

PhD Dissertation

**Spatio-temporal dynamics of spontaneous
neural activity during the transition from
slow oscillations to an awake-like state in
isolated cortical networks**

Román Arango ¹

Faculty of Science and Technology, Bournemouth University

August 2024

¹ **Supervisors:**

Dr Emili Balaguer-Ballester

Prof Mavi Sanchez-Vives

Für Tanja

Acknowledgements

I would like to thank my supervisors and Sanchez-Vives Lab. A debt of gratitude is also owed to my friends and family for their encouragement. Special thanks to Helena, without whom this thesis may never have been finished.

Abstract

Slow oscillations (SO) of neural activity occur spontaneously in the neocortex during functionally disconnected brain states, such as deep anaesthesia or slow-wave sleep. These consist in a near-1-Hz alternation of high- (Up) and low-responsive (Down) periods, and spatially propagate as a travelling wave. During the transition towards wakefulness, globally synchronous SO break down, giving rise to asynchronous awake-like states or micro-arousals. There is currently a shortage of quantitative techniques to detect such state transitions in neurophysiological recordings. Notably, the disentanglement of non-stationary oscillatory signals embedded in background colour noise bears significant challenges. Thus, the emergence of asynchrony from the SO is poorly understood; in particular, the role of the local cortical network in orchestrating this transition spatially.

We devised a new statistical procedure based on the time-frequency decomposition of multivariate signals to investigate the network activity dynamics during this liminal regime. Our method, built upon an original mathematical framework, proposes a new normalisation of the wavelet power components that allows quantifying the level of asynchronicity at every given instant.

We then applied the novel methodology to extracellular, multi-electrode-array recordings on acute brain slices. These slices exhibited robust SO and were subsequently subjected to neurochemical modulations for eliciting asynchronous awake-like states, emulating a transitional awaking regime. Our analyses show how, on the brink of wakefulness, an excited SO-like activity cohabits with periods of near-asynchrony. Their spatio-temporal interplay depends on the structure of the cortical network: the cortical layer dictates the firing rate intensity, while the cortical column sustains and confines the oscillation. Furthermore, we found that the surges of asynchrony are spatio-temporally correlated, suggesting that the switch from synchrony to asynchrony pervades the whole network.

Overall, our findings reveal for the first time that the local circuits have the mechanisms for alternating spontaneously between synchronous and asynchronous states, confirming the recently attributed role of the cortex in the transitions between sleep and wakefulness. Also importantly, our methodological approach enables future investigations on state transitions in a variety of neurophysiological contexts.

Contents

Acknowledgements	iii
Abstract	v
List of Acronyms	x
Overview and Objectives	1
1 Introduction	5
1.1 Slow oscillations in local cortical networks	6
1.1.1 Cortical oscillations and slow wave activity	6
1.1.2 Slow wave dynamics and functional implications	7
1.1.3 Functional anatomy of the neocortex	8
1.1.4 The transition from slow oscillations to wakefulness	10
1.2 Wavelet analysis on stochastic processes	12
1.2.1 The continuous wavelet transform (CWT)	12
1.2.2 The wavelet power spectrogram (WPSg)	14
1.2.3 Wavelet components of Gaussian processes	15
1.2.4 Inverse power law and fractional Gaussian noises (fGn)	17
2 Materials and Methods	19
2.1 Experimental model and set-up	20
2.2 Data and tools	22
2.3 Multi-unit activity estimation	24
2.3.1 Decomposing extracellular recordings	24
2.3.2 A refined estimation of the MUA	25
2.4 Wavelet analysis: practical aspects	28
2.4.1 Numerical estimation of the CWT	28
2.4.2 The cone of influence	30
2.4.3 An energy-preserving definition of the WPSg-coefficients	30
2.4.4 Local power spectra and the global power spectrum	32
3 Methodological Developments and Theoretical Results	35
3.1 Normalisation of the wavelet power spectrogram	36
3.1.1 Inverse power law of the background spectrum	36
3.1.1.1 Robust estimation of the background power law	37
3.1.1.2 fGn as a null model for the logMUA	40
3.1.1.3 Bibliographical review and discussion	42

3.1.2	Probability distribution of fGn's spectrogram	44
3.1.2.1	Variances of complex random variables	44
3.1.2.2	Delta-variance of fGn's wavelet components	45
3.1.2.3	fGn's WPSg-coefficients as a linear combination of χ^2	51
3.1.2.4	Approximation of the distribution	52
3.1.2.5	Summary and discussion	53
3.1.3	Definition of the normalised spectrogram coefficients	56
3.1.3.1	Mean-normalised WPSg-coefficients of fGn processes	56
3.1.3.2	Normalised coefficients of logMUA signals	57
3.1.3.3	Theoretical interpretation	58
3.1.3.4	Application and practical consequences	59
3.2	Detection of synchronous and asynchronous states	62
3.2.1	Definition of asynchronous states	63
3.2.1.1	A test to discard asynchrony	63
3.2.1.2	Interpretation of the test statistic	65
3.2.1.3	Parameters of the test and Monte Carlo procedure	67
3.2.1.4	Raw asynchronous periods	70
3.2.2	Band-restricted detection of oscillatory activity	72
3.2.2.1	A test to detect synchrony	72
3.2.2.2	Test parameters and Monte Carlo implementation	74
3.2.3	Multi-channel optimisation of significance levels	75
4	Data Analyses and Biological Results	83
4.1	Alternation of synchrony and asynchrony in the awaking-like state	84
4.1.1	Experimental model	84
4.1.2	A new method to discern synchrony from asynchrony	86
4.2	Spatio-temporal pattern of the alternation	90
4.2.1	Spatial correlation of the synchronicity tenor	90
4.2.2	Co-occurrence of asynchronous states	93
4.3	Columnar SO-like activity during AS synchronous periods	95
4.3.1	Columnar coordination in the frequency domain	95
4.3.1.1	Normalisation of the wavelet power spectrogram	95
4.3.1.2	Peaks and shape of the global spectra	98
4.3.2	Decomposition of the oscillation into sub-states	104
4.3.3	Duration of the synchronous sub-states	108
4.3.4	Variability of the synchronous sub-states	110
4.3.5	Profile asymmetry of sync-Up states	115
4.3.6	Propagation of sync-Up states during synchronous periods	117
4.3.6.1	Phase-coupling strength between oscillations	117
4.3.6.2	Inter-laminar propagation lag	119
4.3.6.3	Influence of the Up onset lag on the oscillation	121
4.4	Laminar dependence of asynchronous MUA on synchronous sub-states	126
4.4.1	Inter-laminar maximisation of MUA during asynchrony	126
4.4.2	Conservation of mean MUA across states	127
4.4.3	Sync sub-states' bounds for the asynchronous MUA	129

5 Discussion	135
5.1 Summary of findings	136
5.2 Methodology: theory and application	137
5.2.1 Estimation of the MUA	137
5.2.2 Limitations of previous state-detection approaches in AS	138
5.3 Biological results	140
5.3.1 Awakening-like state and infra-slow alternation	140
5.3.2 Oscillatory activity sustained by the cortical column	142
5.3.3 SO-like synchronous activity during the AS	143
Conclusions	147
A Supplementary Mathematical Details	149
A.1 On the admissibility condition of the reduced Morlet wavelet	149
A.2 Gaussian wavelet components: proofs	150
A.3 Kolmogorov-Smirnov distance	151
B Research Dissemination	155
Bibliography	157

List of Acronyms

acvf	auto-covariance function
a.e.	almost everywhere
AS	awaking-like state
CV	coefficient of variation
(e)cdf	(empirical) cumulative distribution function
CL	confidence level
COI	cone of influence
CQV	coefficient of quartile variation
CWT	continuous wavelet transform
ECR	extra-cellular recording
EEG	electroencephalogram
fBm	fractional Brownian motion
FC	fold-change
FFT	fast Fourier transform
fGn	fractional Gaussian noise
fMRI	functional magnetic resonance image
IG	infra-granular
i.i.d.	independent and identically distributed
IRLS	iteratively reweighted least-squares
KS	Kolmogorov-Smirnov
LFP	local field potential
MEA	multi-eletrode array
MUA	multi-unit activity
pdf	probability density function
PLV	phase-locking value
SD	standard deviation
SEM	standard error of the mean
SG	supra-granular
SO	slow oscillations
WPSg	wavelet power spectrogram

Overview and Objectives

While the cortical network dynamics of both slow wave sleep and wakefulness have been well studied, the transition between these two global brain states remains poorly understood. In this thesis we developed a novel methodology to investigate the local network dynamics on the verge of an awake-like state. For doing this, we analysed spatially distributed multi-unit activity signals recorded from cortical slices: an isolated system *in vitro*. Our findings show, for the first time, that the local circuits have the mechanisms for alternating spontaneously between synchronous and asynchronous states. Moreover, we reveal and describe their spatio-temporal profile.

The introductory Chapter 1 presents the biological context and provides some technical background. The chapter begins with an overview of cortical oscillations, with a focus on the slow wave activity in cortical networks, their dynamics and functional implications. We outline the cortical network anatomy relevant to our findings, which will involve the spatial transition of slow oscillations towards asynchrony in relation to the cortical organisation into layers and columns. We then summarise some underlying neuromodulatory mechanisms known to drive the state transition from slow oscillations to wakefulness. In a second part, we introduce the fundamentals of wavelet analysis on stochastic processes, since a central part of our methodological developments will build upon the theoretical investigation of the statistical properties of the wavelet components. Alongside, we also introduce the fractional Gaussian noise, a useful family of processes with an spectral inverse power law.

The second chapter summarises experimental materials and technical methods. This includes information about the experimental model used by Sanchez-Vives Lab which provides the data: namely, multi-array recordings in acute cortical slice preparations from ferret and the pharmacological manipulations the slices were subjected to. The Methods Section presents theoretical material of a technical nature to which we are frequently referring throughout this thesis. We have also there included some original developments of a lesser scope, such are a novel refinement to the multi-unit activity estimation, as well as practical aspects of wavelet analysis that we have deemed important.

In the following Chapter 3 we present our methodological developments, which constitute the core theoretical result of this thesis. We devised a theoretically motivated framework for the detection of synchronous and asynchronous periods of a certain type of non-stationary signals. This includes an investigation of the probability distribution of the wavelet power spectrogram of coloured Gaussian noises,

leading to our own definition of the normalised spectrogram coefficients. The meaningful extraction and normalisation of such coefficients is what ultimately allows us to develop the statistical procedure for decomposing the signals into synchronous and asynchronous states. Thus, our work draws on techniques coming from various fields, such as digital signal processing, wavelet analysis, probability theory and applied statistics; the whole in a biological context with the aim to analyse multi-channel extra-cellular electrophysiological recordings of cortical brain tissue. Owing to their technical nature, some sections of this chapter contain their own literature review and discussion—when they are not directly connected to the overarching thread of the thesis. Finally, it should be noted that, while in the present work we used our new methodology to investigate the dynamics of slow brain oscillations on the verge of wakefulness, our technique can virtually be applied to any sort of electrophysiological recording in order to capture transient oscillatory activity.

In Chapter 4 we apply our new analysis tools to novel electrophysiological recordings acquired in Sanchez-Vives Lab, and then report our biologically relevant findings. We describe the alternation of synchrony and asynchrony in what we have called the *awaking-like state*, with attention to its spatio-temporal pattern. In particular, we identified a columnar organisation for the synchronous states and characterised their oscillatory activity. Moreover, we revealed the laminar dependence of asynchronous MUA constrained by the synchronous sub-states. Alongside, in this chapter we also provide details concerning the methods specific to each particular data analysis and visualisation.

Finally, in Chapter 5 we discuss our results—both methodological and biological—in the context of the current literature. While key findings will have already been discussed within each of their respective chapters, here, their impact and the perspective they offer is reviewed in relation to the literature and our current understanding. First, we outline in which aspects our new method advances currently available methods. In a second part we discuss how the detection of an infra-slow alternation in an isolated *in vitro* system expands on current knowledge, as it demonstrates that the cortex is sufficient to sustain such oscillatory patterns without external inputs. In the last part, we discuss potential underlying mechanisms of the proposed mechanism for the generation of such oscillatory patterns.

Note to the reader An interdisciplinary thesis as this one hardly makes for an easy structure. The difficulty stems from having up to four kinds of *methods*: a) computational materials, b) techniques of quantification and visualisation to present the biological results, and c) technical background on which to develop d) our novel analysis tools. We have nonetheless striven to keep introductory, methodological, results and discussion sections apart. Where it has not been possible to achieve this—mostly for the sake of enabling a linear reading—, we have clearly indicated to which category the paragraph should be ascribed. Thus, the reader will find bits of methodology presented in the biological results part (so called *Methodological Notes* on pages 91, 102, 106, 109, 113 and 123); short literature reviews within the methodological development part (*eg* Sections 3.1.1.3 and 3.1.2.5); and even small original results enclosed in the Methods Section or the Appendices (as the refinement to the MUA estimation or the KS embedding of densities).

Objectives

- O1)** Propose a probabilistic framework for assessing the power of time-frequency components of non-stationary electrophysiological signals.
- O2)** Develop a statistical procedure for the detection of synchronous and asynchronous periods of activity in multi-channel neurophysiological recordings.
- O3)** Investigate whether awaking-like recordings of isolated cortical slices can be fully decomposed into synchronous and asynchronous time periods.
- O4)** Identify the spatial dynamics of both synchronous and asynchronous periods during the awaking-like state, according to the columnar and laminar organisation of the cortex.
- O5)** Characterise the oscillatory content and spatio-temporal profile of the synchronous periods of the awaking-like state.

Chapter 1

Introduction

1.1 Slow oscillations in local cortical networks

1.1.1 Cortical oscillations and slow wave activity

Neuronal networks in the brain are formed by closely interacting neurons that are spatially and functionally distributed across various anatomical areas. In the cortex, these networks or micro-circuits mostly consist of excitatory (*eg*, pyramidal) and inhibitory (*eg*, GABAergic interneurons) neuronal populations in varying heterogeneous proportions, and arranged in a recurrent network architecture [Somogyi *et al.* (1998)].

The coordinated electrical activity of the neurons within these cortical networks generates extracellular electro-magnetic fields which can rhythmically fluctuate. A major contributor to the production of such cortical oscillations is the interaction between excitatory and inhibitory neurons in a delicate balance (the so-called E/I balance) [McCormick *et al.* (2003), Shu *et al.* (2003), Denève and Machens (2016)].




Figure 1.1: The six major frequency bands recorded by EEG. The brain waves are classified by their different frequency ranges, in slow (<1 Hz), delta (1–4 Hz), theta (4–10 Hz), alpha (8–12 Hz), beta (15–30 Hz), and gamma (30–90 Hz) waves. (Illustration from [Pavan *et al.* (2022)].)

Brain oscillations are associated with brain states and underlie cognitive processes [Buzsáki (2006)]. Their existence was first demonstrated in the early 20th century by electroencephalography (EEG) recordings [Berger (1929)]. In fact, the coordinated and synchronous activity of neurons in certain brain areas with a specific anatomical structure (*eg*, neocortex, hippocampus) creates electrical fields susceptible of being detected with extracellular electrodes [Buzsáki *et al.* (2012)] [Brette and Destexhe (2012)]. Several frequency bands of cortical oscillations were subsequently discovered and were traditionally classified as delta (1 to 4 Hz), theta (4 to 8 Hz), alpha (8 to 15 Hz), beta (15 to 30 Hz) and gamma (>30 Hz) [Buzsáki (2006)], with their respective attributed functional role to brain processes. In 1993, Steriade and colleagues described in four seminal papers a new cortical oscillatory pattern with a frequency of less than 1 Hz: the *slow oscillations* [Steriade *et al.*, 1993a–d]. Also known as slow wave oscillations, they were soon observed in human EEG recordings during sleep [Contreras and Steriade (1995)] (Figure 1.1).

Slow oscillations (SO) emerge spontaneously in absence of stimulation during functional disconnection of the network. They take place during behavioural quiescence of non-rapid eye movement (NREM) sleep, the so-called slow-wave sleep [Chauvette *et al.* (2011), Steriade *et al.* (2001)], or deep general anaesthesia [Chauvette *et al.* (2011), Lewis *et al.* (2012)]; but are also observable in less physiological situations as in deafferented cortical slabs [Timofeev *et al.* (2000)]. Notably, they have also been reproduced in cortical slice preparations *in vitro*, in which the self-organising properties of the SO have been extensively studied [Sanchez-Vives and McCormick (2000), Shu *et al.* (2003), Compte *et al.* (2003)]. At the cellular level, the SO is characterised by a step-like change in membrane potential showing spontaneous bistable fluctuations between two distinct states: the *Up* and *Down* state. The Up state shows periods of membrane depolarization and sustained firing bursts of action potentials, which alternates with the hyperpolarization of the membrane potential during the quiescent Down state [Sanchez-Vives and Mattia (2014)] (Figure 1.2).




Figure 1.2: Simultaneous LFP (**top**) and intracellular (**bottom**) recordings from the auditory cortex of the anaesthetized rat, exhibiting slow oscillations. Traces from [Sanchez-Vives and Mattia (2014)].

In short, SO are a multi-scale phenomenon that can be observed at the neuronal level (membrane potential), at the local network level (local field potentials and multi-unit activity) and at the whole brain scale (slow waves detected by EEG).

1.1.2 Slow wave dynamics and functional implications

The slow oscillations propagate across the cortex as waves, revealing functional properties of the underlying neuronal network structure. As for its spatial properties, the SO has long been known to behave as a slow travelling wave [Massimini *et al.* (2004)] (Figure 1.3), but only recently has a precise description of its spatio-temporal propagation been given from a local network perspective [Capone *et al.* (2019)]. Interestingly, the travelling Up/Down wavefront helps to reveal properties of the underlying network (*eg*, its laminar structure or its local level of excitability).

From a dynamical point of view, during the SO regime the network broadly behaves as a relaxation oscillator [Mattia and Sanchez-Vives (2012)]. That is, two alternating metastable attractors (Up and Down states) coexist and intrinsically fluctuate from one another, whether by means of deterministic activity-dependent

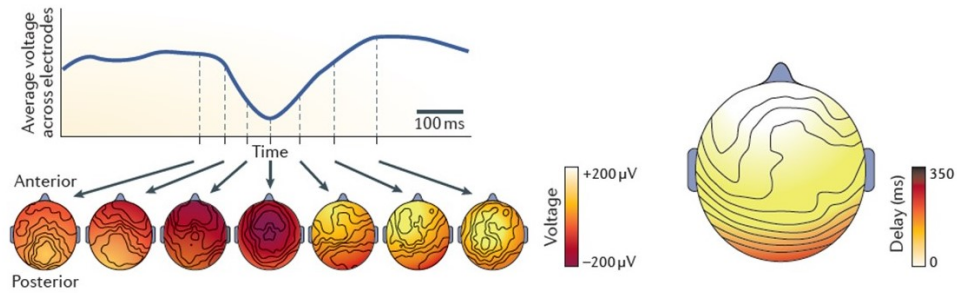


Figure 1.3: The slow oscillation during human sleep is a travelling wave moving along the anterior to posterior axis. The upper panel represents the time course of the SO dynamics, captured by EEG averaged across channels, showing a delay between the SO from anterior to posterior by several hundreds of milliseconds. Illustration from [Massimini *et al.* (2004)] (Copyright 2004 Society for Neuroscience).

adaptation, or elicited stochastically by endogenous noise [Sancristóbal *et al.* (2016)]. Furthermore, the Up/Down alternance seems to be broadly regular (≈ 1 Hz). However, the relaxation-oscillator approach has recently been challenged. In particular, the lack of robust regularity in the Up/Down's alternance undermines the role of the fatigue parameter in the termination of the Up states [Jercog *et al.* (2017)].

In summary, SO act as a low complexity state integrating neuronal, synaptic, and connectivity properties of the cortex. This offers a robust unifying paradigm for studying the implications of such properties in cortical functions.

Important cognitive functions have been assigned to the SO. It has long been known that sleep contributes to the consolidation of new memories. In particular, slow wave activity (SWA) after learning has been associated with enhanced retention of declarative memories. It has been shown that sleep homeostasis, which is regulated by an increasing amount of SWA after wakefulness which returns to baseline during sleep, can be triggered by a learning task involving specific brain regions. This suggests that the homeostatic process reflects synaptic changes that underlies a cellular need for sleep [Huber *et al.* (2004)] [Vyazovskiy and Harris (2013)]. Furthermore, an increase in SO-like potentials in humans by transcranial brain stimulations showed to enhance the retention of memories [Marshall *et al.* (2006)]. Moreover, alterations in sleep-related brain oscillations have been implicated in a number of brain disorders, including Alzheimer's disease [Byron *et al.* (2021)]. Thus investigation of the dynamics of SO and the underlying network mechanisms are not only crucial for revealing fundamental biological insights but may also be useful in future for developing translationally relevant treatment strategies.

1.1.3 Functional anatomy of the neocortex

In this work we are analysing recordings made with a multi-electrode array, placed on cortical slices across two spatial dimensions. One of the advantages of using this kind of recordings is that it enables us to analyse some dynamical aspects of the local network taking into account the anatomical organisation of the neocortex. Here we briefly summarize the main anatomical features of the cerebral cortex

which are relevant for our data analyses and interpretation of our results.

The structure and function of neuronal networks in the cortex have been widely studied [Douglas and Martin (2004)] [Thomson and Lamy (2007)]. The neocortex is organised horizontally in six well-defined layers and vertically into a myriad of columns.

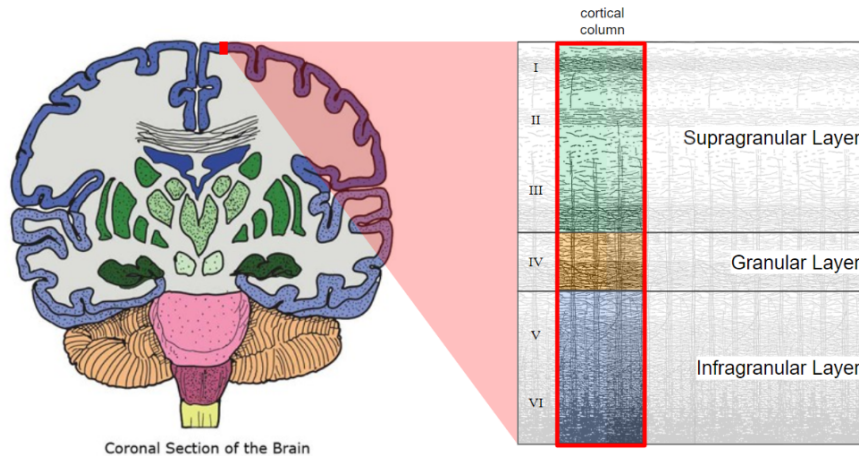


Figure 1.4: Organisation of the cortex into layers and columns. Layers I, II and III are together referred to as supra-granular (SG) layers, layer IV as granular layer, and layers V and VI as infra-granular (IG) layers. Neurons within a column often exhibit shared functional properties. Illustration from [Remmelzwaal *et al.* (2019)]

Shortly, somatosensory cortices, such as the primary visual cortex (V1), comprise layers I, II/III, IV, V, and VI, each with unique cellular compositions and functional roles, which contribute to various aspects of information processing. Layers I, II and III are together referred to as supra-granular layers (SG), layer IV as granular layer, and layers V and VI as infra-granular layers (IG) [Arion *et al.* (2007)] (Figure 1.4). The layers also vary in the inputs they receive and there are numerous inter-layer feedback projections. The layers receive distinct inter-laminar inputs with numerous inter-layer feedback projections, and endow long-range input and output projections with other brain areas. This connected hierarchical layer structure enables cortical information processing, which has been suggested to increase computational efficiency [Larkum *et al.* (2018)].

Furthermore, the cerebral cortex features an organization into anatomical columns [Mountcastle (1997)], where ensembles of neurons are arranged from the surface to the white matter. Neurons within a column often exhibit shared functional properties, such as stimulus orientation selectivity. Columns with varying properties, like different preferred orientations, form organized patterns across the cortical surface. While conventionally the identification of this functional architecture relied on the sparse sampling of microelectrode recordings (*cf* classical studies on receptive fields, *eg*, [Hubel and Wiesel (1963)]), a study using functional imaging with cellular resolution revealed the precise micro-architecture in the visual cortex, showing that columnar borders are organised at a fine scale of up to only one to two cells wide [Ohki *et al.* (2005)].

1.1.4 The transition from slow oscillations to wakefulness

SO are considered as the default global state of the cortical networks [Sanchez-Vives *et al.* (2017)]. Although SO are strikingly robust against external perturbations, the network can be driven into more complex dynamical states, inducing the transition from sleep to wakefulness. Thus the functional disconnection during the SO regime makes the network resilient to perturbances, yet facilitates the transition towards more complex, awake-like states [Massimini *et al.* (2005)] [D'Andola *et al.* (2017)].

Specifically, in the intact brain the transition from slow-wave sleep to wakefulness underlies a shift in the excitatory-inhibitory balance, facilitated by neuromodulatory changes and inputs from other brain areas. The awaking from anaesthesia *in vivo* has been used as an experimental model to investigate the process in which the UP/Down bi-stable configuration breaks down [Tort-Colet *et al.* (2021)]. In this study the transition from deep anesthesia towards wakefulness was described, in which slow oscillations gradually increase in frequency, and were followed by a sudden onset of infra-slow alternations between SO and active states. This transition was replicated with a population rate model, in which increased excitability drives the network to cross a critical point and gives rise to dynamical complexity through the two metastable attractor states.

Another recent model propose that UP states are akin to a circuit attractor that might implement functional computations, and thus could act as windows into consciousness [Destexhe *et al.* (2007)]. Indeed, action potential firing during Up-states resembles the temporal structure of wakefulness states in the range of beta (15 to 30 Hz) and gamma (>30 Hz) oscillations [Steriade *et al.* (1996)] [Destexhe *et al.* (2007)] [Compte *et al.* (2008)].

When the synchronised SO breaks down, a shift in global activity towards asynchrony can be observed. The awake state is characterised by irregular and decorrelated firing [Steriade *et al.* (2001)] [Poulet and Crochet (2019)] (see Fig. 1.5, right). More broadly, the investigation of SO propagation and its transition to wakefulness aims ultimately at revealing network mechanisms linked to the emergence of consciousness, and could give useful insights for better understanding of cognitive processes such as plasticity and memory consolidation.

SO in cortical slices and neuromodulation for state transition

It is generally thought that the Up-Down states are effectively generated inside the local cortical networks. This is because spontaneous Up and Down activity can be observed in the isolated network, such as cortical slices *in vitro*, and also *in vivo* under thalamic lesions [Steriade *et al.* (1993d)]. However, recent evidence also suggests a contribution of thalamic inputs for the Down-state transition via GABA_B-receptor-targeting interneurons *in vivo* [Hay *et al.* (2021)], which could bestow the SO with better regularity and resilience to external inputs [Briggs and Usrey (2008), Destexhe (2009)]. This should thus be taken into account when interpolating results that arise from isolated cortical slices.

Nevertheless, spontaneous SO are observed in isolated slice preparations, and neuromodulation can be used to drive the SO into an asynchronous *awake-like* state targeting synaptic dynamics [Steriade *et al.* (1993c)] [Jones (2005)]. Acetylcholine (ACh) is a neurotransmitters present at low concentrations during slow-wave sleep




Figure 1.5: Transitional Awaking-like State (AS) induced by neuromodulation. Extracellular recording traces from ferret cortical slices showing regular SO (left) and a desynchronized awake-like state (right). Note the downward deflection of the Up state in the extracellular recording. In the experimental protocol, the addition of neuromodulators facilitates the transition from one regime to the other, journeying through what we have called the Awaking-like State (AS).

and increases during wakefulness and REM sleep. ACh neurons discharge during the waking process, and release or application of acetylcholine, or carbachol (a cholinergic agonist), disrupts the SO and mimics arousal [Steriade *et al.* (1993c)] [Jones (2005)]. Furthermore, neurons expressing noradrenaline, histamine or orexin also fire during behavioural arousal, but are silent during REM sleep and show decreased activity during slow-wave sleep [Jones (2005)].

In summary, the SO regime and its transition to a more complex network state are in this sense an ideal workbench to explore some of the defining features of local cortical circuits. Although a number of mechanisms underlying these transitions are known, current analysis tools fail to discern empirically how cortical state dynamics precisely evolve when the SO breaks down. Thus, the development of new analysis methods is needed to give a better and more complete understanding into brain state dynamics (Fig. 1.5).

1.2 Wavelet analysis on stochastic processes

In this thesis we devise a new statistical procedure based on the time-frequency decomposition of multivariate signals in order to characterise the dynamics of brain state discussed in the previous section. Our method, which requires the development of an original mathematical framework, is based on the continuous wavelet transform and its power spectrogram. Wavelet analysis has been largely applied to analyse biophysical systems, including those in neuroscience. Due to its capacity to detect and analyse non-stationary signals, it has useful applications for data compression and visualisation, signal processing and feature detection [Hramov *et al.* (2015)]. While the mathematical developments of the first wavelets (from the French, *ondelettes*) is due to the seminal contributions of Grossmann, Morlet, Meyer, Morlet, Mallat and Daubechies [Daubechies (1992), Mallat (1998)].

Conventional wavelet analysis have found their main use in neuroscience as a mere visualisation tool. It is one aim of this thesis to go further and investigate the statistical properties of the wavelet components of certain neurophysiological signals (Chapter 3, section 3.1). We then propose a new normalisation of the power spectrogram, which allows us to quantify the level of synchrony or asynchrony of the signal at every given instant (Chapter 3, section 3.2).

In this introductory section, we give a minimal mathematical background of some basic concepts of wavelet analysis and stochastic processes upon which our methodological developments will rest. Alongside, we take the opportunity to introduce notational conventions that will be intensively used in subsequent chapters.

1.2.1 The continuous wavelet transform and the Morlet wavelet

The continuous wavelet transform (CWT) is, from a mathematical point of view, one of the most natural tools for the analysis of non-stationary signals. The ability of the CWT to act as a *zoom* both in the time and the frequency domains, makes it useful for detecting regime changes and extracting transient features of the signal.

The reduced analytical Morlet wavelet is given by:

$$\Psi(\tau) = \frac{1}{\sqrt{2\pi}} e^{-\frac{\tau^2}{2}} e^{i\omega_0\tau}, \quad \tau \in \mathbb{R}$$

That is, a complex exponential windowed by a Gaussian density function. As Ψ is not unit-energy,

$$\|\Psi\|_{L^2}^2 = \int_{-\infty}^{+\infty} |\Psi(\tau)|^2 d\tau = \frac{1}{\sqrt{4\pi}},$$

we will be working with its L^2 -normalisation version Ψ_0 instead:

$$\Psi_0(\tau) := \frac{\Psi(\tau)}{\|\Psi\|_{L^2}} = \pi^{-\frac{1}{4}} e^{-\frac{\tau^2}{2}} e^{i\omega_0\tau} \quad (1.1)$$

The angular frequency parameter ω_0 will be once and for all set to $\omega_0 = 6$, in order not to infringe badly (at least not numerically) the wavelet admissibility condition (see [Farge (1992)] p. 411 and Appendix A.1).

Dilations and translations of Ψ_0 give birth to its family of *daughter* wavelets

$$\psi_{s,t}(\tau) := \frac{1}{\sqrt{|s|}} \Psi_0\left(\frac{\tau-t}{s}\right), \quad t, s \in \mathbb{R}, \quad s \neq 0$$

Note that with the \sqrt{s} -scaling, we have ensured $\|\psi_{s,t}\|_{L^2} = \|\Psi_0\|_{L^2} = 1$.

Given a signal $x(t)$ of finite energy (ie, $\|x\|_{L^2} < +\infty$), its continuous wavelet transform with respect to Ψ_0 , at *time instant* t and *time scale* $s > 0$, is defined as:

$$\mathcal{W}x(s, t) := \langle x, \psi_{s,t} \rangle_{L^2} = \int_{-\infty}^{+\infty} x(\tau) \overline{\psi_{s,t}(\tau)} d\tau \in \mathbb{C}, \quad (1.2)$$

where $\langle \cdot, \cdot \rangle_{L^2}$ is the inner product of the underlying $L^2(\mathbb{R})$ Hilbert space and $\overline{\psi_{s,t}(\tau)}$ denotes the complex conjugate of the daughter wavelet.

Now, by simply noticing that $\psi_{s,t}(\tau) = \psi_{-s,0}(t - \tau)$, and calling $\phi_s := \overline{\psi_{-s,0}}$ it will be useful for computational purposes to read (1.2) as a convolution:

$$\mathcal{W}x(s, t) = \int_{-\infty}^{+\infty} x(\tau) \phi_s(t - \tau) d\tau = (x * \phi_s)(t) \quad (1.3)$$

Remark 1.2.1 (The CWT as a phase-space representation of the signal).

The CWT may be thought of as a localization of the signal x in the following sense:

- In the time domain, for each instant t we have

$$t = \int_{-\infty}^{+\infty} \tau |\psi_{s,t}(\tau)|^2 d\tau$$

- In the frequency domain, the centre frequency of the mother wavelet [Mallat (1998)] is given by

$$\eta_0 := \frac{1}{2\pi} \int_0^{+\infty} \omega |\widehat{\Psi}_0(\omega)|^2 d\omega \quad (1.4)$$

where $\widehat{\Psi}_0(\omega) := \int_{-\infty}^{+\infty} \Psi_0(t) e^{-i\omega t} dt$ denotes the Fourier transform of Ψ_0 .

Likewise, the centre frequencies $\xi(s)$ of the daughters (corresponding to scales $s > 0$) follow. In fact, by noticing that $\widehat{\psi}_{s,t}(\omega) = \sqrt{s} \widehat{\Psi}_0(s\omega) e^{-i\omega t}$, it results:

$$\begin{aligned} \xi(s) &:= \frac{1}{2\pi} \int_{-\infty}^{+\infty} \omega |\widehat{\psi}_{s,t}(\omega)|^2 d\omega = \frac{s}{2\pi} \int_{-\infty}^{+\infty} \omega |\widehat{\Psi}_0(s\omega)|^2 d\omega \\ &= \frac{s}{2\pi} \int_{-\infty}^{+\infty} \frac{\omega'}{s} |\widehat{\Psi}_0(\omega')|^2 \frac{d\omega'}{s} = \frac{s}{s^2} \frac{1}{2\pi} \int_{-\infty}^{+\infty} \omega' |\widehat{\Psi}_0(\omega')|^2 d\omega' \\ &= \frac{\eta_0}{s}. \end{aligned}$$

And since the mother can also be viewed as its own daughter, $\Psi_0 = \psi_{1,0}$, her centre frequency plays the role of a scale-frequency conversion factor, $\eta_0 = \xi(1) \cdot 1$, with which to relate time scales s to Fourier frequencies ξ . In our case, for the Morlet wavelet and our choice of angular frequency parameter $\omega_0 = 6$, we will have $\eta_0 = \frac{\omega_0}{2\pi} \approx 0.955$.

Hence, $\mathcal{W}x(\frac{\eta_0}{\xi}, t)$ can be understood as a phase-space representation of x around the instant t and the frequency ξ (see also [Louis et al. (1997)], p.15).

1.2.2 The wavelet power spectrogram (WPSg)

As with Fourier analysis, one of the cornerstones of wavelet analysis is the ability to recover the analysed signal from its decomposition and this, without loss of energy in the way back. Namely, given an L^2 -normalised wavelet Ψ , for every signal x of finite energy, the following *wavelet reproducing formula*¹ holds (see eg [Daubechies (1992)] p. 24–25 for an elegant proof):

$$x = \frac{1}{C_\Psi} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} \langle x, \psi_{s,t} \rangle \psi_{s,t} \frac{1}{s^2} dt ds \quad (1.5)$$

provided the existence of the so-called *admissibility constant* C_Ψ , defined as

$$C_\Psi := 2\pi \int_{-\infty}^{+\infty} \frac{|\widehat{\Psi}(\omega)|^2}{|\omega|} d\omega < +\infty \quad (1.6)$$

A wavelet version of the *Plancherel's identity* also ensues:

$$\|x\|_{L^2}^2 = \frac{1}{C_\Psi} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} |\mathcal{W}x(s,t)|^2 \frac{1}{s^2} dt ds \quad (1.7)$$

That is, the CWT conserves energy (*modulo* a constant), in the sense that the energy of the signal is the *sum* of the energies of its localised wavelet components.

In our case, we will be dealing with an analytic wavelet (i.e., $\widehat{\Psi}_0(\omega) = 0$, for $\omega < 0$) and with real signals (thus, symmetrical Fourier energy spread). Therefore, by considering only positive scales $s > 0$, and with the change of variable from scales to equivalent frequencies $s = \frac{\eta_0}{\xi}$, Plancherel's identity reads:

$$\begin{aligned} \|x\|_{L^2}^2 &= \frac{2}{C_\Psi} \int_0^{+\infty} \int_{-\infty}^{+\infty} |\mathcal{W}x\left(\frac{\eta_0}{\xi}, t\right)|^2 \left(\frac{\xi}{\eta_0}\right)^2 dt \frac{\eta_0}{\xi^2} d\xi \\ &= \frac{2}{C_\Psi \eta_0} \int_0^{+\infty} \int_{-\infty}^{+\infty} |\mathcal{W}x\left(\frac{\eta_0}{\xi}, t\right)|^2 dt d\xi \end{aligned}$$

That is,

$$\|x\|_{L^2}^2 = \frac{2}{C_\Psi \eta_0} \int_0^{+\infty} \int_{-\infty}^{+\infty} \mathcal{P}_x(\xi, t) dt d\xi, \quad (1.8)$$

where $\mathcal{P}_x(\xi, t) := |\mathcal{W}x(\frac{\eta_0}{\xi}, t)|^2$ can thus be interpreted as a phase-space energy distribution of f and is therefore called its *wavelet power spectrogram* (see [Mallat (1998)] p. 87 for further details).

¹Equality in (1.5) can be understood as *weak convergence* in Hilbert spaces, or simply as *pointwise convergence*.

1.2.3 Wavelet components of Gaussian processes

If we are to draw statistical assessments over the wavelet components of an electrophysiological signal, it is essential to provide a probabilistic frame for the generation of these signals. In fact, by their seemingly “noisy” nature, electrophysiological signals are naturally prone to be modelled by stochastic processes.

Stochastic processes

A stochastic process is a family $\{X_t\}_{t \in I}$ of random variables on the same probability space $(\Theta, \mathcal{A}, \mathbb{P})$, indexed by the ordered set of indices I (in our case, as we will favour continuous processes, $I \subseteq \mathbb{R}$ will be an interval of time points). Thus, viewed as an application from $\mathcal{A} \times I$ into \mathbb{R} , a process maps each event $\theta \subseteq \Theta$ onto a signal $X(\theta, \cdot) = \{x_t\}_{t \in I}$, called a *path* or a *realisation*.² Depending on the context, an observed recorded signal will be thought of as a continuous realisation of an underlying process, or simply as a discrete sample of that realisation.

Wavelet-integrable processes

Some theoretical background is required for contextualising the mathematical developments that follow in ulterior sections of this thesis. The first step is to notice that there is a natural extension of the definition of the continuous wavelet transform so it can be applied to a stochastic process $X = \{X_t\}_{t \in I}$ over the time interval $I \subseteq \mathbb{R}$ and the probability space $(\Theta, \mathcal{A}, \mathbb{P})$. Thus, for each time point $t_0 \in I$ and time-scale $s > 0$, we may consider the corresponding wavelet component:

$$\mathcal{W}X(s, t_0) := \int_I X_t \psi_{s, t_0}(t) dt$$

where the equality is meant to hold in the L^2 sense. So, $\mathcal{W}X(s, t_0)$ will be (in principle) a complex random variable on that same probability space. The ensuing power components $|\mathcal{W}X(s, t)|^2$ will each then be real random variables.

However, let us note that the existence of a stochastic integral as that underlying the CWT is not straightforward. We need to impose conditions on the process so that, a) *almost surely*, its paths are “wavelet-integrable”; and b) the resulting stochastic integral is measurable. In our case, since the analysing wavelets are square integrable, for the claim to hold we need at least to assume that the process is jointly measurable and L^2 (ie, $\mathbb{E}(X_t^2) < +\infty, \forall t$). To further ease the case, we may also ask the process to be stationary and to admit a spectral representation with bounded density (see eg [Ash and Gardner (1975)] and [Gallager (2013)] for some context).

We begin stating a technical lemma which, among other things, will allow us to commute integrals over the sample space (ie, expectations) with integrals over the time interval.

²For a mathematical background on probability theory and stochastic processes, the reader may refer to standard textbooks on the topic, such as [Ash and Gardner (1975), Dudley (2002)].

Lemma 1.2.2. *Let $X = \{X_t\}_{t \in I}$ be an L^2 process with mean and auto-covariance functions μ and γ , respectively. Let $\psi, \phi : I \rightarrow \mathbb{C}$ be square-integrable functions such that the processes ψX and ϕX are L^2 -integrable. Then, if γ is continuous a.e. and $\mu\psi$ is integrable,*

$$\mathbb{E} \left(\int_I X_t \psi(t) dt \right) = \int_I \mu(t) \psi(t) dt$$

In addition, if $\mu \equiv 0$ and the functions $\sqrt{\gamma(t,t)}\psi(t)$ and $\sqrt{\gamma(t,t)}\phi(t)$ are integrable, then

$$\mathbb{E} \left(\int_I X_t \psi(t) dt \overline{\int_I X_t \phi(t) dt} \right) = \iint_I \gamma(t,t') \psi(t) \overline{\phi(t')} dt dt' < +\infty$$

In particular, by providing their resulting second order moments, the lemma gives necessary conditions for the integral defining the CWT to be an L^2 process. We will henceforth be only interested in stochastic processes that behave well enough to admit a wavelet decomposition. Thus, we say that a process is Ψ -wavelet-integrable whenever its CWT (with respect to a given mother wavelet Ψ) determines, for each available time and scale, well-defined random variables on the probability space underlying the process³.

Next, when considering the CWT of Gaussian processes, we have the following crucial result. We recall that a process $\{X_t\}_{t \in I}$ is said to be Gaussian if for any n natural and any $(t_1, \dots, t_n) \in I^n$, the random variables X_{t_1}, \dots, X_{t_n} are jointly Gaussian (see [Ash and Gardner (1975), p.19]).

Lemma 1.2.3. *Let $X = \{X_t\}_{t \in I}$ be a zero-mean Gaussian process with auto-covariance function γ ; and let Ψ be a complex wavelet. If X is Ψ -wavelet-integrable and the function $t \mapsto \sqrt{\gamma(t,t)}\Psi(t)$ is integrable, then for every time $t \in I$ and scale $s > 0$, the random variables $\Re(\mathcal{W}X(s,t))$ and $\Im(\mathcal{W}X(s,t))$ are jointly normally distributed with mean 0.*

That is, the wavelet components of a zero-mean Gaussian process are (complex) normal variables with mean zero.

³This we will assume to be the case for stationary fractional Gaussian noises.

1.2.4 Inverse power law and fractional Gaussian noises

Much of the signals arising in a variety of different biophysical contexts happen to show a $1/\xi$ -shaped power spectrum. In particular, the case is well known for neurophysiological signals (whether for LFP [Milstein *et al.* (2009)], EEG [Voytek *et al.* (2015)], fMRI [Bullmore *et al.* (2001), He *et al.* (2010)], or MEG [La Rocca *et al.* (2018)]), in what seems to hint at an underlying *scale-free* phenomenon generating them (see *eg* [He *et al.* (2010)], [Buzsáki and Mizuseki (2014)]). The use *a priori* of white noise's flat spectrum as the reference background against which to assess spectral components (see *eg*, [Chatfield (2004)]) would therefore be ill-advised for this kind of signals [Bullmore *et al.* (2001), La Rocca *et al.* (2018), Donoghue *et al.* (2020)].

In this thesis, we will say that a signal's power spectrum \mathcal{S} follows an *inverse power law*, whenever there exist a constant $\beta \geq 0$, a constant $K > 0$ and a broad-enough union of accessible frequency bands Ξ_β , wherein

$$\mathcal{S}(\xi) \approx K \xi^{-\beta}, \quad \forall \xi \in \Xi_\beta \quad (1.9)$$

The parameter β , known as *spectral slope*, controls for the power-per-frequency decay of the spectrum; whereas the constant K , that we may want to call *spectral intercept*, mostly reflects the total background variance or power of the signal, assuming a given spectral slope. Observe that this inverse power law model includes an implicit offset parameter C , *ie* $\mathcal{S}(\xi) = K \xi^{-\beta} + C$, that we are assuming to be zero, since high frequencies of physiological signals will asymptotically carry no power. In fact, $0 \approx \lim_{\xi \rightarrow \infty} \mathcal{S}(\xi) = C$, for $\beta > 0$.

Finally, let us point out that in our previous definition, we have introduced the region Ξ_β in order to allow for some flexibility when it comes to capture the background spectrum of the signal, *eg* the case of frequency bands where the spectrum would hardly be $1/\xi$ -shaped, or even the case of distinct regions where the spectrum would *follow* different power laws. Nevertheless, we will aim in principle at a unique maximal region within the interval of available frequencies.

A suitable model for coloured noise

Following [Wornell (1993)] and [Kasdin (1995)], in the present work by β -coloured noise we are ideally referring to a zero-mean, second-order stationary stochastic process whose power spectrum follows an inverse power law with spectral slope β (see *supra* section 3.1.1). A useful model for processes with such behaviour is the *fractional Gaussian noise* (fGn), first introduced in [Mandelbrot and van Ness (1968)]. The fGn is heuristically defined as the derivative of the so-called fractional Brownian motion (fBm), which is a zero-mean Gaussian process with stationary increments.

Since the fBm is nowhere differentiable in the L^2 sense, this definition of fGn is met with the same difficulties as that of the usual *white noise*: fBm being a generalization of Brownian motion, in the same way as fGn is one of (Gaussian) white noise. Even when following the phenomenological approach here favoured, the existence of a covariance function or a theoretical spectrum for fGn would not be guaranteed. In practice, most of the common inconsistencies can be dispelled when considering discrete representations of it, or ideally band-pass filtered processes resulting from it (for a good review on the topic see [Kasdin (1995)] and references therein).⁴

⁴Although there are consistent formalisations of time-continuous fGn grounded on the

Importantly, fBm are canonical instances of *self-similar* processes. A process is said to be statistically self-similar if there exists $H > 0$ such that for all $\lambda > 0$, $X_{\lambda t} \equiv \lambda^H X_t$. H is called the *self-similar* or *Hurst* parameter. (Conversely, for each $0 < H < 1$ there exists a fBm that is self-similar with parameter H .) Their associated stationary fGn processes will each have an auto-covariance function (*acvf*) which can be proven to behave, for lags big enough, as:

$$\gamma(\tau) \approx C_{\sigma^2, H} |\tau|^{2H-2} \quad (1.10)$$

where $C_{\sigma^2, H}$ is a constant depending on H and the variance $\sigma^2 < +\infty$ of the underlying Gaussian law the process is pointwise distributed to (see [Barton and Poor (1988), Wornell (1993), Abry *et al.* (1995)] for the details). Such *acvf* then satisfies the following scale-preserving property:

$$\forall \lambda > 0, \quad \gamma(\lambda\tau) = \lambda^{2H-2} \gamma(\tau) \quad (1.11)$$

Now, restricting ourselves to positively correlated fGn processes ($\frac{1}{2} < H < 1$), the Fourier spectrum $\mathcal{S}^{\mathcal{F}}$ can be obtained as:

$$\mathcal{S}^{\mathcal{F}}(\xi) := \int_{-\infty}^{+\infty} \gamma(t) e^{-i2\pi t\xi} dt = C_{\sigma^2, H} \frac{2 \sin(\pi(1-H)) \Gamma(2H-1)}{|2\pi\xi|^{2H-1}} \quad (1.12)$$

whence it can be seen that the spectral slope β of the fGn is directly related to the self-similarity parameter, namely $\beta = 2H - 1$ (see eq. 1.9).

In this study, we will only be concerned with processes whose slope would go from the purely uncorrelated *white* noise ($\beta = 0$) presenting a flat spectrum, to the so-called *pink* noise ($\beta = 1$). The case of anti-correlated coloured noises ($H < \frac{1}{2}$) will not be treated; nor that in which the spectral slopes would be greater than 1 (eg *red* noise with $\beta = 2$). Indeed, failing to be stationary, all the latter would even lack a properly defined theoretical spectrum.⁵

However, the remaining “*infra-pink*” noises ($0 < \beta < 1$), although stationary, present also problems of their own. As integrating over their $1/\xi^\beta$ -law spectrum attests,

$$\lim_{\xi_{\max} \rightarrow +\infty} \int_0^{\xi_{\max}} \xi^{-\beta} d\xi = +\infty \quad (1.13)$$

the spectral variance—or total power—of these processes is ill-defined: a feature known as *ultraviolet catastrophe*. Equivalently, this means there is no easy way to circumvent the divergence towards infinity of their *acvf* for small lags.

Taken together, the fGn is a model to handle with care. In the remainder we will nonetheless assume, all caveats considered, that for each $\beta \in]0, 1[$, there exists a stationary zero-mean Gaussian process with $1/\xi^\beta$ global wavelet spectrum and scale-preserving *acvf*. In our treatment, white noise ($\beta = 0$ or $H = \frac{1}{2}$) is considered as a limit case of fGn; as such, most of the results that follow will, directly or asymptotically, also apply to it.

theory of tempered distributions [Kuo (1996)], the present work falls short of resorting to them.

⁵Again, a definition involving low-pass estimates of the spectrum may be considered in practice. Interestingly, an encompassing theoretical definition does exist that makes use of the wavelet global spectrum [Flandrin (1989)].

Chapter 2

Materials and Methods

2.1 Experimental model and set-up

In this work, we analyse unpublished data obtained from Sanchez-Vives Lab. The experimental methodology and data acquisition done by the collaborator lab are akin to those described in previous studies [D’Andola *et al.* (2017), Capone *et al.* (2019), Barbero-Castillo *et al.* (2021)]. We shortly summarise them below.

Cortical Slices Preparation

Extracellular recordings were carried out on acute coronal cortical slices of the ferret’s primary visual cortex (V1). Under appropriate medium conditions [Sanchez-Vives and McCormick (2000)], the cortical slices were able to exhibit highly synchronized slow oscillations (SO). Then, a precise and timely combination of three pharmacological modulations together with the reduction of the medium temperature [Reig *et al.* (2010)] were applied in order to drive the SO towards a novel, largely asynchronous, *awaking-like state* (AS) [D’Andola *et al.* (2017), Barbero-Castillo *et al.* (2021)]:

- (i) addition of $0.5 \mu\text{M}$ Carbachol (CCh) + $50 \mu\text{M}$ Norepinephrine (NE)
- (ii) reduction of extracellular Calcium (to $0.8\text{--}0.9 \text{ mM}$)
- (iii) reduction of temperature (to $31\text{--}32^\circ\text{C}$)

Recording Set-Up




Figure 2.1: Experimental set-up. **Left:** 16-channel flexible multi-electrode array (MEA) used for the recordings [Illa *et al.* (2015)]. **Right:** MEA placed over a cortical slice. (A) cortex, (B) white matter, (C) infra-granular layers and (D) supra-granular layers (image from [D’Andola *et al.* (2017)], by permission of Oxford University Press).

