

APPLICATIONS OF SONAR SIMULATION IN OBJECT CLASSIFICATION.

J.M. Bell, S. Reed, Y. Petillot

School of Engineering and Physical Sciences, Heriot-Watt University, Riccarton, Edinburgh,
J.Bell@hw.ac.uk

Abstract

Accurate models of the sonar process that can provide realistic synthetic data have a wealth of applications in the design and testing of sonar systems and the provision of ground truth data for the validation of processing algorithms. Models can also play an integral role as part of the processing of the sonar data. This paper will present one such application with the development of a model-based system for the classification of mine-like objects from sidescan sonar imagery. This integration of the model removes the need for training data, which is often the limiting factor in classification systems since their success can be highly dependant on the similarity of the training data to the test data. The paper will outline the framework for the detection and classification system within which this is integrated, before discussing the classification procedure and presenting results on real data.

1. INTRODUCTION

A significant problem in the validation of algorithms for automated processing of sonar imagery lies in obtaining fully ground truthed data for which all parameters are accurately known. One solution to this problem is the use of synthetic data provided by sonar simulation models which accurately model the underlying physical processes producing realistic synthetic data. These models allow the operating scenarios to be easily controlled and individual parameters altered in isolation to assess their impact on the subsequent processing. Bell et al. [1] have used synthetic data for testing seabed classification systems, reconstruction and object detection algorithms.

However modelling can also play a role as an integral part of the actual processing of the data. This paper will highlight one such application with a model-based technique for the classification of objects in sidescan 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 sidescan sonar. 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. This model-based technique attempts to overcome this problem with a three-stage process. 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 will this allow a basic classification of mine/not-mine but it will also provide details of the shape and dimensions of the object casting the shadow.

Traditionally, object classification models in MCM applications have been concerned with labelling an object simply as mine or not-mine. A common approach is to compare a set of extracted features from the mine-like object (MLO) [2,3] to a set of pre-determined training data. These feature-based approaches work well when the test data is similar to the training data but can provide poor results when this criteria is not met. This can occur frequently since sidescan

Vol.26. Pt.5. 2004 (Sonar Signal Processing)

Vol.26. Pt.6. 2004 (Bio Acoustics)

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. This can be counter-acted by fusing the results from multiple classifiers [4]. 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.

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 [5]. 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 [6]. 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.

The technique presented here overcomes these problems since no training data is required and the simulation will take into account the orientation of the object relative to the sonar as well as the range. This ensures that no normalisation or deformation of the shadow is required. The paper will initially briefly summarise the overall detection and classification system before concentrating on the model based aspects of the classification.

2. 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 [7]. 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. An example of the output of the detection stage is shown in figures 2(a) and 2(b).

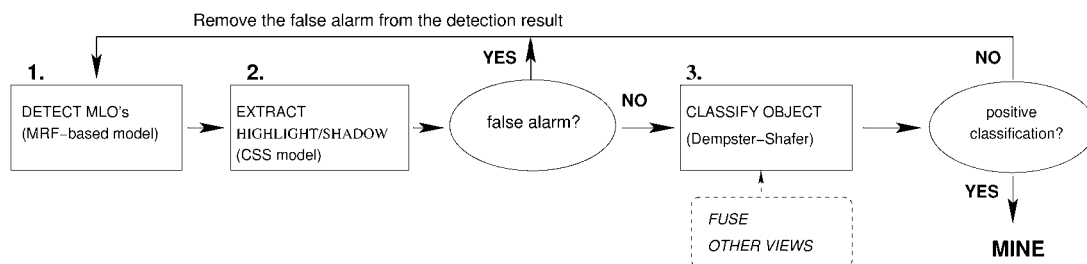


Figure 1: Overview of Detection and Classification System

The highlight and shadow regions of the detected objects are then extracted using a Cooperating Statistical Snake technique (CSS) [7]. 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. This is illustrated in figure 3 for the detected objects from figure 2, where the shadow has been accurately extracted even when the

object is lying on ripples which themselves cast shadows.




Figure 2: (a) Sidescan image (b) Image after processing by MRF model (highlight shown as white, shadow black and background grey) (c) Image after removal of false alarms

The CSS model can also be used to eliminate false alarms from the detection stage. Within figure 2b, particularly in the region of ripples, several areas without objects have been indicated as object highlight. 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 MLO can be identified as a false alarm and removed [8]. The final image with false alarms removed is shown in figure 2c.




Figure 3: Objects from Figure 2 with shadow and highlight contours extracted using CSS

3. OBJECT CLASSIFICATION

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 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.

