

A Power Variance Test for Nonstationarity in Complex-Valued Signals

Thomas E. Bartlett
and Adam M. Sykulski
and Sofia C. Olhede
Department of Statistical Science
University College London, UK
Email: thomas.bartlett@ucl.ac.uk

Jonathan M. Lilly
and Jeffrey J. Early
NorthWest Research Associates
Seattle, USA

Abstract—We propose a novel algorithm for testing the hypothesis of nonstationarity in complex-valued signals. The implementation uses both the bootstrap and the Fast Fourier Transform such that the algorithm can be efficiently implemented in $\mathcal{O}(N \log N)$ time, where N is the length of the observed signal. The test procedure examines the second-order structure and contrasts the observed power variance—i.e. the variability of the instantaneous variance over time—with the expected characteristics of stationary signals generated via the bootstrap method. Our algorithmic procedure is capable of learning different types of nonstationarity, such as jumps or strong sinusoidal components. We illustrate the utility of our test and algorithm through application to turbulent flow data from fluid dynamics.

© 2015 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

I. INTRODUCTION

In this paper we introduce a new automated method of detecting nonstationarity in observed time signals. Identifying and dealing with nonstationarity is a problem of fundamental importance in machine learning [1]–[3]. Nonstationarity can be hard to detect, because while there is only one way for a signal to be stationary, there are many ways for this assumption to be violated [4]. Our approach is based on examining the variance of the observed signal. As we assume the signal is zero mean, or that the mean is removed, this is the simplest way that the signal may depart from stationarity.

We focus our attention on bivariate signals, that are an important class of observations. Such data are commonly represented as complex-valued signals, as discussed in [5], and as performed in numerous applications including fMRI [6], blood flow [7], neural networks [8], oceanography [9], seismology [10], and meteorology [11]. The complex-valued representation is particularly useful for modelling trajectories of fluid particles in oceanography and related applications [12], also referred to by the name “Lagrangian data.” In Fig. 1 we display three trajectories of particles propagated in a simulation of forced-dissipative two-dimensional fluid turbulence. These trajectories exhibit various degrees and types of nonstationarity, and we will test our procedures on this data in this paper.

To develop our test procedure we draw inspiration from existing methods to detect nonstationarity in time signals. Here tests have been developed for real-valued data, focussing on different structural behaviour of the observed signal. Standard procedures are formulated in the frequency domain [13]–[16] or the wavelet domain [17], [18]. We propose to develop theory from the frequency domain, however, to capture nonlinear characteristics of the observed time signal, we develop a test in the time domain. Specifically our test procedure examines the power variance of the observed signal, i.e. the variability of the instantaneous variance over time.

This is, to our knowledge, the first test for nonstationarity designed specifically for complex-valued signals. We will demonstrate that testing the power variance can both identify *heteroscedastic* signals, where the variance of the signal changes over time, such as will be seen shortly in Fig. 1(b), and signals with strong sinusoidal components that appear ‘phase-locked’ as in Fig. 1(c). The latter is a strongly nonlinear feature and is associated with relatively low power variance, as we shall demonstrate. We will also derive the expectation of the test statistic under the null hypothesis, and show how the test can be efficiently performed in $\mathcal{O}(N \log N)$ time, where N is the length of the observed signal, using the bootstrap and Fast Fourier Transforms.

II. SECOND-ORDER PROPERTIES AND THE POWER VARIANCE

We first formulate theory for continuous signals, before detailing how approximations are made with discrete observations. Consider a continuous bivariate signal $\{x_t, y_t\}$ represented as a complex-valued quantity $z_t = x_t + iy_t$ at each time point t , where $i \equiv \sqrt{-1}$. If the complex-valued signal $\{z_t\}$ is second-order stationary then $\{z_t\}$ has a finite constant mean, and an autocovariance sequence $E(z_t z_{t-\tau}^*)$ that is finite and does not depend on t , where z_t^* denotes the complex conjugate of z_t . To fully describe the properties of the complex-valued signal, we also need to describe its relation sequence [5], given by $E(z_t z_{t-\tau})$, which is also independent of time for a second-order stationary signal.

