Fractal dimension, wavelet shrinkage, and anomaly detection for mine hunting*

J. D. B. Nelson[†] and N. G. Kingsbury[‡]
james@stats.ucl.ac.uk, ngk@eng.cam.ac.uk

Abstract

An anomaly detection approach is considered for the mine hunting in sonar imagery problem. We exploit previous work that used dual-tree wavelets and fractal dimension to adaptively suppress sand ripples and a matched filter as an initial detector. Here, lacunarity inspired features are extracted from the remaining false positives, again using dual-tree wavelets. A one-class support vector machine is then used to learn a decision boundary, based only on these false positives. The approach exploits the large quantities of ‘normal’ natural background data available but avoids the difficult requirement of collecting examples of targets in order to train a classifier.

1 Introduction

For more than a decade supervised approaches have received much attention in the mine hunting literature [2, 3, 19]. A common approach outlined in Figure 1 requires that the ground truth (location of the targets in the training data) is available. Firstly, a detector is applied to eliminate any areas of the image which can be easily distinguished from the targets. A good example of such a detector is the matched filter, first proposed for mine hunting by Dobeck et al [4]. This initial screening or detection stage is designed to accept a large number of false positives (natural background) in order to capture all of the true positives (the mines). The ground truth is then used to select and label the positives as either target or background. After extracting features, a proportion of these positives are then used to train a classifier. The remaining positives are then used to test the classifier. A common complaint against supervised approaches is that, although they might perform well on the data used to train the system, there is no guarantee that they will be able to generalise well to unseen data and especially unseen targets.

*This research was supported by the EPSRC (grant number EP/H012834/1) and the UK Ministry of Defence, as part of the University Defence Research Centre on Signal Processing.

[†]Department of Statistical Science, University College London, London, UK

[‡]Signal Processing and Communications Group, University of Cambridge, Cambridge, U.K.

Figure 1: Common detector/classification scheme. Detector finds (true and false) positives. **Training phase:** a binary classifier is trained on features extracted from both true and false positives. **Testing phase:** the detector is applied, features are extracted from positives, which are then classified by a learned decision function.

Unsupervised approaches have also been attempted. Amongst the most promising is the work of Reed et al. [16]. They used unsupervised Markov random field based detection to segment the image into shadow, seabed, and object highlight regions. Cooperating statistical snakes were used to extract highlight and shadow regions for object classification.

The approach explored in this paper falls somewhere in-between the supervised and unsupervised methods and is summarised in Figure 2. Like the supervised approach, a (matched filter) detector is used to prune the data. However, only false positives are used to train the classifier—in this case, a one-class support vector machine. We also implement previous work [13] that uses fractal dimension to adaptively suppresses any sand ripples that may be present prior to the matched filter.

Figure 2: Proposed detector/classification scheme. Pre-processing suppresses ripples. Detector finds (true and false) positives. **Training phase:** a unary classifier is trained on features extracted from false positives only. **Testing phase:** Pre-processing suppresses ripples. Detector is applied, features are extracted from positives, which are then classified by a learned decision function (true positives are anomalies).

Unlike the fully supervised approach, which uses binary or multi-class learning, the training stage does not require any examples of targets. Instead it only requires examples of background data, and thus avoids the major difficulty of capturing examples of real threats in any significant quantity. It also potentially means that unanticipated, or even hitherto unknown, types of threats can be detected and will in theory allow the system to evolve according to the environmental conditions. The approach regards the background, or natural seabed, as normal whereas the unknown targets are treated as anomalies.

In the feature extraction phase, some novel features will be presented that are loosely based on the idea of lacunarity. Mandelbrot [10] first introduced the concept of lacunarity to measure deviations from translational statistical invariance. Small values indicate the presence of translational invariance in the surface texture of interest. It is assumed here that the presence of a mine will cause an anomalous increase in local lacunarity. Unlike the usual box counting approach, we compute a lacunarity based measure using the dual-tree complex wavelet transform. This allows us to efficiently compute and compare deviations in the statistics of regions within each scale level and direction.

Section 2 summarises previous work [13] on the first two steps, ripple suppression and detection, of the approach explored here. Section 3 describes the lacunarity-based features. Section 4 discusses the results.

