

Non-linear Auto-Regressive Models for Cross-Frequency Coupling in Neural Time Series

Tom Dupré la Tour^{1*}, Lucille Tallot^{2,3}, Laetitia Grabot^{4,5}, Valérie Doyère³, Virginie van Wassenhove^{4,5}, Yves Grenier¹, Alexandre Gramfort^{1,5,6}

1 LTCI, Telecom ParisTech, Université Paris-Saclay, Paris, France

2 Neuroscience Paris Seine, CNRS, INSERM, Sorbonne Universités, UPMC, Paris, France

3 Neuro-PSI, Université Paris-Sud, Université Paris Saclay, CNRS, Orsay, France

4 Cognitive Neuroimaging Unit, CEA/DRF/Joliot, INSERM, Université Paris-Sud, Université Paris-Saclay, NeuroSpin center, Gif-sur-Yvette, France

5 CEA, Université Paris-Saclay, Gif-sur-Yvette, France

6 Inria, Université Paris-Saclay, Saclay, France

* tom.duprelatour@telecom-paristech.fr

Abstract

We address the issue of reliably detecting and quantifying cross-frequency coupling (CFC) in neural time series. Based on non-linear auto-regressive models, the proposed method provides a generative and parametric model of the time-varying spectral content of the signals. As this method models the entire spectrum simultaneously, it avoids the pitfalls related to incorrect filtering or the use of the Hilbert transform on wide-band signals. As the model is probabilistic, it also provides a score of the model “goodness of fit” via the likelihood, enabling easy and legitimate model selection and parameter comparison; this data-driven feature is unique to our model-based approach. Using three datasets obtained with invasive electrophysiological recordings in humans and rodents, we demonstrate that these models are able to replicate previous results obtained with other metrics, but also reveal new insights such as the influence of the amplitude of the slow oscillation. Using simulations we demonstrate that our parametric method can reveal neural couplings with shorter signals than non-parametric methods. We also show how the likelihood can be used to find optimal filtering parameters, suggesting new properties on the spectrum of the driving signal, but also to estimate the optimal delay between the coupled signals, enabling a directionality estimation in the coupling.

Author Summary

Neural oscillations synchronize information across brain areas at various anatomical and temporal scales. Of particular relevance here, slow fluctuations of brain activity have been shown to affect high frequency neural activity by regulating the excitability level of neural populations. Such cross-frequency-coupling can take several forms. The most frequently observed type is when the power of high frequency activity is time-locked to a specific phase of slow frequency oscillations, yielding phase-amplitude-coupling (PAC). Even when readily observed in the neural recordings, such non-linear coupling is particularly challenging to formally characterize. Typically, neuroscientists use custom band-pass filtering or Hilbert transforms with cross-frequency correlations. Here, we explicitly address current limitations and propose an alternative probabilistic signal modeling approach for which statistical inference is fast and well-posed. To statistically model PAC, we propose the use of non-linear auto-regressive models which estimate the spectral modulation of a signal conditionally to a driving signal. This conditional spectral analysis enables easy model selection and clear hypothesis-testing by using the likelihood of a given model. We demonstrate the advantage of the model-based approach on three datasets acquired in rats and in humans. We further provide novel neuroscientific insights on previously reported PAC phenomena capturing two possible mechanisms for PAC: influence of amplitude and directionality estimation.

Introduction

The characterization of neural oscillations, which are observed in the mammalian brain at different temporal and spatial scales, have given rise to important mechanistic hypotheses regarding their functional role in neurosciences (e.g. [Buzsáki, 2006, Fries, 2015]). One working hypothesis suggests that the coupling across neural oscillations may regulate and synchronize multi-scale communication of neural information within and across neural ensembles [Buzsáki, 2010, Fries, 2015]. The coupling across different oscillatory activity is generically called cross-frequency-coupling (CFC) and has started to receive much attention [Jensen and Colgin, 2007, Lisman and Jensen, 2013, Canolty et al., 2006, Canolty and Knight, 2010, Hyafil et al., 2015]. The most frequent instance of CFC consists in the observation that the power of high frequency activity is modulated by fluctuations of low-frequency oscillations, resulting in phase-amplitude-coupling (PAC). Other instances of CFC include phase-phase coupling [Tort et al., 2007, Malerba and Kopell, 2013], amplitude-amplitude coupling [Bruns and Eckhorn, 2004, Shirvalkar et al., 2010], and phase-frequency coupling [Jensen and Colgin, 2007, Hyafil et al., 2015]. By far, PAC is the most reported CFC in the literature.

Seminally, PAC was described in local field potential (LFP) of rodents displaying a modulation of gamma band power (40–100 Hz) as a function of the phase of their hippocampal theta band (5–10 Hz) [Bragin et al., 1995, Tort et al., 2008]. Parallel recordings in different brain areas in behaving rodents have also highlighted differences in PAC between brain areas (e.g., hippocampus and striatum) at specific moments during a goal-oriented behavior, both in terms of which high-frequency range and how narrow-band the low frequency is [Tort et al., 2008]. PAC may promote cellular plasticity underlying memory formation [Axmacher et al., 2006]. In humans, theta (4–8 Hz)/gamma (80–150 Hz) PAC was described in human auditory cortex during speech perception [Canolty et al., 2006]. In more recent work, theta/gamma PAC was reported during the processing of auditory sequences in both humans and monkeys [Kikuchi et al., 2017], during working memory maintenance in human hippocampus [Axmacher et al., 2010], and during serial memory recall using non-invasive human magnetoencephalography (MEG) [Heusser et al., 2016]. PAC has been proposed to support the maintenance of information and to play an important role in long distance communication between different neural populations, considering that slow oscillations can propagate at larger scales than fast ones [Jensen and Colgin, 2007, Khan et al., 2013, Lisman and Jensen, 2013, Hyafil et al., 2015, Bonnefond et al., 2017]. Consistent with this notion, PAC has also been reported across distinct brain regions [Sweeney-Reed et al., 2014]. In sum, PAC has been proposed as a canonical mechanism for neural syntax [Buzsáki, 2010].

Given the growing interest in CFC, and in PAC more specifically, developing adequate and unbiased tools to quantify the posited signatures of neural computations has motivated a number of contributions [Canolty et al., 2006, Penny et al., 2008, Tort et al., 2010]. While these works have met some success, they have also pointed out numerous methodological challenges [Aru et al., 2015, van Wijk et al., 2015] illustrated and discussed in several simulation studies [Tort et al., 2010, Özkurt and Schnitzler, 2011]. Some important limiting factors and potential pitfalls of current used techniques are related to the impact of incorrectly choosing the bandwidth of bandpass filters [Dvorak and Fenton, 2014, Aru et al., 2015], the consequences of applying Hilbert transform on wide-band signals [Chavez et al., 2006, Dvorak and Fenton, 2014], and the potential misidentification of CFC [Hyafil, 2015, Kramer et al., 2008]. The direction we advocate to address these issues is to use a signal model that would fit the data in the sense of the mean squared error, and from which a measure of PAC or CFC could be derived. What all existing metrics have in common is that they do not make use of a signal model, and therefore setting the parameters of the method, such as filtering parameters, can only be driven by how much they lead to a strong PAC/CFC metric. As a consequence, even though current metrics give reasonable estimates of PAC, a legitimate and controlled comparison of methods and parameters, and therefore of the results, is impossible. Additionally, while simulations provide better control, they do not fully solve this issue faced by experimentalists since a simulation may approximate at best, or miss at worst, the real structure of neurophysiological signals. Hence, with the present contribution, we hope to propose a major improvement in CFC/PAC estimation by adopting a principled modeling approach.

To initiate this statistical model-based approach, we propose to consider a signal model which is rich enough to capture time-varying phenomena such as PAC. A model-based approach allows computing the likelihood of a recorded neural signal, which can be interpreted as a measure of the goodness of fit of the

model. In the present case, the goodness of fit corresponds to the classical measure of explained variance. Such an evaluation metric is a natural criterion to compare models, and a first step towards an automatic model parameters selection on brain data. Additionally, with a model-based approach, it becomes possible to answer a large variety of practical questions: From a statistical standpoint, one can define which parametrization or which hyper-parameter should be selected, while from a neuroscientific standpoint, one could ask, for instance, whether the instantaneous amplitude of slow oscillations contributes to PAC, or whether the amplitude modulation affects the entire neural spectrum identically. Importantly, our statistical signal modeling approach is not a biophysical one and is thus distinct from previous work [Hyafil et al., 2015, Chehelcheraghi et al., 2017]. Our goal here is to better explain and describe the empirical data themselves in the absence of any assumption regarding the neural mechanisms that have generated them.

To capture PAC in univariate time-series, we propose to use auto-regressive (AR) models, which are stochastic signal models as opposed to deterministic. A canonical deterministic signal model consists in modeling a time-series as a pure sinusoid corrupted by some additive noise. As recently discussed, considering neural oscillations as pure sinusoidal brain responses is an oversimplification of the complex neurophysiological reality [Cole and Voytek, 2017]. AR models do not make the strong assumption of sinusoidality and are rather based on the ability of the model to forecast the signal values. AR models have been successfully used to address multiple problems in neurophysiology such as spectral estimation [Spyers-Ashby et al., 1998], temporal whitening [Mahan et al., 2015], and connectivity measures like Granger causality [Granger, 1988, Valdés-Sosa et al., 2005, Haufe et al., 2010] or coherence [Baccalá and Sameshima, 2001].

In the present work, we show how driven auto-regressive (DAR) models [Grenier, 2013, Dupré la Tour et al., 2017] can be extended in order to study PAC. Standard AR models are statistically efficient given their low number of parameters, but they are linear, and therefore they cannot directly model non-linear phenomena like PAC. In order to extend AR models to cope with such situations of non-linearity and non-stationarity in signals, various advanced AR models have been proposed in other research fields such as audio signal processing and econometrics [Tong and Lim, 1980, Chan and Tong, 1986, Grenier and Omnes-Chevalier, 1988] (see related non-linear AR models section below). Here, we consider the slow oscillation as the exogenous driver so as to allow the coefficients of the AR model to be time varying, thereby making the instantaneous power spectral density (PSD) of the signal be a function of a slow exogenous time series. By doing so, DAR models are statistical signal models (and not biophysical ones) that do not model PAC explicitly but which are able to capture it.

In what follows, we first present how PAC is commonly analyzed in neural time series while pointing out limitations. We then review the current literature on non-linear AR models beyond neuroscience. We then detail the methodological aspects of DAR models first generically and then specifically for brain data. We describe how we can derive from a DAR model a power spectral density parametrized by an external signal and how it allows to compute so-called comodulograms and estimate directionality of coupling; this enables a fine-grained analysis of spectral modulations observed in PAC. Finally, results show how the comparison of goodness of fit between different models allow to conclude on optimal parameters (e.g. filtering parameters), leading to new neuroscientific insights, such as the wideband or the asymmetric spectral properties of the driver. We also report that taking into account the instantaneous amplitude of the driver improves the goodness of fit, which suggests that this information is relevant to the coupling phenomenon and should not be discarded as in most PAC metrics. Last but not least, directionality estimation results reveal some delays between high frequency bursts and the driving low frequency oscillation.

Notations

We note y the signal containing the high-frequency activity, and x the signal with slow frequency oscillations, also called the exogenous driver. When a signal x results from a band-pass filtering step, we note the central frequency of the filter f_x and the bandwidth Δf_x . The value of the signal x at time t is denoted $x(t)$.

PAC quantification methods and current limitations

To estimate PAC, the typical pipeline reported in the literature consists in four main processing steps:

1. Bandpass filtering is performed in order to extract the narrow-band neural oscillations at both low frequency f_x and high frequency f_y ;
2. A Hilbert transform is applied to get the complex-valued analytic signals of x and y ;
3. The phase ϕ_x of the low frequency oscillation and the amplitude a_y of the high frequency signals are extracted from the complex-valued signals;
4. A dedicated approach is used to quantify the correlation between the phase ϕ_x and the amplitude a_y signals.

The Modulation Index (MI) described in the pioneering work of [Canolty et al., 2006] is the mean over time of the composite signal $z = a_y e^{i\phi_x}$. The stronger the coupling between ϕ_x and a_y , the more the MI deviates from zero. This index has been further improved by Ozkurt *et al.* with a simple normalization [Özkurt and Schnitzler, 2011]. Another approach [Lakatos et al., 2005, Tort et al., 2010] has been to partition $[0, 2\pi]$ into smaller intervals to get the time points t when $\phi_x(t)$ is within each interval, and to compute the mean of $a_y(t)$ on these time points. PAC was then quantified by looking at how much the distribution of a_y differs from uniformity with respect to ϕ_x . For instance, a simple height ratio [Lakatos et al., 2005], or a Kullback-Leibler divergence as proposed by Tort *et al.* [Tort et al., 2010], can be computed between the estimated distribution and the uniform distribution. Alternatively, it was proposed in [Bruns and Eckhorn, 2004] to use direct correlation between x and a_y . As this method yielded artificially weaker coupling values when the maximum amplitude a_y was not exactly on the peaks or troughs of x , this method was later extended to generalized linear models (GLM) using both $\cos(\phi_x)$ and $\sin(\phi_x)$ by Penny *et al.* [Penny et al., 2008]. Other approaches employed a measure of coherence [Colgin et al., 2009] or the phase-locking value [Lachaux et al., 1999]. All these last three approaches offer metrics which are independent of the phase at which the maximum amplitude occurs. The methods of Tort *et al.* [Tort et al., 2010], Ozkurt *et al.* [Özkurt and Schnitzler, 2011], and Penny *et al.* [Penny et al., 2008] will be considered for comparison in our experiments.

As one can see, there is a long list of methods to quantify CFC in neural time series. Yet, a number of limitations which can significantly affect the outcomes and interpretations of neuroscientific findings exist with these approaches. For example, in typical PAC analysis, a systematic bias rises where one constructs the so-called *comodulogram*. A comodulogram is obtained by evaluating the chosen metric over a grid of frequency f_x and f_y . This bias emerges from the choice of the bandpass filter, which involves the critical choice of the bandwidth Δf_y . It has been reported several times that to observe any amplitude modulation, the bandwidth of the fast oscillation Δf_y has to be at least twice as high as the frequency of the slow oscillations f_x : $\Delta f_y > 2f_x$ [Berman et al., 2012, Dvorak and Fenton, 2014]. As a comodulogram uses different values for f_y , many studies have used a variable bandwidth, by taking a fixed number of cycles in the filters. The bandwidth is thus proportional to the center frequency: $\Delta f_y \propto f_y$. This choice leads to a systematic bias, as it hides any possible coupling below the diagonal $f_y = 2f_x/\alpha$, where $\alpha = \Delta f_y/f_y$ is the proportionality factor. Other studies have used a constant bandwidth Δf_y ; yet this also biases the results towards the low driver frequency f_x , considering that it hides any coupling with $f_x > \Delta f_y/2$. A proper way to build a comodulogram would be to take a variable bandwidth $\Delta f_y \propto f_x$, with $\Delta f_y > 2f_x$. However, this is not common practice as it is computationally very demanding, because it implies to bandpass filter y again for each value of f_x .

