Unsupervised Segmentation of Object Shadow and Highlight using Statistical Snakes

Scott Reed, Judith Bell, Yvan Petillot

June 25, 2001

Abstract

This paper deals with the detection and analysis of objects in side-scan sonar images for target identification. The initial step for any identification is a good segmentation of the scene and a reliable extraction of the shadow of the object. The literature reports various techniques for detecting the shadow of objects on the seabed, the most popular of which is based on the use of Markov Random Field (MRF) models. We introduce in this paper, a new technique based on a statistical snake to perform this operation. This technique relies on the modeling of the shadow and non-shadow regions as probability density models expressed as exponential laws (Gaussian, Rayleigh, Poisson,..) which are well suited to the physics of sonar sensing. The algorithm determines the snake shape that optimally separates the scene into regions described by the same probability density function utilising a Monte Carlo Markov Chain (MCMC) method to drive the snake. Moreover, as in classical snake techniques, the shape can be constrained via an intrinsic energy term enabling a model-based search for man-made objects in the image. The problems introduced by directional seabeds are discussed and the use of directional filters to extract further information investigated. This technique has been tested on both real and synthetic images and its effectiveness clearly demonstrated. This work is the first building block of a complete target identification system currently under development.

1 Introduction

Side-scan Sonar Imagery has long been a field of intense research interest with military and civilian applications. One such application is in the area of mine detection and classification where traditionally a human operator would be required to carry out the analysis based on their expert knowledge. With the development of Autonomous Underwater Vehicles (AUV's), there is now an abundance of side-scan images to be studied underlining the need for an automatic system to either replace or aid the human expert in the analysis . Unsupervised techniques will also allow analysis to be carried out while the data is being collected, enabling mission plans to be changed depending on the data collected. Possible techniques are considered in [1] where the geometrical features of the objects are extracted using the Hough Transform to be later used in a classification process. The classification is then conducted by considering a variety of techniques such as an expert system, a neural network and statistical analysis. Dobeck et al. [2] implement a detection and classification system which uses a non linear matched filter to detect mine like regions after which both a neural network based classifier and an optimal discriminating filter classifier are used together to classify the mine. A wide variety of features are measured and an optimal subset chosen to improve the probability of classification. Doherty et al. in [3] use an adaptive thresholding algorithm coupled with an adaptive averaging technique to detect mine-like objects after which an attribute-based decision tree is used to classify the objects. In [4] an expert system based approach is implemented which detects and classifies mine like targets based on the property of the elastic response of the return sonar signal where information such as the time delay and frequency are used to discriminate between actual mines and mine-like objects. A wavelet based classification scheme is used in [5] where wavelet packets are used to decompose the acoustic backscattered signal into several sub-bands. A linear predictive coding (LPC) model is then fitted to each sub-band signal to extract the features which are used in a neural network classifier. Another detection and classification system is implemented in [6] where background clutter is suppressed using an adaptive filtering technique whilst preserving the target signature. This improved the detection algorithm allowing classification to be carried

out on a set of features taken from pre and post filtered data. Mignotte et al. argue in [7] and [8] that the echo features are often difficult to use due to speckle noise and so utilise an unsupervised MRF model to segment the image into shadow and non-shadow regions after which a template [9] is used to highlight the object shadow. Based on the degree of fit of the template, the object is labelled as man-made or not man-made. The approach described in this paper uses an iterative snake algorithm driven by an MCMC model that highlights the shadow of the object based on statistical considerations of the distributions inside and outside the snake. This method is successful in that it is unsupervised, very fast and robust to noise. The use of a snake based method instead of a template allows the model to be generalised for wide spread applications while the MCMC technique ensures that the model can escape false maxima when searching for the optimal solution. This will form the basis of a classification system which will use the extracted shadow to decide if the object is man-made or not. The report is split into the following sections. Section 2 will describe the statistical snake and MCMC models, section 3 will show the results gained on real and synthetic data and section 4 will lay out the plans for future research.

2 Mathematical Theory

2.1 Introduction

The technique described in this section is one based on the research conducted in [10] where snakes are used under a statistical framework to segment noisy images. Whilst many region based techniques are often computer intensive, the technique described here is very fast, allowing much of the computer intensive calculations to be computed before the segmentation commences. This is possible when the statistical laws present in the image can be described by the exponential family, making the approach applicable in a wide range of physical situations. For example, in the field of side-scan imagery, it has been shown that the Rayleigh distribution offers a reasonable estimate of the observed signal for a variety of seabed types. Regions of shadow are also often seen to be described by a a Gaussian distribution where the only noise present is assumed to be due to electronic noise[7].