The recording device consisted of a specifically-designed 16-channel SU-8-based flexible microprobe [Illa *et al.* (2015)], placed over an acute cortical slice so that half of the microelectrodes would lie at infra-granular cortical layers (resp. supra-granular). Thus, population activity will ideally be recorded at different layers across various cortical columns (see details in Figures 2.1 and 2.2; see also [D’Andola *et al.* (2017), Capone *et al.* (2019), Rebollo *et al.* (2021)], where the same set-up was used).




Figure 2.2: Nissl-stained ferret’s V1 cortical slice. Histologically determined cortical layers are shown together with the location of the microprobe electrodes after the recording experiment. The electrodes cover at least three distinct cortical columns and ideally sample infra- and supra-granular layers (figure from [Capone *et al.* (2019)], by permission of Oxford University Press).

Extracellular electrical field potentials were recorded, the raw signal being amplified (x100), high-pass filtered (≥ 0.1 Hz) and digitised at a sampling rate of 5 or 10 kHz (for all the details, see [Barbero-Castillo *et al.* (2021)]).

2.2 Data and tools

Dataset selection

We have analysed recordings carried out in 13 distinct slices coming from as many different animals. As indicated above, the slices first exhibited spontaneous slow oscillations before being subjected to the manipulations that induced the awaking-like state (AS). Recordings during the AS condition lasted at least 300 s, up to more than 3000 s. Local field potentials (LFP) and multi-unit activity (MUA) signals were obtained from the raw extra-cellular recordings (ECR) (See below in Section 2.3.) LFPs and MUAs of 16-channel recordings were visually inspected to check for possible artefacts or faulty channels. Wherever feasible, we chose from each recording a stable period as long as 354 s to ensure we could extract edge-effect-free wavelet spectrograms of 300 s. (This is to take into account the cone of influence associated to the wavelet; see details in Section 2.4.2, p. 30.) From the MUA signals, only the central 300 s of the above-mentioned periods were kept for downstream analyses. That is, overall we analysed 13 processed MUA signals, one per slice in the AS condition; each of 300 s duration¹, plus the corresponding margins for their time-frequency decomposition.

Structural grouping of the MEA channels

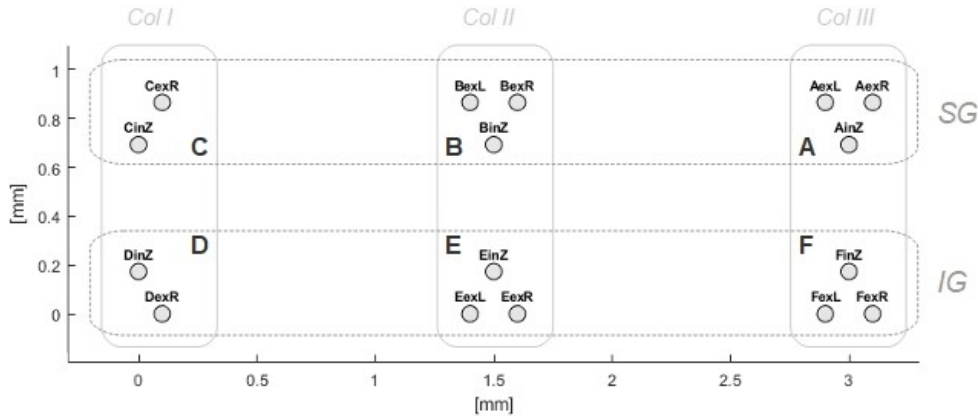


Figure 2.3: Channels' nomenclature and grouping. Channels are grouped according to the ideal spatial layout of the MEA probe on the slice. Three levels were considered for the analyses: the node (*A-F*), the column (*I-III*) and the relative position of the layer (IG or SG).

In order to comparatively characterise the recorded neuronal activity across the span of the microprobe, we grouped the channels according to the layout of the MEA. For this, we took into account the ideal location of the electrodes on the

¹*NB*: throughout the present work we will usually refer to these 300-s extracted periods as *the whole recording time* of the corresponding slice. No analysis carried out over the rejected recording periods was reported nor taken into consideration for multi-slice quantifications.

slice, as established by histological analysis in previous works using the same set-up [Capone *et al.* (2019), Rebollo *et al.* (2021)]. At the simplest level, electrodes are arranged into 6 groups comprising 3 or 2 electrodes each (2.1, left). These we have called *nodes*, denoted by *A* to *F*. Each electrode (*eg*, *CexR*) is named after the letter of the node to which it belongs, and its relative position within that node (*interior* or *exterior*; *Left*, *Right* or *Zentrum*). Next, owing to the anatomy of the cortex, those nodes that are *ca* 1.5-mm longitudinally apart will certainly lie on distinct cortical columns, and chances are that each group formed by transversally opposed nodes (*eg*, *E&B* or *F&A*) will sample activity from the same column. We have called them the MEA *columns* (from *I* to *III*). Finally, each MEA column will cover various depths of the anatomical cortical column, thereby recording at different cortical layers. In the experiment, care was taken to place the transversally lower electrodes at infra-granular (IG) layers (*ie*, below L4), and the upper ones at supra-granular (SG) layers (see Fig. 2.2). Even if that were not always achieved with the utmost exactitude, we could nonetheless be certain that, restricted to a single MEA column, each transversally opposed node does represent two distinct depths of the cortex laminar structure. With all, for simplicity we have further grouped the MEA nodes into two broad laminar categories: nodes *A* to *C* were deemed SG, whereas nodes *E* to *F*, were deemed IG. (See Fig. 2.3.)

Computational Tools

In this work, analysis and visualisation of data, numerical implementations and simulations were all performed in MATLAB (R2018b) [The MathWorks Inc. (2018)]. Some of our own functions and scripts made use or adapted the next specialised functions: `cwt` and `icwt` from the *Wavelet Toolbox*; `robustfit` from the *Statistics and Machine Learning Toolbox*; `randn` and `ColoredNoise` (with default `rng` parameters: `Mersenne Twister` and `seed 0`) from the *DSP System Toolbox*.

2.3 Multi-unit activity estimation

Throughout this work we will be analysing a special kind of signal called the Multi-Unit Activity (MUA). It is a processed signal obtained from raw extracellular recordings, and, for the purpose of our study, enjoys a number of advantages over the more commonly considered LFP. In this section, we present the MUA and describe a novel estimation procedure of it. It is one original contribution of this thesis to have provided such refined estimation, which, in keeping a reportable measure unit, has proven crucial for the analyses presented in Section 4.4.

2.3.1 Decomposing extracellular recordings

Extracellular recordings (ECR) of neuronal electrical activity are broadband rich. They can be decomposed into up to three signal types that are complementary and complexly interdependent: from the Local Field Potential (LFP), to the Multi-Unit Activity (MUA) and the Single-Unit Activity (SUA) [Teleńczuk and Destexhe (2014)].

LFP: electromagnetic field that results from the weighted average of slow neural events generated by current sinks and sources in the extracellular space. It is thought to mostly represent *afferent* neuronal activity (e.g. the summation of excitatory post-synaptic potentials) [Buzsáki *et al.* (2012)].

MUA: average of all the spikes available in the vicinity of a recording electrode. Ideally, it represents the local population firing rate and is, thereby, a proxy for the combined *effluent* neuronal activity [Pettersen *et al.* (2012)].

SUA: mere collection of those singled-out neuron's spikes directly distinguishable in the ECR.

(See Figures 2.4 and 2.5.)

The MUA as a proxy for the population firing rate. Whereas, on one hand, there are a wealth of spike-sorting procedures to obtain single units from high-frequency-sampled recordings [Gonzalo Rey *et al.* (2015)], and on the other hand, the LFP is assumed to be well captured by the low-frequency range of the ECR (*ca* <150 Hz), the MUA is a more elusive phenomenon and its estimation is subject to various approaches [Stark and Abeles (2007), Rasch *et al.* (2008), Mattia *et al.* (2010), Burns *et al.* (2010)]. The common ground of all these methods is to extract the MUA as a suitable average of the high-frequency band of the ECR (between 100 Hz and 10 KHz, see in figs. 2.4 and 2.5 what we have called *raw MUA*), under the assumption that these fast dynamics of the signal are mostly due to the firing of nearby neurons [Pettersen *et al.* (2008)]. Moreover, theoretical studies have argued that high-frequency Fourier power spectral components of the population firing rate are asymptotically proportional to the individual firing rates of the neurons constituting the local ensemble [Mattia and Del Giudice (2002)]. The MUA can thus be estimated as the relative change in power of the ECR's high-band [Mattia *et al.* (2010), Reig *et al.* (2010), Mattia *et al.* (2013)].

2.3.2 A refined estimation of the MUA

Taking the latter as a starting point, and building up on Mattia's procedure, we propose a slight modification to the MUA estimation. Its steps are as follows:

1. *Short-Time Discrete Fourier Transform.* For each non-overlapping sliding window I_t of width 5 ms, we compute

$$c_k(t) := |\hat{x}_k|^2, \quad k = 1, 2, \dots, k_{\max}$$

the discrete Fourier power coefficients of the ECR signal x over I_t . Only those coefficients corresponding to available ξ_k frequencies up to a high cut-off frequency ξ_{\max} are kept. That is, $\xi_k = k\Delta\xi$, where $\Delta\xi = 200$ Hz is the sampling step in the frequency domain, and so, $k_{\max} := \lfloor \frac{\xi_{\max}}{\Delta\xi} \rfloor$. Following [Reig *et al.* (2010)], we have taken $\xi_{\max} = 1500$ Hz.

2. *Quantile Normalisation.* In order to produce, for a given time t_0 , a balanced average of the $c_k(t_0)$, we first need to compensate for the tendency of the ECR's lower frequencies to carry more power (ECR's power spectrum follows indeed a so-called *inverse power law* [Milstein *et al.* (2009)]). To that aim, we normalise the coefficients $c_k(t)$ taking into account their relative behaviour over time, through what is called a *quantile normalisation*². Namely, for each k , the empirical cumulative distribution function (cdf) F_k is drawn from the values taken by each c_k over the whole recording period. Then, we obtain a consensus quantile function G , computed as the point-wise median of the family $\{F_k^{-1}\}$. That is, for each p -th quantile, $0 \leq p \leq 1$,

$$G(p) := \text{median}\{F_k^{-1}(p) : k = 1, 2, \dots, k_{\max}\}$$

Finally, the normalised coefficients are defined, for each instant t_0 , as

$$\tilde{c}_k(t_0) := G(F_k(c_k(t_0))) \quad (2.1)$$

thus all sharing a common cdf $\tilde{F} := G^{-1}$. Note that the quantile normalisation ensures that, in terms of physical units, $[\tilde{c}_k(t)] = [c_k(t)] = [x]^2$.

3. *Averaging.* Finally, the MUA is estimated as the average of the quantile-normalised power coefficients,

$$\nu(t) := \langle \tilde{c}_k(t) \rangle_k \quad (2.2)$$

In short, our procedure amounts to averaging, at each sampled time point t , the rank of each Fourier power coefficient relatively to its intrinsic range. Since averaging Fourier power coefficients amounts to estimating the variance of the demeaned windowed signal, our MUA estimate can be interpreted in terms of relative power allocated in the 200–1500 Hz band of the ECR signal; its units being in μV^2 .

²The quantile normalisation is widely used for the data analysis of gene-expression micro-arrays, where the disparity of each sample to be compared has prompted the development of this and similar homogenisation techniques (see e.g. [Bolstad *et al.* (2003)] for a review).




Figure 2.4: Signals obtained from a 15s sample of an extracellular recording from an infra-granular electrode during the SO regime. From **top to bottom**: ECR, LFP (low-pass filtered ≤ 200 Hz ECR), raw MUA (band-pass filtered 200–1500 Hz ECR), and MUA (as estimated according to our method)




Figure 2.5: Idem as above showing a recording during the *awaking-like* regime.

Advantages of our procedure. The MUA is a measure of an intensive property of the network, being an estimate of the local neuronal ensemble’s firing rate. It is therefore a suitable readout for the analysis of network states [Mattia *et al.* (2010), Pettersen *et al.* (2008)]. The theoretical basis of our estimation of the MUA closely follows that of other similar studies [Reig *et al.* (2010), Capone *et al.* (2019)]. However, contrary to theirs, in which MUA’s values were normalised across channels and deprived from a natural physical unit, our estimation procedure endeavours to make the comparison between different channels’ MUA more meaningful. This in particular is ensured by not overriding the original power unit when averaging over the local Fourier components. Thus, a detailed description of the values taken by the MUA across different electrode locations and during different time periods may prove highly informative (see more in the Discussion Section 5.2.1).

MUA vs logMUA For time-frequency analyses we rather considered the natural logarithm of the MUA (which is in practice strictly positive). We simply denote it by logMUA. In fact, scaling logarithmically the MUA condensed the information carried from high-frequency events compared to the plain MUA (as in *eg.*, [Reig *et al.* (2010), D’Andola *et al.* (2017), Capone *et al.* (2019)]). It was therefore more amenable to be analysed through its spectrogram. Furthermore, the logMUA is quasi normally distributed, which proved crucial for the normalisation of the wavelet power spectrogram (see below Sections 2.4 and 3.1). It is important to note that both the sync/async decomposition of the AS recordings (see Section 3.2 in Chapter 3) and the sync-Up sub-states detection (see methodological notes 4.3.2 in Chapter 4, Section 4.3.2) were done using the logMUA. However, all the analyses dealing with the intensity and variability of the signal were directly carried out over the plain MUA. When this distinction is not essential (*eg.*, for studying event propagation), we may sometimes refer indifferently to the MUA or the logMUA.

2.4 Wavelet analysis: practical aspects

The uses of the WPSg among the neuroscience community have been typically constrained to visualisation purposes and qualitative analyses, and seldom have there been any attempts to numerically exploit or statistically assess its yields.

We have therefore tried to compensate for this shortcoming by proposing a precise way of extracting wavelet power coefficients using the well-known *Morlet* wavelet, whereby these are easily interpretable as instantaneous time-frequency power components. Additionally, in this section we also provide technical aspects that will be required for subsequent methodological developments (see Sections 3.1, 3.2), as well as the data analyses carried out in Chapter 4.

2.4.1 Numerical estimation of the CWT

Signal sampling

In the present work we estimate wavelet transforms of logMUA signals $\nu(t)$ (see previous Section 2.3). These signals have been sampled, so they are finite in length (N samples) and discrete (sampling period $\Delta t = 5$ ms). Therefore, although essentially *power signals*, they will in practice be considered of finite energy \mathcal{E}

$$\mathcal{E}(\nu) = (\sigma_\nu^2 + \mu_\nu^2) N\Delta t$$

where σ_ν^2 and μ_ν denote their variance and mean, respectively. In order to minimise the impact of Ψ_0 having a non-vanishing mean, our signals will systematically be demeaned ($\mu_\nu = 0$). If $t_k = k\Delta t$, $k = 1, \dots, N$ are the time samples, let us call $\nu_k := \nu(t_k)$, the sampled logMUA signal during the interval $[0, T]$

Sampling in the frequency domain for efficient computation

There remains to select those time scales whereat to estimate the CWT components of the time series $\{\nu_k\}$. Not that the CWT could not be considered at whatever scale $s \neq 0$, but, on the one hand, we aim at interpreting the obtained CWT and the ensuing spectrogram in terms of classical Fourier power coefficients; and on the other hand, we want to avoid incurring burdensome computational cost that might lead to extracting redundant phase-space information of the signal.

The former will be ensured by the wavelet Plancherel's identity (eq. 1.8), as long as we obtain wavelet components at scales whose equivalent frequencies ξ lie within the range of theoretically-accessible non-zero Fourier frequencies:

$$0 < \xi_{\min} := \frac{1}{N\Delta t} \leq |\xi| \leq \xi_{\text{Nyquist}} = \frac{1}{2\Delta t}$$

In addition, by imposing an upper bound over the time scales, the approximation to the admissibility criterion should be fulfilled (see Remark A.1 below). Indeed, owing to the length N of our signals ($N \sim 10^5$), the left-tail truncation of $\widehat{\Psi}_0$ will be carried out at $\omega_{\min} := 2\pi\eta_0 \left(2^{\frac{1}{24}} s_{\max}\right)^{-1} \gtrsim 2\pi\xi_{\min} = 2\pi\frac{N}{\Delta t}$. With $\omega_0 = 6$ and the order of magnitude of N , this results in $\widehat{\Psi}_0(0) < 1.15 \cdot 10^{-8}$, which, for our numerical applications, will amount to having an admissible wavelet [Farge (1992)], specially when dealing with demeaned signals [Hramov *et al.* (2015), p. 32].

As for the efficient computation, by applying the Fourier convolution theorem to eq. (1.3) for each s , and by noticing that

$$\widehat{\phi}_s(\omega) = \widehat{\psi}_{-s,0}(\omega) = \overline{\widehat{\psi}_{-s,0}(-\omega)} = \sqrt{s} \overline{\widehat{\Psi}_0(s\omega)} = \sqrt{s} \widehat{\Psi}_0(s\omega)$$

where we had defined $\phi_s := \overline{\psi_{-s,0}}$ (see p. 13), we will have

$$\mathcal{F}[\mathcal{W}x(s, \cdot)] : \omega \mapsto \widehat{x}(\omega) \widehat{\phi}_s(\omega) = \sqrt{s} \widehat{x}(\omega) \widehat{\Psi}_0(s\omega)$$

(where \mathcal{F} denotes the Fourier transform operator, used only here for convenience instead of the *hat* notation). This allows for an efficient implementation using the FFT algorithm. However, its computational cost will still increase linearly with the number of scales required ($\sim J \cdot N \cdot \log_2(N)$, where J is the number of scales) [Hramov *et al.* (2015)]. Besides, since it transforms a signal of length N into $J \cdot N$ components, the CWT will necessarily yield redundant (*ie* correlated) information. Furthermore, the precision of the discrete decomposition at a given scale is determined by the window width of $\widehat{\psi}_{s,t}$, not the number of scales used to compute the components. These considerations have customary led to consider a logarithmically spaced set of time scales over which to extract the CWT components. Typically, powers-of-2-based geometric progressions are taken, with a certain number of steps (called *voices*) within each scale doubling (or *octaves*). Interestingly, it turns out to be particularly appropriate for analysing the logMUA signals, whose spectrum tends to follow an inverse power law, whereby only dynamic information at the octave-level precision shall be relevant (see [Percival (1995), Hramov *et al.* (2015)] and Section 3.1.1 in Chapter 3 below).

Time-frequency sampling mesh

According to the above considerations, we have estimated the wavelet components of the logMUA signal ν_k at the following time scales:

$$\begin{cases} s_0 := \eta_0 \xi_{\text{Nyquist}}^{-1} = \eta_0 2\Delta t \\ s_j := s_0 2^{\frac{j}{\mathfrak{v}}}, \text{ for } j = 0, \dots, J-1, \text{ with } J \leq \lfloor \mathfrak{v} \log_2 \left(\frac{N}{2} \right) \rfloor \end{cases} \quad (2.3)$$

where η_0 is the scale-frequency conversion factor of the wavelet (see remark 1.2.1 on p. 13) and the condition upon J ensures $2^{\frac{1}{\mathfrak{v}}} s_{J-1} \leq \eta_0 \xi_{\text{min}}^{-1}$. The *voices per octave* parameter \mathfrak{v} merely controls for the resolution in the frequency domain and has been set to $\mathfrak{v} = 24$ throughout our analyses.

For a given T , Δt and \mathfrak{v} as above, we will refer by *time-frequency sampling mesh* to the following set of sampled time points and frequencies:

$$\left\{ (\xi_j, t_k) : (j, k) \in \llbracket 0, J-1 \rrbracket \times \llbracket 1, N \rrbracket \right\} \subset \left[\frac{1}{T}, \frac{1}{2\Delta t} \right] \times [0, T] \quad (2.4)$$

where $\xi_j := \eta_0 s_j^{-1}$ are the sampled frequencies.

Finally, let us mention that any measure involving an integration across the frequency domain (*eg*, the estimation of a signal band power or the computation of the L^2 -distance between spectral densities) will be approximated as a sum over the support of logarithmically spaced time scales or frequencies. We will denote the frequency bins as

$$\Delta\xi_j := \xi_j - \xi_{j+1} = \xi_j \left(1 - 2^{\frac{1}{\mathfrak{v}}} \right) \quad (2.5)$$

2.4.2 The cone of influence (COI)

In practice, as in every frequency decomposition analysis, the lower the component's frequency aimed at, the longer the signal required. For finite signals, beyond the obvious bound on the minimum theoretically accessible frequency ξ_{\min} (see previous paragraph), there are in addition edge effects introduced by the analysing technique at the signal's abrupt vanishing ends.

In the case of the CWT, this comes to the study of the *cone of influence* (COI) of the wavelet. The COI depends on the particular wavelet used and is defined in the time-scale space (see [Mallat (1998)] p. 174), at each time instant t , as:

$$COI_{\Psi}(t) := \{ (\tau, s) \in \mathbb{R} \times \mathbb{R}_+^* : \tau \in \text{Supp}(\psi_{s,\tau}) \}$$

where $\text{Supp}(\psi_{s,\tau})$ denotes the compact support of daughters $\psi_{s,\tau}$ of the wavelet Ψ . For the Morlet wavelet, in numerical applications it is usual to assume that $\text{Supp}(\Psi_0) = [-\sqrt{2}, \sqrt{2}]$ (see [Hramov *et al.* (2015)] p. 34), and so

$$COI_{\Psi_0}(t) = \{ (\tau, s) \in \mathbb{R} \times \mathbb{R}_+^* : |\tau - t| < \sqrt{2}s \}$$

This allows us to determine a safe region where the CWT components will not be affected by the COI of end-of-signal singularities. For a finite signal sampled over a time interval I of length $T > 0$ and beginning at t_0 , such region $\mathcal{R}(I)$ can be described in the time-frequency space as:

$$\mathcal{R}(I) = \left\{ (t, \xi) \in I \times [\xi_{\min}, \xi_{\text{Nyquist}}] : \xi \geq \frac{\sqrt{2}\eta_0}{\frac{T}{2} - |t - (t_0 - \frac{T}{2})|} \right\}$$

This way, in order to extract CWT components of a signal that are free from edge effects over a time period J and up to the lowest desired frequency $\xi_0 \geq \xi_{\min}$, it is enough to compute the CWT over an interval \tilde{J} containing J , such that $J \times [\xi_0, \xi_{\text{Nyquist}}] \subseteq \mathcal{R}(\tilde{J})$. Taking $\tilde{J} = [\inf(J) - \tau_{\max}, \sup(J) + \tau_{\max}]$ does the job when $\tau_{\max} \geq \sqrt{2}\eta_0\xi_0^{-1}$. In our analyses we are considering time periods of length $T = 300$ s sampled at $\Delta t \approx 5$ ms (then $\xi_{\min} = 0.0033$ Hz, $\xi_{\text{Nyquist}} \approx 100$ Hz). As we wish edge-effect-free components down to frequencies $\xi_0 > 0.05$ Hz, we have set margins of $\tau_{\max} = 27.01$ s around.

2.4.3 An energy-preserving definition of the WPSg-coefficients

With these tenets, we are now in a position to propose our own working definition for the wavelet power spectrogram (WPSg), which is readily interpretable in terms of a variance decomposition of the signal.

At each sampled scale s_j ($j = 0, \dots, J - 1$) and instant t_k ($k = 1, \dots, N$), we define the *WPSg-coefficients* (WPSg-coeffs) of the signal ν as:

$$p_{j,k}^\nu := \frac{2}{C' \eta_0} \mathcal{P}_\nu \left(\frac{\eta_0}{s_j}, t_k \right) \quad (2.6)$$

where $\mathcal{P}_\nu \left(\frac{\eta_0}{s_j}, t_k \right)$ is the squared modulus of the wavelet transform estimate at the corresponding scale and instant (as in eq. 1.8); the modified admissibility constant C' and the wavelet scale-frequency conversion factor η_0 having both been previously defined (see eq. 1.4 and Appendix A.1).

Thus defined, our coefficients $p_{j,k}^\nu$ follow a straightforward sampling version of the *conservation of energy* formula:

$$\sum_{k=1}^N |\nu_k|^2 \Delta t \approx \sum_{k=1}^N \sum_{j=0}^{J-1} p_{j,k}^\nu \Delta \xi_j \Delta t \quad (2.7)$$

where $\Delta \xi_j$ are as in eq. (2.5). The left part is an estimation of the demeaned signal's energy $\mathcal{E}(\nu) = \sigma_\nu^2 N \Delta t$. The right part is a discrete approximation of the double integral of the *wavelet Plancherel's identity* (eq. 1.8) over the time-frequency sampling mesh (eq. 2.4). The relative error of the approximation—usually an underestimation of the variance—hovers at 1.5 – 3% for simulated Gaussian white noise of the same length and sampling frequency as the logMUA signals.

Although there have been similar attempts to define *ready-to-use* wavelet power coefficients in other fields [Terrence and Compo (1998), Shyun and Sun (2002), Büssov (2007)], they have tended to overlook the corrections required due to the reduced Morlet wavelet's non-compliance with the admissibility condition. Above all, the formulation we have here proposed seems in our view easier to interpret in terms of the time-frequency decomposition of the signal's variance and, thereby, of classical Fourier power components.

Energy, Power and Local Powers

From the conservation of energy equation (2.7), it follows that the WPSg-coeffs can be considered as *instantaneous frequency-local* power components, whose units would be in power-units per frequency unit. It will be useful at this point to clarify what we understand by *power* and *local power* of a signal in this context.

Power is *mean instantaneous energy*, or equivalently, variance—the signals having been systematically demeaned (see paragraph 2.4.1 above). Power is expressed in $[\nu]^2$ units (where $[\nu]$ denotes the physical unit of signal ν) and is usually estimated as an average of instantaneous energies across a time interval. Ideally, the instantaneous or *time-local* power at a given time is the energy carried by the signal at that instant by the whole broad band of available frequencies.

In the frequency domain, power can be naturally decomposed as *band-restricted* power allocated to a given frequency band. In the limit of such decomposition, *frequency-local* power can be defined as the power allocated to a single frequency, having thus power-units-per-frequency unit, or $[\nu]^2 \cdot \text{Hz}^{-1}$. Additionally, a *mean frequency-local* power over a frequency band of interest can be obtained by dividing the corresponding band-restricted power by that band-width in Hz.

Finally, WPSg-coeffs are to be understood as power components localized in time and frequency, and are therefore expressed in $[\nu]^2 \cdot \text{Hz}^{-1}$. Thus, averaging WPSg-coeffs over their time or frequency dimensions will provide different power summaries of the signal. On the one hand, by taking a time-average of WPSg-coeffs associated to a given frequency ξ_j , we end up having a frequency-local power estimate of ν at ξ_j : $\langle p_{j,\cdot}^\nu \rangle_k$ (expressed in frequency-local power units $[\nu]^2 \cdot \text{Hz}^{-1}$). On the other hand, by integrating the WPSg-coeffs at time t_k over a frequency band of interest Ξ_0 , we end up having the instantaneous power of ν at t_k over the band Ξ_0 : $\sum_{\xi_j \in \Xi_0} p_{j,k}^\nu \Delta \xi_j$ (expressed in power units $[\nu]^2$).

2.4.4 Local power spectra and the global power spectrum

It is customary to consider the *instantaneous* or *local power spectra* as the scale-wise slices of the wavelet-based energy decomposition. Namely, for each instant $t_k = k \Delta t$,

$$\mathcal{S}_\nu^{\text{loc}}(t_k) := (p_{0,k}^\nu, \dots, p_{J-1,k}^\nu) \in \mathbb{R}^J \quad (2.8)$$

from whose scale-wise integration, the *instantaneous power* of the signal results:

$$\Pi_\nu^{\text{loc}}(t_k) := \sum_{j=0}^{J-1} p_{j,k}^\nu \Delta \xi_j \quad (2.9)$$

expressed in power units $[\nu]^2$.

It is also usual to consider the *global wavelet power spectrum*, defined for each frequency ξ_j as the time-average of the local spectra's $(j+1)$ -th component,

$$\mathcal{S}_\nu(\xi_j) := \langle p_{j,k}^\nu \rangle_k = \frac{1}{N} \sum_{k=1}^N p_{j,k}^\nu \quad (2.10)$$

expressed in frequency-local units $[\nu]^2 \cdot \text{Hz}^{-1}$.

Interestingly, the global spectrum turns out to be a smooth approximation to the Fourier spectrum of the signal [Percival (1995), Perrier *et al.* (1995)].

Now, the time-averaged frequency-wise integration of the wavelet power coefficients provides the *total wavelet power* of the signal associated to the underlying sampling mesh:

$$\Pi_\nu := \frac{1}{T} \sum_{k=1}^N \sum_{j=0}^{J-1} p_{j,k}^\nu \Delta \xi_j \Delta t = \langle \Pi_\nu^{\text{loc}}(t_k) \rangle_k = \sum_{j=0}^{J-1} \mathcal{S}_\nu(\xi_j) \Delta \xi_j \quad (2.11)$$

From the conservation of energy formula, eq. (2.7), it follows that

$$\Pi_\nu \approx \sigma_\nu^2 \quad (2.12)$$

Thus, the total wavelet power may in some contexts also be called the *wavelet variance* [Flandrin (1992), Percival (1995)].

Let us finally mention that, by a slight abuse in the notation, we may often consider the WPSg-coeffs as a function p^ν of both scale (or frequency) and time, $p^\nu(s_j, t_k) = p^\nu(\eta_0 \xi_j^{-1}, t_k) = p_{j,k}^\nu$. This should not enclose any ambiguity: on the one hand, the definition of the coefficients is essentially independent from the sampling period Δt and the number of voices per scale; and on the other hand, frequency and scale are univocally related through the wavelet's scale-frequency conversion factor η_0 . It will thus indistinctly be to these coefficients or to the function p_ν that we will refer to when alluding to the wavelet power spectrogram.

For convenience, we may also drop the need to distinguish between *time-* and *frequency-local* when referring to powers or spectra whenever the context provides enough clarity: *eg*, if *local* is to be applied to a spectrum, it can only be *local in time*, since a spectrum by definition spans several frequencies; the same applies to a band power, restricted or not. Only will the distinction be held for WPSg-coeffs, since these can be averaged both in their time and frequency dimensions.

***NB:** A further number of methods specific to the data analyses and quantifications of Chapter 4 will be conveniently found therein under the banner **Methodological Notes** on pages 91, 102, 106, 109, 113 and 123.*

Chapter 3

Methodological Developments and Theoretical Results

In order to extract relevant time-frequency components from the logMUA signal, the natural question arises of how significant each of the components $p'_{j,k}$ of the WPSg (see definition in Section 2.4.3) is to be deemed. This we address through a novel normalisation of the WPSg against what we have called its *intrinsic background spectrum*. The purpose of such normalisation is two-fold: to draw a statistical assessment of the prominence of each WPSg-coefficient, independently of its time- and frequency-centre; to measure how far a given instantaneous power spectrum departs from *asynchrony* (in a sense that will be made clear later). Whereas the former will allow to compare different signals' spectrograms under sunder conditions, the latter will prove crucial to decompose the signal into synchronous and asynchronous periods (see Section 3.2 below).

3.1 Normalisation of the wavelet power spectrogram

In the intricate sections that follow, we give the theoretical foundations that motivate and justify our normalisation of the spectrogram. We first begin by observing that the logMUA signals under scrutiny follow an inverse power law, motivating, *inter alia*, their being modelled as fractional Gaussian noises (fGn). Alongside, we provide a simple and useful method to estimate their spectral parameters from the wavelet global spectrum. Then, we investigate how the wavelet power spectrogram of stationary fGn processes is marginally distributed for each time point and scale, offering a complete proof thereof. This will finally motivate, under certain modelling hypothesis we strive to state with clarity, our proposed normalisation for the WPGg-coefficients of logMUA signals.

3.1.1 Inverse power law of the background spectrum

In a crucial first step towards the spectrogram normalisation, a decision must be taken upon the reference distribution of power coefficients against which to compare the observed WPSg. For this, a reasonable way of gauging how the WPSg-coeffs are distributed is to inspect the corresponding global power spectrum \mathcal{S} of the signal, *ie*, the time-average of the coefficients across scales (see *supra* eq. 2.10 in section 2.4.4 of chapter 2). In fact, it can be shown that under mild stationarity assumptions, the global spectrum provides an ergodic estimate of the expectation of each coefficient's sample mean. Besides, \mathcal{S} being a smooth approximation of the signal's Fourier spectrum, it presents desirable features that in this context give it an edge over other classical non-parametric estimators [Wornell (1993), Abry *et al.* (1995), Perrier *et al.* (1995), Veitch and Abry (1999)] (see *infra*, section 3.1.1.1).

The inspection of the logMUA signals' spectra revealed that, notwithstanding the possible presence of prominent peaks—which eventually may or may not be attributable to persistent oscillatory activity—, these appear to follow an *inverse power law*.

In particular, we found that logMUA signals under the *awaking-like* regime typically exhibit background spectral slopes $\beta \approx 0.1\text{--}0.4$ (see full details in section 4.3.1 of chapter 4). This suggests that a mildly correlated noise could well serve as the underlying asynchronous stock out of which the signal is woven (see section 3.1.1.2 below).

3.1.1.1 Robust estimation of the background power law

In practice, we will estimate the underlying inverse power law of the logMUA signals by fitting the global wavelet spectrum \mathcal{S} in the log-log space to a simple linear model by means of a robust regression method.

That is, we begin by assuming that certain points of the spectra agree with the model

$$\log(\mathcal{S}(\xi_j)) = \beta_0 + \beta_1 \log(\xi_j) + \epsilon_j \quad (3.1)$$

where ϵ_j are *i.i.d.* normal error terms, and j is indexing consecutive sampled frequencies ξ_j within an *a priori* broad subinterval $\Xi_{\text{Reg}} \subseteq [\xi_{\min}, \xi_{\text{Nyquist}}]$. In our analyses, we let $\Xi_{\text{Reg}} = [0.1, 85]$ Hz, so as to avoid potentially problematic end-of-spectrum points¹ that would otherwise have a high regression leverage².

We then fit the parameters β_0, β_1 by *iteratively reweighted least-squares* (IRLS), a robust regression technique that prevents outlying observations from having a strong influence [Kutner *et al.* (2005)]. In short, for each iteration $k = 0, 1, \dots$, the model to be fitted is:

$$\log(\mathcal{S}(\xi_j)) = \beta_0 + \beta_1 w_{j,k} \log(\xi_j) + \epsilon_j \quad (3.2)$$

where $w_{j,k} \in [0, 1]$ are the weights for the k^{th} iteration. For $k = 0$, we set $w_{j,0} = 1$, and the initialisation model is given by eq. (3.1). Each iteration is then adjusted by ordinary-least-squares.

After each iteration, weights are updated according to a weight function of the residuals issued from the fitting. The role of the weight function is to penalise outliers, so that the greater the residual of an observation, the smaller the weight to be assigned to that observation at the next iteration. For our analyses, we have used the *Tukey bisquare* weight function w_{bsq} with tuning constant $c = 4.685$, defined as:

$$w_{\text{bsq}}(r^*) := \begin{cases} \left(1 - \left(\frac{r^*}{c}\right)^2\right)^2, & \text{if } |r^*| \leq c \\ 0, & \text{if } |r^*| > c \end{cases}$$

where r^* denote the scaled residuals of a fitting (see [Kutner *et al.* (2005)] sections 11.1 and 11.3 for the details).

Thus, after a number of iterations (the maximum number of which we have set to 50), both the weights and the residuals usually stabilise and a terminal pair of least-square estimates $\hat{\beta}_0, \hat{\beta}_1$ is kept. We take then

$$K = e^{\hat{\beta}_0} \quad \text{and} \quad \beta = -\hat{\beta}_1 \quad (3.3)$$

as the parameters governing the background spectral inverse power law in eq. (1.9), that we have called spectral intercept and spectral slope, respectively. Alongside them, terminal weights \hat{w}_j are also retrieved. Weights close to zero would typically correspond to frequencies where the power departs the most from the linear model (see Figure 3.1). This way, with a desired margin of tolerance $\varepsilon > 0$, a practical definition of Ξ_β could be obtained by means of

$$\Xi_\beta \cap \Xi_{\text{Reg}} = \{\xi_j \in \Xi_{\text{Reg}} : \hat{w}_j > \varepsilon\}$$

¹cf discussion on the cone of influence in section 2.4.2 of the previous chapter.

²cf *infra* remark 3.1.1 in this chapter.

In short, this simple strategy prevents eventual power spectral peaks from having an undue influence over the fitting of the underlying inverse power law tendency of the broad-band spectrum.




Figure 3.1: Inverse power law of logMUA signal's background spectrum estimated by iteratively reweighted least-squares (IRLS). **Top left:** global wavelet power spectrum of a logMUA signal plotted in a log-log space (grey curve). Dashed red line fits the background tendency of the spectrum. **Bottom left:** weights applied to the observations at each iterated regression. The darker and thicker the trace, the higher the weight. Observations corresponding to conspicuous spectral peaks have their weights rapidly converging to zero, indicating they strongly depart from the background dominant tendency. **Right:** Idem as top left for the global spectrum with its background fit showing its characteristic negative-exponential shape when plotted in a log-linear space: $\mathcal{S} \propto e^{-\beta \log(\xi)}$.

Remark 3.1.1 (A note on the balance of the regression leverage). *A simple measure of the influence a single observation has upon the fitting of a linear model over a set of observations is its leverage. The leverage values are defined as the diagonal elements of the so-called hat matrix of the model [Kutner et al. (2005)]. Let us recall that for a simple regression on J observations of the predictor-predicted variable pair (y_j, z_j) , $j = 1, \dots, J$, with design matrix*

$$Y = \begin{bmatrix} 1 & y_1 \\ 1 & y_2 \\ \vdots & \vdots \\ 1 & y_J \end{bmatrix}$$

the hat matrix H would be $H := Y(Y'Y)^{-1}Y'$. Thus, the leverage is solely expressed in terms of the predictor variable. Namely, it is easy to check that in this case the

leverage h_{jj} of the j^{th} observation satisfies:

$$h_{jj} = \frac{1}{J} \left(\frac{(y_j - \bar{y})^2}{y^2 - \bar{y}^2} + 1 \right)$$

where \bar{y} denotes the arithmetic mean of all the J observations y_k and \bar{y}^2 the mean of their squares y_k^2 . That is h_{jj} is a second order polynomial of y_j with vertex $(\bar{y}, \frac{1}{J})$. Observe also that the sum of the leverage values will here always be $\sum h_{jj} = 2$, which is the number of parameters that are being estimated. This results in a mean leverage value of $2/J$.

In our case, the role of predictor variable is played by the logarithm of the sampled frequencies $y_j = \log(\xi_j)$, whereas $z_j = \log(\mathcal{S}(\xi_j))$, the logarithm of the power spectrum at these frequencies, is the predicted variable (see eq. 3.1). The leverage measures how far the log-frequency of an observation would be of the (arithmetic) mean of the log-frequencies of the whole sample. It comes therefore as no surprise that, in order to fit the log-log linear model above (eq. 3.1), considering frequencies not logarithmically sampled would result in unbalanced leverage values.

Unbalanced in two senses: on the one hand, if $\bar{y} \neq \log(\sqrt{\xi_{\min} \xi_{\max}})$, the leverage polynomial would not reach its minimum near the centre of the span of the predictor variable observations (ie the mean log-frequency). This would entail having far from admissible leverage values³, not only on the fringes, but also on most of the available frequency band. On the other hand, and even if \bar{y} were the mean log-frequency, non-logarithmically sampled frequencies would be unevenly distributed with respect to the log-scaled frequency dimension. This would necessarily allocate more influence to one of the spectrum ends. For instance, having equispaced sampled frequencies from the frequency band Ξ_{Reg} will result in a fitting that would be highly biased by the low frequencies powers (see fig. 3.2).




Figure 3.2: Comparison between the leverage polynomial of two sampling supports for the linear model pitting the log-spectrum against the log-frequencies (as in eq. 3.1). In **blue**, leverage of observations issued from equispaced sampled frequencies; *idem* in **red** for log-spaced frequencies. Both examples are plotted onto two different coordinate spaces: \log_2 -scaled frequency axis on the **left**; linear frequency axis on the **right**. The examples are computed for $\Xi_{\text{Reg}} = [0.1, 100.2]$ Hz with $J = 239$ samples (24 voices per octave for the logarithmic case). Mean leverage value is then $2/239$.

³A leverage value 2 times the mean leverage is commonly deemed high, indicating the associated observation to have an unduly influence upon the regression [Kutner *et al.* (2005)].

3.1.1.2 fGn as a null model for the logMUA

The overarching motivation of the following theoretical developments is to determine the probability distribution of the WPSg coefficients of the logMUA under a certain null hypothesis \mathcal{H}_0 . Consequently, it is also one endeavour of the present work to provide a proper statement of such null hypothesis. We are heuristically guided by the assumption that the logMUA must harbour activity which is responsible for the inverse power law grossly followed by its spectrum. Thus, whether in the frequency or in the time domains, we seek to identify a component of the signal which would be a component of maximal stationarity. Maximal stationarity in the sense that the component might ideally enclose all the background activity that is not directly related to generate observable oscillations (transient or otherwise); and hence our naming it the *asynchronous* or *background* component. Therefore, in a first instance, it all comes down to finding a stochastic process that could faithfully render the asynchronous behaviour of the logMUA.

As we have seen (Introduction, Section 1.2.4), fractional Gaussian noises (fGn) are a canonical family of processes, which, for a certain range of their spectral slope parameter, are stationary and present a spectrum with an inverse power law. As such, one could think of them as maximally uncorrelated processes for a given spectral power law. In addition, their gaussianity makes them analytically and computationally tractable. So thus far, the fGn choice satisfies three important aspects to serve as a suitable null model to the logMUA: stationarity, spectral power law and tractability.

Now, it happens to be the case that logMUA values are more or less normally distributed. Although the conditions for a process to be Gaussian are rather restrictive⁴, in a first approach it can be indicative enough to check how collections of random samples of logMUA values of a given signal are distributed. Our preliminary investigations showed that logMUA sample densities were particularly well fitted by gamma distributions with very high shape parameters ($\alpha > 500$), thus approximating a normal distribution. This was specially verified when the samples were drawn from the same type of sub-state (*ie*, Down or Up in SO; sync-Down, sync-Up or asynchronous periods in AS⁵).

Wavelet-based accessible variance

As we previously noticed, any stationary process strictly observing an spectral inverse power law will be beset by the ultra-violet catastrophe, inherent to their non-integrable spectra for high frequencies (see Introduction, eq.(1.13) in Section 1.2.4). The opposite would occur for processes with very long-range auto-correlation dependencies, or, more broadly, with non-stationary ones. However, for our purposes, since the limiting high and low frequencies will respectively be constrained by the sampling period $\Delta t > 0$ and the total length T of the signals to be modelled, we will resort to what we call their *accessible variance*: a wavelet-based estimate of the variance in the time-frequency domain, and denoted, for a process X by $\hat{\sigma}_X^2$. In practice, an approximate value of it can always be obtained through the total

⁴There are a number of advanced statistical procedures to assess whether an observed realisation of a process may be deemed Gaussian, *eg*, the Epps test or the Lobato&Velasco test (see [Nieto-Reyes *et al.* (2014)] and references therein).

⁵The definition and characterisation of these sub-states will be extensively treated in Chapter 4.

wavelet power of any finite realisation $x = \{x_k\}_{1 \leq k \leq N}$ of the process over a given time-frequency mesh⁶. That is,

$$\hat{\sigma}_X^2 \approx \Pi_x = \sum_{j=0}^{J-1} \mathcal{S}_x(\xi_j) \Delta \xi_j \quad (3.4)$$

where Π_x and \mathcal{S}_x denote, respectively, the wavelet total power and the wavelet global spectrum of x , associated to the mesh (see Methods, Section 2.4.4).

Notice that $\hat{\sigma}_X^2 \approx \sigma_x^2$, by virtue of the formula of conservation of energy (eq. (2.7) in section 2.4.3).

Remark 3.1.2 (Accessible variance and spectral parameters of a fGn).

Given that for a fGn X with spectral slope β , $\mathcal{S}^{\mathcal{F}}(\xi) \propto \xi^{-\beta}$ (eq. (1.12)), we can draw a useful relation from the above approximation. In fact, the global spectrum \mathcal{S}_x of a realisation x provides estimates β_x and K_x of the inverse-power-law parameters β and K of the process (see section 3.1.1.3).

Since $\mathcal{S}_x(\xi) \approx K_x \xi^{-\beta_x}$, continuing from eq. (3.4) we can simply relate the accessible variance in terms of these slope and intercept:

$$\hat{\sigma}_X^2 \approx K_x \sum_{j=0}^{J-1} \xi_j^{-\beta_x} \Delta \xi_j = K_x \frac{1-\rho}{1-\rho^{1-\beta_x}} \left(\frac{1}{2\Delta t} \right)^{1-\beta_x} \left(1 - \rho^{J(1-\beta_x)} \right)$$

where we have denoted $\rho := 2^{-1/\mathfrak{v}}$, \mathfrak{v} being the number of voices per octave. Unsurprisingly, when the sampling in the frequency domain is refined, as $\lim_{\mathfrak{v}} \rho(\mathfrak{v}) = 1$, the right term of the previous equation will approach:

$$\frac{K_x}{1-\beta_x} \left(\left(\frac{1}{2\Delta t} \right)^{1-\beta_x} - \xi_J^{1-\beta_x} \right) = K_x \int_{\xi_J}^{1/2\Delta t} \xi^{-\beta_x} d\xi \approx K \int_{1/T}^{1/2\Delta t} \xi^{-\beta} d\xi$$

The final right term is in fact a band-restricted circumvention of the ultraviolet catastrophe (eq. 1.13).

In turn, for a small low cut-off frequency ξ_J , we get the next useful relation between the accessible variance of a fGn and the estimated spectral parameters of any of its finite realisations:

$$\hat{\sigma}_X^2 \approx \frac{K_x}{1-\beta_x} \left(\frac{1}{2\Delta t} \right)^{1-\beta_x} \quad (3.5)$$

Finally, let us also observe that even if the theoretical spectral variance of a fGn might remain elusive (ie, $\gamma(0) = +\infty$), we can still compare the total power of two fGn having the same spectral slope β through approximations of their respective spectral intercepts K, K' ; and this, whatever their high-frequency cut-off:

$$\forall \xi_{max} > 0, \quad \frac{\int_0^{\xi_{max}} K' \xi^{-\beta} d\xi}{\int_0^{\xi_{max}} K \xi^{-\beta} d\xi} = \frac{K'}{K}$$

This is how in practice we will control the targeted variance of a simulation of a β -fGn (see figure 3.3 below and section 3.2).

⁶In fact, we shall see that $\hat{\sigma}_X^2 := \mathbb{E} \Pi_X$, ie the expectation of the wavelet total power of the process.




Figure 3.3: Example of a logMUA signal (blue) and a simulation of a fGn (red) with spectral parameters estimated from the logMUA’s background ($\beta = 0.456$ and $\hat{\sigma}^2 = 0.358$).

3.1.1.3 Bibliographical review and discussion

Let us first point to a theoretical aspect of our approach that could potentially hamper the results. That is, the ability of the global wavelet spectrum to capture faithfully the inverse power law followed by the signal. In fact, owing to the correlations introduced by the CWT (both in time and between scales), it is not straightforward that the wavelet global spectrum should provide an accurate estimation of the theoretical spectrum of the underlying process generating the signal. As such, global spectra would tend to underestimate the spectral slope β and inflate the variance of the signal, unless the analysing wavelet has at least $\lceil \frac{\beta-1}{2} \rceil$ vanishing moments⁷ [Perrier *et al.* (1995)]. However, the original reduced Morlet wavelet has not even a 0th-order vanishing moment (this amounts to saying that it would not even be an admissible wavelet, and hence the correction proposed in remark A.1). Fortunately, with our choice of angular frequency parameter ($\omega_0 = 6$), higher orders of the Morlet wavelet (corrected or not) will remain very small in the context of our numerical applications. This would allow for the correct estimation of spectral slopes up to $\beta < 7$ (see [Farge and Schneider (2006)]), which in any case seems to be a bound tall enough for the observed logMUA spectral slopes.

Only a modest amount of previous studies in the field have also strived at inferring the intrinsic background spectra of neurophysiological signals, but with other purposes in mind and through a variety of methods different from ours; none of them were applied to MUA signals [Bullmore *et al.* (2001), Milstein *et al.* (2009), Voytek *et al.* (2015), La Rocca *et al.* (2018), Donoghue *et al.* (2020)]. For instance, to review the more relevant to us, in [Donoghue *et al.* (2020)] the Welch’s estimation of the power spectrum is considered to fit what the authors term the *a-periodic component* of LFP, EEG and MEG signals. Thus estimated, such spectra present a bias-variance trade-off that is particularly ill-adapted to the estimation of inverse power laws, resulting in a biased estimate of the spectral slope [Abry *et al.* (1995)]. Furthermore, and related to this last issue, the equispaced frequency support of those spectra will allocate a disproportionate leverage onto lower frequencies (see remark 3.1.1 *supra*). Besides, their technique also requires a great amount of supervised parametrization to the modelling of what they call the *periodic components* of the signals, and, in certain cases, resorts to *ad hoc* inter-trial smoothing of the spectra. On the contrary, the wavelet approach offers a natural advantage over Fourier-based estimates upon these three issues, *ie*, noise, bias and unbalanced leverage—the latter, owing to its log-scaled frequency

⁷Shortly, a wavelet Ψ is said to have $M \in \mathbb{N}$ *vanishing moments* if it is orthogonal to polynomials of degree $k < M$; or equivalently, if all the first $M - 1$ derivatives of its Fourier transform vanish at the origin (see *eg* [Mallat (1998)] p. 166).

support. Notably, it has also proven particularly robust for the estimation of a signal's inverse power law [Abry *et al.* (1995), Veitch and Abry (1999)] and is therefore a long-established approach in such fields as the physics of turbulence [Farge (1992), Farge and Schneider (2006)] or the theory of fractal stochastic processes [Flandrin (1992), Wornell (1993)]. Hence in another work [La Rocca *et al.* (2018)], the authors pointedly use the wavelet spectrum to infer the free-scale parameters of MEG signals. In this particular study, since the signals under evaluation are nearly scale-free (*ie*, they lack prominent spectral peaks), spectral slopes could be derived through simple linear regressions. In our case however, our richer signals with persistent intermingled oscillatory activity asked for a robust estimation, which nonetheless provides us with straightforward estimates without the need of much prior information.