Another common issue arises with the use of the Hilbert transform to estimate the amplitude and the phase of real-valued signals. Such estimations rely on the hypothesis that the signals x and y are narrow-band, i.e. almost sinusoidal. However, numerous studies have used this technique on very wide-band signals such as the entire gamma band (80-150 Hz) [Canolty et al., 2006] (see other examples in [Chavez et al., 2006]). The narrow-band assumption is debatable for high frequency activity and, consequently, using the Hilbert transform may yield non-meaningful amplitude estimations, and potentially poor estimations of PAC [Chavez et al., 2006, Dvorak and Fenton, 2014]. Note also that, in this context, wavelet-based filtering is equivalent to the Hilbert transform [Quiroga et al., 2002, Bruns, 2004], and therefore does not provide a more valid alternative option.

Besides these issues of filtering and inappropriate use of Hilbert transforms, Hyafil [Hyafil, 2015] also warned that certain choices of bandwidth Δf_y might mistake phase-frequency coupling for PAC, or create spurious amplitude-amplitude coupling; see also the more recent work in [Aru et al., 2015] for discussion and more practical recommendations for PAC analysis.

Here we advocate that the DAR models detailed in the next sections address a number of the limitations just mentioned. They do *not* use bandpass filter or Hilbert transform on the high frequencies y . They introduce a measure of *goodness of fit*, through the use of a probabilistic signal model whose quality can be assessed by evaluating the *likelihood* of the data under the model. In practice, the likelihood quantifies how much variance of the signal can be explained by the model, and is similar to the R^2 coefficient in generalized linear models (GLM). To the best of our knowledge, the only related model-based approach to measure PAC used GLM [Penny et al., 2008]. With GLM, however, the modeling part is done independently on each signal y_f , which is the band-pass filtered version of y around frequency f . For each of these frequencies f a different model is fitted. By doing so, a GLM approach cannot model the wide-band signal y as it is limited to multiple estimations frequency by frequency bin. This largely limits the use of the likelihood to compare models or parameters. On the contrary, we propose to model y globally, without filtering it in different frequency bands.

To conclude this section, and to position this work in the broader context of modeling approaches for neuroscience data, we would like to stress that our proposed method can be considered as an encoding model for CFC, as opposed to a decoding model [Kay et al., 2008, Naselaris et al., 2011, Huth et al., 2016]. Indeed, our model reports how much empirical data can be explained and by doing so enables us to test neuroscience hypotheses in a principled manner [Naselaris et al., 2011].

Non-linear AR models beyond Neuroscience

The literature on the use of non-linear auto-regressive (AR) models is quite large and covers fields such as audio signal processing and econometrics. For instance, AR models with conditional heteroskedasticity (ARCH [Engle, 1982], GARCH [Bollerslev, 1986]) are extremely popular in econometrics where they are used to model signals whose overall amplitude varies as a function of time. Here, however, in the context of CFC and PAC, one would like to model variations in the spectrum itself, such as shifts in peak frequencies (a.k.a. frequency modulations) or changes in amplitude only within certain frequency bands (a.k.a. amplitude modulations). To achieve this, one idea is to define a linear AR model, whose coefficients are a function of time and change slowly depending on a *non-linear* function of the signal.

The first models based on this idea are SETAR models [Tong and Lim, 1980], which switch between several AR models depending on the amplitude of the signal with respect to some thresholds. To get a smoother transition between regimes, SETAR models have inspired other models like EXPAR [Haggan and Ozaki, 1981] or STAR [Chan and Tong, 1986], in which the AR coefficients change continuously depending on a non-linear function of the past of the signal.

These models share the same underlying motivation as the DAR models described below but, crucially, DAR models can be designed and parametrized to capture PAC phenomena independently of the phase in the driving signal at which the high frequency content is the strongest. In other words, DAR models can work equivalently well if the high frequency peaks are in the troughs, the rising phase, the decreasing phase or the peaks of the low frequency driving signal. Moreover, as DAR models do not require to infer the driving behavior from the signal itself and rather rely on the prior knowledge of the slow oscillation, the inference is significantly faster and more robust.

Methods

In this section, we first define our statistical signal model, explain how we estimate its parameters on a signal y and its driver x , and demonstrate how one can infer hyper-parameters by comparing the likelihood of several models. Then, we detail how to use this model on neurophysiological signals by presenting the preprocessing steps and showing how to make the model invariant to the phase of the coupling. We also detail how to express the power spectral density (PSD) conditionally to the driver's phase, which allows for a fine-grained analysis of the signal's spectral properties and to build comodulograms. Finally, we present our protocol to simulate signal with CFC and the empirical datasets used to validate our method.

DAR models - General presentation

Model definition

An auto-regressive (AR) model specifies that a signal y depends linearly on its own p past values, where p is the *order* of the model:

$$\forall t \in [p+1, T] \quad y(t) + \sum_{i=1}^p a_i y(t-i) = \varepsilon(t) . \quad (1)$$

Here, T is the length of the signal and ε is the *innovation* (or *residual*) modeled with a Gaussian white noise: $\varepsilon(t) \sim \mathcal{N}(0, \sigma(t)^2)$. To extend this AR model to a non-linear model, one can assume that the AR coefficients a_i are non-linear functions of a given exogenous signal x , here called the *driver*:

$$\forall i \in [1, p] \quad a_i(t) = g_i(x(t)) . \quad (2)$$

What was proposed in [Grenier, 2013, Grenier, 1983] is to consider that the non-linear functions g_i are polynomials:

$$\forall t \in [p+1, T] \quad g_i(x(t)) = \sum_{k=0}^m a_{ik} x(t)^k , \quad (3)$$

where $x(t)^k$ is $x(t)$ to the power k . Note that one could imagine more complex models for the g_i , yet empirical evidence confirms that polynomials are rich enough for the different neurophysiological signals considered in this work (see **S1 Fig**).

Inserting the time-dependent AR coefficients (2) into the AR model (1), we obtain the following equation:

$$\forall t \in [p+1, T] \quad y(t) + \sum_{i=1}^p \sum_{k=0}^m a_{ik} x(t)^k y(t-i) = \varepsilon(t) . \quad (4)$$

This can be simplified and rewritten as:

$$\forall t \in [p+1, T] \quad y(t) + A^\top \tilde{Y}(t) = \varepsilon(t) , \quad (5)$$

where $A \in \mathbb{R}^{p(m+1)}$ is a one-dimensional vector composed of the scalars a_{ik} and $\tilde{Y}(t) \in \mathbb{R}^{p(m+1)}$ is a one-dimensional vector composed of the regressors $x(t)^k y(t-i)$. We note \hat{A} the estimated value of A and A^\top its transpose.

We also consider a time-varying innovation variance $\sigma(t)^2$ driven by x . This corresponds to the assumption that the power of the signal at time t depends on the driver at this same instant. Since the standard deviation is necessarily positive, we use the following polynomial model for its logarithm:

$$\forall t \in [p+1, T] \quad \log(\sigma(t)) = \sum_{k=0}^m b_k x(t)^k = B^\top X(t) , \quad (6)$$

where $B \in \mathbb{R}^{m+1}$ is a one-dimensional vector composed of the scalars b_k .

This model is called a driven auto-regressive (DAR) model [Dupré la Tour et al., 2017]. DAR models are thus a natural extension of well-known AR models and enable analyzing non-stationary or non-linear spectral phenomena observed in neurophysiological time-series. One methodological concern is that the AR coefficients cannot be arbitrary as they should define stable models (see [Proakis and Manolakis, 1996], section 3.6). In our case there is an extra difficulty because they are changing over time. To guarantee that each set of coefficients defines a stable filter, it is possible to re-parametrize the model. This requires significantly more complex derivations and leads to a slightly slower optimization, while providing equivalent empirical observations. For simplicity, this manuscript will focus on the simplest model dedicated to CFC analysis, and we refer to [Dupré la Tour et al., 2017, Grenier and Omnes-Chevalier, 1988] for readers interested in more details.

Model estimation

DAR model equation (4) is non-linear with respect to the given signals x and y , yet it is linear with respect to the regressors $x(t)^k y(t-i)$. Therefore after computing the regressors, it is possible to obtain an analytical expression of the parameters in A .

As the innovation $\varepsilon(t)$ is assumed to be a Gaussian white noise, the model likelihood L is obtained via:

$$L = \prod_{t=p+1}^T \frac{1}{\sqrt{2\pi\sigma(t)^2}} \exp\left(-\frac{\varepsilon(t)^2}{2\sigma(t)^2}\right) \quad (7)$$

or

$$-2\log(L) = T\log(2\pi) + \sum_{t=p+1}^T \frac{\varepsilon(t)^2}{\sigma(t)^2} + 2 \sum_{t=p+1}^T \log(\sigma(t)) .$$

DAR models are estimated by likelihood maximization. Here, if the innovation variance $\sigma(t)^2$ is considered fixed, maximizing this function boils down to minimizing the sum of squares of the $\varepsilon(t)$. Since $\varepsilon(t)$ is the residual, the inference of the parameters amounts to maximizing the variance explained by the model. To start the maximization, we first assume that the innovation variance is fixed and equal to the signal's empirical variance. The likelihood then reads:

$$-2\log(L) = C + \sum_{t=p+1}^T \frac{\varepsilon(t)^2}{\sigma(t)^2} = C + \sum_{t=p+1}^T \frac{1}{\sigma(t)^2} \left(y(t) + A^\top \tilde{Y}(t)\right)^2 , \quad (8)$$

where C is a constant. We can thus estimate the DAR coefficients $A \in \mathbb{R}^{p(m+1)}$ by solving the following linear system (a.k.a. normal equations):

$$\left(\sum_{t=p+1}^T \frac{1}{\sigma(t)^2} \tilde{Y}(t)\tilde{Y}(t)^\top\right) \hat{A} = -\left(\sum_{t=p+1}^T \frac{1}{\sigma(t)^2} \tilde{Y}(t)y(t)\right) . \quad (9)$$

Given \hat{A} , one can then estimate the vector $B \in \mathbb{R}^{m+1}$ from the residual $\varepsilon(t)$, maximizing the likelihood using off-the-shelf optimization techniques adapted to differentiable problems. We use a Newton-Raphson procedure (see more details in [Grenier, 2013]). The instantaneous variance is then updated with B using (6). One can then iterate between the estimation of the coefficients A and B . In our experiments, we observed that one or two iterations are sufficient.

Model selection

DAR models are probabilistic, which offers a significant advantage over traditional PAC metrics. Indeed, one can evaluate the likelihood of the model L (7), which quantifies the *goodness of fit* of the data given the model, and select the hyper-parameters that maximize it. To compare models with different number of hyper-parameters (i.e. degrees of freedom), different information criteria can be used such as the Akaike Information Criterion (AIC) [Akaike, 1998] or the Bayesian Information Criterion (BIC) [Schwarz et al., 1978], which read respectively $AIC = -2\log(L) + 2d$ and $BIC = -2\log(L) + d\log(T)$. The integer d is the number of degrees of freedom of the model, which for DAR models equals to $d = (p+1)(m+1)$. The BIC has been used extensively for hyper-parameter selection in AR models, and enables the joint estimation of both p and m (see S2 Fig for simulation results).

As the statistical inference with a DAR model is very fast, one can easily estimate models on a full grid of hyper-parameters (p and m) and keep the model that leads to the lowest BIC. Another alternative used in our experiments is the use of cross-validation (CV), consisting of splitting the recorded signals into two sets: a training signal to estimate parameters, and a test signal to evaluate the likelihood of new data. In practice, CV requires that a large amount of data be left for model testing, and that the two datasets be separated by a minimum delay to ensure data independence [Arlot et al., 2010]. In S3 Fig, we present model selection based on cross-validation, applied on real signals.




Fig 1. Example of a preprocessed signal. The driver x is extracted from the raw signal z with a bandpass filter, and then subtracted from z to give y . A PAC effect is present as we see stronger high frequency oscillations at the peaks of the driver. The signal y is presented here as it looks before the temporal whitening step.

DAR models for CFC estimations on brain data

Data preprocessing: Extracting the driver

In this section, we describe the preprocessing applied to the raw signal z to obtain the signal y and the driver x .

First, the raw signal was downsampled to an appropriate sampling frequency f_s (we used 333 Hz and 625 Hz in our examples), as we only consider frequencies up to $f_s/2$. Power line noise (50 or 60 Hz for French and US datasets, respectively) and its harmonics were removed. As power line noise fluctuates over time, the frequency of the power line noise was estimated by projecting the signal onto sinusoids, and the maximum-power projection was subtracted from the raw signal. When necessary, a high-pass filter was used to detrend the signal (typically at 1 Hz) [Bigdely-Shamlo et al., 2015].

A bandpass filter was then used to extract the driver x . The center frequency f_x and bandwidth Δf_x for this filter will be discussed in details in the Results section. The filter equation is:

$$w(t) = b(t) \cos(2\pi f_x t) , \quad (10)$$

where $b(t)$ is a Blackman window of order $[1.65 * f_s / \Delta f_x] * 2 + 1$, chosen to have a bandwidth of Δf_x at -3 dB (the filter attenuation is 50% at $f_x \pm \Delta f_x/2$). The filter is zero-phase since it is symmetric.

Then, the driver x is subtracted from the raw signal z to create the signal $y = z - x$. Note that the signal y now contains a frequency gap around f_x , which can be a nuisance for the AR estimate that provides a compact model for the broad band power spectrum density of the signal. To solve this issue, we fill this gap by adding a Gaussian white noise filtered with the filter w , and adjusted in energy to have a smooth power spectral density (PSD) in y . Such a preprocessing step is also commonly used when working with multivariate auto-regressive models (MVAR) on neurophysiological signals corrected for power line noise.

Finally, we whiten y with a linear AR model, by applying the inverse AR filter to the signal. This temporal whitening step is not necessary, yet it reduces the need for high order p in DAR and therefore reduces both the computational cost and the variance of the model. After this whitening step, the PSD of y is mostly flat, and DAR models only contain the modulation of the PSD. Fig. 1 shows a time sample of a preprocessed signal, taken from a rodent striatal LFP signal (see below for a description of the data). One can see how the slowing varying driving signal follows the original raw signal and how the high frequencies bursting on the peak of the slow oscillation remain in the processed signal y . The general processing pipeline is summarized in Fig. 2(a).