2 Ripple Suppression and Initial Mine Detection

2.1 Statistical self-similarity

Of particular interest for this application are statistically self-similar processes. That is $f : \mathbb{R}^2 \mapsto \mathbb{R}$, such that

$$\mathbb{E}[f(\gamma \mathbf{x})f(\gamma \boldsymbol{\xi})] = \gamma^{2H} \mathbb{E}[f(\mathbf{x})f(\boldsymbol{\xi})] , \quad (1)$$

where $H \in [0, 1]$ for some $\gamma \in \mathbb{R}$. Statistical self similarity is one of the key properties that informally defines a (stochastic) fractal process [15]. It is satisfied by processes with power spectra $\mathcal{P}f$ that follows the power law decay:

$$(\mathcal{P}f)(\boldsymbol{\omega}) \propto \|\boldsymbol{\omega}\|^{-2(H+1)} . \quad (2)$$

In this case, the fractal dimension of f is $\phi \triangleq 3 - H$. Statistical self-similarity is also satisfied by fractional Brownian surfaces. These are stochastic processes $B_H : \mathbb{R}^2 \mapsto \mathbb{R}$, with initial condition $B_H(\mathbf{0}) = 0$, such that the local increments:

$$(\Delta B_H)(\mathbf{x}) \triangleq B_H(\mathbf{x} + \Delta \mathbf{x}) - B_H(\mathbf{x}) ,$$

are stationary Gaussian random fields with variance

$$\mathbb{E}[|(\Delta B_H)(\mathbf{x})|^2] \propto \|\Delta \mathbf{x}\|^{2H} .$$

Again, the fractal dimension is $(3 - H)$, and the power law (2) holds in an average power spectrum sense [15].

2.2 Wavelet transforms and fractal dimension

The wavelet transform of a surface $f : \mathbb{R}^2 \mapsto \mathbb{R}$ can be written as

$$(\mathcal{W}f)(k, m; \mathbf{x}) = 2^{-k} \int_{\mathbb{R}^2} f(\boldsymbol{\xi}) \overline{\psi_m(2^{-k}(\mathbf{x} - \boldsymbol{\xi}))} d\boldsymbol{\xi} ,$$

where ψ_m is a zero-mean mother wavelet with orientation indexed by m , and where k denotes the k th finest scale level ($k = 0$ being the scale of the original ‘pixel’ coordinates), and where $\mathbf{x}$ is the spatial location. We have (e.g. [13]) that

$$\log_2 \mathbb{E} [|(\mathcal{W}f)(k, m; \cdot)|^2] = 2k(H + 1) + C_{\psi_m, f} \quad (3)$$

holds for statistically self similar processes. The $C_{\psi_m, f}$ term is independent of scale k and the right-hand-side is a linear function of scale k . Hence, the exponent H can be computed by measuring the average slope, over each of the directional subbands m , of (3) via linear least squares regression. In practice the variance is approximated by the sample variance $\frac{1}{N} \sum_n |(\mathcal{W}f)(k, m; \mathbf{x}_n)|^2$. In [12], the authors proposed the use of the dual-tree complex wavelet transform (DTCWT) [9, 18] to estimate local and directional fractal dimension. The DTCWT is fast (decimated) and, with 6 strongly directional subbands, it has good directional selectivity. The stripe directions of the filters are oriented at $\{(30m - 15)^\circ\}_1^6$ in an anticlockwise direction from the horizontal (see [9, 18] for more details). Moreover, unlike other fast wavelet transforms, the DTCWT also has good shift invariance which ensures that the magnitudes of its complex coefficients remain stable and large near any singularities.

