

Coherence Modified for Sensitivity to Relative Phase of Real Band-Limited Time Series

William Menke

Lamont-Doherty Earth Observatory of Columbia University, Palisades NY 10964 USA

MENKE@LDEO.COLUMBIA.EDU

Abstract. As is well known, coherence does not distinguish the relative phase of a pair of real, sinusoidal time series; the coherence between them is always unity. This behavior can limit the applicability of coherence analysis in the special case where the time series are band-limited (nearly-monochromatic) and where sensitivity to phase differences is advantageous. We propose a simple modification to the usual formula for coherence in which the cross-spectrum is replaced by its real part. The resulting quantity behaves similarly to coherence, except that it is sensitive to relative phase when the signals being compared are strongly band-limited. Furthermore, it has a useful interpretation in terms of the zero-lag cross-correlation of real band-passed versions of the time series.

Introduction

The purpose of time series analysis is to extract useful knowledge from time series. The word *knowledge* is used advisedly, and implies that the data are being distilled down into a more useful form that is capable of addressing a specific question posed by an analyst. The choice of time series analysis technique should be guided by this principle. One analyzes time series in a way designed to best extract knowledge from them. One should always be willing to adapt an analysis method to achieve this goal.

The issue considered here is how best to quantify the similarity between time series that are 1) real (as contrasted to complex) and 2) band-limited (in the sense of being nearly

monochromatic). Such time series constitute an important special case, because most natural phenomena are described using real numbers and many are dominated by a single period of oscillation. The daily period often contributes strongly to physiological and meteorological signals, the annual period to environmental and climatic signals, the precessional period (25.7 ka) (e.g. Olsen & Kent, 1995) to sedimentary and paleontological signals, and so forth. An important property of nearly-monochromatic signals is their relative phase. Whether two time series that are in-phase (as in Fig. 1A) or out-of-phase (as in Fig. 1B) may be important, for example, from the perspective of an analyst trying to unravel the dynamics of the underlying causative processes.




Fig. 1. A) Two nearly-monochromatic time series (black and red curves) with the relative phase, $\varphi = 0$. B) Two nearly-monochromatic time series (black and red curves) with relative phase, $\varphi = \pi/2$.

Traditional coherence analysis (e.g. Bendat & Pierson, 2010) has very limited application in this case, because of the well-known insensitivity of coherence to relative phase. The

coherence of two sinusoidal time series of the same period is always unity, irrespective of their relative phase. Simply put, coherence does not distinguish a sine from a cosine. Given the general usefulness of coherence in other settings, it is well to ask why it “fails” in this special case and whether it can be modified to produce what may, in some circumstances, be a more useful measure of similarity.

When asking why any quantity encountered in time series analysis, such as coherence, behaves in a certain way, one must contend with the fact that most, if not all, such quantities can be derived from several different perspectives. Any answer will probably make sense only from one of these points of view. Consider, for example, the estimated mean of a time series. This deceptively simple quantity can be understood, alternately, as arising through the minimizing of *error* (a deterministic derivation) (e.g. Lawson & Hanson, 1974), or through the maximizing of *likelihood* (a probabilistic derivation) (e.g. Millar, 2011) or through the maximization of *importance* (an informational derivation) (Menke 2012), to name just a few. The answer to a question concerning the mean, say for example, whether it should always be bounded by the smallest and largest datum, will necessarily refer to one of these perspectives. The same is true for coherence. We adopt here a deterministic perspective:

The coherence between two time series, at frequency, ω_0 , is closely-related to the zero-lag cross-correlation of band-passed versions of those time series, where the band-pass filter is one-sided and has center-frequency, ω_0 . In fact, the former is merely a normalized and squared version of the latter.

This is but one perspective among many, but one we find helpful because it brings out a relationship to the cross-correlation, another quantity useful in assessing the similarity between

two time series. Cross-correlation is defined in the time-domain, as contrasted to coherence, which is defined in the frequency-domain, so the link provides complimentary information.

