Image Cover Sheet

TITLE

Image-based classification of sidescan sonar detections

This page is left blank

This page is left blank

Copy No. 25

DEFENCE DÉFENSE

IMAGE-BASED CLASSIFICATION OF SIDESCAN SONAR DETECTIONS

J.A. Fawcett
Defence Research Establishment Atlantic

Defence R&D Canada

Technical Memorandum
DREA TM 2001-026
April 2001

National Defence Défense nationale

Canada

Copy No:________

Image-based classification of sidescan sonar detections

John A. Fawcett
Defence Research Establishment Atlantic

Defence Research Establishment Atlantic

Technical Memorandum
DREA TM 2001-026
APRIL 2001

Author

John Fawcett

John A. Fawcett

Approved by

Ron H. Kuwahara

Head Electromagnetics Section

Approved for release by

K. Foster

Chair DRP

Abstract

In this report it is shown that the small images about target detections in sidescan sonar imagery can be used to classify the targets. In other words, the image itself is considered as a feature. This approach avoids the problems that traditional image-segmentation, feature-extraction and classification may entail. In particular, a situation where the background seabed imagery interferes with image-segmentation techniques is considered. The small images are analyzed using principal component decomposition, discriminate analysis and clustering. As well, a commercial software package is also used.

Résumé

Le présent rapport montre que de petites images relatives à des détections de cibles en imagerie sonar à balayage latéral peuvent être utilisées pour classer les cibles. Cette façon de procéder permet de contourner les problèmes que peuvent entraîner la segmentation d'image, l'extraction de caractéristiques et la classification. On traite, en particulier, d'une situation dans laquelle l'imagerie du fond marin d'arrière-plan nuit aux techniques de segmentation d'image. Les petites images sont analysées au moyen de la décomposition des éléments principaux, de l'analyse discriminante et du regroupement. De plus, on utilise un progiciel commercial.

This page intentionally left blank.

Executive summary

Background:

The ability to classify detections in Mine-Countermeasures as manmade or non-manmade is very important. Otherwise, the time spent investigating detections can quickly become prohibitive. Once a detection has been made on a sidescan sonar image, then typically a small image containing the detection is extracted (mugshot). Traditionally, this mugshot is then segmented into highlight/shadow and background regions and various features of the highlight and shadow computed. For example, the area of the shadow, the shadow's dimensions, the area of the highlight, and many more. These features are then fed into some classification algorithm. However image-segmentation, especially for complex seabed backgrounds, can be a difficult process and any resulting classification is only as good as the previous segmentation and feature extraction algorithms. In this report, the possibility of using the mugshot itself, without any segmentation and feature extraction, as the feature for classification is examined.

Principal Results:

It is shown that it is possible to obtain a good-classifier based upon imagery alone The classifier performs very well, even in the case of a significantly rippled seabed background

Future Research:

The concepts described in this report should be investigated with real sidescan sonar imagery for a wide-variety of background and operating conditions

Fawcett, J A., 2001, Image-based classification of sidescan sonar detections, DREA TM 2001-026, Defence Research Establishment Atlantic

Sommaire

Contexte

Dans les mesures de lutte contre les mines, il est très important de pouvoir classer les cibles détectées en cibles artificielles et en cibles non artificielles, sans quoi le temps consacré à l'étude des cibles détectées peut rapidement devenir excessif Une fois qu'une cible a été détectée sur une image sonar obtenue par balayage latéral, on extrait généralement une petite image contenant cette cible (photo d'identité). Habituellement, cette photo d'identité est ensuite segmentée en zones brillantes/sombres et en zones d'arrière-plan, et diverses caractéristiques des zones brillantes et sombres font l'objet de calculs, par exemple la superficie de la zone sombre, les dimensions de la zone sombre et la superficie de la zone brillante. Les valeurs obtenues sont ensuite passées à un algorithme de classification. Cependant, la segmentation d'image, en particulier dans le cas des fonds marins d'arrière-plan complexes, peut être une opération difficile et la qualité de toute classification résultante ne peut être supérieure à la qualité offerte par les algorithmes précédents de segmentation et d'extraction de caractéristiques. Dans ce rapport, on étudie la possibilité d'utiliser la photo d'identité elle-même, sans segmentation ni extraction de caractéristiques, pour la classification. L'avantage de cette méthode est qu'elle permet de contourner le processus complet de segmentation/extraction de caractéristiques.

Principaux résultats

On montre qu'il est possible d'obtenir un bon classificateur basé sur l'imagerie seulement. Le classificateur est très efficace, même dans le cas d'un fond marin d'arrière-plan très ondulé

Recherche future

Les concepts décrits dans ce rapport devraient être étudiés en imagerie sonar réelle à balayage latéral pour une grande variété de conditions d'arrière-plan et de conditions d'utilisation.

Fawcett, J.A., 2001, Image-based classification of sidescan sonar detections, CRDA TM 2001-026, Centre de recherches pour la défense, Atlantique

Table of contents

Abstract i

Résumé i

Executive summary iii

Sommaire iv

Table of contents v

List of figures vi

INTRODUCTION 1

THEORY 2

Principal Component Decomposition 2

Discriminant Analysis 3

Classifiers 4

NUMERICAL EXAMPLES 6

Summary and discussion of results 18

References 19

List of figures

Figure 1 Eight representative images from Manta-class 7

Figure 2 Eight representative images from cylinder-class 7

Figure 3 Eight representative images from rock-class 8

Figure 4 Mean image from 303 images 9

Figure 5 Variation of singular values of training set; dashed line indicates the mean value..... 9

Figure 6 First 4 principal components interpreted as images 10

Figure 7 Clusters (blue is manta, red is rocks, green is cylinders) for training set 10

Figure 8 Clusters for testing set 11

Figure 9 Values of test set output coordinate 1 (blue), 2 (red), 3 (green) from HNET.... 12

Figure 10 Example images from training set, 3 "mantas", 2 rocks and 3 cylinders 13

Figure 11 First discriminating image for noisy training set 14

Figure 12 Second discriminating image for noisy training set 14

Figure 13 Clusters for noisy training set 15

Figure 14 Clusters for noisy testing set 15

Figure 15 The six misclassified images from cluster analysis 16

Figure 16 Normalized clustering for noisy training set 17

Figure 17 Normalized clustering for testing set 17

This page intentionally left blank

INTRODUCTION

In computer-aided classification of sidescan sonar images, various "features" of the shadow and perhaps highlight of the image are used as input to a statistical or neural-network classifier. For example, after an initial rough detection of a mine-like object has been made, a small image (mugshot) about the detection centre can be extracted. A detailed shadow/highlight segmentation (see, for example [1-3] for a discussion of possible segmentation algorithms) can then be performed yielding (hopefully) accurate shadow and highlight regions. Various features, such as elliptical parameters for the shadow and highlight, ratios of perimeter/area, Fourier descriptors, etc can be computed and used in clustering, discriminant, and neural network classification. One problem with this approach is that the performance of the classifiers depend to a significant extent upon the earlier operations of image segmentation and feature extraction. Thus, for example, if the computed features are very inaccurate due to, say, a seabed background with ripples or interfering shadows, then the output of the classifier will probably also be in error. In reference [2] a hybrid approach is proposed (i.e., a joint classification/segmentation approach) In this report the possibility of using the images themselves as the feature is considered. This alleviates the need for any prior segmentation or feature extraction. It is assumed, however, that this mugshot has been obtained on the basis of some prior detection method.

