MODEL BASED CLASSIFICATION – A POSSIBLE APPROACH FOR TARGET CLASSIFICATION IN SAS AND SAR?

JM Bell Oceans Systems Group, Heriot-Watt University, Riccarton, Edinburgh, EH14 4AS
J.Bell@hw.ac.uk
Y Petillot Oceans Systems Group, Heriot-Watt University, Riccarton, Edinburgh, EH14 4AS
Y.R.Petillot@hw.ac.uk
S Reed Seebyte Ltd, 6a Canaan Lane, Edinburgh, EH10 4SY, scott.reed@seebyte.com

1 INTRODUCTION

With the increasing deployment of Synthetic Aperture systems for sonar and radar the computer aided detection and classification of targets in their imagery is an important issue. The majority of existing techniques employ supervised classification systems which are reliant on training data. The success of these systems can be highly dependant on the similarity of the test data to the training data, which includes the effect of the background region on which the target was located.

A model based technique for the automated detection and classification of objects in sidescan sonar imagery has been developed. This employs a three stage process to attempt to overcome the limitations of trained systems. The first stage detects possible mine like objects before extracting the highlight and shadow regions from the detected objects. The final stage is the object classification which iteratively compares the shadows to synthetic shadows generated using a simulation model to determine the most likely object to have cast such a shadow. Not only does this allow a basic classification of mine/not-mine but it also provides details of the shape and dimensions of the object casting the shadow. The process has also incorporated a final post-processing stage which can exploit multiple views of the one target, combining the classification results in a fusion system to increase the probably of correct classification.

This technique is also applicable to Synthetic Aperture Sonar (SAS) and Synthetic Aperture Radar (SAR) data. This paper will present initial work considering the potential for the extension of the technique to SAS imagery, showing results of the detection and extraction stages. Although the paper will illustrate the technique using only SAS imagery, it will also discuss how this technique could be exploited for SAR imagery.

2 TARGET DETECTION AND CLASSIFICATION

2.1 Sonar Imagery

With the recent advances in autonomous underwater vehicle (AUV) technology for mine-countermeasures (MCM) the need has arisen for automated techniques for object identification from sonar data. Research carried out in developing MCM tools is generally split into Computer Aided Detection (CAD)[1-5] to detect all possible mine-like objects, and Computer Aided Classification (CAC) models [6-11] to classify whether the detected object is a mine or not. A common approach is to compare a set of extracted features from the mine-like object (MLO) [6,3] to a set of pre-determined training data (supervised feature-based approaches). The system is trained using a set of ground-truthed data before being run on the unknown "test" data. These approaches work well when the test data is similar to the training data but can provide poor results when this criteria is not met[3]. This can occur frequently since sonar imagery is very dependent on the sensor to target azimuth, ensuring images of the same underwater scene can look very different depending on the particular conditions. In addition, trained systems essentially learn the combination of the target and the context, and since the image formation process is non-linear it can be difficult to classify the same target in another context. This can be counter-acted by fusing the results from multiple classifiers[11]. However, while this process produces improved results, it does not confront the true underlying problem that the MLO and its features vary greatly depending on the specific sonar conditions. These systems offer a blackbox solution to the problem, where it is difficult to ascertain why a particular result is obtained.

Obtaining information other than the basic mine or not-mine label is usually referred to as object identification. Information such as the shape and dimension of the mine can allow the mine type to be determined and can help detail how best to neutralise the threat. Man-made objects such as mines generally have regular shapes and so leave regular shaped shadows in sidescan images. The shape of these shadows can be used to identify the objects by extracting relevant features from the shadows and comparing these to known training data[8]. The non-linear nature of the shadow-formation process ensures a shadow normalisation step is required for these approaches to be widely applicable. Another approach is to fit template approximations of the shadows produced by known mine types to the MLO's shadow[9]. These models are useful in discriminating between different mine types but often have to assume that the detected MLO is a mine to begin with. The templates are also generally deformed using linear operators and are therefore not always accurate in modelling the non-linear shadow formation process.

