



HAL
open science

Uncertainty and information in physiological signals: Explicit physical trade-off with log-normal wavelets

Alexandre Guillet, Françoise Argoul

► **To cite this version:**

Alexandre Guillet, Françoise Argoul. Uncertainty and information in physiological signals: Explicit physical trade-off with log-normal wavelets. *Journal of The Franklin Institute*, 2024, 361 (18), pp.107201. 10.1016/j.jfranklin.2024.107201 . hal-04700580

HAL Id: hal-04700580

<https://hal.science/hal-04700580>

Submitted on 17 Sep 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



Uncertainty and information in physiological signals: Explicit physical trade-off with log-normal wavelets

Alexandre Guillet^{a,b,*}, Françoise Argoul^a

^a Univ. Bordeaux, CNRS, LOMA, UMR 5798, F-33400 Talence, France

^b Max Planck Institute for the Physics of Complex Systems, Dresden, Germany

ARTICLE INFO

Keywords:

Time–frequency analysis
Log-normal wavelet
Uncertainty principle
Polysomnography
Physiological signals
Coherence
Statistical significance
Mutual information
Single trial time–frequency statistics

ABSTRACT

Physiological recordings contain a great deal of information about the underlying dynamics of Life. The practical statistical treatment of these single-trial measurements is often hampered by the inadequacy of overly strong assumptions. Heisenberg's uncertainty principle allows for more parsimony, trading off statistical significance for localization. By decomposing signals into time–frequency atoms and recomposing them into local quadratic estimates, we propose a concise and expressive implementation of these fundamental concepts based on the choice of a geometric paradigm and two physical parameters. Starting from the spectrogram based on two fixed timescales and Gabor's normal window, we then build its scale-invariant analogue, the scalogram based on two quality factors and Grossmann's log-normal wavelet. These canonical estimators provide a minimal and flexible framework for single trial time–frequency statistics, which we apply to polysomnographic signals: EEG representations, HRV extraction from ECG, coherence and mutual information between heart rate and respiration.

1. Introduction

Physiological signals have been extensively studied over the last two decades [1,2], for their complex behaviours, ranging from periodic to stochastic and spanning a wide range of timescales. The formal identification and quantification of distant organ interactions (e.g. between respiration, heart and brain) has since been integrated into the field of “Network Physiology” [3–6]. Pairwise interactions have been approached from different angles [7], such as synchronization, coherence, time delay, mutual information or transfer entropy, and then extended to higher-dimensional structures [6,8]. Such developments require dedicated analytic tools to deal with simultaneous signals from different sources, with non-stationary dynamics spanning a broad range of timescales. The probabilistic approach is a key element to model the variability of these signals, considered as a realization of a stochastic process. The considerable difficulty of conditioning living systems into well-controlled and reproducible initial states limits their repeatability, so that in practice we are often forced to consider a physiological recording as a single trial experiment. The assumption of stationarity, at least locally within well-chosen timescales, is a common simplification for the statistical treatment of such unique signals. However, their study is complicated by the multiple uncertainties associated with the variety of recorded events, in particular their timescales (of relaxation, correlation, evolution) and their possible separation. Time–frequency representation is another key component to investigate complex dynamics [9]. Many types of kernels, windows, filters, wavelets and other atomic shapes have been developed over the years to decipher rich signals. Dynamical information unfolds in the time and frequency domains as more or less local features. Again, several uncertainty relations are available to characterize localization properties [10].

* Corresponding author at: Max Planck Institute for the Physics of Complex Systems, Dresden, Germany.

E-mail address: aguillet@pks.mpg.de (A. Guillet).

<https://doi.org/10.1016/j.jfranklin.2024.107201>

Received 18 December 2023; Received in revised form 23 June 2024; Accepted 20 August 2024

Available online 26 August 2024

0016-0032/© 2024 The Author(s). Published by Elsevier Inc. on behalf of The Franklin Institute. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

The diversity of these techniques and their lack of integration with the statistical approach requires expert choices that we intend to simplify in this paper. With this article, we pay tribute to Alain Arneodo for his pioneering work on the intersection between wavelet theory and physics [11,12]. His late interest in biomedical images [13] has been a great stimulus in initiating the research developed here.

How can we make statistical estimates from individual non-stationary signals without knowing the underlying timescales? Is it possible to manage our uncertainty in terms of statistical significance and localization in time and scales? In this article, we present an elegant framework that succinctly answers these questions. It embeds statistical estimators in a time–frequency or time-scale analysis, to jointly handle the mix of dynamical and statistical information contained in biophysical signals. In particular, we design a flexible and versatile scalogram estimator based on log-normal wavelets, which we apply to a variety of signals. Our uncertainty and assumptions are given a precise shape by the choice of two quality factors, that control both the time–frequency resolution and the statistical significance of the analysis. The classical idea of the uncertainty principle or scale separation, between correlation and evolution timescales, is here revisited in a scale-invariant formalism constructed from time–frequency uncertainty atoms, available as statistical samples in a time–frequency analysis. Exact in the proposed formalism, this straightforward uncertainty relation unlocks the practical management of probabilistic and dynamical uncertainty (or information) in the vicinity of Heisenberg’s uncertainty bound.

This article is divided into seven sections. Section 2 introduces concepts arising from a joint statistical and time–frequency analysis. We discuss, illustrate and formalize the relationships between the assumption of timescale (or quality factor) separation, Heisenberg’s uncertainty principle and the statistical approach to unique signals. The continuous numbering of effective statistical degrees of freedom or time–frequency uncertainty atoms is given an information-theoretic interpretation. Section 3 summarizes these concepts in an explicit parametric framework. Two paradigmatic time–frequency estimators are introduced, one for the spectrogram and the other for the scalogram, called canonical for their elementary expressions in terms of a pair of physical parameters, respectively timescales or quality factors. In particular, their ratio gives the number of uncertainty atoms. The first paradigm is based on the Gabor transform and the normal window, and the second on the wavelet transform and the (Altes–Grossmann) log-normal wavelet. Sections 4, 5 and 6 contain the application of this analytical framework to human physiological recordings of the neural, cardiac and respiratory activity. These non-stationary and multiscale signals are part of the same polysomnography (*shhs2-200901*), downloaded from the National Sleep Research Resource (Sleep Heart Health Study) [14,15] and used after ethics review board approval. In Section 4, the intermittent and noisy neural dynamics captured by an EEG signal illustrates the comparison of the two time–frequency paradigms. Their limiting cases and their physical parameters are discussed. In contrast, the ECG signal introduced in Section 5 contains the highly rhythmic activity of the heart. The extraction of its fluctuating rate is devised within the log-normal wavelet framework; a quadratic estimation of the heart rate signal is formulated and its three parameters are optimized against the standard estimation. In Section 6, we analyse the time–frequency coherence between the resulting heart rate signal and the respiratory oscillations. Its spurious statistics are characterized in terms of the number of atoms and its significance is interpreted in terms of a time–frequency mutual information density. This completes the picture of the trade-off between statistical significance and time–frequency localization in the analysis of single trial recordings. We conclude in Section 7.

2. Single trial statistics in an uncertain dynamical context

In order to obtain valuable statistical information from single trial (unique) measurements of biophysical dynamics, it is often necessary to make assumption about their stationarity and ergodicity. The first moments of their distribution can be estimated at different times if the living system is assumed to be stationary in the short term. However, in the absence of prior dynamical information, the distinction between short and long term processes is uncertain. In discussing this practical issue, we consider two types of temporal scale that appear in the Heisenberg uncertainty relation, between duration and frequency, and the need to separate them in order to create local statistical degrees of freedom. We then present the possibility of designing statistical estimators based on such temporal scale separation, which do not fix any arbitrary timescale parameter. Such a scale-free statistical approach is proposed as a parsimonious way to handle dynamical assumptions in biophysical contexts. Finally, we introduce the formalism associated with non-negative quadratic operators and their time–frequency and statistical interpretations, upon which the framework proposed in the next section is built.

2.1. Statistics and timescales of evolution and correlation

In a statistical approach, the data is divided into samples assumed to have the same probability distribution. Statistical moments can then be estimated from an empirical mean over these samples. For a signal of duration T sampled at the frequency f_s , we can access at most $N = Tf_s$ samples at the cost of the strongest dynamical assumption of stationarity. The ratio of the macro and micro timescales T and $dt = f_s^{-1}$ represents the number of data points N and the maximum possible number of degrees of freedom available for a statistical estimation.

In order to reveal a possible time evolution of the statistics, the system can be assumed to be stationary in the short term, *i.e.* for durations shorter than a meso-timescale τ_0 , and non-stationary for longer durations. This weaker dynamical assumption comes at the cost of a smaller number of statistical samples. Correlated dynamics can also be demonstrated by taking signal blocks of duration τ_0 as samples if the process is stationary. A time delay $\tau < \tau_0$ is introduced into the estimator of the correlation (second order moment), or equivalently, a power spectrum with frequency $f > \tau_0^{-1}$ is estimated (*e.g.* Welch’s method). The spectral domain also facilitates the identification of periodic oscillations in the data.

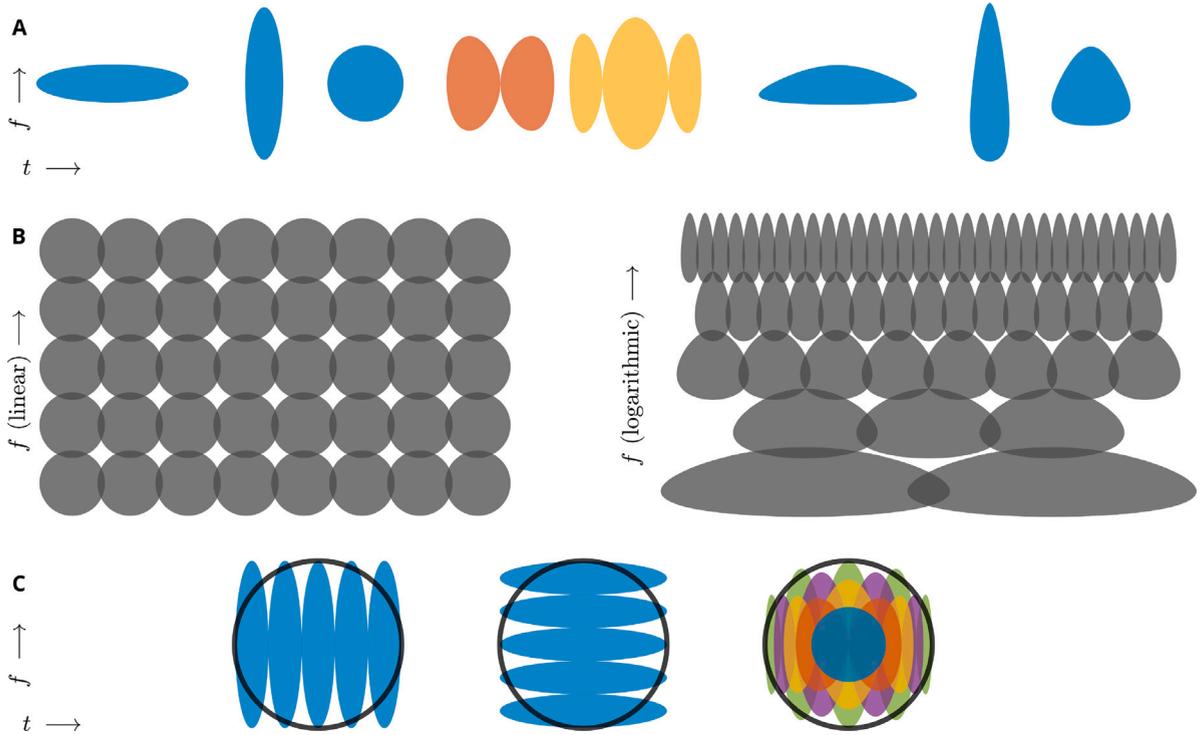


Fig. 1. Schematic representation of information quanta (A) as atoms of time–frequency uncertainty, (B) their even dispositions in the time–frequency plane for Gabor (left) and wavelet (right) transforms, and (C) equivalent ways to compose them to produce a single-trial time–frequency statistical estimate.

The analysis of both the temporal evolution on periods longer than τ_0 and the correlation on intervals shorter than τ_0 have incompatible dynamical requirements. Statistical estimation is indeed performed either within or across these intervals, based respectively on $n = \tau_0/\Delta t$ or on $n = T/\tau_0$ statistical samples. Fortunately, the probabilistic treatment of a signal interpreted as a single realization of an evolving and correlated stochastic process is still possible. This is achieved by further restricting the range of dynamical timescales and reducing the number of statistical degrees of freedom: $n = \tau_+/\tau_- \geq 1$ samples are obtained as signal blocks of duration τ_- , assumed to evolve on timescales longer than τ_+ . Statistical estimates on unique signals with complex dynamics are only possible by separating evolution and correlation timescales.

2.2. Time and frequency uncertainty versus statistics

The smallest distinguishable evolution timescale τ_+ represents a time resolution or uncertainty $\Delta t = \tau_+$ in the analysis, and the longest correlation timescale τ_- corresponds to a spectral bandwidth or frequency resolution $\Delta f = \tau_-^{-1}$. Here, the lower bound of Heisenberg’s uncertainty principle, $\Delta t \Delta f \geq 1$, appears as the condition for a statistical time–frequency analysis of the signal. The case of a unique timescale $\tau_+ = \tau_-$, thus $\Delta t \Delta f = 1$, has been described by Gabor as a decomposition of the signal into information quanta called “logons” [16]. These atoms of time–frequency uncertainty are sketched in Fig. 1. As shown in (A), different aspect-ratios of the atoms correspond to the resolution trade-off between time and frequency. Shapes of order zero (blue) are more compact than shapes of higher orthogonal order (other colours). The even sampling of the time–frequency plane with Gabor’s oscillating atoms (based on a normal window), as sketched in Fig. 1(B, left), endows it with the Euclidean or flat geometry, in contrast with the wavelet approach (B, right).

Fig. 1(C) represents a single-trial statistical estimate in time and frequency based on a local average over several atoms, here $n = 5$, corresponding to both to the number of statistical samples and the area of the resulting time–frequency uncertainty (black circle). Equivalent estimates can be obtained from different compositions and shapes of atoms, from left to right: contiguous along time or frequency, or superimposed orthogonal shapes.

The notion of statistical degrees of freedom has been introduced into the spectral or time–frequency analysis with orthogonal atoms called “tapers”, yielding multi-taper estimators [17–19]. Single-taper spectrograms play the role of statistical samples for the multi-taper estimation of the spectrogram. Schematized in Fig. 1(C, right), the multi-taper approach is implicitly related to the uncertainty relation, since the number of tapers is optimally chosen from the uncertainty product $n = \Delta t \Delta f$. The zeroth order taper (a simple window) gives the best time–frequency resolution, but a single sample is insufficient for a statistical estimation.

The probabilistic approach to evolution and correlation in single complex biophysical recordings relies on the decomposition of the signals into time–frequency atoms. The separation of timescales to collect enough samples, $n = \Delta t \Delta f = \frac{\tau_+}{\tau_-}$, is a sine qua non.

These uncertainty atoms can be further recomposed into a statistical estimator. The existence of such a range of scale separation cannot be checked a priori in living systems, but it appears to be a necessary trade-off between statistical information and dynamical uncertainty. The choice of relevant timescales in the analysis is therefore an important practical issue if we want to measure them. We then discuss how to avoid an arbitrary choice of timescales as control parameters.

2.3. Scale-free time–frequency statistics with constant quality factors

By choosing arbitrarily two timescales for a spectrogram estimation, the information about long-term correlation and short-term evolution may be lost. The analysis of the complex biological systems would benefit from adapting the crucial assumption of local stationarity to any selected timescale.

This goal can be achieved with a scale-independent approach, using an affine time–frequency (or time-scale) representation, also called constant- Q representation such as the wavelet transform and the scalogram. The timescale is treated as a variable (inverse frequency) whose relative precision increases with a parameter Q : the quality factor. A constant Q determines the time–frequency resolution as a fixed relative bandwidth $\Delta f/f \propto Q^{-1}$ and dimensionless duration $f\Delta t \propto Q$ (number of oscillations). Therefore, the resulting analysis has no arbitrary reference timescale. A hyperbolic geometry is known to be associated with the resulting time–frequency representation [20,21], see Fig. 1(B, right). The logarithmic frequency axis improves the flat visualization of wavelets at multiple scales, although their apparent area is no more representative of their atomic (unit) resolution.

Analogous to the multi-taper spectrogram, a multi-wavelet scalogram has been formulated from orthogonal analytic wavelets [18] to control the degrees of freedom for a statistical estimation. In practice, the simpler smoothing approach is performed as an empirical statistical mean [22–24]. This local average in time or frequency provides a clearer control of the time–frequency uncertainty of the analysis.

In this article, we demonstrate that a full statistical interpretation of the scalogram is based on the distinction of two quality factors, $Q_+ = f\Delta t$ controlling durations, and $Q_- = \Delta \log f$ controlling bandwidths. Most importantly, we also identify their ratio $\frac{Q_+}{Q_-} \geq 1$ as the number n of uncertainty atoms or statistical degrees of freedom, playing the role of a scale separation in a scale-free analytic framework.

2.4. Continuous numbering of orthogonal atoms

The consistency of the probabilistic perspective on individual trajectories of complex dynamics derives from the assumption of an incompressible separation between correlation and evolutionary timescales. This scale separation leaves room for degrees of freedom for a statistical estimation, also interpreted as atoms of time–frequency uncertainty. We reformulate these concepts by relying on a result in [19] that relates quadratic Hermitian (bilinear) operators on signals to the spectral multi-taper approach.

Let x denote a signal or a spectrum. The global energy of the data x , possibly divergent for infinite signals, can be written as a scalar product along time or frequency: $x^*x = \int x(t)x^*(t)dt = \int \hat{x}(f)\hat{x}^*(f)df$, see the definition Eq. (B.1) of the Fourier transform $\hat{x}(f)$. The following development is written for one-dimensional data x in order to focus on time and frequency. It also applies to multidimensional data $x = [x_1 \ x_2 \ \dots]$ such as a multichannel simultaneous recording (a set of time series). The complex conjugate x^* is thus interpreted as the conjugate transpose, turning the scalar product x^*x into an energy matrix.

A time–frequency statistical estimator operates on the data x . In particular, quadratic estimators are defined from an operator K that can take the form of a bivariate temporal kernel \check{K} or spectral kernel \hat{K} :

$$x^*K^*x = \iint x(t_1)x^*(t_2)\check{K}^*(t_1, t_2)dt_1dt_2 = \iint \hat{x}(f_1)\hat{x}^*(f_2)\hat{K}^*(f_1, f_2)df_1df_2 \quad (1)$$

The first notation refers to the matrix product, which applies to a discrete time or frequency variable. The following expressions are explicit integrals over the temporal or spectral continuum. If the operator K is non-negative, the estimate can be interpreted as a K -localized energy.

In the discrete and finite context of multi-taper estimators, the kernel is a Hermitian symmetric matrix (chosen to be positive-definite), whose eigenvectors are the multiple orthogonal tapers. They take the statistical interpretation of complex degrees of freedom in the estimation operator K , and their effective number, n , is readily computed from the corresponding (positive) eigenvalues. This computation is equivalently expressed, directly from the kernel, as an effective number of eigenstates (orthogonal atoms), given by a beautiful trace formula [19]:

$$n = \frac{(\text{Tr } K)^2}{\text{Tr } K^*K} = \frac{\left(\int \check{K}(t, t)dt\right)^2}{\iint |\check{K}(t_1, t_2)|^2 dt_1 dt_2} = \frac{\left(\int \hat{K}(f, f)df\right)^2}{\iint |\hat{K}(f_1, f_2)|^2 df_1 df_2} \quad (2)$$

Again, the second and third equal signs clarify the correspondence between the linear algebra notation and integration. The Hermitian symmetry writes $\check{K}^*(t_2, t_1) = \check{K}(t_1, t_2)$ in time, and similarly in frequency. The matrix product and the linear algebraic trace Tr each correspond to an integral.

In a continuous context, the analytical computation of the number of atoms n can be intractable (especially the integrals $\text{Tr } K^*K$), favouring the discrete approach. An important contribution of this work is the continuous formulation of spectrogram and scalogram estimators for which the effective number of statistical degrees of freedom is not only tractable analytically, but also intuitive and simple. Beyond these solvable cases, an accurate and analytical approximation of n can be obtained using Laplace’s method, as showcased in Appendix E.

