A Biased Look at Phase Locking:
Brief Critical Review and Proposed Remedy

Christopher K. Kovach*, *Department of Neurosurgery, University of Iowa Hospitals and Clinics, Iowa City, IA

Abstract—A number of popular measures of dependence between pairs of band-limited signals rely on analytic phase. A common misconception is that the dependence revealed by these measures must be specific to the spectral range of the filtered input signals. Implicitly or explicitly, obtaining analytic phase involves normalizing the signal by its own envelope, which is a nonlinear operation that introduces broad spectral leakage. We review how this generates bias and complicates the interpretation of commonly used measures of phase locking. A specific example of this effect may create spurious phase locking as a consequence of nonzero circular mean in the phase of input signals, which can be viewed as spectral leakage to 0 Hz. Corrections for this problem which recenter or uniformize the distribution of phase may fail when the amplitudes of the compared signals are correlated. To address the more general problem of spectral bias, a novel measure of phase locking is proposed, the amplitude-weighted phase locking value (awPLV). This measure is closely related to coherence, but it removes ambiguities of interpretation that detract from the latter.

Index Terms—coherence; phase locking; cross-frequency coupling; spectral analysis

I. INTRODUCTION

MEASURES of phase locking are rapidly becoming a standard tool in the analysis of biophysical signals, especially so in applications to neurophysiology, where they are used to study interactions between anatomically separated neural populations [1]–[3] as well as cross-frequency coupling within responses from the same population [4], [5]. Their burst in popularity has been followed by a more gradually dawning awareness of the many pitfalls that accompany their use [6]–[16]. For example, recent entries to the body of cautionary literature [17]–[19] point out how non-uniformity of the distribution of phase in the input signals may create spurious results in many of the most commonly used indices of phase locking. This problem may be identified as one of a more general family of biases related to the spectral interpretation of phase.

Viewed from the spectral domain, it becomes clear that all members of this family share the same culprit: the transformation through which phase is obtained, which divides an analytic signal by its own envelope. The nonlinearity of this transformation induces spectral leakage across a broad range of frequencies, and the resulting unit-amplitude representation of phase is in general not a bandlimited function [20], [21]. Moreover, consequences of this are not confined to measures that explicitly invoke envelope normalization, but also extend to those that extract phase from the argument of the analytic

Corresponding author: C.K. Kovach

signal, for which one may consider normalization as implicit. The measures in question include phase-locking values (PLV), also known as phase coherence [2], phase-amplitude coupling [4] and information-theoretic measures derived from the empirical distribution of phase angles [5], along with others that correct for small-sample bias [22].

One class of measure is immune to spectral biases of this type: those obtained directly from signal cross spectra, of which coherence is the prime example. By avoiding amplitude normalization, coherence circumvents the underlying spectral distortions. However, the ambiguity of its interpretation, which depends on both phase locking and amplitude correlation, has dampened its popularity, motivating the aforementioned alternatives that discard amplitude [2].

After reviewing the origin of spectral leakage in analytic phase, another alternative is considered, the amplitude-weighted phase locking value (awPLV), which combines some of the favorable qualities of coherence with those of PLV. Unlike coherence, awPLV always yields a value of 1 for perfectly phase-synchronized signals (meaning here two signals whose phase difference remains constant over time), regardless of amplitude correlations. It improves on PLV and other phase-only measures by avoiding the spectral distortion responsible for the aforementioned biases. It also addresses another potential drawback of PLV, which is that PLV weights all samples equally, including those in which the signal amplitude approaches zero, for which phase is either poorly defined or likely to be dominated by background noise. Like coherence, awPLV weights the phase of each sample according to the product of the input signal amplitudes; its numerator therefore contains a cross-spectral estimate. It diverges from coherence by normalizing with the sum of amplitude products rather than the product of separate root-mean-square amplitudes. Weighting this way allows the resulting measure to be viewed, like PLV, as an average of unit-length phase vectors, but a weighted average.

A. Nonlinearity and Phase

Throughout the following discussion, spectral broadening related to amplitude normalization will be described as "spectral leakage," and its effect as a form of "bias." Some readers may rightly object to this terminology: spectral leakage bias normally describes a set of artifacts extrinsic to the signal, related to multiplying the signal with an externally determined window function, not, as with amplitude normalization, a "window" pulled from the signal itself. Our use (or rather abuse) of this terminology calls for more justification. It should

be emphasized at the outset that spectral broadening from amplitude normalization is neither an artifact nor inherently a flaw in conventional measures of phase. Though phase synchronization may arise through linear interactions, "pure" phase-phase interactions, which are not affected by signal amplitude, are nonlinear, so it would be a mistake to assume that the spectral constraints of linear interactions must apply generally to phase locking. In many settings, this nonlinearity is what makes phase interesting [23]. On the other hand, it is very difficult to resist the urge to think of phase as a spectrally localized quantity, given that the underlying signals are often band-limited and the rate of change of phase defines the instantaneous frequency.

A detailed consideration of the spectral interpretation of instantaneous phase is not the primary goal of the present work, however. A large body of work already addresses this topic [20], [24]–[27], which is also closely related to the longstanding question of the relationship between instantaneous frequency and and the signal spectrum [21], [28], [29]. The present aim is instead to describe conditions under which phase locking might be safely viewed as a measure of frequency-specific dependence between pairs of signals, in the same way as measures of linear dependence derived from the cross-spectrum. For this purpose, spectral broadening in analytic phase will be treated as a form of spectral leakage bias, although readers who prefer more nuanced terminology might wish to replace "bias" with "misinterpretation of the spectral nature of phase locking."

B. Organization

Sections II-C and III are intended as a tutorial introduction. Section IV presents the basic problem of bias related to non-uniform phase distributions, while section V explains how and why previously described solutions to this problem may fail in the presence of correlated amplitude modulation. Section VI considers some methods for overcoming the problem, with emphasis on the awPLV.

II. MEASURES OF DEPENDENCE IN COMPLEX SIGNALS

A. Complex and Analytic Signal Representations

Most commonly used measures of phase locking are derived from complex signal representations of the type obtained with a combination of band-limiting filter and Hilbert transform or related procedures such as a continuous wavelet transform [30], [31], short-time Fourier transform [32] or complex demodulation [33]–[35]. In all cases, the resulting complex signal can be understood as the outcome of a band-limiting filter applied to the original real-valued signal.

It is common to think of the result as analytic, meaning that it contains no negative frequencies, but this restriction is not essential. The important feature is rather that the filter be of a bandpass type, attenuating energy outside a single spectral band. Because its spectrum must be symmetric about 0 Hz, a real-valued signal fulfills this condition only if it is lowpass filtered; for this reason, the result is in general complex-valued, even when not strictly analytic. In the following discussion, the term "analytic signal" will be used loosely to refer to any

complex-valued signal obtained with a single-band filter. The conventional procedure combines an ordinary bandpass filter, g, with a Hilbert-transform-approximating quadrature filter, which suppresses energy in the negative band, so that the complex-valued result is the filtered signal in the real part and its (approximate) Hilbert transform in the imaginary part:

$$X(t) = g * x(t) + i\mathscr{H}\{g * x(t)\} \tag{1}$$

In all cases, the result may be understood as the outcome of a simple band-limiting filter, which will therefore contain no spectral component lacking in the original signal.

The same nonlinearity of phase that complicates its spectral interpretation also makes its meaning unclear when the signal is the sum of many components. One purpose of bandpass filtering in this setting is to isolate components in the hope of obtaining an interpretable result. While bandpass filtering is the simplest and most common approach the problem, it is by no means the only one [36], [37]. To avoid questions tangential to the present aim it will therefore be assumed that the signal of interest, x, falls within the passband of g, so that $g * x(t) \approx x(t)$, except where otherwise specified. Finally, it is very often useful to represent the analytic signal in a form that makes phase explicit, as the product of a positive real-valued envelope function $A(t)$ and a complex unit-magnitude valued phase:

$$X(t) = A_x(t)e^{i\phi_x(t)} \tag{2}$$

with $A_x = |X(t)|$ and $\phi_x = \arg(X)$.

