Analysis and Detection of Neural Synchrony in the Prefrontal Cortex

Master Thesis

Jack Tchimino

TUDelft
Delft
University of
Technology

Challenge the future

ANALYSIS AND DETECTION OF NEURAL SYNCHRONY IN THE PREFRONTAL CORTEX

MASTER THESIS

by

Jack Tchimino

Born in Thessaloniki, Greece

in partial fulfillment of the requirements for the degree of

Master of Science
in Biomedical Engineering

at the Delft University of Technology

TUDelft
Delft
University of
Technology
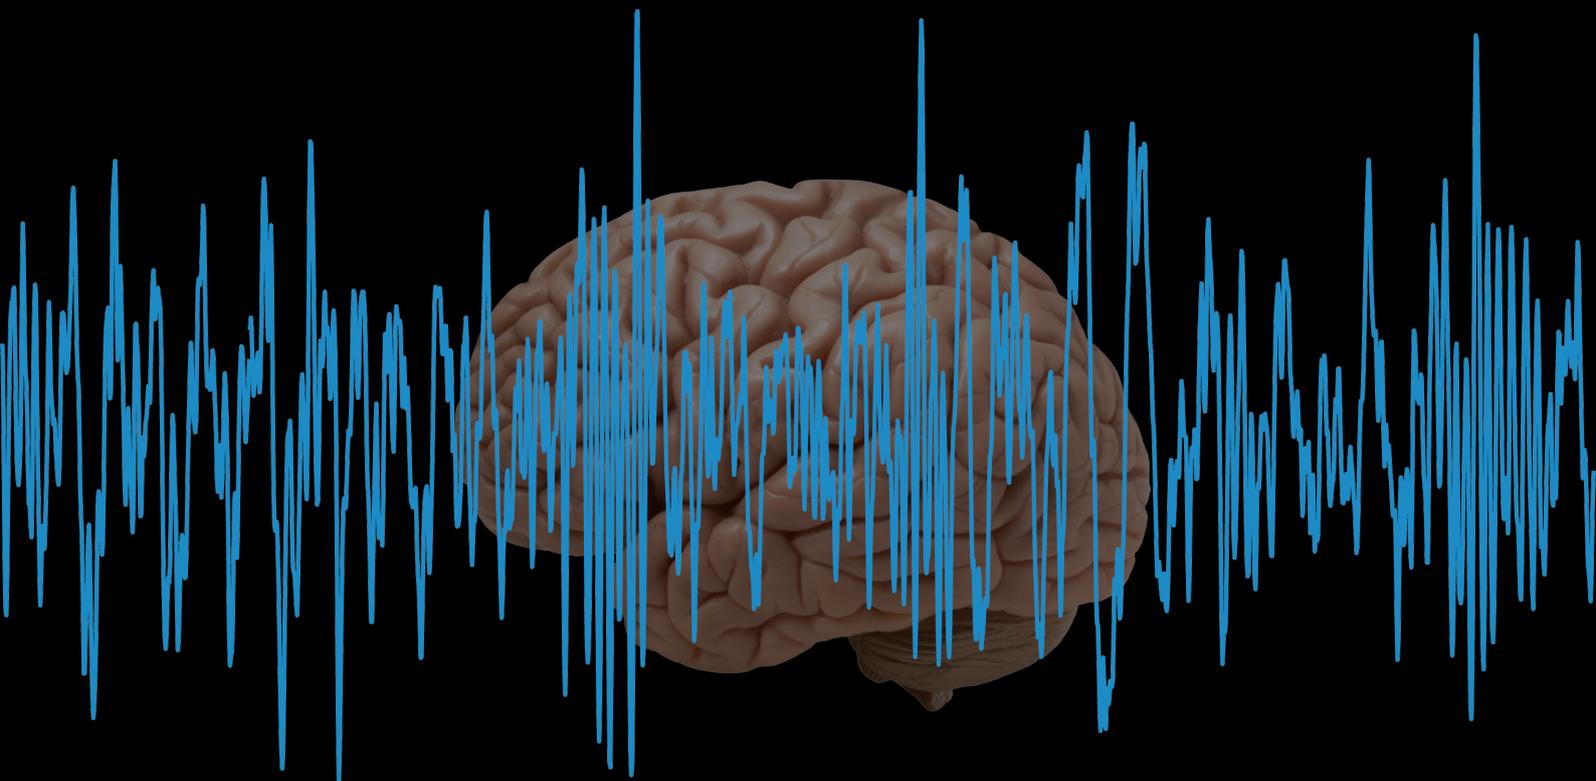
como a vida é frágil, se a abandonam

— José Saramago, "Ensaio Sobre a Cegueira"

Abstract

Signal Analysis techniques are routinely used in Neuroscience to interpret raw signals harvested from the Nervous System. From a simple Fourier analysis to more complicated methods such as multiresolution wavelet analysis, such techniques must be used for signal manipulation in order to reach informed conclusions on the measurements taking place.