3.1.2 Probability distribution of fGn's WPSg-coefficients

The principal aim of this section is to show that, under certain conditions which will be verified for our practical applications, the wavelet power spectrogram coefficients of *infra-pink* Gaussian noises are approximately χ_2^2 -distributed. For the case of the analytic Morlet wavelet we have been treating so far, an elementary and self-contained proof of this result is given. Importantly, we do so by introducing the delta-variance, a new measure associated to complex random variables. The delta-variance proves to be an invariant for Morlet components of fGn, on which the validity of the χ_2^2 approximation exclusively depends. Then we investigate aspects of its numerical approximation, and finally discuss the practical implications of this important result.

3.1.2.1 Variance, pseudo-variance and delta-variance of complex random variables

For a complex random variable Z with mean $\mu_Z \in \mathbb{C}$, two second central moments are customarily defined: the variance $\text{Var}(Z) := \mathbb{E}(|Z - \mu_Z|^2)$ and the *pseudo-variance* $\text{pVar}(Z) := \mathbb{E}((Z - \mu_Z)^2)$. It can easily be checked that:

$$\begin{cases} \mathbb{E}(\Re Z) = \Re(\mathbb{E}(Z)) & \text{and} & \mathbb{E}(\Im Z) = \Im(\mathbb{E}(Z)) \\ \text{Var}(Z) = \text{Var}(\Re Z) + \text{Var}(\Im Z) \\ \text{pVar}(Z) = \text{Var}(\Re Z) - \text{Var}(\Im Z) + i 2 \text{Cov}(\Re Z, \Im Z) \end{cases} \quad (3.6)$$

$\text{Cov}(\Re Z, \Im Z)$ being the covariance between the real and imaginary parts of Z . Notice that

$$\text{pVar}(Z) = 0 \iff (\text{Var}(\Re Z) = \text{Var}(\Im Z) \wedge \text{Cov}(\Re Z, \Im Z) = 0) \quad (3.7)$$

In particular, for Z zero-mean and normal, the latter amounts to $\Re Z$ and $\Im Z$ being independent, identically distributed normal variables.⁸ In such case, it is well-known that the squared module of Z will be proportionately χ^2 -distributed with 2 degrees of freedom. Namely,

$$\frac{|Z|^2}{\sigma^2} = \frac{(\Re Z)^2}{\sigma^2} + \frac{(\Im Z)^2}{\sigma^2} \sim \chi_2^2 \quad (3.8)$$

where $\sigma^2 = \frac{1}{2} \text{Var}(Z) > 0$ denotes the variance of $\Re Z$ or $\Im Z$ (which are the same). Were these variances not to coincide, the previous identity could nonetheless be generalised. For that purpose, let us introduce a quantity that will play an important role throughout our development.

Definition 3.1.3 (Delta-variance of a complex random variable). *We define the delta-variance of a complex variable Z with non-zero finite variance, as the absolute relative difference between the variances of its real and imaginary parts. We denote it by $\Delta\text{Var}(Z)$. That is:*

$$\Delta\text{Var}(Z) := \frac{|\text{Var}(\Re Z) - \text{Var}(\Im Z)|}{\text{Var}(\Re Z) + \text{Var}(\Im Z)}$$

⁸For a zero-mean complex normal variable Z , the condition $\text{pVar}(Z) = 0$ is also equivalent to Z being *circular symmetric*; that is, $e^{i\theta}Z \sim Z$, for all $\theta \in \mathbb{R}$.

Observe that, by definition, $\Delta\text{Var}(Z) \leq 1$; the equality being reached when $\Re Z$ or $\Im Z$ are *a.e.* constant. That will be the case if Z is real, of course. Note also that the delta-variance will be zero *iff* $\Re Z$ and $\Im Z$ have the same variance. The delta-variance must thus be understood as an index that measures how far away are the variances of the real and imaginary parts of a random variable.

We can then go on stating our generalisation of (3.8) as:

$$\frac{|Z|^2}{\frac{1}{2}\text{Var}(Z)} = (1 + \epsilon \cdot \Delta\text{Var}(Z)) \frac{(\Re Z)^2}{\text{Var}(\Re Z)} + (1 - \epsilon \cdot \Delta\text{Var}(Z)) \frac{(\Im Z)^2}{\text{Var}(\Im Z)} \quad (3.9)$$

with $\epsilon = \text{sgn}(\text{Var}(\Re Z) - \text{Var}(\Im Z))$. That is, for the particular case of a zero-mean complex normal variable with real pseudo-variance, its squared modulus can be expressed as the linear combination of two independent χ^2 variables with 1 degree of freedom. Furthermore, this sum will converge in distribution to a χ^2_2 variable when $\Delta\text{Var}(Z)$ tends to zero.

3.1.2.2 Study of the delta-variance of fGn's wavelet components

However, for wavelet components of fGn the delta-variance will never be zero, as we shall see. This will prevent the corresponding WPSg-coeffs from being χ^2_2 -distributed *sensu stricto*. It is nonetheless the aim of the remaining part of this section to show that, under certain conditions, the pseudo-variance of such components is real and their delta-variance can be kept reasonably small. For that to be the case, we will merely require that the spectral slope β of the fGn be not close to 1, as long as their components are computed over a long time interval and using a Morlet wavelet with enough vanishing moments (*ie*, parameter ω_0 big enough; on this, see *supra* footnote 7 on p. 42).

First, we need a lemma that relates the variances and pseudo-variances of fGn's wavelet components between different scales and time points. For convenience, from then on processes and their CWT will be considered over the whole real line ($I = \mathbb{R}$).

Lemma 3.1.4. *Let X be a fGn of Hurst parameter H ($1/2 < H < 1$). For any time points $t, t' \in \mathbb{R}$, and any scale $s > 0$,*

$$\begin{aligned} \text{Var}(\mathcal{W}X(s, t)) &= s^{2H-1} \text{Var}(\mathcal{W}X(1, t')) \\ \text{pVar}(\mathcal{W}X(s, t)) &= s^{2H-1} \text{pVar}(\mathcal{W}X(1, t')) \end{aligned}$$

Proof. It suffices to show the equalities for $t' = 0$.

From lemma 1.2.2, $\text{Var}(\mathcal{W}X(s, t)) = \iint_{\mathbb{R}^2} \gamma(\tau - \tau') \overline{\psi_{s,t}(\tau)} \psi_{s,t}(\tau') d\tau d\tau'$.

Recalling that $\psi_{s,t}(\tau) = s^{-1/2} \psi_{1,0}(\frac{\tau-t}{s})$, the variable change $(u, u') = (\frac{\tau-t}{s}, \frac{\tau'-t}{s})$ of Jacobian $|\mathcal{J}| = s^{-2}$ applied to the last integral results in $\iint_{\mathbb{R}^2} \gamma(s(u - u')) s^{-1/2} \overline{\psi_{1,0}(u)} s^{-1/2} \psi_{1,0}(u') s^2 du du'$.

Finally, owing to eq. (1.11), we have $\gamma(s(u - u')) = s^{2H-2} \gamma(u - u')$, and the equality follows. The proof for the pseudo-variance is analogous. \square

It follows from the previous lemma that both the variance and the pseudo-variance are scale-preserving with power $2H - 1$, which is precisely the spectral slope of the process ($\beta = 2H - 1$, as in eq. 1.12 p. 18).⁹ Notably, we can see that

⁹When referring to a fGn, we will hereafter favour specifying its spectral slope over its Hurst parameter.

the component's time-centre is irrelevant, as expected in light of the stationarity of the process. Besides, let us point that this result would hold for any wavelet-integrable, scale-preserving process, be it Gaussian or not. (Nor has the nature of the analysing wavelet played a role yet.) Here however, the case of zero-mean Gaussian processes would be more relevant: in fact, the lemma ensures that the distributions of the wavelet components of each of such processes are all the same up to a constant which is determined by their scale.¹⁰

Finally, the cornerstone technical result of this section.

Proposition 3.1.5 (Variances and delta-variance of fGn's CWT-components). *Let $X = \{X_t\}_{t \in \mathbb{R}}$ be a fGn of spectral slope β ($0 < \beta < 1$) and acvf γ . Let $W = \mathcal{W}X(s, t_0)$ be the CWT of X at $(s, t_0) \in \mathbb{R}_+^* \times \mathbb{R}$ with respect to Ψ_0 , the analytical Morlet wavelet of angular frequency parameter $\omega_0 > 0$. Then, the pseudo-variance of W is real and strictly positive. Furthermore, the delta-variance satisfies:*

$$\Delta \text{Var}(W) = \frac{\text{pVar}(W)}{\text{Var}(W)} = e^{-\omega_0^2} \frac{\int_{-\infty}^{+\infty} h(\tau) d\tau}{\widehat{h}(2\omega_0)} \quad (3.10)$$

where $h := \gamma|\Psi_0|^2$ and $\widehat{h}(\omega) := \int_{-\infty}^{+\infty} h(\tau)e^{-i\omega\tau} d\tau$ denotes its Fourier transform.

Proof. We proceed with $s = 1$ and $t_0 = 0$, as the general case will directly ensue from lemma 3.1.4. Since X has mean 0 and h is integrable, by applying both parts of lemma 1.2.2, we first show that:

$$\text{pVar}(W) = \mathbb{E}(W^2) = \iint_{\mathbb{R} \times \mathbb{R}} \gamma(t-t') \overline{\Psi_0(t)\Psi_0(t')} dt dt'$$

Now, noting that $\overline{\Psi_0(t)\Psi_0(t')} = \frac{1}{\sqrt{\pi}} e^{-\frac{1}{2}(t^2+t'^2)} e^{-i\omega_0}$, the variable change $(u, v) = (t, t-t')$ gives:

$$\text{pVar}(W) = \frac{1}{\sqrt{\pi}} \int_{-\infty}^{+\infty} \gamma(v) e^{-\frac{1}{4}v^2} A(v) dv$$

where $A(v) := \int_{-\infty}^{+\infty} e^{-(u-\frac{1}{2}v)^2} e^{-i2\omega_0(u-\frac{1}{2}v)} du$. For each $v \in \mathbb{R}$, the change $z = u - \frac{1}{2}v$ allows to see that the complex integral $A(v)$ is in fact real and does not depend on v :

$$A(v) = e^{-\omega_0^2} \int_{-\infty}^{+\infty} e^{-(z+i\omega_0)^2} dz = e^{-\omega_0^2} \sqrt{\pi}, \quad \text{and so, } \text{pVar}(W) = e^{-\omega_0^2} \int_{-\infty}^{+\infty} \gamma(v) e^{-\frac{1}{4}v^2} dv$$

Finally, from the scale-preserving property of γ and with the change $v = 2t$ we get:

$$\text{pVar}(W) = e^{-\omega_0^2} 2^\beta \sqrt{\pi} \int_{-\infty}^{+\infty} \gamma(\tau) |\Psi_0(\tau)|^2 d\tau < +\infty$$

Hence, $\text{pVar}(W)$ is real. Furthermore, owing to the definition of the acvf γ of a positively correlated fGn, $h = \gamma|\Psi_0|^2 > 0$, so will $\text{pVar}(W)$ be too.

¹⁰May we recall that a complex normal variable is fully determined by its mean, variance and pseudo-variance.

Now, using the identities in (3.6) we get $\text{pVar}(W) = |\text{Var}(\Re W) - \text{Var}(\Im W)|$, and the first equality in (3.10) follows.

For the expression of $\text{Var}(W)$ we proceed analogously and show that

$$\text{Var}(W) = 2^\beta \sqrt{\pi} \int_{-\infty}^{+\infty} h(\tau) e^{-i2\omega_0 \tau} d\tau = 2^\beta \sqrt{\pi} \widehat{h}(2\omega_0) > 0$$

which, after simplification of the common factor, completes the proof. \square

Corollary 3.1.6. *$\Re W$ and $\Im W$ are zero-mean, independent normal variables, such that $\text{Var}(\Im W) < \text{Var}(\Re W)$.*

Proof. Lemma 1.2.3 ensures that $\Re W$ and $\Im W$ are zero-mean and jointly Gaussian. Now, from $\text{pVar}(W)$ being real it follows that $\text{Cov}(\Re W, \Im W) = 0$, which together with the joint Gaussianity proves the independence. The strict positiveness of $\text{pVar}(W)$ justifies the final claim. \square

That is, although the Morlet-components of fGn are 0-mean complex normal variables, their delta-variance will not be zero. As we are about to see, how much different from zero will have consequences on how far the power-spectrogram coefficients will fail to be χ^2 -distributed. But first, let us work through a more practical expression of the delta-variance that will lead to a crucial observation.

Corollary 3.1.7. *The delta-variance of any Morlet-wavelet component W of a given β -fGn ($0 < \beta < 1$) can be expressed as:*

$$\Delta\text{Var}(W) = e^{-\omega_0^2} \frac{\Gamma(\frac{1-\beta}{2})}{\int_{-\infty}^{+\infty} |\omega|^{-\beta} e^{-(\omega-\omega_0)^2} d\omega} \quad (3.11)$$

where $\omega_0 > 0$ is the Morlet parameter and Γ denotes the Gamma function.

Proof. Let be h as in proposition 3.1.5, and let us obtain its Fourier transform. Recalling that $|\Psi_0(t)|^2 = \frac{1}{\sqrt{\pi}} e^{-t^2}$, some calculations lead to:

$$\widehat{h}(\omega') = \frac{1}{2\pi} (\widehat{\gamma} * |\widehat{\Psi_0}|^2)(\omega') = \frac{1}{\pi} \int_{-\infty}^{+\infty} \mathcal{S}^{\mathcal{F}}(2\omega) e^{-(\omega - \frac{1}{2}\omega')^2} d\omega$$

where we have denoted by $\mathcal{S}^{\mathcal{F}} = \widehat{\gamma}$ the Fourier spectrum of the fGn process. Now, simply noticing that $\int_{-\infty}^{+\infty} h(\tau) d\tau = \widehat{h}(0)$ and owing to $\mathcal{S}^{\mathcal{F}}(\omega) \propto |\omega|^{-\beta}$, we get:

$$\frac{\widehat{h}(0)}{\widehat{h}(2\omega_0)} = \frac{\int_{-\infty}^{+\infty} |\omega|^{-\beta} e^{-\omega^2} d\omega}{\int_{-\infty}^{+\infty} |\omega|^{-\beta} e^{-(\omega-\omega_0)^2} d\omega}$$

Finally, by virtue of the variable change $t = \omega^2$, we can well see that

$$\int_{-\infty}^{+\infty} |\omega|^{-\beta} e^{-\omega^2} d\omega = 2 \int_0^{+\infty} \omega^{-\beta} e^{-\omega^2} d\omega = \int_0^{+\infty} t^{\frac{1-\beta}{2}-1} e^{-t} dt = \Gamma(\frac{1-\beta}{2})$$

And we apply eq. (3.10) to conclude the proof. \square

Remark 3.1.8 (Delta-variance associated with the Morlet-CWT of a fGn). *As already pointed out in proposition 3.1.5 and its proof, since $\Delta\text{Var}(W) = \frac{\text{pVar}(W)}{\text{Var}(W)}$, by virtue of lemma 3.1.4 the delta-variance remains unchanged irrespectively of the scale and time-centre of the component W . Nor will it depend, in light of eq. (3.11), on the variance of the fGn. Indeed, $\Delta\text{Var}(W)$ does solely depend on the Morlet parameter ω_0 and the spectral slope β of the analysed fGn, not on the particular component W considered.*

We shall therefore define the delta-variance associated with the Morlet-CWT of a fGn, which we denote by $\delta\mathcal{W}$, as follows:

$$\forall \omega_0 > 0, \forall 0 < \beta < 1, \quad \delta\mathcal{W}(\omega_0, \beta) := e^{-\omega_0^2} \frac{\Gamma(\frac{1-\beta}{2})}{\int_{-\infty}^{+\infty} |\omega|^{-\beta} e^{-(\omega-\omega_0)^2} d\omega} \quad (3.12)$$

That is, the function that for each $\omega_0 > 0$ and $0 < \beta < 1$ gives the delta-variance of any wavelet component (whatever its time and scale) of an arbitrary fGn process of spectral slope β , with respect to the complex Morlet wavelet of angular frequency parameter ω_0 .

Proposition 3.1.9 (Bounds and limits of $\delta\mathcal{W}$). *The delta-variances associated with the Morlet-CWT of a fGn are bounded by:*

$$e^{-\omega_0^2} < \delta\mathcal{W}(\omega_0, \beta) < 1, \quad \forall \omega_0 > 0, \forall 0 < \beta < 1$$

Although strict, the bounds are reached for white noises ($\beta = 0$) to the left, and pink noises ($\beta = 1$) to the right.

Proof. Since, for each ω_0 and β , $\delta\mathcal{W}(\omega_0, \beta)$ is the delta-variance of a complex random variable, the upper bound merely follows by definition of the delta-variance (def. 3.1.3). For the lower bound, it suffices to apply the general property $|\hat{h}| \leq \int_{-\infty}^{+\infty} |h|$ to the identity (3.10) in proposition 3.1.5 (as both h and \hat{h} are positive).

Now, given $\omega_0 > 0$, let us call D_{ω_0} the function over $[0, 1[$ defined as:

$$D_{\omega_0}: \beta \mapsto \int_{-\infty}^{+\infty} |\omega|^{-\beta} e^{-(\omega-\omega_0)^2} d\omega \quad (3.13)$$

That is, $\delta\mathcal{W}(\omega_0, \beta) = e^{-\omega_0^2} \Gamma(\frac{1-\beta}{2}) / D_{\omega_0}(\beta)$ (corollary 3.1.7).

For $\beta = 0$, as $D_{\omega_0}(0) = \sqrt{\pi} = \Gamma(\frac{1}{2})$, the lower bound is reached: $\delta\mathcal{W}(\omega_0, 0) = e^{-\omega_0^2}$.

For $\beta = 1$, neither $\Gamma(\frac{1-\beta}{2})$ nor D_{ω_0} are defined (they both diverge to $+\infty$). We can nevertheless extend $\delta\mathcal{W}(\omega_0, \cdot)$ by continuity to 1, by showing that

$$D_{\omega_0}(\beta) \underset{\beta \rightarrow 1^-}{\sim} e^{-\omega_0^2} \Gamma(\frac{1-\beta}{2})$$

and the proof will be done.

D_{ω_0} can be simply decomposed into the sum $D_{\omega_0} = D_{\omega_0}^- + D_{\omega_0}^+$, where:

$$D_{\omega_0}^-(\beta) := \int_{-\infty}^0 |\omega|^{-\beta} e^{-(\omega-\omega_0)^2} d\omega$$

$$D_{\omega_0}^+(\beta) := \int_0^{+\infty} |\omega|^{-\beta} e^{-(\omega-\omega_0)^2} d\omega$$

Let us first turn to $D_{\omega_0}^+$. After some calculations, it can be proven that:

$$D_{\omega_0}^+(\beta) = \frac{1}{2} \Gamma\left(\frac{1-\beta}{2}\right) M\left(\frac{\beta}{2}, \frac{1}{2}, -\omega_0^2\right) + \omega_0 \Gamma\left(1 - \frac{\beta}{2}\right) M\left(\frac{\beta+1}{2}, \frac{3}{2}, -\omega_0^2\right) \quad (3.14)$$

where M denotes the Kummer's confluent hypergeometric function¹¹. Since, on the one hand, $M\left(\frac{1}{2}, \frac{1}{2}, -\omega_0^2\right) = e^{-\omega_0^2}$, and on the other hand $0 < M\left(1, \frac{3}{2}, -\omega_0^2\right) < +\infty$, it follows that

$$D_{\omega_0}^+(\beta) \underset{\beta \rightarrow 1^-}{\sim} \frac{1}{2} e^{-\omega_0^2} \Gamma\left(\frac{1-\beta}{2}\right) + \omega_0 \sqrt{\pi} M\left(1, \frac{3}{2}, -\omega_0^2\right) \underset{\beta \rightarrow 1^-}{\sim} \frac{1}{2} e^{-\omega_0^2} \Gamma\left(\frac{1-\beta}{2}\right)$$

Now, although for $D_{\omega_0}^-$ we have found no closed analytical expression, we can however bound it quite tightly. Indeed, we can show that

$$\frac{1}{2} \Gamma\left(\frac{1-\beta}{2}\right) \mathbf{m}(\beta) < D_{\omega_0}^-(\beta) < \frac{1}{2} \Gamma\left(\frac{1-\beta}{2}\right) e^{-\omega_0^2} \quad (3.15)$$

where $\mathbf{m}(\beta) := e^{-(\omega_0^2+1-\beta)} \left(\sqrt{\frac{1-\beta}{\omega_0^2+1-\beta}}\right)^{1-\beta}$

Finally, as $\mathbf{m}(\beta) \xrightarrow{\beta \rightarrow 1^-} e^{-\omega_0^2}$, we will have, also for $D_{\omega_0}^-$, that

$$D_{\omega_0}^-(\beta) \underset{\beta \rightarrow 1^-}{\sim} \frac{1}{2} e^{-\omega_0^2} \Gamma\left(\frac{1-\beta}{2}\right)$$

which terminates the proof. \square

Proof sketch of inequality (3.15). First, by letting $\omega_0 = 0$ in (3.13), it is convenient to denote $D_0(\beta) := \int_{-\infty}^{+\infty} |\omega|^{-\beta} e^{-\omega^2} d\omega$. And with the same notation used for the sum decomposition of D_{ω_0} , we will then have $D_0^-(\beta) = \frac{1}{2} \Gamma\left(\frac{1-\beta}{2}\right)$.

Now, the inequality on the right simply follows from comparing the integrands of D_0^- and $D_{\omega_0}^-$.

The left inequality can be worked through by checking that, for all $\varepsilon > 0$, the polynomial $p_\varepsilon(\omega) = -(\omega_0^2 + \varepsilon) - \left(\frac{\omega_0^2}{\varepsilon} + 1\right) \omega^2 + (\omega - \omega_0)^2$ is negative with a double root.

It ensues that, $e^{-(\omega_0^2+\varepsilon)} e^{-\left(\frac{\omega_0^2}{\varepsilon}+1\right)\omega^2} \leq e^{-(\omega-\omega_0)^2}$,

which applied to $D_{\omega_0}^-$, and through the change of variable $\omega \leftrightarrow \omega / \sqrt{\frac{\varepsilon}{\omega_0^2+1}}$ ensures that:

$$\forall \varepsilon > 0 \text{ and } \forall \beta \in [0, 1[, \quad D_0^-(\beta) e^{-(\omega_0^2+\varepsilon)} \left(\sqrt{\frac{\varepsilon}{\omega_0^2+1}}\right)^{1-\beta} < D_{\omega_0}^-(\beta)$$

Finally, taking $\varepsilon = 1 - \beta$ in particular, will do. \square

Remark 3.1.10 (Numerical approximation of the delta-variance). *The bounds given for $D_{\omega_0}^-$ in the previous proposition are reasonably tight, allowing for a good numerical approximation of the delta-variance for $0 < \beta < 1$, without the need to resort (explicitly) to fiendish methods of numerical integration. Indeed, from the*

¹¹The Kummer's function is defined as the power series $M(a, b, z) := \sum_{k=0}^{\infty} \frac{(a)_k}{(b)_k} \frac{z^k}{k!}$, where $(a)_k$ here denotes the Pochhammer symbol [Slater (1964)]




Figure 3.4: Delta-variance of wavelet components of fractional Gaussian noise as a function of its spectral slope β . Example plots computed for theoretically determined components corresponding to the analytical Morlet wavelet of parameter $\omega_0 = 6$. The delta-variances have been computed using the upper bound of $D_{\omega_0}^-$ as per inequality (3.15). **Left:** $\delta\mathcal{W}(\omega_0, \beta)$ as a function of β compared to its theoretical lower and upper bounds for $0 < \beta < 1$, as in proposition 3.1.9. The delta-variance being a ratio between quantities of the same unit, has no unit. **Right:** Limit of the delta-variance in the vicinity of $\beta = 1^-$, compared with its upper and lower bounds.

inequalities (3.15) and by calling, respectively, the lower and upper bound

$$a(\omega_0, \beta) := \frac{D_0(\beta)}{\frac{1}{2}D_0(\beta)e^{-\omega_0^2} + D_{\omega_0}^+(\beta)} \quad \text{and} \quad A(\omega_0, \beta) := \frac{D_0(\beta)}{\frac{1}{2}D_0(\beta)\mathbf{m}(\beta) + D_{\omega_0}^+(\beta)}$$

we have

$$a(\omega_0, \beta) < \delta\mathcal{W}(\omega_0, \beta) < A(\omega_0, \beta)$$

The relative error of the bound is therefore

$$Err_{Rel} = \frac{A - a}{a} = \frac{e^{-\omega_0^2} - \mathbf{m}}{D_{\omega_0}^+}$$

which reaches 0 only for $\beta = 1$.

Now, owing to eq. (3.14), a rough upper bound for that error could for example be:

$$Err_{Rel} < \frac{e^{-\omega_0^2}}{\frac{1}{2}\sqrt{\pi}e^{-\omega_0^2} + \omega_0 M(1, 3/2, -\omega_0^2)}$$

For our selection of Morlet parameter $\omega_0 = 6$ this amounts approx. to $2.74 \cdot 10^{-15}$. That is, we can have fGn's delta-variances determined with ca. $3 \cdot 10^{-13}$ % accuracy over the whole range of infra-pink slopes (see fig. 3.5 below).




Figure 3.5: Relative error between the theoretical bounds used to numerically approximate the delta-variance of fGn's wavelet components. The relative errors are shown as a function of the spectral slope β , for a fixed parameter $\omega_0 = 6$.

3.1.2.3 WPSg-coefficients of fGn as a linear combination of χ^2 variables

That is, for the case of the Morlet-based CWT, an elementary proof has just been given of the fact that, although the real and imaginary parts of the wavelet components of a stationary fGn are uncorrelated, their variances will fail to be equal. This prevents the wavelet power spectrogram coefficients to be strictly distributed as a χ_2^2 (not even for Gaussian white noise). Rather, the coefficients are distributed as a linear combination of two independent χ_1^2 variables. We give their actual expression in the following proposition.

Proposition 3.1.11 (Probability law of fGn's WPSg-coeffs). *For any sample scale $s > 0$ and time $t \in I$, the WPSg-coeff¹² $p^X(s, t)$ of a β -fGn X ($0 < \beta < 1$) is distributed as:*

$$p^X(s, t) \sim \frac{\mathbb{E} p^X(s, t)}{2} \left((1 + \delta) Z_1^2 + (1 - \delta) Z_2^2 \right) \quad (3.16)$$

where $Z_i \sim \mathcal{N}(0, 1)$ i.i.d., and δ is a strictly positive quantity that only depends on β and the Morlet parameter ω_0 .

Proof. Let us write $p = p^X(s, t)$, and let $W = \mathcal{W}X(s, t)$ be the corresponding wavelet component of the fGn. As random variables, $p \propto |W|^2$, and since $\mathbb{E}W = 0$ (corollary 3.1.6)

$$\frac{p}{\mathbb{E}p} = \frac{|W|^2}{\text{Var}(W)}$$

¹²cf section 2.4.3 on p. 30 for their definition.

Using our general formula (3.9, p. 45) for the distribution of the squared module of a complex variable, we get:

$$\frac{2p}{\mathbb{E}p} = (1 + \Delta\text{Var}(W))\widetilde{W}_{\Re}^2 + (1 - \Delta\text{Var}(W))\widetilde{W}_{\Im}^2 \quad (3.17)$$

where

$$\widetilde{W}_{\Re} := \frac{\Re W}{\sqrt{\text{Var}(\Re W)}}, \quad \widetilde{W}_{\Im} := \frac{\Im W}{\sqrt{\text{Var}(\Im W)}}$$

are independent standard normal variables (corollary 3.1.6).

Thus, taking $Z_1 := \widetilde{W}_{\Re}$, $Z_2 := \widetilde{W}_{\Im}$ and $\delta := \Delta\text{Var}(W)$ will do. From corollary 3.1.7 and remark 3.1.8, we know that $\Delta\text{Var}(W) = \delta\mathcal{W}(\omega_0, \beta)$ only depends on β and ω_0 . Finally, the lower bound of $\Delta\text{Var}(W)$ is owed to proposition 3.1.9. \square

Interestingly, from the latter it can readily be deduced that the WPSg-coeffs of a given fGn of spectral slope β are all identically distributed up to a constant that will solely depend on their scales. Namely, this we state in the next corollary:

Corollary 3.1.12 (Scale-preserving property of the laws of fGn's WPSg-coeffs). *With the assumptions of proposition 3.1.11, any two WPSg-coeffs sampled at scales $s, s' > 0$ and times $t, t' \in I$, respectively, have distributions satisfying:*

$$p^X(s, t) \sim \left(\frac{s}{s'}\right)^\beta p^X(s', t') \quad (3.18)$$

Proof. Let us denote by $p = p^X(s, t)$, $p' = p^X(s', t')$ any such two coefficients; and by $W = \mathcal{W}X(s, t)$, $W' = \mathcal{W}X(s', t')$ the underlying wavelet components.

From the previous proposition's proof, in eq. 3.16 we have $\delta = \Delta\text{Var}(W)$.

Now, since $\Delta\text{Var}(W) = \delta\mathcal{W}(\omega_0, \beta) = \Delta\text{Var}(W')$ (remark 3.1.8), to prove the lemma it suffices to show that

$$\mathbb{E}p = \left(\frac{s}{s'}\right)^\beta \mathbb{E}p'$$

Yet, this follows from noticing that $\frac{\mathbb{E}p}{\mathbb{E}p'} = \frac{\text{Var}(W)}{\text{Var}(W')}$ and from applying lemma 3.1.4 (p. 45), whence it can easily be derived that $\text{Var}(W) = (s/s')^\beta \text{Var}(W')$, recalling that $\beta = 2H - 1$. \square

As we will see in the next section (3.1.3), the previous corollary, together with equation 3.20 below (p. 54) will be crucial to motivate the normalisation of the WPSg-coeffs of a given signal.

3.1.2.4 Approximating the distribution of fGn's WPSg-coefficients

Let us now turn to studying how short the WPSg-coeffs of fGn fall from being χ_2^2 -distributed. Keeping with the hypotheses and notation of proposition 3.1.11 and its proof, for an $\omega_0 \gg 1$ fixed once and for all¹³, we define the random variable U_β , parametrised by $\beta \in]0, 1[$, as:

$$U_\beta := (1 + \delta(\beta))Z_1^2 + (1 - \delta(\beta))Z_2^2 \quad (3.19)$$

¹³The angular frequency parameter ω_0 determines the choice of the Morlet wavelet to use, which is to remain the same for our numerical applications. As mentioned earlier, we will take $\omega_0 = 6$.

Recall that $\delta(\beta)$ is nothing else than $\delta\mathcal{W}(\omega_0, \beta)$ the delta-variance of the complex wavelet components, that depends neither on the component scale nor on any other feature of the fGn except its spectral slope. Therefore, the law of U_β is shared by all WPSg-coeffs up to a proportionality constant (as in eq. 3.16). Let us compare U_β to $Z_1^2 + Z_2^2 \sim \chi_2^2$. While both variables have the same mean:

$$\mathbb{E}(U_\beta) = \mathbb{E}(Z_1^2 + Z_2^2) = 2, \quad \forall \beta \in]0, 1[$$

the ratio of their variances satisfy:

$$1 + \frac{e^{-\omega_0^2}}{2} < \frac{\text{Var}(U_\beta)}{\text{Var}(Z_1^2 + Z_2^2)} = 1 + \frac{\delta(\beta)^2}{2} < \frac{3}{2}$$

(see prop. 3.1.9).

Thus, owing to the behaviour of $\delta(\beta)$ (see fig. 3.4), the variances of these two variables will remain very close in most cases, but will non-negligibly diverge (up to a 50%) for values of β approaching 1 (see fig. 3.6, left). Accordingly, their distribution functions (call them F_{U_β} and $F_{Z_1^2 + Z_2^2}$, respectively) will also be the more dissimilar the closer β to 1. Such divergence between distributions can be conservatively captured by the Kolmogorov-Smirnov distance between two random variables (see Appendix A.3):

$$d_{\text{KS}}(U_\beta, Z_1^2 + Z_2^2) := \sup_{u \in \mathbb{R}} |F_{U_\beta}(u) - F_{Z_1^2 + Z_2^2}(u)|$$

(See fig. 3.6 on the right for a simulation.) Within the framework of the two-sided KS-test, this would mean that for a sample size of $N_{\text{sim}} = 10^5$ and a level of significance 0.05, the null hypothesis of two independent samples of U_β and $Z_1^2 + Z_2^2$ coming from the same distribution (*ie*, a χ_2^2) would only be rejected for values of $1 - \beta < 3.83 \cdot 10^{-14}$.

Remark 3.1.13. *There has been relatively recent attempts to approximate analytically and compute efficiently the distribution of the linear combination of independent χ^2 [Moschopoulos and Canada (1984), Bausch (2013)]. Although not admitting a closed expression in general, such distribution becomes rather simple for our case of two 1-dof variable terms. For example, adapting [Bausch (2013)] we can obtain a direct formula for the density function of U_β . It reads:*

$$f_{U_\beta}(u) \propto \frac{1}{1 - \delta(\beta)^2} e^{-\frac{u}{2(1 - \delta(\beta)^2)}} \mathcal{J}_0\left(-\frac{\delta(\beta)}{2(1 - \delta(\beta)^2)}u\right)$$

where \mathcal{J}_0 denotes the Bessel function of the first kind with parameter 0. Notice that, were $\delta(\beta)$ ever to tend towards 0, f_{U_β} would converge to a χ_2^2 distribution.

3.1.2.5 Summary of our theoretical findings and discussion

Let us conclude with a summary of the most relevant findings above presented. These are to be directly applied to the normalisation of the WPSg, and subsequently, also to the *sync/async* decomposition procedure.

In this section we have demonstrated that the Morlet wavelet power-spectrogram coefficients of fractional Gaussian noises are approximately distributed as a Chi-square with two degrees of freedom. Namely, putting together propositions 3.1.9 and 3.1.11, and remark 3.1.8, this main result can be summarised as follows:




Figure 3.6: Statistical comparison between the distribution U_β of the scaled WPSg-coeffs and a χ_2^2 , as a function of the fGn spectral slope β in the vicinity of 1^- (see eqs. 3.17 and 3.19). Example plots shown for a simulation of $N_{\text{sim}} = 10^5$ pairs of independent standard normals, using theoretically determined values of $\delta(\beta)$ for the Morlet wavelet of parameter $\omega_0 = 6$, computed as in fig. 3.4. **Left:** Ratio between the variances of U_β and a χ_2^2 . **Right:** Kolmogorov-Smirnov statistic between two independent samples of U_β and a χ_2^2 (sample sizes: $N_{\text{sim}} = 10^5$).

Corollary 3.1.14 (Overall summary). *Given the Morlet wavelet of angular frequency $\omega_0 > 1$, any resulting WPSg-coeff p^X of a β -fGn X ($0 < \beta < 1$) has a probability law satisfying, independently of its time and scale,*

$$\frac{2}{\mathbb{E}p^X} p^X \sim U_\beta \quad (3.20)$$

where $U_\beta = (1+\delta)Z_1^2 + (1-\delta)Z_2^2$ is a linear combination of two independent squared normal variables, and $\delta = \delta\mathcal{W}(\omega_0, \beta)$ is the relative inter-variance difference of any arbitrary wavelet component W of X .

Furthermore, δ only depends on ω_0 and β , and can be determined as:

$$\delta = e^{-\omega_0^2} \frac{\Gamma(\frac{1-\beta}{2})}{\int_{-\infty}^{+\infty} |\omega|^{-\beta} e^{-(\omega-\omega_0)^2} d\omega} \quad (3.21)$$

whence it follows that δ is strictly positive and, provided $\beta \ll 1$, is of the order of magnitude of $e^{-\omega_0^2}$.

We have also seen that, although $\delta > 0$, in most cases delta-variances of fGn wavelet components may be deemed negligible (proposition 3.1.9). For instance, with $\omega_0 = 6$ and $\beta < 0.999$, these would remain below $1.7 \cdot 10^{-12}$ (fig. 3.4); below the magnitude order of precision we are requiring in our numerical applications to biological analyses. Furthermore, as we have observed in the final part of the section, the impact of δ failing to be zero both on the variance and, more in general,

on the distribution of the resulting WPSg-coeffs, would be of small consequence for spectral slopes slightly smaller than one (fig. 3.6). Hence, and every other condition permitting, we will in practice assume that:

$$\forall \beta, \beta' < 1, \quad U_\beta \stackrel{d}{\approx} \chi_2^2 \stackrel{d}{\approx} U_{\beta'} \quad (3.22)$$

where $\stackrel{d}{\approx}$ stands for *approximately equal in distribution*. Consequently, all WPSg-coeffs of fGn will be assumed proportionally χ_2^2 -distributed, no matter their time scale nor the colour of the noise they are extracted from.

This fundamental result has already been alluded to by some authors, albeit only in passing [Percival (1995), Abry *et al.* (1995), Terrence and Compo (1998), Veitch and Abry (1999), Maraun *et al.* (2007)]. To our knowledge, no proof has ever been provided, nor its precise approximation made explicit either. (And when so, in [Ge (2007)], an incomplete proof marred some of their claims.) Although the study of the variance of the wavelet components of fractional Brownian motions has been the subject of much investigation [Flandrin (1989), Vergassola and Frisch (1991), Flandrin (1992), Tewfik and Kim (1992), Wornell (1993)], it is less so for the fractional Gaussian noise; and to our knowledge, no work has ever mentioned the pseudo-variance and the particularities that arise when dealing with complex wavelets components. In fact, for convenience it has usually been assumed that the real and imaginary parts of wavelet components of Gaussian noises have equal variance. This is not the case in general, as we have just proved. On the whole we believe that, by introducing a clear formalisation, our contribution may help to dispel ambiguities that have encumbered this topic and its implications for the significance assessment of transient spectral peaks in inverse power law signals.

3.1.3 Definition of the normalised WPSg-coefficients

We are finally well placed to define our normalisation of the wavelet power spectrogram. We will first describe a natural normalisation for the WPSg-coefficients of a process and point to its desirable effect when applied to coloured noises. Then, we will proceed with the definition to a normalisation of sampled logMUA signals.

3.1.3.1 Mean-normalised WPSg-coefficients of fGn processes

From the vantage point of corollary 3.1.14 we can simply define:

Definition 3.1.15 (Mean-normalised WPSg-coeffs of a process). *Let p^X be any WPSg-coeff of a given stochastic process X . We define its mean-normalised form, denoted by $\natural p^X$, as the random variable:*

$$\natural p^X := \frac{2}{\mathbb{E}p^X} p^X \quad (3.23)$$

By definition, all the normalised coefficients will have the same expectation $\mathbb{E}(\natural p^X) = 2$. Let us also note that this normalisation is natural in the sense that the operator \natural is idempotent: $\natural(\natural p^X) = \natural p^X$. In the case of fGn, the mean-normalised WPSg-coefficients (μ -nWPSg-coeffs) additionally enjoy a number of useful properties.

Proposition 3.1.16 (Key statistical properties of fGn's μ -nWPSg-coeffs).

Let X and X' be fGn over the same time interval I , with spectral slopes β and β' ($0 < \beta, \beta' < 1$), respectively. Then, for any times $t, t' \in I$ and scales $s, s' > 0$, the corresponding μ -nWPSg-coeffs¹⁴ of X and X' have probability laws satisfying:

$$\begin{aligned} (N1) \quad & \natural p^X(s, t) \sim \natural p^X(s', t') \\ (N2a) \quad & \natural p^X(s, t) \sim U_\beta \stackrel{d}{\approx} U_{\beta'} \sim \natural p^{X'}(s', t') \\ (N2b) \quad & \natural p^X(s, t) \stackrel{d}{\approx} \chi_2^2 \end{aligned}$$

where the approximations in distribution withstand as long as $\beta, \beta' < 1$.

Proof. (N2a) and (N2b) are mere restatements of (3.20) in corollary 3.1.14, having taken into account the approximations in (3.22). (N1) is ultimately a consequence of the scale-preserving property of fGn and can be easily derived from corollary 3.1.12. In fact, let us write $p = p^X(s, t)$ and $p' = p^{X'}(s', t')$, and $\natural p, \natural p'$ their respective normalised forms. Then, owing to relation (3.18) we have:

$$\natural p = \frac{2}{\mathbb{E}p} p \sim \frac{2}{\mathbb{E}\left(\left(\frac{s}{s'}\right)^\beta p'\right)} \left(\frac{s}{s'}\right)^\beta p' = \frac{2}{\left(\frac{s}{s'}\right)^\beta \mathbb{E}p'} \left(\frac{s}{s'}\right)^\beta p' = \frac{2}{\mathbb{E}p'} p' = \natural p'$$

□

¹⁴As mentioned earlier, we assume that all the underlying CWT components result from the same Morlet wavelet with given parameter $\omega_0 \gg 0$.

Importantly, when $\beta \rightarrow 0$, the properties hold. This entails that mean-normalised coefficients of any fGn are marginally distributed as WPSg-coefficients of white Gaussian noise. In particular, the resulting mean-normalised global and local wavelet spectra will be as flat as that of purely uncorrelated noise. Therefore, these normalised spectra have all an even allocation of their variance across the frequency range. Additionally, our simulations suggest that the whole correlation structure of fGn's μ -nWPSg-coeffs is statistically indistinguishable from that of white noise's. In summary, the mean normalisation might be understood as a *whitening* of the wavelet power spectrogram.

3.1.3.2 Normalised WPSg-coefficients of logMUA signals

In view of the previous observations, we propose an applied version of the WPSg mean-normalisation to sampled signals of the kind we will be analysing in this work. Put it shortly: we are going to normalise the spectrogram against the inverse power law of the signal's spectral background.

Let ν be a logMUA signal, and let $p_{j,k}^\nu$ be its WPSg-coefficients computed over the usual time-frequency sampling mesh (ξ_j, t_k) (see sections 2.4.1 and 2.4.3). From its global power spectrum \mathcal{S}_ν , let us assume we have been able to estimate the background inverse power law the signal follows (as described in sections 3.1.1 and 3.1.1.1). Two estimates then resulted: the signal's spectral slope β_ν and intercept K_ν . Thus, for each time-frequency index (j, k) we define the normalised WPSg-coefficients of ν , denoted by $\tilde{p}_{j,k}^\nu$, as

$$\tilde{p}_{j,k}^\nu := \frac{2}{K_\nu \xi_j^{-\beta_\nu}} p_{j,k}^\nu \quad (3.24)$$

Definition's rationale. The previous formula (3.24) is a mere reformulation of eq. (3.23), where the expectation term on the denominator has been substituted by an estimate of the ergodic mean of the coefficient $p_{j,k}^\nu$, under the assumption that the signal is, at that precise time and frequency, statistically indiscernible from its underlying background noise. To clarify this delicate point, let us put it the other way around: in the special case of ν being a realisation of a coloured noise X with spectral parameters β_X and K_X , for an arbitrary sample time t and frequency ξ_j , we would have:

$$\mathbb{E}p^X(\eta_0 \xi_j^{-1}, t) \stackrel{(1)}{\approx} \frac{1}{T} \sum_{k=1}^N p_{j,k}^\nu \Delta t = \mathcal{S}_\nu(\xi_j) \stackrel{(2)}{\approx} K_\nu \xi_j^{-\beta_\nu} \quad (3.25)$$

where η_0 is the wavelet's scale-frequency conversion factor (see eq. 1.4 from remark 1.2.1), and $p^X(\eta_0 \xi_j^{-1}, t)$ —here a random variable—is the WPSg-coefficient of X at time t and scale $s_j = \eta_0 \xi_j^{-1}$. That is, approximation (1) refers to the ergodic estimate of the expectation of the WPSg-coefficient of the stationary process X at the given scale s_j . This average is, by definition, the global spectrum component $\mathcal{S}_\nu(\xi_j)$ of the realisation of ν corresponding to that scale (eq. 2.10 section 2.4.3). Now, the latter can be well approximated (2) by the inferred inverse power law $K_\nu \xi_j^{-\beta_\nu}$ the signal follows (eq. (1.9) section 3.1.1).

Notice also that, at any given scale s , the process $p^X(s, \cdot)$ is stationary [Flandrin (1989)], and that $\mathbb{E}p^X(s, t)$ is in fact a component of what we can call the *expected wavelet*

global spectrum of X . That is, for any low-enough frequency ξ and an arbitrary time-point t :

$$\overline{\mathcal{S}}_X(\xi) := \mathbb{E}p^X(\eta_0\xi^{-1}, t) \quad (3.26)$$

$\overline{\mathcal{S}}_X$ is a smooth and band-restricted estimate of the theoretical spectrum $\mathcal{S}_X^{\mathcal{F}}$ (see sections 1.2.4 and 3.1.1.2) and provides good estimates of its spectral parameters [Perrier *et al.* (1995), Abry *et al.* (1995)]. The approximations of eq. (3.25) are nothing but observations of these estimates through a realisation of the process.

Remark 3.1.17 (WPSg normalisation: processes *vs* signals).

Although both similar and stemming from the same motivation, the two preceding normalisations are not interchangeable.

To begin with, if the mean-normalisation formula (eq. 3.23) were to be applied to the WPSg-coeffs of an observed signal (ie, fixed real numbers), the expectation term would not be very informative.

The other way around, applying directly eq. 3.24 to a stochastic process would be problematic if, for example, the process did not follow an inverse power law.

All in all, if ν is a realisation of the fGn X , then \tilde{p}^ν , the normalised WPSg-coeffs of the signal ν , can be deemed as observations of $\mathfrak{W}p^X$, the mean-normalised WPSg-coeffs of the process X .

3.1.3.3 Theoretical interpretation

Although originally devised for logMUA signals, nothing in the definition prevents the signal from being an arbitrary one. For the normalisation to make sense, the main requirement is that the signal follows, in a broad-enough frequency band, an inverse power law that can arguably be well estimated by means of its wavelet global spectrum.¹⁵ Additionally, if the spectrum's estimated spectral slope lies between 0 and 1 (ie, background noise in *infra-pink* regime), and if we are also willing to assume the signal's background noise to be Gaussian, then we will be under the hypotheses of the summarising corollary 3.1.14 (p. 54).¹⁶ In that case, interpreting the normalised coefficients in light of the above approximations (eq. (3.25) on p. 57), will prove very useful.

For that, let us consider a family of null hypotheses concerning single WPSg-coefficients of a given signal, where each of these coefficients is viewed as an observation drawn from a population of plausible time-frequency components of the said signal. For each $k \in \llbracket 1, N \rrbracket$ and $j \in \llbracket 0, J - 1 \rrbracket$ we state:

¹⁵This in turn presupposes some degree of stationarity for the signal, at least if restricted to a frequency band of interest. Otherwise, neither the global spectrum could possibly estimate a non-existent theoretical spectrum, nor the inverse power law inferred this way could be interpreted in terms of stationary background noise.

¹⁶Throughout the development of the previous theoretical explorations (sections 3.1.1.1, 3.1.2), there are also a number of additional hypotheses we have already taken for granted. Such are: 1) a preliminary choice of the Morlet parameter ω_0 ensuring, on the one hand, that the wavelet has enough vanishing moments to properly estimate the spectral slope, and on the other hand, that the lower bound in eq. 3.15 is kept small; 2) a signal's time interval I long enough wherein to compute the wavelet components, in order to avoid edge-effects up to a low frequency cut-off, and to allow for the assumption of even integrands in the proof of proposition 3.1.5.

$\mathcal{H}_0(j, k)$: *The time-frequency component of the signal at time t_k and frequency ξ_j is solely accounted for by the signal's own intrinsic background noise, which is fractional Gaussian.* (3.27)

In fact, assuming $\mathcal{H}_0(j, k)$ is equivalent to saying that:

(A) the process generating the signal is, at that particular time-frequency point, statistically indiscernible from a fGn;

and

(B) such fGn is precisely that which models the best the inferred power law followed by the signal.

This two components of the hypothesis play a different role: statement (B) is a ground modelling hypothesis or *prior* that we are assuming throughout the present work, whereas statement (A) is a conditional one that we may want it subjected to hypothesis testing (see below in section 3.2.1.1).

Thus, under $\mathcal{H}_0(j, k)$, a signal's (j, k) -th WPSg-coefficient will be an observation from a population that is distributed as a certain linear combination of two independent normal random variables (see proposition 3.1.11 for the precise formulation involving the delta-variance of the wavelet components of the underlying process). And, as per proposition 3.1.16, the ensuing mean-normalised WPSg-coefficient of the underlying process will then be approximately χ_2^2 -distributed.

In particular, for any collection of indices $J' \times K' \subseteq \llbracket 0, J \rrbracket \times \llbracket 1, N \rrbracket$, the normalised (j', k') th WPSg-coefficients of signals under $\mathcal{H}_0(j', k')$, $(j', k') \in J' \times K'$, will all be comparable—comparable insofar as their null distributions are very close (see section 3.1.2.4). By extension, for any such Gaussian signal that follows an inverse power law, it will be useful to interpret their normalised coefficients as χ_2^2 -scores, even if the null hypothesis may not be assumed at particular time-frequency points.

3.1.3.4 Application and practical consequences

A useful way of interpreting the normalisation of the WPSg is to see it as a *whitening* of the background noise of the signal—again, provided there exist sound reasons for that background to be modelled by a coloured noise, as it is the case here for the logMUA, a near Gaussian, inverse-power-law-following signal. This interpretation can easily be apprehended through the effect the normalisation has upon the global power spectrum. We recall that the global spectrum is defined as the time average of the spectrogram for a given long enough period (see section 2.4.4, eq. 2.10).

In fact, the global spectrum of a normalised spectrogram follows the flat power law distinctive of white noises (see fig. 3.7). Importantly, not only does the normalisation flattens the spectra, but it also ensures that all normalised spectrograms end up with the same amount of energy allocated to their backgrounds. Notice here that, since the *background* WPSg-coefficients are (almost) χ_2^2 -distributed, their ergodic mean (of which the global spectrum provides an estimation) matches the mean of the said random variable, *ie* 2. And this, independently of the time scale of the coefficient. However, in fig. 3.7 the units of the global spectra must be interpreted cautiously: these are not properly χ_2^2 -scores, but rather *means* of χ_2^2 -scores, from which no straightforward conclusion in terms of significance testing should

directly be extracted.¹⁷

In summary, our normalisation procedure allows for normalised WPSg-coefficients of logMUA signals to be fruitfully compared between themselves, irrespectively of their corresponding frequency or whether they come from distinct logMUAs following different power laws (see fig. 3.8). Finally, and owing to the approximations of eq. (3.22), we can safely convene that in practice the null distribution of all these normalised coefficients is a χ^2 with 2 degrees of freedom.¹⁸

¹⁷In fact, for that we ought to study the distribution of the sample mean of certain collections of non-independent χ^2_2 variables. It could be there the case that the correlation between them does depend on their scales. We address this and other closely related issues in the next section 3.2.