2.2 SAS Detection and Classification

The first commercially produced Synthetic Aperture Sonar (SAS) systems are now under development, bringing the promise of higher resolution surveys. In the area of Mine Countermeasures this offers the capability to detect mines at longer ranges and provide higher resolution images of targets for subsequent classification. However, little work has been done to investigate the automatic detection and classification of mines from SAS imagery, since previous research in this area concentrated on developing the sensor rather than the automated processing of the images, since it was always assumed that their higher resolution would simplify this task. However although providing higher resolution, the images contain a significant level of speckle due to the construction of the image. Filtering methods have been employed, but these can degrade either the shadow or the highlight. Higher order statistics have been employed for detection of both buried and proud targets, however this required some a priori knowledge of the characteristics of the echoes[12]. Much of the literature concerning detection and classification of targets from SAS imagery has concentrated on the detection of buried objects or looking at novel techniques, including bistatic approaches, to improve further the potential for the detection of buried targets.

2.3 SAR detection and Classification

Automatic Target Recognition (ATR) in SAR images is a difficult problem which has been under study since the early 70s and is significantly more advanced than SAS detection and classification. Correlation and optical correlation have been widely used to perform template matching on large databases of stored images[13]. Neural Nets and, more recently, support vector machines have also been used[14]. Flexible histograms[15] have recently been introduced and provide a good alternative to template matching as it reduces the amount of training data required. Hidden Markov Models (HMM) have been proposed as a means to model articulated objects[16] and deal with occlusions[17]. These techniques have proven relatively successful on known targets and simple backgrounds. However, performances degrade rapidly on more challenging backgrounds or when confronted with targets with incomplete training data. Most authors recognise that pose estimation is critical for successful recognition using learning techniques[15,18]. Recently, model-based detection and classification techniques have been evaluated[17]. Critical to model-based techniques is the ability to accurately simulate data and extract characteristics from the real image to perform the model matching. Although some authors have studied the first problem and accurate simulators exist (DARPA MSTAR project[15,18]), the second problem has been overlooked.

2.4 Model Based Detection and Classification

As discussed many techniques for classification are reliant on training data and their success can be highly dependant on the similarity of the training data to the test data. The authors have previously developed a model based technique[4,10] for the detection and classification of mine-like objects (MLOs) in sidescan imagery to overcome this problem with a three-stage process. This paper will initially briefly summarise the overall detection and classification system before discussing its application to SAS imagery.

3 OVERVIEW OF DETECTION AND CLASSIFICATION SYSTEM

The three stage process for detection and classification is summarised in figure 1. The first stage is the detection of mine like objects (MLOs) within the sidescan sonar image. This stage uses a Markov Random Field (MRF) model to directly segment the image into regions of object highlight, shadow and background[4]. Unlike many previous models for object detection this requires no training and the structure of the MRF model also allows known information to be modelled and included through the use of priors that take into account the characteristics of sidescan data. This is simple information such as the fact that a target is generally displayed as an area of object highlight followed by shadow, and that the highlight tends to form small clusters. This stage, as with all other parts of the process, requires no training and the Markovian parameters to represent the three regions of object highlight, shadow and seabed reverberation are estimated from the image under consideration.

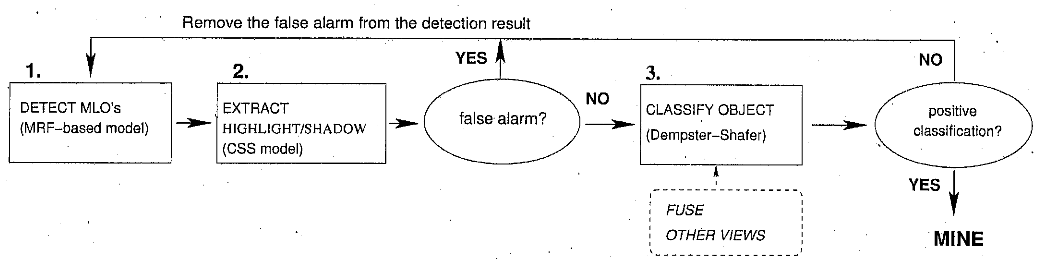

Figure 1: Overview of Detection and Classification System

The detected targets from this first stage are then passed to stage two, where the highlight and shadow regions of the detected objects are extracted using a Cooperating Statistical Snake technique (CSS)[4]. The CSS model approximates the background as three homogenous regions – object highlight, background and shadow and so uses two statistical snakes, one to segment the highlight and one the shadow. The a priori information between the highlight and shadow is used to constrain the movement of the snakes so as to achieve accurate segmentation results regardless of the seabed type involved.

