

# Fractal dimension based sand ripple suppression for mine hunting with sidescan sonar

J. D. B. Nelson and N. G. Kingsbury

**Abstract**—Sand ripples present a difficult challenge to current mine hunting approaches. We propose a robust and adaptive method that suppresses sand ripples prior to the detection stage. The method exploits a fractal model of the seabed and the connection between: dual-tree wavelets and local, directional fractal dimension; interscale energy ratios, scale invariant frequency localised fractal dimension, and a novel wavelet shrinkage approach. Tests on a reasonably large, real synthetic aperture sonar imagery dataset show that the ripple suppression method preserves detection performance of the matched filter on non-rippled data and significantly increases the detection performance on data that contain ripples.

## I. INTRODUCTION

THERE is growing interest in the analysis of sonar imagery for automatic mine hunting. Manual inspection of such data can be a time consuming task that requires significant and constant concentration. It is hoped that future computer aided systems will help eliminate distraction or fatigue and perform faster, more consistent, processing. A practicable mine hunting system must overcome some difficult challenges including clutter and seabed variability. Sand ripples, which can share some of the characteristics of mine-like objects, present one of the major challenges for automatic mine hunting applications [1], [2].

Amongst the most promising techniques that have the potential to overcome the effects of sand ripples is the work of Reed et al. [3]. 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. Some robustness to sand ripples, compared with previous methods, was reported.

In the following, we present a method that automatically suppresses sand ripple effects. Our scale and rotation invariant method exploits ideas from fractal analysis to distinguish rippled from non-rippled seabed regions. The information obtained is used to design a space, direction, and scale adaptive wavelet shrinkage operator that targets and attenuates the ripple energy. Although our method is not necessarily intended as an alternative to that of Reed et al. [3], it can potentially facilitate the use of simpler detection/classification methods which were previously ineffectual in the presence of sand ripples. As an example, we will use the matched filter detector as a validation tool to compare the detection performance with and without ripple suppression.

Williams and Coiras [4] recently proposed a filter bank with 6 directions and 6 scale levels, which resembled a Gabor basis, to detect sand ripples in synthetic aperture sonar. After applying the 36 (undecimated) convolutions with the input image, a ‘ripplicity’ measure was computed by taking the difference between filter responses at mutually orthogonal directions. Given that the location and orientation of the ripples could be found robustly, they proposed that the heading of the sensor should be changed to a direction parallel to the ripple propagation in order to moderate the ripple effects. In a sense, they ventured a hardware, as opposed to our proposed software, ripple suppression approach. Unfortunately, owing to variations in wave frequency and amplitude of currents, sand ripples are not guaranteed to result in regular, linear patterns. Rather, the ripples, as seen from above, form pseudo sinusoidal, bifurcated, and braided patterns [5] (see Fig. 3(a)). In practice, this could potentially mean that the hardware solution requires that the sensor re-inspect the same region of seabed from multiple angles each time a ripple field was encountered.

Using wavelets to preprocess sonar imagery prior to mine detection has been attempted before. However, these methods have only been used to denoise the sonar speckle rather than to suppress sand ripples. Chen and Nguyen [6] employ an undecimated, overcomplete Haar wavelet transform before a (wavelet decomposed) matched filter is applied to each scale level. A scale dependent threshold is then applied to the correlation surfaces at each scale level and the resulting binary maps are intersected to form the final binary detection map. Unfortunately, an undecimated wavelet transform, together with multiple instances of fully sampled matched filter operations (at each scale level) is computationally expensive. It is also unclear as to how the vertical, horizontal, and diagonal subbands are combined, and how the scale dependent thresholds are predetermined or generalised to unseen data.

Huynh et al. [7] used several different denoising methods including Gaussian, difference of Gaussian, and standard wavelet shrinkage. After some experimentation with real sidescan sonar data, they reported that wavelet shrinkage gave rise to the best matched filter detector performance and that shrinking only the 2nd finest scale level gave the best detection results. Being non-adaptive purely linear operators which blur the highlight/shadow regions associated with mines, the Gaussian and difference of Gaussian filters in [7] gave worse results than not performing any preprocessing.

