Nonlin. Processes Geophys. Discuss., 2, 1227–1273, 2015
www.nonlin-processes-geophys-discuss.net/2/1227/2015/
doi:10.5194/npgd-2-1227-2015

Nonlinear Processes in Geophysics
Discussions
Open Access

Cumulative areawise testing in wavelet analysis and its application to geophysical time series

J. A. Schulte

Department of Meteorology, The Pennsylvania State University, University Park, Pennsylvania, USA

Received: 14 May 2015 – Accepted: 12 July 2015 – Published: 29 July 2015

Correspondence to: J. A. Schulte (jas6367@psu.edu)

Published by Copernicus Publications on behalf of the European Geosciences Union & the American Geophysical Union.

Abstract

Statistical significance testing in wavelet analysis was improved through the development of a cumulative areawise test. The test was developed to eliminate the selection of two significance levels that an existing geometric test requires for implementation. The selection of two significance levels was found to make the test sensitive to the chosen pointwise significance level, which may preclude further scientific investigation. A set of experiments determined that the cumulative areawise test has greater statistical power than the geometric test in most cases, especially when the signal-to-noise ratio is high. The number of false positives identified by the tests was found to be similar if the respective significance levels were set to 0.05. The new testing procedure was applied to the time series of the Atlantic Multi-decadal Oscillation (AMO), North Atlantic Oscillation (NAO), Pacific Decadal Oscillation (PDO), and Niño 3.4 index. The testing procedure determined that the NAO, PDO, and AMO are consistent with red-noise processes, whereas significant power was found in the 2–7 year period band for the Niño 3.4 index.

1 Introduction

In many research fields, it is of interest to understand the behavior of time series in order to achieve a deeper understanding of physical mechanisms or relationships. Such a task can be formidable given that time series are composed of oscillations, nonstationarities, and noise. Fortunately, many tools have been developed to extract information from time series, including singular spectrum analysis (Vautard et al., 1992), Fourier analysis (Jenkins and Watts, 1968), and wavelet analysis (Meyers et al., 1993; Torrence and Compo, 1998). The goal of each of these tools is to assess whether deterministic features are embedded in a time series. Fourier analysis, as an example, is a method by which a time series is decomposed into frequency components so that embedded oscillations can be detected. However, the underlying assumption in Fourier

NPGD
2, 1227–1273, 2015

analysis is that time series are stationary. The limitation can be circumvented by using a windowed Fourier analysis but with the caveat that the window width is fixed, which can lead to poor resolution at low-frequencies (Lau and Weng, 1995). Wavelet analysis, relaxing the assumption of stationarity, offers an alternative method to Fourier analysis in which the window width is no longer fixed, minimizing aliasing (Meyers et al., 1993; Torrence and Compo, 1998). Wavelet analysis has been demonstrated to be useful in the understanding of the North Atlantic Oscillation (NAO; Higuchi et al., 2003; Olsen et al., 2012), applications to oceanographic problems (Meyers et al., 1993; Lee and Lwiza, 2008; Whitney, 2010; Wilson et al., 2014), assessments of historical hydroclimate variability (Labat, 2004, 2008), and many other geophysical applications (Grinsted et al., 2004; Velasco and Mendoza, 2008).

When using any time series extraction procedure it is important to assess the statistical significance of the computed test statistic against some null hypothesis. In geophysical applications, for example, red noise is typically chosen as the null hypothesis. Torrence and Compo (1998) were the first to apply wavelet analysis in a statistical framework using pointwise significance testing, allowing deterministic features to be distinguished from stochastic features. In a pointwise significance test, one tests each estimated wavelet power coefficient against a stationary theoretical red-noise background spectrum. Despite the insights gained from the statistical procedure, it has many deficiencies, as noted by Maraun and Kurths (2004), who showed that it can lead to many spurious results simply due to multiple testing. Addressing the multiple-testing problem, Maraun et al. (2007) developed an areawise test that assesses the significance of so-called pointwise significance patches, contiguous regions of pointwise significance in a wavelet power spectrum. The areawise test, though dramatically reducing the number of spurious results, is computationally inefficient, involving a root-finding algorithm to estimate a critical area of the reproducing kernel corresponding to the desired significance level of the test. Furthermore, the critical area needs to be computed for different analyzing wavelets and for their associated parameters, such

NPGD
2, 1227–1273, 2015

as the central frequency for the Morlet wavelet and the order in the case of the Paul wavelet.

A simpler procedure for addressing multiple testing problems is the geometric test developed by Schulte et al. (2015). Like the areawise test, the test statistic for the procedure is based on patch area, or more specifically, the normalized area of the patch, which allows patches at different periods to be compared simultaneously. The calculation of the critical level for the geometric test is much simpler than that for the areawise test, involving the computation of the normalized area for a large ensemble of patches under a null hypothesis that results in a null distribution from which the desired critical level can be calculated.

Both the geometric and areawise tests, however, suffer from a binary decision: one must choose both a pointwise and areawise or geometric significance level. The problem with such a statistical construction is that the outcomes of the testing procedure may depend on the chosen pointwise significance level. For an ideal test, there is a single significance level that is chosen and the results of the testing procedure depend only on that significance level so that a test statistic, for example, that is 1 % significant is guaranteed to be 5 % significant. In the present case, however, there is no such guarantee: a 1 % geometrically significant patch at one pointwise significance level may not be 5 % significant at another pointwise significance level. In such cases, the statistical significance of patches is ambiguous and may preclude further scientific investigation. This sensitivity problem underscores the need to develop a computationally efficient testing procedure free of binary decisions. The approach taken here will consider the areas of patches over all pointwise significance levels, and hence the method is called the *cumulative* areawise test. This test has the important feature that the significance of the wavelet power coefficients is a monotonic increasing function of the pointwise significance level. In other words, a wavelet power coefficient, for example, that is 1 % significant under the new procedure will be guaranteed to be 5 % significant, a consistent statistical construction.

NPGD
2, 1227–1273, 2015

The paper is organized as follows: in Sect. 2, a brief description of wavelet analysis is provided together with a discussion of existing statistical testing procedures, including the sensitivity of the geometric test to the chosen pointwise significance, motivating the construction of the cumulative areawise test. Before proceeding to the development of the new testing procedure, the topological properties of red noise are analyzed in Sect. 3. In Sect. 4, the cumulative areawise test is developed and is followed by a comparison of the test in terms of statistical power to the existing geometric test. Applications of the test to prominent climate indices are presented in Sect. 5, followed by concluding remarks in Sect. 6.

2 Existing wavelet analysis significance tests

2.1 Wavelet analysis

The wavelet transform of a time series is defined as the convolution of the time series with a wavelet function ψ_0. The wavelet transform of a time series $x_n (n = 1, \ldots, N)$ with a wavelet function ψ_0 is given by

$$W_n(s) = \sqrt{\frac{\delta t}{s}} \sum_{n'=1}^{N} x_{n'} \psi_0 \left[(n' - n) \frac{\delta t}{s} \right] \quad (1)$$

where s is the wavelet scale, δt is a time step determined by the data, and N is the length of the time series. There are many kinds of wavelets, but perhaps the most common are the Morlet, Paul, and Dog wavelets. For geophysical applications, the Morlet wavelet is often used and is given by

$$\psi_0(\eta) = \pi^{-1/4} e^{i\omega_0 \eta} e^{-\frac{1}{2}\eta^2} \quad (2)$$

where ω_0 is the dimensionless frequency, $\eta = s \cdot t$, t is time, and the wavelet scale is related to the Fourier period by $\lambda = 1.03s$ if $\omega_0 = 6$. This particular wavelet balances

NPGD
2, 1227–1273, 2015

both frequency and time-localizations. Thoughout the paper, $\omega_0 = 6$. The Paul wavelet, which is also a complex wavelet, is more localized in time, less localized in frequency space, and is given by

$$\psi_0(\eta) = \frac{(2i)^m m!}{\sqrt{\pi(2m)!}}(1 - i\eta)^{-(m-1)}, \tag{3}$$

where m is the order of the Paul wavelet, which controls the localization properties of the analyzing wavelet. In this case, the Fourier period is related to the wavelet scale by the following equation:

$$\lambda = \frac{4\pi s}{2m+1}. \tag{4}$$

If even more time-localization is desired, one can use the Dog wavelet, a real wavelet given by

$$\psi_0(\eta) = \frac{(-1)^{m+1}}{\sqrt{\Gamma\left(m+\frac{1}{2}\right)}}\frac{d^m}{d\eta^m}e^{-\eta^2/2}, \tag{5}$$

where m represents the order of the derivative and Γ is the gamma function. For the Dog wavelet, the Fourier period is related to the wavelet scale by the equation

$$\lambda = \frac{2\pi s}{\sqrt{m+\frac{1}{2}}}. \tag{6}$$

In this paper, the Paul wavelet is used with $m = 4$ and the Dog wavelet is used with $m = 2$.

The wavelet power is given by

$$|W_n(s)|^2 \tag{7}$$

NPGD
2, 1227–1273, 2015

and represents the wavelet power spectrum of the time series. Inherent in the wavelet transform are edge effects due to the finite time series. In particular, in a wavelet power spectrum there exists a region called the cone of influence, which is defined as the *e*-folding time of the autocorrelation for wavelet power at each scale. The *e*-folding time is defined as the point at which the wavelet power for a discontinuity at the edge drops by a factor of e^{-2} (Torrence and Compo, 1998).

2.2 Pointwise significance test

In spectral analysis, it is important to assess the statistical significance of spectral power against a noise background. In geophysical applications of wavelet analysis, one often tests each individual wavelet power coefficient against a stationary red-noise background to determine their statistical significance (Torrence and Compo, 1998). For a first-order autoregressive (Markov) process

$$x_n = \rho x_{n-1} + w_n, \tag{8}$$

where ρ is the lag-1 autocorrelation coefficient and w_n is Gaussian white noise with $x_0 = 0$, the normalized theoretical stationary red-noise power spectrum is given by

$$p_f = \frac{1-\rho^2}{1+\rho^2-2\rho\cos(2\pi f/N)}, \tag{9}$$

where $f = 0,\ldots,N/2$ is the frequency index (Gilman et al., 1963). To obtain, for example, the 5 % pointwise significance level ($\alpha = 0.05$), one must multiply Eq. (9) by the 95th percentile of a chi-square distribution with two degrees of freedom and divide the result by 2 to remove the degree-of-freedom factor (Torrence and Compo, 1998). The result of the so-called pointwise testing procedure is a subset of wavelet power coefficients whose values exceed the specified background noise spectrum. Recall from Sect. 1 that significant wavelet power coefficients often occur in clusters or contiguous regions

NPGD
2, 1227–1273, 2015

of pointwise significance called pointwise significance patches (referred to as patches, hereafter).

