

INTRODUCING THE DUAL-TREE COMPLEX WAVELET LOCALLY STATIONARY MODEL

J. D. B. Nelson and A. J. Gibberd

Department of Statistical Science, University College London
{j.nelson, alexander.gibberd.12@ucl.ac.uk}@ucl.ac.uk

ABSTRACT

This paper reconciles Kingsbury’s dual-tree complex wavelets with Nason and Eckley’s local stationary model. We here establish that the dual-tree wavelets admit an invertible debiasing matrix and that this matrix can be used to invert the covariance relation. We also show that the added directional selectivity of the proposed model adds utility to the standard two-dimensional local stationary model. Non-stationarity detection on random fields is used as a motivating example. Experiments confirm that the dual-tree model can distinguish anisotropic non-stationarities significantly better than the current model.

Index Terms— wavelet, locally stationary, anisotropic

1. INTRODUCTION

First proposed by Nason et al. [1] in the one-dimensional setting, the locally stationary wavelet (LSW) model offers a translation-invariant, statistically well-principled means to capture the local covariance and local spectrum of a large class of, possibly non-stationary, stochastic processes. Extension of the theory to two-dimensional processes was recently pioneered by Eckley et al. [2, 3] and was followed up with methodology applied to non-stationarity detection in image textures [4] and segmentation of imagery into stationary regions [5].

Alternative wavelet approaches to modelling random fields often involve modelling the raw wavelet coefficients directly [6] or passing them as features to machine learning tools [7]. This is in contrast to the LSW where the process is explicitly described with a parameterised, generative model where the parameters are estimated in a statistically rigorous manner to ensure that: the asymptotic bias is understood and corrected; and that consistency is characterised and, where appropriate, obtained. Furthermore, the LSW model offers a natural way to extend the notion of wavelet auto-covariance to the local case— Prop. 1 and Cor. 1 in [2] show respectively that the local auto-covariance converges uniformly to auto-covariance as sample size increases and that one can be derived from the other via explicit linear mappings. Although other wavelet models, such as HMTs [8, 9], capture dependency structures between wavelet coefficients, no other wavelet-based model other than the LSW appears to satisfy these fundamental covariance properties. Perhaps the closest model to LSW is [10, 11] which

was applied to paper and mamographic imagery; however their wavelet variance-based framework is restricted to the smaller class of (intrinsically) stationary fields.

The LSW model is translation invariant thanks to the use of what Nason terms ‘discrete wavelets’, formed by an inverse, undecimated, à trous algorithm using finite impulse response (FIR) wavelet filters. However, the current design currently only accommodates wavelet bases with the standard ‘horizontal’, ‘vertical’, and ‘diagonal’ orientations. Such poor directionality is a serious drawback when the data of interest presents anisotropic behaviour. Some non-stationarity detection experiments in Section 4 illustrate examples of this point.

Motivated by this lack of directionality we here move towards reconciling two significant contributions to the field of wavelet analysis, namely: the LSW model [1, 2] and Kingsbury’s dual-tree complex wavelet transform (DTCWT) [8, 12–14]. The former evolved from the motivation to build a statistically rigorous notion of localised auto-covariance using wavelets. The latter was motivated by the desire to design a wavelet basis with good shift-invariance which could be implemented in a decimated pyramidal scheme— a consequence of which is that, in two-dimensions, the DTCWT furnishes wavelets with superior directionality than standard real, FIR wavelets. Since the DTCWT is implemented with an FIR filterbank our proposed DTCWT-LSW model also uses ‘discrete wavelets’ in a similar fashion to the LSW case.

By extending some key notions such as the auto-correlation wavelet and the biasing matrix to the complex-valued DTCWT-LSW case we then: (i) establish some mild conditions under which the biasing matrix can be inverted (Theorem 6); (ii) prove that the mapping between local auto-covariance and standard auto-covariance is invertible (Prop. 8); (iii) demonstrate that the added directional selectivity of the proposed model adds utility to the standard two-dimensional local stationary model in that it is more general and, unlike the standard LSW, can be used to detect anisotropic non-stationarities.

