A test of stationarity for textured images

Sarah L. Taylor*
Idris A. Eckley
and
Matthew A. Nunes†

Department of Mathematics and Statistics, Lancaster University

July 1, 2013

Abstract

This article proposes a test of stationarity for random fields on a regular lattice motivated by a problem arising from texture analysis. Our approach is founded on the locally stationary two-dimensional wavelet (LS2W) process model for lattice processes which has previously been used for standard texture analysis tasks such as texture discrimination and classification. We propose two variants of our stationarity test, both of which can be performed on a single realisation – a feature of particular practical importance within texture analysis. We illustrate our approach with pilled fabric data, demonstrating that the test is capable of identifying visually-subtle changes in stationarity. Supplementary material for this article is available online.

Keywords: bootstrapping; fabric analysis; lattice processes; non-stationarity; random field.

*The authors gratefully acknowledge the financial support of Unilever Research Ltd. Taylor also acknowledges funding from EPSRC. We thank Robert Treloar and Eric Mahers of Unilever Research for discussions which have helped stimulate this research. We would like to thank the editorial team for their encouragement to develop the ideas in an earlier version of this paper which has led to the current, much improved manuscript. We are also grateful to Claudia Kirch for helpful comments during the revision of this article.

†Corresponding author: m.nunes@lancaster.ac.uk

1 Introduction

This article is motivated by an application emerging from the analysis of textured images. When one thinks about *texture*, a typical example that comes to mind is that of a woven material, straw or a brick wall. More formally, *image texture* is the visual property of an image region with some degree of regularity or pattern: it describes the variation in the data at smaller scales than the current perspective (Petrou and Sevilla, 2006). Texture structure can be thus considered to exist on several different *scales* of an image. Moreover it is well-documented that image processing within the mammalian visual system is performed in such a manner to preserve local and global information, see for example Daugman (1990) or Field (1999). It is therefore desirable that texture analysis tools reflect these two important properties of texture, namely that (i) it has a location-dependent structure and (ii) it is *multiscale* in nature.

The image set which we consider (Figure 1) arises from work with an industrial collaborator. It comprises six fabrics buffed to varying degrees in an attempt to simulate different levels of garment wear. The effect of this abrasive process is to induce pilling – a build up of fibrous clusters on the surface of the material. By its very nature pill is a localised (short memory) phenomenon with the amount of pill in any region dependent on the amount of wear in that particular area (see Chan and Pang (2000) and Palmer et al. (2011) for further details). It is therefore of considerable interest to be able to accurately assess the pill level, for example by means of classification against a reference set of data.

Within the field of texture analysis researchers are typically interested in three broad activities, namely texture discrimination, classification and segmentation. As such, appropriate and efficient modelling of the second-order properties of an image can often be an important consideration. Many established techniques for texture analysis have an underlying assumption of (second-order) stationarity, see for example Gonzalez and Woods (2001). In other words, the process has a constant mean, but the covariance between two spatial locations is a function of the vector difference between them: $\text{cov}(X_{\mathbf{r_1}}, X_{\mathbf{r_2}}) = \gamma(\mathbf{r_1} - \mathbf{r_2})$.

Conversely, to account for the inherent multiscale structure of such images, others have proposed the use of wavelet-based approaches, see for example Laine and Fan (1993), Unser (1995) or Eckley et al. (2010). In particular the approach proposed by Eckley et al. (2010)

provides a mechanism for modelling and estimating the non-stationary structure of textured images.

Wavelets are a form of localised basis function which provide a scale based decomposition of an image's structure (Vidakovic, 1999; Nason, 2008). Recent research by Eckley et al. (2010) indicated that when a textured image appeared visually to be stationary then, as one might anticipate, Fourier-based classification approaches provide improved classification performance when compared against their wavelet counterparts. Conversely when an image is non-stationary, a wavelet based approach is more appropriate.

In practice one will not generally know *a priori* whether or not the images being analysed are (second-order) stationary. It is therefore difficult to identify whether a multiscale or stationary approach should be adopted to analyse a given set of images. One way of resolving this issue would be to develop a test of (second-order) stationarity for such short memory processes based on a single realisation of each image. This is the question which we address in this paper, developing a new test of stationarity for random fields, highlighting its application with textured images.

Existing tests of spatial stationarity. Several tests of stationarity for spatial processes have been proposed in the literature in recent years, with contributions coming from statistics, geology and the environmental sciences. We now consider the suitability of these to the texture analysis application described above. In particular we seek to identify an approach capable of using single process realisations and, ideally, which takes into account the scale-based nature of images.

Ephraty and collaborators have suggested a number of tests of spatial stationarity. Initial work, described in Ephraty et al. (1996), proposed a test statistic for stationarity calculated using the l_2 norm of the off-diagonal elements of the second-order cumulant spectrum (the Fourier transform of the image cross-correlation function). Under the null hypothesis of stationarity, only the diagonal elements of the matrix should be zero. A relatively large sample size is needed to ensure the accuracy of the method. Further work introduced by Ephraty et al. (2001) introduced a likelihood-based test and also a test using spectral estimation methods under the assumption of a low degree of stationarity. However, unfortunately this approach also requires multiple realisations of the process – a requirement that is often not

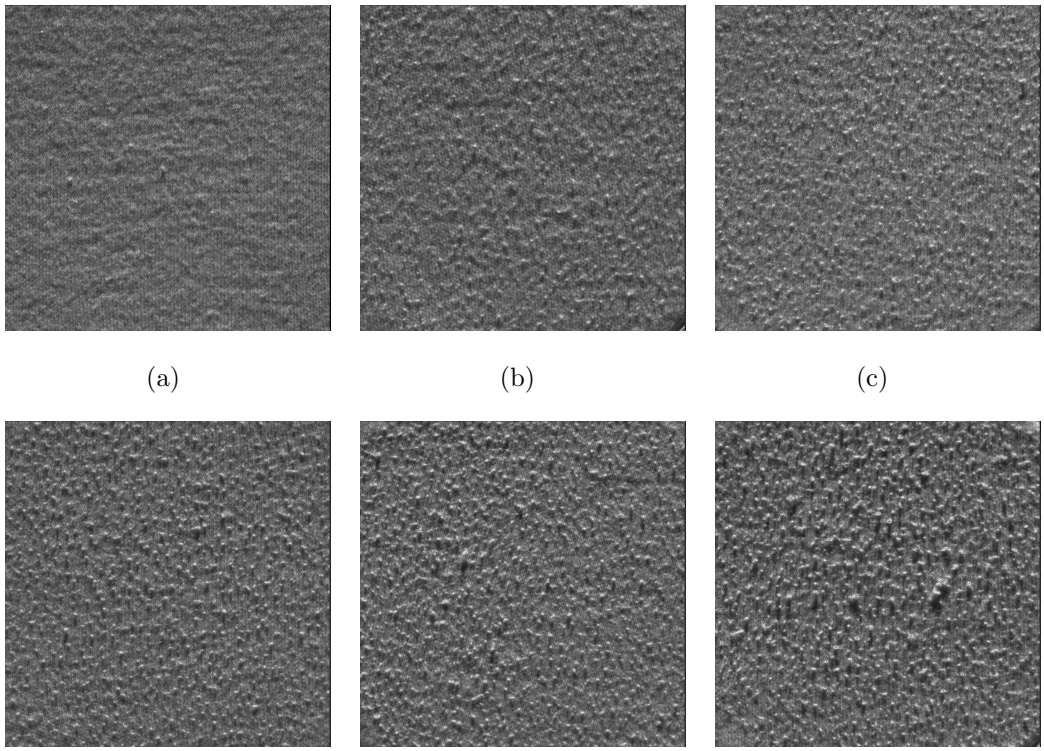

(a) (b) (c)

(d) (e) (f)

Figure 1: A sequence of six pilled fabric images. The amount of pilling increases across the images from (a) (lowest pill) to (f) (highest pill). The first five images clearly show a high degree of stationarity across the pills; due to the increased pilling, the sixth image shows small areas of uneven bobbling.

possible in texture analysis and remote sensing settings.