In Section II the connection between wavelets and fractal dimension is briefly outlined for self similar processes. Section III introduces a fractal based model to distinguish between rippled and non-rippled seabed regions in a scale invariant

J. D. B. Nelson and N. G. Kingsbury are with the Signal Processing and Communications Group of the University of Cambridge, UK.

This work was supported by the EPSRC (grant number EP/H012834/1) and the UK MoD University Defence Research Centre on Signal Processing.

manner as well as a wavelet shrinkage method to suppress the ripples. The same model motivates a modification of a matched filter, proposed by Dobeck et al. [8], which is presented in Section IV. In Section V, validation is carried out by comparing the receiver operating characteristic (ROC) curves obtained from applying a matched filter with and without the ripple suppression preprocessing step.

## II. FRACTAL DIMENSION

### A. 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 [9]. It is satisfied by processes with power spectra  $\mathcal{P}f$  that satisfy 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 [9].

### B. 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  $\psi_m$  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. Appendix A shows 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 [10], the authors proposed the use of the dual-tree complex wavelet transform (DTCWT) [11], [12] 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 [11], [12] 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.

## III. RIPPLES, FRACTAL DIMENSION, AND COMPLEX WAVELET SHRINKAGE

### A. 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 1 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. [13] 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  $\phi$  (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.

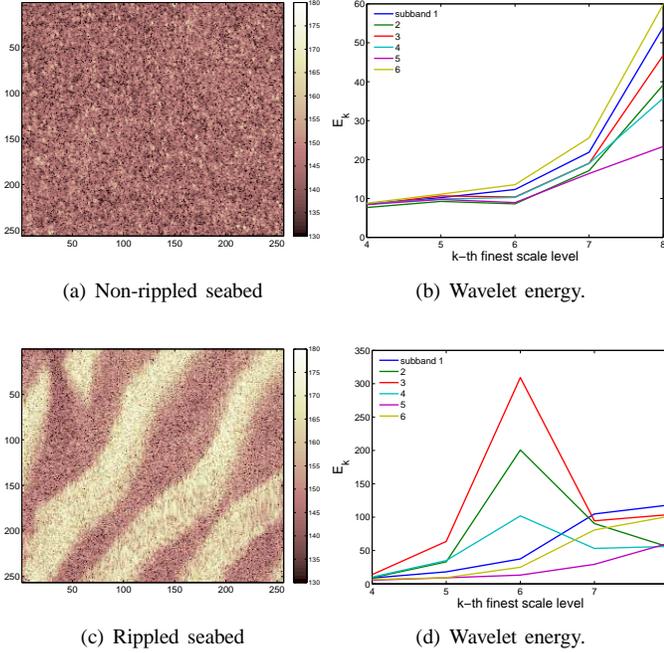


Fig. 1. A patch of non-rippled seabed (a), gives rise to a wavelet amplitude spectrum (b) that follows a power law decay with respect to finer scales (it decays as frequency increases). Conversely, a rippled seabed patch (c) will invalidate the power law at one or more scales and subband directions (d).

### B. 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 2. 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.

### C. Ripple suppression algorithm

To summarise, the ripple suppression algorithm proceeds as follows.