A second-order nonstationary signal does not satisfy one or more of the constraints of $E(z_t)$, $E(z_t z_{t-\tau}^*)$ and $E(z_t z_{t-\tau})$ being independent of t . There are a multitude of ways that this can happen. We focus on the simple case where $\text{Var}(z_t) \equiv E(|z_t|^2)$ varies with time, such that the signal is




Fig. 1: The trajectories of three different particles propagated in a large two-dimensional fluid turbulence simulation.

heteroscedastic. One simple class of such models is that of *uniformly modulated processes* [4], where $z_t = \sigma_t u_t$ where u_t is a stationary process, and σ_t is a deterministic signal. Hence, a measure of nonstationarity during the time interval $(0, T)$ is the time mean of the squared deviation of the variance from its time-mean value, which we term the *power variance*, given by

$$\Omega(T) = \frac{1}{T} \int_0^T (\text{Var}\{z_t\} - \bar{\sigma}^2)^2 dt, \quad (1)$$

where the *sample variance*, or *average power*, is given by

$$\bar{\sigma}^2 = \frac{1}{T} \int_0^T \text{Var}\{z_t\} dt. \quad (2)$$

It is clear that for a stationary signal, the variance of z_t is constant over time such that $\Omega(T) \equiv 0$. However for nonstationary signals it would commonly be the case that $\Omega(T) \neq 0$, depending on the form of nonstationarity.

If we now consider discrete observations of $\{z_t\}$, then we do not observe $\text{Var}\{z_t\}$ directly, so instead we calculate a test statistic. For a regularly sampled complex-valued signal, $\mathbf{z} = (z_0, z_1, z_2, \dots, z_{N-1})$, we define the *observed power variance* as

$$\hat{\Omega}(\mathbf{z}) = \frac{1}{N} \sum_{n=0}^{N-1} (|z_n|^2 - \hat{\sigma}^2)^2, \quad (3)$$

where

$$\hat{\sigma}^2 = \frac{1}{N} \sum_{n=0}^{N-1} |z_n|^2. \quad (4)$$

For a discretely observed stationary signal, it is *not* typically the case that $\hat{\Omega}(\mathbf{z}) = 0$, in contrast to $\Omega(T)$ in (1). This is because of the natural stochasticity in z_t , which yields $|z_n|^2 \neq \hat{\sigma}^2$ at each observed time point. It is therefore possible for a nonstationary signal to have a smaller observed power variance than a stationary signal—a feature which we account for in our algorithm by performing two-sided statistical tests.

III. BOOTSTRAP TEST PROCEDURE

To test whether an observed signal is stationary or not, we first generate a number of randomised signal replicates that are constructed to be stationary, where each replicate has the same overall variability as the observed signal. We then contrast the power variance of the observed signal and that of the replicates. If the power variance of the observed signal is

significantly different from the typical power variance found in the replicates, then we can reject the null hypothesis of stationarity, and identify the signal as nonstationary.

We generate the signal replicates using the bootstrap method [19]. We do this in the frequency domain by following the exact procedure of [20]. This simple procedure takes any complex-valued signal and calculates its Discrete Fourier Transform (DFT) given by

$$|Z_k| e^{i\phi_k} = \sum_{n=0}^{N-1} z_n e^{-2\pi i k n / N}. \quad (5)$$

The Discrete Fourier Transform is complex-valued and we have decomposed this in terms of an amplitude $|Z_k|$ and corresponding phase ϕ_k at each frequency k . The idea of [20] is to fix the amplitudes $|Z_k|$ but uniformly generate randomised phases, $\tilde{\phi}_k \sim \mathcal{U}(-\pi, \pi)$, and then perform the inverse DFT into the time domain to generate a randomised replicate of the signal $\tilde{\mathbf{z}} = (\tilde{z}_0, \tilde{z}_1, \tilde{z}_2, \dots, \tilde{z}_{N-1})$ given by

$$\tilde{z}_n = \frac{1}{N} \sum_{k=0}^{N-1} |Z_k| e^{i\tilde{\phi}_k} e^{2\pi i k n / N}. \quad (6)$$

By definition, the frequency components of a stationary signal have uncorrelated phases, and hence the replicated signal $\tilde{\mathbf{z}}$ (which has completely random phases) is stationary, even if the observed signal \mathbf{z} is not. See also [20] in relation to this.