Bose and Steinhardt (2002) formulate a hypothesis test using the generalised likelihood ratio statistic under the assumption of a centrosymmetric form of the spatial covariance. The covariance of stationary processes is known to have this property. This in turn means that subspaces spanned by particular eigenvectors can be inspected for orthogonality under the null hypothesis. However, the test also assumes multiple realisations of the spatial process and, as noted by Fuentes (2005), the proposed test is likely to be sensitive to the form of covariance.

A hierarchical Bayesian approach is proposed by Fuentes (2005) who models the continuous spatial process using a parametric form for the covariance structure and estimating the spectral density of a process via weight functions evaluated on windows. The author tests for

stationarity by extending the 'analysis of variance' approach of Priestley and Rao (1969) to spatial processes. The dependence of the approach on the choice of window (weight function bandwidth) is an obvious computational drawback. Additionally we seek a *discrete-spatial* approach which can encapsulate the scale-based nature inherent within textured images.

Finally, Blanc et al. (2008) investigate the large sample behaviour of the empirical mean and variance statistics over a class of spatial processes for which the theoretical asymptotic behaviour is known. The rates of convergence of the statistics for an observed spatial process are estimated by image subsampling and then fitting a nonparametric estimator. Non-stationarity of the mean or variance is detected by looking for 'anomalies' in the behaviour of the empirical statistics compared with theoretical rates. Evaluating the large sample trend of the statistics is obviously computationally intensive. In addition, there is no clear measure of a sufficiently convergent trend and thus automatic implementation of the method is not considered. Hence, detection of non-stationarity is achieved through visual inspection. This approach is therefore not suitable for use with textured images where typically one may have a large number of candidate textures which need to be (automatically) analysed to identify whether they are stationary.

A wavelet-based approach? Each of the above methods suffer various disadvantages for the application under consideration. The two tests of stationarity which we propose in Section 3 adopt the recently proposed wavelet-based model of Eckley et al. (2010). Specifically, the locally stationary framework of Eckley et al. (2010) provides the flexibility to accommodate realistic non-stationary behaviour whilst also being able to model the inherent multiscale structure of texture. It is therefore natural to consider whether this framework can be used to develop a test of stationarity for short memory random fields, such as the pill images, particularly since the wavelet-based framework permits estimation of the local spectrum with a single realisation of the spatial process. As such our proposed tests do not suffer from the limitation of other tests of stationarity in requiring multiple process realisations. Moreover we find that they are able to detect quite subtle locally non-stationary behaviour of the spatial process and demonstrate how this can be applied in the texture context. The first test of stationarity we introduce extends the work of Cardinali and Nason (2011) from univariate time series to the spatial setting, examining average spectral variation

to assess the stationarity of an image. Since, in the one-dimensional setting, this approach has been observed to be conservative in nature, we also consider a reformulation of the test as a multiple hypothesis test.

The article is organised as follows. We begin, in Section 2, by providing a brief introduction to wavelets and two-dimensional locally stationary wavelet processes. We then propose our two tests of stationarity in Section 3, providing assessments of their performance through simulation in Section 4. Our tests are then applied to several texture examples provided by an industrial collaborator (Sections 4.2 and 4.3).

2 Wavelets and 2D locally stationary wavelet (LS2W) processes

Briefly, wavelets are oscillatory basis functions which provide efficient (sparse) multiscale representations of signals. For example, for a function $f \in L^2(\mathbb{R})$, we have the expression $f(x) = \sum_{k\in\mathbb{Z}} c_{0,k}\phi_{0,k}(x) + \sum_{j\leq J}\sum_{k\in\mathbb{Z}} d_{j,k}\psi_{j,k}(x)$, where the wavelet $\psi_{j,k}(x) = 2^{-j/2}\psi(2^{-j}x - k)$ is a dilated and translated version of a (mother) wavelet ψ and similarly for the father wavelet $\phi(x)$. The coefficients $d_{j,k}$ at location k and scale j represent the oscillatory behaviour of the signal f at a particular frequency, whereas the coefficients $c_{j,k}$ give information about the mean behaviour of the signal at different scales j.

Wavelets have received considerable attention within the statistics community during the last 20 years, not least because of their ability to provide an efficient location-scale decomposition of signals (see Vidakovic (1999), Percival and Walden (2006) or Nason (2008) for accessible introductions to this area). Below we provide a brief overview of the pertinent theory associated with locally stationary two-dimensional wavelet processes.

2.1 Discrete wavelets

We begin by providing a formal definition of the key building blocks within the LS2W framework, namely discrete wavelets.

Let ψ be a (compactly supported) wavelet, such as those introduced by Daubechies (1992), and denote by $\{h_k, g_k\}$ the low- / high-pass filter pair associated with ψ. Furthermore,

let N_h denote the number of non-zero coefficients of $\{h_k\}$, and define $L_j = (2^j - 1)(N_h - 1) + 1$.

The one-dimensional *discrete wavelets* at a given scale $j \in \mathbb{Z}^+$ as introduced by Nason et al. (2000) are defined to be the vectors $\psi_j = \left(\psi_{j,0}, \ldots, \psi_{j,L_j-1}\right)$, with $\psi_{-1n} = \sum_k g_{n-2k}\delta_{0k} = g_n$ and $\psi_{(j-1)n} = \sum_k h_{n-2k}\psi_{jk}$ for $n = 0, \ldots, L_{j-1} - 1$, where δ_{0k} is the Kronecker delta. The discrete father wavelet is defined similarly using the associated low-pass filter $\{h_k\}$. As Eckley et al. (2010) note, this construction can easily be extended to two dimensions as follows:

Definition 1. Let $\mathbf{k} = (k_1, k_2)$ where $k_1, k_2 \in \mathbb{Z}$. The 2D discrete wavelet filters, $\{\psi_j^l\}$, are defined as the finite square matrices (of dimension $L_j \times L_j$):

$$\psi_j^l = \begin{bmatrix} \psi^l_{j,(0,0)} & \cdots & \psi^l_{j,(0,L_{j-1})} \\ \vdots & \vdots & \vdots \\ \psi^l_{j,(L_{j-1},0)} & \cdots & \psi^l_{j,(L_{j-1},L_{j-1})} \end{bmatrix},$$

for horizontal, vertical or diagonal directions $l = h, v$ and d, where the elements of the wavelets are the tensor products of the corresponding 1D discrete wavelets: $\psi^h_{j,\mathbf{k}} = \phi_{j,k_1}\psi_{j,k_2}$; $\psi^v_{j,\mathbf{k}} = \psi_{j,k_1}\phi_{j,k_2}$ and $\psi^d_{j,\mathbf{k}} = \psi_{j,k_1}\psi_{j,k_2}$.

Discrete father wavelets $\phi_{j,\mathbf{k}}$ can be defined similarly in two dimensions by taking the tensor product of 1D discrete father wavelets.

Finally, a family of *nondecimated* discrete wavelets is formed via translations in $\mathbb{Z}^2$ as $\psi^l_{j,\mathbf{u}}(\mathbf{r}) = \psi^l_{j,\mathbf{u}-\mathbf{r}}$ for every scale j, direction l and locations $\mathbf{u}, \mathbf{r} \in \mathbb{Z}^2$. It is these which we use in the spatial process model introduced in the next section.

2.2 The LS2W spatial model

The test of stationarity which we introduce in Section 3 extends ideas presented in a time series context by Cardinali and Nason (2011) to a spatial setting. To achieve this, we adopt the locally stationary spatial modelling framework introduced by Eckley et al. (2010). This introduced a new class of multiscale lattice processes with a location-dependent second-order structure. Instead of assuming a stationary process behaviour, these processes are assumed to have a *locally stationary* character. In other words, the covariance is assumed to vary across (pixel) locations of an image, as typically seen in many everyday examples of texture, e.g. wear on a garment made from woven fabric. Eckley et al. (2010) refer to spatial processes

constructed under such a model as *locally stationary wavelet fields* (LS2W). We now provide a brief introduction to the main elements of the LS2W modelling approach.