In neuroscientific research, it is common practice to scan a neural recording manually to label the areas of the signal that are relevant to the research at hand. This can, obviously, be very time-consuming for the researchers. What is more, this method can prove imperfect, seeing as two different researchers can disagree on the labeling of the data.

In every experiment, signal epochs are isolated because they stand out from the rest of the recording due to a special characteristic, which is, in the previous case, visible with the unaided eye. In signal processing terms, this means that the signal displays specific spectrotemporal characteristics during these epochs. Thus, these characteristics can be isolated, quantified and studied independently, while a detection algorithm can be developed so that the detection and labeling of the significant signal epochs can be carried out automatically.

In this project, the spontaneous activity of the neurons in the Prefrontal Cortex was analyzed in relation to neural synchrony, using time-varying AR Models. It was concluded that the signal epochs of neural synchrony display common characteristics besides being visually similar. This allowed the isolation of the synchrony epochs based on model parameters.

However, the training of the models is very computationally intensive, so a detection algorithm was developed, based on a matched filter which made use of one of the isolated epochs as a template. The detection scheme was then validated using a recording harvested during electrical stimulation of the deep brain, evaluating the quality of the scheme was evaluated.

Finally, the detection scheme was applied to stimulation recordings to study the electrical behavior of the Prefrontal Cortex during electrical stimulation of deep brain structures.

Keywords: prefrontal cortex, neural synchrony, local field potential, biomedical signal analysis, VTA stimulation, change detection, AR modeling, matched filters

Contents

LIST OF FIGURES

LIST OF TABLES

1

INTRODUCTION

Biomedical Signal Analysis is located at the juncture between Medicine and Engineering. It constitutes a prime example of how engineering principles can be applied to a completely different field, providing new methods and points of view and being a large factor in the advancement of medical research.

Signal analysis methods are routinely used in the field of Neuroscience, to interpret electrical signals harvested from the Nervous System. The study of these signals have provided researchers with valuable insight into the complex function of the brain, paving the way for important breakthroughs in our knowledge of how the brain works.

Of course, the methods applied in this kind of experiments have been developed by engineers and signal processing experts who analyzed the signals provided by neuroscientists. The heterogeneity among the signals that can be harvested from the nervous system has led to a multitude of different signal analysis methods to be tailored for their application on this type of time series, creating fertile ground on further research on these methods.

Despite decades of studying the brain, the scientific community agrees that there is still a long way to go before we manage to unlock all its secrets. Research might focus on some areas or functions of the brain, while others have not been so closely studied, or not discovered at all.

The relation between the Prefrontal Cortex and a deep brain structure called the Ventral Tegmental Area will be studied in this project. These two structures communicate through the mesocortical dopaminergic pathway, and their communication is mediated by the release of the neurotransmitter dopamine and monitor a small part of the activity of the brain.

The goal of this project is twofold. It is primarily dedicated to the analysis of the spontaneous activity of the Prefrontal Cortex, which is shown to display short events of neural synchrony, manifesting themselves as oscillatory activity at certain frequencies. Their occurrence spans the entire length of the recording and following no discernible pattern. An analysis and detection scheme for this type of neural synchrony is, eventually, developed.

What is more, it was deemed useful to observe and study the electrical behavior of the Prefrontal Cortex when the Ventral Tegmental Area was electrically stimulated. Since these two are connected, the stimulation is expected to elicit some kind of response at the Prefrontal Cortex. In this case, the quantification of this response would be the fluctuation of the number of neural synchrony events detected in the signal around the stimulation pulses.

The experiments and recordings were carried out using b-6 mice at the Erasmus MC Neuroscience Department, in Rotterdam and the results of the research are presented in this report, which is structured as follows: Chapter 2 is an overview of the function of the nervous system and the structures that are important for this project, Chapter 3 is a presentation of the signal analysis methods that will be used for the study of the neu-

ral recordings, Chapter 4 presents the experimental procedure for the harvesting of the signals, Chapter 5 contains the analysis that was carried out, Chapter 6 presents the conclusions drawn from the analysis and recommendations for future work are provided in Chapter 7.

2

THE NERVOUS SYSTEM: AN OVERVIEW

Since this project is carried out in a medical setting, despite being primarily engineering- and mathematically oriented, it is useful to provide the reader with a brief overview of the biological structures and functions that will be mentioned in the entire length of this report. At the end of this chapter, the motives for the research that follows will have been made clearer.

2.1. BASIC STRUCTURE AND FUNCTION OF THE CENTRAL NERVOUS SYSTEM

The Central Nervous System constitutes the control center of the entire organism. With few exceptions, every operation taking place in the body, either voluntary or involuntary, is controlled by the CNS. The basic functions of the CNS will be analyzed in this section.