B. Cross Spectra and Coherency

Instantaneous linear relationships between complex signals are quantified in the same way as for real-valued signals through measures derived from the second moment at zero lag:

$$R_{XY}(t) = \mathrm{E}\left[X(t)Y^*(t)\right] \tag{3}$$

When X and Y are stationary, this expectation assumes a constant value, which may then be estimated by integrating over time:

$$\hat{S}_{XY} = \frac{1}{T}\int_T X(t)Y^*(t)\,dt = \frac{1}{T}\int_T A_x A_y e^{i(\phi_x - \phi_y)}\,dt \tag{4}$$

Because the analysis filter in the present case, g introduces timing uncertainty, the relationship revealed by this is not properly "instantaneous," rather its timing is ambiguous according to the uncertainty relationship. Eq. (3) instead may be regarded as an estimate of a cross-spectral coefficient within the band of the analysis filter, g [38].

Measures of dependence derived from $\hat{S}_{xy}$ in essence reflect the spectral overlap between signals within the originating band. This fact follows from Plancherel's theorem:

$$\int_{-\infty}^{\infty} X(t)Y^*(t)\,dt = \int_A^B \tilde{X}(\omega)\tilde{Y}^*(\omega)\,d\omega \tag{5}$$

where the support of $\tilde{X}(\omega)\tilde{Y}^*(\omega)$ is restricted to the band $[A, B]$. Two signals with non-overlapping spectra will therefore always have a cross-spectrum of 0. Signals whose spectra

overlap but are otherwise uncorrelated will yield cross-spectral estimates subject to small-sample bias tending to 0 with increasing sample size.

Just as one quantifies the degree of instantaneous dependence between real-valued sequences with Pearson's correlation coefficient, a measure of dependence for analytic signals is obtained by scaling the cross-spectral estimate, *coherency*:

$$C_{xy} = \frac{S_{xy}}{\sqrt{S_{xx}S_{yy}}} \tag{6}$$

Coherency differs from its real-valued analog in one critical respect: its nature is composite, blending two distinct aspects of the dependence between analytic signals. The first relates to the correlation of envelopes, and the second to the consistency of phase difference between the signals, $\Delta\phi_{xy} = \phi_x(t) - \phi_y(t)$, that is, their degree of phase locking. The overall degree of dependence may be quantified by the magnitude of the coherency, $|C_{xy}|$ (referred to as coherenc*e*) while the characteristic difference of phase between signals is given by the argument (phase angle), $\theta_{xy} = \arg C_{xy}$.

Because of its composite nature, the meaning of coherence is inherently ambiguous. A given value might reflect either correlated amplitudes or phase locking. Moreover, the contribution of each term depends on an interaction with the other. When the signals remain in a perfectly constant phase relationship ($\Delta\phi_{xy}(t) = c$), coherence reduces to an uncentered correlation of the amplitudes. In such cases, it will tend to vary around some intermediate value when amplitudes are uncorrelated, but it can assume any other value depending on the variance and skewness of the amplitude distributions. On the other hand, two signals that have no phase relationship will give coherence that tends to 0, regardless of amplitude correlation. This ambiguity represents a significant drawback to an otherwise appealingly simple measure.

C. Measures on Phase Only

The problems with coherence motivate alternative measures that are more easily interpreted. A seemingly straightforward way to address the issue is to examine dependence between the phases of analytic signals in isolation from amplitude. Measures of pure phase dependence discard information about signal amplitudes by dividing each signal by its own modulus, thus normalizing away amplitude fluctuations:

$$\Phi_x(t) = \frac{X(t)}{|X(t)|} = e^{i\phi_x(t)} \tag{7}$$

The simplest of these pure phase measures, phase-locking value (PLV), or phase coherency, computes coherency as in (6) from the amplitude normalized signals, giving the expectation of the phase difference vector, $E\left(e^{i\Delta\phi_{xy}}\right)$ [2]. The magnitude of this vector remains at 1 when the phase difference between the signals is perfectly stable and tends to zero when there is no consistent phase relationship. This measure also has an intuitive geometric interpretation as the degree to which a sample of unit length phase vectors, representing the angular difference, align in the same direction.

One drawback of PLV is that it addresses only second-order dependence between the phase vectors. More general forms

of dependence can be revealed through information-theoretic metrics applied to the bivariate distribution of phase angles, such as Kullback–Leibler divergence or the Kolmogorov–Smirnov test. However, both PLV and these alternatives are fundamentally measures on the same distribution of phase angles, meaning that differences between them reflect the nature of the measure rather than any property of the quantity measured. For this reason, one may view amplitude normalization as implicit in measures based purely on phase, along with any spectral consequence.

D. Nonstationary Extensions

Discussion has so far been confined to stationary and ergodic signals, but it is often of interest to consider applications in which stationarity is not assumed. One simple approach to the nonstationary case accounts for time dependence by introducing time-weighting to the estimator. Formulating Eq. (4) as a Riemann-Stieltjes integral over the relevant weighting measure allows the extension to apply both to continuous and atomic weighting. Such weighting may, for example, describe a subsampling of the signal used to reveal phase-locking at specific delays with respect to some externally defined series of events, allowing for event-related phase-locking analyses. This point will be considered again briefly in section VI-C.

III. SPECTRAL BIAS FROM AMPLITUDE NORMALIZATION

We have just seen that measures derived from cross-spectral estimates represent the dependence between signals within their originating bands. The same is not true for measures based purely on phase. This fact follows from the nonlinearity of amplitude normalization, which may induce spectral components outside the original signal bandwidth. Two signals that initially shared no common spectral features, with coherence strictly zero, may therefore exhibit nonzero coherence after the normalization (i.e. phase locking) should any of the induced spectral features happen to overlap. Next, we consider how this comes about.

A. Origin in the Fourier Domain

Normalization entails multiplying the original band-limited analytic signal, $X(t)$, by a function obtained from the inverse of its norm, $N_x(t) = |X(t)|^{-1}$. The spectral consequences of this follow from the convolution theorem, according to which multiplying two signals in the time domain is equivalent to convolving them in the frequency domain, and vice versa. As a result, the spectrum of Φ is that of the original signal smeared with the spectrum of its inverse envelope, creating a form of spectral leakage bias.

One worrisome aspect of this is that the inverse envelope tends to have a very broad spectrum because it contains large spikes wherever the envelope approaches zero; as a result, amplitude normalization tends to smear spectral energy across a wide swath of the spectrum. The effect of such spectral smearing can be appreciated most easily in phase "jumps" or "slips" within the normalized signal. An example of this effect is shown in Fig. 1 for an amplitude modulated signal composed

Fig. 1. Spectral consequence of amplitude normalization. *Top panels*: the test signal, $X(t)$ is composed of two sinusoids at 10 Hz and 12 Hz, whose summation creates an amplitude modulated signal with center frequency 11 and modulation frequency of 2. *Middle panels*: the inverse amplitude function, $N_x(t)$, exhibits spikes where the amplitude envelope approaches 0, for which reason its spectrum, $\tilde{N}(\omega)$, includes harmonics of the modulation frequency. *Bottom panels*: As a result of the convolution property of the Fourier transform, multiplying X by its inverse modulus to obtain Φ gives a signal that includes harmonics and subharmonics of the test signal spaced at 2 Hz. These features create the potential for spurious phase locking if they overlap with features of spectrally unrelated signals.

of the sum of a 10 Hz and a 12 Hz sinusoid, giving an 11 Hz signal modulated at 2 Hz. The inverse envelope contains energy at all harmonics of the modulation frequency, as a result of which the normalized signal also contains energy at harmonics of 2 Hz in addition to 10 Hz and 12 Hz. These new harmonics, which were not present in the original signal, may create spurious phase locking if they overlap with components of a compared signal.

B. Some Nasty Statistics, or, the Futility of Filtering

If the example in Fig. 1 seems contrived to produce broad spectral artifacts, one may consider the somewhat more relevant cases of random noise. For an analytic signal whose real and imaginary parts are independent and Gaussian, amplitude follows a Rayleigh distribution [39]. For this distribution one finds that the inverse amplitude has a finite expected value,

$$\langle |X|^{-1} \rangle = \frac{1}{\sigma^2} \sqrt{\frac{\pi}{2}}$$

but unbounded variance,

$$\langle |X|^{-2} \rangle \to \infty$$

meaning in practice that sample variance for the inverse amplitude will increase with sample size as the number of intermittently appearing large outliers grows, creating a signal whose energy comes increasingly to be dominated by such outliers, whose spectrum therefore becomes broader as its duration increases.