2.5. Localization uncertainty as time–frequency entropy

The time–frequency interpretation of Eq. (1) is introduced with the Wigner–Ville representation:

$$W_{xx}(t, f) = \int x(t + \frac{\tau}{2})x^*(t - \frac{\tau}{2})e^{-i2\pi f \tau} d\tau = \int \hat{x}(f + \frac{\eta}{2})\hat{x}^*(f - \frac{\eta}{2})e^{i2\pi\eta t} d\eta \tag{3}$$

$$x^*K^*x = \iint W_{xx}(t, f)K^*(t, f)dtdf \tag{4}$$

A Hermitian operator K translates into a real kernel $K(t, f) = K^*(t, f)$, which we choose to be non-negative. Relations to temporal and spectral representations \check{K} and \hat{K} are given in Eqs. (C.1), (C.2), clarifying the choice of these notations. The spectrogram and the scalogram of x can be estimated from the Wigner–Ville representation in terms of certain convolution operators K . In the case of the square of the Gabor or wavelet transform (introduced in the next section), the kernel covers a single uncertainty atom, $n = 1$, which is insufficient for a statistical interpretation. The Wigner–Ville representation itself corresponds to the degenerate case $n = 0$: no smoothing or atom of uncertainty can yield infinite variance and negative values [25,26].

In Wigner–Ville’s quadratic time–frequency representation, the logarithm of the number of atoms:

$$\log n = -\log \frac{\iint K(t, f)^2 dtdf}{\left(\iint K(t, f) dtdf\right)^2} = H_2[K] \tag{5}$$

turns out to be the order 2 Rényi entropy $H_2[K]$ of the non-negative time–frequency kernel. Since an exponential entropy counts an effective number of “states” in the distribution $K(t, f)$ [27], we arrive at the following interpretation: to gain statistical significance, n quanta of information (Gabor’s logon) are averaged in the smoothing operation, but the distinct time–frequency localization of these n atoms of uncertainty is lost.

The Rényi entropy has previously been applied to time–frequency representations of a signal [28–32], as a way of quantifying their concentration. Interestingly, the case of the Rényi entropy of order 2 was rejected because of its unsuitable handling of cross terms. This is irrelevant here since the kernel $K(t, f)$ is non-negative; n can be interpreted as the time–frequency area effectively covered by the kernel.

Although the continuous time–frequency integral expression for statistical degrees of freedom has already appeared in the literature [26], its interpretation and practical usefulness remained unrecognized so that explicit eigendecompositions seemed inevitable. In the following section, we present two time–frequency statistical estimators, whose number of atoms is exactly derived from resolution parameters, shedding a new light on the formulation and the application of the uncertainty principle.

3. Canonical quadratic time–frequency estimators

In this section, we present a concise perspective on the prolific and mature field of continuous time–frequency representations by focusing on their single trial statistics. In what follows, two quadratic estimators in the Cohen and the affine classes of time–frequency representations [33] are revisited: the spectrogram and the scalogram. They are built from a decomposition of the signal into paradigmatic time–frequency atoms, Gabor’s normal oscillatory windows w and Grossmann’s log-normal wavelets ψ . The recomposition of these atoms into the estimator’s kernel is denoted K_w and K_ψ respectively. These quadratic estimators are referred to as *canonical*, each in its own paradigm, for their elementary and exact parameterization of the statistical and time–frequency properties discussed in Section 2 in terms of two essential physical quantities. The time–frequency resolution is controlled by either a pair of timescales (τ_+, τ_-) or a pair of quality factors (Q_+, Q_-), and their scale separation (ratio) corresponds exactly to the effective number n of statistical degrees of freedom, Eqs. (2), (5). Heisenberg’s uncertainty principle is strikingly reformulated as the requirement $n \geq 1$ for a statistical description of single-trial signals, in both the fixed-scale and scale-free approaches.

3.1. Normal window paradigm

We start with the Gaussian or normal window of width (duration) τ :

$$w_\tau(t' - t) = e^{-\pi\left(\frac{t'-t}{\tau}\right)^2} \tag{6}$$

whose spectrum $\hat{w}_\tau = \tau w_{\frac{1}{\tau}}$ has bandwidth $\frac{1}{\tau}$. The short-time Fourier transform (STFT) of a signal x under this normal window decomposes the spectrum \hat{x} in the time domain, by shifting the window to any time t . This STFT is also known as the Gabor transform:

$$G_x(t, f; \tau) = \int x(t')w_\tau(t' - t)e^{-i2\pi f t'} dt' = \int \hat{x}(f')w_{\frac{1}{\tau}}(f' - f)\tau e^{i2\pi(f' - f)t} df' \tag{7}$$

The pure wave multiplied by the normal window constitutes Gabor’s oscillating atom of uncertainty, or logon [16], shifted at any time–frequency location (t, f) . From this decomposition, schematized in Fig. 1(B, left), the spectrogram of a stochastic process x can be defined as a second order moment:

$$S_{xx}(t, f; \tau) = \frac{\sqrt{2}}{\tau} \mathbb{E} [G_x G_x^*] \tag{8}$$

normalized by $\frac{\tau}{\sqrt{2}}$, the energy of the atomic window w_τ . It is interpreted as a time-dependent power spectral density, in other words the time–frequency density of energy in x .

A statistical estimate of the spectrogram is obtained by smoothing the product $G_x G_x^*$, either along time or along frequency, assuming its probability distribution locally invariant. This recomposition of atoms from a local averaging operation again involves a normal window of different duration. The choice of a Gaussian smoothing leads to two equivalent definitions of the canonical spectrogram estimator:

$$S_{xx}(t, f; \tau_+, \tau_-) = \frac{2}{\tau_+ \tau_-} \int G_x(t', f; \tau_-) G_x^*(t', f; \tau_+) w_{\tau_+}^2(t' - t) dt' \quad , \quad \tau_+^2 = \tau_+^2 - \tau_-^2 \quad (9)$$

$$= \frac{2\tau_+}{\tau_+} \int G_x(t, f'; \tau_+) G_x^*(t, f'; \tau_+) w_{\tau_+}^2(f' - f) df' \quad , \quad \tau_+^{-2} = \tau_-^{-2} - \tau_+^{-2} \quad (10)$$

The corresponding compositions of atoms are schematized in Fig. 1(C), the left and middle ones respectively. A time-smoothing extent τ_+ is required to obtain a composition of duration τ_+ from atoms of duration τ_- (similarly for $\frac{1}{\tau_+}$ in the frequency domain). Two different timescales $\tau_+ \geq \tau_-$ are thus introduced in these decomposition and recomposition operations, as summarized in the time–frequency kernel (Wigner–Ville representation) of this estimator $x^* K_w^* x$:

$$K_w(t', f'; t, f) = \frac{2\tau_-}{\tau_+} (w_{\tau_+}(t' - t) w_{\tau_-}(f' - f))^2 \quad (11)$$

This density function provides the exact shape of the composition of atoms in Gabor’s normal window paradigm. See Appendix C.2 for purely temporal or spectral explicit expressions. This averaging operator covers a time–frequency box located at (t, f) with widths $\Delta t = \tau_+$ and $\Delta f = \tau_-^{-1}$, which contains:

$$n = \frac{\tau_+}{\tau_-} \quad (12)$$

uncertainty atoms, or equivalently n statistical degrees of freedom. The scale ordering $\tau_+ \geq \tau_-$ is thus equivalent to Heisenberg’s uncertainty relation, here $n = \Delta t \Delta f \geq 1$.

3.2. Log-normal wavelet paradigm

To change of paradigm, it could be sufficient to let the duration of the atom adapt to the frequency, $\tau = Q|f|^{-1}$, and instead fix a number of oscillations or a quality factor Q . The resulting time–frequency decomposition, proportional to $G_x(t, f; Q|f|^{-1})$, is known as the wavelet transform with Morlet’s Gabor-like atom [34] or the closely related S-transform [35]. The relationships between these conventions are discussed further in Appendix B. However, the previous development would fail at Eq. (10). Moreover, this normal (Gabor–Morlet) wavelet would lack important properties, useful in a constant- Q signal analysis and synthesis: it is neither analytic (only approximately so for $Q \gg 1$), nor admissible (without adjustment).

Instead, like others before us, we propose to use a “particularly good” wavelet in this respect [36], already anticipated in [37]:

$$\hat{\psi}_Q\left(\frac{f'}{f}\right) = e^{-\pi(Q \log \frac{f'}{f})^2} \Theta\left(\frac{f'}{f}\right) \quad , \quad (13)$$

where the Heaviside step function Θ handles the sign of the frequency domain (relevant for complex-valued signals). This is the log-normal (Altes–Grossmann) wavelet with quality factor Q , a symmetric limit case in the 2-parameter Morse family of analytic wavelets [38], also known as the log-Gabor filter (its radial part) in image processing [39–41]. Defined from the frequency domain, this uncertainty atom is scaled at frequency f and shifted at time t once multiplied by $e^{-i2\pi f' t}$. There is no closed-form temporal expression for ψ_Q , but an effective number of oscillations $Q = |f| \Delta t$ is expected [42]. The log-normal and the normal shapes are similar for $Q \gg 1$: in this regime, the Altes–Grossmann wavelet is indistinguishable from the Gabor–Morlet wavelet, but it retains its analyticity and admissibility with a small quality factor. It has other interesting properties, such as invariance with respect to derivation, orthogonality to any polynomial trend in the signal and a faster than polynomial decay in time and in frequency on both sides, which makes it particularly regular and localized [42].

The log-normal wavelet has a relative bandwidth $\Delta \log f = Q^{-1}$ around the peak frequency $f' = f$, so that we define the corresponding wavelet transform as a time–frequency decomposition:

$$X(t, f; Q) = \int \hat{x}(f') \hat{\psi}_Q^*\left(\frac{f'}{f}\right) e^{i2\pi f' t} df' = \int x(t') \psi_Q^*(f(t' - t)) |f| dt' \quad (14)$$

Alternative spectral interpretations of the scale variable differ by a factor $\lambda^\alpha = \exp \frac{\alpha}{2\pi Q^2}$ [42], for instance $\alpha = \frac{3}{2}$ for the instantaneous frequency (phase derivative) of ψ_Q . From the atomic decomposition Eq. (14) schematized in Fig. 1(B, right), the scalogram can be defined as a second order moment:

$$S_{xx}(t, f; Q) |f| = \sqrt{2} Q \mathbb{E} [X X^*] \quad (15)$$

normalized by $(\sqrt{2} Q)^{-1}$, the admissibility coefficient of the wavelet ψ_Q . The scalogram is interpreted as a time-dependent power log-frequency density: the explicit factor $|f|$ is compensated when integrated with respect to $\frac{df}{|f|}$ (equal to $d \log f$ for $f > 0$). This notation as a product $S_{xx} |f|$ conserves the physical unit of the evolutionary spectrum S_{xx} in both paradigms.

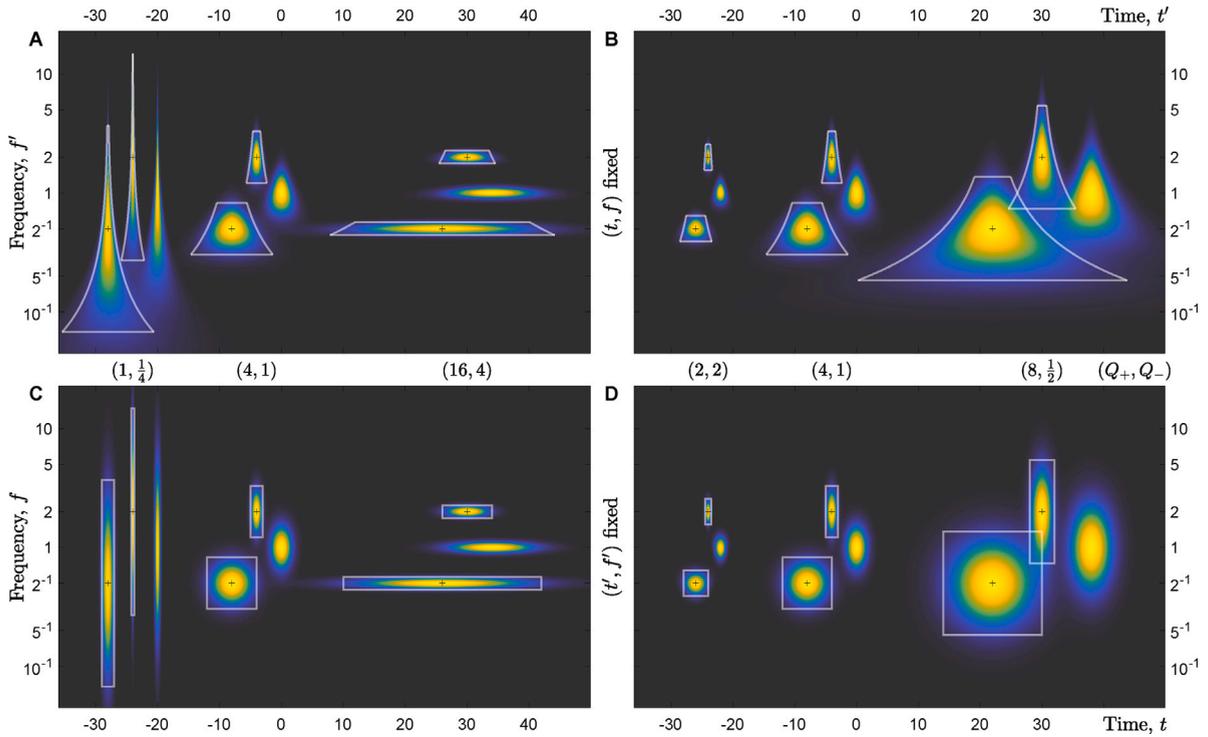


Fig. 2. Log-normal kernels $K_\psi(t', f'; t, f)\Delta t\Delta \log f$ (nondimensionalized): (A, C) composed of a fixed number of atoms $n = \frac{Q_+}{Q_-} = 4$ with different mean quality factors $Q_0 = \sqrt{Q_+Q_-} = \frac{1}{2}, 2, 8$, (B, D) with a fixed $Q_0 = 2$ and increasing number of atoms $n = 1, 4, 16$ (from left to right). Each kernel is represented in 3 copies at distinct time–frequency locations, and both in (A, B) Wigner–Ville’s (t', f') and in (C, D) the wavelet’s (t, f) planes. White boxes, of time–frequency area equal to n , materialize the uncertainty $\Delta t = Q_+|f'|^{-1}$, $\Delta \log f = Q_-^{-1}$.

A statistical estimate of the scalogram is obtained by smoothing the product XX^* . The use of the log-normal wavelet with a different quality factor to compute the local average along log-frequencies (or scales) leads to the canonical scalogram estimator:

$$S_{xx}(t, f; Q_+, Q_-)|f| = 2Q_+Q_- \int X(t, f'; Q_+)X^*(t, f'; Q_+)\hat{\psi}_{Q_+}^2\left(\frac{f'}{f}\right)\frac{df'}{|f'|} , \quad Q_+^{-2} = Q_-^{-2} - Q_+^{-2} . \quad (16)$$

This smoothing is formally analogous to Eq. (9), where time is replaced by log-frequency, but conceptually similar to Eq. (10) and Fig. 1(C, middle), where the recomposition step is performed along the vertical (spectral) direction. The change in (relative) bandwidth Q_+^{-1} is the amount of smoothing required to achieve a bandwidth Q_-^{-1} , for a composition of atoms with bandwidth Q_+^{-1} . Again, the two different quality factors $Q_+ \geq Q_-$ introduced here are visible in the kernel of the canonical scalogram estimator, when written in the spectral domain:

$$\hat{K}_\psi(f_1, f_2; t, f) = \sqrt{2}Q_-\left(\hat{\psi}_{Q_+}\left(\frac{f_1}{f_2}\right)\hat{\psi}_{Q_-}\left(\frac{f_1f_2}{f^2}\right)\right)^{\frac{1}{2}}\theta\left(\frac{f_1}{f}\right)\theta\left(\frac{f_2}{f}\right)e^{-t2\pi(f_1-f_2)t} . \quad (17)$$

The time and time–frequency kernels for this operator are intractable integral forms, see Appendix C.1, as for the log-normal wavelet ψ_Q in the time domain. Nevertheless, we compute $K_\psi(t', f'; t, f)$ numerically from Eqs. (17), (C.1) to confirm the extension of the time–frequency uncertainty, and note its non-negativity (up to numerical precision). The kernels are represented for different quality factors in Fig. 2 as dimensionless functions $K_\psi(t', f'; t, f)\Delta t\Delta \log f$. In contrast to Gabor’s window paradigm (Eq. (11)), they have distinct shapes in Wigner–Ville’s time–frequency plane (t', f') (A, B) and in the wavelet’s (target) time–frequency plane (t, f) (C, D). They represent the area of the time–frequency uncertainty, well summarized by the widths $\Delta t = Q_+|f'|^{-1}$ and $\Delta \log f = Q_-^{-1}$ (white boxes), thus confirming the interpretation of the quality factor, even for small values. The case $Q_+ = Q_-$ corresponds to the Wigner–Ville distribution of the log-normal wavelet, which is a quadratic time–frequency representation of the uncertainty atom.

The kernel of the canonical scalogram estimator contains a number of uncertainty atoms:

$$n = \frac{Q_+}{Q_-} , \quad (18)$$

which corresponds precisely to the time–frequency area of the illustrated uncertainty boxes, fixed in Fig. 2(A, C) or increasing from $n = 1$ (single atoms, left) to $n = 16$ (right) in Fig. 2(B, D). For a uniform visualization of these various time–frequency densities on linear-logarithmic axes, we have plotted their dimensionless form obtained by multiplying K_ψ by $\Delta t\Delta \log f = n|f'|^{-1}$.

As before, there is also a temporal smoothing derived from the log-normal wavelet that corresponds to the horizontal composition of atoms schematized in Fig. 1(C, left). Instead of acting on a single time variable in the product of wavelet transforms, it acts on

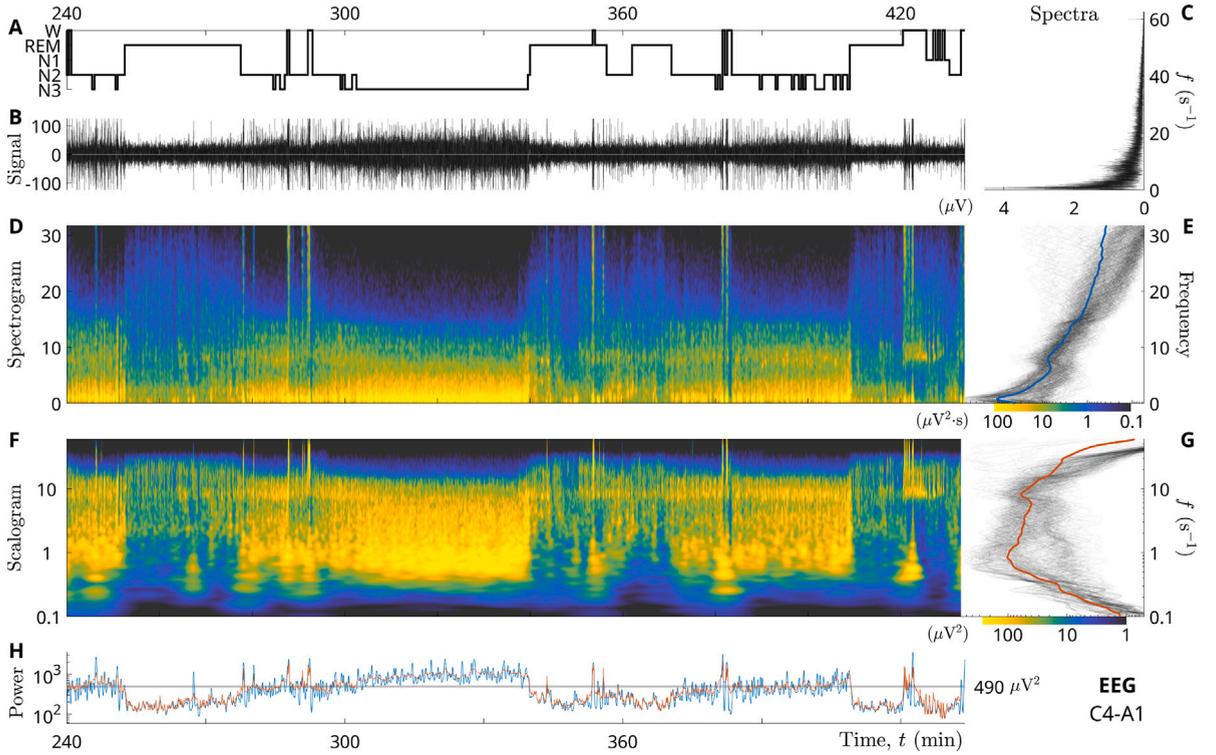


Fig. 3. Electroencephalogram (EEG) during sleep. (A) Hypnogram: EEG scoring of the sleep stages “wake”, “rapid eye movement” (REM) and “non-REM” stages 1, 2 and 3 (light to deep sleep). (B) EEG signal (μV) from electrodes C4-A1. (C) Fourier amplitude spectrum $|\hat{x}(f)|/T$. (D) Spectrogram $S_{xx}(t, f; \tau_+, \tau_-)$ with $(\tau_+, \tau_-) = (20, 1)$ s, and (E) power spectral density (blue). (F) Scalogram $S_{xx}(t, f; Q_+, Q_-)|f|$ with $(Q_+, Q_-) = (100, 5)$, and (G) power log-frequency density (red). The colour bars for (D, F), between quantiles [0.1, 0.99], are aligned with the corresponding axes in (E, G) and persistence spectra are shown (grey). (H) Time-dependent power estimated from (D, blue) and (F, red), and mean power (grey line). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

two time variables:

$$S_{xx}(t, f; Q_+, Q_-)|f| = 2Q_+Q_- \iint X(t_1, f; Q_-)X^*(t_2, f; Q_-)\chi_{Q_+}^*(t_1 - t, t_2 - t)dt_1dt_2$$

$$\text{where } \chi_{Q_+}(t, t') = \int \psi_{Q_+}(f't)\psi_{Q_+}^*(f't')|f'|df', \quad Q_+^2 = Q_+^2 - Q_-^2. \quad (19)$$

This remarkable expression is specific to the log-normal wavelet paradigm.