Our test procedure generates many such replicated signals and compares the power variance with that of the observed signal, to see if there is a significant difference. We detail the exact test procedure in Algorithm 1, where FFT denotes the Fast Fourier Transform and $\mathbb{I}(\cdot)$ denotes the indicator function. If the null is rejected then there is sufficient evidence to suggest that the observed signal is nonstationary, otherwise the null should not be rejected using this test procedure. Note that this is a two-sided test procedure—we have tested whether the observed power variance is significantly lower or higher than that found in the distribution of stationary replicates. One-sided tests can be performed by using the values of $q(\mathbf{z})$ or $r(\mathbf{z})$ in Algorithm 1, instead of $p(\mathbf{z})$, to respectively test for significantly high or low power variance.

By using the FFT to carry out the Fourier transforms, we can achieve $\mathcal{O}(N \log N)$ computational cost, for each bootstrap iteration. For example, performing a test on one complex-valued signal with $N = 1000$ and $B = 1000$ on a 2.6

Algorithm 1 Bootstrap Power Variance Test

INPUTS: $\mathbf{z} = (z_0, \dots, z_{N-1})$, B

$$\widehat{\Omega}(\mathbf{z}) \leftarrow \frac{1}{N} \sum_{n=0}^{N-1} \left(|z_n|^2 - \frac{1}{N} \sum_{n=0}^{N-1} |z_n|^2 \right)^2$$

$$\{Z_k\} \leftarrow \text{FFT}(\mathbf{z})$$

for $b = 1$ to B **do**

$$\tilde{\phi}_k \stackrel{iid}{\sim} \mathcal{U}(-\pi, \pi), \text{ for } k = 0, \dots, N-1$$

$$\tilde{\mathbf{z}}^{(b)} \leftarrow \frac{1}{N} \text{FFT} \left(\left\{ |Z_k| e^{i\tilde{\phi}_k} \right\} \right)$$

$$\widehat{\Omega}(\tilde{\mathbf{z}}^{(b)}) \leftarrow \frac{1}{N} \sum_{n=0}^{N-1} \left(\left| \tilde{z}_n^{(b)} \right|^2 - \frac{1}{N} \sum_{n=0}^{N-1} \left| \tilde{z}_n^{(b)} \right|^2 \right)^2$$

end for

$$q(\mathbf{z}) \leftarrow \frac{1}{B} \sum_{b=1}^B \mathbb{I}[\widehat{\Omega}(\tilde{\mathbf{z}}^{(b)}) > \widehat{\Omega}(\mathbf{z})]$$

$$r(\mathbf{z}) \leftarrow \frac{1}{B} \sum_{b=1}^B \mathbb{I}[\widehat{\Omega}(\tilde{\mathbf{z}}^{(b)}) < \widehat{\Omega}(\mathbf{z})]$$

$$p(\mathbf{z}) \leftarrow 2 \min \{q(\mathbf{z}), r(\mathbf{z})\}$$

if $p(\mathbf{z}) < 0.05$ **then**

Reject null

end if

GHz 2012 Macbook Pro takes 0.072s. As the bootstraps are generated independently, the method is straightforward to run in parallel with larger datasets. We note that bootstrapping, phase scrambling, and similar procedures for generating time series replicates are discussed further in [21]–[24].

IV. ILLUSTRATIVE EXAMPLES

By way of illustrative examples, we present three canonical cases for investigating our test procedure. All code in this paper is generated in Matlab[®] and is available online from <http://www.ucl.ac.uk/statistics/research/spg/software>, and all results are exactly reproducible. Also free to download is an algorithm for applying the test procedure to any observed complex-valued signal.

The first example we present is an autoregressive AR(1) model, where the sequence $\{z_n\}$ is defined by

$$x_n = 0.9x_{n-1} + 0.1\sigma_n^{(x)}, \text{ where } \sigma_n^{(x)} \stackrel{iid}{\sim} \mathcal{N}(0, 1),$$

$$y_n = 0.9y_{n-1} + 0.1\sigma_n^{(y)}, \text{ where } \sigma_n^{(y)} \stackrel{iid}{\sim} \mathcal{N}(0, 1),$$

$$z_n = \frac{x_n + iy_n}{\sqrt{2}}.$$

This model generates stationary signals. We now check if our test procedure correctly does not reject the null for realisations from this model. We set $N = 1000$ and $B = 1000$ in Algorithm 1 and compute $p(\mathbf{z})$ for a signal generated from this model. We perform a Monte Carlo simulation and repeat this procedure 10,000




Fig. 2: Distribution of p -values calculated from $p(\mathbf{z})$ in Algorithm 1 for AR(1) process.

times with independently generated signals from the AR(1) model. We then get a distribution of p -values which are shown in Fig. 2, with the 5% significance level indicated with a vertical red line. It is clear that there is no tendency to reject the null for this process, with 5.21% of null hypotheses rejected, which is consistent with a type I error level of 5%.