The locally stationary two-dimensional wavelet process model introduced by Eckley et al. (2010) is defined as

$$X_{\boldsymbol{r};\boldsymbol{R}} = \sum_{l} \sum_{j=1}^{\infty} \sum_{\boldsymbol{u}} w_{j,\boldsymbol{u}}^{l} \psi_{j,\boldsymbol{u}}^{l}(\boldsymbol{r}) \xi_{j,\boldsymbol{u}}^{l}, \tag{1}$$

for horizontal, vertical and diagonal directions $l = h, v$ or d and spatial locations $\boldsymbol{r} = (r, s) \in \{0, \ldots, R-1\} \times \{0, \ldots, S-1\} = \mathcal{R}$, where $R = 2^m$, $S = 2^n \geq 1$, with $n, m \in \mathbb{N}$ and we denote the lowest common scale as $J(R, S) \equiv \log_2\{\min(R, S)\}$. In (1),

(i) $\{\xi_{j,\boldsymbol{u}}^{l}\}$ is a zero-mean random orthonormal increment sequence;

(ii) $\{\psi_{j,\boldsymbol{u}}^{l}\}$ are a set of discrete nondecimated wavelets (see Definition 1);

(iii) $\{w_{j,\boldsymbol{u}}^{l}\}$ are a collection of amplitudes;

(iv) and the rate of evolution of the second-order properties of the image are controlled by smoothness constraints (see Eckley et al. (2010) for further details).

Note that the locally stationary wavelet process has a dependence on the dimension of the image, $\boldsymbol{R} = (R, S)$. However, for notational convenience we drop this explicit dependence and denote such a process as $X_{\boldsymbol{r}}$, though the dependence should be assumed.

Analogous to Fourier-based spectral theory, one can define the *local wavelet spectrum* (LWS) associated with an LS2W process. The LWS for a given location $\boldsymbol{z} \in (0, 1)^2$, at scale j in direction l is $S_j^l(\boldsymbol{z}) \approx w_j^l(\boldsymbol{u}/\boldsymbol{R})^2$, where $\boldsymbol{z} = \boldsymbol{u}/\boldsymbol{R} := (u/R, v/S)$. As such the LWS quantifies the contribution to the process variance at *rescaled* spatial locations $\boldsymbol{z} \in (0, 1)^2$, decomposition directions l, and scales, j (see Eckley et al. (2010) for a discussion of why we use rescaled spatial locations).

To assess the stationarity, or otherwise, of a textured image, the tests proposed in Section 3 use an estimate of the LWS. Drawing parallels with estimation theory associated with the Fourier spectral density, Eckley et al. (2010) propose the following estimator for the LWS:

$$I_{j,\boldsymbol{u}}^{l} = |d_{j,\boldsymbol{u}}^{l}|^2 = \left(\sum_{\boldsymbol{r}} X_{\boldsymbol{r}} \psi_{j,\boldsymbol{u}}^{l}(\boldsymbol{r}) \right)^2, \tag{2}$$

where $d^l_{j,\boldsymbol{u}} = \sum_{\boldsymbol{r}} X_{\boldsymbol{r}} \psi^l_{j,\boldsymbol{u}}(\boldsymbol{r})$ denotes the empirical wavelet coefficients of the LS2W process, $X_{\boldsymbol{r}}$, at a particular location, scale and direction. Defining $\boldsymbol{z} = \boldsymbol{u}/\boldsymbol{R}$ as above, the array $\boldsymbol{I}(\boldsymbol{z}) = \{I^l_{j,\boldsymbol{u}}\}$ for $j = 1, \ldots, J$, $l \in \{h, v, d\}$, and locations $\boldsymbol{u} \in \mathcal{R}$ in (2) is referred to as the *local wavelet periodogram* (LWP). It is biased as an estimator for the LWS. However Eckley et al. (2010) established that the periodogram estimator can be bias-corrected using the inverse of the inner product matrix of discrete autocorrelation wavelets, i.e.

$$\hat{S}(\boldsymbol{z}) = \boldsymbol{L}(\boldsymbol{z}) = A_J^{-1} \boldsymbol{I}(\boldsymbol{z}), \tag{3}$$

where A_J is the two-dimensional discrete autocorrelation wavelet matrix (see Eckley and Nason (2005) for further details). It is this bias-corrected version of the LWP which we incorporate within our tests of stationarity, introduced in the next section.

3 Testing for stationarity in LS2W processes

We now introduce our tests of stationarity for textured images. The tests take the form of a hypothesis test for which a particular statistic is used to measure the degree of non-stationarity under the null hypothesis of stationarity. In particular, Section 3.1.1 introduces a test based on a statistic of average spectral variation; in Section 3.1.2 we consider a multiple testing framework using spectral test statistics specific to each scale-direction pair.

To formalise the hypothesis test we note that a (spatial) process is stationary if and only if its spectrum is constant across locations for all scale-direction pairs. In other words, we can test a process spectrum for constancy in order to determine whether the process is stationary. Thus given an observed process, $X_{\mathbf{r}}$, with associated wavelet spectrum $\boldsymbol{S} = \{S^l_j(\boldsymbol{z})\}_{j,l}$, our hypothesis test is

$$H_0 \colon S^l_j(\boldsymbol{z}) \text{ is a constant function of } \boldsymbol{z} \text{ for all } j > 0 \text{ and } l \in \{h, v, d\}$$
$$H_A \colon S^l_j(\boldsymbol{z}) \text{ is not a constant function of } \boldsymbol{z} \text{ for some scale } j \text{ and direction } l.$$

In view of the above observation, the null hypothesis corresponds to an assumption of stationarity of a spatial process. We therefore look for departures from constancy within each scale and direction of the local wavelet spectrum to indicate non-stationarity. Note that this

departure from constancy for a fixed scale-direction pair can be measured by:

$$T\{S_j^l(\boldsymbol{z})\} = \int \{S_j^l(\boldsymbol{z}) - \bar{S}_j^l\}^2 \, d\boldsymbol{z} \quad \text{with } \boldsymbol{z} \in (0,1)^2. \tag{4}$$

Here $\bar{S}_j^l = \int S_j^l(\boldsymbol{z}) \, d\boldsymbol{z}$ for a particular direction $l \in \{h, v, d\}$ and scale level $j = 1, \ldots, J$. The hypothesis test above can thus be performed by examining the spectrum for departures from constancy in *any* scale-direction pair.

Similarly, the *average* variation in the wavelet spectrum can be quantified using the average measure:

$$T_{ave}\{\boldsymbol{S}(\boldsymbol{z})\} = (3J)^{-1} \sum_l \sum_{j=1}^{J} \int \{S_j^l(\boldsymbol{z}) - \bar{S}_j^l\}^2 \, d\boldsymbol{z} \quad \text{with } \boldsymbol{z} \in (0,1)^2. \tag{5}$$

Note that the quantity in (5) is zero if, and only if, the spectrum $S_j^l(\boldsymbol{z})$ is constant across locations for each scale-direction pair.

Since the spectrum $\boldsymbol{S}(\boldsymbol{z})$ is unknown, in practice it is replaced by an estimator, $\hat{\boldsymbol{S}}(\boldsymbol{z})$, for example the corrected LWS $\boldsymbol{L}(\boldsymbol{z})$ described in Section 2.2. Note that, for a given realisation, $\hat{S}_j^l(\boldsymbol{z})$ can be denoted $\hat{S}_{j,\boldsymbol{u}}^l$. In addition, the integral $\bar{S}_j^l$ can be estimated by

$$\widetilde{S}_j^l = (RS)^{-1} \sum_{\boldsymbol{u} \in \mathcal{R}} \hat{S}_{j,\boldsymbol{u}}^l. \tag{6}$$

Taking $\hat{S}_{j,\boldsymbol{u}}^l = L_{j,\boldsymbol{u}}^l$ (the corrected LWP) in (6) results in a consistent estimator for the spectrum under the assumption of stationarity (see supplementary material for a proof which establishes this property).

We can also estimate the integral in (4) and (5) using the empirical variance (across locations) for a scale-direction spectrum pair. Hence our test statistic for a fixed scale-direction pair can be calculated as follows

$$T\{\hat{S}_j^l(\boldsymbol{z})\} = \mathrm{var}_{\boldsymbol{u}}(\hat{S}_{j,\boldsymbol{u}}^l), \tag{7}$$

where the variance is taken over the lattice $\mathcal{R}$. The alternative *average variation* test statistic (5) can also be calculated as

$$T_{ave}\{\hat{\boldsymbol{S}}(\boldsymbol{z})\} = (3J)^{-1} \sum_l \sum_{j=1}^{J} \mathrm{var}_{\boldsymbol{u}}(\hat{S}_{j,\boldsymbol{u}}^l). \tag{8}$$

The test statistic in (8) can be seen as the mean empirical variance of the spectrum estimate, where the average is taken over all scale-direction pairs.