2.2 The Statistical Snake

It is assumed that the observed scene s is composed of two areas, the target zone which we wish to segment and the background, which includes the rest of the image. If we consider the image $s = s(x, y)$ which is composed of $N_x \times N_y$ pixels. The target's grey levels $\mathbf{t}$ and the background pixels $\mathbf{b}$ are assumed uncorrelated and have N_t and N_b pixels respectively. Both regions are described by a probability density function(pdf) p^{μ_t} and p^{μ_b} where μ_t and μ_b are the parameters of the two pdf's. We define a binary window function $\mathbf{w} = \{w(x, y) \mid (x, y) \in [1, N_x] \times [1, N_y]\}$ which defines the shapes of the snake at any given time. Defining $w(x, y)$ to be equal to 1 or 0 everywhere, the image becomes composed of two regions $\Omega_t = \{(x, y) \in w(x, y) = 1\}$ and $\Omega_b = \{(x, y) \in w(x, y) = 0\}$ so that the observed image can be viewed as the sum of two components

$$s(x, y) = t(x, y)w(x, y) + b(x, y)[1 - w(x, y)] \tag{1}$$

where $t(x, y)$ and $b(x, y)$ are values drawn from their respective probability distributions. Without any additional a priori knowledge, $\mathbf{w}$ is chosen by maximising the likelihood

$$P[\mathbf{s} \mid \mathbf{w}, \mu_t, \mu_b] = P(\chi_t \mid \mu_t)P(\chi_b \mid \mu_b) \tag{2}$$

where

$$P(\chi_u \mid \mu_u) = \prod_{(x,y)\in\Omega_u} p^{\mu_u}[s(x, y)] \tag{3}$$

where the likelihood is expressed as a product of probabilities as the distributions are assumed uncorrelated and $\chi_u = s(x, y) \mid (x, y) \in \Omega_u (u = t \ or \ u = b)$.

The likelihood function in 2 depends on the parameters of the probability functions as well as the template $\mathbf{w}$. The parameters μ_u where $(u = t \ or \ u = b)$ are a priori unknown and are in general of no interest. However, for segmentation to be possible, it is necessary to compute these values. These parameters can be calculated using a variety of techniques such as a Maximum A

Posteriori approach but for simplicity, a Maximum Likelihood approach is utilised as no a priori knowledge on the values of the parameters is available.

2.3 The Rayleigh Distribution

The authors in [10] showed that for the Rayleigh distribution one can obtain the maximum likelihood segmentation by maximising

$$l(\mathbf{s}, \mathbf{w}) = N_t log[F_t(\chi_t, N_t)] + N_b log[F_b(\chi_b, N_b)] \tag{4}$$

where $\chi_u = s(x,y) \mid (x,y) \in \Omega_u (u = t \; or \; u = b)$ as before and F_u is a function depending only on χ_u and N_u, the number of pixels in region u and not on the pdf parameters. Determining the form of function F_u is carried out by considering the log of the probability function being used and injecting the maximum likelihood estimates of the pdf parameters back into the formula.

The Rayleigh Distribution is often used to model the sonar return for side-scan sonar. The distribution is described by

$$p(s(x,y)) = \frac{s(x,y)}{\sigma^2} exp(-\frac{s(x,y)^2}{2\sigma^2}) \tag{5}$$

where σ^2 is the variance of the distribution. If we consider two regions, t the target region describing the area inside the snake and b, the background region outside the snake, we can determine a log likelihood expression

$$\begin{aligned} l(\mathbf{s}, \mathbf{w}, \sigma_t^2, \sigma_b^2) \quad = \quad & \sum_{(x,y)\in\Omega_t} log[s(x,y)] - \sum_{(x,y)\in\Omega_t} log[\sigma_t^2] - \sum_{(x,y)\in\Omega_t} \frac{s(x,y)^2}{2\sigma_t^2} + \\ & \sum_{(x,y)\in\Omega_b} log[s(x,y)] - \sum_{(x,y)\in\Omega_b} \log[\sigma_b^2] - \sum_{(x,y)\in\Omega_b} \frac{s(x,y)^2}{2\sigma_b^2} \end{aligned} \tag{6}$$

Using the maximum likelihood estimates of σ_u^2 where $u = t$ or $u = b$, the authors in [10] show that the log likelihood simplifies to

$$l(\mathbf{s}, \mathbf{w}) = -N_t log[\frac{1}{N_t} \sum_{(x,y)\in\Omega_t} s(x,y)] - N_b log[\frac{1}{N_b} \sum_{(x,y)\in\Omega_b} s(x,y)] \tag{7}$$

An iterative scheme was implemented where the likelihood was computed at each iteration. This process was very fast as the bulk of the calculations for $F_u(\chi_u, N_u)$ were computed before the iterative process began. For details on this procedure see [10].

For the purpose of extracting the object shadow, both the shadow regions and the non-shadow regions were modeled as being rayleigh distributed.