Model selection: Optimizing the filtering parameters

During the preprocessing, the driver x_c is extracted from the preprocessed signal z through a band-pass filter, with center frequency f_x and bandwidth Δf_x . To select automatically these two parameters, we perform a grid-search and select the driver which yields the highest likelihood. Importantly, the highest likelihood does

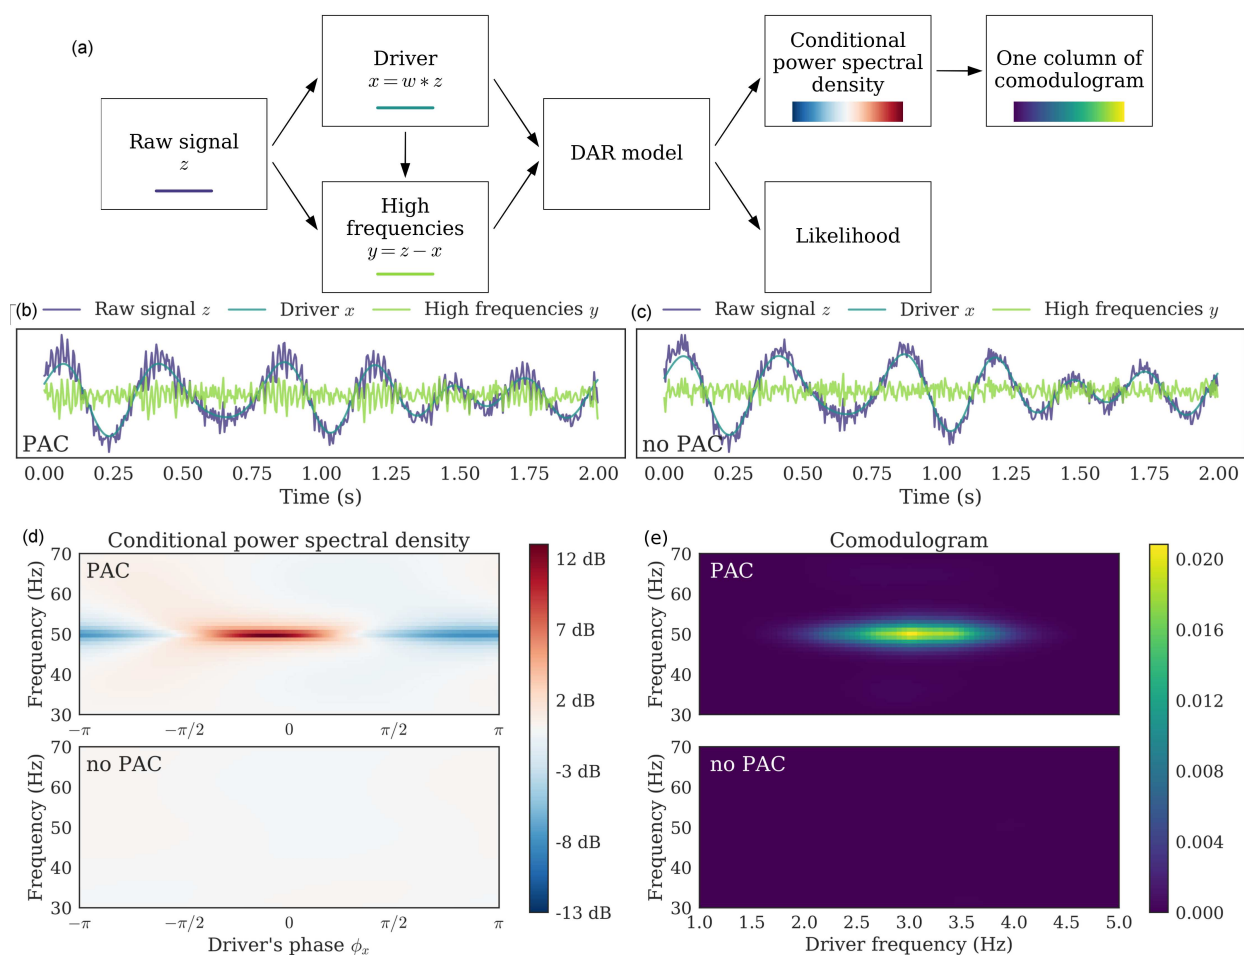


Fig 2. Example on simulated signals: Pipeline, signal, conditional PSD, and comodulogram.

(a) Pipeline of the method. We applied it with $(p, m) = (10, 1)$ on two simulated signals: (b) Simulated signal with PAC and (c) Simulated signal without PAC. (d) From a fitted model, we computed the PSD conditionally to the driver's phase ϕ_x ; each line is centered independently to show amplitude modulation. PAC can be identified in the fluctuation of the PSD as the driver's phase varies: around 50 Hz, the PSD has more power for one phase than for another. This figure corresponds to a single driver's frequency $f_x = 3.0$ Hz. (e) Applying this method to a range of driver's frequency, we build a comodulogram, which quantifies the PAC between each pair of frequencies.

not necessarily correspond to the highest PAC score, but rather to the model with the highest goodness of fit or explained variance. This detail is crucial, as optimizing for the best fit is a legitimate data-driven approach, whereas optimizing for the highest PAC score is statistically more questionable. One additional benefit of our approach is that other types of filters can be compared using a common methodological approach.

For a fair comparison of the drivers, we need to evaluate the model fitting on the exact same signal y . To do so, we remove from y all the possible drivers tested on this grid-search using a high-pass filter above maximum center frequency f_x . Once again, we fill the frequency gap by adding a Gaussian white noise filtered with a low-pass filter complementary to the high-pass filter.

Phase invariance: Removing the bias using a complex driver

We consider PAC where high frequencies are coupled with a low frequency driving signal. We denote by ϕ_0 the preferred phase of the coupling, i.e. the phase of the driver that corresponds to the maximum amplitude of the high frequencies. When $\phi_0 = 0$, the high frequency bursts happen in the peaks of the driver, whereas when $\phi_0 = \pi$, they happen in the troughs.

The driver extracted as described in previous subsection is real-valued, and the filter used to extract it is based on a cosine. As the cosine has the same value for ϕ and $\pi - \phi$, the PAC estimation is biased and underestimated when ϕ_0 is not equal to 0 or π . Indeed, the model mixes the contribution of ϕ and $\pi - \phi$, which attenuates the effect. This bias was already reported in PAC estimation by [Bruns and Eckhorn, 2004], and the solution proposed by [Penny et al., 2008] was to use both cosine and sine to have PAC estimates that are invariant to the preferred phase ϕ_0 of the low frequency oscillation.

The same technique can be used on the DAR models. We not only filter the raw signal with $w(t) = b(t) \cos(2\pi f_x t)$ to obtain x , but also with $\bar{w}(t) = b(t) \sin(2\pi f_x t)$ to obtain \bar{x} , creating a complex-valued driver $x_c = x + j\bar{x}$ where j denotes the complex number, $j^2 = -1$. With a complex-valued driver, DAR models are naturally extended by adding more regressors:

$$\forall i \in [1, p] \quad a_i(t) = \sum_{0 \leq k+l \leq m} a_{ikl} x(t)^k \bar{x}(t)^l . \quad (11)$$

Instead of $p(m+1)$ regressors, we now have $p(m+1)(m+2)/2$ regressors, and the number of degrees of freedom of the model is now $d = (p+1)(m+1)(m+2)/2$. Typical values for m are below 3, so the number of parameters stays within a reasonable range despite the squared dependence in m . The number of parameters in the model remains much lower than the number of time samples and the estimation problem stays therefore well-posed. The innovation variance model is also updated accordingly:

$$\log(\sigma(t)) = \sum_{0 \leq k+l \leq m} b_{kl} x(t)^k \bar{x}(t)^l \quad (12)$$

A validation result is presented in the supporting information (S4 Fig), showing that indeed it removes the preferred-phase bias.

Conditional spectrum: Deriving the PSD from the model

After preprocessing and model estimation, we can compute the *conditional* power spectral density (PSD) of the model, noted $\text{PSD}_y(x_{c0})$ since it is here a function of the driver's value x_{c0} . For a given complex driver's value x_{c0} , we compute the AR coefficients $a_i(x_{c0})$ using Eq. (11), along with the innovation's standard deviation $\sigma(x_{c0})$ using Eq. (12). Since AR models with time-varying coefficients are locally stationary [Dahlhaus, 1996], we can compute the PSD at a frequency f with:

$$\text{PSD}_y(x_{c0})(f) = \sigma(x_{c0})^2 \left| \sum_{i=0}^p a_i(x_{c0}) e^{-j2\pi f i} \right|^{-2} \quad (13)$$

where $j^2 = -1$ and $a_0(x_0) = 1$. We thus have a different PSD for each driver value x_c , i.e. for each time instant, as the driver $x_c(t)$ fluctuates in time. Note that the PSD is a function of the driver value and not really its phase, unless the driver has been normalized in amplitude. The varying AR coefficients $a_i(t)$ parametrize the fluctuations of the shape of the PSD, whereas the varying innovation's standard deviation $\sigma(t)$ is responsible for the absolute fluctuations in power over all frequencies.

In Fig. 2(b and c), one can see two simulated signals: one with PAC on the left (b) and one without PAC on the right (c). The simulation process, which does not make use of DAR model, will be detailed and discussed in a following section. In Fig. 2(d), one can see an example of $\text{PSD}_y(x_c)$ as a function of the driver's value $x_c = \rho \exp(j\phi)$. In practice we set ρ to the median of the driver's amplitude, $\rho = \text{median}(|x_c|)$, and ϕ spans the entire interval $[-\pi, \pi]$. The PAC is identified by the modulation of the spectrum with respect

to the driver’s phase. Note that with DAR models, the amplitude modulation is modeled jointly on the entire spectrum, as opposed to frequency band by frequency band in most other PAC metrics.

Using the fast Fourier transform (FFT) in Eq. (13), we are able to compute the PSD for a broad range of frequencies $f \in [0, f_s/2]$. As model estimation is also very fast, we obtain the conditional PSD very quickly: With $(p, m) = (10, 1)$ on a signal with $T = 10^5$ time points, the entire DAR estimation and PSD computation lasts 0.12s.

Comodulogram: Quantifying the modulation at each frequency

We now detail how comodulograms can be derived from DAR models. To quantify the coupling at a given frequency f , we compute $\text{PSD}_y(x_c)(f)$ for a range of n artificial driver’s values $x_c(k)$ located on a circle of radius $\rho = \text{median}(|x_c(t)|)$, and we normalize it to sum to one:

$$\forall k \in [1, n] \quad p_f(k) = \frac{\text{PSD}_y(\rho e^{j2\pi k/n})(f)}{\sum_{k=1}^n \text{PSD}_y(\rho e^{j2\pi k/n})(f)} . \quad (14)$$

Then, we use the same method as [Tort et al., 2010] to measure the fluctuation of $p_f(k)$: we compute the Kullback-Leibler divergence between $p_f(k)$ and the uniform distribution $q(k) = 1/n$, and we normalize it with the maximum entropy $\log(n)$:

$$M(f) = \frac{1}{\log(n)} \sum_{k=1}^n p_f(k) \log \left(\frac{p_f(k)}{q(k)} \right) . \quad (15)$$

This metric M is between 0 and 1 and allows us to quantify the non-flatness of $p_f(k)$, i.e. the amplitude modulation, frequency by frequency. We can then build a *comodulogram*, a figure commonly used in the literature which depicts the PAC metric for a grid of frequency (f_x, f) . When derived from DAR models as described here, the comodulogram is not affected by the systematic biases presented in the introduction.

Fig. 2(e) shows the comodulogram of the two simulated signals, one with and one without PAC. As expected the DAR model reports no coupling when there is none, and reports here a strong coupling between 3 and 50 Hz when it is actually present.

Directionality: Estimating a delay between the signal and the driver

The DAR model in Eq. (4) uses the driver at time t to parametrize the AR coefficients with the assumption that, in PAC, high frequency activity is modulated by the driving signal with no delay. However, one could legitimately ask whether slow fluctuations in oscillatory activity precedes the amplitude modulation of the fast activity or, conversely, whether high frequency activity precedes low frequency fluctuations. This is what we will refer to as *directionality estimation*.

To assess the problem of directionality estimation with DAR models, we capitalize on the likelihood function. After observing some PAC between y and x , we introduce a delay τ in the driver so as to estimate the goodness of fit of a DAR model between y and x_τ , where $x_\tau(t) = x(t - \tau)$. Using a grid of values for the delay, one can report the value τ_0 that corresponds to the maximum likelihood. In practice, as we use non-causal filters to extract the driver, and because the AR model uses samples from the past of the y , one might wonder if there is any positive or negative bias in our delay estimation. To address this issue, we apply our analysis on both forward and time-reversed signals, and sum the two models log-likelihood. In this way, the filtering bias is strongly attenuated because it similarly affects both forward and time-reversed models. If the best delay is positive, it means that the past driver yields a better fit than the present driver, i.e. that the slow oscillation precedes the amplitude modulation of the fast oscillation. Inversely, a negative delay means that the amplitude modulation of the fast oscillation follows changes in the slow oscillation.

It is worth noting that the estimation of the delay τ_0 shall not be considered as an alternative way to estimate the preferred phase ϕ_0 defined in a previous section. Although the preferred phase and the delay would be identical should the driver be a perfect stationary sinusoid, in non-stationary neural systems in which the driver’s instantaneous frequency may fluctuate with time, the preferred phase and delay will likely