It is natural to suppose that bandpass filtering does away with this problem by smoothing out the dips in amplitude

responsible for spectral broadening of the inverse. Casual inspection of the signal might lead one to take this point for granted, as a noisy signal initially full of phase jumps and amplitude kinks looks smooth and oscillatory after filtering. Unfortunately, any such improvement is an illusion: filtering only decreases the effective sample size of the signal, thereby decreasing sample variance of the inverse amplitude, but it does not change the underlying distribution, which remains as poorly behaved as before. To see why this must be true, one need only note that decimating a stationary narrowband Gaussian signal to a sampling rate that matches its bandwidth should yield nearly independent Gaussian samples; but decimation does not change the ergodic distribution of sampled amplitudes.

A similar conclusion may be reached in an entirely different way [21]: consider that the autocorrelation of the spectrum of a signal is the Fourier transform of the square of the signal's time envelope. The constant envelope of Φ entails a spectrum composed of a single delta spike at 0 Hz. If a discrete finitely sampled spectrum is band-limited, its autocorrelation must contain a nonzero value at the lag separating the outermost nonzero samples. This precludes the possibility that a band-limited spectrum of a finite signal can generate the required autocorrelation, unless the bandwidth is infinitesimal, in which case the signal must be a pure sinusoid. This argument leaves no escape from the conclusion that amplitude normalization causes spectral broadening of the normalized signal, except in trivial cases.

Fig. 2. Example of decentering bias. Before amplitude normalization a test signal with non-uniform phase distribution (*inset*) has a mean of zero (*top panel*). After normalization (*bottom panel*) the mean is offset from zero (*dashed line*). This effect can be understood as an instance of spectral leakage from the original signal bandwidth to 0 Hz following normalization. Before normalization, the analytic amplitude (*gray line*) is modulated with the same period as the original signal. This fact induces the 0 Hz component following multiplication by the inverse amplitude.

IV. DECENTERING BIAS

Spectral leakage from amplitude normalization increases the opportunity for bias in measures of phase locking by spreading energy in the Fourier domain, but this effect is not sufficient by itself to create such bias. We will now consider a specific example of how spectral leakage may cause the appearance of a phase relationship between signals that are, in fact, entirely independent. As touched on in the discussion of Eq. (4), cross-spectral estimation generally does not involve mean centering because the input signals are band-limited by assumption and so should already have an expectation of 0. Clearly, spectral leakage means this condition may no longer apply following amplitude normalization.

An instance of this effect was recently highlighted by van Driel et al. [19], who describe it as resulting from non-uniformity of the distribution of phase angles in the input signals. If this distribution is not at least radially symmetric, any non-uniformity will cause a deviation of the expected phase vector from the origin. One obtains a nonzero value in the absence of any true correlation if both expectations are nonzero. One solution may seem obvious at this point: subtract the mean phase vector from at least one of the signals before computing any statistics related to phase locking. From our Fourier domain perspective, this correction amounts to applying a maximally high-pass filter, eliminating only the 0 Hz component. There are immediate problems with this solution, however, which will be considered in the following section.

First, to understand how energy comes to be leaked to 0 Hz, consider that a peak in the distribution of a signal's phase implies the presence of amplitude and frequency modulation whose spectrum overlaps with that of the original signal (Fig. 2). This is easiest to understand in terms of frequency

modulation, as a peak at a given phase value means that the progression of phase in the signal must linger around the peak value, implying a drop in average instantaneous frequency. For the signal to be zero-mean, such a dip must be accompanied by a drop in amplitude to compensate for the corresponding bias in the direction of phase. These spectral features in the amplitude carry over to its inverse, and normalization therefore induces a 0 Hz spectral component through the cancellation of the dominant component in the original signal and the matching (phase-locked) component in amplitude modulation.

Decentering bias therefore occurs whenever the phase and magnitude of the underlying signal are correlated. This can also be shown from the fact that the covariance of amplitude and phase vector yields the expectation of the original signal, which here is assumed to be 0:

$$\mathrm{E}\left[|X|\Phi_X\right] = \mathrm{E}[X] = 0 \tag{8}$$

so that the two are second-order independent only if phase is centered:

$$\mathrm{E}\left[|X|\Phi_X\right] = \mathrm{E}\left[|X|\right]\mathrm{E}\left[\Phi_X\right] \quad \text{iff} \quad \mathrm{E}\left[\Phi_X\right] = 0 \tag{9}$$

For this to happen in a stationary signal, the signal bandwidth, which governs the spectral content of amplitude fluctuations, generally must approach the spectral range of the signal, that is, $\Delta W \approx f_1$, where f_1 is the lower bound of the signal spectrum; however, because spectral leakage from amplitude normalization may be infinitely extended, this condition should be taken as a rule of thumb rather than a guarantee. The example of two interfering sinusoids in Fig. 1 gives a simple case in which spectral leakage extends to 0 Hz, even though bandwidth condition is met.

It should be noted here that in spite of the examples considered in the previous sections, stationary Gaussian signals are not susceptible to this type of bias because they are completely characterized by their first two moments (mean and cross-covariance). Any dependence between the signal envelope and phase must involve moments of order 3 or higher. For this reason, if the originating signals are strictly Gaussian, pure phase-locking measures must still be directly related to second order statistics, which has recently been confirmed by Aydore et al. [12]. But many if not most biological signals are manifestly non-Gaussian; this fact is exploited by widely-used blind-source separation techniques, which equate sources with maximally non-Gaussian signal components [40]. Here it will likewise be taken for granted that the signals in question are non-Gaussian.

A. The Problem with Recentering

Centering the normalized signal removes the bias resulting from an off-center mean. It has been suggested that this correction effectively serves to uniformize the phase distribution, implying that it deals more generally with problems related to non-uniform phase [19]. Such a conclusion is premature because it neglects another important consequence of recentering: following mean subtraction, phase vectors no longer have uniformly unit length. It is therefore mistaken to suppose that mean centering of the normalized signal implies a balancing

of the distribution of phase angles. Instead, the correction depends in part on a re-injection of some of modulation that was removed through amplitude normalization.

The centered signal is:

$$\Phi^*(t) = h(t)e^{i\phi^*(t)}$$
$$= e^{i\phi(t)} - |\mu_\phi|e^{i\theta_\mu} \qquad (10)$$

where the amplitude-normalized signal before centering is $\Phi(t) = e^{i\phi(t)}$ and the centering term is given as $\langle \Phi(t) \rangle = |\mu_\phi|e^{i\theta_\mu}$. It is easy to see that this results in an amplitude modulated signal with envelope:

$$h(t) = \sqrt{1 + |\mu_\phi|^2 - 2|\mu_\phi| \cos\left(\phi(t) - \theta_\mu\right)} \qquad (11)$$

This mean-centered signal is no longer the pure representation of phase we had hoped to obtain, and we find ourselves back in the position from which we started.

B. Corrections to Recentering

The foregoing example illustrates the whack-a-mole character such corrections often assume in the spectral domain, where compensating for one artifact introduces others elsewhere. To preserve the desired properties of the phase vector, we briefly point out two improvements to mean-subtraction.

1) Iterated recentering: The first of these simply repeats recentering and normalization through multiple iterations:

$$\Phi_{\text{ic}}^{(n+1)}(t) = \frac{\Phi_{\text{ic}}^{(n)}(t) - \mu_\phi^{(n)}}{h^{(n)}(t)} \qquad (12)$$

We have observed this to converge rapidly on a phase representation with the desired properties: unit amplitude and zero mean. The resulting distribution of phase is radially symmetric but not generally uniform.

2) Phase Uniformization: The second correction enforces a strictly uniform distribution by applying a suitable transformation to the phase angle [17]. This is easily done by passing the angular representation of phase through its cumulative distribution over $[0, 2\pi]$

$$\phi_{\text{u}}(t) = 2\pi\hat{F}\left(\phi(t)\right)$$

Such a correction is a more general fix for any problem that might arise from a lack of uniformity. A possible drawback in some settings is the need to estimate the cumulative distribution function, $\hat{F}$, which is a more complex problem than recentering, involving more degrees of freedom. It may therefore not suit cases in which a reliable estimate is difficult to obtain, such as with small samples.