2.4 Justification in Using the Rayleigh Distribution

For side-scan sonar images, one can assume that the signal is dominated by seabed reverberation and that this signal is composed primarily of diffuse scattered energy, an assumption which only breaks down at near normal incidence when the seabed is directly below the tow-fish. In these circumstances, it is expected that the side-scan sonar images can be approximated as having Rayleigh statistics. Research into the distributions of sonar reverberation has been investigated by [11] and [12] where it was observed that the observed probability density function was dependent on the sea floor roughness. After investigating a range of sediment types and sea floor structures, the probability density functions were observed to follow Rayleigh distributions for isotropic regions of the sea bed. This was again confirmed by [13] where the Rayleigh distribution was shown to describe the statistics of relatively flat sea beds very closely. [14] also showed, using authentic and synthetic side-scan sonar data, that the Rayleigh distribution could be used to accurately describe a variety of sediment types and sea bed roughness. For more complex sea bed types, such as sand ripples, [13] showed that multi modal Rayleigh pdf's were needed to accurately describe the distribution. The author successfully showed a 2-mode and a 4-mode distribution accurately fitting the observed statistics. This assumption was tested on a real data image with multiple mines as seen in figure 1. To assess the statistics of the shadow/non

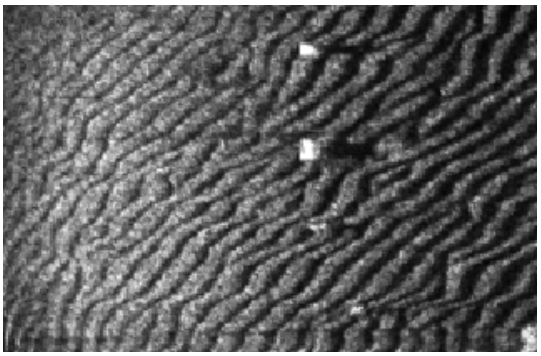

Figure 1: An area of seabed with multiple objects

Figure 2: 2 class segmentation using hierarchical MRF model

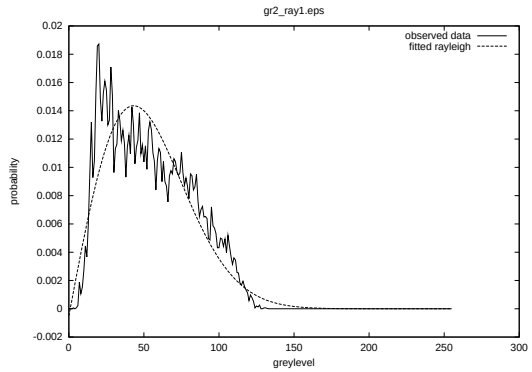

Figure 3: Fitting a Rayleigh Distribution to the Shadow pixels

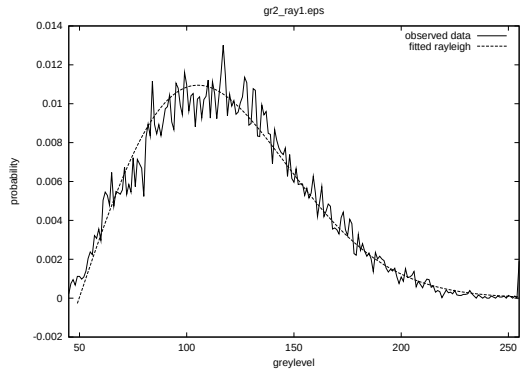

Figure 4: Fitting a Rayleigh Distribution to the Non-Shadow pixels

shadow regions it was necessary to first segment the image as seen in figure 2. This was achieved by using a hierarchal Markov Random Field model where the model parameters were determined in an unsupervised manner. Once partitioned, an assessment of the statistics was possible.

As figures 3 and 4 show, the shadow and non-shadow regions are described well by the Rayleigh distribution. This distribution was therefore used to model the statistics of the side-scan sonar images.

2.5 Monte Carlo Markov Chain Methods

Bayesian techniques are often used in analysis techniques where we wish to deduce $\mathbf{X}$, the unknown, often multidimensional 'ideal' information given the observed noisy data $\mathbf{Y}$. As stated in the previous chapter, Bayes theorem allows us to perform a maximum a posterior calculation of $P(X \mid Y)$ using both the likelihood and any a priori information that is available. Theoretically, this calculation can be carried out directly by considering the expectation value as shown in equation 8.

$$E[f(\mathbf{X}) \mid \mathbf{Y}] = \frac{\int f(\mathbf{X}) P(\mathbf{Y} \mid \mathbf{X}) P(\mathbf{X}) d\mathbf{X}}{\int P(\mathbf{Y} \mid \mathbf{X}) P(\mathbf{X}) d\mathbf{X}} \tag{8}$$

where $f(\mathbf{X})$ is a function of the unknown 'true' data. These integrations are in practice very difficult and for data of high dimensionality, usually analytically impossible therefore requiring the use of numerical techniques. This section will be concerned with MCMC sampling techniques and will discuss how these methods can be used successfully in our specific application.