Section 2 offers a brief overview of the LSW model and the DTCWT. We propose the DTCWT-LSW model in Section 3, establish some fundamental results, and use a toy example to illustrate that its enhanced directionality represents a more general model space than the standard LSW. Experiments are conducted in Section 4 which demonstrate the superiority of the DTCWT-LSW at non-stationarity detection.

J. D. B. Nelson and A. J. Gibberd were supported by Dstl. J. D. B. Nelson was also partially supported by EPSRC grant EP/N508470/1

2. BACKGROUND

The (2-d) LSW and DTCWT are here reviewed.

2.1. Locally stationary wavelet model

The local stationary wavelet (LSW) model represents a stochastic process $X = \sum_{j,\ell} \sum_u w_j^\ell(u) \xi_j^\ell(u) \psi_j^\ell(u - \cdot)$ on the finite lattice $\llbracket 1, T \rrbracket^2$ as a weighted filtering (with weights w) of a zero-mean orthonormal increment sequence ξ and an overcomplete basis comprising ‘undecimated discrete wavelet functions’ $\psi_j^\ell(\cdot) \in \ell_2(\mathbb{Z}^2)$ [1, 2]. The wavelet functions are defined/derived by performing the usual inverse decimated wavelet transform (with upsampling) of Delta functions placed at the j th scale level, cf. Daubechies p204 [15] and the ℓ th orientation. Note that, in contrast to the decimated case, a collection of integer translated wavelets are used. This not only offers a translation invariant basis but also ensures that traditional stationary processes are captured. Asymptotically, the weights are associated with a Lipschitz continuous function $W_j: (0, 1)^2 \mapsto \mathbb{R}$, defined over rescaled space $z = t/T$, $t \in \llbracket 1, T \rrbracket^2$ which restricts the variation of the weights over space. In turn this captures the fact that the local structure becomes evermore stationary with respect to decreasing spatial neighbourhoods. For notational convenience, the scale and orientation indexes are amalgamated into one index $\eta = \eta(j, \ell)$ —e.g. $\eta = j + (\ell - 1)J$, $j \in \llbracket 1, J \rrbracket$ and we note that the model can be written in terms of convolutions as: $X = \sum_\eta w_\eta \xi_\eta * \bar{\psi}_\eta$, where the $\bar{\cdot}$ denotes a left-right and up-down spatial flip: $\bar{f} := f(-\cdot)$. A key thrust of the LWS is that it affords a natural description of the notion of local covariance.

Definition 1 (local covariance, Eckley et al. [2], Nason et al. [1]). *The local covariance (LCV) $c: (0, 1)^2 \times \mathbb{Z}^2 \mapsto \mathbb{R}$ of an LSW process with LWS $S_\eta(\cdot)$ is defined by*

$$c(z, t) := \sum_\ell \sum_{j \in \mathbb{N}} S_\eta(z) \Psi_\eta(t),$$

with local wavelet spectrum (LWS) $S_\eta(\cdot) = |W_\eta(\cdot)|^2$ and autocorrelation wavelet (ACW)

$$\Psi_\eta = \psi_\eta * \bar{\psi}_\eta = \sum_t \psi_\eta(t) \psi_\eta(\cdot + t). \quad (1)$$

A naive estimate of the LCV would be the (undecimated) local wavelet periodogram (LWP) $X_\eta^\sim := |X * \psi_\eta|^2$. However, the following result suggests that this would be far from ideal.