Finally, the convenient intuition of the time–frequency statistics in the normal window paradigm is largely preserved in this log-normal wavelet paradigm. The separation of the correlation scale (frequency) and the evolution scale (duration) is achieved for each frequency of interest f by the distinction $Q_+ \geq Q_-$, which corresponds to an uncertainty relation $n = |f|\Delta t \Delta \log f \geq 1$, allowing a statistical interpretation of the scalogram estimate.

4. Visualizing an electroencephalogram

During sleep, the physiological interaction between organs is orchestrated by the autonomic nervous system; the EEG recording of brain cortex potential fluctuations provides valuable information about these interactions, their regulation or their imbalance [4–6]. Polysomnographic signals (including EEG signals) are now used in routine clinical practice to identify not only sleep disorders but also their interaction with other metabolic, organic or cognitive functions [43].

An EEG recording provides a rich and natural illustration for our discussion. Time–frequency statistics are particularly well suited to the analysis of such “unique” (in the sense of strictly non-reproducible) signals, which intermittently switch from highly stochastic (scale invariant) to more regular dynamics during the night. Canonical estimates of the EEG spectrogram and scalogram are used to illustrate the brain activity. A fair comparison of these two paradigms of continuous atomic decomposition and recomposition is given in terms of their resolution, highlighting their strengths and limitations.

The EEG shown in Fig. 3 is part of the polysomnography *shhs2-200901* from the Sleep Heart Health Study [14,15] that were recorded on a 54-year-old sleeping woman. The cardiac (ECG) and respiratory signals from this polysomnography are also presented in the next figures, in the same selected time interval: from the fourth hour after sleep onset to awakening. Further characteristics of these electrophysiological signals are detailed in Appendix A.

4.1. Single trial time–frequency statistics without orthogonal atoms

Time–frequency statistical estimators such as the spectrogram and the scalogram are ideal tools for approaching and deciphering the complex mixture of ever-changing fluctuations and oscillations that make up the brain activity. A time–frequency statistical analysis of the EEG based on the multi-taper spectrogram has already been advocated [44]. We intend to refine this perspective, starting with a criticism: the explicit use of orthogonal atoms such as multi-tapers is unnecessary and do not provide an explicit control of the time–frequency resolution. The continuous numbering of statistical degrees of freedom Eq. (2) is straightforward for any (non-negative) quadratic operator, and those we introduce in this article are explicitly parameterized in terms of time and frequency resolutions. Some general considerations are worth mentioning when choosing these analytical tools. As we have already pointed out in previous sections, this choice builds on compromises.

On the one hand, it is generally accepted that the best quadratic time–frequency representations are those with the highest resolution. On the other hand, an averaging based on orthogonal windows (tapers) is often pushed forward as optimal for estimating statistics from these representations. This implies a systematic trade-off between time–frequency resolution and the number of statistical degrees of freedom. The discrete orthogonal approach explicitly details the shape of each independent atom, but the resulting resolution is implicit. The converse is true for the continuous (smoothing) approach: the resolution is explicitly parameterized, but the exact shape of each uncertainty atom is implicit. We further argue that these shapes (see Fig. 1C) are contingent since the effective number of atoms is known exactly.

Indeed, for the same resulting composition of atoms, alternative continuous decompositions (into windows or wavelets) and recompositions (into a statistical estimator) are available: see Eqs. (9), (10) for the spectrogram and Eqs. (16), (19) for the scalogram. Thus, the duration of a window can be thought of as belonging to $[\tau_-, \tau_+]$, or the number of wavelet oscillations to $[Q_-, Q_+]$, but only the final composition of atoms matters (characterized by the bounds of these intervals). Moreover, we have shown that both the resulting time–frequency resolution and the effective number of statistical degrees of freedom can be parameterized exactly as two timescales or quality factors and their ratio respectively.

Therefore, the spectrogram of the EEG $S_{xx}(t, f; \tau_+, \tau_-)$ in Fig. 3(D) provides the same quality of time–frequency statistical estimation as the multi-taper spectrogram [44], without relying on orthogonal windows and with an explicit control of the resolution: here $\tau_+ = 20 \text{ s} = \Delta t$ and $\tau_- = 1 \text{ s} = \Delta f^{-1}$. The choice between the spectrogram and the scalogram is discussed next.

4.2. Resolution in the spectrogram and in the scalogram

The choice of a time–frequency paradigm is another critical aspect of the EEG analysis. In Fig. 3(D, F), we try to compare the spectrogram (D) and the scalogram (F) as fairly as possible: both estimates are composed of $n = 20$ uncertainty atoms and the spectral ranges are chosen to have the same apparent resolution (along their respective linear and logarithmic axes). The colour scale is further bounded by magnitudes chosen at the same quantiles: 10% and 99% of the magnitude distribution in the images. The use of a perceptual colour coding [45] (on a logarithmic axis) is also an important aspect to prevent visual artefacts from hindering a quantitative assessment of the EEG power densities. The flat representation of the time–frequency plane, adapted to the geometry of the spectrogram, constitutes a drawback for the scalogram, which can be circumvented by zooming on shorter time ranges to examine higher frequencies.

The scalogram $S_{xx}(t, f; Q_+, Q_-)|f|$ in Fig. 3(F) is similar to the spectrogram in Fig. 3(D) once the factor $|f|$ is omitted: $S_{xx}(t, f; Q_+, Q_-) \sim S_{xx}(t, f; \tau_+, \tau_-)$, up to their different time–frequency resolution. The spectrogram and the scalogram are time–frequency densities of the signal energy, and the factor $|f|$ appears for the scalogram to keep this interpretation on a log-frequency axis.

The frequency dependence of the time resolution in the scalogram is well visible for impulsive events. The large upper quality factor chosen in Fig. 3(F), $Q_+ = 100$, is not the most suitable to visualize single oscillations (as it averages over 100 of them), especially at low frequencies where single events are spread over several minutes. Nevertheless, the time resolution at high frequencies is greater in the scalogram (F) than in the spectrogram (D): at $f = 20 \text{ s}^{-1}$, singular events spread over $\Delta t = 5 \text{ s}$ in the former compared to $\Delta t = 20 \text{ s}$ in the latter. The inversion of the temporal resolution between the scalogram and the spectrogram occurs at the frequency $Q_+/\tau_+ = 5 \text{ s}^{-1}$, and the same inversion occurs for the spectral resolution, $Q_-/\tau_- = 5 \text{ s}^{-1}$.

Focusing now on the spectral resolution, the spectrogram representation appears to be a natural extension of the traditional discretization of the EEG into (linearly spaced) frequency bands, approximately every 4 s^{-1} [44]. But the existence of a neural activity on multiple timescales [46], supported by the similarity between the EEG and a pink noise (self-similar process), clearly points to the scale-invariant wavelet paradigm and its scalogram. Indeed, the observation of the EEG fluctuations shows that the spectral range is only limited by the instrumental characteristics, especially at low frequencies, where the effect of the high-pass filter is particularly marked in Fig. 3(F, G). This limitation is hidden in the spectrogram, since the spectral range below its frequency resolution $\Delta f = 1 \text{ s}^{-1}$ is not resolved.

The brain waves in the bands β - γ (fast) and δ (slow) are particularly irregular compared to the intermediate ones (θ and especially α and σ), which are more rhythmic. The characteristics of bands and their relation to sleep stages are given in Appendix A.1. In the scalogram, the irregularity of brain waves can be observed independently of their frequency as their relative bandwidth, visible as their spread in log-frequency. We cannot see this aspect from the spectrogram, which relies heavily on its fixed timescales, in particular $\tau_- = \Delta f^{-1}$. The difference is striking when comparing the extension of weaker high-frequency activity with the narrowing of the stronger low-frequency activity in the spectrogram. This picture is more balanced in the scalogram.

The frequency-localized structures in the intermediate band (between 8 and 16 s^{-1}) are composed of σ waves (sleep spindles, in the non-REM stage) and α waves (strongest in the wake stage). In Fig. 3, the spectrogram benefits from a better spectral resolution in this band than the scalogram. The latter can, at best, distinguish between frequencies 9 and 11 s^{-1} with the quality factor $Q_- = 5$.

Our conclusion is nuanced: the scalogram appears to be the preferred representation for a broad perspective on the diversity of timescales and on the regularity of the recorded oscillations, whereas the spectrogram is useful for focusing on neural phenomena near the reference timescales.

4.3. Temporal and spectral densities of the EEG power

We finally comment on the representation of the EEG as simpler temporal, spectral, fully global or local estimates, by following the description of marginal and limit cases of the canonical spectrogram and scalogram estimators detailed in Appendix C.3.

Spectral estimates are shown in Fig. 3(E, G). The EEG spectrogram (D) is aligned with its power spectral density $S_{xx}(f; \tau_-)$ (E) and the EEG scalogram (F, blue) with its power log-frequency density $S_{xx}(f; Q_-)|f|$ (G, red). Persistence spectra are also shown (time-sampled spectrogram and scalogram), giving a glimpse of the temporal variability associated with sleep stages and more room to explore higher order moments, or other statistical quantities such as quantiles (used for colour-coding). The $1/f$ trend of the EEG, visible in the spectrogram and in the power spectral density, is compensated in the scalogram and in the power log-frequency density. Switching to the wavelet scale-invariant paradigm therefore provides a better contrast to EEG time–frequency representations. For instance, the narrowband activity of the $\alpha - \sigma$ waves is clearer in Fig. 3(G) than in (E), in spite of a lower resolution: $\Delta f \approx 2$ to 3 s^{-1} versus 1 s^{-1} .

Temporal estimates of the local EEG power or variance are compared in Fig. 3(H). Contrary to the Gabor estimate of the time-dependent power $P_{xx}(t; \tau_+)$ (blue), the wavelet estimate $P_{xx}(t; Q_+)$ (red) has an unspecified resolution that varies with the timescale of the fluctuation under consideration; see definition in Eq. (C.6). Both estimates may differ locally, but their trend is very similar. The switch between NREM and REM (see the hypnogram, A) has a clear effect on the local variance, which can drop from 1000 μV^2 during deep sleep (N3 stage) down to 100 μV^2 during REM.

At maximal statistics, no time–frequency resolution is left: the mean power (global variance) of the EEG in Fig. 3 is about 500 μV^2 (grey line in H), related to the electrical power by an impedance assumed constant or slowly varying.

At maximal time–frequency resolution, no room is left for a statistical estimation. Besides the degenerate case $n = 0$ of the Wigner–Ville distribution, $n = 1$ leads to the common conception of the spectrogram and the scalogram as squares of the corresponding time–frequency decomposition. On top of the amplitude, the phase information in linear representations can be of interest for studying the micro-architecture of sleep and its rhythmic brain waves, such as sleep spindles (σ) and slow oscillations (δ) in deep sleep (N3), or α waves. A quadratic statistical estimator with a few degrees of freedom can also enhance the analysis of rhythmic behaviours, as applied to the cardiac activity in the next section.

5. Quadratic decoding of the heart rate fluctuations

As the central conductor of the dual circulatory system, the contractile activity of the heart adapts in real time to the needs of the entire organ system and is finely regulated by the sympathetic and parasympathetic branches of the autonomic nervous system. The heart rate variability (HRV) has been proposed as a health resilience marker more than three decades ago [47]; HRV is still the subject of a long-lasting interest [48,49]. The further improvement and comparison of HRV identification methods remains an important issue. Non-invasive techniques have been developed to monitor the rhythmic activity of the heart, such as electrocardiography, photoplethysmography, and mechanocardiography.

Here, we focus on the electrocardiogram (ECG) from the same polysomnography as before. The standard discrete estimation of the heart rate from inter-beat time intervals is reinterpreted within the continuous log-normal wavelet framework. After an ECG-specific preprocessing step illustrated in Fig. 4, we then take a statistical perspective to propose a stable rate signal estimator derived from the phase of the cardiac oscillation. This quadratic method is arguably the simplest non-linear estimation of the heart rate signal. We discuss its advantages and limitations compared to existing non-linear methods.

5.1. Wavelet-based preprocessing of the ECG

The electrical activity of the heart, as recorded by the ECG, is shown in Fig. 4. The sharp and prominent oscillations in this rhythm, called QRS complexes, are a clear marker of a heartbeat, visible in the time domain, either in the signal (A, zoom in E top) or in time–frequency representations (wavelet transforms B and C), but not in the spectral domain (D). QRS refers to a sequence of 3 local extrema in the ECG waveform, in the time domain, including the sharpest peak. Its detection (red circles in E, F) is routinely used to extract the heart rate (black dots in B, C, E, F) from the time intervals between beats, as described in Appendix A.2. Due to the frequent polarity reversal of the ECG signal (depending on the placement of the electrodes on the chest), a proxy for the QRS amplitude is often preferred to the actual voltage peak, after band-pass filtering baseline drift and high frequency noise.

This ECG-specific preprocessing is simply the modulus of the wavelet transform:

$$y(t) = |X(t, f_{\text{QRS}}; Q_{\text{QRS}})| \quad (20)$$

at the characteristic scale of the QRS complex. Aiming at “softening” the cardiac oscillation, this step can be skipped for other types of cardiac recording. Note that $y(t)$ is specified by both a frequency location f_{QRS} and a resolution: the quality factor Q_{QRS} .

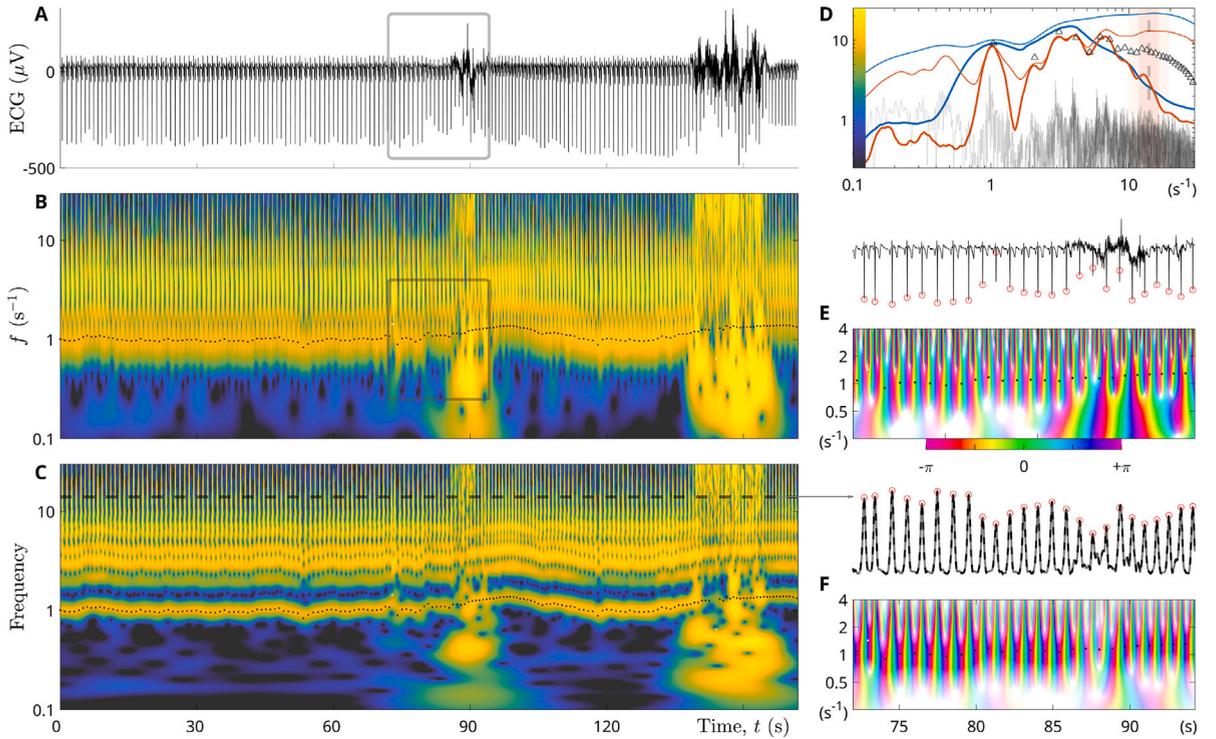


Fig. 4. (A) ECG signal $x(t)$ (μV). (B) $|X(t, f; Q)|$ with $Q = 1$. (C) $|X(t, f; Q_{\text{qrs}})|$ with $Q_{\text{qrs}} = 3$. The colour scale (μV) is aligned with (D). (D) Amplitude spectra: $|\hat{x}(f)|/T$ over one RR interval (black triangles) or the selected time interval (grey line), median (thick lines), and root-mean-square (thin lines) over time of (B, blue) and (C, red). (E) $x(t)$ and $X(t, f; Q)$ (zoom in the boxes in A, B), with modulus-angle colour coding near the heart rate. (F) Amplitude of the QRS complexes $y(t) = |X(t, f_{\text{qrs}}; Q_{\text{qrs}})|$ at $f_{\text{qrs}} = 14 \text{ s}^{-1}$ (grey dashed line in C, D) and $Y(t, f; Q)$ with $Q = 1$. Heart rate (dots in B, C, E, F) estimated from detected QRS complexes (red circles in E, F). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

The modulus of two wavelet transforms with different quality factors $Q = 1$ and $Q_{\text{qrs}} = 3$ are shown in Fig. 4(B and C) in the time–frequency domain, and a time slice of the latter is shown at the selected frequency $f_{\text{qrs}} = 14 \text{ s}^{-1}$ (F top). The non-negative signal $y(t)$ represents the QRS amplitude, whose local maxima are a key element in the detection of the QRS oscillations and the estimation of the heart rate (inverse of the RR time intervals). The standard discrete estimation described in Appendix A.2 is shown on all time–frequency representations in Fig. 4(B, C): it clearly follows the frequency modulations of the fundamental mode of the ECG.

The angular components of the wavelet transforms X (ECG) and Y (QRS amplitude) are compared in Fig. 4(E and F bottom). It can be seen that the cardiac cycle, as detected by the discrete reference method, follows the angle at the cardiac frequency. This is particularly robust in the phase of the QRS amplitude (F) compared to the ECG (E), which can also encode occasional oscillations of the baseline (motion artefacts). This suggests an alternative and continuous description of the heart rate, as the derivative of this cycling heart phase.