2.3 Measuring ripples with fractal dimension

In order to help distinguish between rippled and non-rippled patches of seabed, we consider a measure of local dual-tree wavelet root-energy, namely $E_{k,m}(\mathbf{x}) \triangleq |(\mathcal{W}f)(k, m; \mathbf{x})|$, at scale k , directional subband m , and location $\mathbf{x}$. The root-energy is plotted against the k th finest scale level in Figure 3 for a rippled and non-rippled seabed region. It can be observed that the non-rippled region gives rise to an approximate power law spectrum as in (2) whereas the rippled region contains spikes in more than one directional subband that invalidate the power law. That the relatively flat seabed follows a power law relationship corroborates the model of Pailhas et al. [14] who used fractional Brownian surfaces, discussed briefly in the previous section, to synthesise sonar imagery data of seabeds. From (3), a surface with a power law spectrum satisfies:

$$E_{k,m}(\mathbf{x}) \propto 2^{(4-\phi)k},$$

where ϕ (nominally $\in [2, 3]$) is fractal dimension. Hence, for a relatively flat seabed region, we arrive at the following bound condition for a non-rippled seabed:

$$\frac{1}{4} \leq \frac{E_{k,m}(\mathbf{x})}{E_{k+1,m}(\mathbf{x})} = 2^{\phi-4} \leq \frac{1}{2}, \quad \forall k. \quad (4)$$

Here, the root-energy at scale k is divided by the root-energy at scale $k + 1$ for each location and subband direction. This interscale wavelet energy ratio results in a value independent of scale level. Since the bounds are independent of k , this is a scale invariant condition. Since the ratio only involves two scale levels, it is akin to a frequency

localised measure of fractal dimension. Because of the space, direction, frequency localised (approximate) nature of this measure, and because we only want to shrink a wavelet coefficient if we have strong evidence that it contributes to a ripple region, the upper and lower bounds are relaxed somewhat in practice to some $\lambda_1 < \lambda_2$.

Conversely, at scale k , subband m , location $\mathbf{x}$, the wavelet coefficient $w_{k,m}(\mathbf{x}) \triangleq (\mathcal{W}f)(k, m; \mathbf{x})$ is deemed to contribute to rippled seabed region if, for some $\lambda_1 < \lambda_2 \in \mathbb{R}^+$:

$$R_{k-1,m}(\mathbf{x}) \triangleq \frac{E_{k-1,m}(\mathbf{x})}{E_{k,m}(\mathbf{x})} < \lambda_1 \quad (5)$$

or

$$R_{k,m}(\mathbf{x}) \triangleq \frac{E_{k,m}(\mathbf{x})}{E_{k+1,m}(\mathbf{x})} > \lambda_2. \quad (6)$$

In this way, we can distinguish between rippled and non-rippled regions. By considering the distance between the interscale energy ratio and the bounds, we also have a measure of how certain we are that a given region contains ripples.

2.4 Fractal-dimension-based complex wavelet shrinkage

Given an image f , a general wavelet shrinkage procedure can be summarised by:

- (i) Take wavelet transform: $w = \mathcal{W}f$
- (ii) Shrink wavelet coefficients: $w^\sim = \mathcal{S}w$
- (iii) Take inverse wavelet transform: $f^\sim = \mathcal{W}^{-1}w^\sim$

More specifically, for the ripple suppression application the shrinkage operation in the second step is designed to reduce or threshold any coefficients $w_{k,m}(\mathbf{x})$ that contribute the ripples. To this end, the shrinkage operator is applied to the wavelet coefficients via:

$$w_{k,m}^\sim(\mathbf{x}) = S_{k,m}(\mathbf{x})w_{k,m}(\mathbf{x}), \quad S: (k, m; \mathbf{x}) \mapsto [0, 1]. \quad (7)$$

The shrinkage function $S_{k,m}(\mathbf{x}) \triangleq S_{k,m}^-(\mathbf{x})S_{k,m}^+(\mathbf{x})$ adapts to scale, direction, and location according to the minimax functions:

$$S_{k,m}^-(\mathbf{x}) = \min \left(1, \max \left(0, \frac{R_{k-1,m}(\mathbf{x}) - \lambda_0}{\lambda_1 - \lambda_0} \right) \right) \quad (8)$$

$$S_{k,m}^+(\mathbf{x}) = \min \left(1, \max \left(0, \frac{\lambda_3 - R_{k,m}(\mathbf{x})}{\lambda_3 - \lambda_2} \right) \right) \quad (9)$$

These functions are plotted and explained in Figure 4. Note that this strategy differs from the standard wavelet shrinkage approach which shrinks wavelet coefficients that have small absolute values or are uncorrelated with respect to space or scale.

