

Untangling cross-frequency coupling in neuroscience

Juhan Aru^{1,2}, Jaan Aru^{1,3,4,5,10}, Viola Priesemann^{3,6}, Michael Wibral⁷, Luiz Lana^{1,3,4,8}, Gordon Pipa⁹, Wolf Singer^{1,3,4}, Raul Vicente^{1,3,4,7,10*}

¹Dept. Neurophysiology, Max-Planck Institute for Brain Research, Frankfurt am Main, 60528 Germany

²École Normale Supérieure de Lyon (UMPA), Lyon, 69364 France

³Frankfurt Institute for Advanced Studies, Frankfurt am Main, 60438 Germany

⁴Ernst Strüngmann Institute, Frankfurt am Main, 60528 Germany

⁵Institute of Public Law, Tartu University, Tallinn, 10119 Estonia

⁶Neural Systems and Coding, Max-Planck Institute for Brain Research, Frankfurt am Main, 60438 Germany

⁷MEG Unit, Brain Imaging Center, Goethe University, Frankfurt am Main, 60528 Germany

⁸Brain Institute, Federal University of Rio Grande do Norte, Natal, 59056, Brazil

⁹Institute of Cognitive Science, University of Osnabrueck, 49069, Germany

¹⁰Institute of Computer Science, University of Tartu, Tartu, 50409, Estonia

*Corresponding author: raulvicente@gmail.com

Abstract

Cross-frequency coupling (CFC) has been proposed to coordinate neural dynamics across spatial and temporal scales. Despite its potential relevance for understanding healthy and pathological brain function, critical evaluation of the available methods for CFC analysis has been lacking. Here we aim to unravel several fundamental problems with the current standard analysis for CFC that make it non-specific and difficult to interpret physiologically. We show that apparent CFC can arise because of spectral correlations due to common non-stationarities and non-linear effects that might be unrelated to a true interaction between frequency components. After reviewing common problems in CFC analysis, we discuss how different statistical/modeling approaches to CFC can be conceptually organized according to their biophysical interpretability and statistical inference approach. This classification provides a possible road-map towards mechanistic understanding of cross-frequency coupling. We end with a list of practical recommendations to avoid common errors and enhance the interpretability of the analysis.

Table of Contents

Cross-frequency coupling: How much is that in real money?.....	3
Caveats and confounds of the CFC analysis.....	5
Instantaneous phase and amplitude: when are they meaningful and when not?.....	6
The importance of the bandwidth.....	6
Non-stationarity and spectral correlations: two sides of the same coin.....	7
Surrogate data: none are perfect but some are better than others.....	11
CFC modulation across conditions.....	12
Organization of modeling/statistical approaches to CFC	13
Practical recommendations.....	17
Conclusions.....	18
Acknowledgements.....	19
References.....	19
Supplementary Materials.....	22
Materials and Methods.....	22
Analysis of CFC.....	22
Data collection for Figure 3 in the main text.....	22
Data collection for the Supplementary Figure 3.....	23
Supplementary Results.....	24
Literature review.....	24
Examples of spurious cross-frequency coupling (CFC).....	26
Effect of input on CFC in human intracranial recordings.....	28
Phase-amplitude coupling for event-related potentials (ERPAC).....	29
Atmospheric noise shows CFC after squaring the signal.....	31
Small static non-linearity in ECoG data generates CFC.....	32
Mathematical example of a non-linearity.....	33
Supplementary Discussion.....	35
Conditions for a meaningful phase.....	35
The importance of the bandwidth.....	36
Different model/statistical approaches to assess phase-amplitude CFC	37
The transitivity of correlation between phase and amplitude.....	39
Supplementary discussion on causality methods.....	40
References for the supplementary material.....	41

Cross-frequency coupling: How much is that in real money?

One of the central questions in neuroscience is how neural activity is coordinated across different spatial and temporal scales. An elegant solution to this problem could be that the activity of local neural populations is modulated according to the global neuronal dynamics. As larger populations oscillate and synchronize at lower frequencies and smaller ensembles are active at higher frequencies (Buzsaki, 2006), cross-frequency coupling would facilitate flexible coordination of neural activity simultaneously in time and space. In line with this proposal, many studies have reported such cross-frequency relationships (Bragin et al., 1995; Canolty et al., 2006; Canolty and Knight, 2010). Especially phase-amplitude CFC, where the phase of the low frequency component modulates the amplitude of the high frequency activity, has been claimed to play important functional roles in neural information processing and cognition, e.g. in learning and memory (Axmacher et al., 2010; Canolty and Knight, 2010; Lisman and Idiart, 1995; Lisman and Jensen, 2013). Furthermore, changes in CFC patterns have been linked to certain neurological and mental disorders such as Parkinson's disease (Lopez-Azcarate et al., 2010), schizophrenia (Allen et al., 2011; Moran and Hong, 2011) and social anxiety disorder (Miskovic et al., 2011). Therefore, CFC is potentially essential for normal brain function and understanding of CFC patterns can be crucial for diagnosing and eventually treating various disorders.

The classical analysis of CFC seems very straightforward (Figure 1) and is widely used. However, not all signatures of CFC as detected by this analysis method need to be due to interactions between different physiological processes occurring at different frequencies, as is commonly reported. It has been previously shown that signals with abrupt changes lead to spurious CFC results (Kramer et al., 2008) (see *Supplementary results: Examples of spurious CFC* for a related example). The roots of this problem are much more general. Let us take as an example the Van der Pol oscillator, which is a very simple non-linear relaxation oscillator. Conducting the CFC analysis on this oscillator would indicate that the phase of the low frequency components modulates the activity of the higher frequencies. However, despite strong CFC signal there is no simple physical interpretation for the different frequency components of the oscillator, and even less for their interaction. Indeed, any interpretation in terms of modulating or causally interacting frequencies is misleading as the spectral correlations are related to the non-linear characteristics of a single oscillator (see *Supplementary results: Examples of spurious CFC* for a thorough description of this example).

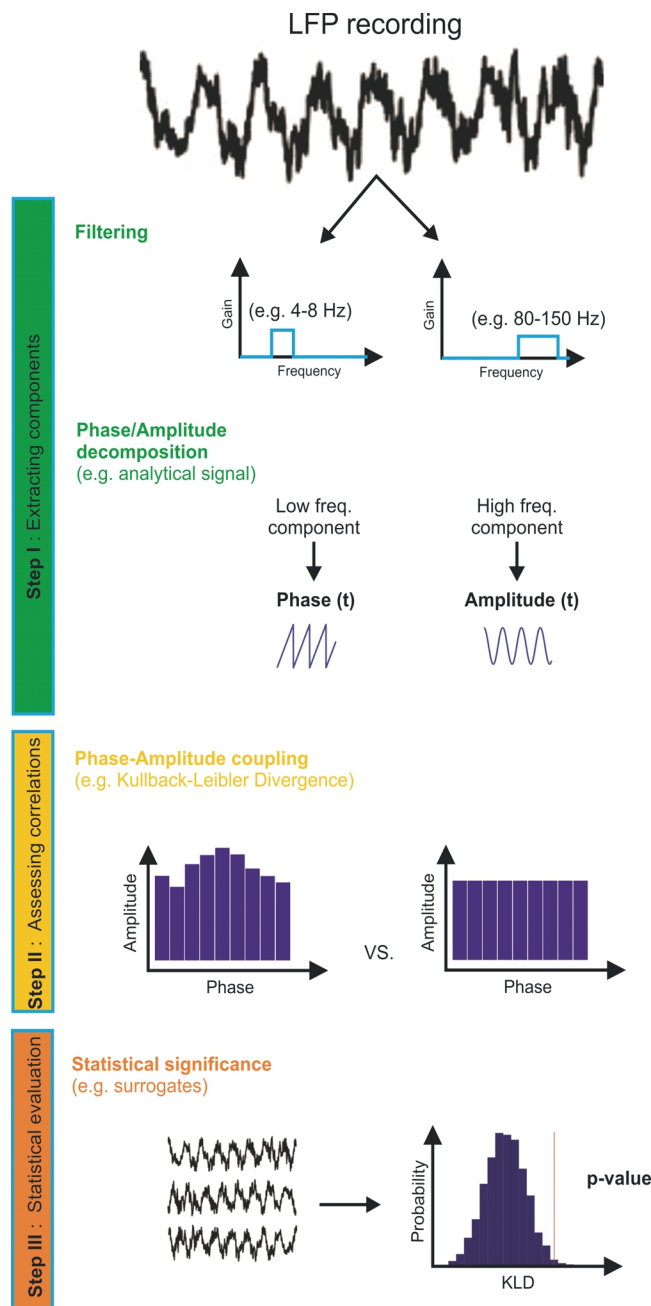


Figure 1. Typical approach to analyze phase-amplitude cross-frequency coupling. Step 1: Extracting the relevant components. This step is implemented by band-pass filtering and extraction of phase and amplitude dynamics for the relevant frequency bands. Step 2: Assessing correlations between components. This stage requires the computations of appropriate correlation or dependency measures between amplitude and phase. Therefore, a general measure of phase-amplitude coupling is to precisely quantify how much the histogram of mean amplitude versus phase deviates from a uniform distribution. Step 3: Statistical evaluation. Parametric or non-parametric approaches comparing to suitable surrogate data can be used to assign a p-value to the observed coupling strength.

This hints that the current analysis of CFC is inherently ambiguous regarding the nature and origin of the observed correlations between the frequency components. A significant CFC measure can be observed in case there are true modulations between subsystems oscillating at different frequencies. However, it can be also observed under very generic conditions that imply no coupling. Similarly to

the example above, any non-linear response where fast components are short-lived compared to the slow components of the signal would produce a significant CFC. In particular this means that current CFC measures of phase-to-amplitude coupling are not specific enough for one to automatically conclude, as it is almost invariably done in the literature, that the phase of a low frequency oscillation modulates the power of high-frequency activity. The same holds for CFC measures of amplitude-to-amplitude, or phase-to-phase coupling.

We wish to emphasize that we do not question the possible importance of CFC as a phenomenon. In fact, we believe that such a mechanism would be an elegant solution to several computational demands the brain has to cope with (Buzsaki, 2006). However, precisely because of the potential relevance of CFC for understanding the healthy and the pathological brain it is necessary to be aware of the pitfalls and misinterpretations in the methodology currently applied. We hope that the careful assessment of concerns will eventually strengthen the power of CFC analysis as an experimental tool.

The outline of the rest of the paper is as follows: First, we shall point out fundamental caveats and confounds in the current methodology of assessing CFC. Some of these points are original and some others have been known in other fields for years, yet all share the characteristic of being unattended in many of the current studies. Our literature review of the phase-amplitude CFC studies from the years 2010-2014 shows that these issues are relevant and timely (see *literature review* in the *Supplementary material*). Second, we propose an organization of different approaches to CFC analysis according to their biophysical interpretability and statistical inference approach. Finally, we outline some practical recommendations for CFC analysis.

Caveats and confounds of the CFC analysis

In this section we concentrate on what we call the classical CFC analysis – it is illustrated and explained in Figure 1. Any result of this analysis can be used to classify different conditions but only as a marker that is devoid of concrete and clear physiological interpretation. To give a physiological interpretation to CFC, one needs to know the set of potential mechanisms responsible for neural coupling. This set of mechanisms is only beginning to emerge (discussed below). We now discuss some main methodological confounds that make it difficult to build connections between the CFC measure and the underlying neurophysiological processes. More caveats and confounds and spurious examples can be found in the *Supplementary results*.

Instantaneous phase and amplitude: when are they meaningful and when not?

Standard phase-amplitude CFC analysis proceeds by first selecting two frequency bands followed by the computation of some index for the correlation or dependency between the phase of one band and the amplitude of the other (Fig. 1). The phase and amplitude values extracted from filtered signals can unfortunately only be interpreted in a meaningful way, i.e. as representing physiological oscillations, if a number of basic requirements are met. The same holds for CFC analyses based on them. In *Supplementary discussion: Conditions for a meaningful phase* we present a short but rather thorough review of the conditions that must be met for a meaningful interpretation of phase and amplitude values. The main conclusion is – not that surprisingly - that a clear peak in the power spectrum of the low frequency component is a prerequisite for a meaningful interpretation of any CFC pattern. Our literature review shows that even these well-known conditions were and are not always met in the literature, resulting in a strong over-interpretation of phase and amplitude (see *literature review* in the *Supplementary material*).

The importance of the bandwidth

The two components entering a phase-amplitude CFC analysis after filtering the signal, are determined by the center frequencies and bandwidths of the filters used to isolate them. Our literature review shows that majority of studies proceed by scanning the center frequencies for the phase and amplitude components while keeping a fixed bandwidth of a few Hz. However, this choice of bandwidth is important because it defines what is considered as a component and how the component's power or group phase changes in time (Figure 2). Thus, it is not the same thing to scan a center frequency from 20 Hz to 60 Hz with a bandwidth of 2 Hz or to consider at once the band centered at 40 Hz with a width of 21 Hz – different effects will be observed. Unfortunately, little or no justification is given to the choice of parameters in most analyses. The choice of bandwidth for the phase component is constrained by the condition of having a meaningful phase and is therefore often correctly chosen to be narrow. However, one also needs to be careful with the bandwidth size for amplitude - if the bandwidth of the higher frequency component (f_2) does not include the side peaks produced by the lower frequency (f_1), then CFC cannot be detected even if it is present (Figure 2A). Thus certain parameter values usually chosen in the literature can bias the CFC measures towards obtaining false negative results (see *Supplementary Discussion: The importance of the bandwidth*).