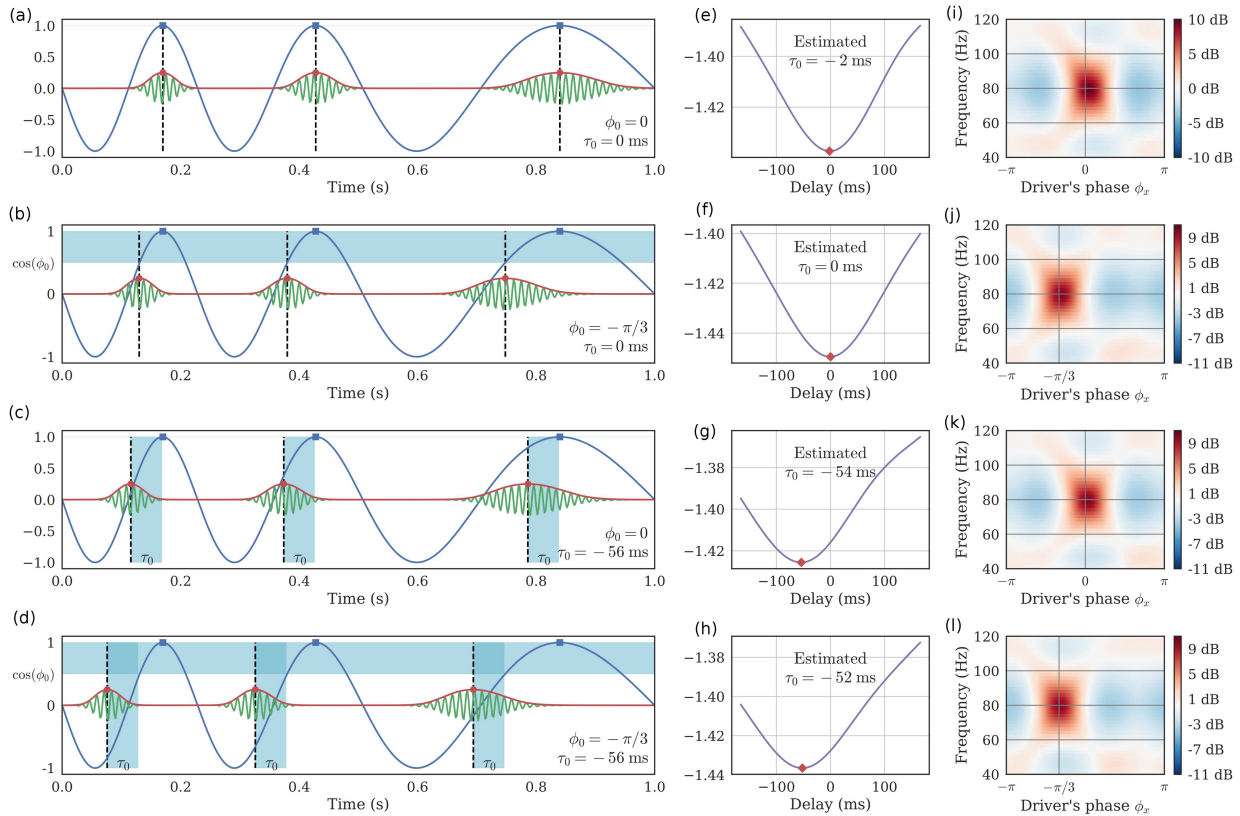


Fig 3. Preferred phase and temporal delay are different. The temporal delay τ is distinct from the preferred phase ϕ_0 . (a) When both are equal to zero, the high frequency bursts happen in the driver's peaks. (b) When $\tau = 0$ and $\phi \neq 0$, the bursts are shifted in time with respect to the driver's peaks, and this shift *varies* depending on the instantaneous frequency of the driver. (c) When $\tau \neq 0$ and $\phi = 0$, the bursts are shifted in time with respect to the driver's peaks, and this shift is *constant* over the signal. In this case, note how the driver's phase corresponding to the bursts varies depending on the instantaneous frequency of the driver. (d) τ and ϕ_0 can also be both non-zero. (e-h) Negative log-likelihood of DAR models, fitted with different delays between the driver and the high frequencies. The method correctly estimates the delay even when $\phi_0 \neq 0$. (i-l) PSD conditional to the driver's phase, estimated through a DAR model with the best estimated delay. The maximum amplitude occurs at the phase ϕ_0 .

differ. Such a scenario can occur either because signal waveforms are not perfect sinusoids [Cole et al., 2016] or because the driver actually changes [Tort et al., 2008]. In other words, τ_0 is a time delay which is identical at all time and corresponds to different phase shifts $\phi(t) = \tau_0 * (2\pi f(t))$ which depend on the instantaneous frequency $f(t)$. On the contrary, ϕ_0 is a phase shift that is constant over time and which corresponds to different time delays $\tau(t) = \phi_0 / (2\pi f(t))$ which also depend on the instantaneous frequency.

Fig. 3 illustrates these specific points and disentangles the two distinct notions.

Like most PAC metrics, DAR models are invariant with respect to the preferred phase. Both the strength of the coupling and the model likelihood are unchanged with respect to ϕ_0 . On the contrary, all PAC metrics including DAR models are strongly affected by a time delay τ when the driver is not a perfect sinusoid. This attenuates the coupling and can potentially modify artificially the preferred phase. When the delay is too large, all metrics would measure zero coupling. This is in fact what justifies the use of surrogate techniques that introduce a large time shift to quantify the variance of the measure in the absence of coupling [Canolty et al., 2006, Tort et al., 2010, Aru et al., 2015]. Fig. 3(e-h) provides the delay estimated with DAR models, obtained by maximizing the likelihood over a grid of delays, as explained previously. In Fig. 3(i-l), we show

the conditional PSD of the model obtained with the best estimated delay as well as the estimated preferred phase ϕ_0 . Using the likelihood, our method is thus able to estimate this delay, thus improving PAC estimation and extending PAC analysis.

The question of delay estimation has been previously addressed in the literature with measures borrowed from information theory [Besserve et al., 2010]. For instance, Transfer entropy (TE) [Schreiber, 2000, Wibral et al., 2013, Park et al., 2015] was adapted to PAC by directly computing the envelop of both the slow and the fast oscillations. When variables follow a Gaussian distribution, TE is equivalent to Granger causality [Barnett et al., 2009], but Granger causality can fail when the signal-to-noise ratios of the two signals are different [Nolte et al., 2010], which is often the case for the driver and the envelop of the high frequencies. Hence, it remains to be investigated how TE compares with DAR models in terms of performance. It has also been addressed with another method called Cross-Frequency Directionality (CFD) [Jiang et al., 2015], which is based on the phase slope index (PSI) [Nolte et al., 2008]. CFD is not designed to measure a delay; thus, we modified it so as to compare it *qualitatively* to our approach (see Results). If such delay estimation results may not reflect pure causality [Aru et al., 2015], for instance because of the transitivity of correlation, they nevertheless improve the analysis of PAC going one step further by estimating the delay between the coupled components.

Data

Simulated signals

In our experiments, methods comparisons and model evaluations, we make an extensive use of simulated signals showing some PAC. The simulations are generated as follows: we simulate a coupling between $f_x = 3$ Hz and $f_y = 50$ Hz, with a sampling frequency $f_s = 240$ Hz, during T time points (T depends on the experiment). We do not use a perfect sinusoidal driver, as using such an ideal oscillatory signal is over-simplistic. Drivers can indeed have a larger band as suggested next by our experiments. Finally strong empirical evidence suggests that neurophysiological signals have complex morphologies that can be overlooked when studied as ideal sinusoids [Cole and Voytek, 2017].

To construct a time-varying driver peaking at 3 Hz, we bandpass filter a Gaussian white noise at a center frequency $f_x = 3$ Hz and with a bandwidth $\Delta f_x = 1$ Hz, using the same filter detailed in the preprocessing section. The smaller the bandwidth, the closer the driver is to a perfect sinusoid. We normalize the driver to have unit standard deviation $\sigma_x = 1$. We then modulate the amplitude a_y of a sinusoid at, for example, $f_y = 50$ Hz, using a sigmoid on the driver x :

$$a_y(t) = \frac{1}{1 + \exp(-\lambda x(t))} , \quad (16)$$

with a sharpness set to $\lambda = 3$. By doing so, the amplitude varies between 0 and 1 depending on the driver's value. We normalize the signal y to have a standard deviation $\sigma_y = 0.4$. We then add to the signal y both the driver x and a Gaussian white noise with a standard deviation $\sigma_\epsilon = 1$. Note that this simulation procedure is not based on a DAR model. In other words, we do not validate our model using signals that fall perfectly into the category of stochastic signals that are synthesized with a DAR model.

Empirical datasets

We tested DAR models on three different datasets, with invasive recordings in humans and rodents:

1. A rodent local field potential (LFP) recording in the dorso-medial striatum from [Dallérac et al., 2017] (supplementary figure 2), (1800 seconds down-sampled at 333 Hz)
2. A rodent LFP recording in the hippocampus from [Khodagholy et al., 2015] collected during rapid eye movement (REM) sleep (100 seconds down-sampled at 625 Hz)
3. A human electro-corticogram (ECoG) channel in auditory cortex from [Canolty et al., 2006] (730 seconds down-sampled at 333 Hz)

We refer the reader to corresponding articles for more details on the recording modalities of these neurophysiological signals.




Fig 4. Driver's frequency and bandwidth selection (simulations). (a,b) Two examples of grid-search over both f_x and Δf_x , over simulated signals with (a) $\Delta f_x(\text{simu}) = 0.4$ Hz and (b) $\Delta f_x(\text{simu}) = 1.6$ Hz. For all bandwidths (except 6.4 Hz), the negative log-likelihood is minimal at the correct frequency $f_x = 4$ Hz. (c) To see more precisely the bandwidth estimation, we plot the negative log-likelihood (relative to each minimum for readability), for several ground-truth bandwidth $\Delta f_x(\text{simu})$, with $f_x = 4$ Hz. For each line, the minimum correctly estimates the ground-truth bandwidth (depicted as a diamond), showing empirically that the parameter selection method gives satisfying results.

Results

We now illustrate the benefits of DAR models using simulated data and neurophysiological recordings in rats and humans. Our primary goal is to demonstrate that assessing the validity of the model through its likelihood is a statistically principled approach that allows to reveal new insights on empirical data with interesting neuroscientific consequences.

Model selection and data-driven estimation of the driving signal

In this section, we present the outcome of using the model selection procedure to estimate the best filtering frequency f_x and bandwidth Δf_x to extract the driver x . We first describe the outcome on simulated signals (ground truth) and then on empirical datasets.

Simulation study: Validating the approach

To validate our approach, we tested on simulated signals whether both center frequency f_x and bandwidth Δf_x were correctly estimated. For this, we simulated signals using $f_x(\text{simu}) = 4$ Hz and $\Delta f_x(\text{simu}) \in [0.2, 0.4, 0.8, 1.6, 3.2]$ Hz. To mimic the duration of our three real signals, we used $T = [100f_s]$ time points, for a duration of 100 seconds.

As the noise mainly affected the amplitude modulation but barely altered the driver, we also added a Gaussian white noise low-pass filtered at 20 Hz, and scaled to have a PSD difference of 10 dB at $f_x = 4$ Hz between the driver and the noise. In this way, the driver was also altered by noise.

Given the simulated signal, we performed a grid-search over f_x and Δf_x , extracting the driver, fitting the model and computing the likelihood of the model. Results are presented in Fig. 4. As expected, the center

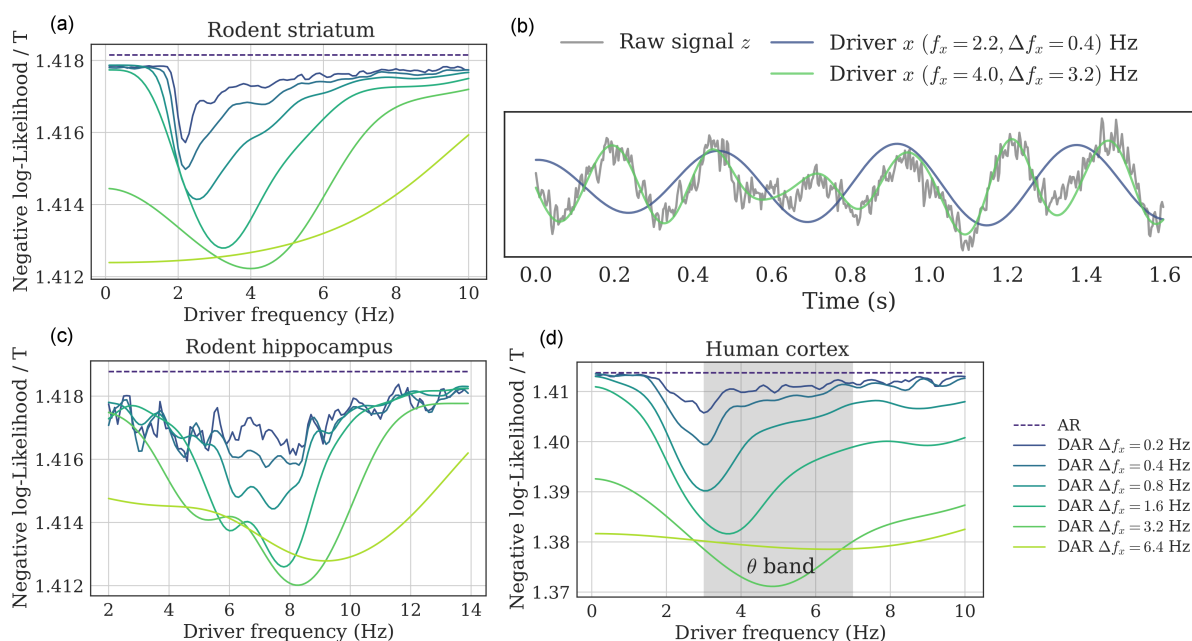


Fig 5. Driver's frequency and bandwidth selection (real signals). (a,c,d) Negative log-likelihood of the fitted model for a grid of filtering parameters f_x and Δf_x : (a) rodent striatum, (c) rodent hippocampus, (d) human auditory cortex. The optimal bandwidth was very large (3.2 Hz), and the optimal center frequency changed as the bandwidth increased, suggesting that the optimal driver had a wide asymmetrical spectrum. (b) This portion of the rodent striatal signal shows two examples of driver with different bandwidths: The wide-band driver better follows the raw signal, and *independently* also leads to a better fit in DAR models.

frequency f_x was correctly estimated for all bandwidths Δf_x . More importantly, the negative log-likelihood was minimal at the correct simulated bandwidth.

This simulation confirms that the likelihood can be used to estimate the correct parameters for the driver's filtering step, and that does not present any obvious bias in the estimation.

Driver estimates on human and rodent recordings

The outcome of the model selection procedure on the three real neurophysiological signals are reported in Fig. 5. Two general observations can be made. First, for all three signals, the optimal bandwidth was relatively large (3.2 Hz). Second, the optimal center frequency changed as we increased the bandwidth. Interestingly, this phenomenon was not observed in the simulation study, suggesting that in real data the optimal driver is wide-band and has an asymmetrical spectrum. In other words, the driver's frequency is not precisely defined, and a large band-pass filter should be preferred to extract the driver.

These observations thus question which parameters should be used in practice. The classical approach is to build a comodulogram to select the best driver frequency f_x , choosing arbitrarily the bandwidth Δf_x . This bandwidth is typically quite narrow, e.g. 0.4 Hz, and chosen to clearly isolate the maximum frequency in the comodulogram. However, this approach relies on the assumption that the driver is nearly sinusoidal, which is unrealistic [Cole and Voytek, 2017]. On the contrary, our data-driven approach selects the frequency and bandwidth that lead to the highest goodness of fit of our model.

