

# MULTIRESOLUTION MARKOV RANDOM FIELD WAVELET SHRINKAGE FOR RIPPLE SUPPRESSION IN SONAR IMAGERY

*J. D. B. Nelson*

Dept. of Statistical Science  
University College London, UK  
j.nelson@ucl.ac.uk

*N. G. Kingsbury*

Signal Processing and Communications Laboratory  
Dept. of Engineering, University of Cambridge, UK  
ngk@eng.cam.ac.uk

## ABSTRACT

A recent dual-tree wavelet shrinkage method to suppress sand ripples in sonar imagery is extended with a Markov random field framework. Markov chain Monte Carlo sampling is used to estimate the posterior marginal ripple state in the wavelet domain. Ripple suppression is realised by multiplying the dual-tree wavelet coefficients by the conditional probabilities of the non-ripple state. Tests on real data confirm that this extended method significantly further reduces false positives.

**Index Terms**— Wavelet transforms, Markov random fields, image denoising, sonar detection

## 1. INTRODUCTION

Mine detection in the presence of background clutter is an open problem in automatic sonar image processing. Sand ripples, for example, exhibit highlight and shadow features that can cause detectors to give many false positives. In [1], the authors developed an adaptive sand ripple suppression preprocessing step which gave significant reduction, of around 63%, in the number of false positives. We extend this wavelet shrinkage ripple suppression method here to incorporate a Markov random field (MRF) model which exploits the fact that rippled and non-rippled locations are spatially clumped.

MRFs have been used in sonar image processing before but their use has been restricted to seabed classification rather than background suppression. For example, Mignotte and Collet [2] used MRFs to separate shadow regions from the seabottom and constrain the output of a fuzzy seabed classifier. Although wavelet methods to preprocess sonar imagery have also been reported, the objective was to denoise speckle [3]. In the wider literature, MRF-based wavelet shrinkage has proved to be an effective way to denoise natural imagery. Examples include Malfait and Roose [4] and Pižurica [5] who both used the marginal posterior to shrink the unwanted signal; in their

cases, this meant suppressing Gaussian white noise. In our case this means suppressing sand ripples.

## 2. STOCHASTIC MODEL

The wavelet transform of a stochastic field  $g: \mathbb{R}^2 \mapsto \mathbb{R}$ , is  $(\mathcal{W}g)(k, m; r) := 2^{-k} \int_{\mathbb{R}^2} g(r') \overline{\psi_m(r - 2^{-k}r')} dr'$ , where  $\psi_m$  is a mother wavelet with orientation indexed by  $m$ , and where  $k$  denotes the  $k$ th finest scale level ( $k = 0$  being the original ‘pixel’ coordinates), and where  $r \in \mathbb{R}^2$  is the spatial location. It can be shown (e.g. [1]) that

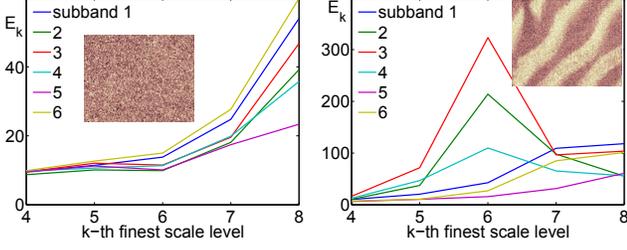
$$\mathbb{E}(|(\mathcal{W}g)(k, m; \cdot)|^2) \propto 2^{2k(H+1)} \quad (1)$$

holds for weak statistically self similar fields (WSSS), i.e.  $g$  with finite variance, such that  $\mathbb{E}(g(\gamma \cdot)) = \gamma^H \mathbb{E}(g(\cdot))$ , and  $\mathbb{E}(g(\gamma r)g(\gamma \cdot)) = \gamma^{2H} \mathbb{E}(g(r)g(\cdot))$ , with Hurst parameter  $H \in (0, 1]$ . An extended version of this model, considered in [1] and here, allows the Hurst parameter to vary over space and direction, whence  $H = H(m; r)$ . In practice, the 2nd moment in (1) is approximated by the sample 2nd moment, a local average of magnitude squared coefficients over some finite square neighbourhood (for each scale  $k$  and direction  $m$ ). Denoting the square root of the sample 2nd moment by  $E_{k,m}(r)$ , it follows that WSSS regions satisfy  $E_{k,m}(r) \propto 2^{k(H(m;r)+1)}$ . Since  $H \in (0, 1]$ , the interscale wavelet energy ratio gives us the scale invariant inequality:

$$\frac{1}{4} \leq \frac{E_{k,m}(r)}{E_{k+1,m}(r)} = 2^{-(H(m;r)+1)} \leq \frac{1}{2}, \quad \forall k. \quad (2)$$

