

Reconstructing Time-Dependent Dynamics

Philip Clemson, Gemma Lancaster and Aneta Stefanovska

Abstract—The usefulness of the information contained in biomedical data relies heavily on the reliability and accuracy of the methods used for its extraction. The conventional assumptions of stationarity and autonomicity break down in the case of living systems because they are thermodynamically open, and thus constantly interacting with their environments. This leads to an inherent time-variability and results in highly nonlinear, time-dependent dynamics. The aim of signal analysis usually is to gain insight into the behaviour of the system from which the signal originated. Here, a range of signal analysis methods is presented and applied to extract information about time-varying oscillatory modes and their interactions. Methods are discussed for the characterization of signals and their underlying non-autonomous dynamics, including time-frequency analysis, decomposition, coherence analysis and dynamical Bayesian inference to study interactions and coupling functions. They are illustrated by being applied to cardiovascular and EEG data. The recent introduction of chronotaxic systems provides a theoretical framework within which dynamical systems can have amplitudes and frequencies which are time-varying, yet remain stable, matching well the characteristics of life. We demonstrate that, when applied in the context of chronotaxic systems, the methods presented facilitate the accurate extraction of the system dynamics over many scales of time and space.

Index Terms—Biomedical signal analysis, time-frequency analysis, dynamical Bayesian inference, time-dependent dynamics, wavelet bispectrum, coupling function, phase coherence, cardiovascular system, brain dynamics.

I. INTRODUCTION

Continuous technological advances allow the acquisition of an increasing number of biomedical signal types. These signals may arise from simple mechanical effects, such as the movement of the thorax during breathing, electrical effects, such as the synchronization of firing neurons in the brain, as measured during an electroencephalogram (EEG), optical effects, as utilised in near infra-red spectroscopy (NIRS) and laser Doppler flowmetry (LDF), or from any other measurable biological process. Improvements in the temporal resolution of these techniques allows accurate recording of the time-dependent dynamics inherent to all biomedical signals. Although systems on a microscopic scale may initially appear to be very complicated, there are cases when simple macroscopic behaviour may still arise from these systems [1], [2]. Decomposition of these macroscopic effects recorded by experimental signals can now be considered within already well established theoretical frameworks, based on dynamical systems which are nonlinear, non-autonomous, and far from equilibrium, as

has repeatedly been shown to be the case in living systems. Information obtained from the analysis of these signals has led to a greater understanding of fundamental physiology and is also contributing to important advances in medicine. More specifically, biological oscillations exhibiting a wide range of characteristic frequencies have been observed in biomedical data [3], spanning from very high frequencies, e.g. in EEG data [4], to very low frequencies, e.g. in cerebral hemodynamics [5], [6], microvascular blood flow [7], intracellular calcium levels [8], and individual mitochondria [9].

Although observed in many living systems [10], the importance of biological oscillators, and their interactions, is often overlooked, despite the fact that the extraction of their dynamics at different time scales could bring new insights and understanding of the function of living systems. In fact, these oscillations have been shown to be of great importance in many systems, such as cellular signalling [11], [12], cellular energy metabolism [13] and neural networks [14].

The coupled nonlinear oscillators approach is marked by two major milestones: the introduction of the entrainment of collective oscillators by Winfree [15] and its analysis using the phase dynamics approach of Kuramoto [16], [17]. The identification of the underlying mechanisms of some of these oscillations allows their use in the characterization of different physiological states. Observing changes in these oscillations and their interactions then yields valuable information about the underlying system, for example during epileptic seizures [18], or in skin microvascular blood flow, where changes in oscillatory behaviour have been demonstrated in pathological states with impaired microvasculature, such as hypertension [19], diabetes [20] and skin melanoma [21]. Not only do these changes in the dynamics of biological oscillators provide important physiological insights, they can also be directly utilised in medicine. Potential applications include the identification of the depth of anaesthesia [22], [23], monitoring of intracranial pressure [24], and detection of impaired cerebrovascular reactivity after acute traumatic brain injury [25].

Frequency analysis of biomedical signals first became feasible with the widespread availability of computational resources combined with fast Fourier transforms (FFTs). In this way, biological oscillations were observed on many scales. However, whilst very useful in the first instance, assumptions made in the application of frequency domain analysis techniques are often not applicable when considering signals arising from living systems. The Fourier transform assumes stationarity of the signal, i.e. that the frequency content does not change over time. In living systems this is never the case due to the openness of the system, resulting in time-varying amplitudes and frequencies of oscillations. The introduction of time-frequency analysis techniques provided an optimal solution to this problem, allowing full characterisation of the underlying

This work was supported by the Engineering and Physical Sciences Research Council (UK) (Grant No. EP/100999X1), the EU-NEST BRACCIA Project (No. 517133), the Action Medical Research (UK) MASDA project (1963), and by the Slovenian Research Agency (Program No. P20232). The authors are with the Department of Physics, Lancaster University, Lancaster, LA1 4YB, UK. (e-mail: aneta@lancaster.ac.uk).

dynamics of an oscillatory biological system in time, with no prior assumptions. Now established as almost mandatory in biomedical signal analysis [26], time-frequency analysis methods are continually being developed for the investigation of biological systems in terms of their oscillatory components and the nature of their interactions.

Analysis based on time-frequency methods is often like peeling an onion. Fig. 1 shows how different classes of methods can be used at each level of analysis to bring new information. The initial time-frequency representation of a signal is ideal for decoding the complexity caused by combinations of nonstationary oscillations at different frequencies. These can be extracted by decomposition, which form the next class of methods. After the individual components of the signal have been separated they can be characterised by another set of methods, giving general information about the underlying system (e.g. the frequency range of the components, whether the components exchange information or are coherent). Another set of methods allow the direct physical interpretation of a modelled system from the data observed, such as whether its oscillations and interactions are stable or unstable.

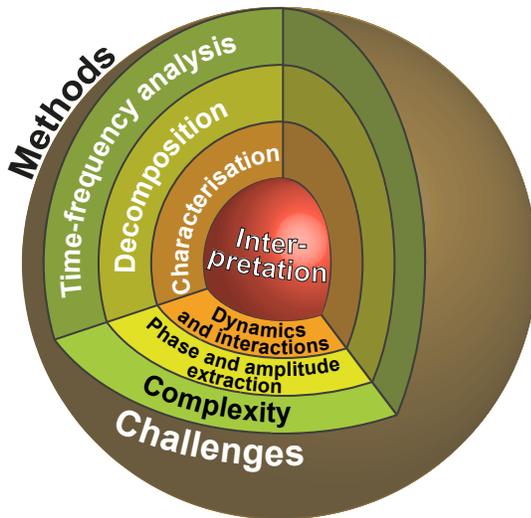


Fig. 1. An illustration showing the challenges related to each level of analysis and the corresponding methods used to tackle them. For biomedical time series, the challenge of the signal's complexity must first be overcome with time-frequency analysis. This allows the identification and extraction of the phase and amplitude of individual oscillatory components using decomposition methods. Information from these modes can be used to characterise the dynamics of the modes and detect how they interact with each other. Finally, the properties of an explicit physical model of the dynamics provide an interpretation of the system that generated the signal.

Depending on the nature of the signal, not all of these levels of analysis may be available. For example, if the system has dominant stochastic properties, with a homogeneous amplitude distribution in the time-frequency representation, then it cannot be decomposed into separate oscillatory components, excluding most of the methods that would be used for the characterisation or interpretation as well. Similarly, if the system is strongly deterministic and its components could be extracted from a signal, but the signal was perhaps too short

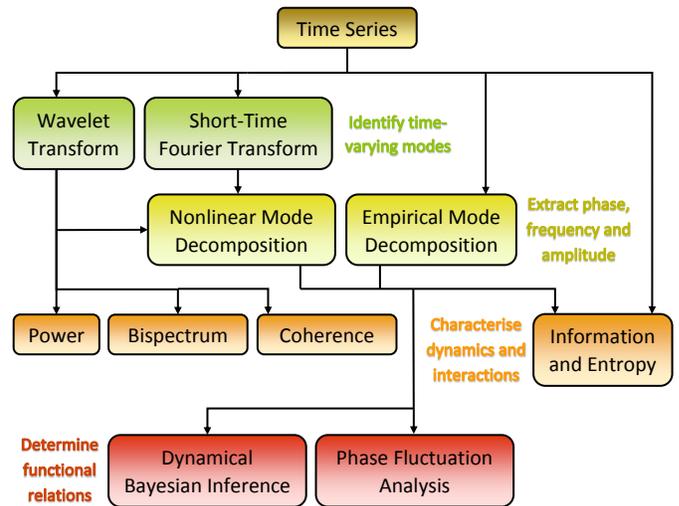


Fig. 2. A workflow chart of the methods discussed in this paper, organised into the levels of analysis illustrated in Fig. 1. Some methods are able to extract information directly from the original time series data. However, the methods further down the chart which are used to characterise or interpret the underlying system often depend on the outcome of time-frequency analysis and decomposition.

relative to the inherent time scales, then there is no way any meaningful interpretation can be made either.

Fig. 2 provides an overview of the methods discussed in this paper within the context of this framework. It can be seen that while some methods can be applied to the time series directly, others rely on either one or two additional steps. However, even for those that do not depend on prior analysis, there still exists a great deal of overlap between the information provided by each method. For example, if the initial time-frequency analysis shows only random noise fluctuations then the corresponding power spectrum should have a smooth continuous distribution. Similarly, the characterisation of the dynamics and interactions should match the information from the methods used to determine the actual functional relations (such as the direction of coupling between two modes).

II. RELATION TO DYNAMICAL SYSTEMS THEORY

Living systems require special treatment when their dynamics is analysed. This comes from the features discussed in the following sections.

A. Nonlinearity

If one is given all of the information about a linear system at one point in time it is trivial to find its state at any other point in time, thus making analytical solutions possible. Essentially, this means that in a linear model it is possible to understand all of the properties of the system without actually observing its dynamics. In contrast, *nonlinear* systems such as living systems cannot be analysed in this way. While some of their properties can still be derived analytically, it is not possible to know what trajectory a nonlinear system will take without the aid of computer simulations or by observing the dynamics of real systems. It is also worth noting that this is not a

question of the complexity of the system; a linear system can be incredibly complex or a nonlinear system can be very simple, but these fundamental rules still apply to their analysis.

The effects of nonlinearity can be quite profound. Not only does it cause mathematical headaches, but it results in phenomena such as *hysteresis*. This describes the effect where the trajectory that a system takes from one state to another is different from the trajectory it takes in the reverse direction between the same two states, making the arrow of time important in the analysis of nonlinear systems. Nonlinearity also causes the effect of *harmonics* which are modes that can be detected when nonlinear oscillations are analysed using methods based on linear systems.

B. Openness

The properties described above can in fact be thought of as manifestations of a single feature of living systems: the fact that they are open and exchange energy and matter with their environment. The main theories of dynamical systems typically assume that the system is closed, meaning that it is autonomous and completely described by its state in space. In contrast, living systems are open and *non*-autonomous, which means they are described by both their state in space *and* time. Consequently, the inclusion of time-dependent variables is vitally important in the analysis of living systems.

Note that the statistical properties of closed systems can still vary with time and the dynamics of such systems is said to be *nonstationary*. Complex nonstationary dynamics in closed systems is usually modelled by chaotic behaviour, where small perturbations in a system's trajectory grow exponentially over time [27]. Complex dynamics that do not conform to the chaotic approach are often instead modelled by stochastic systems, where the nonstationarity arises from the influence of external random variables. Both of these approaches still fit into the framework of autonomous systems and as such time-dependent variables are traditionally not included in the analysis. However, due to the fact that living systems are not only nonstationary but non-autonomous, neither of these approaches can be applied [28].

C. Chronotaxic systems

Recently, a new class of systems has been developed which more closely captures the properties of living systems. The new model follows from the theory of self-sustained limit cycle oscillators, which have been used to describe oscillations with stable amplitude dynamics [29]. Named *chronotaxic systems*, they now add to this theory by combining it with the theory of non-autonomous systems [30] to provide a mechanism for stable oscillations with time-varying frequencies [31], [32]. A new framework of analysis has since been developed to detect such chronotaxic behaviour [33], [34].