To further describe the implications of our observations, we focused on the rodent striatal LFP signal. We compared an arbitrary choice of driver bandwidth (0.4 Hz) and the optimal choice with respect to the model likelihood (3.2 Hz). For these two bandwidths, we selected the optimal center frequencies, 2.2 Hz and 4.0 Hz respectively. Fig. 5b shows a time sample of the raw signal and the two extracted drivers. The wide-band driver (in green) followed very well the raw signal, which was not a perfect sinusoid. On the contrary, the

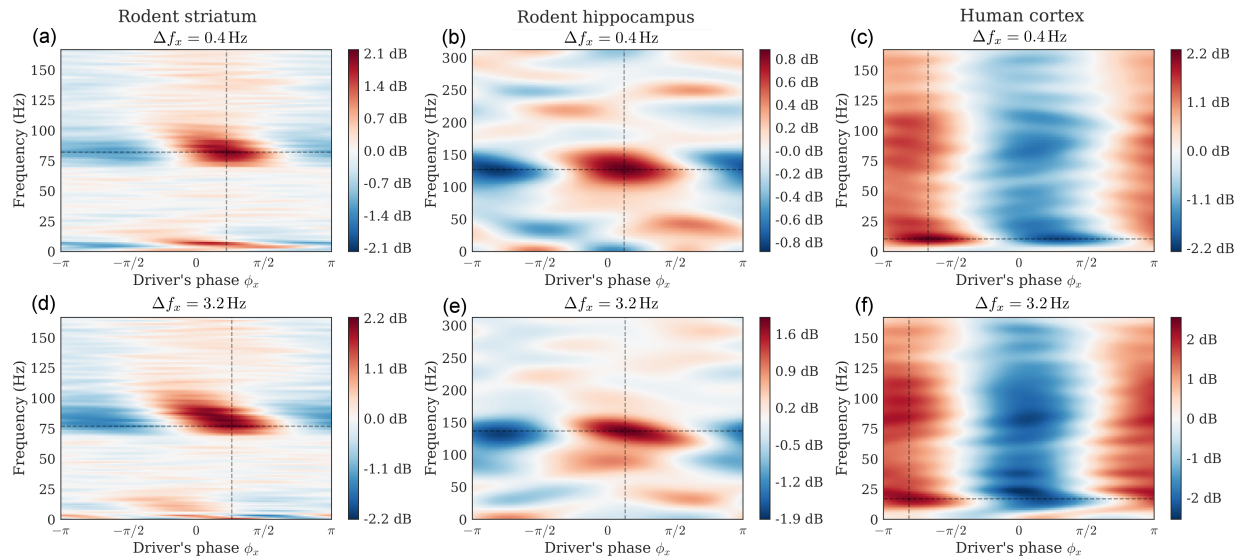


Fig 6. PSD conditional to the driver's phase. Dataset: Rodent striatum (a,d), rodent hippocampus (b,f), human auditory cortex (c,g). We derive the conditional PSD from the fitted DAR models. In one plot, each line shows at a given frequency the amplitude modulation with respect to the driver's phase. The driver bandwidth Δf_x is 0.4 Hz (top row) and 3.2 Hz (bottom row). Note that the maximum amplitude is not always at a phase of 0 or π (i.e. respectively the peaks or the troughs of the slow oscillation). In figure (d), we can also observe that the peak frequency is slightly modulated by the phase of the driver (phase-frequency coupling).

narrow-band driver (in blue) seemed poorly related to the slow oscillation of the raw-signal. To know which of these two slow varying signals best captured the temporal amplitude modulations of the high frequencies, we fitted a model using each of these two drivers. We used the *same* high-pass filtered signal y , i.e. where we removed all frequencies below 16 Hz.

We found that the model with the wide-band driver explained more variance in the high frequencies than the model with the narrow-band driver, that is, that the model better explained the amplitude fluctuations of the high frequencies. As we used the same signal y for both models and as the driver lied in different frequency intervals, it should be noted that the fitting to the slow oscillation visible in Fig. 5(b) and the quality of the fit of the amplitude modulation in the high frequencies were two independent observations. By optimizing for the latter, one observe in Fig. 5b that it leads to a more realistic extraction of the driving neural oscillation.

More generally, our results emphasize the complexity of the choice of parameters in PAC analysis. With our method, we were able to easily compare different parameters, even on non-simulated data, which offered a principled way to set them and helped avoiding misinterpretation due to bad choices. We now investigate the conditional PSDs and the comodulograms estimates on these three datasets.

PSDs and comodulograms estimates of human and rodent neurophysiological recordings

Fig. 6 shows for each signal the PSD (in dB) depending on the driver's phase, estimated through DAR models. The hyper-parameters (p, m) for the rodent striatal, the rodent hippocampal and the human cortical data were (90, 2), (15, 2), and (24, 2), respectively. These parameters were chosen by cross-validation with an exhaustive grid search: $p \in [1, 100]$ and $m \in [0, 3]$. The filtering parameters $(f_x, \Delta f_x)$ were chosen to maximize the likelihood as described in previous section, and are respectively (8.2, 3.2) Hz, (5.2, 3.2) Hz and (4.0, 3.2) Hz (Fig. 6 bottom).

For comparison, we also show the PSD obtained with a bandwidth $\Delta f_x = 0.4$ Hz, with the center frequency chosen to maximize the likelihood. We used (6.4, 0.4) Hz, (3.2, 0.4) Hz and (2.2, 0.4) Hz (Fig. 6




Fig 7. Comodulograms with parameters optimizing the likelihood. Dataset: rodent striatum (top), rodent hippocampus (middle) and human auditory cortex (down). Methods, from left to right: [Ozkurt *et al.* 2011] [Özkurt and Schnitzler, 2011], [Penny *et al.* 2008] [Penny *et al.*, 2008], [Tort *et al.* 2010] [Tort *et al.*, 2010], DAR models (our). The DAR model parameters and the driver bandwidth are chosen to be optimal with respect to the likelihood.

top), respectively.

We observed two kinds of phenomena. In the rodent striatal and hippocampal data, the coupling was mainly concentrated around a given high frequency (i.e. 80 Hz and 125 Hz, respectively), whereas in the human cortical data, the coupling was observed at all frequencies, with a maximum around 20 Hz. The smoothness of the figures depends on the parameter p : a low value leads to a smoother PSD. Please note that the interpretation of the results depends on the filtering parameters.

The two phenomena can also be visualized in Fig. 7, where we plotted comodulograms for all three signals, computed with four different PAC metrics: the mean vector length first proposed in [Canolty *et al.*, 2006] and updated by [Özkurt and Schnitzler, 2011], the parametric model of [Penny *et al.*, 2008] based on GLM, the modulation index of [Tort *et al.*, 2010] that uses the Kullback-Leibler divergence, and our method based on DAR models.

In DAR models, the hyper-parameters used are the cross-validated values described above. With other methods, high frequencies were extracted with a bandwidth Δf_y twice the highest driver frequency used: 14 Hz, 28 Hz, and 20 Hz respectively. For all methods, the drivers were extracted with the optimal bandwidth $\Delta f_x = 3.2$ Hz.

The resulting comodulograms did not look like typically reported comodulograms, as we used a much larger bandwidth (3.2 Hz) than what was found in the literature. With an arbitrary choice of parameters, we could have obtained more classical comodulograms, as shown in Fig. 8. However, these results could be misleading because they suggest a coupling between sinusoidal oscillators, although the corresponding drivers are not perfectly sinusoidal (see Fig. 5(b)).

On these comodulograms, the four methods yielded comparable results as the data we used were very long: 1800, 100, and 730 seconds for the rodent striatal, the rodent hippocampal and the human cortical data, respectively. However, as we will see in the next section, differences emerge between methods when the signals are shorter.

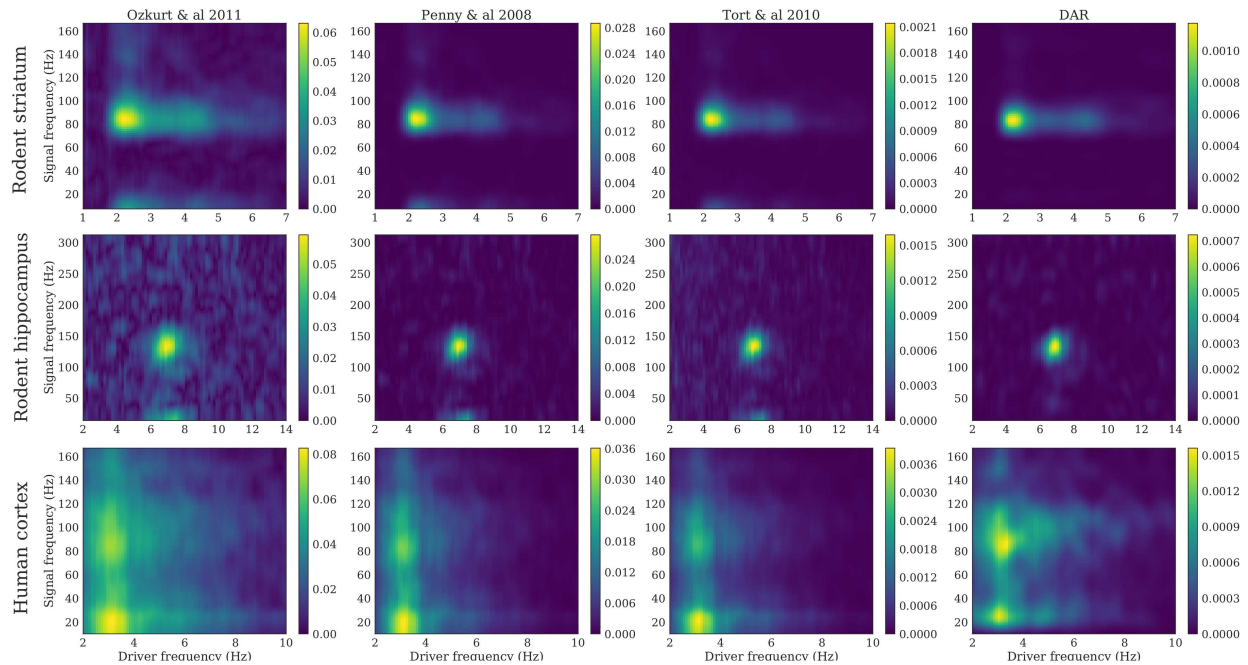


Fig 8. Comodulograms with parameters optimizing the comodulogram sharpness. Dataset: rodent striatum (top), rodent hippocampus (middle) and human auditory cortex (down). Methods, from left to right: [Ozkurt *et al.* 2011] [Özkurt and Schnitzler, 2011], [Penny *et al.* 2008] [Penny *et al.*, 2008], [Tort *et al.* 2010] [Tort *et al.*, 2010], DAR models (our). The driver bandwidth is chosen to have a well defined maximum in driver frequency: $\Delta f_x = 0.4$ Hz. The DAR model parameters are chosen to give similar results than the other methods: $(p, m) = (10, 2)$.

Statistical power and robustness to small sample sizes

Given that DAR models are parametric with a limited number of parameters to estimate, less time samples may be needed to estimate PAC as compared to non-parametric methods. We tested this assumption using simulated signals of varying duration. We computed their comodulograms (as in Fig. 7) and selected the frequencies of maximum coupling. For each duration, we simulated 200 signals, and plotted the 2D histogram showing the fraction of time each frequency pairs corresponded to a maximum. We then compared the same four methods: DAR models with $(p, m) = (10, 1)$, the GLM-based model [Penny *et al.*, 2008], and two non-parametric methods [Tort *et al.*, 2010, Özkurt and Schnitzler, 2011]. Results shown in Fig. 9 show that parametric approaches provided a more robust estimation of PAC frequencies with short signals ($T = 2$ sec) than non-parametric methods.

The robustness to small sample size is a key feature of parametric models, as it significantly improves PAC analysis during shorter experiments. When undertaking a PAC analysis across time using a sliding time window, parametric models should therefore provide more robust PAC estimates. Note that the specific time values in these simulations should not be taken as general guidelines as they depend on the simulation parameters such as the signal-to-noise ratio. However, across all tests, parametric methods consistently provided more accurate results than non-parametric ones.

The amplitude fluctuations of the driver matter

One can note that in DAR models, the driver contains not only the phase of the slow oscillation, but also its amplitude. As the driver is not a perfect sinusoid, its amplitude fluctuates with time. On the contrary, most PAC metrics discard the amplitude fluctuations of the slow oscillation and only consider its phase. To evaluate these two options, we compared two drivers using DAR models: the original driver $x_c(t)$, and the

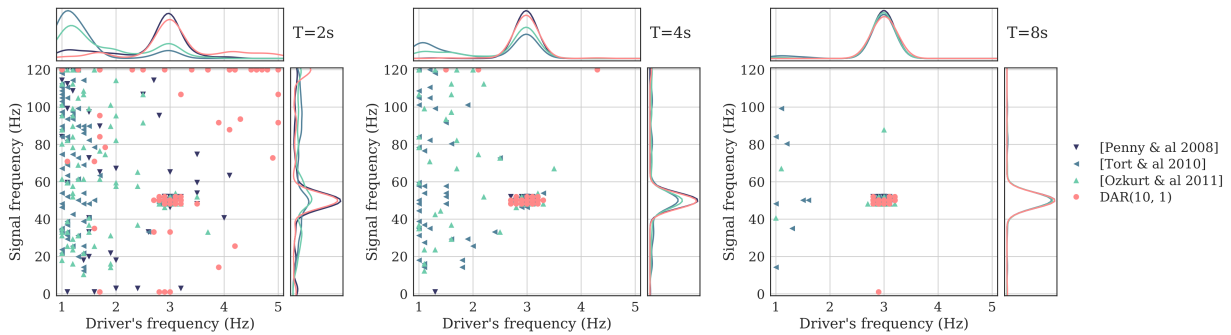


Fig 9. Robustness to small samples. Frequencies of the maximum PAC value, with four methods: a DAR model with $(p, m) = (10, 1)$, the GLM-based model from Penny *et al.* [Penny *et al.*, 2008], and two non-parametric models from Tort *et al.* [Tort *et al.*, 2010] and Ozkurt *et al.* [Özkurt and Schnitzler, 2011]. Each point corresponds to one signal out of 200. Kernel density estimates are represented above and at the right of each scatter plot. From left to right, the simulated signals last $T = 2, 4$ and 8 seconds. The signals are simulated with a PAC between 3 Hz and 50 Hz. The DAR models correctly estimate this pair of frequency even with a short signal length, as well as the GLM-based metric [Penny *et al.*, 2008], while the two other metrics [Tort *et al.*, 2010, Özkurt and Schnitzler, 2011] are strongly affected by the small length of the signals, and do not estimate the correct pair of frequency.

normalized driver $\tilde{x}_c(t) = x_c(t)/|x_c(t)|$. This normalized driver only contains the phase information, as in most traditional PAC metrics. Using cross-validation, we compared the log-likelihood of four fitted models, and found a difference always in favor of the non-normalized driver $x_c(t)$, as it can be visualized in S3 Fig. This result shows that the coupling phenomenon is associated with amplitude fluctuations, a kind of phase/amplitude-amplitude coupling, as it was previously observed in [van Wijk *et al.*, 2015]. Indeed, the GLM parametric method [Penny *et al.*, 2008] was improved when taking into account the amplitude of the slow oscillation. Here, we use our generative model framework to provide an easy comparison tool through the likelihood, to validate this neuroscientific insight from the signals.

Directionality and delay estimation

In this section, we report the results of the directionality estimates using both simulations and neurophysiological signals. It is noteworthy that in DAR models, we arbitrarily call driver the slow oscillation although the model makes no assumption on the directionality of the coupling.

Simulations and comparisons

To validate the directionality estimations approach, we simulated signals with $T = 1024$ time points (4 seconds). We introduced a delay between the slow oscillation and the amplitude modulation of the fast oscillation, and verified that our method could correctly estimate the delays. We did not use a perfect sinusoidal driver, as the delay would only end up in a phase shift, and would not change the strength of the PAC. We used a bandwidth $\Delta f_x = 2$ and a noise level $\sigma_\varepsilon = 0.4$.