5.2. Fluctuating rates: from the analytical signal to a time–frequency statistical estimation

Given a single-channel oscillatory signal x , its amplitude and frequency modulations (AM and FM) can be measured from the modulus and angle of the analytic version of the signal, x^+ , defined in Eq. (B.11). A complex rate can be obtained from it as a time derivative (D_t):

$$F_x(t) = \frac{1}{i2\pi} D_t \log x^+(t) \quad , \quad (21)$$

whose real part estimates the fluctuating frequency of the oscillation (the imaginary part is a magnitude rate).

In practice, the resulting modulated rate signal is corrupted by random noise and non-linearities (non-circular waveforms), which can be mitigated by band-pass filtering the original signal. This is precisely the purpose of a time–frequency decomposition such as the wavelet transform: the band is determined by a fixed frequency f and its bandwidth, absolute (Δf) or relative ($\Delta \log f$). The log-normal wavelet representation $X(t, f; Q)$ has the advantage of being an analytic signal for any fixed f , $x^+(t)$ being its limit $Q \rightarrow 0$ as well as its marginal case, as detailed in Eq. (B.12). For this reason, the complex rate $F_x(t)$ can be considered as the limit $Q \rightarrow 0$ of the more general partial time derivative (∂_t):

$$F_x(t, f; Q) = \frac{1}{i2\pi} \partial_t \log X(t, f; Q) = \lambda^{\frac{1}{2}} f \frac{X(t, \lambda f; Q)}{X(t, f; Q)} \quad , \quad \lambda = \exp \frac{1}{2\pi Q^2} \quad . \quad (22)$$

The second equality highlights the practical way to derive with respect to time, as a slight frequency scaling specific to the log-normal wavelet framework. Note that the limit $Q \rightarrow \infty$ corresponds to $F_x \rightarrow f$ as expected. $F_x(t)$ inherits the time–frequency resolution of the wavelet transform. Here, the quality factor counts the number of oscillations of period $|f|^{-1}$ used to estimate $F_x(t)$.

If the oscillatory signal x has a modulated frequency $F(t)$, then the real part $\Re\{F_x(t, f; Q)\}$ can be used to estimate this varying frequency as a phase derivative, provided that the frequency band selected with f and Q contains only this mode. The imaginary part captures the modulation rate of the magnitude of the oscillation.

Considering the scale continuum f at a fixed time t , there exist basins of attraction that yield the same frequency estimate. This “squeezing” phenomenon can be used to reassign time–frequency atoms (or their intensity) in the same basin of attraction to their estimated frequency: the real part of Eq. (22). This process is known as frequency reassignment for quadratic evolutionary spectra or as synchrosqueezed transform [50,51] for time–frequency decompositions such as the wavelet and Gabor transforms.

In practice, we intend to use $F_x(t, f; Q)$ as a signal, after selecting a convenient band (f, Q) . For natural recordings x , however, this complex rate both vanishes and diverges at certain locations related to zeros (phase defects) of the time–frequency decomposition. This singular behaviour has been described in terms of a determinantal point process in the case of the white Gaussian noise [21]. Their perturbation on reassigned or synchrosqueezed representations is mitigated by an additional parameter: these regions with modulus below a certain threshold are excluded [51,52].

We propose here to average the complex rate Eq. (22) over more than one time–frequency atom, in order to preserve the continuity of the estimator in situations with a noisy oscillator, quite common in natural signals. Instead of applying a local mean directly to $F_x(t, f; Q)$, we introduce a quadratic weight $|X(t, f; Q)|^2$ to favour higher amplitudes, and we normalize it with the average local intensity, which is the scalogram estimator. This leads to the generalized definition of the complex rate as a statistical estimator, involving both quality factors (Q_+, Q_-) , the signal $x(t)$ and its derivative $\dot{x}(t) = D_t x(t)$:

$$F_x(t, f; Q_+, Q_-) = \frac{S_{\dot{x}\dot{x}}(t, f; Q_+, Q_-)}{i2\pi S_{xx}(t, f; Q_+, Q_-)} \tag{23}$$

The newly introduced control parameter corresponds to the number of atoms, $\frac{Q_+}{Q_-} \geq 1$, taken as statistical samples. The statistical estimator $F_x(t, f; Q_+, Q_-)$ is free from the poles and zeros of $F_x(t, f; Q)$, thus stabilizing the complex rate estimate. A multi-taper reassignment procedure has been proposed in [53]; the frequency reassignment of the canonical scalogram estimator based on Eq. (23) is expected to achieve a similar goal of improved concentration and stability.

The broadband limit $Q_- \rightarrow 0$, which can be denoted $F_x(t; Q_+)$, was introduced in [54] as a more stable alternative to the $F_x(t)$ limit (reached for $Q_+ = Q_- \rightarrow 0$). When several distinct modes are contained within a bandwidth, for example with frequencies f_1, f_2 satisfying $Q_+^{-1} < \log \frac{f_2}{f_1} < Q_-^{-1}$, then Eq. (22) behaves as the intensity-weighted version of the arithmetic mean $\frac{f_1+f_2}{2}$. For indistinguishable modes such as $\log \frac{f_2}{f_1} < Q_-^{-1}$, the rate estimation is affected by the beating phenomenon. Alternative rate estimates can be developed based on geometric or harmonic averages; they all agree for a sufficiently large Q_- (narrow band). Finally, the degree of squeezing can be expressed in terms of the coherence between the rhythmic signal and its derivative [42]. Such a measure of the accuracy of rate estimation could take over the role of the signal quality index used in practice.

We believe that the explicit balance between the dynamical uncertainty and the number of statistical degrees of freedom in this type of rate estimator makes it quite generic and flexible. It can introduce a statistical perspective into more sophisticated non-linear methods of rate estimation [55–58]. In the following, we apply it to the case of the heart rate in the most simple and direct way, which consists in fixing the band frequency f .

5.3. Statistical estimation of the heart rate signal from the QRS amplitude

We now estimate the complex rate Eq. (23) of the selected cardiac signal: the QRS amplitude $y(t)$ extracted from the ECG, plotted in Fig. 5(A top). The estimation can be applied to other recordings of the cardiac oscillation such as the photoplethysmogram, or directly to the raw ECG signal, although with inferior results.

In this statistical estimation of the rate signal, the squeezing phenomenon is also observed at the cardiac frequency in Fig. 5(B). In a neighbouring range of the order of the bandwidth Q^{-1} , the wavelet transform has a nearly constant angle, as can be observed in Fig. 4(E, F), so that the real part of the complex rate is also nearly constant in this band. We use this flexibility to fix the frequency variable in this range at $f_1 = 1 \text{ s}^{-1}$ for any time, as the simplest estimate of the heart rate.

The pair of quality factors remains to be determined: we want a stable estimator (non-vanishing, non-diverging, hence $Q_- < Q_+$) as close as possible to achieve a time resolution of one heartbeat ($Q_+ \sim 1$). The bandwidth Q_-^{-1} should be chosen wide enough to contain the heart rate variations, but narrow enough to separate the fundamental mode from other influences. In particular, the presence of a second harmonic mode implies the constraint $Q_-^{-1} < \log 2$, or $Q_- > 1.5$ using Eq. (B.13), to avoid its interference (beating).

Although the choice of the quality factors is already quite limited for this application, we propose to optimize their values by minimizing the mean squared relative error of the estimated heart rate as compared to the reference discrete estimator. In fact, there are two ways of doing this: either we interpolate the reference discrete heart rate as a continuous variable, or we discretize the continuous heart rate signal at discrete times (midpoints of the RR intervals). In both cases, the error is about the same: 3% on average over the selected time interval in Fig. 5 this signal, or even 1% if we exclude the time intervals for which the reference estimation is considered unreliable (using NN intervals). A noticeable difference is the observation of a bias (about -0.2%) with the discrete method as compared to no bias with the continuous method. This is caused by a slight oscillation in the heart rate signal at

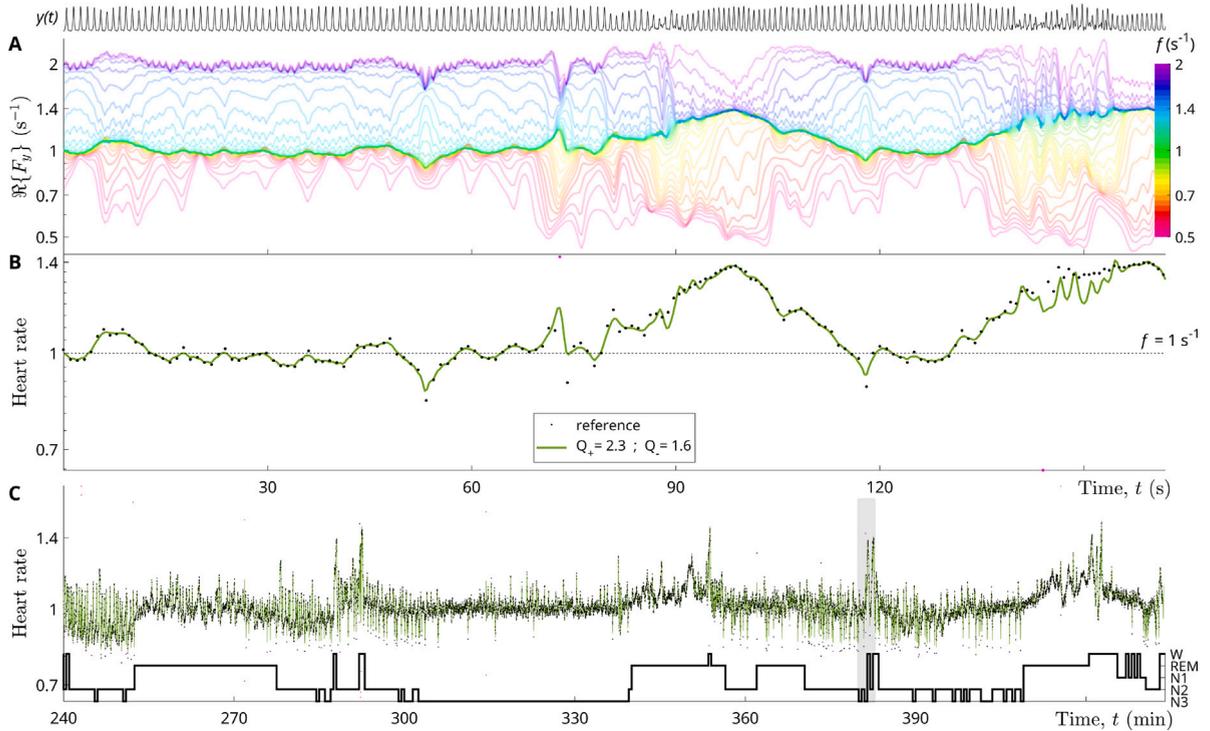


Fig. 5. Estimation of the heart rate from the real part of the complex rate Eq. (23) of the heart signal $y(t)$ (QRS amplitude, top). (A) Rate signals $F_j(t) = \Re\{F_y(t, f_j; Q_+, Q_-)\}$ for fixed frequencies $f_{j=1,2,\dots}$, see colour bar, with $(Q_+, Q_-) = (2.3, 1.6)$. (B) Continuous heart rate estimate for fixed $f = 1.0 \text{ s}^{-1}$ (green line) and discrete reference estimation from RR intervals (dots). (C) Heart rate estimates and hypnogram (bottom); the time interval selected in (A, B) and in Fig. 4 is highlighted in grey. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

the cardiac frequency, interpreted as a residual beat from unresolved higher harmonic orders. This phenomenon is amplified when decreasing Q_- , especially for the broadband marginal considered in [54].

The error is minimal for all these different methods near $Q_+ = 2.3$ and $Q_- = 1.6$, which corresponds to a mean quality factor $Q_0 = \sqrt{Q_+ Q_-} \approx 1.9$ and an effective number of statistical degrees of freedom $n \approx 1.4$. The frequency variable f_1 and the band $(f_{\text{qrs}}, Q_{\text{qrs}})$ defining the signal $y(t)$ from the ECG recording $x(t)$ can also be optimally selected. The error is very sensitive to f_1 , nearly optimal at 1.0 s^{-1} , but less sensitive to the QRS band, fixed to $f_{\text{qrs}} = 14 \text{ s}^{-1}$ and $Q_{\text{qrs}} = 3$ as in Fig. 4(C, F).

5.4. Discussion of the resulting heart rate estimation

Fig. 5(C) illustrates the evolution of the heart rate variability (HRV) with sleep stages. Non-REM sleep is characterized by a noticeably lower rate, with particularly small fluctuations during deep sleep (N3). The heart rate and its variability increase during REM sleep, with bursts during brief awakenings. Compared to the reference discrete estimation, both the time resolution and the value of the estimated heart rate prove to be accurate. This HRV signal is therefore used in the next section to compare cardiac and respiratory activity.

However, the accuracy of the simplest quadratic estimation of the heart rate signal, based on a fixed scale f , a mean quality factor Q_0 and an effective number of atoms n , decreases for large or sudden heart rate modulation events. This limitation suggests a refinement from fixed f to variable $f(t)$ as for a ridge detection approach [55,59]. From an iterative procedure based on the self-consistent relation $f(t) \approx \Re\{F_y(t, f(t))\}$, the heart rate signal estimate can be refined. The fixed scale $f = 1 \text{ s}^{-1}$ would thus only remain as an initial guess. Remarkably, no such extension was necessary to demonstrate the usefulness of the canonical time–frequency statistical estimator in this context.

The rate estimation over several atoms degrades the time–frequency resolution, but a single atom is not sufficient for stability: this problem has also been addressed in [58,60], where only $n = 2$ orthogonal tapers were used, and the averaging was performed over numerous representations, each based on a random linear combination of these two atoms. Our continuous approach allows to further reduce this number to a non-integer optimum $n = \frac{Q_+}{Q_-} \approx \sqrt{2}$. As an extension, the number n could be made adaptive, increasing as the regularity of the rhythm is lost.

The phase of the fundamental mode of the cardiac oscillation is at the core of our quadratic approach. Other non-linear but non-statistical methods have made use of the specific shape of the waveform [57,61]. A future development in this direction may

benefit from the comparison with the maximum likelihood approach of smoothly deformed Gaussian stationary processes developed in [62–64].

6. Coherence and mutual information of cardio-respiratory signals

Cardio-respiratory coupling is one of the most important factors in the complex dynamics of the cardiovascular system and has aroused great interest of both clinicians and theoreticians. The interaction of the heart with other organs ensures the normal functioning of the body [65], while a deterioration of this coordination may indicate health impairment [66]. Some authors have even argued that the irregularity of breathing and the nonlinear nature of heart-respiration coupling could be potential sources of temporal variability in the heart rate [67,68]. This interaction results in a frequency modulation of the heart rate at the breathing frequency, one of the component of the HRV known as respiratory sinus arrhythmia (RSA) [69]. The distinction between cardio-respiratory synchronization and coordination has been the subject of a debate [70]. The complexity of heart rate variability and cardio-respiratory coupling has previously been estimated from RR intervals using correlation dimension, Lyapunov exponents [71,72], spectral coherence [73–75] or information-theoretic measures [76].

We investigate here the time–frequency coherence of cardio-respiratory signals and we relate its statistical significance to the Gaussian estimate of the mutual information. We elaborate on the adequate choice of the quality factors for a localized and significant analysis of the cardio-respiratory interaction.

6.1. Time–frequency coherence: generic definition and single trial estimation

Our aim here is to characterize the common information shared between two distinct and simultaneous signals. Sitting somewhere between pairwise correlations and mutual information, we revisit the concept of time–frequency coherence [22,77–79], an approach that offers both dynamical and statistical insights.

Given two processes y, z , an evolutionary cross-spectrum $S_{yz}(t, f)$ can be defined as in Eqs. (8) or (15), once a time–frequency decomposition is specified. This cross-spectrum can be normalized using the auto-spectra, resulting in a dimensionless coefficient called the time–frequency coherence between the processes [80,81]:

$$\gamma_{yz}(t, f) = \frac{S_{yz}(t, f)}{\sqrt{S_{yy}(t, f)S_{zz}^*(t, f)}}. \quad (24)$$

Like the cross-spectrum, the coherence has a Hermitian symmetry: $\gamma_{zy} = \gamma_{yz}^*$. It characterizes the similarities between two signals over time and frequency, in terms of the strength (modulus) and phase difference (angle) of their linear relationship. Coherence reaches a maximum modulus of 1 in the case of a perfectly linear relationship. In other words, it is a complex correlation coefficient between the statistics of the oscillating atoms at location (t, f) in each signal.

The above definition of the time–frequency coherence is generic, and applies to various definitions of the evolutionary spectra, such as the spectrogram or the scalogram as statistical expectations, or as single trial estimators. When applied to the spectral, temporal or global marginal cases described in Appendix C.3, Eq. (24) reduces to the spectral coherence, the time-dependent or the global Pearson correlation coefficient, respectively.

The estimation of the time–frequency spectra and coherence should require several statistical samples, such as independent and identically distributed realizations of the signal pairs. However, single trial experiments provide unique realizations of the processes (which cannot be identically repeated). In this case, nearby atoms from the time–frequency decomposition of the signals are used as distinct statistical samples. Such estimation comes at the cost of a loss of time–frequency resolution: around the location (t, f) , the original atomic uncertainty is extended to a continuous composition of n atoms, as defined in Eqs. (2) or (5). The case of the wavelet-based canonical estimator, $n = \frac{Q_+}{Q_-}$, is represented in Fig. 2.

The number of atoms plays a crucial role in the interpretation of the coherence value: the larger the number n , the more significant the statistical estimate, but the lower the time–frequency resolution and the stronger the assumption of local stationarity or self-similarity. The limit case of the coherence with all statistical degrees of freedom and no resolution is the (global) Pearson correlation coefficient, or its complex version if restricted to positive frequencies. Conversely, no statistical averaging corresponds to the limit case $n = 1$ with maximum resolution (first described in [77]): whatever the pair of signals, their atomic decompositions are always perfectly linearly related (unit modulus of coherence). The case $n = 0$ of an ideal time–frequency localization, directly based on Wigner–Ville distributions, leads to pathological properties [80,82]. The interference of several oscillating atoms is therefore crucial to obtain a statistical estimate of the coherence [22].

The trade-off between dynamical uncertainty and statistical significance is a crucial aspect for the single trial estimation of the time–frequency coherence. The control of the significance of the estimate as n is varied is examined next.