Theorem 2 (Eckley et al. [16]). *Let X be an LSW process with LWP X^\sim ; then*

$$\mathbb{E}X_\eta^\sim([zT]) = \sum_\nu A_{\eta,\nu} S_\nu(z) + O(T^{-1}),$$

where the biasing matrix $A := A_{\eta,\nu}$ is defined as the Gram matrix with entries

$$A_{\eta,\nu} = \langle \Psi_\eta, \Psi_\nu \rangle = \sum_t \Psi_\eta(t) \Psi_\nu(t). \quad (2)$$

Hence the redundancy of the non-decimated LWP results in a biased estimator of the LCV. This mixes LWS content between other scales and orientations. However, Nason et al. [1] showed, in the one-dimensional Shannon and Haar wavelet case, that this matrix is non-singular. Hence the LWP can be debiased with the inverse biasing matrix to arrive at $\widehat{S}(z) := A^{-1}X^\sim(z)$, i.e.

$$\widehat{S}_\eta(z) = \sum_\nu A_{\eta,\nu}^{-1} X_\nu^\sim(z) \quad (3)$$

with the property that $\mathbb{E}\widehat{S} = S(z) + O(T^{-1})$. Furthermore, numerical experimentation has shown, in practice, that finite order versions of A are invertible for other bases such as Daubechies Extremal Phase [2].

In essence, the LSW model describes the multiresolution, local covariance structure of the process of interest. In broad terms, one expects the amplitude of $w_{\eta(j,\ell)}(t)$ to be large if the image at point t is locally well-correlated at the j th finest scale level and ℓ th orientation.

2.2. Dual-tree complex wavelets

It is remarkable that the dual-tree complex wavelet achieves near translation-invariance and directional selectivity whilst, at the same time, it can be realised with a decimated wavelet transform, with finite impulse response filters, that satisfy the perfect reconstruction property [14]. As such the DTCWT also admits a multiresolution analysis in the true and full sense in that the resulting two-scale relations are satisfied with finite sequence filters. One interpretation of the q-shift solution is that it represents a multiwavelet which comprises two orthonormal wavelets. In any guise, the DTCWT is without doubt a singular member of what many refer to as the wavelet zoo.

In contrast to the LSW model, the DTCWT is non-decimated. An immediate means to harmonise both frameworks is to incorporate a non-decimated version of the DTCWT into the LSW. Whilst the extra computation required does very little in terms of improving translation invariance—from near-invariance to exact invariance—non-decimation does greatly enhance the poor directional description suffered by the LSW model. Indeed, undecimated versions of the DTCWT have recently been proposed by Hill et al. [17] who, in addition, note the added virtues that the resulting wavelet coefficients are truly co-located over the scale levels rather than lying on the usual dyadic grid.

The DTCWT uses complex-valued basis functions to achieve near translation invariance by designing the real and imaginary parts of the filters such that they approximately form a Hilbert transform pair (90° out of phase with each other). There exist several practicable means to achieve this but perhaps the most appealing is via q-shift filters [12]. These result in a pair of wavelets $\{\psi_{:,0}, \psi_{:,1}\}$ such that:

$$\psi_{:,1}(t) = \psi_{:,0}(n - 1 - t), \quad (4)$$

where n denotes the FIR filter length. The resulting quasi-analytic wavelet function $\psi_{:,0} + i\psi_{:,1}$ will then not only be

(approximately) shift invariant but, in two-dimensions, it can also be used to form directionally selective wavelets in the twelve directions $\{(30k - 15)^\circ\}_{k=1}^{12}$. In contrast, the DWT only offers space-scale content over ‘two and a half’ directions, namely vertical, horizontal, and a ‘diagonal’ direction which mixes both the $+45^\circ$ and -45° orientations together.

3. THE DTCWT-LSW MODEL

With complex-valued wavelets, the DTCWT-LSW model reads similarly to the real-case, namely

$$X = \sum_{\eta} w_{\eta} \xi_{\eta} * \bar{\psi}_{\eta},$$

but now the weights $w_{\eta}: \mathbb{N}^2 \mapsto \mathbb{C}$ can be complex-valued. The ACW, Eqn. (1) is hereby generalised to the complex case using complex-valued autocorrelation and the biasing matrix, Eqn. (2), is extended to the complex case by simply adopting the more general definition of inner products for complex-valued vectors. We hence propose the following definitions as natural generalisations of those found in the LSW framework.