In order to evaluate our approach a synthetic data set was generated using a MATLAB simulation code [4]. Three classes (manta-like, cylinder, and rock) were used and 151 images (64x32 pixels) of each were computed. These were generated with a random orientation, a somewhat random centring of the target within the mugshot, and random parameters (within a specified range) for objects. The so-called rock class was generated by simply allowing its description to be more random than the other two classes. This, in general, resulted in a highlight/shadow structure which was more irregular than the other 2 classes, although there are cases, where the rock-class image appears like one of the other two classes. The background in these initial images was simply set to a constant but we will examine the effect of different types of background, by superimposing different backgrounds.

The basic input to the classifiers considered is simply the mugshot itself. The data set of images will be broken down into training (101 images for each class) and testing (51 images for each class); a principal component analysis will, in fact, reduce the dimensionality of the input data. The basic approach we use is to apply a principal-component analysis and then use a discriminate analysis to determine the vectors of principal component coefficients which best discriminate the object-classes. These vectors are then used to cluster the images Simple classification methods can then be used. Another approach which will be examined is to simply input the set of images to the commercial software package HNET[5].

THEORY

Principal Component Decomposition

We will consider the image to be in two basic forms: (1) as a matrix of pixels **X** or (2) re-arranged as a linear vector **x**. Let us denote the mean image vector as μ. Then the covariance matrix of the image vector can be written as

$$R = \sum_{i} (x_i - \mu)^T (x_i - \mu) \qquad (1)$$

The principal components are the simply defined as the singular vectors corresponding to the leading singular values of the system. Recall that R is a symmetric, positive definite (or possibly singular) matrix. Numerically, we will perform the singular-value decomposition (SVD) of the matrix R to obtain

$$R = V^T S^2 V \qquad (2)$$

Then, for example, if we take the leading 5 singular values (from the diagonal matrix S) and the corresponding columns from the matrix U we have a 5-term principal component expansion. By summing up the squares of all the singular values and taking the ratio of the partial sum of squares to the full sum, we get an estimate of the fraction of the total variance represented by a specified number of components.

In fact, in the following work, we don't explicitly form the covariance matrix; instead if there are M images in the training set, each with N pixels then this set of data can be written as the MxN matrix **A**. The mean image of this matrix is computed and subtracted from each of the images of **A** and then we take the SVD of this matrix obtaining

$$A = USV^T \qquad (3)$$

where **S** and **V** are as above. In fact, even if we retain all the coefficients of the principal components of A corresponding to non-zero singular values, the dimensionality of the feature space has been reduced from N to M, which can be a significant reduction if the number of training images is much less than the number of image pixels (and this is the case in this report)

Discriminant Analysis

Given a set of labelled images or sets of features, the question arises as to what are the best combination of features to discriminate between the classes. In the principal component analysis described above, we sought the eigenvectors which accounted for most of the variation of the entire data set; however, this analysis does not take into account the known classes of the object. In this case, for the training set there is a known classification for each of the images or feature sets. We now seek the vectors which minimize the variation within the groups, but maximize the variation between the groups . In particular, denote the variance between the groups, for a particular set of feature vectors as

$$R_m = \frac{1}{L}\Sigma(\mu_i - m)^T(\mu_i - m) \tag{4}$$

and the mean within group variance as

$$R_w = \frac{1}{L}\sum_{i=1}^{L} E(x-\mu_i)^T(x-\mu_i). \tag{5}$$

There are various measures based upon these 2 matrices which could be used to measure the performance of various feature vectors.. A common choice [6] is to find those combinations of features which maximize

$$J = trace(R^{-1}{}_w R_m). \tag{6}$$

This effectively minimizes inner class variation while maximizing inter mean variation The largest (L-1) eigenvectors of

$$R_w^{-1} R_m \tag{7}$$

(where the matrices now refer to the original set of features) can be shown to be the combination of features which maximize J.

In our case, one could calculate these ideal features from the images themselves or one could first reduce the space by performing a principal component analysis and then performing a discriminant analysis based on the principal component coefficients. This is the approach that is taken below.

Classifiers

Above we discussed two possible mappings of features into a new feature space. However, once the features have been settled upon there are a number of classifiers which can be applied to a new set of features for classification. Below, we discuss a few possible options.

Linear Classification

Let us suppose that there are a number, N, of distinct classes of images (or their principal component coefficients) and that the subset of images which are used for training are labelled. Further, let us assume that in feature space, the classes have a distribution about their respective class mean which can be approximated by a normal distribution. Each of these distributions would, in general, have a different covariance matrix. However, if it is assumed that they each have the same covariance matrix, say the mean within-group covarince, then some simplifying mathematics occurs. In particular, if a vector of features **x** is to be classified then its normalized distance from the i th mean μ_i is

$$d_i = (x - \mu_i)^T R_m^{-1} (x - \mu_i) \qquad (8)$$

where R_m is the mean covariance matrix.
Expanding this out, it is found that

$$d_i = x^T R_m^{-1} x - 2\mu_i^T R_m^{-1} x + \mu_i^T \mu_i \qquad (9)$$

From this it can be seen that the first term does not depend upon the class, so that in comparing the "distance" of X from the various class means, we need consider only the linear functional,

$$\ell_i = -2\mu_i^T R_m^{-1} x + \mu_i^T \mu_i \qquad (10)$$

Another, even, simpler classifier results if the covariance matrix is ignored and one simply looks for the closest mean.

Quadratic Classification

The analysis for this classifier is identical to that of the linear classifier, except now the individual covariance matrices are used in the definition of the normalized distance, so that the distance

$$d_i = x^T R_i^{-1} x - 2\mu_I^T R_i^{-1} x + \mu_i^T \mu_i + \ln |\det(R_i)| \tag{11}$$

is now used

Although it would appear that the quadratic classifier should be more accurate, there are two mitigating factors (1) The underlying assumption of normal distributions may be incorrect in which case the linear classifier may be more robust (2) the statistical estimation of the individual within-group covariance matrices are erroneous, so that, once again, the linear estimator (based upon using the mean within-group covariance) may be more robust

HNET Classifier

The two above classifiers have assumed that the distributions of the features around the class means are fairly regular and unimodal in nature (i..e, there is one distinct class-centre for each class) Other classifiers attempt to model more complicated feature space behaviour There are higher-order polynomial approximations to the probability distribution, Parzen classifiers, neural networks, etc. In this report, we chose to use the Holographic Neural Network (HNET[5]) package. It is an estimator which attempts to find a set of linear weights which best match the input features to the desired outputs. More importantly it also iterates over the input features and various orders of cross-products of the input features to determine which inputs are important in the resulting output predictions for the training set. There are also a variety of conversions (Fourier, wavelet) and normalizations, and also various parametrizations, amplitude, phase, etc, of the input which can be used It is also possible to specify various ranges of parameters to the program and allow the HNET software itself to optimize the settings.