The appearance of a *one-sided* filter (Fig. 2A), may seem counter-intuitive, because such filters are almost never used in practice, or at least not when the data are real, for they turn a real time series into a complex one. All the band-pass filters that an analyst would commonly use are two-sided (Fig. 2B), and so have real output. The reason for its appearance here is that the usual definition of coherence is *completely general*. It does not presume that the signals being compared are real, and so builds in the possibility that negative and positive frequency components of the time series may behave completely differently from one another. This is contrast to real time series, where they are complex conjugate pairs. However, in being general, it cannot exploit an important property of real signals: that sines and cosines are distinguishable from one another. As we show below, substituting a two-sided filter produces a version of coherence that distinguishes sines from cosines; that is, one that is sensitive to the relative phase of band-limited signals.

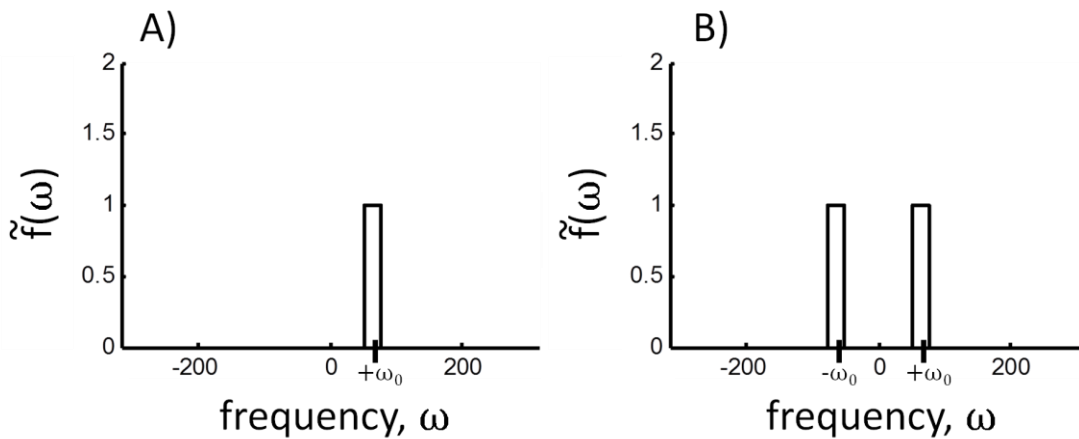


Fig. 2. A) Fourier transform of a one-sided band-pass filter, consisting of a single boxcar function centered at frequency, ω_0 , and with width, $2\Delta\omega$. B) Fourier transform of a two-sided band-pass filter, also centered at frequency, ω_0 .

Coherence-Like Measure of Similarity Based of Cross-Correlation.

The problem we consider is how to quantify the similarity of two real, transient time series, $u(t)$ and $v(t)$, in the vicinity of a specified frequency, ω_0 . The strategy we adopt is to band-pass filter these time series and then to compute their zero-lag cross-correlation. The filter selects out frequencies near ω_0 and the cross-correlation quantifies similarity, since it attains its largest value when $u(t) = v(t)$ (ignoring, for the moment, the issue of normalization). We denote the filtered time series as, $f(t) * u(t)$ and $f(t) * v(t)$ where the symbol $*$ denotes convolution. We require the filtered time series to be purely real, so that the filter, $f(t)$, has a two-sided Fourier transform with the symmetry, $\tilde{f}(\omega) = \tilde{f}^*(-\omega)$, where the tilde denotes Fourier transformation and the asterisk denotes complex conjugation. We choose a filter with a purely-real Fourier transform, built from two unit-amplitude boxcar functions, one centered at $-\omega_0$ and the other at $+\omega_0$, each of width $2\Delta\omega$. This filter does not affect Fourier components within the pass-band and completely rejects those outside of it.

The cross-correlation of two real time series is defined as:

$$c(t) = u \star v = \int_{-\infty}^{+\infty} u(\tau) v(t + \tau) d\tau \tag{1}$$

and has zero-lag value:

$$c(t = 0) = (u \star v)_{t=0} = \int_{-\infty}^{+\infty} u(\tau) v(\tau) d\tau \quad (2)$$

The zero-lag cross-correlation of the filtered time series is:

$$\begin{aligned} c(t, \omega_0, \Delta\omega) &= \{f(t, \omega_0, \Delta\omega) \star u(t)\} \star \{f(t, \omega_0, \Delta\omega) \star v(t)\} \\ &= f(-t, \omega_0, \Delta\omega) \star f(t, \omega_0, \Delta\omega) \star u(-t) \star v(t) \end{aligned} \quad (3)$$

