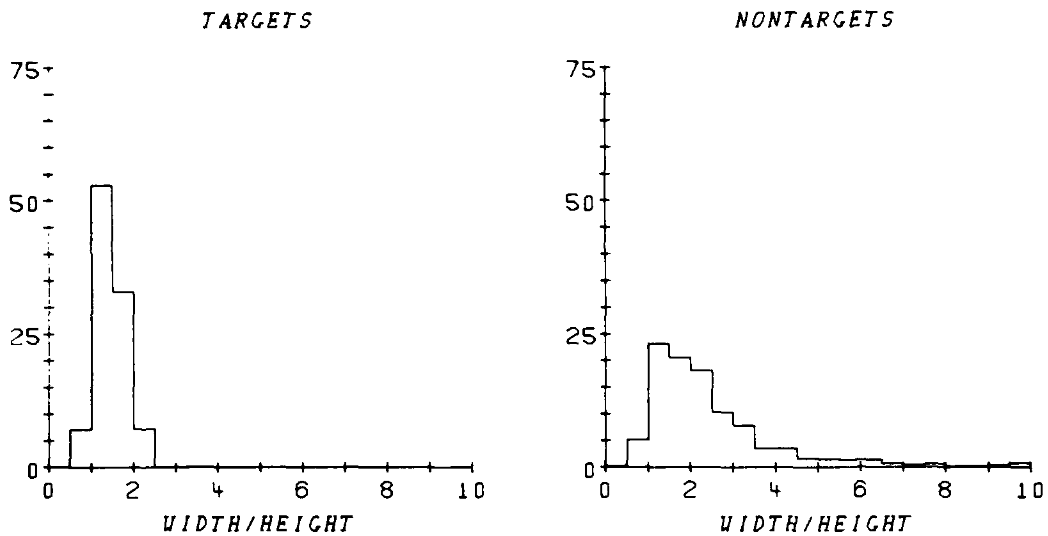

SIT GENERATOR

HISTOGRAMS: 0 100 20

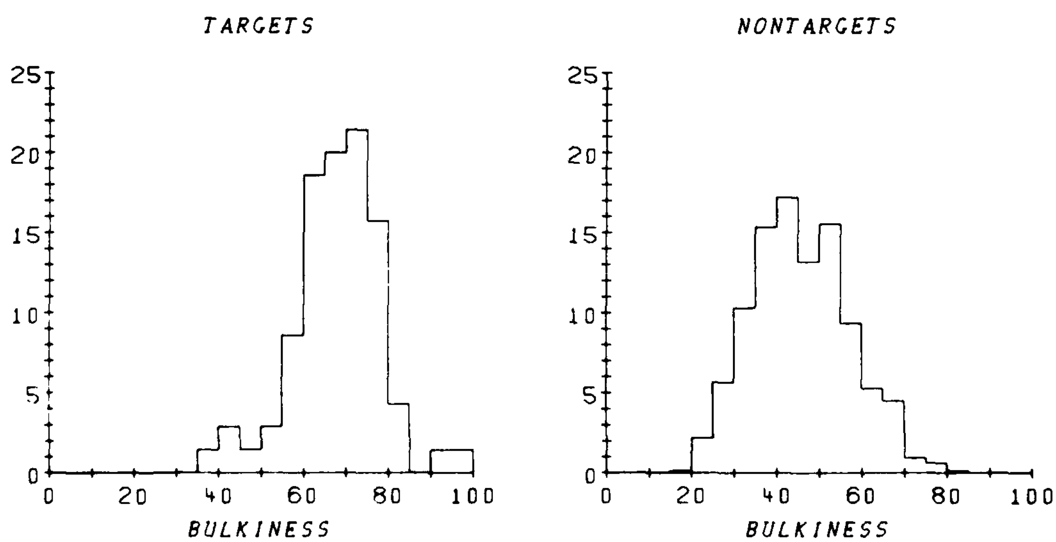

SIT GENERATOR

HISTOGRAMS: 0 20 20

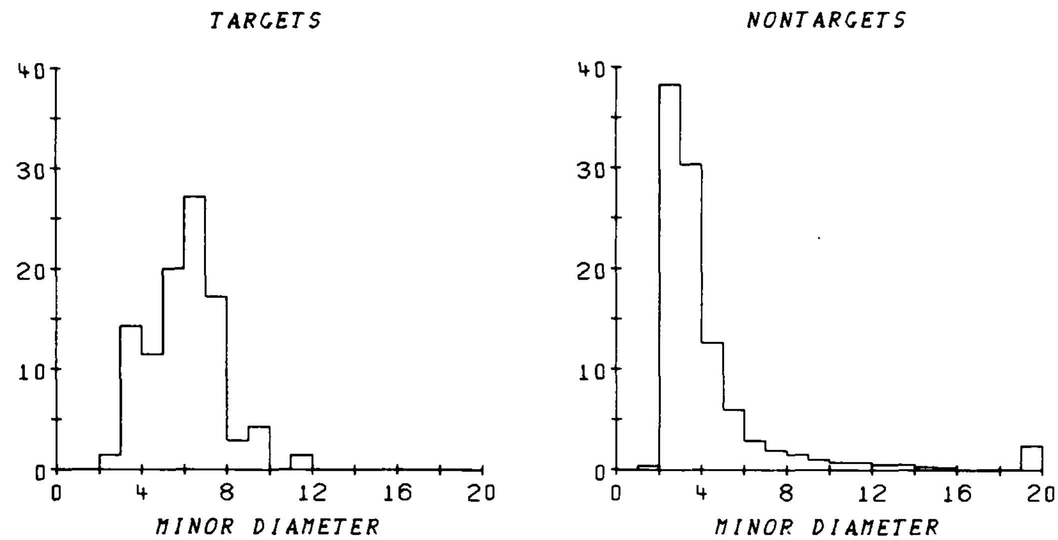

SIT GENERATOR

HISTOGRAMS: 0 20 20

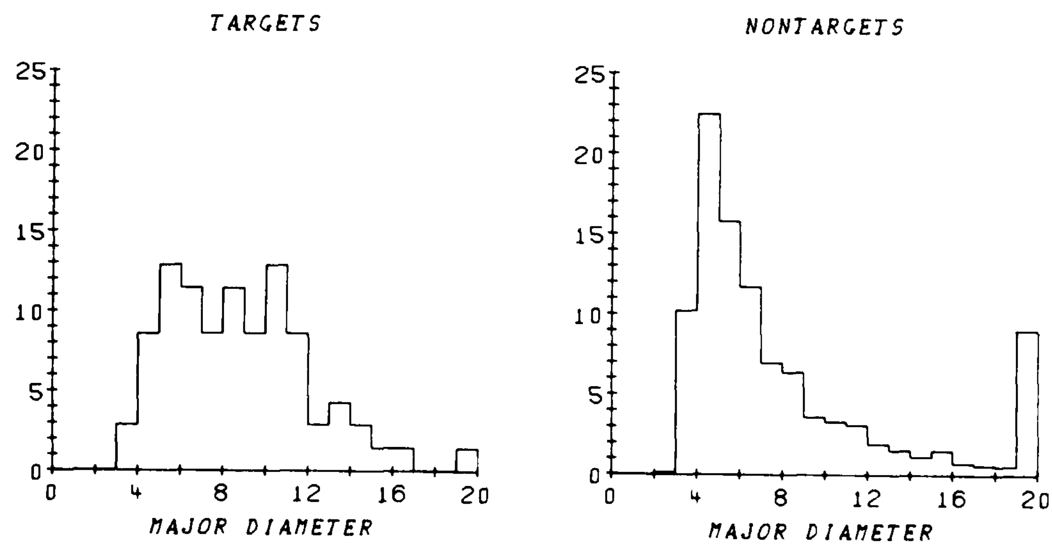

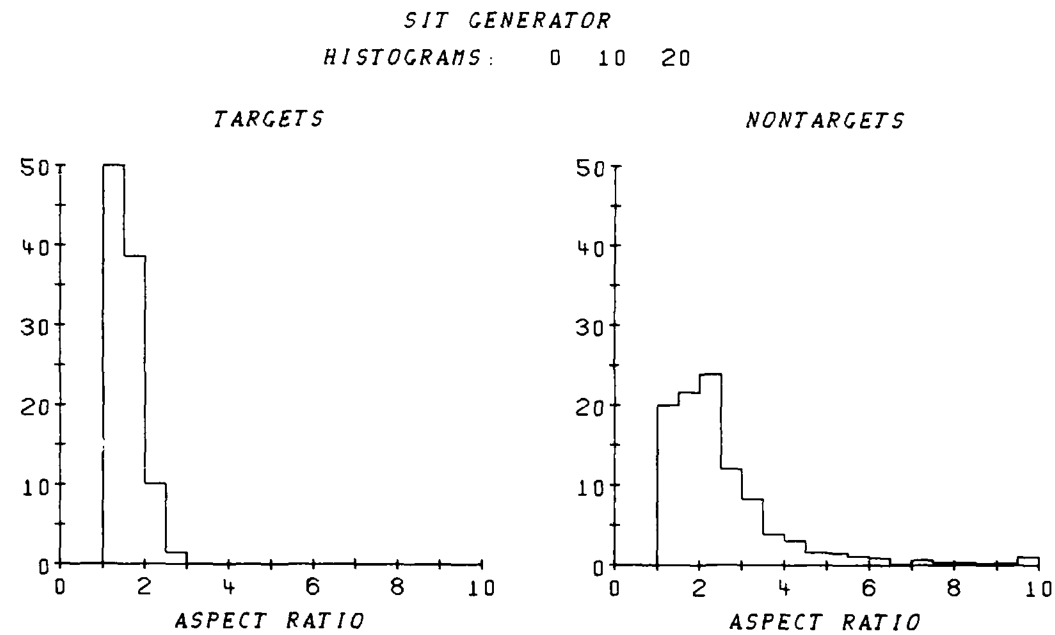
SIT GENERATOR
HISTOGRAMS: 0 10 20
TARGETS
NONTARGETS
50
40
30
20
10
0
0 2 4 6 8 10
ASPECT RATIO
ASPECT RATIO

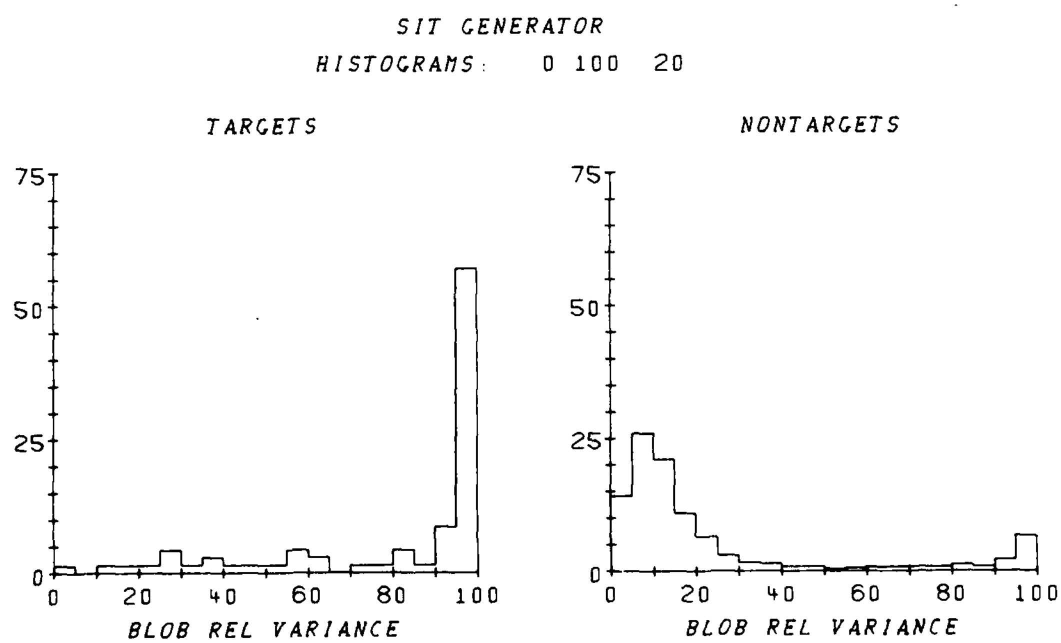
SIT GENERATOR
HISTOGRAMS: 0 100 20
TARGETS
NONTARGETS
75
50
25
0
0 20 40 60 80 100
BLOB REL VARIANCE
BLOB REL VARIANCE

FIGURE 8 - The 13 histogram pair derived from segmenter No. 6. The associated features are the same as in Fig. 7.