2.1.1. ANATOMY AND PHYSIOLOGY

The CNS comprises of two main organs: the brain and the spinal cord. The brain is safely enclosed in the skull, while the spinal cord is enclosed in the vertebral column, running its entire length. Processing of information takes place in the brain, while the spinal cord relays signals between the brain and the rest of the body. These signals can be either *afferent* or *efferent*, with direction to the CNS from body tissues or vice versa respectively.

At their lowest level, all living tissues are made from functional units called *cells*. The cells of the Nervous System are called *neurons*, about 100 billion of which make up the CNS. A typical neuron can be seen in Figure 2.1, from which its intricate structure is apparent. The neuron comprises of two main parts: the *soma* (the large part of the cell, on the left of the image) and the *axon* (the long tail at the bottom), connected at the *axon hillock*.

The soma itself contains most of the cellular organelles necessary for the function of the cell, as well as a large number of tree-like processes, aptly named *dendrites* (from the Greek word for tree: *déntron*). The axon is a long process (as long as one meter in larger neurons), which connects one neuron to the next. The most distal part of the axon, the axon terminal, connects to one of the dendrites of the next neuron, forming a *neural synapse*.

A special species of neurons is that of motor neurons, which innervate muscles. The axons of these neurons terminate on muscles at a different kind of synapse, the *neuromuscular junction*, but these belong to the Peripheral Nervous System and are out of the scope of this work.

The most important function of neurons is the relay of messages from one neuron to the next, which occurs unidirectionally in every neuronal cell, from the soma, through the axon, to the axon terminal and the synapse, to dendrites of the next neuron and so on [1].

2.1.2. THE SYNAPSE AND NEUROTRANSMITTERS

The connection between two neurons, that is, between an axon terminal and a dendrite, is called a neural synapse. Synapses can be either electrical, the rarer kind, or chemical. A schematic of a typical chemical
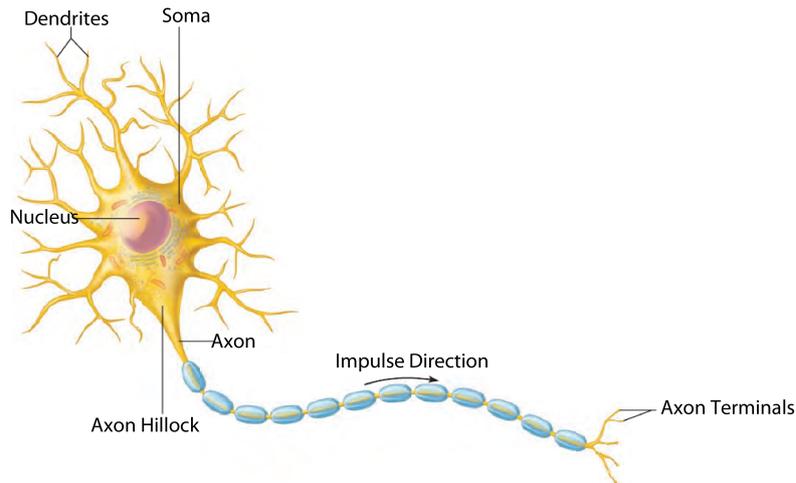
Figure 2.1: A typical neuron. Image modified from [1]

synapse can be seen in Figure 2.2
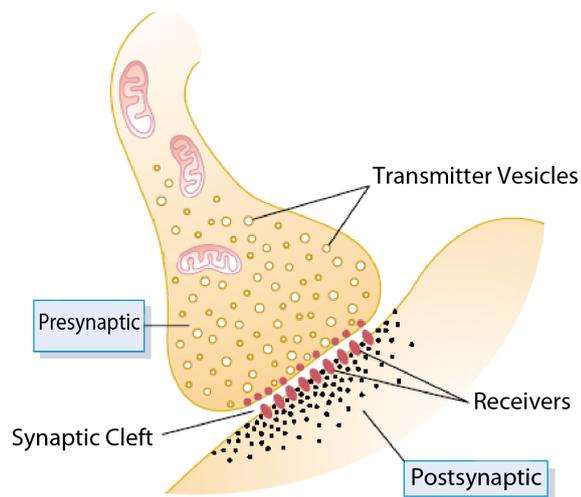
Figure 2.2: A typical chemical synapse. Image courtesy of [2]

The neuron whose axon terminal is part of the synapse is called the pre-synaptic neuron, while the one who receives the message is called the post-synaptic neuron. The narrow extracellular space between the two is called the *synaptic cleft*. In very simple terms, when the pre-synaptic neuron delivers a message to the synapse in the form of an *action potential* (which will be analyzed in 2.2.1), vesicles from the cell interior release a *neurotransmitter* into the synaptic cleft, which binds to receptor proteins on the post-synaptic neuron. The uptake of the neurotransmitter by the post-synaptic receptors elicits a response in the post-synaptic neuron that can be either excitatory or inhibitory.