Here, we have relied on the identity, $a(t) \star b(t) = a(t) \star b(-t)$. At zero lag, the cross-correlation is proportional the integral of its Fourier transform, $\tilde{c}(\omega)$:

$$c(t = 0) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} \tilde{c}(\omega) \exp(i\omega t)|_{t=0} d\omega = \frac{1}{2\pi} \int_{-\infty}^{+\infty} c(\omega) d\omega \quad (4)$$

Inserting (3) into (4) and using the rule that the Fourier transform of a convolution is the product of the transforms and the rule that the transform of $a^*(-t)$ is $\tilde{a}^*(\omega)$, yields:

$$\begin{aligned} c(t = 0, \omega_0, \Delta\omega) &= \frac{1}{2\pi} \int_{-\infty}^{+\infty} \tilde{f}(\omega, \omega_0, \Delta\omega) \tilde{f}(\omega, \omega_0, \Delta\omega) \tilde{u}^*(\omega) \tilde{v}(\omega) d\omega \\ &\approx \frac{1}{2\pi} \int_{-\omega_0-\Delta\omega}^{-\omega_0+\Delta\omega} \tilde{u}^*(\omega) \tilde{v}(\omega) d\omega + \frac{1}{2\pi} \int_{+\omega_0-\Delta\omega}^{+\omega_0+\Delta\omega} \tilde{u}^*(\omega) \tilde{v}(\omega) d\omega \\ &= \frac{1}{\pi} \int_{\omega_0-\Delta\omega}^{\omega_0+\Delta\omega} \text{Re} \{ \tilde{u}^*(\omega) \tilde{v}(\omega) \} d\omega = \frac{2\Delta\omega}{\pi} \overline{\text{Re}\{\tilde{u}^*(\omega_0) \tilde{v}(\omega_0)\}} \end{aligned} \quad (5)$$

Here $\overline{\tilde{a}(\omega_0)}$ denotes the mean value of $a(\omega)$ in the frequency band, $\omega_0 \pm \Delta\omega$. Note that $c(t = 0, \omega_0, \Delta\omega)$ is defined for $\omega_0 > 0$, only. The quantity $\tilde{u}^*(\omega) \tilde{v}(\omega)$ is the cross-spectrum. Thus, the zero-lag cross-correlation of the real band-pass filtered time series depends upon the average

value of the real part of their cross-spectrum in that pass-band. The amplitude of $c(t = 0, \omega_0, \Delta\omega)$ depends on the amplitude of two time series, as well as upon their degree of similarity. We remove this dependence by normalizing by the energy E_u and E_v in the two time series, defined as:

$$E_u = (u \star u)_{t=0} = \int_{-\infty}^{+\infty} u^2(\tau) d\tau \text{ and } E_v = (v \star v)_{t=0} = \int_{-\infty}^{+\infty} v^2(\tau) d\tau \quad (6)$$

The normalized measure of similarity, say S , is:

$$S = \frac{c(t = 0, \omega_0, \Delta\omega)}{E_u^{1/2} E_v^{1/2}} = \frac{\overline{\text{Re}\{\tilde{u}^*(\omega_0)\tilde{v}(\omega_0)\}}}{|\tilde{u}(\omega_0)| |\tilde{v}(\omega_0)|} \text{ with } 0 \leq \omega_0 < +\infty \quad (7)$$

Note that the quantity, S^2 , which we nickname here *similarity*, varies between zero and unity. It has almost exactly the functional form of the quantity called *coherence*, except for the taking of the real part. The imaginary part cancelled from (5) precisely because the time series are real and the filter is two-sided.

Coherence Related to Zero-Lag Cross-Correlation

As asserted in the Introduction, the usual formula for coherence can be obtained simply by switching to a one-sided filter, a single unit step function of width $2\Delta\omega$ centered at frequency ω_0 (where $-\infty < \omega_0 < +\infty$). The filtered time series $f \star u$ and $f \star v$ are complex, so that one must define a cross-correlation appropriate for complex signals; that is, replace $u(\tau)$ with $u^*(\tau)$ in (1). These modifications lead to a version of (7) that is exactly the usual formula for the coherence:

$$C^2(\omega_0, \Delta\omega) = \frac{c^2(t=0, \omega_0, \Delta\omega)}{E_u E_v} = \frac{|\tilde{u}^*(\omega_0)\tilde{v}(\omega_0)|^2}{|\tilde{u}(\omega_0)|^2 |\tilde{v}(\omega_0)|^2} \quad \text{with } -\infty < \omega_0 < +\infty \quad (8)$$

As an aside, we note that our derivations of $C^2(\omega_0)$ and $S^2(\omega_0)$ hide an inconsistency in the interpretation of $|\tilde{u}(\omega_0)|^2$ as the power in the time series $u(t)$ at frequency, ω_0 . It represents power for a complex time series but only half the power for a real one, owing to the different intervals over which frequency, ω_0 , is defined. This factor of two compensates for the apparent loss of power when the real part is taken in (5).