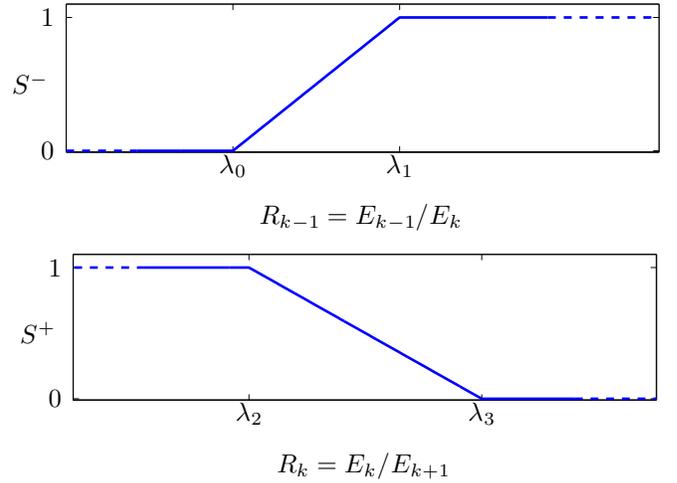


Fig. 2. The dual-tree wavelet adaptive shrinkage functions. The function  $S^-$ , given by (8) is designed to threshold or shrink wavelet coefficients that give rise to a wavelet energy ratio  $R_{k-1}$  that is low enough to satisfy (5). On the other hand, the shrinkage function  $S^+$  given by (9) thresholds or shrinks coefficients that give rise to an energy ratio  $R_k$  that is large enough to satisfy (6).

- Compute the DTCWT coefficients of the input  $f$ :

$$w_{k,m}(\mathbf{x}_n) \leftarrow (\mathcal{W}f)(k, m; \mathbf{x}_n),$$

- at scale  $k$ , in subband (direction)  $m$ , at location  $\mathbf{x}_n$ .
- Estimate local root-energy  $E_{k,m}(\mathbf{x}) \leftarrow \sum_{\chi} |w_{k,m}(\mathbf{x}_n)|$ , in some spatial neighbourhood  $\chi$  of  $\mathbf{x}$ .
- Compute ratio  $R_{k,m}(\mathbf{x}) \leftarrow E_{k,m}(\mathbf{x})/E_{k+1,m}(\mathbf{x})$ . In practice, the sampling rate is proportional to  $2^k$ ; therefore  $E_{k+1,m}$  is up-sampled by a factor of 2 to compute  $R_{k,m}$ .
- Compute ratio  $r_{k,m}(\mathbf{x}) \leftarrow$  down-sampled (by a factor of 2) version of  $R_{k-1,m}$ .
- Compute adaptive shrinkage functions:

$$S_{k,m}^-(\mathbf{x}) \leftarrow \min \left( 1, \max \left( 0, \frac{r_{k,m}(\mathbf{x}) - \lambda_0}{\lambda_1 - \lambda_0} \right) \right)$$

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

- Combine functions:  $S_{k,m}(\mathbf{x}) \leftarrow S_{k,m}^-(\mathbf{x})S_{k,m}^+(\mathbf{x})$ .
- Apply shrinkage:  $w_{k,m}^\sim(\mathbf{x}) \leftarrow S_{k,m}(\mathbf{x})w_{k,m}(\mathbf{x})$ .
- Take inverse DTCWT  $f^\sim \leftarrow \mathcal{W}^{-1}w_{k,m}^\sim$ .
- $f^\sim$  is a ripple suppressed version of input image  $f$ .

In the results provided later, the shrinkage thresholds were set to  $\{\lambda_n\} = \{0, 0.25, 2.5, 3.5\}$  and the size of the spatial neighbourhood  $\chi$  was set to 256-by-256 pixels. In practice, we found that the success of the ripple suppression method was not unduly sensitive to these parameters choices; improvements (compared to no suppression) in detection performance were realised even when the parameters were varied significantly.

## IV. 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 [8]. Depending on seabed elevation, the shadow length will vary significantly with respect to range. Consequently the shadow 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. [6], [7], [14].

To validate the ripple suppression method described in the previous sections, 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. A classifier can then be trained and tested on positive responses to the detector. The matched filter operation can be described as the 2-dimensional cross-correlation

$$g(\mathbf{x}) \triangleq \sum_n h(\mathbf{x}_n) f(\mathbf{x} + \mathbf{x}_n), \quad (10)$$

between image  $f$ , and filter  $h$ . Putting  $(x, y) \triangleq \mathbf{x}$ , the matched filter is constructed with

$$h(\mathbf{x}) \triangleq h_0(x) h_1(y).$$

That is, it is a separable filter and it follows that (10) can be implemented as a 1-dimensional correlation down the columns followed by another across the rows of the image. Along the range  $x$ , the filter is constructed as a superposition of shifted raised cosines, viz.:

$$h_0(x) \triangleq \sum_{n=1}^3 \alpha_n h_{0,n}(x),$$

where  $n = 1, 2$  and  $3$ , correspond to the highlight, dead-zone, and shadow regions respectively. Unlike Dobeck et al. [8], 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 1, 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 filter regions are described by

$$h_{0,n}(x) \triangleq \begin{cases} \cos^2 \frac{\pi}{\ell_n} (x - \beta_n), & \text{if } |x - \beta_n| \leq \frac{\ell_n}{2} \\ 0, & \text{otherwise} \end{cases}$$