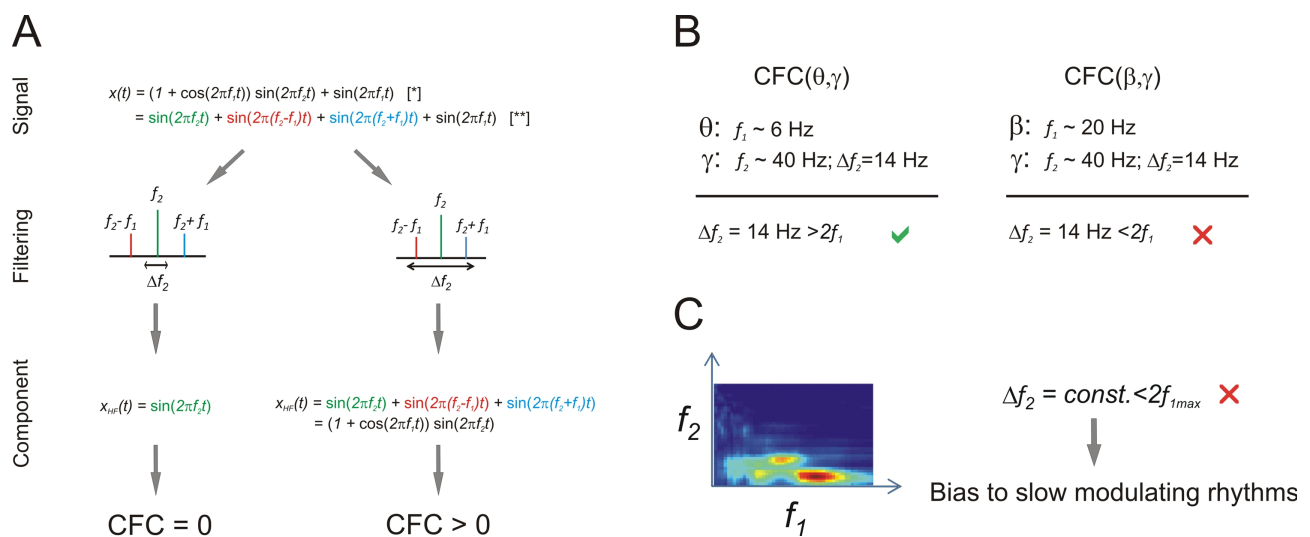


Figure 2. Mathematical decomposition and filtering bandwidth are key parameters to infer and interpret the presence of CFC. A) (Signal) The very same signal can be decomposed into different mathematically equivalent representations [*] or [**]. The choice of the representation leads to different interpretations regarding the interactions of the components (Filtering and Component). Different filtering bandwidths around the same frequency can lead to different results depending on whether the bandwidth includes modulating sidebands or not. B) For a fixed bandwidth of the modulated frequency, only a range of modulating frequencies can be captured. For example, in the simplified harmonic case, a bandwidth of 14 Hz around a frequency of 40 Hz would allow detection of a potential modulation from a 6 Hz rhythm, but not from a 20 Hz oscillation. C) When scanning the modulating (f_1) and modulated (f_2) frequencies, a fixed bandwidth biases CFC analysis and favors low frequencies for f_1 .

Non-stationarity and spectral correlations: two sides of the same coin

Most neuronal signals that we measure are non-stationary. Time-varying sensory stimuli, top-down influences, neuromodulation, endogenous regulatory processes and changes in global physiological states render neuronal dynamics non-stationary. In contrast to a stationary process, a non-stationary process in general exhibits spectral correlations between components of its Fourier expansion (Lii and Rosenblatt, 2002). These correlations may be misinterpreted as CFC. The underlying reason for these spectral correlations is that in constructing the spectrum, we decompose a process which is by definition not time-invariant (non-stationary) into the eigenvectors of the time-shift operator, i.e., the complex exponentials in the Fourier expansion. Therefore, in the non-stationary case, there are two possible scenarios leading to positive CFC measures:

One scenario is that physiological processes indeed interact. This interaction then leads to non-stationarities, and at the same time we observe spectral correlations in the Fourier representation. For example, if the phase of a neural input oscillating at theta frequency modulates the amplitude of local gamma oscillations, both obtained from the same LFP recording, the statistical properties of the gamma oscillation amplitude series will change in time, as does theta phase. Specifically, their

properties will vary in time only to be repeated after a full cycle of the slow oscillation, and thus exhibiting a particular type of non-stationarity called cyclo-stationarity (Gardner et al., 2006).

The other and problematic scenario is that *unspecific* non-stationarities (that is, *any kind of* change of the statistical properties of the signal), not related to or caused by coupling of neural processes, will also be reflected in spectral correlations which could be over-interpreted as the result of causal interactions among frequency specific neuronal processes. This second scenario can occur if non-stationary input to a given area simultaneously affects the phase of a low frequency component and increases high-frequency activity (common drive to frequency components of the same signal). For example, typical evoked potentials affect a broad range of frequency components (Makeig et al., 2002). In this case, high-frequency amplitude increases occur preferentially for certain phases of slow oscillations even without any need of interaction between the two rhythms.

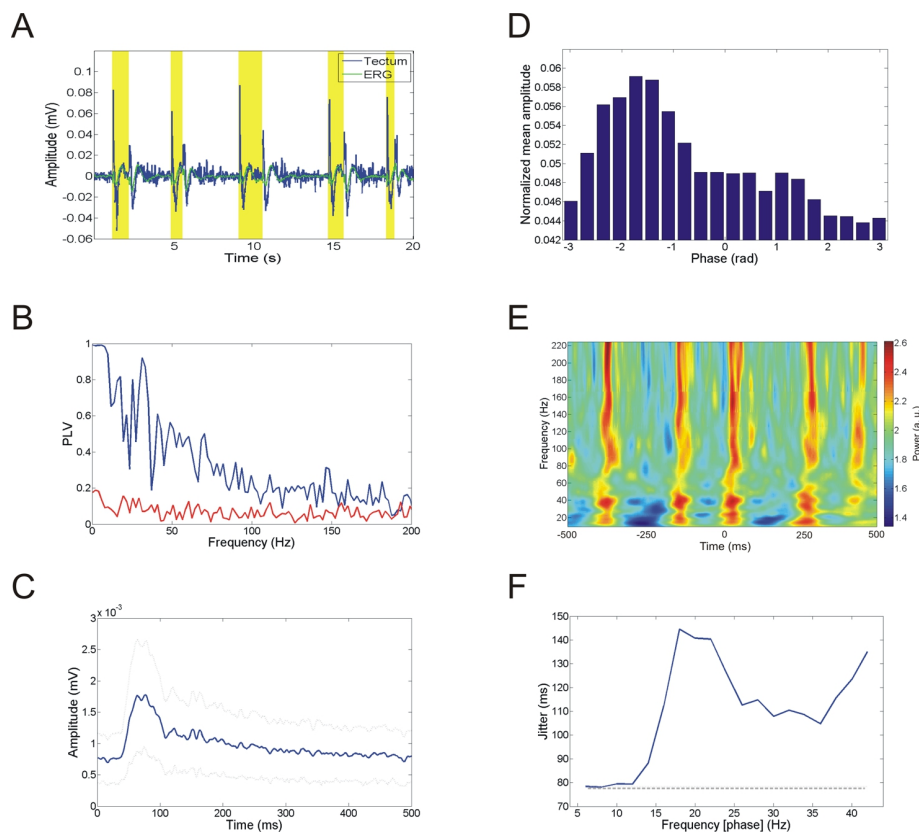


Figure 3. A) Time series of stimulus and recorded activity in the retina and optic tectum. The rest of the panels are based on the electroretinogram only but similar results hold for the optic tectum. B) Phase locking values as a function of frequency for a period following (blue) or preceding (red) the onset of visual stimulation. C) Amplitude of high-frequency activity (80-200 Hz) locked to stimulus onset (mean is represented in blue while mean plus/minus standard deviation are in grey). D) Amplitude-phase histogram for the whole time series (100 seconds). E) Power locked to the trough of a low-frequency band (4-8 Hz). F) Jitter from the peak of high-frequency (80-200 Hz) power to stimulus onset (71 ms; grey line) is always smaller than to the trough of the phase of lower frequency bands closest to their maximal power (blue line). This means that most likely input simultaneously affected both the phase of the low frequency component and the amplitude of the high frequency component, thus generating a false-positive CFC.

To illustrate the latter case, we analyzed the CFC patterns of the electroretinogram (ERG) and LFP recorded in the optic tectum of a turtle (*Pseudemys scripta elegans*) during visual stimulation (Figure 3A and *supplementary methods*). The comparison of phase locking values for windows placed before and after the stimulus onset revealed an increase of phase locking preferentially for lower frequencies (Figure 3B). At the same time, the stimulus elicited an increase in power, including the power of high-frequency bands, 80-200 Hz (Figure 3C). The combination of increased activity at high-frequencies and phase-locking to the stimulus of lower frequencies is sufficient to obtain significant measures in standard CFC analysis (Figure 3D-E). However, we can gain more insight into this issue by studying the difference in relevant locking strengths: if the high-frequency amplitude were indeed modulated by the phase of slower oscillations, we should expect that the locking between these frequency components exceeds the one between the amplitude and the stimulus onset. For our recorded data it turned out that the jitter between the peak of the amplitude of high-frequency activity and stimulus onset was always smaller than the jitter between the peak amplitude and any phase of the slower rhythm (Figure 3F; we also replicated these results in a human intracranial recording: see *Supplementary Figure 3*). This indicates that it was not the phase of the slower rhythm that controlled the amplitude of the high-frequency activity, but simply the input which affected both the phase of the low frequency oscillation and the amplitude of the high frequency activity.

Therefore, the key issue is to distinguish whether the observed phase-amplitude correlation between two bands is due to common effect, generated by external or internal input or whether the correlation is due to a causal interaction between rhythms (which, of course, could also be triggered by the input). Recently, a new approach (Voytek et al., 2013) has been developed to measure transient phase-amplitude coupling directly in an event-related manner. Whereas ideally, their approach of analyzing phase-amplitude relations with respect to the stimulus onset should avoid some event-related artifacts, it is questionable whether the marker actually works as intended (see *Supplementary results: Phase-amplitude coupling for event-related potentials*). Ultimately solving these questions requires a formal causal analysis between the spectral variables of different bands (see also *Supplementary discussion: Causality methods*).

The above observations suggest that non-stationary input to a given area can generate correlations between bands, which are not necessarily a signature of interactions between these bands. The argument goes well beyond the relationship between sensory stimulation and CFC in sensory areas: if a brain area under a recording electrode receives time-varying input from any other brain area,

this input might generate similar dependencies across frequency components (Figure 4A). The problem is that usually one has no control over the timing of the internal input to the examined brain area (Figure 4B). Thus, phase-amplitude coupling measured anywhere in the brain can be potentially explained by common influence on the phase and amplitude, without the phase of a low frequency oscillation modulating the power of high frequency activity.

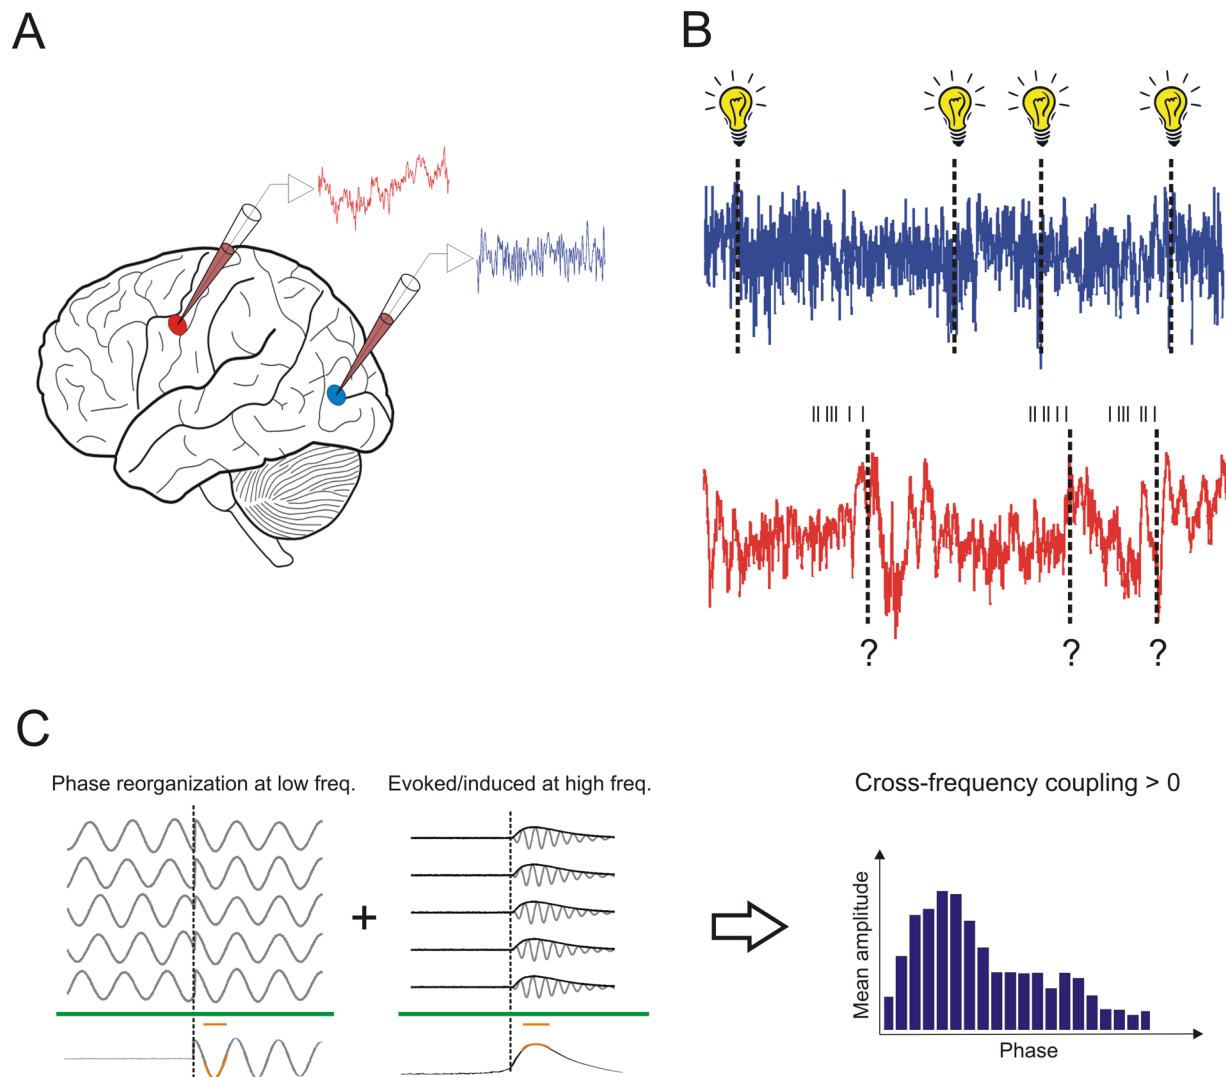


Figure 4. Illustration of how time-varying input can lead to false-positive CFC. A) Illustration of recordings in a cortical sensory area (blue) and a higher cortical area (red). B) Importantly, both can be subject to non-stationary neuronal input but its timing can only be determined for the sensory area. C) Knowing the input timing is necessary to perform phase locking analysis, distinguish between evoked/induced responses and to disambiguate the origin of CFC. Without this additional information the results of the CFC analysis remain ambiguous.