It should be pointed out that we are only using the "Supervised Learning Platform" of the HNET package – the entire package allows one to build much more complicated learning structures if desired.

The HNET learning platform is, in fact, more a sophisticated interpolator than strictly a classifier. Thus, for example, if the class-labels are 1,2,3 then HNET will provide an interpolator in feature space based upon the training set to predict the class-label at any other combination of input-features; the output will, hopefully be close to 1,2 or 3 in order to make a classification but could, in fact, be 2.47 for example. One possible problem with this approach is that it seems to be label-dependent; For example, will one obtain the same results if the labels on the first 2 classes are switched? In out trials, the answer appears to be no, particularly for those values which are not well determined.

However, one can allow for a vector of output responses. Thus instead of using labels 1,2,3 one can use labels (1,0,0) for class 1, (0,1,0) for class 2 and (0,0,1) for class 3. It was found that this was a useful way of making the interpolation/classification process label-independent.

NUMERICAL EXAMPLES

In order to test out various image classifiers, 3 classes of images (153 in each class were generated). The 3 basic classes were truncated-cone class (loosely representing Manta-shaped mines), cylinder class, and a rock class. These images were generated for a flat seabed using a ray-trace code [4]. In summary, this ray-tracing code discretizes the object into triangular facets. From discrete points along the seabed, corresponding to a beamwidth strip from the along-track position of the sidescan sonar, rays are projected to the sonar. If they intersect the object, the amplitude and travel time of the ray corresponds to the closest of the intersection points on the object. This code can also generate the response from structured (for example, rippled) seabeds, however, this takes more computation time. Hence for these examples, where many images needed to be generated, only a flat seabed was used in the numerical generation. Synthetic backgrounds which roughly correspond to a rippled seabed will simply be superimposed

The target shapes and orientations were generated somewhat randomly, the basic resolution of the images is taken to be 10x10 cm – this corresponds to a Klein 5500 smallest along-track resolution and roughly to the SQS-511 sonar's high resolution of 12 5 cm For the truncated-cone class, the target had a bottom base diameter which varied randomly between 0.8 to 1.2 m, a height which varied between 0.4 and 0.6m, and an upper diameter between 0.3 and 0.7 m. Due to the azimuthal symmetry of this target there was no aspect variation. The cylinder had a length varying from 1.0 m to 1.6 m and a diameter of 0.4 to 0.6 m. The cylinder is lying horizontally on the seabed; its aspect is allowed to vary randomly between -45 and +45 degrees. Finally, the rock class consists of objects whose length varies from 0 5 to 1 1 m. The length of the rock is broken up into 7 segments; At each of these segment boundaries 6 angles are defined and each of these points has a random radius varying from 0 2m to 0.3 m All these points together define a triangulation of the object. The aspect of the rock is randomly varied from –45 to +45 degrees. By its construction the profile of the rock is more irregular than the Manta or cylinder class. However, there is a chance that some of the rocks will appear to be either Manta-like or cylinder-like. In reality, there would probably be less variation than this, particularly for the "Manta-class" where the parameters are fairly precise; however, we wished to define fairly generic classes to see if the classifier could classify when the various dimensions are not precise. In Figs.1-3, 8 representative mugshots for each class are shown A small amount of pre-processing is applied to the images: first the highlight region is clipped at 3 standard deviations from the mean – this makes the highlight region more uniform, we wished to avoid the situation where the classification is triggering on a few very bright pixesl; in realistic sidescan sceanarios the highlight region is often rather diffuse and unpredictable. Second the 64 x 32 mugshot is Fourier-transformed; two phase factors are defined,

$$\psi_x = F(\Delta k_x, 0) / |F(\Delta k_x, 0)|, \qquad \psi_y = F(0, \Delta k_y) / |F(0, \Delta k_y)| \qquad (12)$$

The inverse Fourier transform is then defined

$$I(x,y) = \sum_{-N/2}^{N/2-1} \sum_{-M/2}^{M/2-1} F(n\Delta k_x, j\Delta k_y)\psi_x^{-n}\psi_y^{-j} w(n,j)\exp(-i(\,n\Delta k_x x + j\Delta k_y y)\,) \qquad (13)$$

where w(i,j) is a low-pass filter The multiplication by the powers of the phase-factors is analagous to the one used in Fourier descriptors [2] and is meant to make the images shift-invariant