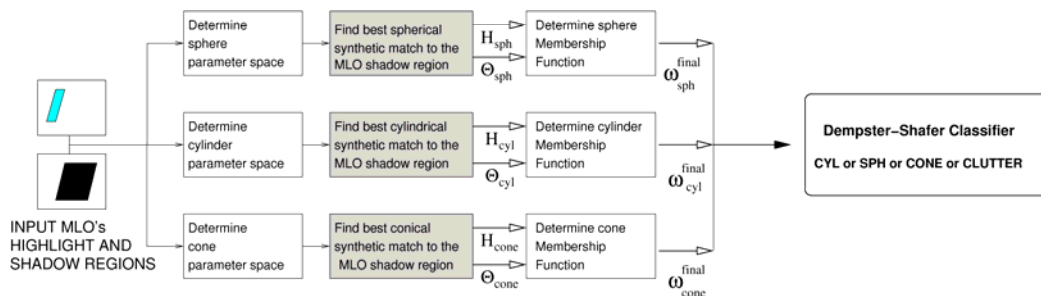


Figure 4: Overview of Classification Model

The general plan for the classification model can be seen in Figure 4. 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 is contained within the shaded box of Figure 4. A more detailed overview of this section is shown in Figure 5.

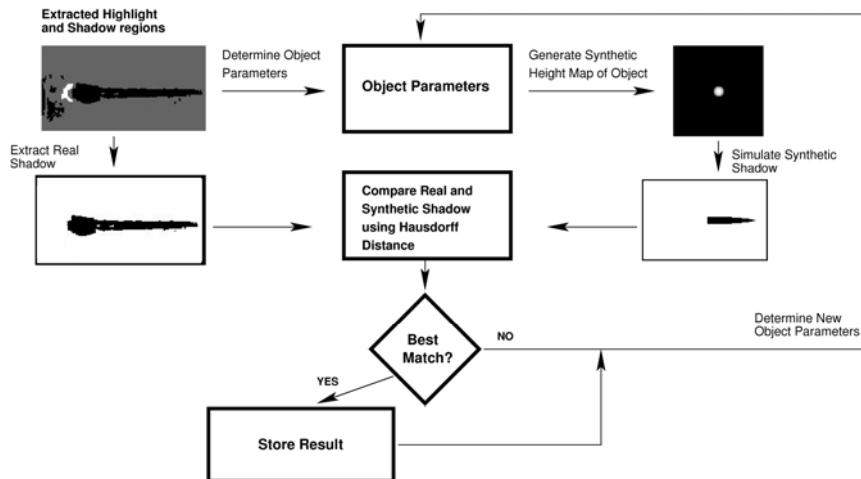


Figure 5: Determination of best synthetic shadow match to real shadow

For each object class first approximations of the object parameters are calculated from the objects highlight and shadow dimensions. This is used to create a synthetic height map of this object lying on a flat seabed which can be input to the simulation system to create the synthetic shadow. This procedure is repeated iteratively using amended dimensions for the object until the best match is obtained. 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. Each of these stages of the process will now be considered in detail before looking at the Dempster-Shafer classifier and some sample results.

3.1 Object Parameters

The tested objects were limited to cylinder, sphere and truncated cone classes, since these simple shapes closely represent the shape of most mines. The cylinder class is assumed completely described by 4 parameters representing its radius and length, the angle of its major axis with respect to the along track direction and the depth to which it is submerged into the seafloor (with a value of 0 relating to a cylinder sitting proud on the seabed). The sphere class is described by only two parameters – the radius and depth of submergence. For simplicity, only the base-radius was used to describe the truncated cone class. The radius at the top of the truncated cone was assumed to be 0.5 of the base radius while the height was set at 0.8 of the radius. This is in good agreement with currently tested truncated cone objects and allowed a generic truncated cone class to be specified rather than generating a specific class for each known mine type. Using these parameters, object height maps with the object lying on a completely flat seabed with the same resolution as the real image can be easily generated.

3.2 Shadow Simulation

The height map together with the slant range, fish height and image resolution information recorded from the real data is used to generate the synthetic sonar image. The simulation is only required to produce the synthetic sonar image of the shadow of the object rather than the complete synthetic image of the object, background and shadow. As a result of this a very simple line of sight ray tracing based calculation can be used.