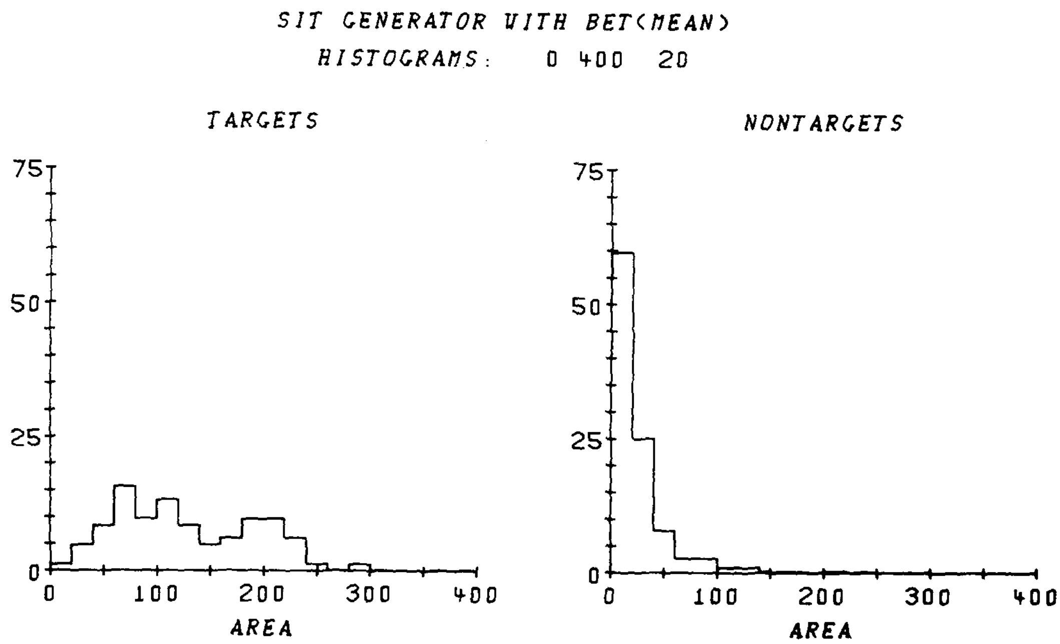

SIT GENERATOR WITH BET(MEAN)

HISTOGRAMS: 0 200 20

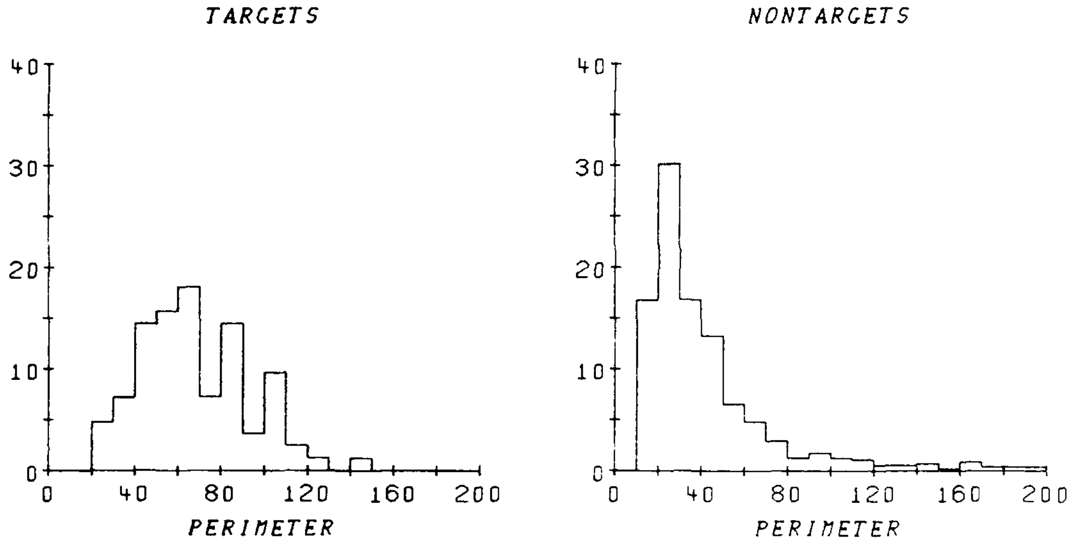

SIT GENERATOR WITH BET(MEAN)

HISTOGRAMS: 0 100 20

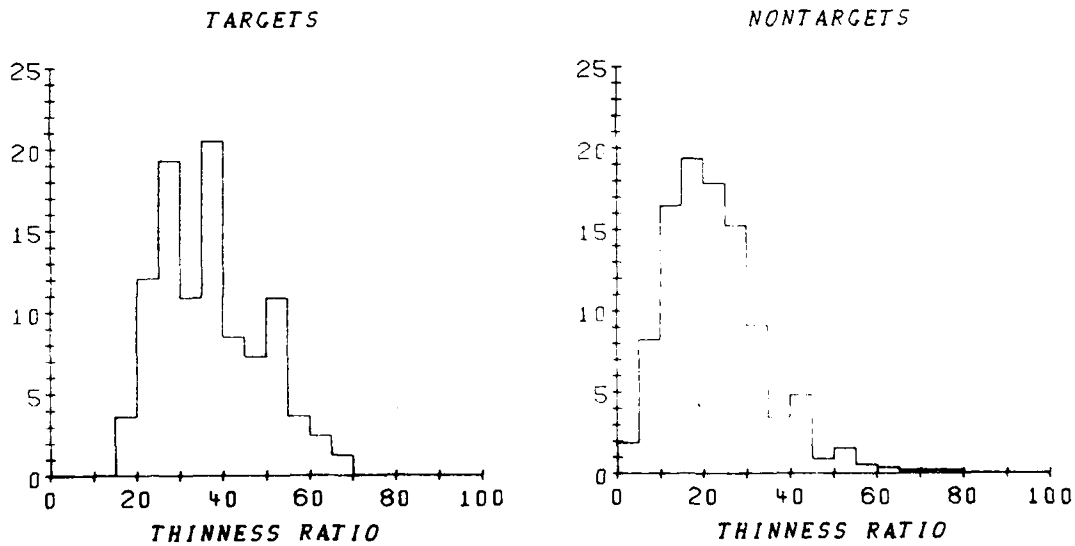

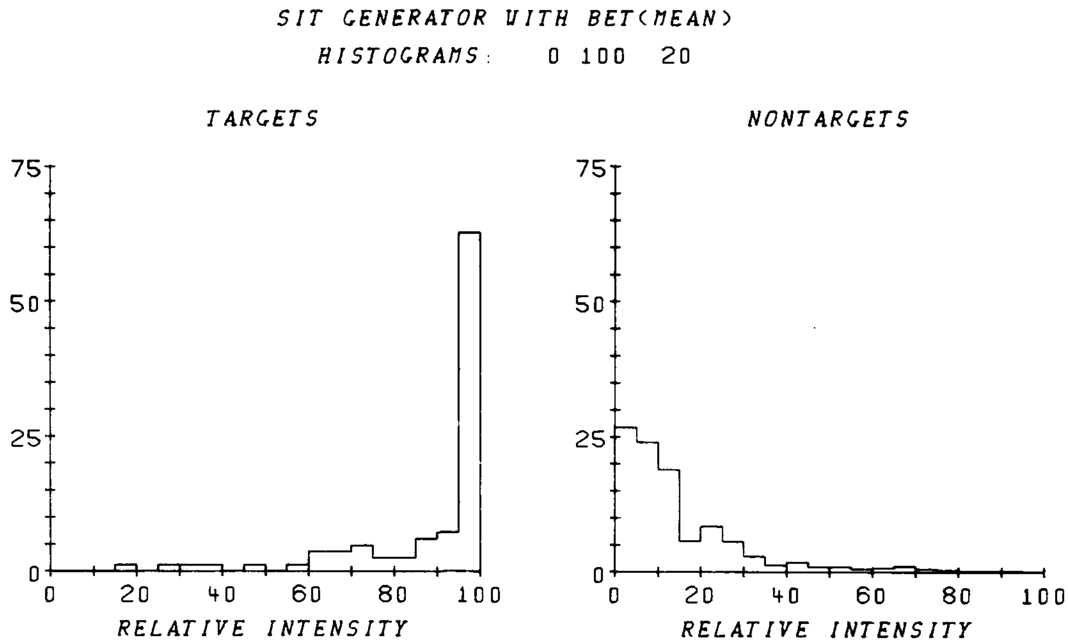
SIT GENERATOR WITH BET(MEAN)
HISTOGRAMS: 0 100 20
TARGETS
NONTARGETS
75
50
25
0
0 20 40 60 80 100
RELATIVE INTENSITY
RELATIVE INTENSITY

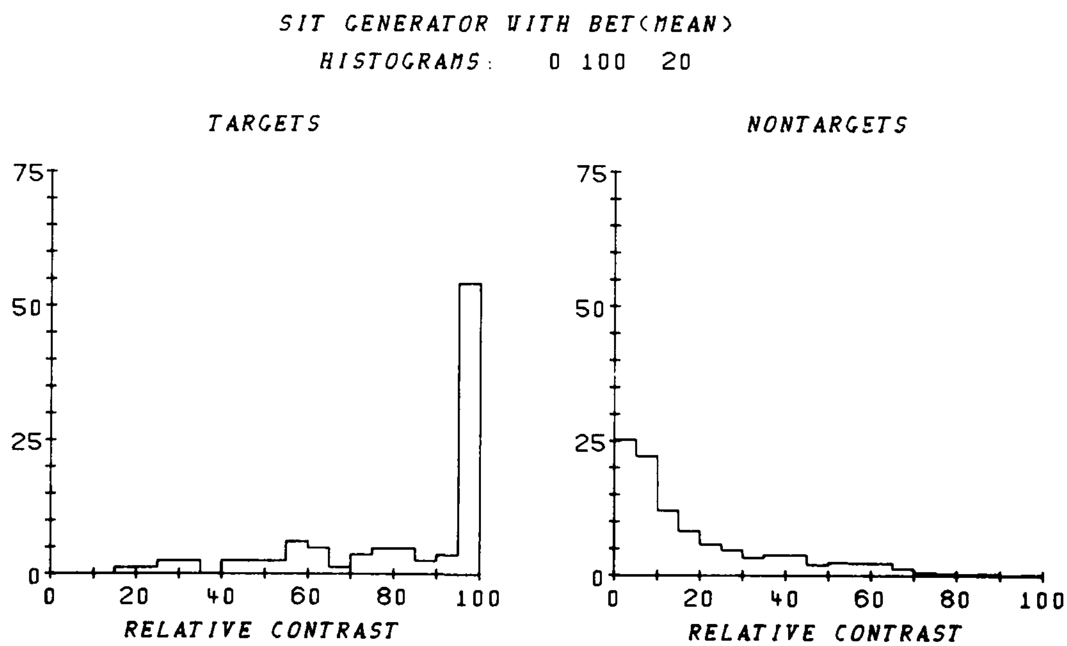
SIT GENERATOR WITH BET(MEAN)
HISTOGRAMS: 0 100 20
TARGETS
NONTARGETS
75
50
25
0
0 20 40 60 80 100
RELATIVE CONTRAST
RELATIVE CONTRAST

SIT GENERATOR WITH BET(MEAN)