¹⁸As we have strived to demonstrate in previous sections, the claim is true for realisations of infra-pink fractional Gaussian noises. However, it would not necessarily be so for other Gaussian signals that, although following an inverse power law, nonetheless present a richer spectral content (*eg*, conspicuous oscillations). It is precisely this observation that lays the ground for our method to detect asynchronous periods amidst *non-stationary* signals (see next section 3.2).




Figure 3.7: Effect of the WPSg-coefficients' normalisation on the global power spectrum. From two distinct logMUA signals simultaneously recorded at close location (electrodes BinZ and BexR), wavelet global spectra over 300s period are considered. **Left:** global spectra computed over energy-preserving wavelet WPSg-coefficients are shown. For each signal, dashed lines correspond to the estimated inverse power law of their intrinsic background spectrum (spectral slopes $\beta_1 = 0.2718$ and $\beta_2 = 0.0628$ for BinZ and BexR, respectively). **Right:** global spectra computed over normalised WPSg-coefficients. Here the dashed line indicates the flat power law (slope 0) that both normalised background spectra have in common, and which corresponds to that of white noise's.




Figure 3.8: Effect of the normalisation on the power spectrogram. From the same two logMUA signals as in fig. 3.7 broad band, full-time period power spectrograms are considered (only 40s shown). **Top row:** energy-preserving WPSg. **Centre row:** normalised WPSg. **Bottom row:** p-values associated to the marginal tests corresponding to each normalised WPSg-coefficient under the null hypothesis (3.27).

3.2 Detection of synchronous and asynchronous states

In the previous section we have studied how the coefficients of the wavelet power spectrogram of a canonical family of stochastic processes marginally distribute. There we argued that these processes—fractional Gaussian noises in the infra-pink regime—can usefully be assumed to underlie the logMUA signals, thereby providing their “asynchronous substance”. Any time-frequency component of a signal that do not statistically behave as its intrinsic background noise would therefore point to a marginal departure in asynchrony. However, for the analysis of logMUA signals to be biologically relevant, we wish to detect time-frequency patches of asynchrony that are broad-band in frequency and persistent in time. Complementarily, we also want to assess whether any departure from asynchrony is really due to the existence of genuine synchronous activity attributable to biologically functional oscillations.

With that aim in mind, we will now be interested in studying instantaneous broad-band *slices* of the spectrogram. Following our notation (see in particular sections 2.4.4 and 3.1.2), at each instant t , the local spectrum $\mathcal{S}_\nu^{\text{loc}}(t)$ of a signal ν may be thought of as a sample of J random variables, where J is the number of time-scales of the spectrogram (see eq. 2.8 on p. 32)¹⁹. This sample is far from being composed of independent and identically distributed drawings. In fact, the unavoidable inter-scale correlations that the CWT time-frequency decomposition brings about preclude the underlying variables from being mutually independent. In addition, if the signal follows an inverse power law, the power coefficients of different scales will not be identically distributed, as their expectation will indeed depend on the scale of the component and the spectral slope.

However, although not independent, there is still an opportunity for the sample of power coefficients to be identically distributed: that will be the case if they are extracted from a fractional Gaussian noise, and when their normalised form is considered. As we have shown in the previous section, the mean-normalised coefficients of a given fractional Gaussian are all identically distributed, following a probability law U_β that only depends on the spectral slope β of the fGn and the Morlet parameter used for the CWT (see corollary 3.1.14 and proposition 3.1.16 on pages 54 and 56 respectively). Crucially, these variables U_β happen to be extremely close to each other, and all of them almost χ_2^2 -distributed.

Over this population of local spectra, we propose a time-dependent family of statistical tests to assist in the classification of synchronous and asynchronous signal periods. First, we present a test that provides a definition of asynchrony by default; then a test that aims at detecting oscillations in a restricted band of interest. The null distributions for the statistics of both tests are computationally generated by a Monte Carlo procedure. Finally, the effective classification is carried out by optimising a pair of significance levels jointly using the entire recording period and all the channels of the slice. We have called the whole procedure: Monte Carlo Coupled Testing (MCCT).

¹⁹There is here a subtlety that was missing when dealing with observed signals, but which has been implicitly arising since then. Local and global wavelet spectra of a process defined over a time-frequency mesh are, in purity, random vectors. We will then tend to consider also their expectations (see eq. (3.26)). Under certain circumstances, they are good estimators of the theoretical spectrum or the spectral variance (*eg*, for fGn or for stationary processes; see section 3.1.1.2 and eq. (3.25)). .

3.2.1 Definition of asynchronous states

Now, as we previously pointed out, normalised WPSg-coeffs of a Gaussian signal that follows an inverse power law may be interpreted as observations of the mean-normalised coefficients of a certain fGn. Therefore, in light of the mesh of null hypotheses $\mathcal{H}_0(j, k)$ introduced in section 3.1.3.3 (p. 58), it follows that, at a given instant t_k and under the composite null hypothesis $\bigwedge_j \mathcal{H}_0(j, k)$, the *normalised local spectrum* of such a signal at t_k can in practice be assumed to be a sample of J identically χ_2^2 -distributed variables.

3.2.1.1 A test to discard asynchrony

Let ν be an observed logMUA signal of length T and sampling period Δt , with WPSg-coefficients $p_{j,k}'$ corresponding to the usual time-frequency sampling mesh points $(t_k, \xi_j) \in [0, T] \times [\xi_{\min}, \xi_{\text{Nyquist}}]$ (see Methods, section 2.4.1). As we have seen, its global spectrum *follows* an inverse power law. Let us call β_ν and K_ν its estimated spectral parameters. It is here crucial to recall that, for a global spectrum, the concept *following a power law* with slope β was defined as the existence of a maximal frequency band $\Xi_\beta \neq \emptyset$ on which the spectrum does observe an inverse β -power law (see section 3.1.1). This definition is still valid for the realised global spectrum of a process over a given time-frequency sampling mesh: thus, we will also say of a process that it follows a power law. Let be X such a process. As per our logMUA modelling hypothesis, we assume that X is Gaussian and follows the same inverse power law as ν , with *estimated* spectral parameters $\hat{\beta}_X = \beta_\nu$ and $\hat{K}_X = K_\nu$.

Now, assuming our modelling hypothesis²⁰ that the power-law behaviour of ν is solely accounted for by a fGn Φ with²¹ spectral slope $\beta_\Phi = \beta_\nu$ and intercept $K_\Phi = K_\nu$, we are interested in assessing how similar is the signal ν to its fGn background at a given instant t_k . To that aim, we can devise a statistical test over a population of local spectra that can be formalised as follows:

Population: Local spectra at instant t_k of any such Gaussian process X that could have generated the observed signal ν .

The elements of the population are thus random vectors with components

²⁰That is, the component (B) of the hypothesis (3.27) on section 3.1.3.3 p. 59; see also section 3.1.1.2

²¹Let us note again the essential difference between

- i) the fGn Φ has spectral slope β_Φ , and
- ii) a process X (not necessarily a fGn) follows an inverse power law with estimated slope $\hat{\beta}_X = \beta_x$ (where x would be a given realisation of the process).

In fact, the spectral slope of a fGn is intimately related to its Hurst parameter, and as such, is constitutive of the process. A fGn trivially follows a power law too (at least for low-enough frequencies). However, the matter is different for general processes which might not even be stationary. Here, the sampled global spectrum provides a band-limited approximation from which the power-law-following parameters can be estimated.

the WPSg-coeffs²² of X at the given t_k :

$$\mathcal{S}_X^{\text{loc}}(t_k) = (p_{0,k}^X, \dots, p_{J-1,k}^X)$$

Note that β_ν and K_ν will be acting as hyperparameters of the test.

Feature of interest: Of the joint distribution of $\mathcal{S}_X^{\text{loc}}(t_k)$ we know little. Only that i) it is marginally Gaussian, and ii) the marginal ergodic expectation of its components is an inverse power law function of ξ_j with parameters β_ν and K_ν ; that is:

$$\mathbb{E} \left(T^{-1} \sum_{k=1}^N p_{j,k}^X \Delta t \right) = \mathbb{E} \mathcal{S}_X(\xi_j) = S_\nu(\xi_j) \approx K_\nu \xi_j^{-1}$$

However, its covariance structure remains largely unknown and depends on both the wavelet transform and the auto-covariance function of the process X at time t_k (which might very well change across times).

We can nonetheless restrict our interest to a relevant and observable feature of the population: *the joint deviation from its ergodic expectation*. That is, in the context of our modelling hypothesis, the aberration between the local spectrum and its intrinsic fGn background.²³

Sample: The local spectrum of ν at instant t_k :

$$\mathcal{S}_\nu^{\text{loc}}(t_k) = (p_{0,k}^\nu, \dots, p_{J-1,k}^\nu)$$

That is, a sample consisting of a single observed vector.

Statistic: For $\alpha_0 \in]0, 1]$ and $J_0 < J$, we define the statistic R_{J_0, α_0}^ν (or simply R when the context is clear) as:

$$R_{J_0, \alpha_0}^\nu : (p_0, \dots, p_{J-1}) \mapsto \frac{1}{J_0} \sum_{j=0}^{J_0-1} \mathbb{1}_{]F_U^{-1}(1-\alpha_0), +\infty[} \left(\frac{2}{K_\nu \xi_j^{-\beta_\nu}} p_j \right) \quad (3.28)$$

where F_U^{-1} is the quantile function of the random variable $U = U_{\beta_\nu}$ with parameter β_ν (as defined in eq. (3.19), section 3.1.11); and $\mathbb{1}_A$ denotes the $\{0, 1\}$ -valued *indicator* function of any set $A \subseteq \mathbb{R}$. Thus,

$$\mathbb{1}_{]F_U^{-1}(1-\alpha_0), +\infty[} \left(\frac{2}{K_\nu \xi_j^{-\beta_\nu}} p_j \right) = 1 \iff p_j > \frac{K_\nu \xi_j^{-\beta_\nu}}{2} F_U^{-1}(1-\alpha_0)$$

²²Henceforth, when there is no ambiguity regarding the time-frequency sampling mesh, we will for convenience use the sub-index notation also for the WPSG-coeffs of a process. That is $p_{j,k}^X := p^X(\eta_0 \xi_j^{-1}, t_k)$.

²³We are thus implicitly acknowledging the existence of a *latent parameter* ρ of the population, which would be responsible for the *ergodic aberration* of the process at a given instant. In practice, we will estimate ρ by a statistic $R \propto \hat{\rho}$ that we are going to provide, and whose null distribution will in turn be estimated by a simulated empirical distribution. This way of proceeding allows to accommodate our computational approach to the classical formalisation of hypothesis testing (see *eg* [Robert and Casella (2004)]).

R simply gives the proportion of the first J_0 components of a vector that are above certain cut-off values $\frac{K_\nu \xi_j^{-\beta_\nu}}{2} F_U^{-1}(1 - \alpha_0)$. The range of R is therefore $\{\frac{l}{J_0} : 0 \leq l \leq J_0\} \subset [0, 1]$.

Null hypothesis:

$$H_0^{\text{async}} : (p_{0,k}^X, \dots, p_{J_0-1,k}^X) \sim (p^\Phi(\xi_0), \dots, p^\Phi(\xi_{J_0-1})) \quad (3.29)$$

Or, in plain words: *the first J_0 components of the local spectrum at instant t_k of the process generating ν , $\mathcal{S}_X^{\text{loc}}(t_k)$, are statistically indiscernible from those of any local spectrum $\mathcal{S}_\Phi^{\text{loc}}$ of a fGn Φ with spectral slope β_ν and intercept K_ν . That is,*

$$H_0^{\text{async}} \implies \bigwedge_{0 \leq j < J_0-1} \mathcal{H}_0(j, k)$$

where $\mathcal{H}_0(j, k)$ are the mesh of marginal null hypotheses²⁴ (3.27) stated in 3.1.3.3. Note that for a random vector, since *marginal* does not generally imply *joint*, the converse of the previous implication is not necessarily true.

Sampling distribution: $F_R(r | H_0^{\text{async}}) = \mathbb{P}(R_{J_0, \alpha_0}^\nu \leq r | H_0^{\text{async}})$

The null distribution of R will be estimated through the empirical distribution $\widehat{F}_{R, N_{\text{sim}}}$ of the observed values of R for N_{sim} independent local spectra of a process simulated under the *async* null hypothesis (*ie*, the Monte Carlo experiment).

Monte Carlo test: Then, for a fixed significance level $\alpha^{\text{async}} > 1/N_{\text{sim}}$, we will reject H_0^{async} whenever

$$\widehat{F}_{R, N_{\text{sim}}}(R_{J_0, \alpha_0}^\nu(\mathcal{S}_\nu^{\text{loc}}(t_k))) > 1 - \alpha^{\text{async}} \quad (3.30)$$

For a given signal ν , and a fixed significance level α^{async} , the test may be repeatedly carried out at every instant t_k , yielding a collection of rejected and accepted local spectra. By *async-test* we will refer to this time-dependent procedure, or to any single test carried out at an isolated instant.

3.2.1.2 Interpretation of the test statistic

Before we can further proceed, we need to see that, under the *async* null hypothesis, the *async-test* statistic admits a useful interpretation.

Proposition 3.2.1 (Null distribution of the statistic R).

Let $R = R_{J_0, \alpha_0}^\nu$ be as in (3.28). Then, for every time sample t_k , under the hypothesis H_0^{async} :

$$R(\mathcal{S}_X^{\text{loc}}(t_k)) \underset{H_0^{\text{async}}}{\sim} \frac{1}{J_0} \sum_{j=0}^{J_0-1} B_j^{\alpha_0, \beta_\nu} \quad (3.31)$$

²⁴Here we are precisely interested in testing the (A) sub-hypotheses, all in assuming the ground modelling hypothesis (B) (see p. 59).

where for each j , $B_j^{\alpha_0, \beta_\nu} := \mathbb{1}_{]F_U^{-1}(1-\alpha_0), +\infty[}(\mathfrak{h}p^X(\xi_j, t_k))$, and $\mathfrak{h}p^X$ denote the mean-normalised WPSG-coefficients of X .

Furthermore, the $B_j^{\alpha_0, \beta_\nu}$ are identically distributed Bernoulli variables with parameter α_0 .

Proof. H_0^{async} means that $p^X(\xi_j, t_k) \sim p^\Phi(\xi_j, t)$, for all $0 \leq j < J_0$ and an arbitrary instant t . Therefore,

$$\mathfrak{h}p^X(\xi_j, t_k) \sim \mathfrak{h}p^\Phi(\xi_j, t) \sim U_{\beta_\nu}$$

(by virtue of proposition 3.1.16 on p. 56).

Now, under each marginal null hypothesis $\mathcal{H}_0(j, k)$,

$$K_\nu \xi_j^{-\beta_\nu} \approx \mathbb{E}p^X(\xi_j, t_k)$$

so the normalised $\tilde{p}_{j,k}^\nu$ could be considered as observations of the mean-normalised $\mathfrak{h}p^X(\xi_j, t_k)$ (see eq. (3.25) on p. 57).

Finally, recalling that $\tilde{p}_j^\nu := 2K_\nu^{-1} \xi_j^{\beta_\nu} p_j^\nu$, the relation (3.31) follows.

For the remaining part of the proof, let us simply write $\mathbb{P}_0 := \mathbb{P}(\cdot | H_0^{\text{async}})$

and $\mathfrak{h}p^X = \mathfrak{h}p^X(\xi_j, t_k)$

Under H_0^{async} , $\mathfrak{h}p^X \sim U_{\beta_\nu}$, and so $\mathbb{P}_0(\mathfrak{h}p^X \leq c) = F_U(c)$

We will then have:

$$\begin{aligned} \mathbb{P}_0 \left(B_j^{\alpha_0, \beta_\nu} = 1 \right) &= \mathbb{P}_0 \left(\mathbb{1}_{]F_U^{-1}(1-\alpha_0), +\infty[}(\mathfrak{h}p^X) = 1 \right) \\ &= \mathbb{P}_0 \left(\mathfrak{h}p^X > F_U^{-1}(1 - \alpha_0) \right) \\ &= 1 - F_U \left(\mathfrak{h}p^X \leq F_U^{-1}(1 - \alpha_0) \right) \\ &= 1 - (1 - \alpha_0) = \alpha_0 \end{aligned}$$

where the penultimate equality merely follows from the definition of the quantile function of a random variable.

Finally, noticing that $B_j^{\alpha_0, \beta_\nu} \in \{0, 1\}$ ends the proof. \square

Therefore, under H_0^{async} the statistic R is the success rate of J_0 simultaneous—but not necessarily independent—Bernoulli trials with parameter α_0 . In particular, R can be interpreted as the positive test rate, at an α_0 significance level, of a multiple right-tailed test assessing whether each of the normalised power coefficients is likely to have been marginally drawn from a U_{β_ν} probability distribution²⁵. As such, R provides a measure of how much an observed local spectrum deviates *by excess* from its background power law above the low cut-off frequency ξ_{J_0} .

Importantly, let us briefly observe that the case of deviation *by default* would be of no concern. In the first place, U_{β_ν} , the marginal distribution of the normalised WPSG-coefficients under the null hypothesis, is very close to a χ_2^2 (see section 3.1.2.4), which has a positive support and is strongly positively skewed. Namely, more than 60% of the normalised coefficients accumulate at small values between 0 and 2, their mean. That is, only a high value can make coefficients stand out.

In the second place, however, there is a more profound reason from which the first argument stems. We recall that a signal—or its underlying process—is said to

²⁵That is why in subsequent sections we will often refer to the statistic R as *the positive test rate at α_0 level*.

follow an inverse power law if the background tendency of its global spectrum can be fitted by that power law (see section 3.1.1). As we argued in section 3.1.1.2, that behaviour is the hallmark of maximally uncorrelated coloured noises. Therefore, any local deviation from the background spectrum must be caused by a noticeable event in the covariance structure of the signal, and hence imply a surge in local power. In short: that is why the background spectrum got its name.

3.2.1.3 Parameters of the test and Monte Carlo procedure

Now, if we want our statistic to be informative about the population's feature we are striving to asses (*ie*, its departure from its background fGn's spectral regularity), we need to select the test parameters J_0 and α_0 accordingly.

On the one hand, it directly follows from the previous proposition that:

$$\mathbb{E}(R_{J_0, \alpha_0}^\nu \mid H_0^{\text{async}}) = \alpha_0$$

So, in principle, a small marginal significant level α_0 might shed light on how wildly a fGn's local spectrum departs above its power law. On the other hand, we want J_0 close to J so that $\xi_{J_0-1} \gtrsim \xi_{\min}$, in order to capture the broadest possible frequency band.

However, owing to R being bounded, select too small an α_0 , and R will saturate at 0, hardly discriminating between different deviation levels. The opposite would happen for a moderate α_0 value, even for a relatively small J_0 . To illustrate this point, it could help to assume for one moment that the Bernoulli trials $B_j^{\alpha_0, \beta_\nu}$ are independent. The null distribution of the ensuing statistic (call it R^*) would therefore be completely determined. In fact, R^* would be the sample distribution of the sample proportion estimating the Bernoulli rate parameter $1 - \alpha_0$. Therefore

$$J_0 R^* \sim \text{Bi}(J_0, \alpha_0)$$

a binomial distribution with parameters $n = J_0$ and $p = \alpha_0$. That is:

$$\mathbb{P}_0 \left(R^* = \frac{j}{J_0} \right) = \binom{J_0}{j} (\alpha_0)^j (1 - \alpha_0)^{J_0 - j} \quad \text{for } j \in \llbracket 0, J_0 \rrbracket$$

where \mathbb{P}_0 denotes the probability law under the null hypothesis (as defined in the proof of proposition 3.2.1 on p. 65). In particular, we would have

$$\mathbb{P}_0(R^* = 0) = (1 - \alpha_0)^{J_0} \quad \text{and} \quad \mathbb{P}_0 \left(R^* \geq \frac{1}{J_0} \right) = 1 - (1 - \alpha_0)^{J_0}$$

let alone the perks that would come with such assumption of independence: known variance, and the possibility of benefiting from a Normal approximation to R^* whenever $J_0 \cdot \alpha_0$ and $J_0 \cdot (1 - \alpha_0)$ be jointly big enough.

Another aspect to bring into consideration is the computational cost involved in simulating the empirical distribution $\hat{F}_{R, N_{\text{sim}}}$. As we shall see, the choice of α_0 and J_0 tightly depends on the parameters of the Monte Carlo experiment.

In fact, without any further information about the non-trivial covariance structure of the local spectrum of a fGn, in order to generate N_{sim} independent samples

of such random vectors, one has to resort to generating N_{sim} independent realisations of the said process. Then, the continuous wavelet transform can be computed for each sample path to obtain the WPSg-coefficients of a single local spectrum which will finally provide an observed value for the R statistic. Below we point to important aspects of the whole procedure.

Practical considerations for the Monte Carlo experiment Let us first expose a number of constraints to take into account:

- (A1) For the samples to be independent, they must indeed be drawn from independent realisations of the process, and not only from distinct sampled time points, however distant. This is because of the long-range auto-correlation of the fGn and the unavoidable temporal correlations introduced by the wavelet transform.
- (A2) In order to ensure that the low-frequency components of the sampled local spectra are free from edge-effect artefacts (refer to Methods, section 2.4.2 on the cone of influence), they will best be extracted at an instant located towards the centre of the time period of the realisation.
- (A3) The realisations must ideally share the sampling parameters of the logMUA signal under consideration and its power spectrogram: sampling period Δt , voices per octave \mathbf{v} and sampled time-scales s_j (see Methods, section 2.4.1).
- (A4) Because of the last two points, the lower the choice of low cut-off frequency $\xi_{\text{sim}} := \xi_{J_0-1}$, the longer the time length of the realisation T_{sim} must be.
- (A5) In theory, the async-test depends on the hyper-parameters of the test that the observed signal ν imposes upon: its variance σ_ν^2 , its spectral parameters β_ν and K_ν , and the $(1 - \alpha_0)$ -quantile of the distribution U_{β_ν} . Consequently, the null distributions of the statistic R_{J_0, α_0}^ν are different for each set of parameters, and will require as many Monte Carlo experiments in order to simulate their empirical distributions.

Fortunately, there are some relaxations that would make the simulation computationally less burdensome.

- (B1) If ξ_{sim} is the targeted low cut-off frequency, it will suffice to generate realisations of length $T_{\text{sim}} \geq 2\sqrt{2} \frac{\eta_0}{\xi_{\text{sim}}}$ seconds and uniquely compute and draw the local spectrum corresponding to an available time sample $t_{\text{sim}} = k_{\text{sim}}\Delta t$ the closest to the mid-period $T_{\text{sim}}/2$. The wavelet component at time-frequency $(\xi_{\text{sim}}, t_{\text{sim}})$ shall thus be out of reach of end-point artefacts (see Methods section 2.4.2).
- (B2) Constraint (A5) can be addressed by virtue of propositions 3.2.1 (p. 65) and 3.1.16 (p. 56). On the one hand, $U_{\beta_\nu} \stackrel{d}{\approx} \chi_2^2$ and then

$$B_j^{\alpha_0, \beta_\nu} \stackrel{d}{\approx} \mathbb{1}_{]c(\alpha_0), +\infty[} (\mathbb{1}p^\Phi(\xi_j, t_k))$$

where $c(\alpha_0)$ is the $(1 - \alpha_0)$ -quantile of a χ_2^2 distribution and Φ is a fGn underlying the null hypothesis of the async-test, *ie* having spectral slope β_ν . On the other hand, as we pointed out on section 3.1.3.1, all mean-normalised WPSg-coefficients of any fGn are almost equally distributed to each other,

and that also applies to white noise ($\beta = 0$). Therefore, it will suffice to simulate realisations of a Gaussian white noise. The targeted variance of the simulations can be arbitrary, because the *whitening* that brings about the normalisation ensures that all normalised local spectra of fGn have the same expected band power.

(B3) Finally, owing to eq. (3.25) and remark 3.1.17 (p. 58), the drawings or realisations of mean-normalised coefficients $\mathbb{I}p^\Phi$ of the simulated white noise Φ can be simply retrieved as $2K_\phi^{-1}p_{k_{\text{sim}}}^\phi$, where $p_{k_{\text{sim}}}^\phi$ are the WPSg-coefficients of each simulation ϕ of Φ corresponding to time t_{sim} . To estimate the accessible spectral intercept K_ϕ , three options are available.

i) If the whole spectrogram of ϕ has been computed, we can use its global spectrum $\mathcal{S}_\phi(\xi_j) = \left\langle p_{j,k}^\phi \right\rangle_k$ and compute

$$K_\phi \approx \sum_{j=0}^{J_0-1} \mathcal{S}_\phi(\xi_j) \Delta \xi_j$$

ii) If not, owing to the stationarity of ϕ , we can simply resort to a fast and dirty solution by taking

$$K_\phi \approx \sum_{j=0}^{J_0-1} p_{j,k_{\text{sim}}}^\phi \Delta \xi_j$$

that is, the instantaneous power $\Pi_\phi^{\text{loc}}(t_{\text{sim}})$ of the realisation ϕ at t_{sim} (see Methods, section 2.4.4).

iii) Alternatively, since the accessible variance and spectral intercept are easily related for white noise as

$$K_\phi \approx \frac{\sigma_\phi^2}{1/2\Delta t - 1/T_{\text{sim}}}$$

any time-domain estimation of the variance of the realisation will also do (see remark 3.1.2 on p. 41)

Parameter selection Taking into account the previous considerations, we have set the following parameters for our numerical implementation:

- Low cut-off frequency: $\xi_{\text{sim}} = 0.5$ Hz, which for $\Delta t = 5$ ms results in $J_0 - 1 = 184$ (see section 2.4.1 on p. 28 for details on the available time scales and equivalent pseudo-frequencies).
- Length of the realisation: $n_{\text{sim}} = 2^{12}$, that is, $T_{\text{sim}} = 20.48$ s. This allows enough margin to absorb the edge-effect of half-width $\sqrt{2} \frac{\eta_0}{\xi_{\text{sim}}} = 2.7$ s existing at both time end-points of the spectrogram at frequency 0.5 Hz.
- Parameter of the Bernoulli trial: $\alpha_0 = 0.001$, which on the one hand will give a reasonable significance level for each of the underlying J_0 marginal tests implicit in the definition of the statistic R . And on the other hand, with together with the choice of J_0 , the probability of R being 0 under the null hypothesis would *a priori* be close to $(1 - \alpha_0)^{J_0} = 0.83$. This seems a good trade-off for R to be both sensitive and specific to discern if a local spectrum departs from its coloured background.

Numerical implementation Putting the pieces together: each simulation will consist of generating a time-series ϕ of length $n_{\text{sim}} = 2^{12} = 4096$ of purely uncorrelated (discret) Gaussian noise with targeted mean 0 and standard deviation 1 (relaxation B2); then extracting its local power spectrum at index $k_{\text{sim}} = \frac{n_{\text{sim}}}{2} = 2048$; and finally drawing the value r that returns the relaxed statistic R , computed as

$$r := \frac{1}{J_0} \sum_{j=0}^{J_0-1} \mathbb{1}_{]c(\alpha_0), +\infty[} \left(\frac{2}{\Pi_{\phi}^{\text{loc}}(t_{\text{sim}})} p_{j, k_{\text{sim}}}^{\phi} \right)$$

(as per relaxations B2 and B3.ii).

Before computing the the WPSG-coefficients, each simulated time-series is *standardized*; that is, they are re-scaled so that their mean and variance are 0 and 1, respectively, and their sampling frequency is that of the logMUA signals ($\Delta t = 5$) ms. Whereas the demeaning has been a usual preprocessing step before computing the wavelet transform throughout this work, the variance homogenisation is to ensure that the variability of the sampled local spectra will only be attributable to the frequency domain of the time-series.

With these parameters we have generated $N_{\text{sim}} = 5 \cdot 10^5$ observations r . This allows to carry a Monte Carlo test with a minimum significance level of at least $1/N_{\text{sim}} = 2 \cdot 10^{-6}$. The saturation values of the distribution were found as: $\hat{F}(0) = 0.9667$ and $\inf(\hat{F}^{-1}(1)) = \frac{25}{J_0} = 0.1351$. That is, 0.9677 is the empirically estimated probability that $R = 0$ (note that this is sensibly bigger than the *a priori* value of 0.83 obtained from the independent Bernoulli trials). The other saturation value means that for simulated local spectra of white noise, no more than 25 out of $J_0 = 185$ WPSg-coefficients (the 13.5%) within the band $[0.5, 100]$ Hz are significantly high for a level $\alpha_0 = 0.001$ (see Figure 3.9).

Note that, in strict sense, this procedure generates $\hat{F}_{R, N_{\text{sim}}}$, the empirical null distribution of R , but solely for the spectral slope hyper-parameter $\beta = 0$. However, the relaxations above discussed justify that it can also be a reasonable approximation of the empirical distributions corresponding to different values of $\beta < 1$. For the sake of computational economy, it is this distribution that we have used to provide the p-values associated with the Monte Carlo test for every logMUA signal.

3.2.1.4 Raw asynchronous periods

In summary, rejecting H_0^{async} for an observed local spectrum at t_k means that, with a significance level α^{async} , we consider its frequency components too extreme as to have possibly been drawn from the asynchronous background that underlies the signal. That is, we are merely discarding asynchrony at that given instant, without any further assumption for now of what we are accepting as an alternative.

On the contrary, accepting H_0^{async} with a confidence level $\text{CL}^{\text{async}} := 1 - \alpha^{\text{async}}$, entails that, insofar as the statistic R is empirically able to measure, our local spectrum would be as regular and "featureless" as any one drawn from the fGn that fits the best the signal. Importantly: as such, our test provides a phenomenological definition of what asynchrony can mean for an electrophysiological signal as the logMUA.

Finally, a collection of contiguous time instants where the *async* null hypothesis has not been rejected, is what we will be calling a *raw asynchronous period*.




Figure 3.9: Empirical null distribution $\hat{F}_{R, N_{\text{sim}}}$ of the statistic R with hyper-parameters $\beta = 0$, $\alpha_0 = 0.001$, simulated by means of a Monte Carlo experiment with parameters $N_{\text{sim}} = 5 \cdot 10^5$, $J_0 = 185$ and $n_{\text{sim}} = 4096$. Dashed black line indicates the saturation value $\hat{F}(0)$; red one the saturation value $\inf(\hat{F}^{-1}(1))$.

Although unlikely to occur, because of the WPSg auto-correlations in time and frequency, a raw asynchronous period can be degenerated, *ie* of duration Δt . Later on, we will refine these states by imposing restrictions on their minimum duration, as well as on the minimum inter-state duration between them. This will be carried out taking into account the physical context of the analysed signal and the nature of the synchronous states. Namely, for each case (asynchronous or synchronous), any inter-state interval shorter than 100 ms will be merged with the surrounding states. Likewise, asynchronous (resp. synchronous) periods will not be allowed to last less than 100 ms. This refinement step will be carried out the last, once optimal detection thresholds have been determined (see section 3.2.3).

3.2.2 Detection of oscillatory activity in a band of interest

If the previous async-test allowed us to discard asynchrony, now the question arises of what ought to be understood by synchrony for a logMUA signal of the kind we are analysing. We would like to pin down what has made certain local spectra fail the async-test. That is, we are in search of a more specific test, one that would narrow the alternative hypothesis to H_0^{async} .

A local spectrum $\mathcal{S}_\nu^{\text{loc}}(t)$ of a signal ν fails the async-test whenever enough of its components are significantly high compared to the asynchronous background spectrum of the signal (see section 3.2.1.1 above). This implies that $\Pi_\nu^{\text{loc}}(t)$, the instantaneous wavelet power of the signal at time t (see Methods eq. (2.9)), must be high²⁶. Whether this is due to a broad-band surge in power, or to a localised one, is what the test we are next proposing aims at discerning. Whilst the latter would indicate that an instantaneous spectral peak is present—the hallmark of transient oscillatory activity—the former would on the contrary point towards the existence of a raised-in-power local spectrum lacking peaks, which should therefore not be deemed synchronous.

3.2.2.1 A test to detect synchrony

Let be ν the observed signal under the same modelling assumptions that we made for the async-test: a Gaussian signal which follows the inverse power law with spectral parameters β_ν and K_ν . In addition, we are now also interested in its wavelet total power or wavelet variance $\Pi_\nu \approx \sigma_\nu^2$ (see eq. (2.11) in Section 2.4.4). We would like to assess whether the instantaneous power of ν at time t_k over a (narrow) frequency band of interest $\Xi_0 \subsetneq [\xi_{J-1}, \xi_0]$, is high enough as to be worth of hinting at an oscillatory peak. Or put it other way: it is high enough as to discard its having been generated by any kind of stationary process with a peak-less spectrum. In this context, and under our modelling hypothesis of the logMUA (see mesh of hypotheses \mathcal{H}_0 (3.27) in Section 3.1.3.3, the latter might be assumed to be no other but a certain fractional Gaussian noise acting as a *transient* background noise: with same spectral slope and total wavelet power as ν . The rationale being that a signal might well experience a transient surge in broad band power, while its whole correlation structure, which is endowed by the spectral slope decay of the power law, would remains less affected.

Because we are going to deal with band-restricted powers, a word on notation first. When the frequency band Ξ_0 is fixed (as it will be), we will simply denote instantaneous wavelet power at time t_k of ν restricted to a Ξ_0 as

$$\Pi_\nu^{\text{loc}}(t_k) := \Pi_{\nu|_{\Xi_0}}^{\text{loc}}(t_k) = \sum_{\xi_j \in \Xi_0} p_{j,k}^\nu$$

and the restricted total wavelet power as

$$\Pi'_\nu := \left\langle \Pi_\nu^{\text{loc}}(t_k) \right\rangle_k = \sum_{\xi_j \in \Xi_0} \mathcal{S}_\nu(\xi_j)$$

²⁶The possibility of a local spectrum failing the test even though presenting a low power must be discarded. In fact, having enough below-average components will never compensate for the very high power of the ones triggering the async rejection. This stems from the definition of the intrinsic background spectrum and the heavily positive skewness of the WPSg-coefficients distributions (see comments on p. 66).

(see *ibid* eq. 2.9)).

Their random-variable equivalents would naturally result for a process X , in which case, we will rather consider their expectations, denoted for convenience as:

$$\overline{\Pi_X^{\text{loc}}}(t_k) := \mathbb{E} \left(\Pi_X^{\text{loc}}(t_k) \right) = \sum_{\xi_j \in \Xi_0} \mathbb{E} p_{j,k}^X$$

and

$$\overline{\Pi'_X} := \mathbb{E} \Pi'_X = \sum_{\xi_j \in \Xi_0} \mathbb{E} \mathcal{S}_X(\xi_j)$$

Notice that the latter is its band-restricted accessible variance $\hat{\sigma}_X^2 |_{\Xi_0}$ over the underlying time-frequency sampling mesh.

The test we describe below targets the same population as before, that is, the random local spectra at a given instant of a process X among whose finite realisations could ν be found. We assume that $\hat{\beta}_X = \beta_\nu$ and $\hat{K}_X = K_\nu$ are its estimated spectral parameters. Let $\hat{\sigma}_X^2$ be its accessible variance. We propose:

Population: $\mathcal{S}_X^{\text{loc}}(t_k)$

Feature of interest: $\overline{\Pi_X^{\text{loc}}}(t_k)$

Sample: $\mathcal{S}_\nu^{\text{loc}}(t_k)$

Statistic:

$$Q_{\Xi_0} : (p_0, \dots, p_{J-1}) \longmapsto \sum_{\xi_j \in \Xi_0} p_j \Delta \xi_j \quad (3.32)$$

That is, Q_{Ξ_0} is simply a discrete restricted-band integration. When applied to a local spectrum, it results:

$$Q_{\Xi_0}(\mathcal{S}_\nu^{\text{loc}}(t)) = \Pi_\nu^{\text{loc}}(t)$$

$$Q_{\Xi_0}(\mathcal{S}_\nu) = \Pi'_\nu$$

We will then simply call this statistic the band-power over Ξ_0 or instantaneous band-power if applied to local spectra.

Null hypothesis:

$$H_0^{\text{sync}} : \overline{\Pi_X^{\text{loc}}}(t_k) = \overline{\Pi'_\Phi} \quad (3.33)$$

versus the composite alternative

$$H_a^{\text{sync}} : \overline{\Pi_X^{\text{loc}}}(t_k) > \overline{\Pi'_\Phi} \quad (3.34)$$

where Φ is the fGn with spectral slope $\beta_\Phi = \beta_\nu$ and accessible variance $\hat{\sigma}_\Phi^2 = \hat{\sigma}_X^2$.

Sampling distribution: $F_Q(q | H_0^{\text{sync}}) = \mathbb{P}(Q_{\Xi_0} \leq q | H_0^{\text{sync}})$

The null distribution of Q will be estimated through the empirical distribution $\hat{F}_{Q, N_{\text{sim}}}$ of the observed values of Q for N_{sim} independent local spectra of a process simulated under the *sync* null hypothesis (*ie*, the Monte Carlo experiment).

Monte Carlo test: Then, for a fixed significance level $\alpha^{\text{sync}} > 1/N_{\text{sim}}$, we will reject H_0^{sync} whenever

$$\hat{F}_{Q, N_{\text{sim}}}(Q_{\Xi_0}(\mathcal{S}_\nu^{\text{loc}}(t_k))) > 1 - \alpha^{\text{sync}} \quad (3.35)$$

That is, under the null hypothesis, by introducing Φ , we are evenly spreading the excess of variance of X not attributable to its spectral background over the whole band. If in addition, there are reasons to assume that the expected global spectrum of X has at most one peak, and then H_a^{sync} means that this peak must be located in Ξ_0 . In fact, under this extra at-most-one-peak assumption:

$$H_0^{\text{sync}} \iff \mathcal{S}_X^{\text{loc}}(t_k) \sim \mathcal{S}_\Phi$$

And by virtue of the relation between spectral intercept and accessible variance (see Remark 3.1.2 in Section 3.1.1.2), the test equivalently consist to assess whether $K_\Phi = K_\nu$.

Therefore, if for t_k , the local spectrum $\mathcal{S}_\nu^{\text{loc}}(t_k)$ fails the test, we will accept the alternative: that local spectrum must be participating in a localised power surge, hinting to the existence of a, possibly transient, oscillation in the time domain.

3.2.2.2 Test parameters and implementation of the Monte Carlo experiment

Most of the considerations raised for the *async*-test are still valid. In this case we will also allow the same relaxations. We will again generate a sample of white noise local spectra simulations. The trick will be here too, to relate the observed empirical statistic \tilde{q} generated by a white simulation with the parameters β_ν and σ_ν^2 of our, background-coloured signal. Calling \tilde{p}_j the WPSg-coeffs of the white simulation, it is easy to check that taking:

$$q = \frac{1}{2} \sigma_\nu^2 (1 - \beta_\nu) (2\Delta t)^{1-\beta} \sum_{\xi_j \in \Xi_0} \tilde{p}_j \xi_j^{-\beta_\nu}$$

will do (see eq. (3.5) on p. 41).

Apart from the parameters used for the simulation, that will be the same as those used for the *sync*-test, we only need to decide the critical frequency band of interest. After inspection of the global spectra of the logMUA signals and some preliminary explorations, we have taken $\Xi_0 = [0.54]$ Hz. This band not only targets a meaningful reduced range of brain oscillations, but is also broad enough to allow for variability between channels and slice experiments.

Finally, let us note that the test could be carried recurrently for different target bands Ξ_0 . Interestingly, it can easily accommodate different priors for the slope of the fGn Φ or its spectral variance. This would allow to sweep across as many as scenarios.

As for the async-test, the sync-test will be repeatedly carried out along all the sample times t_k and a decision can be met for each instant. Collections of contiguous rejected instants will be called *raw synchronous periods*.

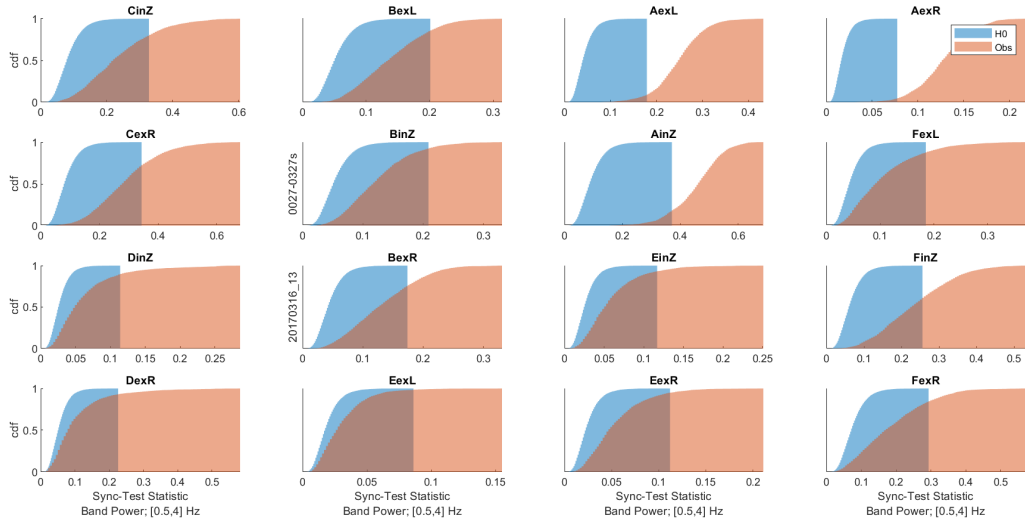


Figure 3.10: Simulated empirical null distributions (blue) compared to the observed population distribution (red) of the sync-test for each 16 channels of a given slice.

3.2.3 Multi-channel optimisation of the significance-level thresholds

Under provision of a reasonable simulation size for the empirical null distributions, the two tests above defined will each offer an objective assessment for each time-instant to be deemed either asynchronous, or synchronous. However, it is the combination of the two that will give them strength. In fact, if the async-test aims at being sensitive (*ie*, discarding non-asynchronous local spectra), the sync-test was proposed to enhance the specificity of the procedure (*ie*, detecting a synchronous instant).

Another crucial observation is that each test will be repeatedly carried out across the whole recording epoch, over samples that are by no means independent: neither experimentally independent (samples from the same signal of the same slice), nor statistically so (non-negligible time auto-correlation of the WPSg). Since sweeping the tests across times amounts to carrying out a massive multiple-test procedure, the interpretation of significance levels α^{async} and α^{sync} should not be taken *prima facie* or be accordingly adapted.

In our case, with our objective in mind of decomposing the loMUA in the time domain into synchrony and asynchrony, we are rather interested in exploiting the

complementarity of the two tests, and the size of the multiple testing, in order to find significance values *a posteriori* that would maximise the classification procedure. In addition, we will also resort to the information that all signals of different channels but simultaneously recorded offer. The aim is to provide, for each experiment (*ie*, slice) a pair $(1 - \alpha^{\text{async}}, \alpha^{\text{sync}})$ that would optimise the classification procedure for all recordings.

For each set of 16-signals of one experiment, we propose the following strategy to obtain what we have called its classification coverage:

- For each signal and for each value $0 < \alpha^{\text{async}} < 1$, we carry out the N async-tests. The same is done for the α^{sync} levels for the *sync - test*. A collection of *posterior p-values* is obtained for each threshold (Fig. 3.11).
- For each pair $(1 - \alpha^{\text{async}}, \alpha^{\text{sync}})$, each instant t_k is flagged according to the results of the tests: *correctly classified* (asynchronous if both tests are negative, or synchronous if both are positive), *non-classified* (async-test positive, sync-test negative) or *over-classified* (async-test negative, sync-test positive).
- For each signal, its posterior *classification coverage* is then considered, each pair $(1 - \alpha^{\text{async}}, \alpha^{\text{sync}})$ giving the percentage of instants falling in either three classification cases above defined (Fig. 3.12).

The idea is to optimise a threshold pair so that the classification coverage is maximised for all 16 signals together. The optimisation procedure sets an extra number of rules that seem desirable: it is better to have under-classified instants than over-classified ones; the theoretical entropy (Fig. 3.13) of the correctly-classified must be jointly-maximised (*ie*, allowing for the maximal variety) for all channels (Fig. 3.14)

These considerations allowed us to decomposed unambiguously each signal into its sync-async periods, according to a pair $(1 - \alpha^{\text{async}}, \alpha^{\text{sync}})$ that will be shared by all the channels of the slice (the experimental unit).

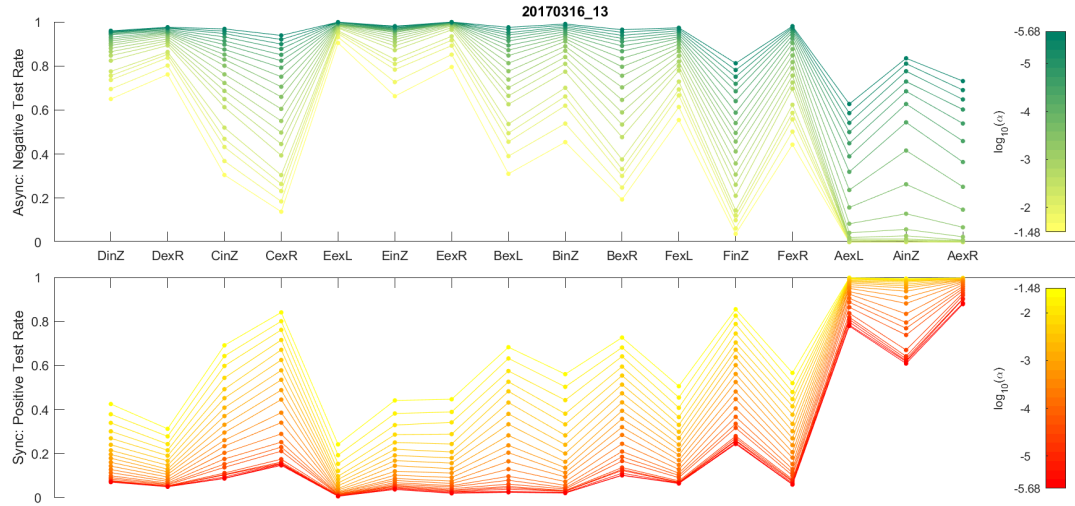


Figure 3.11: Resulting posterior detection rates for each $1 - \alpha^{\text{async}}$ (top in greenish) and α^{sync} (bottom, redish) values

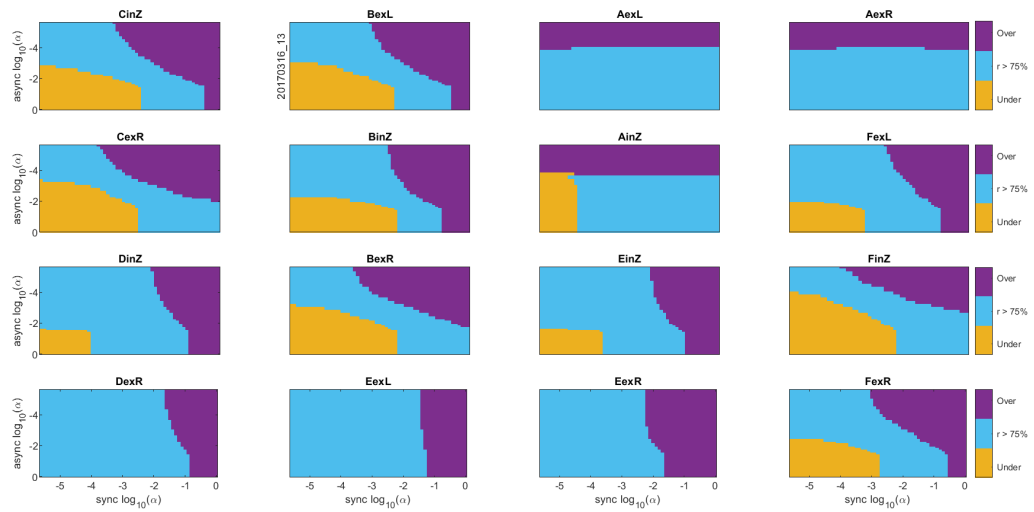


Figure 3.12: Discretisation of the classification coverage for each 16 channel of the same slice. Light blue region means over 75% of instants are correctly classified for a given $(1 - \alpha^{\text{async}}, \alpha^{\text{sync}})$.




Figure 3.13: Probability entropy resulting from the correctly classified instants for each $(1 - \alpha^{\text{async}}, \alpha^{\text{sync}})$ value.

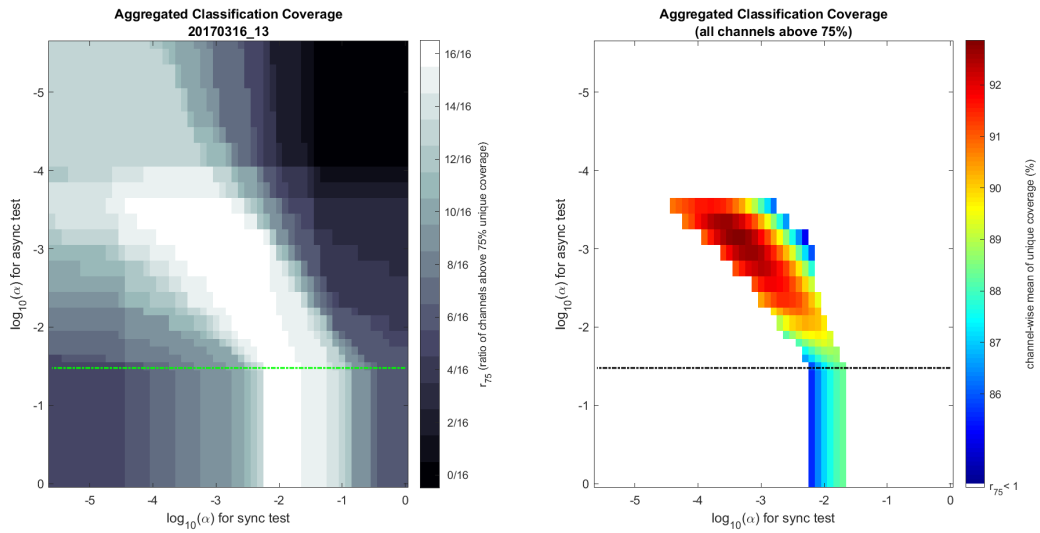


Figure 3.14: Joint optimisation of the classification coverage, for maximal entropy and minimal over classification. Observe the saturation value of the null distribution for the async-test.