2.6 Markov Chains

The Markov Chain is a very flexible method for generating random variables from a prescribed density. The chain is is defined as a set of random variables $x_1, x_2, x_3, ...x_N$ where variable x_{N+1} is influenced by x_N only

$$p(x_{N+1}|x_1, x_2, x_3, ...x_N) = p(x_{N+1}|x_N) \tag{9}$$

For a Markov Chain to be useful in sampling techniques it must first demonstrate two properties : there must be an unique invariant distribution and the Markov Chain must be ergodic. The ergodic property is dependent on the Markov Chain being irreducible which states that it is possible for the chain to move to any other state in parameter space within a finite period. If $p_{ij}^{(n)}$ is the probability of going from state i to state j in n steps, it is required that $\exists n : p_{ij}^{(n)} > 0 \; \forall i, j$. The chain should also be aperiodic in that it is possible to return to any state at any time rather than with a fixed period. The ergodic property ensures that regardless of the initial starting point of the chain, it will eventually converge onto the invariant distribution as time $N \to \infty$.

2.7 Monte Carlo Integration

Once the Markov chain has converged onto the invariant distribution it is possible to approximate the integral described in equation 4 to provide a maximum a posterior solution to the problem. The invariant distribution of the Markov chain is assumed to be the same as the posterior distribution of the problem and so by drawing N samples of $\mathbf{X}_i$, $0 < i \leq N$ from the posterior distribution, the integral can be approximated by

$$E[f(\mathbf{X}) \mid \mathbf{Y}] \approx \frac{1}{N} \sum_{i=1}^{N} f(\mathbf{X}) \tag{10}$$

where the complex integral has been replaced by a summation where the number of samples summed over depends on the accuracy to which the solution is required. To ensure the chain has successfully converged onto the invariant distribution it is necessary to allow a 'burn in' period where the first M samples of the chain are discarded. There are a number of different sampler methods possible but in this instance, the Metropolis Hastings algorithm was implemented.

2.8 The Metropolis Hastings Algorithm

The Metropolis Algorithm is a technique for producing a random sequence from a given density and was first developed by Metropolis for computing the properties of interacting particles. The general form of the algorithm assumes the probability density is described by $y = p(x)$ where $x \in \Theta$ and Θ is the parameter space which the algorithm searches through by means of a random walk. It is also important to note that the density need not be normalized. If X_i is the i^{th} element, the next element is generated by adding a random perturbation ς_i so that

$$Y_i = X_i + \varsigma_i \tag{11}$$

where ς_i is drawn from a proposal distribution $s(\varsigma)$. For this report, the proposal distribution was modeled as a Gaussian distribution with a decreasing standard deviation. This was used to encourage large changes at the start of the iterative scheme so that a rough segmentation could be reached. As the variance decreases, the changes to the model become smaller to allow the model to 'fine tune' its parameter values to obtain the maximum a posterior solution. The algorithm can either accept the new value or reject it after which it returns to the old value X_i so that the process can be repeated. The Metropolis Algorithm decided on whether to accept or decline a new position by considering the acceptance function