Consider the time series of the Atlantic Multi-decadal Oscillation (AMO), North Atlantic Oscillation (NAO), Pacific Decadal Oscillation (PDO), and Niño 3.4 indices shown in Fig. 1. The PDO index describes detrended sea surface temperature (SST) variability in the North Pacific poleward of 20° N latitude (Mantua and Hare, 2002) and the AMO index captures the detrended SST variability in the Atlantic Ocean basin (Kerr, 2002). As shown in Fig. 1a and b, the PDO and AMO indices exhibited multi-decadal variability, with periods of 20–60 years, from, respectively, 1856 to 2014 and 1900 to 2014. The reason for the low-frequency variability of the PDO is subject to debate. Some studies suggest it is the reddened response to white noise atmospheric forcing, whereas other studies hypothesize that it is also the integrated response of the El-Niño/Southern Oscillation (ENSO) signal (Newman et al., 2004). The NAO index, an atmospheric index, quantifies the difference in sea-level pressure of the Icelandic Low and the Azores High and is related to the strength and position of the jet stream across the North Atlantic (Hurrell et al., 2003). As shown in Fig. 1b, the NAO mainly operated on time scales of months and seasons, and the raw time series is rather noisy. The Niño 3.4 index is an oceanic metric for quantifying the strength of ENSO and is defined as SST anomalies in the Equatorial Pacific in the region bounded by 120–170° W and 5° S–5° N (Trenberth, 1998). The Niño 3.4 index time series exhibited variability on an array of time scales, especially in the 2–7 year period band. Various physical interpretations for the 2–7 year oscillation have been proposed, including the Unified Oscillator, Delayed Oscillator, and the Recharge Oscillator (Wang et al., 2004).

Shown in Fig. 2 are the wavelet power spectra of the AMO, NAO, Niño 3.4, and PDO indices. The wavelet power spectrum of the AMO detected enhanced variance at a period of 512 months, as indicated by the thin contour that encloses a region of 5 % pointwise significance. All of the other patches are located at periods less than 32 months. The wavelet power spectrum of the NAO indicated that the NAO exhibited enhanced variability on an array of time scales. For example, there is a patch located

NPGD
2, 1227–1273, 2015

at a period of 64 months and 1910. Like the wavelet power spectrum of the AMO, numerous patches were also found at periods less than 32 months. Large regions of enhanced variance were found in the wavelet power spectrum of the Niño 3.4 index. The largest of these regions was located in the time period 1950–2014 and the period band 16–32 months, consistent with how the ENSO varies with periods of 2–7 years. In the same patch, there are holes as described by Schulte et al. (2015) that may indicate the presence of nonlinearities. Holes are defined formally as follows: for the closed unit interval $I = [0\,1]$, let $f : I \to P$ be a continuous closed path in a significance patch P. A patch is said to contain a hole if there exists a path that cannot be continuously deformed into a point, where the feature obstructing such a deformation is the hole. Two patches in the same period band were also identified from 1870 to 1890. For the wavelet power of the PDO index, a large patch centered at a period of 512 months extending from 1910 to 2013 was detected. Most of the patches, however, were located at periods less than 8 months, time scales not typically associated with the PDO.

2.3 Geometric significance test

To determine if the results from the pointwise test are artifacts of multiple testing, a geometric test was applied to the patches located in the wavelet power spectra (Schulte et al., 2015). The test statistic for the geometric test is given by a normalized area

$$A_n = \frac{A}{A_R}, \tag{10}$$

where A is the area of the patch and A_R is the area of the reproducing kernel dilated and translated according to the centroid of the patch. Regarding patches as polygons with vertices (x_k, y_k) with $k = 1, \ldots, m-1$, the area of the patch is determined by a simple formula given by

$$A = \frac{1}{2}\left|\sum_{k=0}^{m-1}(x_k y_{k+1} - x_{k+1} y_k)\right|, \tag{11}$$

NPGD
2, 1227–1273, 2015

where $y_0 = y_m$ and $x_0 = x_m$. Similarly, the centroids of the polygons are given by

$$C_t = \frac{1}{6A} \sum_{k=0}^{m-1} (x_k + y_{k+1})(x_k y_{k+1} - x_{k+1} y_k) \quad (12)$$

and

$$C_s = \frac{1}{6A} \sum_{k=0}^{m-1} (y_k + x_{k+1})(x_k y_{k+1} - x_{k+1} y_k), \quad (13)$$

where C_t and C_s are the time and scale coordinates, respectively, of the centroid (Schulte et al., 2015).

To determine the critical level of test, a large ensemble of patches under a noise model is generated and A_n is computed for each patch, resulting in a null distribution from which the desired critical level of the test can be obtained.

The application of the geometric test to the wavelet power spectra of the AMO, NAO, PDO, and Niño 3.4 time series is also shown in Fig. 2. For the AMO index 5 % geometrically significant patches were identified in the period band 2–16 months prior to 1880, after which no patches were identified as geometrically significant. Note that the large patch centered around the period of 512 months was not found to be geometrically significant at the 5 % level, suggesting that the multi-decadal variability is stochastic. A few geometrically significant patches were identified in the wavelet power spectrum of the NAO: one centered at a period of 64 months and 1910, a second one centered at a period of 8 months and 1910, and several others centered at a period of 4 months throughout the record length. Many geometrically significant patches were identified in the wavelet power spectrum of the Niño 3.4 index. For example, a rather large geometrically significant patch is located in the 16–64 month period band from 1950 to 2014. There are also many geometrically significant patches located in the 2–8 month period band from 1900 and 1950. Numerous geometrically significant patches were identified in the wavelet power spectrum for the PDO index, all of which were located in

NPGD
2, 1227–1273, 2015

the 2–8 month period band. There was also a large patch centered on a period of 512 months, but the patch was not found to be geometrically significant, suggesting that, like the AMO, the multi-decadal variability arose from stochastic processes.

2.4 Sensitivity of the geometric test to the chosen pointwise significance level

To show that the geometric test is sensitive to the pointwise significance level chosen, it will be useful to compute the quantity

$$r = \frac{N_{\alpha_1,\alpha_2}}{N_{\alpha_1}}, \tag{14}$$

where N_{α_1,α_2} is the number of geometrically significant patches at pointwise significance level α_1 that are also geometrically significant at pointwise significance level α_2 and N_{α_1} is the number of patches at α_1 that are geometrically significant, where the geometric significance level for all cases is fixed α_{geo}. In the ideal situation, $r = 1$, indicating that geometrically significant patches never lose their geometric significance as the pointwise significance level is increased. This case, however, is optimistic, as the calculation of geometric significance is rather stochastic. To demonstrate the stochastic nature of the geometric test, r was computed for 1000 wavelet power spectra of red-noise processes with lengths 1000 and lag-1 autocorrelation coefficients equal to 0.5 under four different scenarios, the first of which (Scenario 1) is the case in which $\alpha_1 = 0.1$, $\alpha_2 = 0.05$, and $\alpha_{\text{geo}} = 0.05$ (Fig. 3a). With the mean of r (denoted by $\overline{r}$ hereafter) being 0.3, it can hardly be expected that a geometrically significant patch at $\alpha_1 = 0.1$ to remain significant when the pointwise significance level is changed to $\alpha_2 = 0.05$, at least in the case of red-noise processes. Scenario 2, shown in Fig. 3b, is the same as Scenario 1 except that $\alpha_{\text{geo}} = 0.01$. In this case, $\overline{r} = 0.15$, suggesting that the geometric test is even more sensitive to the chosen pointwise significance level for smaller α_{geo}. Also note that, unlike the distribution shown in Fig. 3a, the distribution is skewed, favoring lower values and supporting the idea that the geometric test is more sensitive to the chosen pointwise significance level for $\alpha_{\text{geo}} = 0.01$.

NPGD
2, 1227–1273, 2015

In Scenario 3, $\alpha_1 = 0.05$ and $\alpha_2 = 0.01$, with $\alpha_{geo} = 0.01$. The distribution shown in Fig. 3c is even more skewed than that corresponding to Scenario 2, with $\overline{r} = 0.05$. Also note that in many cases $r = 0$, indicating that there are patches that are not geometrically significant for both $\alpha_1 = 0.05$ and $\alpha_2 = 0.01$. The reason is that some patches existed at $\alpha_1 = 0.05$ but did not exist at $\alpha_2 = 0.01$ so that their normalized areas are zero.

Scenario 4 was similar to Scenario 3 except that $\alpha_{geo} = 0.05$. Although Scenarios 3 and 4 used the same pointwise significance levels, the results differ, with $\overline{r} = 0.22$, suggesting that the geometric test is less sensitive for larger α_{geo}. The results are similar to that of Scenarios 1 and 2, where increasing the pointwise significance level increased the sensitivity of the geometric test to the chosen pointwise significance level.

3 Persistent topology

3.1 Persistent homology

Before developing the cumulative areawise test, it will be necessary to understand the topology of features found in a typical wavelet power spectrum. It will be especially important to understand how the features evolve as the pointwise significance level is increased or decreased. Such information can be obtained using persistent homology, a tool in applied algebraic topology (Edelsbrunner and Harer, 2008). Persistent homology will provide a formal setting for calculating the lifetimes of patches and holes, where lifetimes describe when features first appear and when they disappear. For example, a patch appearing at the 5 % pointwise significance level and vanishing at the 1 % pointwise significance level would have a lifetime of 4. In this paper, the intuitive foundation of persistent homology will be described rather than giving a formal mathematical exposition.

NPGD
2, 1227–1273, 2015

Suppose that A is a patch at the pointwise significance level $\alpha = \alpha_1$ as shown in Fig. 4a. One can increase the size of the patch A by increasing α to α_2, which lowers the threshold for significance, resulting in the new geometric realization A' of A shown in Fig. 4a. The evolution of the patch can be monitored using a barcode (Ghrist, 2008), which is a collection of horizontal lines representing the birth and death of features. Following the convention of persistent homology, the y axes of barcodes will be denoted by H_0 for patches. In algebraic topology, H_0 are called homology groups and measure the path-connectedness (Appendix A) of sets. The patch A, being created at α_1, results in the line segment beginning at α_1 in Fig. 4e. Furthermore, the patch neither vanishes nor merges with another patch at α_2 so that the horizontal line continues to α_2. Note that a new patch B is created at α_2 because the pointwise significance test is less stringent. The creation of the patch results in a new line starting from α_2. A more complicated situation occurs at α_3 where the patches A' and B merge and result in a new patch C. In this case, it is unclear if C is a geometric realization of A' or B. It will therefore be necessary to use the so-called Elder rule from persistent homology (Edelsbrunner and Harer, 2009). According to this rule, the oldest patch will continue to live and the younger patch will die entering the merger point. In the present case, A' is the older patch because it first appears as A at α_1 and B is the younger patch, being created after A at α_2. Therefore, according to the Elder rule, the horizontal line corresponding to A in the barcode continues to α_4, but the line corresponding to B terminates at α_3, as it dies entering α_3. Also note the creation of a new patch D at α_3 and the corresponding line segment in the barcode. Another merger occurs at α_4 and the Elder rule determines that the line segment for D ends and the horizontal line for A continues, where the arrow indicates that A never dies.