In conclusion, the above considerations imply that the current phase-amplitude CFC measure is constitutive for the non-stationary responses of driven systems and therefore is not a very specific marker of biophysical coupling. From a mathematical perspective the key aspect is that any consistent response to input, whatever its shape, implies a certain phase locking between its

different Fourier components (Gardner, 1986; Mitra & Bokil, 2007). Thus, if the power of any of the fast components lasts a bit more or a bit less than the period of a slow component, then its amplitude will accumulate preferably at certain phases of the slow component. This is all that is needed to give rise to phase-amplitude CFC, as measured for example by the modulation index (Tort et al., 2010). It is therefore necessary to recognize that beyond the phase-amplitude CFC index, which is just an index at the signal level, additional information about how the process under study reacts to input and the statistics of the input itself are needed to better resolve the origin of such correlations. Analyses of surrogate data can help to remove some of the ambiguity but only offers partial remedies to the problem, as we shall discuss next.

Surrogate data: none are perfect but some are better than others

After some index of CFC has been estimated one needs to rely on statistical inference to reach a conclusion about the statistical significance of the measure. Currently, most studies of CFC rely on the frequentist approach of using surrogate data to estimate a p-value. Some issues related to generation of the surrogate data are discussed here. For the state of some Bayesian approaches see Figure 5 in the section “*Organization of modeling/statistical approaches to CFC*”.

A suitable surrogate construction should only destroy the specific cyclo-stationarities related to the hypothesized CFC effect, while keeping all the unspecific non-stationarities and non-linearities of the original data. Often, it is impossible to construct perfect surrogates that selectively destroy the effect of interest but some approaches are more conservative than others. In our context a sensible requirement is to construct surrogates that minimize the distortion of both phase and amplitude dynamics for each frequency component. If data are organized in detectable repetitive events such as trials locked to an external stimulus or to saccades, shuffling the full/intact phase or amplitude components between different events seems the most straightforward approach. Unfortunately, the very existence of an event-related potential implies that some frequency components are both locked to the event and between themselves. Consequently, this strategy to obtain surrogates alone cannot discern the source of modulation. Future developments of methods to partial out the common drive effect of the event could help to test the significance of a direct phase-amplitude modulation.

Finding appropriate surrogate data for a single continuous stream of data comes with its own challenges. For example, phase scrambling does not meet the minimal distortion criterion, as the generated surrogate data are fully stationary after scrambling, i.e. the non-stationarities of interest

and the unspecific non-stationarities are both destroyed alike (Nakamura et al., 2006). A significantly larger CFC in the original data may in this case be due to the removal of non-stationarities not specifically related to physiological CFC. Another approach has been using block re-sampling (Canolty et al., 2006) where one of the continuous time series (i.e., the instantaneous phase) is simultaneously cut at several points and the resulting blocks permuted randomly. This method suffers again from destroying in excess the non-stationary structure of the original series. More conservative surrogates can be obtained by minimizing the number of blocks by cutting at single point at a random location and exchange the two resulting time courses (Canolty et al., 2006). Repeating this procedure leads to a set of surrogates with a minimal distortion of the original phase dynamics.

Thus, while perfect surrogate data that selectively disrupt phase-amplitude coupling might be impossible to build (as is the case for most types of non-linear interactions) conservative approaches that minimize distortion of phase and amplitude dynamics can reduce the number of false positives.

CFC modulation across conditions

Several studies have reported significant changes in phase-amplitude CFC with variations of experimental parameters or across two different conditions. The modulation of CFC by the task or experimental condition has then been taken as an indication of its physiological role (Axmacher et al., 2010; Canolty and Knight, 2010; Tort et al., 2009; Voytek et al., 2010). However, for now there is only little reason to believe that this modulation could not be due to side-effects of more basic changes between the conditions.

Since the power of bands directly influences the range within which they can modulate or be modulated, it is possible that changes in CFC correlations are a direct consequence of changes in power spectra. For example, changes in the observed CFC can have their origins in the fact that power changes affect the signal-to-noise ratio of phase and amplitude variables and their correlations (e.g.- Muthukumaraswamy & Singh, 2011). It is thus necessary to control whether correlations between CFC and other behavioral or physiological variables might be simply due to changes in, for example, the strength or frequency of oscillations. Unfortunately, our literature review shows that in around half of the reviewed studies where conditions are compared, changes in the power spectrum across the conditions are not considered. If the data permit, it is therefore highly recommended to rely on stratification techniques (e.g. Schoffelen et al., 2005) to obtain a subset of matched trials in which the distribution of power across trials is identical for both the phase and amplitude frequency

bands in the two conditions to be compared.

In general, as CFC is a statistic based on the correlation of certain variables, it is necessary to control for the explanatory power of these variables themselves, before a specific role for the correlation might be distilled.

Organization of modeling/statistical approaches to CFC

Until now we have focused on what we call the classical approach (see Figure 1) to assess phase-amplitude CFC effects consisting of: i) isolating frequency components, ii) assessing their dependencies, and iii) computing p-values based on surrogates. We have described how difficult it is at this stage to draw any conclusions about the biophysical mechanisms underlying these measures. However, different frameworks exist to assess relationships among rhythmic processes from experimental time series. A short description of some frameworks can be found in the *Supplementary discussion*. For the purpose of understanding the role of different frameworks in gaining physiological understanding of CFC, we have found it useful to organize them according to their biophysical interpretability and statistical inference approach (Figure 5).

We believe the location of the method along those axes (Figure 5) has to be taken into account to avoid over-interpretations of the results of a CFC analysis. The first section of this paper can be in fact seen as an explanation as to why some models including the classical approaches are positioned very close to the “marker” section. For example, a simple correlation-based quantifier as provided by classical approaches is probably all that is needed, if the sole purpose of CFC analysis is to have a marker to classify different conditions (e.g. disease states). If one insists that this marker should be more specific than say just changes in the power-spectrum, already more work is needed, as discussed above. Finally, if the aim is to attach a well-defined physiological meaning to the observed CFC pattern, it is imperative to have either a generative model or additional external information, such as obtained from a direct perturbation of the putative physiological CFC mechanism, to link signal and underlying processes. Both of these latter approaches require a biophysical theory to be put forward as to how a neuron or an ensemble of neurons physically implements the coupling. For the moment, potential biological mechanisms of cross-frequency coupling are only starting to be discovered.