Without loss of generality, we assume that the data is gathered from the starboard side of the sensor platform. We require that the raised cosine for the highlight region is centred on  $\ell_1/2$ ; i.e.  $\beta_1 = \ell_1/2$ . The dead-zone is constructed to intersect zero at the same location that the highlight takes its maximum; hence  $\beta_2 = (\ell_1 + \ell_2)/2$ . The shadow region is designed such that it only overlaps the dead-zone at a single point and we have that  $\beta_3 = \ell_2 + (\ell_1 + \ell_3)/2$ . In order to eliminate any zero-frequency bias, the matched filter is required to have zero mean. That is we want  $\int h_0(x) dx = 0$ , or  $\alpha_n$  such that

$$\sum_{n=1}^3 \alpha_n \int h_{0,n}(x) dx = 0.$$

This implies  $\sum_{n=1}^3 \alpha_n \ell_n = 0$ . We choose  $\alpha_1 = 1, \alpha_2 = 1/3$ , which assumes the highlight region gives a response 3 times stronger than the deadzone. We also assume that the highlight and deadzone lengths in the range direction are similar. I.e.  $\ell_1 = \ell_2 \Rightarrow \alpha_3 = -4\ell_1/(3\ell_3)$ . In the track, or cross-range direction, the filter component  $h_1(y)$  is defined by another raised cosine with width  $\ell_0$ . The precise choices of  $\{\ell_n\}$  should be tailored according to scenario dependent parameters, which are usually known, such as sensor altitude, resolution, and approximate target size. The shadow is also lengthened, by increasing  $\ell_3 = \ell_3(x)$ , with respect to range in a piecewise manner.

## V. RESULTS

Figure 3 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  $75^\circ$  anticlockwise from the horizontal), and that the associated shrinkage function ( $S_{6,3}$ ) 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.

TABLE I

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

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  $(p_n^-, p_n^+)$  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 4 shows the ROC curves for the suppressed and unsuppressed methods for the rippled, non-rippled, and combined dataset. On the rippled and combined dataset, 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 I 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%.

## VI. CONCLUSION

We have introduced the technique of fractal dimension based dual-tree wavelet shrinkage to suppress sand ripples in sonar imagery. Results on a reasonably large dataset indicate that the method can act as a useful preprocessing stage, prior to mine detection. It is worth pointing out that the matched filter detector discussed here is a separate development to the ripple suppression step. It is anticipated that our ripple suppression method could also enhance other previous and current mine detection strategies. This method is only intended as the first phase in a larger detection/classification system. As in previous detector/classifier mine hunting approaches, it is anticipated that a classification phase will further reduce the number of false positives incurred by the detection phase.

## ACKNOWLEDGMENTS

This work was supported by the EPSRC (grant number EP/H012834/1) and the UK MoD University Defence Research Centre on Signal Processing. The authors would like to