This inequality can be used to test whether a region about location  $r$  is statistically self similar at scale  $k$ , along direction  $m$ . Since non-ripples broadly follow a power law decay with finer scales, and the oscillatory component of rippled regions cause a spike in the directional wavelet spectra (see Fig. 1), the general premise in what follows is that non-rippled regions satisfy inequality (2) whereas rippled regions do not (along the ripple propagation direction). This observation agrees with Pailhas et al [6]. They synthesised sonar imagery of non-rippled seabeds with fractional Brownian surfaces, which have statistically self similar 2nd moments and thus satisfy (2). They also synthesised rippled regions with frequency shifted fractional Brownian surfaces.

This work was supported by the EPSRC (grant number EP/H012834/1) and the UK MoD University Defence Research Centre on Signal Processing. All data used here was provided by the NATO Centre for Maritime Research and Experimentation.



**Fig. 1.** Left: A patch of non-rippled seabed (inset) has a wavelet spectrum that follows a power law decay with respect to finer scales. Right: conversely, a rippled seabed patch (inset) will invalidate the power law at one or more scales and subband directions.

### 3. WAVELET SHRINKAGE FOR RIPLE SUPPRESSION

Since the objective is to suppress ripples adaptively, we should avoid suppressing the wavelet energy of a non-rippled seabed (which can result in poor detection results). Hence, the upper and lower bounds in inequality (2) are relaxed. The method proposed in [1] proceeds as follows: first the dual-tree wavelet transform is taken, then the shrinkage functions  $S: (k, m; r) \mapsto [0, 1]$ , with  $S_{k,m}(r) = S_{k,m}^1(r)S_{k,m}^2(r)$ , are computed, where (see fig. 2)

$$S_{k,m}^1(r) := \min \left( 1, \max \left( 0, \frac{y_{k,m}^-(r) - \theta_0}{\theta_1 - \theta_0} \right) \right) \quad (3)$$

$$S_{k,m}^2(r) := \min \left( 1, \max \left( 0, \frac{\theta_3 - y_{k,m}^+(r)}{\theta_3 - \theta_2} \right) \right) \quad (4)$$

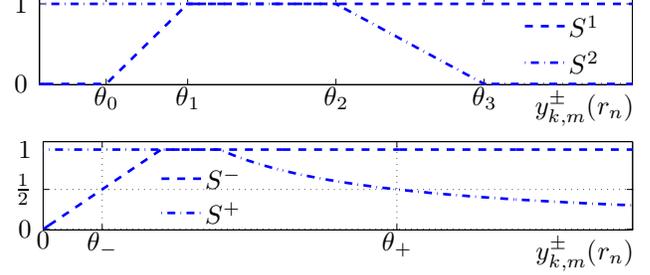
$$y_{k,m}^-(r) := \frac{E_{k-1,m}(r)}{E_{k,m}(r)}, \quad y_{k,m}^+(r) := \frac{E_{k,m}(r)}{E_{k+1,m}(r)} \quad (5)$$

where  $\theta_0, \dots, \theta_3$  are suitably chosen thresholds. Since dyadic sampling is used, note here that the  $E_{k+1}$  used to compute  $y_k^+$  is upsampled by a factor of 2 and  $y_k^-$  is a downsampled (by a factor of 2) version of  $y_{k-1}^+$ . The wavelet coefficients are multiplied by the shrinkage functions and the inverse wavelet transform is taken to arrive at the ripple suppressed image. However, when seen as functions of energy ratios:  $S^1 = S^1(E_{k-1}/E_k)$  and  $S^2 = S^2(E_k/E_{k+1})$ , it is apparent that these functions do not treat scalings ( $\geq 1$ ) of  $E_k$  (which could indicate a ripple) in the same way:  $S^1$  is scale invariant (for  $\theta_0 = 0$ , and  $E_{k-1}/E_k \leq \theta_1$ ) but  $S^2$  is not (even for  $E_k/E_{k+1} \geq \theta_2$ ). This motivates the following departure from [1] (see also Fig. 2):