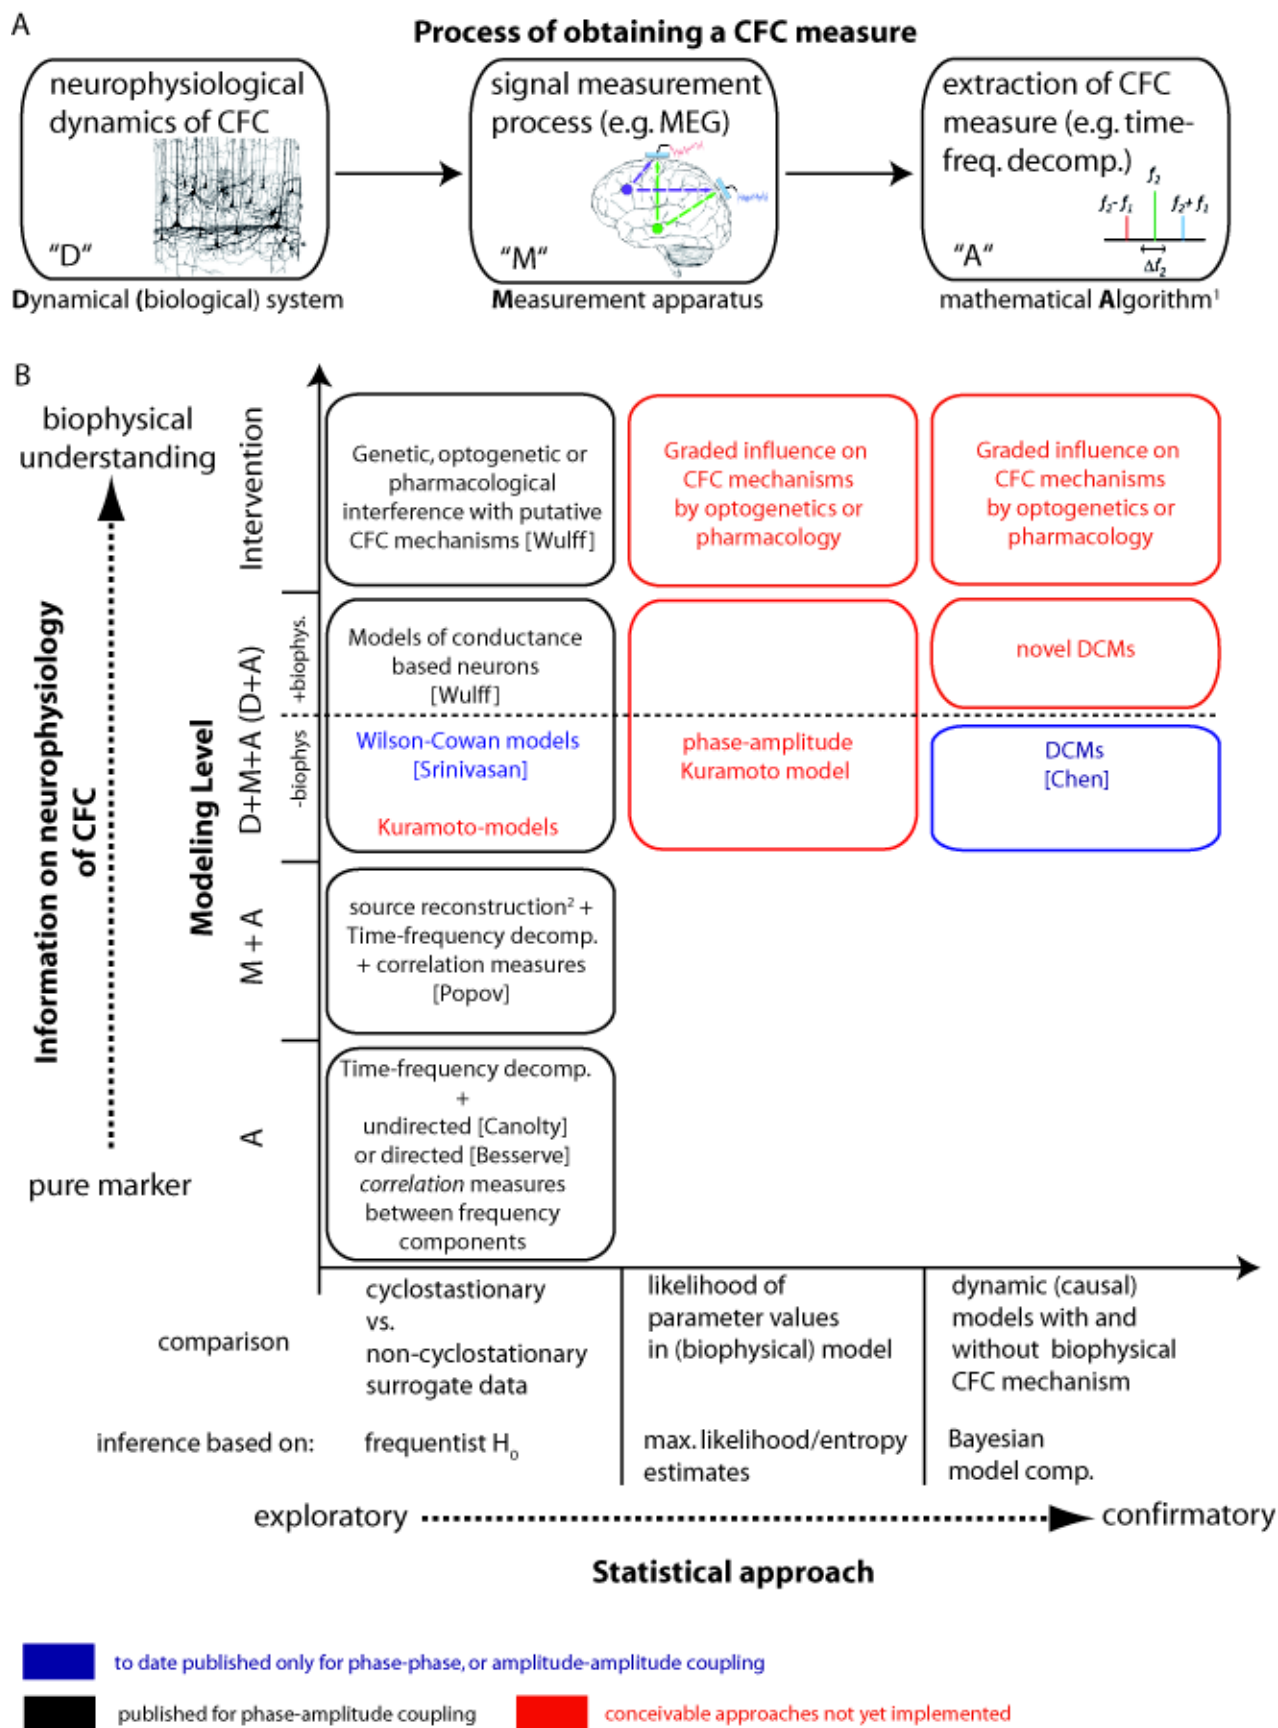


Figure 5. Organization of approaches to CFC. (A) The process by which a researcher obtains a CFC measure. “D” signifies the biological dynamical system (that may or may not have biological CFC). “M” signifies the physical transduction and measurement process, including all physical (unavoidable) filtering distortion, instrument noise and potentially mixing processes. The measurement process can generate CFC in absence of biological CFC. “A” signifies the mathematical algorithm applied to the measured data to obtain a figure of merit for CFC. Typically this step involves filtering or a time-frequency decomposition, and a linear or nonlinear correlation measure. As shown in the main text, also this process can give rise to CFC in absence of biological coupling, e.g. by ignoring the limits of time-frequency analysis in the face of non-stationarities. (B) Two-dimensional organization of CFC approaches. The x-axis sorts the approaches by the statistical inference technique that is used. Frequentist H_0 based approaches just test for presence of absence of CFC in the measured data, while maximum likelihood or Bayesian approaches perform inference on coupling parameters in models – and hence only appear for dynamic models with coupling parameters. The y-axis indicates the part(s) of the process in (A) that are modelled: “A” – just extracting a CFC measure from measured data is equivalent to modelling this extraction process itself. “M+A” – the generative model now comprises the measurement process and the measure extraction. “D+M+A” – the modelling comprises a dynamic model of the biological process, either via a dynamic proxy process that models only the bare essentials of the dynamics (like a Kuramoto model for phase-phase coupling), or has biological detail (like a Hodgkin-Huxley model with some explicit mechanisms implementing CFC). Approaches that were only published for phase-phase or amplitude-amplitude coupling are highlighted in blue text, conceivable approaches, that have to our knowledge not been implemented at all yet are highlighted in red.

(¹) The fact that we use some numbers (measured data) and feed them to a mathematical algorithm to obtain other numbers (the CFC measure) can be modelled by just executing the algorithm. Nevertheless it is a part of the CFC process model.

(²) By source reconstruction we mean any inversion of the measurement process, e.g. unmixing via ICA, electromagnetic source reconstruction in electro- or magnetoencephalography, or removal of measurement noise.

We provided some representative references on the figure: Besserve, M., Scholkopf, B., Logothetis, N.K. & Panzeri, S. Causal relationships between frequency bands of extracellular signals in visual cortex revealed by an information theoretic analysis. *Journal of Computational Neuroscience* **29**, 547-566 (2010). Canolty, R.T., *et al.* High gamma power is phase-locked to theta oscillations in human neocortex. *Science* **313**, 1626-1628 (2006). Chen, C.C., *et al.* A dynamic causal model for evoked and induced responses. *Neuroimage* **59**, 340-348 (2012). Popov, T., Steffen, A., Weisz, N., Miller, G.A. & Rockstroh, B. Cross-frequency dynamics of neuromagnetic oscillatory activity: Two mechanisms of emotion regulation. *Psychophysiology* **49**, 1545-1557 (2012). Srinivasan, R., Thorpe, S. & Nunez, P.L. Top-down influences on local networks: basic theory with experimental implications. *Front Comp Neuroscience* **7**, 1:15 (2013). Wulff, P., *et al.* Hippocampal theta rhythm and its coupling with gamma oscillations require fast inhibition onto parvalbumin-positive interneurons. *Proc Natl Acad Sci U S A* **106**, 3561-3566 (2009).

Indeed, whilst there is extensive knowledge about the physiological mechanisms responsible for different frequency components (Buzsaki, 2006), not much is known about the cellular and network mechanisms of the *interactions* between these components (Canolty and Knight, 2010). Only recently some evidence about concrete mechanisms of interaction has been obtained from intervention studies in physiological systems and computational models. For example, by using transgenic mice it has been shown on the level of LFPs that hippocampal theta-gamma coupling depends on fast synaptic inhibition (Wulff *et al.*, 2009) and NMDA receptor-mediated excitation of parvalbumin-positive interneurons (Korotkova *et al.*, 2011). CFC at the LFP level has also been observed between alpha oscillations at the infragranular layer and gamma activity at the supragranular layers (Spaak *et al.*, 2012). Recently, it was also shown that feedback inhibition enables CFC at the level of membrane potential fluctuations (Pastoll *et al.*, 2013). Thus, biological

mechanisms for CFC can occur at the population level, at the single neuron level, or both. The most parsimonious explanation to account for these findings is that the low frequency oscillation reflects periodic fluctuations of the membrane potential and thus excitability, which in turn gate the occurrence of higher frequency activity in a phase specific manner (Pastoll et al., 2013). Typically this higher frequency activity reflects spikes which could display a rhythmic pattern (possibly reflected as gamma oscillations).

Along plausible biological mechanisms, one also needs to take into account possible alternative explanations for non-zero CFC measures. In this perspective we have tried to aid the interpretation of CFC by demonstrating that although CFC patterns are typically interpreted as reflecting physiological coupling, they can be also generated 1) by biological processes unrelated to direct coupling between different neural processes, and 2) by methodological pitfalls. In the following section we compiled some practical recommendations to help to avoid some of these errors. Taking these alternative explanations into account and controlling for them experimentally will help towards a clear interpretation of CFC results.

When intervention is not possible another principled approach to test for the presence of biophysical CFC is a formal comparison of computational models that do or do not incorporate biophysical CFC mechanisms with respect to their ability to explain the observed data. This could be done for example using sufficiently detailed Dynamical Causal Models and Bayesian model comparison (Friston and Dolan, 2010). If neither intervention nor formal model comparison are feasible, the researcher would have to limit the interpretation of observed CFC patterns (see Fig. 5) to that of a marker.

To further exemplify the hierarchy of approaches illustrated in Figure 5, we turn to a different, well-known measure -spectral power analysis. For example, measures of LFP power in the gamma band can simply be used as markers to classify different conditions (lower left corner in Fig. 5B). However, we also have several reasonable biophysical models (conceptual and computational) about the generating mechanisms of hippocampal and cortical gamma oscillations. These mechanisms were ultimately identified by interventional approaches (pharmacological, genetic, lesion) in a variety of physiological systems, as well as in computational models (Buzsaki and Wang, 2012). This means the field of spectral power analysis can draw on interventional approaches as well as formal model comparisons (upper row / right column of Fig. 5B). Nevertheless, we note that even for the mature field of spectral power analysis a change in LFP power in the gamma band can be due to several biophysical mechanisms (change in the number of neurons engaged in

oscillations, their synchrony, etc) and these are not mutually exclusive. However, we can still map changes in gamma activity to a limited number of mechanistic options, each of them relatively well understood. Furthermore, since we are aware of several alternative explanations such as eye-movement artifacts that might contribute to gamma-band power changes we can design experiments to control for them (Yuval-Greenberg et al., 2008). We believe that similar steps will be required before CFC patterns in a signal can be confidently linked to any concrete biophysical mechanism.

Practical recommendations

As previously discussed we will need progress in several directions to establish phase-amplitude or other types of CFC as a fundamental mechanism in coordinating neuronal activity. Together with experimental and modeling advances, stricter standards in the use of CFC metrics are also necessary. Below we list practical recommendations to avoid some of the mentioned caveats and increase the specificity of the most popular phase-amplitude CFC metric (see Fig. 1). Rather than a comprehensive algorithm this list should be thought as a check list that should help to minimize technical pitfalls and over-interpretation of phase-amplitude CFC measures in macroscopic signals.

- 1 *Presence of oscillations.* Signatures of oscillatory processes with clear peaks in a time-resolved power spectrum are indispensable prerequisites. The frequency component for defining the instantaneous phase should include one of the peaks.
- 2 *Selection of bandwidths.* The frequency band used to define the instantaneous phase should isolate energy associated with the oscillatory component of interest. If the center frequency is relatively stable a natural choice for the bandwidth can be directly obtained from the width of the corresponding peak in the power spectrum. The latter can be estimated by subtracting from the real power spectrum that of the background (defined by either a fit or proper baseline). Note that the band defining the instantaneous amplitude at the higher frequency must be large enough to fit the sidebands caused by the assumed modulating lower frequency band (figure 2A) and the lower frequency band should be narrow enough to define a meaningful phase. Therefore, adaptive rather than fixed bandwidths might be necessary when scanning the modulating frequency in explorative analyses.
- 3 *Interpretation of instantaneous phase.* A meaningful interpretation of instantaneous phase requires its monotonic growth in time. The presence of phase slips or reverses (also observed as negative instantaneous frequencies) must be checked and justified.
- 4 *Precision.* The precision of the method used to assign an instantaneous phase and amplitude

to a signal should be determined for each analysis. The precision of the computation of the Hilbert transform of a signal $s(t)$ can be estimated from the variance of $s(t)+H^2(s(t))$, which analytically should be identical to zero. Given the non-locality of Hilbert (or wavelet) transforms, edge effects can be severe. It is recommended to discard at least a few characteristic periods of the signal at the beginning and end of each segment of interest.

5 *Testing for non-linearities.* Non-linear responses to input or nonlinearities during the signal transduction can contribute to phase-amplitude CFC. The presence of harmonics in the signal should be tested by a bicoherence analysis and its contribution to CFC should be discussed. Partialization of phase-phase and amplitude-amplitude coupling is necessary to assess the role of non-linearities in generating spectral correlations (see also *Supplementary results*).

6 *Testing for input-related non-stationarities.* When the timing of neuronal input to the recorded area is available, an analysis of relative locking between phase, amplitude and input can inform about the origin of the correlations.

7 *Temporal structure.* Information about the temporal structure of the putative interaction (e.g., sustained during many cycles versus a transient coupling) can be helpful to better characterize a presumed CFC and to disambiguate its origins. The modulation index only offers an average measure of CFC by computing the distance from a phase-amplitude histogram to a uniform distribution. However, such histogram can be used to identify the phase at which the average amplitude of high-frequency activity is maximal. The time series obtained by sampling the amplitudes of high-frequency activity at that particular phase can be used to provide some information about the temporal dynamics of the coupling.

8 *Surrogates.* Surrogate data should be created that minimally interfere with the phase and amplitude dynamics. For continuous recordings random point block-swapping is preferred over phase scrambling or cutting at several points.

9 *Specificity of effects.* Differences of CFC indices across conditions should be controlled for the differences in power at the presumed bands of interaction. The specific role of the coupling can be better assessed once the explanatory level of the power spectrum has been accounted for: When trial-based measures are available, stratification techniques should be used to compare subsets of trials for which the distributions of power at the bands of interest are identical across the conditions.

Conclusions

Cross-frequency coupling (CFC) might be a key mechanism for the coordination of neural dynamics. Several independent research groups have observed cross-frequency coupling and related

it to information processing, most notably to learning and memory (Canolty and Knight, 2010; Lisman and Jensen, 2013). Recently, CFC has also been used to investigate neurological and psychiatric disorders (Allen et al., 2011; Lopez-Azcarate et al., 2010; Miskovic et al., 2011; Moran and Hong, 2011). Thus, CFC analysis is potentially a promising approach to unravel brain function and some of their pathologies.

In the present manuscript we have reviewed some confounds that hamper phase-amplitude CFC analysis. Importantly, these confounds have not been considered in a significant percentage of recent publications and may have contributed to over-interpretations. This is a serious issue that needs to be resolved because CFC analysis is potentially a powerful tool to reveal fundamental features of neural computations. An obvious first step is to adopt stricter standards and canonical procedures for CFC analysis. To this end we suggested a – probably incomplete – list of controls that should be routinely checked. We have also attempted to organize the current modeling/statistical approaches to CFC in order to better identify their respective advantages and pitfalls and to point out where further methodological advances are required.

Acknowledgements

This work was supported by the Max Planck Society. RV, VP and MW were supported by LOEWE grant "Neuronale Koordination Forschungsschwerpunkt Frankfurt (NeFF)". RV also acknowledges the support from the Hertie Foundation and the projects PUT438 and SF0180008s12 of the Estonian Research Council.