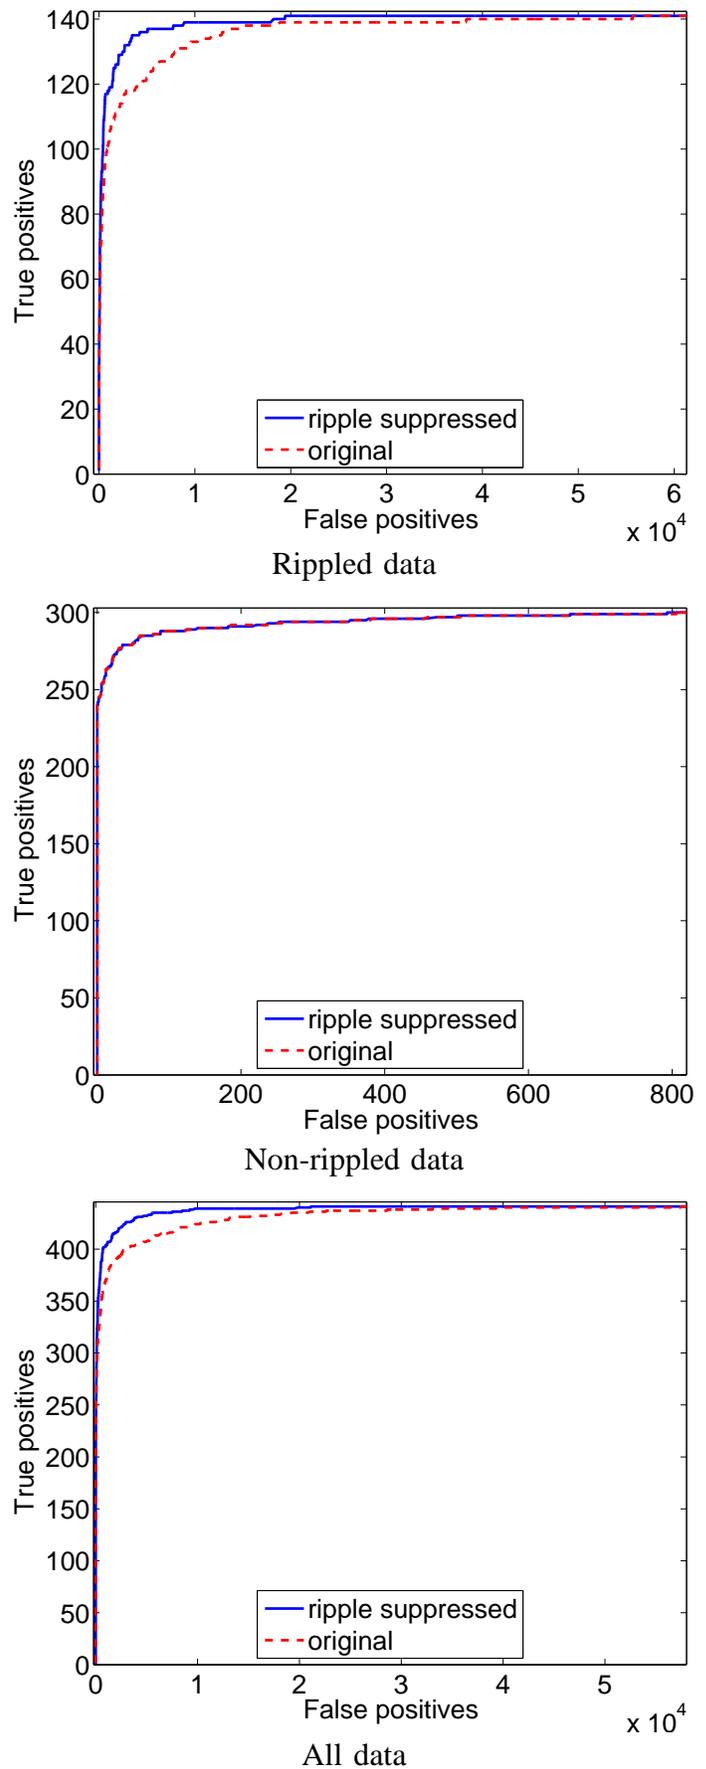


Fig. 4. ROC curves compare numbers of true and false positives for the suppressed and unsuppressed methods on rippled, non-rippled, and both data types combined.