$$S_{k,m}^-(r) := \min \left( 1, \frac{1}{2} \frac{y_{k,m}^-(r)}{\theta_-} \right) \quad (6)$$

$$S_{k,m}^+(r) := \min \left( 1, \frac{1}{2} \frac{\theta_+}{y_{k,m}^+(r)} \right). \quad (7)$$

A further shortcoming of the deterministic method in [1] is that the knowledge that rippled or non-rippled regions tend to be spatially clumped is not fully exploited. Intuitively, given that a surrounding neighbourhood of local regions are deemed



**Fig. 2.** The dual-tree wavelet adaptive shrinkage functions: top panel shows (3) and (4) used in [1]; the lower panel shows those introduced here. The function  $S^1$  (and  $S^-$ ) shrinks coefficients which break the lower bound of (2) and  $S^2$  (and  $S^+$ ) shrinks those that break the upper bound of (2).

strongly to be ripples (non-ripples), it is reasonable to assume that the central region of the neighbourhood is likely to be rippled (non-rippled). In other words, the ripple suppression approach of [1] can benefit by incorporating contextual, spatial constraints. To this end, we consider extending the method by incorporating a Markov random field model.

#### 3.1. Markov Random Field based Shrinkage

For some fixed scale level and subband direction, let  $x_n$  denote the state (ripple =  $-1$  or non-rippled =  $1$ ) of the wavelet coefficient at spatial location  $r_n$  (i.e.,  $n$ th site or node of an undirected graph). Let  $y = (y_1, \dots, y_N)$  denote the corresponding observed measures of non-rippled state. The joint posterior probability  $p(x|y) \propto p(y|x)p(x)$  is a distribution defined over the  $2^N$ -dimensional configuration label space  $x = \{\pm 1\}^N$ . Assuming the labels constitute a Markov random field, the joint posterior probability takes the form  $p(x|y) \propto \exp(-U(x|y))$ , where the posterior energy is  $U(x|y) = U(y|x) + U(x)$ , the sum of the prior and likelihood energies.

The aim is to estimate the marginal posterior probabilities  $p(x_n = 1|y_n)$  for each wavelet coefficient site,  $n$  (and over different scale levels and subband directions). For then, a ripple suppressed image can be formed by multiplying the wavelet coefficients by  $p(x_n = 1|y_n)$  followed by the inverse dual-tree wavelet transform.

Here, we use the Ising model as the prior, with energy  $U(x_n) := -\beta \sum_{n' \sim n} x_n x_{n'}$  where  $n' \sim n$  denotes that the site  $n'$  is in the 4-neighbourhood of site  $n$ . The parameter  $\beta$  controls the influence of the prior which favours spatially clumped states. The heuristic non-rippled marginal energy measure is defined by  $U(y_n|x_n) := -\eta \log(S_{k,m}^*(r_n|x_n))$ , where  $S_{k,m}^*(r_n|1) := S_{k,m}^+(r_n)S_{k,m}^-(r_n)$  is constructed via the modified shrinkage functions (6) and (7). This means, subject to the influence of the prior, it is likely that wavelet coefficients which fail to satisfy only one of the bounds in inequality (2) are proportionately shrunk and those that fail to satisfy both bounds are shrunk twice. Similar to the wavelet shrinkage denoising models of [4] and [5], we assume conditional independence on the likelihood  $p(y|x) = \prod_n p(y_n|x_n)$ . Under the reasonable assumption that the image does not take

constant values over large regions, the local mean wavelet energy will always take finite, non-zero values. Hence, the interscale ratios will take on finite, non-zero values and the density  $p(y_n)$  will have bounded support, say  $[\theta_\infty^{-1}, \theta_\infty]$ . Similar to [4], we make the assumption that  $y$  is uniformly distributed:  $p(y_n|x_n=1) + p(y_n|x_n=-1) = \text{const}$ , over the support of  $y$  and leave other possible constructions (empirical or otherwise) for further work. The ratio between the parameters  $\eta$  and  $\beta$  balances the influence of the prior and likelihood.