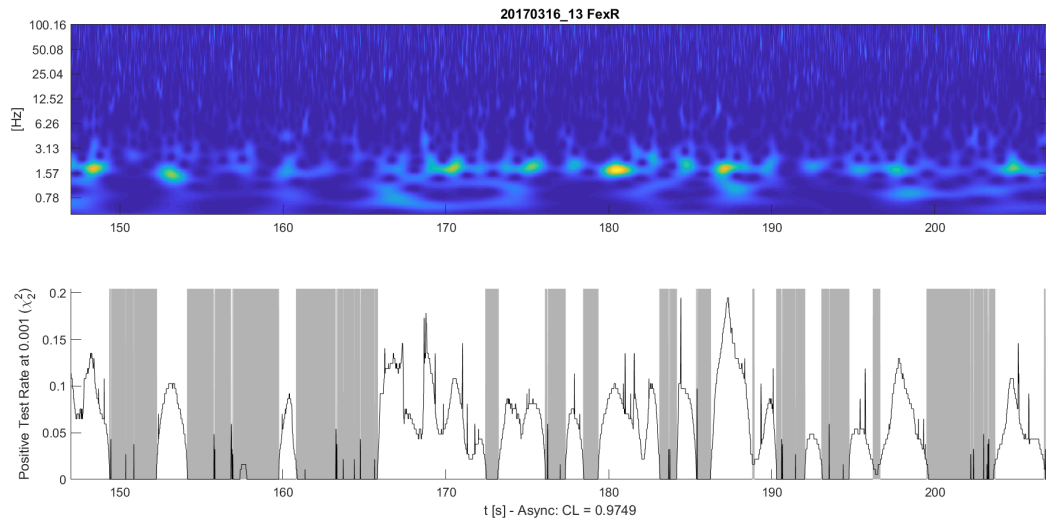


Figure 3.15: Example of the resulting async-detection for a logMUA WPSg.

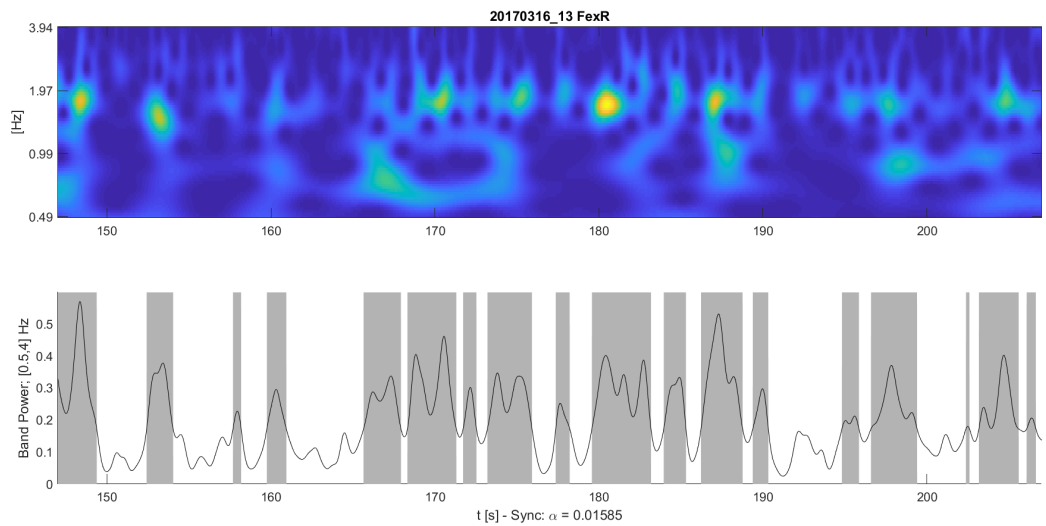


Figure 3.16: Idem as above, for the sync-detection.




Figure 3.17: Example of resulting classification for the whole set of 16 and the optimised pair of significance-levels.




Figure 3.18: Final classification coverage for the optimised pair.

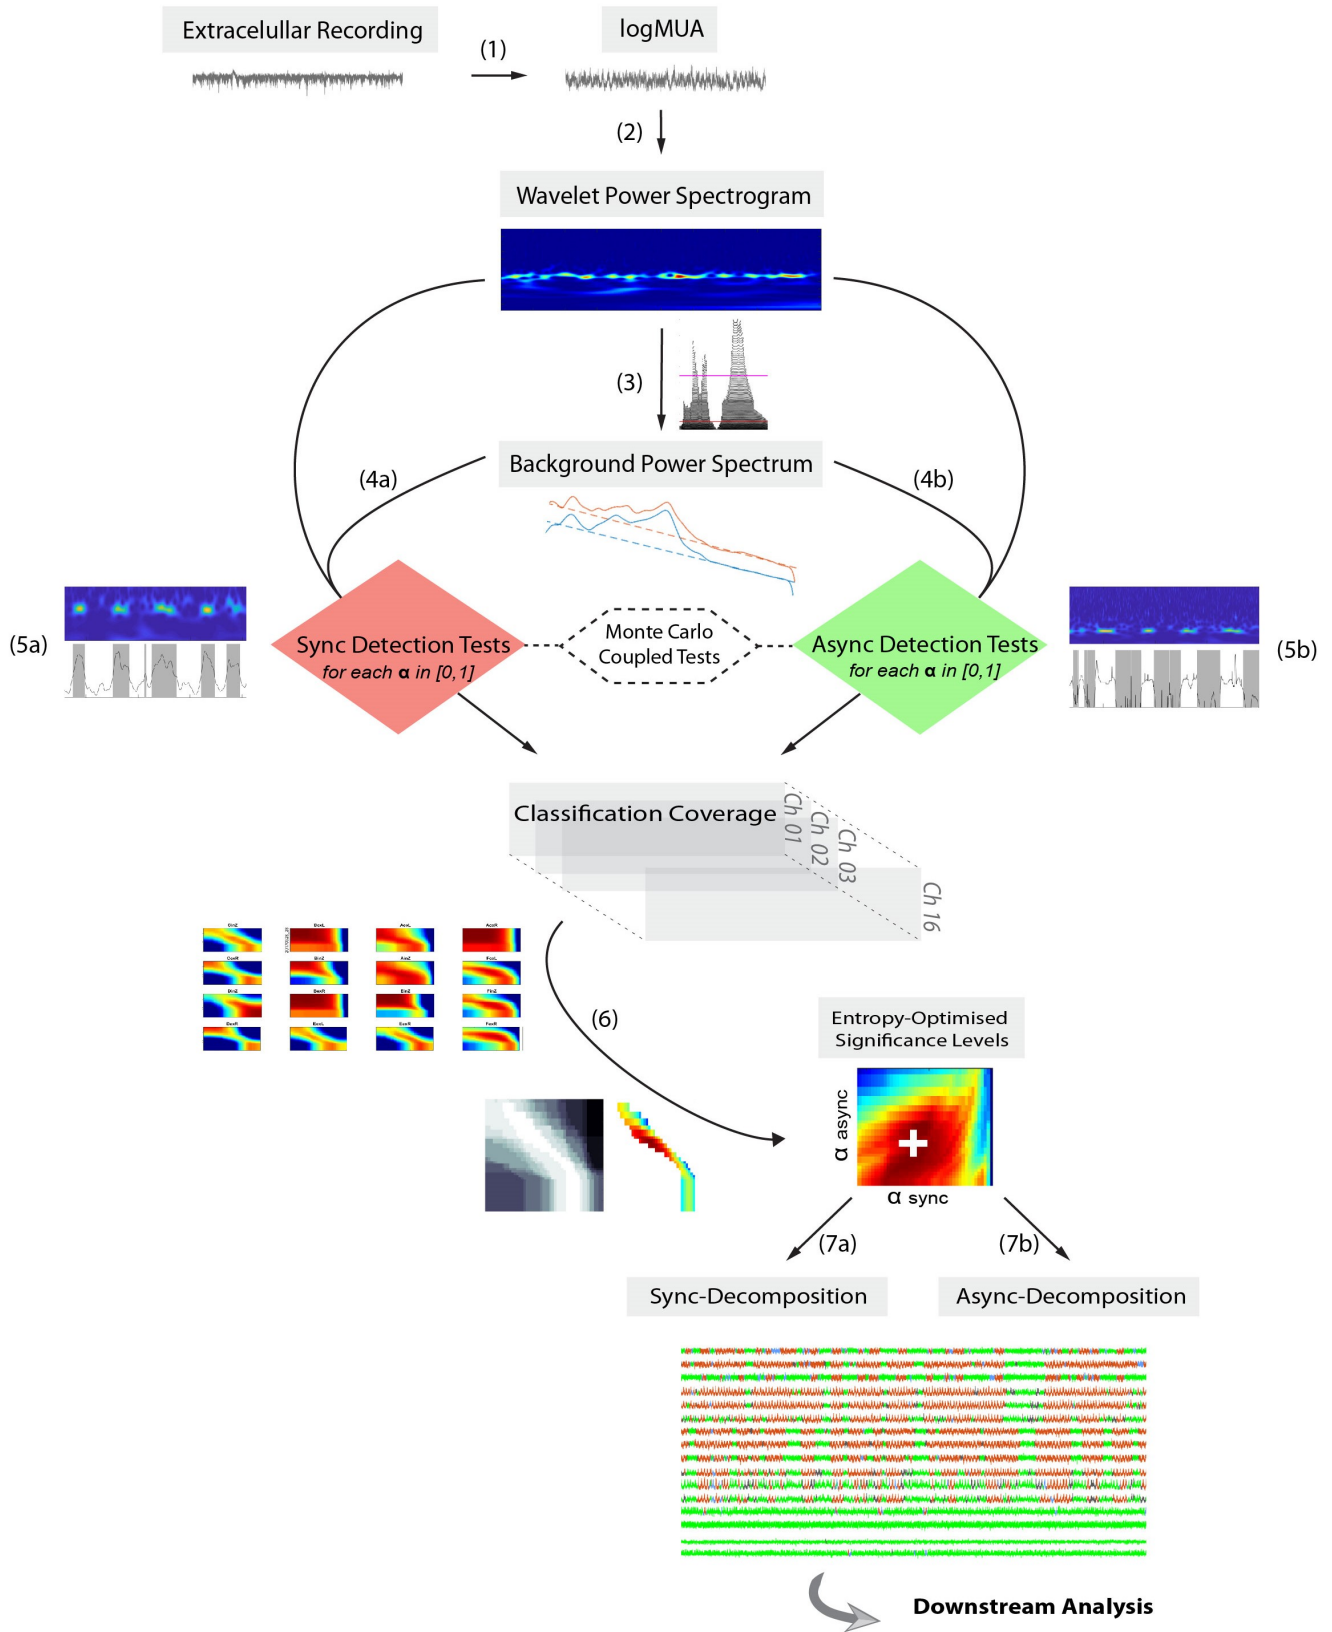


Figure 3.19: Sync/Async decomposition pipeline: our Monte Carlo Coupled Testing (MCCT)

Chapter 4

Data Analyses and Biological Results

4.1 Alternation of synchrony and asynchrony in the awaking-like state

Key message: The Awaking-like State (AS) can be decomposed into synchronous and asynchronous periods

4.1.1 Experimental model

We analysed the multi-unit activity (MUA) of signals recorded using a multi-electrode array (MEA) in ferret visual cortex slices to investigate the transition from a default state of slow oscillations (SO) towards a higher complexity regime, the awaking-like state (AS); which is elusive to current methods [Jercog *et al.* (2017), Tort-Colet *et al.* (2021)]. To this end, we devised a novel approach based on the wavelet power spectrogram (WPSg) of the MUA logarithm ($\log\text{MUA}$) (see Methods, sections 2.3 and 2.4 for details).

Cortical slices exhibited robust, time-persistent and highly synchronized slow oscillations (SO) in the range 0.5–1 Hz, with low variability (Figure 4.1).




Figure 4.1: Representative signals obtained from a 50-second sample of an extracellular recording (ECR) from a single electrode during the SO regime. From top to bottom: broad-band ECR; band-pass filtered ECR at 200–1500 Hz frequency band (or raw MUA); $\log\text{MUA}$ (see Methods section 2.3 for details on our own estimation of the MUA); normalised wavelet power spectrogram of the $\log\text{MUA}$ (see section 3.1 for details on our own normalisation of the spectrogram; in short, the inverse-law-background normalisation of the WPSg provides a measure of each WPSg coefficient’s significance, expressed in χ^2 -scores, allowing for a meaningful comparison between frequencies, channels, slices and conditions).

Application of cholinergic and noradrenergic agonists (see materials, section 2.1 and methods in [Barbero-Castillo *et al.* (2021)]) evoked a largely asynchronous regime, in which oscillatory (*ie*, synchronous) periods appear at a frequency *ca* 1.5–2 Hz and are intertwined with non-rhythmic (*ie*, mostly asynchronous) ones (figs. 4.2 and 4.3 top).




Figure 4.2: Idem as previous figure for a recording during the AS regime.

Such highly synchronized slow oscillation periods alternating with asynchronous ones have been described in rat cerebral cortex when coming out from deep anaesthesia [Tort-Colet *et al.* (2021)]. This alternation was replicated in a computational model using the brain stem as the subcortical pacemaker that would drive such infra-slow rhythm [Tort-Colet *et al.* (2021)], resembling mechanisms to model irregular SO regimes (*eg*, [Sanchez-Vives *et al.* (2017), Sancristóbal *et al.* (2016), Jercog *et al.* (2017)]). However, the identification of this same pattern in an isolated cortical tissue *in vitro*, indicates that such infra-slow rhythm is intrinsic to the local network, in line with some recent models of the transition *in vivo* from deep to light anaesthesia [Pazienti *et al.* (2022)]. Thus, to understand and better characterise this phenomenon in cerebral cortex slices, we developed a novel method to decompose the awake-like state (AS) recordings into periods of synchronous and asynchronous dynamical regimes.

4.1.2 A new method to discern synchronous from asynchronous periods

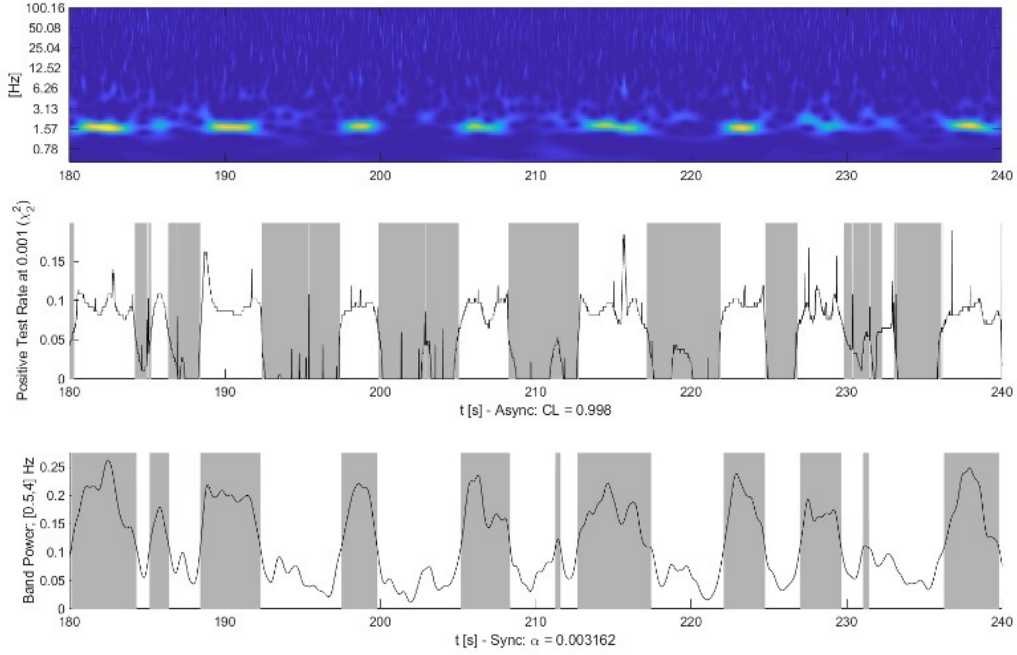


Figure 4.3: Representative example of the wavelet-analysis-based Sync/Async decomposition of a single-channel 60-second logMUA sample. **Top:** normalised wavelet power spectrogram of the logMUA. **Centre:** time-series of the R statistic specifically defined to test for asynchrony (R is the positive test rate at 0.001 marginal significant level of the multiple comparison of the local wavelet spectrogram coefficients against the asynchronous background). The distribution of this statistic under the null hypothesis that the local spectrogram is drawn from an asynchronous coloured noise, has been empirically derived by means of a Monte Carlo experiment ($N_{\text{Sim}} = 10^5$). We define asynchronous states (grey bars) as the time intervals where the probability of the values taken by R remains within a certain level of confidence (in this example $CL = 0.998$). **Bottom:** time-series of the instantaneous wavelet band-restricted power over the band $\Xi_0 = [0.5; 4]$ Hz. (Ξ_0 has been determined as that whose support provides the most divergent—and as such, the richest in information—spectral density when compared to that of the channel’s intrinsic background noise.) The significance of this instantaneous power is assessed against the distribution of the instantaneous wavelet Ξ_0 -restricted powers of a coloured Gaussian noise with inverse power law that of the channel’s intrinsic background, and variance that of the broad-band logMUA signal. This null distribution has been empirically estimated by means of a Monte Carlo experiment ($N_{\text{Sim}} = 10^5$). Synchronous states (grey bars) are then defined as the time intervals where the Ξ_0 -restricted instantaneous power is marginally significant at a certain level α (in this example, $\alpha = 0.003162$). (For details on how the pair (CL, α) has been determined, see section 3.2.)

We resorted to a three-step statistical procedure (see previous chapter, section 3.2 and summary figure 3.19). First, we evaluated how much the normalised power coefficients deviate from the mean background noise, by assessing the significance, at each instant, of the ratio of significant wavelet power spectral coefficients over a broad band of available frequencies ($\Xi = [0.5; 100]$ Hz). Thus, for a given confidence level, a period of asynchrony consists of consecutive time points whose instantaneous power spectrum is in average as flat as the instantaneous spectra derived from the signal's inferred background noise (the null hypothesis for discarding asynchrony detection). Empirical null distributions were obtained through Monte Carlo experiments (see section 3.2.1 for details). Shaded areas in Figure 4.3 (central panel) indicate these detected asynchronous periods (see third step below for the optimisation of the significance threshold).

Second, we investigated whether the remaining periods (non-shaded time windows in Figure 4.3, central panel) would map meaningful oscillations, that is, they are indeed synchronous. We concurrently tested the strength of power surges in a critical band of interest $\Xi_0 = [0.5; 4]$ Hz (see section 3.2 for the details). To this end, we compared, at each time point, the band-restricted instantaneous power against the distribution of powers that the channel's intrinsic background noise would have carried should it have the same variance as the broad-band signal (the null hypothesis for synchrony detection). Shaded areas in Figure 4.3 (bottom panel) depict these surges of significant periodicity.

In the third and final step, we validated whether these two detected regimes (asynchronous and synchronous) are capable of covering multiple channel recordings over the entire slice. If their coverage is high, this would suggest that the AS is effectively a combination of asynchronous and synchronous periods, reminiscent of *in vivo* anaesthetised settings [Tort-Colet *et al.* (2021)]. We determined optimal significance thresholds by jointly maximising the entropy of the classification for the whole set of channels in each slice, constrained to a space of maximal coverage and minimal overlapping (see details in Section 3.2 of Chapter 3).




Figure 4.4: Resulting sync/async decomposition over the logMUA signal of the same example showcased in previous fig 4.3. The colours indicate its classification coverage: in red, uniquely classified synchronous states; in green, uniquely classified asynchronous ones; in blue, overlapping synchronous and asynchronous states (over-classified). Finally, in dark grey appear the remaining non-classified time intervals (under-classified).

Consistently, the classification accuracy was high. We detected both synchronous and asynchronous periods with 89.20 (2.34) % uniquely classified periods; only 4.55 % (1.42) periods overlapped and 6.24 (1.85) % were non-classified (results reported as mean (SD); $n = 13$ slices). That is, AS recordings can indeed be decomposed into states of synchrony and asynchrony that seldom overlap and cover most of the recording (see Figure 4.4 for an example of the classification coverage).




Figure 4.5: Resulting sync/async decomposition over the logMUA signal of the same example showcased in the previous Figure 4.3. Colours indicate its classification coverage: in red, uniquely classified synchronous states; in green, uniquely classified asynchronous ones; in blue, overlapping synchronous and asynchronous states (over-classified). Finally, in dark grey appear the remaining non-classified time intervals (under-classified).

Moreover, most channels in all slices show an alternation between synchronous and asynchronous activity: the inter-quartile range of asynchrony content per recording is 41.33 (13.01) percentage points [p.p.], and the percentage of alternating-rich channels per slice (defined as those having at least a 20–80 % of both synchrony and asynchrony content) is 62.50 (13.26) % (see Figures 4.5 and 4.6). Furthermore, the regularity observed in such state alternation points to the existence of an infra-slow rhythm (see *eg* Figures 4.4 and 4.6).




Figure 4.6: Representative examples of sync/async decomposition over the multi-channel logMUA signals. A 120-second sample from two different slices are shown (slice #03 **top**, slice #05 **bottom**). The slices exhibit various levels of synchronicity throughout their channels (refer to Figure 4.5 for their overall synchronicity content). Synchronous periods, shown in red, were detected at the stated significance level $\alpha = 3.981 \cdot 10^{-2}$, and $\alpha = 6.31 \cdot 10^{-3}$, respectively. Asynchronous periods, in green, were detected within a confidence level $CL = 97.49\%$ and $CL = 99.68\%$, respectively. Over-classified periods are in blue; the non-classified, in black.

4.2 Spatio-temporal pattern of the alternation

4.2.1 Spatial correlation of the synchronicity tensor

Key message: The synchronicity tensor correlates spatially across the MEA, with columnar organisation having the strongest effect.

We then investigated whether the synchronicity tensor of the AS recording is spatially organised. We refer to the asynchrony (resp. synchrony) content as the proportion of time a signal spends in an asynchronous (resp. synchronous) state, with respect to the total recording time, at a given spatial location. The term *synchronicity tensor* merely refers to the pair of complementary sync/async contents.

The channels of the multi-electrode array (MEA) span at least three cortical columns. Transversally, two main laminar regions (infra- and supra-granular layers) are considered, and groups of two or three closely located electrodes constitute nodes (Fig 2a).




Figure 4.7: Schematic of a ferret's V1 cortical slice depicting cortical layers and the location of the multi-electrode array (MEA). Electrodes will ideally lie on different layers, across various cortical columns. The 16 channels can thus be grouped onto 6 nodes (of 2 or 3 electrodes each), 3 columns (along the horizontal or longitudinal axis), and 2 laminar levels (infra- or supra-granular, along the vertical or transversal axis).

We found that the asynchrony (resp. synchrony) content is positively correlated in the MEA space (see Figure 4.8). That is, the further apart two channels are located, the higher the differences in synchronicity between them, as shown by the regressions of longitudinal coordinates distance against the excess of asynchrony content (slope = 8.86 p.p./mm (SD: 7.37); see Figure 4.9). This deterministic, strong linear correlation also suggests beyond-chance differences in asynchrony content.

At the structural level, node and column expressed asynchrony content differences stronger than layers. That is, electrodes located in the same column have a

similar asynchrony content which tends to differ between different columns. The differences between electrodes located in the infra- or supra-granular layers, although significant, were far less marked than their inter-column differences counterpart, suggesting that the columnar structure is the most determinant organisational level affected by the pharmacological manipulations to those slices in regards to overall synchronicity levels. On the contrary, the layer structure appears only weakly affected by it. (Differences in asynchrony content: intra-node vs inter-node, $p = 1.22 \cdot 10^{-4}$; intra-column vs inter-column, $p = 3.66 \cdot 10^{-4}$; intra-layer vs inter-layer, $p = 0.0133$; paired Wilcoxon signed rank test, $n = 13$ slices) (Fig 2d).

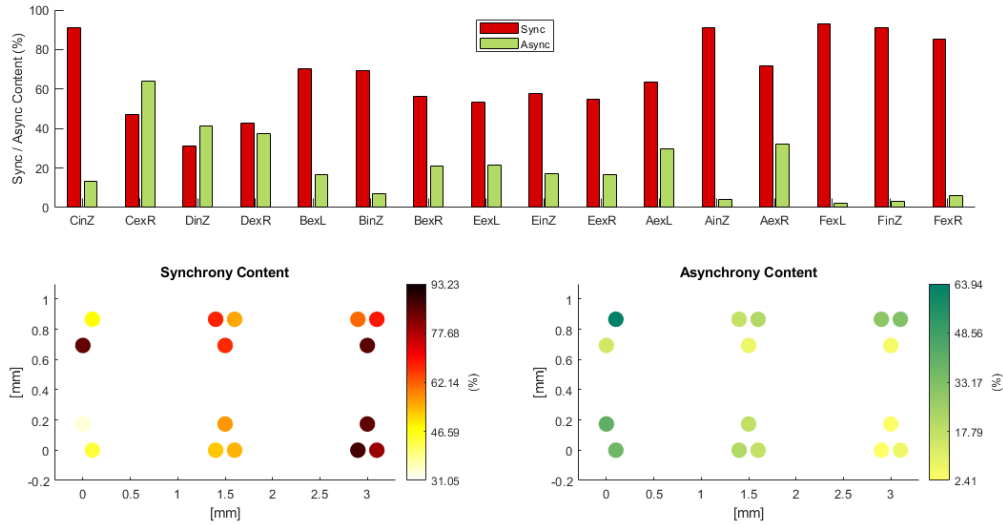


Figure 4.8: Representative example of the spatial distribution of asynchrony content for all 16 channels of a single slice. Asynchrony (resp. synchrony) content is defined as the proportion of time a signal spends in the asynchronous (resp. synchronous) state, with respect to the total recording time. **Bottom left:** The location of each dot gives the physical spatial layout of the multi-electrode array. **Bottom right:** Idem as above for the synchrony content of the same representative slice.

Since synchrony and asynchrony content are complementary, Comparable results arise when correlating spatially the synchrony content, since it is complementary to the asynchrony content for a given channel. Interestingly, the spatial correlation of the synchronicity tensor remains largely independent from other signal features, notably from its mean firing rate (see *infra* section 4.4).

Methodological Notes 4.2.1.

- (i) Against the use of the standard error of the mean (SEM) for merely descriptive results: although it may well be used in the context of an estimate or a test for the contrast of the means (see below), the SEM is not a descriptive statistics. The standard deviation (SD) is to be preferred in such cases as here.
- (ii) Classification coverage given with the SD of whole-slice means, not of pooled individual channels.




Figure 4.9: Asynchrony content's spatial correlation. **Left:** relative differences in asynchrony content (in percentage points) against the longitudinal distance between channels (ie, along the horizontal axis of the MEA layout). Each dot represents the mean relative difference in asynchrony for all such channels within a given slice that are equally set apart along the horizontal direction. Linear regression lines for each slice's spatial correlation are also included (dashed lines). Dot and line colour helps to identify each distinct slice. **Right:** factor-level marginal comparisons of the mean relative difference in asynchrony content for each MEA's structural level (from left to right: node, column and layer), segregated by differences between channels that lay within (intra) and between (inter) the factor of interest. Units in the ordinate axis and dot colour as in figure 3c. $n = 13$; (***) $p < 0.001$, (*) $p < 0.05$.

- (iii) Relative differences in async-content ($asyncR$) are computed as the $[0, 1]$ -scale-normalised absolute differences between each pair of electrodes' $asyncR$, with respect to all the absolute differences between electrodes of a given slice.
- (iv) In order to avoid giving more regression-weight to overrepresented coordinate-distances, only unique distances are taken into account as predictor variables, with the mean of relative $\Delta asyncR$ per unique coordinate as the corresponding response variable.
- (v) When comparing intra- vs inter-node (resp. column or layer), differences of asynchrony content between each pair of channels were considered, pooled together irrespective of their location in the slice. For the inter-node and inter-column comparison, only differences between pairs of channels located in adjacent columns were considered, to prevent giving different treatment to channels located in the inner columns.
- (vi) Let us finally note that an ANOVA-type test would be ill-suited for this experimental design. In fact, considering the recording per slice as the elemental experimental independent replicate, the three organisational structures (layer, column and node) can be neither viewed as levels of the same factor, nor are themselves factors from which to draw independent observations. For example, for the col vs lam contrast, each slice provides 4 measures (intra-col, inter-col, intra-lam and inter-lam). These observations would of course not be mutually independent. Although other test designs could address this shortcoming (eg, a randomised sampling from different slices for every condition to be tested), our number of independent replicates ($n = 13$) would fall short of providing enough statistical power to them.

4.2.2 Co-occurrence of asynchronous states

Key message: The sync/async alternation pattern is well conserved within columns, but can also occasionally propagate across columns.

Asynchronous periods tend to appear simultaneously in channels in close proximity (nodes). Interestingly, they are also well replicated within columns and only occasionally in channels located far apart in different columns (see Figures 4.6 and 4.10 for an overview). To have a better grasp of this phenomenon, we isolated the asynchronous periods from different channels and aggregated them in function of time. This revealed a high degree of time co-occurrence among certain channels, depending on their location (see figure 4.11).




Figure 4.10: Representative example of sync/async decomposition over a 60-second period of all 16 logMUA signals. Following the colour code of fig. 4.4, asynchronous periods are shown in green.

In order to quantify this co-occurrence in asynchrony, we computed the Jaccard similarity index of various grouping aggregations of asynchronous states. We were interested in two main organisational levels of the spatial layout: node and column. Because of the disparity in time structure for each of those asynchrony-aggregated time series, we tested the resulting Jaccard coefficients against the empirical distribution of the Jaccard coefficients yielded by surrogates of each asynchrony-aggregated time-series. (The surrogates were obtained by keeping the asynchrony time-structure of each series, that is, only permuting the aggregated periods of asyn-

chrony, all in keeping the inter-interval period duration). In addition, this bootstrap procedure provided us with a significance assessment for each comparison.




Figure 4.11: Representative example of asynchronous periods spatio-temporal co-occurrence from the same multi-channel signals shown in previous figure (4.10). Async states are aggregated by node (as their colour indicates) and overlapped for all 6 nodes. The *co-occurrence ratio* in the y-axis is the ratio of channels (spatially weighted by columns) that find themselves simultaneously in an asynchronous state.

We found that asynchronous periods do occur simultaneously in distinct channels. Channels in the same node, and those located in the same column show the highest degree of co-occurrence in asynchronous periods (see Figure 4.12). Strikingly, there were also some scant, but insightful, evidences of transcolumnar asynchrony co-occurrence.




Figure 4.12: Single-slice quantification of the spatio-temporal coherence of asynchronous states' occurrence. Similarity matrices are shown for each level of aggregation: channel (left), node (centre) and column (right). For each level (*ie*, channel, node, etc.) and each instance couple (*eg*, channel *FinZ* against channel *BexR*, or node *A* against node *B*), the significance of their async states' co-occurrence was obtained from the Jaccard coefficient of surrogate simulations of their intersections ($N = 10^4$, for column- and node-aggregates' interactions; $N = 10^3$ for the individual channels' interactions).

Overall, this is the first step towards the evidence that transitions between synchronous and asynchronous states underlie a column-wise organisation, and seem also able to propagate across the whole slice.

4.3 Columnar SO-like activity during AS synchronous periods

Key message: The synchronous periods consist of highly regular alternating Up- and Down-like states that are dynamically coordinated within cortical columns.

We next investigate whether these spatio-temporally fragmented synchronous periods of the awaking-like state are akin to SO. Although fragmented—temporally, into periods of synchrony; spatially, into segregated cortical columns—the synchronous states share some of SO’s local features. Such spatio-temporal fragmentation offers the unique perspective brought about by the present study and its experimental model, which allows the investigation of the synchronous activity coordination in greater detail. Importantly, the spontaneous emergence of SO in AS recordings would challenge the accepted hypothesis stating that, at least in *in vivo* setting, alternating SO are a global phenomenon requiring an external drive of the network [Destexhe and Sejnowski (2003), Crunelli and Hughes (2010), Sanchez-Vives *et al.* (2017), Tort-Colet *et al.* (2021)].

4.3.1 Columnar coordination in the frequency domain

4.3.1.1 Normalisation of the wavelet power spectrogram

After having decomposed the MUA of the AS recordings into synchronous and asynchronous periods, we were interested in describing the former: what these periods are made of, and how the oscillations of the MUA in the AS recordings are spatially coordinated. A way of gauging the overall oscillatory activity of a signal is to consider its wavelet global power spectrum: the time-average of the wavelet power spectrogram over a long-enough recording time (here, 300s), which gives a smooth estimate of the signal’s Fourier spectrum (Figure 4.13, left).

The logMUA spectra, for all recordings (all channels and slices) followed an inverse power law of similar, but different, spectral slopes $0 < \beta < 1$ (median slope per slice 0.227 (SD: 0.119); inter-quartile range per slice 0.220 (SD: 0.074), with a maximum slope per slice of 0.449 (SD: 0.133)). This indicates that an infra-pink background noise pervades the logMUA signals. In addition, the overall power of each signal seems to be trumped by the power of the aperiodic background, which may vary considerably from one signal to another (Figure 4.13, left). Thus, in order to draw meaningful comparisons between each spectrum’s profile, an offset of both the total power and the background noise is required. To achieve precisely this, we have devised a method of our own that normalises the wavelet power spectrogram WPSg, and *a fortiori*, the global spectrum (Figure 4.13, right, and Figure 4.14). In particular, we have demonstrated that for signals following an infra-pink inverse power law, this normalisation provides Morlet power components that are approximately distributed as χ^2 random variables with two degrees of freedom (see section 3.1 on chapter 3).




Figure 4.13: Wavelet global power spectra of the logMUA of three channels from the same node of a given slice and their normalisation. **Left:** the global spectrum is the time-average of the wavelet power spectrogram (WPSg) of a signal during a long-enough period (here, 300s of recording time) and thus, captures the mean power-per-frequency profile of the signal. (Its units are in logMUA squared units per Hz). The dashed curves describe the inverse power law of each channel’s intrinsic background noise. Background inverse power laws (*ie*, $\mathcal{S}(\xi) \propto \xi^{-\beta}$, β denoting the spectral slope) have been estimated through a robust regression procedure. Even though all three spectra in this example correspond to channels closely located, that of channel *BinZ* appears to be much powerful than the others’; it also exhibits a higher spectral slope. **Right:** Global spectra resulting from the normalisation of the WPSg of the same channels pictured left. The normalised global spectra are now more similar, both in terms of power and shape. By offsetting the intrinsic coloured background noise of each logMUA, our normalisation aims at allowing meaningful comparisons between different signal’s spectra, as well as between powers allocated at different frequencies. Furthermore, we have demonstrated that normalised WPSg components of infra-pink coloured noises are approximately distributed as chi-squared random variables with two degrees of freedom, thereby providing a useful statistical framework to objectively assess surges in power at a given frequency. Thus, normalised WPSg-coefficients of logMUA signals can grossly be expressed as χ_2^2 -scores, and so too their time-averages (if a slight abuse is permitted). Finally, let us notice how all three different inverse power laws in the previous figure have now merged into a horizontal line of normalised power 2 (the expectation of a χ_2^2 , precisely). That is, the background has been reduced to a flat spectrum, which is the hallmark of white noise.




Figure 4.14: Effect of the WPSg normalisation on the signal's power. One of the motivations for normalising the spectra is to capture the effective power of significant synchronous activity in the signal, irrespectively of the strength and colour of its background noise. **Left:** strong positive correlation between power and background power shows that the signal's total power may indeed be trumped by the power of its asynchronous background (as inferred through the inverse power law of the global spectrum). Grey hollow dots represent individual channels' power. Regressions are fitted over slice means, here depicted by blue big dots. **Centre:** if normalised, power spectral components do not longer correlate with the total observed power of the signal, as desirable. Here, normalised mean local power refers to the time- and frequency-average of normalised WPSg-coefficients. However, as the very strong correlation on the **right** plot suggests, these normalised components do make a good job as proxies of what we have called the normalisation power gain. The latter, defined as the ratio between the total power of the signal and the power of its background (expressed in dB), provides a straightforward measure of the strenght of the signal's oscillatory component (*ie*, the peaks of the spectrum) relatively to its aperiodic background.

4.3.1.2 Peaks and shape of the global spectra

The normalised global spectra of logMUA in the AS recordings consistently show a main prominent peak of frequency 1.86 Hz (SD: 0.31 Hz) and normalised power 33.67 (SD: 29.84) expressed in mean χ_2^2 -score units, to showcase significance. That is, across different slices ($n = 13$), the synchronous periods of the AS recordings are dominated by a conspicuous oscillation of *ca* 1.5–2 Hz (see Figure 4.13 for an representative example of normalised global spectra).

Isolation of the synchronous component of the spectrum. However, for the case of non-stationary signals as these, a description of the oscillatory activity that were solely based on the mean global spectrum would arguably be inappropriate. In fact, the spectral peaks above considered not only capture the oscillations present in the logMUA, but they also reflect the synchronicity tenor of each recording. Namely, their power and shape (although not their frequency) are a compounded assessment of the relative intensity of the synchronous activity alongside the sync/async ratio: in the sense that the longer the asynchronous activity, the more flattened the global spectra will appear. Therefore, in order to have a more accurate picture of the overall synchronous activity, we have proceeded to decompose each spectrum into its synchronous and asynchronous components. This is done taking advantage of the flexibility of the wavelet analysis and the decomposition already carried out in the time domain (see section 4.1.2). Bearing in mind that a snapshot of the WPSg at a given time point gives an estimate of the instantaneous spectral content of the signal at that instant, we can obtain synchronous (resp. asynchronous) components of the spectrum by averaging instantaneous spectra over synchronous (resp. asynchronous) periods (Figure 4.15). (See methodological notes 4.3.1 for more details). The synchronous global spectrum (sync-spectrum) will thus render an image of the oscillatory activity of the recording fully.




Figure 4.15: **Left:** Representative normalised global power spectrum of the logMUA of a single channel, shown as the superimposed mean of instantaneous spectra (blue). **Right:** Corresponding synchronous (red) and asynchronous (green) spectra, shown as the superimposed mean of instantaneous spectra following the sync/async decomposition in the time domain. Care was taken to average over the same number of instantaneous spectra, to avoid artefacts resulting from variability in synchronicity tenor. Dashed horizontal lines represent the flat background spectrum resulting from the normalisation.

Comparison of spectral peaks in frequency and power. We then asked whether the structural organisation of the slice played a role in the frequency and power of this observed dominant oscillation. Namely, we assessed to which extent the intra-slice spatial variability of the sync-spectrum, in regards with the peak frequency and prominence, was attributable to the electrode’s location in the MEA (Figure 4.16).

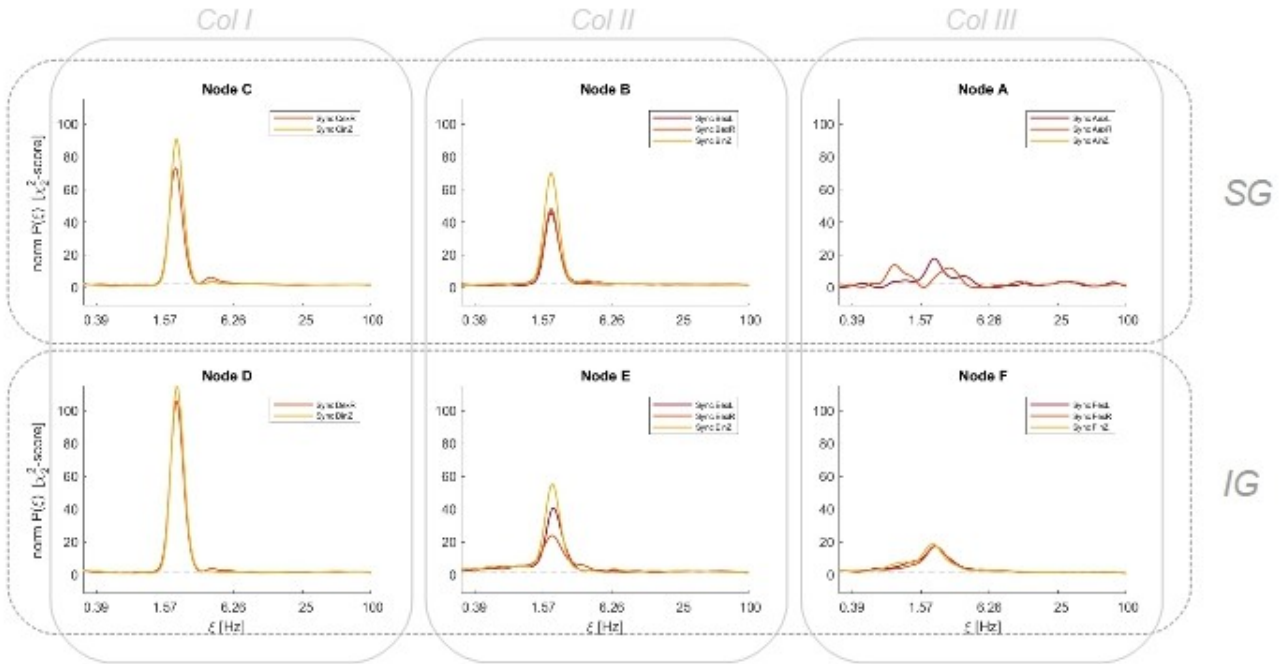


Figure 4.16: Normalised synchronous spectra of individual channels for a representative slice. Plots are grouped by nodes of the MEA (channel names correspond to those of Figure 4.18 below), organised according to distinct columns and supragranular (SG) and infragranular (IG) layers. Note the absent spectra in node A due to its lack of synchronous periods. Dashed horizontal lines represent the flat background spectrum.

Not unexpectedly, the main factor for channels to exhibit a similar spectral peak was their belonging to the same node. Indeed, signals recorded from channels this close may to some extent be treated as statistical replicates. More interestingly however was that both frequency and normalised power show a significant effect of belonging to the same column. That is, electrodes located in the same column but in different nodes tend to have similar peaks in power and frequency, as shown by the similarity between the first (SG) and second (IG) rows in Figure 4.16. However, these differ significantly between electrodes located in different columns. For instance, the mean intra-columnar difference in frequency (or beat frequency) is 0.0830 ± 0.0160 Hz, against 0.2150 ± 0.0360 Hz for the inter-columnar one (results given in mean \pm SEM, $n = 13$ slices). Accordingly, inter-laminar responses are undistinguishable in power and frequency, and do not permit to segregate between infra- and supra-granular layers. This effect shows in all relative differences between normalised peak power (inter-node *vs* intra-node, $p = 3.66 \cdot 10^{-4}$; inter-column *vs* intra-column, $p = 0.0012$; inter-layer *vs* intra-layer, $p = 0.863$; see Figure 4.17, left) and frequency

(inter-node *vs* intra-node, $p = 2.44 \cdot 10^{-4}$; inter-column *vs* intra-column, $p = 1.22 \cdot 10^{-4}$; inter-layer *vs* intra-layer, $p = 0.997$; Figure 4.17, right; paired Wilcoxon signed-rank tests, $n = 13$ slices). Consistently, an orthogonal contrast pitting intra-columnar against intra-laminar differences gave also significant results ($p = 0.0023$ for the relative difference in power, and $p = 1.22 \cdot 10^{-4}$ for the mean beat frequency; paired Wilcoxon signed-rank tests, $n = 13$ slices). For instance, spectral peak's normalised powers of channels laying in the same column but in different layers were $95.3 \pm 11.6\%$ different, whereas those of channels located in the same layer but in different columns were $178.2 \pm 23.9\%$ (relative differences computed as the exponential of absolute *logratios*; results given in mean \pm SEM, $n = 13$ slices).

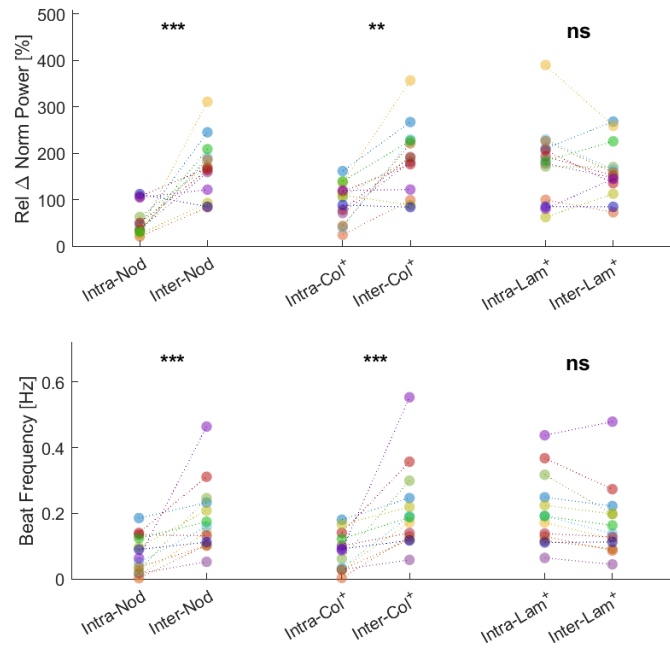


Figure 4.17: Influence of the cortical structural organisation on the power and frequency of the main oscillation of the logMUA's synchronous activity. **Top:** the relative difference between the main peak's normalised power of each pair of synchronous spectra was considered. [For each pair of synchronous spectra, we considered the relative difference between their main peak's normalised power.] (Relative differences computed as the exponential of the absolute *log* ratio between two powers.) For each MEA layout factor (node, column and layer) we compared these mean relative differences between two levels (ie, belonging or not to the given layout factor). Each dot represents the slice mean of the relative differences between channels belonging (intra) or not (inter) to a given location factor (node, column or layer). (+) Because of the strong effect exerted by the nodes, intra-columnar and intra-laminar means do not comprise intra-node differences. (***) $p < 0.001$, (**) $p < 0.01$, (*) $p < 0.05$, (ns) not significant. Dot colour codes for slice ($n = 13$). **Bottom:** same as left for the influence on the frequency of the main oscillation. In this case, spectral peaks's beat frequencies (ie, absolute differences between the frequencies) were considered.

Shape comparison of the spectra. To conclude, we investigated whether the previous result held for the whole spectral range, beyond the two features of the main spectral peak considered above. A functional distance between power spectra (the L^2 distance, see methodological notes below) allowed us to compare the overall shape between each pair of normalised sync-spectra. Again, the laminar level proved not significant to account for spectral similarity, in contrast to the columnar level and the nodal level (inter-node vs intra-node, $p = 0.0039$; inter-column vs intra-column, $p = 0.0195$; inter-layer vs intra-layer, $p = 0.6289$; paired Wilcoxon signed-rank tests, $n = 13$ slices.) (Figure 4.19). A hierarchical clustering using the L^2 similarity distance between sync-spectra (see methodological notes below) illustrates this effect. In line with the testing results, for all slices we found that spectra tend to form clusters broadly encompassing cortical columns (see Figure 4.18 for an example in one slice).




Figure 4.18: Similarity of the normalised global power spectra for a representative slice. Dendrogram representation of the channel similarity according to the L^2 -distance between normalised spectra. Four non-degenerated clusters were selected, their colour matching that of the dots in the MEA schematic below. The name of each electrode (eg, *AinZ*) encodes for the node to which it belongs (capital letters *A-F*) and its relative position within this node (**i**nterior or **e**xterior; **L**eft, **R**ight or **Z**entrum).




Figure 4.19: Multi-slice quantification of the spatial organisation of the global power spectra. For each organisational factor (*ie.* node, column, layer) a test on the marginal means of the z-normalised L^2 -distances was carried out. Dots represent the mean per slice and condition (intra- *vs* inter-node, resp. column, resp. layer) of the z-normalised L^2 -distances between spectra. Again, dot colour codes for slice. (**) $p < 0.01$, (*) $p < 0.05$, (ns) not significant. $n = 13$ slices.

Taken together, results show that the synchronous spectral component of the awaking-like state (sync-spectrum) consists of a single main oscillation in the range of a regular SO, whose frequency, intensity and persistence over time are well-maintained within cortical columns, but not between different cortical columns. That is, the distinction between infra- and supra-granular layer plays to his respect no role at all. This effect is reminiscent to the decrease in spatio-temporal correlation of slow waves for decreasing levels of anaesthesia *in vivo* [Pazienti *et al.* (2022)], not observed cortical slices to the date.

Methodological Notes 4.3.1.

- (i) *Synchronous/Asynchronous global spectra: in order to prevent the global components to be unduly watered down by each other (because of the non-negligible cross-correlations between neighbouring local spectra), each decomposed spectrum is only made up of the 10% most synchronous (resp. asynchronous) instants or local spectra, following our sync/async decomposition procedure. Roughly speaking, the more powerful the peak of the normalised local spectra, the more synchronous the spectra is considered. Complementarily, the flatter a normalised spectrum appears, the more asynchronous it is deemed.*
- (ii) *Since the node grouping has proven the best to conserve the synchronous spectral properties, to asses the influence which the column (resp. the layer) have on these properties, the inter- vs intra-column (resp. layer) contrasts are carried out node-free. That is, we compare channels belonging to the same column but in different nodes, against channels belonging to a different column (and a fortiori, a different node too). Analogously for the intra- vs inter-laminar comparisons, as well as for the intra-column vs intra-layer contrast. Otherwise, the columnar grouping would be more influenced by the node effect than are the laminar ones. In fact, as each column is composed of two nodes of three (resp. two) channels each, 6 out of 15 (resp. 2 out of 6 for the 4-channel column) pairwise intra-column channel comparisons are*

between channels belonging to the same node. That is, 40% (resp. 0.33%) intra-column comparisons are intra-node. For the intra-laminar comparisons however, only 7 out of 28 pairwise channel comparisons are between intra-node channels; that is, a 25%.

- (iii) We define relative differences for two strictly positive quantities a and b as,

$$\Delta Rel := e^{|\log(a/b)|}$$

This can be interpreted as a unitless > 1 ratio and can therefore be readily expressed in % as $(\Delta Rel - 1) \cdot 100$. For example, a difference of 0% would mean both quantities are equal; a difference of 100%, that one is the double of the other. Thus defined, relative differences provide a meaningful measure to symmetrically compare two positive quantities where the original unit of their difference would lack of interest, as it was here the case for normalised powers.

- (iv) L^2 -distance between power spectra is simply defined as the functional L^2 -norm of the difference between the two real valued functions with broad-band support describing the spectrum of each signal. That is, for two spectra p, q defined over a shared frequency support Ξ ,

$$d_{L^2}(p, q) := \int_{\Xi} |p - q|^2$$

- (v) **Hierarchical clustering:** From the matrix of dissimilarities between spectral densities based on the L^2 -distance, we followed an agglomerative clustering algorithm to build a hierarchy of densities' clusters leading to a dendrogram. The complete or farthest-neighbour linkage criterion was used to calculate distances between clusters, as reflected by the node heights of the resulting tree (see eg [Everitt et al. (2001)]). This amounts to considering the supremum distance, noted d_{∞} , induced by the underlying L^2 -distance over sets of spectral densities. Namely,

$$d_{\infty}(C, D) := \max\{d_{L^2}(p, q) \mid p \in C, q \in D\}$$

where C and D are two clusters of densities. Heuristically, this method maximises differences between clusters and thus provides us with a way to discriminate groups of densities [Everitt et al. (2001)]. Two other standard linkage functions suitable for non-euclidean metric spaces were also considered (nearest-neighbour and UPGMA), both yielding similar dendrograms.