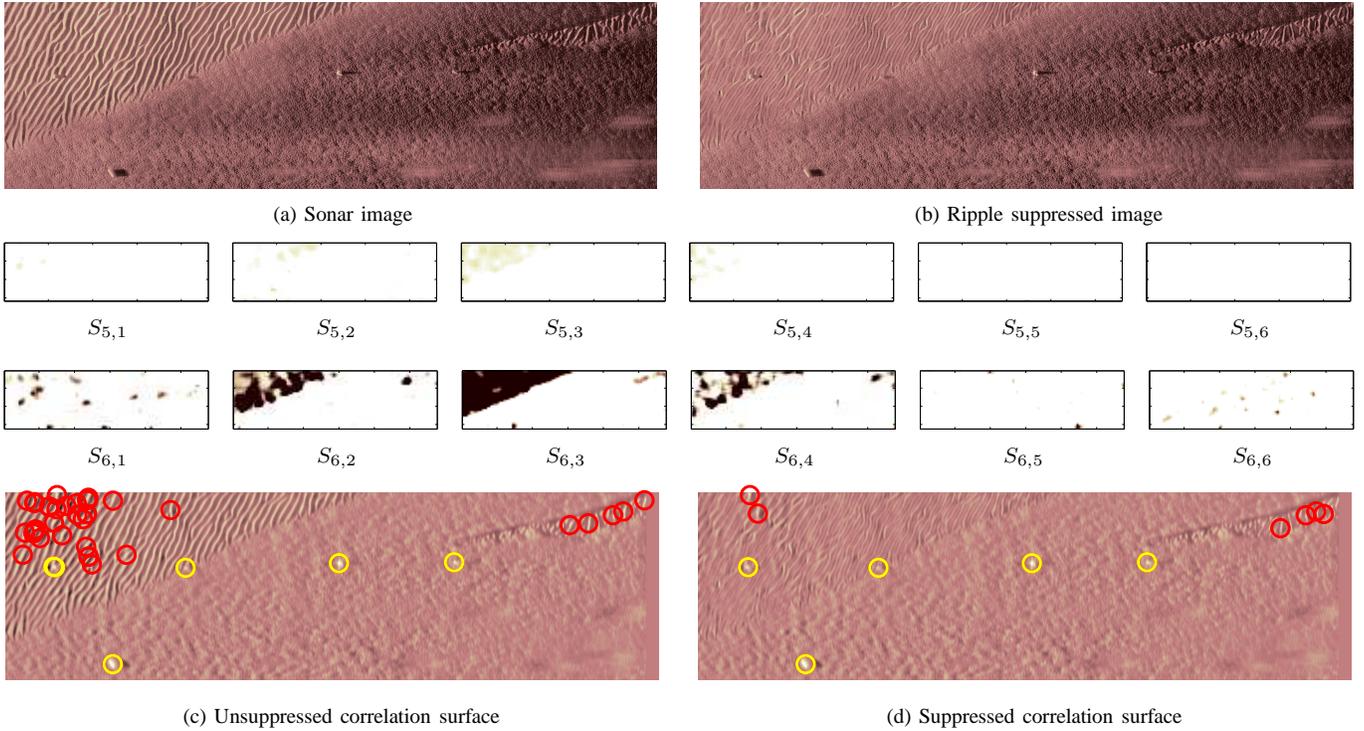


Fig. 3. The shrinkage functions  $S_{k,m}$  at scale  $k$ , subband direction  $m$  are computed adaptively from the synthetic aperture sonar image (a) using (8) and (9). The shrinkage operation results in the ripple suppressed image (b). If no suppression is applied prior to application of the matched 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.

thank: William Oxford and Daniel Pearce from the DSTL data centre for their efforts in securing the release of the sidescan sonar data; Duncan Williams and Samantha Dugelay from DSTL for discussions on the data and mine hunting literature; and Dave Bull and Paul Hill at Bristol University for providing ground truth of the data and other useful information.