HISTOGRAMS: 0 100 20

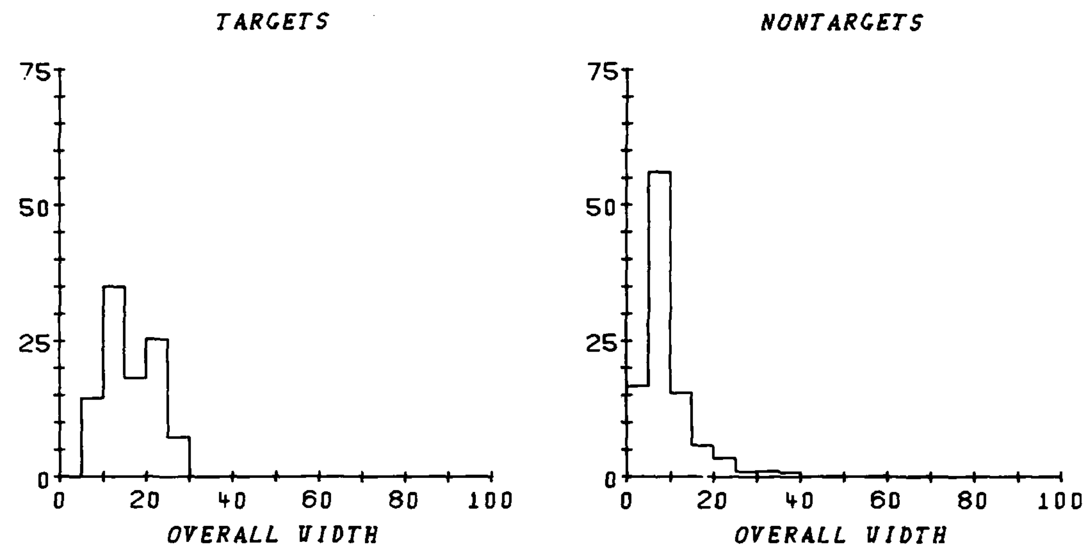

SIT GENERATOR WITH BET(MEAN)

HISTOGRAMS: 0 100 20

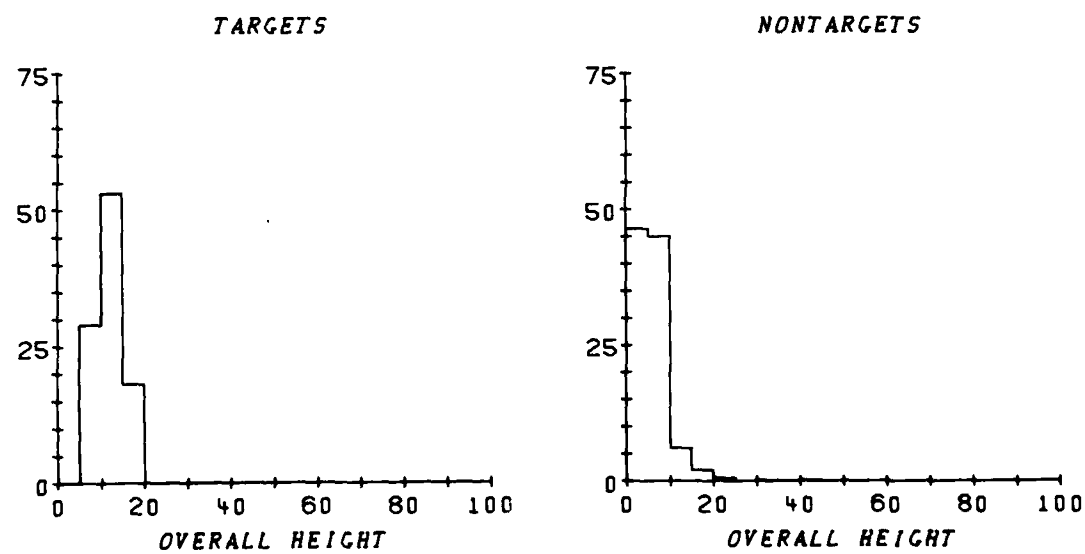

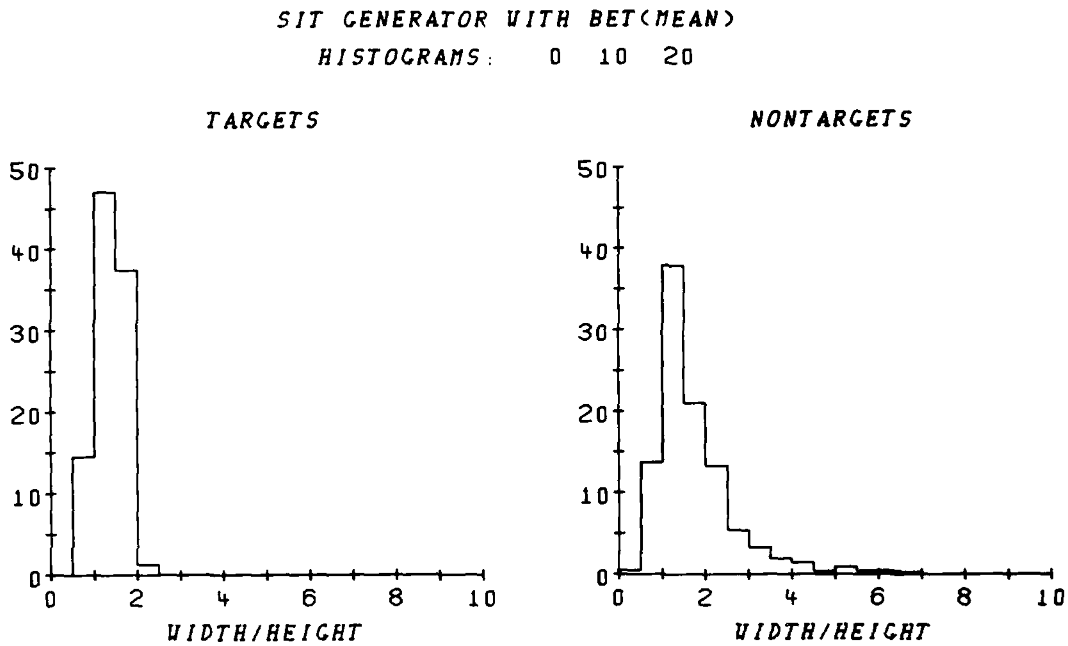
SIT GENERATOR WITH BET(MEAN)
HISTOGRAMS: 0 10 20
TARGETS
NONTARGETS
0
10
20
30
40
50
0
2
4
6
8
10
WIDTH/HEIGHT
WIDTH/HEIGHT

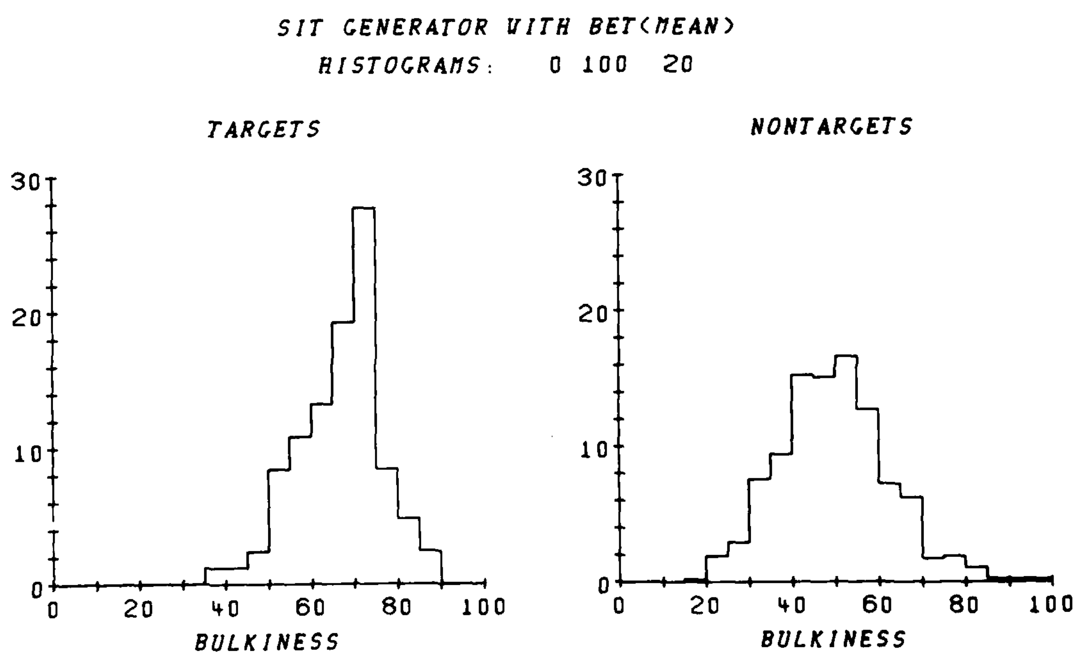
SIT GENERATOR WITH BET(MEAN)
HISTOGRAMS: 0 100 20
TARGETS
NONTARGETS
0
10
20
30
0
20
40
60
80
100
BULKINESS
BULKINESS