$$A(X_i, Y_i) = min(1, Q(X_i, Y_i) \tag{12}$$

where

$$Q(X_i, Y_i) = \frac{p(Y_i)T(Y_i|X_i)}{p(X_i)T(X_i|Y_i)} \tag{13}$$

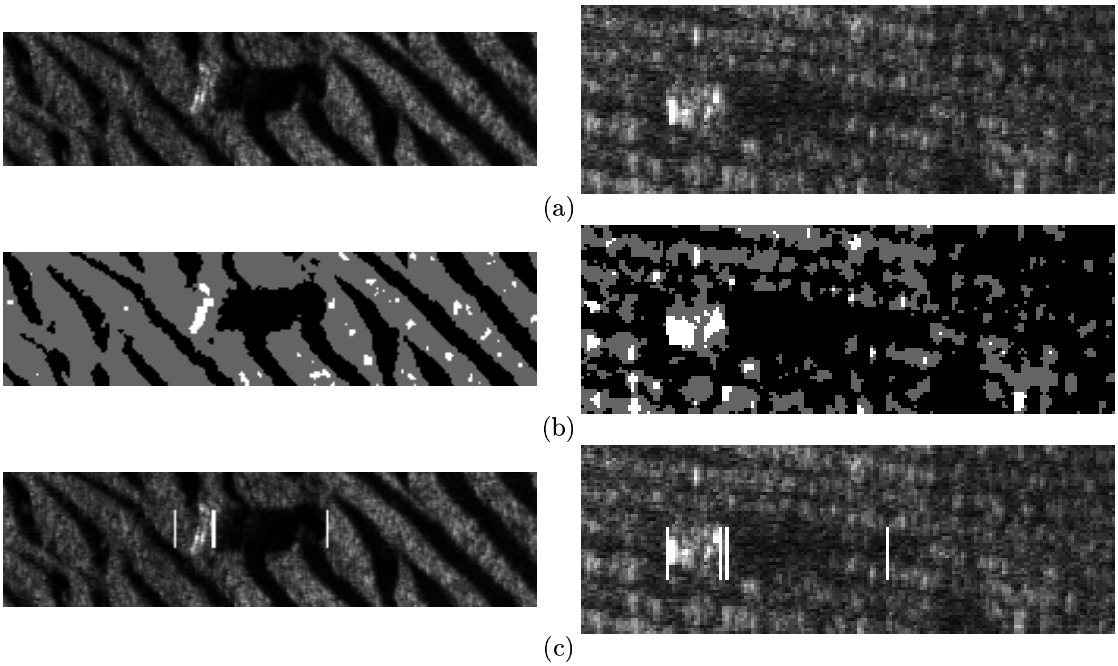

Figure 5: (a) sonar images of 2 mines(b) mrf image of images in (a), split into 3 classes (c)detection algorithm looking for highlight and shadow

where the conditional probability density $T(Y_i|X_i)$is identical to the proposal density $s(\varsigma)$. This means $T(X_i|Y_i)$ is given by $s(-\varsigma)$. If the proposal distribution is symmetrical about the origin, the the acceptance equation 13 reduces to

$$Q(X_i, Y_i) = \frac{p(Y_i)}{p(X_i)} \tag{14}$$

where the probability of acceptance is independent of the proposal density $s(\varsigma)$. The value of $A(X_i, Y_i)$ is then compared to a random number generated from the range $[0, 1]$ to assess whether the new proposal is accepted or rejected.

2.9 Initialising the Markov Chain

Although the ergodic nature of the Markov chain assures convergence as time $t \to \infty$, for practical purposes a good starting point is essential for quick convergence. The results shown in this paper assume that the object has already been detected using an object detection algorithm therefore allowing the snake to search through a small image centered on the mine. An unsupervised detection system is being developed that first segments the image into 3 classes - object, background and shadow by using a simple Markov Random Field model. Based on this image, a detection algorithm is implemented which searches through the image using MCMC techniques, looking for clusterings of object pixels followed by clusterings of shadow pixels, therefore making use of the a priori knowledge of the geometry of the highlight/shadow regions. The model is flexible in that the clustering of shadow and object regions can be any size and distance apart therefore making it robust to a wide variety of conditions. Details of this model are not described here but the initial promising results can be seen in figure 5 where the shadow and object 'boxes' were initialised in the top left corner of the images with dimensions 5×5 pixels. Obviously it will be necessary for this algorithm to work on very large images and so it was deemed necessary to implement harsh initial conditions to test the algorithm. As can be seen, the results are very promising, offering a good starting point for segmentation of both the shadow and the highlight.

2.10 Operating within a Bayesian structure

Bayes theorem states

$$P(X \mid Y) = \frac{P(Y \mid X) \cdot P(X)}{P(Y)} \tag{15}$$

where $P(X \mid Y)$ is the posterior distribution and $P(Y \mid X)$ is the likelihood. The posterior distribution assesses the probability of the true data given the observed data Y based on a combination of the likelihood function and the prior probabilities $P(X)$ and $P(Y)$ [15]. These prior probabilities can have any form and are a method for adding additional information to the system i.e. imposing limits on the model parameter values. 2 priors are used in this report. The first P_1 is a regulating prior as in [10] whilst the second limits the angles between the snake segments to be $> 45^o$ and has the form

$$P_2(\mathbf{w}) = \begin{cases} c & \text{if } \theta_{i,i+1} > \frac{\pi}{4} \text{ and } \theta_{i,i+1} < \frac{7\pi}{4} \\ 0 & \text{otherwise } \forall i \end{cases}$$

where $\theta_{i,i+1}$ is the angle between segments i and $i+1$.

The optimal solution is therefore obtained by maximising $J(\mathbf{s}, \mathbf{w})$ which is the log of the posterior probability.

$$J(\mathbf{w}, \mathbf{s}) = (1-\lambda) l(\mathbf{s}, \mathbf{w}) + \lambda log[p_1(\mathbf{w})] + log[P_2(\mathbf{w})] \tag{16}$$

and λ describes the weighting between the internal regulating energy and the external energy due to the statistics of the image.

2.11 Using Filters

Filter provide a method for gaining additional information about the properties of an image. In particular, Sobel Filters provide a simple method for providing textural information about the image by producing edge fields which describe the variation in the grey level with respect to changes in the x and y direction. Convolving the original image with the two directional Sobel filters produced two edge fields, whose statistics were also considered to segment the objects shadow assuming that the shadow and non-shadow regions were described by different edge fields. The Sobel statistical snake had to maximise the following posterior distribution

$$J(\mathbf{w}, \mathbf{s}) = \frac{(1-\lambda)}{2} l_{Sobel_X}(\mathbf{s}, \mathbf{w}) + \frac{(1-\lambda)}{2} l_{Sobel_Y}(\mathbf{s}, \mathbf{w}) + \lambda log[p_1(\mathbf{w})] + log[P_2(\mathbf{w})] \tag{17}$$

where l_{Sobel_X} and l_{Sobel_Y} are the log likelihoods of the two edge fields which are assumed to be Gaussian distributed. The exact form of the Gaussian likelihood term can be seen in [10]. It should also be noted that the relative weights of the two likelihood terms are kept equal in the absence of any a priori information.