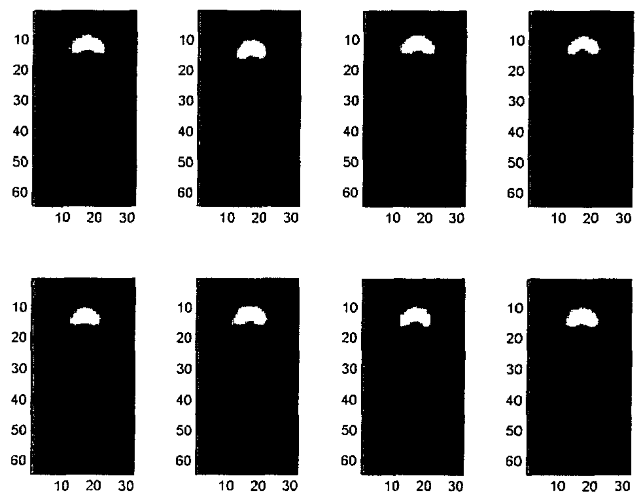

Figure 1 Eight representative images from Manta-class

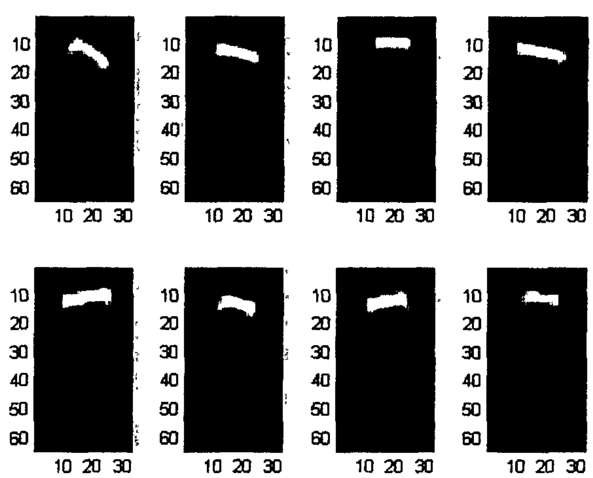

Figure 2 Eight representative images from cylinder-class

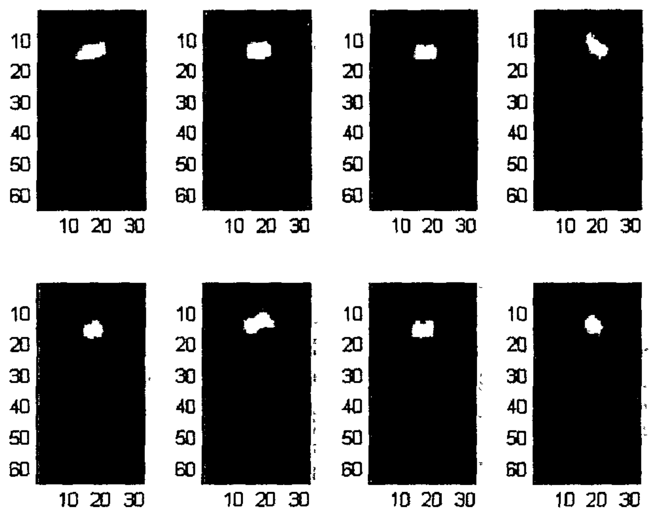

Figure 3 Eight representative images from rock-class

The images in these 3 figures can for the most part be distinguished visually; however, it is clear that there are "rocks" which had some similarity in appearance to the "mantas" and some to the "cylinders" The "mean image" over the first 303 images is shown below (we conside the first 303 images as this will constitute our "training" set)

From the first 303 images, we can consider the 303X2048 matrix constructed from the 303 image vectors, with the mean image subtracted off. Taking the singular value decomposition of this matrix then yields the singular values and the corresponding principal components. In Fig 5 the values of the singular values (squared) in descending order is shown As well the mean value of the singular values is shown; from this it can be that by about singular value 25 the mean values has been reached.

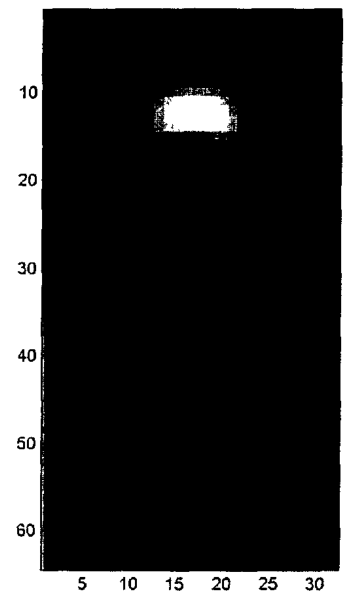

Figure 4 Mean image from 303 images

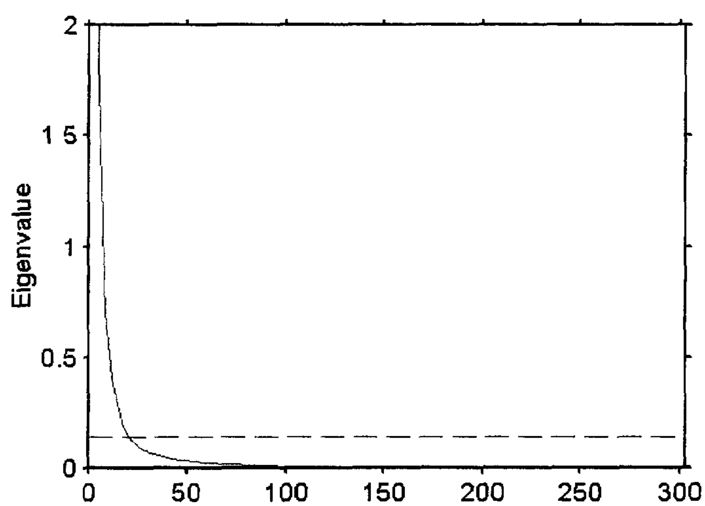

Figure 5 Variation of singular values of training set; dashed line indicates the mean value.

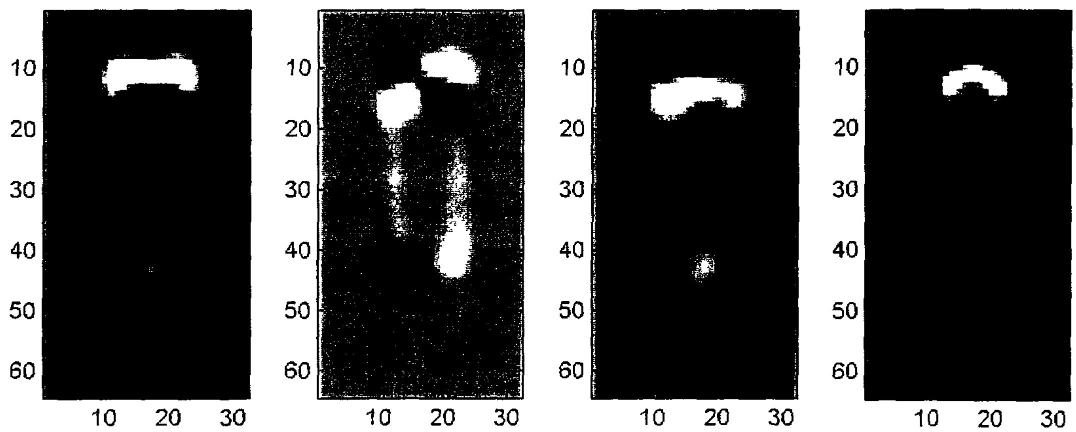

Figure 6 First 4 principal components interpreted as images

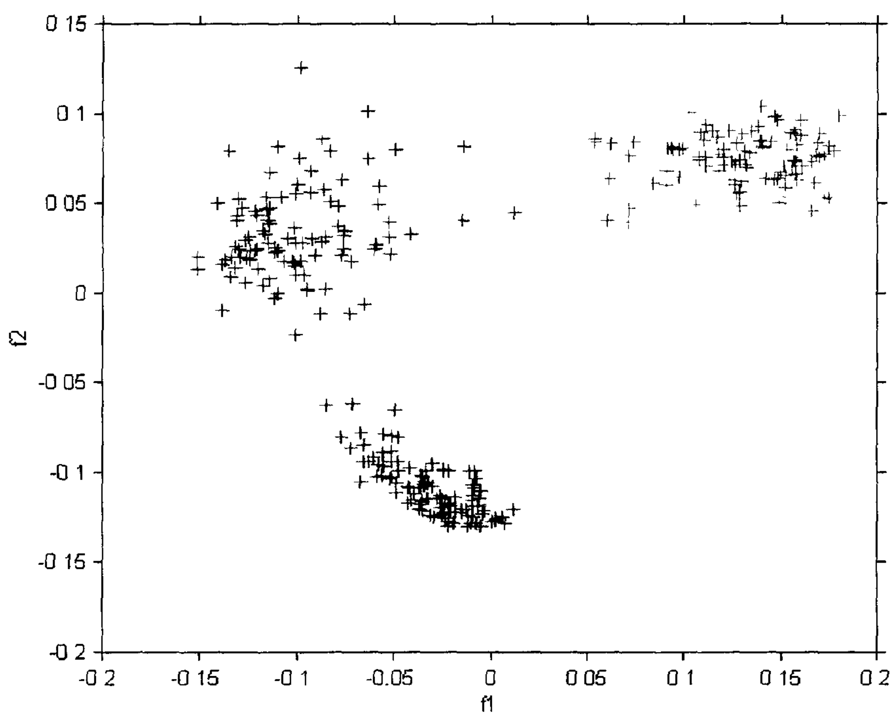

Figure 7 Clusters (blue is manta, red is rocks, green is cylinders) for training set