The CSS model can also be used to eliminate false alarms from the detection stage. Areas can be indicated as object highlight and false alarms produced, especially from more complex regions such as sand ripples. If the CSS model is applied to these regions, since they do not have the expected characteristics of MLOs, the shadow snake will expand in an uncontrolled manner. If the snakes expand beyond mine-like dimensions the detection can be identified as a false alarm and removed.

After a detected MLO has passed through both the Detection and CSS modules, stage 3 of the system will classify the object. To do this, the system has the extracted shadow and object-highlight region of the MLO, as well as some simple information extracted from the navigation data, such as height of the sonar above the seabed and the range to the object at the time of ensonification. Although the highlight is generally not considered for classification purposes since it is generally unpredictable and difficult to model, basic information can be extracted from it and used together with the information from the shadow region.

The model represents possible mine-like shapes (cylinder, sphere, truncated-cone) using parametric models which allow a sonar simulator to generate the resulting shadow region from such an object. Each shadow region is specific to the particular parameters of each shape and is generated under the same sonar conditions as the MLO was detected. As the model searches through the different parameter options, the resulting synthetic shadows are compared to the real MLO shadow to find the best match for each considered class. This section of the model is the most computationally intensive and a detailed overview of this section is shown in Figure 2.

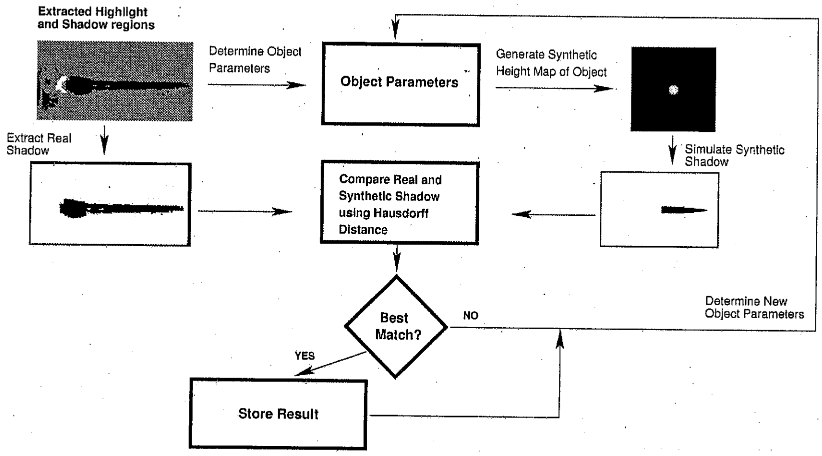

Figure 2: Determination of best synthetic shadow match to real shadow

Once this has been completed, the degree of match and the shape parameters used to obtain it are output to the classifier to define a class membership function. These membership functions are entered into a Dempster-Shafer classifier which identifies the belief that the object is a cylinder, sphere, truncated-cone or clutter object.

As indicated in figure 1, this Dempster-Shafer classification can also be extended to include multiple views[10] if the region has been traversed in a number of directions. This often occurs due to the 'lawn-mower' nature of surveys which ensures that the same object often appears in multiple images. The ability to consider multi-view analysis allows the classification system to use more of the available information before providing a classification result and increases the accuracy.

4 RESULTS

The first two stages of the classification have been tested using SAS data from both a rail based system and an AUV. This used directly the software developed for the detection and extraction of targets in sidescan imagery with no modifications. The third stage which then undertakes the model-based classification could not be tested since this requires knowledge of the exact resolution of the images and slant ranges to targets to enable dimensions in metres to be extracted from the data.

The first step was the MRF based detection and results are illustrated in figure 3. In this case, both of the targets highlights were clearly identified and the shadow region of each target is apparent. However, the shadow region was not detected as "pure" shadow and contained regions labelled as seabed reverberation within the general shadow zone of the target. This was particularly apparent for the cylindrical target.