3 Results

Figure 6 lays out a summary of the model used to segment the objects shadow. This model was then tested on real and synthetic data.

3.1 Results with Real Images

Both the standard statistical snake and the Sobel-based snake where tested on real images containing mines. Very good segmentation of the shadow was achieved every time. The visual results for 3 of these mines can be seen in figure 7.

As figure 7 shows, very good segmentation results were obtained with both techniques. However, it can be seen that the snake sometimes moves into the sand dune shadows due to the posterior probability having no term which can distinguish between the two shadow types. Whilst this is obviously not a problem for objects which lie on flat seabeds, a completely general unsupervised segmentation system should be able to deal with any sediment type. This problem is discussed further in the last section. Based on the results shown, it is very difficult to differentiate between the efficiency of the two models and indeed it is likely that a choice of model would depend on the case involved.

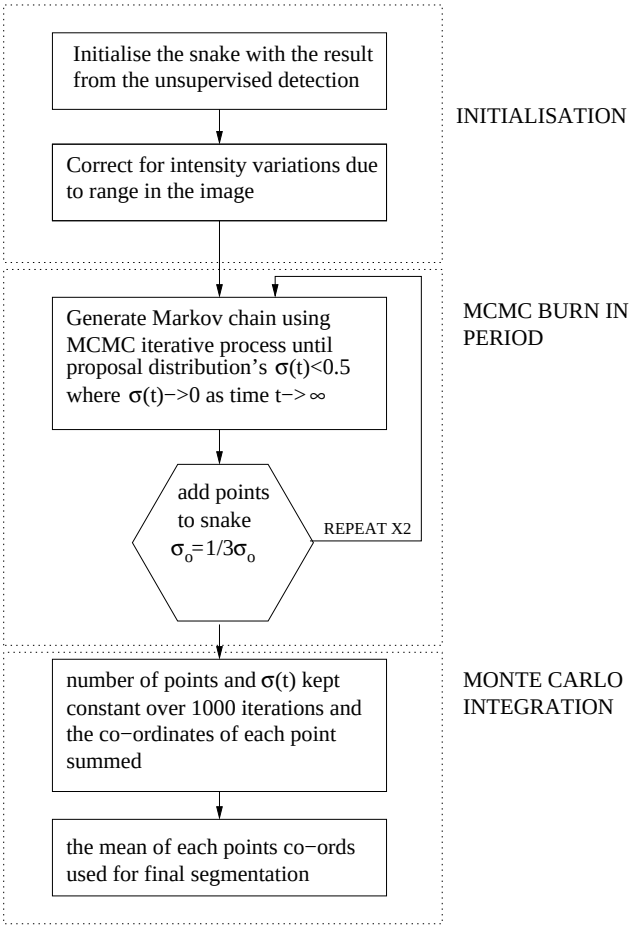

Figure 6: The Metropolis-Hastings Algorithm

3.2 Ground truth Results on Synthetic Data

A quantitative analysis of the snake algorithm was carried out using a side scan sonar simulator developed in [14] which allowed ground truth results of the objects shadow to be obtained. Ground truth results were considered for 4 objects (2 cylinder and 2 spheres) lying on both sandy and sand dune seabeds. The results for 1 of the cylinders on both of the sediment types are shown here in the form of two measured probabilities which are defined below.

$$prob\ of\ detection(PD) = \frac{number\ of\ shadow\ pixels\ inside\ snake}{total\ number\ of\ shadow\ pixels} \tag{18}$$

$$prob\ of\ false\ alarm(PFA) = \frac{number\ of\ non-shadow\ pixels\ inside\ snake}{total\ number\ of\ pixels\ inside\ snake} \tag{19}$$

The two images were corrupted with varying degrees of noise where the noise was drawn from a Gaussian distribution of specified mean m and standard deviation σ. An example of the cylinder on the two sediment types and the resulting segmentations can be seen in figure 8.

The probability of detection(PD) and probability of false alarm(PFA) were computed for differing amounts of noise for the same images and the results shown in table 1 where the results shown are the average over 10 experiment runs.