Definition 3 (Complex ACW). *Let the superscript \cdot^* denote complex conjugation. Define the ACW associated with the possibly complex-valued wavelet ψ by*

$$\Psi_{\eta} = \psi_{\eta} * \bar{\psi}_{\eta}^* = \sum_t \psi_{\eta}^*(t) \psi_{\eta}(\cdot + t). \quad (5)$$

Definition 4 (DTCWT-LSW biasing matrix). *Let Ψ be the complex ACW as above. Then the associated biasing matrix is*

$$A_{\eta,\nu} = \langle \Psi_{\eta}, \Psi_{\nu} \rangle = \sum_t \Psi_{\eta}(t) \Psi_{\nu}^*(t). \quad (6)$$

3.1. Properties

With the DTCWT-LSW defined, we now establish some key results which confirm that the biasing matrix of the proposed model can be inverted under very mild conditions.

Lemma 5. *Let Ψ be the ACW (Defn. 1). The entries of the biasing matrix (Defn. 4) are positive, real, and symmetric.*

Proof. By construction, A is a Gram matrix. Hence, symmetry follows immediately. We then invoke Plancherel to write $A = \langle \psi_{\eta} * \bar{\psi}_{\eta}^*, \psi_{\nu} * \bar{\psi}_{\nu}^* \rangle = \langle |\psi_{\eta}^{\wedge}|^2, |\psi_{\nu}^{\wedge}|^2 \rangle$. ■

Theorem 6 (Invertibility of the biasing matrix). *Let the biasing matrix A' associated with an appropriately chosen real-valued wavelet $\psi_{\cdot,0}$ be non-singular. Then, the biasing matrix A (Defn. 4) associated with the dual-tree wavelet (Eqn. 4) $\psi = \psi_{\cdot,0} + i\psi_{\cdot,1}$ is non-singular with $A = 2A'$.*

Proof. We use the same arguments as the proof Lemma 5 and then note that the Fourier magnitudes of the real and imaginary parts of the dual-tree wavelets are identical by design: $A = \langle \psi_{\eta} * \bar{\psi}_{\eta}^*, \psi_{\nu} * \bar{\psi}_{\nu}^* \rangle = \langle |\psi_{\eta}^{\wedge}|^2, |\psi_{\nu}^{\wedge}|^2 \rangle = \langle |\psi_{\eta;0}^{\wedge}|^2 + |\psi_{\eta;1}^{\wedge}|^2, |\psi_{\nu;0}^{\wedge}|^2 + |\psi_{\nu;1}^{\wedge}|^2 \rangle = 2 \langle |\psi_{\eta;0}^{\wedge}|^2, |\psi_{\nu;0}^{\wedge}|^2 \rangle$. ■

Remark 7. *Thm. 6 also holds for analytic wavelets since the real and imaginary components have equal magnitude.*

Proposition 8 (DTCWT-LSW covariance inversion). *Assume that the debiasing matrix A is non-singular; then*

$$S_{\eta}(z) = \sum_{\nu} A_{\eta,\nu}^{-1} \sum_t c(z, t) \Psi_{\nu}^*(t).$$

Proof. As in [1], we follow a proof by verification and then note that A is both symmetric and, by Lemma 5, real:

$$\begin{aligned} & \sum_{\nu} A_{\eta,\nu}^{-1} \sum_t \sum_{\eta'} S_{\eta'}(z) \Psi_{\eta'}(t) \Psi_{\nu}^*(t) \\ &= \sum_{\eta'} S_{\eta'}(z) \sum_{\nu} A_{\eta,\nu}^{-1} \sum_t \Psi_{\eta'}(t) \Psi_{\nu}^*(t) \\ &= \sum_{\eta'} S_{\eta'}(z) \sum_{\nu} A_{\eta,\nu}^{-1} A_{\nu,\eta'} = \sum_{\eta'} S_{\eta'}(z) \delta_{\eta,\eta'} = S_{\eta}(z). \blacksquare \end{aligned}$$