To demonstrate the analytical framework discussed in this paper and its application to signals from living systems the following time series from a chronotaxic phase oscillator system are used,

$$x(t) = \cos(\alpha_{x,1}t) + \cos(\alpha_{x,2}t) + \cos(\alpha_{x,3}t) + \eta_1(t), \quad (1)$$

where $\alpha_{x,i}$ are the phases of three chronotaxic modes and $\eta_1(t)$ is a 1/f noise signal. In addition, a second time series $p(t)$ containing the external modes which drive the x modes is defined as,

$$p(t) = \cos(\alpha_{p,1}t) + \cos(\alpha_{p,2}t) + \cos(\alpha_{p,3}t) + \eta_2(t), \quad (2)$$

where $\alpha_{p,i}$ are the phases of the modes driving the system and $\eta_2(t)$ is a separate 1/f noise signal. The chronotaxic modes were generated using the equations,

$$\begin{cases} \dot{\alpha}_{\mathbf{p}} = \omega_0(t) , \\ \dot{\alpha}_{\mathbf{x}} = \varepsilon\omega_0(t) \sin(\alpha_{\mathbf{x}} - \alpha_{\mathbf{p}}) + \xi(t) , \end{cases} \quad (3)$$

where ε is the coupling strength from the external variable \mathbf{p} to the observed variable \mathbf{x} . The function $\xi(t)$ is white Gaussian noise with standard deviation $\eta = \sqrt{2E}$, where $\langle \xi(t) \rangle = 0$, $\langle \xi(t)\xi(\tau) \rangle = \delta(t - \tau)E$. The frequency of $\alpha_{\mathbf{p}}$ is given as

$$\omega_0(t) = \omega_1 [1 + A \sin(\omega_2 t)]. \quad (4)$$

The modes were given the parameters [$\omega_1 = 2\pi$, $\omega_2 = 0.016\pi$, $A = 1$] for mode 1, [$\omega_1 = 0.3\pi$, $\omega_2 = 0.005\pi$, $A = 0.15$] for mode 2 and [$\omega_1 = 0.05\pi$, $\omega_2 = 0.001\pi$, $A = 0.025$] for mode 3. In each case the mode was made chronotaxic by setting the coupling strength $|\varepsilon| = 1.5$.

III. TIME-FREQUENCY ANALYSIS

Complexity is the first challenge that is encountered when dealing with biomedical signals. Not only do these signals often comprise of a mixture of oscillations at different frequencies, but these oscillations each have their own time-dependent dynamics. In the time domain, these oscillations cannot be easily separated. We must therefore transform the signal to the *time-frequency* domain.

A. Continuous transforms

The Short-time Fourier transform (STFT) (also known as the windowed Fourier transform), was developed as a solution to the shortcomings of the Fourier transform when dealing with nonstationary signals [35]. While the Fourier transform provides a representation of a signal in the frequency domain, the STFT transforms the data to the *time-frequency* domain.

The STFT is computed by calculating the Fourier transform of a sliding window which moves over the signal. The Fourier spectrum of the window is assigned to the central point. This process is defined mathematically as

$$STFT(\tilde{\omega}, t) = \int_{-L/2}^{L/2} g(u - t)f(u)e^{-\frac{2\pi i \tilde{\omega} u}{t}} du, \quad (5)$$

where $f(u)$ is a signal of length L and $g(u)$ is a rectangular function of length l that is zero outside u , $-l/2 \leq u \leq l/2$. The variable $\tilde{\omega}$ is directly related to the frequency ω by $\omega = \frac{\tilde{\omega}}{l\Delta t}$, while t is the time.

One immediate problem with this form of the STFT is time-localisation. By using a rectangular window function this means that data at the edges of the window carries the same weight as the data at the centre. Ideally only the data at the centre of the window would contribute to the Fourier

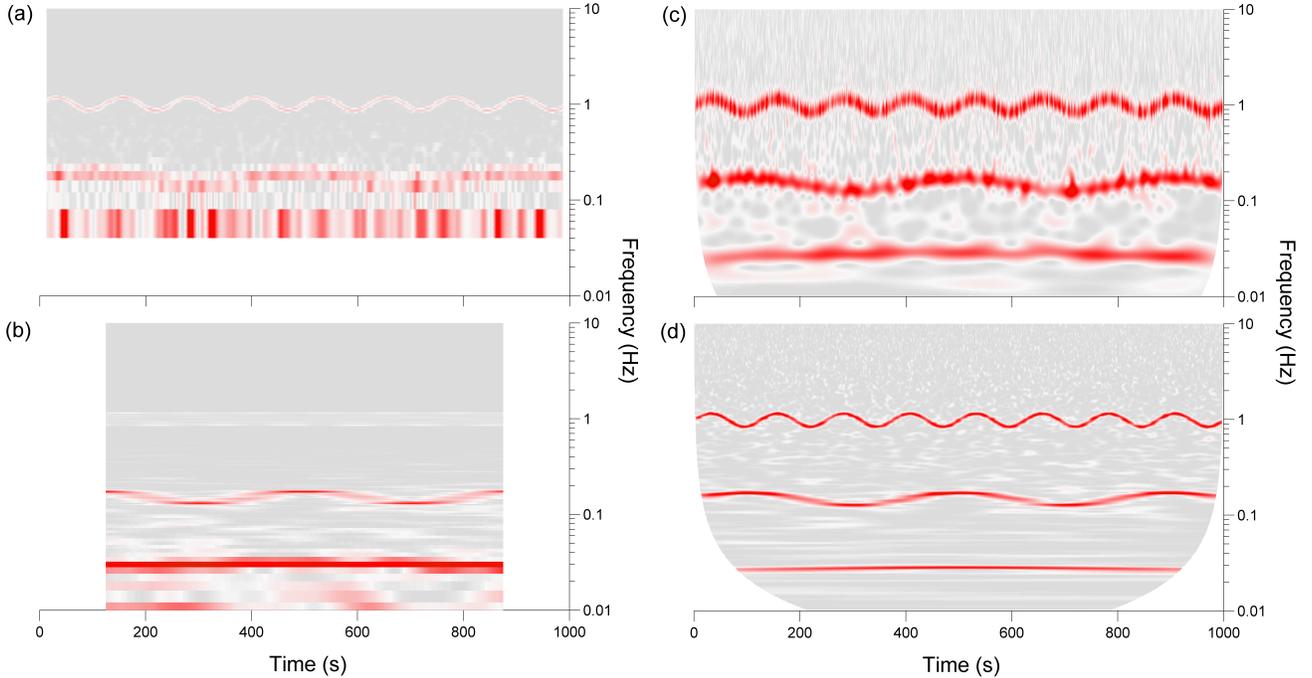


Fig. 3. Time-frequency analysis of the time series (1). STFTs of this time series are shown in (a) and (b) for a 25 s and 250 s window respectively. Continuous Morlet wavelet transforms of the same time series are shown in (c) and (d) where the central frequencies $f_0 = 1$ and $f_0 = 5$ were used respectively. White spaces indicate the limit of the cone of influence where the transform is not defined.

transform. However, in Fourier transforms the frequency resolution is proportional to the length of the data. Therefore, reducing the window size to improve time-localisation also reduces the frequency resolution and makes it more difficult to separate oscillations of different frequencies in the time-frequency domain. This limitation comes from the *uncertainty principle*: one cannot determine the exact frequency of an oscillation at an exact time. The window size also determines the lowest possible frequency that can be detected, so that the amplitude of oscillations that have frequencies below this value are merged together into the same Fourier coefficient at $\omega = 0$. A quick fix for this problem is to use a Gaussian window, which provides optimal time resolution [35]. However, this still only provides the optimal resolution for the lowest observable frequencies. Higher frequencies can still be observed in smaller windows and the frequency resolution relative to the frequency of these oscillations is much better. Other solutions therefore tried to make an adaptive transform that took the frequency of the oscillation into account [36]. One such idea was to use windows of different sizes to compute each frequency in the Fourier spectrum, resulting in the *wavelet transform* [37], [38].

The wavelets that form the basis of the wavelet transform are distinct from the Fourier transform in that they are defined in time as well as frequency. The continuous wavelet transform (CWT) given by

$$W_T(s, t) = \int_{-L/2}^{L/2} \Psi(s, u - t) f(u) du, \quad (6)$$

where $\Psi(s, t)$ is the mother wavelet, which defines all wavelets

by being scaled according to the scale s to change its frequency distribution and time-shifted according to t . Instead of computing a “stand-alone transform” for each time window, the wavelet transform performs a different calculation depending on both time *and* frequency (or more specifically, s). This makes it possible to define an adaptive window size that is small for high frequencies and large for low frequencies. The time resolution at high frequencies is therefore no longer limited by the condition of needing a large window to detect low frequencies.

The Morlet wavelet provides a basis which is closest to the Fourier basis and is defined as [39],

$$\Psi(s, t) = \frac{1}{\sqrt[4]{\pi}} \left(e^{\frac{2\pi i \omega_c t}{s}} - e^{-\frac{2\pi \omega_c^2}{2}} \right) e^{-\frac{t^2}{2s^2}}, \quad (7)$$

where $s = 1/\omega$. The parameter ω_c is the central frequency, which determines the time-frequency resolution of the wavelet; high values ($\omega_c > 2$) give good frequency but poor time resolution while low values ($\omega_c < 1$) give good time but poor frequency resolution. At very small values ($\omega_c < 0.2$) the wavelet transform becomes meaningless as the wavelets become smooth functions with no defined cycles, while at very high values relative to the length of the time series the wavelet transform has a distribution similar to the Fourier transform. A more in depth review on the technical aspects of the wavelet transform can be found in Ref. [40].

Fig. 3 shows a comparison of the STFT with the CWT for the time series (1). Neither transform is defined for all times and frequencies due to either not being able to observe a full cycle in a given window (resulting in the white space

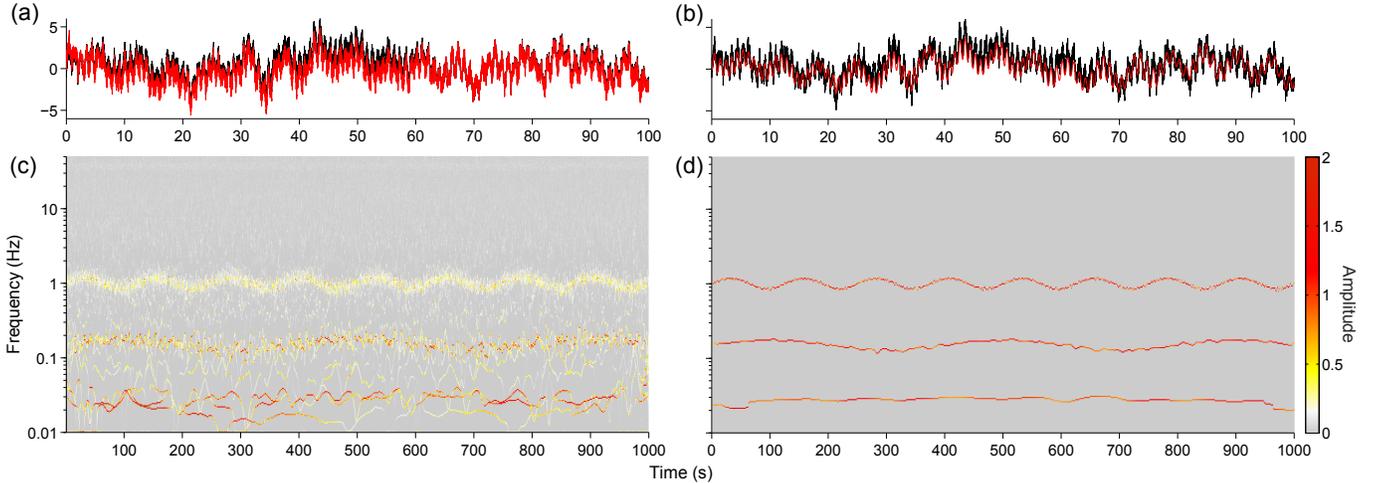


Fig. 4. Decomposition of the time series defined in (1). The signal reconstructed from the first 19 IMFs from EEMD is shown by the red line in (a), while the reconstruction from NMD is shown by the red line in (b). In both plots the black line corresponds to the original signal. Plots (c) and (d) show the amplitude of the modes transformed to the time-frequency domain using the Hilbert transform for EEMD and NMD respectively. The Hilbert transform generates the analytic signal of a sinusoidal oscillation, which can then be used to calculate its instantaneous frequency. This offers a direct comparison with the time-frequency analysis shown in Fig. 3.

below 0.04 Hz in (a)) or the window running over the edge of the time series. This region defines the *cone of influence* for the times where oscillations of certain frequencies can be observed. However, the limitations of the time-frequency resolution are much more apparent in the STFT. In (a) the short windows allow the highest-frequency mode to be distinguished but the second mode is blurred and the third is not visible due to the low-frequency limit. In (b) the two lower-frequency modes are resolved but the highest-frequency mode fades into the background as its frequency variability cannot be tracked by the large window. In contrast, the adaptive resolution of the CWT in (c) makes it possible to resolve all 3 components simultaneously. In (d) the use of a higher central frequency improves the frequency resolution of each of the components while still allowing the variation in these frequencies to be tracked.