We compared our method, using a DAR model with $(p, m) = (10, 1)$, to the cross-frequency directionality (CFD) approach described in [Jiang *et al.*, 2015]. This method makes use of the phase slope index (PSI) [Nolte *et al.*, 2008] for PAC estimation. In [Jiang *et al.*, 2015], the PSI over a frequency band F is defined as:

$$\text{PSI}(F) = \Im\left(\sum_{f \in F} C_{xy}(f) * C_{xy}(f + \delta_f)\right) \quad (17)$$

where C_{xy} is the complex coherence, \Im is the imaginary part, and δ_f is the frequency resolution. As the PSI

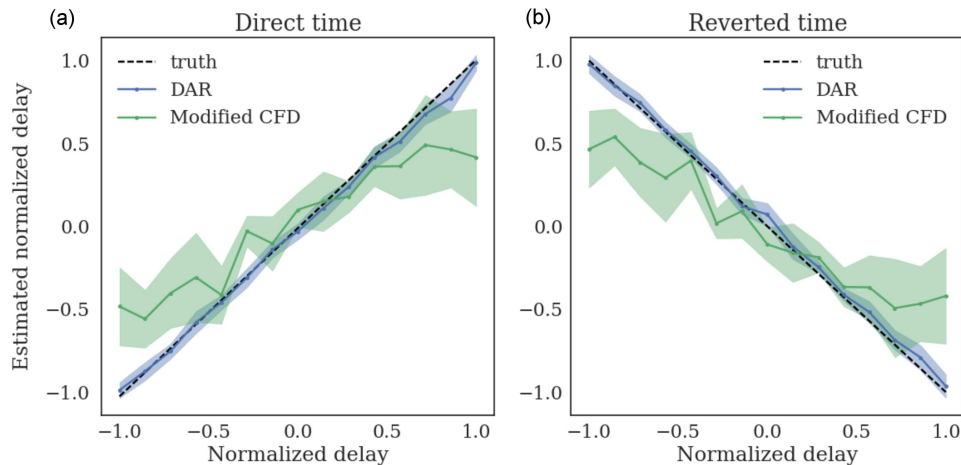


Fig 10. Estimating delays over 20 simulations for each delay. The delays are estimated on both direct time (a) and reverted time (b). The delays are normalized: $\tau = 1$ corresponds to one driver oscillation, i.e. $1/f_x$ sec. The modified CFD shows a bias and only serves for qualitative comparison. The delay estimation based on DAR models correctly estimates the delays, with no visible bias.

was not designed to provide an estimation of the delay, we *modify* it into:

$$\tau_{\text{PSI}}(F) = \frac{\text{PSI}(F)}{2\pi\delta_f n_F}, \quad (18)$$

where n_F is the number of frequencies in set F . Note that this delay estimator was correct only if the coherence was almost perfect $C_{xy}(f) \approx 1$ and if the phase slope was small enough to have $\sin(\phi(f + \delta_f) - \phi(f)) \approx \phi(f + \delta_f) - \phi(f)$. As these assumptions are rarely met in practice, we only used this estimator to compare *qualitatively* with our own estimator.

Results presented in Fig. 10 show the mean estimated delay with ± 1 standard deviation computed over 20 simulations. The delay was correctly estimated in both time directions, with no visible bias. As expected, the modified CFD was biased, and showed a higher variance than the DAR-based approach.

Delay estimates of human and rodent data

We applied this delay estimation method on the three experimental recordings, with both direct time and reverted time. Once again, note that the comparison with the CFD [Jiang et al., 2015] is only qualitative, since the CFD was not designed as a delay estimator. The temporal delay was estimated with DAR models with the parameters obtained from the cross-validation selection. The driver we used was obtained with the best filtering parameters found in our grid-search: $(f_x, \Delta f_x)$ are respectively (4.0, 3.2) Hz, (8.2, 3.2) Hz, and (5.2, 3.2) Hz.

The results are presented in Fig. 11, where we display the delay estimated with maximum likelihood over multiple DAR models. We also display error bars which correspond to the standard deviation obtained with a block-bootstrap strategy [Carlstein, 1986]. This strategy consists in splitting the signal into $n = 100$ non-overlapping blocks of equal length, to draw at random (with replacement) n blocks, and to evaluate the delay on this new signal. We repeated this process 20 times and computed the standard deviation of the 20 estimated delays. Such strategy is used to estimate empirically the variance of a general statistic from stationary time series.

First, one can observe that the results were qualitatively similar between the two methods. The CFD was biased toward zero for large delays, as was previously shown in the simulations. One can also note that the directionality was not always the same. For the rodent striatal and the human cortical data, the delay was negative and the best model fit happened between the signal y and the *future* driver x . Inversely, in the

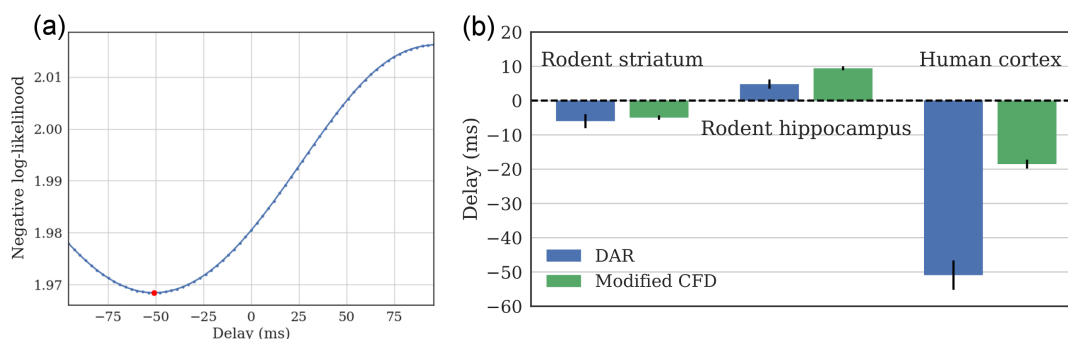


Fig 11. Estimated delays on rodent and human datasets. (a) Negative log likelihood for multiple delays, with the human cortical signal. (b) Optimal delay for the three signals, computed with model selection on DAR models, and with a CFD method modified to provide a delay. The error bars indicate the standard deviation obtained with a block-bootstrap strategy. For the rodent hippocampal data, the delay is positive: the low frequency oscillation precedes the high frequency oscillation. For the human cortical data and the rodent striatal data, the delay is negative: The high frequency oscillation precedes the low frequency oscillation.

rodent hippocampal data, the delay was positive, and the best model fit was between the signal y and the *past* driver x . Further experiments need to be performed to better understand the origin of such delays but we demonstrate here the use of DAR models to estimate them.

Note that our method does not select the delay that leads to the maximum PAC, which would not be statistically valid. On the contrary, we select the delay that leads to the best fit of the model on the data. This approach is more rigorous since we maximize the variance explained by our model, and not the effect of interest.

Discussion

Cross-frequency coupling (CFC) and phase-amplitude coupling (PAC) more specifically have been proposed to play a fundamental role in neural processes ranging from the encoding, maintenance and retrieval of information [Buzsáki, 2010, Jensen and Colgin, 2007, Lisman and Jensen, 2013, Axmacher et al., 2010, Fell and Axmacher, 2011, Kaplan et al., 2014, Hyafil et al., 2015], to large-scale communication across neural ensembles [Canolty and Knight, 2010, Jirsa and Müller, 2013, Khan et al., 2013, Florin and Baillet, 2015]. While a steady increase in observations of PAC in neural data has been seen, how to best detect and quantify such phenomena remains difficult to settle. We argue that a method using DAR models, as described here, is rich enough to capture the time-varying statistics of brain signals in addition to provide efficient inference algorithms. These non-linear statistical models are probabilistic, allowing the estimation of their goodness of fit to the data, and allowing for an easy and fully controlled comparison across models and parameters. In other words, they offer a unique principled data-driven model selection approach, an estimation strategy of phase/amplitude-amplitude coupling based on the approximation of the actual signals, a better temporal resolution of dynamic PAC and the estimation of coupling directionality.

One of the main features of PAC estimation through our method is the ability to compare models or parameters on non-synthetic data. On the contrary, traditional PAC metrics *cannot* be compared on non-synthetic data, and two different choices of parameters can lead to different interpretations. There is no legitimate way to decide which parameter shall be used with empirical data using traditional metrics. The likelihood of the DAR model that can be estimated on left-out data offers a rigorous solution to this problem.

We presented here results on both simulated signals and empirical neurophysiological signals. The simulations gave us an illustration of the phenomenon we want to model, and helped us understand how to visualize a fitted DAR model. They also served a validation purpose for the bandwidth selection approach that we performed on real data. Using the data-driven parameter selection on non-synthetic signals, we

showed how to choose sensible parameters for the filtering of the slow oscillation. All empirical signals are different, and it was for example reported in the neuroscience literature that peak frequencies vary between individuals [Haegens et al., 2014] and that this should not be overlooked in the analysis of the data. The parameter selection based on fitted DAR models makes it possible to fit parameters on individual datasets. Our results also shed light on the asymmetrical and wide-band properties of the slow oscillation, which could denote crucial features involved in cognition [Cole and Voytek, 2017].

The second novelty of our method stands in considering the amplitude fluctuations of the slow oscillation in the PAC measure and not only its phase. Using the rodent and human data, we showed that the instantaneous amplitude of the slow oscillation influences the coupling in PAC, as it was previously suggested in [van Wijk et al., 2015]. The amplitude information should therefore not be discarded as it is done by existing PAC metrics. For instance, the measure of alpha/gamma coupling reported during rest [Osipova et al., 2008, Roux et al., 2013] should incorporate alpha fluctuations when studied in the context of visual tasks [Voytek et al., 2010], as an increase of alpha power is often concomitant with a decrease of gamma power [Fries et al., 2001]. The comparison between DAR models considering or not these low-frequency power fluctuations would inform on the nature of the coupling: purely phase-amplitude, or rather phase/amplitude-amplitude. In Tort *et al.* [Tort et al., 2008], both theta power changes and modulation of theta/gamma PAC were reported in rats having to make a left or right decision to find a reward in a maze. The use of our method could decipher whether the changes in coupling were related to the changes in power, informing on the underlying mechanisms of decision-making. Moreover, as our method models the entire spectrum simultaneously, a phase-frequency coupling could potentially be captured in our models. Therefore, our method is not limited to purely phase-amplitude coupling, and extends the traditional CFC analysis.

Furthermore, in those types of experiments, changes in PAC can be very fast depending on the cognitive state of the subject. Therefore, the need for dynamic PAC estimates is growing [Tort et al., 2008]. We showed with simulations that DAR models are more robust than non-parametric methods when estimating PAC on small time samples. This robustness is critical for time-limited experiments and also when analyzing PAC across time in a fine manner, typically when dynamic processes are at play.

Last but not least, likelihood comparison can also be used to estimate the delay between the coupled components, which would give new insights on highly debated questions on the role of oscillations in neuronal communication [Fries, 2005, Bastos et al., 2015]. For example, a delay close to zero could suggest that the low and high frequency components of the coupling might be generated in the same area, whereas a large delay would suggest they might come from different areas. As an alternative interpretation, the two components may come from the same area, but the coupling mechanism itself might be lagged. In this case, a negative delay would suggest that the low frequency oscillation is driven by the high frequency oscillations, whereas a positive delay would suggest that the low frequency oscillation drives the high frequency amplitude modulation. In any case, this type of analysis will provide valuable information to guide further experimental questions.

A recent concern in PAC analysis is that all PAC metrics may detect a coupling even though the signal is not composed of two cross-frequency coupled oscillators [Kramer et al., 2008, Lozano-Soldevilla et al., 2016, Amiri et al., 2016]. It may happen for instance with sharp slow oscillations, described in humans intracranial recordings [Cole et al., 2016]. Sharp edges are known not to be well described by a Fourier analysis, which decomposes the signal in a linear combination of sinusoids. Indeed, such sharp slow oscillations create artificial high frequency activity at each sharp edge, and these high frequencies are thus artificially coupled with the slow oscillations. This false positive detection is commonly referred to as “spurious” coupling [Jensen et al., 2016]. Even though our method does not use filtering in the high frequencies, it does not solve this issue and is affected in the same way as other traditional PAC metrics. Indeed, our work shed light on the wide-band property of the slow oscillations, but DAR models cannot cope with full-band slow oscillations, which contain strong harmonic components in the high frequencies. However, we consider that such “spurious” PAC can also be a relevant feature of a signal, as stated in [Cole et al., 2016]. In their study, they show that abnormal beta oscillations (13-30 Hz) in the basal ganglia and motor cortex underlie some “spurious” PAC, but are actually a strong feature associated with Parkinson’s disease. A robust way to disentangle the different mechanisms that lead to similar PAC results remains to be developed.

The method we presented in this paper uses univariate signals obtained invasively in rodents or humans.

As a lot of neurophysiological research uses non-invasive MEG or EEG recordings containing multiple channels, a multivariate analysis could be of high interest. One way to use data from multiple channels is to estimate a single signal using a spatial filter such as in [Cohen, 2017]. Such a method is therefore complementary to univariate PAC metrics like ours which can be applied to the output of the spatial filter. The method from [Cohen, 2017] builds spatial filters that maximize the difference between, say, high-frequency activity that appears during peaks of a low-frequency oscillation *versus* high-frequency activity that is unrelated to the low-frequency oscillation. Again, from the signal obtained with the spatial filter, it is straightforward to adapt most PAC metrics such as our method.

Neurophysiological signals have all the statistical properties to make them a challenge from a signal processing perspective. They contain non-linearities, non-stationarities, they are noisy and they can be long, hence posing important computational challenges. Our method based on DAR models offer novel and more robust possibilities to analyse neurophysiological signals, paving the way for new insights on how our brain functions via spectral interactions using local or distant coupling mechanisms.

Our method is fully available as an open source python package that comes with documentation, tests, and examples: <https://pactools.github.io>.