Figure 3: A patch of non-rippled seabed (a), gives rise to a wavelet amplitude spectrum (b) that follows a power law decay with respect to finer scales (it decays as frequency increases). Conversely, a rippled seabed patch (c) will invalidate the power law at one or more scales and subband directions (d).

2.5 Matched filter

Since objects protruding above the seabed tend to be more reflective than the sediment they return a higher intensity signal back to the sensor. Moreover, such objects will also block the signal from reaching the seabed behind them, thus creating a shadow region. This motivated Dobeck et al. to construct a matched filter that comprises a highlight region, dead-zone, and shadow region [4]. Depending on seabed elevation, the shadow length will vary significantly with respect to range. Consequently the shadow

Figure 4: The dual-tree wavelet adaptive shrinkage functions. The function S^- , given by (8) is designed to threshold or shrink wavelet coefficients that give rise to a wavelet energy ratio R_{k-1} that is low enough to satisfy (5). On the other hand, the shrinkage function S^+ given by (9) thresholds or shrinks coefficients that give rise to an energy ratio R_k that is large enough to satisfy (6).

component of the matched filter is varied in length as a function of range. In studies where sand ripples are not present or considered, the matched filter detector of Dobeck et al. has received some attention, c.f. [1, 8, 19]. To validate the ripple suppression method we compare detection results obtained from using a matched filter, similar to that of Dobeck et al, with and without the ripple suppression preprocessing step. Although the matched filter is too simplistic to accurately and uniquely represent the wide variety of potential targets, it is merely intended as an initial detection step to discard any regions of the data that are very unlikely to contain mine-like-objects.

Unlike Dobeck et al. [4], our matched filter is constructed as a superposition of raised cosines rather than their step functions. Our reasoning is motivated by the fact that the sand ripples in our data cannot be well approximated by pure sinusoidal plane waves. Observing Figure 3, we see that the ripples are more like square waves. In the frequency domain, the ripples will therefore contain higher frequency harmonics in addition to the fundamental frequency. The ripple suppression method will usually only suppress the fundamental frequency. A matched filter constructed from a superposition of square waves will also have higher frequency harmonics which will be excited by the harmonics left over from the ripple suppression method. On the other hand, by shaping the matched filter into raised cosines, the matched filter will attenuate the left-over ripple harmonics. From experimentation, we have also observed that the raised cosines give better ROC curves than the step functions with or without ripple suppression. The raised cosine matched filter used here is illustrated in Figure 5 and is fully described in [13].

2.6 Suppression/detector results

Figure 6 shows a typical result of applying a matched filter with and without the ripple suppression step, together with the computed shrinkage functions and resulting ripple suppressed image. Note that most of the ripples have an orientation that is roughly aligned with the 3rd subband direction of the DTCWT (stripe direction at 75° anticlockwise from the horizontal), and that the associated shrinkage function ($S_{6,3}$)

Figure 5: The mine template is constructed with a superposition of 3 shifted raised cosines. Along track is plotted vertically; cross track, or range, is plotted horizontally. The shadow length is constructed to increase with respect to range in a piecewise manner.

takes low values in a region that coincides with the ripples. The shrinkage functions in the neighbouring directions also take low values in the same region; this correctly captures the fact that the ripples manifest bifurcated and braided behaviour rather than a perfectly parallel pattern. We can see that the ripples are indeed suppressed by the ripple suppression method and that the non-rippled regions remain largely the same as the input image. In this example, for both suppressed and unsuppressed cases, the lowest (local maxima) correlation score associated with the true positives was chosen as a threshold. All scores above that threshold were labelled as true positives if they were in a neighbourhood of the ground truth. Otherwise, they were labelled as a false positive. We can see that most, if not all, of the false positives lie in the ripples and that the ripple suppression method gives rise to fewer false positives.