Most importantly, neither of these approaches addresses the general problem of spectral bias. They only correct the problem as it affects mean phase. We next consider how and when such corrections fail.

V. Bias and Correlated Amplitude Modulation

We have just seen a concrete example of how spectral leakage might affect measures of phase locking through decentering of the amplitude normalized signal. We also reviewed some simple techniques to correct that particular incarnation of bias.

Next we will describe conditions under which decentering bias can escape any correction designed to recenter the distribution of phase. The key to understanding how this comes about lies, once again, in the spectral effect of amplitude modulation.

Modulating (i.e. multiplying) a signal, $X(t)$ by a pure complex sinusoid, $e^{i\omega_c t}$, has the effect of shifting the entirety of the signal's spectrum over by an amount equal to the frequency of the sinusoid. This follows directly from the convolution theorem. The result, $\tilde{X}(\omega - \omega_c)$, is therefore in every respect identical to the original spectrum of $\tilde{X}(\omega)$, except for having been shifted along the frequency axis by ω_c.

This property brings up an important point that will prove conceptually useful: the invariance of the spectrum under frequency translation carries over to all measures of phase locking we have encountered so far. This follows from the fact that the modulating term $\omega_c t$ enters into the complex argument (phase angle) of the modulated signal and will therefore cancel when computing the phase difference between two signals that have been modulated by the same sinusoid, resulting in the same outcome as if no modulation had been applied to either. It should also be apparent that this property extends more generally to any complex-valued modulating signal with unit amplitude, $e^{i\beta(t)}$. Modulation of the original signal with a complex unit-amplitude function will therefore cause some of the 0 Hz component of the normalized signal, responsible for decentering bias, to shift to modulating frequency(-ies). In this scenario, it is possible for the mean phase vector to be zero and the distribution of phase angles perfectly uniform in the input signals. Yet, if the two signals under comparison share amplitude modulation, the cancellation of the common component in the argument reestablishes decentering bias.

The relevance of this to real-valued signals becomes clear by constructing the modulating signal from a zero-mean analytic signal $B(t) = |B(t)|e^{i\beta(t)}$ representing amplitude fluctuations around a constant offset μ_b:

$$b(t) = |B(t)| \cos\left(\beta(t)\right) + \mu_b$$
$$= \frac{1}{2}|B(t)| \left(e^{i\beta(t)} + e^{-i\beta(t)}\right) + \mu_b \qquad (13)$$

Provided the spectrum of $b(t)$ is confined to a lower range of frequencies than the signal it modulates, $x(t)$, the Hilbert transform preserves the amplitude modulation with b (Bedrosian's theorem) [41]:

$$\mathscr{H}\{b(t)x(t)\} = b(t)\mathscr{H}\{x(t)\} \qquad (14)$$

Consequently, the analytic representation of the modulated signal is

$$A_{bx}(t)e^{\phi_{bx}(t)} = b(t)X(t)$$
$$= \left[|B(t)| e^{i\beta(t)} + |B(t)| e^{-i\beta(t)} + \mu_b\right] X(t) \qquad (15)$$

From Eq. (15), one can see that modulation with a low-frequency real-valued sinusoid (e.g. $b(t) = \cos(\omega_c t)$, $\mu_b = 0$) has the effect of reduplicating and summing the original spectrum, translating one copy by $+\omega_c$ and the other by $-\omega_c$ along the frequency axis. More generally, a smeared out version of the original spectrum will be translated in both directions by an amount reflecting the spectral range of B,

Fig. 3. Comparison of decentering bias in different measures. Test data are generated by skewing the magnitude of independent Gaussian samples drawn from two processes exhibiting phase dependence. For both processes, amplitude is skewed according to an ellipse (example given in upper inset panel), with the degree of skew related to eccentricity (ϵ). The phase differences ($\Delta\phi$) between samples from each process follow a von Mises distribution with scale parameter, κ. Compared measures are phase-locking value (PLV), centered phase-locking value (cPLV), iteratively centered phase locking (icPLV) and amplitude-weighted phase locking (awPLV). *Left columns*, $\kappa = 1.0$, giving expected unbiased PLV of 0.45; *Right columns*, $\kappa = 0.0$, giving expected unbiased PLV of 0 (i.e. no phase locking). The directional dependence of amplitude results in decentering bias in measures of phase locking that normalize by magnitude; this bias is shown as the deviation from the proper phase locking value for the von Mises distribution used to generate phase differences. Displayed values are averaged over 100 draws of 100 samples each and shown as a function of ϵ and $\Delta\phi$. Four cases are addressed: *Top rows*: amplitude skew aligned in the same ($\epsilon > 0$) or opposite ($\epsilon < 0$) directions; *Bottom rows*: skew at right angles. In the presence of actual phase locking (*left columns*), PLV and cPLV underestimate phase locking when the difference matches the skew directions. In the absence of phase locking (*right columns*), PLV overestimates phase locking at large ϵ (note different color scales for left and right columns). Other measures are not affected. The differences in awPLV are a consequence of increasing small sample bias rather than decentering (see section VI-A1 in the text).

creating "sidebands". If μ_b is nonzero, then a third copy at the original location, the "baseband", will reflect the presence of an offset (non-vanishing 0 Hz component) in the spectrum of b.

The consequence this will have for measures of phase locking depends on the analysis filter used to obtain the relevant statistics. If the passband covers the full spectral range of the modulated signal, $b(t)X(t)$, and if b is positive, the distribution of phase remains unaffected because the modulation is retained within the amplitude of the result. If the passband instead isolates one of the side bands, $|B|e^{\pm i\beta}$, then the effect of amplitude modulation enters into the phase of the result by way of the complex argument.

Appendix A reviews three cases: (1) an analysis filter whose passband covers the entire spectrum of the modulated signal (broadband case); (2) a passband that covers only one of the shifted sidebands (sideband case); and (3) a passband that covers the unshifted band induced by any nonzero offset of the modulating signal, as occurs when the modulating signal is strictly non-negative (baseband case). In all cases, pure phase-locking measures remain unaffected by the modulating signal because any influence on the phase of the separate signals cancels from the difference. This means, however that any decentering bias in the unmodulated case remains in effect after modulation, even though the distribution of phase may be uniform within each side band. The off-center mean may be obscured as a result of phase dispersion introduced by amplitude modulation, making recentering ineffective. This effect is illustrated in Fig. 5, which also compares some of the alternative measures discussed in the following section,

8

Test Signal

Fig. 4. Test signal used to illustrate shifted decentering bias in Fig. 5. *Top panel:* Five second example from a random test signal. Test signals contained random and independent uncorrelated Gaussian transients modulated by a common 0–6 Hz signal. The modulating signal (dashed line) was obtained by squaring 1.5–3 Hz filtered white noise with mean of 1 and standard deviation of 0.5 after filtering, to which a further positive 0.25 offset was added after squaring. Power spectral density is shown in the bottom panel. The common envelope was applied to 100 random uncorrelated trains of 25 ms std. width Guassians with center times determined by positive peaks in 8–12 Hz filtered white noise. A common 50 Hz pure sinusoid served as a positive comparison. Total signal duration was 80 s. *Middle panel:* Signal after application of an analysis filter centered at 12 Hz. The analysis filter applied a frequency-domain cosine window over 0–24 Hz. The presence of the common low-frequency amplitude modulation creates spurious phase locking in the range of the uncorrelated train of gaussians at 12 Hz (see Fig. 5).

using test signals described in Fig. 4.

More generally, it is not necessary to assume that the bands induced by amplitude modulation are fully separated by the analysis filter. When the bands overlap so that the modulated signal lacks distinct side bands or the analysis filter otherwise provides an incomplete separation, the filtered signal contains an unequally weighted sum of components arising from each band. In such cases, the effect of amplitude modulation on the phase of the result is will be intermediate to the side- and broadband cases, depending on relative weighting, so that recentering is still incompletely effective.