There remains the question as to which are the best principal components to use for classification purposes. One way to reolve this is to use the discriminant analysis discussed before – we will seek the combination of principal components which minimize the ratio of inner class variance with respect to intra-class variance For example, taking 20 principal components we can find the 2 vector combinations of these coefficients which yield the best discrimination. The resulting clustering for the "training" and "testing" sets are shown in Figs. 7 and 8, where f_1 and f_2 are the projections of the images onto these 2 vectors.

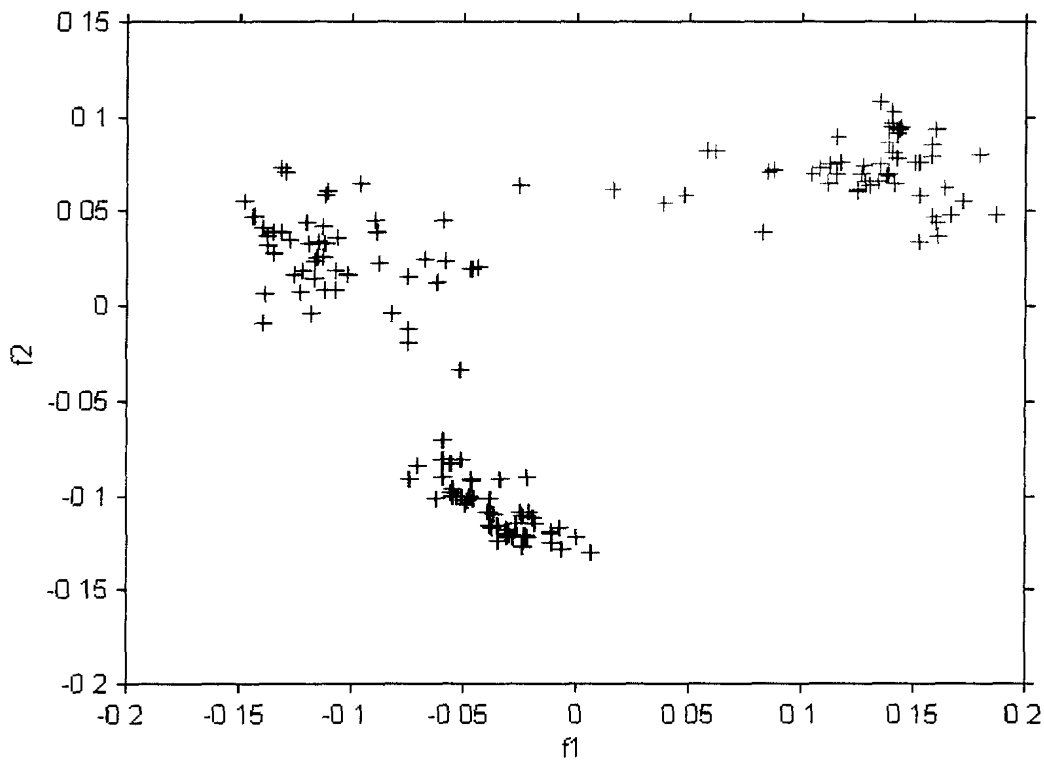

Figure 8 Clusters for testing set

It can be seen that there is very good cluster discrimination and using a linear classifier 100% results are obtained for both the training and testing classes.

Using HNET with 248 cortical memory elements on the image itself (after the pre-processing on the raw image) the following outputs are obtained for the "testing" class. Here it can be seen that if we choose the maximum number as the class 100% results are obtained. Thus in this noiseless case both the clustering (after discriminant analysis) and HNET worked very well.

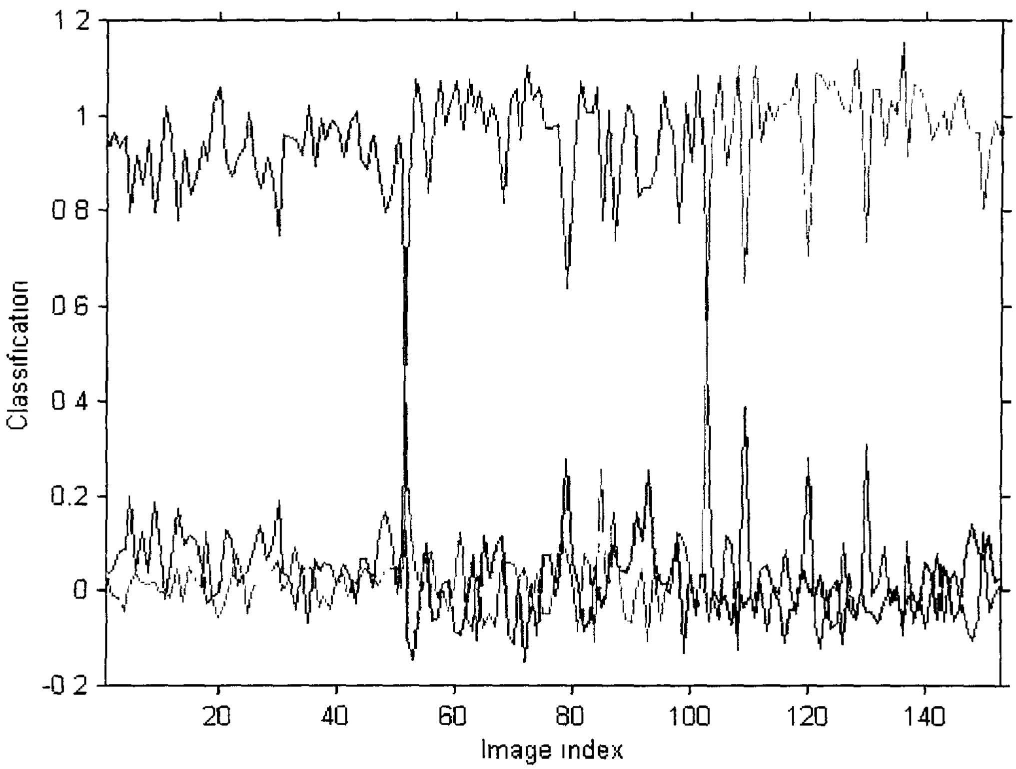

Figure 9 Values of test set output coordinate 1 (blue), 2 (red), 3 (green) from HNET