Validation of the ripple suppression preprocessing step was carried out by comparing the receiver operating characteristic (ROC) curves obtained from applying a matched filter with and without the ripple suppression preprocessing step over 61 rippled images and 140 non-rippled images, most of which comprised 14 megapixels. The data is real synthetic aperture sonar data and was acquired by the NATO Undersea Research Centre (NURC) and provided to us by the DSTL Data Centre.

To construct the ROC curves, a series of thresholds indexed by t , say, were applied to the correlation surface and the number of true positives $p^+(t)$, and false positives $p^-(t)$, recorded. Then, $p^+(t)$ was plotted against $p^-(t)$. To combine curves $(p_n^-, p_n^+)_1^N$ over $n = 1, \dots, N$ images, we simply plot $\sum_{n=1}^N p_n^+(t)$ against $\sum_{n=1}^N p_n^-(t)$.

Figure 7 shows the ROC curves for the suppressed and unsuppressed methods for the rippled, non-rippled, and combined dataset. On the rippled and combined dataset,

Table 1: Number of false positives incurred in order to recover 90%, 95%, 98%, 99%, and 100% of the total number of true positives.

Seabed type	Number of true pos.	Number of false positives	
		No suppression	Suppression
Non-rippled	270	22	21
	285	59	60
	294	259	253
	297	460	473
	300	806	793
Rippled	127	6320	2029
	134	10468	3345
	138	15032	7753
	140	38348	18209
	141	55719	19427
Both	397	2585	676
	419	7854	2256
	432	15332	4683
	437	22640	8494
	441	57457	21143

we observe that the matched filter achieves better detection results when the ripple suppression method is used. For the non-rippled data, the suppressed and unsuppressed methods give very similar results, as expected and required.

Table 1 records the number of false positives incurred in order to recover a certain percentage of the true positives. For example, the table shows that all the mines (targets) in the rippled data can be detected at a cost of 55,719 false positives with no suppression and 19,427 false positives with suppression: a reduction of some 65%.

3 Feature extraction

The motivation for using a lacunarity-based feature is the simple observation that the targets are easier to see by eye when viewed in the context of the surrounding background. It is assumed here that the presence of a man-made object, like a mine, will disrupt the statistical properties of the background. Under this premise, the translational statistical invariance in the neighbourhood of the mine will be less than that of a mine-free region.

Due to its connection to fractal dimension (see e.g. [6]), computation of lacunarity often follows a box counting approach. A window of size k -by- k pixels is moved over a region of interest, pixel-by-pixel. At each location in the region, mass is computed; in

Figure 6: The shrinkage functions $S_{k,m}$ at scale k , subband direction m are computed adaptively from the synthetic aperture sonar image (a) using (8) and (9). The shrinkage operation results in the ripple suppressed image (b). If no suppression is applied prior to application of the matched filter, the result is the correlation surface shown in (c). The yellow (resp. red) rings show the location of the true (resp. false) positives. Sub-figure (d) shows the resulting correlation surface when suppression is applied before the filter.

Figure 7: ROC curves compare numbers of true and false positives for the suppressed and unsuppressed methods on rippled, non-rippled, and both data types combined.

practice this usually means summing absolute values. Denoting the mean and variance of these masses in a k -by- k region centred at pixel $\mathbf{x}$ by $\mu(k; \mathbf{x})$ and $\sigma^2(k; \mathbf{x})$, the usual lacunarity definition used in practice is

$$\Lambda(k; \mathbf{x}) \triangleq \left(\frac{\sigma(k; \mathbf{x})}{\mu(k; \mathbf{x})} \right)^2.$$

In other words, the ratio of local variance to the square of the local mean. In some work, this is only computed at a single window size; for example Meyers [11] proposed the use of lacunarity as a feature in a mine hunting system. More commonly, it is computed at more than one scale (window size k) to give a lacunarity signature or vector. For example, Du and Yeo [5] compare the segmentation performance of several competing ways to measure lacunarity; all are based on the box counting approach and are computed at different scale levels by using different sizes of local regions. Partly inspired by the mono- to multi- generalisation of fractal dimension, Vernon et al [20] generalise the lacunarity measure by using higher order moments; again, the computation is based on box counting.