As expected, the probability of detection(PD) and the probability of false alarm(PFA) remain virtually constant over the range of noise for the case of the sandy seabed where a confident segmentation is achieved in every case. For the scenario where the cylinder lies on the sand dunes, the model often becomes confused and attempts to segment both the object shadow and the sand dune shadow as can be seen in the decrease in the PD values and the increase in the PFA values. Figure 9 shows the snake successfully segmenting the shadow on a noisy image

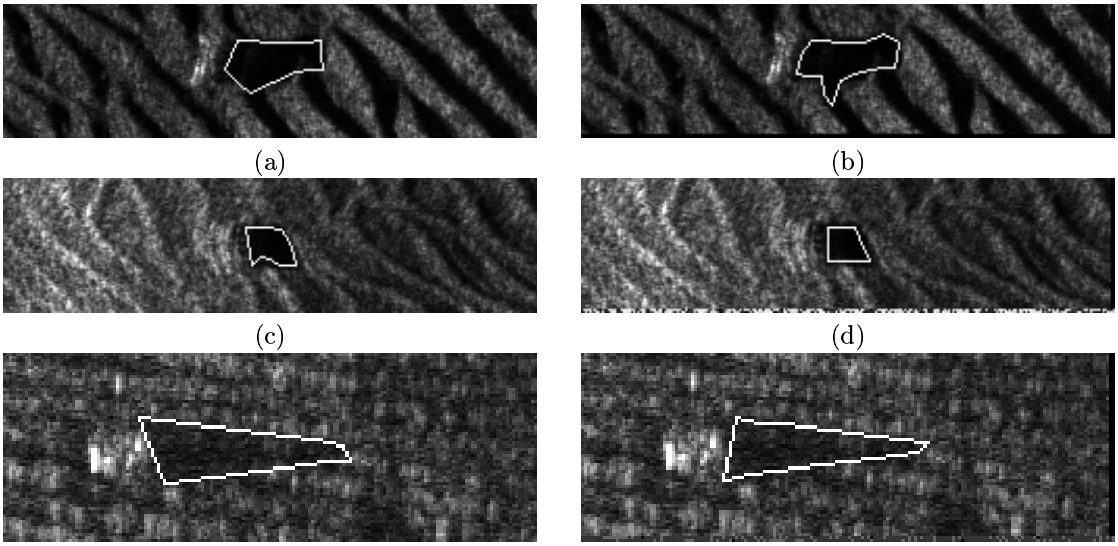

Figure 7: (a)mine1 shadow segmentation with stat snake. (b)mine1 shadow segmentation with sobel snake. (c)mine2 shadow segmentation with stat snake. (d)mine2 shadow segmentation with sobel snake. (e)mine3 shadow segmentation with stat snake. (e)mine3 shadow segmentation with sobel snake.

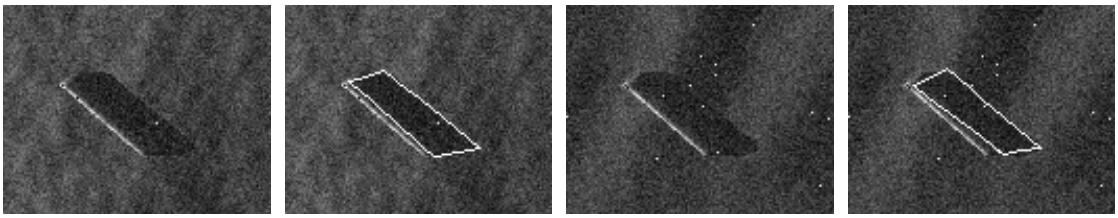

Figure 8: (a) cylinder on a sandy seabed corrupted with noise added drawn from distributon with m=40 and $\sigma = 15$. (b) segmentation of shadow from (a). (c) cylinder on a sand dune seabed with noise added drawn from distributon with m=40 and $\sigma = 15$. (d) segmentation of shadow from (c).

noise values		*sandy sea bed*		*sand dune sea bed*	
m	σ	PD	PFA	PD	PFA
20	10	0.96	0.22	0.82	0.34
20	20	0.96	0.19	0.66	0.55
20	30	0.95	0.18	0.52	0.41
20	40	0.92	0.18	0.53	0.42
40	10	0.96	0.22	0.72	0.54
40	20	0.97	0.20	0.64	0.60
40	30	0.95	0.22	0.72	0.43
40	40	0.95	0.20	0.50	0.35
40	50	0.96	0.15	0.55	0.43

Table 1: Quantitative assessment of the Robustness to Noise

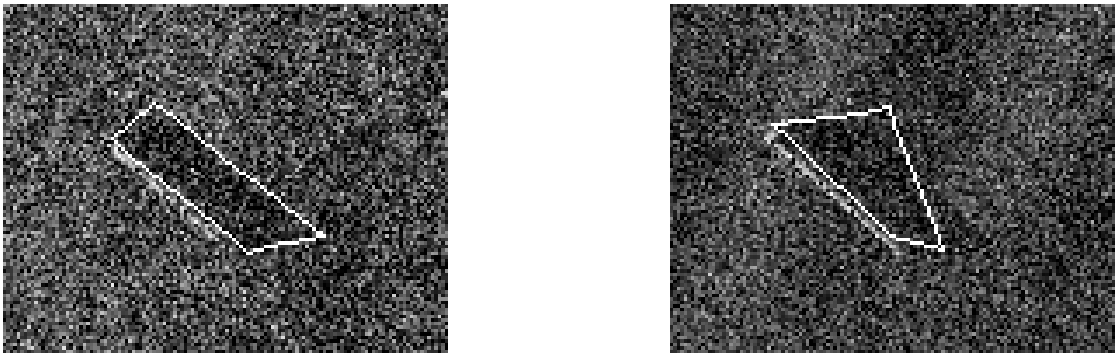

(a) (b)