3.1 Two bootstrap tests of stationarity

To perform the above hypothesis test we require knowledge of the distribution of the test statistic $T_{(\cdot)}\{\hat{S}_j^l(\mathbf{z})\}$, under the null hypothesis. Unfortunately, in most practical cases this distribution will in general be unknown. However, the spectrum $\boldsymbol{S}$ determines the distribution of the observed random field (assuming stationarity). We therefore perform (parametric) bootstrap resampling on the model innovations $\{\xi_{j,\boldsymbol{u}}^l\}$ together with the assumed stationary spectral structure to establish the distribution of the test statistic under H_0.

As noted in Cardinali and Nason (2011), the estimator $\hat{S}_j^l(\boldsymbol{z})$ in Section 3 needs to be a consistent estimator for the stationary spectrum (under the null hypothesis) for the parametric bootstrap implementation to be valid. To this end, we establish the consistency of our estimator $\bar{\boldsymbol{L}} = (RS)^{-1}\sum_{\boldsymbol{u}\in\mathcal{R}}\boldsymbol{L}_{\boldsymbol{u}}$ in the supplementary material accompanying this article. This result holds for both bootstrap implementations outlined below.

3.1.1 Testing for stationarity using average spectral variation

We now outline how our test of stationarity is performed using the test statistic T_{ave} in (8). Our approach extends earlier work by Cardinali and Nason (2011) from a time series to a two-dimensional setting. The bootstrap test works by first computing the spectral estimate $\boldsymbol{L}$ for the observed image and then calculating the test statistic t_{ave}^{obs} (using (8)). We then obtain simulated realisations by the following procedure (Algorithm 1). In essence we simply simulate the underlying innovations $\{\xi_{j,\boldsymbol{u}}^l\}$ (see Remark 1) and feed these into our process model under the null hypothesis, inverting to obtain a realisation in the spatial domain. For each simulated stationary process, we compute the test statistic (8). The significance of the test statistic for the observed image can then be assessed by appealing to Monte Carlo arguments (see Davison and Hinkley (1997) for more details). Our bootstrap approach, which we call $\texttt{Bootstat}_{\texttt{LS2W}}$, is summarised in Algorithm 1.

The $\texttt{Bootstat}_{\texttt{LS2W}}$ test can be interpreted as evaluating how unlikely the value of t_{ave}^{obs} is compared to realistic (bootstrap) values of T_{ave} assuming a *stationary* spectral structure based on the observed process. Thus the p-value of the test can be seen as a measure of how non-stationary the observed process is.

`Bootstat`$_{\texttt{LS2W}}$`:`

1. Compute the estimate of the LWS for the observed image, $\hat{S}_j^l(\boldsymbol{z})$.
2. Evaluate T_{ave} on $\hat{\boldsymbol{S}}(\boldsymbol{z})$ (using (8)), call this value t_{ave}^{obs}.
3. Compute the pixel average stationary spectrum $\widetilde{S}_j^l$ by taking the average of spectrum values for each scale and direction.
4. **Iterate** for i in 1 to B bootstraps
   (a) Simulate $X_{\boldsymbol{r}}^{(i)}$ from the stationary LS2W model using squared amplitudes given by $\widetilde{S}_j^l$ and Gaussian process innovations.
   (b) Compute the test statistic T_{ave} on the simulated realisation, call this value $t_{ave}^{(i)}$.
5. Compute the p-value for the test as $p = \frac{1+\#\left\{ t_{ave}^{obs} \leq t_{ave}^{(i)} \right\}}{B+1}$.

Algorithm 1: The bootstrap algorithm for testing the stationarity of locally stationary images.

3.1.2 Testing for stationarity using a multiple testing framework

In the one-dimensional setting, the bootstrap procedure outlined in Section 3.1.1 can be conservative in some cases (Cardinali and Nason, 2011). In the time series setting, multiple testing approaches have recently been shown to ameliorate conservative stationarity tests (Nason, 2013). As an alternative to using the average variation statistic (8), the test of stationarity can also be formulated in a multiple hypothesis framework and thus we explore the potential of this approach to testing spatial stationarity. We base our approach to this multiple testing problem on the `Bootstat`$_{\texttt{LS2W}}$ bootstrap test in Section 3.1.1.

Recall from Section 3 that we can test each scale-direction spectral plane for constancy. We wish to reject the null hypothesis of stationarity if there is a significant departure from constancy in *any* of the scale-direction planes of the local wavelet spectrum. In other words,

for each given direction-scale pair, (j, l), we must test to see whether

$$H_0 \colon S_j^l(\boldsymbol{z}) \text{ is a constant function of } \boldsymbol{z}$$
$$H_A \colon S_j^l(\boldsymbol{z}) \text{ is not a constant function of } \boldsymbol{z}.$$

This test must then be repeated for all possible scale-direction pairs. In other words we are in a multiple hypothesis testing scenario, and draw on the the work of Davison and Hinkley (1997, Chapter 4.4) for multiple testing in a bootstrapping context. Let $\eta = \eta(j, l) \in \{1, \ldots, 3J\}$ represent an index over all scale-direction pairs. To account for possible differences in the distribution of the statistics $T_{\eta(j,l)} := T\{S_l^j(\boldsymbol{z})\}$ (see (7)), as well as possible dependence between them, Davison and Hinkley (1997) suggest examining the quantity

$$Q = \min\{P_1, \cdots, P_{3J}\}, \tag{9}$$

where P_η denotes the true tail probability associated to the test statistic T_η for a fixed scale-direction spectral plane. In other words, to test the stationarity or otherwise of an observed image, we regard Q as the test statistic associated to the null hypothesis.

The algorithm for assessing the significance of Q is described in Algorithm 2. We denote this multiple hypothesis testing procedure as $\texttt{Bootstat}^{\texttt{mh}}_{\texttt{LS2W}}$.

Remark 1. Within both the bootstrap procedures described above, the realisations of the LS2W process under the null hypothesis are simulated by assuming the innovations of the process take a given distributional form. In particular we have assumed that the innovations are Gaussian; however, other distributions can be used in the model representation (1).

4 Examples

We now consider the performance of the two proposed tests of stationarity proposed in Section 3.1, focussing in particular on data generated from an industrial application. The analyses were performed using the LS2W spectral estimation implementation in the R package *LS2W* (Eckley and Nason, 2011). For simulating the Gaussian random fields in the study below, we used the *RandomFields* R package (Schlather, 2012); the spatial moving average

$\texttt{Bootstat}^{\texttt{mh}}_{\texttt{LS2W}}$:

1. Compute the estimate of the LWS for the observed image, $\hat{S}_j^l(\boldsymbol{z})$.
2. Evaluate T_η on the observed image for all j and l (using (7)) , call these values t_η^{obs}, $\eta = 1, \ldots, 3J$.
3. Compute $\widetilde{S}_j^l$, $j = 1, \ldots, J$ for each of the decomposition directions for the observed image using (6).
4. **Iterate** for i in 1 to B bootstraps
   (a) Simulate $X_{\boldsymbol{r}}^{(i)}$ from the stationary LS2W model using squared amplitudes given by $\widetilde{S}_j^l$ and Gaussian process innovations (see Section 3.1).
   (b) Compute the test statistics T_η on the simulated realisation, for all j and l, call these values $t_\eta^{(i)}$.
   (c) Compute $\widetilde{S}^{(i)}$, for the *simulated realisation* image $X_{\boldsymbol{r}}^{(i)}$ using (6).
   (d) **Iterate** for k in 1 to M (secondary) bootstraps
      i. Simulate $X_{\boldsymbol{r}}^{(i)(k)}$ from the LS2W model using squared amplitudes given by $\widetilde{S}^{(i)}$.
      ii. Compute the test statistics T_η on the second-level simulated realisation, for all j and l, call these values $t_\eta^{(i)(k)}$.
   (e) Compute the test statistic $q^{(i)} = \min\left\{ \frac{1+\#\left\{ t_1^{(i)} \leq t_1^{(i)(k)} \right\}}{M+1}, \cdots, \frac{1+\#\left\{ t_{3J}^{(i)} \leq t_{3J}^{(i)(k)} \right\}}{M+1} \right\}$.
5. Compute the observed test statistic $q^{obs} = \min\left\{ \frac{1+\#\left\{ t_1^{obs} \leq t_1^{(i)} \right\}}{B+1}, \cdots, \frac{1+\#\left\{ t_{3J}^{obs} \leq t_{3J}^{(i)} \right\}}{B+1} \right\}$
6. Compute the p-value for the test as $p = \frac{1+\#\left\{ q^{obs} \leq q^{(i)} \right\}}{B+1}$.