SIT GENERATOR WITH BET(MEAN)
HISTOGRAMS: 0 20 20

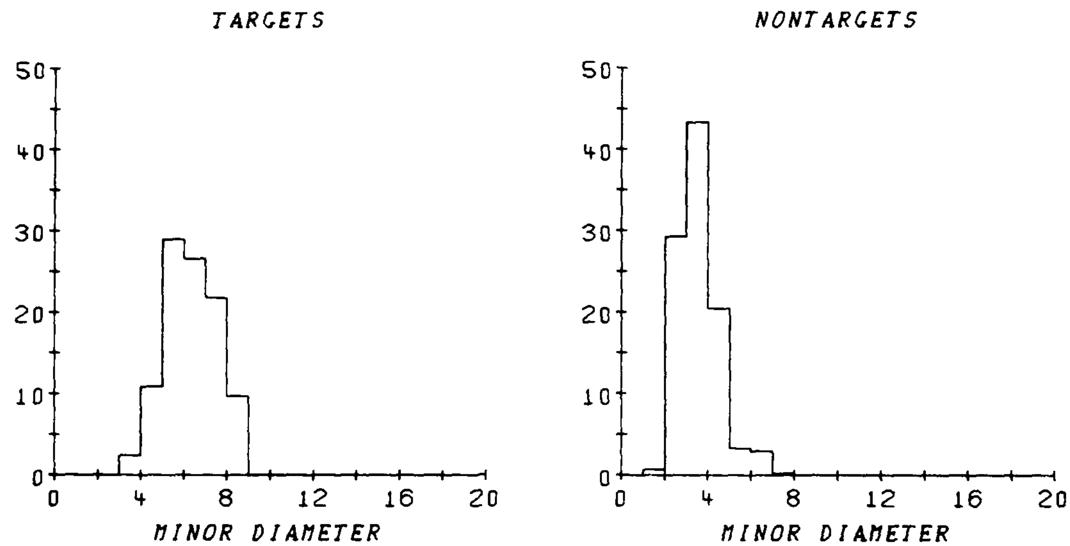

SIT GENERATOR WITH BET(MEAN)
HISTOGRAMS: 0 20 20

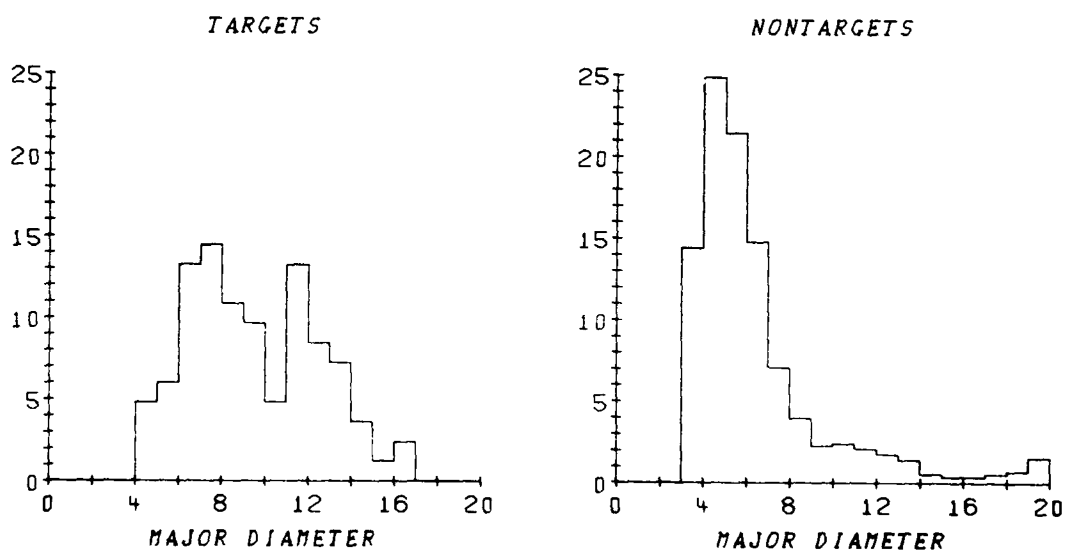

SIT GENERATOR WITH BET(MEAN)
HISTOGRAMS: 0 10 20

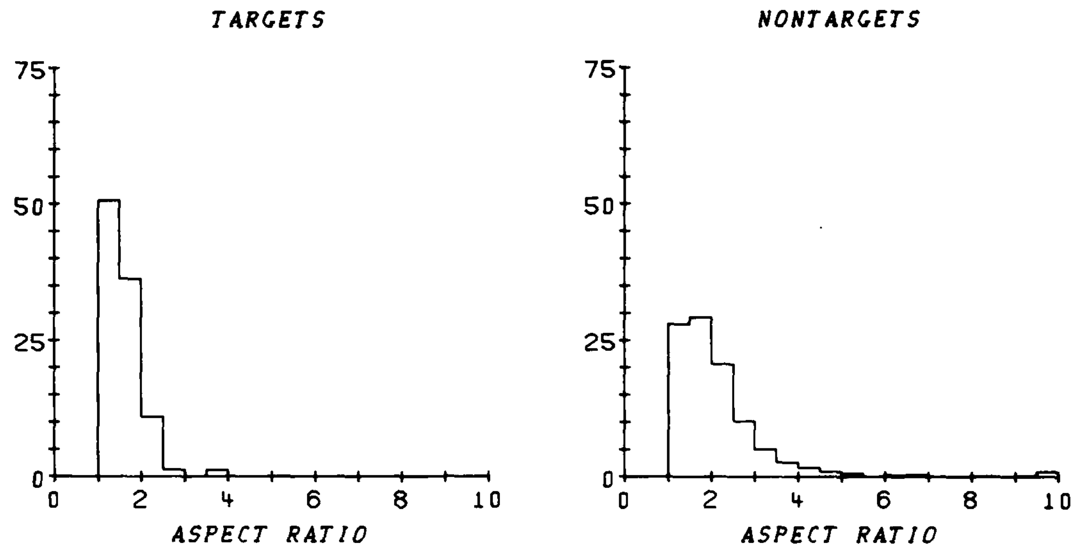

SIT GENERATOR WITH BET(MEAN)
HISTOGRAMS: 0 100 20

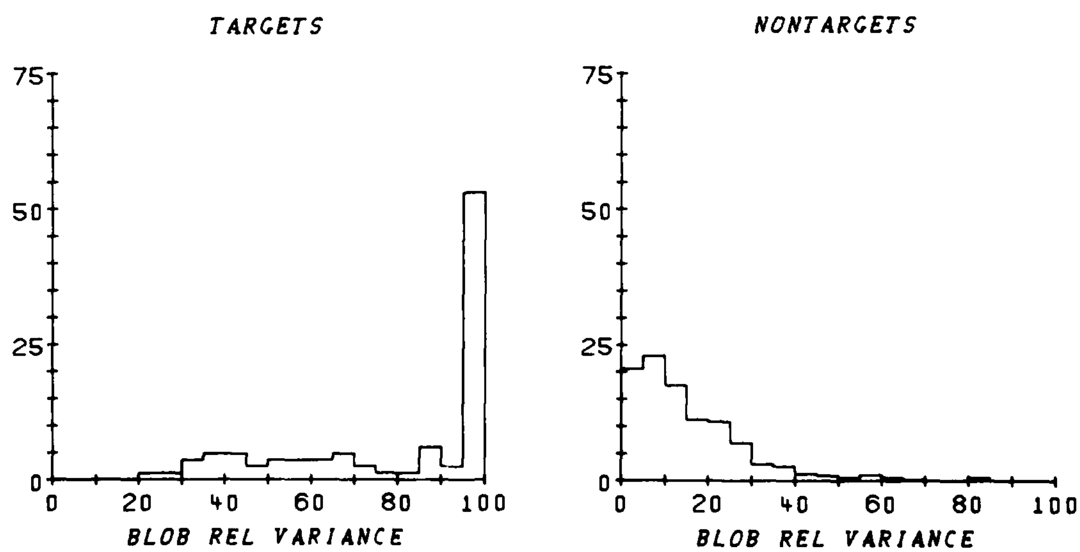

FIGURE 9 - This figure, derived from segmenter No. 1, shows scatter plots of the following features:

a) Relative Intensity versus Relative Contrast
b) Relative Intensity versus Relative Blob Variance
c) Relative Contrast versus Relative Blob Variance

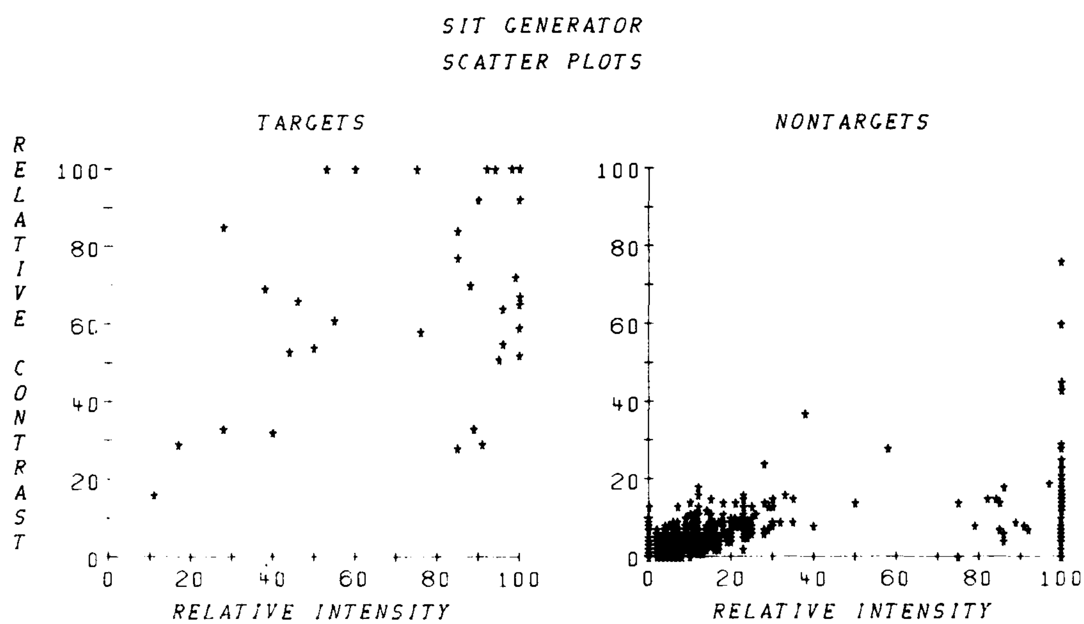

SIT GENERATOR
SCATTER PLOTS

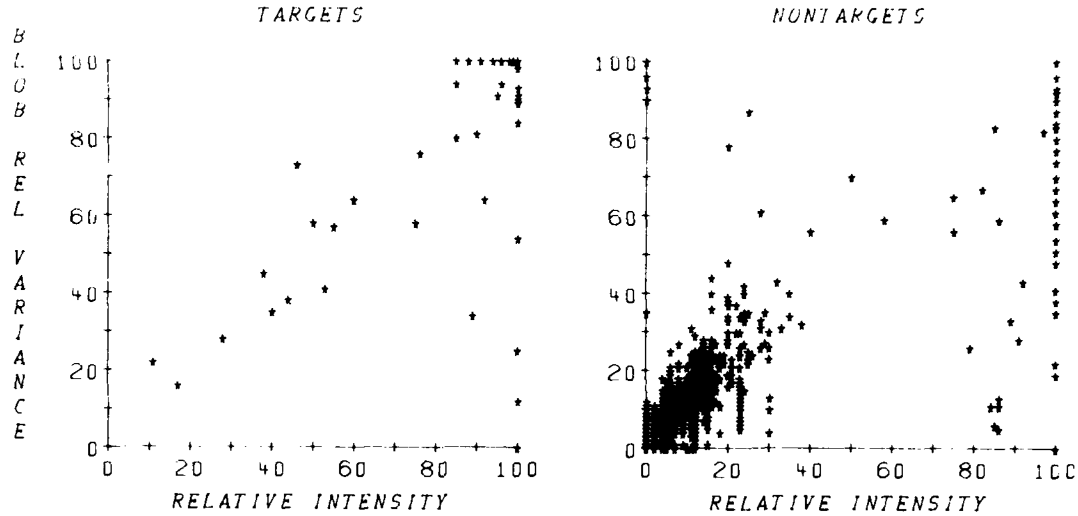

SIT GENERATOR
SCATTER PLOTS

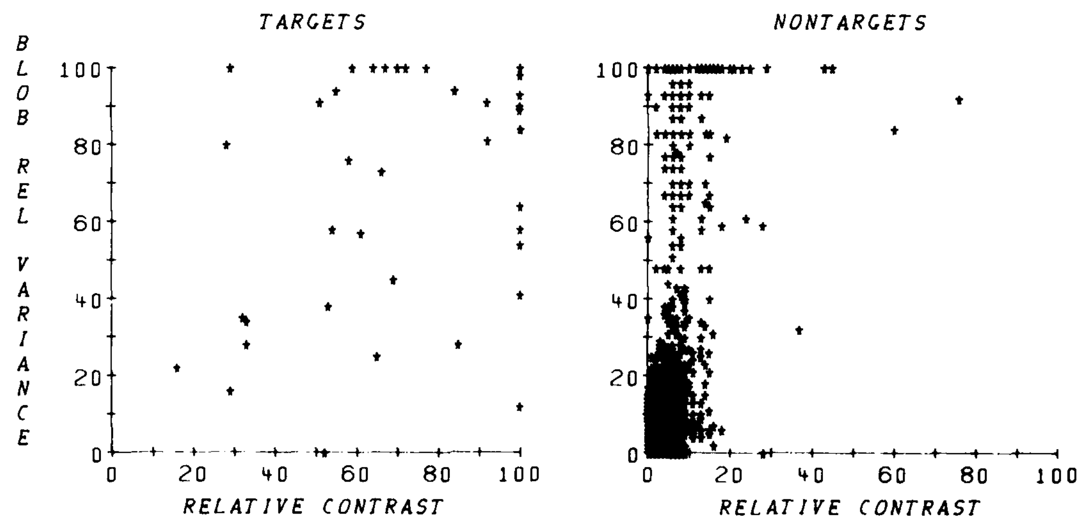

FIGURE 10 - This figure, derived from segmenter No. 6, is the equivalent of Fig. 9.