3.2. Simulation

If the local wavelet spectrum of a process is known one can readily sample fields directly from the LSW and DTCWT-LSW models. Furthermore, as described in Algorithm 1, one can also simulate fields from data using the estimated LWS \hat{S} . A telling comparison between LSW and DTCWT-LSW can be seen in Fig. 1. We first simulate some textures using the LSW model. The field ‘dwt, band 1’ is created by setting $S_{\eta} = 0$ at all bands except the vertical band at one of the scale levels which is set to 1. Likewise ‘dwt, band 3’ uses only the diagonal band. These fields are then ‘resimulated’ with the DTCWT model using the estimated LWS. The top row of Fig. 1 labels these resimulations as ‘dwt sim of dwt...’. In the second row of Fig. 1, the roles are reversed so that the DTCWT is used to simulate and the DWT is used to resimulate. It should be clear that the DTCWT model can reproduce the behaviour of the DWT model much better than the DWT can reproduce the DTCWT. Clearly, the DWT cannot simulate the range of directionality inherent in the DTCWT-LSW model.

Algorithm 1 Simulation algorithm, cf. Taylor [4]

```

for (level, orientation) =  $\eta$ ; do
  compute LWP  $X_{\eta}^{\sim}$ , LWS  $\hat{S}_{\eta}$ , and  $\mu_{\eta} := T^{-2} \sum_t \hat{S}_{\eta}(t)$ 
  simulate:  $\hat{X} = \sum_{\eta} \hat{w}_{\eta} \xi_{\eta} * \bar{\psi}$ , using  $\hat{w}_{\eta}(\cdot) \equiv \sqrt{\mu_{\eta}}$ 
end for

```

4. EXPERIMENTS

Non-stationarity testing was carried out using the bootstrapping scheme proposed in [4], namely for a given field X we test the null hypothesis that $S_{\eta}(z)$ is a constant function for all z and η using the test statistic $\tau_X := (JK)^{-1} \sum_{\eta} \hat{\sigma}^2(\hat{S}_{\eta})$ which is the mean, over J scales and K subbands, of the LWS sample variance computed over pixels. When compared to bootstrapped samples cf. Alg. 2, this affords a statistically rigorous means to evaluate how likely an observed statistic τ_X is compared to bootstrapped values $\tau_{\hat{X}_i}$ assuming stationarity. One would expect a non-stationary field to register a small

p-value. As such this would indicate that the null hypothesis could be rejected. As in [4], experiments were performed on both stationary and non-stationary fields. Unlike [4], we also include experiments on anisotropic fields.

Algorithm 2 Stationarity detection, cf. Taylor [4]

```

for (level, orientation) =  $\eta$ ; do
  compute LWS  $\hat{S}_\eta$  from (3) and test statistic  $\tau_{\text{obs}}$ 
  compute mean stationary spectrum  $\mu_\eta := T^{-2} \sum_t \hat{S}_\eta(t)$ 
end for
for  $i$  in  $1 : B$  bootstraps do
  simulate  $\hat{X}^i$  using Alg. 1 and compute test statistic  $\tau_{\hat{X}^i}$ 
end for
compute p-value  $p := (B + 1)^{-1} (1 + \#(\tau_X \leq \tau_{\hat{X}^i}))$ 

```

4.1. Stationary processes

Stationarity tests were performed on an isotropic first order spatial moving average field with weight parameter 0.9 (process ‘S2’ in [4]). One hundred realisations of this process were simulated. For each realisation, bootstrapping was performed over one hundred simulations, cf. Alg 2. Both the LWS and DTCWT-LWS model registered p-values greater than 0.05 for all realisations. As such, this experiment corroborates findings in [4] and [18] that the test is conservative.