Algorithm 2: The multiple testing bootstrap algorithm for testing the stationarity of locally stationary images.

processes were simulated with the R package *spdep* (Bivand, 2012) using modifications to code featured in Anselin (2007).

4.1 Simulated performance of the LS2W tests of stationarity

In order to investigate the performance of our tests of stationarity proposed in Section 3.1, we performed a simulation experiment focussing in particular on size and power properties.

Size assessment. To explore the empirical size of the stationarity tests, we chose a number of different process types. Each process represents a different form of second order *stationary* structure. S1 is a two-dimensional white noise process, i.e $X_{\boldsymbol{r}} \sim N(0,1)$. S2 is a spatial moving average process with parameter $\rho = 0.9$, i.e.

$$X(\boldsymbol{u}) = 0.9W\varepsilon + \varepsilon,$$

where ε is a random component with variance σ^2, W is a spatial weight matrix associated to a chosen neighbour structure. A realisation of this process is shown in Figure 2(b). We also consider an isotropic Gaussian random Field with a Matérn covariance (Matérn, 1960) with shape parameter $\nu = 1$, specified by

$$C(\boldsymbol{u}, \boldsymbol{v}) = \sigma^2\sqrt{2}\|\boldsymbol{u} - \boldsymbol{v}\|K_1\Big(\sqrt{2}\|\boldsymbol{u} - \boldsymbol{v}\|\Big), \tag{10}$$

where $K_1(\cdot)$ denotes the modified Bessel function of order 1. See for example Stein (1999) for more details of this covariance function. We denote this process by S3 (see Figure 2(c)).

The fourth stationary process, S4, is an exponential covariance with range parameter $\phi = 2$:

$$C(\boldsymbol{u}, \boldsymbol{v}) = \sigma^2 \exp\left\{\frac{-\|\boldsymbol{u} - \boldsymbol{v}\|}{2}\right\}. \tag{11}$$

This process is similar to that considered for continuous spatial processes in Fuentes (2005). A realisation of this process is shown in Figure 2(d).

For the processes considered, we simulated $K = 1000$ spatial realisations for different square image sizes $R = S = 2^J$, with $J = 6, \ldots, 9$. We then examined each realisation with each test of stationarity as follows. We evaluated the Monte Carlo significance test in Section 3.1 using $B = 250$ bootstrap simulations, each time treating the realisation as observed. In other words, we perform the $\texttt{Bootstat}_{\texttt{LS2W}}$ (respectively $\texttt{Bootstat}^{\texttt{mh}}_{\texttt{LS2W}}$) hypothesis test and record whether the realisation is stationary or not at a 5% significance level. For each process, we then note the number of simulated realisations which resulted in rejecting the null hypothesis of stationarity. We then compare the proportion of those rejecting the null with the nominal size.

Table 1 explores the size properties of the $\texttt{Bootstat}_{\texttt{LS2W}}$ and $\texttt{Bootstat}^{\texttt{mh}}_{\texttt{LS2W}}$ tests on the stationary specifications expressed as a percentage of the $K = 1000$ images rejecting stationarity, for $\sigma = 1$. We observed similar results for all values of the standard deviation, σ; these more detailed results can be found in tabular form in the supplementary material.

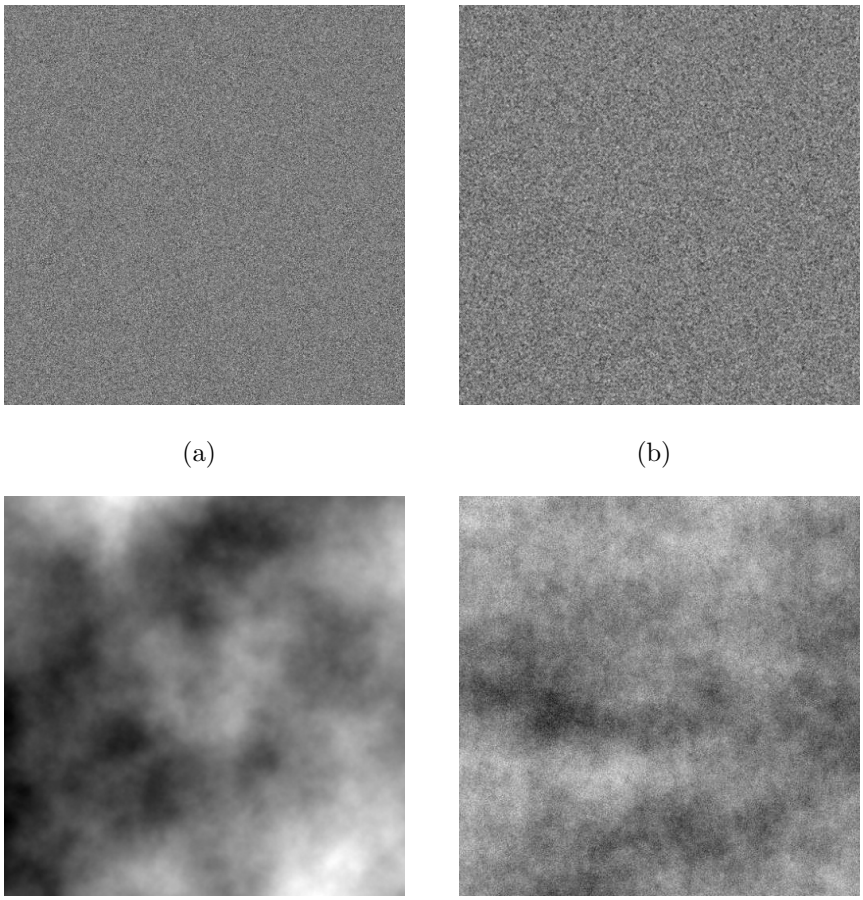

(a) (b)

(c) (d)

Figure 2: Process realisations from the stationary models for the simulations in Section 4.1 with $\boldsymbol{R} = 2^9 \times 2^9$. (a) S1. A two-dimensional Gaussian field; (b) S2. A moving average process; (c) S3. Matérn Gaussian random field with covariance specified by (10); (d) S4. An exponential Gaussian random field with covariance (11).

Focussing first on results of the $\texttt{Bootstat}_{\texttt{LS2W}}$ test of stationarity: for all model types, the percentage of images rejecting the null hypothesis of stationarity (i.e. judged as non-stationary) was well below the nominal size of 5%. These results mirror those obtained by Cardinali and Nason (2011) for the time series setting – namely that this test approach is conservative. Conversely, the multiple testing bootstrap procedure, $\texttt{Bootstat}^{\texttt{mh}}_{\texttt{LS2W}}$ achieves results much nearer the nominal size of 5%, particularly for large image dimensions. This is a desirable feature of the nested bootstrap approach, though this comes at a considerable added computational cost.

Next we explore the power of our test to identify whether power is lost or maintained as

Table 1: Results from a simulation experiment to assess the empirical size of the LS2W test of stationarity. The table indicates the size (rounded to the nearest 0.1%) using a 2D white noise stationary process (S1); a spatial moving average process (S2); a Gaussian random field with Matérn covariance (S3); an exponential Gaussian random field (S4). See text for details of the process used.

	$\texttt{Bootstat}_{\texttt{LS2W}}$				$\texttt{Bootstat}^{\texttt{mh}}_{\texttt{LS2W}}$			
	Image dimension				Image dimension			
Model	64	128	256	512	64	128	256	512
S1	0.1	0.3	0.3	0.4	2.4	2.8	5.2	4.8
S2	1.0	1.3	0.4	2.3	2.8	3.2	2.4	3.6
S3	1.1	0.4	1.4	2.1	0.8	0.4	2.0	2.1
S4	0.4	0.3	0.2	0.2	0.3	0.4	0.8	1.2

a result of this conservativeness in the tests.