Our approach is to first apply the dual-tree complex wavelet transform to decompose the image over different scale levels and directional subbands. Root energy is then computed, in each directional subband and over a selection of scale levels, in the immediate region of the positives and in a larger surrounding region. A normalised dot product between the two energies then approximates the amount of correlation between the inner and outer region. In this way, the dot product can be seen as a measure of lacunarity. The main advantage of performing the computation in the wavelet domain is that any differences in the influence of wavelet coefficients between the inner and outer regions is preserved.

Figure 8 shows an example of a false and true positive located in a ripple field. The solid yellow box in figures 8(a) and 8(b) delineates the inner region. These are automatically centred on the location of each positive; no manual adjustments have been made. The size of the region is chosen to be the same size as the matched filter.

Note that, like the matched filter, the length is increased with respect to range in a piecewise manner to accommodate longer shadow regions further away from the sensor. The region between the solid and dotted yellow lines defines the outer region. For some scale level k , subband direction m , the root sum of squares of the inner region X_0 , and outer region X_1 , is computed:

$$v_\ell(k, m) \triangleq \left(\sum_{\mathbf{x} \in X_\ell} |w_{k,m}(\mathbf{x})|^2 \right)^{1/2}, \quad \ell = 0, 1. \quad (10)$$

(Note that $\ell = 0$ denotes inner region and $\ell = 1$ denotes outer region.) These values are shown in figures 8(c)-(e) for scale levels 4 and 5. Note that, for illustrative purposes, the root mean squared energies are plotted rather than the root sum squared energy defined in 10; this is because the outer region is five times the size of the inner region. The plots correctly reflect the fact that the inner and outer regions of the false positive are better correlated than those of the true positive. A similar result can be seen in Figure 9 for an example of a false and true positive located on a non-rippled patch of seabed.

The sums in (10) are computed for each directional subband and a selection of scale levels (3rd to 5th finest in our case) and the numbers are assembled into the column vector $\mathbf{v}_0$ for the inner region and $\mathbf{v}_1$ for the outer region. Our first lacunarity based feature is the normalised measure of correlation:

$$\gamma \triangleq \frac{\mathbf{v}_0^T \mathbf{v}_1}{\max(\|\mathbf{v}_0\|^2, \|\mathbf{v}_1\|^2)}. \quad (11)$$

This measures the statistical translational invariance of the local region (and is therefore a measure of inverse lacunarity). Another similar feature is formed by carrying out the above using the left and right half of the inner region and no outer region. This measures lack of correlations between the highlight and shadow region of the candidate positive. Both these features are computed using the ripple suppressed and unsuppressed data. Together with the result of the matched filter, we arrive at 5 features. In future work, we would hope to add more features to help discriminate target and background further but it is instructive to investigate the performance of these 5 features alone using a one-class support vector machine.

4 Classification experiment

For classification we need to form a functional relationship between the training data and the class labels. The function must also be chosen such that it generalises well to unseen data. The support vector machine (SVM) addresses the curse of dimensionality problem with the addition of two main terms: the first penalises the error between the solution and the training set; and the second regularises, or smooths in some way, the solution. More precisely, for two or more classes of objects, the SVM fits hyperplanes at

(a) False positive in the ripple field

(a) True positive in the ripple field

(c) Inner and outer region RMS energy of the false positive at the 4th finest scale level.

(d) Inner and outer region RMS energy of the true positive at the 4th finest scale level.

(e) Inner and outer region RMS energy of the false positive at the 5th finest scale level.

(f) Inner and outer region RMS energy of the true positive at the 5th finest scale level.

Figure 8: Comparison of the root sum of square energies of a true and false positive that lie inside a ripple field. (For illustrative purposes, the root mean sum is depicted.) Note that the false positive energies of the inner and outer regions are better correlated than in the true positive case.

a maximum distance from, and parallel to, the convex hulls of each class of trained data. Because the optimal hyperplanes only depend upon the support vectors, which are data points lying on the convex hulls, the SVM attains good generalisation. Moreover, the

(a) False positive in a non-rippled region

(b) True positive in a non-rippled region

(c) Inner and outer region RMS energy of the false positive at the 4th finest scale level.

(d) Inner and outer region RMS energy of the true positive at the 4th finest scale level.