A second test was performed on anisotropic Gaussian noise, simulated using a simple directional mask oriented at 45° with angular width 30° together with a Fourier synthesis method. The LWS again registered p-values great than 0.05 for all realisations. The DTCWT-LWS resulted in 2 realisations with p-values greater than 0.05 which is closer to the 5% figure one would expect.

4.2. Non-stationary processes

Further experiments were carried out on non-stationary random fields. An isotropic field was generated where pixels in one half were drawn from unit-variance Gaussian white noise. The other half was formed from a Gaussian random field with an order 1 Bessel-Matérn covariance function (cf. process ‘NS2’ from [4]). This experiment was repeated for three different scale values. Table 1 records the number of p-values below 0.05 that resulted from 100 realisations (i.e. the number of times that the null-hypothesis was correctly rejected). The DTCWT-LSW appears to detect the non-stationarity better than the LSW. The extra redundancy may add some value to the estimation or simulation process but this requires further work to fully explain.

A set of fields with covariance $\exp(-(t-t')^\top DR_\theta(t-t'))$ were used with $D = \text{diag}(4, 1)$, where the rotation matrix R_θ is piecewise constant with $\theta = 90^\circ$ for $(t-t') \in t_0$ and $\theta \in \{0^\circ, 20^\circ, 40^\circ, 60^\circ\}$ for $(t-t') \notin t_0$ and where t_0 is a central square region, cf. Fig. 2. Note that these fields are anisotropically non-stationary. The number of instances where the null-hypothesis was correctly rejected is recorded in Table 2. Thanks to its superior directionality, the proposed




Fig. 1. Top row: dtcwt simulations of dwt simulation. Bottom row: dwt simulations of dtcwt simulation.




Fig. 2. Anisotropically non-stationary textures. Processes comprise inner and outer region with different orientations. Top row: spectral density functions. Bottom row: realisations of the processes.

Table 1. Proportion (%) of bootstrap samples with p-value below 0.05 for non-stationary isotropic random field.

Method/scale	0.4	0.5	0.6
DTCWT-LSW	1	51	94
DWT-LSW	0	25	28

Table 2. Proportion (%) of bootstrap samples with p-value below 0.05 for non-stationary anisotropic random field.

Method/ θ	0	20	40	60
DTCWT-LSW	99	99	98	66
DWT-LSW	9	10	5	1

method offers a clear advantage here. Furthermore, as the orientations of the two regions become more similar the non-stationary detection becomes worse for both methods.

5. CONCLUSION

Kingsbury’s dual-tree complex wavelet basis was successfully harmonised with Nason and Eckley’s locally stationary wavelet model. We have established some initial properties of the resulting proposed DTCWT-LSW model and have confirmed the enhanced utility afforded by the extra directionality. The pervasiveness of anisotropic non-stationary random fields in image processing and the elementary description that covariance provides means that there is much scope for further investigation including: quality inspection via machine vision (e.g. texture pilling [4, 19]); segmentation [5]; sonar/denoising [20, 21], change detection [22]. It is also anticipated that a more complete range of anisotropy testing machinery (e.g. Thon et al. [11]) would provide greater statistical power.