References

- Allen, E.A., Liu, J., Kiehl, K.A., Gelernter, J., Pearlson, G.D., Perrone-Bizzozero, N.I., and Calhoun, V.D. (2011). Components of cross-frequency modulation in health and disease. *Front Syst Neurosci* 5, 59.
- Axmacher, N., Henseler, M.M., Jensen, O., Weinreich, I., Elger, C.E., and Fell, J. (2010). Cross-frequency coupling supports multi-item working memory in the human hippocampus. *Proc Natl Acad Sci U S A* 107, 3228-3233.
- Bragin, A., Jando, G., Nadasdy, Z., Hetke, J., Wise, K., and Buzsaki, G. (1995). Gamma (40-100 Hz) oscillation in the hippocampus of the behaving rat. *J Neurosci* 15, 47-60.
- Buzsaki, G. (2006). *Rhythms of the Brain* (New York: Oxford University Press).
- Buzsaki, G., and Wang, X.J. (2012). Mechanisms of gamma oscillations. *Annu Rev Neurosci* 35, 203-225.

- Canolty, R.T., Edwards, E., Dalal, S.S., Soltani, M., Nagarajan, S.S., Kirsch, H.E., Berger, M.S., Barbaro, N.M., and Knight, R.T. (2006). High gamma power is phase-locked to theta oscillations in human neocortex. *Science* 313, 1626-1628.
- Canolty, R.T., and Knight, R.T. (2010). The functional role of cross-frequency coupling. *Trends Cogn Sci* 14, 506-515.
- Friston, K.J., and Dolan, R.J. (2010). Computational and dynamic models in neuroimaging. *Neuroimage* 52, 752-765.
- Gardner, W.A. (1986). The spectral correlation theory of cyclostationary time series. *Signal Processing*, 11, 13–36.
- Gardner, W.A., Napolitano, A., and Paura, L. (2006). Cyclostationarity: Half a century of research. *Signal Processing* 86, 639-697.
- Korotkova, T., Fuchs, E.C., Ponomarenko, A., von Engelhardt, J., and Monyer, H. (2011). NMDA receptor ablation on parvalbumin-positive interneurons impairs hippocampal synchrony, spatial representations, and working memory. *Neuron* 68, 557-569.
- Kramer, M.A., Tort, A.B., and Kopell, N.J. (2008). Sharp edge artifacts and spurious coupling in EEG frequency comodulation measures. *J Neurosci Methods* 170, 352-357.
- Lii, K.S., and Rosenblatt, M. (2002). Spectral analysis for harmonizable processes. *Annals of Statistics* 30, 258-297.
- Lisman, J.E., and Idiart, M.A. (1995). Storage of 7 +/- 2 short-term memories in oscillatory subcycles. *Science* 267, 1512-1515.
- Lisman, J.E., and Jensen, O. (2013). The theta-gamma neural code. *Neuron* 77, 1002-1016.
- Lopez-Azcarate, J., Tainta, M., Rodriguez-Oroz, M.C., Valencia, M., Gonzalez, R., Guridi, J., Iriarte, J., Obeso, J.A., Artieda, J., and Alegre, M. (2010). Coupling between beta and high-frequency activity in the human subthalamic nucleus may be a pathophysiological mechanism in Parkinson's disease. *J Neurosci* 30, 6667-6677.
- Makeig, S., Westerfield, M., Jung, T. P., Enghoff, S., Townsend, J., Courchesne, E., & Sejnowski, T. (2002). Dynamic brain sources of visual evoked responses. *Science*, 295(5555), 690-694.
- Miskovic, V., Moscovitch, D.A., Santesso, D.L., McCabe, R.E., Antony, M.M., and Schmidt, L.A. (2011). Changes in EEG cross-frequency coupling during cognitive behavioral therapy for social anxiety disorder. *Psychol Sci* 22, 507-516.
- Mitra, P., & Bokil, H. (2007). *Observed Brain Dynamics*. Oxford University Press.
- Moran, L.V., and Hong, L.E. (2011). High vs low frequency neural oscillations in schizophrenia. *Schizophr Bull* 37, 659-663.
- Muthukumaraswamy, S. D., & Singh, K. D. (2011). A cautionary note on the interpretation of phase-locking estimates with concurrent changes in power. *Clinical Neurophysiology*, 122(11), 2324-2325.
- Nakamura, T., Small, M., & Hirata, Y. (2006). Testing for nonlinearity in irregular fluctuations with long-term trends. *Physical Review E*, 74(2), 026205.
- Pastoll, H., Solanka, L., van Rossum, M.C.W., and Nolan, M.F. (2013). Feedback Inhibition Enables Theta-Nested Gamma Oscillations and Grid Firing Fields. *Neuron* 77, 141-154.
- Schoffelen, J. M., Oostenveld, R., & Fries, P. (2005). Neuronal coherence as a mechanism of effective corticospinal interaction. *Science*, 308(5718), 111-113.
- Spaak, E., Bonnefond, M., Maier, A., Leopold, D.A., and Jensen, O. (2012). Layer-Specific Entrainment of Gamma-Band Neural Activity by the Alpha Rhythm in Monkey Visual Cortex. *Current Biology* 22, 2313-2318.
- Tort, A.B., Komorowski, R., Eichenbaum, H., and Kopell, N. (2010). Measuring phase-amplitude coupling between neuronal oscillations of different frequencies. *J Neurophysiol* 104, 1195-1210.
- Tort, A.B., Komorowski, R.W., Manns, J.R., Kopell, N.J., and Eichenbaum, H. (2009). Theta-gamma coupling increases during the learning of item-context associations. *Proc Natl Acad Sci U S A* 106, 20942-20947.
- Voytek, B., Canolty, R.T., Shetyuk, A., Crone, N.E., Parvizi, J., and Knight, R.T. (2010). Shifts in gamma phase-amplitude coupling frequency from theta to alpha over posterior cortex during visual

tasks. *Front Hum Neurosci* 4, 191.

Voytek, B., D'Esposito, M., Crone, N., and Knight, R.T. (2013). A method for event-related phase/amplitude coupling. *Neuroimage* 64, 416-424.

Wulff, P., Ponomarenko, A.A., Bartos, M., Korotkova, T.M., Fuchs, E.C., Bahner, F., Both, M., Tort, A.B., Kopell, N.J., Wisden, W., and Monyer, H. (2009). Hippocampal theta rhythm and its coupling with gamma oscillations require fast inhibition onto parvalbumin-positive interneurons. *Proc Natl Acad Sci U S A* 106, 3561-3566.

Yuval-Greenberg, S., Tomer, O., Keren, A.S., Nelken, I., and Deouell, L.Y. (2008). Transient induced gamma-band response in EEG as a manifestation of miniature saccades. *Neuron* 58, 429-441.

Supplementary Materials

Materials and Methods

Analysis of CFC

All analyses were performed using custom built MATLAB routines. To extract the frequency components, all signals were band-filtered with a two-way least-squares FIR filter (eegfilt.m from the EEGLAB toolbox [1]). For each component of interest, instantaneous phase and amplitude were estimated by the analytical signal approach. The Hilbert transform was applied to the filtered signals in order to define the imaginary part of the complex-valued time series (with real part being the filtered signal). The polar coordinates of these analytical signals define the instantaneous amplitudes and phases.

The power plots locked to phase-troughs of slow components used to create Figure 3 (panel E), Supplementary Figure 1 (bottom panel), Supplementary Figure 2 (panel C), and Supplementary Figure 4, were implemented as described in the supporting information of Ref.[2].

The modulation index used in Figure 5 (bottom panels), Supplementary Figure 2 (panel D), and Supplementary Figure 4 (bottom panels) followed the original formulation by Tort et al. [3]. All modulation indices and histograms of mean amplitudes were computed for 20 equally-sized bins for the phase variable.

Data collection for Figure 3 in the main text

We recorded neuronal activity from the retina and the tectum of the turtle (*Pseudemys scripta elegans*) to analyze the effect of stimulus input on cross-frequency coupling.

Preparation. Experiments were approved by the German local authorities (Regierungspraesidium, Hessen, Darmstadt). One turtle (*Pseudemys scripta elegans*) was anesthetized with 15 mg Ketamine, and 2 mg Medetomidinhydrochloride and decapitated. The entire brain with the eyes attached was removed as described in Ref.[4]. The brain was placed in a petri dish and superfused with oxygenated ringer. The ringer consisted of (in mM) 96.5 NaCl, 2.6 KCl, 2.0 MgCl₂, 31.5 NaHCO₃, 20 D-glucose, 4 CaCl₂ at pH 7.4 and was administered at room temperature (22 C).

Electrophysiological recordings. The electroretinogram was recorded with a chlorided silver wire in a Vaseline well that was built around the right eye. The tectal signal was recorded in a superficial layer at the center of the left tectum with a quartz/platinum-tungsten electrode (Thomas Recordings, Giessen, Germany) with impedance 1 M Ω at 1 kHz. Data were amplified and filtered (1 Hz to 6

kHz) before being digitized at 32 kHz. For the analysis, data were low-pass filtered with 240 Hz, down sampled to 500 Hz and cut into 60 trials with 50 s each.

Visual stimulation. A sequence of red LED light pulses with random duration (uniform distribution between 1 ms and 2 s) and random inter pulse interval (uniform distribution between 1 ms and 5 s) was triggered via the parallel port using MATLAB and the Psychophysics Toolbox extension [5,6]. A light guide projected the full field flashes onto the retina.

Data collection for the Supplementary Figure 3

Recordings. We analyzed electrocorticograms from 2 subjects with pharmaco-resistant epilepsy who had implanted strip electrodes (AD-Tech) on their visual cortex for diagnostic purposes. The electrodes were referenced to linked mastoids, amplified (Schwarzer GmbH), and recorded at a sampling rate of 1000 Hz. The location of electrode contacts was ascertained by MRI.

Visual stimulation. The subjects were presented with brief (150 ms) noisy images containing or not containing a person on them. The subjects had to indicate via a button press whether they had perceived a person on the image or not and whether the person was male or female. The pre-stimulus time window was from -1000 ms to 0 as aligned to the stimulus onset and post-stimulus epoch corresponded to the time interval 0 to 1000 ms. The response screen appeared only after that, thus our analysis time window did not contain motor responses. Supplementary Figure 2 shows the results of the CFC analysis for one of the representative electrodes in the vicinity of the extrastriate body area (MNI coordinates -56, -67, -8).

Supplementary Results

Literature review

To assess the prevalence of the critical issues presented in the current manuscript we conducted a literature review. In particular, we evaluated papers that appeared over the recent years (2010-2014) to demonstrate that the caveats discussed in the main text are timely. To show that the problems are prevalent also in the key journals of the field of neuroscience, we analyzed papers from journals Science, PNAS, Neuron, Journal of Neuroscience and Neuroimage (our search terms did not yield papers in Nature or Nature Neuroscience). We searched for articles in PubMed with terms “cross-frequency coupling”, “phase-amplitude coupling”, “cross frequency interactions” and “nested phase amplitude”. We added manually one high-ranking paper that did not come up with those searchers. We excluded three papers that came up in the search but focused on phase-phase or amplitude-amplitude coupling. A total of 22 articles were evaluated according to five criteria, each related to one of the main issues discussed in the main text, and detailed as follows.

- *Phase interpretability.* As stated in the main manuscript and the supplementary discussion, a meaningful interpretation of a phase variable, and thus of a phase-amplitude CFC pattern, requires a clear peak in the power spectrum. Therefore, we first assessed whether the paper provides evidence for a spectral peak for the modulating frequency component. We concluded that the paper did that if there either was a remark about a respective spectral peak in the main text or in the supplementary material or if any of the figures presented either a power spectrum or a time-frequency representation which allowed us to conclude that there is a spectral peak. We categorized the paper as not providing evidence for a spectral peak when there was no remark about a peak and when there was no figure depicting either power spectrum or time-frequency decomposition or when the respective figure did not contain evidence for a spectral peak. We also categorized the paper as not having evidence for a spectral peak, when more than 50 different channels were analyzed and it was not stated in the text that the authors took care that each channel which had or was part of a significant coupling had a spectral peak for the modulating frequency.
- *Selection of bandwidth.* We studied how the bandwidth selection or scanning was done. This was usually straightforwardly reported in the methods section. In particular, we evaluated whether the bandwidth met the condition that should allocate sidebands of the modulating (phase) frequency. When the paper just scanned the whole frequency band with some fixed frequency step without further discussion, the paper was classified as not providing a justification for the chosen bandwidth.

- *Non-stationarity.* As explained in the main text, non-stationary input to a given area can create CFC patterns even when there is no physiological coupling. We assessed whether the paper discussed the effects of non-stationary inputs potentially leading to the observed CFC measure.
- *Surrogates.* We considered how the surrogate distribution was created. As reported in the main text, most conservative surrogates are obtained by minimizing the number of blocks by cutting at single point at a random location. The computation of the surrogates was usually reported in the methods section. We classified the paper as having tried to construct most conservative surrogates if for event-related data trial-shuffling was used and if for continuous data the data were cut in one point. Otherwise, e.g. if phase scrambling was used or if continuous data was cut in many points, we classified the paper as not constructing the most conservative surrogates. If the paper did not use a surrogate distribution, e.g. when the effect was computed by comparing two conditions, we did not count the paper.
- *Control for spectral changes across conditions.* We assessed whether the effects of potential differences of the power spectrum were taken into account or controlled for while analyzing CFC differences between conditions. As explained in the main text, differences in spectral power can lead to differences in CFC without a true change in coupling strength. We classified the paper as not having a control for the effects of the power spectrum if it was reported that the power spectrum differed between the analyzed conditions but it was not evaluated the impact of such difference on the CFC measure. When the paper did not analyze differences between conditions or the conditions did not differ in the magnitude of CFC, we did not count the paper. We classified the paper as taking care of the differences in the power spectrum when the authors acknowledge that differences in the power spectrum could lead to differences in CFC and take steps to deal with the issue.