Power assessment. To evaluate the power of the $\texttt{Bootstat}_{\texttt{LS2W}}$ and $\texttt{Bootstat}^{\texttt{mh}}_{\texttt{LS2W}}$ stationarity tests, we consider three contrasting processes which exhibit different non-stationary behaviour. In particular, we consider the following forms. NS1 describes a piecewise white noise spatial process, in which the left half-plane has unit variance, whereas the second (vertical) half-plane has variance σ^2. In other words, we simulate K images, each constructed from $2^{(n-1)} \times 2^n$ i.i.d. $N(0,1)$ samples for $\boldsymbol{z} \in (0,1/2) \times (0,1)$ concatenated with $2^{(n-1)} \times 2^n$ values sampled independently from $N(0,\sigma^2)$ for $\boldsymbol{z} \in (1/2,1) \times (0,1)$. Similarly we also consider a process in which the left half-plane is a unit variance white noise process, with the second half being a realisation from the process S3 with a Matérn covariance (NS2). NS3 represents an image which is a *montage* of four stationary processes in the four quadrants of the image (see Figure 3(b)); this is an LS2W process with spectral structure given by

$$w^d_{j,[2^J\mathbf{z}]} = \begin{cases} \sigma & \text{if } j = 1 \text{ and } \mathbf{z} \in (0,1/2)\,\mathrm{x}\,(0,1/2); \\ \sigma & \text{if } j = 2 \text{ and } \mathbf{z} \in (1/2,1)\,\mathrm{x}\,(0,1/2); \\ \sigma & \text{if } j = 3 \text{ and } \mathbf{z} \in (0,1/2)\,\mathrm{x}\,(1/2,1); \\ \sigma & \text{if } j = 4 \text{ and } \mathbf{z} \in (1/2,1)\,\mathrm{x}\,(1/2,1); \\ 0 & \text{otherwise.} \end{cases} \tag{12}$$

In particular, the process is a montage of four *diagonal Haar moving average* processes with different orders (Eckley et al., 2010).

The specification of σ in realisations of NS1 and NS3 controls the degree of non-stationarity in the simulated images: for low values of σ, the processes describe behaviour which is approximately stationary (i.e. the values are closer to unit variance across the whole image); for higher variance values the boundaries are more marked. For this study, we investigate the performance of the stationarity test with values in the range $\sigma = \{1.2, \ldots, 1.6\}$.

To investigate more subtle changes in second order structure we also define a fourth process, specified in the spatial domain by $X_{\boldsymbol{r}}(\boldsymbol{z}) \sim N(0, \sigma^2_{a;\tau;\delta}(\boldsymbol{z}))$, where $\sigma_{a;\tau;\delta}(\boldsymbol{z})$ is constructed such that it varies spatially across an image. For the simulations below we use a particular choice of the standard deviation function which changes across the horizontal coordinate of the image according to the parameters a, τ and δ in the following way:

$$\sigma_{a;\tau;\delta}(\boldsymbol{z}) = \tau + \frac{\delta - \tau}{1 + \exp(-10a(2z_1 - 1))} \quad \text{for } z_1 \in (0, 1). \tag{13}$$

In other words, the standard deviation of the stochastic process varies smoothly from τ to δ (see Figure 3(c)). Changing the parameter a has the effect of changing the shape of the deviation curve with lower values of a giving a more subtle behaviour across the x-axis. We note that if $\tau = 1$ in (13), as a increases the second order structure of NS4 process realisations will resemble that of realisations from model NS1. A realisation of NS4 can be seen in Figure 3(d).

A similar analysis to the size assessment was performed using 1000 realisations from the non-stationary processes NS1 – NS4 described above. In each case, the 1000 realisations were assessed for stationarity and the number which rejected the null recorded.

Table 2 shows the statistical power results (expressed as a percentage) for the two tests of stationarity when applied to simulated images from processes NS1 – NS3. The results illustrate that both tests are unable to distinguish between the two noise variances in NS1 for the lowest values of σ when the image dimension is small, especially when implementing the $\texttt{Bootstat}_{\texttt{LS2W}}$ stationarity test. However the results improve dramatically for both tests for larger image sizes. This is not particularly surprising since, for low values of σ, the non-stationary behaviour is more difficult to detect because the boundary in the simulated images will appear blurred. However, for intermediate and high values of σ and moderate image

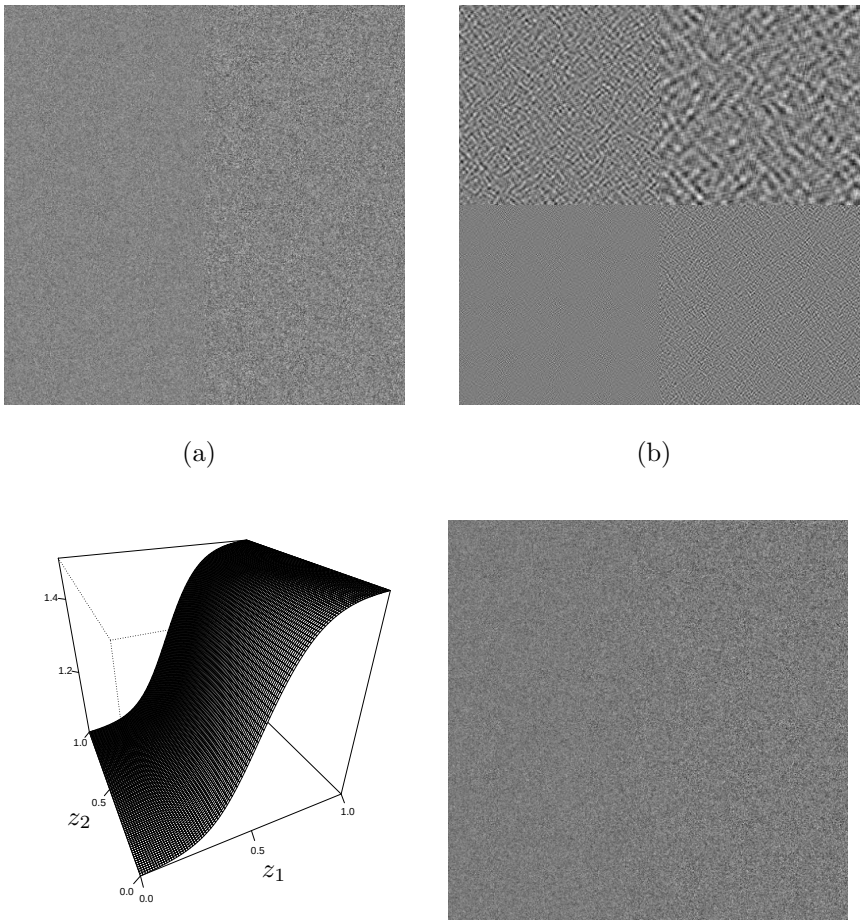

Figure 3: Process realisations from the non-stationary models for the simulations in Section 4.1 with $\boldsymbol{R} = 2^9 \times 2^9$. (a) NS1. The image consists of two Gaussian half planes; (b) NS3 (see (12)). Each quadrant of the image is a $2^8 \times 2^8$ subimage sampled from a different diagonal Haar MA process. The texture ranges from finest detail (bottom-left) to coarsest (top-right); (c) A two-dimensional representation of the second order structure of the process NS4 specified in the spatial domain; (d) A realisation of the model NS4 (see (13)). The power displayed in the image varies smoothly across the horizontal plane according to Figure 3(c).

sizes, both tests attain good performance, rejecting the null hypothesis of stationarity for the majority of $K = 1000$ images (Table 2). The uniformity across σ values for processes NS2 and NS3 is reassuring, indicating that the tests are insensitive to the severity of boundary non-stationarities, even for small variances. Finally we turn to consider NS4. Again the power

Table 2: Results from a simulation experiment to assess the empirical power of the LS2W test of stationarity. The table indicates the power (rounded to the nearest 0.1%) under different non-stationary spectral specifications: piecewise Normal plane (NS1); Gaussian-Matérn process (NS2); Haar Montage (NS3). See text for details of the processes used.