8

VI. REMEDIES

We have seen that measures of dependence taken directly from the cross spectrum are not affected by the potentially severe forms of spectral leakage produced by amplitude normalization (Fig. 5). But the canonical measure in this category, coherency, has other drawbacks, which diminish its attractiveness as a proxy for phase locking; in particular, it combines phase locking and amplitude correlation in a way that complicates interpretation. In the following we consider how the simple average used by PLV can be replaced with a weighted average. A variant of this approach, amplitude-weighted (awPLV), avoids amplitude normalization altogether, instead retaining the original cross-spectral amplitude weighting but with a rescaling that overcomes some of the interpretational difficulties of coherence. These measures avoid any effects of spectral broadening that follow from amplitude normalization. We also briefly consider more general extensions of this scheme to other types of weighted averages.

A. Amplitude-weighted Phase Locking

The covariance of two analytic signals that are perfectly phase locked, meaning $\Delta\phi_{xy} = \phi_x(t) - \phi_y(t)$ remains constant, is given by:

$$E\left[S_{xy}\right] = \int E\left[A_x(t)A_y(t)e^{i\Delta\phi_{xy}(t)}\right] dt$$
$$= e^{i\Delta\phi_{xy}} \int E\left[A_x(t)A_y(t)\right] dt \tag{16}$$

One might therefore consider a measure of "pure" phase locking that factors out the influence of amplitude covariance simply by normalizing the cross spectrum with $\int E\left[A_x(t)A_y(t)\right] dt$. Equivalently, such a measure is obtained from the normalization of coherency by the uncentered amplitude correlation in the presence of phase locking.

$$\Upsilon_{xy} = \frac{\langle XY^*\rangle}{\sqrt{\langle|X|^2\rangle\langle|Y|^2\rangle}} \bigg/ \frac{\langle|X||Y|\rangle}{\sqrt{\langle|X|^2\rangle\langle|Y|^2\rangle}}$$
$$= \frac{\langle XY^*\rangle}{\langle|X||Y|\rangle} = \frac{\langle A_x A_y e^{\Delta\phi_{xy}}\rangle}{\langle A_x A_y\rangle} \tag{17}$$

The result of this normalization has a natural interpretation as a weighted-average phase coherence:

$$\hat{\Upsilon}_{xy} = \sum_t w_t e^{\Delta\phi_{xy}[t_i]} \tag{18}$$

with weights given by normalized amplitude products:

$$w_t = \frac{A_x A_y}{\sum_t A_x A_y} \tag{19}$$

The meaning of this measure is therefore much more closely aligned with that of phase-only measures, despite its close relationship to coherence. Like PLV, its magnitude becomes 1 if the signals are perfectly phase locked, regardless of how the amplitudes are related.

This measure will differ from standard PLV in three ways; first, as already noted, it avoids spectral biases of the type reviewed above. Second, like coherence, awPLV weights the estimate towards samples with relatively greater amplitude,
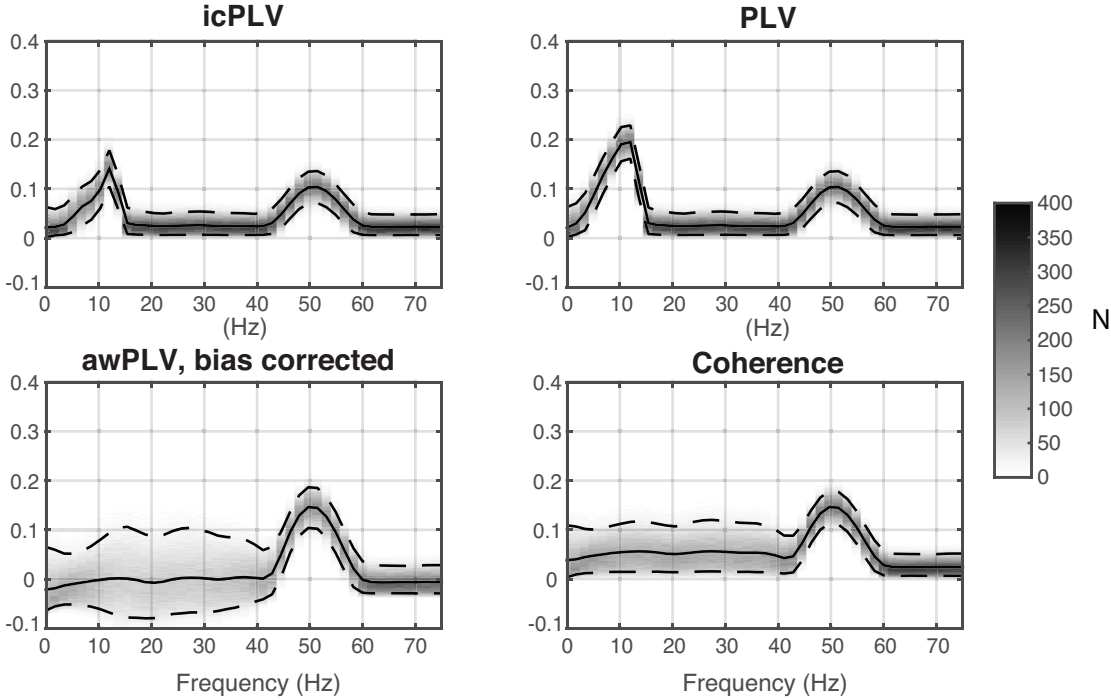
Fig. 5. Spectral bias in four measures of phase dependence applied to test signals exemplified in Fig. 4. Phase-locking measures were computed between all 10,000 pairs of signals drawn from two sets of 100, and the figure color scale indicates the distribution of values for each frequency bin. Lines show the median, 5th and 95th percentiles of these distributions. In all cases, the peak at 50 Hz reflects the 50 Hz positive comparison signal component. *Top panels:* Mean-centered (icPLV) and conventional phase-locking values (PLV) both show spurious phase locking around 12 Hz, in the frequency range of the uncorrelated train of gaussians, as a result of a common 1.5–3 Hz modulating envelope. Mean centering creates relatively greater suppression of the bias at low frequencies (*top left panel*), which reflects the contribution of the unmodulated constant component of the envelope. *Bottom panels:* awPLV and coherence are not affected by this bias. Measures were computed within overlapping cosine-windowed 16 Hz frequency bands using the DBT method (Kovach and Gander, 2016).

and therefore will diverge from PLV if the phase difference exhibits dependence on signal amplitudes. Third, as reviewed in the following section, statistical properties of awPLV will depend on effective sample size, which is always smaller for a non-uniformly weighted average than for an unweighted average.

1) Effective Sample Size: For a non-uniformly weighted average, sample size is no longer an accurate reflection of the degrees of freedom in the average, which is more accurately reflected by the effective sample size, ν. For independent samples, effective sample size is approximated as

$$\nu = \frac{\left(\sum_t w_t\right)^2}{\sum_t w_t^2} \qquad (20)$$

In general, $\nu \leq N$. In the absence of phase locking, the expected value of PLV is related to sample size; this small-sample bias is given by the inverse square root of sample size, $1/\sqrt{N}$ (assuming independent samples). Small-sample bias is approximated for awPLV in the same way, substituting effective sample size:

$$\beta = \frac{1}{\sqrt{\nu}} \qquad (21)$$

Using this relationship one may define a bias-corrected measure, which is centered at zero in the absence of phase locking and otherwise uses the full dynamic range between 0 and 1:

$$\hat{\Upsilon}_{xy}^* = \frac{\hat{\Upsilon}_{xy} - \beta}{1 - \beta} \qquad (22)$$

This correction addresses another shortcoming of standard coherence estimates, which is that the degree of bias varies across frequency, depending on the amplitude distributions within the respective analysis bands [16], potentially creating the false impression of frequency-dependent changes in coherence.

While inflated small-sample bias is clearly an important consequence and possible disadvantage of non-uniform weighting, the apparent loss of statistical power in awPLV and coherence relative to PLV is very likely to be offset by the fact that samples with amplitude approaching zero, for which phase is either poorly defined or unlikely to reflect any signal beyond the noise floor, are granted equal weight in PLV. To the extent that phase locking is associated with the amplitudes of the carrier signals, awPLV will therefore tend to improve the signal to noise ratio. On the other hand, for a signal contaminated with high-amplitude transients, such as movement artifacts, awPLV will perform poorly, raising the importance of careful artifact rejection.