Obviously, the brain is an extremely complex structure, so only one type of neurotransmitter would not suffice for the execution of all of its functions. There is a large number of different neurotransmitters, controlling different areas or functions of the brain. Some of the most well known are *dopamine, serotonin, epinephrine*, and *gamma-aminobutyric acid*, while *acetylcholine* was the first neurotransmitter to be identified.

This project will focus on the action of one particular neurotransmitter, namely **dopamine** (DA). It was first identified in the human brain in 1957 [9] and its role as a neurotransmitter was recognized in 1958 by Carls-

son and Hillarp, who won the Nobel Prize in Physiology and Medicine for their findings [10]. Its chemical formula is $C_8H_{11}NO_2$ and a model of the dopamine molecule can be seen in Figure 2.3. Besides being defined as a neurotransmitter, dopamine is also regarded as a neuromodulator without inhibitory or excitatory capabilities [11], and is synthesized in the brain, by neurons originating in the substantia nigra (SN) and ventral tegmental area (VTA) and terminating in the striatum [2]. There are also projections of dopaminergic pathways to other brain structures, such as the cerebral cortex, the amygdala and the hippocampus [12].
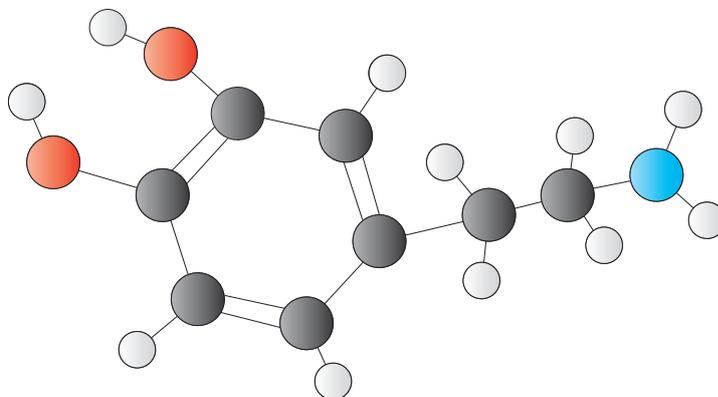
Figure 2.3: The Dopamine molecule. (Atom Coloring: Carbon - Black, Oxygen - Red, Nitrogen - Blue, Hydrogen - White)

As is the case with all neurotransmitters, dopamine does not float freely in the brain, but follows very specific paths within it, called *dopaminergic pathways*, each of which controls a different behavioral process, such as mood, cognition and reward. Dopamine deficiencies can lead to diseases like Parkinson's disease, while cognitive disorders like schizophrenia are a result of a dopamine surplus [3]. The main dopaminergic pathways are shown in Figure 2.4.
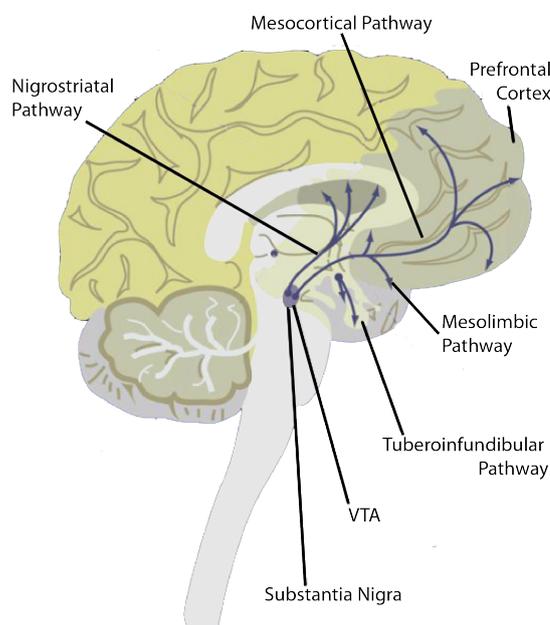
Figure 2.4: The main dopaminergic pathways. Image modified from [3]

Later on, the effect of electrical stimulation of the Ventral Tegmental Area on the electrical activity of the Prefrontal Cortex will be studied (See Chapter 5).

2.2. THE ELECTRICAL ACTIVITY OF THE BRAIN

The study of the electrical activity of nerves has evolved dramatically since the empirical experiments conducted by Galvani in the 18th century, thanks to whom the electrical basis for the function of neural tissue was established. Existing modern methods accurately measure the electrical potentials created by cells, enabling the in-depth study of their behavior.