Persistent indices of patches can also be obtained, which will be defined as the difference between the pointwise significance level for which the features die entering and the level at which the features were born. For patches that never die, their persistent indices, by convention, will be set to infinity.

NPGD
2, 1227–1273, 2015

3.2 Persistent homology of red-noise

To understand the topology of patches generated from red-noise processes, it is useful to use Monte Carlo methods to determine, for example, the number of patches at a particular pointwise significance level, or similarly, the number of holes. Shown in Fig. 5a is the ensemble mean of the number of patches (denoted by β_0, hereafter) as a function of α for the Moret, Paul, and Dog wavelets obtained from generating 100 wavelet power spectra of red-noise processes of length 300 and computing β_0 for each of the wavelet power spectra at each pointwise significance level.

For the the Morlet and Paul wavelets, the number of patches reached minima at $\alpha = 0.01$ and $\alpha = 0.99$ and maxima at $\alpha = 0.18$. The minima in the number of patches for the Dog wavelet was the same as the Morlet and Paul wavelets, but the maximum occurred at $\alpha = 0.5$. Perhaps more interesting are the number of holes for all three wavelets: very few holes existed at low pointwise significance levels and the number increased rapidly until $\alpha = 0.95$, at which point β_1 no longer increased.

To understand more fully the curves shown in Fig. 5a, the persistent homology of patches generated from red-noise processes of length 150 was computed as α varies from 0.01 to 0.99, and barcodes representing the evolution of patches in the wavelet power spectra were computed. In each case the lag-1 autocorrelation coefficients were set to 0.5, but the results are identical for other autocorrelation coefficients. Shown in Fig. 6a is a barcode corresponding to a wavelet power spectrum obtained using the Morlet wavelet. Recalling that the beginning of the line segment represents the birth of patches, the barcode indicates that a few patches were born at $\alpha = 0.02$. As α increases to $\alpha = 0.3$ more patches are born, consistent with the fact that more spurious results occur for larger pointwise significance levels. Note that, for $\alpha > 0.2$, patches begin to die, representing the merger of smaller patches into larger patches. The merging process occurs until $\alpha = 0.7$, at which point all patches have merged into a single patch, which is called the essential class. To show that the distribution of the lifetimes for patches generated from red-noise processes is not random, 100 wavelet power spec-

NPGD
2, 1227–1273, 2015

tra of red-noise processes were generated and the persistence indices for all patches in each wavelet power spectra were computed (Fig. 6b). The resulting distribution indicates that lifetimes of patches is typically 0.1 and relatively few patches live longer than 0.6. Overall, the distribution characterizes patches generated from red-noise processes as short-lived.

A typical barcode corresponding to a wavelet power spectrum of a red-noise process whose wavelet power spectrum was obtained using the Paul wavelet is shown in Fig. 6c. The barcode is similar to that of the Morlet wavelet, with many patches being born before $\alpha = 0.3$ and generally merging for $\alpha > 0.3$. The distribution of persistence indices shown in Fig. 6d obtained from Monte Carlo methods as described above suggests that the lifetime patches are typically longer than that generated from the Morlet wavelet; the distribution has smaller negative skewness, and there are more persistence indices within the range 0.6–0.85.

A barcode corresponding to a wavelet power spectrum of a red-noise process obtained using a Dog wavelet is shown in Fig. 6e. Consistent with Fig. 6, the barcode differs from that of the other wavelets, with many patches merging for $\alpha > 0.4$. In fact, unlike the case for the Morlet wavelet, many patches were found to have merged after $\alpha > 0.8$. The merging for large α suggests that patches generated using the Dog wavelet tend to be smaller so that it takes longer to form the essential class. In other words, larger patches represent a larger fractional area of the wavelet domain so that fewer of them can reside in the wavelet domain, whereas a greater number of small patches can exist in a wavelet domain of equal size. The results from the Monte Carlo methods shown in Fig. 6f show that the distribution of persistence indices is indeed different. Compared to the Morlet and Paul wavelets, fewer patches were found to have persistence indices less than 0.15, and persistence indices are more uniformly distributed in the range 0.2–0.7.

The non-random evolution of patches and holes for red-noise processes suggests that a test can be developed that uses the information of patches at many pointwise significance levels. One would expect that patches arising from signals would behave

NPGD
2, 1227–1273, 2015

differently, breaking apart less frequently when the pointwise significance level is increased. In the extreme case that a patch is the result of taking the wavelet transform of a pure sinusoid, there would be a single patch at all pointwise significance levels and that patch would not break apart. Suppose, however, that white noise was added to the sinusoid such that the patch breaks apart for small α. In this case, the area of the patch would be smaller with respect to the pure sinusoid case, the difference in area arising from the splitting of the patch. On the other hand, if α were increased, then the patch would become a single patch by merging, and the area will be closer to that of the pure case. These facts suggests that a test can be constructed whose test statistic is calculated over a set of pointwise significance levels, capturing the behavior of patches as they evolve. In particular, the test should make use of how the evolution of patches and holes under the null hypothesis of red noise is not random and how the evolution of patches arising from signals may differ from that of the null hypothesis.

4 Development

4.1 Geometric pathways

Unlike the geometric test that assesses the significance of patches at a single pointwise significance level, the cumulative areawise test (referred to as the areawise test, hereafter) will assess the significance of a patch as it evolves under a changing pointwise significance level. The goal of the method is to remove the binary decision from which the geometric test suffers. The idea behind the test statistic will be that a patch that is consistently geometrically significant for different pointwise significance levels is more significant than a patch that is only geometrically significant at a single pointwise significance level. In other words, the geometric test simply assesses the significance of patches at a single pointwise significance level but does not take into account that the patch could have been significant at that particular pointwise significance level by chance (Fig. 3).

NPGD
2, 1227–1273, 2015

The first step of the areawise test is to select a finite set of pointwise significance levels that remains fixed throughout the testing procedure. Although there are infinite number of pointwise levels to choose from, it will be shown through empirical arguments that the selection need only be limited to a finite set of pointwise significance levels. Furthermore, it will be shown that one needs only to limit the test to a certain range of pointwise significance levels, making the test more computationally feasible. The selection of pointwise significance levels will be discussed in depth in Sects. 4.2. and 4.5. For the theoretical development of the testing procedure, the pointwise significance levels will be left unspecified.

The evolution of a patch under a changing pointwise significance level will be made precise by introducing the notion of a geometric pathway. A geometric pathway will be defined as a collection P of r patches at the corresponding pointwise significance levels $\alpha_1 < \alpha_2 < \ldots < \alpha_r$ such that

$$P_1 \subset P_2 \subset P_3 \subset \ldots \subset P_r \tag{15}$$

and

$$g_1 < g_2 < g_3 \ldots < g_r, \tag{16}$$

where each g_j is a normalized area corresponding to the patch P_j. For this testing procedure, the normalized area will be calculated by dividing the patch area by the scale coordinate of the centroid squared. The inequalities (Eq. 16) are guaranteed to hold for any nested sequence (Eq. 15) (Appendix B). Viewing the α_j's as time parameters, one can think of the patch being in its initial configuration at α_1 and at its final configuration at α_r. The length of a pathway will be given by r, the number of elements in the pathway. If the computation of the geometric pathways is limited to an interval $I = [\alpha_{\min}, \alpha_{\max}]$, then all pathways will end at the same pointwise significance level but need not begin at the same pointwise significance level. The reason why all pathways end at $\alpha_{\max}$ is because, once a pathway is generated, it can never die, as elements of P grow in size relative to its initial element.

NPGD
2, 1227–1273, 2015

To illustrate the idea of a pathway, it is perhaps best to consider an ideal case (Fig. 7). Consider, for example, three pathways X, Y, and Z whose lengths are, respectively, $r_x = 4$, $r_y = 4$, and $r_z = 5$. The pathway X can be written explicitly as

$$X_1 \subset X_2 \subset X_3 \subset X_4, \tag{17}$$

indicating that the patch exists at $\alpha_1^X = \alpha_2$, $\alpha_2^X = \alpha_3$, $\alpha_3^X = \alpha_4$, and $\alpha_4^X = \alpha_5 = \alpha_{\max}$. There also exists another pathway Y such that $X_2 = Y_2$, $X_3 = Y_3$, and $X_4 = Y_4$ but whose initial element Y_1 is distinct from X. The pathways are thus distinct until α_3, at which point the pathways merge, resulting in the remaining elements being identical (Fig. 7b). The third pathway Z, on the other hand, shares only one element with X and Y, merging at α_5, so that Z represents a distinct pathway. Unlike the other pathways, the pathway Z begins at $\alpha_1 = \alpha_{\min}$, implying that the pathway exists over the entire interval I.

The development of the areawise test will require the calculation of a test statistic for all pathways in the wavelet domain that end at the same pointwise significance level for a selected interval I. The test statistic used in this procedure will be the total sum of normalized areas

$$\gamma = \sum_{j=1}^{r} g_j, \tag{18}$$

one for each pathway.

The calculation of the critical level for the test can be computed using Monte Carlo methods as follows: (1) fix I and a partition of the interval with uniform spacing and generate red-noise processes with the same autocorrelation coefficients as the input time series, (2) for each red-noise process generate synthetic wavelet power spectra and all corresponding pathways; and (3) for every pathway, compute γ, resulting in a null distribution from which the desired critical level of test can be obtained. The critical level corresponding to the 5 % significance level of the test, as an example, is the 95th percentile of the null distribution.