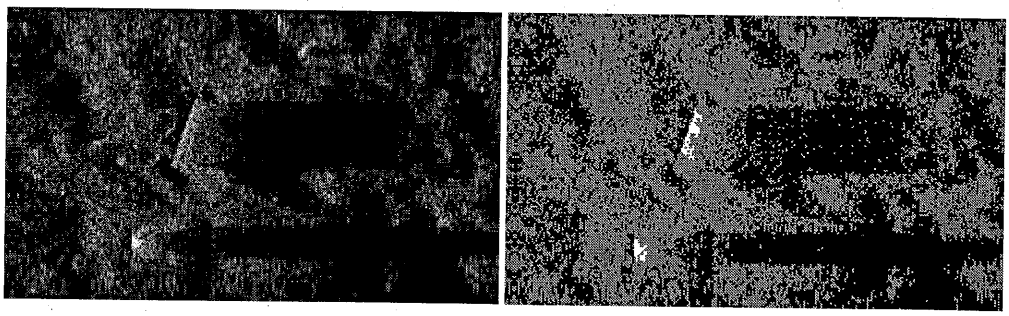

Figure 3: (a) SAS image of cylindrical and spherical target (b) MRF detection result (white represents target, black shadow and grey background reverberation)

However, the 2nd stage of the process, where the contour of the highlight and shadow are extracted using Cooperating Statistical Snakes, performed well. This technique is searching for areas with similar statistical properties and uses the original image and not the segmented image. The detection result is used purely to identify possible target regions, which are then to be examined for classification. The CSS has successfully been able to extract the contours of the shadow and the increased speckle within the shadow has not degraded the result.

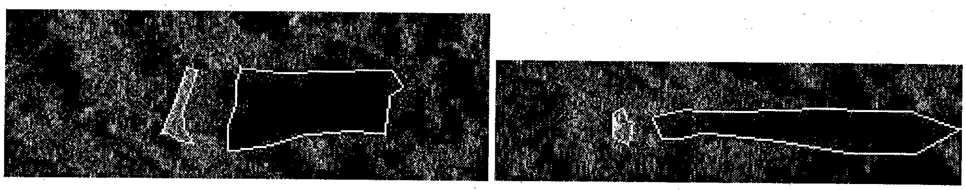

Figure 4: Extracted highlight and shadow regions using Cooperating Statistical Snakes (a) cylindrical target and (b) spherical target from figure 3

Results are also shown in figures 5 and 6 for a range of other target shapes. The same characteristics were also noted, that the detection stage was able to successfully detect target regions containing the characteristic highlight and shadow, but that parts of the shadow region were often incorrectly labelled as seabed reverberation. The CSS was however able to successfully extract the contours of the highlight and the shadow. This was particularly noticeable in figure 6, where even although the detection model only labelled the bright bands on the target as highlight, the CSS was able to extract the entire region of target and not just the brightest reflectors on it.

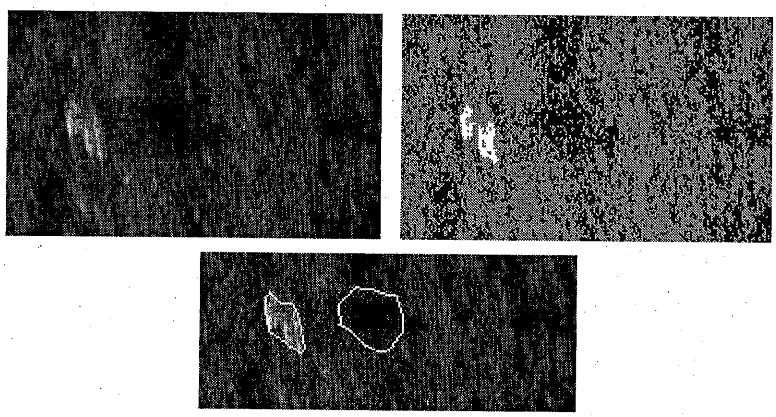

Figure 5: (a) SAS image of spherical target with two cylindrical anchors (b) MRF detection result (c) extracted highlight and shadow

Figure 7 illustrates a SAS image of a more complex target, where there appears to be a shadow region in front of the target, probably caused by scouring. This shadow region as well as the shadow behind the target has been detected by the MRF segmentation, and appears as a more consistent region of shadow. This has then led to a failure of the CSS since it has extracted the contour of the shadow in front of the target instead of the shadow behind. (The system has been configured to work for both port and starboard sidescan images and will look for shadow both to the right and left of a target). Improvement of the MRF detection of the shadows should help to eliminate this problem.