4.3.2 Decomposition of the oscillation into sync-Up and sync-Down states

The oscillation that dominates the synchronous component of the awaking-like states consists of two regularly alternating states. The main reason prompting us to consider them as distinctive states is the peak-valley imbalance of the oscillation, which resembles that of a very regular and accelerated form of SO [San Cristóbal *et al.* (2016), Dasilva *et al.* (2021), Pazienti *et al.* (2022)], and can thus be further decomposed into Up- and Down-like states (named here sync-Up and sync-Down sub-states). (See Figure 4.20.) We next resorted to an unbiased approach to characterize this alternating regime without assuming SO dynamics *a priori*. Hence, existing methods to detect Up/Down states are not necessarily suitable in the context of this identification [Chen *et al.* (2009), Reig *et al.* (2010), Jercog *et al.* (2017), D’Andola *et al.* (2017)]. Thus, we devised an analytical method to detect the sub-states of the synchronous periods, based on finding the extrema of a second-order differential expansion of the inverse-wavelet-transform–reconstructed logMUA signal on the [0.5; 4] Hz band (see Figures 4.21 and 4.23, and methodological notes 4.3.2 on p.106). Crucially, this allowed us to capture sync-Up and sync-Down states in a way that does not directly involve either the distribution of the MUA nor the duration of the states, thereby ensuring there is no bias towards those properties we were subsequently interested in analysing.




Figure 4.20: Representative example of detected sync-Up sub-states within synchronous periods in multichannel logMUA signals. The 60-second long signals were obtained from six simultaneously recorded channels located in two distinct cortical columns across different layers (see *eg* figs. 4.16 or 4.18 for the MEA layout channel-code). *logMUA* asynchronous periods appear in green, whereas, over the grey synchronous periods, the crenelated maroon trace indicate the sync-Up/Down alternation.




Figure 4.21: **Top:** Illustration of the procedure to detect sync-Up/sync-Downs sub-states over a short sample of logMUA signal (in grey in the background) during a synchronous period. The synchronous period is partitioned into time subintervals according to the $[-\pi; \pi]$ -rad cycles described by the instantaneous phase of the WPSg's main ridge in $\Xi_0 = [0.5; 4]$ Hz. The inverse-wavelet-transform f of the logMUA over the band Ξ_0 is superimposed in black. The ridge phase may be here interpreted as the phase of f . Note that f 's main peaks precisely occur around the centre of each cycle-interval (*ca* 0 rad). Short extracts of $T_2f(t_k, \cdot)$, the local Taylor polynomials of order 2, are plotted in red (for the Up onset) and blue (for the Up termination) at those time points t_k that mark the onset or termination of the Up state. The gap parameter h_0 is the time increment at which the maximisation (resp. minimisation) of the vertical increment of $T_2f(\Delta, h_0)$ is sought (see Fig S2b below). **Bottom:** Determining the extrema that allow to define the limits of the sync-Up states. In a grey scale are plotted the functions $f' + hf''/2$ for different values of h ranging from 0 ms (dark grey) to 200 ms (white). For each fixed h , these functions are proportional to $f - T_2f(\cdot, h)$, the increment of the 2nd order Taylor polynomial of the wavelet-filtered logMUA signal f). The function corresponding to our selection of parameter $h_0 = 50$ ms appears singled out in black. In each subinterval matching a $[-\pi; \pi]$ -rad cycle of the ridge phase (see Fig S2a above), the last maximum of $f - T_2f(\cdot, h_0)$ before f 's main peak defines the onset of the Up state, whereas the first minimum of that increment function after the peak defines its termination (or the start of the next Down state).

Methodological Notes 4.3.2 (Detection of sync-Up sub-states).

- (i) **Partition of each synchronous period** into time subintervals that map the cycle of the main oscillation: Let us remember that the synchronous periods of the signal are precisely those characterised by the presence of a conspicuous and sustained main oscillation in the $\Xi_0 = [0.5; 4]$ Hz band. The hallmark of this oscillation is a unique, continuous, single-component ridge in the normalised WPSg. For such a ridge, a phase function can naturally be defined by associating instantaneous phases from the corresponding complex wavelet components to the frequency-time points constituting the power ridge. Then, owing to its sawtooth nature, the phase function of the ridge induces a partition of the synchronous period into time intervals, each mapping (up to) a whole $[-\pi; \pi]$ -rad cycle of the signal's main oscillation.
- (ii) **Definition of sync-Up's onset and termination points** for each cycle-interval: In a second stage, we considered f , the analytical function resulting from the wavelet-filtered logMUA on the Ξ_0 band (ie, its inverse continuous wavelet transform over that band). At each time point t , its Taylor expansion of order 2 reads:

$$T_2f(t, h) = f(t) + f'(t)h + \frac{f''(t)}{2}h^2$$

Within each of the cycle-intervals previously defined, the main peak of f (ie, the maximum of its local maxima therein) is always reached near the centre of the interval (phase close to 0 rad), which broadly matches the upper section of each oscillation cycle. Thus, in this case, the ridge phase just defined, may be understood as the instantaneous phase of the band-filtered signal f . Now, for a fixed time gap h_0 , the last local maximum of the increment function $T_2f(t, h_0) - f(t)$ just before the main peak of f in each cycle-interval served us to define the Up onset. Symmetrically, the first local minimum of that expansion after the main peak of f defined the Up termination. Finally, sync-Down sub-states were dually defined (see Figure 4.21).

- (iii) **Selection of parameter h_0** : The parameter h_0 represents the forward time horizon at which the potential increase/decrease of the 2nd order polynomial approximation to f is to be evaluated. In practice, h_0 could be thought of as the relative weight given to the first and second derivatives in the Taylor expansion (Figure 4.21b).

Different choices of the gap parameter h_0 would allow for different Up-state definitions, appropriate in principle for different analyses. For example, for h_0 small or zero the first derivative component of the Taylor expansion would dominate, providing centred Ups detected around their peak. This would be most appropriate for the analysis of the mean MUA value or of MUA densities of Up states, but rather inadequate for studying their onset profile and shape (in this case, only the upper part of the oscillation around its peak would be detected as Up state, leaving any transitional activity out). In another case, for h_0 big or ∞ , the second derivative component would dominate, providing early Up onset detections long before its peak, and Up terminations just at their peak. This kind of detection would hardly be valid for the study of MUA densities but seemingly interesting for the study of precise state-change timings occurring just at the onset of their build-up.

In this work we have heuristically taken $h_0 = 50$ ms, which has proven a good compromise between the former two extreme cases. We have deemed desirable to anticipate the Up occurrence right from its onset, all in keeping it long enough to allow for its lasting inertia to come off just at the beginning of the following Down state.

- (iv) The novelty of our detection procedure rests on the use of the wavelet-based filter and, crucially, on the induced partition into cycles stemming from the power ridge. The detection based on a differential analysis is however more or less standard. If our method has proven so robust and seamless is because it operates on the smoothly filtered logMUA, within conveniently delimited alternation cycles. Importantly, through the selection of its parameters, our method is transparent about its influence on those features of the detected states that will be of interest for subsequent downstream analyses (features such as their MUA density, onset/termination timing and shape).*

4.3.3 Duration of the synchronous sub-states

Key message: Duration of Sync-Ups in supra-granular layers is robustly conserved accross columns and slices




Figure 4.22: Correlations between the intra-channel median period of the oscillation (abscissa) and the intra-channel median duration of each sub-state (ordinate), according to the laminar position of the channel (IG or SG) and the type of sub-state (whether sync-Up or sync-Down). For each case, the corresponding regression line is shown with its 95% confidence bands. Points close to the thin dotted line of equation $y = \frac{1}{2}x$ represent channels where the median sub-state duration makes up around half of the Up-Down alternation duration. Such is generally the case for IG channels. Normally, the longer the oscillation period, the longer the sub-states; except for sync-Ups in SG, whose duration appears to be strongly conserved.

oscillation, the proportionally longer one would naturally expect the sync-Ups and sync-Downs to be (filled dots in Figure 4.22; $T_{Up} = 0.016 + 0.397 \cdot T_{Osc}$, $p < 10^{-14}$,

We then quantified the median duration of each of the sub-states of the synchronous component of the AS regime, in different locations on the slice. The alternation rhythm between synchronous sub-states is at the basis of the main oscillatory activity detected in the frequency domain and described in a previous section (see section 4.3.1). As we have there shown, this rhythm is only guaranteed to be maintained within columns. Therefore, the comparison is *a priori* only pertinent between infra-granular (IG) and supra-granular (SG) channels belonging to the same column (see *eg* Figure 4.20 for an example of synchronous periods of distinct columns having different rhythms).

We found that sync-Up states were $31.60 \pm 9.14\%$ longer in IG than in SG. In fact, sync-Ups in IG account for $45.08 \pm 1.34\%$ of the corresponding channel's oscillation period, against $34.87 \pm 1.94\%$ for SG channels (results given in mean \pm SEM, $n = 13$ slices; $p = 1.22 \cdot 10^{-4}$, paired Wilcoxon signed-rank test). Strikingly, despite the differences in frequency and power observed for this oscillation between different columns, we found that the duration of sync-Ups in SG was very well conserved across columns. What is more, sync-Up's duration in SG was almost constant even among different slices, with a median duration per slice of 0.182 s (SD: 0.012 s). By contrast, this was not the case for the sync-Ups in IG.

Moreover, in IG both sync-Up and sync-Down their duration (T_{Up} , T_{Down}) depends linearly on the oscillation period (T_{Osc}): that is, the slower the oscillation,

$R^2 = 0.482$; $T_{\text{Down}} = -0.021 + 0.568 \cdot T_{\text{Osc}}$, $p < 10^{-29}$, $R^2 = 0.742$). Specifically, in IG each sub-state makes up for around half (40–60%) of the oscillation period duration, the sync-Ups being slightly, but consistently, shorter. However, in SG only the sync-Down duration linearly correlates with the oscillation period (violet circles in *ibidem* Figure $T_{\text{Down}} = -0.1827 + 0.9735 \cdot T_{\text{Osc}}$, $p < 10^{-63}$, $R^2 = 0.971$), with the sync-Up duration appearing to be uncorrelated (see Figure 4.22). Indeed, in SG the constant model ($T_{\text{Up}} = 0.182$ s, $p < 10^{-91}$) strongly outperforms the non-significant linear one ($p = 0.60$).

Taken together, in SG the median duration per channel of the sync-Ups remains strikingly constant regardless of the alternation frequency between synchronous sub-states. Consequently, the inter-channel variability of the oscillation frequency can be solely explained by differences in length of the sync-Down states in supra-granular channels, but not in IG layers.

Methodological Notes 4.3.3.

- (i) For analyses carried out in this section and ulterior ones (sections 4.3.4, 4.3.5 and 4.3.6), only channels with at least 3% of synchronous content were considered (that is, channels with at least 10 s of detected synchronous activity). This ensured having approximately no less than 15–20 sync-Up and sync-Down sub-states per channel. From the 13 slices analysed throughout, all but 4 had more than 14 out of 16 channels thus rich in synchrony (10 being the minimum number of admissible channels that only one slice presented). Importantly, all slices had synchronous-rich channels in both layers (SG and IG), and simultaneously in at least two out of their three cortical columns.
- (ii) T_{Up} and T_{Down} are measured as the median per channel of the durations of each sub-state; T_{Osc} as the intra-channel median of the duration of a Up-Down alternation (or equivalently, as the length of the cycle-intervals defined for the sync-Up detection, see Methods). Therefore, T_{Osc} closely coincides with the period of the main oscillation that was inferred from the power spectrum.
- (iii) Linear regressions of each sub-state (sync-Up or sync-Down) median duration are carried out over all the IG (resp. SG) channels from all the slices, pooled but spatially weighted (so that each node is given the same weight during the fitting).

4.3.4 Variability of the synchronous sub-states

Key message: The study of the intra-channel variability of Sync-Up/Down states in duration and mean MUA further characterises the SO-like activity between layers. In SG: variably long and stably silent sync-Downs, with well-conserved short sync-Ups. In IG: variably short and unstable sync-Downs, that are as noisy as the sync-Ups.

Although mostly regular, the frequency of the sync-Up/Down alternation slightly fluctuates in time depending on the duration of each instance of a sub-state, which in turn depended on their laminar location in the slice (Figure 4.22). Thus, the study of the intra-channel variability of these fluctuations may help to further characterise the synchronous sub-states with respect to their laminar location in the slice. Alongside, we also studied the variability in mean MUA of these states within individual channels, as well as the MUA variability within each occurring sub-state, here termed MUA roughness (see table 4.1 below, and Figures 4.24 and 4.25). Combined, these three analyses will give information on how stereotypical the occurrence of each sub-state is.




Figure 4.23: Representative example of six 30-second long logMUA signals (grey traces) from simultaneously recorded channels. The channels were located in different layers (IG: infra-granular, SG: supra-granular) from the same cortical column. The sample shown contains a number of asynchronous period that have been further decomposed into sync-Up (red) and sync-Down (blue) sub-states. In black, superimposed to the logMUA, are plotted the inverse-wavelet-reconstructed signals over the band [0.5, 4] Hz, which are instrumental in the detection of the sync-Up states (see methods in section 4.3.2). Note that, for the duration of this sample, channel *AexR* does not show any synchronous period.

Variability in duration. First, we measured the variability in duration with the coefficient of quartile variation (cq_v) of the durations of each sync-state type occurring within each channel. (The cq_v is a robust version of the coefficient of variation, see [Bonett (2006)]). We found that the duration stability of the sync-Ups in SG was also very well maintained at the intra-channel level. In the one hand, SG sync-Ups were much less variable in duration than SG sync-Downs (\log_2 fold-change from Down to Up of -1.080 ± 0.184 ; or equivalently, Ups are in average 46.24–58.36 % less variable in duration than Downs; $p = 1.22 \cdot 10^{-4}$ Wilcoxon signed-rank test). That is, in SG fluctuations in the oscillation’s instantaneous frequency are mostly accounted for by the variability of the sync-Down states, as found *in vivo* when anaesthesia levels are reduced [Pazienti *et al.* (2022)]. In IG, however, both sub-states have a similar duration variability; if so, slightly bigger for sync-Ups (\log_2 FC of 0.242 ± 0.065 , or equivalently: Ups are in average 13.12–23.71 % more variable in duration than sync-Downs). In the other hand, sync-Ups were also considerably more stable in IG than in SG: their duration’s cq_v being 52.37–63.68 % smaller ($p = 1.22 \cdot 10^{-4}$, Wilcoxon signed-rank test); whereas no significant difference in variability was observed for sync-Downs duration between layers ($p = 0.6848$, two-tailed Wilcoxon signed-rank test). (See Figure 4.24 left, and Figure 4.25).

	CQV Duration		CQV Mean MUA		Median CV MUA	
	Sync-Down	Sync-Up	Sync-Down	Sync-Up	Sync-Down	Sync-Up
SG	0.202 (0.097)	0.095 (0.046)	0.080 (0.038)	0.107 (0.037)	0.478 (0.062)	0.656 (0.094)
IG	0.171 (0.039)	0.210 (0.062)	0.131 (0.042)	0.084 (0.015)	0.615 (0.130)	0.613 (0.090)

Table 4.1: Intra-channel variabilities of synchronous sub-states. Group averages expressed as mean (SD). CQV: coefficient of quartile variation. CV: coefficient of variation.

Variability in mean MUA. The opposite occurs for the variability of the mean MUA exhibited by each synchronous sub-state instance: in SG, well conserved mean MUA for sync-Downs with respect to that of sync-Ups; whereas in IG, sync-Downs appearing the most variable (see Figure 4.24 centre, and Figure 4.25). More precisely, in SG the Down-to-Up \log_2 FC is 0.458 ± 0.183 (equivalently: the cq_v of the mean MUA is in average 20.94–55.93 % higher in Up than in Down; $p = 0.0164$, Wilcoxon signed-rank test); whereas in IG, this \log_2 FC is negative (-0.578 ± 0.124 , that is, the cq_v of the mean MUA is in average 26.96–38.51 % lower in Up than in Down; $p = 0.0012$, Wilcoxon signed-rank test). Notably, SG’s sync-Down are in average 38.96–38.51 % less variable in mean MUA than IG’s ($p = 1.22 \cdot 10^{-4}$, Wilcoxon signed-rank test).

MUA roughness. Finally, we also studied the median intra-state variability of the MUA values taken by each state type, according to its laminar location. This gives an idea of how rough or jagged each state occurrence is in average. We found that, whereas in IG, both sub-states are equally rough, in SG sync-Downs are significantly smoother than sync-Ups (see Figure 4.24 right, and Figure 4.25). Specifically, the median cv of sync-Up's MUA is similar (between 0.58 % lower and 3.58 % higher) to that of sync-Downs ($p = 0.273$, two-tailed Wilcoxon signed-rank test), and both measures appear strongly correlated with each other (slope estimate 1.008, $p = 3 \cdot 10^{-33}$ for a linear model without intercept, fitted on spatially weighted pooled infra-granular node averages).




Figure 4.24: Intra-channel variabilities of synchronous sub-states. **From left to right:** variability in duration, expressed as the intra-channel coefficient of quartile variation (cq_v) of sub-state durations; variability in mean MUA, expressed as the intra-channel cq_v of mean MUA per sub-state; and MUA roughness, measured as the intra-channel median of the MUA coefficient of variation per sub-state. For each intra-channel measure of variability, a plot illustrates the relation between the observed values for sync-Downs (abscissa) and sync-Ups (ordinate), segregated by laminar location (IG or SG). Each dot represents a node-average, that is, the mean of values measured for channels located in the same node. Thin dotted lines connect dots representing nodes located in the same column. Values of all 13 slices shown. The red (resp. blue) region delimits the subset where Up variabilities are higher (resp. lower) than those of Down substates.

In summary, in SG, we have short sync-Ups of well-conserved duration and long sync-Downs, which present a smooth and stable mean firing rate. In IG, rather the opposite occurs, with shorter, more unstable sync-Downs, as rough in value as sync-Ups. This indicates that SG are relatively more silent layers, as their comparatively long sync-Down states display almost no neuronal activity, while in IG, sync-down states do contain non-negligible activity. This seems consistent with longer IG sync-Downs (see previous section) containing less coordinated—and hence more stochastic—activity, specially towards the end of the Down state, as recently described for classic SO [Camassa *et al.* (2022)].

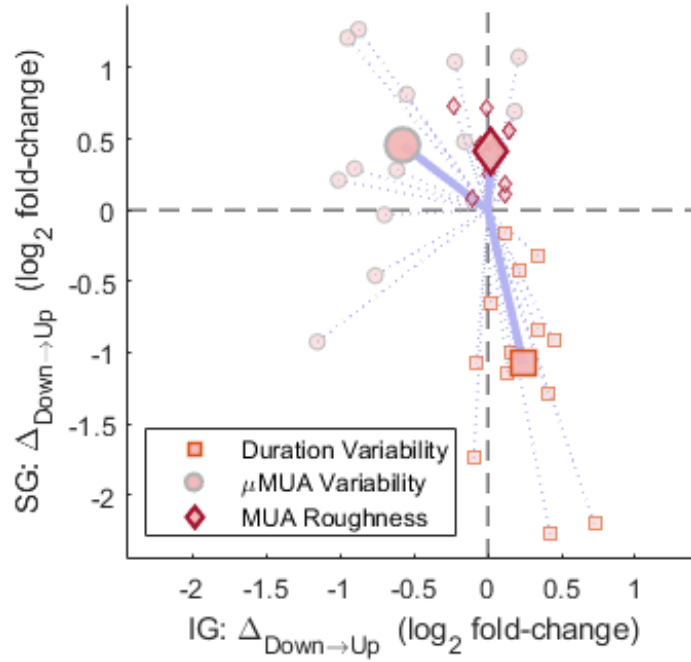


Figure 4.25: Characterisation of the synchronous activity according to its laminar location, by means of the Down-to-Up change in substate variability. For each measure of variability (in duration, in mean MUA, and MUA roughness; see Figure 4.24 above), the intra-channel increase from sync-Down to sync-Up are expressed as \log_2 fold-changes (eg, 0: same variability; 1: sync-Up variability is 2 times that of sync-Down; -1: sync-Up variability is half time that of sync-Down). Faint small dots represent the two laminar components (IG in abscissa, SG in ordinate) of balanced averages by slice ($n = 13$); big dots are the overall means (see methods for details).

Methodological Notes 4.3.4.

- (i) *Intra-channel variability in duration calculated as the coefficient of quartile variation of the durations of each sync-state type occurring within each channel. Analogously for the intra-channel variability of the mean MUA of each state occurrence.*
- (ii) *As for the MUA roughness, we measured it as the median per channel of the cv of the MUA of each state occurrence during the recording period of a channel.*

- (iii) As always, layer (resp. slice) averages are spatially balanced, meaning that each node (resp. column) is given the same weight, irrespectively of the number of available channels to average for that node (resp. column).
- (iv) For each variability feature, the intra-channel difference of Down to Up is expressed in $\log_2 FC$, that is, the \log_2 of the ratio between the values of that feature measured for Up and Down (in that order).
- (v) **Comparisons of features between states** were carried out as follows: for each channel we computed the \log_2 ratio between the value measured for the Up states with those for the Downs. This provides, for each channel, a Down-to-Up increase of the feature of interest expressed in $\log_2 FC$. These values are then averaged in a spatially balanced way: that is, first per node (up to three channels per node, allowing for missing values resulting from very poorly synchronous channels), then per layer (up to 3 nodes per layer). We end up thus with one summary value per slice ($n = 13$) for each layer. Then, a Wilcoxon signed-rank test is independently carried out for each layer against the null hypothesis that the Down-to-Up $\log_2 FC$ is 0.
- (vi) **Comparisons of features between layers** were carried out as follows: for each channel and each state (sync-Up or sync-Down), the feature of interest (cqv of duration, cqv of mean MUA, or median cv of MUA) was measured. Then, these values were averaged per node. Next, for each state and each pair of layer-opposite nodes belonging to the same column, we computed the \log_2 ratio between the SG and the IG nodes. This provides, for each column, the inter-laminar increase of the feature of interest for each state. These columnar replicates were then averaged per slice, resulting in one summary value per slice ($n = 13$) for each state. Finally, a Wilcoxon signed-rank test was independently carried out, for each state, against the null hypothesis that the inter-laminar $\log_2 FC$ is 0.
- (vii) The equivalence between the reported Mean \pm SEM ($m \pm s$) expressed in $\log_2 FC$ and the interval percentage of increase/decrease [p_1 ; p_2] is computed as follows:

$$p_1 = (2^{m-s} - 1) \cdot 100 \quad \text{and} \quad p_2 = (2^{m+s} - 1) \cdot 100$$

Positive $\log_2 FC$ values give positive percentage increases, whereas negative $\log_2 FC$ values give negative percentage increases (ie, decreases).

4.3.5 Profile asymmetry of sync-Up states

Next, we studied the dynamics of sync-Up's onset and termination events during each oscillation cycle. To this end, we will use their phase as a proxy measure of state duration as a fraction of the cycle, and so, readily comparable between columns or slices, without the need of having to take into account differences between the oscillation frequencies of various signals—or even the fluctuations in frequency within the same channel over different synchronous periods.

Sync-Up's onset (termed here UpON) occurs earlier in the ridge cycle in IG than in SG (-0.51π (SD: 0.04π) rad in IG vs -0.39π (SD: 0.09π) rad in SG), the column-wise interlaminar contrast being significant ($p = 1.22 \cdot 10^{-4}$). The opposite occurs for the sync-Up termination (hereafter UpOFF), which comes significantly later for IG relatively to the ridge cycle (0.43π (SD: 0.04π) rad in IG vs 0.34π (SD: 0.05π) rad in SG; ($p = 1.22 \cdot 10^{-4}$). (For both cases, Wilcoxon signed-rank tests, $n = 13$). (See Figure 4.26 for an illustration.) That Up states in IG be broader within the Up-Down alternation cycle than those in SG comes as no surprise, owing to their being also longer in duration (see *supra* section 4.3.3).




Figure 4.26: Representative example of circularly aligned sync-Up/Down state alternations from two MUA signals of the same MEA column (**top**: SG channel, **bottom**: IG channel). For each signal, the alignment is done over *ca* 100 sync-Up states occurring during a 90s-long, simultaneously recorded, period. Note that the states are aligned with respect to their ridge-phases (modulo 2π in the interval $[-\pi, \pi[$, hence circular), not time. Each alternation sample is plotted in a different shade of grey. The crenellated maroon traces indicate the corresponding circular mean of the states' phases over the whole period. For the sync-Up onset and termination events (UpON and UpOFF, respectively), their intra-channel phase-locking value (PLV) was computed. These are illustrated by the clock-like representations above the MUA traces. In short, the phases (angles in radians) of each single sync-Up occurrence event appear in grey (matching the colour of the MUA trace samples); their circular mean is depicted as the thick clock-hand in the unit-radius circle. While the angle of such mean matches that of the corresponding Up event in the crenellated trace, its module (*ie* the length of the clock-hand) is precisely the PLV of the Up event. Note that the highest PLV (*ie* that closest to length 1) is reached for the SG UpOFF event.

However, the present analysis also reveals that the lag in intrinsic phase is longer in the onset than in the termination of the Up. That is, the onset delay in phase between layers is bigger than the termination delay in phase: namely, $0.022\pi \pm 0.010\pi$ rad bigger; this difference being significant ($p = 0.0402$, Wilcoxon signed-rank tests, $n = 13$). This means there is a lack of symmetry between SG

and IG Ups relatively to their occurrence in their own oscillation cycle. To be justly assessed, this result should be interpreted in the context of our method to detect sync-Up sub-states within the synchronous periods. The method rested on finding certain differential properties of the reconstructed logMUA signal via the inverse wavelet transform in the $[0.5, 4]$ Hz frequency band. Filtered logMUA Ups thus tend to peak around 0 rad at each alternation cycle in all cases, driven by the dominant frequency of the power spectrogram’s main ridge in that band—which served to their definition. A lack of symmetry between two states would indicate a different composition in the weights and phases of the frequencies making up the reconstructed oscillation.

Here, the lack of inter-laminar UpON/UpOFF relative symmetry shows in sync-Ups appearing more negatively skewed in SG than in IG: that is steeper peak-to-valley than valley-to-peak oscillation in SG compared to IG. This suggests that the MUA surge in intensity during the sync-Up is relatively more concentrated towards the state’s termination (and hence, more spread at the onset) in SG than IG (see again Figure 4.26 for an example).

Intrigued by the sync-Ups terminating more abruptly in SG layers, we investigated whether they also had a preferred termination phase. To this end, we computed the intra-channel phase-locking value (PLV) of the UpON and UpOFF events. We found that, while onset PLVs are similar between layers—if so, slightly higher (0.52–4.26 %) in SG ($p = 0.3396$)—termination PLV’s are significantly higher in SG (6.13–9.83 % higher, $p = 6.10 \cdot 10^{-4}$). That is, within each channel, SG sync-Ups show a stronger preference to the relative occurrence of their termination with respect to the oscillation cycle, than any of the rest of events (see Figures 4.26 and 4.27). Since the PLV is but a measure of phase dispersion (adapted to circular quantities modulo 2π , [Lachaux *et al.* (1999)]), this result complements the findings on state’s duration variabilities (see section 4.3.4). Namely, for IG sync-Up states, the relatively low PLV in both onset and termination underlies their high duration variability (Figure 4.25 in previous section 4.3.4). By contrast, well-conserved SG sync-Ups may be anchored to their reliable termination, given its high PLV.




Figure 4.27: Characterisation of the sync-Up events according to the increase of their intra-channel PLV from infra-granular (IG) to supra-granular (SG) layers. For each slice, the mean IG-to-SG increase in intra-channel event PLV (expressed in percentage points) is represented by the coordinates of a faint-coloured arrow (UpON event in abscissa, UpOFF in ordinate); the thick darker arrow being their overall mean. Only the increase in ordinate (UpOFF) is significantly bigger than 0.

4.3.6 Propagation of sync-Up states during synchronous periods

In the previous section we focused on the phase-delays and variabilities of Up events relatively to their intra-channel oscillation cycles (as captured by the event's PLVs). However, these analyses did not provide a way of directly comparing distinct channel oscillations. In this section we delve into the analysis of inter-channel phase-synchrony, which will give insights into the sync-Ups' propagation during synchronous states of the AS.

Now, two methodological difficulties arise. First, to compare oscillation naively, these should have the same frequency; or if not, frequencies ought to be at least related by simple rational factors [Rosenblum *et al.* (2001)]. However, this seems not to be granted for the synchronous periods of the AS, which present main oscillatory activities varying from one channel to another in seemingly incongruous beat-frequencies (see section 4.3.1.2). Second, and crucial: signals must be simultaneously oscillating. This constraint appears to be here even more problematic, since the AS fluctuates from synchronous to asynchronous periods across all channels of the slice, with different co-occurring time-windows (see section 4.2.2). Thus, assessing sync-Up dynamics embedded in a transitional state—as here the AS—can be challenging for existing approaches on SO propagation (see *eg* [D'Andola *et al.* (2017), Pazienti *et al.* (2022)]; see also [Kharas *et al.* (2022)] for an example on non-spontaneous oscillatory activity different from SO, but comparable to a certain extent). Fortunately, in our case synchronous periods tend to co-occur in channels of the same MEA column (section 4.2). Furthermore, the intensity and frequency of the oscillation proved well conserved at the column level too (see section 4.3.1). Thus, the two previous caveats can be overcome if we constrain our study to the intra-columnar propagation of the oscillation.

4.3.6.1 Phase-coupling strength between oscillations

We then conducted a whole-slice-spanning analysis that circumvented the aforementioned challenges. For each pair of channels, we computed the PLV of the difference of the ridge-phases over the intersection of their synchronous periods. Built on the classical PLV measure [Lachaux *et al.* (1999)], we call our specific variant the inter-ridge PLV (see Methodological Notes 4.3.5 below and Figure 4.28a for an example in one slice). This simple method gives an overview of the phase-coupling strength between each pair of channels, even when the possibility exists of having different oscillation frequencies and non-co-occurring synchronous periods. In fact, on the one hand, a circular mean of the phase differences from a pair of channels with different main oscillation frequencies over a long enough period will be close to zero. This will happen because the instantaneous phase differences will end up spreading across the whole angle range, hence returning a negligible PLV. On the other hand, pairs of channels with a poor and fragmented overlap of synchronous periods will likely present a highly variable collection of phase differences between them, resulting again in a low PLV. Thus, only pairs of signals with sustained co-occurrent synchronous periods and a very close oscillation frequency are susceptible to return high PLVs, indicating that their difference in phase is reasonably conserved along time. (In such case, the value of the mean phase lag will also be meaningful.)

In agreement with sections 4.2 and 4.3.1.2, the highest inter-ridge PLVs corresponded to pairs of channels located in the same MEA node or column, and tended

to be typically negligible for channels located in different columns (see Figures 4.28). Namely, the mean inter-ridge PLV between channels belonging to the same MEA node was significantly higher than for pairs of channels belonging to different nodes (0.821 ± 0.039 , against 0.462 ± 0.065 ; $p = 1.22 \cdot 10^{-4}$). Likewise, pairs of channels located in the same column but in different nodes (to offset the strong node effect), showed a significantly higher PLV than pairs of channels in distinct columns (0.772 ± 0.058 , against 0.381 ± 0.072 ; $p = 1.22 \cdot 10^{-4}$). As for the laminar effect, by contrast, the mean PLV between channels in the same MEA layer (but again, in different nodes) was low and very similar to the mean inter-ridge PLV found between channels located in distinct layers and columns (to avoid the significant columnar effect): 0.388 ± 0.073 vs 0.374 ± 0.072 , $p = 0.2072$ (for all three previous contrasts, paired Wilcoxon sign-rank tests, $n = 13$). See Figure 4.28.




Figure 4.28: **Left:** representative example of the collection of inter-ridge phase-locking values (PLV) computed for every pair of channels of the same slice over the whole recording period. We have defined the inter-ridge PLV between two channels as the module of the circular mean of the instantaneous differences between the phases of the power ridges of these two channels' WPSg, that these channels describe over the mutual intersection of their synchronous periods (see Methodological Notes 4.3.5; for more details on the ridge and its phase, see also Methodological Notes 4.3.2 of section 4.3.2). The PLV thus defined offers then a measure of how coupled two single oscillations are in average over co-occurrent synchronous periods. 0 means totally uncoupled and 1, totally coupled. Note that, by definition it is a symmetrical function, hence the symmetry of the matrix shown. Thick green lines help to demarcate channels belonging to the same MEA column; dotted green lines, those belonging to the same MEA node. **Right:** influence of the cortical structural organisation on the coupling-strength between the logMUA's oscillatory activity, as captured by the inter-ridge PLV between every pair of channels. For each MEA layout factor (node, column and layer) we compared the mean inter-ridge PLV averaged over two levels: intra *vs* inter (*ie*, belonging or not to the given layout factor). Each dot represents, for a given slice, the mean of the inter-ridge PLV between channels belonging (intra) or not (inter) to a given location factor (node, column or layer). (+) Because of the strong effect exerted by the nodes, intra-columnar means do not comprise intra-node PLVs. (#) The same for intra-laminar means; in addition, because of the strong effect exerted by the columns, inter-laminar means do not comprise intra-column differences either. (***) $p < 0.001$, (ns) not significant. Dot colour codes for slice ($n = 13$).

This finding is consistent with our previous analyses on the global spectra (see section 4.3.1) and the spatio-temporal co-occurrence of synchronous states (see sections 4.2.1 and 4.2.2), where we also found a dominant columnar coordination of the synchronous activity.

4.3.6.2 Inter-laminar propagation lag between sync-Up events

Once the phase-coupling strength between synchronous activities established all across the slice, we can safely evaluate the phase-lags between these oscillations. And provided we restrict our attention to channels located in the same MEA column, the circular mean of their instantaneous phase differences over co-occurrent synchronous periods will give a measure of the mean phase-lag between them. We found that the intra-columnar signed phase lag between the IG and SG oscillation was $+0.375\pi \pm 0.047\pi$ rad ($p = 1.22 \cdot 10^{-4}$). That was significantly higher than the mean absolute phase lag between intra-node oscillations: 335.23–556.28 % higher ($p = 1.22 \cdot 10^{-4}$). (For both contrasts, Wilcoxon sign-rank tests, $n = 13$; figures not shown.) That is, the oscillation generated by the sync-Up/Down alternation during synchronous periods appears to propagate from IG to SG channels within the MEA columns.

However, given the inter-laminar asymmetry of the Up sub-states' occurrence within each oscillation cycle (see section 4.3.5), the previous analysis does not yield information about the precise delays between sync-Up events of different layers. To address this, we designed a measure to quantify phase lags between reasonably phase-coupled channels (see Methodological Notes 4.3.5 below). In short, for a given pair of channels, we projected the instants of one channel's Up events (onset and termination) onto the phase of the other channel (that we called the reference channel). Note that this analysis only takes into account periods of co-occurrent synchrony between the two channels. Then, the circular mean of this collection of projected phases was compared to the intra-channel circular mean of those reference channel's event phases taking place during the same intervals of shared synchrony. The whole operation returns the difference in mean phase delays of the sync-Up's onset and termination between these two channels (see Figure 4.29 left for an example on UpON, and 4.29 right for UpOFF). We found that sync-Up onsets (UpON) in SG channels tend to lag those in IG channels of the same column by $0.812\pi \pm 0.048\pi$ rad with respect to the oscillation of the IG channels ($p = 1.22 \cdot 10^{-4}$). Strikingly, this advantage was almost lost for the sync-Up terminations (UpOFF), which were only slightly ahead for IG channels ($0.106\pi \pm 0.063\pi$ rad), and no significantly so ($p = 0.2439$). Furthermore, this difference between UpON's and UpOFF's inter-laminar phase lags was stable across columns: UpOFF's mean lag per column being 73.46–84.13 % smaller than UpON's ($p = 2.44 \cdot 10^{-4}$) (For the four preceding tests: Wilcoxon sign-rank tests, $n = 13$.)

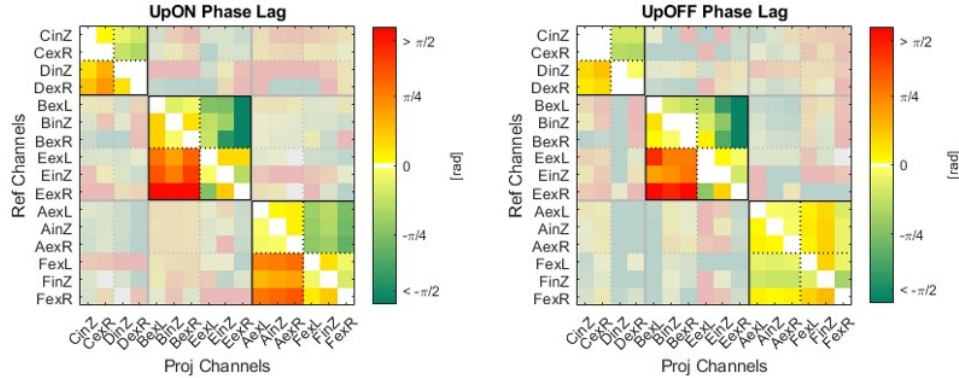


Figure 4.29: **Left:** phase lags of the sync-Up onset events (UpON) between channels of a representative slice (same as depicted in panel c). Only intra-columnar interactions between pairs of channels are taken into account, the rest appearing masked because of their lack of reliability owing to low inter-ridge PLVs (see Figure 4.28). Given a pair of channels, UpON’s phase lag of one channel (projected channel) with respect to the other (reference channel) is computed as the circular mean of the phases of the projected channel’s UpON events relative to the reference channel ridge-phase. Thus, lags between channels in the same MEA node hover at 0 rad (yellowish hues); lags from IG channels to SG ones are negative (green), meaning that UpON events occur earlier in IG than in SG; the opposite occurs from SG to IG (red). Note that the average of phases is only taken for UpON events happening at periods of co-occurring synchrony between the pair of channels. Because of that, the matrix of lags is not perfectly antisymmetric. Only IG-to-SG lags (*ie*, those taking IG as reference) were used for subsequent analyses. **Right:** same as left panel for sync-Up termination events (UpOFF).

Next, we were interested in converting these phase lags into time terms. This we did by taking into account the mean period of the corresponding columnar oscillation (*ie*, the inverse of the oscillation frequency). We found that sync-Up started in average 79.2 ± 2.8 ms earlier in IG channels than in the SG ones located in the same column ($p = 1.22 \cdot 10^{-4}$). However, the sub-states ended almost at the same time—if so, negligibly earlier in IG (13.1 ± 6.1 ms), and not significantly different from a 0 ms lag ($p = 0.3054$). Again, the difference in inter-laminar lag size between onset and offset was substantial: UpOFF’s lag being in average 73.34–84.11 % smaller than UpON’s ($p = 2.44 \cdot 10^{-4}$). (See Figure 4.30.) Finally, taking into account the separation between laminar-opposite electrodes in the MEA, this gives an estimate of $8.47\text{--}9.09$ mm·s $^{-1}$ for the speed of the sync-Up onset propagation, in line with those found in comparable previous studies [Sanchez-Vives *et al.* (2010), Capone *et al.* (2019)].

In summary, this analysis supports that sync-Up onset propagates within cortical columns from deeper to shallower cortical layers, consistent with Ups initiating in deeper layers layers (V or VI) during SO regimes both *in vitro* and *in vivo* (*eg*, [Sanchez-Vives and McCormick (2000), Wester and Contreras (2012), Capone *et al.* (2019), Pazienti *et al.* (2022), Mattia *et al.* (2021)]). In addition, the intra-columnar propagation of stimulus from deeper to shallower layers in a rest-to-awake transition has also been documented *in vivo* [Kharas *et al.* (2022)]. Interestingly, we also found that sync-Up seem to terminate abruptly and simultaneously in the column, thus supporting a deterministic phase of the Down states,

as argued in [Camassa *et al.* (2022)]. However, no study provides a quantification which carefully differentiates between onset and termination of the Up state for this elementary intra-columnar IG-to-SG propagation mode.




Figure 4.30: All-slice collection of sync-Up event's inter-laminar time lags depicted in context with SG sync-Ups duration. Each dotted-line-connected pair of markers represents the average inter-laminar time lag between the channels within one MEA column; hence, two pairs of 3 lags per slice (slices with some column missing indicate that all the channels of at least one node of that column had not a synchronous content above 3%, as per our admissibility criterion; see methodological notes for section 4.3.3). For convenience, relative timings in the figure appeared aligned with respect to SG UpON events (time 0 ms). For each slice and column, the time shown between SG UpON and UpOFF events corresponds to the mean duration of SG sync-Ups in that column. Grey shades give the overall mean size of the inter-laminar lags at the onset (79.2 ms) and the termination (13.1 ms). The two shades are separated by the overall, very well-conserved, median duration of SG sync-Ups (182 ms). (**UpON**: onset ; **UpOFF**: termination; **IG**: infra-granular; **SG**: supra-granular).

4.3.6.3 Influence of the inter-laminar Up-onset lag on the oscillation frequency

Let us note that the two previous analyses (*ie*, lag in phase and time) are not *a priori* redundant. In fact, both results might very well not have agreed, had the phase lag depended negatively on the frequency of the columnar oscillation. That is why we also controlled for the dependence between phase lag and oscillation period. Strikingly, the sync-Up onset phase lag correlated positively with the mean period of the oscillation ($\theta_{\text{UpON}} = 2.098 \cdot T_{\text{Osc}} - 0.399$, $p < 10^{-3}$, $R^2 = 0.2912$; figure not shown). *A fortiori*, this relation is reinforced in the time domain: since, if the longer the period, the longer the UpON phase lags, then, with all the more reason, the longer the UpON lag also in time ($p < 4 \cdot 10^{-9}$, $R = 0.6569$). Interestingly,

nothing of the sort happened for inter-laminar termination events. UpOFF lags are conspicuously uncorrelated with the oscillation's mean period ($p = 0.15$, $R^2 = 0.0616$). (See Figure 4.31.) That gives more ground to the hypothesis that sync-Up terminates simultaneously within the same column. Furthermore, owing to that very notable lack of dependence between UpOFF lag and the columnar oscillation frequency, we propose that the relation between UpON lags and oscillation can be seen the other way around: it is not that the oscillation frequency influences the inter-laminar onset delay, but rather that inter-laminar Up onset delay might mostly be responsible for the oscillation rhythm, and the latter according to an approximate quadratic relation: $T_{\text{Osc}} \approx 3.56 \cdot \theta\text{UpON}^2 + 1.00 \cdot \theta\text{UpON} + 0.46$ ($R^2 = 0.6591$). (See Figure 4.31 left.) Such claim crucially rests on our finding that SG sync-Ups have a very well conserved duration, no matter their column and hence the frequency of the sub-state alternation (see previous section 4.3.3).




Figure 4.31: **Left:** Relation between the inter-laminar time lag of the sync-Up onset (UpON) and the main oscillation rhythm. Each dot maps the mean infra-to-supra-granular time lag of a MEA column (abscissa) with the mean period of that column's oscillation (ordinate). Groups of connected dots form triangles representing the slices to which the corresponding columns belong. Oscillation mean periods strongly and positively correlated with time lags ($p < 4 \cdot 10^{-9}$). A quadratic polynomial was linearly fitted over the centres of gravity of the triangles (*ie*, the slice means), which fared slightly better than the linear model counterpart in terms of coefficient of determination.

Right: Same as left panel for the sync-Up termination events. In this case, time lags hover at 0 and no correlation was found between them and the oscillation period.

Methodological Notes 4.3.5 (relative to sections 4.3.5 and 4.3.6).

- (i) **Notation:** Given $Z = (z_k)_{k \in K}$ a sequence of complex numbers indexed by the set K of cardinality $|K| < \infty$, it will be convenient to denote its arithmetic mean as

$$\langle Z \rangle = \langle z_k : k \in K \rangle := \frac{1}{|K|} \sum_{k \in K} z_k$$

- (ii) **Circular mean:** Given any finite sequence of angles $\Theta = (\theta_k)_{k \in K}$ in radians, by its circular mean we refer to the angle of the mean vector resulting from averaging the unit-norm vectors with angle θ_k . Formally, with complex arithmetic we can simply denote the circular mean of the θ_k by

$$\langle \Theta \rangle_{\text{circ}} = \langle \theta_k : k \in K \rangle_{\text{circ}} := \arg \left(\langle e^{i\theta_k} : k \in K \rangle \right)$$

We report the circular mean in $[-\pi, \pi[$ modulo 2π .

Let us also denote the modulus associated to this circular mean as

$$|\Theta|_{\text{circ}} = |\theta_k : k \in K|_{\text{circ}} := \left| \langle e^{i\theta_k} : k \in K \rangle \right| \in [0, 1]$$

- (iii) **Intra-Channel Event PLV:** Let $\nu(t)$ be a logMUA signal and $\phi(t)$ the phase function of its oscillation ridge in the Ξ_0 band (see supra methodological notes 4.3.2 on p. 106). Note that while ν is defined over the whole recording interval I , ϕ is only defined in I_{sync} , the union of the synchronous periods of ν . If $t_k \in I_{\text{sync}}$, $k \in K$, are the instants at which an event of interest occurs (UpON, say), we call the intra-channel event phase-locking value (PLV) of UpON, the modulus associated to the circular mean of the phases of UpON in ν . That is:

$$PLV_{\text{intra}}(\text{UpON}) := |\phi(\theta_k) : k \in K|_{\text{circ}} = \left| \langle e^{i\phi(\theta_k)} : k \in K \rangle \right|$$

Thus defined, the PLV is a real number between 0 and 1. The closer the PLV to 1, the closer the averaged phases are modulo 2π . On the contrary, the smaller the PLV, the higher the variability of the event phases.

- (iv) **Inter-ridge Phase Lag and PLV:** For two channels ν_1 and ν_2 , with associated oscillation's ridge phase functions ϕ_1 and ϕ_2 , defined over I_{sync}^1 and I_{sync}^2 , respectively, we define the inter-ridge phase lag from ν_1 to ν_2 as the circular mean of the phase differences of the oscillations over the co-occurrent synchronous periods. That is, calling $Z = (\phi_2(t) - \phi_1(t))_{t \in I_{\text{sync}}^1 \cap I_{\text{sync}}^2}$ the sequence of phase differences, the inter-ridge phase lag will be:

$$\theta_{\text{lag}}(\nu_1, \nu_2) := \langle Z \rangle_{\text{circ}}$$

(Notice that, since we are dealing with sampled signals, both I_{sync}^1 and I_{sync}^2 are finite, so is Z).

We can also consider the inter-ridge PLV simply as

$$PLV_{\text{inter}}(\nu_1, \nu_2) := |Z|_{\text{circ}}$$

While the latter is symmetric: $PLV_{\text{inter}}(\nu_1, \nu_2) = PLV_{\text{inter}}(\nu_2, \nu_1)$ the former is anti-symmetric: $\theta_{\text{lag}}(\nu_1, \nu_2) = -\theta_{\text{lag}}(\nu_2, \nu_1)$

- (v) **Sign of the inter-ridge phase lag:** The sign of the inter-ridge phase lag is informative for pairs of oscillations with strong phase-coupling, that is, with PLV_{inter} close to 1, indicating that the instantaneous differences in phase have low variability over time. In order to give a summary average of the phase lags computed between pairs of channels belonging to certain groups (eg, the mean phase-lag from IG to SG channels located in the same column), some precautions must be taken. a) Since the phase lags are angular quantities modulo 2π , their mean should be circular. b) For the sign of a circular mean to be meaningfully reported, the mean must have been taken over phase lags of the same sign, and the sample must be significantly different from zero. That is why we first assessed the quality of the lag signs. Indeed, we found that 12 out of 13 slices had all their admissible columns showing positive IG-to-SG phase lags.
- (vi) **Inter-Channel Event Phase Lag:** For an event of interest (UpON, say) and a pair of channels ν_1 and ν_2 , we can consider the collection of instants at which the event occurs. Call them T_{UpON}^1 and T_{UpON}^2 , respectively. Notice that $T_{UpON}^j \subset I_{sync}^j$, $j = 1, 2$. To measure the mean phase lag between the occurrences of the event UpON in two channels, we have defined the inter-channel event phase-lag of UpON from ν_1 to ν_2 as:

$$\theta_{\nu_1 \rightarrow \nu_2}(UpON) := \langle \phi_1(t) : t \in T_{UpON}^2 \cap I_{sync}^1 \rangle_{circ} - \langle \phi_1(t) : t \in T_{UpON}^1 \cap I_{sync}^2 \rangle_{circ}$$

where ϕ_1 is the oscillation's ridge phase function of ν_1 . We respectively call ν_1 and ν_2 , the reference and the projected channel. The phase lags $\theta_{\nu_1 \rightarrow \nu_2}$ are reported modulo 2π in $[-\pi, \pi[$ rad.

The antisymmetry is here no longer enjoyed. Mainly because

$$\phi_1(T^2 \cap I_{sync}^1) \neq -\phi_2(T^1 \cap I_{sync}^2)$$

However, for pairs of channels with a high overlap of their synchronous periods (ie, $I_{sync}^1 \approx I_{sync}^2$), and high inter-ridge PLV (ie, strong phase-coupling), we usually do have $\theta_{\nu_1 \rightarrow \nu_2} \approx -\theta_{\nu_2 \rightarrow \nu_1}$. This would typically be the case for channels located in the same MEA column.

- (vii) **Inter-Channel Event Time Lag:** In order to interpret the inter-channel event phase lag in time terms, we need to convert the phase lags into time lags. For each pair of channels (ν_1, ν_2) we simply do this by considering the period of the oscillation of the reference channel ν_1 , which we approximate as the mean duration of the Up-Down alternation cycle during the co-occurrent synchronous periods: call it $\langle \mathcal{T}_{Osc}^1 \rangle_{I_{sync}^1 \cap I_{sync}^2}$. The inter-channel event time lag is thus computed as

$$\tau_{\nu_1 \rightarrow \nu_2} := \frac{\langle \mathcal{T}_{Osc}^1 \rangle_{I_{sync}^1 \cap I_{sync}^2}}{2\pi} \theta_{\nu_1 \rightarrow \nu_2}$$

- (viii) **Interpretation:** Thus defined, our inter-channel event time lag from ν_1 to ν_2 provide a mean approximation to $\langle t_k^1 - t_k^2 \rangle$, where (t_k^1, t_k^2) would denote the timings of the "aligned" k^{th} event observed at channels ν_j , $j = 1, 2$. The main advantage of our method is that it does not require any preliminary

alignment of the events, nor makes any assumption on the propagation nature of the events: events are phenomena occurring a priori in a channel, beyond the need to assume beforehand that close occurrences of an event in different channels may, or may not, correspond to the very “same” event. In fact, for an arbitrary pair of channels, determining a collection of (t_k^1, t_k^2) event-aligned instants implies considering almost each event co-occurrence on a case-by-case basis, thereby introducing a number of biases and resulting in a poorly robust analysis tool. On the contrary, our method, although approximate, is objective, extremely easy to implement, and does not require additional parameters.