IV. DECOMPOSITION

The Morlet wavelet transform provides the best compromise when dealing with signals in the time-frequency domain and can be used to track time-varying oscillations. It can be interpreted in the same way as the STFT, but while this compatibility is an advantage it also means that the wavelet transform inherits the same problem of generating harmonics for nonlinear oscillations since it is still a linear method.

Further work is therefore needed to deal with the problem of decomposing and extracting individual nonlinear modes. One way to do this is by using Empirical Mode Decomposition (EMD) [41]. In this method the components are extracted by marking all of the peaks and troughs in a time series and interpolating between these two sets of points using splines. The average of these two ‘margins’ is used to define the trend of the time series, which contains dynamics relating to all but the highest-frequency component in the time series.

Subtracting this trend leaves the highest-frequency component. If there are still trends the process is repeated until

$$N_p + N_t - N_z = 0 \text{ or } \pm 1, \quad (8)$$

where N_p , N_t and N_z are the number of peaks, troughs and zero-crossings respectively. Once this condition is met the extracted component is subtracted from the original time series and the method continues by attempting to extract the next highest-frequency component.

One of the problems with EMD is mode mixing, which happens when the amplitude of a mode falls to zero. The result is that the next mode replaces the one being extracted, which can result in errors if the difference in amplitudes is large. This problem has been tackled by repeating the procedure on the same time series with different iterations of additive noise and taking the average of the result, a technique known as *ensemble* EMD (EEMD) [42]. However, the iterative nature of EMD is still susceptible to error propagation, causing small errors in the extraction of high-frequency components to affect the extraction of the lower-frequency components.

An alternative to EMD is to use information from the time-frequency domain to decompose the time series. Specifically, identifying all of the harmonics of a nonlinear mode in the time-frequency domain makes it possible to separate and reconstruct this mode in the time domain. This can be done in the wavelet transform by making use of the time-dependent phase information of the oscillations, which is given by $\phi(s, t) = \arg[W_t(s, t)]$. Assuming the waveform of the oscillation keeps the same shape, the phases of harmonics will share the same dynamics. This means that two harmonics at scales s_1 and s_2 will have the relation $\phi(s_1, t) = (s_1/s_2)\phi(s_2, t)$.

Therefore, despite the fact that the wavelet transform is a linear transformation, methods have now been developed that use it to find and extract nonlinear oscillations [43]. Given that the noise fluctuations for the harmonics are different from the

noise fluctuations in the fundamental mode, the real dynamics of the mode can be more easily separated from the noise this way. The most recent method, Nonlinear Mode Decomposition (NMD), uses information from the harmonics to improve the extraction of the mode as a whole [44]. Specifically, the method is based on ridge extraction where the highest peak over a defined frequency range in the wavelet transform is traced over time. This line can then be used to extract the oscillation at the defined times and frequencies. The algorithm then looks at the phase and amplitude variations of several oscillations together in order to distinguish harmonics (which share the same variations) from genuine independent modes.

Fig. 4 compares EEMD with NMD in the analysis of the example time series $x(t)$. While the oscillatory modes are observable in both decompositions, NMD is more selective and does not extract modes relating to the $1/f$ noise. The continuous noise distribution means that the modes calculated using EMD suffer from mixing, which means it is not possible to isolate the three main oscillatory modes. Instead, most of the EMD modes are associated with high-frequency components which try to fit the original time series as closely as possible and as such do not have much physical meaning.

V. CHARACTERISATION

After transforming a signal to the time-frequency domain and/or extracting its oscillatory modes, the next step is to use this representation to characterise the dynamics of the system. This means connecting what is seen in the signal with properties that can be related to a physical system.

A. Power

The power spectrum of a time series is defined in the frequency domain as the integral of the square of the amplitude. For the Fourier transform this is a straightforward process since the frequency scale is linear, causing the square of the Fourier transform to be directly proportional to the power spectrum. Similarly, the wavelet power spectrum can be found using

$$P_W(\omega', t) = \int_{\omega' - \frac{d\omega}{2}}^{\omega' + \frac{d\omega}{2}} |W_T(\omega, t)|^2 d\omega. \quad (9)$$

However, in the case of the wavelet transform the frequency scale is logarithmic, which means that the components at higher frequencies correspond to larger frequency intervals. Obtaining the power spectrum from the Morlet wavelet transform is also not as simple because the transform is continuous. This means that although the wavelet amplitude is analogous to the Fourier amplitude, for finite data the integration of the squared amplitude to find the power is always an estimate (a continuous curve cannot be integrated discretely).

Taking the average of P_W in time provides a good starting point in the analysis of any time series data. Specifically, it is used to identify the frequency range of the main oscillatory components. Once this is known, the variation in the power of each individual component over time can be found using one of the two methods shown in Fig. 5(b). The first method involves summing over the wavelet transform in the frequency

intervals defined in the time-averaged power spectrum for each point in time (black line). The second method is to instead follow the peak in the power spectrum for the given frequency interval in time (grey line), which is known as ridge extraction [45]. As can be seen, the two methods show similar fluctuations caused by the noise in the system although the second method is less susceptible to these variations.

B. Bispectrum

The bispectrum is a frequency-frequency domain method that arises from high-order statistics [46]. Specifically, the bispectrum is a third order statistic, in the same sense that the skewness of a data series is of the third order, which comes after the mean (first order) and variance (second order) [47]. In this case the bispectrum is the next order measure after the frequency domain spectrum of a time series.

The bispectrum provides information about the quadratic properties of the time series, which makes it ideal for investigating nonlinear couplings between oscillations. However, the frequency-frequency domain is still unable to track time-variability. Therefore, similar to the need for time-frequency analysis, a need for time-frequency-frequency analysis lead to a proposal of wavelet-based bispectral analysis [48]. The wavelet bispectrum is given by,

$$B^W(s_1, s_2) = \int_L W_T(s_1, t) W_T(s_2, t) W_T^*(s_3, t) dt, \quad (10)$$

where $s_3 = 1/(\frac{1}{s_1} + \frac{1}{s_2})$. It is also possible to define an *instantaneous* bispectrum with amplitude $A(s_1, s_2, t) = |W_T(s_1, t) W_T(s_2, t) W_T^*(s_3, t)|$ and phase $\phi(s_1, s_2, t) = \phi(s_1, t) + \phi(s_2, t) - \phi(s_3, t)$.

Couplings between two oscillations at s_1 and s_2 can be identified by peaks in the amplitude of the bispectrum or by observing the dynamics of the phase $\phi(s_1, s_2, t)$, where if the phase is constant a coupling exists [48]. In the case of the amplitude though, the value is also dependent on the amplitude of the oscillations in the wavelet transform. To remove this effect a normalised version can be defined as

$$b(s_1, s_2) = \frac{|B^W(s_1, s_2)|}{\sqrt{\int_L |W_T(s_1, t) W_T(s_2, t)|^2 dt \int_L |W_T(s_3, t)|^2 dt}}, \quad (11)$$

where $b(s_1, s_2)$ is known as the bicoherence and takes values between 0 and 1. However, even with this normalisation it is important to note that the bispectrum / bicoherence will still be non-zero for Gaussian white noise. These random peaks are biased towards lower frequencies, with a chi-squared distribution [49], [50], [51].

An additional complication comes from dealing with the scale s_3 , which causes the bispectrum to become meaningless as both $f_1 = 1/s_1$ and $f_2 = 1/s_2$ approach the Nyquist frequency $f_s/2$. This is because f_3 starts to take amplitude and phase information at frequencies which are outside the observable range. Couplings at the highest frequencies are therefore not detectable, meaning an 'effective' Nyquist frequency for the bispectrum is defined as the line from $f_1 = f_s/4$, $f_2 = f_s/4$ to $f_1 = f_s/2$, $f_2 = 0$ (and vice versa for when f_1 and f_2 are switched).

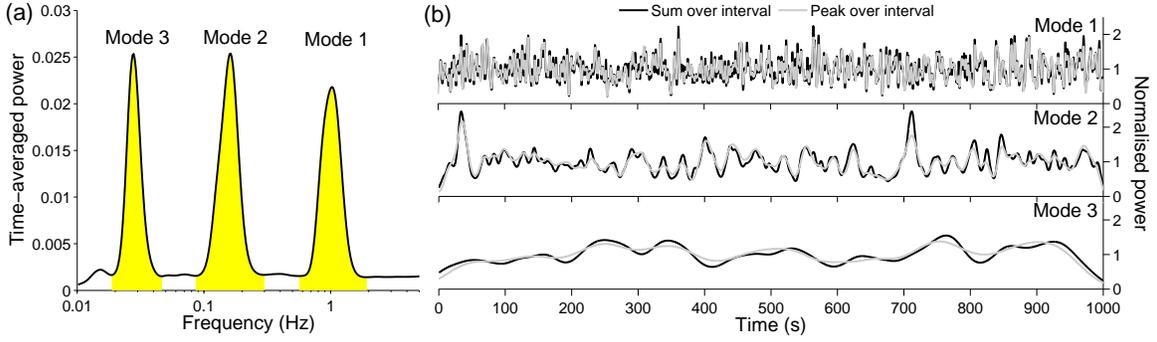


Fig. 5. Wavelet power spectrum of the time series (1). The time-averaged power spectrum is shown in (a), where the frequency ranges of the three modes are defined using the minima at each side of the peak (yellow shaded regions). In (b) the power of the three modes is traced in time using the sum and peak methods. The power is normalised by dividing by the sum or peak of the time-averaged spectrum respectively.

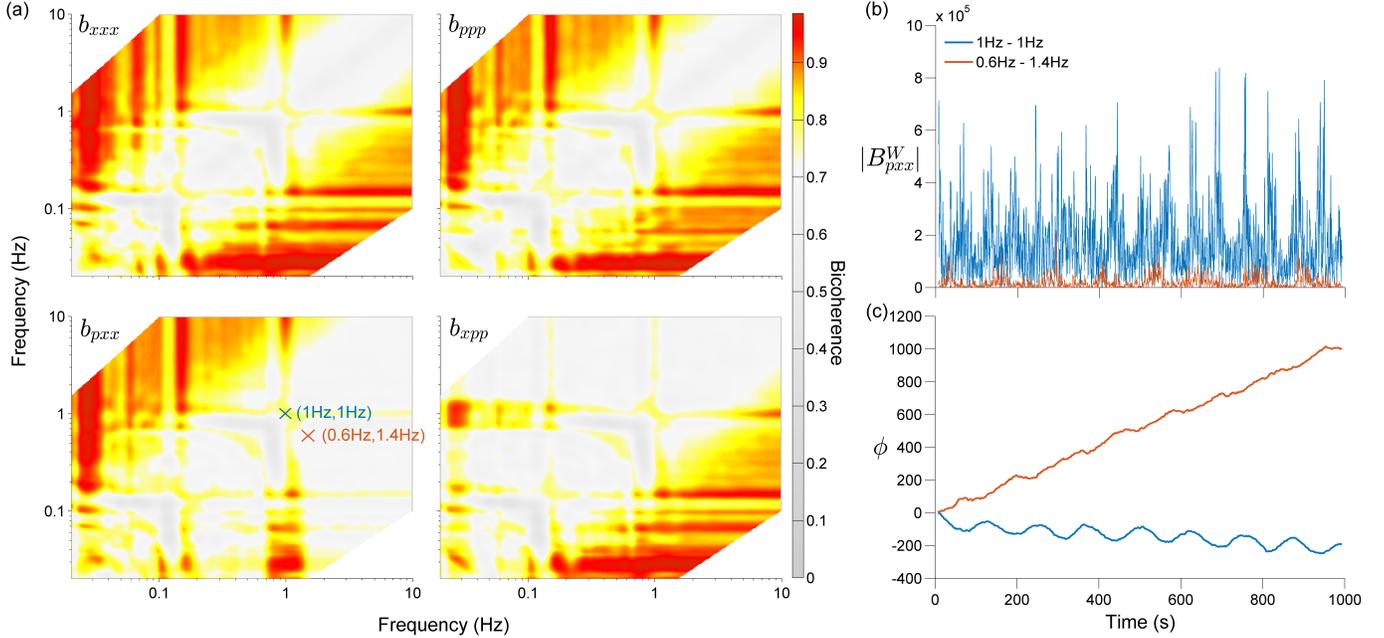


Fig. 6. Wavelet bispectrum analysis for the time series $x(t)$ and $p(t)$. The bicoherence for the various combinations of the cross-bispectrum are shown in (a), where without the time axis it is difficult to distinguish the amplitude due to noise fluctuations from the amplitude contributions of genuine couplings. In (b) and (c) the bispectral amplitude and phase are shown for the points marked in (a) which correspond to the coupled frequency pair (1 Hz, 1 Hz) and unrelated frequency pair (0.6 Hz, 1.4 Hz). While the amplitude is higher for the coupled pair, the coupling is also indicated by the fact that the phase does not grow over time.