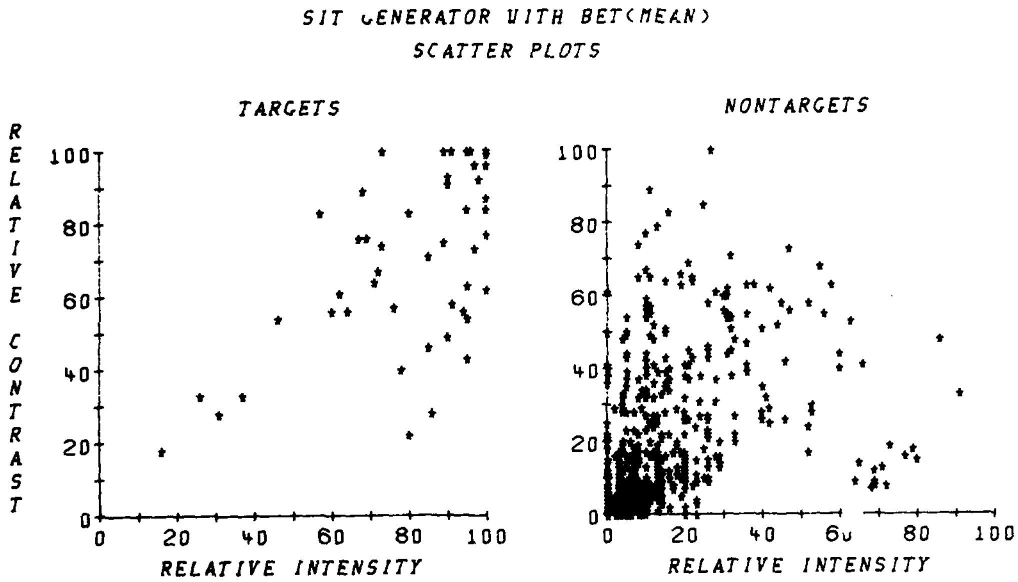

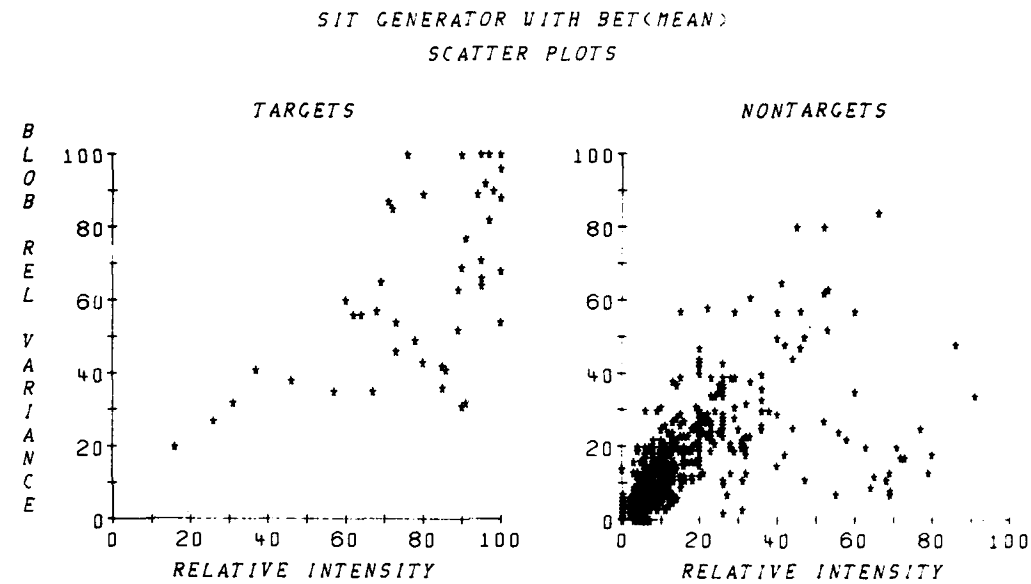
SIT GENERATOR WITH BET(MEAN)
SCATTER PLOTS
TARGETS
NONTARGETS
BLOB REL VARIANCE
RELATIVE INTENSITY
RELATIVE INTENSITY

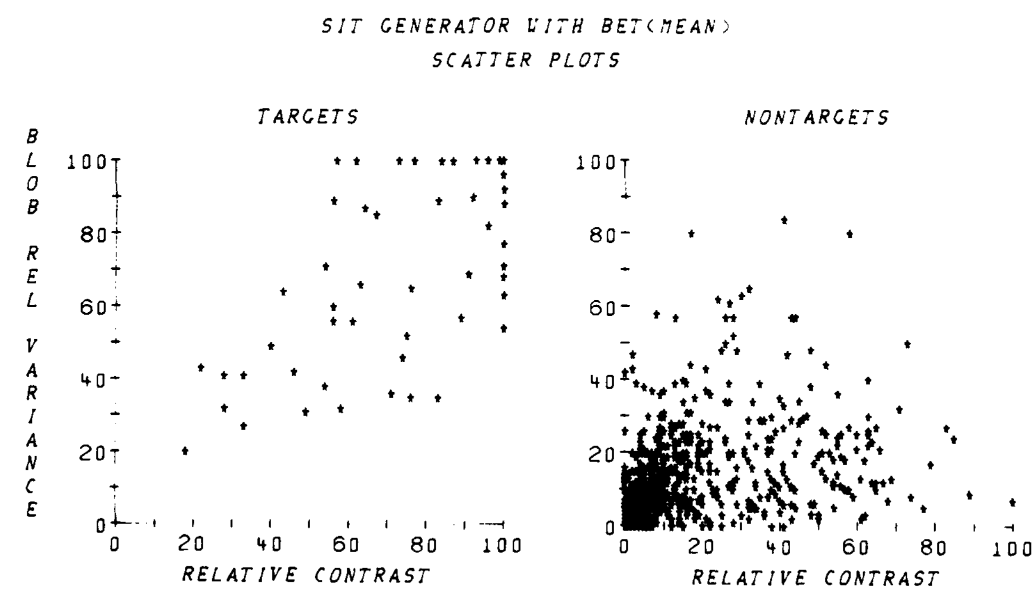
SIT GENERATOR WITH BET(MEAN)
SCATTER PLOTS
TARGETS
NONTARGETS
BLOB REL VARIANCE
RELATIVE CONTRAST
RELATIVE CONTRAST

6.3 Comments

We will often refer, in the remainder of this section, to specific images of the Alabama Data Base. Readers interested in viewing these images are directed to Ref. 1 where histogram-equalized pictures of the 43 images that made up the database are given along with the relevant ground truth.

We mentionned in Sec. 4 that a given segmenter is first valued according to its ability to segment just about all the targets liable to be perceived in the pictured scene. In Table I, the objects generated by all 6 segmenters, in relation to the Alabama Data Base, are sorted by object type through the agency of the provided ground truth. As mentioned before, the SIT Generator produces the largest number of nontargets. This might be an indication that this segmenter is more prone to false alarms than the others. However, there are really no grounds for believing that the number of false alarms is generally directly proportional to the number of nontargets. The only thing we know for sure is that the number of false alarms will be equal to 0 if the number of nontargets is equal to 0. On the other hand, a small number of nontargets is no guarantee of efficiency as one can see from the figures for the SCIT Generator. Of the first 3 segmenters, the ISCIT Generator is the best at extracting targets. Nevertheless, its extraction rate is not as good as the one of a segmentation algorithm incorporating, in one way or another, the Background Elimination Technique described in detail in Ref. 1, and outlined in Sec. 2 of this report. This is then a good point for this technique. On the sole basis of their extraction rate, 2 segmenters, segmenters No. 4 and No. 6, surpass the rest. The 2 targets missed by both these segmenters are APC's: one in image 13 and one in image 41. Interestingly enough the first 3 segmenters do extract the APC in image 13. However, the APC in image 41 eludes all 6 extraction schemes. Finally, Table I shows that the tank is probably the easiest target to extract.

TABLE I

NUMBER OF SEGMENTED OBJETS PER CLASS

SEGMENTATION ALGORITHM	TANK	APC	JEEP	TARGETS	NONTARGETS
1. SIT GENERATOR	40	19	14	73	1011
2. SCIT GENERATOR	32	18	10	60	318
3. ISCIT GENERATOR	39	24	14	77	404
4. SIT GENERATOR WITH BET(HFS)	40	27	16	83	581
5. SIT GENERATOR WITH BET(MAX)	39	25	16	80	838
6. SIT GENERATOR WITH BET(MEAN)	40	27	16	83	584
GROUND TRUTH DATA	40	29	16	85	

Table I does not give a complete picture of the segmentation results. It leaves out the problem of repeated detections, that is, of targets split up into several blobs and hence likely to be construed as forming a group of distinct targets. This problem is not too important as far as segmenter No. 6 is concerned for the only multiblob targets (in Table I a multiblob target is classified as 1 target of the appropriate type) are the tank in image 7 and the bus in image 33, 2 relatively large-size targets. Segmenter No. 4, however, is much more affected with nearly 10 multiblob targets. As this serious flaw might greatly reduce the usefulness of this segmenter, segmenter No. 6 emerges here as the best. No multiblob target arises from the SIT Generator whereas segmenters 2, 3 and 5 produce only one such target.

Another useful criterion to assess the practicality of a segmenter is the fidelity of the segmentation process with regard to the geometrical properties of the targets. From this point of view, segmenter No. 1 and No. 6 may be rated as the best. Segmenter No. 4 exhibits a marked tendency to narrow the targets (Fig. 3d) but this is to be expected since it is based on the HFS images. The other 3 segmenters (2, 3 and 5) generally exaggerate the size of the targets even to the point of, sometimes, merging 2 neighboring targets into a single blob (in Table I such a blob was classified as a nontarget). Also, in a few instances, although the target was not connected to another target blob its shape was so distorted as to be unrecognizable. This is the case, for example, in relation to segmenter No. 5, of the tank in image 17 and of the APC's in images 9 and 17. These distorted targets were classified as nontargets, in Table I. The last 3 segmenters in this table would otherwise have the same extraction rate.

It should be obvious from the preceding paragraphs that the segmenters 1 and 6 are the most interesting ones. This explains why, not to mention more prosaic reasons, the only feature histograms appearing in this report pertain to these 2 segmenters. However, the general conclusions drawn from Figs. 7 to 10 apply equally well to all 6 segmenters.

By examining Figs. 7 and 8, we readily conclude that only 3 out of the 13 plotted features are really peculiar to a target blob. These are:

a) Relative Intensity

b) Relative Contrast

c) Relative Blob Variance

It turns out quite naturally that these are relative features. Indeed, this is the only way to eliminate the constraints imposed by the experimental conditions and also by the fact that we are dealing with a discontinuous sequence of pictures. Although the other features are no good at discriminating targets from nontargets, they might well be very useful to classify the targets themselves. However, this is something we will not attempt to do in the present report. It is not surprising that intensity and contrast features are distinguishing target characteristics since we are dealing with IR imagery. Nevertheless, it is amazing to observe that so is doing the variance. Given the size of the targets, we would rather intuitively expect the variance to be insignificant, but the experimental results show that this is not the case. Another important aspect should be emphasized here. In Fig. 7, the numerical values assigned to the various features were derived from the original raw images, whereas in Fig. 8 these numerical values originate from the Mean Fine Structure images, that is, images that bear little resemblance to the original ones. So one may wonder, in the case of segmenter No. 6, what the feature values would have been had their evaluation been based on the original images. It makes no difference for shape features (e.g. area, perimeter, overall width, etc.) but this should normally affect moment features such as the intensity, contrast, minor diameter etc. that depend on the gray level of the pixels involved. Fig. 11 is meant to elucidate the question. As we can see, the moment features in Figs. 8 and 11 exhibit the same trends except for the relative intensity that is obviously not a distinctive target feature when evaluated from the original images. There is then no point in going back to the original images insofar as segmenter No. 6 is concerned. This, in fact, confirms that BET saves all the useful information about the targets.

Once features peculiarly belonging to the targets have been identified, one can assess the degree of distinctiveness imparted to the targets as opposed to the nontargets. That quantity is proportional to the extent of overlap of the relevant pair of histograms (Figs. 7 and 8), and then can be determined accordingly. However, it serves our purpose better to give here some examples of the classification results one can expect from the aforementioned subset of 3 features. To this end, the confusion matrix:

$$\begin{pmatrix} \text{target} & \text{miss} \\ \text{false alarm} & \text{nontarget} \end{pmatrix}$$

has been determined for various combinations of the 3 features. The decision rule for each feature (the t's in Fig. 12) is simply a fixed threshold whose level corresponds to the 5th percentile of the feature's target histogram. Hence, all the objects associated with a feature whose value is greater than the specified threshold are discarded as nontargets (Fig. 12). The classification results that ensue for each feature alone, for a combination of 2 out of the 3 features, and for all 3 features together are shown in Table II. In this table, these features are identified as follows (Sec. 4): 6: Relative Intensity; 7: Relative Contrast; and 17: Relative Blob Variance. It is important to note that the order of the features in a combination is not immaterial for, given the structure of the decision tree (Fig. 12), the results are not necessarily the same if the features undergo a permutation. Also, for the same reason, the probability of detection (number of targets classified as such) of any combination of features cannot exceed that of its least effective member. However, by combining features one can greatly reduce the number of false alarms. To convince oneself that this is indeed the case it suffices to compare the confusion matrix for feature 7 (Table II) to that for features 7 and 17. Clearly, a trade-off has to be made between the detection rate one would like to obtain and the false alarm rate that can be tolerated. In any way, Table II shows that it should be possible to obtain with segmenter No. 6 a detection rate in excess of 90% with a false alarm rate not greater than 3%.