The second canonical example is a basic jump process, with $\{z_n\}$ defined by

$$x_n = \sigma_n^{(x)}, \text{ where } \sigma_n^{(x)} \stackrel{iid}{\sim} \mathcal{N}(0, 1),$$

$$y_n = \sigma_n^{(y)}, \text{ where } \sigma_n^{(y)} \stackrel{iid}{\sim} \mathcal{N}(0, 1),$$

$$z_n = \begin{cases} 1 + (x_n + iy_n)/\sqrt{2}, & \text{if } n \leq N/2, \\ 3 + (x_n + iy_n)/\sqrt{2}, & \text{if } n > N/2. \end{cases}$$

This model generates nonstationary signals with a significant change-point midway through the series. We check if our test procedure correctly rejects the null for realisations from this model. Again setting $N = 1000$ and $B = 1000$, we generate 10,000 independent instantiations of such a jump process, and compute p -values using Algorithm 1. This time we perform a one-sided test and report values for $q(\mathbf{z})$ in Algorithm 1 to check for high power variance. The one-tailed test p -values are shown in Figure 3, with the 5% significance level similarly indicated with a vertical red line. In this Monte Carlo study, there is a tendency to reject the null for this process, with an alternative of increased power variance: 71.8% of null hypotheses are rejected at the 5% significance level. This occurs because the process is nonstationary and has a significant change-point, which stationary replicates will not have. This yields a significantly higher observed power variance in most cases. The null is not always rejected however as the signal-to-noise ratio is set reasonably low.




Fig. 3: Jump process with p -values calculated from $q(\mathbf{z})$ in Algorithm 1.

The final illustrative example is a cyclo-stationary process, with $\{z_n\}$ defined by

$$z_n = ae^{i\omega n/N} + \frac{\sigma_n^{(x)} + i\sigma_n^{(y)}}{\sqrt{2}}, \text{ where } \sigma_n^{(x)}, \sigma_n^{(y)} \stackrel{iid}{\sim} \mathcal{N}(0, 1).$$

We set $\omega = 10$ and $a = 1$ such that the signal-to-noise ratio is relatively low. This model generates nonstationary signals because the sinusoid is ‘phase-locked’ over time. For our simulations we again set $N = 1000$ and $B = 1000$, and perform 10,000 Monte Carlo replicates. We perform a one-sided test and report values for $r(\mathbf{z})$ in Algorithm 1 to check for low power variance. The one-tailed test p -values are shown in Figure 4, with the 5% significance level again indicated with a vertical red line. It is clear that there is a tendency to reject the null for this process, with an alternative of decreased power-variance: 82.3% of null hypotheses are rejected at the 5% significance level. We observe a decreased power-variance

here because the fixed phase of the sinusoid implies that the amplitude of the signal will be less variable than the stationary replicates—which will not be ‘phase-locked’ and exhibit more variability in terms of their instantaneous variance.

The illustrative examples given above (Figures 2-4) are all based on signal length $N = 1000$. Performance decreases as signal length decreases, however the null is still rejected in many more cases in the nonstationary examples than would be expected for a stationary signal, for even the shortest signals. Table I summarises these findings.




Fig. 4: Cyclo-stationary process with p -values calculated from $r(\mathbf{z})$ in Algorithm 1.

N	AR	Jump	CS
1000	5.21%	71.8%	82.3%
500	4.81%	57.2%	60.6%
200	5.11%	39.2%	36.1%
100	5.16%	29.0%	24.5%
50	5.51%	21.5%	17.1%
20	5.94%	14.8%	11.7%
10	5.53%	11.5%	9.50%

TABLE I: Percentages of null hypotheses rejected at the 5% level for various signal lengths N , for auto-regressive (AR), jump and cyclo-stationary (CS) processes. All settings except N are identical to those used to generate Figures 2-4.