One other disadvantage is that wherever the difference in the frequencies of a pair of oscillations is large, the adaptive frequency resolution of the wavelet transform means that the combined frequency $f_{\text{HIGH}} + f_{\text{LOW}} \approx f_{\text{HIGH}}$, making these couplings undetectable. On the other hand, the logarithmic frequency scale of the wavelet transform means that couplings between pairs of low frequencies can be identified.

The same method can also be used to detect couplings between components from different time series. The cross-bispectrum can be defined in several ways [48], [52], [53] using different combinations of the three wavelet components in (10), i.e.

$$B_{122}^W(s_1, s_2) = \int_L W_1(s_1, n) W_2(s_2, n) W_2^*(s_3, n) dn, \quad (12)$$

where W_1 and W_2 are the wavelet transforms of the corresponding time series. The wavelet cross-bicoherence can similarly be defined as [54]:

$$b_{122}^W(s_1, s_2) = \frac{|B_{122}^W(s_1, s_2)|}{\sqrt{\int_L |W_2(s_1, t) W_2(s_2, t)|^2 dt \int_L |W_1(s_3, t)|^2 dt}}. \quad (13)$$

By comparing the cross-bispectra from different combinations it is also possible to deduce some information about the direction of coupling between the oscillations in two separate time series.

Fig. 6 shows the bispectral analysis of the time series $x(t)$ and $p(t)$. The time-averaged bicoherence shows interactions between the modes but also across a much wider range of

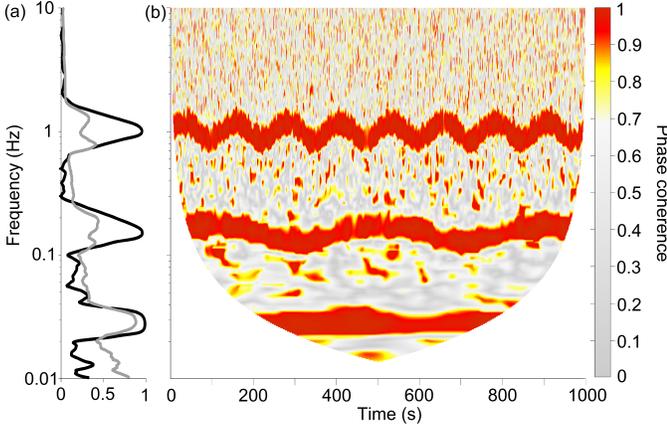


Fig. 7. Wavelet phase coherence between the time series $x(t)$ and $p(t)$. In (a) significant phase coherence is shown when the coherence (black line) is greater than the 95th percentile of 100 pairs of IAAFT surrogates (grey line). The windowed phase coherence is shown in (b), which reveals the time-variability of the modes but at the cost of losing information about lower frequencies.

frequencies. These interactions appear due to the $1/f$ noise present in the time series, which makes the time-averaged bispectrum difficult to interpret. However, when the phase of a coupled frequency pair is compared in time with that of two unrelated frequencies it is clear that the phase corresponding to the coupled frequencies is locked and stays close to its original value.

C. Coherence

Waves can be coherent in space and oscillations can be coherent in time. Generally, coherence in time describes all properties of correlation between physical quantities of a single oscillation, or between several oscillations. Like power, bispectral amplitude and phase, coherence is defined for specific frequencies. If at a certain frequency the changes in amplitude and phase of an oscillatory component are the same for the same oscillation observed in different time series, then they are said to be coherent at that frequency. While this is the general definition of coherence, it is also possible to define separately *amplitude* and *phase* coherence, which consider only matching amplitude or phase dynamics of the components respectively.

A phenomenon related to coherence is synchronization. However, while oscillations can be coherent without direct coupling, synchronization is a property of the underlying system which results from a coupling between two oscillations. This phenomena is not as trivial as one might expect, with multiple ways to both define and detect synchronization [55]. Specifically, oscillators can be phase synchronized, phase and amplitude synchronized or Lyapunov synchronized (also known as generalised synchronization) and also have $n : m$ relations where there are n cycles of one oscillator in m cycles of the other. Phase coherent oscillations can result in 1:1 phase synchronization.

In the case of biomedical signals, measures of phase synchronization are often used as a simple way to observe interactions between two oscillations. Methods such as the

phase synchronization index typically rely on the detection of phase locking, where the phase shift between two oscillations remains constant [56]. However, if only 1:1 synchronization between two signals is of interest then it is more straightforward to consider the phase coherence instead, which is defined as [57],

$$\Pi(s) = \frac{1}{N} \left| \sum_{n=1}^N e^{i(\phi_1(s, t_n) - \phi_2(s, t_n))} \right|, \quad (14)$$

where $\phi_1(s, t_n) - \phi_2(s, t_n)$ is the phase difference between the oscillatory components of the same frequency from two signals at time t_n . If the oscillations remain phase locked for all time (i.e. the oscillations are coherent) then $\Pi(s) = 1$, whereas if $\Pi(s) = 0$ there is no tendency to preserve a particular phase difference. A more general definition is also given by [58],

$$\Pi = \left\{ \left[\frac{1}{2} \sum_{n=1}^N w_1(t_n) w_2 * (t_n) \right] \left[\frac{1}{2} \sum_{n=1}^N w_1(t_n) * w_2(t_n) \right] \right\}^{1/2} \quad (15)$$

where $w_i(t)$ represents any time-frequency representation with complex values corresponding to the analytic signals of the oscillations. In this paper Π will represent the phase coherence of the power-normalised wavelet transform, which is identical to the first definition shown above.

The phase coherence can be calculated systematically for all wavelet scales s to provide a graph of coherence vs. frequency [59], [60], [61]. However, even in the phase coherence of two noise signals there is some level of coherence. This means that the coherence rarely approaches 0 and significant coherence is usually close to 1. Additionally, this baseline coherence is not constant for all scales but increases when moving lower in frequency to the point where at the lowest observable frequencies $\Pi(s) \approx 1$ even if the dynamics is unrelated.

This bias towards lower frequencies can be accounted for by using *surrogates* of the signals [62], [63]. These are designed to preserve all of the properties of the original signals apart from the property relating to the hypothesis that is being tested. In the case of phase coherence, the null hypothesis is that the phases in the signals are independent for all frequencies, which means that it is the time-phase information that needs to be randomised in the surrogates. In this case, iterative amplitude adjusted Fourier transform (IAAFT) surrogates fit this null hypothesis [64]. The phase coherence between these surrogates at each frequency can then be used as a baseline above which the coherence is said to be significant.

In order to track time-variations in the phase coherence, the calculation can be performed over time by using a sliding window along each scale. The time-localized phase coherence is able to avoid the issue of frequency bias since the windows can be scaled appropriately with frequency so that the window will always contain the same number of cycles [58]. However, this successive windowing of the signals can result in a loss of low-frequency information. Firstly, preprocessing is often required to detrend the data and remove the effect of high harmonics of oscillations that are too slow to be observed, usually resulting in the loss of data at the edges. The cone of influence from the continuous wavelet transform also causes

a significant reduction in the part of the time series that is observable at low frequencies. Finally, to observe significant coherence at least 5 cycles need to be observed [36], which further raises the low-frequency limit of the analysis.

Fig. 7 shows the phase coherence of the example time series $x(t)$ and $p(t)$. It can be seen that the coherence is only greater than the surrogate level at the frequencies corresponding to the common oscillatory modes. Note however that the surrogate level is inversely proportional to the frequency, which means it is more difficult to detect significant coherence in low-frequency oscillations. The windowed wavelet phase coherence also reveals the shared time-variability of these modes.

D. Information and entropy

Wavelets are not the only available tool for the analysis of interactions in complex systems. Another way to detect couplings is by using statistics based on information theory, such as transfer entropy [65] and Granger causality [66], [67]. In the latter case, a coupling is said to exist if one system gives information about the state of the other system at some point in the future [48], [62], [68], [69]. Starting with the probability distributions of the two time series, $p(x_1(t))$ and $p(x_2(t))$, the Shannon entropy for each can be defined as

$$H(x_i) = - \sum p(x_i) \log p(x_i), \quad (16)$$

which gives a measure of the uncertainty or ‘randomness’ in x_i . The *joint* entropy can also be defined as

$$H(x_i, x_j) = - \sum \sum p(x_i, x_j) \log p(x_i, x_j), \quad (17)$$

where $p(x_i, x_j)$ is the 2-dimensional joint probability distribution. The amount of common information contained in x_i and x_j , which is analogous to the inverse of the joint entropy, is given by the mutual information:

$$I(x_i; x_j) = H(x_i) + H(x_j) - H(x_i, x_j). \quad (18)$$

Finally, the conditional entropy is defined as

$$H(x_i|x_j) = - \sum \sum p(x_i, x_j) \log p(x_j|x_i), \quad (19)$$

where $p(x_j|x_i)$ is the probability distribution for x_j if the value for x_i is given. The dependence between x_i and x_j without the possible influence of another variable x_3 can then be defined using the *conditional* mutual information (CMI),

$$I(x_1; x_2|x_3) = H(x_1|x_3) + H(x_2|x_3) - H(x_1, x_2|x_3). \quad (20)$$

Consider now two time series $x(t)$ and $y(t)$. The information flow from x to y is given by $I(x; y_d|y)$, where y_d is the delayed time series $y(t + \tau)$ with $l\Delta t = \tau$. This quantity excludes information from both the history of $y(t)$ on itself and the common history of $x(t)$ and $y(t)$ [48]. Similarly, the information flow from y to x is given by $I(x; y_d|x)$. Therefore, the strength of coupling from one time series to another should be indicated in the amount of information flow in the corresponding direction.

In reality, there is always going to be some baseline mutual information contained within even two completely unrelated time series. This is why the method requires the use of

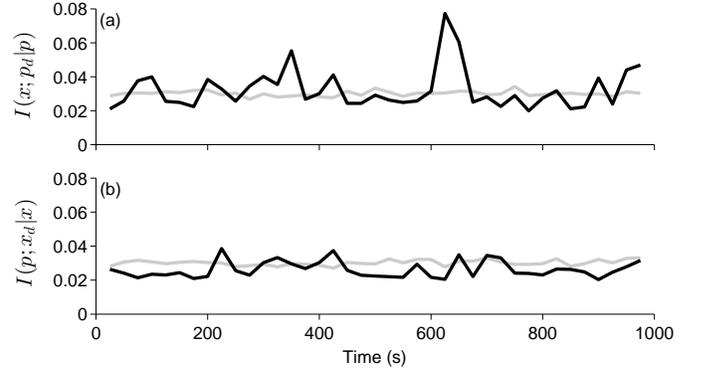


Fig. 8. Information analysis of the time series $x(t)$ and $p(t)$. The solid black lines show the average CMI calculated for delays ranging from 0.05 s to 5 s. The grey lines show the 75th percentiles of the CMI calculated between 100 pairs of IAAFT surrogates. The information transfer in the direction $p \rightarrow x$ is shown in (a), while (b) shows the information transfer in the direction $x \rightarrow p$.

surrogate data to determine whether there is a *significant* amount of information being transferred either from $x_1 \rightarrow x_2$ or $x_2 \rightarrow x_1$. It is also dependent on estimates of the probability distributions of the time series, which require careful consideration. However, the main advantage of this approach is that it is not restricted by frequencies; the CMI gives a measure of the information transfer between two arbitrary sets of data, rather than being localised in any one domain. It is also worth noting that Granger causality can be calculated using other methods that do not rely on CMI [70].

An advantage of information and entropy based measures is that they are dimensionless, meaning that they can be applied to any type of signal. In particular, they can be applied both to the raw signal but also to the extracted phases of the oscillations in order to determine specific phase-phase interactions [48], [68].

Fig. 8 shows the CMI of the time series defined in (1) and (2). It can be seen that there is much more significant information transfer above the surrogate level for $I(x; p_d|p)$ as opposed to $I(p; x_d|x)$, which suggests a coupling from the modes in $p(t)$ to the modes in $x(t)$. However, it is also worth noting the times when there appears to be no transfer of information. This can be explained by the fact that the chronotaxic modes are very close to being phase synchronized with the driving modes in p . When they do become fully synchronized there is no transfer of information, which means that even if there is a coupling it is not possible to detect one.