Similarity and Coherence of Real Band-Limited Signals

Suppose that time series $u(t)$ and $v(t)$ are monochromatic, with equal frequency, ω_0 , but with different amplitudes, u_0 and v_0 , and relative phase, φ :

$$u(t) = u_0 \sin(\omega_0 t) \quad \text{and} \quad v(t) = v_0 \sin(\omega_0 t - \varphi) \quad (9)$$

The similarity, $S^2(\omega_0)$, is most easily calculated using its time-domain definition. Taking, without loss of generality, the window of observation to be $0 < \tau < 2\pi$, we have:

$$E_u = u_0^2 \int_0^{2\pi} \sin^2(\omega_0 \tau) d\tau = \frac{\omega_0 u_0^2}{2} \quad \text{and} \quad E_v = v_0^2 \int_0^{2\pi} \sin^2(\omega_0 \tau - \varphi) d\tau = \frac{\omega_0 v_0^2}{2}$$

$$\text{and } c(t=0) = u_0 v_0 \int_0^{2\pi} \sin(\omega_0 t) \sin(\omega_0 t - \varphi) d\tau = \frac{\omega_0 u_0 v_0}{2} \cos(\varphi)$$

$$\text{so } S^2 = \frac{c^2(t=0)}{E_u E_v} = \cos^2(\varphi)$$

(10)

Thus, S^2 is unity when the two sinusoids are in-phase ($\varphi = 0$) and declines monotonically to zero when they are out-of-phase ($\varphi = \pi/2$).

The coherence, $C^2(\omega_0)$, is calculated by recognizing that a sine function is built up of two complex exponentials of frequency $+\omega_0$ and $-\omega_0$ and that the one-sided filter selects only the one with positive frequency:

$$f(t) * u(t) = U_0 \exp(i\omega_0 t) \quad \text{and} \quad f(t) * v(t) = V_0 \exp(i\omega_0 t)$$

$$\text{with } U_0 = \frac{u_0}{2i} \quad \text{and} \quad V_0 = \left\{ \frac{u_0}{2i} \cos(\varphi) - \frac{v_0}{2i} \sin(\varphi) \right\}$$
(11)

We then find:

$$E_u = U_0^2 \int_0^{2\pi} \exp(-i\omega_0 t) \exp(i\omega_0 t) d\tau = 2\pi U_0^2 \quad \text{and} \quad E_v = 2\pi V_0^2 \quad \text{and} \quad c(t=0) = 2\pi U_0 V_0$$

$$\text{so } C^2 = \frac{c^2(t=0)}{E_u E_v} = 1$$
(12)

Thus, C^2 is unity irrespective of the relative phase of the two sinusoids. This behavior is a consequence of the one-sided filter, which turns both $\sin(\omega_0 t)$ and $\cos(\omega_0 t)$ into functions proportional to the same complex exponential, $\exp(i\omega_0 t)$.

Examples and Conclusions

We consider the example of a sequence of nearly-monochromatic wavelets, formed by taking the product of a phase-shifted sinusoid of frequency, ω_0 , and a Normal envelope function of half-width, σ :

$$\sin(\omega_0 t - \varphi) \exp(-(t - t_0)^2 / (2\sigma^2))$$

(13)

Figure 3A-C illustrates pairs of these wavelets with different phase relationships. Note that the wavelets are not merely time-shifted versions of one another, since the position of the zeros crossings of the sinusoid (parameterized by φ) can and the position of the center of the envelope (parameterized by t_0) can be independently varied. One might imagine a time series analysis scenario where $u(t)$ represents the external forcing applied to some dynamical system, and $v(t)$ represents the response. In such a context, the distinction between these different wavelet shapes is important, say for detecting whether or not some anticipated interaction has occurred. In this case, the similarity, $S^2(\omega_0)$ (red curves in Figure 3D-E) is a more useful quantity than the coherence, $C^2(\omega_0)$ (black curves), since it varies strongly with the phase-relationships, whereas coherence does not.