V. THEORY

The bootstrap signal replicates, $\tilde{\mathbf{z}}$, generated in Algorithm 1 are used to approximate a distribution of power variances, under the null hypothesis that the observed signal is stationary. In this section we derive the expectation of this distribution theoretically. We denote $E(\hat{\Omega}(\tilde{\mathbf{z}}))$ as this expectation and analytically derive its form in terms of the Fourier amplitudes $|Z_k|$ found in (5).

Proposition 1: For a given complex-valued signal \mathbf{z} then the power variance of bootstrap replicates $\tilde{\mathbf{z}}$ generated using (6) has expectation given by

$$E(\hat{\Omega}(\tilde{\mathbf{z}})) = \frac{1}{N^4} \left\{ \left(\sum_{k=0}^{N-1} |Z_k|^2 \right)^2 - \sum_{k=0}^{N-1} |Z_k|^4 \right\}, \quad (7)$$

conditional on the observed signal $\mathbf{z} = (z_0, z_1, z_2, \dots, z_{N-1})$, where $|Z_k|$ is found using (5).

This proposition is particularly useful when performing one-sided tests, in which one is looking for evidence that $\hat{\Omega}(\mathbf{z})$ is either unreasonably high or low. This is because the expectation can be pre-computed before implementing the bootstrap. Then if the expectation is greater than the observed power variance when testing for a high power variance—or vice-versa for a low power variance—there is no need to

perform the bootstrap as the null can already conclusively not be rejected at this stage. The proof of the proposition is as follows.

Proof: First we directly compute from (3) and (4) that the expectation for the replicates takes the form

$$\begin{aligned} E(\hat{\Omega}(\tilde{\mathbf{z}})) &= E \left[\frac{1}{N} \sum_{n=0}^{N-1} |\tilde{z}_n|^4 - \left(\frac{1}{N} \sum_{n=0}^{N-1} |\tilde{z}_n|^2 \right)^2 \right] \\ &= E(|\tilde{z}_n|^4) - E \left[\left(\frac{1}{N} \sum_{n=0}^{N-1} |\tilde{z}_n|^2 \right)^2 \right]. \quad (8) \end{aligned}$$

Examining the first term of (8) and using (6) we see that

$$\begin{aligned} E(|\tilde{z}_n|^4) &= \frac{1}{N^4} E \left(\sum_{k=0}^{N-1} \sum_{l=0}^{N-1} \sum_{p=0}^{N-1} \sum_{q=0}^{N-1} |Z_k| |Z_l| |Z_p| |Z_q| \times \right. \\ &\quad \left. e^{2\pi i(k+l-p-q)n/N} e^{i(\tilde{\phi}_k + \tilde{\phi}_l - \tilde{\phi}_p - \tilde{\phi}_q)} \right) \\ &= \frac{1}{N^4} \sum_{k=0}^{N-1} |Z_k|^4 + \frac{2}{N^4} \sum_{k=0}^{N-1} \sum_{k \neq l}^{N-1} |Z_k|^2 |Z_l|^2, \quad (9) \end{aligned}$$

in which the Fourier amplitude coefficients $|Z_k|$ are regarded as deterministic. This last expression follows as the phases are generated randomly meaning

$$E \left[e^{i(\tilde{\phi}_k + \tilde{\phi}_l - \tilde{\phi}_p - \tilde{\phi}_q)} \right] = \begin{cases} 1, & \text{if } p = k \text{ and } q = l, \\ & \text{or if } q = k \text{ and } p = l, \\ 0, & \text{otherwise,} \end{cases}$$

and thus the factor of two in the second term of (9) follows from the two possibilities by which the indices in the summation can pair together ($p = k$ and $q = l$, or $q = k$ and $p = l$). The first term at the right-hand side of (9) arises from the special case $q = p = l = k$. Equation (9) can then be rewritten as

$$E(|\tilde{z}_n|^4) = \frac{2}{N^4} \left(\sum_{k=0}^{N-1} |Z_k|^2 \right)^2 - \frac{1}{N^4} \sum_{k=0}^{N-1} |Z_k|^4. \quad (10)$$