2.2.1. GRADED AND ACTION POTENTIALS OF NEURONS

Every cell in an organism is surrounded by a double semi-permeable membrane, the *cell membrane*. A resting potential of typically around $-90mV$ is sustained across the membrane by ionic concentrations (mainly of potassium and sodium ions), which are, in turn, sustained by the action of ionic pumps peppered across the surface of the membrane. This ionic equilibrium and the voltages created by various ionic concentrations on either side of the membrane are described by the Nerst Equation [4] which is, however, outside the scope of this work.

Usually, the transmembrane voltage does not deviate from the resting potential. A stimulus, however, can locally alter the ionic configuration across the membrane or affect the function of the ionic pumps, creating a localized change in the transmembrane potential, which propagates along the membrane and usually degrades soon after the onset of the stimulus, returning the voltage to its resting value. The voltage created during this process is called a *graded potential* and is proportional to the strength of the stimulus.

Some cells, however, such as nerve cells, are *excitable*. This means that when their transmembrane potential surpasses a certain *threshold*, the propagation of the stimulus becomes self-sustaining, creating a different form of response, known as an *action potential*. Action potentials follow the *all-or-nothing* principle, which states that whenever a cell is properly stimulated, that is to say when the voltage threshold is surpassed, the amplitude of the elicited action potential will be exactly the same, regardless of the strength of the stimulus. In nerve cells, action potentials propagate from the soma along the length of the axon toward the pre-synaptic terminal. Frequently, action potentials of nerve cells are called *nerve impulses*.

Besides exciting the cell, in which case they are called *excitatory*, stimuli can also be *inhibitory* and suppress the creation of action potentials. In the latter case, the amount of excitatory stimulus required for the activation of the cell is increased, making the creation of action potentials more difficult [1, 4]. Figure 2.5 displays this behavior in detail.
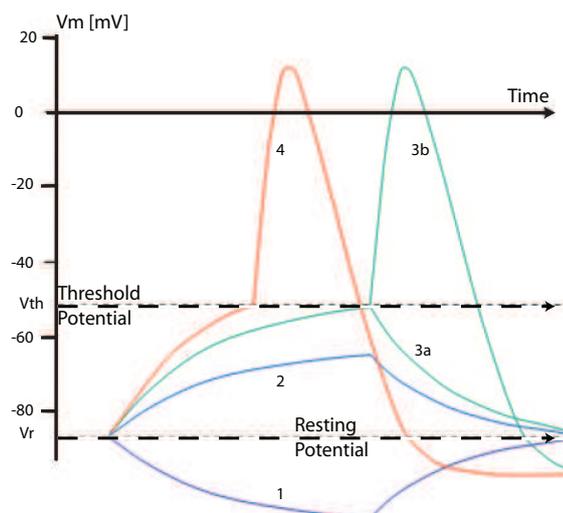
Figure 2.5: Various stimuli of a cell. (1) is an inhibitory stimulus, moving the transmembrane voltage towards more negative values and away from the threshold. (2) is a graded potential. In (3), the threshold is marginally reached, which is why it could either degrade (3a) or elicit an action potential (3b). (4) is an action potential. Image courtesy of [4]

The exact mechanism of ionic transportation across the membrane and propagation of the action potentials are outside the scope of this work, but [4] provides important insight to the nature of all bioelectrical mechanisms. It is noted that this type of electrical activity is ubiquitous within the CNS. The constant movement of ions through the cell membranes of the CNS neurons creates an electromagnetic field in and around the brain, albeit weak, which can be detected and recorded with modern equipment.

2.2.2. RECORDING METHODS

There is a number of methods that are used for the harvesting of neural signals. Examples include the Electroencephalogram (EEG), the Electrocorticogram (ECoG) and the Magnetoencephalogram (MEG). There is also the option of probing deep into the brain with specially made electrodes and measure *Local Field Potentials* or the activity of single neurons. An overview of these methods is displayed in Table 2.1.

Table 2.1: Neural recording methods

Method	Invasiveness	Level of Detail	Equipment
EEG	Non-invasive	Low	Scalp electrodes
MEG	Non-invasive	Low	Magnetic field sensor
ECoG	Invasive	High	Cortical electrodes
Single-Cell or LFP	Invasive	Very High	Glass pipette or metal probe

The EEG is recorded with an array of scalp electrodes, most often arranged in the standardized 10/20 configuration (see Figure 2.6b and [13]). Usually, an EEG headset houses the electrodes. An example is given in Figure 2.6a. The amplitude of the EEG signal is relatively low, while its frequency content is limited, due to the low-pass effect of the skull and soft tissue surrounding it. At any rate, it offers a very good trade-off between invasiveness and signal quality, which is why it is widely used in research and clinical applications.