- (ix) **Averages for statistical summaries and contrasts:** Throughout this work, the slice recording experiment is considered the elemental statistically independent replicate, of which we have $n = 13$. For each measure we want to assess statistically, different strategies have been used to come up with a summary per slice (see Methodological Notes 4.2.1, 4.3.1, 4.3.3 and 4.3.4). In all these sections, as well as in the present one, most of the measures involve inter-laminar comparisons within the MEA column. Thus, each column returns a summary value which is in turn averaged with the other two columns’ values. However, the particularity for some of the measures of this section is that, for each column, inter-laminar comparisons are made of channel-to-channel comparisons (not node-to-node, as it was typically the case in the aforementioned sections). For example, in order to give an intra-columnar average of the IG-to-SG event time-lag, we computed the mean of every available inter-channel event time lag $\tau_{\nu_1 \rightarrow \nu_2}$, where the reference channel is in IG and the projected channel in SG, and in the same column. Hence, at most 32 observations for columns with 3-channel nodes and 22 for columns with 2-channel nodes. (We recall that it may happen for some channels not to be deemed synchronous-rich enough to take part in these analyses; see Methodological Notes 4.3.3 on p. 109.)
- (x) For the **linear regressions**, all the available values per column were pooled, but weighted so that each slice is given the same importance. Both the linear and quadratic models were fitted through ordinary least square estimates.

4.4 Laminar dependence of asynchronous MUA constrained by the synchronous sub-states

4.4.1 Inter-laminar maximisation of mean MUA during asynchrony

Key message: mean MUA is higher in IG than in SG for every state; the asynchronous state maximises this inter-laminar difference.

In contrast to what happens during the synchronous periods described above, during the asynchronous periods of the awaking-like state (henceforth here called Async states) there is little cross-spectral information between channels to be drawn. In fact, our detection method for asynchronous periods provided a working definition for asynchrony: a time window where the signal is statistically undistinguishable from its underlying background noise, which is a maximally uncorrelated process following the inverse power law (see Methodological Developments, sections 3.1.1.2 and 3.2.1). Therefore, owing to very poor cross-correlations between co-occurring asynchronous periods, the inter-channel coordination between Async states cannot be investigated in the same way as done for the synchronous periods. Instead, we resorted to a “synchronicity-agnostic” clustering approach of the MUA probability densities, in order to characterise the Async states spatially and in relation to the surrounding sync-Up and sync-Down states. The firing rate of neural ensembles is highly dependent on the cortical layer [Einevoll *et al.* (2007), Senzai *et al.* (2019)], thus the intensity and the probability density of the firing rate is expected to be similar between electrodes located in closer layers.




Figure 4.32: Representative example for one slice of the inter-channel similarities between MUA densities corresponding to three different sub-states: sync-Down, sync-Up and Asynchronous. Each sub-state is represented by a graph with vertices matching the spatial location of the channels in the MEA (see Methods, section 2.2). Similarity between MUA densities are given by the KS-distance (0: most similar; 1: least similar; see Appendix A.3), and are represented by the edges of the graph (more similar: thick red lines; less similar: thin blue lines).

In our case, a sample-independent statistical distance (the Kolmogorov-Smirnov distance, see Methods) allowed us to compare MUA densities from different locations and states. For each of the 3 sub-states of the AS (Async, sync-Down and sync-Up), we found that MUA densities tend to cluster across layers, irrespective of the column and hence of their synchronicity levels and oscillatory coordination (Figure 4.32, via the Kolmogorov-Smirnov statistical distance, see Appendix A.3).

Additionally, the mean MUA (μMUA) was always higher in IG than SG ($145.18 \pm 17.32\%$ higher in Async; $86.55 \pm 15.47\%$ in sync-Down; and $94.98 \pm 20.58\%$ in sync-Up). Interestingly, the Async state significantly maximises the inter-laminar difference in mean firing rate ($p = 0.0458$ for the contrast against sync-Down; $p = 0.0152$ against sync-Up; the difference of the inter-laminar increase between sync-Up and sync-Down being negligible; paired t-tests, where non-departure from normality was controlled with Lilliefors tests, $n = 5$ slices). (See Figure 4.33). Thus, this result suggests that the asynchronous state enhances the functional differentiation between neuronal ensembles located in different layers compared to the lower-complexity synchronous state [Renart *et al.* (2010), Hansen *et al.* (2012), Kharas *et al.* (2022)].




Figure 4.33: Multi-slice quantification of the inter-laminar increase of the mean MUA (μMUA) from the three sub-states of the AS. Each dot represents the slice average of the SG-to-IG increase in MUA ($n = 5$ slices). SG-to-IG average increases are computed between laminar-opposite nodes of the same MEA column.

4.4.2 Conservation of mean MUA across states

Key message: mean MUA is conserved across synchronous and asynchronous periods; important implications on fundamental properties of the network, and the relation between sync-Downs and sync-Ups in relation to Async ensue.

Next, we compared the mean firing rates of each state within each channel. Interestingly, we found that the mean firing rate of the Async states ($\mu\text{MUA}_{\text{Async}}$) matches that of the synchronous periods ($\mu\text{MUA}_{\text{Sync}}$, sync-Ups and sync-Downs comprised) strikingly well. That is, for every channel, $\mu\text{MUA}_{\text{Async}} \approx \mu\text{MUA}_{\text{Sync}}$, as shown by the linear regressions without intercept fitted for each slice; specif-

ically: $\mu\text{MUA}_{\text{Async}} = (1.0265 \pm 0.0273) \cdot \mu\text{MUA}_{\text{Sync}}$ ($R^2 = 0.9348 \pm 0.0281$, and all p-values under 10^{-9}) (Figure 4.34). This result was somewhat surprising, as the synchronous periods, which are composed of alternating sync-Up and sync-Down states (with higher firing rates at sync-Ups and lower at sync-Downs), are qualitatively very different to the asynchronous firing mode (homogeneous firing without dominant oscillation). Thus, the average neuronal activity is maintained constant between the two states, indicating that the shift between synchronous and asynchronous states does not entail a change in mean firing rate, but only a mere rearrangement of the temporal coordination of the intensities. Since the net firing output of the network (as captured by the MUA) is well maintained when transitioning between the two states, this would suggest that the structural E/I balance remains constant during the AS. This is consistent with recent results *in vivo*, where different levels of anaesthesia having an influence on the complexity state of the network but not on the overall balance between excitation and inhibition [Pazienti *et al.* (2022)].




Figure 4.34: Linear correlations between the μMUA of synchronous and asynchronous periods. (Synchronous periods comprise both sync-Up and sync-Down sub-states.) Each dot represent the mean MUA of a given channel: filled dots for IG, hollow ones for SG. Dots of the same colour belong to the same slice. Regression lines for each slice given as thin dotted lines. All the dots are close to the thick dashed identity line. Regressions are spatially weighted, so as to give the same importance to each node of the MEA within the slice.

It also follows from this result that, within a channel, the Async firing rate is an average of the sync-Down and sync-Up firing rates weighted by the relative durations of the sync-Ups and sync-Downs. Thus, for SG channels, whose sync-Up duration is highly conserved (as previously shown in section 4.3.3), this also implies that $\mu\text{MUA}_{\text{Async}}$ is strongly correlated with the oscillation frequency of the synchronous periods within each channel. In short, the faster the oscillation, the higher the mean firing rate of the asynchronous states (assuming sync-Up states of similar mean MUA).

4.4.3 Bounds of the asynchronous MUA in terms of the synchronous sub-states

Key message: in terms of their MUA, Async states are closer to sync-Ups in IG, but closer to sync-Downs in SG. Restricted to the column, an interesting inequality ensues that combines all six previous quantities.

Since the firing rate of the synchronous periods is, for each channel, a weighted average of the mean firing rates of their sync-Downs and sync-Ups, from the conservation of mean MUA across states it comes as no surprise that the Async mean MUA must lie between the mean MUA of sync-Down ($\mu\text{MUA}_{\text{syncD}}$) and the mean MUA of sync-Up ($\mu\text{MUA}_{\text{syncU}}$). Indeed, for every channel,

$$\mu\text{MUA}_{\text{syncD}} \leq \mu\text{MUA}_{\text{Async}} \leq \mu\text{MUA}_{\text{syncU}} \quad (4.1)$$

($p = 0.0023$ for the first inequality, $p = 0.0006$ for the second; paired t-tests, with Lilliefors tests to control for departure from normality, $n = 5$ slices).




Figure 4.35: Representative example traces of the MUA signal of laminar-opposite channels from the same column. Async periods in green; synchronous ones in grey. Crenellated maroon traces superimposed to the grey synchronous periods indicate the overall mean MUA values of the sync-Up and sync-Down states. Dashed green line indicates the overall mean MUA of the asynchronous states.

Mean MUA inequalities. Next, we asked whether there is a laminar effect upon them. We found that in IG channels, the μMUA of Async is closer to their sync-Up than they are in SG channels (in IG, $\mu\text{MUA}_{\text{Async}}$ is $72.40 \pm 12.04\%$ higher than $\mu\text{MUA}_{\text{syncD}}$; in SG, only $38.36 \pm 14.57\%$ so; and this difference is significant, $p = 0.0459$, $n = 5$). The opposite occurs for the relation between the μMUA of Async and sync-Up (in IG $\mu\text{MUA}_{\text{Async}}$ is $77.57 \pm 3.30\%$ as high as $\mu\text{MUA}_{\text{syncU}}$, whereas in SG, only $59.68 \pm 4.87\%$ so; again, this difference is significant, $p = 0.0100$, $n = 5$). This is naturally in accordance with the previous results (see sections 4.3.3 and 4.3.5), where we showed that sync-Ups are relatively longer in IG than in SG. Hence, the mean MUA of sync-Up weights more in IG than in SG for the time-average over synchronous periods (see Figures 4.35 and 4.36). We were also able to draw a more specific relation of the effect the laminar difference inflicts upon the mean MUA of the three states for a given channel. Namely, within a column, Async mean MUA for IG channels is larger than the geometric mean of

sync-Down mean MUA and sync-Up mean MUA; whereas the opposite holds for the SG channels on this column. That is:

$$\frac{\mu\text{MUA}_{\text{Async}}}{(\mu\text{MUA}_{\text{SyncU}} \cdot \mu\text{MUA}_{\text{SyncD}})^{\frac{1}{2}}}[\text{IG}] > \frac{\mu\text{MUA}_{\text{Async}}}{(\mu\text{MUA}_{\text{SyncU}} \cdot \mu\text{MUA}_{\text{SyncD}})^{\frac{1}{2}}}[\text{SG}] \quad (4.2)$$

($p = 0.0184$, paired t-test, with Lilliefors test to control for departure from normality, $n = 5$). (See Figure 4.36 for a summary of all previous inequalities.)

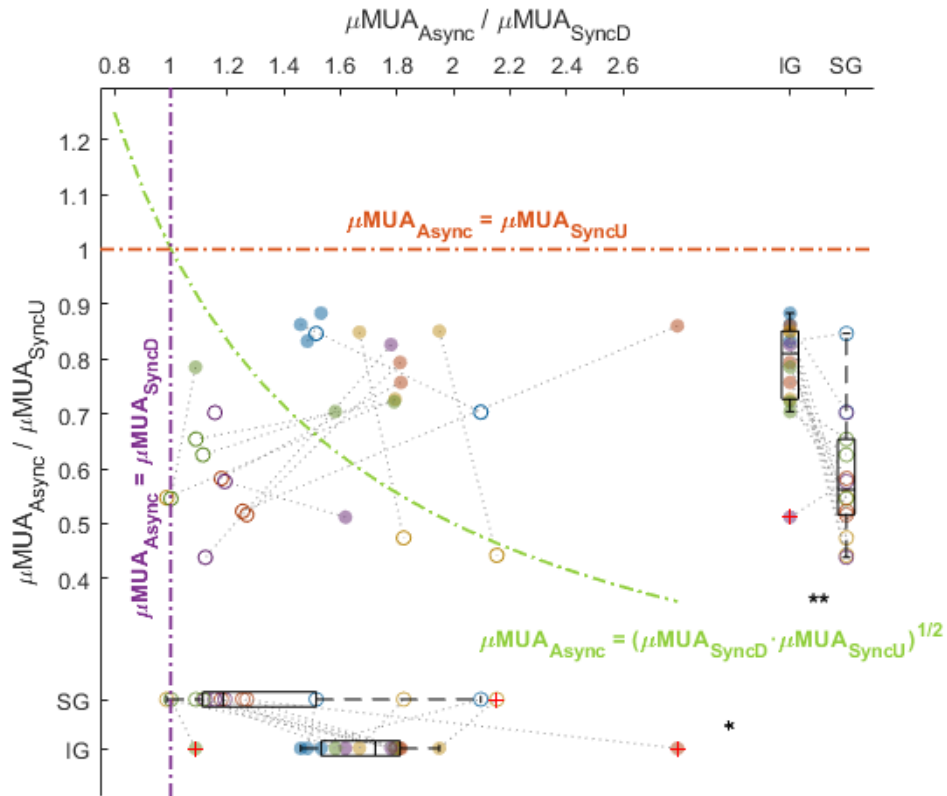


Figure 4.36: Relation between the Async-over-Down μMUA ratio ($\mu\text{MUA}_{\text{Async}}/\mu\text{MUA}_{\text{SyncD}}$, in abscissa) and the Async-over-Up μMUA ratio ($\mu\text{MUA}_{\text{Async}}/\mu\text{MUA}_{\text{SyncU}}$, in ordinate). Each dot represents a node average (dots of the same colour belong to the same slice). Pairs of laminar-opposite dots from the same column are connected by thin dotted lines (filled dots for IG, hollow ones for SG). Thick dashed lines represent the boundaries of some regions of interest as indicated by their equations. For example: all the dots are below the red line $\mu\text{MUA}_{\text{Async}} = \mu\text{MUA}_{\text{SyncU}}$, which means that for every node, the Async mean MUA is lower than the sync-Up mean MUA. Another example: dots located above the green curve $\mu\text{MUA}_{\text{Async}} = (\mu\text{MUA}_{\text{SyncD}} \cdot \mu\text{MUA}_{\text{SyncU}})^{1/2}$ indicate nodes whose Async-over-Up ratio is smaller than their Async-over-Down ratio; that is, nodes whose Async mean MUA is closer to their sync-Up mean MUA than to their sync-Down mean MUA (this is the case for most IG nodes). Box-plots on the fringes show the marginal inter-laminar paired comparisons. (Although all the nodes are depicted, only slice-means were taken into account for the multi-slice quantification.) (**) $p < 0.01$, (*) $p < 0.05$; $n = 5$ slices.

Inequalities in distribution. More generally, the previous result also holds for the whole MUA probability densities, not only for their means. Namely, the Async firing rate is a mixture of sync-Down and sync-Up firing rates in distribution; suggesting that the asynchronous state results rather from a change in the temporal correlation structure of the signal, than from a change in its underlying probability density, which remains almost unperturbed.




Figure 4.37: Representative example of logMUA cumulative distribution functions (cdf) of all AS sub-states for channels of the same MEA column. The cdf is univocally determined by the probability density function (pdf)—and *vice versa*. In blue, sync-Down cdfs; in red, sync-Ups; and in green, Async cdfs. Dotted faint curves represent the cdfs of the synchronous periods (*ie* the weighted average between sync-Up and sync-Down sub-states). BexL and BinZ channels are infra-granular (IG); EexL and EinZ, supra-granular (SG) (see Materials, section 2.2). Note that here for clarity, logMUA densities are depicted instead of MUA ones; nonetheless, the distribution-free property of the KS-distance ensures results described below will be the same irrespective of having considered MUA or logMUA densities (see Appendix A.3 eq. (A.8)).

Furthermore, over the deformation path that goes from the sync-Down to the sync-Up MUA densities, the Async MUA density lies closer to sync-Down in SG, but closer to sync-Up in IG (see Figures 4.37 and 4.38).

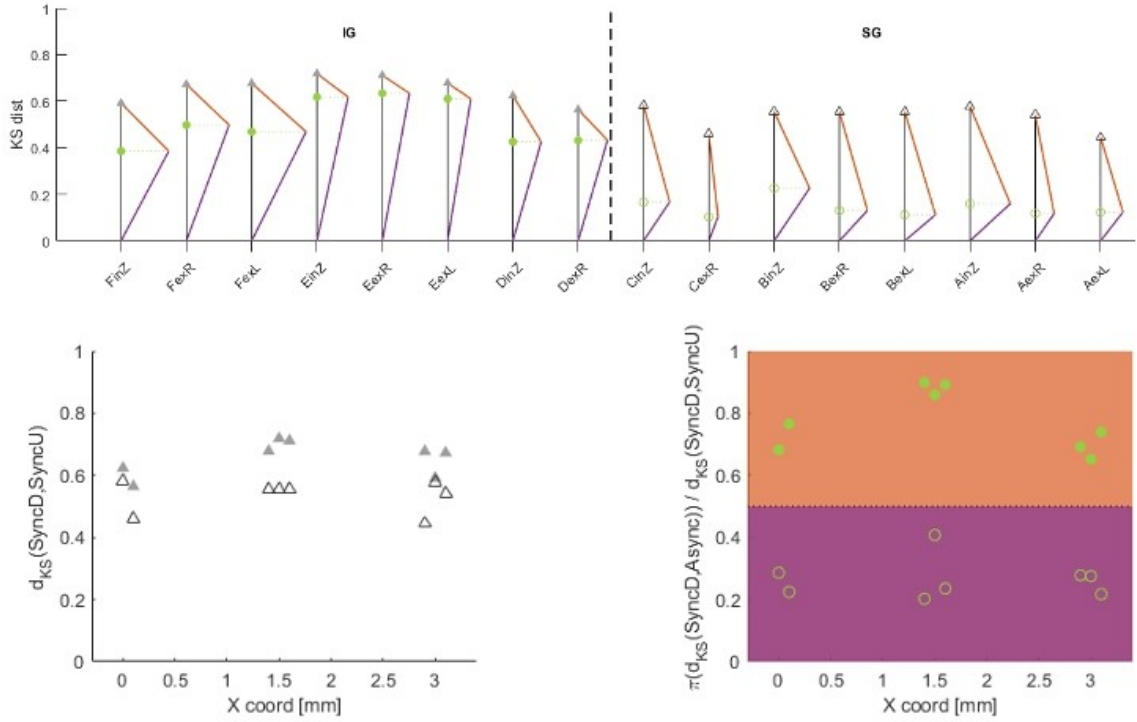


Figure 4.38: Intra-channel similarity between the MUA densities of AS sub-states for a representative slice (same slice as in previous figure). **Top:** Planar projections of AS sub-states' MUA densities through their Kolmogorov-Smirnov distances (d_{KS}). Because of the d_{KS} abiding to the triangular inequality, distances between three densities can be conveniently represented in 2-D (see Appendix A.3). For each triangle, which corresponds to one channel, the grey vertical arrow represents the oriented distance from the sync-Down to the sync-Up MUA densities $d_{KS}(\text{syncD}, \text{syncU})$. The triangle's blue edge is the distance between sync-Down and Async, whereas the red one is the distance between Async and sync-Up. The green dot represents the orthogonal projection of the Async density over the linear path stretching from sync-Down to sync-Up. Filled arrow spears are for IG channels, hollow ones for SG. **Bottom left:** Relation between sync-Down-to-Up distances $d_{KS}(\text{syncD}, \text{syncU})$ and the columnar location of the channels (as given by the MEA horizontal dimension, see *eg* Figure 4.7 on p. 90). Each marker matches one of the arrow spears depicted in panel above (**top**): filled markers for IG, hollow ones for SG. **Bottom right:** Relation between KS projection-ratios of Async over sync-Down-to-Up distances $\Pi_{KS}(\text{Async} | \text{syncD} \rightarrow \text{syncU})$ and the columnar location of the channels. KS projection-ratios are simply the proportion between the orthogonal projection ordinate (green dots of **top** panel) and the sync-Down-to-Up oriented distances (arrow of **top** panel) (See also eq. (A.7) in Appendix A.3.) Thus, a projection-ratio bigger than 0.5 means that the Async density is closer to sync-Up (red region); if smaller than 0.5, Async closer to sync-Down (purple region). Filled dots for IG channels, hollow ones for SG.

Finally, an analogous inequality to that presented above for the mean MUA (4.2) appears empirically verified for the the entire MUA densities, expressed in terms of the orthogonal projections of Kolmogorov-Smirnov distances between sampling distributions (see Methods section A.3 and Figure 4.37 bottom right):

$$\Pi_{\text{KS}}(\text{Async} \mid \text{syncD} \rightarrow \text{syncU})_{[\text{IG}]} > \Pi_{\text{KS}}(\text{Async} \mid \text{syncD} \rightarrow \text{syncU})_{[\text{SG}]} \quad (4.3)$$

($p = 0.0105$, paired t-test, with Lilliefors test to control for departure from normality, $n = 5$ slices).

Because it imposes particular constraints upon the firing rates densities exhibited across various network global states, this result may be highly relevant from a computational modelling stance.

Chapter 5

Discussion

5.1 Summary of findings

In this thesis we developed new analysis tools based on the time-frequency decomposition of multivariate signals to investigate the neuronal network dynamics of the transition from slow wave oscillations (SO), typically occurring during slow wave sleep, towards a more asynchronous activity during wakefulness. We elucidated the network dynamics of the awaking-like state (AS) from electrophysiological recordings from cortical slice preparations which exhibits a transitional regime driving the SO to a state mimicking the purely decorrelated activity found in awake states.

In order to disentangle the complex dynamics arising in the AS, we developed a method that detects periods of synchronous and asynchronous activity (see Methodological Developments, chapter 3 section 3.2). Our method is a three-step statistical procedure carried out on the wavelet power spectrogram (WPSg) of the logMUA signal. Crucially, it draws its statistical validity from our detailed theoretical work on the normalisation of the WPSg (see section 3.1 in Methodological Developments).

In short, we first detected for each channel which periods could be categorised as asynchronous, defined as the time windows in which the spectral content is statistically indistinguishable from the mean background spectrum (the null hypothesis for the *async* detection, see section 3.2.1). The background activity was modelled as a fractional Gaussian noise with spectral slope intrinsically fitted from each channel spectrum (see sections 3.1.1 and 3.1.2).

Second, we investigated whether the remaining non-asynchronous periods mapped meaningful oscillations, that is, that they were genuinely synchronous. We consequently tested the strength of power surges in a fixed band of interest (0.5–4 Hz). To this end, we compared, at each time point, the band-restricted instantaneous power against the distribution of powers that the channel’s intrinsic background power would have carried should it be of the same variance as the broad-band signal (the null hypothesis for *sync* detection).

Finally, we determined optimal significance thresholds by maximising the coverage of the two complementary tests, and in particular, by jointly maximising the entropy of the classification for the whole set of channels in each slice. Our procedure resulted in a high percentage of uniquely classified periods, under very reasonable significance levels for each slice. This strongly suggests that the AS recordings are indeed essentially composed of states of synchronous and asynchronous activity that seldom overlap and cover most of the recording period, reminiscent of *in vivo* anaesthetised settings [Tort-Colet *et al.* (2021)].

One key finding resulting from this work is the identification of this same pattern in an isolated cortical tissue *in vitro*, which indicates that such an alternation, which gives rise to an infra-slow rhythm is intrinsic to the local network, in line with some recent models of the transition *in vivo* from deep to light anaesthesia [Pazienti *et al.* (2022)]. Furthermore we found that the synchronicity tenor correlates spatially across the MEA, with the columnar organisation having the strongest effect. There is a co-occurrence of asynchronous states, whereby the sync/async alternation pattern is well conserved within columns, and it can also occasionally propagate across columns. When analysing the synchronous periods, we found that they consist of highly regular alternating Up and Down-like states that are dynam-

ically coordinated within cortical columns. We then decomposed the oscillation into sync-up and sync down states and, interestingly, we found that the duration of sync-ups in supra-granular layers is robustly conserved across columns and slices. Furthermore, the firing dynamics by analysing MUA intensity show that the mean MUA is higher in IG than in SG for every state, whereby the MUA intensity of the asynchronous periods is constrained by the MUA intensities of the synchronous sub-states in a way that differs between laminar layers. Strikingly, we found that the mean MUA is conserved between synchronous and asynchronous periods, which has important implications on fundamental properties of the network, and the relation between sync-Downs and sync-Ups in relation to async. In terms of their MUA, async states are closer to sync-Ups in IG, but closer to sync-Downs in SG.

The remaining of this chapter starts by discussing some of our methodological contributions. Commencing with our adaptation of the MUA estimation, we then discuss the limitations of previous approaches for analysing awaking-like states and highlight the advantages of our new technique. Next, we focus on the interpretation of our biological results in the context of the current literature. This includes the presence of the infra-slow rhythm in the *in vitro* slice model, and the role of the cortical column as the unit sustaining the oscillatory activity. We also put in context our findings regarding the dynamics of the SO-like activity during the AS. Finally, we propose a potential mechanism for the generation of SO-like oscillations during the synchronous periods embedded in the AS.

5.2 Methodology: theory and application

5.2.1 Estimation of the MUA

In this work we developed a slight enhancement to the procedure for estimating the multi-unit activity (MUA). Our three-step procedure departs from the high-frequency band of the extra-cellular recording (ECR) and consists of a short-time discrete Fourier transform, followed by a quantile normalisation of the resulting power coefficients, to finish with an averaging of these previously homogenised coefficients (see Methods, Section 2.3 on p. 24). While the theoretical basis of our estimation is the same as used in other studies [Reig *et al.* (2010), Capone *et al.* (2019), Tort-Colet *et al.* (2021)] (all stemming from the work of Mattia and Del Giudice [Mattia and Del Giudice (2002)]) the advantage of our approach is the preservation of the original power unit.

The main difference between our procedure to that of other authors consists in the homogenisation of the local power spectra brought about by the quantile normalisation (step 2 of the described procedure). Without it, the MUA estimation can be strongly biased by the inverse-power law of the high-pass spectrum of the ECR [Milstein *et al.* (2009)]. This would mean giving more weight to slower, but more powerful, activity (thought to be less local), potentially neglecting the activity of the highest frequencies (ca. 1000 Hz), which is likely to represent the single units closer to the electrode.

In addition, another particularity of our procedure is that it does not ground the MUA to a reference value, as it is usually done in works using the MUA in the context of SO recordings [Reig *et al.* (2010), Sancristóbal *et al.* (2016), D'Andola *et al.* (2017), Capone *et al.* (2019), Barbero-Castillo *et al.* (2021)]

[Camassa *et al.* (2022)]. The reference value therein taken is typically an average value of the Down states, by which the MUA is divided, so that the logMUA of the Down state would hover at 0. This procedure seems appropriate for the analysis of SO recordings, but could pose a number of problems for the kind of analysis we have carried out throughout the present work:

- (i) It implicitly requires a previous identification of Up/Down states. This makes that procedure ill-suited for transitional regimes involving altered forms of SO—as the AS here investigated—, let alone for those recordings lacking any form of SO altogether.
- (ii) It assumes that Down states are essentially activity-silent states (from the local network point of view, at least; and not in [Camassa *et al.* (2022)]). That was precisely a property we were interested in investigating.
- (iii) It prevents a meaningful comparison between the intensity of MUA signals coming from different channels in a multi-channel set-up as ours—and *a fortiori*, from different slices or experiments.

Thus, our approach circumvents these potential problems all in enabling a meaningful comparison between the MUA of different channels. In particular, our characterisation of the MUA densities, depending both on the electrode’s spatial location and the temporal network state (see Biological Results, Chapter 4, Section 4.4), would not have been possible without our objective estimation of the MUA.

5.2.2 Limitations of previous approaches to detecting states in similar AS contexts

Our novel method addresses some of the shortcomings besetting previous approaches to detect the breakdown of the SO regime—or comparable scenarios as this one—, in which bi-stable dynamics (*eg*, the SO) cannot be exclusively assumed [Mochol *et al.* (2015), Jercog *et al.* (2017), Tort-Colet *et al.* (2021)]. For example, in [Tort-Colet *et al.* (2021)] there was a need to detect so-called micro-arousals amidst interrupted SO during the fading of anaesthesia *in vivo*. To discern between Up/Down alternations and longer windows of sustained firing—the latter identified there as micro-arousals—the authors relied on an observational method based on the 1-s-moving variance of the LFP signal. An *ad hoc* 15% threshold on the cumulated moving variance helped them separate between periods of oscillations (higher variance) and micro-arousals (lower variance). Since the variance is equivalent to the spectral power of the signal, their method amounts to defining the micro-arousal periods as LFP periods of low energy. To a certain extent, this goes partially in line with one step of our own approach, in which we detect periods of synchrony as those beyond a certain significant level of instantaneous band-restricted power (but in our case, of the logMUA, see Section 3.2.2). However, in the referred paper, such threshold level of variance was not statistically assessed. Moreover, the authors did not take into account the steep inverse-power law of the LFP signal ¹, which would disproportionately allocate high variances to the lower frequencies, and thereby would risk to inflate the chance of falsely rejecting micro-arousal periods.

¹Our investigations on the spectrum of LFP signals yielded spectral slopes of *ca.* 0.5–1 Hz, well above the 0.1–0.4 of the logMUA.

In [Mochol *et al.* (2015)], a mechanistic approach was followed which was applied to spiking activity. There, the authors distinguished between synchronized and desynchronized local brain states in anaesthetised mice by means of what they called the *silence density*: define as the ratio of steadily silent units during every 20-ms periods. Depending on fixed thresholds set upon the mean values of this silence density, they went on classifying various recording periods from synchronized (with the highest silence density) to desynchronized states (with the lowest one), passing through what they called an intermediate state. Thus, their method consisted in identifying desynchronized states as those which mostly lacked silent periods, or Down phases in their terminology. (Notice that owing to their experimental conditions—involving light anaesthesia and auditory cortex recordings, *inter alia*—, no trace of sustained stereotypical SO was there apparent.) However principled, their method relied on a relatively scant number of spike trains, and the silence density was only estimated during interleaved short 1.5-s periods that were deemed spontaneous (approximately 50% of the total recording time, which also included external auditory stimulation). Furthermore, no objective justification was provided for the choice of thresholds that distinguished between intermediate and synchronized states. [Mochol *et al.* (2015)]-method was also used in [Jercog *et al.* (2017)], at the preprocessing stage, and solely to identify long-enough periods of recognisable SO. As in the former, their resort to rather *ad hoc* thresholds in their doing.

Advantages of our method

This method helps to quantify the alternating pattern between different levels of synchronization in the local cortical networks, which, for lack of an appropriate methodology, had remained an elusive phenomenon until now. By unequivocally decomposing the signal into time periods corresponding to two states, the method facilitates the investigation of the spatio-temporal distribution of the network's synchronicity and the rhythm of its alternation. Moreover, it indirectly provides a measure of the likelihood that each level of synchronization arise for a given recording, and could thus be adapted for uses in future studies aiming at identifying subtler intermediate states. Importantly, our method proposes a working definition of phenomenological asynchrony at the network level, and in doing so, strives to be transparent in its modelling assumptions and statistical hypotheses.

In addition to the dataset in our study, the method can be applied to other broad-band extracellular recording modalities or model-simulated signals. It seems particularly promising for the detection and quantification of transient synchronous events arising in a number of classical experimental settings such as sharp-wave ripples in the hippocampus or spindle oscillations in the cortex during sleep. Whereas most of the analyses usually resort to *ad hoc* thresholds parameters and seldom provide an objective quantification of the oscillatory intensity, our normalisation technique provides a nuanced criterion for threshold selection, which allows to compare directly the power of the oscillation even between recordings from different subjects or preparations. Background noise levels often vary between individual recordings, and our method controls for that. Although not the current focus of this method, examples beyond *in vitro* recordings include detection of theta oscillations in the hippocampus during spatial navigation [Hasselmo *et al.* (2002)], or detection of re-activation events during sleep, such as sharp wave ripples [Joo and Franck (2018)].

Furthermore, our procedure has been thoroughly conceived to be easily scalable to a wide range of settings in which competing oscillatory activity in different bands could arise, or even a mixture of various background noises could be inferred. Thus, an interesting future line of research might be to adapt our method to provide additional insights in such scenarios.

NB: Discussions concerning technical aspects of the methodology can be found in Chapter 3, Sections 3.1.1.3 on p. 42 and 3.1.2.5 on p. 53.

5.3 Biological results

5.3.1 The awaking-like state as the spontaneous infra-slow alternation of sync/asynchronous dynamics

Our newly devised method enabled us to investigate the spatial and temporal dynamics of the transition from a default SO state towards a completely decorrelated activity that would experimentally simulate an state of wakefulness. It is the transitional regime between these two dynamical extremes—synchrony and asynchrony—that we have called the awaking-like state (AS). We analysed data from multi-unit activity (MUA) recorded in ferret visual cortex slices that had been subjected to pharmacological treatments for the induction of the state transition. The key findings discussed in this section revealed that (1) the AS can be effectively decomposed into synchronous and asynchronous periods, and (2) the two complementary identified states tend to alternate temporally (key example in Figure 4.6 on Section 4.1). Interestingly, the regularity observed in such state alternation points to the existence of an infra-slow rhythm *in vitro*, in line with mounting evidence *in vivo* [Leopold *et al.* (2003), Tort-Colet *et al.* (2021), Gonzalo Cogno *et al.* (2024)].

Awake states are typically associated with a higher degree of temporally and spatially asynchronous firing [van Vreeswijk and Sompolinsky (1996), Destexhe (2009), Renart *et al.* (2010)] compared to sleep, and in particular the SO. Therefore, the transition from sleep to wakefulness brings the brain into a dynamical regime of an enriched repertoire of activity, including the appearance of micro-arousal (MA) periods [Schartner *et al.* (2017)]. Several studies based on computational models suggest that the transition between brain states underlies changes in local, non-linear, dynamics of neuronal assemblies [Hill and Tononi (2005), Curto *et al.* (2009), Destexhe (2009), Liang *et al.* (2023)]. Others have suggested that the activity patterns that appear during such state transition must be driven by an external input [Tort-Colet *et al.* (2021), Jercog *et al.* (2017), Paziienti *et al.* (2022)]; or both mechanisms [Camassa *et al.* (2022)]. To date, the underlying mechanisms are still under debate and the network dynamics remain relatively unexplored.

A recent study investigated the transition of the network dynamics *in vivo* as anaesthesia diminishes in the rat visual cortex, in which highly synchronized slow oscillation periods alternating with asynchronous ones was described [Tort-Colet *et al.* (2021)]. They found that, while initially slow oscillations increase in frequency with a peak of maximum regularity, this is followed by a sudden onset of an infra-slow alternation (≈ 0.2 Hz) between sleep-like oscillations and activated states. This finding suggests that the dynamical richness in awake cortical networks arises from a competition between two metastable attractor states. Furthermore, in that same work,

such alternation was replicated in a computational model using the brain stem as the subcortical pacemaker that would drive the resultant infra-slow rhythm.

Our results are in line with those from this study regarding the observation of the occurrence of an infra-slow alternation. Interestingly, however, our findings are obtained from an isolated cortical tissue—as it is the cerebral cortex slice *in vitro*—which strongly suggests that local cortical networks are on their own sufficient to generate such a rhythmic alteration. Thus, the spontaneous intrinsic emergence of asynchronous periods interleaved with an SO-like activity, as observed here in AS recordings, would be in contrast with the hypothesis that this alternation is a global phenomenon mostly facilitated by an external drive of the network [Contreras and Steriade (1995), Destexhe and Sejnowski (2003), Destexhe *et al.* (2007), Crunelli and Hughes (2010), Tort-Colet *et al.* (2021)]. It is worth stressing that this thesis not only shows the existence of such alternation, but also provides insights into detailed dynamics of this regime (*eg*, sync-Up/Down variability and inter-laminar phase-locking, or columnar spectral coordination). These aspects of the AS dynamics were not studied before to the best of our knowledge.

Slow oscillation wave propagation dynamics have been extensively studied at multiple spatial scales in SO [Massimini *et al.* (2004), Capone *et al.* (2019)], [Sanchez-Vives *et al.* (2017), Muller *et al.* (2018)]. Substantial evidence suggests that the SO propagates as a travelling wave front that moves globally from anterior to posterior regions of the brain [Massimini *et al.* (2004), Muller *et al.* (2018)]. However, how this SO propagation breaks down and whether asynchronous states propagate during this transition period remains scarcely studied and thus largely unknown, and, to our knowledge, has not been studied before in isolated cortical networks.

Nevertheless, a recent study investigating this activity pattern *in vivo* across different anaesthesia levels provided relevant insights [Pazienti *et al.* (2022)]. This study introduces a quantitative measure called the wave entropy index (WEI) to assess the complexity of slow waves. The findings reveal a correlation between WEI, the number of spontaneously expressed propagation modes, and the progressive loss and recovery of sequential memory between consecutive waves. A model based on these findings suggests that changes in local excitability can drive the dynamical transition in the global cortical network during the arousal process from deep anaesthesia, which was also in agreement with findings from [Dasilva *et al.* (2021)]

In our study, we provide evidence that would partially support the hypothesis of a locally mediated transition, as we consistently observe periodic infra-slow alternations in recordings from isolated cortical tissue. However, our analyses suggest that the cross-columnar propagation proper to standard SO activity could hardly occur for a temporally and spatially segregated synchronous regime as the SO-like observed in the transitional AS regime, which seem to confine propagation vertically within columns which (see below, Section 5.3.2).

It should be noted, however, that the comparison of network dynamics between deep and light anaesthesia is not necessarily equivalent to an “awaking-like” state. While deep anaesthesia may be comparable in terms of the network dynamics of the SO, light anaesthesia is qualitatively different to a natural waking process, due to effects of the anaesthetic agent itself, such as positive allosteric modulation of GABA-A receptors and blocking of NMDA receptors [Topf *et al.* (2003), Brosnan (2011)].

These receptors are essential for neurotransmission and synaptic plasticity, thus their modulation can dramatically alter network dynamics.

5.3.2 The oscillatory activity during synchronous substates is sustained within the cortical column

In our study, not only did we detect the presence in of near asynchronous states cohabiting in alternation with SO-like oscillations, we also set out to determine some of their precise dynamical features in time and space. In particular, our analyses expand on previous studies by combining the laminar and columnar dimensions of the cortical structure to the temporal and spectral characterization of metastable states occurring during the AS; an approach which was hitherto largely unexplored.

First, we investigated whether the proportional time spent in any one of the two complementary states of the AS was spatially organised along columns and layers. We found that the proportional synchrony content is positively correlated in the MEA space: similar within the same column (even between IG and SG layers), and tending to diverge between different columns. Thus, in regards to overall synchronicity levels, the columnar structure seems the most determinant organisational level influencing the state transitions during the AS (Figures 4.8 and 4.9 in Section 4.2.1).

This was further validated in a following analysis with a better time resolution, where we found that the asynchronous periods occur simultaneously along the same column, irrespectively of the laminar depth. This is seldom the case when columns are different (Figure 4.11 in Section 4.2.2.)

Then, when analysing the spectral properties of the synchronous component of the AS, we found that the same trend is followed. Namely, our results show that the synchronous spectral component of the AS state consists of a single main oscillation whose frequency, intensity, and persistence over time are well maintained within cortical columns, but not between different ones. Thus, we found no evidence of global coordination of the synchronous activity across different column. In other words, the distinction between infra- and supra-granular layer appears again to play no role (see Figure 4.16 in Section 4.3.1).

Taken together, the three previous results indicate that the interleaved oscillations of the AS appear spatially segregated across different cortical columns. This locality can be put in context of the observation that interlaminar propagation *in vivo* is much weaker in awake [Kharas *et al.* (2022)], which may be due to a lower E/I ratio in awake. However, it is interesting to note that the lack of symmetry between IG/SG layers reported in our study has, to the date, not been found *in vivo* [Kharas *et al.* (2022)]. Thus, our methodology may enable future studies to further the understanding of synchronous intervals in the awaking state.

Strongly interconnected neuronal assemblies within the same cortical column could facilitate emerging oscillatory synchrony in the AS state. Along these lines, Krueger and colleagues [Krueger *et al.* (2008)] suggested that an intrinsic property of such local neuronal assemblies may be to mediate sleep [Krueger *et al.* (2008)]. Specifically, their theoretical model showed that the sleep-like states synchronise within individual cortical columns, which, they propose as the minimal processing units that can manifest a sleep-like state [Roy *et al.* (2008)]

This hypothesis is based on experimental findings of rats *in vivo*, in which

evoked responses of potentials were measured to infer wake-like and sleep-like states in cortical columns [Rector *et al.* (2005)]. The authors found that, while most columns showed “sleep-like” responses during sleep, as expected, some individual columns were in an awake-like state, suggesting that sleep may be a property of individual cortical columns [Rector *et al.* (2005)].

In summary, by showing that the column maintains the spatio-temporal co-occurrence and the spectral properties of the synchronous states, our findings would also support the central role of the column in the transition from sleep to awake [Krueger *et al.* (2008), Rector *et al.* (2005)]. For example, we found that while some columns show to be predominantly kept in a synchronous state, others have much higher proportion of asynchronous activity. This scenario would entail that the cortical column might functionally determine each of the transitions between the SO periods and the asynchronous ones. To the best of our knowledge, we present here the first direct evidence in an *in vitro* model supporting that the cortical column is the structural unit mediating the sleep-wake transition. Furthermore, our results suggest that the experimental slice model is a valid approach for studying sleep-to-wakefulness state transitions.

5.3.3 SO-like synchronous activity during the AS

In section results 4.3 we proposed features that help characterise the synchronous activity embedded in the AS transitional regime. Synchronous periods consist of a main oscillation with a frequency of *ca* 1.5–2 Hz which appears conserved between channels located in the same MEA column, but not otherwise (Figure 4.17). Furthermore, the oscillations are strongly coupled in phase between intra-columnar signals (and only between them), and propagate from deeper to shallower layers within the columns (Figure 4.29). The oscillation is akin to the stereotypical SO in that it is generated by the regular alternation of two sub-states: a higher-firing sync-Up and a rather quiescent sync-Down (for this, see also Section 4.4). However, the properties of these sub-states vary considerably depending on the relative laminar location of the electrodes wherein they are observed. Whereas in IG, sync-Ups and sync-Downs appear equally variable in duration and MUA “noisiness”, and sync-Downs are more unstable in mean firing value; in SG, sync-Downs are relatively silent, and longer and more variable in duration. Interestingly, the duration of SG sync-Ups has proven very well conserved (*ca* 180 ms) irrespectively of their column and across all the slices.

Our result is partially consistent with studies supporting that Down states become longer with decreasing levels of global excitation [Sanchez-Vives *et al.* (2010), Sanchez-Vives *et al.* (2017)]. Previous findings *in vivo* also show a systematic and pronounced shortening of Down states and increase oscillation frequency as anaesthesia fades out [Pazienti *et al.* (2022)]. The particularity of our experimental setup, which allows probing distinct cortical columns in the novel awaking-like state, is that it can simultaneously give access to neuronal ensembles with different local levels of excitation. Complementarily, in [Jercog *et al.* (2017)], they also found shorter Up than Down states in average. However, they also reported that Up states were highly variable in duration, as much as Downs. This may seem at odds with our findings. In fact, in that study, no differentiation was made between laminar location of the states, and thus, they may have missed the well-maintained conservation only specific to SG Up states we observed. (Figure 4.22).

Moreover, our reported increase of sync-Down duration with decreasing oscillation frequency may well align with an enhanced “deterministic” network coordination for shorter Down periods *in vitro* under direct current stimulation [Camassa *et al.* (2022), Barbero-Castillo *et al.* (2021)] and under light anaesthesia [Dasilva *et al.* (2021), Camassa *et al.* (2022)]. Thus, the afterhyperpolarization mechanism driving Down states onset might underlie the linear mapping between sync-Down and the oscillation period in IG and SG (Figure 4.22).

We have also shown that the transitions from Down to Up, and *vice versa*, are not symmetrical between them, depending again on the laminar depth. SG sync-Ups appeared to be anchored to their termination, where the most of the firing intensity tended to concentrate. Notably, their termination phase was well maintained relatively to the oscillation cycle. This suggests a deterministic mechanism governing the shut-down of the sync-Up states in SG layers, consistent with the hypothesis of a stochastic generation of Up states in SO regimes *in vitro* followed by deterministic termination through adaptation mechanisms (*eg*, [Sanchez-Vives and Mattia (2014), Sancristóbal *et al.* (2016), Mattia *et al.* (2021), Sanchez-Vives *et al.* (2017)]).

Finally, we also demonstrated that the lag of the sync-Up onset is what drives the inter-laminar propagation of the oscillation during the synchronous periods. Strikingly, whereas sync-Up occurs *ca* 80 ms earlier in IG, there seems to be only a negligible delay between layers at their termination. Interestingly, the onset lag strongly and positively correlated with the frequency of the oscillation.

A proposed mechanism for the generation of SO-like oscillation during synchronous periods

Therefore, the dynamics of the transiently synchronous sub-states we found and characterised suggest a possible mechanism for their accelerated SO-like oscillation, as follows:

Prone-to-fire IG assemblies easily transition from noisy unstable sync-Downs to robustly firing sync-Ups. Meanwhile, the local inter-laminar connectivity within the columnar network may foster after-hyperpolarization and thus determine how long the mostly silent SG assemblies must wait until they are sufficiently stochastic [Camassa *et al.* (2022), Capone *et al.* (2019)] to transition to a short, snappy sync-Up. Owing to the low variability of their duration and to their abrupt and coordinated termination, SG sync-Up’s sustainment might be determined by well-conserved parameters of the local assembly that forces the network to shut down once a certain level of excitability has been reached, for example.

Now, the grossly simultaneous termination of Up states across the whole column (common for standard SO, *eg*, [Mochol *et al.* (2015)]) seems therefore to be mediated by elements of the network mainly dependent on shallower layers. For example, a type of layer-I GABAergic interneurons called NDNF interneurons seems a fitting candidate. In fact, NDNF interneurons play an important role in inhibiting and orchestrating the whole cortical column, and this in a state-dependent manner: their influence becoming more decisive with increased arousal levels [Cohen-Kashi *et al.* (2021)].

Finally, known fatigue parameters of the IG local networks will determine

how intense and lasting the subsequent sync-Down state in IG will be until they can again transition to an Up state [Mattia and Sanchez-Vives (2012), Mattia *et al.* (2021)]. That is, in IG, the longer the sync-Up, the stronger the fatigue, and hence the longer the deterministic part of the subsequent sync-Down.

In summary, the frequency of the whole alternation might be mainly explained by similar mechanisms to the ones driving the standard SO regime, that is, the ability of IG assemblies to elicit a sync-Up state in shallower layers, and hence the crucial role of the inter-laminar time lag of the sync-Up onset.

Conclusions

- C1)** We proposed a theoretically motivated normalisation of the wavelet power spectrogram (WPSg), which allows time-frequency power components from distinct logMUA signals to be statistically compared against their intrinsic background when modelled as fractional Gaussian noise.
- C2)** We developed a multi-step Monte Carlo statistical procedure based on the normalisation of the WPSg (**C1**) to decompose multi-channel logMUA signals into synchronous and asynchronous periods of activity.
- C3)** Our method (**C2**) was able to detect alternating periods of synchronous and asynchronous activity in isolated cortical slices during the awaking-like state (AS).
- C4)** Our time-frequency analyses of spatially segregated AS recordings revealed that, although the alternation between synchrony and asynchrony (**C3**) pervades the whole network, the cortical layer dictates the firing rate intensity and the cortical column sustains and confines the oscillatory activity.
- C5)** The columnar AS oscillations (**C4**) consist of two alternating sub-states whose duration, variability and inter-laminar phase-coupling remind those of an accelerated form of slow oscillations.

Appendix A

Supplementary Mathematical Details

A.1 On the admissibility condition of the reduced Morlet wavelet

For all its convenience, the reduced Morlet wavelet does not comply with the wavelet *admissibility condition* (1.6). In fact,

$$\widehat{\Psi}_0(\omega) = \begin{cases} \pi^{-\frac{1}{4}} e^{-\frac{(\omega-\omega_0)^2}{2}}, & \text{if } \omega \geq 0 \\ 0, & \text{if } \omega < 0 \end{cases}$$

and since $\widehat{\Psi}_0(0) \neq 0$, $\lim_{a \rightarrow 0} \int_a^{+\infty} \frac{|\widehat{\Psi}_0(\omega)|^2}{|\omega|} d\omega = +\infty$. Fortunately, the subset of admissible wavelets (*ie*, square-integrable functions that verify the admissibility condition) is dense in $L^2(\mathbb{R})$ [Louis *et al.* (1997)]; that means we would be able to approximate Ψ_0 by an admissible wavelet as close as desired. In practice, because we will be dealing with finite signals and will therefore have a lower bound ω_{\min} on the accessible circular frequencies, we propose such an approximation (call it $\Psi_0^{\omega_{\min}}$) to be locally defined through its Fourier transform on a neighbourhood of the angular centre frequency ω_0 :

$$\widehat{\Psi}_0^{\omega_{\min}} = \begin{cases} \kappa \widehat{\Psi}_0(\omega), & \text{if } \omega \in \left[\frac{\omega_{\min}}{2}; 2\omega_0 - \frac{\omega_{\min}}{2} \right] \\ 0, & \text{if not} \end{cases} \quad (\text{A.1})$$

where κ is a normalising factor to ensure unit energy. This way, the corresponding admissibility constant of such modified Morlet wavelet,

$$C' := C_{\Psi_0^{\omega_{\min}}} \quad (\text{A.2})$$

could be estimated accordingly. With our usual set of parameters N , ω_0 , we will use $C' \approx 1.0624$.

Successful theoretical attempts to loosen the admissibility condition allow the reduced Morlet wavelet to equivalently abide by most of the results on wavelet analysis we are concerned with (see e.g. [Lebedeva and Postnikov (2014)]). In any case, it must be mentioned that, although it is to a modified Morlet wavelet that we are implicitly resorting to for some theoretical developments, in practice, when dealing with stationary processes (theoretically), or with finite signals (whether theoretically or for numerical applications), under an amenable selection of parameters, one can nonetheless confidently use the original reduced Morlet [Hramov *et al.* (2015)].

A.2 Wavelet components of Gaussian processes: proofs of two lemmata

Lemma 1.2.2

Indications for a proof. The result may be seen as a special version of Fubini's theorem, applied to a stochastic context. For a partial proof involving Riemann sums and convergence in L^2 see [Ash and Gardner (1975)] p. 34–36. It can then be adapted to include a broader class of covariance functions. Note that the interval I needn't be bounded, but for the last part we do require that the process admit a spectral representation. \square

Lemma 1.2.3

Proof. Let be $t_0 \in I$, $s > 0$; we call $W = \mathcal{W}X(s, t_0)$ and $\psi = \psi_{s, t_0}$. Let us consider $\mathcal{L} := L^2(\{X_t : t \in I\})$, the closure in $L^2(\Theta)$ of the \mathbb{C} -linear span of the random variables constituting X . Since any \mathbb{C} -linear combination of jointly normal variables is complex normal (and dwells *a fortiori* in L^2), and the limit of any L^2 -convergent sequence of such variables will also be so (L^2 -convergence implies convergence in distribution), \mathcal{L} only contains complex normal variables.