TABLE II

EXEMPLES OF CLASSIFICATION RESULTS FOR SEGMENTER NO. 6

FEATURE NUMBER	FEATURE THRESHOLD	CONFUSION MATRIX (ABSOLUTE)	CONFUSION MATRIX (PERCENTAGE)	TARGET MISCLASSIFICATIONS TYPE(IMAGE)
6	45	79 4 31 553	95 4 5 94	A(28); T(36); T(37) A(38)
7	30	79 4 130 454	94 4 22 77	T(13); A(28); T(37) A(29)
17	30	81 2 60 524	97 2 10 89	A(28); T(36)
6 7	45 30	77 6 15 569	92 7 2 97	T(13); A(28); T(36) T(37); A(38); A(39)
6 17	45 30	79 4 13 571	95 4 2 97	A(28); T(36); T(37) A(38)
7 17	30 30	78 5 23 561	93 6 3 95	T(13); A(28); T(36) T(37); A(39)
6 7 17	45 30 30	77 6 9 575	92 7 1 98	T(13); A(28); T(36) T(37); A(38); A(39)

A classification experiment using the Fisher linear discriminant (Ref. 5) was attempted on the scatter plots (Figs. 9 and 10) but the results merely point out that one is entitled to use the features independently. The Fisher linear discriminant attempts to find the optimum linear projection of the feature vectors onto a line, and the optimum partition of this line, such that the ratio of between-class scatter to within-class scatter is maximized (Ref. 6).

FIGURE 11 - The 6 histogram pairs derived from segmenter No. 6. The associated features listed below were evaluated from the original raw images instead of the Mean Fine Structure images:

a) Relative Intensity
b) Relative Contrast
c) Minor Diameter
d) Major Diameter
e) Aspect Ratio
f) Blob Relative Variance

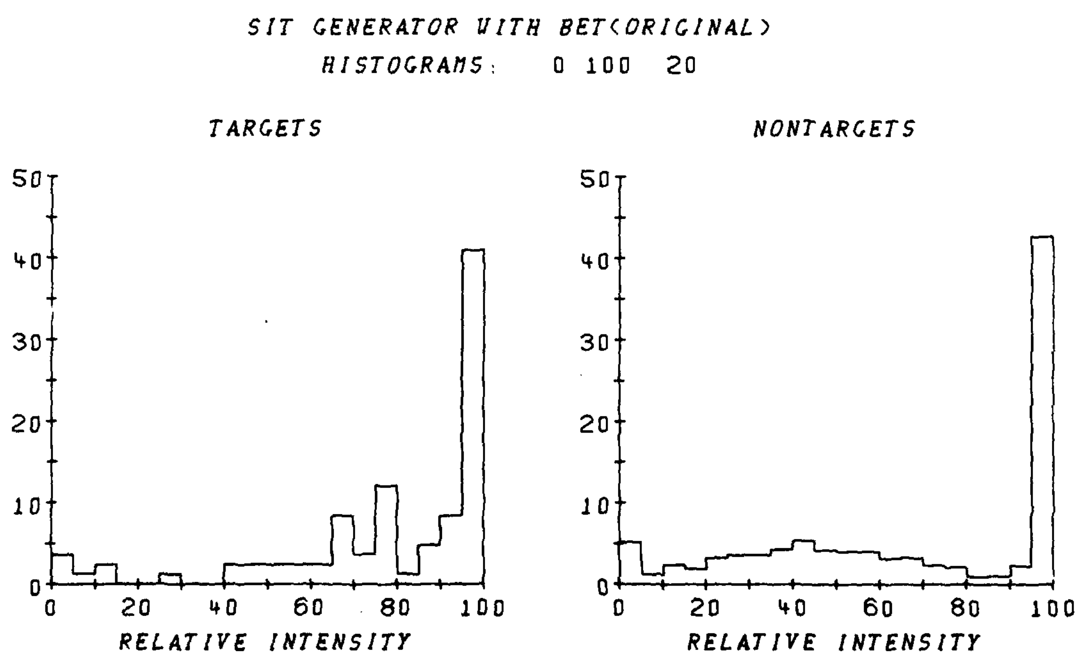

SIT GENERATOR WITH BET(ORIGINAL)
HISTOGRAMS: 0 100 20

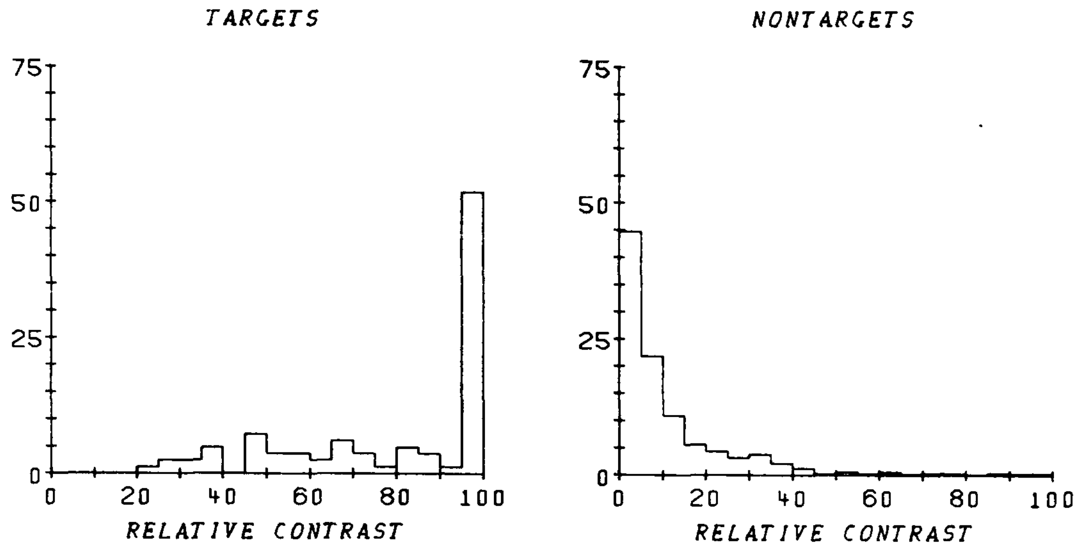

SIT GENERATOR WITH BET(ORIGINAL)
HISTOGRAMS: 0 20 20

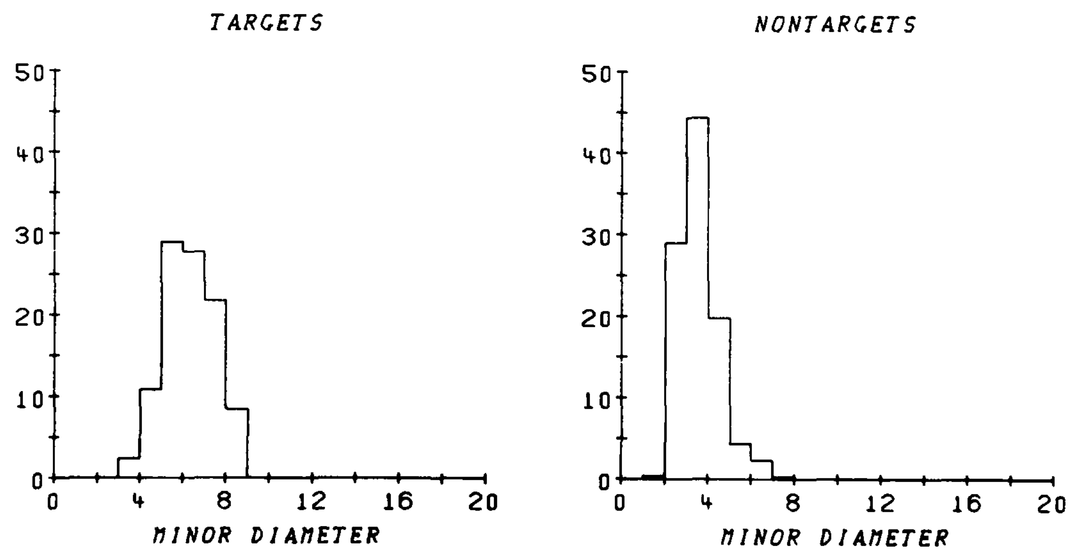

SIT GENERATOR WITH BET(ORIGINAL)
HISTOGRAMS: 0 20 20

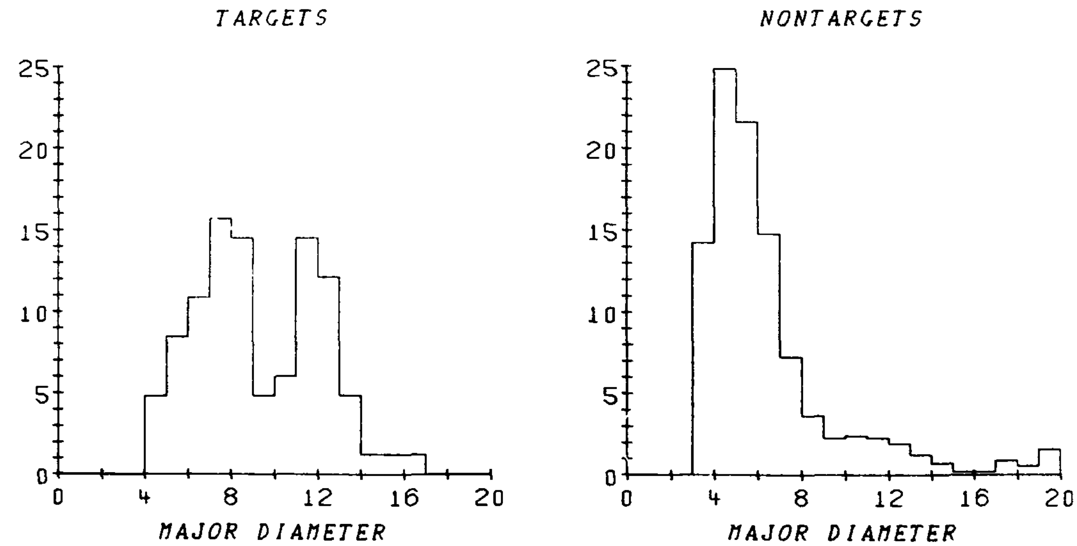

SIT GENERATOR WITH BET(ORIGINAL)
HISTOGRAMS: 0 10 20

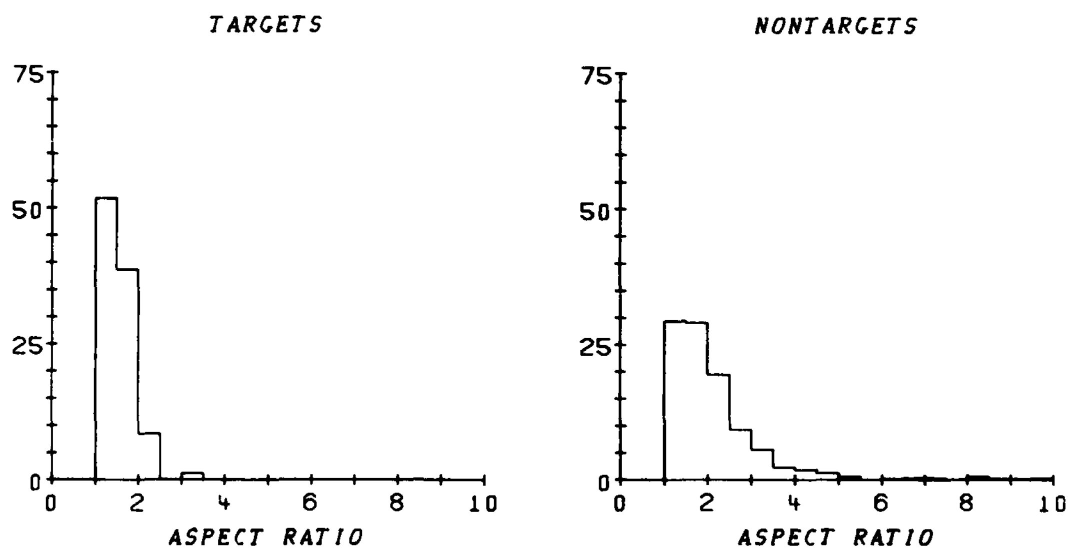

SIT GENERATOR WITH BET(ORIGINAL)