6.2. Spurious statistics and significance

To assess the significance of a time–frequency coherence estimate when statistical samples are limited, we propose to compare the value of the coherence modulus with what would be obtained from completely independent (*i.e.* incoherent) signals. The null

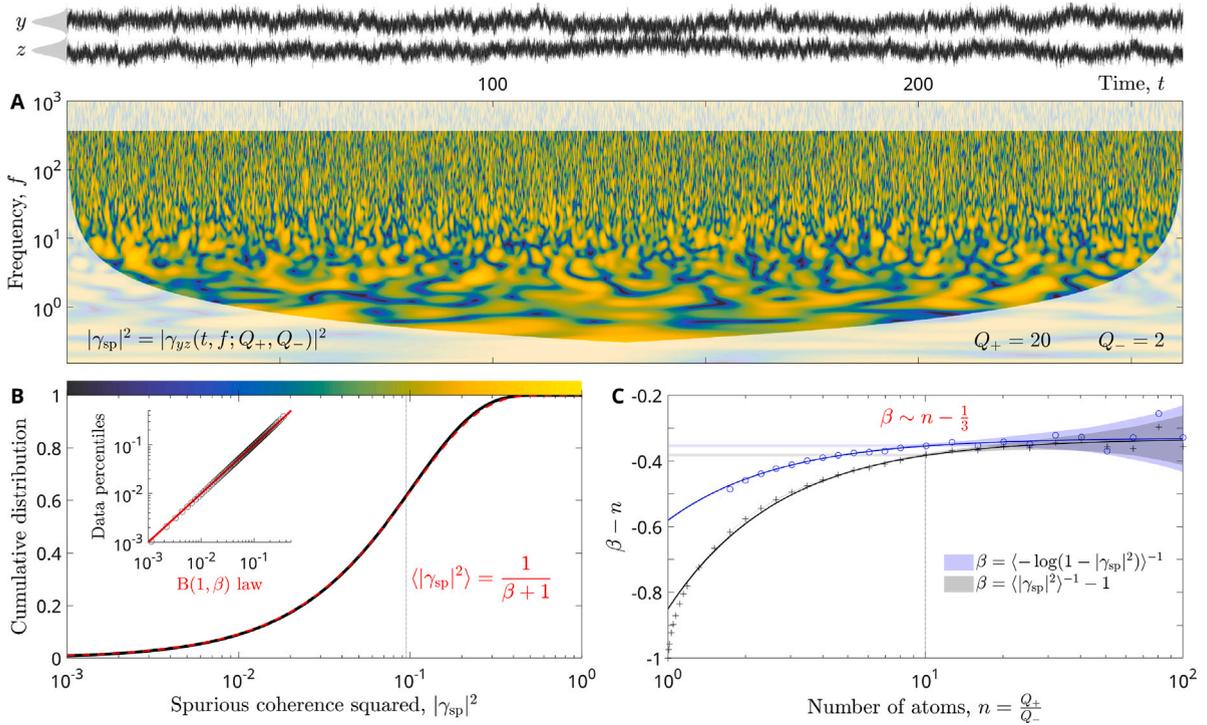


Fig. 6. Spurious statistics for the canonical time-frequency coherence estimator, associated with a pair of independent self-similar stationary real Gaussian processes. (A) Spurious coherence squared $|\gamma_{sp}|^2 = |\gamma_{yz}(t, f; Q_+, Q_-)|^2$ with quality factors $Q_+ = 20$ and $Q_- = 2$, computed from pink noises y, z of length $N = 2^{19}$ (top). (B) Cumulative distribution of $|\gamma_{sp}|^2$ (black solid line) compared to the beta distribution $B(1, \beta)$ with the same mean (red dashed line), aligned with the colour bar. Inset: quantile–quantile plot (close agreement). (C) Difference $\beta(n) - n$ for two estimations (makers) relying on the beta law, fits (lines) and 95% confidence intervals. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

hypothesis is thus constructed from the residual value of the coherence estimated between two independent stochastic trajectories. In Fig. 6, the spurious statistics are characterized from extensive simulations of the canonical estimate for the coherence between two independent stationary self-similar Gaussian processes. In particular, the spurious coherence shown in Fig. 6(A) is a wavelet-based estimate computed for two pink noises.

In the case of Gaussian stationary processes, the spurious spectral coherence squared $|\gamma_{sp}|^2$ is known to follow the Goodman distribution [83–85]. In the independent case, this probability distribution simplifies to a certain beta law with a single parameter, $B(1, \beta)$:

$$\mathbb{P}[|\gamma_{sp}|^2 \leq \gamma^2] = 1 - (1 - \gamma^2)^\beta \quad , \quad (25)$$

which can be easily transformed into a significant coherence threshold $\gamma^2(p) = 1 - p^{1/\beta}$ given the complementary probability $p = \mathbb{P}[|\gamma_{sp}|^2 > \gamma^2]$ or p -value. Fig. 6(B) shows that this distribution accurately matches the empirical distribution of the spurious coherence squared $|\gamma_{sp}|^2 = |\gamma_{yz}(t, f; Q_+, Q_-)|^2$ obtained from the simulation (A) for $n = 10$. The value of β is close to n and can be read from (C); as explained later, the distinction of the two curves indicates a departure from the beta law for small values of n .

The number $\beta + 1$ is interpreted as the number of (effective and complex) degrees of freedom of the statistical distribution, and it is identified in [85] with Eq. (2), which we call the (effective) number of atoms n . For the canonical estimations of the coherence, the number of atoms is known exactly: $n = \frac{Q_+}{Q_-}$ or $n = \frac{\tau_+}{\tau_-}$ in the wavelet or the window paradigm.

Assuming a beta law for $|\gamma_{sp}|^2$, we point out in Fig. 6(C) a small difference between the actual number of atoms n and the estimated statistical degrees of freedom $\beta + 1$. By varying n in the spurious coherence simulation, we estimate the parameter β , either from the mean of the beta-distributed squared spurious coherence $\langle |\gamma_{sp}|^2 \rangle = \frac{1}{\beta + 1}$, or from the alternative expression $\langle -\log(1 - |\gamma_{sp}|^2) \rangle = \frac{1}{\beta}$. The average is estimated over the time–frequency domain:

$$\langle |\gamma_{yz}|^2 \rangle = \|\Omega\|^{-1} \iint_{\Omega} |\gamma_{yz}(t, f; Q_+, Q_-)|^2 dt df \quad . \quad (26)$$

Here, Ω is the inner region in Fig. 6(A), free of boundary effects, with area $\|\Omega\| = \iint_{\Omega} dt df = a - b(1 + \log \frac{a}{b})$, $a = \frac{N}{2} e^{-\frac{k}{Q_-}}$, $b = kQ_+ e^{\frac{k}{Q_-}}$, where $N = \frac{T}{dt}$ is the length of the time series, and k is the distance from the boundaries (in resolution units). The outer region (faint colours) that is excluded from the estimation in Fig. 6 corresponds to $k = 2$.

Even a very long single simulation of the spurious coherence is not sufficient to obtain an accurate estimate of β , especially for a high value of n . Therefore, we average Eq. (26) over numerous simulations $j = 1, 2, \dots$, so that the cumulative averaging area $\sum_j \|\Omega_j\|$, exceeds $2^{21} \times n$ for each value of n .

Any location (t, f) in the inner time–frequency region Ω (away from the boundaries) yields the same estimate: the spurious coherence is homogeneously distributed. A distance of twice the unit resolution, $k = 2$, is sufficient to separate the excluded outer region from Ω . We could decrease the sampling along log-frequency down to two samples per resolution unit without observing alteration of the estimation.

The difference $\beta - n$ is plotted in Fig. 6(C) (markers) and linearly fitted with respect to n^{-1} (lines). The parameter $\beta(n)$ grows from zero (for $n = 1$) to an asymptotic value:

$$\beta(n) \sim n - \frac{1}{3} - \frac{c}{n} \quad , \quad (27)$$

in spite of the growing uncertainty for large n . The accuracy of assuming a beta law $B(1, \beta)$ for $|\gamma_{sp}|^2$ is visible in (C) by comparing the term of order $\mathcal{O}(n^{-1})$ for both estimators of β : its coefficient is compatible with $c = \frac{1}{2}$ for the mean-based estimate (plus markers and black line), and compatible with $c = \frac{1}{4}$ for the alternative estimate (circle markers and blue line) exploiting the relationship between beta and exponential random variables. Small values ($n < 1.5$) have been discarded for the latter due to numerical imprecision.

We have tested the robustness of this asymptotic behaviour by varying different choices, such as the Hurst exponent of the independent signals (0 for pink or $-\frac{1}{2}$ for white noises), the geometric paradigm for the canonical estimation, the value of the associated parameters (Q_+, Q_-) or (τ_+, τ_-) , and numerical parameters, without observing discrepancies beyond the numerical uncertainty. Only the number of atoms n matters.

In practice, the case $n = 10$ in Fig. 6(A, B) yields $\beta = 9.618 \pm 0.004$ (black) and 9.647 ± 0.005 (blue), or a very small difference in the significance test. The beta law is not the exact distribution of the spurious coherence squared, but it may be considered as a very good approximation for a sufficiently large n . Eqs. (25), (27) provide a simple way to assess the significance of the time–frequency coherence estimator γ_{yz} , as the p -value $p = (1 - |\gamma_{yz}|^2)^\beta$.

6.3. Mutual information and its time–frequency distribution

In the case of two Gaussian processes y and z , the modulus of their coherence gives the time–frequency distribution of their mutual information in the following form:

$$i_{yz}(t, f) = -\frac{1}{2} \log (1 - |\gamma_{yz}(t, f)|^2) \quad . \quad (28)$$

The relationship between coherence and rate of mutual information, $I_{yz} = \int i_{yz}(f)df$, is indeed proved in the case of Gaussian stationary processes [86]. In the simplest case (for white noises), the mutual information of two Gaussian random variables is obtained directly by applying Eq. (28) to the Pearson correlation coefficient (as a global version of the coherence).

This provides a natural non-stationary estimator of the mutual information between two Gaussian processes y and z , here measured in natural units (nats), and possibly restricted to a time–frequency domain Ω :

$$I_{yz}[\Omega] = \iint_{\Omega} i_{yz}(t, f) dtdf \quad . \quad (29)$$

As for the coherence, in situations with limited time–frequency atoms (n small), it is necessary to assess the significance and the spurious expected value of this statistical estimator. A beta law $B(1, \beta)$ for the spurious coherence squared $|\gamma_{sp}|^2$ implies an exponential law $\text{Exp}(2\beta)$ for the spurious density of mutual information, expected for independent Gaussian processes:

$$\langle i_{sp} \rangle = \frac{1}{2\beta} \quad , \quad \langle I_{sp}[\Omega] \rangle = \frac{\|\Omega\|}{2\beta} \quad . \quad (30)$$

This explains the alternative way of estimating the parameter β in Fig. 6(C, blue circles). Furthermore, it allows us to interpret the magnitude of significance of the estimated coherence as the ratio of the estimated and spurious density of mutual information: $-\log p = i_{yz} / \langle i_{sp} \rangle$. Evenly sampled, this quantity is useful to colour significant values of a coherence estimate (see Fig. 7).

The difference $I_{yz} - \langle I_{sp} \rangle$ can be used as a “corrected” estimator of the mutual information [87], which corresponds to the “debiased” coherence estimator $(|\gamma_{yz}|^2 - \langle |\gamma_{sp}|^2 \rangle) / (1 - \langle |\gamma_{sp}|^2 \rangle)$ [88]. Here, $\langle |\gamma_{sp}|^2 \rangle = \frac{1}{\beta+1}$, but the direct use of $\frac{1}{n}$ (without correction) was shown to be accurate for an effective number of atoms as low as $n = 8$ [88].

These time–frequency coherence and mutual information estimators based on Gaussian processes are applied to physiological signals in the next section. We assume that they are locally stationary over Q_+ oscillations and locally self-similar over a relative bandwidth Q_-^{-1} . Due to the non-stationarity of these signals, we argue that a time–frequency estimator of the mutual information based on the Gaussian hypothesis may be more informative than alternative non-parametric estimators based on a stronger assumption of stationarity.

6.4. Different regimes of cardio-respiratory coherence

To study the modulation of the heart rate at the breathing frequency, called respiratory sinus arrhythmia (RSA), and other forms of cardio-respiratory coherence, we estimate in Fig. 7 the coherence and mutual information between the heart rate and respiratory

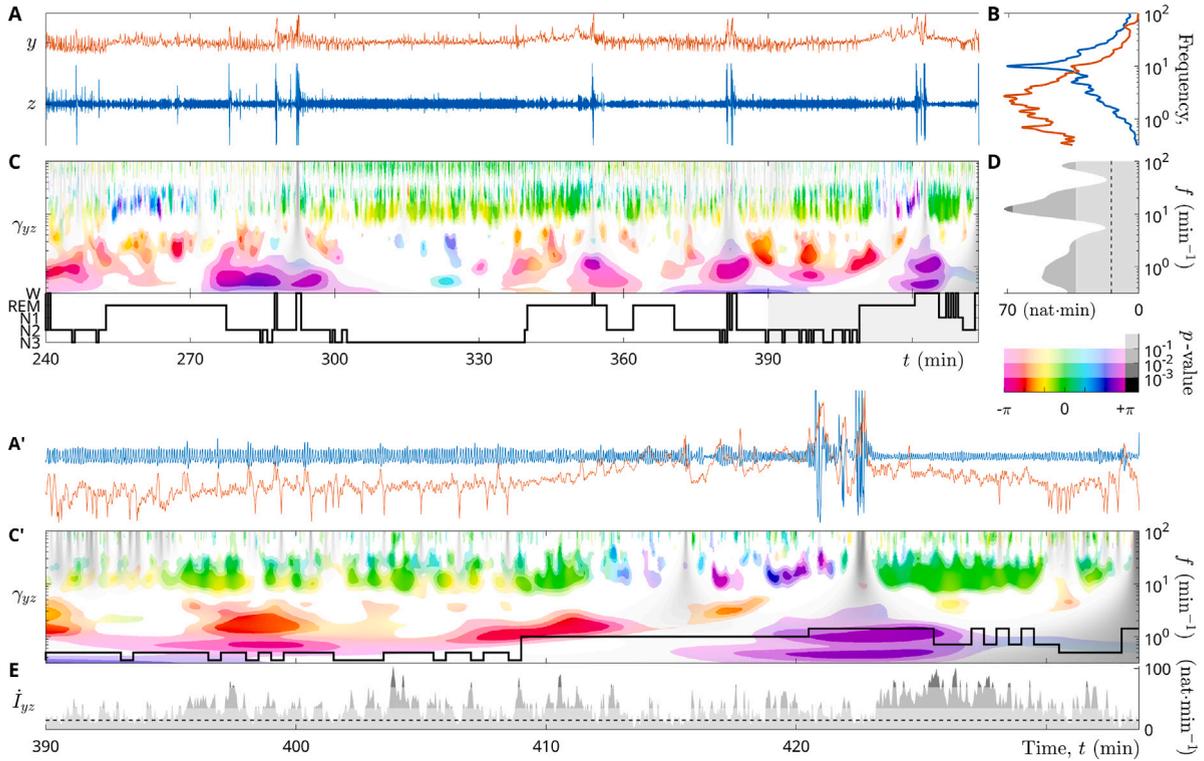


Fig. 7. Coherence between heart rate and breathing, and associated mutual information estimates. (A) Signals of the heart rate y (red) and the respiratory effort z (blue), with arbitrary units. (B) Amplitude spectra for y and z , with frequency in beats per minute (min^{-1}). (C) Time–frequency coherence $\gamma_{yz}(t, f; Q_+, Q_-)$ for $Q_+ = 7$ and $Q_- = 1$, coloured according to its phase (hue), significance (saturation) and signals’ lack of quality (shades). Hypnogram (black stairs). (A’, C’) Focus at the end of sleep. Mutual information: (D) spectral density ($\text{nat}\cdot\text{min}$) and (E) rate (temporal density, $\text{nat}\cdot\text{min}^{-1}$) with expected spurious value (black dashed line). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

signals. The former is the rate signal previously extracted from the ECG, and the latter is a breathing effort signal defined as the sum of recordings from thoracic and abdominal belt transducers. The sign of the respiratory oscillation is set to positive for inhalation.

The oscillatory content in the signals Fig. 7(A) –the heart rate is denoted $y(t)$ (red) for simplicity and the respiratory signal $z(t)$ (blue)– is summarized in their respective amplitude spectra (B). Quite regular, the respiratory cycle is concentrated in a relatively narrow frequency band around 10 min^{-1} . In contrast, the heart rate variability covers a wider band centred around $2\text{--}3 \text{ min}^{-1}$. The spectrum of the respiratory recording is attenuated at frequencies below 1 oscillation per minute, as a result of instrumental filtering. There is no such high-pass filtering for the heart rate signal, which is extracted from modulations in the ECG recording. Instead, most of the amplitude at the cardiac (carrier) frequency (60 min^{-1}) has disappeared from the estimated heart rate signal. Note that the slow modulations of the respiratory frequency and amplitude could also be estimated as a complex rate signal. Cardio-respiratory coherence with both types of respiratory signals, oscillations or fluctuating rate, have been analysed extensively for all subjects in the *shhs2* dataset in [42]. Breathing rate signals would enhance the study of very slow modulations, at the cost of deteriorating the observation of RSA at the carrier (respiratory) frequency.

The time–frequency coherence γ_{yz} between these signals (A) is then estimated in Fig. 7(C, zoom in C’), with quality factors $Q_+ = 7$ and $Q_- = 1$, and colour-coded as follows. The cyclic hue corresponds to the phase, and the 4 levels of saturation are delimited by 3 significance thresholds for the modulus: $p = 10^{-a}$, $a = 1, 2, 3$, thus highlighting only the relevant (density of mutual) information. Regions affected by boundary effects, recording saturation or higher than average error (3%) in the heart rate estimation are shaded.

We can distinguish three bands (timescales) of coherent cardio-respiratory activity, as summarized in the spectral marginal density (time integral) of mutual information (D). Overall, the intermediate band close to and above the respiratory frequency is the most significant: this is the RSA. Patches of coherence are observed in a lower and wider band, associated with slow intermittent and coordinated cardio-respiratory events. Finally, a cardiac mode (at 60 min^{-1} and above) indicates that the respiratory recording can capture slight heart pulses. This coherent mode is absent when the reference (discrete interpolated) estimator of the heart rate is used, since it does not contain the residual cardiac oscillations present in the continuous estimator.

The coherence of the respiratory (RSA) mode often reaches high levels of significance, $p < 10^{-3}$, as shown in Fig. 7(C, C’). This is the case not only around the fundamental respiratory frequency (10 min^{-1}) but also beyond 20 min^{-1} , implying that the specific shape of the respiratory oscillation (carried by higher order partials, here indistinguishable with $Q_- = 1$) is coherent in both signals. Part of the intermittent loss of RSA coherence coincides with regions potentially affected by a poor heart rate estimation (shaded), either caused by motion artefacts or by strong heart rate variability. The RSA coherence is otherwise quite persistent.

The phase difference between the RSA mode in the heart rate and the oscillation of the breathing effort can vary significantly. Mostly in-phase (angle of coherence 0 in green: the heart rate increases when breathing in), the RSA mode can be lagged of $-\frac{\pi}{3}$ (yellow); this tends to happen during deep sleep (N3 stage, see the hypnogram). In contrast, the RSA phase tends to increase during rapid-eye-movement (REM) stage, occasionally reaching positive phase quadrature ($+\frac{\pi}{2}$ in blue) or even phase opposition (π in magenta). The role of sleep stages in RSA changes would require further statistical confirmation across many subjects. However, these observations are consistent with a weaker control of the cardio-respiratory inter-regulation during REM, than during non-REM sleep.

The spectral range is sufficient to observe longer timescales of cardio-respiratory coherence, intermittently reaching high levels of significance ($p < 10^{-3}$) in a broad band below 5 min^{-1} . The angle of coherence facilitate the distinction of two types of coherent cardio-respiratory events in this band, that were previously hard to separate in a low and a very low-frequency band [75].

Micro-awakenings are easily detected in the respiratory signal, as a deep inhalation, followed by a long exhalation (or more). The empty lungs then need several breathing cycles to refill, while the heart rate remains high to pump the oxygenated blood, hence the phase opposition (π in magenta) in this slow mode. Since the rather singular events that generate it are averaged over $Q_+ = 7$ oscillations, they can interfere with (or even dominate) less intense coherent events. Lower quality factors would be better suited to focus on these broadband time-localized structures.

Coherent events of comparatively lower amplitude, such as hyperpnea and hypopnea, are associated with patches of coherence close to negative phase quadrature ($-\frac{\pi}{2}$ to $-\frac{2\pi}{3}$ in orange-red). They disappear during the long episode of deep (N3) sleep, and may be associated with a decrease in RSA coherence. A strong and persistent occurrence of this coherent mode, as a rhythm near 1 to 2 min^{-1} , is associated with severe sleep apnea [54], hence called the apneic mode.

This time–frequency map of the cardio-respiratory coherence allows a concise visualization of local and significant statistical similarities between physiological fluctuations and oscillations, together with their phase. It should be remembered that the time–frequency plane has a hyperbolic metric in the wavelet paradigm, so that its flat projection as an image should be interpreted with caution. In fact, low frequency atoms look more important than their actual spread, and vice versa at high frequencies, as schematized in Fig. 1(B, right). We may think of it as a perspective view: close-by (low frequency) wavelets look larger than the ones near the horizon (maximum frequency), as best noticed in the spurious coherence Fig. 6(A).

6.5. Quality factors and trade-off between localization and significance of mutual information

The pictures of the significant values of the coherence estimate in Fig. 7(C, C') combine the angle of coherence to the empirical time–frequency distribution of mutual information in the Gaussian approximation $i_{yz}(t, f; Q_+, Q_-)$. In Fig. 7(D and E), we plot the spectral and temporal marginal densities of mutual information by integrating it over time or frequency respectively. The corresponding significance thresholds are shown with grey levels. From the visible region in (D), we can estimate a corrected value of the total mutual information $I_{yz}[\Omega] - \langle I_{sp}[\Omega] \rangle = 4.2 \text{ knat}$, within a range $\|\Omega\| = 2f_{\max}T = 3.9 \times 10^3$ (the factor 2 accounts for negative frequencies), or on average just above 1 nat per unit area.