NPGD
2, 1227–1273, 2015

4.2 Pointwise significance level selection: maximization method

It will often happen that pathways under consideration are of different lengths so that it is not obvious what pointwise significance patches to report after the implementation of the testing procedure. The problem can be circumvented, however, using the following procedure: let

$$\gamma_j = \sum_{i=0}^{r-j} g_{r-i} \tag{19}$$

be the cumulative sum associated with the jth element P_j of a pathway with length r, and let γ_{crit} be the critical level of the test; then, the appropriate pointwise significance level to use for a pathway is determined by the following rule:

$$\gamma_{\max} = \max_{j=1,2\ldots,r} \gamma_j > \gamma_{\text{crit}}, \tag{20}$$

where the statistic satisfying the above rule is denoted by $\gamma_{\max}$. The element of the pathway corresponding to $\gamma_{\max}$ is the output of the testing procedure. Note that the output elements need not be located at the same pointwise significance levels, contrasting with the pointwise and geometric tests. If, for a given pathway the statistic $\gamma_{\max}$ does not exist, then the pathway is deemed insignificant. It is also important to note that using the inequality (Eq. 20) guarantees that the appropriate number of pathways will be deemed insignificant if the null hypothesis is known to be true, the reason for which is that if M pathways have statistics γ such that $\gamma > \gamma_{\text{crit}}$, then there also must exist M statistics $\gamma_{\max}$ satisfying the inequality (Eq. 20). In other words, the sums γ must have crossed the critical level at some point in the pathways, or, otherwise, they would have not been deemed significant.

NPGD
2, 1227–1273, 2015

4.3 Application to ideal pathways

The testing procedure is now demonstrated using an ideal case. For the pathway X shown in Fig. 7 the test statistic at α_5 is equal to the area of a single patch in the pathway, i.e.

$$\gamma_x = A_4^x, \tag{21}$$

where A_j^x denotes the area of a pathway element at α_j^x. The test statistic at α_1^x, on the other hand, is the sum of four areas

$$\gamma_x = A_1^x + A_2^x + A_3^x + A_4^x \tag{22}$$

so that the test statistic reaches its maximum value at the smallest pointwise significance level for which the pathway exists. In fact, all pathways will satisfy this property. The evolution of the patch can be viewed as a shape changing with time as shown in Fig. 7b. Viewing α_5 as a time parameter, one can say that X merges with a large patch at time α_5 and first appears at time α_2. The evolution of the test statistic is shown in Fig. 7c, in which the test statistic increases with decreasing α_j such that its maximum value is attained at α_2.

Now consider the pathway Y, whose maximal test statistic is similar to that of X except that A_2^y differs from A_2^x. According to Fig. 7b, the pathways evolved identically from α_3 to α_5, merging at α_3. As shown in Fig. 7c, the evolution of γ_y only differs slightly from that of γ_x and also represents a significant pathway. For both X and Y, the condition $\gamma_{\max} > \gamma_{\text{crit}}$ is satisfied at α_3 so that a single patch at α_3 will be the output of the testing procedure, not the two patches representing distinct elements of X and Y at α_2. As will be shown in Sect. 4.6, the overall effect of this merging process is to enhance the detection of signals.

For the pathway Z, the maximum value of the test statistic can be decomposed into five summands, i.e.

$$\gamma_z = A_1^z + A_2^z + A_3^z + A_4^z + A_5^z. \tag{23}$$

NPGD
2, 1227–1273, 2015

The pathway Z, being longer than X and Y, therefore has an additional pointwise significance level to allow the test statistic to exceed the critical level of the test. However, the elements of the pathway have smaller normalized areas so that the pathway is not significant. The length of a pathway is thus not the only factor influencing the significance of a pathway, as the size of elements is also important.

4.4 The null distribution

Recall from Sect. 3.2 that, for patches generated from red noise, the merging of patches is not random, with typical lifetimes of patches following a non-uniform distribution, favoring shorter lifetimes. Therefore, if the test statistic is proportional to the lifetime of patches, one can expect the test statistic to follow a similar distribution to that of persistence indices, where the smallest values of γ are preferred. To test this hypothesis, 1000 wavelet power spectra of red-noise processes with fixed autocorrelation coefficients were generated and the cumulative area of all pathways in each wavelet power spectra was computed. In the experiments, $\alpha_{\text{min}} = 0.02$, $\alpha_{\text{max}} = 0.82$, and $\Delta\alpha$, the discrete spacing between adjacent pointwise significance levels, is set to 0.02 to make calculations less computationally expensive. The results for the Morlet, Paul, and Dog wavelets are shown in Fig. 8.

The distribution of γ for the Morlet wavelet is generally similar to the shape of the distribution for the persistence indices for H_0 (Fig. 6b), where the smallest values of γ are preferred. It turns out that the distribution of γ for the Morlet wavelet can be well described by an exponential distribution. Using the method of maximum likelihood (Weerahandi, 2003), a theoretical exponential distribution was fitted to the empirical distribution, where the empirical distribution was found to be best described by an exponential distribution with mean 6.5. To show that the theoretical distribution models the empirical distribution, the percentiles of a theoretical exponential distribution with mean 6.5 were plotted as a function of the percentiles of the empirical distribution (Fig. 8b). The linear relationship between the percentiles shown in Fig. 8b indicates

NPGD
2, 1227–1273, 2015

that the theoretical distribution well models the empirical distribution, with the 95 % percentiles only differing by 1.0.

The results for the Paul wavelet are generally similar except that the mean value of the test statistic was found to be smaller with a value of 6.2. The distribution of the test statistic, unlike that for the Morlet wavelet, was found to be poorly modeled by an exponential distribution, with the empirical distribution generally having larger values than those predicted by a theoretical exponential distribution; this difference was more pronounced for larger γ.

The results for the Dog wavelet are shown in Fig. 8e. The null distribution differs from that for the Morlet and Paul wavelets, with the mean value of the distribution being 3.4. This result is perhaps not surprising; Fig. 6 shows that the lifetimes of patches generated using the Dog wavelet have larger persistence indices attributed to how patches die at larger α compared to that of the Morlet and Paul wavelets, resulting from the lack of merging with the essential class. The delayed merging implies that the area of patches must be smaller relative to that of the Morlet and Paul wavelets for large α to prevent merging with other patches. The effect is to generate a smaller cumulative sum. The null distribution, unlike that for the Morlet wavelet, was not found to be well-described by a theoretical exponential distribution.

As will be shown in Sect. 4.6, the implementation of the proposed testing procedure will allow small patches that are seemingly indistinguishable from noise to emerge, allowing the recovery of a signal that has been contaminated by noise.

4.5 Computational remarks

An important parameter in the areawise test is $\Delta\alpha$, the spacing between pointwise significance levels. It is critical that $\Delta\alpha$ is chosen to be small enough to sample the merging of patches so that, for example, the null distribution can be representative of the null hypothesis. If $\Delta\alpha$ is set too large, then a pathway may be missed entirely if it is born between two adjacent pointwise significance levels. On the other hand, if $\Delta\alpha$ is too small, then the test will become computationally expensive, requiring the calculation of

NPGD
2, 1227–1273, 2015

more normalized areas. The distribution of the persistence indices for H_0 suggests that $\Delta\alpha = 0.01$ because the statistical modes of the distributions are 0.01. However, it will be shown in Sect. 4.6 that the results of the areawise test do not differ if one chooses $\Delta\alpha = 0.02$.

Another important parameter is α_{max}. It would be computationally inefficient if the testing procedure had to be performed on pathways whose end points are, for example, $\alpha_{max} = 0.99$. Moreover, as shown in Fig. 5b, the number of holes increases for increasing α, which would increase the computational costs associated with the calculation of patch areas for large α. The reason for the increase in computational costs is that the areas of the holes need to be subtracted from the area the patches would have if they did not contain holes. To circumvent the problem, properties of the areawise test can be assessed at peaks of the curves shown in Fig. 5a so that results for other end points can be inferred. For the Morlet wavelet, the peak occurs at $\alpha = 0.18$ so that $\alpha_{max} = 0.18$. To the left of the maximum, the average number of patches are equal to that of some pointwise significance level to the right of the maximum. However, patches to the right of the maximum have larger areas so that choosing $\alpha_{max} > 0.18$ would result in the testing procedure detecting larger patches as significant. The same reasoning holds for the Paul and Dog wavelets except that the maximum of the curve for the Dog wavelet is $\alpha_{max} = 0.5$. The sensitivity of the testing procedure to α_{max} is discussed in Sect. 4.6.

Figure 6 also suggests that α_{max} need not be any larger than 0.7 for the Morlet wavelet because all patches have merged with the essential class before that point, at least for red-noise processes. Thus, for $\alpha_{max} > 0.7$, patches arising from signals cannot be distinguished from those generated from noise because all patches have merged into a single, large patch. A similar argument holds for the Paul wavelet, but, for the Dog wavelet, all patches merge at a larger α_{max}.

NPGD
2, 1227–1273, 2015

4.6 Comparison with geometric test

With the areawise test now developed, it will be useful to assess the statistical power of the test relative to that of the geometric test. The first aspect of the assessment will be to quantify how well both tests detect true positive results. To do so, let

$$x(t) = A\sin(2\pi f t) + w(t) \tag{24}$$

be a sinusoid with amplitude A, frequency f, and additive Gaussian white noise $w(t)$. The goal will be to evaluate the ability of both tests to detect true positives within a particular period band. A theoretical patch to which the ability of the geometric and areawise tests were compared was constructed as follows: (1) the time series $x(t)$ for all $t \in [0, 500]$ was generated but with no additive white noise, (2) the wavelet power spectra of $x(t)$ was computed and the 5 % pointwise significance test was performed on the wavelet power spectrum; and (3) the width of the significance patch in the wavelet power spectrum was calculated at $t = 250$ where edge effects are negligible. The theoretical patch calculated using the Morlet wavelet is indicated by dotted lines in Fig. 9, where the theoretical patch is a rectangle of fixed width extending from $t = 0$ to $t = 500$. In all experiments, $\alpha_{\max} = 0.18$ and $\Delta\alpha = 0.02$, but implications of other choices are discussed at the end of the section.

To assess the ability of the tests to detect true positives, the area of patches deemed significant by the tests were compared to the total area of the theoretical patch. More specifically, if P_{geo} is the union of all pointwise significance patches at α that are geometrically significant at the α_{geo} level and P_{theory} is the theoretical patch, then

$$r_{\text{a}} = \frac{A_{P_{\text{geo}} \cap P_{\text{theory}}}}{A_{P_{\text{thoery}}}} \tag{25}$$

represents the areal fraction of P_{theory} detected by the geometric test, where $A_{P_{\text{geo}} \cap P_{\text{theory}}}$ denotes the area of $P_{\text{geo}} \cap P_{\text{theory}}$ and $A_{P_{\text{thoery}}}$ denotes the area of P_{theory}. If $r_{\text{a}} = 1$, then the

NPGD
2, 1227–1273, 2015

test detected all of the true positive results that are known by construction. Small values of r_a indicate that the tests performed poorly, detecting only a fraction of the theoretical patches to be significant. A similar construction can be made for the areawise test by replacing P_{geo} with P_c. Figure 9a illustrates the procedure for the areawise test when f = 0.8, A = 0.8, and the signal-to-noise ratio (defined below) equals 1.0. As indicated by the thick black contours, the areawise test was able to detect 30 % of the true positives comprising the theoretical patch, whereas Fig. 9b shows that the geometric test was only able to detect 20 % of the true positives. It will be necessary to compute N = 1000 values of r_a for different values of f and signal-to-noise ratios of the Gaussian white noise to determine if the tests truly perform differently. In the following experiments, the signal-to-noise ratio is defined as

$$\sigma = 10 \log \left(\frac{p_{signal}}{p_{noise}} \right), \tag{26}$$

where

$$p_{signal} = \frac{A^2}{2}, \tag{27}$$

p_{noise} is the average power of the Gaussian white noise, and σ is measured in decibels (DB). In the discussion below the results for the Morlet are presented first; results for the Paul and Dog wavelets are discussed at the end of the section. It is also noted that because σ and A do not vary independently there is no need to perform different experiments for different values of A. For the experiments, A was set to 1.0.