The second term of (8) can be immediately found in terms of $|Z_k|$ using the Parseval-Rayleigh relationship

$$E \left[\left(\frac{1}{N} \sum_{n=0}^{N-1} |\tilde{z}_n|^2 \right)^2 \right] = \left(\frac{1}{N^2} \sum_{k=0}^{N-1} |Z_k|^2 \right)^2. \quad (11)$$

The proposition then follows by substituting in (10) and (11) into (8). \blacksquare

The full analytic distribution of $\hat{\Omega}(\tilde{\mathbf{z}})$ cannot be easily found for a given signal \mathbf{z} . Even if the observed signal is Gaussian, the distribution of power variances will in general not be Gaussian for finite N . For this reason we do not proceed with an analytic approach of performing the power variance test, and instead the bootstrap method of Section III is preferred.




Fig. 5: The real part of the velocity signal from the trajectories plotted in Fig. 1. The time length of the signal is 1000 days. The imaginary part of the velocity signal is similar in shape and magnitude, but the rapid oscillations seen in the real part typically coincide with oscillations in the imaginary part that are ninety degrees out of phase, leading to looping motions on the complex plane, which can be seen in Fig. 1.

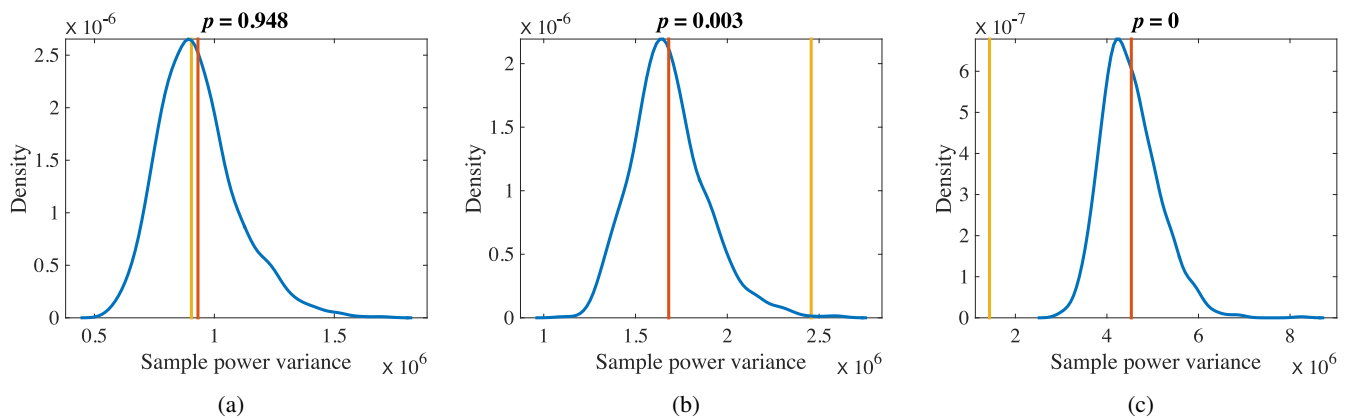


Fig. 6: Results from bootstrap power variance tests for the trajectories plotted in Fig. 1. Each plot shows the observed power variance (\bullet), as well as the distribution of power variances from the bootstrap replicates (---). The theoretical expectation of this distribution is also given (---). The p -value of the observed power variance is reported above the figure, and values less than 0.05 lead to rejection of the null hypothesis that the signal is stationary.

VI. PARTICLE TRAJECTORIES IN TWO-DIMENSIONAL TURBULENCE

Quasigeostrophic (QG) turbulence [25] is a model for large-scale fluid motions in the ocean and atmosphere that is predominantly two-dimensional and is dominated by a balance between the pressure force and the apparent Coriolis force. Qualitatively, the fluid can be divided into two regimes—regions of swirling masses of fluid known as eddies, and regions in between the eddies. The positions of particles transported by the flow, referred to as drifters, are tracked over time, leading to trajectory curves. In this simulation there are some drifter trajectories which experience nonstationarity due to crossing eddy boundaries, thereby transitioning between eddy interiors and the ambient fluid, and others which do not. This section uses our test procedure to classify the trajectories into different types.

As in Section IV, all code and data is available online at <http://www.ucl.ac.uk/statistics/research/spg/software> and all results are exactly reproducible. Furthermore, the QG turbulence simulation can also be reproduced using the software available at jeffreyearly.com/numerical-models/.