We now move onto a harder case – rippled, noisy backgrounds will be added to the original images. In this case, the same target images as above were used; however, the uniform background for this case is replaced by a sinusoidally-varying one. The orientation of this background is generated randomly Representative images from this new set are shown in Fig. 10. As before, we use the first 303 images as the "training" set. This time 160 of the principal components are used and the discriminating combinations of these vectors are found. As before, the 2 vectors of the coefficients which provide the best discrimination are computed. These vectors of coefficients can be interpreted in terms of images by multiplying them by the corresponding eigenvector matrix. The discriminating images are shown below in Figs.11 and 12 As one would expect most of the large discriminating pixels are concentrated in the regions of the edges of the highlights and shadows of the objects, most of the pixels in the rest of the image are close to zero.

The resulting clusters for the training and testing sets are shown in Figs. 13 and 14 In Fig.13 a fairly tight clustering is obtained, however, for the testing set (Fig.14), the clustering is more diffuse and it can be seen that there will be some ambiguous classifications near the overlap regions of these distributions; using a linear classifier, there are, in fact, 6 incorrect classifications. The program HNET was also run on this set, producing a total of 8 incorrect classifications; This should not be deemed significant as it is obvious that there is an area where the distributions overlap where a number of the classifications are uncertain and, in practice, a classifier should flag a whole region of overlap for operator analysis However, it is promising that for this difficult image set with very challenging backgrounds, there were only 6-8 incorrect classifications The six incorrect classifications from the cluster analysis are shown below in Fig.15 For the first 5 of these images, it is not surprising that a classifier would have problems, in these cases, the target highlight is merging with the ripple highlight and the target shadow is merging with the seabed shadow. The sixth misclassification, cylinder as a rock, is somewhat surprising, as certainly to the eye it appears to be a cylinder. It should be re-emphasized that since all these misclassifications and the some of the correct classifications are from an area where the clusters are overlapping they should be flagged for operator intervention.