In the first experiment, the areawise significance (denoted by α_c, herreafter) was set to 0.05, α_{geo} = 0.05, and α = 0.01, 0.05, 0.1. The signal-to-noise ratio for the Gaussian white noise was varied from −5 to 5 DB. The results are shown in Fig. 10. For both tests, the ability to detect true positives increased with increasing signal-to-noise level. At low signal-to-noise ratios, the tests performed similarly, detecting on average 10 % of true positives. On the other hand, differences between the test performances became

NPGD
2, 1227–1273, 2015

larger as the signal-to-noise ratio was increased; in fact, the areawise test outperformed the geometric test regardless of the chosen pointwise significance levels when $\sigma \geq -2.5$ DB. The results indicate that the areawise test is particularly useful in low noise situations but one can expect the test to detect more true positives even in high noise conditions. It also worth noting that the performance of the geometric test depended strongly on the chosen pointwise significance level, especially when the signal power was high.

A second experiment was conducted where $\alpha_c = 0.01$ and $\alpha_{geo} = 0.01$. The same pointwise significance levels as the first experiment were chosen and the range of signal-to-noise ratios was also the same. The results are presented in Fig. 11. Note that the shape of the curves for both tests are the same as the first experiment, where greater true positive detection is preferred for large signal-to-noise ratios. However, both tests detected fewer true positives, consistent with how significance levels of the tests were increased. The relationships between the tests, like for experiment 1, depended on the signal-to-noise ratio; for low signal-to-ratios the performance of the tests are similar, whereas for high signal-to-noise ratios differences between r_a for the areawise test and r_a for the geometric increases. All the results were found to be statistically significant.

Additional experiments were performed using different values of f to determine if the frequency at which patches are located affects the performances of the areawise and geometric tests. True positive detection, for a fixed σ, was generally found to increase for larger f, though the areawise test was still found to detect more true positives. The behavioral assessments of both tests for $f < 0.8$ required the use of synthetic time series whose lengths were greater than 500, as patches lengthened in the time direction for lower frequencies and the COI impacts became more pronounced. The array of experiments concluded that the areawise test should be the preferred method for signal detection.

Another set of experiments were conducted to evaluate how the areawise test detects false positives relative to the geometric test. In the first experiment, 1000 wavelet

NPGD
2, 1227–1273, 2015

power spectra arising from red-noise processes of length 1000 with equal lag-1 autocorrelation coefficients were generated. For each wavelet power spectra, the geometric and areawise tests were applied at the 0.05 level and the pointwise significance level was set to 0.05. The results were found to be independent of the chosen lag-1 autocorrelation coefficient, so only the results for the case when the lag-1 autocorrelation coefficient was set to 0.5 are presented. The ratio between the number of false positive results for the geometric and areawise tests for each wavelet power spectra was then computed. The average ratio was found to be 0.08, implying that one can expect more false positive results for the areawise test relative to the geometric test. However, in most cases, the increase in the number of false positive results with respect to the geometric test was small compared to the increase in true positive detection (Figs. 10 and 11), suggesting that the areawise test had greater overall statistical power. Note that spurious patches arising from the areawise test will have generally larger areas so that results from the test should be interpreted carefully.

The above experiment was repeated for different pointwise significance levels. The experiments indicated that the difference between false positive detection between both tests decreased as α increased, the reason for which is that the number of false positive results for the geometric test for a fixed α_{geo} will increase until $\alpha = 0.18$, the point at which the number of patches generally peaks (Fig. 5a). For $\alpha > 0.18$, the number of spurious patches may decrease, but the areas of the spurious patches resulting from the geometric test will be larger. A similar argument would hold if α_c and α_{geo} were increased or decreased by the same amount.

All the above experiments were performed with the Paul and Dog wavelets and the theoretical 5 % pointwise significance patch was adjusted to account for the different scale- and time-localization properties of the wavelets, where the corresponding theoretical patches were found to be wider in scale compared to the Morlet wavelet. The results from the experiments were qualitatively similar to that of the Morlet wavelet, as the areawise test detected more true positives than the geometric test for both the Paul and Dog wavelets. Like for the Morlet wavelet, the differences in performances were

NPGD
2, 1227–1273, 2015

more pronounced for larger σ. The array of experiments provides evidence that the areawise test has greater statistical power than the geometric test regardless of the chosen analyzing wavelet.

All the above experiments were conducted for different values of $\Delta\alpha$. For $\Delta\alpha < 0.02$, the results of the experiments were virtually identical. On the other hand, if $\Delta\alpha > 0.03$, the percentage of true positives detected by the areawise test decreased. The results suggest that $\Delta\alpha$ should be chosen to be no larger than 0.02 to ensure that true positive detection is maximized.

The parameter α_{max} was found to strongly influence the performance of the areawise test. For $\alpha_{max} > 0.18$ and $\alpha_c = 0.05$, the number of true positives detected increased and the number of false positives decreased relative to the geometric test with $\alpha_{geo} = 0.05$. However, the normalized areas of the spurious patches resulting from the areawise test with $\alpha_c = 0.05$ or $\alpha_c = 0.01$ were found to be 2–10 times larger than those resulting from the geometric test, making spurious features appear significant. The problem was remedied by decreasing α_c, but this adjustment was found to be the same as increasing α_{max}. For $\alpha_{max} < 0.18$, the true and false positive detection of the areawise test for $\alpha_c = 0.05$ approached that of the geometric test until $\alpha_{max} = 0.05$, at which point they were approximately equal.

5 Climate applications

To determine if any features in the wavelet power spectra of the AMO, NAO, Niño 3.4, and PDO time series are deterministic or stochastic, the areawise test was performed on geometric pathways in the wavelet power spectra at the 0.01 level. A red-noise background spectrum was used for each, with $\alpha_{max} = 0.18$, $\alpha_{min} = 0.02$, and $\Delta\alpha = 0.02$. The wavelet power spectrum for the AMO index is shown in Fig. 12a. Although the AMO is usually characterized by its multi-decadal variability, no areawise significant wavelet power coefficients were detected at periods greater than 24 months, suggesting that the variability results from stochastic processes. In contrast, from 1860 to 1940, areawise-

NPGD
2, 1227–1273, 2015

significant wavelet power coefficients were detected, primarily in the 2–16 month period band. After 1940, no areawise-significant wavelet power coefficients were identified at any time scale.

The wavelet power spectrum of the NAO index is shown in Fig. 12b. Only a few regions of areawise significance were detected: one at a period of 2 months and 1900, a second one at a period of 6 months and 1880, and a third one located at a period of 100 months and 1870. The lack of areawise significance suggests that the NAO is a red-noise process with no preferred time scale, consistent with Feldstein (2002).

The wavelet power spectrum for the Niño 3.4 index, on the other hand, does indicate potential predictive capabilities (Fig. 12c). There are two notable features, one extending from 1870 to 1920 in the 16–64 month period band and another one extending from 1960 to 2014 in the 8–64 month period band. Perhaps just as interesting is the deficit in areawise significance from 1920 to 1960 in the 8–64 month period band. The deficit could be the result of the 2–7 year mode being modulated by a decadal ENSO mode, a nonlinear paradigm (Timmermann, 2002). Such a modulation would imply that the behavior of the 2–7 year mode is determined by the phase of the decadal mode, where, for example, more extreme El Niño phases would be favored if the decadal mode is in a positive regime. On the other hand, results shown in Fig. 11c show that neither the decadal nor the multi-decadal variability exceed a red-noise background so modulations would be difficult to predict.

The wavelet power spectrum of the PDO index is shown in Fig. 12d. Like the AMO index, there is enhanced variance at multi-decadal time scales but the variance does not exceed a red-noise background. Areawise-significant regions, however, were detected in the 2–8 month period band from 1900 to 1960. The results indicate that the PDO is a red-noise process, consistent with prior work showing that the PDO results from the oceanic integration of atmospheric white-noise stochastic forcing (Newmann et al., 2003).

NPGD
2, 1227–1273, 2015

6 Conclusions

An areawise test was developed for assessing the significance of features in wavelet power spectra. The test was generally found to have greater statistical power than the geometric test except possibly under high-noise situations, in which case the tests were found to perform similarly. The main advantage of the new testing procedure is that the results are no longer dependent on the chosen pointwise significance level. In contrast, the geometric test results were found to be very sensitive to the chosen pointwise significance level, making it difficult for researchers to decide what patches are significant and what patches are not significant. In particular, the cumulative areawise test was found to detect more true positives relative to the geometric test for some common pointwise and geometric significance levels. The large increase in true-positive detection of the cumulative areawise test was also accompanied by a small increase in false-positive detection compared to the geometric test performed at the 5 % level, resulting in the areawise test having greater statistical power. The difference between the tests, however, was found to decrease for low signal-to-noise ratios, indicating that there are still deterministic features that are, in principle, indistinguishable from background noise.

The results from the climate-mode analysis suggest that the predictability of the AMO, PDO, and NAO is limited and that the multi-decadal variability of the AMO and PDO is the result of a stochastic process. The Niño 3.4 index, by contrast, was found to have deterministic features, implying that future states of ENSO may be predictable. Such predictability is important given that ENSO has regional- to global-scale impacts on precipitation and temperature. The ability to predict future changes of regional climate thus, to some extent, depends on the ability to predict ENSO. However, currently, the future state of ENSO as determined by climate models is uncertain, with some climate models suggesting large changes and others indicating no change at all (Latif and Keenlyside, 2008).

NPGD
2, 1227–1273, 2015

A Matlab software package written by the author to implement the cumulative areawise test is available at justinschulte.com.

Appendix A

A path in a set X is defined as a continuous function $f : I \rightarrow X$. A set X is said to be path-connected if any two points x and y in X can be joined by a path. The path-component of a topological space X is the maximal path-connected subset of a set. Intuitively, one can think of a path-component as the largest isolated piece of the set. For example, the set could be the disjoint union of a square and a disc, in which case both the square and the disc are path-components.