(e) Inner and outer region RMS energy of the false positive at the 5th finest scale level.

(f) Inner and outer region RMS energy of the true positive at the 5th finest scale level.

Figure 9: Comparison of the root sum of square energies of a true and false positive that do not lie inside a ripple field. (For illustrative purposes, the root mean sum is depicted.) Note that, at the 4th finest scale level, the false positive energies of the inner and outer regions are better correlated than in the true positive case.

problem of finding the hyperplanes is reduced to a tractable quadratic programming problem with a unique, optimal solution.

On the other hand, unlike the binary or multi-class version, the training stage of

the one-class Support vector machine (OC-SVM), proposed by Schölkopf et al [17] does not require examples of targets. It only requires examples of normal background input data, and thus avoids the major difficulty of capturing examples of real threats in any significant quantity. It also potentially means that unanticipated and/or unknown types of threats can be detected.

The one-class support vector machine can be seen as a two class SVM where the origin is defined as the only member of one of the classes; in our case, the target class is defined as the origin in the training phase. The training algorithm learns the boundary which encloses most of the dataset subject to minimising the volume of the space within the boundary in feature space. Like the binary or multi-class SVM, a kernel feature mapping can be efficiently adopted to construct non-linear decision boundaries in the input feature space. For our experiment, we chose the radial basis kernel. We thus have two parameters to optimise: the kernel parameter which control the smoothness of the boundary and a regularisation parameter (often referred to simply as ν) which balances the trade-off between volume enclosed by the boundary and the number of points allowed outside the boundary.

The experimental setup was chosen similar to that of Hill et al [7], who applied log-Gabor, matched and shaped filters, some heuristic morphological constraints, together with a two-class SVM to the same UDRC NURC data set considered here. Half of the 180 sonar images were used to train the data and the other half held back for testing.

In our experiment, 9,186 false positives (in the set of training images) left over from the detection phase (with detection threshold chosen to accept 100% true positives) were used to train a one-class support vector machine. This was then tested on the true and false positives in the hold-out set. This can be contrasted with Hill et al. [7] chose 1000 randomly sampled non-target points from each 2000-by-7000 image and then used a two-class SVM trained on both false and true positives. They also reported the results of two human experts. Figure 10 shows the ROC curve of our one-class SVM test and the two-class SVM ROC curve approximated from Hill et al., together with results of the two operators reported by Hill et al. Around the 95% true positive region, the result approaches the performance of the two-class learning version of Hill et al. [7]¹: in our approach, 337 false positives are incurred in order to recover 95% of true positives whereas Hill et al incur 310. However, the left-hand-side (higher specificity) of the ROC curve of Hill et al. is noticeably superior; it beats operator 1 and intersects operator 2. As further comparison, we also carried out a two-class SVM version of our method which required true positive training examples in addition to the false positives. Again, like the one-class SVM experiment, we use the radial basis kernel and half of the sonar images were used to train the data with the other half held back for testing. The resulting ROC curve is shown in Figure 11. In this two-class version, 95% of true positives were recovered at the cost of 201 false positives which appears to be a better result than Hill et al. However, the ROC curve of Hill et al shows higher specificity for less than around 80% of true positives.

¹comparisons of results with Hill et al. [7] are drawn under the caveat that the experimental setups were somewhat different

Figure 10: ROC curve of one class support vector machine. Also plotted is the 2-class SVM method ROC curves of Hill et al. and the results of two expert operators (taken from [7]).

Figure 11: ROC curve of two class support vector machine. Also plotted is the 2-class SVM method ROC curves of Hill et al. and the results of two expert operators (taken from [7]).

5 Conclusion and further work

We have shown that anomaly detection has potential applications to mine hunting problems. Ripple suppression, together with a matched filter, discards the most ob-

vious background regions. The lacunarity-based features help to distinguish between the remaining background and target examples. A one-class support vector machine is trained only on the false positives from the matched filter, so that precise target characteristics need not be specified or known.