### 3.2. Inference

A Metropolis sampler is used to estimate the posterior marginal inference. The algorithm proceeds as follows (denoting the number of coefficients at scale level  $k$  by  $N_k$ ).

```

for scale  $k$  and subband direction  $m$  do
   $\forall n, x_n^{(1)} \leftarrow 1$ , if  $S_{k,m}^*(r_n|1) > \frac{1}{2}$ , and 0 otherwise
   $\nu_n \leftarrow 0$ 
  for  $t = 1 : \text{number\_iterations}$  do
     $\ell \leftarrow$  random permutation of  $1 : N_k$ 
    for  $n = \ell(1 : N_k)$  do
      compute  $\alpha \leftarrow p(-x_n^{(t)}|y_n) / p(x_n^{(t)}|y_n)$ 
      if  $\alpha > \alpha' \sim \text{uniform}([0, 1])$  then
        set  $x_n^{(t+1)} \leftarrow -x_n^{(t)}$ 
      end if
      if  $t > \text{burn\_in\_period}$  then
         $\nu_n \leftarrow \nu_n + \delta_{x_n^{(t+1)}, 1}$ , ( $\delta_{\cdot, \cdot} =$  Kronecker's delta)
      end if
    end for
  end for
   $p^\sim(x_n = 1|y_n; k, m) \leftarrow \nu_n / (\text{number\_iterations} - \text{burn\_in\_period})$ 
end for

```

The posterior marginal ratio  $p(-x_n^{(t)}|y_n)/p(x_n^{(t)}|y_n)$  is computed via the unnormalised marginal priors and likelihoods:

$$\frac{p^*(-x_n|y_n)}{p^*(x_n|y_n)} = \frac{1 - S_{k,m}^*(r_n|x_n)^\eta}{S_{k,m}^*(r_n|x_n)^\eta} \exp\left(-2\beta \sum_{n' \sim n} x_n x_{n'}\right).$$

The result  $p^\sim(x_n = 1|y_n; k, m)$  is an estimate of the marginal posterior probability that the wavelet coefficient at location  $n$ , scale  $k$ , and subband direction  $m$  is a non-ripple. The ripple suppression algorithm is completed by multiplying the dual-tree wavelet coefficients by  $p^\sim$  before taking the inverse dual-tree wavelet transform.

## 4. EXPERIMENTS & CONCLUSIONS

Validation was carried out as follows. Real synthetic aperture sonar data was acquired by the NATO Centre for Maritime Research and Experimentation. This data comprises 61 rippled images and 140 non-rippled (mostly 14 megapixel) images. The resulting images from the deterministic and MRF ripple suppression methods were passed to a matched filter (c.f. [1]). Receiver operating characteristic (ROC) curves were constructed and some key values recorded in Table 1.