We have not performed an exhaustive analysis of the differences between $S^2(\omega_0)$ and $C^2(\omega_0)$, when they are applied to broad-band signals. The key difference is the effect of the taking of the real part:

$$\begin{aligned}
 S^2 &\propto (\overline{\tilde{u}_R \tilde{v}_R} + \overline{\tilde{u}_I \tilde{v}_I})^2 \\
 C^2 &\propto (\overline{\tilde{u}_R \tilde{v}_R} + \overline{\tilde{u}_I \tilde{v}_I})^2 + (\overline{\tilde{u}_R \tilde{v}_I} - \overline{\tilde{u}_I \tilde{v}_R})^2
 \end{aligned}
 \tag{14}$$

where the Fourier transforms are written in terms of their real and imaginary parts, $\tilde{u} = \tilde{u}_R + i\tilde{u}_I$ and $\tilde{v} = \tilde{v}_R + i\tilde{v}_I$. Since S^2 and C^2 differ by a manifestly positive amount, we are guaranteed that $C^2 \geq S^2$. However, without further specification of the behavior of \tilde{u} and \tilde{v} , no further characterization is possible. In the special case where both time series contain a common function $w(t)$, so that $u(t) = w(t) + x(t)$ and $v(t) = w(t) + y(t)$ and where $w(t)$, $x(t)$ and $y(t)$ are all broad-band, we find:

$$S \propto (\overline{\tilde{w}_R \tilde{w}_R} + \overline{\tilde{w}_I \tilde{w}_I}) + \text{crossterms like } \overline{\tilde{x}_R \tilde{y}_R}$$

$$C \propto (\overline{\tilde{w}_R \tilde{w}_R} + \overline{\tilde{w}_I \tilde{w}_I}) + \text{crossterms like } \overline{\tilde{x}_R \tilde{y}_I}$$

(15)

We might expect in the case that $C^2 \approx S^2$, since the crossterms are averages of functions that oscillate around zero and therefore likely to be small. Numerical tests (Figure 4) support this idea, at least for non-transient broad-band time series with a moderate degree of correlation.

In summary, we recommend this simple modification of coherence in cases where the time series that are being compared are narrow-band and where phase relationships between them are considered important. For pure sinusoids differing by phase, φ , it obeys the rule $S^2 = \cos^2 \varphi$, that is, similarity monotonically decreases from unity, when $\varphi=0$, to zero, when $\varphi = \pi/2$. In other respects, it behaves very similarly to coherence. Finally, it has a very intuitive time-domain interpretation: $S(\omega_0)$ gives you exactly what you would get if you normalized each time series by the square-root of its energy, band-pass filtered each with a two-sided boxcar filter centered around frequency, ω_0 , and computed their zero-lag cross-correlation.