For a bandpass-filtered signal, Eq. (21) is clearly not valid when summing over contiguous samples because filtering creates correlations among samples over time, violating the assumption of independence. Using (21) will give a greatly inflated estimate of the effective sample size in such cases. One way to correct the problem is by downsampling the signal so that samples are more nearly independent; this is conveniently addressed by spectral estimation techniques that combine bandpass filtering with downsampling, adjusting sam-

pling to the bandwidth of the analysis filter. We have recently described such an approach, the demodulated band transform (DBT) [35], which uses a frequency-domain implementation of complex demodulation [33].

Effective sample size is strictly valid only when weights are independent of phase, which may often fail to hold in the presence of phase locking. It is, however, a valid assumption under the null hypothesis of no phase locking. Null hypothesis tests on awPLV address the question how likely one is to obtain the observed value for two signals with the respective amplitude envelopes in the absence of phase locking.

B. Permutation Tests

Statistical validation of phase locking frequently resorts to a surrogate null distribution generated by shuffling data over samples. The rationale is that shuffling abolishes any true phase locking while leaving biases intact. The considerations reviewed here show the limits of this assumption. In the case of pure measures of phase locking, we have seen that bias may depend on correlations of the signal amplitudes. Clearly, shuffling destroys any amplitude correlation along with phase locking, meaning that the surrogate distribution will fail to account for any related bias. Although awPLV avoids spectral biases, it is not immune to this problem because shuffling may systematically change effective sample size, likewise ruining the validity of the comparison. By increasing the variance of the product of amplitudes, positive correlations depress effective sample size; shuffling may therefore produce a systematic decrease in small sample bias, raising type I error in any test that relies on the comparison to shuffled data.

There seem to be few perfectly satisfactory options for addressing this problem. One is to shuffle phase independently of amplitude, but this creates distortions similar to those arising from amplitude normalization. A better alternative might be to shuffle amplitude and phase in tandem but reduce (or increase) the number of samples in the averages drawn from the surrogate distribution in proportion to

$$N^{(k)} = N \frac{\nu}{\nu^{(k)}} \tag{23}$$

where $\nu^{(k)}$ is the effective sample size of the k^{th} resampling including N samples. Such an adjustment at least removes the tendency for systematic changes in effective sample size, although it does not exactly recreate the amplitude distribution of the original signals. Another more complicated approach might be to stratify the permutation by amplitude, so that surrogate data maintain a similar amplitude profile.

C. Generalized Weighting

One might also consider alternative weighting schemes other than simple amplitude weighting as used by awPLV; for example, Ehm et al. suggest discarding samples whenever amplitude drops below some threshold [42]. In general, reweighting will take the form:

$$\hat{\Psi}_{xy} = \frac{\int \Phi_x \Phi_y^* \, dW(t, |X(t)|, |Y(t)|)}{\int dW} \tag{24}$$

The hat symbol, used here to indicate an estimator, is deliberately retained in (24) because this notion of weighting also encompasses sampling of the signal at specific points in time within some window of observation. So that it applies to this range of cases, the notation in Eq. (24) expresses the measure as a Riemann–Stieltjes integral, which is a formality valid for both discrete atomic and continuous weighting. In general, cross-spectral measures are optimal with respect to spectral leakage, so any form of weighting other than a rescaling of the amplitude product will create some additional spectral broadening. Nevertheless, reweighting may improve the measure in other ways, such as in moderating small-sample bias and other non-spectral-leakage-related biases imparted by amplitude weighting.

Such reweighting is also implicit when samples are drawn from selected time points, as arises generally in the context of digital signal processing, but also more specifically when measuring phase locking at certain time delays, τ, from some event, such as trial onset. This case is represented by the following estimator:

$$\hat{\Psi}_{xy}(\tau) = \int w(t) \Phi_{xy} \sum_i \delta(t - t_i - \tau) \, dt$$
$$= \sum_i w(t_i + \tau) \Phi_{xy}(t_i + \tau) \tag{25}$$

for which weighting is the composite of a continuous weighting function w and the subsampling represented by the delta train. To understand the spectral consequences of such subsampling, one may note that the Fourier transform of a train of delta functions with period ΔT is also a periodic train of spikes in the frequency domain, with period $2\pi/\Delta T$. One can observe that this form of "windowing" has the effect of reduplicating the original signal spectrum at frequency intervals of $2\pi/\Delta T$. If the period of the original train is less than the inverse bandwidth of the signal, expected cross-spectral estimates remain unchanged because the reduplicated spectra do not overlap. More generally, the cross-spectral estimate will be relatively unaffected if the minimum interval between samples is less than the inverse bandwidth of the signal. This point simply restates the well-known Whittaker–Nyquist sampling theorem.

In event-related analyses the goal is to reveal nonstationarities associated with some event. In this context, sampled points are widely spaced and the reduplicated spectral windows overlap, which implies convolutional smearing within the signal bandwidth. Any systematic change that results from this is not necessarily cause for concern, however, as it captures the nonstationarities that motivated the sampling in the first case.

VII. DISCUSSION

A. Some Advantages of Linearity

We have treated spectral broadening related to the nonlinearity of phase locking as a form of bias, though we also noted at the outset that this might fairly be considered an abuse of terminology. The nonlinearity of phase extraction is in fact faithful to the nature of analytic phase; what we have taken to calling bias is a genuine property of the signal, not the result of

bias in the usual meaning of the word. More accurately stated, the problem is a choice between approximations: whether to approximate phase locking with measures that preserve linearity or to treat nonlinear measures as approximately linear in our interpretation of them (wittingly or unwittingly).

One key advantage of the first option is that the underlying measures behave in a more consistent and predictable fashion. In particular, nonlinear effects are sensitive to the parameters of the analysis filter in ways that are difficult to expect, making the second approximation unreliable. A good example of this is given in section IV, where we observed that decentering bias depends on some overlap between the spectrum of the signal envelope and the signal itself. Decreasing the analysis bandwidth tends to eliminate or attenuate the overlap, making the bias disappear. On the other hand, decentering bias masked by correlated amplitude modulation occurs when the analysis filter was sufficiently narrow to isolate individual sidebands arising from the modulation. Both of these effects show the sensitivity of the result to the analysis filter. In contrast, for stationary cross-spectral estimates, changing the analysis filter bandwidth affects frequency resolution but does not fundamentally change the result. While broader filters provide a less highly resolved spectral estimate, the decomposition into more finely grained bands remains additive: the picture becomes sharper (and noisier), but it does not otherwise change.

Some authors have cautioned against narrow filters because they artificially force the signal to take on an oscillatory character [43] and create timing ambiguities [44]. A rejoinder to the first point is that techniques which couple filtering with complex demodulation and downsampling matched to bandwidth, don't yield an inherently oscillatory signal representation [35]. The equivalence between these techniques and the more standard approach to filtering shows that the oscillations in filtered data are really an artifact of the choice of representation; the errors associated with narrow filtering come from treating the signal as though it might occupy a wider bandwidth than what has been forced upon it, but these errors may be forestalled through downsampling. With respect to the second point, detailed timing information also remains encoded in the relationship of phase across multiple bands, from which it may be recovered using, for example, an inverse Fourier transform computed across bands.

B. Extensions

Amplitude-weighted PLV joins a host of other measures designed to address various shortcomings of PLV and coherence. Alternative measures have been proposed to better cope with small sample bias [22], volume conduction [45], as well as decentering bias [19] and combinations of these problems [46]. Each of these proposed alternatives has its strengths and weakness, making it suited to some applications but not others, which is certainly also the case for awPLV. Most of these measures are meant to address problems with analytic signals considered in isolation. One important advantage of linearity-preserving measures derived from cross spectra is that the array of tools afforded by spectral analysis remain applicable; these might serve to isolate signal components of

interest or to more clearly identify salient properties of the dependence. For example, the time-domain representation of cross-spectral estimators as impulse response functions may reveal useful information about timing that is not reliably conveyed by phase lags considered in isolation. This point becomes especially relevant whenever within-band time resolution is lost to narrowband filtering. It is easy to forget that detailed timing information is nevertheless preserved in the relative phase across multiple bands, which may be recovered through the inverse Fourier transform.