Now, from the previous lemma, $\mathbb{E}(W) = 0$ and $\mathbb{E}(W\overline{W}) < \infty$ (W is in L^2).

Finally, if we prove that $W \in (\mathcal{L}^\perp)^\perp$ the result will follow. In fact, let be $Y \in \mathcal{L}^\perp$; using again the previous lemma we get:

$$\mathbb{E}(Y\overline{W}) = \mathbb{E}\left(Y \int_I X_t \psi(t) dt\right) = \mathbb{E}\left(\int_I Y X_t \psi(t) dt\right) = \int_I \mathbb{E}(Y X_t) \psi(t) dt = 0$$

\square

A.3 Kolmogorov-Smirnov distance

In this work, we are often interested in evaluating how different two given probability densities functions (*pdf*) f and g are: the MUA densities of two states or two different channels, for instance. Since such densities are continuous real functions with bounded Lebesgue integral, one could in principle turn to one of the usual functional distances based on the L^p -norms, [Dudley (2002)]. That is:

$$d_{L^p}(f, g) := \|f - g\|_p = \int |f - g|^p$$

defined whenever the right-term integral exists. Or the special case for bounded functions (L^0):

$$d_\infty(f, g) := \|f - g\|_\infty = \sup |f - g|$$

where \sup denotes the least upper bound of the range. For all cases, the implicit definition interval is the real line.

However, in practice these MUA densities can only be known from their estimates, given by random samples of their values. This introduces two difficulties: (1) the construction of a finite-sample estimate of a pdf with continuous support is a delicate operation, usually involving a number of parameters such as the size and number of bins, or the smoothing procedure; (2) even when the latter is sorted out, it will remain to assess how likely it is that the estimated densities of two different samples corresponding to the same density be different—and by how far.

As we shall see, both difficulties can be easily addressed by resorting to a distance based on the two-sample Kolmogorov-Smirnov test's statistic (KS statistic) [Feller (1948)]. This will provide a simple way of comparing two densities through their samples.

A.3.1 From the KS statistic to a distance between densities

We recall that for a (X_1, \dots, X_n) random sample of the population f , the empirical distribution function (ecdf) or sampling function is, defined as:

$$\hat{F}_n(x) = \frac{1}{n} \sum_{i=0}^n \mathbb{1}_{X_i \leq x} \quad , \text{ for any } x \in \mathbb{R} \quad (\text{A.3})$$

where $\mathbb{1}_{X_i \leq x} = \mathbb{1}_{]0, x]}(X_i)$ and $\mathbb{1}$ denotes the indicator function.

We will loosely refer by *ecdf* to the random variable above defined or to any of its observations. Now, the statistic of the two sample Kolmogorov-Smirnov test:

Definition A.3.1 (KS Statistic). *Given $\hat{F}_{1,n}$ and $\hat{F}_{2,m}$, empirical cumulative distributions functions of two random samples with underlying population densities f_1, f_2 , and sample size n, m , respectively, the KS statistic $D_{n,m}$ is defined as:*

$$D_{n,m}(\hat{F}_{1,n}, \hat{F}_{2,m}) := \max |\hat{F}_{1,n} - \hat{F}_{2,m}| \quad (\text{A.4})$$

Thus, for two samples, an observation of $D_{n,m}$ is simply the $\|\cdot\|_\infty$ -norm distance between two bounded step functions: $D_{n,m}(\hat{F}_{1,n}, \hat{F}_{2,m}) = d_\infty(\hat{F}_{1,n}, \hat{F}_{2,m})$

It follows that, for samples the same size, $n = m$, the statistic $D_{n,n}$ is a metric over

observed sampling distributions. That is, for any three samples $\hat{F}_{1,n}$, $\hat{F}_{2,n}$, $\hat{F}_{3,n}$, it satisfies

$$\begin{aligned} D_{n,n}(\hat{F}_{1,n}, \hat{F}_{2,n}) &= 0 \Leftrightarrow \hat{F}_{1,n} = \hat{F}_{2,n} \quad (\text{reflexivity}) \\ D_{n,n}(\hat{F}_{1,n}, \hat{F}_{2,n}) &= D_{n,n}(\hat{F}_{2,n}, \hat{F}_{1,n}) \quad (\text{symmetry}) \\ D_{n,n}(\hat{F}_{1,n}, \hat{F}_{2,n}) &\leq D_{n,n}(\hat{F}_{1,n}, \hat{F}_{3,n}) + D_{n,n}(\hat{F}_{3,n}, \hat{F}_{2,n}) \quad (\text{triang. ineq.}) \end{aligned} \tag{A.5}$$

Notice also that $D_{n,n} \leq 1$ is bounded.

The KS as a statistical distance

Such distance has so far only been defined over samples through their empirical distributions, and not over their underlying (theoretical) population densities. We can nonetheless go on analogously defining a distance over continuous densities.

Definition A.3.2 (KS distance between two densities). *Given f and g two densities with absolute continuous distribution functions F and G , respectively, we define the KS distance between them as:*

$$d_{KS}(f, g) := \sup |F - G| \tag{A.6}$$

Again, because of the properties of the functional distance d_∞ , the latter will well behave as a distance between this specific class of functions. Furthermore, we can relate the d_{KS} to the KS-statistic of any corresponding ecdfs \hat{F}_n and \hat{G}_n of f and g , respectively. In fact, the Glivenko-Cantelli theorem ensures that $D_{n,n}(\hat{F}_n, \hat{G}_n) \xrightarrow[n \rightarrow +\infty]{\text{a.s.}} d_{KS}(f, g)$ (see eg [Dudley (2002)], p.400). Thus, we could in principle assimilate the KS distance between two densities to the *KS distance between two any samples* from each of them.

However, such assimilation would compromise treating d_{KS} as a distance. In fact, for two distinct samples of the same f with ecdfs $\hat{F}_{1,n}$ and $\hat{F}_{2,n}$, respectively, it will be the case that:

$$D_{n,n}(\hat{F}_{1,n}, \hat{F}_{2,n}) \neq 0 \quad \text{whereas obviously} \quad d_{KS}(f, f) = 0$$

breaching the reflexivity.

Therefore, it remains to be seen how well, in practice, distances taken over two distinct random samples of the same distribution relate to each other. Interestingly, owing to the asymptotic distribution of the KS-statistic—upon which the two-sample KS test rests—, for a fixed n , a level of confidence $1 - \alpha$ can be given for observing the value of $D_{n,n}$ under the hypothesis that the two samples $\hat{F}_{1,n}$ and $\hat{F}_{2,n}$ are drawn from the same density. Namely, with $1 - \alpha$ confidence level,

$$D_{n,n} \leq \sqrt{\frac{2}{n}} q_{1-\alpha}$$

where $q_{1-\alpha}$ is the $(1 - \alpha)$ -quantile of the theoretical distribution of the KS statistic $D_{n,n}$ [Smirnov (1948)].

From the latter and the triangular inequality (A.5), a confidence level can also be established for the discrepancy between the KS statistic of two samples, each

drawn from two populations.

Let \hat{F}_i and \hat{G}_i , $i = 1, 2$ be two pairs of sampling distributions from populations f, g , respectively. All of them with the same sample size n . Then, with $1 - \alpha$ confidence level, we will have:

$$|D_{n,n}(\hat{F}_1, \hat{G}_1) - D_{n,n}(\hat{F}_2, \hat{G}_2)| \leq 2\sqrt{\frac{2}{n}} q_{1-\alpha}$$

(For example, with $n = 1000$, and a confidence level of 90%, the error between the KS statistics of two pair of samples drawn from the same respective densities will be smaller than 0.011.)

Formally, this means that the KS statistic induces a *statistical distance* between probability densities. In practice, we will loosely refer by d_{KS} to the observed statistic $D_{n,n}$ between empirical distributions, rather than to the theoretical statistical distance, implying that it is giving us a reasonable and assessable approximation to the d_{KS} properly defined for the theoretical densities. And when the implicit sample sizes are big enough, this will not be a major concern for our numerical applications.

A.3.2 Two simple properties

Planar representation

A straightforward consequence of the triangular inequality of d_{KS} is that any three given empirical distributions F_1, F_2 and F_3 , as a 3-element set, will always admit a trivial isometric embedding in \mathbb{R}^2 equipped with its usual Euclidean distance d_2 .

Proof. Let us denote by $a_{ij} := d_{\text{KS}}(F_i, F_j)$, $1 \leq i, j \leq 3$, the KS-distance between the F_i . Now the map

$$\begin{aligned} \Gamma : (\{F_1, F_2, F_3\}, d_{\text{KS}}) &\longrightarrow (\mathbb{R}^2, d_2) \\ \{F_1 \mapsto (0, 0), F_2 \mapsto (0, a_{12}), F_3 \mapsto (h_1, h_2)\} \end{aligned}$$

with

$$h_1 := \sqrt{a_{13}^2 - h_2^2} \quad \text{and} \quad h_2 := \frac{1}{2a_{12}} (a_{12}^2 + a_{13}^2 - a_{23}^2) \quad (\text{A.7})$$

will do (see Figure A.1). So defined, Γ will indeed satisfy:

$$d_2(\Gamma(F_i), \Gamma(F_j)) = d_{\text{KS}}(F_i, F_j)$$

□

From the latter we can derive a simple way of comparing one density with respect to two other densities over one single linear dimension.

Definition A.3.3 (KS projection-ratio of one density over two others' directed distance). *Given three empirical distributions F_1, F_2 and F_3 , we define the KS projection-ratio of F_3 over F_1 -to- F_2 directed distance, as the real number*

$$\Pi_{\text{KS}}(F_3 | F_1 \rightarrow F_2) := a_{12}^2 + a_{13}^2 - a_{23}^2$$

where $a_{ij} := d_{\text{KS}}(F_i, F_j)$, $1 \leq i, j \leq 3$.

That is, $\Pi_{\text{KS}}(F_3 | F_1 \rightarrow F_2) = \frac{h_2}{d_{\text{KS}}(F_1, F_2)}$, where h_2 has been defined in eq. (A.7).

The idea behind is to get a proxy of the relation between the three distributions, would F_3 have lied half-way through a minimal deformation path from F_1 to F_2 .




Figure A.1: Planar representation of the KS-distances between three empirical densities and orthogonal projection. $\Gamma_i = \Gamma(F_i)$ are the images in \mathbb{R}^2 of the F_i , with coordinates depending on the $d_{\text{KS}}(F_i, F_j)$, so that $d_2(\Gamma(F_i), \Gamma(F_j)) = d_{\text{KS}}(F_i, F_j)$. For instance, $\Gamma_3 = (h_1, h_2)$, with h_2 the ordinate of the orthogonal projection of Γ_3 over $(\Gamma_1\Gamma_2)$. The KS projection-ratio of F_3 between F_1 and F_2 can be read accordingly.

Distribution-free distance

Finally, as a by-product of the distribution-free property of the KS statistic [Everitt *et al.* (2010), Feller (1948)], for every one-to-one function φ we will have:

$$d_{\text{KS}}(f_1 \circ \varphi, f_2 \circ \varphi) = d_{\text{KS}}(f_1, f_2) \quad (\text{A.8})$$

In practice, because of the injectivity of the *log* transform, this entails that KS-distances between MUA samples or between logMUA ones will be the same.

Appendix B

Research Dissemination and Complementary Formation

Conferences and Workshops

- **Barcelona Computational, Cognitive and Systems Neuroscience (BARCCSYN).**
Centre de Recerca Matemàtica (CRM), Barcelona, Spain. 15th–16th June 2017.
Poster presentation: *Inference of Global States Stability in Cortical Networks*.
- **EITN Workshop: "Brain States: Models and Experiments."**
European Institute for Theoretical Neuroscience (EITN), Paris, France. 1st–2nd March 2018.
- **Analysis and Modeling of Complex Oscillatory Systems (AMCOS) Conference.**
Universitat Pompeu Fabra (UPF), Barcelona, Spain. 19th–23rd March 2018.
Poster presentation: *Transient Dynamics of Cortical Networks Underpinning Slow Oscillations*.
- **Bernstein Conference 2018.**
The Bernstein Network Computational Neuroscience. Berlin, Germany. 25th–28th September 2018.
Poster presentation: *Identifying Propagation Modes of Spontaneous Activity in Cortical Slices* [Arango et al. (2018)].
- **The 28th Computational Neuroscience Conference (CNS).**
The Organization for Computational Neuroscience. Barcelona, Spain. 13th–17th July 2019.
Poster presentation: *Propagating densities of spontaneous activity in cortical slices* [Arango et al. (2019)].
- **The 30th Computational Neuroscience Conference (CNS).**
The Organization for Computational Neuroscience. Online. 3rd–7th July 2021.
Poster presentation: *Emergence and propagation of asynchronous states of spontaneous cortical activity* [Arango et al. (2021)].

Competitively Funded Courses

- **First HBP Curriculum Workshop Series. *Understanding the Brain: Neurobiology for Non-Specialists***, which included passing an exam on the HBP Curriculum Online Course.
Human Brain Project (HBP) Education Programme. Medizinische Universität Innsbruck, Austria.
3rd–5th July 2017.
- **8th G-Node Short Course on Neural Data Analysis.**
German Neuroinformatics Node. Ludwig-Maximilians-Universität (LMU). Munich, Germany.
31st July – 4th August, 2017.
- **Advanced Computational Neuroscience School (aCNS): *Neural Circuit Theories for Primate Cognition*.**
Bernstein Center for Computational Neuroscience (BCCN), Max Planck Institute for Dynamics and Self-Organization. Göttingen, Germany.
2nd–10th September 2017.

Bournemouth University's Outreach Activities

- **SciTech PGR Conference.**
24th May 2017.
Poster presentation: *Inferring Stability of Global States in Cortical Networks*.
- **The 10th Annual Postgraduate Conference.**
7th March 2018.
Poster presentation: *Transient Dynamics of Cortical Networks Underpinning Slow Oscillations*.
- **SciTech PGR Conference.**
9th May 2018.
Poster presentation: *Transient Dynamics of Cortical Networks Underpinning Slow Oscillations*.
- **Postgraduate Research Live Exhibition.**
5th December 2018.
Poster presentation: *Propagating Densities of Spontaneous Activity in Cortical Slices*.
- **SciTech PGR Conference.**
22th May 2019.
Oral presentation: *Propagating Densities of Spontaneous Activity in Cortical Slices*.

Bibliography

- [Abry *et al.* (1995)] Abry P, Gonçalvès P and Flandrin P. (1995). *Wavelets, spectrum analysis and 1/f processes*. In: *Wavelets and Statistics. Lecture Notes in Statistics, vol 103, pp.15–30*. Ed. Antoniadis A and Oppenheim G. Berlin: Springer-Verlag.
- [Arango *et al.* (2018)] Arango R, Mateos-Aparicio P, Sanchez-Vives MV and Balaguer-Ballester E. (2018). *Identifying Propagation Modes of Spontaneous Activity in Cortical Slices*. Bernstein Conference 2018. doi: 10.12751/nncn.bc2018.0020
- [Arango *et al.* (2019)] Arango R, Mateos-Aparicio P, Balaguer-Ballester E and Sanchez-Vives MV. (2019). *Propagating densities of spontaneous activity in cortical slices*. In: *The 28th Computational Neuroscience Conference (CNS). Barcelona, 13th–17th July 2019*. BMC Neurosci, 20(Suppl 1):56(P276).
- [Arango *et al.* (2021)] Arango R, Mateos-Aparicio P, Sanchez-Vives MV and Balaguer-Ballester E. (2021). *Emergence and propagation of asynchronous states of spontaneous cortical activity*. In: *30th Annual Computational Neuroscience Meeting: CNS*2021–Meeting Abstracts*. J Comput Neurosci, 49(Suppl 1):33(P8).
- [Arion *et al.* (2007)] Arion D, Unger T, Lewis DA and Mirnic K. (2007). *Molecular markers distinguishing supragranular and infragranular layers in the human prefrontal cortex*. Eur J Neurosci, 25(6):1843–54.
- [Ash and Gardner (1975)] Ash RB and Gardner MF. (1975). *Topics in stochastic processes*. London: Academic Press.
- [Barbero-Castillo *et al.* (2021)] Barbero-Castillo A, Mateos-Aparicio P, Porta LD, Camassa A, Perez-Mendez L and Sanchez-Vives MV. (2021). *Impact of GABA_A and GABA_B inhibition on cortical dynamics and perturbational complexity during synchronous and desynchronized states*. J Neurosci, 41:5029–44.
- [Barton and Poor (1988)] Barton RJ and Poor VH. (1988). *Signal detection in fractional Gaussian noise*. IEEE Trans Inf Theory, 34:943–59.
- [Bausch (2013)] Bausch J. (2013). *On the efficient calculation of a linear combination of chi-square random variables with an application in counting string vacua*. J Phys A: Math Theor, 46:505202.
- [Bazhenov *et al.* (2002)] Bazhenov M, Timofeev I, Steriade M and Sejnowski TJ. (2002). *Model of thalamocortical slow-wave sleep oscillations and transitions to activated states*. J Neurosci, 22(19):8691–704.
- [Berger (1929)] Berger H. (1929). *Über das Elektroencephalogramm des Menschen*. Arch Psychiatr Nervenkr, 87(1):527–70.
- [Bharioké *et al.* (2022)] Bharioké A, Munz M, Brignall A, Kosche G, Eizinger MF, Ledergerber N, Hillier D, Gross-Scherf B, Conzelmann KK, Macé E and Roska B. (2022). *General anesthesia globally synchronizes activity selectively in layer 5 cortical pyramidal neurons*. Neuron, 110:1–17.

- [Bolstad *et al.* (2003)] Bolstad BM, Irizarry RA, Astrand and Speed TP. (2003). *A comparison of normalization methods for high density oligonucleotide array data based on variance and bias*. *Bioinformatics*, 19(2):185–193.
- [Bonett (2006)] Bonett DG. (2006). *Confidence interval for a coefficient of quartile variation*. *Comput Stat Data Anal*, 50:2953–7.
- [Brette and Destexhe (2012)] Brette R and Destexhe A (ed.). (2012). *Handbook of neural activity measurement*. Cambridge, UK: Cambridge University Press.
- [Briggs and Usrey (2008)] Briggs F and Usrey WM. (2008). *Emerging views of corticothalamic function*. *Curr Opin Neurobiol*, 18(4):403–7.
- [Brosnan (2011)] Brosnan RJ. (2011). *GABA(A) receptor antagonism increases NMDA receptor inhibition by isoflurane at a minimum alveolar concentration*. *Vet Anaesth Analg*, 38(3):231–9.
- [Bullmore *et al.* (2001)] Bullmore E, Long C, Suckling J, Fadili J, Calvert G, Zelaya F, TA Carpenter and Brammer M. (2001). *Colored noise and computational inference in neurophysiological (fMRI) time series analysis: resampling methods in time and wavelet domains*. *Hum Brain Mapping*, 12:61–78.
- [Burns *et al.* (2010)] Burns SP, Xing D and Shapley RM. (2010). *Comparisons of the dynamics of the local field potential and multiunit activity signals in macaque visual cortex*. *J Neurosci*, 30(41):13739–49.
- [Büssow (2007)] Büssow R. (2007). *An algorithm for the continuous Morlet wavelet transform*. *MSSP*, 21:2970–9
- [Buzsáki (2006)] Buzsáki G. (2006). *Rhythms of the brain*. Oxford, UK: Oxford University Press.
- [Buzsáki *et al.* (2012)] Buzsáki G, Anastassiou CA and Koch C. (2012). *The origin of extracellular fields and currents – EEG, ECoG, LFP and spikes*. *Nat Rev Neurosci*, 13(6):407–20.
- [Buzsáki and Mizuseki (2014)] Buzsáki G and Mizuseki K. (2014). *The log-dynamic brain: how skewed distributions affect network operations*. *Nat Rev Neurosci*, 15(4):264–78.
- [Byron *et al.* (2021)] Byron N, Semenova A and Sakata S. (2021). *Mutual interactions between brain states and Alzheimer’s disease pathology: a focus on gamma and slow oscillations*. *Biology*, 10(8):707.
- [Capone *et al.* (2019)] Capone C, Rebollo B, Muñoz A, Illa X, Del Giudice P, Sanchez-Vives MV and Mattia M. (2019). *Slow waves in cortical slices: how spontaneous activity is shaped by laminar structure*. *Cereb Cortex*, 29(1):319–335.
- [Chatfield (2004)] Chatfield C. (2004). *The analysis of time series: an introduction, 6th ed.* Boca Raton, FL: CRC Press.
- [Camassa *et al.* (2022)] Camassa A, Galluzzi A, Mattia M and Sanchez-Vives MV. (2022). *Deterministic and stochastic components of cortical down states: dynamics and modulation*. *J Neurosci*, 42(50):9387–400.
- [Chauvette *et al.* (2011)] Chauvette S, Crochet S, Volgushev M and Timofeev I. (2011). *Properties of slow oscillation during slow-wave sleep and anesthesia in cats*. *J Neurosci*, 31:14998–5008.
- [Chen *et al.* (2009)] Chen Z, Vijayan S, Barbieri R, Wilson MA and Brown EN. (2009). *Discrete- and continuous-time probabilistic models and algorithms for inferring neuronal UP and DOWN states*. *Neural Computation*, 21:1797–862.
- [Cohen-Kashi *et al.* (2021).] Cohen-Kashi Malina K, Tsivourakis E, Kushinsky D, Apellat D, Shtiglitz S, Zohar E, Sokoletsky M, Tasaka GI, Mizrahi A, Lampl I and Spiegel I. (2021). *NDNF interneurons in layer 1 gain-modulate whole cortical columns according to an animal’s behavioral state*. *Neuron*, 109(13):2150–64.e5.

- [Compte *et al.* (2003)] Compte A, Sanchez-Vives MV, McCormick DA and Wang XJ. (2003). *Cellular and network mechanisms of slow oscillatory activity (μ 1 Hz) and wave propagations in a cortical network model*. J Neurophysiol, 89:2707–25.
- [Compte *et al.* (2008)] Compte A, Reig R, Descalzo VF, Harvey MA, Puccini GD and Sanchez-Vives MV. (2008). *Spontaneous high-frequency (10–80 Hz) oscillations during Up states in the cerebral cortex in vitro*. J Neurosci, 28(51):13828–44.
- [Contreras and Steriade (1995)] Contreras D and Steriade M. (1995). *Cellular basis of EEG slow rhythms, a study of dynamic corticothalamic relationships*. J Neurosci, 15:604–22.
- [Crunelli and Hughes (2010)] Crunelli V and Hughes SW. (2010). *The slow (μ 1 Hz) rhythm of non-REM sleep: a dialogue between three cardinal oscillators*. Nat Neurosci, 13(1):9–17.
- [Curto *et al.* (2009)] Curto C, Sakata S, Marguet S, Itskov V and Harris KD. (2009). *A simple model of cortical dynamics explains variability and state dependence of sensory responses in urethane-anesthetized auditory cortex*. J Neurosci, 29(34):10600–12.
- [D’Andola *et al.* (2017)] D’Andola M, Rebollo B, Casali AG, Weinert JF, Pigorini A, Villa R, Massimini M, Sanchez-Vives MV. (2017). *Bistability, causality, and complexity in cortical networks: an in vitro perturbational study*. Cereb Cortex, 28(7):2233–42.
- [Dasilva *et al.* (2021)] Dasilva M, Camassa A, Navarro-Guzman A, Pazienti A, Perez-Mendez L, Zamora-López G, Mattia M and Sanchez-Vives MV. (2021). *Modulation of cortical slow oscillations and complexity across anesthesia levels*. Neuroimage 224:117415.
- [Daubechies (1992)] Daubechies, I. (1992). *Ten lectures on wavelets*. Philadelphia, PA: SIAM.
- [Davis *et al.* (2021)] Davis ZW, Benigno GB, Fletterman C, Desbordes T, Steward C, Sejnowski TJ, Reynolds JH and Muller L. (2021). *Spontaneous traveling waves naturally emerge from horizontal fiber time delays and travel through locally asynchronous-irregular states*. Nat Comm, 12:6057.
- [Denève and Machens (2016)] Denève S and Machens CK. (2016). *Efficient codes and balanced networks*. Nat Neurosci, 19(3):375–82.
- [Destexhe and Sejnowski (2003)] Destexhe A and Sejnowski TJ. (2003). *Interactions between membrane conductances underlying thalamocortical slow-wave oscillations*. Physiol Rev, 83(4):1401–53.
- [Destexhe *et al.* (2007)] Destexhe A, Hughes SW, Rudolph M, and Crunelli V. (2007). *Are corticothalamic ‘up’ states fragments of wakefulness?* Trends Neurosci, 30:334–42.
- [Destexhe (2009)] Destexhe A. *Self-sustained asynchronous irregular states and Up-Down states in thalamic, cortical and thalamocortical networks of nonlinear integrate-and-fire neurons*. J Comput Neurosci, 27(3):493–506.
- [Donoghue *et al.* (2020)] Donoghue T, Haller M, Peterson EJ, Varma P, Sebastian P, Gao R, Noto T, Lara AH, Wallis JD, Knight RT, Shestyuk A and Voytek B. (2020). *Parameterizing neural power spectra into periodic and aperiodic components*. Nat Neurosci, 32:1655–65.
- [Douglas and Martin (2004)] Douglas RJ and Martin KAC. (2004). *Neuronal circuits of the neocortex*. Annu Rev Neurosci, 27:419–51.
- [Dudley (2002)] Dudley RM. (2002). *Real analysis and probability*. Cambridge, UK: Cambridge University Press.
- [Einevoll *et al.* (2007)] Einevoll GT, Pettersen KH, Devor A, Ulbert I, Halgren E and Dale AM. (2007). *Laminar population analysis: estimating firing rates and evoked synaptic activity from multielectrode recordings in rat barrel cortex*. J Neurophysiol, 97(3):2174–90.

- [Everitt *et al.* (2001)] Everitt BS, Landau S, Leese M and Stahl D. (2001). *Cluster analysis, 5th ed.* Chichester, UK: John Wiley & Sons Ltd.
- [Everitt *et al.* (2010)] Everitt BS and Skrdonal A. (2010). *The Cambridge dictionary of statistics, 4th ed.* Cambridge, UK: Cambridge University Press.
- [Farge (1992)] Farge M. (1992) *Wavelet transforms and their applications to turbulence.* Annu Rev Fluid Mech, 24:395–457.
- [Farge and Schneider (2006)] Farge M and Schneider K. *Wavelets: Application to Turbulence.* (2006). In: *Encyclopedia of Mathematical Physics*, ed. Françoise JP, Naber GL and Tsun TS. London, UK: Academic Press. 2006.
- [Feller (1948)] Feller W. (1948). *On the Kolmogorov-Smirnov limit theorems for empirical distributions* Ann Math Stat, 19:177–89.
- [Flandrin (1989)] Flandrin P. (1989). *On the spectrum of fractional Brownian motions.* IEEE Trans Inf Theory, 35:197–9.
- [Flandrin (1992)] Flandrin P. (1992). *Wavelet analysis and synthesis of fractional Brownian motion.* IEEE Trans Inf Theory, 38:910–7.
- [Gallager (2013)] Gallager RG. (2013). *Stochastic Processes: Theory for Applications.* Cambridge, UK: Cambridge University Press.
- [Ge (2007)] Ge Z. (2007). *Significance tests for the wavelet power and the wavelet power spectrum.* Ann Geophys, 25:2259–69.
- [Gonzalo Cogno *et al.* (2024)] Gonzalo Cogno S, Obenhaus HA, Lautrup A, Jacobsen RI, Clopath C, Andersson SO, Donato F, Moser MB and Moser EI. (2024). *Minute-scale oscillatory sequences in medial entorhinal cortex.* Nature, 625(7994):338–44.
- [Gonzalo Rey *et al.* (2015)] Gonzalo Rey H, Pedreira C and Quiñan Quiroga R. (2015). *Past, present and future of spike sorting techniques.* Brain Research Bulletin, 119:106–17.
- [Hansen *et al.* (2012)] Hansen BJ, Chelaru MI and Dragoi V. (2012). *Correlated Variability in Laminar Cortical Circuits* Neuron, 76(3):590–602.
- [Hasselmo *et al.* (2002)] Hasselmo ME, Hay J, Ilyn M and Gorchetchnikov A. (2002). *Neuromodulation, theta rhythm and rat spatial navigation.* Neural Netw, 15(4–6):689–707.
- [Hay *et al.* (2021)] Hay YA, Deperrois N, Fuchsberger T, Quarrell TM, Koerling AL and Paulsen O. (2021). *Thalamus mediates neocortical Down state transition via GABAB-receptor-targeting interneurons.* Neuron, 109(17):2682–90.e5.
- [He *et al.* (2010)] He BJ, Zempel JM, Snyder AZ and Raichle ME. (2010). *The temporal structures and functional significance of scale-free brain activity.* Neuron, 66:353–69.
- [Hill and Tononi (2005)] Hill S and Tononi G. (2005). *Modeling sleep and wakefulness in the thalamocortical system.* J Neurophysiol, 93(3):1671–98.
- [Hramov *et al.* (2015)] Hramov AE, Koronovskii AA, Makarov VA, Pavlov AN and Sitenikova E. (2015). *Wavelets in neuroscience.* Berlin: Springer-Verlag.
- [Hubel and Wiesel (1963)] Hubel DH and Wiesel TN. (1963). *Shape and arrangement of columns in cat's striate cortex.* J Physiol, 165:559–68.
- [Huber *et al.* (2004)] Huber R, Ghilardi MF, Massimini, M and Tononi G. (2004). *Local sleep and learning.* Nature, 430:78–81.
- [Illa *et al.* (2015)] Illa X, Rebollo B, Gabriel G, Sánchez-Vives MV and Villa R. (2015). *A SU-8-based flexible microprobe for close and distal recordings from the cortical network.* In: SPIE Microtechnologies. International Society for Optics and Photonics. p. 951803.
- [Istas (1992)] Istas J. (1992). *Wavelet coefficients of a Gaussian process and applications.* Ann Inst Henri Poincaré, 28(4):537–556.

- [Jercog *et al.* (2017)] Jercog D, Roxin A, Barthó P, Luczak A, Compte A and de la Rocha J. (2017). *UP-DOWN cortical dynamics reflect state transitions in a bistable network*. eLife, 6:e22425.
- [Jones (2005)] Jones BE. (2005). *From waking to sleeping: neuronal and chemical substrates*. Trends Pharmacol Sci, 26(11):578–86.
- [Joo and Franck (2018)] Joo HR and Frank LM. (2018). *The hippocampal sharp wave-ripple in memory retrieval for immediate use and consolidation*. Nat Rev Neurosci, 19(12):744–57.
- [Kasdin (1995)] Kasdin NJ. (1995). *Discrete Simulation of Colored Noise and Stochastic Processes and $1/f^\alpha$ Power Law Noise Generation*. Proc IEEE, 83(5):802–27.
- [Kharas *et al.* (2022)] Kharas N, Andrei A, Debes SR and Dragoi V. (2022). *Brain state limits propagation of neural signals in laminar cortical circuits*. Proc Natl Acad Sci, 119(30):e2104192119.
- [Krueger *et al.* (2008)] Krueger JM, Rector DM, Roy S, Van Dongen HP, Belenky G and Panksepp J. (2008). *Sleep as a fundamental property of neuronal assemblies*. Nat Rev Neurosci, (12):910–9.
- [Kutner *et al.* (2005)] Kutner MH, Nachtsheim CJ, Neter J and Li W. (2005). *Applied Linear Statistical Models, 5th ed.* New York, NY: McGraw-Hill/Irwin.
- [Kuo (1996)] Kuo H-H. (1996). *White Noise Distribution Theory*. Boca Raton, FL: CRC Press.
- [Lachaux *et al.* (1999)] Lachaux JP, Rodriguez E, Martinerie J and Varela FJ. (1999). *Measuring phase synchrony in brain signals*. Hum Brain Mapp, 8(4):194–208.
- [Larkum *et al.* (2018)] Larkum ME, Petro LS, Sachdev RNS and Muckli L. (2018). *A Perspective on Cortical Layering and Layer-Spanning Neuronal Elements*. Front Neuroanat, 12:56.
- [La Rocca *et al.* (2018)] La Rocca D, Zilber N, Abry P, van Wassenhove V and Ciuciu P. (2018). *Self-similarity and multifractality in human brain activity: A wavelet-based analysis of scale-free brain dynamics*. J Neurosci Meth, 309:175–187.
- [Lebedeva and Postnikov (2014)] Lebedeva EA and Postnikov EB. (2014). *On alternative wavelet reconstruction formula: a case study of approximate wavelets*. R Soc open sci, 1:140124.
- [Leopold *et al.* (2003)] Leopold DA, Murayama Y and Logothetis NK. (2003). *Very slow activity fluctuations in monkey visual cortex: implications for functional brain imaging*. Cereb Cortex, 13(4):422–33.
- [Lewis *et al.* (2012)] Lewis LD, Weiner VS, Mukamel EA, Donoghue JA, Eskandar EN, Madsen JR, Anderson WS, Hochberg LR, Cash SS, Brown EN and Purdon PL. (2012). *Rapid fragmentation of neuronal networks at the onset of propofol-induced unconsciousness*. Proc Natl Acad Sci, 109(49):E3377–86.
- [Liang *et al.* (2023)] Liang, Y., Liang, J., Song, C. et al. Complexity of cortical wave patterns of the wake mouse cortex. Nat Commun 14, 1434 (2023).
- [Louis *et al.* (1997)] Louis A, Maaß P and Rieder A. (1997). *Wavelets, Theory and Applications* Chichester, UK: John Wiley & Sons Ltd.
- [McCormick *et al.* (2003)] McCormick DA, Shu Y, Hasenstaub A, Sanchez-Vives MV, Badoual M and Bal T. (2003). *Persistent cortical activity: mechanisms of generation and effects on neuronal excitability*. Cereb Cortex, 13(11):1219–31.
- [Mallat (1998)] Mallat S. (1998). *A Wavelet Tour of Signal Processing, 2nd ed.* San Diego, CA: Academic Press.

- [Mandelbrot and van Ness (1968)] Mandelbrot BB and van Ness JW. (1968) *Fractional Brownian motions, fractional noises and applications*. SIAM Rev, 10:422–437.
- [Maraun *et al.* (2007)] Maraun D, Kurths J and Holschneider M. (2007). *Nonstationary Gaussian processes in wavelet domain: Synthesis, estimation, and significance testing*. Phys Rev E, 75:016707.
- [Marchini and Ripley (2000)] Marchini JL and Ripley BD. (2000). *A new statistical approach to detecting significant activation in functional MRI*. NeuroImage, 12:366–80.
- [Marshall *et al.* (2006)] Marshall L, Helgadóttir H, Mölle M and Born J. (2006). *Boosting slow oscillations during sleep potentiates memory*. Nature, 444(7119):610–3.
- [Massimini *et al.* (2004)] Massimini M, Huber R, Ferrarelli F, Hill S, and Tononi G. (2004). *The sleep slow oscillation as a traveling wave*. J Neurosci, 24:6862–70.
- [Massimini *et al.* (2005)] Massimini M, Ferrarelli F, Huber R, Esser SK, Singh H and Tononi G. (2005). *Breakdown of cortical effective connectivity during sleep*. Science, 309:2228–32.
- [The MathWorks Inc. (2018)] The MathWorks, Inc. (2018). *MATLAB version: 9.5.0 (R2018b)*. Natick, MA.
- [Mattia and Del Giudice (2002)] Mattia M and Del Giudice P. (2002). *Population dynamics of interacting spiking neurons*. Phys Rev E, 66 051917.
- [Mattia *et al.* (2010)] Mattia M, Ferraina S and Del Giudice P. (2010). *Dissociated multi-unit activity and local field potentials: A theory inspired analysis of a motor decision task*. NeuroImage, 52:812–23.
- [Mattia and Sanchez-Vives (2012)] Mattia M and Sanchez-Vives MV. (2012). *Exploring the spectrum of dynamical regimes and timescales in spontaneous cortical activity*. Cogn Neurodyn, 6(3):239–50.
- [Mattia *et al.* (2013)] Mattia M, Pani P, Mirabella G, Costa S, Del Giudice P and Ferraina S. (2013). *Heterogeneous attractor cell assemblies for motor planning in premotor cortex*. J Neurosci, 33(27):11155–68.
- [Mattia *et al.* (2021)] Mattia M, Perez-Zabalza M, Tort-Colet N and Sanchez-Vives MV. (2021). *Multiscale dynamics underlying neocortical slow oscillations*. bioRxiv (preprint), doi: 10.1101/2021.06.09.447804
- [Mazzoni *et al.* (2013)] Mazzoni A, Logothetis NK and Panzeri S. (2013). *The information content of Local Field Potentials: experiments and models*. In: *Principles of Neural Coding*, ed. Quian Quiroga R and Panzeri S. Boca Raton, FL: CRC Press. 2013.
- [Milstein *et al.* (2009)] Milstein J, Mormann F, Fried I and Koch C. (2009). *Neuronal shot noise and Brownian $1/f^2$ behavior in the local field potential*. PLoS One, 4(2):e4338.
- [Mizuseki and Buzsáki (2013)] Mizuseki K and Buzsáki G. (2013). *Preconfigured, skewed distribution of firing rates in the hippocampus and entorhinal cortex*. Cell Rep, 4(5):1010–21.
- [Mochol *et al.* (2015)] Mochol G, Hermoso-Mendizabal A, Sakata S, Harris KD and de la Rocha J. (2015). *Stochastic transitions into silence cause noise correlations in cortical circuits*. PNAS, 112:3529–34.
- [Moschopoulos and Canada (1984)] Moschopoulos PG and Canada WB. (1984). *The distribution function of a linear combination of chi-squares*. Comput Math with Appl, 10(4/5):383-6.
- [Mountcastle (1997)] Mountcastle VB. (1997). *The columnar organization of the neocortex*. Brain, 120(4):701–22.
- [Muller *et al.* (2018)] Muller L, Chavane F, Reynolds J and Sejnowski TJ. (2018). *Cortical Travelling Waves: Mechanisms and Computational Principles*. Nat Rev Neurosci, 19:255–68.

- [Nieto-Reyes *et al.* (2014)] Nieto-Reyes A, Cuesta-Albertos JA and Gamboa F. (2014). *A random-projection based test of Gaussianity for stationary processes*. *Comput Stat Data Anal*, 75:124–41.
- [Ohki *et al.* (2005)] Ohki K, Chung S, Ch'ng YH, Kara P and Reid RC. (2005). *Functional imaging with cellular resolution reveals precise micro-architecture in visual cortex*. *Nature*, 433(7026):597–603.
- [Pavan *et al.* (2022)] Pavan B, Bianchi A, Botti G. (2022). *In vitro cell models merging circadian rhythms and brain waves for personalized neuromedicine*. *iScience*, 25(12):105477.
- [Pazienti *et al.* (2022)] Pazienti A, Galluzzi A, Dasilva M, Sanchez-Vives MV and Mattia M. (2022). *Slow waves form expanding, memory-rich mesostates steered by local excitability in fading anesthesia*. *iScience*, 25(3):103918.
- [Percival (1995)] Percival DP. (1995). *On estimation of the wavelet variance*. *Biometrika*, 82(3):619–31.
- [Perrier *et al.* (1995)] Perrier V, Philipovitch T and Basdevant C. (1995). *Wavelet spectra compared to Fourier spectra* *J Math Phys*, 36(3):1506–19.
- [Pettersen *et al.* (2008)] Pettersen KH, Hagen E and Einevoll GT. (2008). *Estimation of population firing rates and current source densities from laminar electrode recordings*. *Comput Neurosci*, 24(3):291–313.
- [Pettersen *et al.* (2012)] Pettersen KH, Lindén H, Dale AM and Einevoll GT. *Extracellular spikes and CSD*. (2012). In: *Handbook of neural activity measurement*, ed. Brette R and Destexhe A. Cambridge, UK: Cambridge University Press. 2012.
- [Poulet and Crochet (2019)] Poulet JFA, Crochet S. (2019). *The cortical states of wakefulness*. *Front Syst Neurosci*, 12:64.
- [Rasch *et al.* (2008)] Rasch MJ, Gretton A, Murayama Y, Maass W and Logothetis NK. (2008). *Inferring spike trains from local field potentials*. *J Neurophysiol*, 99:1461–76.
- [Rebollo *et al.* (2021)] Rebollo B, Telenczuk B, Navarro-Guzman A, Destexhe A and Sanchez-Vives MV. (2021). *Modulation of intercolumnar synchronization by endogenous electric fields in cerebral cortex*. *Sci Adv*, 7:eabc7772.
- [Rector *et al.* (2005)] Rector DM, Topchii IA, Carter KM and Rojas MJ. (2005). *Local functional state differences between rat cortical columns*. *Brain Res*, 1047:45–55.
- [Reig *et al.* (2010)] Reig R, Mattia M, Compte A, Belmonte C and Sanchez-Vives MV. (2010). *Temperature modulation of slow and fast cortical rhythms*. *J Neurophysiol*, 103:1253–61.
- [Rommelzwaal *et al.* (2019)] Rommelzwaal LA, Mishra A and Ellis GFR. (2019). *CTNN: Corticothalamic-inspired neural network*. arXiv:1910.12492v1 (preprint).
- [Renart *et al.* (2010)] Renart A, de la Rocha J, Bartho P, Hollender L, Parga N, Reyes A and Harris KD. (2010). *The asynchronous state in cortical circuits*. *Science*, 327(5965):587–90.
- [Reyes-Puerta *et al.* (2016)] Reyes-Puerta V, Yang JW, Siwek ME, Kilb W, Sun JJ and Luhmann HJ. (2016). *Propagation of spontaneous slow-wave activity across columns and layers of the adult rat barrel cortex in vivo*. *Brain Struct Funct*, 221(9):4429–49.
- [Robert and Casella (2004)] Robert CP and Casella G. (2004). *Monte Carlo Statistical Methods*, 2nd ed. New York: Springer-Verlag.
- [Rosenblum *et al.* (2001)] Rosenblum M, Pikovsky A, Kurths J, Schäfer C and Tass PA. (2001). *Phase synchronization: from theory to data analysis*. In: *Handbook of Biological Physics, vol 4 Neuroinformatics*, pp. 279–321. Ed. Moss F and Gielen S. Amsterdam: Elsevier.

- [Rousseeuw (1987)] Rousseeuw PJ. (1987). *Silhouettes: a graphical aid to the interpretation and validation of cluster analysis*. J Comp App Math, 20:53–65.
- [Roxin *et al.* (2011)] Roxin A, Brunel N, Hansel D, Mongillo G and van Vreeswijk C. (2011). *On the distribution of firing rates in networks of cortical neurons*. J Neurosci, 31(45):16217–26.
- [Roy *et al.* (2008)] Roy S, Krueger JM, Rector DM and Wan Y. (2008). *Network models for activity-dependent sleep regulation*. J Theor Biol, 253:462–8.
- [Sanchez-Vives and McCormick (2000)] Sanchez-Vives MV and McCormick DA. (2000). *Cellular and network mechanisms of rhythmic recurrent activity in neocortex*. Nat Neurosci, 3(10):1027–34.
- [Sanchez-Vives *et al.* (2010)] Sanchez-Vives MV, Mattia M, Compte A, Perez-Zabalza M, Winograd M, Descalzo VF and Reig R. (2010). *Inhibitory modulation of cortical up states*. J Neurophysiol, 104(3):1314–24.
- [Sanchez-Vives and Mattia (2014)] Sanchez-Vives MV, Mattia M. (2014). *Slow wave activity as the default mode of the cerebral cortex*. Arch Ital Biol. 152(2-3):147–55.
- [Sanchez-Vives *et al.* (2017)] Sanchez-Vives MV, Massimini M, Mattia M. (2017). *Shaping the default activity pattern of the cortical network*. Neuron, 94(5):993–1001.
- [Sancristóbal *et al.* (2016)] Sancristóbal B, Rebollo B, Boada P, Sanchez-Vives MV and Garcia-Ojalvo J. (2016). *Collective stochastic coherence in recurrent neuronal networks*. Nat Phys. 12:881–7.
- [Schneider and Farge (2006)] Schneider K and Farge M. *Wavelets: Mathematical Theory*. (2006). In: *Encyclopedia of Mathematical Physics*, ed. Françoise JP, Naber GL and Tsun TS. London, UK: Academic Press. 2006.
- [Schartner *et al.* (2017)] Schartner MM, Pigorini A, Gibbs SA, Arnulfo G, Sarasso S, Barnett L, Nobili L, Massimini M, Seth AK and Barrett AB. (2017). *Global and local complexity of intracranial EEG decreases during NREM sleep*. Neurosci Conscious, 2017(1):niw022.
- [Senzai *et al.* (2019)] Senzai Y, Fernandez-Ruiz A and Buzsáki G. (2019). *Layer-Specific Physiological Features and Interlaminar Interactions in the Primary Visual Cortex of the Mouse*. Neuron, 101(3):500–13.e5.
- [Shu *et al.* (2003)] Shu Y, Hasenstaub A and McCormick D. (2003). *Turning on and off recurrent balanced cortical activity*. Nature, 423:288–293.
- [Shyun and Sun (2002)] Shyun HC and Sun YS. (2002). *Construction of a Morlet wavelet power spectrum* MSSP, 13:101–11.
- [Slater (1964)] Slater LC. *Confluent Hypergeometric Functions*. (1964). In: *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*, ed. Abramowitz M and Stegun IA. Washington DC: United States Department of Commerce, National Bureau of Standards. 1972.
- [Smirnov (1948)] Smirnov N. (1948). *Table for estimating the goodness of fit of empirical distributions*. Ann Math Stat, 19(2):279–81.
- [Somogyi *et al.* (1998)] Somogyi P, Tamás G, Lujan R and Buhl EH. (1998). *Salient features of synaptic organisation in the cerebral cortex*. Brain Res Rev, 26(2–3):113–35.
- [Stark and Abeles (2007)] Stark E and Abeles M. (2007). *Predicting movement from multiunit activity*. J Neurosci, 27(31):8387–94.
- [Steriade *et al.* (1993a)] Steriade M, Nuñez and Amzica F. (1993). *A novel slow (<1 Hz) oscillation of neocortical neurons in vivo: depolarizing and hyperpolarizing components*. J Neurosci, 13(8):3252–65.

- [Steriade *et al.* (1993b)] Steriade M, Nuñez A, Amzica F. *Intracellular analysis of relations between the slow (≈ 1 Hz) neocortical oscillation and other sleep rhythms of the electroencephalogram.* J Neurosci, 13(8):3266–83.
- [Steriade *et al.* (1993c)] Steriade M, Amzica F and Nuñez A. (1993). *Cholinergic and noradrenergic modulation of the slow (approximately 0.3 hz) oscillation in neocortical cells.* J Neurophysiol, 70(4):1385–400.
- [Steriade *et al.* (1993d)] Steriade M, McCormick DA and Sejnowski TJ. (1993). *Thalamocortical oscillations in the sleeping and aroused brain.* Science, 262(5134):679–85.
- [Steriade *et al.* (1996)] Steriade M, Contreras D, Amzica F and Timofeev I. *Synchronization of fast (30–40 Hz) spontaneous oscillations in intrathalamic and thalamocortical networks.* J Neurosci, 16(8):2788–808.
- [Steriade *et al.* (2001)] Steriade M, Timofeev I and Grenier F. (2001). *Natural waking and sleep states: a view from inside neocortical neurons.* J Neurophysiol, 85:1969–85.
- [Teleńczuk and Destexhe (2014)] Teleńczuk B and Destexhe A. *Local field potential, relationship to unit activity.* (2014). In: *Encyclopedia of Computational Neuroscience*, ed. Jaeger D and Jung R. New York: Springer-Verlag. 2015.
- [Terrence and Compo (1998)] Torrence C and Compo GP. (1998). *A practical guide to wavelet analysis.* Bull Am Meteorol Soc, 79(1):61–78.
- [Tewfik and Kim (1992)] Tewfik AH and Kim M. (1992). *Correlation structure of the discrete wavelet coefficients of fractional Brownian motion.* IEEE Trans Inf Theory, 38(2):904–9.
- [Thomson and Lamy (2007)] Thomson AM and Lamy C. (2007). *Functional maps of neocortical local circuitry.* Front Neurosci, 1(1):19–42.
- [Timofeev *et al.* (2000)] Timofeev I, Grenier F, Bazhenov M, Sejnowski TJ and Steriade M. (2000). *Origin of slow cortical oscillations in deafferented cortical slabs.* Cereb Cortex, 10:1185–99.
- [Topf *et al.* (2003)] Topf N, Jenkins A, Baron N and Harrison NL. (2003). *Effects of isoflurane on gamma-aminobutyric acid type A receptors activated by full and partial agonists.* Anesthesiology, 98(2):306–11.
- [Torao-Angosto *et al.* (2021)] Torao-Angosto M, Manasanch A, Mattia M and Sanchez-Vives MV. (2021). *Up and Down states during slow oscillations in slow wave sleep and different levels of anesthesia.* Front Syst Neurosci, 15:609645.
- [Tort-Colet *et al.* (2021)] Tort-Colet N, Capone C, Sanchez-Vives MV and Mattia M. (2021). *Attractor competition enriches cortical dynamics during awakening from anesthesia.* Cell Rep, 35:109270
- [van Vreeswijk and Sompolinsky (1996)] van Vreeswijk C and Sompolinsky H. (1996). *Chaos in neuronal networks with balanced excitatory and inhibitory activity.* Science, 274(5293):1724–6.
- [Veitch and Abry (1999)] Veitch D and Abry P. (1999). *A wavelet based joint estimator of the parameters of long-range dependence.* IEEE Trans Inf Theory, 45(3):878–97.
- [Vergassola and Frisch (1991)] Vergassola M and Frisch U. (1991). *Wavelet transforms of self-similar processes* Physica D, 54:58–64.
- [Voytek *et al.* (2015)] Voytek B, Kramer MA, Case J, Lepage KQ, Tempesta ZR, Knight RT and Gazzaley A. (2015). *Age-Related Changes in 1/f Neural Electrophysiological Noise.* J Neurosci, 35(38):13257–65.
- [Vyazovskiy and Harris (2013)] Vyazovskiy VV and Harris KD. (2013). *Sleep and the single neuron: the role of global slow oscillations in individual cell rest.* Nat Rev Neurosci, 14(6):443–51.

- [Wester and Contreras (2012)] Wester JC and Contreras D. (2012). Columnar interactions determine horizontal propagation of recurrent network activity in neocortex. *J Neurosci*, 32(16):5454–71.
- [Wornell (1993)] Wornell GW. (1993). *Wavelet-based representations for the 1/f family of fractal processes*. *Proc IEEE*, 81:1428–50.