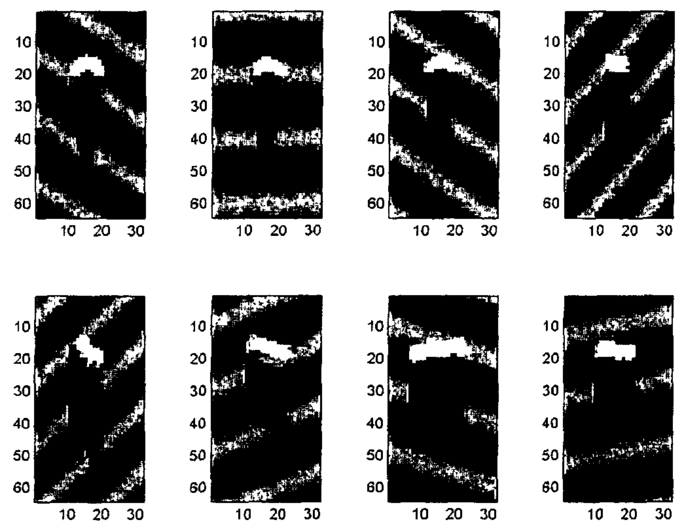

Figure 10 Example images from training set, 3 "mantas", 2 rocks and 3 cylinders

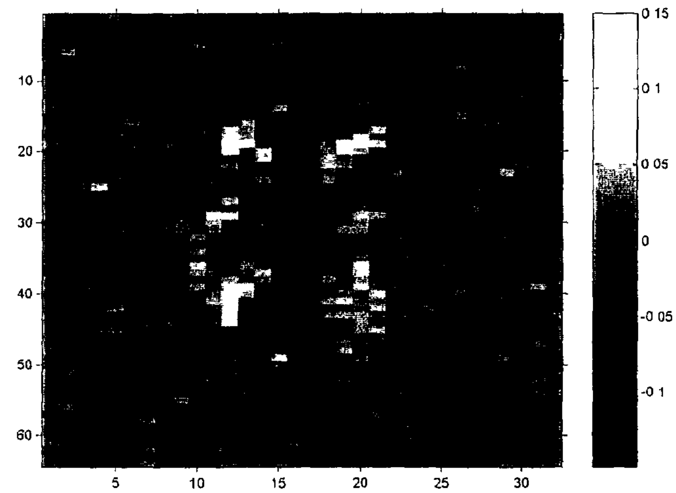

Figure 11 First discriminating image for noisy training set

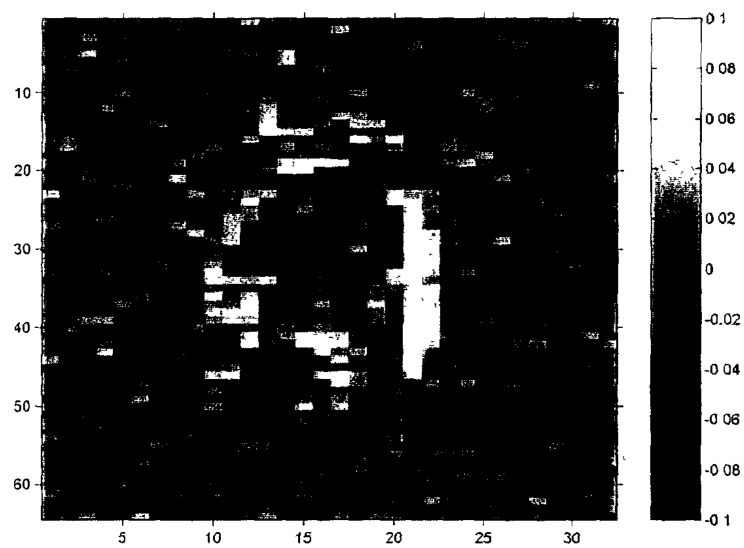

Figure 12 Second discriminating image for noisy training set

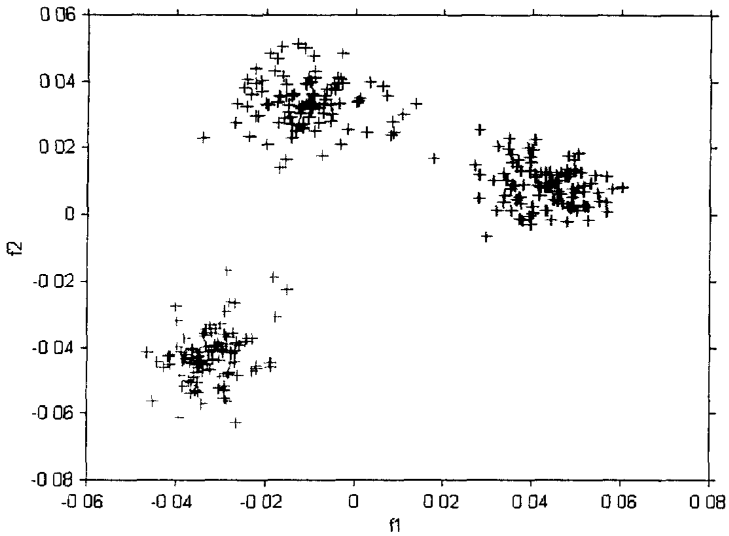

Figure 13 Clusters for noisy training set

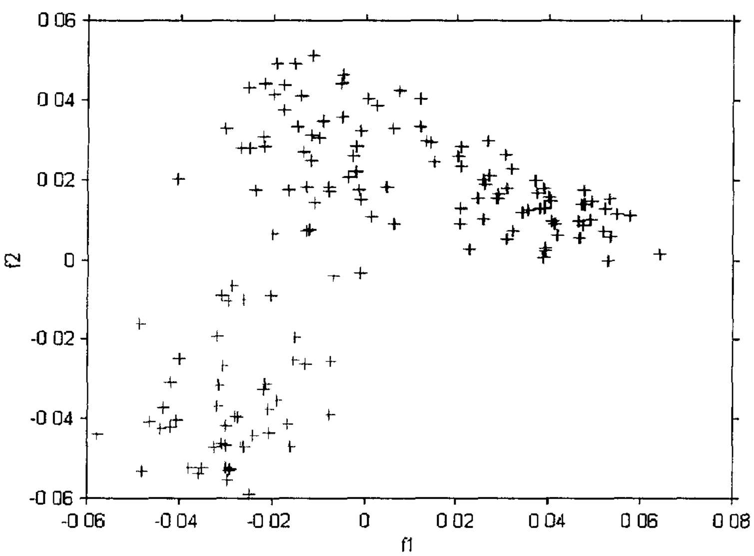

Figure 14 Clusters for noisy testing set

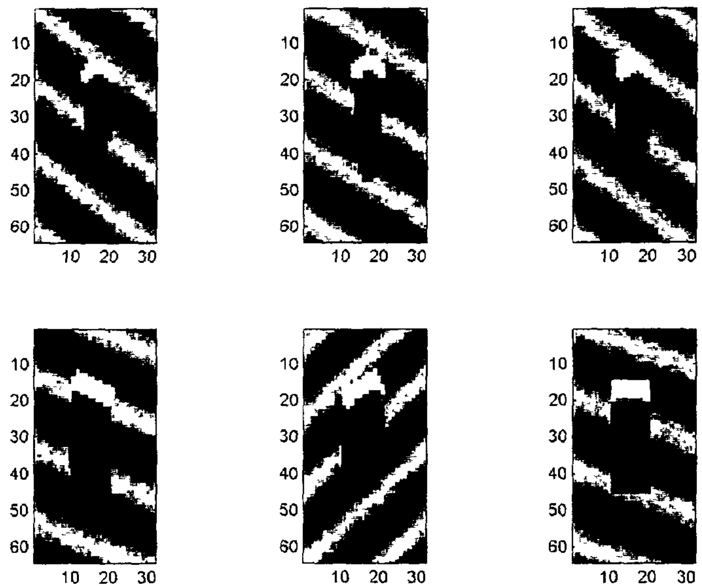

Figure 15 The six misclassified images from cluster analysis

We now consider the question: suppose we train under one set of seabed conditions (uniform background) Can we classify with the other seabed type (sinusoidal) and vise versa? We first train with 303 images from the homogeneous background and then test on the rippled background. The classification results are not shown, they are very poor. This is probably not surprising; the linear relationships which are determined on the basis of a homogeneous background to optimally discriminate between targets do not work when the images have a significant and interfering background. On the other hand, if the training set has difficult backgrounds, perhaps the determined discriminators will work well for another background (the homogeneous case) This is found to be the case. Using 160 principal components for the training set and then finding the optimal discriminators, the resulting cluster for the training (and this corresponds to the clustering for the rippled seabed example) and the test set (homogeneous background) are shown below in Figs.16 and 17 This example illustrates another problem; because the background statistics are significantly different between the training and test set, the statistics of the clusters are not the same In order to try to equate the two cluster sets, we normalize each by their standard deviations in the f1 and f2 coordinate directions, respectively. In this way, the 2 distributions correspond fairly closely; however, close examination reveals that the mean of the clusters are slightly offset from each other; this implies that a classification scheme base upon the mean-values from the training set may not work

It can be seen that there is significant clustering for this case; the means of these clusters are close to those of the training set but are different. This indicates that if the means from the training set are used for classification in the test set there maybe some misclassifications due to the difference in the means A possible correction to this situation can be obtained by reclustering the new data, to do this, we assume means with the values from the training set–

cluster (nearest-mean) according to these means, recompute the means, recluster, etc In this manner the initial means.

(1.481,.339), (- 434,1 157), (-1 04,-1 496) become (1 277,.715),(- 116,.708), and (-1 36,-1.66)

and there are a total of 11 misclassifications We also ran Hnet on this case and obtained 10 misclassifications, many of which are the same ones from the cluster-analysis