VI. INTERPRETATION

The information learned from the characterisation of the signals can only go so far. For example, chaotic, stochastic and non-autonomous systems can all generate similar continuous power spectra [28]. The bispectral couplings also provide basic information about the direction and strength of the coupling but little information about the form of the function, while phase coherence does not provide any information about the means by which the signals become coherent. To obtain more

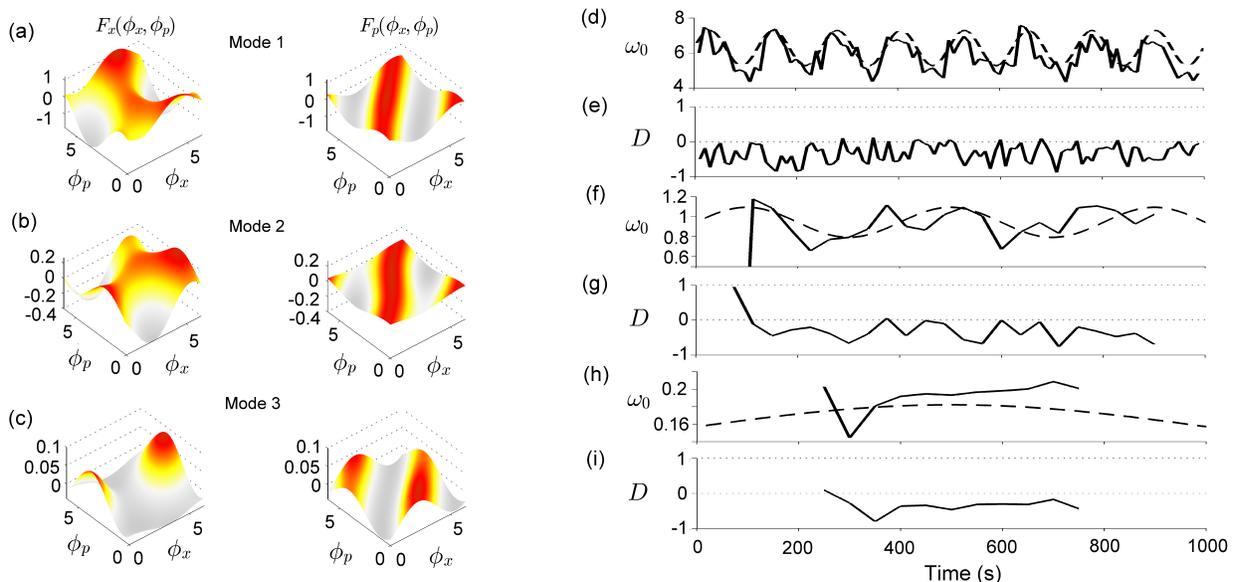


Fig. 9. Dynamical Bayesian inference analysis of the phases extracted using NMD from the time series $x(t)$ and $p(t)$. Plots (a), (b) and (c) show the coupling functions for the pairs of phases extracted from each mode. The inferred value of $\omega_0(t)$ is shown in (d), (f) and (h) by the solid black line, while the dotted line is the actual value. Plots (e), (g) and (i) show the direction of coupling calculated by taking the ratio of the amplitudes for the terms dependent on the other phase for each coupling function. $D > 0$ for a coupling in the direction $x \rightarrow p$ and $D < 0$ for a coupling in the direction $p \rightarrow x$. The model parameters were inferred using a 20 s moving window with 50% overlap for mode 1, a 150 s window with 75% overlap for mode 2 and a 500 s window with 90% overlap for mode 3. In each case the propagation constant had the value $p = 0.2$

information it is necessary to infer the properties of a physical model of the system based on the observed dynamics.

A. Phase oscillator model

Living systems are characterised by a multitude of oscillations over a broad range of timescales. However, the real complexity arises from the interactions between these oscillations, and there have been many attempts to model this accurately [3]. The simplest models of coupled oscillators focus purely on the changes that occur in an oscillator's phase over time and neglect any amplitude variations. This simplification is justified in models of the heart (or similar oscillators with all-or-nothing responses), because it is only changes in the timing of the periodic features that carry significance. Even when a system does vary in amplitude, many oscillators can still remain close to an attracting *limit cycle*, which again causes these variations to be negligible.

An example of a phase oscillator model is given by $d\phi/dt = \omega$, which describes a phase increasing at a rate of ω , which is the natural frequency of the oscillator. A model of two interacting oscillators is given by

$$\begin{aligned} \frac{d\phi_1}{dt} &= \omega_1 + F_1(\phi_1, \phi_2), \\ \frac{d\phi_2}{dt} &= \omega_2 + F_2(\phi_1, \phi_2), \end{aligned} \quad (21)$$

where $F_1(\phi_1, \phi_2)$ and $F_2(\phi_1, \phi_2)$ are *coupling functions* which allow the dynamics of one oscillator to be dependent on the other. Such coupling functions are expected to be periodic on

the phases ϕ_1 and ϕ_2 , which means that they can be modelled by a Fourier series [71],

$$\begin{aligned} F_{1,2}(\phi_1, \phi_2) &= \sum_{l,m} [a_{1,2(l,m)} \cos(m\phi_1 + l\phi_2) \\ &\quad + b_{1,2(l,m)} \sin(l\phi_1 + m\phi_2)], \end{aligned} \quad (22)$$

where $a_{1,2(l,m)}$ and $b_{1,2(l,m)}$ are the parameters which describe the function. By inferring these parameters from the observed dynamics of the system it is possible to gain an in-depth understanding of the oscillations, such as whether they exhibit synchronization, or how they respond to perturbations.

B. Dynamical Bayesian inference

A problem with the model shown above is that the parameters are stationary. This makes it more difficult for this model to reveal information about the coupling functions in non-autonomous systems, which are expected to be time-dependent. However, there is no straightforward way to apply a moving time window because windowing means that a smaller data series goes into the algorithm used to estimate of the couplings, increasing the uncertainty.

The Bayesian theorem offers a solution to this windowing problem. When applied to inverse problems where one would like to infer parameters related to the generation of a data set [72], [73], [74], [75], [76] it is known as *dynamical Bayesian inference*. The theorem is summarised in

$$P(\mathcal{M}|\mathcal{X}) = \frac{P(\mathcal{X}|\mathcal{M})P_{\text{pr}}(\mathcal{M})}{P(\mathcal{X})}, \quad (24)$$

where $P(\mathcal{X}|\mathcal{M})$ is the conditional probability of observing the data \mathcal{X} given the hypothesised parameters \mathcal{M} . $P_{pr}(\mathcal{M})$ is the probability of \mathcal{M} before observing the data \mathcal{X} and

$$P(\mathcal{X}) = \int P(\mathcal{X}|\mathcal{M})P_{pr}(\mathcal{M})d\mathcal{M} \quad (25)$$

is the marginal probability of \mathcal{X} . $P(\mathcal{M}|\mathcal{X})$ is known as the *posterior* probability – the probability that the hypothesised parameters are correct given \mathcal{X} and the *prior* probability $P_{pr}(\mathcal{M})$.

The most likely combination of values for the parameters for a single window of data is inferred by locating the stationary point in the negative-log likelihood function, known as maximum likelihood estimation. In this case the likelihood function is specified for the phases of two systems [75] defined by the following stochastic differential equations,

$$\frac{d\phi_{1,2}}{dt} = \omega_{1,2} + F_{1,2}(\phi_{1,2}) + G_{1,2}(\phi_1, \phi_2) + \xi_{1,2}(t), \quad (26)$$

where $F_{1,2}(\phi_{1,2})$ and $G_{1,2}(\phi_1, \phi_2)$ are coupling functions which, as in the previous methods, are modelled using a Fourier basis. The parameters c_k for this basis are eventually inferred in a covariance matrix denoted Ξ . By making use of Bayes' theorem, the posterior covariance matrix for the previous window can exploit information from the prior covariance matrix Ξ_{prior} for the current window. Hence, information is allowed to propagate between windows, enabling the inferred parameters to become more accurate with time [75].

However, the inference only improves if the parameters do not vary in time. To account for changes in the values of the parameters, the prior can take the form of a convolution between the posterior of the previous window and a diffusion matrix which describes the change in c_k [75]. The standard deviation corresponding to the diffusion of the parameters is assumed to be a known fraction of the parameters themselves, $\sigma_k = pc_k$, where p is known as the propagation constant. This modification allows the method to track the change in the couplings over time.

A tutorial for the implementation of this Bayesian-based approach is provided in [77], which includes a Matlab toolbox.

Fig. 9 shows the method applied to the extracted phases of the example time series $x(t)$ and $p(t)$. The coupling functions for the modes in $x(t)$ have a much higher amplitude than the modes in $p(t)$, which suggests a strong coupling term such as the one present in the chronotaxic modes. The method also reconstructs the time-variability of ω_0 for the first mode. However, for the other modes the frequency variation is not traced as well because the window sizes are much larger in order to cover enough cycles of the low-frequency oscillations. Despite this, in all cases the method correctly identifies a coupling in the direction from the modes in $p(t)$ to those in $x(t)$.

C. Phase fluctuation analysis

Let us now return to chronotaxic systems, which were introduced in Sec. II-C. They are non-autonomous systems and have stable dynamics relative to a time-dependent point

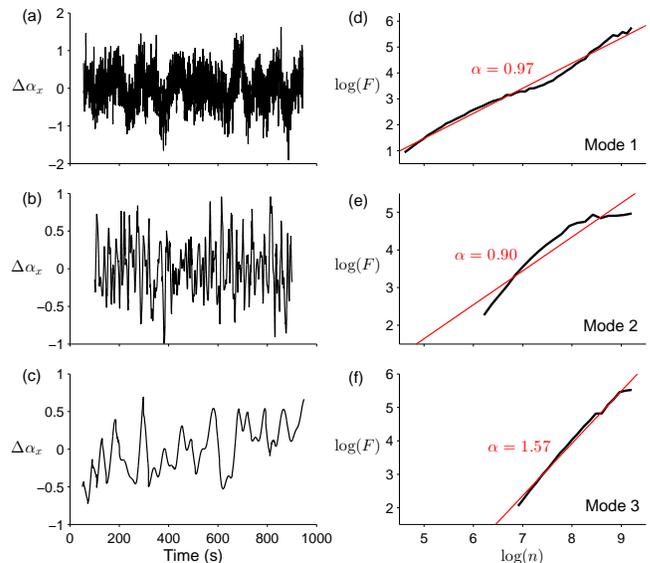


Fig. 10. Phase fluctuation analysis of the modes in the time series $x(t)$. Estimates of the phase α_x for mode were obtained from the wavelet transform using ridge extraction and $\omega_c = 0.5$. Estimates of the phase α^u were found in the same way but using $\omega_c = 2$ and integrating over the smoothed instantaneous frequency. Plots (a), (b) and (c) show the detrended difference between the two phases for each mode. Plots (d), (e) and (f) show the detrended fluctuation analysis of $\Delta\alpha_x$ (solid black lines). Linear least squares fits (red lines) were used to estimate the values of α in each case.

attractor [31]. This property determines how the system responds to perturbations, whether resisting them or allowing them to dictate their dynamics. However, despite this strong dichotomy, the actual effects are not obvious and hidden even in the time-frequency domain. Another approach is therefore needed.

The easiest way to determine whether a system is chronotaxic or not is to observe its fluctuations relative to its unperturbed trajectory. If the original distribution of the perturbations is known, then the stability of the system relative to the unperturbed trajectory (which by definition follows the time-dependent point attractor in a chronotaxic system) can be determined from how these fluctuations grow or decay over time. For example, take the non-chronotaxic phase oscillator

$$\frac{d\alpha_x}{dt} = \omega_0(t) + \eta(t), \quad (27)$$

where $\omega_0(t)$ is the time-dependent natural frequency and the observed phase α_x is perturbed by noise fluctuations $\eta(t)$. Integrating we find

$$\alpha_x = \int \omega_0(t)dt + \int \eta(t)dt. \quad (28)$$

Assuming that $\omega_0(t) > 0$ and $\eta(t)$ is an uncorrelated Gaussian process, this means that the dynamics of α_x will consist of a monotonically increasing phase perturbed by a random walk noise (Brownian motion). However, the situation is different

for a chronotaxic phase oscillator, e.g.

$$\begin{aligned} \frac{d\alpha_p}{dt} &= \omega_0(t), \\ \frac{d\alpha_x}{dt} &= \varepsilon\omega_0(t) \sin(\alpha_p - \alpha_x) + \eta(t), \end{aligned} \quad (29)$$

where α_p is an external phase and $|\varepsilon| > 0$. Here the stability provided by the point attractor causes each noise perturbation to decay over time, preventing $\eta(t)$ from being integrated over to the same extent. The perturbations still do not decay instantly as the system takes time to return to the point attractor, meaning that some integration of the noise still takes place. However, the size of the observed perturbations over longer timescales is greatly reduced, causing a change in the overall distribution from that expected for Brownian motion.²

The phase α_x of the observed system can be estimated using the decomposition methods mentioned in the previous section. However, further work is needed to obtain the unperturbed phase α_x^u . In particular, it is difficult to separate the dynamics corresponding to α_x^u from the effect of the noise perturbations $\eta(t)$. This task is simplified by assuming that the dynamics of α_x^u is confined to timescales larger than a single cycle and that the noise is either weak or comparable in magnitude.