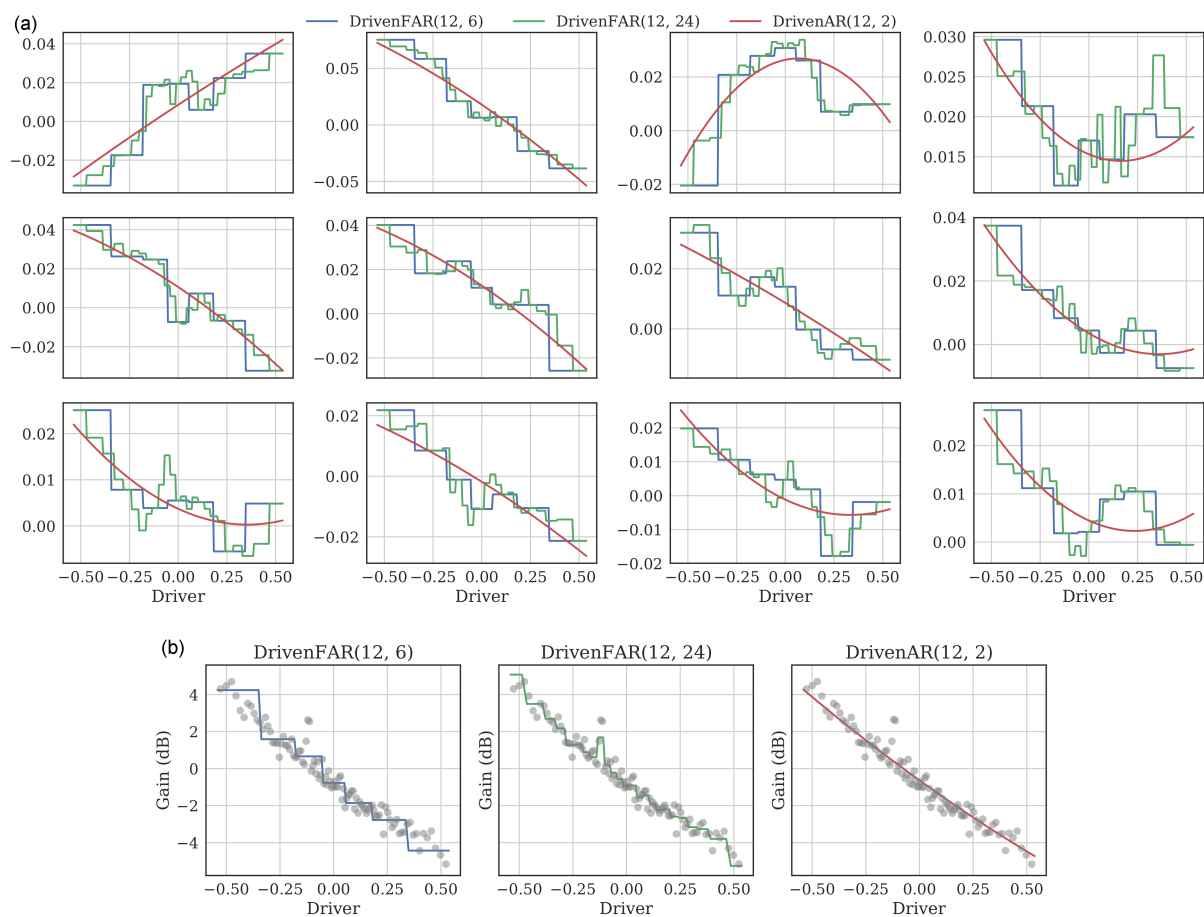
Acknowledgments

We would like to thank Khodagholy *et al.* [Khodagholy et al., 2015] for sharing rodent hippocampal LFP sample data and Canolty *et al.* [Canolty et al., 2006] for sharing human cortical ECoG sample data.

This work was supported by the ERC Starting Grant SLAB ERC-YStG-676943 to Alexandre Gramfort, the ERC Starting Grant MindTime ERC-YStG-263584 to Virginie van Wassenhove, the ANR-16-CE37-0004-04 AutoTime to Valérie Doyère and Virginie van Wassenhove, and the Paris-Saclay IDEX NoTime to Valérie Doyère, Alexandre Gramfort and Virginie van Wassenhove,

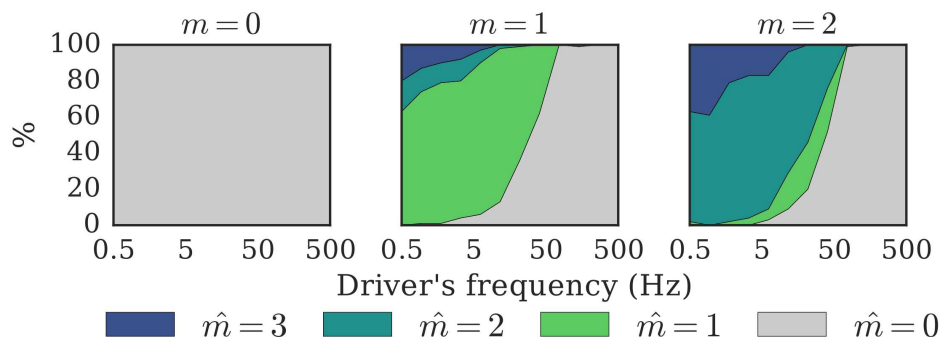
Supporting Information

S1 Fig. Validating the polynomial basis for DAR. (a) AR coefficients $a_i(x)$ as functions of the driver x . (b) Innovation variance $\sigma^2(x)$ as functions of the driver x . The gray circles are the mean squared-residual over 100 bins of the driver values. We compared three parametrizations of these functions: Two staircase functions with respectively $s = 7$ and $s = 25$ steps, and one polynomial function with order $m = 2$, as used in the rest of this work. For a piecewise constant AR parametrization with s -steps, we divided the driver's values into s equally distributed bins, then we fitted s independent linear AR models on the time samples of each bin. This approach is more general than the polynomial approach, yet each AR model uses only T/s samples, whereas the polynomial approach in the DAR model uses all the T samples. The DAR model is therefore more robust. The polynomial parametrization also uses fewer coefficients since a low order m is sufficient. The models are fitted on the human cortical signal, using $p = 12$. We see that an order-2 polynomial is sufficient to approximate the trend of both the AR coefficients and the innovation variance, as estimated by the staircase functions, yet the polynomial uses much fewer parameters.



S1 Fig. Validating the polynomial basis for DAR.

S2 Fig. Model selection with the Bayesian information criterion (BIC). We simulated 100 signals from 100 DAR models with $p = 10$ and $m \in [0, 1, 2]$. We then estimated new DAR models on these signals, with p ranging from 1 to 20, and m ranging from 0 to 3. We selected \hat{p} and \hat{m} that minimized the BIC. The graphs show the proportion of runs that lead to each value of \hat{m} . The hyper-parameter m is correctly estimated in most cases if the driver's frequency is not too high ($f_x < 50$ Hz). The hyper-parameter p is correctly estimated at ± 2 in most cases (not shown).

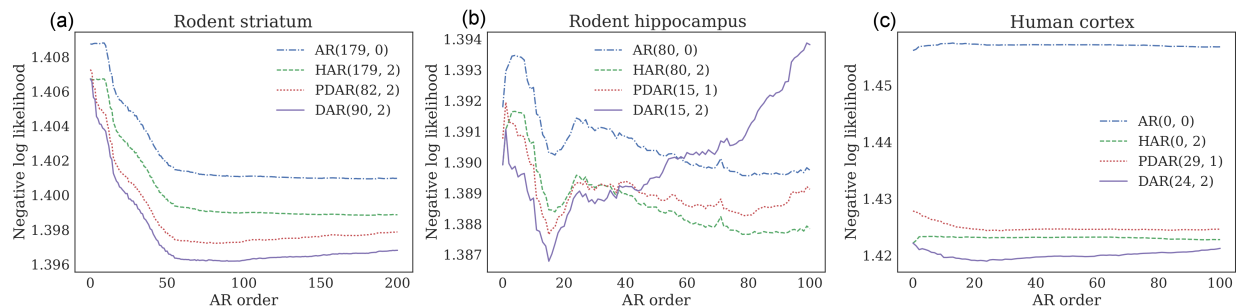


S2 Fig. Model selection with the Bayesian information criterion (BIC).

S3 Fig. Model selection with cross-validation. Cross-validation on our three dataset, (a) rodent striatum, (b) rodent hippocampus, and (c) human auditory cortex, to select the best model and the best parameters. Splitting the signal in half, we fitted the models on the first half, and evaluated the model likelihood on the second half. We compared four different models on a grid of parameter $p \in [0, 100 - 200]$ and $m \in [0, 3]$:

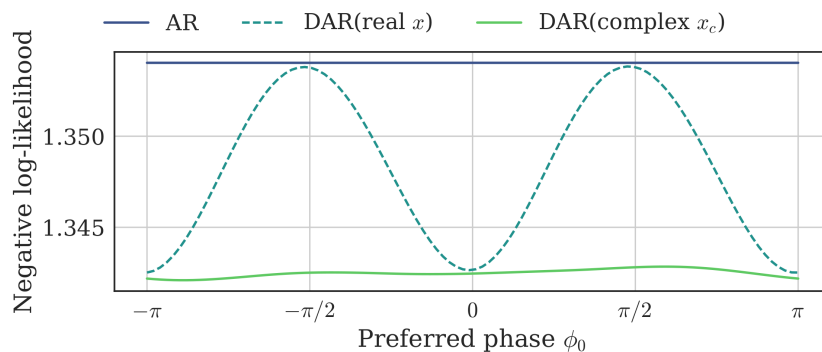
1. AR: a linear AR model
2. Heteroskedastic AR (HAR): an hybrid model between a linear AR model and a DAR model, where the innovation variance σ^2 is driven by x_c , but the AR coefficients are constant in time.
3. Phase DAR (PDAR): a DAR model, with a normalized driver: $x_c/|x_c|$. In this way, we only consider the phase of the slow oscillation, as in most PAC metrics.
4. DAR: a DAR model, where both the innovation variance σ^2 and the AR coefficients are driven by x_c .

The figures present the negative log likelihood (lower is better) by time sample. Each line corresponds to a given model with its best parameter m . The legend shows which orders (p, m) are the best for each model. One can observe that the curves of negative log-likelihood are not convex, yet they exhibit rather clear minima used to define the optimal parameters.



S3 Fig. Model selection with cross-validation.

S4 Fig. Removing the bias with a complex-valued driver. We simulated a signal as described in the Methods section, introducing a phase difference ϕ_0 in the modulation. For each value of ϕ_0 , we fitted three different models, and compared their negative log-likelihood by time sample (the lower the better): an AR, a DAR with a real driver, and a DAR with a complex driver. The parameters were set to $p = 10$ and $m = 1$. A bias is visible around $\phi_0 = \pm\pi/2$, since the real-valued driver DAR model does not fit better than the AR model. As expected, this bias disappears when we update the model to a complex-valued driver DAR model.



S4 Fig. Removing the bias with a complex-valued driver.

References

- Akaike, 1998. Akaike, H. (1998). Information theory and an extension of the maximum likelihood principle. In *Selected Papers of Hirotugu Akaike*, pages 199–213. Springer.
- Amiri et al., 2016. Amiri, M., Lina, J.-M., Pizzo, F., and Gotman, J. (2016). High frequency oscillations and spikes: separating real hfos from false oscillations. *Clinical Neurophysiology*, 127(1):187–196.
- Arlot et al., 2010. Arlot, S., Celisse, A., et al. (2010). A survey of cross-validation procedures for model selection. *Statistics surveys*, 4:40–79.
- Aru et al., 2015. Aru, J., Aru, J., Priesemann, V., Wibral, M., Lana, L., Pipa, G., Singer, W., and Vicente, R. (2015). Untangling cross-frequency coupling in neuroscience. *Current opinion in neurobiology*, 31:51–61.
- Axmacher et al., 2010. 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. *Proceedings of the National Academy of Sciences*, 107(7):3228–3233.
- Axmacher et al., 2006. Axmacher, N., Mormann, F., Fernández, G., Elger, C. E., and Fell, J. (2006). Memory formation by neuronal synchronization. *Brain research reviews*, 52(1):170–182.
- Baccalá and Sameshima, 2001. Baccalá, A. L. and Sameshima, K. (2001). Partial directed coherence: a new concept in neural structure determination. *Biological Cybernetics*, 84(6):463–474.
- Barnett et al., 2009. Barnett, L., Barrett, A. B., and Seth, A. K. (2009). Granger causality and transfer entropy are equivalent for gaussian variables. *Physical review letters*, 103(23):238701.
- Bastos et al., 2015. Bastos, A. M., Vezoli, J., and Fries, P. (2015). Communication through coherence with inter-areal delays. *Current opinion in neurobiology*, 31:173–180.

- Berman et al., 2012. Berman, J. I., McDaniel, J., Liu, S., Cornew, L., Gaetz, W., Roberts, T. P., and Edgar, J. C. (2012). Variable bandwidth filtering for improved sensitivity of cross-frequency coupling metrics. *Brain connectivity*, 2(3):155–163.
- Besserve et al., 2010. Besserve, M., Schölkopf, B., Logothetis, N. K., and Panzeri, S. (2010). Causal relationships between frequency bands of extracellular signals in visual cortex revealed by an information theoretic analysis. *Journal of computational neuroscience*, 29(3):547–566.
- Bigdely-Shamlo et al., 2015. Bigdely-Shamlo, N., Mullen, T., Kothe, C., Su, K.-M., and Robbins, K. A. (2015). The prep pipeline: standardized preprocessing for large-scale eeg analysis. *Frontiers in neuroinformatics*, 9:16.
- Bollerslev, 1986. Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of econometrics*, 31(3):307–327.
- Bonnefond et al., 2017. Bonnefond, M., Kastner, S., and Jensen, O. (2017). Communication between brain areas based on nested oscillations. *eNeuro*, 4(2):ENEURO–0153.
- Bragin et al., 1995. Bragin, A., Jandó, G., Nádasdy, Z., Hetke, J., Wise, K., and Buzsáki, G. (1995). Gamma (40–100 Hz) oscillation in the hippocampus of the behaving rat. *The Journal of Neuroscience*, 15(1):47–60.
- Bruns, 2004. Bruns, A. (2004). Fourier-, hilbert- and wavelet-based signal analysis: are they really different approaches? *Journal of Neuroscience methods*, 137(2):321–332.
- Bruns and Eckhorn, 2004. Bruns, A. and Eckhorn, R. (2004). Task-related coupling from high-to low-frequency signals among visual cortical areas in human subdural recordings. *International Journal of Psychophysiology*, 51(2):97–116.
- Buzsáki, 2006. Buzsáki, G. (2006). *Rhythms of the Brain*. Oxford University Press.
- Buzsáki, 2010. Buzsáki, G. (2010). Neural syntax: cell assemblies, synapsembles, and readers. *Neuron*, 68(3):362–385.
- Canolty et al., 2006. 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(5793):1626–1628.
- Canolty and Knight, 2010. Canolty, R. T. and Knight, R. T. (2010). The functional role of cross-frequency coupling. *Trends in cognitive sciences*, 14(11):506–515.
- Carlstein, 1986. Carlstein, E. (1986). The use of subsample values for estimating the variance of a general statistic from a stationary sequence. *The Annals of Statistics*, pages 1171–1179.
- Chan and Tong, 1986. Chan, K. S. and Tong, H. (1986). On estimating thresholds in autoregressive models. *Journal of Time Series Analysis*, 7(3):179–190.
- Chavez et al., 2006. Chavez, M., Besserve, M., Adam, C., and Martinerie, J. (2006). Towards a proper estimation of phase synchronization from time series. *Journal of Neuroscience methods*, 154(1):149–160.
- Chehelcheraghi et al., 2017. Chehelcheraghi, M., van Leeuwen, C., Steur, E., and Chie, N. (2017). A neural mass model of cross frequency coupling. *PLoS ONE* 12(4): e0173776.
- Cohen, 2017. Cohen, M. X. (2017). Multivariate cross-frequency coupling via generalized eigendecomposition. *eLife*, 6:e21792.
- Cole et al., 2016. Cole, S. R., Peterson, E. J., van der Meij, R., de Hemptinne, C., Starr, P. A., and Voytek, B. (2016). Nonsinusoidal oscillations underlie pathological phase-amplitude coupling in the motor cortex in parkinson’s disease. *bioRxiv*, page 049304.