#### APPENDIX A

##### FRACTAL DIMENSION OF SELF SIMILAR PROCESSES

Since wavelets are zero mean, the variance of the transform is  $\mathbb{E} \left[ |(\mathcal{W}f)(k, m; \cdot)|^2 \right] =$

$$2^{-2k} \int_{\mathbb{R}^2} \int_{\mathbb{R}^2} \mathbb{E} \left[ \psi_m(2^{-k}(\cdot - \mathbf{x})) \overline{\psi_m(2^{-k}(\cdot - \boldsymbol{\xi}))} f(\mathbf{x}) f(\boldsymbol{\xi}) \right] d\mathbf{x} d\boldsymbol{\xi}.$$

After a change of variables, this is

$$2^{2k} \int_{\mathbb{R}^2} \int_{\mathbb{R}^2} \psi_m(\mathbf{x}) \overline{\psi_m(\boldsymbol{\xi})} \mathbb{E} [f(\cdot - 2^k \mathbf{x}) f(\cdot - 2^k \boldsymbol{\xi})] d\mathbf{x} d\boldsymbol{\xi}.$$

Invoking self-similarity (1), and putting  $\mathbf{x}_n = n2^k \Delta \mathbf{x}$ , we have  $\mathbb{E} \left[ |(\mathcal{W}f)(k, m; \mathbf{x}_n)|^2 \right] =$

$$2^{2k(H+1)} \int_{\mathbb{R}^2} \int_{\mathbb{R}^2} \psi_m(\mathbf{x}) \overline{\psi_m(\boldsymbol{\xi})} \mathbb{E} [f(n\Delta \mathbf{x} - \mathbf{x}) f(n\Delta \mathbf{x} - \boldsymbol{\xi})] d\mathbf{x} d\boldsymbol{\xi},$$

and the integrand is independent of  $k$ , which yields (3).

#### REFERENCES

- [1] P. Chapple, "Automated detection and classification in high-resolution sonar imagery for autonomous underwater vehicle operations," *Australian Government Department of Defence Technical Report, DSTO-GD-0537*, 2008.
- [2] J. Groen, E. Coiras, and D. Williams, "Detection rate statistics in synthetic aperture sonar images," *IACM Proc. 3rd International Conference and Exhibition on Underwater Acoustic Measurements*, 2008.
- [3] 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*, vol. 28, no. 1, pp. 90–105, 2003.
- [4] D. Williams and E. Coiras, "On sand ripple detection in synthetic aperture sonar imagery," *IEEE Proc. International Conference on Acoustics, Speech, and Signal Processing*, pp. 1074–1077, 2010.
- [5] J. L. Hansen, M. Van Hecke, A. Haaning, C. Ellegaard, K. H. Andersen, T. Bohr, and T. Sams, "Instabilities in sand ripples," *Nature*, vol. 410, pp. 324, 2001.
- [6] Y. J. Chen and T. Q. Nguyen, "Sea mine detection based on multiresolution analysis and noise whitening," *Technical Report, Brown University*, 1999.
- [7] Q. Q. Huynh, N. Neretti, N. Intrator, and G. Dobeck, "Image enhancement for pattern recognition," *Proceedings of SPIE*, 1998.
- [8] G. J. Dobeck, J. C. Hyland, and L. Smedley, "Automated detection/classification of sea mines in sonar imagery," *SPIE*, vol. 3079, pp. 90–110, 1997.
- [9] B. Pesquet-Popescu and J. L. Vehel, "Stochastic fractal models for image processing," *IEEE Signal Processing Magazine*, vol. 19, no. 5, pp. 48–62, 2002.
- [10] 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.
- [11] N. G. Kingsbury, "Complex wavelets for shift invariant analysis and filtering of signals," *Journal of Applied and Computational Harmonic Analysis*, vol. 10, pp. 234–253, May 2001.
- [12] I. W. Selesnick, R. G. Baraniuk, and N. G. Kingsbury, "The dual-tree complex wavelet transform," *IEEE Signal Processing Magazine*, vol. 22, no. 6, pp. 123–151, 2005.
- [13] Y. Pailhas, Y. Petillot, C. Capus, and K. Brown, "Real-time sidescan simulator and applications," *IEEE Oceans Europe Conference*, 2009.
- [14] J. Tian and C. Zhang, "Automated detection/classification of objects in side-scan sonar imagery," *Proc. International Conference on Intelligent Mechatronics and Automation*, pp. 632–637, 2004.