Appendix B

Let P_1 and P_2 be two subsets of a patch P with area A such that P_2 is the set complement of P_1. Let A_1 and A_2 denote the areas of P_1 and P_2, respectively. One can thus write

$$A = A_1 + A_2, \tag{B1}$$

$$A = r_1 A + r_2 A, \tag{B2}$$

and

$$r_2 = 1 - r_1, \tag{B3}$$

NPGD
2, 1227–1273, 2015

where $r_1 r_2 \in [0, 1]$. The centroid of P can be written as

$$\frac{1}{A}\iint_P s\mathrm{d}s\mathrm{d}t = \frac{1}{A}\iint_{P1} s\mathrm{d}s\mathrm{d}t + \frac{1}{A}\iint_{P2} s\mathrm{d}s\mathrm{d}t \tag{B4}$$

$$= \frac{1}{A}\iint_P s\mathrm{d}s\mathrm{d}t = \frac{r_1}{A_1}\iint_{P1} s\mathrm{d}s\mathrm{d}t + \frac{r_2}{A_2}\iint_{P2} s\mathrm{d}s\mathrm{d}t \tag{B5}$$

or

$$C^s = r_1 C_1^s + r_2 C_2^s, \tag{B6}$$

so that

$$\frac{C^s - r_2 C_2^s}{r_1} = C_1^s, \tag{B7}$$

where C_1^s and C_2^s are the scale coordinates of the centroids for P_1 and P_2. The equation implies that

$$C^S - r_2 C_2^s > 0 \tag{B8}$$

because C_1^s is always positive. The normalized areas of P and P_1 are given by

$$A^N = \frac{A}{(C^s)^2} \tag{B9}$$

and

$$A_1^N = \frac{A_1}{(C_1^s)^2}. \tag{B10}$$

NPGD
2, 1227–1273, 2015

Thus,

$$r_{\text{norm}} = \frac{A_1^N}{A^N} = \frac{r_1^3 (C^s)^2}{(C^s - (1 - r_1) C_2^s)^2}. \tag{B11}$$

At $r_1 = 0$, $r_{\text{norm}} = 0$ because P_1 has no area. At $r_1 = 1$, $r_{\text{norm}} = 1$ because $A_1 = A$. Moreover, the function is monotonically increasing for $r_1 \in [0\,1]$ so that $r_{\text{norm}} \leq 1$. The same arguments hold for P_2 except that r_{norm} decreases monotonically.

Acknowledgements. Support for this research was provided by the National Science Foundation Physical Oceanography Program (award number 0961423) and the Hudson River Foundation (award number GF/02/14). It is a pleasure to thank R. G. Najjar and N. Higson for the helpful advice, which resulted in an improved manuscript.

References

Edelsbrunner, H. and Harer, J.: Computational Topology: an Introduction, Amer. Math. Soc., Providence, Rhode Island, 241 pp., 2009.

Edelsbrunner, H. and Harer, J.:Persistent homology – a survey, in: Surveys on Discrete and Computational Geometry, Contemp. Math., 453, Amer. Math. Soc., Providence, Rhode Island, 257–282, 2008.

Efron, B.: Bootstrap methods: another look at the jackknife, Ann. Stat., 7, 1–26, 1979.

Ghrist, R.: Barcodes: the persistent topology of data, B. Am. Math. Soc., 45, 61–75, 2008.

Grinsted, A., Moore, J. C., and Jevrejeva, S.: Application of the cross wavelet transform and wavelet coherence to geophysical time series, Nonlin. Processes Geophys., 11, 561–566, doi:10.5194/npg-11-561-2004, 2004.

Higuchi, K., Huang, J., and Shabbar, A.: A wavelet characterization of the North Atlantic Oscillation variation and its relationship to the North Atlantic sea surface temperature, Int. J. Climatol., 19, 1119–1129, 1999.

Hurrell, J. W., Kushnir, Y., Ottersen, G., and Visbeck, M. (Eds.): The North Atlantic Oscillation: Climatic Significance and Environmental Impact, Geophys. Monogr. Ser. 134, American Geophysical Union, Washington, D.C., 279 pp., 2003.

NPGD
2, 1227–1273, 2015

Jenkins, G. W. and Watts, D. G.: Spectral Analysis and its Applications, Holden-Day, San Francisco, California, 541 pp., 1968.
Kerr, R. A.: A North Atlantic climate pacemaker for the centuries, Science, 288, 1984–1985, 2000.
Labat, D.: Wavelet analysis of the annual discharge records of the world's largest rivers, Adv. Water Resour., 31, 109–117, 2008.
Labat, D.: Cross wavelet analyses of annual continental freshwater discharge and selected climate indices, J. Hydrol., 385, 269–278, 2010.
Latif, M. and Keenlyside, N. S.: El Niño/Southern Oscillation response to global warming, P. Natl. Acad. Sci. USA, 106, 20578–20583, 2008.
Lau, K. M. and Weng, H.: Climate signal detection using wavelet transform: how to make a time series sing, B. Am. Meteorol. Soc., 76, 2391–2402, 1995.
Lee, Y. J. and Lwiza, K. M. M.: Factors driving bottom salinity variability in the Chesapeake Bay, Cont. Shelf Res., 28, 1352–1362, 2008.
Mantua, N. J. and Hare, S. R.: The Pacific Decadal Oscillation, J. Oceanogr., 58, 35–44, 2002.
Maraun, D. and Kurths, J.: Cross wavelet analysis: significance testing and pitfalls, Nonlin. Processes Geophys., 11, 505–514, doi:10.5194/npg-11-505-2004, 2004.
Maraun, D., Kurths, J., and Holschneider, M.: Nonstationary Gaussian processes in wavelet domain: synthesis, estimation, and significance testing, Phys. Rev. E, 75, 016707, doi:10.1103/PhysRevE.75.016707, 2007.
Meyers, S. D., Kelly, B. G., and O'Brien, J. J.: An introduction to wavelet analysis in oceanography and meteorology: with application to the dispersion of Yanai waves, Mon. Weather Rev., 121, 2858–2866, 1993.
Newman, M., Compo, G. P., and Alexander, M. A.: ENSO-forced variability of the Pacific Decadal Oscillation, J. Climate, 16, 3853–3857, 2003.
Olsen, J., Anderson, J. N., and Knudsen, M. F.: Variability of the North Atlantic Oscillation over the past 5,200 years, Nat. Geosci., 5, 808–812, 2012.
Torrence, C. and Compo, G. P.: A practical guide to wavelet analysis, B. Am. Meteorol. Soc., 79, 61–78, 1998.
Trenberth, K. E.: The definition of El Niño, B. Am. Meteorol. Soc., 78, 2771–2777, 1997.
Vautard, R., Yiou, P., and Ghil, M.: Singular-spectrum analysis: a toolkit for short, noisy chaotic signals, Physica D, 58, 95–126, 1992.

NPGD
2, 1227–1273, 2015

Velasco, V. M. and Mendoza, B.: Assessing the relationship between solar activity and some large scale climatic phenomena, Adv. Space Res., 42, 866–878, 2008.
Schulte, J. A., Duffy, C., and Najjar, R. G.: Geometric and topological approaches to significance testing in wavelet analysis, Nonlin. Processes Geophys., 22, 139–156, doi:10.5194/npg-22-139-2015, 2015.
Wang, C. and Picaut J.: Understanding ENSO physics – a review, in: Earth's climate: the ocean–atmosphere interaction, Geophys. Monogr. Ser., 147, American Geophysical Union, Washington, D.C., 21–48, 2004.
Weerahandi, S.: Exact Statistical Methods for Data Analysis, Springer, New York, 329 pp., 2003.
Whitney, M. M.: A study on river discharge and salinity variability in the Middle Atlantic Bight and Long Island Sound, Cont. Shelf Res., 30, 305–318, 2010.
Wilson, M., Meyers, S. D., and Luther, M. E.: Synoptic volumetric variations and flushing of the Tampa Bay estuary, Clim. Dynam., 42, 1587–1594, 2014.

NPGD
2, 1227–1273, 2015

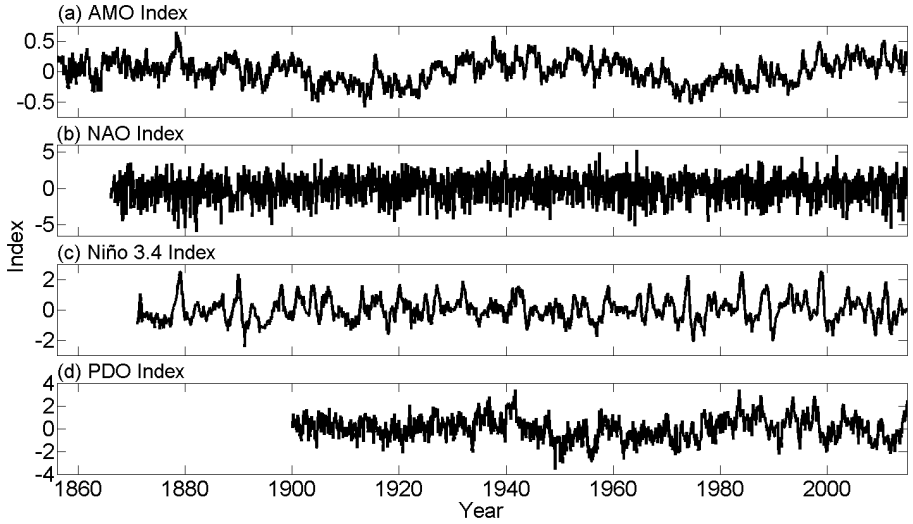

Figure 1. The monthly **(a)** AMO, **(b)** NAO, **(c)** Niño 3.4, and **(d)** PDO indices. Data sources are the Climate Prediction Center for the AMO index (http://www.esrl.noaa.gov/psd/data/climateindices/list/), National Center for Atmospheric Research for the NAO (https://climatedataguide.ucar.edu/data-type/climate-indices) and Niño 3.4 (http://www.cgd.ucar.edu/cas/catalog/climind/TNI_N34/) indices, and University of Washington for the PDO index (http://research.jisao.washington.edu/pdo/PDO.latest). The Niño 3.4 index was converted to monthly anomalies by subtracting off the mean annual cycle.