HISTOGRAMS: 0 100 20

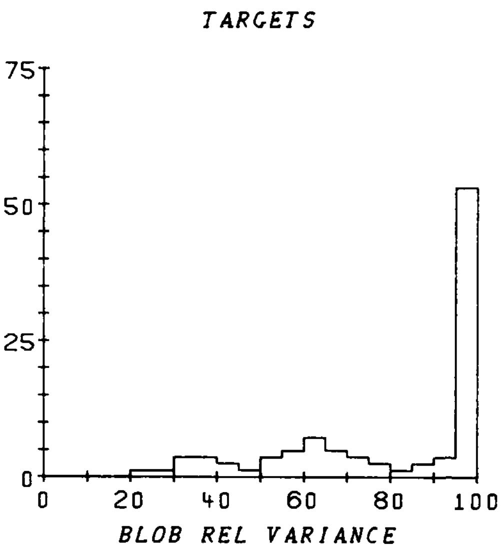

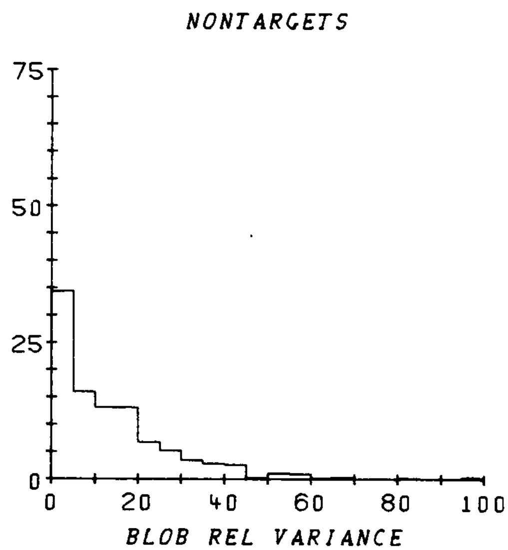

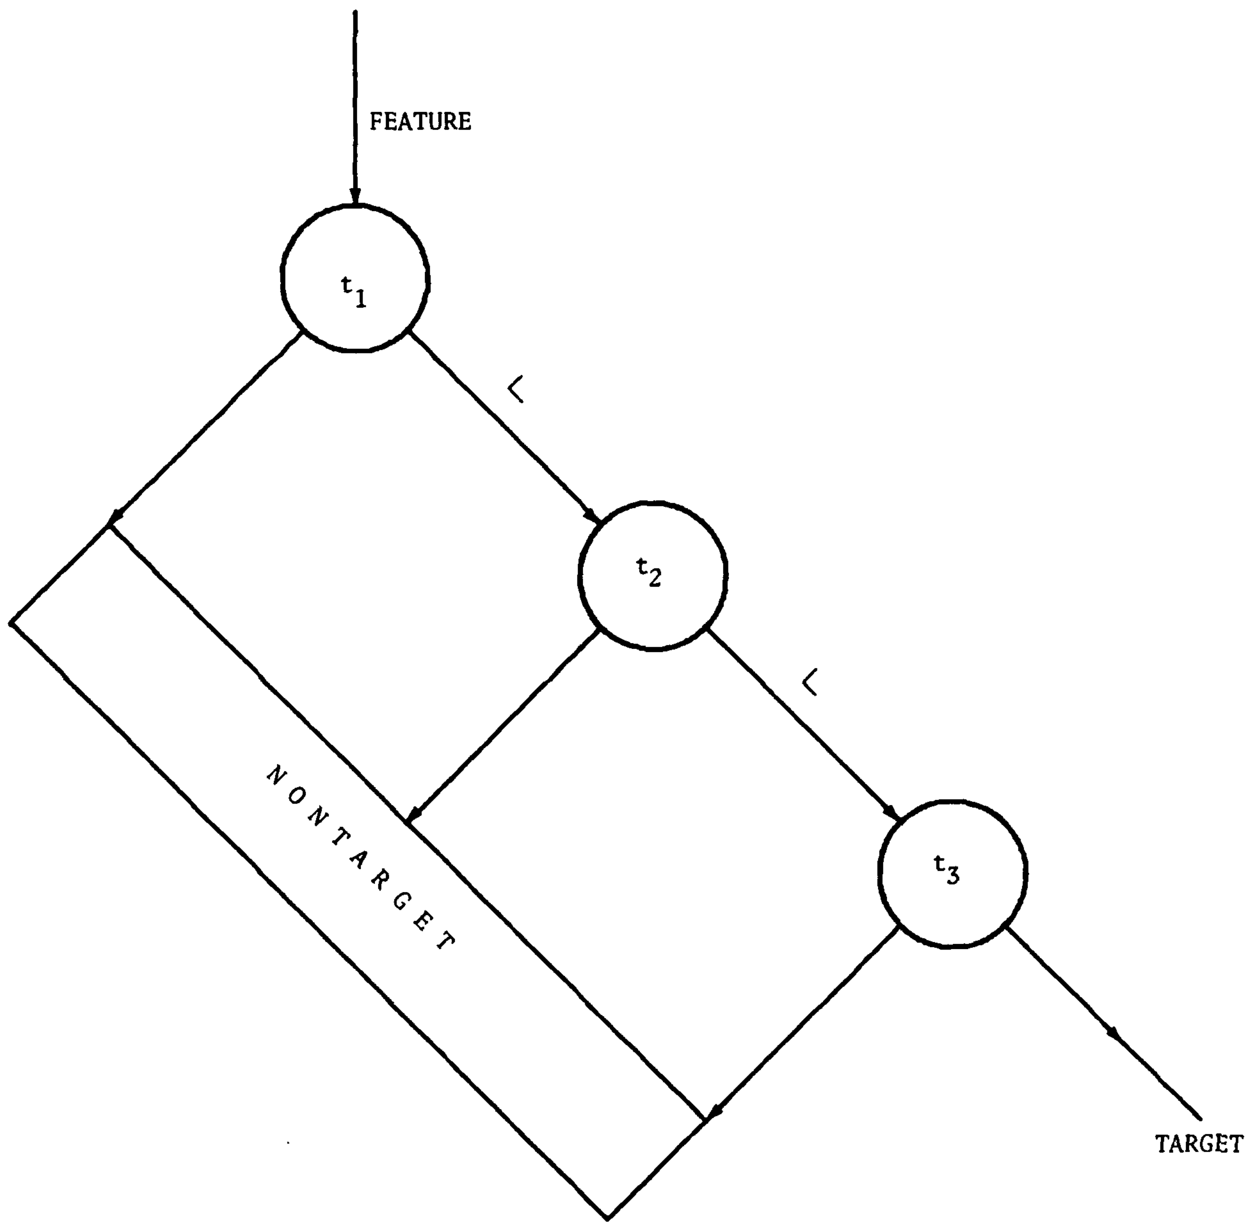

FIGURE 12 - Decision Tree

7.0 CONCLUSION

The purpose of the present report was to evaluate the performance of 6 different segmentation algorithms or segmenters based on the single assumption that the targets present a larger thermal image than the background. The first 3 segmenters considered deal with an image in its entirety, whereas the last 3 incorporate a technique, the Background Elimination Technique or BET, which aims at eliminating wholly or partly the background. The segmenters are judged according to:

a) their extraction rate;

b) the fidelity of the segmentation with respect to the geometrical properties of the extracted targets;

c) the degree of distinctiveness imparted to the extracted targets as opposed to the nontargets.

The 3 segmenters relying on BET have a better extraction rate than the other 3 that try to cope with the background simply by partitioning the image. Most segmenters here distort in one way or another the shape of the targets. The 2 exceptions are segmenters No. 1 (Single Intensity Threshold Generator or SIT Generator) and No. 6 (SIT Generator in conjunction with BET through the Mean Fine Structure image). To determine the degree of distinctiveness, one must first single out the feature or set of features that most characterizes the targets. The experimental results show that the following relative features are the best for this purpose:

a) Relative Intensity

b) Relative Contrast

c) Relative Blob Variance

It is not surprising to note that intensity and contrast features are distinguishing target characteristics since we are dealing with IR imagery. It is amazing, however, to observe that so is doing the variance feature. The classification results one can expect from these features together with the segmenter that proves to be the best (segmenter No. 6) amount to a detection rate in excess of 90% with a false alarm rate not greater than 3%. On the other hand, it should be possible to refine further that segmenter to bring the extraction rate up to 100%, although 97% is a percentage already quite acceptable. However, the next thing to do would rather be to test segmenter No. 6 on more complex imagery.

8.0 REFERENCES

1. Sévigny, L., "Segmentation Algorithms for Detection of Targets in IR Imagery", DREV R-4180/80, UNCLASSIFIED, TO BE PUBLISHED

2. Sévigny, L., "Simulation d'un système d'acquisition automatique d'objectif infrarouge dans un contexte sol-sol", DREV R-4081/77, juin 1977, NON CLASSIFIE

3. Sévigny, L., "La reconnaissance de forme et l'acquisition d'objectif en infrarouge: nouvel algorithme de détection", DREV R-4099/78, mars 1978, NON CLASSIFIE

4. Sévigny, L., "Extracteurs séquentiels pour l'acquisition de cibles sur images", DREV R-4153/79, août 1979, NON CLASSIFIE

5. Rosenfeld, A. and Kak, A.C., "Digital Picture Processing", Academic Press, 1976.

6. Duda, R.O. and Hart, P.E., "Pattern Classification and Scene Analysis", Wiley, 1973.

7. Papoulis, A., "Probability, Random Variables, and Stochastic Processes", McGraw-Hill, 1965.

8. Agrawala, A.K. and Kulkarni, A.V., "A Sequential Approach to the Extraction of Shape Features", Computer Graphics and Image Processing, Volume 6, pp. 538-557, 1977.