		$\texttt{Bootstat}_{\texttt{LS2W}}$				$\texttt{Bootstat}^{\texttt{mh}}_{\texttt{LS2W}}$			
		Image dimension				Image dimension			
Model		64	128	256	512	64	128	256	512
NS1	σ=1.2	1.3	23.7	98.9	100	17.6	48.4	100	100
	σ=1.3	20.2	97.5	100	100	40.8	100	100	100
	σ=1.5	97.0	100	100	100	99.6	100	100	100
	σ=1.6	100	100	100	100	100	100	100	100
NS2	σ=1.2	50.2	99.6	100	100	43.9	98.2	100	100
	σ=1.3	99.5	100	100	100	100	100	100	100
	σ=1.5	100	100	100	100	100	100	100	100
	σ=1.6	100	100	100	100	100	100	100	100
NS3	σ=1.2	100	100	100	100	100	100	100	100
	σ=1.5	100	100	100	100	100	100	100	100

Table 3: Empirical power assessment of the LS2W test of stationarity for process NS4. The table indicates the power (rounded to the nearest 0.1%) under different location-dependent variance structures given by (13).

	$\texttt{Bootstat}_{\texttt{LS2W}}$				$\texttt{Bootstat}^{\texttt{mh}}_{\texttt{LS2W}}$			
	Image dimension				Image dimension			
(a, τ, δ)	64	128	256	512	64	128	256	512
(0.25,1,1.5)	7.6	77.3	100	100	27.2	85.2	100	100
(0.5,1,1.5)	57.7	100	100	100	68.4	100	100	100
(1,1,1.5)	85.7	100	100	100	98.4	100	100	100

results here mirror those of NS1–NS3. In particular, as observed for NS1, the percentage of correctly classified images improves dramatically as the size of the image under analysis

increases for all scenarios, even for the quite subtle change in structure for NS4 described by low a values in $\sigma_{a;\tau;\delta}(\boldsymbol{z})$ (see Table 3). These results are consistent with the findings of Table 2 for NS1 and $\sigma = 1.5$, which can be seen as an extreme case of the process NS4.

These simulations suggest that for the non-stationary processes considered, both the $\texttt{Bootstat}_{\texttt{LS2W}}$ and the $\texttt{Bootstat}^{\texttt{mh}}_{\texttt{LS2W}}$ tests achieve good statistical power for image dimensions greater than $n = 128$. It is reassuring that power is maintained despite the conservative nature of the tests observed in our exploration of empirical size. These results are consistent with those of other recently published tests of stationarity in the time series literature (see Cardinali and Nason (2011) and Dwivedi and Subba Rao (2011)). In other words, for a stationary process, each of these procedures (including our own) will not reject the null hypothesis to indicate non-stationarity, i.e. report a false positive. However for a non-stationary process they would reject in favour of the alternative. Of our two proposed tests, $\texttt{Bootstat}^{\texttt{mh}}_{\texttt{LS2W}}$ demonstrates a more reasonable empirical size but at significantly increased computational cost.

4.2 Analysis of pilled fabric images

We now apply the tests of stationarity proposed in Section 3.1 to a number of real examples of textured images. In particular, we use the tests of stationarity on a series of images of garment material which have been artificially buffed until the fabric displayed increasing degrees of pilling (build up of clumped fibres). A similar set of images, taken under slightly different lighting conditions, were previously analysed by Eckley et al. (2010) in the context of texture classification. The fabrics are illustrated in Figure 1. Note that visually, all the images appear to be spatially stationary. However, it is arguable that the most heavily-pilled image (Figure 1(f)) displays a degree of non-stationarity due to more increased and irregular "bobbling" of the fabric.

As described in Section 1, our interest is to examine whether the images are indeed assessed as stationary by our stationarity tests introduced in Section 3.1. Prior to performing the tests, as outlined in Section 3, the six images were examined to verify that their spatial structure did not have a significant amount of long-range dependence. In addition, a median polish was applied to each image to remove any non-zero trend (so that the images met the zero-mean assumption of an LS2W process). The $\texttt{Bootstat}_{\texttt{LS2W}}$ and $\texttt{Bootstat}^{\texttt{mh}}_{\texttt{LS2W}}$ stationarity

tests were then applied to each of the six individual images. In each case $B = 250$ bootstrap simulations were made. The p-value associated with each test was then recorded for each fabric image. Note that the size of images analysed in this industrial application is typically large, and so we expect both tests to be reliable in view of Section 4.1.

As one would perhaps anticipate, the $\texttt{Bootstat}_{\texttt{LS2W}}$ test judges the first five images as stationary (with p-values in the range (0.21,0.95)). This concurs with our visual perception of the images. The sixth image is judged as non-stationary, with a p-value of 0.03. The multiple testing procedure $\texttt{Bootstat}^{\texttt{mh}}_{\texttt{LS2W}}$ also assesses the pilled images as stationary apart from the last image, with p-values between 0.11 and 0.97. The sixth image has a p-value of 0.04. This reflects what we might visually assess, namely that the sixth fabric has less regular texture structure than the other fabrics. Indeed, one could argue that this image contains regions of heavier and lighter pill.

The tests have also been applied to other textured images, such as a sequence of images of differing hair types and examples taken from reference texture libraries. In all cases we obtained similar statistical results – namely that mono-textures appear to be stationary. For reasons of brevity we do not report these results here.

4.3 Analysis of texture mosaics

We now turn to consider the application of our test of stationarity for a more realistic situation in fabric analysis. In many settings it is useful to be able to detect differing fabric structure, for example to identify whether there is an area of uneven wear within a sample of material. To avoid the subjectivity of human inspection of materials it is thus desirable to develop an automatic detection method for uneven wear.

We begin by noting that regions which contain uneven wear consist of multiple texture types – the majority of one pill level, with some patches of another pill level. In other words, an image of a fabric which contains uneven wear could be considered to be non-stationary. Consequently a test of stationarity could be used as an automated proxy to detect whether such an area exists within a fabric sample. To examine how our proposed testing methods handle this type of textural non-stationarity, we performed an experiment to mimic the situation described above. More specifically, we constructed some pilled fabric mosaics for analysis with the two bootstrap tests. A texture mosaic comprises two or more

different texture subimages combined into one image for analysis. Two such mosaics were constructed by inserting sections of a texture image from Section 4.2 into another texture. Firstly a subimage of the pilled fabric in Figure 1(c) was inserted as the central part of that in Figure 1(a) (Mosaic A); Mosaic B contains a section of the fabrics in Figure 1(b) as well as Figure 1(c) within the lightest pilled fabric (Figure 1(a)). The mosaics can be seen in Figure 4. The inserted subimages are indicated within the figures.

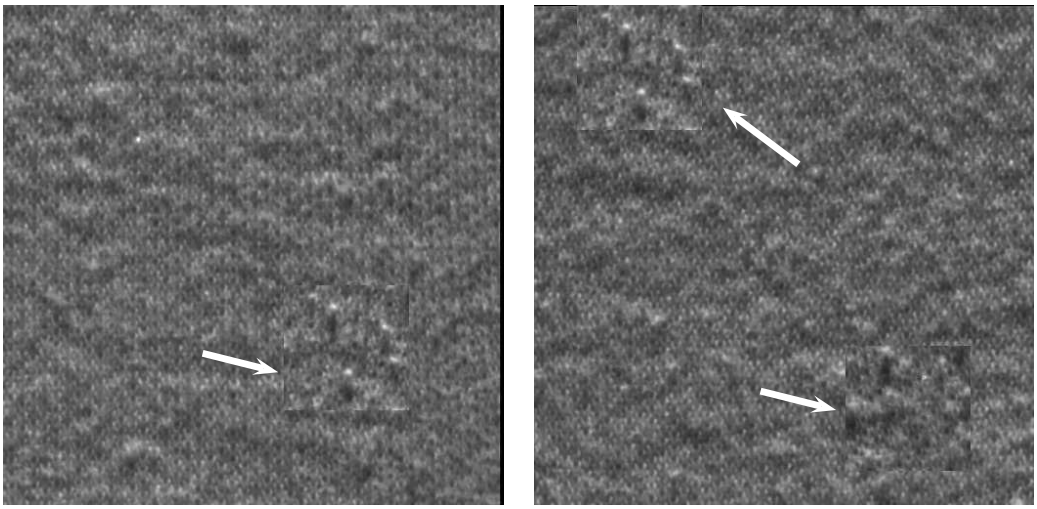

(a) (b)

Figure 4: Examples of texture mosaics of different pilled fabrics. (a) Mosaic A: a portion of Pill 3 inside Pill 1; (b) Mosaic B: a portion of Pills 2 and 3 inside Pill 1. The images display localised changes in texture. They can therefore be considered to be non-stationary, though this structure can be difficult to detect visually. On each image, the arrows show the subimages which were inserted to create the multitextured mosaics.