4.1 Discussion of Results

Although it was possible at this stage to only assess the results of the first two stages of the process, the results appear promising, in that accurate contours of the objects highlight and shadow could be extracted in the majority of the cases and that this suggest no reason for the overall process not to be successful. In the final case, this illustrated a more complex example where the seabed topography was such as to create an additional shadow.

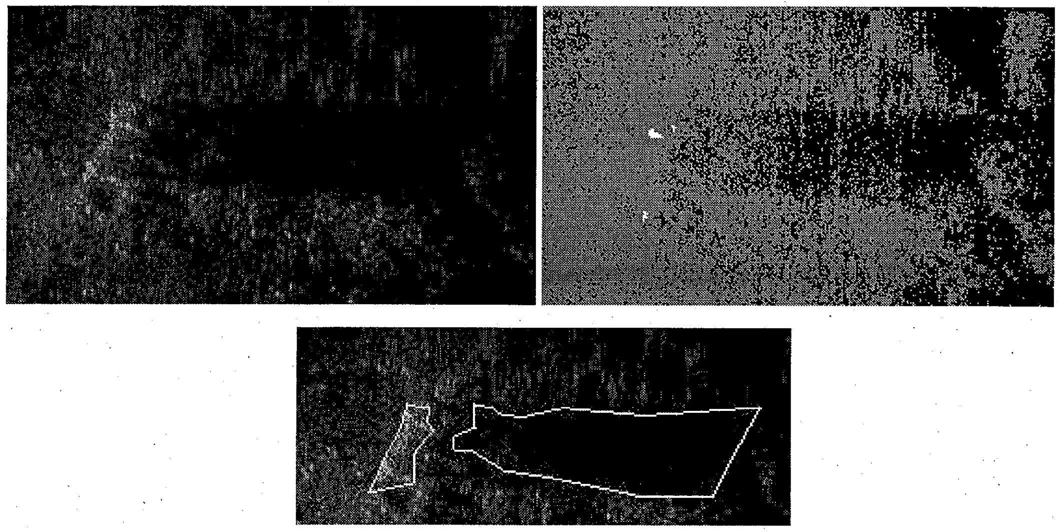

Figure 6: (a) SAS image of cylindrical target (b) MRF detection result (c) extracted highlight and shadow

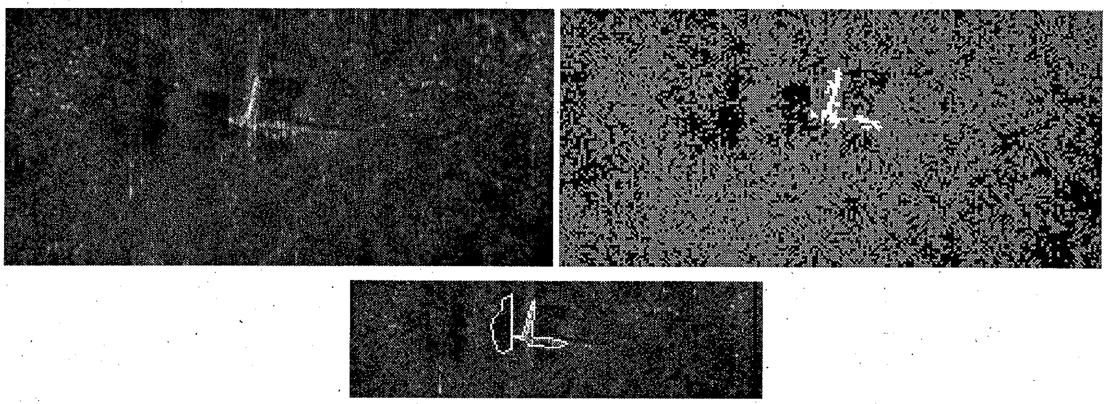

Figure 7: (a) SAS image of target (b) MRF detection result (c) extracted highlight & shadow

The SAS images are significantly higher resolution than the sidescan images which have previously been processed using this technique. In the majority of sidescan images the target highlight and shadow is typically only a few pixels. In these images much greater resolution was available. However, there was significant speckle present within the images.