With these assumptions, an estimate of α_x^u can be found by filtering out high-frequency components of α_x . However, such a filter should not smooth over the dynamics of α_x^u . An optimal way of removing these high-frequency noise fluctuations without affecting the unperturbed dynamics is to instead smooth over the frequency extracted from the wavelet transform [33]. This provides the estimated angular velocity $\dot{\alpha}_x^u$, which can in turn be integrated over time to give α_x^u .

Given the estimates of α_x and α_x^u , the next step is to analyse $\Delta\alpha_x = \alpha_x - \alpha_x^u$ to find the distribution of fluctuations in the system relative to the unperturbed trajectory.

In order to quantify the distribution of fluctuations, detrended fluctuation analysis (DFA) can be performed on $\Delta\alpha_x$ [78], [79]. This method provides an estimate of the fractal self-similarity of fluctuations at different timescales. The scaling of these fluctuations is determined by the self-similarity parameter α , where fluctuations at timescales equal to t/a can be made similar to those at the larger timescale t by multiplying with the factor a^α .

In order to calculate α the time series $\Delta\alpha_x$ is integrated in time and divided into non-overlapping sections of length n . For each section the local trend is removed by subtracting a fitted polynomial – usually a first order linear fit [78], [79]. The root mean square *fluctuation* for the scale equal to n is then given by

$$F(n) = \sqrt{\frac{1}{N} \sum_{i=1}^N Y_n(t_i)^2}, \quad (30)$$

where $Y(t)$ is the integrated and detrended time series and N is its length. The fluctuation amplitude $F(n)$ follows a

²Note that this assumes the noise does not cause phase slips in α_x . This would cause perturbations over large timescales (i.e. greater than one cycle) to not decay even if the system was chronotaxic. In these cases another approach should be used instead [33].

scaling law if the time series is fractal. By plotting $\log F(n)$ against $\log n$, the value of α is simply the gradient of the line. For completely uncorrelated white Gaussian noise (the noise assumed to perturb the system) the parameter for α has a value of 0.5, while integrated white Gaussian noise (expected in non-chronotaxic systems) returns a value of 1.5.

When $\alpha < 1.5$ this therefore suggests that there is some resistance to perturbations (chronotaxicity) which prevents their integration over larger timescales. The actual value is typically dependent on the gradient of the coupling function relative to $\alpha_p - \alpha_x$. If the gradient close to the point attractor is very steep then the system returns to the attractor more quickly after being perturbed and less integration of the noise occurs, resulting in smaller values of α .

Fig. 10 shows the phase fluctuation analysis for the extracted phases of the modes in the time series $x(t)$. The method finds $\alpha < 1$ for the first two modes, which suggests that they are chronotaxic. However, the method identifies the third mode as being non-chronotaxic since $\alpha \approx 1.5$. This is due to the fact that not enough cycles of the oscillation are observed to reliably determine whether the mode is chronotaxic or not.

VII. APPLICATIONS

To demonstrate how the described analytical framework can be used in practice, the methods from the previous sections are now applied to real biomedical signals.

A. Skin microvascular flow evaluated by Laser Doppler flowmetry (LDF)

LDF is a technique applied to measure blood flow in the microvasculature. It involves shining laser light into the microvascular bed, that includes capillaries and small arterioles, and measuring the Doppler shift in the light caused by the movement of the blood. This movement is influenced by a wide range of oscillations of different frequencies ranging from 0.005 Hz to 2 Hz, originating from both systemic and local processes [7], [3], [79]. As each of these oscillations is also time-varying, the result is a very complex signal.

Fig. 11 shows the time-frequency analysis of an LDF signal. The strong cardiac oscillation is easily recognised in the Fourier transform. However, there is also power at low frequencies which appear as a continuous noise-like distribution. The wavelet transform reveals these low-frequency fluctuations to be highly nonstationary oscillations, relating to myogenic, neurogenic and endothelial activity [7]. Without time-frequency techniques, such oscillations can often elude discovery or be discounted as noise [26].

There are limitations in the analysis of such signals, however. While the power spectrum can almost always be used to characterise the dynamics of the underlying processes, attempting to decompose and analyse the low-frequency oscillations in an LDF signal is a very difficult task. The oscillations are simply too nonstationary, and the available part of the time series is too short, due to the cone of influence. If an LDF signal was several hours long then the low-frequency dynamics could be analysed using the other decomposition, characterisation and interpretation methods discussed. However, the length

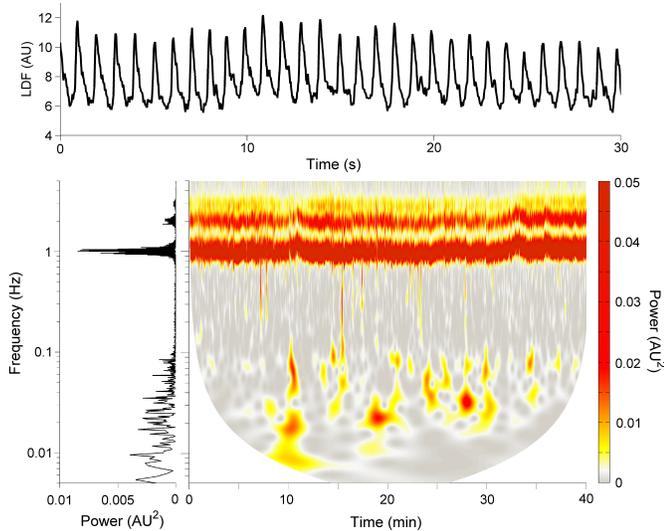


Fig. 11. Analysis of an LDF signal measured from the left arm for 40 minutes. In the time domain (top) the cardiac oscillation is clearly seen, with ~ 1 s pulses. A Fourier transform (bottom left) provides a representation of the signal in the frequency domain, which also shows the mode with the largest power to be at the cardiac frequency ~ 1 Hz with its harmonics at ~ 2 Hz and ~ 3 Hz. There is also power at lower frequencies but this appears only as a continuous, noise-like spectrum with no identifiable modes. In the wavelet transform (bottom right) the power at low-frequencies is revealed to be due to nonstationary oscillations whose frequencies and amplitudes vary in time.

of recordings is limited by the fact that subjects must remain motionless throughout as movement artefacts strongly affect the low-frequency components of the signal [80]. But even when only the power of the oscillations is considered, without attempting to consider their coupling, the insights obtained can be of diagnostic and prognostic use, as shown in a recent study of blood flow in melanoma [21].

On larger scales, the dynamics of the cardiovascular system is no less complex. In particular, the cardiorespiratory interaction has been shown to exhibit time-dependent coupling functions which cause changes in the synchronization between the heart and lungs [75]. This interaction is primarily defined by the phase relationship between the two systems, which means that it is maintained even when the variations in the heart rate or respiration amplitude are small. Using bispectral analysis this coupling has been shown to propagate to the microcirculation [81].

B. Dynamics of brain waves and their interaction in anaesthesia and awake states

The neuronal activity in the brain has long been characterised by existence of brain waves [4] and we will briefly illustrate how interactions between brain waves can be extracted from an EEG signal. The signal was recorded in the BRACCIA project with electrodes attached to the forehead of a patient under anaesthesia [23]. The traditional waves, δ (0.8–4 Hz), θ (4–7.5 Hz), α (7.5–14 Hz), β (14–22 Hz) and γ (22–100 Hz) were studied. Lower frequency oscillations have also been identified [4], [3], but will not be discussed here. The results of bispectral analysis and dynamical Bayesian inference are summarised in Fig. 12.

When analysed using the wavelet bispectrum, the noise in the signal makes it difficult to interpret the interactions from its amplitude alone. However, the phase of the bispectrum for each frequency coupling shows different rates of change, related to the coupling strength between the brain waves. After extracting the phases for each brain wave, Bayesian inference reveals the coupling functions between the oscillations, as well as the magnitude of these couplings. By observing the ratios between the grey and black lines it is possible to infer the direction of coupling between the brain waves. It can be seen that the other waves drive the γ wave, which is also observed in the form of the F_γ coupling functions. While this provides clear evidence of functional interactions between pairs of brain waves, the same techniques have also been used to show the existence of triplet interactions [82].

VIII. DISCUSSION

Biomedical signals, arising from nonlinear, time-dependent living systems, provide an opportunity to monitor the underlying dynamics of the observed system. The time-variability of biomedical data necessitates the application of time-frequency analysis methods in the first instance. If identified as a stochastic process, the signal may be further characterised using statistical methods. If the signal is found to contain distinct oscillatory modes, these may then be extracted and separated using the techniques presented here. The interactions between these modes can be then investigated to provide yet another layer of information about the dynamics of the system.

Living systems appear to possess underlying preferred amplitudes and frequencies to which the system will return when external influences are removed. To bridge the gap between dynamical systems theory and this apparent stability, a new class of nonautonomous system was introduced, known as chronotaxic systems [31], [32]. This led to the development of methods for the detection of chronotaxicity and their application to real data [33], [34]. This provides a framework in which experimentally observed fluctuations, which may previously have been regarded as noise, or arising from chaotic dynamics, may actually be considered as systems with underlying elements of determinism.

Although techniques for the analysis of biomedical time series have greatly improved, there are still some limitations to overcome. One requirement of methods based on the interpretation of oscillatory activity is a sufficient number of periods of oscillation, especially if the frequency is time-varying, usually set at 5 periods of oscillation. For very low frequencies this may necessitate a time series longer than it is possible to record. For short time series, the presented methods will become less reliable, or even impossible to implement. If the time series is long enough, the analysis approach used strongly depends on the characteristics of the data. For example, in chronotaxic systems, time-frequency analysis will highlight any oscillatory components, but will not provide any information about interactions in that system, thus leading to the need for further extraction to identify chronotaxicity. Similarly, wavelet phase coherence may reveal phase relationships between two signals, but cannot provide any information on