This mutual information is divided into three visible frequency bands as follows: 2.0 knat in the cardiac band, 2.0 knat in the respiratory/RSA band, 0.2 knat in the low frequency band. These bands are bounded by the frequencies 44 min^{-1} and 5.5 min^{-1} of two minima close to the spurious value (black dashed line) of the spectral density of mutual information Fig. 7(D). Note that (D) does not represent a log-frequency density: the high-frequency (cardiac) band accumulates much more mutual information than the low-frequency band.

For this reason, the estimated mutual information rate (temporal marginal) in Fig. 7(E) is dominated by the high-frequency coherence, even though the corresponding modulations have a very small amplitude in the signals. The expected spurious value (black dashed line) correctly describes the baseline mutual information rate. However, its value and significance is very sensitive to the choice of f_{\max} , which includes here the cardiac band.

Again, we try to refine the choice of quality factors. Essentially, we want to maximize the information extracted from the coherence of the signals in the selected time interval by tuning the parameters of our local stationarity and self-similarity assumptions. In practice, this is quite difficult to achieve for several reasons. First, we have to deal with the limitations of the information estimators and the resulting artefacts, namely: boundary effects and the imprecision of the statistical model for small n . The significance decreases near the boundary in a way that is not modelled here, and is therefore excluded from the inner region of integration Ω . The upper frequency limit $f_{\max} < \frac{1}{2}f_s e^{-k/Q_-}$ is particularly critical. The corrected estimate of the mutual information $I_{yz}[\Omega] - \langle I_{sp}[\Omega] \rangle$ increases rapidly when n approaches 1, as a consequence of the imprecision of the beta and exponential models of significance (Fig. 6 C). We thus consider only situations with at least $n = \frac{Q_+}{Q_-} \geq 2$ atoms.

The corrected estimate of the mutual information has no clear global maximum in the valid parameter space: for $f_{\max} = 100 \text{ min}^{-1}$, the estimate is nearly optimal for $Q_- \approx 1$, but keeps increasing for $Q_+ \rightarrow Q_-$. Although corrected, the mutual information has very little significance in the limit of a minimal dynamical (time–frequency) uncertainty. The opposite situation occurs when maximizing the mean magnitude of significance $\langle -\log p \rangle = I_{yz}[\Omega]/\langle I_{sp}[\Omega] \rangle$: the dynamical uncertainty is also maximized, $n \rightarrow +\infty$. An optimal bandwidth emerges from these observations, but the number of atoms varies widely depending on our preference for a high corrected value or a high significance of the mutual information.

We measure a trend $\langle i_{yz} \rangle = I_{yz}[\Omega]/\|\Omega\| \approx \alpha/\log n$, which holds for an optimal value $\alpha = 0.36$ over a wide range $n \in [3.5, 70]$, for $f_{\max} = 100 \text{ min}^{-1}$ and $Q_- \approx 1$. We learn that the mutual information is here inversely proportional to the loss of localization information. As n decreases, the spurious contribution increases. In general, α varies with the quality factors and f_{\max} ; we believe

that these variations are informative about the statistical and dynamical properties of the different coherent components in the analysed signals.

The low quality factor $Q_- = 1$ provides an almost optimal bandwidth in the spectral range of Fig. 7, sufficient to distinguish the cardiac, respiratory and apneic activity bands. The upper quality factor $Q_+ = 7$ represents a compromise between a precise localization and a high level of significance, highlighting the richness of the interplay between the cardio-respiratory signals.

Thanks to a careful choice of the quality factors, the time–frequency coherence and the mutual information provide local dynamical insights with a high statistical significance on different forms of cardio-respiratory coherence. The mutual information and its directed extension, the transfer entropy, have also been estimated on short timescales for physiological point processes [87]: the control of their spurious statistics is not straightforward and has been managed from surrogate data. We therefore believe that the detailed discussion of localization and significance in this type of analysis will be a valuable aid for future investigations. Mutual information between heart rate and respiration is also used in other contexts, such as emotional characterization [89], which could constitute another immediate application of this work.

7. Conclusion

From the interplay between respiratory oscillations and heart rate fluctuations, to the multiscale stochastic and intermittent waves of a sleeping brain, a single recording of biophysical observables can contain a wealth of information about the underlying living system dynamics. However, the practical exploitation of this information is hampered by the inadequacy of standard dynamical assumptions: a stationary, self-similar or scale-separated approach wastes some of the richness of the recorded information. The parsimony of our assumptions is only constrained by the uncertainty principle, expressed here in the time–frequency domain.

In the context of a single trial experiment, we have shown how the uncertainty principle binds the localization of the analysis to the significance of the resulting statistical estimate. On Heisenberg's uncertainty bound, signals are continuously decomposed into time–frequency atoms, but the statistical approach is brought to the brink of collapse. However, time–frequency decomposition is a fertile ground for recomposing uncertainty atoms into flexible quadratic statistical estimates, whose significance can be tested when their localization lies away from this lower bound.

More specifically, the numbering of uncertainty atoms (eigenstates) composing a quadratic statistical estimator is given a continuous time–frequency interpretation from the Wigner–Ville representation of the operator. We have proposed a canonical implementation of these generic concepts in a concise analytical framework, based on the choice of a geometric paradigm and two essential physical parameters. This versatile approach has also been illustrated in practical situations.

First, a canonical spectrogram estimator built from Gabor's oscillating normal windows has been defined. The translation invariance of this atom in the time–frequency plane has been associated with a flat (Euclidean) representation with two fixed timescales, one for the decomposition, and a different one for the recombination into a statistical estimate. Their ratio gives the exact number of statistical degrees of freedom, while matching the classical expression of the uncertainty principle.

We have then constructed the scale-invariant analogue, the canonical scalogram estimator, from the log-normal wavelet of Altes and Grossmann. The translation and scaling of this atom is associated with a hyperbolic geometry (hence the delicacy of the representation as a flat time–frequency image). In this paradigm, quality factors are fixed parameters, replacing timescales. This leads to a scale-free formulation of scale separation and of the uncertainty principle.

Using the electroencephalogram (EEG) of a sleeping subject as an illustration, we have discussed more concretely the two geometric paradigms and the choice of the localization parameters. The spectrogram approach can be convenient for studying the brain cortical activity near fixed reference timescales that separate well its correlation time from its evolution, whereas the scalogram allows a broader and clearer study of narrowband (rhythmic) and broadband (noisy) components across scales.

In the application to the heart rhythm, we have proposed an original continuous and quadratic estimator of the heart rate signal. Based on our statistical perspective, we have derived this estimator from the phase of the cardiac oscillation. Interestingly, compared to other estimators, this rate is non-diverging and fine-tuned using three parameters: a timescale, a quality factor and a number of degrees of freedom. We argue that this quadratic method is among the simplest non-linear ones for a continuous heart rate extraction.

Finally, we have measured the coherence between heart rate modulations and respiratory oscillations over time and scales, and we have established the interpretation of its significance as the ratio of the Gaussian estimate of the mutual information to its spurious value. The spurious statistics that arise in the coherence and mutual information estimations has been precisely characterized in terms of the number of atoms, and the relation between mutual information and localization information has been investigated.

This tour of physiological signals completes the picture of the uncertainty trade-off linking the significance of a single trial estimate and its time–frequency localization. Few physical parameters are thus left to tune our dynamical and statistical assumptions about the living system, as applied here to the complexity of polysomnographic information. The proposed canonical framework provides a practical summary of the fundamental concepts at stake in single trial time–frequency statistics, and offers a generic and explicit way of dealing with rich but uncertain dynamics.

CRedit authorship contribution statement

Alexandre Guillet: Conceptualization, Formal analysis, Investigation, Software, Visualization, Writing – original draft, Writing – review & editing. **Françoise Argoul:** Data curation, Funding acquisition, Methodology, Supervision, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We would like to pay tribute to Alain Arneodo, whose enthusiasm has durably stimulated the development of the above ideas. We are very grateful to Jacques Taillard (SANPSY) for insightful discussions. AG acknowledges LOMA and MPI-PKS for their hospitality and EURLight S&T for funding. This work was supported by Agence Nationale de la Recherche (ANR), France through contract: ANR-18-CE45-0012-01.

Appendix A. Characteristics of electrophysiological recordings

A.1. Electroencephalogram: bands and sleep stages

Electroencephalography (EEG) is a recording technique that uses electrodes in contact with the scalp to record the potential variations produced by neurons in the brain. Each electrode locally integrates a collective and spatial neural activity. The resulting signal, shown in Fig. 3(B), exhibits combinations of erratic (scale-free) and rhythmic dynamics over a wide range of frequencies (more than two decades). Unlike cardiac or respiratory signals, no single fundamental mode can be identified in this neural signal. This EEG signal measures the electrical potential between the points C4-A1 with a resolution of $1 \mu\text{V}$, a sampling frequency $f_s = 125 \text{ s}^{-1}$, and instrumental filters (high-pass at 0.15 s^{-1} and notch at 60 s^{-1}).

The complexity of an EEG can be represented in the time–frequency plane to distinguish both its evolutionary and spectral aspects, see Fig. 3(D, F). The traditional way of dealing with this complexity is to divide time into discrete stages and frequency into discrete bands. Wake-sleep stages have been subdivided into wake (W), rapid eye movement (REM) and non-REM (or NREM) stages, and the NREM stage has been further decomposed into N1 (lightest), N2 and N3 (deep sleep) stages. This classification is based on a set of criteria about the relative EEG power in each frequency band (possibly supplemented by the EOG and EMG), within 30-second epochs [90]. The sequence of these stages, available in most polysomnographic databases as a clinician annotation every 30 s, is a simplified representation of sleep called a hypnogram, see Fig. 3(A). With the aim of empirically distinguishing neural waves, the frequency bands are roughly defined in slices of 4 s^{-1} and denoted by Greek letters: δ up to 4 s^{-1} , θ from 4 to 8 s^{-1} , α from 8 to 12 s^{-1} , σ from 12 to 16 s^{-1} , β from 16 to 20 s^{-1} , and γ above 20 s^{-1} . Thus, γ waves are intense during wakefulness whereas they disappear during NREM sleep, which is characterized by intense δ waves (especially during the N3 stage). In contrast, during REM sleep, the EEG power is low in all bands.

A.2. Electrocardiogram: identification of QRS complexes

The electrocardiogram (ECG) is the most common clinical technique to record the cardiac oscillation: the electric potential between electrodes in contact with the skin on the chest and limbs provides a signal with a sharp waveform. The ECG waveform has three main components: the P wave (atrial depolarization), the QRS complex (ventricular depolarization), and the T wave (ventricular repolarization). The detection of RR time intervals from successive R-peaks is a standard and discrete method of measuring the cardiac period from the ECG.

As a reference conventional method, for comparison with the wavelet-based continuous method developed in Section 5, we use the *jqrs* algorithm from the PhysioNet Cardiovascular Signal Toolbox [91], which first computes the QRS intensity using a matched filter (band-pass) and a sequence of signal processing operations. The intensity maxima in the resulting positive pulse train are then detected by a thresholding procedure. To validate the detection method or to reject unreliable RR intervals, a signal quality index can be calculated, for example from a comparison with another detector. Automatic error correction, which interpolates missing beats and removes false beats based on physiological assumptions, can improve the estimation of the heart period (NN intervals for “normal”). Adjustment of sufficient detection parameters (such as thresholds) can give satisfactory results without the need for expert modifications (which are less reproducible).

Appendix B. Linear representations: list of relations

The signal and spectrum representations, as related by the Fourier transform, are considered here as limit cases of time–frequency representations, that are purely localized in one domain and delocalized in the other. We recall their definition here to clarify the convention used, and we provide explicit relations between limit and intermediate cases (as defined in Eq. (14) and (7)) with finite quality factor or timescale. The temporal and spectral interpretations of the quality factor is also described in more detail.

B.1. Signal and spectrum conventions

The Fourier transform and its inverse relate the signal to the spectrum, as follows:

$$x(t) = \int \hat{x}(f)e^{i2\pi ft} df \quad ; \quad \hat{x}(f) = \int x(t)e^{-i2\pi ft} dt \quad . \quad (B.1)$$

The wavelet transform $X(t, f; Q)$ as defined in Eq. (14) is a band-pass filtering, i.e. a spectral decomposition of the signal $x(t)$ (with the same physical unit) into different bands (f, Q). In contrast, the Gabor transform $G_x(t, f; \tau)$, Eq. (7), is a Fourier transform under a sliding window, i.e. a temporal decomposition of the spectrum $\hat{x}(f)$ in different time periods (t, τ). These signal-like or spectrum-like behaviours are best noticed by recovering the actual signal and spectrum in the limit cases:

$$x(t) = \lim_{Q \rightarrow 0} X(t, f; Q) + X(t, -f; Q) = \lim_{\tau \rightarrow 0} \tau^{-1} G_x(t, f; \tau) e^{i2\pi ft} \quad (B.2)$$

$$\hat{x}(f) = \lim_{Q \rightarrow +\infty} Q|f|^{-1} X(t, f; Q) e^{-i2\pi ft} = \lim_{\tau \rightarrow +\infty} G_x(t, f; \tau) \quad . \quad (B.3)$$

The Gabor transform can thus be converted into a spectral decomposition of the signal:

$$X(t, f; \tau) = \tau^{-1} G_x(t, f; \tau) e^{i2\pi ft} \quad . \quad (B.4)$$

Conversely, the wavelet transform becomes a temporal decomposition of the spectrum by reversing this equality with $\tau = Q|f|^{-1}$. The S-transform [35], which would be denoted here $Q|f|G_x(t, f; Q|f|^{-1})$, is a mixture of these two conventions with the phase of the spectrum and the unit of the signal. Based on a Gabor-Morlet wavelet (a scaled normal window), the S-transform behaves as $X(t, f; Q)e^{-i2\pi ft}$ in the log-normal wavelet framework.

B.2. Change of localization

The localization parameter of a time–frequency decomposition can be decreased or increased as follows:

$$X(t, f; Q_-) = \frac{Q_+ Q_-}{Q_-} \int X(t, f'; Q_+) \psi_{Q_-}^* \left(\frac{f'}{f} \right) \frac{df'}{|f'|} \quad , \quad Q_-^{-2} = Q_-^{-2} - Q_+^{-2} \quad (B.5)$$

$$X(t, f; Q_+) = \int X(t', f; Q_-) \psi_{Q_+}^* (f(t' - t)) |f| dt' \quad , \quad Q_+^{-2} = Q_+^{-2} - Q_-^{-2} \quad (B.6)$$

$$G_x(t, f; \tau_-) = \int G_x(t, f'; \tau_+) w_{\tau_-} (f' - f) \tau_- e^{i2\pi(f' - f)t} df' \quad , \quad \tau_-^{-2} = \tau_-^{-2} - \tau_+^{-2} \quad (B.7)$$

$$G_x(t, f; \tau_+) = \frac{\tau_+}{\tau_- \tau_+} \int G_x(t', f; \tau_-) w_{\tau_+} (t' - t) dt' \quad , \quad \tau_+^{-2} = \tau_+^{-2} - \tau_-^{-2} \quad . \quad (B.8)$$

The signal and the spectrum are also recovered by integration over one domain:

$$x(t) = Q \int X(t, f; Q) \frac{df}{|f|} = \int G_x(t, f; \tau) e^{i2\pi ft} df \quad (B.9)$$

$$\hat{x}(f) = \int X(t, f; Q) e^{-i2\pi ft} dt = \frac{1}{\tau} \int G_x(t, f; \tau) dt \quad . \quad (B.10)$$

B.3. Case of the analytic version of the signal

Any signal x can be expressed as the sum $x(t) = x^+(t) + x^-(t)$ of the analytic and anti-analytic versions of the signal. They are complex helical signals defined from the part of the Fourier spectrum $\hat{x}(f)$ that is supported on positive (respectively negative) frequencies:

$$x^\pm(t) = \int \Theta(\pm f) \hat{x}(f) e^{i2\pi ft} df \quad , \quad (B.11)$$

where Θ is the Heaviside step function restricting the integral to positive or negative frequencies.

The time–frequency decomposition of the signal based on the log-normal wavelet, $X(t, f; Q)$, can be understood for any f as the analytic version of the band-pass filtered signal around the frequency f . A similar interpretation is sometimes possible for Eq. (B.4) based on Gabor’s normal filter, although it would not be strictly analytic because of its support on both positive and negative frequencies.

The analytic and anti-analytic versions of the signal x can thus only be expressed as a limit and a marginal case of a wavelet transform, whose wavelet is analytic, such as the log-normal wavelet:

$$x^\pm(t) = \lim_{Q \rightarrow 0} X(t, \pm|f|; Q) = Q \int \Theta(\pm f) X(t, f; Q) \frac{df}{|f|} \quad . \quad (B.12)$$

When the signal x is real, negative and positive frequencies are redundant, so that $x^- = (x^+)^*$, and $x = 2\Re\{x^+\}$.

B.4. Intertwined spectral and temporal interpretations of the quality factor

In the time domain, the quality factor Q is a dimensionless time resolution, $|f|\Delta t$, interpreted as the effective number of oscillations in the analysing wavelet. Oscillations are also localized in the frequency domain, where Q^{-1} is a log-frequency resolution $\Delta \log f$ (alternative expression for the relative bandwidth $\frac{\Delta f}{|f|}$). The quality factor Q in the log-normal wavelet $\psi_Q(f'/f) =$

$e^{-\pi(Q \log(f'/f))^2}$ is smaller than the one calculated from the definition in [92], by a factor of about $\sqrt{\pi}$. The quantities $|f|^{-1}Q$ and Q^{-1} are in fact both very close to the full widths at half maximum of the wavelet [42], respectively in time and in log-frequency.

The specificity of the waveform of a regular rhythm can be distinguished from a purely circular (sine) oscillation in the frequency domain by the presence of multiple harmonic modes with frequencies $f_k = kf_1$ for positive integer orders k . In the wavelet paradigm, the order $k + 1$ can be distinguished from the order k whenever $\Delta \log f = Q^{-1} < \log \frac{k+1}{k}$. Furthermore, wavelets with Q oscillations can only analyse a full rhythmic period $\Delta t = f_1^{-1}$ up to the frequency Qf_1 . As a consequence of the following logarithmic inequalities [93]:

$$\frac{1}{Q + \frac{1}{2}} < \log \frac{Q + 1}{Q} < \frac{1}{Q} < \log \frac{Q + \frac{1}{2}}{Q - \frac{1}{2}} \quad , \tag{B.13}$$

the spectral discrimination of $Q > \frac{1}{2}$ partials requires a wavelet with at least Q oscillations. In other words, a quality factor Q allows the distinction of up to Q harmonic orders. This inequality illustrates two intertwined and interchangeable interpretations of the quality factor: in terms of either time or frequency resolution, expressed as number of oscillations or distinct harmonic orders. We can visualize their relevance in Fig. 4, in the wavelet transforms of the ECG signal (A) with quality factor $Q = 1$ (B) or 3 (C).

These very practical and natural interpretations of the quality factor as the number of wavelet oscillations and the number of distinguishable harmonic orders are silently orchestrated by the presence of the circle constant in the definition of the log-normal wavelet Eq. (13). These realizations have motivated a rescaling of the quality factor, and hence a slight redefinition of the log-normal wavelet, as compared to our previous definition in [42,54,94].

Appendix C. Quadratic representations in time and/or in frequency

The relationships between time, frequency and time–frequency representations for the kernel of a quadratic operator are detailed. As an example, these different representations are provided in the case of the canonical spectrogram estimator. The marginal and limit cases of the canonical spectrogram and scalogram are then discussed.