Results of the literature review

Issue	Yes	No	% Yes	Papers not counted
Are spectral peaks identified for the (modulating) low frequency component?	10	12	45.45%	
Is there justification for the chosen bands?	3	19	13.36%	
Is the possibility of non-	3	19	13.36%	

stationary input leading to observed CFC patterns discussed?				
Are the surrogates constructed in the most conservative way?	10	4	71.43%	8
Are the differences in the power spectrum accounted for?	8	7	53.33%	7


Examples of spurious cross-frequency coupling (CFC)

Cross-frequency coupling analysis is aimed at detecting specific spectral correlations. However, the origins of such correlations are diverse and not always reflect interactions across frequencies. Here we first show that a single oscillator can readily exhibit cross-frequency coupling features simply by virtue of its non-linear properties. We illustrate this case with the well-known Van der Pol system which is a generic model of non-linear relaxation oscillators. Its evolution is given by the following differential equation

$$\frac{d^2 x}{dt^2} - \mu(1 - x^2) \frac{dx}{dt} + x = 0,$$

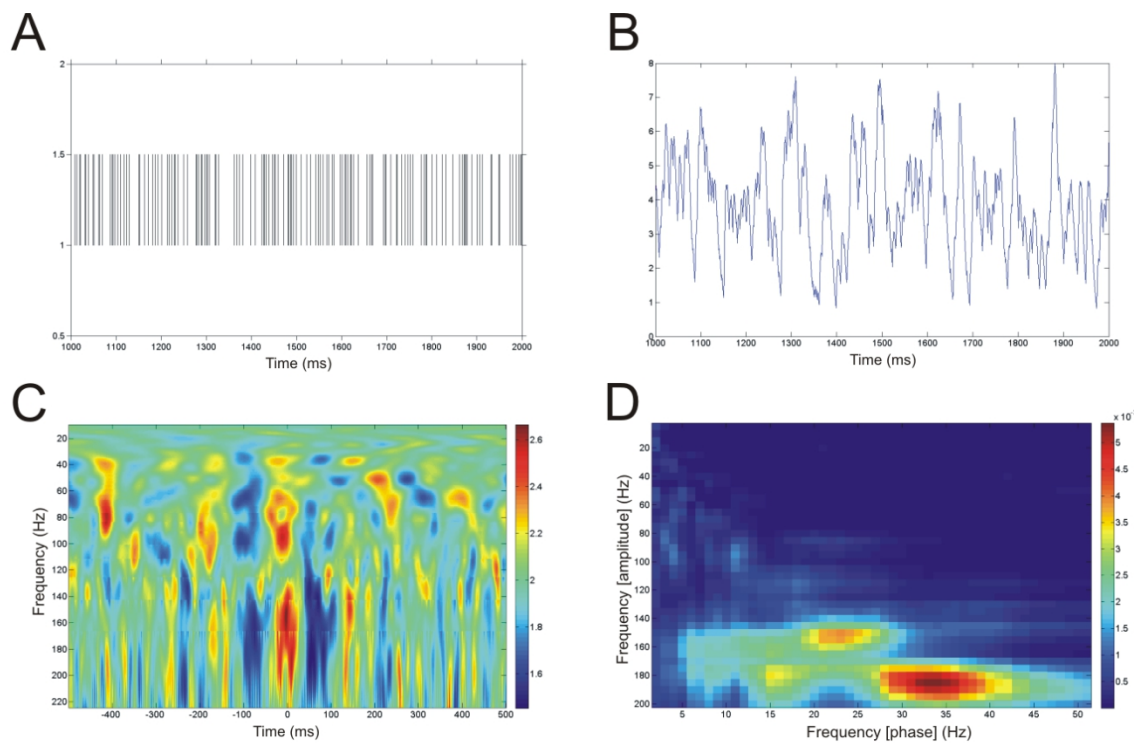
where μ represents a non-linear damping coefficient. Top panel in Supplementary Figure 1 shows the time series of the oscillator for $\mu = 3$. Bottom panel shows the power of high-frequency components when locked to the trough of the phase of a low frequency band around the fundamental period of the oscillation. In both panels a phase dependent modulation of high-frequency power (phase-amplitude CFC) is observed.

However, such systems are not necessarily decomposable into two subsystems oscillating at different frequencies and causally interacting. Instead, the spectral correlations can be simply related to the shape of the oscillatory orbit and the non-linear characteristics of a single oscillator. In particular, the non-linear damping (damping dependent on the state of the oscillator) is responsible for an increasing slope for certain states of the oscillator which is reflected in high-frequency power being associated with certain phases of the fundamental oscillatory frequency.



Supplementary Figure 1. Cross-frequency coupling can arise in a single non-linear oscillator. Top panel: time series of a Van der Pol oscillator with non-linear damping coefficient $\mu = 3$. Bottom Panel: power locked to the trough of the phase of a low frequency component (frequency band (0.005-0.02 cycles/time unit)). CFC analysis can confound non-linear features of oscillations as interaction across frequencies.

A second example of how CFC does not need to imply interactions across frequencies is given by a simulated spiking process. For the time series we have taken numbers 1, 2 ... 30000 and set them to 1 if the number is prime and 0 otherwise. The sampling rate of this signal was nominally set to 1000 Hz. With this procedure we constructed a point process in which spikes occur at the bins that are prime numbers. As CFC analysis is performed for continuous signals and not for point processes as in the example above, we convolved our time series with a smooth kernel (an alpha-function mimicking the conductance response of a synapse to incoming spikes). As evidenced in Supplementary Figure 2 and confirmed by calculating the CFC measures, the artificial series also shows very strong CFC. In this example, the spiking events anchor both the phase of slow frequency bands and the high-frequency components of the kernel function, which automatically leads to CFC.



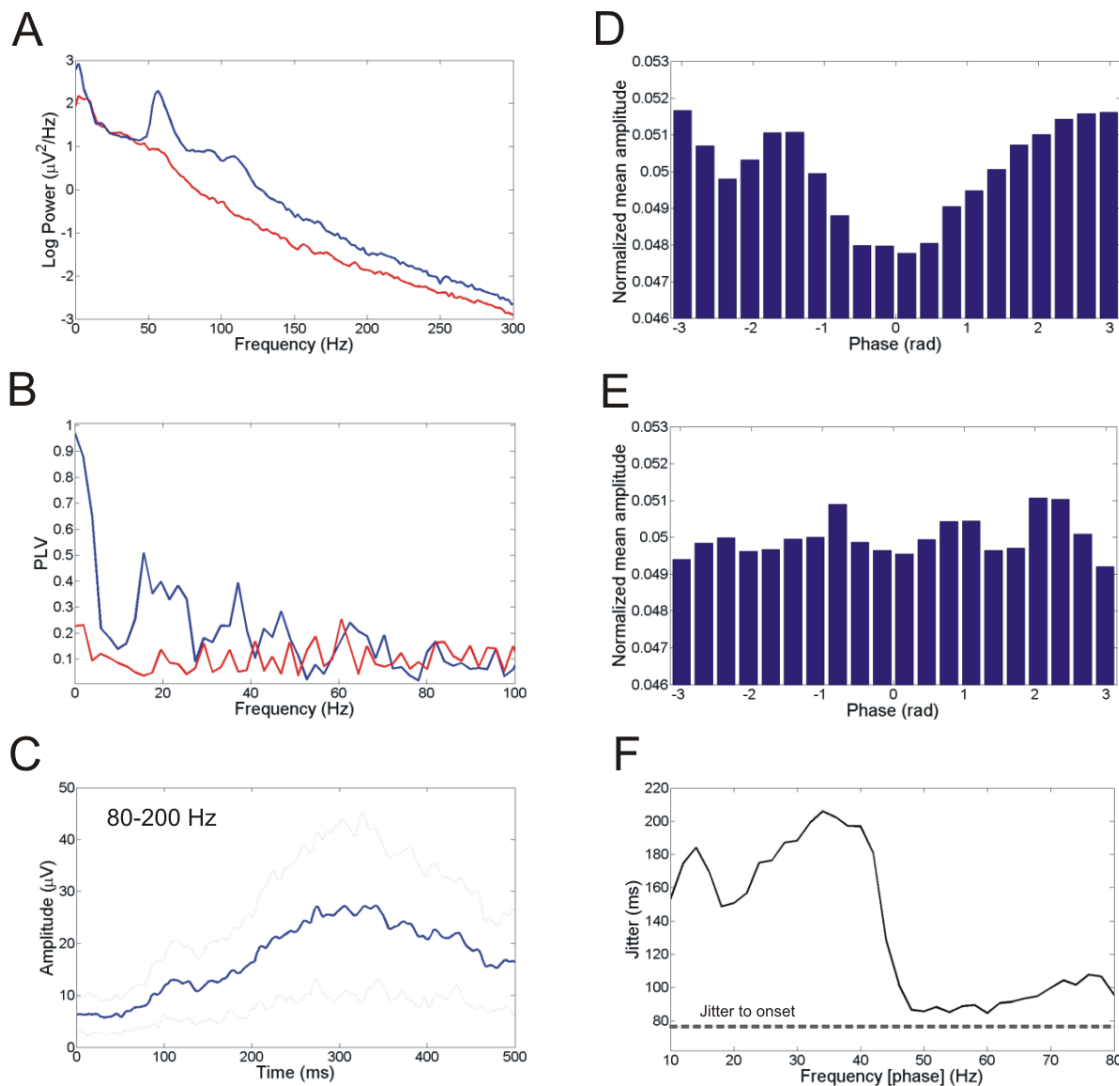
Supplementary Figure 2. Point processes (even when convolved with a smooth kernel) lead to CFC. **A)** Point process where spiking events occur at bins that are prime numbers. **B)** Convolution of this point process with an alpha function with a time constant of 5 ms. **C)** Power of high-frequency activity locked to the trough of low-frequency activity (4-8 Hz). **D)** Kullback-Leibler divergence between the amplitude-phase distribution of the time series and the uniform distribution (modulation index).

Another curious time series for which we found CFC effects are the intervals between consecutive zeros of the Riemann Zeta function (only zeros on the axis with real part $\frac{1}{2}$ were considered). Results not shown.

Effect of input on CFC in human intracranial recordings

As in the turtle case (Figure 3 in the main text), the brief visual stimulation led to an overall increase of power (Supplementary Figure 3, panels A and C). Concomitant with the stimulus phase was reorganized predominantly at low frequencies, as measured by phase locking value, (panel B). As a consequence of power and phase adjustments, the mean amplitude-phase histogram following stimulation is non-uniform (panel D). The pairing amplitude and phase of different trials renders the histogram uniform (panel E). To determine whether the observed cross-frequency coupling is due to an interaction across frequencies or simply an effect of the stimulus, a jitter analysis is presented in panel F. Jitter (as measured by the standard deviation of a series of time differences) from the peak of high-frequency power (80-200 Hz) to stimulus onset is around 80 ms (grey line). The black line represents the jitter from the peak of high-frequency power to the trough of the phase of the slow components closest to their maximal power. For a range of frequencies, the jitter between power and stimulus onset is smaller (more locked) than the one between power and phase (less locked).

This suggests that a common drive from the time-varying input is a more plausible explanation for the correlation than a putative interaction between frequency components.



Supplementary Figure 3. Analysis of human ECoG in a high order visual area showing that non-stationary input leads to positive measures of CFC. **A)** Power Spectra for a period following (blue) or preceding (red) the onset of visual stimulation. **B)** Phase locking values as a function of frequency (colors as in A). **C)** Amplitude of high-frequency activity (80-200 Hz) locked to stimulus onset (mean is represented in blue while mean plus/minus standard deviations are in grey). **D)** Amplitude-phase histogram for the time window of 500 ms following stimulation. **E)** Same as D, but pairing amplitude and phase from different trials. **F)** The jitter from the high-frequency (80-200 Hz) power peaks to the stimulus onset (77 ms; grey line) is always smaller than the jitter measured from the lower frequency trough closest to the maximal high frequency power (blue line).

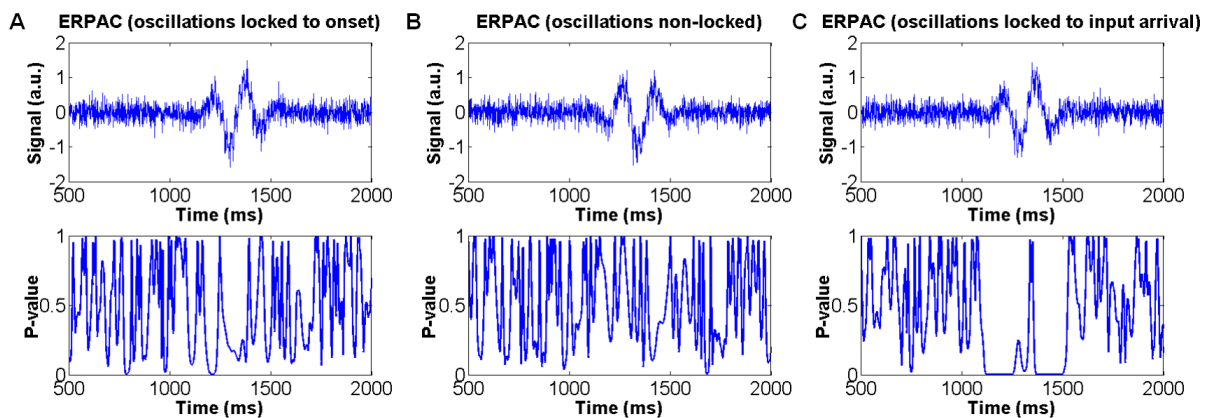
Phase-amplitude coupling for event-related potentials (ERPAC)

Recently, a new approach [7] has been developed to measure transient phase-amplitude coupling directly in an event-related manner. Whereas ideally, this approach of analyzing phase-amplitude relation with respect to the stimulus onset should avoid some event-related artifacts [7], we find that some extra works is needed to show to the marker actually works as it is meant to. Indeed, for the moment we have found that important confounds would need to be controlled for before assigning

any role to the variance explained by a phase variable (see below, Supplementary Figure 4). The extent of the influence of this and other types of confounds is not clear and should be more carefully examined.