For each pixel in the height map image, the model determines if this point is in shadow from any previous point. This assumes isovelocity conditions and a simple point source receiver where the source and receiver are collocated. It also assumes that there is an infinitely thin horizontal beamwidth with no spreading. This is an approximation from an existing more complex sonar simulation model [1]. This is necessary to reduce the computational complexity as a result of the iterative process in which it is embedded. This reduction is a reasonable assumption when only the generated shadow is of interest rather than the full backscattered signal. This can be illustrated in figure 6 which shows the shadow cast by a sphere calculated by both the simplified model and the complex model. The sphere is of radius 60cm at a range of 60m from the sonar which is at a height of 10m above the seabed.

The use of a height map for the objects means that the underside of the sphere and cylinder can not be accurately represented as each point on the height map can have only one value associated with it. The more complex model represents objects in a procedurally defined manner and can therefore represent accurately all points on the surface of the object, including the underside. However as can be seen this approximation does not have a significant effect on the shadow calculation. This can also be partly attributed to the relative size of mine like objects and the resolution of sidescan data which means that small details are not apparent in the image.

The height map also assumes a completely flat seabed. The use of a rough seabed may affect the shadows if there are large roughness variations and will depend on the exact location of the target. For the 2nd example in figure 6 the shadow shape is affected slightly by the ripples.

Another region in which the basic simulation can fail is due to the horizontal beam spreading. This may result in the shadow appearing to get narrower in the along track direction with increasing range or it can be narrower than the highlight and less distinct from the background reverberation around its perimeter. However, as the increased computation time has to be weighed against the improved accuracy, the classification membership functions were relaxed rather than tightening the simulation accuracy (see section 3.5).




Figure 6: (a) Shadow from simple model (b) Simulated image and (c) extracted shadow from complex model (d)(f) Simulated image and (e)(g) shadow when object on rough seabed (h) Simulated image when horizontal beam included.

3.3 Comparing the Shadow Regions

The Hausdorff distance [9], which measures the similarity between two shapes, is used to compare the real and synthetic shadow. Defining $Y = \{y_1, K, y_p\}$ and $V = \{v_1, K, v_q\}$ as points on the perimeter of the real and synthetic shadows respectively, the Hausdorff distance is defined as

$$H(Y, V) = \max(h(Y, V), h(V, Y)) \quad (1)$$

where the *directed* Hausdorff distance $h(Y, V)$ is calculated by first computing the distance between each point in the perimeter of the real shadow to its nearest neighbour in the perimeter of the synthetic shadow. $h(Y, V)$ is the maximum of this set of values. Therefore if $h(Y, V) = u$, each point in real shadow must be within distance u of some point in the synthetic shadow. A similar process is carried out to compute $h(V, Y)$ with the Hausdorff Distance being designated the

maximum value of the two directed distances. $H(Y,V)$ is therefore a measure of mismatch between the real and synthetic shadows. Unlike many other methods for comparing shapes, there is no explicit pairing of points in Y with points in V . The technique is also fast to compute which is necessary when used within an iterative process.

3.4 Iterative Matching Process

An iterative process was used to search through the membership functions for each object to find the best match from each class. To limit the parameter space tested by the model constraints can be imposed by considering the real objects extracted highlight and shadow regions.

Moment analysis was used to define the image ellipse of the object highlight region. This allowed estimates of the major and minor axis of the highlight to be obtained which were used to determine initial values for the radius of each object and the length of the cylinder. The angle of the cylinder was also initialised considering moment analysis. Initial estimates for depth parameters are found by considering the length of the shadow and estimating from knowledge of the fish height and slant range the height of the object. As an estimate for the radius has already been obtained, a value for the depth can then be computed.

Large margins were set on each of the initial parameter estimates, defining a discretized parameter space for each class which had to be searched through to find the best fit to the MLO's shadow. The minimum parameter increment allowed was set at the resolution of the image being tested. The parameter space for both the sphere and cone was 2-dimensional, allowing an quick exhaustive search to be used to find the best Hausdorff Distance. The parameter space for cylinder was 4-dimensional. A simple Maximum Likelihood search was carried out to find the lowest value for the Hausdorff distance.

After the iterative matching process, each object class is allocated its lowest Hausdorff Distance value (H_j^b) obtained with parameters Θ_j for $j \in \{cyl, sph, cone\}$. For cases where multiple sets of parameter values gave the same Hausdorff distance for a given shape, the most mine-like set of dimensions were used. This information is then fused with additional information to produce a membership function for each of the classes.