C.1. Wigner–ville representation as a time–frequency distribution

The Wigner–Ville distribution defined in Eq. (3) can be applied to the temporal or spectral kernel of a quadratic operator, to obtain its time–frequency representation:

$$K(t, f) = \int \check{K}\left(t + \frac{\tau}{2}, t - \frac{\tau}{2}\right) e^{-i2\pi f \tau} d\tau = \int \hat{K}\left(f + \frac{\eta}{2}, f - \frac{\eta}{2}\right) e^{i2\pi \eta t} d\eta \quad . \tag{C.1}$$

The temporal and the spectral representations (related by Fourier transforms) of the quadratic operator are obtained from the time–frequency kernel as follows:

$$\check{K}(t_1, t_2) = \int K\left(\frac{t_1+t_2}{2}, f\right) e^{i2\pi f(t_1-t_2)} df \quad ; \quad \hat{K}(f_1, f_2) = \int K\left(t, \frac{f_1+f_2}{2}\right) e^{-i2\pi(f_1-f_2)t} dt \quad . \tag{C.2}$$

C.2. Kernel of the canonical spectrogram estimator

By denoting the Gaussian function $\phi(u) = e^{-\pi u^2}$ which satisfies $\hat{\phi} = \phi$, we can obtain symmetric expressions for the operator K_w represented in the time and frequency domains:

$$\check{K}_w(t_1, t_2; t, f) = \frac{\sqrt{2}}{\tau_+} \left(\phi\left(\frac{t_1-t_2}{\tau_-}\right) \phi\left(\frac{t_1+t_2-2t}{\tau_+}\right) \right)^{\frac{1}{2}} e^{i2\pi f(t_1-t_2)} \tag{C.3}$$

$$\hat{K}_w(f_1, f_2; t, f) = \frac{\sqrt{2}}{\eta_+} \left(\phi\left(\frac{f_1-f_2}{\eta_-}\right) \phi\left(\frac{f_1+f_2-2f}{\eta_+}\right) \right)^{\frac{1}{2}} e^{-i2\pi(f_1-f_2)t} \quad , \quad \eta_{\pm} = \tau_{\mp}^{-1} \quad . \tag{C.4}$$

The time–frequency representation of this operator is tractable:

$$K_w(t', f'; t, f) = \frac{2}{n} \left(\phi\left(\frac{t'-t}{\tau_+}\right) \phi\left(\frac{f'-f}{\eta_+}\right) \right)^2 \quad , \tag{C.5}$$

it covers $n = \frac{\tau_+}{\tau_-} = \tau_+ \eta_+ = \Delta t \Delta f$ atoms and its trace is $\text{Tr } K_w = 1$.

C.3. Marginal and limit cases of the canonical time–frequency estimates

We detail here the relationships between the canonical time–frequency estimates $S_{xx}(t, f; \tau_+, \tau_-)$ and $S_{xx}(t, f; Q_+, Q_-)|f|$ and simpler temporal, spectral, fully global or local quadratic estimates.

Maximal statistics: no resolution. The time–frequency integration of the spectrogram and the scalogram yields the same global energy: $x^*x = \iint S_{xx}(t, f) dt df$. For the scalogram, the integrator $\frac{df}{|f|}$ cancels the factor $|f|$. Various average densities can be calculated by dividing this energy by the bandwidth, the duration or the area of the time–frequency domain. The global variance of a centred signal $x(t)$ corresponds to the mean power $P_{xx} = \frac{x^*x}{T}$.

The average over all time and frequency corresponds to the limits $\tau_+, Q_+ \rightarrow \infty$ and $\tau_-, Q_- \rightarrow 0$ (in fact, positive and negative frequencies are still distinct for $Q_- \rightarrow 0$ and should be summed). The resulting number of degrees of freedom available for the statistical analysis is maximal, $n \approx \frac{1}{2} f_s T$ for a real signal, and the joint assumption of stationarity and self-similarity is the strongest. In this limit, time–frequency representations are not relevant: there are no degrees of freedom left to obtain localization information.

Spectrally resolved statistics. A spectral marginal density can be obtained by averaging the spectrogram or the scalogram over time. In the window paradigm, it is known as the power spectral density or the periodogram and can be written as $S_{xx}(f; \tau)$. In the wavelet paradigm, the corresponding spectrum $S_{xx}(f; Q)|f|$ can be interpreted as the power log-frequency density. These purely spectral statistical estimators of $S_{xx}(f) = \mathbb{E}[\hat{x}(f)\hat{x}^*(f)]/T$ correspond to the limits $\tau_+, Q_+ \rightarrow \infty$ and to the assumption of a stationary signal. The spectral resolution Δf is still given by either τ_-^{-1} or $Q_-^{-1}|f|$, and the loss of temporal resolution increases the number of statistical degrees of freedom, bounded by the duration of the signal T : $n \approx T/\tau_-$ and $n \approx T|f|/Q_-$ respectively.

In practice, spectral marginal densities are directly computed from the squared time–frequency decomposition, by replacing the ensemble average in Eqs. (8), (15) by the time average, or by smoothing $\hat{x}(f)\hat{x}^*(f)/T$ along the (linear or logarithmic) frequency axis with a normal kernel.

Temporally resolved statistics. The converse temporal marginal density consists in averaging or integrating evolutionary spectra over frequencies. Called the power of the signal $P_{xx}(t; \tau)$ or $P_{xx}(t; Q)$, the integral over all frequencies of $S_{xx}(t, f)$ estimates $P_{xx}(t) = \mathbb{E}[x(t)x^*(t)]$, a time-dependent estimator of the variance for a centred stochastic process. The temporal marginal density is associated with a self-similarity assumption, analogous to the stationarity assumption [95] for the spectral marginal density. More precisely, the average over f of the spectrogram estimates an intensity per s^{-1} , assuming no correlation (white noise); the average over $\log f$ of the scalogram estimates the intensity per octave or decade (depending on the base of the logarithm), assuming it constant (pink noise). Regardless of their interpretation, both averages are proportional to the same variance estimator $P_{xx}(t)$.

This broadband case corresponds to the limits $\tau_-, Q_- \rightarrow 0$, except for the following subtlety. For the scalogram, the limit $Q_- \rightarrow 0$ does not integrate positive and negative frequency components, but keeps them separate. The resulting powers are those of the so-called analytic and anti-analytic signals $x^\pm(t)$, which carry distinct information when $x(t)$ is complex-valued. Their sum is the wavelet estimate of the power:

$$P_{xx}(t; Q) = \iint \hat{x}(f_1)\hat{x}^*(f_2)\hat{\psi}_Q^{\frac{1}{2}}\left(\frac{f_1}{f_2}\right)e^{i2\pi(f_1-f_2)t}df_1df_2, \tag{C.6}$$

with the log-normal kernel $\hat{\psi}_Q^{\frac{1}{2}} = \hat{\psi}_{Q/\sqrt{2}}$ and a number of degrees of freedom $n \approx 2Q$ that is independent of numerical timescales at first order! In contrast, the Gabor estimate $P_{xx}(t; \tau)$ boils down to smoothing $x(t)x^*(t)$ with a normal window of duration $\frac{\tau}{\sqrt{2}}$, yielding $n \approx \frac{1}{2}f_s\tau$ degrees of freedom.

Maximal resolution: no statistics. One could argue that the maximum time–frequency resolution is achieved with the “skeleton” of quadratic representations, namely the Wigner–Ville distribution, Eqs. (4), (3). This extreme case with unspecified atomic uncertainty, $n = 0$, leads to degenerate properties, such as a divergent variance or negative cross terms.

In the present framework, we consider that the highest resolution is achieved by the atomic time–frequency decomposition, $n = 1$. Without recomposition, the corresponding quadratic representations are the non-averaged spectrogram $S_{xx}(t, f; \tau, \tau) = \frac{\sqrt{2}}{\tau}G_x(t, f; \tau)G_x^*(t, f; \tau)$ and the non-averaged scalogram $S_{xx}(t, f; Q, Q) = \sqrt{2}QX(t, f; Q)X^*(t, f; Q)$, for which $\tau_+ = \tau_-$ and $Q_+ = Q_-$. This is the common conception of the spectrogram and the scalogram, but these transforms of the stochastic trajectory x cannot be considered as statistical estimators (at best trivial ones).

Appendix D. Algorithmic considerations

A time–frequency decomposition $X(t, f)$ based on the Gabor or wavelet transform is expressed in terms of convolution integrals, which can be computed numerically from a pair of fast Fourier transforms (FFT). In practice, the continuity of the time–frequency representations is rendered by oversampling with respect to the time–frequency resolution Δt and Δf or $\Delta \log f$. In both paradigms, we implement reflective boundary conditions by extending the time–frequency domain in order to reduce the often unrealistic effects of periodic boundary conditions. Zero padding can sometimes be more realistic, especially in the spectral direction for band-limited experimental recordings. Boundary effects often remain at a distance of $k = 1$ to 3 times the resolution from the initial and final times and spectral boundaries.

The Gabor transform is naturally represented as a matrix that linearly samples the time and frequency domains on a 2-dimensional grid. The spectrogram $S_{xx}(t, f; \tau_+, \tau_-)$ can thus be efficiently estimated, from Eq. (9) or Eq. (10), by a pair of FFT along one or the other dimension.

Since $\Delta \log f$ is constant for wavelets, frequencies are sampled geometrically (regularly on a log-scale). However, the resolution $\Delta t = Q|f|^{-1}$ suggests a time sampling rate that increases with the frequency. For this reason, a single matrix (sometimes called stationary wavelet transform) is not a parsimonious discrete representation. This problem can be mitigated at the cost of splitting the grid into multiple matrices (e.g. one per octave by downsampling time by a factor of 2 when halving frequency). Details of this sampling approach in the similar context of the S-transform can be found in [96].

From the wavelet decomposition, the canonical scalogram estimator is more easily obtained by a smoothing with respect to log-frequencies, Eq. (16) (one pair of FFT), as compared to the double temporal convolution with a non-separable kernel, Eq. (19). Therefore, the discretization on a single grid is convenient to compute this second step, although it relies on the non-parsimonious matrix representation. Unfortunately, this leads to an unreasonable memory consumption when working on a large data set, such as overnight polysomnographic signals.

In a previous paper [54], we have proposed an alternative scalogram estimator, whose numerical implementation is better suited to long time series. The recomposition step performs a smoothing in the time domain rather than in the log-frequency domain, with

a scaled normal window of duration $\tau = mQ|f|^{-1}$, see definition in Eq. (E.1). This scalogram estimator is compatible with a more parsimonious discretization since it implements the horizontal composition of atoms schematized in Fig. 1(C, left). However, the resulting estimator can only be considered as an approximation of the canonical scalogram estimator ($Q_+ \approx mQ, Q_- = Q$): the exact expression for the number of atoms n in terms of the quality factor of the wavelet and the scaling of the smoothing window is intractable; an approximation $n \approx (1 + m^2 \exp \frac{1}{2\pi Q^2})^{\frac{1}{2}}$ is computed in Appendix E using Lagrange’s method.

For the canonical scalogram estimator, the smoothing step along log-frequency complicates the search for a more parsimonious numerical implementation since it has to be repeated for each matrix with a sufficient overlap (several times the resolution $\Delta \log f = Q^{-1}$). There is still much room left for improving parsimony and efficiency with an arbitrary degree of continuity in the numerical implementation of the canonical scalogram (or time-scale representations in general). A discrete wavelet transform (such as the one in [97]) on the dyadic lattice followed by a switch to a continuous (arbitrarily oversampled) wavelet basis, or an analogous algorithm that produces local matrices (rather than coefficients) at each stage of the binary tree, may be potential ways of achievement beyond the FFT algorithmic strategy.

Appendix E. Laplace’s method for an unsolvable number of atoms

Given the operator K of a quadratic estimator x^*K^*x , the associated effective number of atoms (or eigenstates) n , as defined in Eq. (2), is not always as easy to compute as for the estimator of the canonical spectrogram or scalogram: $n = \frac{\tau_+}{\tau_-}$ or $\frac{Q_+}{Q_-}$. The discussion in Appendix D of an alternative scalogram estimator provides an example where the trace (integral) formula for n is intractable. We show in this case how to obtain, nevertheless, an accurate and analytical approximation of n using Laplace’s method.

E.1. Definition of the practical time-smoothing scalogram estimator

The practical scalogram estimator $S_{xx}(t, f; Q, m)|f|$, introduced in [54], is based on a time-smoothing over m times the wavelet duration. It is defined from the product of the wavelet decomposition XX^* similarly to the canonical spectrogram estimator Eq. (9):

$$S_{xx}(t, f; Q, m) = \frac{2}{m} \int X(t', f; Q)X^*(t', f; Q)w_{\tau(f)}^2(t' - t)dt', \quad \tau(f) = mQ|f|^{-1} \quad (E.1)$$

This hybrid estimator is characterized by a kernel that includes both a log-normal wavelet of quality factor Q and a normal window of duration $\tau(f) = mQ|f|^{-1}$. The kernel of the practical scalogram estimator $S_{xx}(t, f; Q, m)|f|$ takes the following form in the frequency domain:

$$\hat{K}(f_1, f_2; f, t) = \sqrt{2Q}\hat{\psi}_Q(\frac{f_1}{f})\hat{\psi}_Q^*(\frac{f_2}{f})(w_{\frac{1}{\tau(f)}}(f_1 - f_2))^{\frac{1}{2}}e^{-i2\pi(f_1 - f_2)t} \quad (E.2)$$

The corresponding number of atoms n , designed to be approximated by m , is defined in Eq. (2) from the ratio of the integral expressions:

$$\left(\int \hat{K}(v, v)dv \right)^2 = 2 \left(Q \int |\hat{\psi}_Q(v)|^2 dv \right)^2 = \lambda^{\frac{1}{2}}, \quad \lambda = e^{\frac{1}{2\pi Q^2}} \quad (E.3)$$

$$\iint |\hat{K}(v, v')|^2 dv dv' = 2Q^2 \iint e^{-2\pi Q^2(\log v)^2 - 2\pi Q^2(\log v')^2 - \pi(mQ)^2(v - v')^2} \Theta(v)\Theta(v')dv dv', \quad (E.4)$$

where we have set $v = f_1/f, v' = f_2/f$ and $t = 0$ for clarity, and without loss of generality. Since the second double integration is intractable, we compute its Laplace approximation.

E.2. Laplace’s approximation

Laplace’s method aims to approximate the integral of a positive complicated function by the integral of a Gaussian which has the same global maximum up to the second order (same peak value and curvature). Provided that the integrated function is doubly differentiable at a unique global maximum, the precision of the approximation depends on the existence of a sufficiently large exponent α , which ensures that significant contributions to the integral come from the neighbourhood of the maximum value of the integrated function. We apply the bivariate form of this approximation:

$$G(\alpha) = \iint e^{\pi\alpha g(x,y)} dx dy \approx \frac{2}{\alpha} \det[-\partial_x \partial_y g(x_0, y_0)]^{-\frac{1}{2}} e^{\pi\alpha g(x_0, y_0)}, \quad (E.5)$$

where (x_0, y_0) is the position of the global maximum of the function $g(x, y)$, $\partial_x \partial_y g(x_0, y_0)$ is its Hessian matrix evaluated at the maximum, and \det refers to its determinant.

Rewriting Eq. (E.4), we can identify the following quantities:

$$\iint e^{\frac{1}{4\pi Q^2} - 2\pi Q^2(\log v - \frac{1}{4\pi Q^2})^2 - 2\pi Q^2(\log v' - \frac{1}{4\pi Q^2})^2 - \pi(mQ)^2(e^{\log v - e^{\log v'}})^2} d \log v d \log v' = \lambda^{\frac{1}{2}} \beta^2 \iint e^{\pi\alpha g(x,y)} dx dy$$

$$\alpha = (mQ)^2 \lambda, \quad \beta = m\sqrt{\frac{\lambda}{2}}, \quad (x, y) = \beta^{-1}(\log(v, v') - \frac{1}{4\pi Q^2}), \quad g(x, y; \beta) = -x^2 - y^2 - (e^{\beta x} - e^{\beta y})^2.$$

The maximum is 0 at $(x_0, y_0) = (0, 0)$ and the Hessian matrix at the maximum is:

$$-\partial_x \partial_y g(x_0, y_0; \beta) = \begin{bmatrix} 2 + 2\beta^2 & -2\beta^2 \\ -2\beta^2 & 2 + 2\beta^2 \end{bmatrix}, \quad (\text{E.6})$$

with determinant $4 + 8\beta^2$. Therefore, we obtain:

$$\iint |\hat{K}(v, v')|^2 dv dv' \approx \lambda^{\frac{1}{2}} \beta^2 G(\alpha; \beta) = \lambda^{\frac{1}{2}} (1 + m^2 \lambda)^{-\frac{1}{2}}, \quad (\text{E.7})$$

so that Laplace's approximation of the number of atoms in the time-smoothing scalogram estimator $S_{xx}(t, f; Q, m)$ is:

$$n \approx (1 + m^2 \lambda)^{\frac{1}{2}}, \quad (\text{E.8})$$

for a sufficiently large value of $\alpha = (mQ)^2 \lambda$. Note the dependence of n on the quality factor Q (through λ).