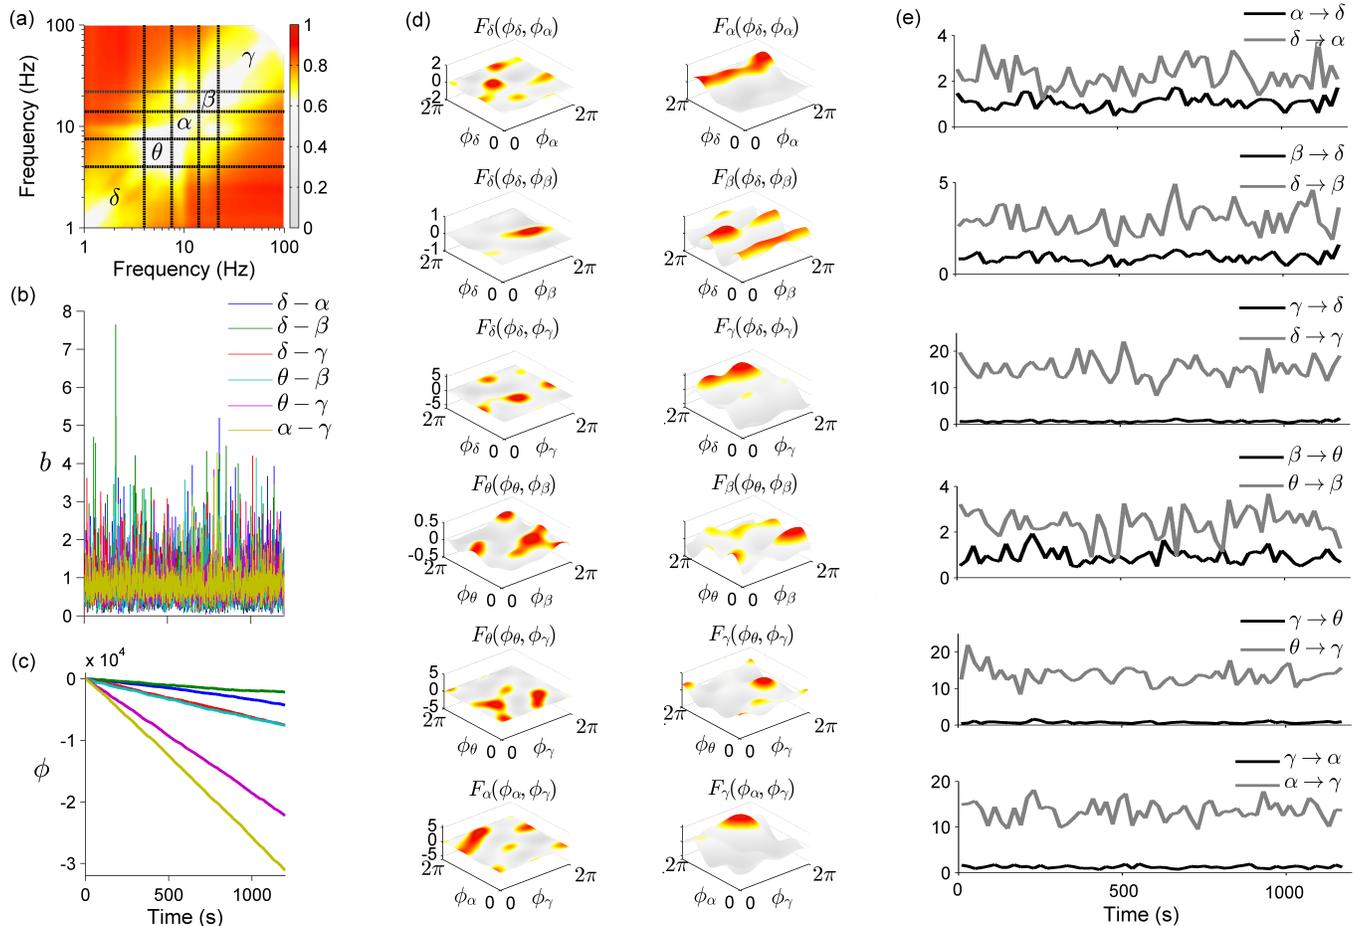


Fig. 12. (a-c) Bispectral and (d,e) dynamical Bayesian inference analysis of an EEG signal. The signal was measured for 20 minutes from the forehead of a subject in anaesthesia. The phases for the δ , θ , α , β and γ waves were extracted using NMD. The plot in (a) shows the bicoherence of the raw EEG signal, while (b) and (c) show the instantaneous bicoherence and phase of the bispectrum respectively for the pairs of brain waves. In (d) the coupling functions for the different pairs of extracted phases are shown (couplings between adjacent bands are not shown due to frequency spillage from imperfect filtering). On the right in (e) are the magnitude of the coupling functions for each point in time, providing an indication of the direction of coupling between the phases. The model parameters were inferred using a 20 s moving window with no overlap and with the propagation constant $p = 0.2$.

their origins. Also important when extracting modes from a time-frequency representation is the frequency variation. If this variation is too fast, or more than one mode is present, it may be difficult to reliably extract a single mode. The frequency resolution in the wavelet transform may be changed to better resolve frequency components, but this comes at the expense of time resolution and may still be insufficient.

In some living systems, only the phase dynamics is considered. For example, in the heart only the changes in beat rate can be directly measured with ease, while the cardiac output (the “amplitude” of the heart) is very difficult to quantify by noninvasive techniques. Whilst some of the presented methods rely on the fact that amplitude variations may be negligible, this is not always a valid assumption. Again using chronotaxic systems as an example, the current inverse approach methods for the detection of chronotaxicity only take into account phase dynamics. However, it is known that in many living systems, both phase and amplitude dynamics are important, as are the interactions between them. In particular, the brain is characterised by both spatial and temporal dynamics [83], [84]. Thus, further work is required to develop methods which are

applicable in all these scenarios, although some have already been proposed [85]. It is clear that to gain as much information as possible from biomedical signals, the optimal solution is to combine different methods according to the information required, as demonstrated here.

IX. CONCLUSION

We have presented the latest techniques for the analysis of signals originating from time-dependent dynamics, with suggestions for their optimal implementation. Using these guidelines, it was shown how these signals can be decomposed, characterised and interpreted to reveal a wealth of information about the underlying system. This is of particular significance in the analysis of biomedical data due to the open nature of living systems. We also discuss that one should not arbitrarily apply these time-dependent methods to any signal and how signals should be analysed in several stages starting with time-frequency analysis. It is often the case that accurate information extracted from the previous level of analysis is necessary to proceed to the next level, such

as the ability to resolve modes in the time-frequency domain before decomposing them, or the determination of the phase and frequency of the modes through decomposition before extracting the coupling functions and chronotaxicity of the underlying dynamics.

Using the numerically-generated time series of a chronotaxic system we have illustrated how to reconstruct the dynamics of noisy, time-dependent systems. We have also shown how these same methods can be applied to real biomedical data where the information obtained has physiological relevance, revealing a wealth of information about the underlying living system.

ACKNOWLEDGMENT

We thank the BRACCIA research team for valuable discussions, and Tomislav Stankovski for helpful comments on the manuscript.

REFERENCES

- [1] H. Haken, "Cooperative phenomena in systems far from thermal equilibrium and in nonphysical systems," *Rev. Mod. Phys.*, vol. 47, pp. 67–121, 1975.
- [2] ———, *Synergetics, An Introduction*. Berlin: Springer, 1983.
- [3] A. Stefanovska, "Coupled oscillators: Complex but not complicated cardiovascular and brain interactions," *IEEE Eng. Med. Bio. Magazine*, vol. 26, no. 6, pp. 25–29, 2007.
- [4] G. Buzsáki and A. Draguhn, "Neuronal oscillations in cortical networks," *Science*, vol. 304, pp. 1926–1929, 2004.
- [5] H. Obrig, M. Neufang, R. Wenzel, M. Kohl, J. Steinbrink, K. Einhaupl, and A. Villringer, "Spontaneous low frequency oscillations of cerebral hemodynamics and metabolism in human adults," *Neuroimage*, vol. 12, no. 6, pp. 623–639, 2000.
- [6] N. K. Logothetis, J. Pauls, M. Augath, T. Trinath, and A. Oeltermann, "Neurophysiological investigation of the basis of the fMRI signal," *Nature*, vol. 412, no. 6843, pp. 150–157, 2001.
- [7] A. Stefanovska, M. Bračić, and H. D. Kvernmo, "Wavelet analysis of oscillations in the peripheral blood circulation measured by laser Doppler technique," *IEEE Trans. Bio. Med. Eng.*, vol. 46, no. 10, pp. 1230–1239, 1999.
- [8] G. Hajnóczky, L. D. Robb-Gaspers, M. B. Seitz, and A. P. Thomas, "Decoding of cytosolic calcium oscillations in the mitochondria," *Cell*, vol. 82, no. 3, pp. 415–424, 1995.
- [9] F. T. Kurz, M. A. Aon, B. O'Rourke, and A. A. Armoundas, "Wavelet analysis reveals heterogeneous time-dependent oscillations of individual mitochondria," *Am. J. Physiol. Heart Circ. Physiol.*, vol. 299, pp. H1736–H1740, 2010.
- [10] A. Goldbeter, C. Gérard, D. Gonze, J.-C. Leloup, and G. Dupont, "Systems biology of cellular rhythms," *FEBS Lett.*, vol. 586, no. 18, pp. 2955–2965, 2012.
- [11] B. N. Kholodenko, "Cell-signalling dynamics in time and space," *Nat. Rev. Mol. Cell Biol.*, vol. 7, no. 3, pp. 165–176, 2006.
- [12] D. G. Spiller, C. D. Wood, D. A. Rand, and M. R. H. White, "Measurement of single-cell dynamics," *Nature*, vol. 465, no. 7299, pp. 736–745, 2010.
- [13] K. Tornheim, V. Andrés, and V. Schultz, "Modulation by citrate of glycolytic oscillations in skeletal muscle extracts," *J. Biol. Chem.*, vol. 266, no. 24, pp. 15 675–15 678, 1991.
- [14] W. Klimesch, "EEG alpha and theta oscillations reflect cognitive and memory performance: a review and analysis," *Brain Res. Rev.*, vol. 29, no. 2, pp. 169–195, 1999.
- [15] A. T. Winfree, "Biological rhythms and the behavior of populations of coupled oscillators," *J. Theor. Biol.*, vol. 16, no. 1, p. 15, 1967.
- [16] Y. Kuramoto, *Chemical Oscillations, Waves, and Turbulence*. Berlin: Springer-Verlag, 1984.
- [17] N. H., "Phase reduction approach to synchronization of nonlinear oscillators," *Cont. Phys.*, *accepted*, 2015.
- [18] K. Lehnertz, G. Ansmann, S. Bialonski, H. Dickten, C. Geier, and S. Porz, "Evolving networks in the human epileptic brain," *Physica D: Nonlinear Phenomena*, vol. 267, pp. 7–15, 2014.
- [19] M. Rossi, A. Bradbury, A. Magagna, M. Pesce, S. Taddei, and A. Stefanovska, "Investigation of skin vasoreactivity and blood flow oscillations in hypertensive patients: effect of short-term antihypertensive treatment," *J. Hypertens.*, vol. 29, no. 8, pp. 1569–1576, 2011.
- [20] K. B. Stansberry, S. A. Shapiro, M. A. Hill, P. M. McNitt, M. D. Meyer, and A. I. Vinik, "Impaired peripheral vasomotion in diabetes," *Diabetes Care*, vol. 19, no. 7, pp. 715–721, 1996.
- [21] G. Lancaster *et al.*, "Dynamic markers based on blood perfusion fluctuations for selecting skin melanocytic lesions for biopsy," *Sci. Rep.*, vol. 5, p. 12825, 2015.
- [22] P. L. Purdon *et al.*, "Electroencephalogram signatures of loss and recovery of consciousness from propofol," *Proc. Nat. Acad. Sci.*, vol. 110, no. 12, pp. E1142–E1151, 2013.
- [23] D. A. Kenwright *et al.*, "The discriminatory value of cardiorespiratory interactions in distinguishing awake from anaesthetized states: a randomised observational study," *Anaesthesia*, p. doi:10.1111/anae.13208, 2015.
- [24] M. Czosnyka and J. D. Pickard, "Monitoring and interpretation of intracranial pressure," *J. Neurol. Neurosurg. Ps.*, vol. 75, no. 6, pp. 813–821, 2004.
- [25] P. Kvandal, L. Sheppard, S. A. Landsverk, A. Stefanovska, and K. A. Kirkeboen, "Impaired cerebrovascular reactivity after acute traumatic brain injury can be detected by wavelet phase coherence analysis of the intracranial and arterial blood pressure signals," *J. Clin. Monitor. Comp.*, vol. 27, no. 4, pp. 375–383, 2013.
- [26] C. Baumgartner, K. J. Blinowska, A. Cichocki, H. Dickhaus, P. J. Durka, P. V. E. McClintock, G. Pfurtscheller, A. Stefanovska, and S. Tong, "Discussion of time-frequency techniques in biomedical signal analysis: A tutorial review of similarities and differences," *Methods. Inf. Med.*, vol. 52, no. 4, pp. 297–307, 2013.
- [27] S. H. Strogatz, "Exploring complex networks," *Nature*, vol. 410, pp. 268–276, 2001.
- [28] P. T. Clemson and A. Stefanovska, "Discerning non-autonomous dynamics," *Phys. Rep.*, vol. 542, no. 4, pp. 297–368, 2014.
- [29] A. A. Andronov, A. A. Vitt, and S. E. Khaikin, *The Theory of Oscillators*. New York: Dover, 2009.
- [30] P. E. Kloeden and M. Rasmussen, *Nonautonomous Dynamical Systems*. American Mathematical Soc., 2011.
- [31] Y. F. Suprunenko, P. T. Clemson, and A. Stefanovska, "Chronotaxic systems: A new class of self-sustained nonautonomous oscillators," *Phys. Rev. Lett.*, vol. 111, no. 2, p. 024101, 2013.
- [32] Y. F. Suprunenko and A. Stefanovska, "Generalized chronotaxic systems: Time-dependent oscillatory dynamics stable under continuous perturbation," *Phys. Rev. E*, vol. 90, no. 3, p. 032921, 2014.
- [33] P. T. Clemson, Y. F. Suprunenko, and A. Stefanovska, "Inverse approach to chronotaxic systems for single-variable time series," *Phys. Rev. E*, vol. 89, no. 3, p. 032904, 2014.
- [34] G. Lancaster, P. T. Clemson, Y. F. Suprunenko, T. Stankovski, and A. Stefanovska, "Detecting chronotaxic systems from single-variable time series with separable amplitude and phase," *Entropy*, vol. 17, no. 6, pp. 4413–4438, 2015.
- [35] D. Gabor, "Theory of communication," *J. IEEE*, vol. 93, pp. 429–457, 1946.
- [36] L. Keselbrener and S. Akselrod, "Selective discrete fourier transform algorithm for time-frequency analysis: Method and application on simulated and cardiovascular signals," *IEEE Trans. Bio. Eng.*, vol. 43, no. 8, pp. 789–802, 1996.
- [37] I. Daubechies, *Ten Lectures on Wavelets*. Philadelphia: SIAM, 1992.
- [38] G. Kaiser, *A Friendly Guide to Wavelets*. Boston: Birkhäuser, 1994.
- [39] P. S. Addison, *The Illustrated Wavelet Transform Handbook: Introductory Theory and Applications in Science, Engineering, Medicine and Finance*. Bristol: IOP Publishing, 2002.
- [40] D. Iatsenko, P. V. E. McClintock, and A. Stefanovska, "Linear and synchroqueezed time-frequency representations revisited: Overview, standards of use, resolution, reconstruction, concentration, and algorithms," *Digit. Signal Process.*, vol. 42, no. 3, pp. 1–26, 2015.
- [41] N. E. Huang, Z. Shen, S. R. Long, M. C. Wu, H. H. Shih, Q. Zheng, N. Yen, C. C. Tung, and H. H. Liu, "The empirical mode decomposition and the Hilbert spectrum for nonlinear and non-stationary time series analysis," *Proc. R. Soc. Lond. A*, vol. 454, pp. 903–995, 1998.
- [42] Z. Wu and N. E. Huang, "Ensemble empirical mode decomposition: a noise-assisted data analysis method," *Adv. Adapt. Data Anal.*, vol. 1, no. 1, pp. 1–41, 2009.
- [43] L. W. Sheppard, A. Stefanovska, and P. V. E. McClintock, "Detecting the harmonics of oscillations with time-variable frequencies," *Phys. Rev. E*, vol. 83, p. 016206, 2011.