The problem seems to lie in the fact that the underlying model is too unspecific: it is looking for one correlation relation, where as it is well possible that both phase and amplitude vary commonly due to some other third factor that is not accounted for. This is indeed a common problem in neuroscience and many other fields. However, in our specific case it seems that generically many other variables of the recorded signal besides the phase of a low-frequency band will have explanatory power of the variance across trials. For example, some variability in the response latency to input together with the phase locking of some frequency components to the response will already give rise to high ERPAC measure. The origin of such jitter can be manifold, and the associated variance across trials can be carried by many other spectral variables than the phase of a low-frequency band. It is highly questionable whether it is helpful to interpret all jittered input with the capability to lock some frequency components as a phase-amplitude coupling.

In Supplementary Figure 4 below we illustrate this case and compared to those of purely evoked and induced ERPs. Note that p-value for panel C in Figure 4 is smaller than 10^{-11} for extended regions near the stimulus response.



Supplementary Figure 4. Event-related phase-amplitude coupling (ERPAC) as detected by the method of Voytek et al. **A)** One representative trial (top) and ERPAC (bottom) measure for a signal with oscillatory components phase locked to a fixed onset (at 1000 ms). **B)** One representative trial and ERPAC measure for a signal with oscillatory components not phase locked over trials. **C)** One representative trial and ERPAC measure for a signal with oscillatory components phase locked to a jittered input response (simulated by a Gaussian amplitude profile). In all cases the signal is composed by white noise plus two sinusoidal components (6 Hz and 50 Hz) that were modulated by a Gaussian profile (mean = 1320 ms; standard deviation = 100 ms) whose center is randomly jittered over trials (jitter = 40 ms). ERPAC measures were obtained from the code available at <http://darb.ketyov.com/professional/publications/erpac.zip>


We also observed that in the concrete example above, as expected, already the amplitude of low frequency component explains the variance of the amplitude of high frequency component. It is thus plausible and possible that replacing the bivariate generalized linear model used in the article with a multivariate model (for example including the amplitude of the low frequency component), could point down a more specific statistical relation between a high-band amplitude and low-band phase at any fixed time-point more specifically.

However, even if the method can be cleaned to avoid some confounds, i.e. made more specific in terms of relating amplitude and phase, any interpretation of this time-point wise amplitude-phase coupling already presupposes that phase-amplitude coupling as a mechanism exists. Ultimately, since causality and mechanisms develop over time not trials the lack of time as a regressor in a model can lead to serious over-interpretations as with the example mentioned above.

Atmospheric noise shows CFC after squaring the signal

Another concern with interpreting cross-frequency interactions is role of non-linearities in creating such correlations. The signal provided by a neurophysiological recording can be influenced by different non-linearities. Some of them are intrinsic to basic neuronal processing such as nonlinear dendritic summation caused by voltage-gated conductances. Others, however, might be of a more mundane origin. These non-linearities can include the electrical properties of the neuronal tissue (e.g., activity-dependent resistivity changes [8]) or small deviations from linearity occurring at any stage of the transduction from the neuronal generators to the final output signal that is subject to analysis.

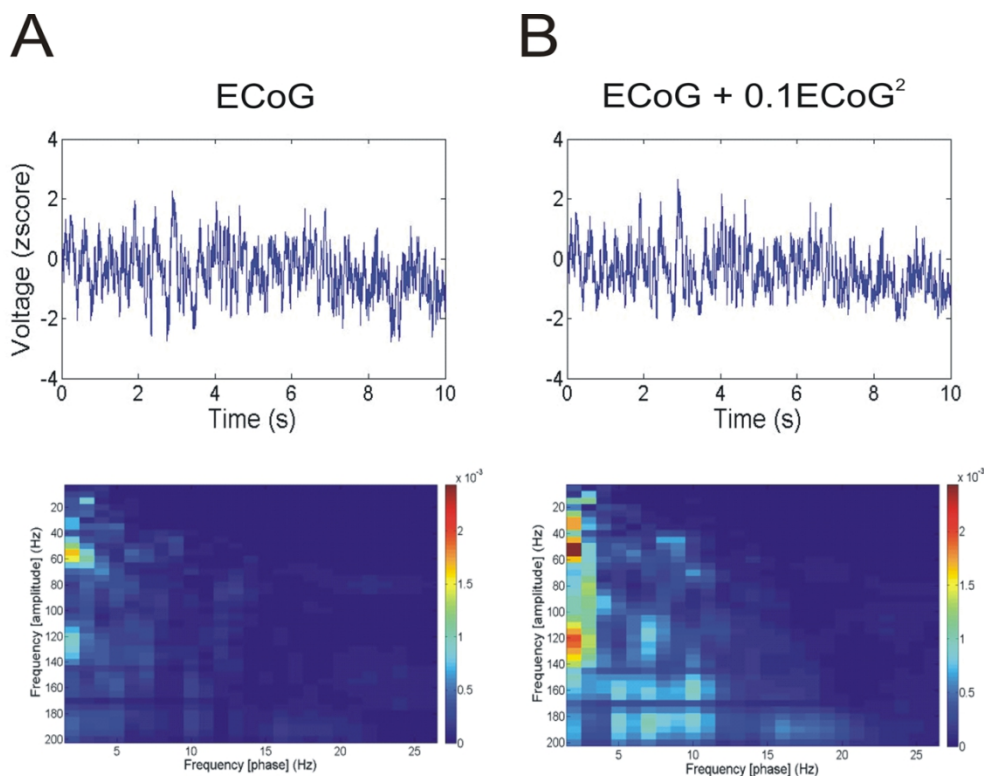
To test the effect of static non-linearities in CFC analyses, we compared the CFC patterns of a random time series and its square. The random time series is composed of uniformly distributed random numbers taken from a physical source (10000 samples of atmospheric noise were taken from www.random.org). As shown in Supplementary Figure 5, random noise does not contain any evident structure of phase-amplitude coupling. However, the square of this random signal displays a clear modulation. Therefore, non-linearities can confound the CFC measures.



Supplementary Figure 5. Non-linearities can generate specific patterns of CFC even for random data (atmospheric noise, 10000 samples, nominal sampling rate 1000 Hz). Top panel: power at different frequencies locked to the trough of phase of a low frequency component (4-8 Hz). Bottom panel: same analysis as in top panel but for the signal obtained by point-wise squaring the atmospheric noise time series.

Small static non-linearity in ECoG data generates CFC

Adding a statistically negligible contribution (Pearson coefficient of correlation between original and distorted signal ~ 0.99) of a squared signal to an ECoG signal not showing any significant CFC, will render CFC highly significant (Supplementary Figure 6).



Supplementary Figure 6. Left panels: original ECoG time series (top) and modulation index from REF. [14] (bottom). Right panels: the same signal after the addition of 10% of its square, and the corresponding modulation index. Notice that the signals (top panels) look very similar, but the modulation index on the right is much stronger. This result shows that even small non-linearities can create significant false-positive CFC patterns.

This result is natural as even small non-linearities readily create harmonics which can spread over the whole spectrum and, thus, generate long-distance spectral correlations for a continuous and large range of frequencies. Importantly, such non-linearities can create specific CFC patterns, where the greatest amplitude of the high-frequency component is related to a specific phase of the low frequency component. For a simple mathematical example of a non-linearity see next section.

Mathematical example of a non-linearity

For a simple mathematical example of a non-linearity consider a harmonic wave with frequency f given by:

$$x(t) = \cos(2\pi ft)$$

and a non-linear version of x that to a second-order approximation is given by:

$$y = x + ax^2 \approx \cos(2\pi ft) + \frac{a}{2} \cos(2\pi(2f)t) + \frac{a}{2},$$

with the approximation holding if $a \ll 1$. Now let us show that in this case the Hilbert amplitude of y is a function of the Hilbert phase of x . Hence if we imagine that we filter y first only narrowly around the frequency f (thereby isolating x) for extracting the phase, and then widely around frequency $2f$ (thereby isolating y) for computing the amplitude, then we will get perfect phase-amplitude coupling.

The proof proceeds as follows. The Hilbert phase of x is by definition seen to be $2\pi ft$ modulus 2π . The Hilbert amplitude of y can be found by constructing the analytic signal of y by just using two facts: 1) analytic signal of harmonic oscillation is given by its complex form, and 2) Hilbert transform is linear. Hence, the analytical signal of y is

$$Y(t) = \exp(i2\pi ft) + \frac{a}{2} \exp(i2\pi(2f)t) + \frac{a}{2}.$$

Now this can be rewritten as

$$\begin{aligned} Y(t) &= \exp(i2\pi ft) + \frac{a}{2} (\exp(i2\pi(2f)t) + 1) = \\ & \exp(i2\pi ft) + \frac{a}{2} \exp(i2\pi ft) (\exp(i2\pi ft) + \exp(-i2\pi ft)) = \\ & \exp(i2\pi ft) (1 + a \cdot \cos(2\pi ft)). \end{aligned}$$

From here we can see that the Hilbert amplitude of y is equal to

$$1 + a \cdot \cos(2\pi ft)$$

and hence an exact function of the phase of x .

Supplementary Discussion

Conditions for a meaningful phase

Intuitively, a phase is a parameter that tells us where we are within a cycle of some repetitive motion or variation. The amplitude informs us about the span of that motion. For a simple harmonic motion ($s(t) = a \cos(\omega t)$), the amplitude is constant and the phase indexes the position along the cycle. Importantly, the phase is useful because it is an index that grows monotonically within a cycle, i.e., the larger the phase the larger the completed fraction of the current cycle. Geometrically, amplitude and phase correspond to the distance to origin and the angle subtended by the complex function $a \cos(\omega t) + i a \sin(\omega t)$, respectively. Notice that the imaginary part is a 90 degrees shifted copy of the original signal (cosine becomes sine).

When faced with a less regular signal $s(t)$, the task of assigning a phase and amplitude is more delicate since there are infinitely many pairs of functions $a(t) > 0$ and $\phi(t)$ that satisfy $s(t) = a(t) \cos(\phi(t))$. Which one to choose? Gabor elegantly proposed the analytical signal approach to provide a unique and unambiguous solution [9]. The idea is the following: since we know how to define a phase and amplitude for a single sinusoidal, we can decompose the signal in its Fourier components and repeat a similar procedure for each component. This leads to the complex analytical signal the real part of which is the original signal $s(t)$ and its imaginary part is formed by 90 degrees shifted copies of the original signal's Fourier components. The Hilbert transform (H) sums up these decomposition and shifting operations, and the analytical signal of $s(t)$ is thus represented as $s_a(t) = s(t) + i H(s(t))$.

However, as several authors have noted [10], a clear interpretation of the phase and amplitude obtained by this or, indeed, any other method is restricted to *narrow-band* signals. Thus, although the analytical approach yields an instantaneous phase and amplitude for any signal, the interpretation is only clear if the signal does not deviate much from being a smooth periodic function. Fortunately, the narrow-band condition covers the presence of moderate noise and smooth frequency fluctuations since in that case the signal will be still narrow-band but with a slowly varying center frequency. The main issue is that for narrow-band signals amplitude and phase can be considered as separate and independent entities, while this is not the case for broad-band signals. Take the example of a modulated signal $s(t) = a(t) \cos(\phi(t))$ in which some physical meaning is attached to $a(t)$ and $\phi(t)$. When $\phi(t)$ is a smooth function of time and $a(t)$ changes slow enough within each cycle (i.e., an increase of $\phi(t)$ by 2π), then $s(t)$ will result in a narrow-band signal where phase and amplitude are separable variables. In this case the estimation of phase and amplitude via

analytical signal analysis or other approaches can recover the exact modulated values of $\phi(t)$ and $a(t)$ from only $s(t)$. The technical reason for this result is known as the Bedrosian theorem [11]. This theorem also implies that for broad-band modulated signals where the spectra of $a(t)$ and $\cos(\phi(t))$ overlap, the phase and amplitude estimated from the analytical signal will be composed by an intricate mixture of $a(t)$ and $\phi(t)$. Moreover, values of the estimated instantaneous frequency (seen as a derivative of the instantaneous phase) could even take negative values. Importantly, this is not a problem of any particular approach but it reflects the fact that irregular signals might have degrees of freedom that cannot be faithfully represented by just a smoothly changing phase-amplitude pair. Indeed, the analytical signal approach can be considered as the optimal two-dimensional delay embedding of a univariate signal to estimate its phase and amplitude. However, many signals need a higher number of dimensions to be fully unfolded [12]. In that case, insisting on the two dimensional description of the signal (by phase and amplitude or any other pair) results in projecting all the degrees of freedom into just two variables and thereby compromising their interpretability.

It is noteworthy then that phase-amplitude CFC is frequently investigated for regions in the spectrum that locally exhibit a power law, i.e., $P(f) = 1/f^\alpha$, and thus do not contain a peak [13]. For these regions, applying a narrow-band filter of a few Hz will render a signal with smooth phase and amplitude dynamics. However, this apparent smoothness is a simple result of filtering, and not a sign that there are processes in this spectral range that indeed have a smooth, or even meaningful, phase. Hence, narrow band filtering in regions exhibiting $1/f^\alpha$ power spectrum does not lead to an interpretable phase, despite its apparent smoothness. Thus, findings of spectral correlations at regions exhibiting $1/f^\alpha$ decays should be interpreted extremely carefully.