3.5 Determining the Class Membership Functions

Fuzzy membership functions are often used when a decision needs to be made by considering multiple sources of information. These functions lie in the real, closed interval $[0,1]$ where values close to 0 describe a low membership (an unlikely outcome) and values near 1 relate to a high membership (highly likely outcome). For classification purposes, the overall membership function for each class is defined by

$$\omega_j^{final}(H_j^b, \Theta_j, \alpha) = \omega_j^{haus}(H_j^b) \times \omega_j^{par}(\Theta_j) \times \omega_j^{high}(\alpha) \quad (2)$$

The function $\omega_j^{haus}(H_j^b)$ considers the best Hausdorff Distance, H_j^b , obtained for each of the classes. If the sonar simulator exactly modelled the real sonar process, the correct object class would be expected to give a Hausdorff Distance value of 0. While accurate, such simulators are computationally heavy precluding their use within an iterative matching process. The simple simulator described in section 3.2 produced Gaussian distributions around non-zero Hausdorff values and the membership function reflects this [8].

The $\omega_j^{par}(\Theta_j)$ membership function describes the degree of match between the synthetic shadow from each class and the real MLO's shadow. While the synthetic shadow may visually appear

Vol.26. Pt.5. 2004 (Sonar Signal Processing)

Vol.26. Pt.6. 2004 (Bio Acoustics)

very similar to the real shadow, no consideration has yet been given to the parameters of the synthetic object. This membership function considers how closely the determined parameters match real-mine dimensions. The form of these functions has been chosen for general mine-like dimensions but could be made case-specific.

The final function considers the elongation, α , of the object's highlight region. A high α value would be expected from objects from the cylinder class and low for other objects.

4. RESULTS

The class membership functions were determined for 50 different objects taken from data sets provided by GESMA and DRDC-Atlantic. The GESMA data was obtained with a Edgetech DF 1000 sonar mounted on an AUV. The DRDC-Atlantic data used a Klein 5500 sonar with a towfish.

The test data extracted from the real sidescan images included 10 cylindrical objects, 10 spherical objects, 10 truncated cone objects and 20 clutter objects. An example of each is provided in Table 1. From the 30 mine-like objects all 10 cylinders and 10 spheres were correctly classified. However, only 3 of the cones could be clearly defined since both the sphere and cone classes often provided non-zero membership functions. This is the case for the example provided in table 1. The cone and sphere classes often provide visually similar shadow regions under certain sonar conditions. Also, a simplified description of the truncated cone class was used with only one parameter when ideally three are required (base-radius, top-radius and height). The test data contains 2 different truncated cone models which have different dimensions and so would be better modelled if they were allocated their own class. It should also be noted that the cylinder class has zero membership for all the cone and sphere images.


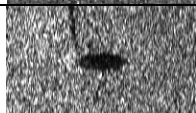


Object Type	Sidescan Images	Membership Functions			Dempster-Shafer Belief			
		ω_{cyl}^{final}	ω_{sph}^{final}	ω_{cone}^{final}	Bel(cyl)	Bel(sph)	Bel(cone)	Bel(clutter)
cylinder		0.84	0	0	0.71	0	0	0.2
sphere		0	1	0	0	0.833	0	0.083
cone		0	1	1	0	0.303	0.45	0.045
clutter		0.4	0	0	0.42	0	0	0.46

Table 1: Example Membership Functions and Dempster-Shafer Results

The clutter image results were dependent on the visual similarity to the considered classes. The very irregular clutter objects provided low membership functions for all 3 classes. For those that closely represent a specific object class, only the relevant membership function was non-zero. This is demonstrated in Table 1 where the clutter object appears visually similar to a cylindrical object. In scenarios such as this where the object is both mine-like in appearance and size, further analysis would be required either through the use of another sensor such as video, or by

human inspection. However, even to the human eye, such objects are not obviously clutter.

4.1 Dempster-Shafer Information Theory

A classification decision could be made on these membership functions by simply imposing a threshold. This is a 'hard' technique where the object is either classified as mine or not-mine. A more subtle approach is to allocate a probability or a 'belief' to each of the classes. This is particularly desirable for Sidescan surveys where the same object is often viewed multiple times offering the opportunity for multi-view analysis to be carried out on the object. This can be achieved by considering Dempster-Shafer (D-S) Information Theory, which has the advantage over Bayesian techniques that it can consider union of classes. This can be employed to improve the separability of the truncated cone and sphere classes.

Dempster-Shafer(D-S) theory allows the representation of imprecision and uncertainty through the definition of 2 functions: plausibility (Pls) and belief (Bel). These are derived from a mass function which is analogous to the well known probability function. Mass functions are defined on the power set of the space of discernment which for classification purposes may be the set of possible classes. Different techniques can be implemented to make a decision based on the belief and plausibility values. The most common, and the one used here, is the maximum belief over all the singleton classes (thus ensuring that the final classification result is not a union of classes). For the mine classification model employed here, the allowed classes were clutter, cyl, sph, cone, sph \cup cone. The non-singleton, sph \cup cone has been defined to remove the degeneracy between the sphere and cone classes on the truncated cone membership results [8].