The MRF detection orientated segmentation requires no training, and estimates the Markovian parameters from the image. It does however assume that the likelihood term for the shadow class is Gaussian and the likelihood term for the seabed reverberation is a shifted Rayleigh[4]. Although, these assumptions are generally regarded as appropriate for sidescan imagery, the increased speckle of the SAS images may mean that these are no longer the most appropriate choices. In the results it can be noted that although the highlight is detected, it is sometimes only the brightest features within the overall highlight region. The shadow region is also detected not as a complete shadow, but with regions of seabed reverberation within it. This is probably due to the increased speckle apparent in the original images. The detection of the shadows was not affected by the choice of parameters, which also suggests that the technique may be converging to a local minimum. The original implementation of the MRF for the sidescan data sacrificed convergence accuracy for speed, since the shadows in sidescan tended to be very clear within the data. This would require further investigation, but suggests that an improvement could be obtained.

The Cooperating Statistical Snakes is seeking to draw contours in the image around regions with similar statistical properties. The fact that the statistics of the SAS images were different from the sidescan images did not degrade the results, since this was not considered by the technique. In this case, the increased speckle did not affect the results, in that the statistics of the shadow with speckle was still different from the seabed with speckle. This clearly shows the importance of the CSS as an integral part of the three stage process, since even when the initial MRF segmentation was not ideal the CSS enables the recovery of the correct highlight/shadow pair ready for classification.

5 FUTURE EXTENSIONS

5.1 Application to SAS

The improved resolution of SAS, particularly of the target echo, provides scope for classification based on this highlight. The majority of classification systems using sidescan tend to rely on the shadow information, whereas it is proposed that exploiting the improved highlight resolution of SAS could improve the classification accuracy. The model based techniques for the classification of sidescan presented here uses an iterative procedure incorporating a simulation model to determine the object most likely to have created the shadow. It is proposed that similar techniques using models of the SAS highlight formation could be investigated. The modular architecture of the Dempster-Shafer fusion system[10] would allow the incorporation of both shadow and highlight based techniques for classification, placing emphasis on the more appropriate algorithm for the particular context.

The flexibility of the adaptive fusion based architecture could also be exploited for the classification of SAS images, since SAS also provides scope for multi-view images from the one sensor. This can be achieved by using SAS coherent beamforming to form both broadside and squint images of the same target using sub-sections of the synthetic aperture, enabling different views of the same target to be obtained. The Dempster-Shafer fusion architecture[10] would permit the combination of the classification results from the sub-apertures to improve further the probability of detection and reduce false alarm rate.

5.2 Application to SAR

SAR Automatic Target Recognition is a difficult problem since although the images are of a higher resolution than SAS, the targets are more complex in shape and located in more complex terrain which may occlude the target. The appearance of the target presents a mono-static response, which is a function of the target-sensor orientation and azimuth. The statistics of the image formed using SAR imagery are complex and multiplicative leading to complex statistical modelling. However, SAR imagery is similar to sonar imagery in terms of statistics (for example through the use of the k distribution). However, since the model based technique does not require training and can extract the target from any background (provided that it is not occluded) it is suggested that this technique may also be applicable to SAR imagery. The added difficulty of SAR data is appreciated and unlike sonar the complex SAR signals could also be exploited.

6 CONCLUSIONS

This paper has shown the feasibility of applying a three stage model based system for detection and classification of objects in sidescan imagery to the detection and extraction of objects within SAS data. The paper showed the segmentation based detection of objects using a Markov Random Field process. Although the detection was not always ideal, the second stage of extracting the contour of the object highlight and shadow using a Cooperating Statistical Snake technique still performed well. However, it was not possible to run the third stage of the classification directly on these images, as the exact details of the images were unknown and the system would require to be tuned for the improved resolution of the SAS images. Given the results of the first two stages there would appear to be no reason for the process not to be successful.

The application to SAS with its improved resolution of the highlight also suggests the extension of the model based classification to consider the highlight as well as shadow. Since the simple line of sight model used to simulate the shadows in the sidescan classification system may prove to be inadequate for the SAS imagery anyway, a more complex model may need to be considered.