6. REFERENCES

- [1] G. P. Nason, R. Von Sachs, and G. Kroisandt, "Wavelet processes and adaptive estimation of the evolutionary wavelet spectrum," *Journal of the Royal Statistical Society, Series B*, vol. 62, pp. 271–292, 2000.
- [2] I. A. Eckley, G. P. Nason, and R. L. Treloar, "Wavelet fields with application to the modeling and analysis of image texture," *Journal of the Royal Statistical Society, Series C*, 2010.
- [3] I. A. Eckley and G. P. Nason, "LS2W: Locally stationary wavelet fields in R," *Journal of Statistical Software*, vol. 43, no. 3, pp. 1–23, 2011.
- [4] S. Taylor, I. A. Eckley, and M. A. Nunes, "A test of stationarity for textured images," *Technometrics*, vol. 56, pp. 291–301, 2014.
- [5] S. Taylor, M. A. Nunes, and I. A. Eckley, "A multiscale test of spatial stationarity for textured images in R," *The R Journal*, vol. 6, pp. 20–30, 2014.
- [6] J. Portilla and E. P. Simoncelli, "A parametric texture model based on joint statistics of complex wavelet coefficients," *International Journal of Computer Vision*, vol. 40, pp. 49–70, 2000.
- [7] F. Yger and A. Rakotomamonjy, "Wavelet kernel learning," *Pattern Recognition*, vol. 44, pp. 2614–2629, 2011.
- [8] C. W. Shaffrey, N. G. Kingsbury, and I. H. Jermyn, "Unsupervised image segmentation via markov trees and complex wavelets," *IEEE International Conference on Image Processing*, 2002.
- [9] G. Zhang and N. G. Kingsbury, "Variational Bayesian image restoration with group-sparse modeling of wavelet coefficients," *Digital Signal Processing*, vol. 47, pp. 157–168, 2015.
- [10] D. Mondal and D. B. Percival, "Wavelet variance analysis for random fields on a regular lattice," *IEEE Transactions on Image Processing*, vol. 21, no. 2, pp. 537–549, 2012.
- [11] K. Thon, M. Geilhufe, and D. B. Percival, "A multiscale wavelet-based test for isotropy of random fields on a regular lattice," *IEEE Transactions on Image Processing*, vol. 24, no. 2, pp. 694–708, 2015.
- [12] N. G. Kingsbury, "A dual-tree complex wavelet transform with improved orthogonality and symmetry properties," *IEEE International Conference on Image Processing*, 2000.
- [13] N. G. Kingsbury, "Complex wavelets for shift invariant analysis and filtering of signals," *Journal of Applied and Computational Harmonic Analysis*, vol. 10, no. 3, pp. 234–253, 2001.
- [14] 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.
- [15] Daubechies, *Ten Lectures on Wavelets*, SIAM, 1992.
- [16] I. A. Eckley, G. P. Nason, and R. L. Treloar, "Technical appendix to LSW fields with application to the modelling and analysis of image texture," Technical Report, Statistics Group, Lancaster University, 2008.
- [17] P. R. Hill, N. Anantrasirichai, A. Achim, M. E. Al-Mualla, and D. R. Bull, "Undecimated dual-tree complex wavelet transforms," *Signal Processing: Image Communication*, vol. 35, pp. 61–70, 2015.
- [18] A. Cardinali and G. P. Nason, "Costationarity of locally stationary time series," *Journal of Time Series Econometrics*, vol. 2, no. 2, pp. 1–33, 2010.
- [19] S. Palmer, J. Zhang, and X. Wang, "New methods for objective evaluation of fabric pilling by frequency domain image processing," *Research Journal of Textile and Apparel*, vol. 13, pp. 11–23, 2011.
- [20] J. D. B. Nelson and N. G. Kingsbury, "Fractal dimension based sand ripple suppression for mine hunting with sidescan sonar," *International Conference on Synthetic Aperture Sonar and Synthetic Aperture Radar*, 2010.
- [21] J. D. B. Nelson and N. G. Kingsbury, "Fractal dimension, wavelet shrinkage and anomaly detection for mine hunting," *IET Signal Processing*, vol. 6, no. 5, pp. 484–493, 2012.
- [22] M. Hussain, D. Chen, A. Cheng, H. Wei, and D. Stanley, "Change detection from remotely sensed images: From pixel-based to object-based approaches," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 80, pp. 91–106, 2013.