(a) The EPOC+ EEG Headset. Image courtesy of EMOTIV at https://www.emotiv.com/product/emotiv-epoc-14-channel-mobile-eeg/

(b) The 10/20 configuration. (Modified combinational nomenclature) Image courtesy of [13]
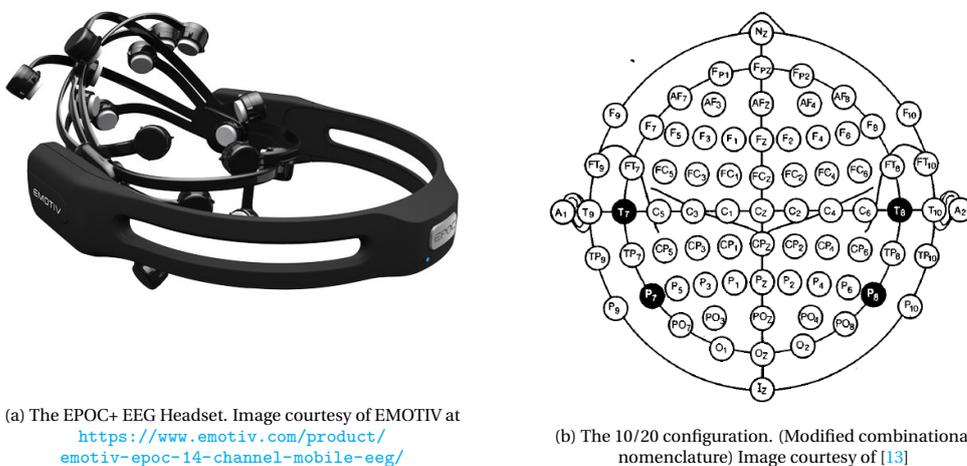
Figure 2.6: EEG headset and electrode configuration

The MEG is measured using an array of extremely sensitive magnetometers, SQUID being the most popular solution [14], which detects the overall neural activity of the brain indirectly, by measuring the magnetic field created by the moving charges, typically between 10 and $1,000\,fT$ [15]. Its advantage over the EEG is that the magnetic field is not as affected by bone and tissue as the EEG. The MEG apparatus and sensors are shown in Figure 2.7

In order to avoid the distortion created by organic structures between the neurons and the recording electrodes, *subdural* electrodes are placed directly on the cerebral cortex for the recording of the ECoG. The procedure is very invasive, as a part of the skull has to be removed, so that the electrodes can be placed onto the

(a) A Magnetoencephalography machine. Image courtesy of I-LABS MEG Brain Imaging Center at http://ilabs.washington.edu/meg-brain-imaging-i-labs

(b) A schematic of a MEG helmet, including a current delivery module for transcranial stimulation. Modified from [16]
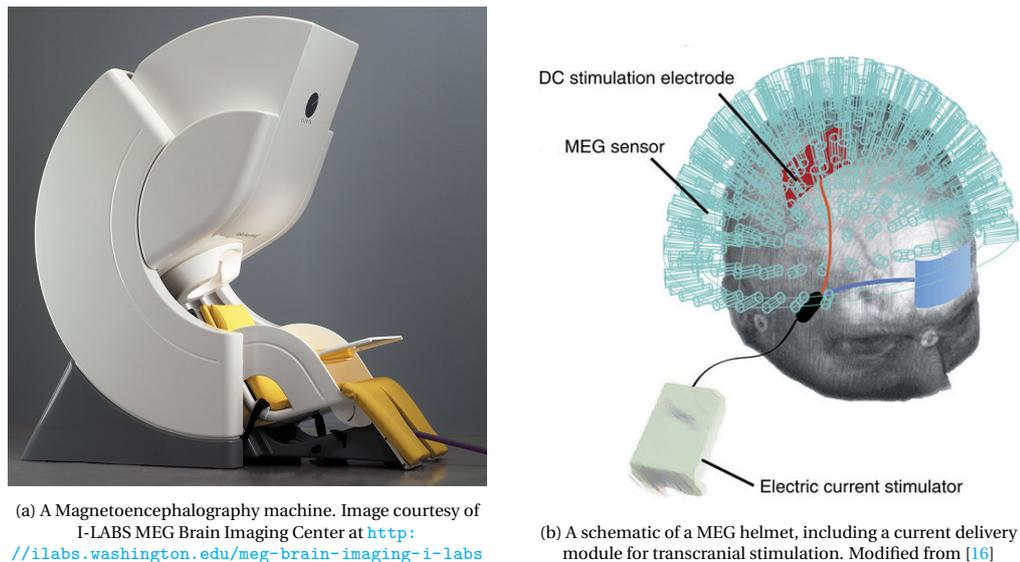
Figure 2.7: MEG equipment
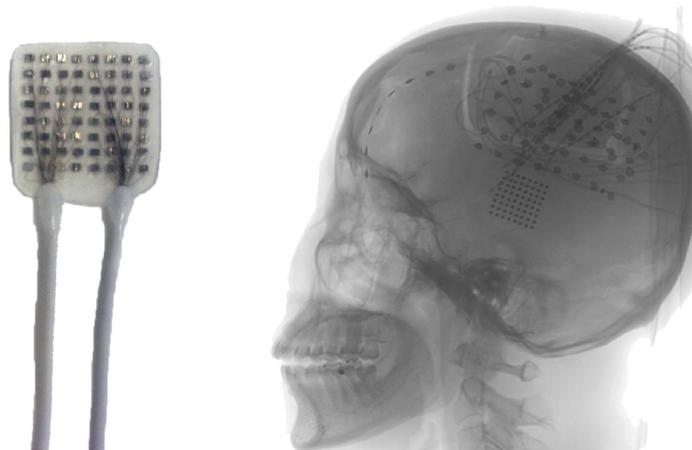
Figure 2.8: An ECoG electrode on the left.. An array of these electrodes was placed on a patient's brain and the result is visible in the X-Ray on the right of the image. Image modified from [5].