The paper has also discussed the feasibility of applying this technique to SAR imagery, although no work has yet been undertaken in this area. The increased complexity of the SAR imagery and the target modelling have implications on the simulation and may prove to be too computationally complex unless simplifying assumptions are applied.

In summary, the paper has presented preliminary ideas for CAD/CAC in SAS and SAR imagery and significant further work is required, although initial results appear promising.

7 ACKNOWLEDGEMENTS

The authors would like to thank T. Sutton and H. Griffiths from UCL for providing the SAS images.

8 REFERENCES

1. T. Aridgides, M. Fernandez and G. Dobeck, "Adaptive 3 Dimensional range-crossrange-frequency filter processing string for sea mine classification in side-scan sonar imagery" Proc SPIE, 3079, 111-122, 1997
2. B.R. Calder, L.M. Linnett and D.R. Carmichael, "Spatial stochastic models for seabed object detection" Proc SPIE, 3079, 1997
3. C.M. Ciany and J.Huang, "Computer aided detection/computer aided classification and data fusion algorithms for automated detection and classification of underwater mines" *Proc. MTS/IEEE Oceans Conf. And Exhibition,* 1:277-284,2000
4. S.Reed, Y.Petillot, J. Bell, "An Unsupervised Approach to the Detection and Extraction of Mine Features in Sidescan Sonar", IEEE J. Oceanic Engineering 28(1), 90-105, Jan 2003
5. B. Stage & B. Zerr, "Detection of objects on the seabottom using backscattering characteristics dependant on the observation point", IEEE J Oceanic Eng.,22(1),40-46,1997
6. G.J. Dobeck, J.C. Hyland and L. Smedley, "Automated detection/classification of sea mines in sonar imagery", Proc SPIE, 3079, pp. 90-110, 1997
7. E. Dura, J. Bell and D. Lane, "Superellipse fitting for the classification of mine-like shapes in side-scan sonar imagery", IEEE/MTS Oceans 2002, pp. 23-28, 2002
8. J.A. Fawcett, "Image based classification of sidescan sonar detections", presented at CAD/CAC Conf. Halifax, Canada, Nov 2001
9. M. Mignotte, C. Collet, P. Perez and P. Bouthemy, "Hybrid genetic optimisation and statistical model based approach for the classification of shadow shapes in sonar imagery", IEEE Trans Pattern Analysis and Machine Intell., 22(2), pp. 129-141, 2000
10. S.Reed, Y.Petillot, J. Bell, "Automated approach to the classification of mine-like objects in sidescan sonar using highlight and shadow information", IEE Proceedings – Radar, Sonar and Navigation, 151(1), pp. 48-56, Feb 2004
11. G.J. Dobeck, "Algorithm Fusion for automated sea-mine detection and classification", Oceans 2001, 1, 130-134
12. E Maussang, J. Chanussot, A. Hetet, "On the use of higher order statistics in SAS imagery", Proc. Int Conf Acoustics, Speech and Signal Processing, ICASSP 2004, 269-272
13. A. Mahalanobis, A. Forman, N. Day, M. Bower, and R. Cherry. Multi-Class SAR ATR using Shift-Invariant Correlation Filters. Pattern Recognition, 27(4):619-626, April 1994
14. Q. Zao and G. Principe, Support Vector Machine for SAR Automatic Target Recognition. IEEE Transactions on Aerospace and Electronics Systems, Vol 37, No 2, April 2001.
15. D. C. Stanford and J. Pitton, Multiresolution Hidden Markov trees for analysis of Automatic Target Recognition Systems. Int. Conference on Image Processing, Vol. 2, 942 -945 , 2000
16. B. Bhanu and B. Tian, Stochastic Models for Recognition of Articulated Objects, IEEE conference on Image Processing, Pages 847-850, vol.2, 1997.
17. B. Bhanu and G. Jones, Object recognition results using MSTAR synthetic aperture radar data, Computer Vision Beyond the Visible Spectrum: Methods & Applications, 55-62, 2000.
18. P. L. Douville, Measured and Predicted Synthetic Aperture Radar Target comparison. IEEE Transactions on Aerospace and Electronics Systems, Vol 38, No 3, January 2002.