It is anticipated that additional carefully designed features, which compliment lacunarity, will improve and generalise the system further. For example, features that describe the spatial distribution and directions of edges, ridges, and points inside the inner and outer regions would be worthy of investigation. Incorporating some of the features and feature handling methods of Hill et al [7] would also be of interest. With suitable complex wavelet domain statistical models it should also be possible to adaptively suppress other background textures too such as vegetation and rocky outcrops.

Similar ideas to those discussed here could also be applied to seabed classification. In this case, a multi-class classifier would be used.

References

- [1] Y. J. Chen and T. Q. Nguyen. Sea mine detection based on multiresolution analysis and noise whitening. *Technical Report, Brown University*, 1999.
- [2] M. Couillard, J. A. Fawcett, V. L. Meyers, and M. Davison. Support vector machines for classification of underwater targets in sidescan sonar imagery. *Technical report, Defence Research and Development, Canada*, 2008.
- [3] G. J. Dobeck. Algorithm fusion for automated sea mine detection and classification. *IEEE OCEANS*, 1:130–134, 2001.
- [4] G. J. Dobeck, J. C. Hyland, and L. Smedley. Automated detection/classification of sea mines in sonar imagery. *SPIE*, 3079:90–110, 1997.
- [5] G. Du and T. Yeo. A novel lacunarity estimation method applied to SAR image segmentation. *IEEE Transactions on Geoscience and remote sensing*, 40(12):2687–2691, 2002.
- [6] R.A. Feagin. Relationship of second-order lacunarity, hurst exponent, brownian motion, and pattern organization. *Physica A*, 328:315–321, 2002.
- [7] P. Hill, A. Achim, and D. Bull. Underwater target detection in synthetic aperture sonar data. *Proceedings of Sensor Signal Processing for Defence*, 2010.
- [8] Q. Q. Huynh, N. Neretti, N. Intrator, and G. Dobeck. Image enhancement for pattern recognition. *Proceedings of SPIE*, 1998.
- [9] N. G. Kingsbury. Complex wavelets for shift invariant analysis and filtering of signals. *Journal of Applied and Computational Harmonic Analysis*, 10:234–253, May 2001.

- [10] B. B. Mandelbrot. *The Fractal Geometry of Nature*. W. H. Freeman and Company, 1982.
- [11] V. Meyers. Decision trees for computer-aided detection and classification (CAD/CAC) of mines in sidescan sonar imagery. *Technical report, Defence Research and Development, Canada*, 2002.
- [12] J. D. B. Nelson and N. G. Kingsbury. Dual-tree wavelets for estimation of locally varying and anisotropic fractal dimension. *IEEE International Conference on Image Processing*, 2010.
- [13] J. D. B. Nelson and N. G. Kingsbury. Fractal dimension based sand ripple suppression for mine hunting with sidescan sonar. *Institute of Acoustics Conference on Synthetic Aperture Sonar and Synthetic Aperture Radar*, 2010.
- [14] Y. Pailhas, Y. Petillot, C. Capus, and K. Brown. Real-time sidescan simulator and applications. *IEEE Oceans Europe Conference*, 2009.
- [15] B. Pesquet-Popescu and J. L. Vehel. Stochastic fractal models for image processing. *IEEE Signal Processing Magazine*, 19(5):48–62, 2002.
- [16] S. Reed, Y. Petillot, and J. Bell. An automatic approach to the detection and extraction of mine features in sidescan sonar. *IEEE Journal of Oceanic Engineering*, 28(1):90–105, 2003.
- [17] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson. Estimating the support of a high-dimensional distribution. *Technical report, Microsoft Research, MSR-TR-99-87*, 1999.
- [18] I. W. Selesnick, R. G. Baraniuk, and N. G. Kingsbury. The dual-tree complex wavelet transform. *IEEE Signal Processing Magazine*, 22(6):123–151, 2005.
- [19] J. Tian and C. Zhang. Automated detection/classification of objects in side-scan sonar imagery. *Proc. International Conference on Intelligent Mechatronics and Automation*, pages 632–637, 2004.
- [20] J. Vernon-Carter, C. Lobato-Calleros, R. Escarela-Perez, E. Rodriguez, and J. Alvarez-Ramirez. A suggested generalization for the lacunarity index. *Physica A*, 388:4305–4314, 2009.