surface of the brain itself. The advantage, however, is much clearer and more detailed signals. Moreover, the ECoG has a higher spatial resolution than the EEG [17], which can be enhanced even further with the use of a thick array of grid or strip electrodes [15]. The materials that are most commonly used for the manufacturing of these electrodes are stainless steel and Platinum-Iridium alloys. An ECoG electrode and an array placement on a patient's brain is shown in Figure 2.8

2.2.3. LOCAL FIELD POTENTIALS

The three methods mentioned above record the collective activity of a large number of neurons. In some applications, it is desirable, however, to record the activity of a very specific group of cells, or even one single neuron. In this case, it is necessary to record the *Local Field Potential* created by these cells, using a very thin probe (made usually of Silicon or a Platinum-Iridium alloy), which is inserted into the brain through a hole in the skull (*cranial window*). This is also a quite invasive procedure, but yields very detailed recordings of the activity of specific cells, which can be used in both research and clinical environments. An example of neural probes are given in Figure 2.9.

The potential in the extracellular space can be recorded with the use of special electrodes which are positioned among a small number of neurons and detect the variations in potential caused by the activity of
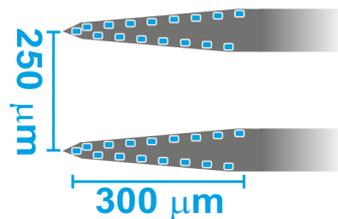
Figure 2.9: Example of a 32-channel neural probe. Image modified from Cambridge NeuroTech at
https://www.cambridgeneurotech.com/silicon-probes

these neurons, commonly known as the Local Field Potential.

The electrodes are placed very close to the neurons and one would expect the signal to comprise of the superposition of a number of pulse trains (like the ones described in 2.2.4). That is not the case, however, because the extracellular medium displays low pass filtering behavior, effectively canceling out the frequencies that make up the spikes that are otherwise visible in single-neuron recordings. Thus, any frequencies over $\approx 100 Hz$ are steeply attenuated [15, 18].

Despite the lack of information beyond $100 Hz$, it is common practice to low-pass filter the LFP recordings at cutoff frequencies between 100 and $300 Hz$ to remove any high frequency noise. The use of a high pass filter is also encouraged for the removal of transient low frequency changes and possible DC offsets [19–24]. Finally, the authors of [25] claim that useful information in the LFP can be found in frequencies up to $500 Hz$, in which case the signal must be filtered accordingly.

There are also indications of the possibility for the detection of rhythmic behavior (see 2.3) in LFP recordings, which, itself can be a powerful diagnostic tool [20, 23].

2.2.4. SINGLE NEURON ACTIVITY

The activity mentioned above represent the cumulative activity of a large number of neurons, *synchronized* in such a way that it manifests itself as almost periodic waveforms. (see below). If one delves deeper into these signals, it becomes apparent that the activity of each individual neuron displays rather different characteristics.

As expected, the recording of a single neuron's activity is, at first sight, little more than a series of APs. However, depending on the function of each neuron, the APs differ in form and pattern. For example, two distinct firing patterns have been recognized in neurons belonging in dopaminergic pathways. These neurons fire either slowly and irregularly, with single or double spikes, or in *bursts* of many APs [26, 27].

The DA neuron APs usually display a long duration and a prominent negative phase. A notch can sometimes be detected in the rising edge of the AP [28]. The spontaneous activity of the neurons is characterized by the appearance of irregularly-firing spikes with a frequency of $3 - 8 Hz$. The cells have a slow depolarization phase, around $78 ms$ until the transmembrane voltage reaches the relatively high AP threshold of around $-42 mV$. It is assumed that this long depolarization is what causes the large duration of inter-spike intervals between spontaneous spikes. Afterhyperpolarization of the cells is measured at $20 ms$ after a single spike and $50 ms$ after a spike burst.