4.2 Classification Results

The 50 objects used previously were classified using the D-S classifier and the results shown in Figure 7 for a range of threshold values. The first line in figure 7 considers the standard mine/not-mine classification while the second considers the classification as correct only if the correct object class has been identified. Results for the four example images are also given in table 1, where the advantage of the D-S theory is clear since the cone can now be correctly identified.




Figure 7: Results for the Dempster-Shafer Classification Model.

A threshold value of 0.18, used within the mass function for the clutter [9], correctly classifies over 85% of the mines by identifying their specific object class. The same threshold correctly identifies over 50% of the clutter objects as false alarms where it is important to note that some of the clutter objects appear very 'mine-like'. This result would allow the number of false alarms detected using any detection model to be dramatically reduced. Often clutter objects were classified as a mine because they visually appeared very like one of the three classes. In these situations, it is

arguably better to mis-classify the clutter object allowing it to be dismissed at a later time. The same threshold value allows around 95% of the mines to be correctly labelled as simply mine. However, obtaining a specific shape classification (as well as parameter information) allows the possibility of identifying the mine type and thus influencing how the specific threat is dealt with.

As indicated in figure 1, this D-S classification can also be extended to include multiple views [8] 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.

5. CONCLUSIONS

This paper has presented a classification model which extends the standard mine/not-mine classification procedure to include shape and parameter information on the object. Results were presented on real mine and clutter images. These results showed how the proposed approach could remove a significant amount of the false alarms picked up by a mine detection model while classifying a large percentage of the real mine objects.

The system could be improved further since describing the truncated cone class with only one parameter was a limiting factor. Another possibility is to remove the generality of the classification process (currently only the general mine shapes are restricted), creating a mission-specific classifier where mines are of a known class and so known shape/dimension could be tested for.

Improvements would also arise with the inclusion of a more complex sonar simulator. The iterative process requires a simple simulator to be used so that a classification decision can be reached quickly. The penalty for this is that the generated synthetic shadows are sometimes over-simplistic, making it necessary to use simple membership functions to accommodate inevitable modelling errors. If a more complex sonar simulator model was considered, these functions could be made much more severe, resulting in a larger class separability and a greater classification rate. Another possibility would be to adapt the model to applications such as Synthetic Aperture Sonar(SAS) where the line-of-sight assumptions used in the simulation model would have more credibility.

ACKNOWLEDGEMENTS

The authors wish to thank NATO Saclant Research Centre, DRDC-Canada and GESMA for the provision of data.

REFERENCES

- [1] J.BELL, G. ELSTON, S.REED, "Sonar Image Synthesis Techniques and Applications", Past, Present and Future Acoustics, Proceedings of Institute of Acoustics , Vol. 24(2), 2002
- [2] G.J.Dobeck, J.C.Hyland and L.Smedley, "Automated detection/classification of sea mines in sonar imagery." *Proc. SPIE-Int. Soc. Optics*, 2079:90-110,1997.
- [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] G.J. Dobeck, "Algorithm fusion for automated sea-mine detection and classification", *Proc MTS/IEEE Oceans 2001*, Vol 1, pp 130-134

Vol.26. Pt.5. 2004 (Sonar Signal Processing)

Vol.26. Pt.6. 2004 (Bio Acoustics)

- [5] J.A. Fawcett, "Image-based classification of side-scan sonar detections", CAD/CAC conference, Halifax, Nov. 2001
- [6] M.Mignotte, C.Collet, P.Perez and P.Bouthemy, "Hybrid Genetic Optimisation and Statistical model-based approach for the classification of shadow shapes in side-scan sonar imagery" *IEEE Trans Patt Anal Mach Int*, 22(2) 129-141
- [7] S.REED, Y.R.PETILLOT and J. BELL, "An Unsupervised Approach to the Detection and Extraction of Mine Features in Sidescan Sonar", *IEEE J. Oceanic Eng.* 28(1), pp. 90-105
- [8] S.REED, Y.PETILLOT, J. BELL, "A Model-Based Approach to the Detection and Classification of Mines in Sidescan Sonar", *Applied Optics*, 23(2), pp 237-246, Jan 2004
- [9] D.Huttenlocher, G.Klanderman and W.Rucklidge, "Comparing images using the Hausdorff distance", *IEEE Trans. Pattern Anal. Mach. Int.*, 15(9), pp. 850-863, Sept. 1993