CRDV R-4172/80 (NON CLASSIFIE)

Bureau - Recherche et Développement, MDN, Canada.
CRDV, C.P. 880, Courcelette, Qué. G0A 1R0

"Evaluation d'une classe de segmenteurs pour images IR"
par L. Sévigny

Dans ce rapport nous évaluons l'efficacité de 6 algorithmes spécialisés dans la segmentation de cibles sur images IR. Ces algorithmes de segmentation, ou segmenteurs, reposent tous sur le principe voulant que la signature thermique d'une cible soit supérieure à celle de tout objet de l'arrière-plan. Les 3 premiers segmenteurs abordent l'image de front, en son entier, tandis que les 3 derniers incorporent la technique de redressement de l'arrière-plan (TRAP), visant à éliminer l'arrière-plan en tout ou en partie en le nivelant. Les divers algorithmes sont jugés d'après a) leur taux d'extraction, b) la fidélité du processus de segmentation en ce qui concerne les propriétés géométriques des cibles et, finalement, d'après c) le degré d'individualisation imprimé aux cibles extraites par rapport aux pseudo-cibles. Les 3 segmenteurs centrés sur TRAP ont une meilleure probabilité d'extraction que les trois autres qui essayent d'éliminer l'arrière-plan simplement en morcelant l'image. Par ailleurs, la plupart des segmenteurs considérés dans ce rapport altèrent d'une façon ou d'une autre la forme des cibles. Les deux exceptions à cette règle sont les segmenteurs No. 1 (générateur de silhouettes à seuil d'intensité unique) et No. 6 (le précédent segmenteur allié à une version particulière de TRAP). Les résultats expérimentaux montrent que l'intensité, le contraste et la variance sont les traits qui permettent le mieux de départager les cibles et les pseudo-cibles. Des expériences de classification réalisées à l'aide du segmenteur No. 6, lequel s'avère le meilleur, en fonction de ces traits caractéristiques indiquent que l'on peut espérer obtenir un taux de détection qui excède 90% avec un taux de fausses alarmes inférieur à 3%. (NC)

CRDV R-4172/80 (NON CLASSIFIE)

Bureau - Recherche et Développement, MDN, Canada.
CRDV, C.P. 880, Courcelette, Qué. G0A 1R0

"Evaluation d'une classe de segmenteurs pour images IR"
par L. Sévigny

Dans ce rapport nous évaluons l'efficacité de 6 algorithmes spécialisés dans la segmentation de cibles sur images IR. Ces algorithmes de segmentation, ou segmenteurs, reposent tous sur le principe voulant que la signature thermique d'une cible soit supérieure à celle de tout objet de l'arrière-plan. Les 3 premiers segmenteurs abordent l'image de front, en son entier, tandis que les 3 derniers incorporent la technique de redressement de l'arrière-plan (TRAP), visant à éliminer l'arrière-plan en tout ou en partie en le nivelant. Les divers algorithmes sont jugés d'après a) leur taux d'extraction, b) la fidélité du processus de segmentation en ce qui concerne les propriétés géométriques des cibles et, finalement, d'après c) le degré d'individualisation imprimé aux cibles extraites par rapport aux pseudo-cibles. Les 3 segmenteurs centrés sur TRAP ont une meilleure probabilité d'extraction que les trois autres qui essayent d'éliminer l'arrière-plan simplement en morcelant l'image. Par ailleurs, la plupart des segmenteurs considérés dans ce rapport altèrent d'une façon ou d'une autre la forme des cibles. Les deux exceptions à cette règle sont les segmenteurs No. 1 (générateur de silhouettes à seuil d'intensité unique) et No. 6 (le précédent segmenteur allié à une version particulière de TRAP). Les résultats expérimentaux montrent que l'intensité, le contraste et la variance sont les traits qui permettent le mieux de départager les cibles et les pseudo-cibles. Des expériences de classification réalisées à l'aide du segmenteur No. 6, lequel s'avère le meilleur, en fonction de ces traits caractéristiques indiquent que l'on peut espérer obtenir un taux de détection qui excède 90% avec un taux de fausses alarmes inférieur à 3%. (NC)

CRDV R-4172/80 (NON CLASSIFIE)

Bureau - Recherche et Développement, MDN, Canada.
CRDV, C.P. 880, Courcelette, Qué. G0A 1R0

"Evaluation d'une classe de segmenteurs pour images IR"
par L. Sévigny

Dans ce rapport nous évaluons l'efficacité de 6 algorithmes spécialisés dans la segmentation de cibles sur images IR. Ces algorithmes de segmentation, ou segmenteurs, reposent tous sur le principe voulant que la signature thermique d'une cible soit supérieure à celle de tout objet de l'arrière-plan. Les 3 premiers segmenteurs abordent l'image de front, en son entier, tandis que les 3 derniers incorporent la technique de redressement de l'arrière-plan (TRAP), visant à éliminer l'arrière-plan en tout ou en partie en le nivelant. Les divers algorithmes sont jugés d'après a) leur taux d'extraction, b) la fidélité du processus de segmentation en ce qui concerne les propriétés géométriques des cibles et, finalement, d'après c) le degré d'individualisation imprimé aux cibles extraites par rapport aux pseudo-cibles. Les 3 segmenteurs centrés sur TRAP ont une meilleure probabilité d'extraction que les trois autres qui essayent d'éliminer l'arrière-plan simplement en morcelant l'image. Par ailleurs, la plupart des segmenteurs considérés dans ce rapport altèrent d'une façon ou d'une autre la forme des cibles. Les deux exceptions à cette règle sont les segmenteurs No. 1 (générateur de silhouettes à seuil d'intensité unique) et No. 6 (le précédent segmenteur allié à une version particulière de TRAP). Les résultats expérimentaux montrent que l'intensité, le contraste et la variance sont les traits qui permettent le mieux de départager les cibles et les pseudo-cibles. Des expériences de classification réalisées à l'aide du segmenteur No. 6, lequel s'avère le meilleur, en fonction de ces traits caractéristiques indiquent que l'on peut espérer obtenir un taux de détection qui excède 90% avec un taux de fausses alarmes inférieur à 3%. (NC)

CRDV R-4172/80 (NON CLASSIFIE)

Bureau - Recherche et Développement, MDN, Canada.
CRDV, C.P. 880, Courcelette, Qué. G0A 1R0

"Evaluation d'une classe de segmenteurs pour images IR"
par L. Sévigny

Dans ce rapport nous évaluons l'efficacité de 6 algorithmes spécialisés dans la segmentation de cibles sur images IR. Ces algorithmes de segmentation, ou segmenteurs, reposent tous sur le principe voulant que la signature thermique d'une cible soit supérieure à celle de tout objet de l'arrière-plan. Les 3 premiers segmenteurs abordent l'image de front, en son entier, tandis que les 3 derniers incorporent la technique de redressement de l'arrière-plan (TRAP), visant à éliminer l'arrière-plan en tout ou en partie en le nivelant. Les divers algorithmes sont jugés d'après a) leur taux d'extraction, b) la fidélité du processus de segmentation en ce qui concerne les propriétés géométriques des cibles et, finalement, d'après c) le degré d'individualisation imprimé aux cibles extraites par rapport aux pseudo-cibles. Les 3 segmenteurs centrés sur TRAP ont une meilleure probabilité d'extraction que les trois autres qui essayent d'éliminer l'arrière-plan simplement en morcelant l'image. Par ailleurs, la plupart des segmenteurs considérés dans ce rapport altèrent d'une façon ou d'une autre la forme des cibles. Les deux exceptions à cette règle sont les segmenteurs No. 1 (générateur de silhouettes à seuil d'intensité unique) et No. 6 (le précédent segmenteur allié à une version particulière de TRAP). Les résultats expérimentaux montrent que l'intensité, le contraste et la variance sont les traits qui permettent le mieux de départager les cibles et les pseudo-cibles. Des expériences de classification réalisées à l'aide du segmenteur No. 6, lequel s'avère le meilleur, en fonction de ces traits caractéristiques indiquent que l'on peut espérer obtenir un taux de détection qui excède 90% avec un taux de fausses alarmes inférieur %. (NC)

DREV R-4172/80 (UNCLASSIFIED)

Research and Development Branch, DND, Canada.
DREV, P.O. Box 880, Courcelette, Que. G0A 1R0

"Evaluation of a Class of Segmenters for IR Imagery"
by L. Sévigny

This report presents an evaluation of the performance of 6 algorithms dedicated to segmentation of targets in IR imagery. These segmentation algorithms or segmenters are based on the single assumption that the targets display a larger thermal signature than the background. The first 3 segmenters deal with an image in its entirety, whereas the last 3 incorporate the Background Elimination Technique (BET), which aims at eliminating wholly or partly the background by levelling it. The segmenters are judged according to a) their extraction rate; b) the fidelity of the segmentation with respect to the geometrical properties of the extracted targets; and c) the degree of distinctiveness imparted to the extracted targets as opposed to the nontargets. The 3 segmenters relying on BET have a better extraction rate than the other 3 that try to cope with the background simply by partitioning the image. Most segmenters here distort in one way or another the shape of the targets. The two exceptions are segmenter No. 1 (Single Intensity Threshold Silhouette Generator or SIT Generator) and No. 6 (SIT Generator in conjunction with a particular version of BET). The experimental results show that the intensity, contrast and variance features are the most effective in discriminating the targets from the nontargets. The classification results one can expect from these features together with the segmenter that proves to be the best (segmenter No. 6) amount to a detection rate in excess of 90% with a false alarm rate not greater than 3%. (U)

DREV R-4172/80 (UNCLASSIFIED)

Research and Development Branch, DND, Canada.
DREV, P.O. Box 880, Courcelette, Que. G0A 1R0

"Evaluation of a Class of Segmenters for IR Imagery"
by L. Sévigny

This report presents an evaluation of the performance of 6 algorithms dedicated to segmentation of targets in IR imagery. These segmentation algorithms or segmenters are based on the single assumption that the targets display a larger thermal signature than the background. The first 3 segmenters deal with an image in its entirety, whereas the last 3 incorporate the Background Elimination Technique (BET), which aims at eliminating wholly or partly the background by levelling it. The segmenters are judged according to a) their extraction rate; b) the fidelity of the segmentation with respect to the geometrical properties of the extracted targets; and c) the degree of distinctiveness imparted to the extracted targets as opposed to the nontargets. The 3 segmenters relying on BET have a better extraction rate than the other 3 that try to cope with the background simply by partitioning the image. Most segmenters here distort in one way or another the shape of the targets. The two exceptions are segmenter No. 1 (Single Intensity Threshold Silhouette Generator or SIT Generator) and No. 6 (SIT Generator in conjunction with a particular version of BET). The experimental results show that the intensity, contrast and variance features are the most effective in discriminating the targets from the nontargets. The classification results one can expect from these features together with the segmenter that proves to be the best (segmenter No. 6) amount to a detection rate in excess of 90% with a false alarm rate not greater than 3%. (U)

DREV R-4172/80 (UNCLASSIFIED)

Research and Development Branch, DND, Canada.
DREV, P.O. Box 880, Courcelette, Que. G0A 1R0

"Evaluation of a Class of Segmenters for IR Imagery"
by L. Sévigny

This report presents an evaluation of the performance of 6 algorithms dedicated to segmentation of targets in IR imagery. These segmentation algorithms or segmenters are based on the single assumption that the targets display a larger thermal signature than the background. The first 3 segmenters deal with an image in its entirety, whereas the last 3 incorporate the Background Elimination Technique (BET), which aims at eliminating wholly or partly the background by levelling it. The segmenters are judged according to a) their extraction rate; b) the fidelity of the segmentation with respect to the geometrical properties of the extracted targets; and c) the degree of distinctiveness imparted to the extracted targets as opposed to the nontargets. The 3 segmenters relying on BET have a better extraction rate than the other 3 that try to cope with the background simply by partitioning the image. Most segmenters here distort in one way or another the shape of the targets. The two exceptions are segmenter No. 1 (Single Intensity Threshold Silhouette Generator or SIT Generator) and No. 6 (SIT Generator in conjunction with a particular version of BET). The experimental results show that the intensity, contrast and variance features are the most effective in discriminating the targets from the nontargets. The classification results one can expect from these features together with the segmenter that proves to be the best (segmenter No. 6) amount to a detection rate in excess of 90% with a false alarm rate not greater than 3%. (U)

DREV R-4172/80 (UNCLASSIFIED)

Research and Development Branch, DND, Canada.
DREV, P.O. Box 880, Courcelette, Que. G0A 1R0

"Evaluation of a Class of Segmenters for IR Imagery"
by L. Sévigny

This report presents an evaluation of the performance of 6 algorithms dedicated to segmentation of targets in IR imagery. These segmentation algorithms or segmenters are based on the single assumption that the targets display a larger thermal signature than the background. The first 3 segmenters deal with an image in its entirety, whereas the last 3 incorporate the Background Elimination Technique (BET), which aims at eliminating wholly or partly the background by levelling it. The segmenters are judged according to a) their extraction rate; b) the fidelity of the segmentation with respect to the geometrical properties of the extracted targets; and c) the degree of distinctiveness imparted to the extracted targets as opposed to the nontargets. The 3 segmenters relying on BET have a better extraction rate than the other 3 that try to cope with the background simply by partitioning the image. Most segmenters here distort in one way or another the shape of the targets. The two exceptions are segmenter No. 1 (Single Intensity Threshold Silhouette Generator or SIT Generator) and No. 6 (SIT Generator in conjunction with a particular version of BET). The experimental results show that the intensity, contrast and variance features are the most effective in discriminating the targets from the nontargets. The classification results one can expect from these features together with the segmenter that proves to be the best (segmenter No. 6) amount to a detection rate in excess of 90% with a false alarm rate not greater than 3%. (U)