References

- [1] M. Costa, A.L. Goldberger, C.-K. Peng, Multiscale entropy analysis of complex physiologic time series, *Phys. Rev. Lett.* 89 (6) (2002) 068102, doi: [10/dwvzzn](#).
- [2] J. Zschocke, R.P. Bartsch, M. Glos, T. Penzel, R. Mikolajczyk, J.W. Kantelhardt, Long- and short-term fluctuations compared for several organ systems across sleep stages, *Front. Netw. Physiol.* 2 (2022) doi: [10/gsnptq](#).
- [3] A. Bashan, R.P. Bartsch, J.W. Kantelhardt, S. Havlin, P.C. Ivanov, Network physiology reveals relations between network topology and physiological function, *Nature Commun.* 3 (1) (2012) 702, doi: [10/f4f8v7](#).
- [4] L. Faes, D. Marinazzo, F. Jurysta, G. Nollo, Linear and non-linear brain–heart and brain–brain interactions during sleep, *Physiol. Meas.* 36 (4) (2015) 683, doi: [10/grs9hk](#).
- [5] P.C. Ivanov, R.P. Bartsch, Network physiology: mapping interactions between networks of physiologic networks, in: G. D'Agostino, A. Scala (Eds.), *Networks of Networks: The Last Frontier of Complexity*, in: *Understanding Complex Systems*, Springer International Publishing, Cham, 2014, pp. 203–222, doi: [10/krst](#).
- [6] R.P. Bartsch, K.K. Liu, A. Bashan, P.C. Ivanov, Network physiology: How organ systems dynamically interact, *PLoS One* 10 (11) (2015) e0142143, doi: [10/f8b4vm](#).
- [7] P. Dayan, L.F. Abbott, *Theoretical neuroscience: computational and mathematical modeling of neural systems*, *Computational Neuroscience*, MIT Press, Cambridge (Mass.) London, 2001.
- [8] F. Battiston, G. Cencetti, I. Iacopini, V. Latora, M. Lucas, A. Patania, J.-G. Young, G. Petri, Networks beyond pairwise interactions: Structure and dynamics, *Networks beyond Pairwise Interactions: Structure and Dynamics*, *Phys. Rep.* 874 (2020) 1–92, doi: [10/ghqt4w](#).
- [9] R. Carmona, W.-L. Hwang, B. Torr sani, *Practical time-frequency analysis: Gabor and wavelet transforms, with an implementation in S*, Academic Press, 1998.
- [10] B. Ricaud, B. Torr sani, A survey of uncertainty principles and some signal processing applications, *Adv. Comput. Math.* 40 (3) (2014) 629–650, doi: [10/gsvpdj](#).
- [11] A. Arneodo, E. Barcy, F. Argoul, J. Elezgaray, J.F. Muzy, *Ondelettes, Multifractales et Turbulences (de l'ADN Aux Crossances Cristallines)*, Diderot  diteur, arts et sciences, Paris, 1995.
- [12] A. Arneodo, B. Audit, E. Bacry, S. Manneville, J.F. Muzy, S.G. Roux, *Scale invariance and beyond: what can we learn from wavelet analysis? in: Scale Invariance and beyond*, Springer, 1997, pp. 37–51.
- [13] E. Gerasimova, B. Audit, S.G. Roux, A. Khalil, O. Gileva, F. Argoul, O. Naimark, A. Arneodo, Wavelet-based multifractal analysis of dynamic infrared thermograms to assist in early breast cancer diagnosis, *Front. Physiol.* 5 (2014) doi: [10/gjvnrh](#).
- [14] G.-Q. Zhang, L. Cui, R. Mueller, S. Tao, M. Kim, M. Rueschman, S. Mariani, D. Mobley, S. Redline, The National Sleep Research Resource: Towards a sleep data commons, *J. Am. Med. Inform. Assoc.* 25 (10) (2018) 1351–1358, doi: [10/gdntmm](#).
- [15] S.F. Quan, B.V. Howard, C. Iber, J.P. Kiley, F.J. Nieto, G.T. O'Connor, D.M. Rapoport, S. Redline, J. Robbins, J.M. Samet, *The sleep heart health study: Design, rationale, and methods*, *Sleep* 20 (12) (1997) 1077–1085.
- [16] D. Gabor, Theory of communication. Part 1: The analysis of information, *J. Inst. Electr. Eng. Radio Commun. Eng.* 93 (26) (1946) 429–441, doi: [10/2rz](#).
- [17] D.B. Percival, A.T. Walden, *Spectral Analysis for Physical Applications: Multitaper and Conventional Univariate Techniques*, Cambridge University Press, Cambridge ; New York, NY, USA, 1993.
- [18] M. Bayram, R. Baraniuk, Multiple window time-frequency analysis, in: *Proceedings of Third International Symposium on Time-Frequency and Time-Scale Analysis, TSTS-96*, IEEE, Paris, France, 1996, pp. 173–176, doi: [10/dzsqxh](#).
- [19] A.T. Walden, A unified view of multitaper multivariate spectral estimation, *Biometrika* 87 (4) (2000) 767–788, doi: [10/d8dxbd](#).
- [20] J.-M. Combes, A. Grossmann, P. Tchamitchian, *Wavelets: time-frequency methods and phase space proceedings of the international conference, marseille, france, december 14–18, 1987*, *Inverse Problems and Theoretical Imaging*, Springer, Berlin Heidelberg, 1989.
- [21] R. Bardenet, A. Hardy, Time-frequency transforms of white noises and Gaussian analytic functions, *Appl. Comput. Harmon. Anal.* 50 (2021) 73–104, doi: [10/gj8mts](#).
- [22] C. Torrence, G.P. Compo, A practical guide to wavelet analysis, *Bull. Am. Meteorol. Soc.* 79 (1) (1998) 61–78, doi: [10/bhbwhf](#).
- [23] B. Cazelles, M. Chavez, A.J. McMichael, S. Hales, Nonstationary influence of El Nino on the synchronous dengue epidemics in Thailand, *PLoS Med* 2 (4) (2005) e106, doi: [10/ftw9z8](#).
- [24] E.A.K. Cohen, A.T. Walden, A statistical study of temporally smoothed wavelet coherence, *IEEE Trans. Signal Process.* 58 (6) (2010) 2964–2973, doi: [10/fcngnq](#).
- [25] L. Stankovi c, S. Stankovi c, Wigner distribution of noisy signals, *IEEE Trans. Signal Process.* 41 (2) (1993) 956–960, doi: [10/brfr8](#).
- [26] J.W. Pitton, The statistics of time–frequency analysis, *J. Franklin Inst.* 337 (4) (2000) 379–388, doi: [10/cpzg52](#).
- [27] V. Susic, N. Saulig, B. Boashash, Estimating the number of components of a multicomponent nonstationary signal using the short-term time-frequency R nyi entropy, *EURASIP J. Adv. Signal Process.* 2011 (1) (2011) 125, doi: [10/dj9v3b](#).
- [28] D.L. Jones, T.W. Parks, A high resolution data-adaptive time-frequency representation, *IEEE Trans. Acoust. Speech Signal Process.* 38 (12) (1990) 2127–2135, doi: [10/bq9xjv](#).
- [29] W.J. Williams, M.L. Brown, A.O. Hero III, Uncertainty, information, and time-frequency distributions, in: F.T. Luk (Ed.), *San Diego, '91*, San Diego, CA, San Diego, CA, 1991, pp. 144–156, doi: [10/dt8pp9](#).
- [30] P. Flandrin, R. Baraniuk, O. Michel, Time-frequency complexity and information, in: *Proceedings of ICASSP '94. IEEE International Conference on Acoustics, Speech and Signal Processing*, iii, IEEE, Adelaide, SA, Australia, 1994, pp. III/329–III/332, doi: [10/bqmd3g](#).
- [31] L. Stankovi c, A measure of some time–frequency distributions concentration, *Signal Process.* 81 (3) (2001) 621–631, doi: [10/dmvq77](#).
- [32] R.G. Baraniuk, P. Flandrin, A.J. Janssen, O.J. Michel, Measuring time-frequency information content using the R nyi entropies, *IEEE Trans. Inform. Theory* 47 (4) (2001) 1391–1409, doi: [10/dcsvvh](#).

- [33] A. Papandreou-Suppappola, F. Hlawatsch, G.F. Boudreaux-Bartels, Quadratic time-frequency representations with scale covariance and generalized time-shift covariance: A unified framework for the affine, hyperbolic, and power classes, *Digit. Signal Process.* 8 (1) (1998) 3–48, doi: [10/d2rtwx](#).
- [34] R. Kronland-Martinet, J. Morlet, A. Grossmann, Analysis of sound patterns through wavelet transforms, *Int. J. Pattern Recognit. Artif. Intell.* 01 (02) (1987) 273–302, doi: [10/bfs8dp](#).
- [35] R. Stockwell, L. Mansinha, R. Lowe, Localization of the complex spectrum: The S transform, *IEEE Trans. Signal Process.* 44 (4) (1996) 998–1001, doi: [10/dj3sx5](#).
- [36] A. Grossmann, J. Morlet, Decomposition of Hardy functions into square integrable wavelets of constant shape, *SIAM J. Math. Anal.* 15 (4) (1984) 723–736, doi: [10/b9ddn3](#).
- [37] R.A. Altes, Sonar for generalized target description and its similarity to animal echolocation systems, *J. Acoust. Soc. Am.* 59 (1) (1976) 97–105, doi: [10/d2c3bb](#).
- [38] J.M. Lilly, S.C. Olhede, Generalized Morse wavelets as a superfamily of analytic wavelets, *IEEE Trans. Signal Process.* 60 (11) (2012) 6036–6041, doi: [10/gjvnrq](#).
- [39] H. Knutsson, C.-F. Westin, G. Granlund, Local multiscale frequency and bandwidth estimation, in: *Proceedings of 1st International Conference on Image Processing*, vol. 1, IEEE Comput. Soc. Press, Austin, TX, USA, 1994, pp. 36–40, doi: [10/fj4x5k](#).
- [40] D.J. Field, Relations between the statistics of natural images and the response properties of cortical cells, *J. Opt. Soc. Amer. A* 4 (12) (1987) 2379, doi: [10/fr7v9c](#).
- [41] P. Kovési, *Invariant Measures of Image Features from Phase Information* (Ph.D. thesis), University of Western Australia Crawley, Australia, 1996.
- [42] A. Guillet, *Revisiting Statistical Physics Tools for the Complex Time-Frequency Characterization of Physiological Signals* Thèse de doctorat, Bordeaux, 2022.
- [43] S. Chokroverty (Ed.), *Sleep Disorders Medicine: Basic Science, Technical Considerations and Clinical Aspects*, Springer New York, New York, NY, 2017, doi: [10/kzrq](#).
- [44] M.J. Prerau, R.E. Brown, M.T. Bianchi, J.M. Ellenbogen, P.L. Purdon, Sleep neurophysiological dynamics through the lens of multitaper spectral analysis, *Physiology* 32 (1) (2017) 60–92, doi: [10/f9gvdh](#).
- [45] P. Kovési, Good colour maps: How to design them, 2015, doi: [arXiv:1509.03700](#).
- [46] G. Buzsáki, A. Draguhn, Neuronal oscillations in cortical networks, *Science* 304 (5679) (2004) 1926–1929, doi: [10/bwtf9q](#).
- [47] R. Pool, Is it healthy to be chaotic? *Science* 243 (4891) (1989) 604–607, doi: [10/c3bdpm](#).
- [48] E. Tobaldini, L. Nobili, S. Strada, K.R. Casali, A. Braghiroli, N. Montano, Heart rate variability in normal and pathological sleep, *Front. Physiol.* 4 (2013) doi: [10/gk42cc](#).
- [49] E.J.C. de Geus, P.J. Gianaros, R.C. Brindle, J.R. Jennings, G.G. Berntson, Should heart rate variability be “corrected” for heart rate? biological, quantitative, and interpretive considerations, *Psychophysiology* 56 (2) (2019) e13287, doi: [10/gfhhcc](#).
- [50] F. Auger, P. Flandrin, Improving the readability of time-frequency and time-scale representations by the reassignment method, *IEEE Trans. Signal Process.* 43 (5) (1995) 1068–1089, doi: [10/dpmrhp](#).
- [51] I. Daubechies, J. Lu, H.-T. Wu, Synchrosqueezed wavelet transforms: An empirical mode decomposition-like tool, *Appl. Comput. Harmon. Anal.* 30 (2) (2011) 243–261, doi: [10/dnnqvt](#).
- [52] G. Thakur, E. Brevdo, N.S. Fučkar, H.-T. Wu, The synchrosqueezing algorithm for time-varying spectral analysis: Robustness properties and new paleoclimate applications, *Signal Process.* 93 (5) (2013) 1079–1094, doi: [10/f4v6wh](#).
- [53] J. Xiao, P. Flandrin, Multitaper time-frequency reassignment for nonstationary spectrum estimation and chirp enhancement, *IEEE Trans. Signal Process.* 55 (6) (2007) 2851–2860, doi: [10/chkh7j](#).
- [54] A. Guillet, A. Arneodo, F. Argoul, Tracking rhythms coherence from polysomnographic records: A time-frequency approach, *Front. Appl. Math. Stat* 7 (2021) 624456, doi: [10/gkdfvx](#).
- [55] D. Iatsenko, P. McClintock, A. Stefanovska, Extraction of instantaneous frequencies from ridges in time–frequency representations of signals, *Signal Process.* 125 (2016) 290–303, doi: [10/gfzgx2](#).
- [56] D. Fourer, F. Auger, K. Czarnecki, S. Meignen, P. Flandrin, Chirp rate and instantaneous frequency estimation: Application to recursive vertical synchrosqueezing, *IEEE Signal Process. Lett.* 24 (11) (2017) 1724–1728, doi: [10/ggnrtb](#).
- [57] C.-Y. Lin, L. Su, H.-T. Wu, Wave-shape function analysis, *J. Fourier Anal. Appl.* 24 (2) (2018) 451–505, doi: [10/gc8dvx](#).
- [58] H.-T. Wu, E.Z. Soliman, A new approach for analysis of heart rate variability and QT variability in long-term ECG recording, *BioMed Eng. OnLine* 17 (1) (2018) 54, doi: [10/gjvpbj](#).
- [59] R. Carmona, W. Hwang, B. Torrèsani, Characterization of signals by the ridges of their wavelet transforms, *IEEE Trans. Signal Process.* 45 (10) (1997) 2586–2590, doi: [10/cdnsc8](#).
- [60] I. Daubechies, Y. Wang, H.-T. Wu, ConceFT: Concentration of frequency and time via a multitapered synchrosqueezed transform, *Phil. Trans. R. Soc. A* 374 (2065) (2016) 20150193, doi: [10/gmv3x6](#).
- [61] L. Su, H.-T. Wu, Extract Fetal ECG from Single-Lead Abdominal ECG by De-Shape Short Time Fourier Transform and Nonlocal Median, *Front. Appl. Math. Stat.* 3 (2017) doi: [10/gj2z39](#).
- [62] H. Omer, B. Torrèsani, Time-frequency and time-scale analysis of deformed stationary processes, with application to non-stationary sound modeling, *Appl. Comput. Harmon. Anal.* 43 (1) (2017) 1–22, doi: [10/f97d6p](#).
- [63] A. Meynard, B. Torrèsani, Spectral analysis for nonstationary audio, *IEEE/ACM Trans. Audio Speech Lang. Process.* 26 (12) (2018) 2371–2380, doi: [10/gh2r36](#).
- [64] A. Meynard, B. Torrèsani, Time-scale synthesis for locally stationary signals, in: *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing*, ICASSP, 2020, pp. 5820–5824, doi: [10/gsv5gh](#).
- [65] P.C. Ivanov, L.A.N. Amaral, A.L. Goldberger, S. Havlin, M.G. Rosenblum, Z.R. Struzik, H.E. Stanley, Multifractality in human heartbeat dynamics, *Nature* 399 (6735) (1999) 461–465, doi: [10/bk6r24](#).
- [66] V. Ponomarenko, M. Prokhorov, A. Karavaev, A. Kiselev, V. Gridnev, B. Bezruchko, Synchronization of low-frequency oscillations in the cardiovascular system: Application to medical diagnostics and treatment, *Eur. Phys. J. Spec. Top.* 222 (10) (2013) 2687–2696, doi: [10/f5ffpf](#).
- [67] N. Wessel, M. Riedl, J. Kurths, Is the normal heart rate “chaotic” due to respiration? *Chaos* 19 (2) (2009) 028508, doi: [10/cgshg3](#).
- [68] Y. Shiogai, A. Stefanovska, P.V.E. McClintock, Nonlinear dynamics of cardiovascular ageing, *Phys. Rep.* 488 (2) (2010) 51–110, doi: [10/dmwhvb](#).
- [69] J.A. Hirsch, B. Bishop, Respiratory sinus arrhythmia in humans: How breathing pattern modulates heart rate, *Am. J. Physiol.-Heart Circ. Physiol.* 241 (4) (1981) H620–H629, doi: [10/gjvpcd](#).
- [70] H. Krause, J.F. Kraemer, T. Penzel, J. Kurths, N. Wessel, On the difference of cardiorespiratory synchronization and coordination, *Chaos* 27 (9) (2017) 093933, doi: [10/gbz8v7](#).
- [71] T. Henriques, M. Ribeiro, A. Teixeira, L. Castro, L. Antunes, C. Costa-Santos, Nonlinear methods most applied to heart-rate time series: A review, *Entropy* 22 (3) (2020) 309, doi: [10/ggtqpg](#).
- [72] Y.M. Ishbulatov, T.S. Bibicheva, V.I. Gridnev, M.D. Prokhorov, M.V. Ogneva, A.R. Kiselev, A.S. Karavaev, Contribution of cardiorespiratory coupling to the irregular dynamics of the human cardiovascular system, *Mathematics* 10 (7) (2022) 1088, doi: [10/gswzfs](#).
- [73] J.P. Saul, R.D. Berger, M.H. Chen, R.J. Cohen, Transfer function analysis of autonomic regulation. II. respiratory sinus arrhythmia, *Am. J. Physiol. Heart Circ. Physiol.* 256 (1) (1989) H153–H161, doi: [10/gnm4tp](#).

- [74] L. Faes, G. Pinna, A. Porta, R. Maestri, G. Nollo, Surrogate data analysis for assessing the significance of the coherence function, *IEEE Trans. Biomed. Eng.* 51 (7) (2004) 1156–1166, doi: [10/dtrzcj](#).
- [75] R.J. Thomas, J.E. Mietus, C.-K. Peng, A.L. Goldberger, An electrocardiogram-based technique to assess cardiopulmonary coupling during sleep, *Sleep* 28 (9) (2005) 1151–1161, doi: [10/gf8nss](#).
- [76] M. Javorka, J. Krohova, B. Czippelova, Z. Turianikova, Z. Lazarova, R. Wiszt, L. Faes, Towards understanding the complexity of cardiovascular oscillations: Insights from information theory, *Comput. Biol. Med.* 98 (2018) 48–57, doi: [10/gdv4dd](#).
- [77] P.C. Liu, Wavelet spectrum analysis and ocean wind waves, in: *Wavelet Analysis and Its Applications*, vol. 4, Elsevier, 1994, pp. 151–166, doi: [10/krsw](#).
- [78] K. Gurley, T. Kijewski, A. Kareem, First- and higher-order correlation detection using wavelet transforms, *J. Eng. Mech.* 129 (2) (2003) 188–201, doi: [10/crqrgt](#).
- [79] A. Grinsted, J.C. Moore, S. Jevrejeva, Application of the cross wavelet transform and wavelet coherence to geophysical time series, *Nonlinear Process. Geophys.* 11 (5/6) (2004) 561–566, doi: [10/d6dhzv](#).
- [80] L. White, B. Boashash, Cross spectral analysis of nonstationary processes, *IEEE Trans. Inform. Theory* 36 (4) (1990) 830–835, doi: [10/bhxr6c](#).
- [81] G. Matz, F. Hlawatsch, Time-frequency coherence analysis of nonstationary random processes, in: *Proceedings of the Tenth IEEE Workshop on Statistical Signal and Array Processing* (Cat. No. 00TH8496), IEEE, Pocono Manor, PA, USA, 2000, pp. 554–558, doi: [10/dfqx44](#).
- [82] M. Orini, R. Bailon, L.T. Mainardi, P. Laguna, P. Flandrin, Characterization of dynamic interactions between cardiovascular signals by time-frequency coherence, *IEEE Trans. Biomed. Eng.* 59 (3) (2012) 663–673, doi: [10/b4rwmn](#).
- [83] N.R. Goodman, Statistical analysis based on a certain multivariate complex Gaussian distribution (an introduction), *Ann. Math. Stat.* 34 (1) (1963) 152–177, doi: [10/cgm7kq](#).
- [84] G. Carter, A. Nuttall, Statistics of the estimate of coherence, *Proc. IEEE* 60 (4) (1972) 465–466, doi: [10/ddpd4f](#).
- [85] E. Cohen, A. Walden, A statistical analysis of morse wavelet coherence, *IEEE Trans. Signal Process.* 58 (3) (2010) 980–989, doi: [10/bt572x](#).
- [86] I.M. Gelfand, A.M. Yaglom, Computation of the amount of information about a stochastic function contained in another such function, *Uspekhi Mat. Nauk* 12 (1) (1957) 3–52.
- [87] G. Mijatovic, R. Pernice, A. Perinelli, Y. Antonacci, A. Busacca, M. Javorka, L. Ricci, L. Faes, Measuring the rate of information exchange in point-process data with application to cardiovascular variability, *Front. Netw. Physiol.* 1 (2022) 765332, doi: [10/gsmrmc](#).
- [88] A.T. Walden, E.A.K. Cohen, Statistical properties for coherence estimators from evolutionary spectra, *IEEE Trans. Signal Process.* 60 (9) (2012) 4586–4597, doi: [10/gjvpmj](#).
- [89] M.T. Valderas, J. Bolea, P. Laguna, Fellow IEEE, R. Bailón, M. Vallverdú, Mutual information between heart rate variability and respiration for emotion characterization, *Physiol. Meas.* 40 (8) (2019) 084001, doi: [10/gtkkg3](#).
- [90] M.M. Grigg-Damberger, The AASM Scoring Manual four years later, *J. Clin Sleep Med.* 8 (3) (2012) 323–332, doi: [10/f998fk](#).
- [91] A.N. Vest, G. Da Poian, Q. Li, C. Liu, S. Nemati, A.J. Shah, G.D. Clifford, An open source benchmarked toolbox for cardiovascular waveform and interval analysis, *Physiol. Meas.* 39 (10) (2018) 105004, doi: [10/gmdh98](#).
- [92] T.-P. Le, P. Argoul, Continuous wavelet transform for modal identification using free decay response, *J. Sound Vib.* 277 (1–2) (2004) 73–100, doi: [10/dkdr47](#).
- [93] F. Topsøe, Some bounds for the logarithmic function, *Inequal Theory Appl.* 4 (2007) 137.
- [94] A. Guillet, A. Arneodo, P. Argoul, F. Argoul, Quantifying the rationality of rhythmic signals, in: P. Flandrin, S. Jaffard, T. Paul, B. Torresani (Eds.), *Theoretical Physics, Wavelets, Analysis, Genomics: an Interdisciplinary Tribute To Alex Grossmann*, in: *Applied and Numerical Harmonic Analysis*, Springer International Publishing, Cham, 2023, pp. 565–593, doi: [10/krsv](#).
- [95] P. Flandrin, P. Borgnat, P.-O. Amblard, From stationarity to self-similarity, and back: Variations on the Lamperti transformation, in: *Processes with Long-Range Correlations*, vol. 621, Springer Berlin Heidelberg, Berlin, Heidelberg, 2003, pp. 88–117, doi: [10/ctsgz8](#).
- [96] R.A. Brown, R. Frayne, A fast discrete S-transform for biomedical signal processing, in: *30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, IEEE, Vancouver, BC, 2008, pp. 2586–2589, doi: [10/cf7pkc](#).
- [97] I. Selesnick, R. Baraniuk, N. Kingsbury, The dual-tree complex wavelet transform, *IEEE Signal Process. Mag.* 22 (6) (2005) 123–151, doi: [10/fwhmtz](#).