- [44] D. Iatsenko, P. V. E. McClintock, and A. Stefanovska, "Nonlinear mode decomposition: a noise-robust, adaptive, decomposition method based on the synchrosqueezed wavelet transform," *Phys. Rev. E*, vol. 92, p. 032916, 2015.
- [45] N. Delprat, B. Escudie, P. Guillemain, R. Kronland-Martinet, P. Tchamitchian, and B. T. esani, "Asymptotic wavelet and Gabor analysis: Extraction of instantaneous frequencies," *IEEE Trans. Inf. Theory*, vol. 38, no. 2, pp. 644–664, 1992.
- [46] K. Hasselmann, W. Munk, and G. MacDonald, "Bispectra of ocean waves," in *Time Series Analysis*. New York: Wiley, 1963, pp. 125–139.
- [47] C. L. Nikias and A. P. Petropulu, *Higher-Order Spectra Analysis: A Nonlinear Signal Processing Framework*. Englewood Cliffs: Prentice-Hall, 1993.
- [48] J. Jamšek, M. Paluš, and A. Stefanovska, "Detecting couplings between interacting oscillators with time-varying basic frequencies: Instantaneous wavelet bispectrum and information theoretic approach," *Phys. Rev. E*, vol. 81, no. 3, p. 036207, 2010.
- [49] J. Chung and E. J. Powers, "The statistics of wavelet-based bicoherence," in *Time-Frequency and Time-Scale Analysis*. Pittsburgh, PA, USA: IEEE, 1998, pp. 141–144.
- [50] S. Elgar and R. T. Guza, "Statistics of bicoherence," *IEEE Trans. Ac. Sp. Sig. Proc.*, vol. 36, no. 10, pp. 1666–1668, 1988.
- [51] R. A. Haubrich, "Earth noise, 5 to 500 millicycles per second," *J. Geophys. Res.*, vol. 70, no. 6, pp. 1415–1427, 1965.
- [52] K. S. Lii and K. N. Helland, "Cross-bispectrum computation and variance estimation," *ACM Trans. Math. Soft.*, vol. 7, no. 3, pp. 284–294, 1981.
- [53] M. J. Hinich and G. R. Wilson, "Time delay estimation using the cross bispectrum," *IEEE Trans. Sig. Proc.*, vol. 40, no. 1, pp. 106–113, 1992.
- [54] B. P. van Milligen, E. Sanchez, T. Estrada, C. Hidalgo, B. Branas, B. Carreras, and L. Garcia, "Wavelet bicoherence – a new turbulence analysis tool," *Phys. of Plasmas*, vol. 2, no. 8, pp. 3017–3032, 1995.
- [55] A. Pikovsky, M. Rosenblum, and J. Kurths, *Synchronization – A Universal Concept in Nonlinear Sciences*. Cambridge: Cambridge University Press, 2001.
- [56] P. Tass, M. G. Rosenblum, J. Weule, J. Kurths, A. Pikovsky, J. Volkmann, A. Schmitzler, and H.-J. Freund, "Detection of $n:m$ phase locking from noisy data: Application to magnetoencephalography," *Phys. Rev. Lett.*, vol. 81, no. 15, pp. 3291–3294, 1998.
- [57] L. W. Sheppard, V. Vuksanović, P. V. E. McClintock, and A. Stefanovska, "Oscillatory dynamics of vasoconstriction and vasodilation identified by time-localized phase coherence," *Phys. Med. Biol.*, vol. 56, pp. 3583–3601, 2011.
- [58] L. W. Sheppard, A. Stefanovska, and P. V. E. McClintock, "Testing for time-localised coherence in bivariate data," *Phys. Rev. E*, vol. 85, p. 046205, 2012.
- [59] M. Le Van Quyen, J. Foucher, J. P. Lachaux, E. Rodriguez, A. Lutz, J. Martinerie, and F. J. Varela, "Comparison of Hilbert transform and wavelet methods for the analysis of neuronal synchrony," *J. Neurosci. Methods*, vol. 111, no. 2, pp. 83–98, 2001.
- [60] J. P. Lachaux, A. Lutz, D. Rudrauf, D. Cosmelli, M. Le van Quyen, J. Martinerie, and F. Varela, "Estimating the time-course of coherence between single-trial brain signals: an introduction to wavelet coherence," *Clin. Neurophysiol.*, vol. 32, no. 3, pp. 157–174, 2002.
- [61] A. Bandrivskyy, A. Bernjak, P. V. E. McClintock, and A. Stefanovska, "Wavelet phase coherence analysis: Application to skin temperature and blood flow," *Cardiovascular Engineering*, vol. 4, no. 1, pp. 89–93, 2004.
- [62] M. Paluš, "From nonlinearity to causality: statistical testing and inference of physical mechanisms underlying complex dynamics," *Contemp. Phys.*, vol. 48, no. 6, pp. 307–348, 2007.
- [63] J. Theiler, S. Eubank, A. Longtin, B. Galdrikian, and J. Farmer, "Testing for nonlinearity in time series: the method of surrogate data," *Physica D*, vol. 58, no. 1–4, pp. 77–94, 1992.
- [64] T. Schreiber and A. Schmitz, "Surrogate time series," *Physica D*, vol. 142, no. 3–4, pp. 346–382, 2000.
- [65] T. Schreiber, "Measuring information transfer," *Phys. Rev. Lett.*, vol. 85, no. 2, pp. 461–464, July 2000.
- [66] C. W. J. Granger, "Investigating causal relations by econometric models and cross-spectral methods," *Econometrica*, vol. 37, no. 3, pp. 424–438, 1969.
- [67] L. Barnett, A. B. Barrett, and A. K. Seth, "Granger causality and transfer entropy are equivalent for Gaussian variables," *Phys. Rev. Lett.*, vol. 103, p. 238701, 2009.
- [68] M. Paluš and A. Stefanovska, "Direction of coupling from phases of interacting oscillators: An information-theoretic approach," *Phys. Rev. E*, vol. 67, p. 055201(R), 2003.
- [69] M. Vejmelka and M. Paluš, "Inferring the directionality of coupling with conditional mutual information," *Phys. Rev. E*, vol. 77, no. 2, p. 026214, 2008.
- [70] D. Marinazzo, M. Pellicoro, and S. Stramaglia, "Kernel method for nonlinear Granger causality," *Phys. Rev. Lett.*, vol. 100, p. 144103, 2008.
- [71] M. G. Rosenblum, L. Cimponeriu, A. Bezerianos, A. Patzak, and R. Mrowka, "Identification of coupling direction: Application to cardiorespiratory interaction," *Phys. Rev. E*, vol. 65, no. 4, p. 041909, 2002.
- [72] K. J. Friston, "Bayesian estimation of dynamical systems: An application to fMRI," *NeuroImage*, vol. 16, no. 2, pp. 513–530, 2002.
- [73] V. N. Smelyanskiy, D. G. Luchinsky, A. Stefanovska, and P. V. E. McClintock, "Inference of a nonlinear stochastic model of the cardiorespiratory interaction," *Phys. Rev. Lett.*, vol. 94, no. 9, p. 098101, 2005.
- [74] U. von Toussaint, "Bayesian inference in physics," *Rev. Mod. Phys.*, vol. 83, no. 3, p. 943, 2011.
- [75] T. Stankovski, A. Duggento, P. V. E. McClintock, and A. Stefanovska, "Inference of time-evolving coupled dynamical systems in the presence of noise," *Phys. Rev. Lett.*, vol. 109, p. 024101, 2012.
- [76] A. Duggento, T. Stankovski, P. V. E. McClintock, and A. Stefanovska, "Dynamical Bayesian inference of time-evolving interactions: From a pair of coupled oscillators to networks of oscillators," *Phys. Rev. E*, vol. 86, no. 061126, 2012.
- [77] T. Stankovski, A. Duggento, P. V. E. McClintock, and A. Stefanovska, "A tutorial on time-evolving dynamical bayesian inference," *Eur. Phys. J. Special Topics*, vol. 223, no. 13, pp. 2685–2703, 2014.
- [78] C. Peng, S. V. Buldyrev, S. Havlin, M. Simons, H. E. Stanley, and A. L. Goldberger, "Mosaic organisation of DNA nucleotides," *Phys. Rev. E*, vol. 49, no. 2, pp. 1685–1689, 1994.
- [79] Y. Shiozai, A. Stefanovska, and P. V. E. McClintock, "Nonlinear dynamics of cardiovascular ageing," *Phys. Rep.*, vol. 488, pp. 51–110, 2010.
- [80] V. Rajan, B. Varghese, T. G. van Leeuwen, and W. Steenbergen, "Review of methodological developments in laser Doppler flowmetry," *Lasers Med. Sci.*, vol. 24, pp. 269–283, 2009.
- [81] J. Jamšek, A. Stefanovska, and P. V. E. McClintock, "Nonlinear cardiorespiratory interactions resolved by time-phase bispectral analysis," *Phys. Med. Biol.*, vol. 49, no. 18, pp. 4407–4425, 2004.
- [82] T. Stankovski, V. Ticcinielli, P. V. E. McClintock, and A. Stefanovska, "Coupling functions in networks of oscillators," *New J. Phys.*, vol. 17, no. 3, p. 035002, 2015.
- [83] C. J. Stam *et al.*, "Graph theoretical analysis of magnetoencephalographic functional connectivity in alzheimers disease," *Brain*, vol. 132, pp. 213–224, 2009.
- [84] V. Jirsa and V. Müller, "Cross-frequency coupling in real and virtual brain networks," *Front. Comput. Neurosci.*, vol. 7, p. 78, 2013.
- [85] C. J. Stam, "Nonlinear dynamical analysis of EEG and MEG: Review of an emerging field," *Clin. Neurophysiol.*, vol. 116, pp. 2266–2301, 2005.



approach to these systems.

Philip Clemson received the MPhys degree in Physics with Astrophysics and Cosmology in 2009 and the PhD degree in Nonlinear and Biomedical Physics in 2013, both at Lancaster University, UK. During his PhD studies and subsequent Research Associate position, also at Lancaster, he became well-versed in the spectrum of time series analysis techniques and their application to time-dependent dynamical systems. He also helped to introduce a new class of oscillators known as chronotaxic systems and has developed methods for the inverse



Gemma Lancaster received the MPhys degree in Physics from the University of Salford, UK, in 2011, and has submitted her PhD thesis in Non-linear Biomedical Physics at Lancaster University, UK, in 2015. Previous areas of interest included the non-invasive diagnosis of melanoma based on characterisation of oscillations in physiological data using optimised time-series analysis techniques, an approach which she is now applying to the study of neurovascular coupling and brain function in dementia as a Research Associate at Lancaster University.



Aneta Stefanovska completed her PhD degree in 1992 at the University of Ljubljana, Slovenia, working in part at the University of Stuttgart, Germany. She is currently professor of Biomedical Physics and head of the Nonlinear and Biomedical Physics Group in the Physics Department at Lancaster University, UK. She has extensive experience in studying the physics of living systems combining measurements, time-series analyses, and modelling based on the phase dynamics approach. She has pioneered several new approaches to the analysis and

modelling of time-varying oscillatory dynamical systems with applications to oscillatory cardiovascular, brain and cell dynamics.