- Cole and Voytek, 2017. Cole, S. R. and Voytek, B. (2017). Brain oscillations and the importance of waveform shape. *Trends in Cognitive Sciences*.
- Colgin et al., 2009. Colgin, L. L., Denninger, T., Fyhn, M., Hafting, T., Bonnevie, T., Jensen, O., Moser, M.-B., and Moser, E. I. (2009). Frequency of gamma oscillations routes flow of information in the hippocampus. *Nature*, 462(7271):353–357.
- Dahlhaus, 1996. Dahlhaus, R. (1996). On the kullback-leibler information divergence of locally stationary processes. *Stochastic Processes and their Applications*, 62(1):139–168.
- Dallérac et al., 2017. Dallérac, G., Graupner, M., Knippenberg, J., Martinez, R. C. R., Tavares, T. F., Tallot, L., El Massioui, N., Verschuere, A., Höhn, S., Bertolus, J. B., et al. (2017). Updating temporal expectancy of an aversive event engages striatal plasticity under amygdala control. *Nature Communications*, 8:13920.
- Dupré la Tour et al., 2017. Dupré la Tour, T., Grenier, Y., and Gramfort, A. (2017). Parametric estimation of spectrum driven by an exogenous signal. In *Acoustics, Speech and Signal Processing (ICASSP), 2017 IEEE International Conference on*, pages 4301–4305. IEEE.
- Dvorak and Fenton, 2014. Dvorak, D. and Fenton, A. A. (2014). Toward a proper estimation of phase–amplitude coupling in neural oscillations. *Journal of Neuroscience methods*, 225:42–56.
- Engle, 1982. Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica: Journal of the Econometric Society*, pages 987–1007.
- Fell and Axmacher, 2011. Fell, J. and Axmacher, N. (2011). The role of phase synchronization in memory processes. *Nature reviews. Neuroscience*, 12(2):105.
- Florin and Baillet, 2015. Florin, E. and Baillet, S. (2015). The brain’s resting-state activity is shaped by synchronized cross-frequency coupling of neural oscillations. *NeuroImage*, 111:26–35.
- Fries, 2005. Fries, P. (2005). A mechanism for cognitive dynamics: neuronal communication through neuronal coherence. *Trends in cognitive sciences*, 9(10):474–480.
- Fries, 2015. Fries, P. (2015). Rhythms for cognition: communication through coherence. *Neuron*, 88(1):220–235.
- Fries et al., 2001. Fries, P., Reynolds, J. H., Rorie, A. E., and Desimone, R. (2001). Modulation of oscillatory neuronal synchronization by selective visual attention. *Science*, 291(5508):1560–1563.
- Granger, 1988. Granger, C. W. (1988). Some recent development in a concept of causality. *Journal of econometrics*, 39(1):199–211.
- Grenier, 1983. Grenier, Y. (1983). Time-dependent arma modeling of nonstationary signals. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, 31(4):899–911.
- Grenier, 2013. Grenier, Y. (2013). Estimating an ar model with exogenous driver. Technical Report 2013D007, Telecom ParisTech.
- Grenier and Omnes-Chevalier, 1988. Grenier, Y. and Omnes-Chevalier, M.-C. (1988). Autoregressive models with time-dependent log area ratios. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, 36(10):1602–1612.
- Haegens et al., 2014. Haegens, S., Cousijn, H., Wallis, G., Harrison, P. J., and Nobre, A. C. (2014). Inter- and intra-individual variability in alpha peak frequency. *Neuroimage*, 92:46–55.
- Haggan and Ozaki, 1981. Haggan, V. and Ozaki, T. (1981). Modelling nonlinear random vibrations using an amplitude-dependent autoregressive time series model. *Biometrika*, 68(1):189–196.

- Haufe et al., 2010. Haufe, S., Tomioka, R., Nolte, G., Müller, K.-R., and Kawanabe, M. (2010). Modeling sparse connectivity between underlying brain sources for EEG/MEG. *Bio. Eng., IEEE Trans.*, 57(8):1954–1963.
- Heusser et al., 2016. Heusser, A. C., Poeppel, D., Ezzyat, Y., and Davachi, L. (2016). Episodic sequence memory is supported by a theta- gamma phase code. *Nature neuroscience*.
- Huth et al., 2016. Huth, A. G., de Heer, W. A., Griffiths, T. L., Theunissen, F. E., and Gallant, J. L. (2016). Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature*, 532(7600):453–458.
- Hyafil, 2015. Hyafil, A. (2015). Misidentifications of specific forms of cross- frequency coupling: three warnings. *Frontiers in Neuroscience*, 9.
- Hyafil et al., 2015. Hyafil, A., Giraud, A.-L., Fontolan, L., and Gutkin, B. (2015). Neural cross-frequency coupling: Connecting architectures, mechanisms, and functions. *Trends in Neurosciences*, 38(11):725–740.
- Jensen and Colgin, 2007. Jensen, O. and Colgin, L. L. (2007). Cross-frequency coupling between neuronal oscillations. *Trends in cognitive sciences*, 11(7):267–269.
- Jensen et al., 2016. Jensen, O., Spaak, E., and Park, H. (2016). Discriminating valid from spurious indices of phase- amplitude coupling. *eneuro*, pages ENEURO–0334.
- Jiang et al., 2015. Jiang, H., Bahramisharif, A., van Gerven, M. A., and Jensen, O. (2015). Measuring directionality between neuronal oscillations of different frequencies. *Neuroimage*, 118:359–367.
- Jirsa and Müller, 2013. Jirsa, V. and Müller, V. (2013). Cross-frequency coupling in real and virtual brain networks. *Frontiers in computational neuroscience*, 7:78.
- Kaplan et al., 2014. Kaplan, R., Bush, D., Bonnefond, M., Bandettini, P. A., Barnes, G. R., Doeller, C. F., and Burgess, N. (2014). Medial prefrontal theta phase coupling during spatial memory retrieval. *Hippocampus*, 24(6):656–665.
- Kay et al., 2008. Kay, K. N., Naselaris, T., Prenger, R. J., and Gallant, J. L. (2008). Identifying natural images from human brain activity. *Nature*, 452(7185):352–355.
- Khan et al., 2013. Khan, S., Gramfort, A., Shetty, N., M., K., Ganesan, S., Moran, J., Lee, S., Gabrieli, J., Tager-Flusberg, H., Joseph, R., Herbert, M., Hämäläinen, M., and Kenet, T. (2013). Local and long-range functional connectivity is reduced in concert in autism spectrum disorders. *Proc. Natl. Acad. Sci.*
- Khodagholy et al., 2015. Khodagholy, D., Gelinias, J. N., Thesen, T., Doyle, W., Devinsky, O., Malliaras, G. G., and Buzsáki, G. (2015). Neurogrid: recording action potentials from the surface of the brain. *Nature neuroscience*, 18(2):310–315.
- Kikuchi et al., 2017. Kikuchi, Y., Attaheri, A., Wilson, B., Rhone, A. E., Nourski, K. V., Gander, P. E., Kovach, C. K., Kawasaki, H., Griffiths, T. D., Howard III, M. A., et al. (2017). Sequence learning modulates neural responses and oscillatory coupling in human and monkey auditory cortex. *PLoS biology*, 15(4):e2000219.
- Kramer et al., 2008. Kramer, M. A., Tort, A. B., and Kopell, N. J. (2008). Sharp edge artifacts and spurious coupling in EEG frequency comodulation measures. *Journal of Neuroscience methods*, 170(2):352–357.
- Lachaux et al., 1999. Lachaux, J.-P., Rodriguez, E., Martinerie, J., Varela, F. J., et al. (1999). Measuring phase synchrony in brain signals. *Human brain mapping*, 8(4):194–208.

- Lakatos et al., 2005. Lakatos, P., Shah, A. S., Knuth, K. H., Ulbert, I., Karmos, G., and Schroeder, C. E. (2005). An oscillatory hierarchy controlling neuronal excitability and stimulus processing in the auditory cortex. *Journal of neurophysiology*, 94(3):1904–1911.
- Lisman and Jensen, 2013. Lisman, J. E. and Jensen, O. (2013). The theta-gamma neural code. *Neuron*, 77(6):1002–1016.
- Lozano-Soldevilla et al., 2016. Lozano-Soldevilla, D., ter Huurne, N., and Oostenveld, R. (2016). Neuronal oscillations with non-sinusoidal morphology produce spurious phase-to-amplitude coupling and directionality. *Frontiers in Computational Neuroscience*, 10.
- Mahan et al., 2015. Mahan, M. Y., Chorn, C. R., and Georgopoulos, A. P. (2015). White noise test: detecting autocorrelation and nonstationarities in long time series after arima modeling. In *Proceedings 14th Python in Science Conference (Scipy 2015), Austin, TX*.
- Malerba and Kopell, 2013. Malerba, P. and Kopell, N. (2013). Phase resetting reduces theta–gamma rhythmic interaction to a one-dimensional map. *Journal of mathematical biology*, pages 1–26.
- Naselaris et al., 2011. Naselaris, T., Kay, K. N., Nishimoto, S., and Gallant, J. L. (2011). Encoding and decoding in fmri. *Neuroimage*, 56(2):400–410.
- Nolte et al., 2010. Nolte, G., Ziehe, A., Krämer, N., Popescu, F., and Müller, K.-R. (2010). Comparison of granger causality and phase slope index. In *NIPS Causality: Objectives and Assessment*, pages 267–276.
- Nolte et al., 2008. Nolte, G., Ziehe, A., Nikulin, V., Schlögl, A., Krämer, N., Brismar, T., and Müller, K.-R. (2008). Robustly estimating the flow direction of information in complex physical systems. *Physical Review Letters*, 100:234101.
- Osipova et al., 2008. Osipova, D., Hermes, D., and Jensen, O. (2008). Gamma power is phase-locked to posterior alpha activity. *PloS one*, 3(12):e3990.
- Özkurt and Schnitzler, 2011. Özkurt, T. E. and Schnitzler, A. (2011). A critical note on the definition of phase– amplitude cross-frequency coupling. *Journal of Neuroscience methods*, 201(2):438–443.
- Park et al., 2015. Park, H., Ince, R. A., Schyns, P. G., Thut, G., and Gross, J. (2015). Frontal top-down signals increase coupling of auditory low-frequency oscillations to continuous speech in human listeners. *Current Biology*, 25(12):1649–1653.
- Penny et al., 2008. Penny, W., Duzel, E., Miller, K., and Ojemann, J. (2008). Testing for nested oscillation. *Journal of Neuroscience methods*, 174(1):50–61.
- Proakis and Manolakis, 1996. Proakis, J. G. and Manolakis, D. G. (1996). *Digital signal processing 3rd edition*. Prentice Hall.
- Quiroga et al., 2002. Quiroga, R. Q., Kraskov, A., Kreuz, T., and Grassberger, P. (2002). Performance of different synchronization measures in real data: a case study on eeg signals. *Physical Review E*, 65(4):041903.
- Roux et al., 2013. Roux, F., Wibral, M., Singer, W., Aru, J., and Uhlhaas, P. J. (2013). The phase of thalamic alpha activity modulates cortical gamma-band activity: evidence from resting- state meg recordings. *Journal of Neuroscience*, 33(45):17827–17835.
- Schreiber, 2000. Schreiber, T. (2000). Measuring information transfer. *Physical review letters*, 85(2):461.
- Schwarz et al., 1978. Schwarz, G. et al. (1978). Estimating the dimension of a model. *Ann. Stat.*, 6(2):461–464.

- Shirvalkar et al., 2010. Shirvalkar, P. R., Rapp, P. R., and Shapiro, M. L. (2010). Bidirectional changes to hippocampal theta–gamma comodulation predict memory for recent spatial episodes. *Proceedings of the National Academy of Sciences*, 107(15):7054–7059.
- Spyers-Ashby et al., 1998. Spyers-Ashby, J., Bain, P., and Roberts, S. (1998). A comparison of fast fourier transform (fft) and autoregressive (ar) spectral estimation techniques for the analysis of tremor data. *Journal of neuroscience methods*, 83(1):35–43.
- Sweeney-Reed et al., 2014. Sweeney-Reed, C. M., Zaehle, T., Voges, J., Schmitt, F. C., Buentjen, L., Kopitzki, K., Esslinger, C., Hinrichs, H., Heinze, H.-J., Knight, R. T., et al. (2014). Corticothalamic phase synchrony and cross-frequency coupling predict human memory formation. *Elife*, 3:e05352.
- Tong and Lim, 1980. Tong, H. and Lim, K. S. (1980). Threshold autoregression, limit cycles and cyclical data. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 245–292.
- Tort et al., 2010. Tort, A. B., Komorowski, R., Eichenbaum, H., and Kopell, N. (2010). Measuring phase-amplitude coupling between neuronal oscillations of different frequencies. *J. Neurophysiol.*, 104(2):1195–1210.
- Tort et al., 2008. Tort, A. B., Kramer, M. A., Thorn, C., Gibson, D. J., Kubota, Y., Graybiel, A. M., and Kopell, N. J. (2008). Dynamic cross-frequency couplings of local field potential oscillations in rat striatum and hippocampus during performance of a t-maze task. *Proc. Natl. Acad. Sci.*, 105(51):20517–20522.
- Tort et al., 2007. Tort, A. B., Rotstein, H. G., Dugladze, T., Gloveli, T., and Kopell, N. J. (2007). On the formation of gamma-coherent cell assemblies by oriens lacunosum-moleculare interneurons in the hippocampus. *Proceedings of the National Academy of Sciences*, 104(33):13490–13495.
- Valdés-Sosa et al., 2005. Valdés-Sosa, P. A., Sánchez-Bornot, J. M., Lage-Castellanos, A., Vega-Hernández, M., Bosch-Bayard, J., Melie-García, L., and Canales-Rodríguez, E. (2005). Estimating brain functional connectivity with sparse multivariate autoregression. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 360(1457):969–981.
- van Wijk et al., 2015. van Wijk, B., Jha, A., Penny, W., and Litvak, V. (2015). Parametric estimation of cross-frequency coupling. *Journal of neuroscience methods*, 243:94–102.
- Voytek et al., 2010. Voytek, B., Canolty, R. T., Shestyuk, A., Crone, N., Parvizi, J., and Knight, R. T. (2010). Shifts in gamma phase–amplitude coupling frequency from theta to alpha over posterior cortex during visual tasks. *Frontiers in human neuroscience*, 4:191.
- Wibral et al., 2013. Wibral, M., Pampu, N., Priesemann, V., Siebenhühner, F., Seiwert, H., Lindner, M., Lizier, J. T., and Vicente, R. (2013). Measuring information-transfer delays. *PloS one*, 8(2):e55809.