Fig. 1 shows three trajectories from the simulation. In Fig. 1(a), we see an example of a stationary trajectory, while Fig. 1(b) shows an example of a nonstationary trajectory comprised of portions with multiple distinct flow characteristics, and Fig. 1(c) is another nonstationary trajectory dominated by a single eddy. Fig. 5 shows the same examples, with the real part of the drifter *velocities* plotted, and we can again observe the same characteristics of stationary and nonstationary behaviour.

Fig. 6 shows the results of applying the bootstrap power variance test of Algorithm 1 (with $B = 1000$) to the complex-valued velocity signals. Included in the figure are the corresponding variances of the null distributions, along with the observed power variances calculated according to (3), and the means of the null power-variance distributions calculated from (7). In Fig. 6(a), the observed power variance is in the middle of the null distribution (and hence the null is not rejected), because this signal appears reasonably stationary: the p -value here is calculated as $p(\mathbf{z})$ in Algorithm 1. In Fig. 6(b), the observed power variance is in the extreme upper tail of the null distribution (leading to rejection of the null), because this signal is clearly nonstationary: the p -value here is calculated as $q(\mathbf{z})$ in Algorithm 1. Finally in Fig. 6(c),

the observed power variance is in the extreme lower tail of the null distribution (and again the null is rejected). This is because this signal is nonstationary in a different way: the sinusoidal oscillations are ‘phase-locked’ in such a way that the instantaneous variance of the signal has *less* variability than stochastic stationary replicates. The p -value here is calculated as $r(\mathbf{z})$ in Algorithm 1.

VII. CONCLUSIONS

In this paper, we have proposed an algorithm for testing for nonstationarity in complex-valued signals. The algorithm is based on calculating the power variance, a measure of the instantaneous variance of power of the observed signal. The algorithm can learn two different types of nonstationarity, this leading to a two-sided test. Both tails are indeed informative: if the power variance is too large, then this indicates that the variance of the signal is too variable to be classified as stationary. On the other hand, if the power variance is too small, then this indicates that the variance of the signal is *less* than that expected under the hypothesis of stationarity.

To perform the test, the algorithm employs the bootstrap method. The bootstrapped signals have the same overall variance, but are stationary by construction and therefore have constant variability. By comparing the original signal with these replicates we can determine if the original signal has the characteristics of a stationary signal. We do this by comparing the power variance. Significantly high power variance indicates a heteroscedastic process which is nonstationary. Significantly low power variance indicates a very small degree of variability; this is indicative of a strong ‘phase-locking’ whereby various Fourier components interact to generate a relatively uniform level of signal power.

Developing fully assumption-free methods of detecting signal nonstationarity is of great current interest. Our focus in this paper has been on coupled pairs of signals; in general we would study signals of arbitrary dimensionality. In such settings it is significantly harder to determine natural test statistics, and extending our methods to such settings remains an outstanding challenge. Finally, while the focus of this paper has been on naturally complex-valued signals, the same basic approach should apply to real-valued signals that have been made complex-valued though taking the analytic part [26], which is a promising extension for future work.

ACKNOWLEDGMENTS

TEB, AMS and SCO acknowledge funding from EPSRC grant no. EP/L025744/1. SCO also acknowledges funding from EPSRC grant no. EP/I005250/1. The research of AMS was supported by a Marie Curie International Outgoing Fellowship within the seventh European Community framework programme. The work of JML and JJE was supported by award #1235310 from the Physical Oceanography Program of the United States National Science Foundation.

REFERENCES

[1] M. Sugiyama and M. Kawanabe, *Machine learning in non-stationary environments*. MIT Press, Cambridge, MA, 2012.
 [2] J. W. Robinson and A. J. Hartemink, “Learning non-stationary dynamic Bayesian networks,” *The Journal of Machine Learning Research*, vol. 11, pp. 3647–3680, 2010.