VIII. CONCLUSION

We have reviewed some pitfalls common to measures of dependence derived from analytic phase and discussed simple approaches to remedying these. The main points are:

1) Obtaining analytic phase involves a nonlinear transformation of the signal that may at times create unexpected biases and complicate spectral interpretation (Section III-V).
2) Pure measures of phase exhibit some poor statistical properties in the presence of noise, including spectral biases that grow with sample size. By reducing effective sample size, filtering to limit signal bandwidth provides only an apparent but not an actual remedy for this problem (Section III-B).
3) Measures of linear dependence taken directly from the cross spectrum, such as coherence, preserve spectral specificity and are often statistically better behaved in comparison to pure-phase statistics (Section VI).
4) One such measure, amplitude-weighted phase locking, avoids the ambiguity that is a major drawback of coherence (Section VI-A).
5) Phase is often more informative considered across multiple bands jointly than in isolation, especially when drawing inferences related to timing. (Section VII-A).

REFERENCES

[1] P. L. Nunez, R. Srinivasan, A. F. Westdorp, R. S. Wijesinghe, D. M. Tucker, R. B. Silberstein, and P. J. Cadusch, "Eeg coherency: I: statistics, reference electrode, volume conduction, laplacians, cortical imaging, and interpretation at multiple scales," *Electroencephalography and clinical neurophysiology*, vol. 103, no. 5, pp. 499–515, 1997.
[2] J.-P. Lachaux, E. Rodriguez, J. Martinerie, F. J. Varela *et al.*, "Measuring phase synchrony in brain signals," *Human brain mapping*, vol. 8, no. 4, pp. 194–208, 1999.
[3] F. Mormann, K. Lehnertz, P. David, and C. E. Elger, "Mean phase coherence as a measure for phase synchronization and its application to the eeg of epilepsy patients," *Physica D: Nonlinear Phenomena*, vol. 144, no. 3, pp. 358–369, 2000.
[4] R. T. Canolty and R. T. Knight, "The functional role of cross-frequency coupling," *Trends in cognitive sciences*, vol. 14, no. 11, pp. 506–515, 2010.
[5] A. B. Tort, R. Komorowski, H. Eichenbaum, and N. Kopell, "Measuring phase-amplitude coupling between neuronal oscillations of different frequencies," *Journal of neurophysiology*, vol. 104, no. 2, pp. 1195–1210, 2010.

[6] H. P. Zaveri, R. B. Duckrow, and S. S. Spencer, "The effect of a scalp reference signal on coherence measurements of intracranial electroencephalograms," *Clinical Neurophysiology*, vol. 111, no. 7, pp. 1293–1299, 2000.

[7] S. J. Schiff, "Dangerous phase," *Neuroinformatics*, vol. 3, no. 4, pp. 315–317, 2005.

[8] R. Guevara, J. L. P. Velazquez, V. Nenadovic, R. Wennberg, G. Senjanović, and L. G. Dominguez, "Phase synchronization measurements using electroencephalographic recordings," *Neuroinformatics*, vol. 3, no. 4, pp. 301–313, 2005.

[9] M. Chavez, M. Besserve, C. Adam, and J. Martinerie, "Towards a proper estimation of phase synchronization from time series," *Journal of neuroscience methods*, vol. 154, no. 1, pp. 149–160, 2006.

[10] P. Celka, "Statistical analysis of the phase-locking value," *IEEE Signal Processing Letters*, vol. 14, no. 9, pp. 577–580, Sept 2007.

[11] M. A. Kramer, A. B. Tort, and N. J. Kopell, "Sharp edge artifacts and spurious coupling in eeg frequency comodulation measures," *Journal of neuroscience methods*, vol. 170, no. 2, pp. 352–357, 2008.

[12] S. Aydore, D. Pantazis, and R. M. Leahy, "A note on the phase locking value and its properties," *NeuroImage*, vol. 74, pp. 231–244, 2013. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1053811913001286

[13] J. Aru, J. Aru, V. Priesemann, M. Wibral, L. Lana, G. Pipa, W. Singer, and R. Vicente, "Untangling cross-frequency coupling in neuroscience," *Current opinion in neurobiology*, vol. 31, pp. 51–61, 2015.

[14] R. Srinath and S. Ray, "Effect of amplitude correlations on coherence in the local field potential," *Journal of neurophysiology*, vol. 112, no. 4, pp. 741–751, 2014.

[15] A. M. Bastos and J.-M. Schoffelen, "A tutorial review of functional connectivity analysis methods and their interpretational pitfalls," *Frontiers in Systems Neuroscience*, vol. 9, no. 175, 2016.

[16] E. Lowet, M. J. Roberts, P. Bonizzi, J. Karel, and P. De Weerd, "Quantifying neural oscillatory synchronization: A comparison between spectral coherence and phase-locking value approaches," *PloS one*, vol. 11, no. 1, 2016.

[17] B. Kralemann, L. Cimponeriu, M. Rosenblum, A. Pikovsky, and R. Mrowka, "Phase dynamics of coupled oscillators reconstructed from data," *Physical Review E*, vol. 77, no. 6, p. 066205, 2008.

[18] R. Cox, J. van Driel, M. de Boer, and L. M. Talamini, "Slow oscillations during sleep coordinate interregional communication in cortical networks," *The Journal of Neuroscience*, vol. 34, no. 50, pp. 16890–16901, 2014.

[19] J. van Driel, R. Cox, and M. X. Cohen, "Phase-clustering bias in phase–amplitude cross-frequency coupling and its removal," *Journal of neuroscience methods*, vol. 254, pp. 60–72, 2015.

[20] N. Papp and P. Ktonas, "Critical evaluation of complex demodulation techniques for the quantification of bioelectrical activity," *Biomedical sciences instrumentation*, vol. 13, p. 135145, 1977. [Online]. Available: http://europepmc.org/abstract/MED/871500

[21] B. Picinbono, "On instantaneous amplitude and phase of signals," *IEEE Transactions on Signal Processing*, vol. 45, no. 3, pp. 552–560, Mar 1997.

[22] M. Vinck, M. van Wingerden, T. Womelsdorf, P. Fries, and C. M. Pennartz, "The pairwise phase consistency: a bias-free measure of rhythmic neuronal synchronization," *Neuroimage*, vol. 51, no. 1, pp. 112–122, 2010.

[23] W. J. Freeman, "Origin, structure, and role of background {EEG} activity. part 1. analytic amplitude," *Clinical Neurophysiology*, vol. 115, no. 9, pp. 2077–2088, 2004. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1388245704001695

[24] L. Mandel, "Interpretation of instantaneous frequencies," *American Journal of Physics*, vol. 42, no. 10, pp. 840–846, 1974. [Online]. Available: http://scitation.aip.org/content/aapt/journal/ajp/42/10/10.1119/1.1987876

[25] B. Boashash, "Estimating and interpreting the instantaneous frequency of a signal. i. fundamentals," *Proceedings of the IEEE*, vol. 80, no. 4, pp. 520–538, 1992.

[26] P. J. Loughlin and B. Tacer, "Comments on the interpretation of instantaneous frequency," *IEEE Signal Processing Letters*, vol. 4, no. 5, pp. 123–125, May 1997.

[27] P. Loughlin and B. Tacer, "Instantaneous frequency and the conditional mean frequency of a signal," *Signal Processing*, vol. 60, no. 2, pp. 153–162, 1997. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0165168497800035

[28] J. d. Ville *et al.*, "Théorie et applications de la notion de signal analytique," *Cables et transmission*, vol. 2, no. 1, pp. 61–74, 1948.

[29] L. Cohen, "Time-frequency distributions—a review," *Proceedings of the IEEE*, vol. 77, no. 7, pp. 941–981, Jul 1989.

[30] M. Le Van Quyen, J. Foucher, J.-P. Lachaux, E. Rodriguez, A. Lutz, J. Martinerie, and F. J. Varela, "Comparison of hilbert transform and wavelet methods for the analysis of neuronal synchrony," *Journal of neuroscience methods*, vol. 111, no. 2, pp. 83–98, 2001.

[31] A. Bruns, "Fourier-, hilbert-and wavelet-based signal analysis: are they really different approaches?" *Journal of neuroscience methods*, vol. 137, no. 2, pp. 321–332, 2004.