In both cases, application of the $\texttt{Bootstat}_{\texttt{LS2W}}$ and $\texttt{Bootstat}^{\texttt{mh}}_{\texttt{LS2W}}$ stationarity tests on the texture mosaics indicated non-stationarity with p-values less than 1%. This is to be expected since the images display different textural properties across the images, but is reassuring nevertheless since the texture boundaries are difficult to pick out visually. We also note that both tests were also able to detect non-stationarity when applied to examples of reference texture mosaics taken from several texture libraries.

5 Concluding remarks

As discussed in Section 1, recent work by Eckley et al. (2010) has indicated that, as a precursor to conducting a texture analytic task such as classification, it advisable to verify the stationarity (or otherwise) of candidate images. The work presented in this paper has addressed this issue by developing two different bootstrap-based tests of stationarity using the LS2W framework. The first test extends the work of Cardinali and Nason (2011) to two dimensions, whilst the second employs a multiple testing bootstrapping framework to assess stationarity. A benefit of these approaches is that they permit the testing of the hypothesis using a single realisation.

When analysing simulated images, the proposed approaches demonstrated good size and power performance for images of reasonable size (i.e. images of at least 128x128), highlighting the insensitivity of the test to image size and localised image variance, including slowly-varying second order structure. Of the two proposed tests, the nested bootstrapping procedure achieved better results on simulated data, but at additional computational cost. We also applied the test procedures to images encountered by an industrial collaborator. The results from these tests are consistent with our visual assessment of the images, namely that these textures are stationary, with heavy pilling introducing non-stationarity. In addition when analysing quite subtle texture mosaics, the tests were correctly able to detect non-stationarity. As such, these approaches could be used to detect, for example, regions of uneven wear within a material.

It is perhaps a little surprising that no *local Fourier* equivalent of this spatial test currently exists. Hence, we note that the development of a spatial analogue of local Fourier time series tests, such as Dwivedi and Subba Rao (2011) would be an interesting avenue for future research. The implementation of such a test could be similar to the $\texttt{Bootstat}_{\texttt{LS2W}}$ test, using bootstrapping in the Fourier domain.

Supplementary Materials

Properties of spectrum estimator: This document contains a proof that the pixel average estimator introduced in (3) is unbiased and consistent as an estimator of the

stationary spectrum. (consistency.pdf)

Size simulations: This document contains tables describing statistical size simulations in Section 4.1 for different values of the standard deviation, σ. (simulations.pdf)

Computer code: R code for performing the hypothesis tests described in the article. It contains a "readme" file describing the details of each function. (rcode.zip)

All of the supplementary material files are contained in a single archive and can be obtained via a single download. (supplementarymaterial.zip)

References

Anselin, L. (2007). *Spatial Regression Analysis in R – A Workbook*. Center for Spatially Integrated Social Science.

Bivand, R. (2012). spdep: Spatial dependence: weighting schemes, statistics and models. R package version 0.5-33.

Blanc, R., J. P. Da Costa, Y. Stitou, P. Baylou, and C. Germain (2008). Assessment of texture stationarity using the asymptotic behavior of the empirical mean and variance. *IEEE Trans. Im. Proc. 17*, 1481–1490.

Bose, S. and A. O. Steinhardt (2002). Invariant tests for spatial stationarity using covariance structure. *IEEE Trans. Sig. Proc. 44*(6), 1523–1533.

Cardinali, A. and G. P. Nason (2011). Costationarity of locally stationary time series. *J. Time Series Econ. 2*(2).

Chan, C. and G. K. H. Pang (2000). Fabric defect detection by fourier analysis. *IEEE Trans. Ind. Appl. 36*(5), 1267–1276.

Daubechies, I. (1992). *Ten Lectures On Wavelets*. Philadelphia:SIAM.

Daugman, J. G. (1990). An information-theoretic view of analog representation in striate cortex. *Computat. Neurosci.*, 403–424.

Davison, A. and D. Hinkley (1997). *Bootstrap Methods and their Application*. Cambridge University Press.

Dwivedi, Y. and S. Subba Rao (2011). A test for second-order stationarity of a time series based on the discrete Fourier transform. *J. Time Series Anal. 32*(1), 68–91.

Eckley, I. A. and G. P. Nason (2005). Efficient computation of the discrete autocorrelation wavelet inner product matrix. *Stat. Comput. 15*(2), 83–92.

Eckley, I. A. and G. P. Nason (2011). LS2W: Locally stationary wavelet fields in R. *J. Stat. Soft. 43*(3).

Eckley, I. A., G. P. Nason, and R. L. Treloar (2010). Locally stationary wavelet fields with application to the modelling and analysis of image texture. *J. Roy. Stat. Soc. C 59*(4), 595–616.

Ephraty, A., J. Tabrikian, and H. Messer (1996). A test for spatial stationarity and applications. In *8th IEEE Sig. Proc. Work. Stat. Sig. Array Proc.*, pp. 412–415. IEEE Computer Society.

Ephraty, A., J. Tabrikian, and H. Messer (2001). Underwater source detection using a spatial stationarity test. *J. Acoust. Soc. Amer. 109*, 1053–1063.

Field, D. J. (1999). Wavelets, vision and the statistics of natural scenes. *Phil. Trans. Roy. Soc. London A 357*(1760), 2527–2542.

Fuentes, M. (2005). A formal test for nonstationarity of spatial stochastic processes. *J. Multivar. Stat. 96*, 30–54.

Gonzalez, R. C. and R. E. Woods (2001). *Digital Image Processing* (2nd ed.). Prentice Hall.

Laine, A. and J. Fan (1993). Texture classification by wavelet packet signatures. *IEEE Trans. Patt. Anal. Mach. Intell. 15*(11), 1186–1191.

Matérn, B. (1960). Spatial variation. Stochastic models and their application to some problems in forest surveys and other sampling investigations. *Meddelanden fran statens Skogsforskningsinstitut 49*(5), 144.

Nason, G. P. (2008). *Wavelet methods in statistics with R*. Springer Verlag.

Nason, G. P. (2013). A test of second-order stationarity and approximate confidence intervals for localized autocovariances for locally stationary time series. *J. Roy. Stat. Soc. B 75*. To appear.

Nason, G. P., R. Von Sachs, and G. Kroisandt (2000). Wavelet processes and adaptive estimation of the evolutionary wavelet spectrum. *J. Roy. Stat. Soc. B 62*(2), 271–292.

Palmer, S., J. Zhang, and X. Wang (2011). New methods for objective evaluation of fabric pilling by frequency domain image processing. *Res. J. Text. Appar. 13*(1), 11–23.

Percival, D. B. and A. T. Walden (2006). *Wavelet methods for time series analysis* (4th ed.). Cambridge University Press.

Petrou, M. and P. G. Sevilla (2006). *Image processing: dealing with texture*. John Wiley & Sons Inc.

Priestley, M. B. and T. S. Rao (1969). A test for non-stationarity of time-series. *J. Roy. Stat. Soc. B 31*(1), 140–149.

Schlather, M. (2012). Randomfields: Simulation and Analysis of Random Fields. R package version 2.0.59.

Stein, M. L. (1999). *Interpolation of spatial data: some theory for kriging*. Springer Verlag.

Unser, M. (1995). Texture classification and segmentation using wavelet frames. *IEEE Trans. Im. Proc. 4*(11), 1549–1560.

Vidakovic, B. (1999). *Statistical Modelling by Wavelets*. Wiley: New York.