NPGD
2, 1227–1273, 2015

Cumulative areawise testing in wavelet analysis

J. A. Schulte

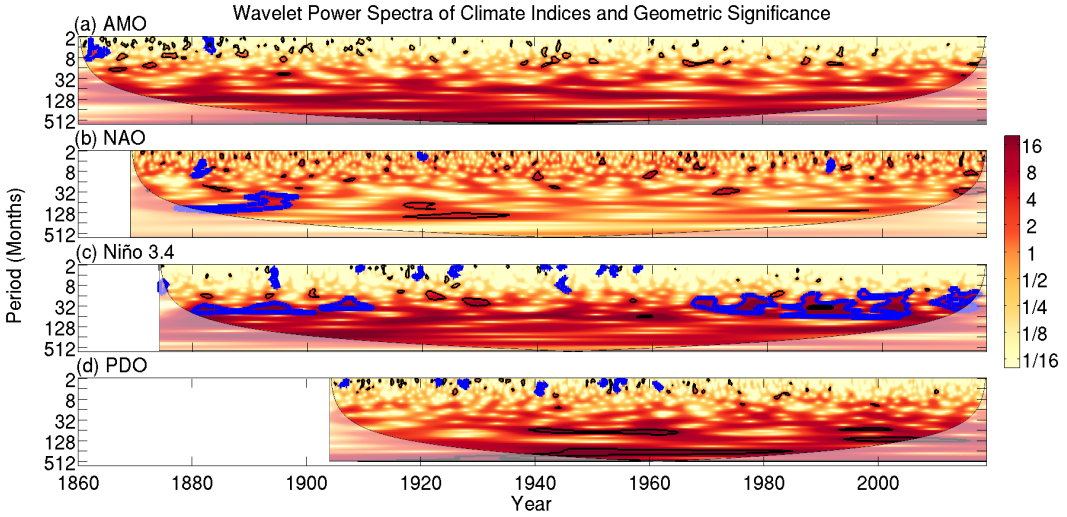

Figure 2. Wavelet power spectra of the **(a)** AMO, **(b)** NAO, **(c)** Niño 3.4, and **(d)** PDO indices. Thin black contours enclose regions of 5 % pointwise significance and thick blue contours indicate those patches that are geometrically significant at the 5 % level. Light shading represents the cone of influence (COI), the region in which edge effects cannot be ignored.

NPGD
2, 1227–1273, 2015

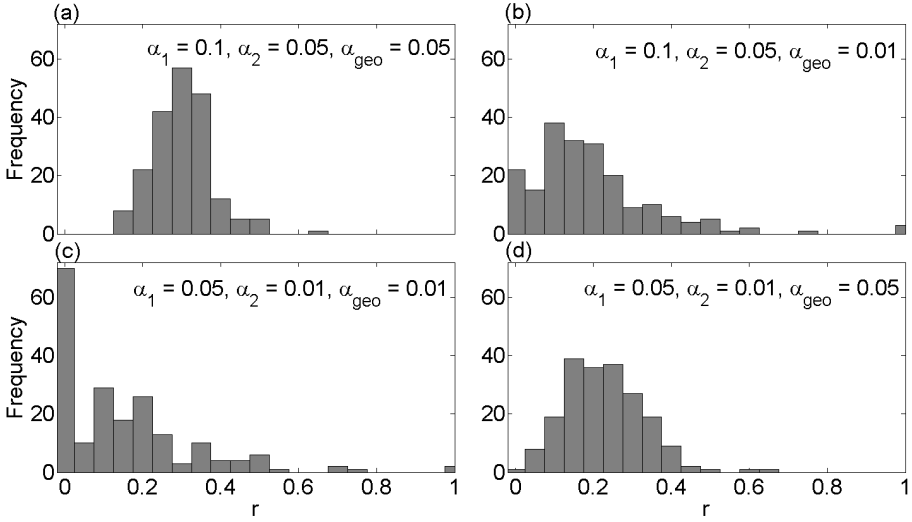

Figure 3. (a) A histogram of $r = \frac{N_{\alpha_1,\alpha_2}}{N_{\alpha_1}}$ for $\alpha_1 = 0.1$, $\alpha_2 = 0.05$, and $\alpha_{geo} = 0.05$ obtained from the generation of 300 wavelet power spectra of red-noise processes of length 1000 with lag-1 autocorrelation coefficients equal to 0.5. **(b)** Same as **(a)** but with $\alpha_{geo} = 0.01$. **(c)** Same as **(a)** but with $\alpha_1 = 0.05$ and $\alpha_2 = 0.01$, and $\alpha_{geo} = 0.01$. **(d)** Same as **(c)** but with $\alpha_{geo} = 0.05$.

NPGD
2, 1227–1273, 2015

Cumulative areawise testing in wavelet analysis

J. A. Schulte

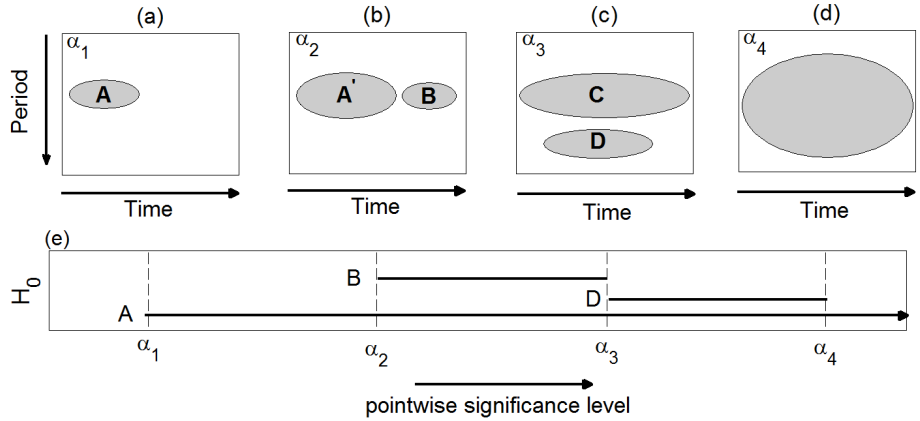

Figure 4. (a–d) The topological evolution of patches across four pointwise significance levels. **(e)** The barcode showing the birth and death of patches throughout the evolution process. Horizontal lines with arrows indicate those patches that never die, the so-called essential classes.

NPGD
2, 1227–1273, 2015

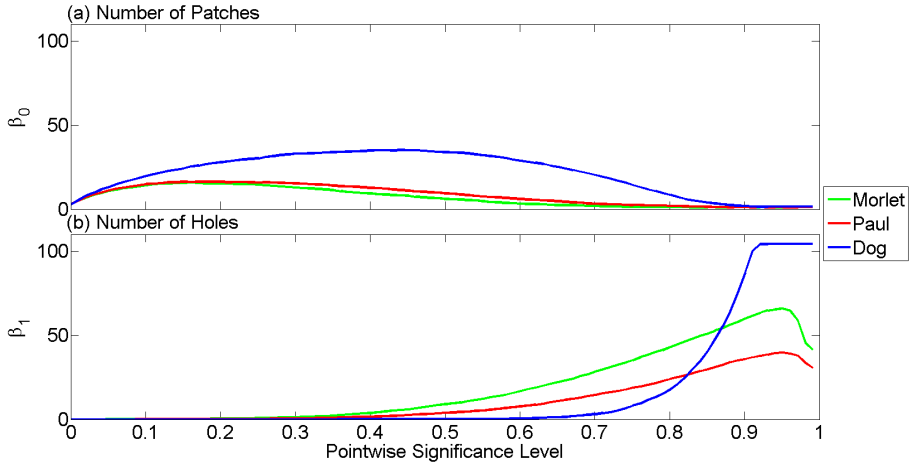

Figure 5. (a) Number of patches and **(b)** the number of holes for three analyzing wavelets as a function of α.

NPGD
2, 1227–1273, 2015

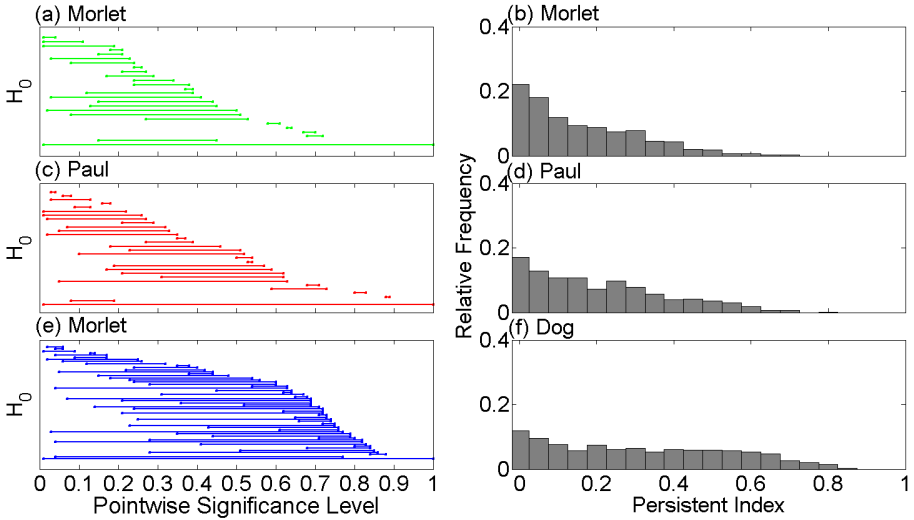

Figure 6. (a) Example barcodes for H_0 corresponding to a wavelet power spectrum obtained from the **(a)** Morlet, **(b)** Paul, and **(c)** Dog wavelets. A red-noise process with length 150 and lag-1 autocorrelation coefficient equal to 0.5 was used to create the barcodes. Distribution of persistence indices representing the lifetimes of patches for the **(b)** Morlet, **(d)** Paul, and **(f)** Dog wavelets. Distribution was obtained by generating 1000 wavelet power spectra of red-noise processes with lengths 500 and lag-1 autocorrelation coefficients equal to 0.5. Essential classes have been removed from the distributions.

NPGD
2, 1227–1273, 2015

Cumulative areawise testing in wavelet analysis

J. A. Schulte

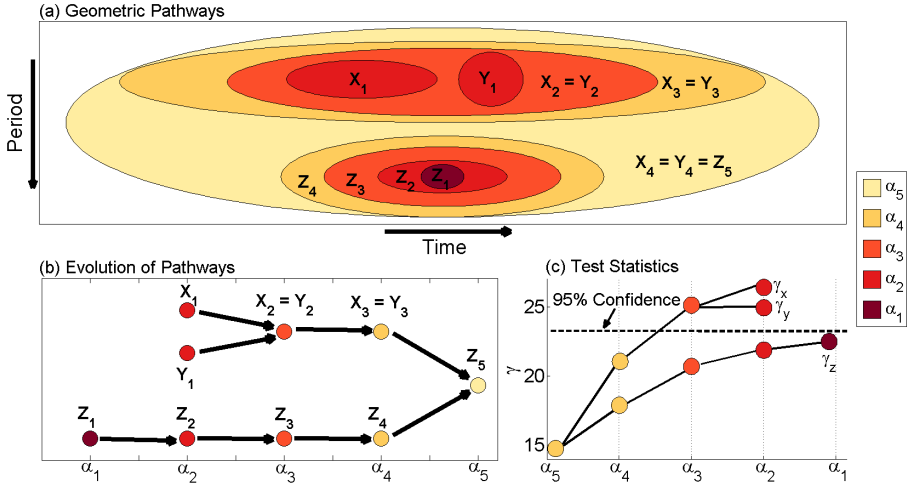

Figure 7. (a) Geometric pathway of three significance patches, X, Y, and Z in the interval $I = [\alpha_1, \alpha_5]$. **(b)** The geometric evolution of the pathways showing how Z_5 was created from the merging of X_3 and Z_4 as α changed from α_4 to α_5. **(c)** The cumulative areas at each step of pathway for each geometric pathway, where the summation begins at α_5, the initial point of the pathway, and the dotted line represents the critical level of cumulative areawise test.