[32] D. Gabor, "Theory of communication. part 1: The analysis of information," *Electrical Engineers-Part III: Radio and Communication Engineering, Journal of the Institution of*, vol. 93, no. 26, pp. 429–441, 1946.

[33] C. Bingham, M. D. Godfrey, and J. W. Tukey, "Modern techniques of power spectrum estimation," *Audio and Electroacoustics, IEEE Transactions on*, vol. 15, no. 2, pp. 56–66, 1967.

[34] J. B. Allen and L. R. Rabiner, "A unified approach to short-time fourier analysis and synthesis," *Proceedings of the IEEE*, vol. 65, no. 11, pp. 1558–1564, 1977.

[35] C. K. Kovach and P. E. Gander, "The demodulated band transform," *Journal of Neuroscience Methods*, vol. 261, pp. 135–154, 2016.

[36] N. E. Huang, Z. Shen, S. R. Long, M. C. Wu, H. H. Shih, Q. Zheng, N.-C. Yen, C. C. Tung, and H. H. Liu, "The empirical mode decomposition and the hilbert spectrum for nonlinear and non-stationary time series analysis," in *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, vol. 454, no. 1971. The Royal Society, 1998, pp. 903–995.

[37] S. G. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Transactions on signal processing*, vol. 41, no. 12, pp. 3397–3415, 1993.

[38] L. H. Koopmans, *The spectral analysis of time series*. Academic press, 1995.

[39] K. Krishnamoorthy, *Handbook of statistical distributions with applications*. CRC Press, 2015.

[40] R. Vigario, J. Sarela, V. Jousmiki, M. Hamalainen, and E. Oja, "Independent component approach to the analysis of eeg and meg recordings," *IEEE Transactions on Biomedical Engineering*, vol. 47, no. 5, pp. 589–593, May 2000.

[41] E. Bedrosian, "A product theorem for hilbert transforms," *Proceedings of the IEEE*, vol. 51, no. 5, pp. 868–869, 1963.

[42] W. Ehm and J. Kornmeier, "Detecting frequency-specific phase synchrony in brain electrical oscillations," *Proceedings of Fechner Day*, vol. 25, no. 1, pp. 411–416, 2009.

[43] N. Yeung, R. Bogacz, C. B. Holroyd, and J. D. Cohen, "Detection of synchronized oscillations in the electroencephalogram: an evaluation of methods," *Psychophysiology*, vol. 41, no. 6, pp. 822–832, 2004.

[44] R. Vanrullen, "Four common conceptual fallacies in mapping the time course of recognition," *Frontiers in Psychology*, vol. 2, p. 365, 2011. [Online]. Available: http://journal.frontiersin.org/article/10.3389/fpsyg.2011.00365

[45] C. J. Stam, G. Nolte, and A. Daffertshofer, "Phase lag index: assessment of functional connectivity from multi channel eeg and meg with diminished bias from common sources," *Human brain mapping*, vol. 28, no. 11, pp. 1178–1193, 2007.

[46] M. Vinck, R. Oostenveld, M. van Wingerden, F. Battaglia, and C. M. Pennartz, "An improved index of phase-synchronization for electrophysiological data in the presence of volume-conduction, noise and sample-size bias," *Neuroimage*, vol. 55, no. 4, pp. 1548–1565, 2011.

Appendix A
Effect of Amplitude Modulation

A. Effect of Modulation on the Cross Spectrum

1) Broadband Case: If the analysis window, within which the cross spectrum is to be estimated, spans both of the positively and negatively translated range of copies of the original signal, then the conjugate product of two identically modulated signals is given by:

$$b^2 XY^* =$$

$$A_x A_y e^{i\Delta\phi_{xy}} \left(\frac{1}{2} + \mu_b^2 + 2\mu_b |B| \cos(\beta) + \frac{1}{2}|B|^2 \cos(2\beta) \right)$$

For simplicity, we consider the case when the envelope of the modulating signal $|B|$ is independent of A_x and A_y, so that

$$
\begin{aligned}
\langle b^2 XY^* \rangle = \\
\left(\frac{1}{2}\sigma_b^2 + \mu_b^2 \right) S_{xy} + 2\mu_b \mu_{|B|} \langle XY^* \cos(\beta) \rangle \\
+ \frac{1}{2}\sigma_b^2 \langle XY^* \cos(2\beta) \rangle
\end{aligned}
\tag{26}
$$

Thus the cross-spectral estimate now includes two 3rd order moments with the modulating phase, such that whenever the phase difference between X and Y, $\Delta\phi_{xy}(t)$, assumes a consistent relationship with any of $\pm 2\beta(t)$ or $\pm\beta(t)$, it will contribute to the estimate. Correlations between the envelopes A_x and A_y and the modulating phase will also influence the result.

Such dependence between the signals and their common source of modulation may influence coherence, but the interpretation is clear: it reflects the fact that the cross-spectral measures are weighted according to amplitude. Such measures will therefore tend to reflect the characteristic phase relationship when the product of the signal amplitudes is large.

On the other hand, if there is no relationship between B and either of the signals or their amplitudes, the outcome is only a rescaling of the original cross-spectrum:

$$
\langle b^2 XY^* \rangle = \left(\frac{1}{2}\left(\sigma_b^2 + \kappa_1 + \kappa_2 \right) + \mu_b \right) S_{xy}
\tag{27}
$$

where $\kappa_j = \langle \cos(j\beta) \rangle$. In the large-sample limit, the effect of such constant scaling will cancel from coherence (although we will later see that B still plays a role in small sample bias).

2) Sideband Case: If the analysis window only spans one of the shifted spectral bands arising from modulation with b, then the outcome of filtering resembles

$$
bX = \frac{1}{2}A_x |B| e^{i\Delta\phi_{xy}+i\beta}
$$

The complex argument of B cancels out in the cross-spectral estimate giving

$$
\langle |B|^2 XY^* \rangle = \frac{1}{4}\sigma_b^2 S_{xy}
\tag{28}
$$

here assuming independence between $|B|$ and the phase difference $\Delta\phi_{xy}(t)$. In calculating coherence, the constant multiplicative terms above will cancel out in the large-sample limit, although as mentioned, modulation may still affect small-sample bias.

3) Baseband case: If the analysis window covers only the baseband created by any nonzero constant offset μ_b, the result is unchanged except for being scaled by μ_b^2.

B. Effect of Modulation on Phase Locking

1) Broadband Case: Again, when the analysis window spans the full range of the modulated signal, we have

$$
\Phi_{bX} = \frac{bX}{|bX|} = \frac{b}{|b|}e^{i\phi_x} = \text{sgn}(b)e^{i\phi_x}
\tag{29}
$$

If b is strictly positive then modulation clearly has no effect on the distribution of phases or phase differences between two signals. If the value of b may become negative, phase undergoes a π rad flip with every sign change. As a result, the phase distribution for each signal will become more symmetric and the mean phase vector will tend towards zero as the probability of b assuming either sign approaches 0.5. However, in this case there is still no effect on the resulting phase difference, as

$$
\Phi_{bX}\Phi_{bY}^* = (\text{sgn}(b))^2\, e^{\Delta\phi_{xy}} = e^{\Delta\phi_{xy}} \quad \text{a.s.}
$$

In that case, recentering will provide no remedy for bias. The $\pm\pi$ jumps occurring at each flip of the sign of B will also contribute to spectral broadening.

2) Sideband Case: As before, when the analysis filter covers only one of the side bands of the modulating signal, phase is approximated by

$$
\Phi_{BX} = \frac{BX}{|BX|} = e^{i\phi_x+i\beta}
\tag{30}
$$

In this case, the modulating term enters only the complex argument of each signal, therefore canceling from the phase difference and leaving measures of phase locking unchanged. However, the distribution of the phase of input signals will come to reflect $\phi_x(t) + \beta(t)$, and so will be dispersed if ϕ_x and β vary independently. Mean centering therefore has little chance of suppressing bias in this case.

3) Baseband Case: If the analysis filter only covers the baseband, the results are unchanged from the unmodulated signal, as the signal within the analysis band is the same as it was before modulation, except for having been scaled by μ_b. Any decentering bias in this band will therefore be corrected by mean subtraction, as before.