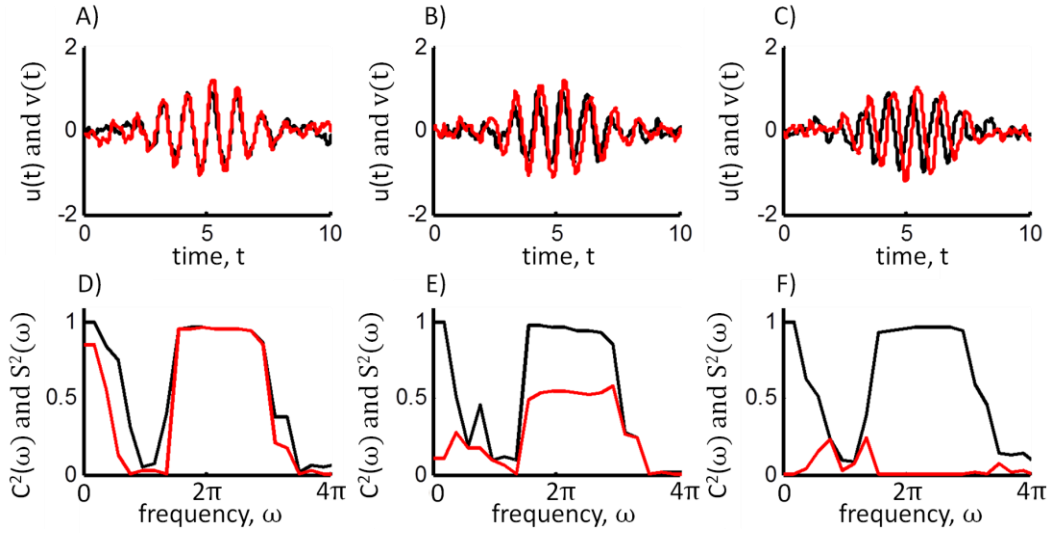


Fig. 3. (A-C) Sequence of three pairs of nearly-monochromatic time series with frequency $\omega_0 \approx 2\pi$ and with relative phase of, $\varphi = 0$, $\varphi = \pi/4$. and $\varphi = \pi/2$, respectively. (D-F) Corresponding coherence, $C^2(\omega)$, and similarity (black curve) and $S^2(\omega)$ (red curve). Note that $C^2(\omega)$, is approximately unity for all three cases, whereas $S^2(\omega)$ decreases as the relative phase increases. In this example $\Delta\omega$ is set to $\omega_0/4$.

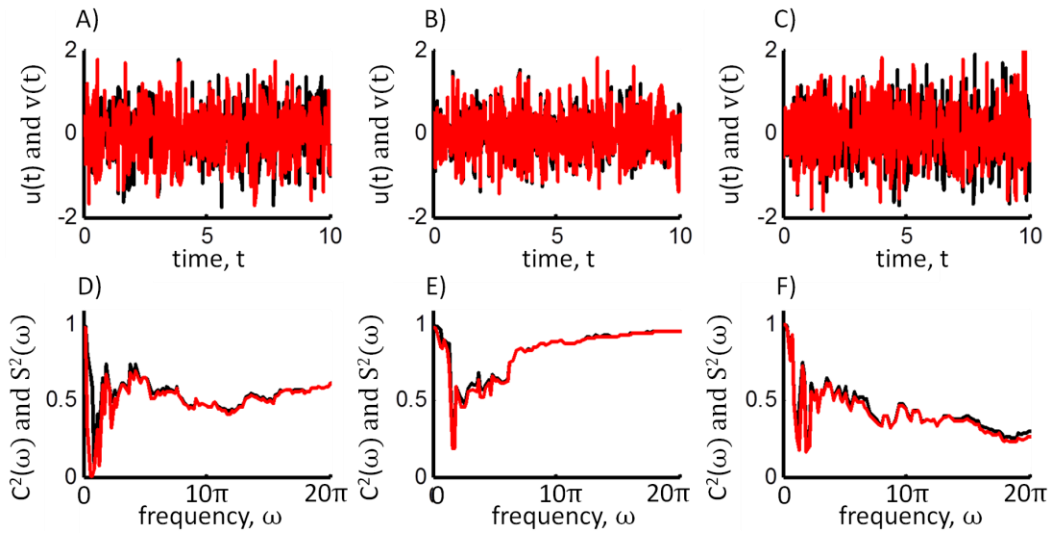


Fig. 4. (A-C) Sequence of three pairs of broad-band time series with an approximately (A) frequency-independent coherence, (B) coherence that increases with frequency and (C) coherence that decreases with frequency. (D-F) Corresponding coherence, $C^2(\omega)$ (black curve), and similarity and $S^2(\omega)$ (red curve), which are approximately equal, although also obeying the rule $C^2 \geq S^2$. In this example $\Delta\omega$ is set to $\omega_0/4$.

References

- Bendat, J.S. and A.G. Piersol, Random Data, Analysis and Measurement Procedures,, John Wiley and Sons, 2010, 604pp.
- Lawson, CL and R.J. and Hanson, Solving Least Squares Problems, Prentice-Hall,, 1974, 340 pp.
- Menke, W., Geophysical Data Analysis: Discrete Inverse Theory, MATLAB Edition, Elsevier, 2012, 293pp.
- Millar, R.B., Maximum Likelihood Estimation and Inference: With Examples in R, SAS and ADMB, John Wiley and Sons, 2011.
- Olsen, P.E. and D.V. Kent, Milankovitch climate forcing in the tropics of Pangea during the Late Triassic, Paleoceanography, Paleoclimatology, Paleoecology 122, 1-26, 1996.