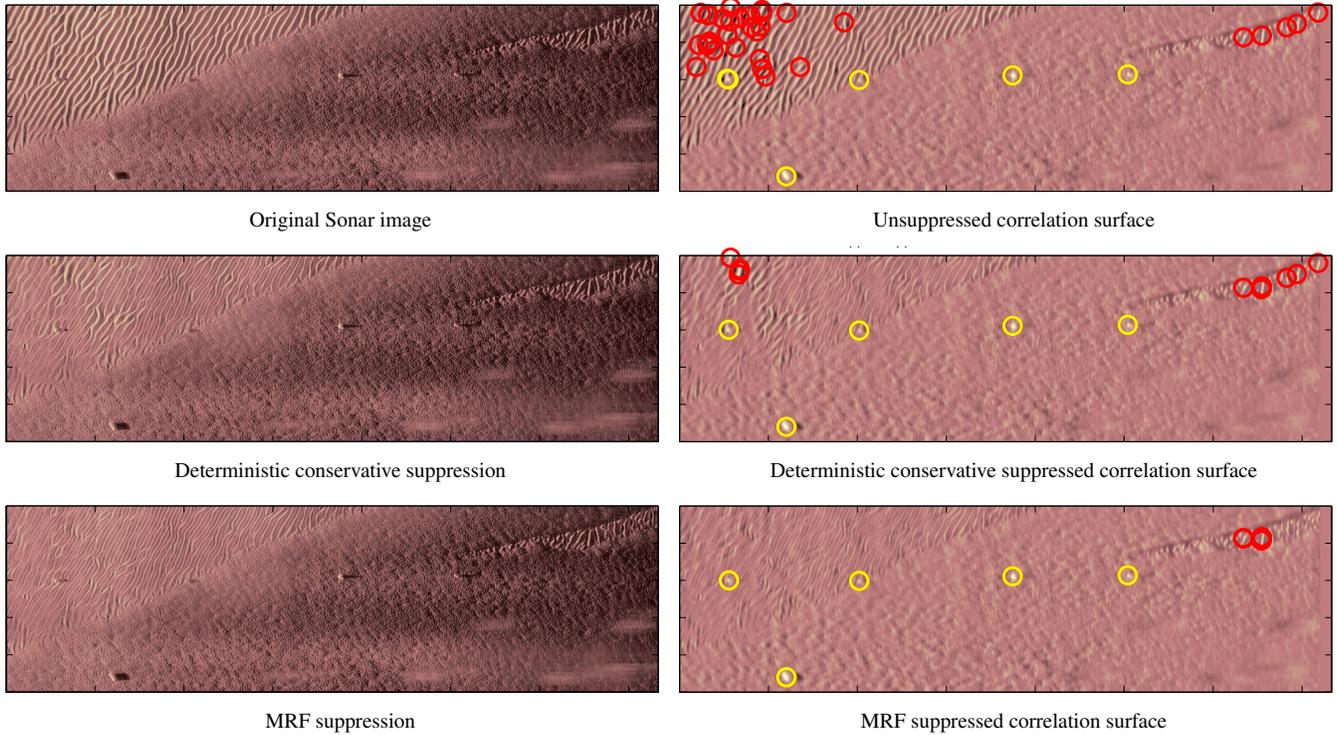
MRF parameter selection was performed by testing various combinations on 30% of the images. The best parameter

values were found to be  $\eta = 1, \beta = 5, \theta_- = 1/4, \theta_+ = 3/2, \text{burn\_in\_period} = 5, \text{number\_iterations} = 15$ . For the deterministic method, various parameter settings were tested on the entire dataset. It was found that values of  $\theta_0 > 0$  or  $\theta_1 > 1/4$  led to an overly aggressive suppressor which gave poorer results than no suppression on the non-rippled seabed datasets. Two different sets of the other values are shown in Table 1. The conservative method has  $\theta_3 = 5/2, \theta_4 = 7/2$  and aggressive method has  $\theta_3 = 3/2, \theta_4 = 5/2$ . The MRF method worked best with the measure  $S^*$  and the deterministic methods worked best with  $S$ . (A likely explanation is that since  $S^*$  is strictly non-zero it allows the Metropolis sampler to flip states more easily whereas it does not allow the deterministic method to set coefficients to zero.) Note that, although the aggressive method incurs fewer false positives than the conservative method on the dataset as a whole, it unduly interferes with the matched filter performance on the non-rippled dataset: e.g. some 43 additional false positives are incurred in order to recover 99% of the true positives. In any case, the MRF based method outperforms both deterministic methods and further reduces the number of false positives reported in [1] by around a third. Crucially, it also appears to have little impact on the non-rippled results.