NPGD
2, 1227–1273, 2015

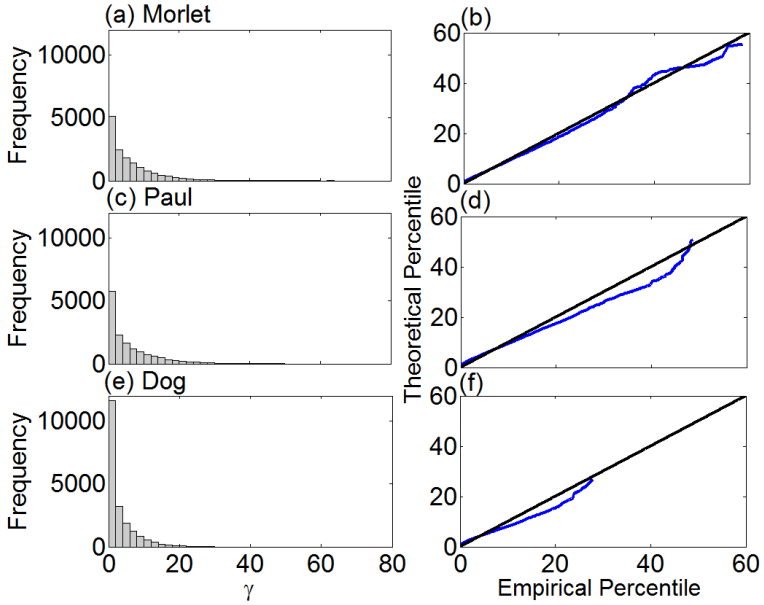

Figure 8. (a) Null distribution of γ for the Morlet wavelet obtained by generating 10 000 geometric pathways under the null hypothesis of red-noise, where the red-noise processes were of length 1000 and had lag-1 autocorrelation coefficients equal to 0.5. **(b)** Percentiles of a theoretical exponential distribution with mean 6.5 plotted as a function of the percentiles calculated from the distribution shown in **(a)**. **(c, d)** Same as **(a, b)** but for the Paul wavelet. **(e, f)** Same as **(a, b)** except for the Dog wavelet.

NPGD
2, 1227–1273, 2015

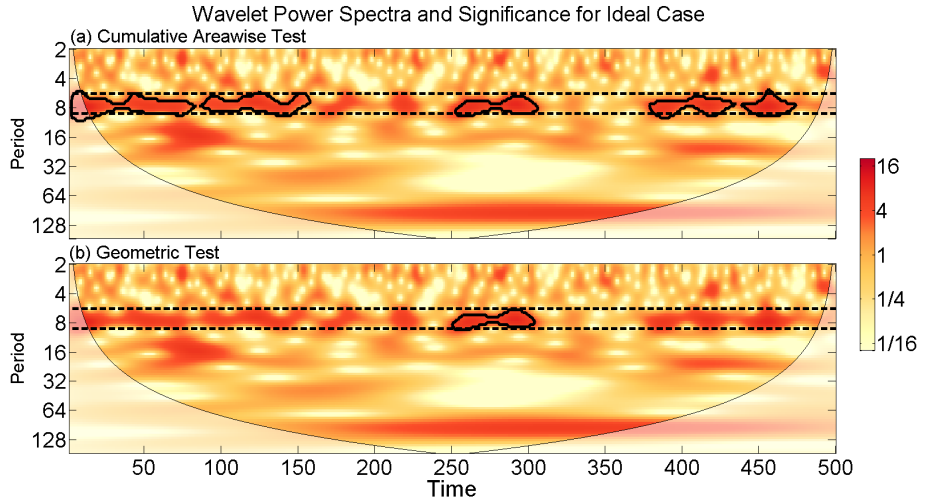

Figure 9. (a) Cumulative areawise test applied to a sinusoid with a frequency of 0.8 and amplitude equal to 0.8. Signal-to-noise ratio is 1.0. Contours represent patches that are elements of 5 % significant pathways. Dotted lines represent the upper and lower boundaries of a theoretical patch obtained by generating the wavelet power spectrum of a pure sine wave and calculating the width of the patch at $t = 250$. **(b)** Same as **(a)** except for the geometric test with $\alpha = 0.05$ and $\alpha_{geo} = 0.05$. Contours represent patches that are geometrically significant.

NPGD
2, 1227–1273, 2015

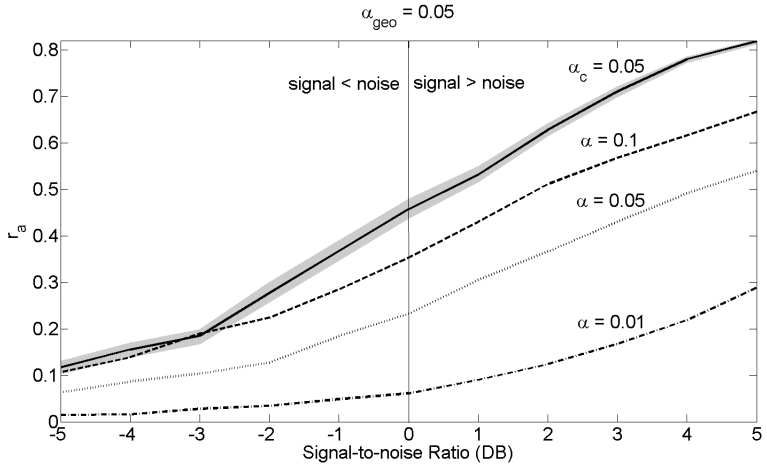

Figure 10. The ensemble mean r_a as a function of the signal-to-noise to ratio for the areawise test with $\alpha_c = 0.05$ and the geometric test with $\alpha_{geo} = 0.05$. Gray shading represents the 95 % confidence interval and all means for the geometric test are significantly different at the 5 % level from the means for the areawise test except for those corresponding to the $\alpha = 0.01$ curve for signal-to-noise ratios less than −1.5. The confidence intervals and statistical significance were obtained by the bootstrap method (Efron, 1979). The data for each signal-to-noise-ratio were sampled with replacement 1000 times to generate a distribution of bootstrap replicates, from which 95 % confidence intervals were obtained. Two ensemble means were said to be significantly different at the 5 % level if their 95 % confidence intervals did not intersect.

NPGD
2, 1227–1273, 2015

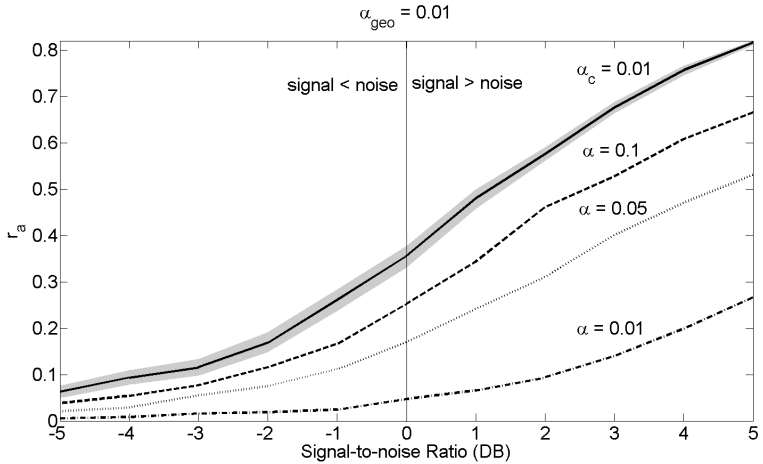

Figure 11. Same as Fig. 10 except with α_c = 0.01 and α_{geo} = 0.01. All means for the geometric test are significantly different at the 5 % level from the means for the areawise test.

NPGD
2, 1227–1273, 2015

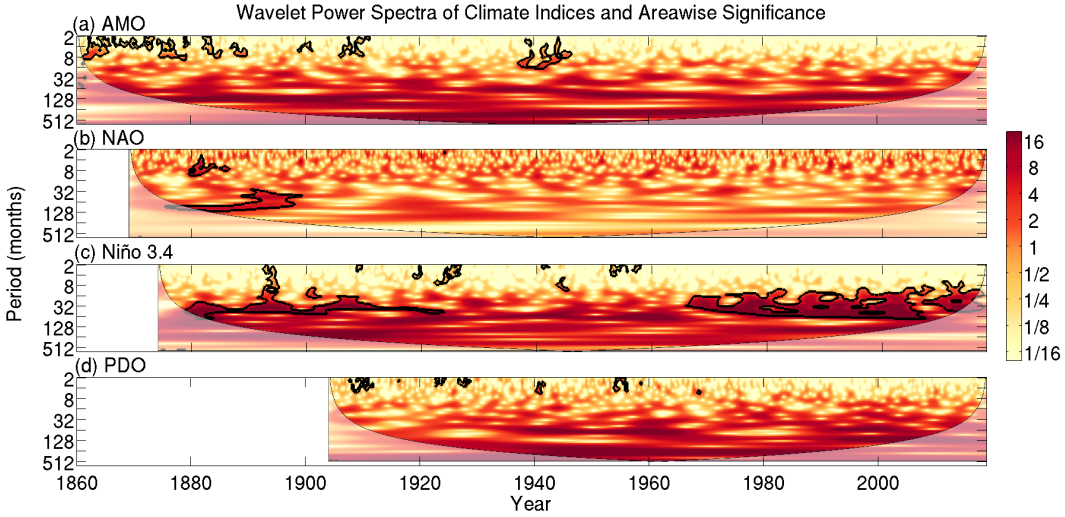

Figure 12. (a) The application of the cumulative areawise test to the AMO index **(a)** the NAO index, **(c)** the Niño 3.4 index, and **(d)** the PDO index. α_c was set to 0.01 in all cases and contours enclose regions of 1 % areawise significance.

NPGD
2, 1227–1273, 2015

Cumulative areawise testing in wavelet analysis

J. A. Schulte