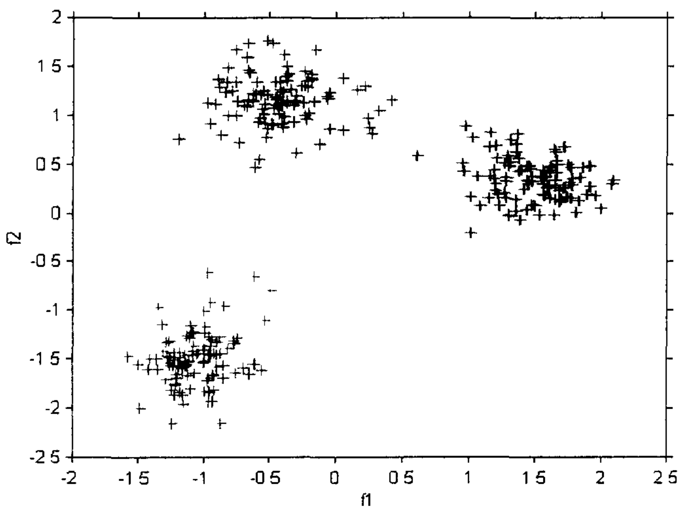

Figure 16 Normalized clustering for noisy training set

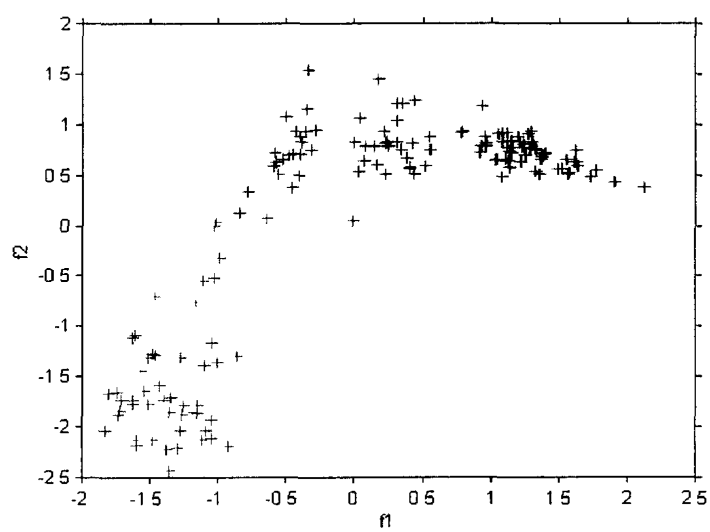

Figure 17 Normalized clustering for testing set

Summary and discussion of results

We have shown in this report that is feasible to use the images themselves for classification It may be that this is not optimal in cases where feature extraction can be accurately used. However, the advantage of this technique is that it avoids the usual steps of highlight/shadow segmentation and feature extraction These processes may be difficult for noisy and structured backgrounds. However, in order to recognize targets against complex backgrounds it seems necessary to train the image classifier with some of these complicated backgrounds. When trained with a wide variety of target aspects; it seemed that the classifier could then work with varying target aspects

Traditional discrimination techniques and a commercial package HNET were used for classification, the 2 methods both gave very similar results The clustering display indicated that there were some images for which a clear-cut classification could not be made. These images should be flagged to the operator.

The examples of this paper would be challenging to any classification method. A wide variation of size parameters and aspects were used and the ripply -background provided an interfering background to overcome.

It may also be possible to further improve the classification methods Additional image-based features could be generated – variances, edge-maps, etc and fed in as additional features to improve classification. In the case of rippled- seabeds some preliminary Fourier, wavelet or predictive filtering could perhaps be used to improve image quality.

References

1 V. Myers, "Image segmentation using iteration and fuzzy logic" DREA TM 2001-010.

2. I. Quidu, J.Ph.Malkasee, G. Burel, and P. Vilbe, "Mine classification based on raw sonar data: an approach combining Fourier descriptors, statistical models, and genetic algorithms", Proceedings of Oceans 2000, Providence , R I, 2001

3. M Kass, A. Witkin and D. Terzopoulos, "Snakes: active contour models", International Journal of Computer Vision, pp.321-331, 1988.

4. J.A. Fawcett, "Simulation of sidescan sonar images using MATLAB", DREA TN 1999-074, 1999

5. HneT2000 Application Development System, Version 1.3, User's Guide, AND Artificial Neural Devices Corporation, 2000

6. K. Fukunaga, *Introduction to statistical pattern recognition.* Academic Press, 1990

UNCLASSIFIED
SECURITY CLASSIFICATION OF FORM
(highest classification of Title, Abstract, Keywords)

DOCUMENT CONTROL DATA

(Security classification of title, body of abstract and indexing annotation must be entered when the overall document is classified)

1 ORIGINATOR (the name and address of the organization preparing the document Organizations for whom the document was prepared, e g Establishment sponsoring a contractor's report, or tasking agency, are entered in section 8)

Defence Research Establishment Atlantic
PO Box 1012
Dartmouth, Nova Scotia, Canada B2Y 3Z7

2 SECURITY CLASSIFICATION (overall security classification of the document including special warning terms if applicable)

UNCLASSIFIED

3 TITLE (the complete document title as indicated on the title page Its classification should be indicated by the appropriate abbreviation (S,C,R or U) in parentheses after the title)

Image-based classification of sidescan sonar detections

4 AUTHORS (Last name, first name, middle initial If military, show rank, e g Doe, Maj John E)

John A. Fawcett

5 DATE OF PUBLICATION (month and year of publication of document)

April 25, 2001

6a NO OF PAGES (total containing information Include Annexes, Appendices, etc)

26 (approx.)

6b NO OF REFS (total cited in document)

6

7 DESCRIPTIVE NOTES (the category of the document, e g technical report, technical note or memorandum If appropriate, enter the type of report, e g interim, progress, summary, annual or final Give the inclusive dates when a specific reporting period is covered)

TECHNICAL MEMORANDUM

8 SPONSORING ACTIVITY (the name of the department project office or laboratory sponsoring the research and development Include address)

Defence Research Establishment Atlantic

9a PROJECT OR GRANT NO (if appropriate, the applicable research and development project or grant number under which the document was written Please specify whether project or grant)

Project 1da12

9b CONTRACT NO (if appropriate, the applicable number under which the document was written)

10a ORIGINATOR'S DOCUMENT NUMBER (the official document number by which the document is identified by the originating activity This number must be unique to this document)

DREA TM 2001-026

10b OTHER DOCUMENT NOs (Any other numbers which may be assigned this document either by the originator or by the sponsor)

11 DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification)

(X) Unlimited distribution
(   ) Defence departments and defence contractors; further distribution only as approved
(   ) Defence departments and Canadian defence contractors; further distribution only as approved
(   ) Government departments and agencies; further distribution only as approved
(   ) Defence departments; further distribution only as approved
(   ) Other (please specify):

12 DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document This will normally correspond to the Document Availability (11) However, where further distribution (beyond the audience specified in (11) is possible, a wider announcement audience may be selected)

UNCLASSIFIED
SECURITY CLASSIFICATION OF FORM

UNCLASSIFIED
SECURITY CLASSIFICATION OF FORM
(highest classification of Title, Abstract, Keywords)

13. ABSTRACT (a brief and factual summary of the document It may also appear elsewhere in the body of the document itself It is highly desirable that the abstract of classified documents be unclassified Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U) It is not necessary to include here abstracts in both official languages unless the text is bilingual)

In this report it is shown that the small images about target detections in sidescan sonar imagery can be used to classify the targets. In other words, the image itself is considered as a feature. This approach avoids the problems that traditional image-segmentation, feature-extraction and classification may entail. In particular, a situation where the background seabed imagery interferes with image-segmentation techniques is considered. The small images are analyzed using principal component decomposition, discriminate analysis and clustering. As well, a commercial software package is also used.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document They should be selected so that no security classification is required Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included If possible keywords should be selected from a published thesaurus e g Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified If it not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title)

Classification, discriminant analysis, principal component analysis, clustering

UNCLASSIFIED
SECURITY CLASSIFICATION OF FORM

Defence R&D Canada
is the national authority for providing
Science and Technology (S&T) leadership
in the advancement and maintenance
of Canada's defence capabilities.

R et D pour la défense Canada
est responsable, au niveau national, pour
les sciences et la technologie (S et T)
au service de l'avancement et du maintien des
capacités de défense du Canada.

DEFENCE DÉFENSE

www.drdc-rddc.dnd.ca