[3] F. Michaud and M. J. Matarić, “Learning from history for behavior-based mobile robots in non-stationary conditions,” *Machine Learning*, vol. 31, no. 1-3, pp. 141–167, 1998.
 [4] M. B. Priestley, *Non-linear and non-stationary time series analysis*. London: Academic Press, 1988.
 [5] P. J. Schreier and L. L. Scharf, *Statistical signal processing of complex-valued data: the theory of improper and noncircular signals*. Cambridge University Press, 2010.
 [6] D. B. Rowe, “Modeling both the magnitude and phase of complex-valued fMRI data,” *Neuroimage*, vol. 25, no. 4, pp. 1310–1324, 2005.
 [7] P. J. Brands, A. P. G. Hoeks, L. A. F. Ledoux, and R. S. Reneman, “A radio frequency domain complex cross-correlation model to estimate blood flow velocity and tissue motion by means of ultrasound,” *Ultrasound in Medicine & Biology*, vol. 23, no. 6, pp. 911–920, 1997.
 [8] D. P. Mandic and V. S. L. Goh, *Complex valued nonlinear adaptive filters: noncircularity, widely linear and neural models*. John Wiley & Sons, 2009, vol. 59.
 [9] A. M. Sykulski, S. C. Olhede, J. M. Lilly, and E. Danioux, “Lagrangian time series models for ocean surface drifter trajectories,” *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, p. doi: 10.1111/rssc.12112, 2015.
 [10] J. C. Samson, “Matrix and Stokes vector representations of detectors for polarized waveforms: theory, with some applications to teleseismic waves,” *Geophysical Journal International*, vol. 51, no. 3, pp. 583–603, 1977.
 [11] J. Gonella, “A rotary-component method for analysing meteorological and oceanographic vector time series,” in *Deep Sea Research and Oceanographic Abstracts*, vol. 19, no. 12. Elsevier, 1972, pp. 833–846.
 [12] J. M. Lilly, R. K. Scott, and S. C. Olhede, “Extracting waves and vortices from Lagrangian trajectories,” *Geophysical Research Letters*, vol. 38, p. L23605, 2011.
 [13] M. B. Priestley and T. S. Rao, “A test for non-stationarity of time-series,” *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, pp. 140–149, 1969.
 [14] T. S. Rao and H. Tong, “On some tests for the time dependence of a transfer function,” *Biometrika*, vol. 60, no. 3, pp. 589–597, 1973.
 [15] K. Sakiyama and M. Taniguchi, “Testing composite hypotheses for locally stationary processes,” *Journal of Time Series Analysis*, vol. 24, no. 4, pp. 483–504, 2003.
 [16] Y. H. Tsao, “Tests for nonstationarity,” *The Journal of the Acoustical Society of America*, vol. 75, no. 2, pp. 486–498, 1984.
 [17] G. Nason, “A test for second-order stationarity and approximate confidence intervals for localized autocovariances for locally stationary time series,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 75, no. 5, pp. 879–904, 2013.
 [18] M. H. Neumann and R. von Sachs, “A wavelet-based test for stationarity,” *Journal of Time Series Analysis*, vol. 21, no. 5, pp. 597–613, 2000.
 [19] B. Efron, “Bootstrap methods: another look at the jackknife,” *The Annals of Statistics*, pp. 1–26, 1979.
 [20] J. Theiler, S. Eubank, A. Longtin, B. Galdrikian, and J. D. Farmer, “Testing for nonlinearity in time series: the method of surrogate data,” *Physica D: Nonlinear Phenomena*, vol. 58, no. 1, pp. 77–94, 1992.
 [21] M. Shinozuka and C. M. Jan, “Digital simulation of random processes and its applications,” *Journal of Sound and Vibration*, vol. 25, no. 1, pp. 111–128, 1972.
 [22] M. Shinozuka and G. Deodatis, “Simulation of stochastic processes by spectral representation,” *Applied Mechanics Reviews*, vol. 44, no. 4, pp. 191–204, 1991.
 [23] D. B. Percival, “Simulating Gaussian random processes with specified spectra,” *Computing Science and Statistics*, pp. 534–538, 1993.
 [24] C. Kirch and D. N. Politis, “TFT-bootstrap: Resampling time series in the frequency domain to obtain replicates in the time domain,” *The Annals of Statistics*, pp. 1427–1470, 2011.
 [25] J. G. Charney, “Geostrophic turbulence,” *Journal of the Atmospheric Sciences*, vol. 28, no. 6, pp. 1087–1095, 1971.
 [26] S. L. Marple Jr, “Computing the discrete-time analytic signal via FFT,” *Signal Processing, IEEE Transactions on*, vol. 47, no. 9, pp. 2600–2603, 1999.