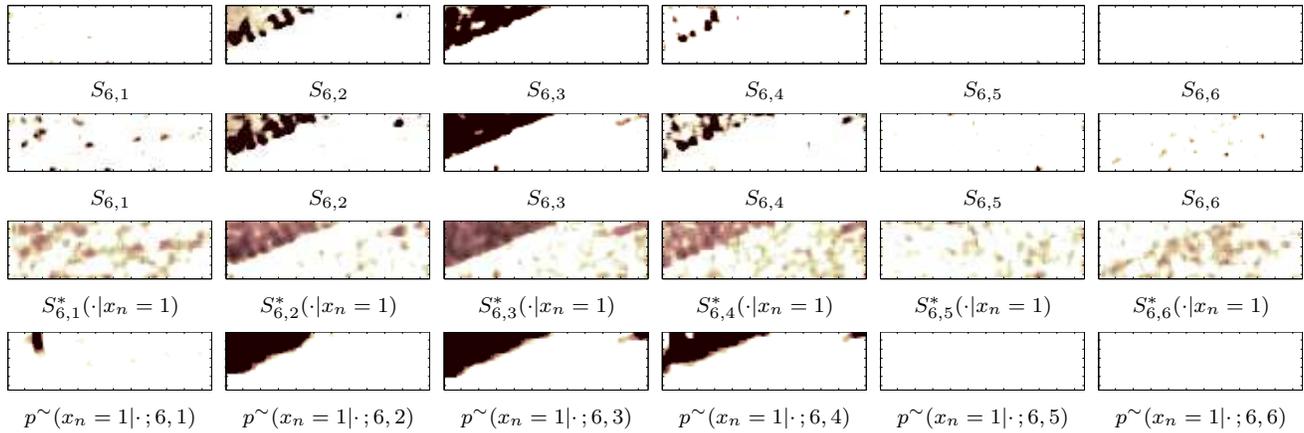
Further work will explore: more principled likelihood models; incorporating scale and directional dependencies into the Markov model; other inference algorithms; and the modelling and suppression of other background clutter.

## 5. REFERENCES

- [1] 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 SAS and SAR*, 2010.
- [2] M. Mignotte and C. Collet, "Markov random field and fuzzy logic modeling in sonar imagery: Application to the classification of underwater floor," *Computer Vision and Image Understanding*, vol. 79, pp. 4–24, 2000.
- [3] Q. Q. Huynh, N. Neretti, N. Intrator, and G. Dobeck, "Image enhancement for pattern recognition," *Proceedings of SPIE*, 1998.
- [4] M. Malfait and D. Roose, "Wavelet-based image denoising using a Markov random field a priori model," *IEEE Transactions on Image Processing*, vol. 6, no. 4, pp. 549–565, 1997.
- [5] A. Pižurica, W. Philips, I. Lemahieu, and M. Acheroy, "A joint inter- and intrascale statistical model for bayesian wavelet based image denoising," *IEEE Trans. Image Processing*, vol. 11, no. 5, pp. 545–557, 2002.
- [6] Y. Pailhas, Y. Petillot, C. Capus, and K. Brown, "Real-time sidescan simulator and applications," *IEEE Oceans Europe Conference*, 2009.



**Fig. 3.** Left: original and ripple suppressed images; right: correlation surfaces; yellow (resp. red) rings show location of true (resp. false) positives after choosing the optimal threshold in each example for illustrative purposes only. Not shown: aggressive deterministic method (for that, see Fig. 3 in [1]).



**Fig. 4.** Non-rippled measures at scale level 6 of the image shown in Figure 3. Top row: conservative deterministic; 2nd row: aggressive deterministic; third row: the new measure used for the MRF method; fourth row: posterior marginal estimates.

**Table 1.** Number of false positives incurred in order to recover 90%, 95%, 98%, 99%, and 100% of the total number of true positives using: no suppression; conservative and aggressive deterministic method; and MRF- based method.

Seabed #tr.pos.	None					Rippled					Both				
	270	285	294	297	300	127	134	138	140	141	397	419	432	437	441
none	22	<b>59</b>	<b>259</b>	460	806	6320	10468	15032	38348	55719	2585	7854	15332	22640	57457
cons.	<b>21</b>	60	253	473	<b>793</b>	2029	3345	7753	18209	19427	<b>676</b>	2256	4683	8494	21143
agg.	<b>21</b>	60	288	503	794	<b>1463</b>	2977	6878	11232	13398	692	<b>1712</b>	4075	5840	15127
MRF	22	<b>59</b>	<b>259</b>	<b>459</b>	806	1933	<b>2945</b>	<b>6630</b>	<b>10514</b>	<b>12308</b>	774	1811	<b>3493</b>	<b>5553</b>	<b>14041</b>