In sum, only a natural concentration of power around some center frequency in a time-frequency decomposition enables a meaningful interpretation of the phase, and, thereby, CFC analysis.

The importance of the bandwidth

The concentration of power around some frequency is a necessary prerequisite for a meaningful interpretation of the phase of a signal. Thus, for the phase variable the center frequency and bandwidth selected should define a range that includes at least one peak of the spectrum. Even under such condition the value of the width severely affects the results. A too narrow filter, at least compared to the natural width of the peak, will result in a smooth and well-behaved phase which, however, is hardly representative of the underlying oscillatory but variable signal. A too broad filter can result in an ill-behaved phase with phase slips or reversals by incorporating $1/f$ components. Importantly, both cases can lead to a loss of sensitivity and interpretability of the analysis. A

possible strategy is to precisely exploit the freedom of the bandwidth to look for a sweet point where the phase dynamics shows maximal robustness against small bandwidths changes. Another heuristic but natural choice is to select the bandwidth as the range around the peak that clearly stands out from the background of the power spectrum. If a good fitting for the background spectrum is available, then subtracting the real and interpolated spectra can deliver an estimate of the width [14]. Finally, if additional physiological information is available it might be possible to use functional criteria to select bands. For example, if interested in CFC of human occipital alpha rhythms it is probably more appropriate to find subject specific alpha bands rather than to select a fixed band ranging from 8 to 12 Hz.

Now we turn to the issue of selecting an adequate bandwidth for the frequency component defining the instantaneous amplitude. Following the example in Figure 2, suppose it is known that the lower frequency component is around f_1 (frequency for extracting the phase) and the higher oscillatory component around f_2 (frequency for extracting the amplitude). If one would filter with a bandwidth smaller than $2f_1$ centered at f_2 , one would come to the representation [**] and the amplitude of the higher component would be constantly 1, hence independent of the phase of the lower frequency component. However, if one filtered with a bandwidth bigger than $2f_1$, the higher component would be defined as the first summand in the representation [*] and hence show CFC. Therefore, our choice of bandwidth defined not only how we should mathematically think of the function, but also whether we can observe CFC. This hints at the possibility of obtaining false-positive and false-negative results in cross-frequency analysis depending on the choice of frequency components. Importantly, to characterize CFC patterns one usually scans the center frequencies f_1 and f_2 while keeping the bandwidth of their components fixed. However, a fixed bandwidth is problematic because the critical bandwidth to observe modulation (detecting the sidebands around f_2) depends on the value of the lower frequency f_1 . In particular, this restricts CFC measures to find low modulating frequencies (f_1) where sidebands fit within the chosen bandwidth, and thus creates an extreme bias against observing potential phase modulations by higher frequencies (Figure 2B,C). A similar case has recently been raised in [15].

Different model/statistical approaches to assess phase-amplitude CFC

Different methods exist to attack the problem of assessing relationships among rhythmic processes from experimental time series. A possible historical classification of different frameworks is as follows: non-linear systems analysis, Fourier based methods, classical stationary time series, non-parametric approaches, and causal statistical modeling. We include below a brief description of each of the mentioned methods.

Notice that in the main text, we find it more helpful to classify concrete methods for CFC analysis using two almost orthogonal dimensions.

I) Fourier/wavelet analyses. These start by projecting the data onto a given basis of oscillatory functions. This first step is indeed common to most of practical approaches to CFC. Fourier/wavelet approaches lack an explicit formulation of a model in terms of equations describing the temporal evolution of relevant variables. However, once a basis or dictionary of functions is selected (oscillatory or not) the projection of data onto such basis amounts to describe the signal as a weighted sum of systems, each of them assuming the dynamics of one basis function. Higher-order spectral quantities, such as bi-coherence, aim to capture correlations between the complex components of an oscillatory basis, and thus, putative CFC relations. The double Fourier transform of a time- or phase-dependent autocorrelation function can also be used to characterize cross-frequency correlations [16]. Unfortunately, higher order functionals are increasingly more difficult to estimate, and multivariate extensions of such approaches are very limited from a practical point of view. The most widely used phase-amplitude CFC measures [2,3] can be thought to measure undirected correlations or dependencies between the frequency components of this class of models.

II) Classical time series techniques. They largely refer to regression techniques such as fitting the data into models such autoregressive moving average processes and its many and important variants both in the time and frequency domain. The strong restrictions of the models (e.g. linearity) are responsible for the extreme data-efficiency and practicality of its multivariate extensions. Models explicitly incorporating cyclo-stationarity (non-stationary components that repeat periodically as used to model seasonal components of financial and geophysical models [17]) can readily serve to quantify certain phase-amplitude CFC effects. Other models exist to regress amplitude variables to non-linear functions of the phase such as the generalized linear models proposed by Penny et al. [18]. Due to the restrictive assumptions of the models, they can easily fall beyond the range of interactions actually occurring in complex systems.

III) Non-linear systems analyses. These analyses typically assume that a low dimensional dynamical system governs the evolution of variables extracted from measured signals. For example, the Kuramoto model can be used to describe the non-linear interaction between the phases (phase coupling) of coupled oscillators. Coupling terms and other coefficients are then fitted/estimated from the data providing a direct interpretation of the interdependencies between variables. For the Kuramoto example, given the time series of the phases from the measured signals, one can estimate the phase-to-phase coupling coefficients between different oscillatory recordings (see [19] for an efficient implementation of maximum-likelihood inference for Kuramoto networks). Similar models are conceivable for the phase-amplitude problem (Stuart-Landau equations). In general, this class of approaches offers an explicit dynamical model of the relations between different variables extracted

from a signal, and can easily account for multivariate descriptions (e.g., by adding more coupling terms), and these approaches are data efficient compared to non-parametric approaches. However, while certain aspects of the model can be justified on theoretical grounds, they usually lack a physiological foundation.

IV) Non-parametric approaches. These aim to estimate relationships between variables without assuming any specific structure of the model, which makes them ideal for explorative analysis. Often they make use of probability or state-space based descriptions of variables and their relationships, which are sensitive to all nonlinear order interactions as it is the case with information theory functionals. Transfer entropy for example, aims to capture Wiener/observational causal relations between variables (possibly related to processes at different frequency bands) by quantifying the increase of predictability about the future states of one variable once the information about the present and past states of another variable are included. However, these approaches are typically highly expensive in terms of data, especially in multivariate settings due to the curse of dimensionality. Thus, practical estimators must rely on certain mild assumptions such as the smoothness of probability distributions and the wide-sense stationarity or ergodicity of the time-series [20].


V) Causal statistical modeling. It offers a scheme to compare how much evidence the data provide about particular models or hypotheses. The hypotheses are usually formalized by a generative model that can incorporate experimental manipulations plus some measuring functions, which relate relevant physical variables to the magnitudes detected by the experimental apparatus. Such procedures typically rely on Bayesian inference to incorporate prior knowledge and estimate the evidence for the different models once new data has been collected. Dynamical causal models (DCM) are being developed to account for biophysical descriptions of neuronal interactions [21]. Unfortunately a practical setting for inferring phase-amplitude cross-frequency interactions has yet to be developed.

Other possibilities include parametric approaches, as for example fitting phase-amplitude histograms to Von Mises distributions [22].

The transitivity of correlation between phase and amplitude

In non-linear oscillators, amplitude and phase are intrinsically coupled and thus simultaneously influenced by generic perturbations (e.g. input). Thus, under natural conditions it is almost impossible to selectively modify the amplitude, without changing the phase and vice versa [23]. Moreover, since amplitude and phase can exhibit different susceptibility to perturbations or inertia it is not trivial to infer causal relationships from the timing of their dynamics. In addition, the

transitivity of correlations makes it hard to distinguish whether the phase of one process and the amplitude of another are directly or indirectly linked, namely via phase-phase or amplitude-amplitude coupling (Supplementary Figure 5A). For example, the phase of the low frequency component (Ψ_L) might influence the phase of the high frequency component (Ψ_H) (phase-phase coupling). Since phase and amplitude of the high frequency process are intrinsically coupled, one will observe phase-amplitude coupling. Consequently, it is advisable to partial out indirect ways (e.g. phase-phase coupling) of phase-amplitude modulation in order to assign a functional role to a specific type of coupling. Notice that also some of the vertical couplings in the Supplementary Figure 7A can also appear as a function of the estimators used to define phase and amplitude. For example, the amplitude and phase defined by the analytical signal approach (using Hilbert transforms) are not fully independent and even a nominal change in one of them induces a perturbation in the other (Supplementary Figure 7B).



Supplementary Figure 7. **A)** For non-linear oscillators phase and amplitude at the same frequency are intrinsically linked. Therefore, there are different ways to obtain phase-amplitude coupling between a low (L) and a high (H) frequency component of different oscillators. **B)** Estimation of instantaneous phase and amplitude can also couple phase and amplitude dynamics. Nominal changes in either the amplitude (left) or phase (right) of a sinusoidal are simultaneously reflected in both the instantaneous phase and amplitude as obtained by the analytical approach (Hilbert transform).

Supplementary discussion on causality methods

Currently applied observational causality approaches meet several difficulties that need to be resolved before they can be applied to cross-frequency interactions of neurophysiological data. The original Granger formalism is restricted to linear relations, disregarding any cross-frequency effects. Some attempts have been made to apply similar techniques to assess linear causality between instantaneous phases and amplitudes of different spectral bands but these interactions may themselves be mediated by highly non-linear processes and thus remain invisible for the Granger formalism. On the other hand, non-parametric approaches such as transfer entropy (accounting for

all orders of non-linear interactions) typically need long stretches of data to be reliably estimated. In addition, the uncertainty principle of harmonic analysis limits the temporal resolution with which a spectrally resolved component can be localized. By this principle low frequency components will be more temporally smeared than high frequency components. Thus, the onset of low frequency components will be advanced more than that of high frequency components if non-causal filtering procedures such as the Hilbert transform are used. The problem arises even when causal filters are applied if unequal signal to noise ratios are encountered. This hampers the interpretation of most causality measures, which rely on temporal asymmetric relations.

References for the supplementary material

1. Delorme, A. & Makeig, S. EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *J. Neuro. Methods* **134**, 9-21 (2004).
2. Canolty, R.T., et al. High gamma power is phase-locked to theta oscillations in human neocortex. *Science* **313**, 1626-1628 (2006).
3. Tort, A., et al. Measuring phase-amplitude coupling between neuronal oscillations of different frequencies. *J. Neurophysiol.* **104**, 1195-1210, (2010).
4. Rosenberg, A.F. & Ariel, M. Visual-response properties of neurons in turtle basal optic nucleus in vitro. *J. Neurophysiol.* **63**, 1033–1045, (1990).
5. Brainard D.H. The psychophysics toolbox. *Spat. Vis.* **10**, 433-436 (1997).
6. Pelli D.G. The VideoToolbox software for visual psychophysics: transforming numbers into movies. *Spat. Vis.* **10**, 437-442 (1997).
7. Voytek, B., D'Esposito, M., Crone, N. & Knight, R.T. A method for event-related phase/amplitude coupling. *Neuroimage* **64**, 416-424 (2013).
8. Lopez-Aguado, L., Ibarz, J.M. & Herreras, O. Activity-dependent changes of tissue resistivity in the CA1 region in vivo are layer-specific: modulation of evoked potentials. *Neuroscience* **108**, 249-262 (2001).
9. Gabor, D. Theory of communication. *Proceedings of the Ieee* **93**, 429-459 (1946).
10. Boashash, B. Estimating and Interpreting the Instantaneous Frequency of a Signal .1. Fundamentals. *Proceedings of the Ieee* **80**, 520-538 (1992).
11. Bedrosian, E. Product Theorem for Hilbert Transforms. *Proceedings of the Ieee* **51**, 868-& (1963).
12. Sauer, T., Yorke, J.A. & Casdagli, M. Embedology. *Journal of Statistical Physics* **65**, 579-616 (1991).
13. He, B.J., Zempel, J.M., Snyder, A.Z. & Raichle, M.E. The temporal structures and functional significance of scale-free brain activity. *Neuron* **66**, 353-369 (2010).
14. Miller KJ, Sorensen LB, Ojemann JG & den Nijs M. Power-Law Scaling in the Brain Surface Electric Potential. *PLoS Comput Biol* **5**(12): e1000609 (2009).
15. Berman, J.I., et al. Variable Bandwidth Filtering for Improved Sensitivity of Cross-Frequency Coupling Metrics. *Brain Connectivity* **2**, 155-163 (2012).
16. Gardner, W. The spectral correlation theory of cyclostationary time series. *Signal Processing*, **11**, 13–36 (1986).
17. Troutman, B.M. Some results in periodic autoregression. *Biometrika*, **66**, 219–228 (1979).
18. Penny, W.D., Duzel, E., Miller, K.J. & Ojemann, J.G. Testing for nested oscillation. *Journal of Neuroscience Methods* **174**, 50-61 (2008).
19. Cadieu, C. & Koepsell, K. Phase coupling estimation from multivariate phase statistics. *Neural*

Computation, **44**, 3107–3126 (2010).

20. Vicente, R., Wibral, M., Lindner, M. & Pipa, G. Transfer entropy- a model-free measure of effective connectivity for the neurosciences. *J Comput Neurosci*, **30**, 45–67 (2011).

21. Friston, K.J. & Dolan, R.J. Computational and dynamic models in neuroimaging. *Neuroimage*, **52**, 752–765 (2010).

22. Canolty, R.T., *et al.* Detecting event-related changes of multivariate phase coupling in dynamic brain networks. *Journal of Neurophysiology* **107**, 2020-2031 (2012).

23. Pikovsky, A., Rosenblum, M. & Kurths, J. *Synchronization* (Cambridge University Press, Cambridge, 2003).