Figure 9: (a) successful segmentation of the object's shadow on a sandy seabed with noise added drawn from a distribution of m=40 and σ=50. (b) unsuccessful segmentation of the shadow when the object lies on a sand dune seabed with noise drawn from a distribution with m=40 and $\sigma = 50$. The snake attempts to highlight both the object and dune shadows.

where the object lies on the sandy seabed but failing in the presence of the dune shadows. This again emphasises the need for the model to include more a priori information so that it can distinguish between the object and background shadows. However, it should be stressed that the example shown in this section is a complex case where the dune shadows are actually larger than the object's shadows.

4 Discussion and Conclusion

This paper has presented an unsupervised statistical snake technique which has been shown to be efficient and robust at segmenting object shadows in side-scan sonar. The algorithm is particularly suited to side-scan sonar images as the distributions can often be described by a Rayleigh distribution allowing most of the required calculations to be computed before the iterative process.

To ensure that the maximum a posterior solution is found, a MCMC algorithm is implemented where a good initial starting value for the Markov chain is produced by an unsupervised object detection system. While this model is still in its infant stages, the initial results look very promising.

A summary of results gained on real and synthetic images is given, with very promising results obtained on all occasions except when the object is seen to lie on large sand dunes where the shadows due to the background can lead to an inaccurate segmentation. Further research into this problem is currently being carried out where it is hoped that the Markov Random Field model used in the detection model can be extended to differentiate between the dune shadows and the object shadows.

In conclusion, a segmentation algorithm is offered here which will become the basis of a completely unsupervised statistical model which will both detect the mine, and extract the object highlight and its shadow.

References

[1] B.Zerr. *Automatic Classification of Objects in Sonar Images, GESMA.*

[2] Gerald J.Dobeck, John C.Hyland, and Le'Derick Smedley. Automated detection/classification of sea mines in sonar imagery. In *Proceedings SPIE-Int. Soc. Optics*, volume 2, 1997.

[3] M.F.Doherty et al. *Side Scan Sonar Object Classification Algorithms, Lockheed Missiles and Space Co.*

[4] John Foster and Myla White. Pattern recognition expert system for mine classification and detection sonar. In *IEEE ProcSouthEast Conf*, 1989.

[5] Mahmood R.Azimi-Sadjadi et al. Underwater target classification using wavelet packets and neural networks. *IEEE Transactions on Neural Networks*, 11(3):784–794, May 2000.

[6] Tom Ardgides, Manuel Frenandez, and Gerald Dobeck. Adaptive 3-dimensional range-crossrange-frequency filter processing string for sea mine classification in side-scan sonar imagery. *SPIE Proc.*, 3079, 1997.

[7] Max Mignotte et al. Sonar image segmentation using an unsupervised hierarchical mrf model. *IEEE Transactions on Image Processing*, 9(7):1216–1231, 2000.

[8] Max Mignotte et al. Three class markovian segmentation of high resolution sonar images. *IEEE Transactions on Computer Vision and Image Understanding*, 76(3):191–203, 1999.

[9] Max Mignotte et al. Hybrid genetic optimization and statistical model-based approach for the classification of shadow shapes in sonar imagery. *IEEE Transactions on Pattern Analysis and Machine Intelligience*, 22(2):129–141, 2000.

[10] Christophe Chesnaud, Philippe Refregier, and Vlady Boulet. Statistical region snake-based segmentation adapted to different physical noise models. *IEEE Transactions on Pattern Analysis and Machine Intelligience*, 21(11):1145–1157, 1999.

[11] T.K.Stanton. Sonar estimates of seafloor microroughness. *Journal of the Acoustical Society of America*, 75(3):809–818, 1984.

[12] D.Alexandrou and C.de Moustier. Evaluation and verification of bottom acoustic reverberation statistics predicted by the point scattering model. *Journal of the Acoustical Society of America*, 91(3):1403–1413, 1992.

[13] Kenneth W.Stewart et al. Quantitative seafloor characterization using a bathymetric sidescan sonar. *IEEE Transactions on Oceanic Engineering*, 19(4):599–610, 1994.

[14] Judith Bell. *A Model for the Simulation of Sidescan Sonar*. PhD thesis, Heriot-Watt University, August 1995.

[15] Brian Calder. *Bayesian Spatial Models for SONAR Image Interpretation*. PhD thesis, Heriot-Watt University, September 1997.