Spike bursts in DA neurons are defined as a series of 2 to 10 spikes whose ISI is smaller than $80 ms$, followed by a quiet period of at least $160 ms$. It has been observed that the amplitude of the spikes decreases as the burst progresses, while the ISI increases. [26, 27]

2.2.5. EVENT-RELATED POTENTIALS

Contrary to the brain's rhythmic activity, which can be characterized as strictly intrinsic, macroscopically the brain's response to external stimuli comes in the form of *event related potentials* or *evoked potentials*.

From this point on, the term *event related potentials* will be used, as it is the most widely accepted one [29]. Depending on the nature of the stimulus, ERPs can be categorized in [7]:

- Auditory

- Somatosensory

- Visual

- Cognition-related

Each type of ERP has its own temporal and spectral features, however these signals are characterized by their very low amplitude, which makes them difficult to distinguish within the electrical signals produced by nearby brain areas, as well as a quite low Signal-to-Noise Ratio. Nevertheless, there are ways to extricate these weak signals and study them independently (see C.2).

2.3. NEURAL SYNCHRONY AND BRAIN RHYTHMS

Under normal conditions, a neural recording will look like a low-amplitude noisy signal to the unaided eye. Significant diagnostic features are hidden in the signal, though, and can be extracted after suitable analysis. There are, of course, some features that are visually obvious, even in raw recordings. When visible, event-related potentials (2.2.5) are a prime example of stimulus-locked features, meaning that they are elicited after a specific excitatory stimulus.

The fundamental difference lies in the level of *neural synchrony* during different parts of the recording. When a group of neurons is excited in a synchronous manner, their electrical responses will align in such a way that their superposition will result in a large amplitude deflection, while continuous synchronous excitation will result in an oscillatory behavior in the signal. The best-known manifestation of this behavior comes in the form of *brain rhythms*. Brain rhythms have been shown to depend largely on the physical and mental state of the subject [1], while it has been proposed that they should be studied in relation to the *neuronal correlates of consciousness*. The NCC are defined in [8] as "the minimal set of neuronal events that are necessary and/or sufficient for perceptual awareness". The main and most widely known brain rhythms are [1, 7, 30]:

- Alpha waves

- Beta waves

- Gamma waves

- Delta waves

- Theta waves

Alpha waves consist of frequencies between 8 and $13Hz$ and have a low amplitude. They are the dominant rhythm of a brain in an awake but resting state, with the eyes closed. Alpha waves are suppressed when the eyes are open and receive visual stimuli and are typically replaced by beta waves.

Beta waves move in the 14 to $30Hz$ frequency band. They are elicited when the brain is alert or busy with a mental task.

Gamma waves are activity in the 30 to $80Hz$ range (or up to $120Hz$ according to some literature [31]). It is assumed that gamma waves are produced as a response to stimuli. Also, they have been correlated to short-term memory functions.

Delta waves are the slowest waves produced, at frequencies of $4Hz$ or lower and are characterized by high amplitude, indicating a high collective synchrony of the brain. They are usually present during sleep or anaesthesia, while their shape can be used as a marker of brain damage or neurological disorder.

Theta waves are irregular waves between 6 and $10Hz$ [31]. They appear mostly in children's EEG but also during concentration in adults. Theta rhythms have also been detected near the hippocampus and are thought

to be related to memory functions.

The basic brain rhythms are shown in Figure 2.10.
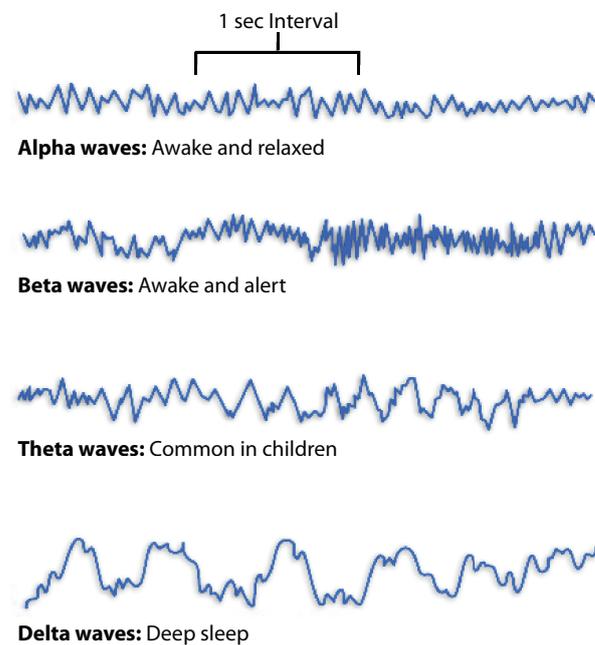
Figure 2.10: The four basic brain rhythms. Modified from [1]

It should be noted that rhythmic activity is more prevalent in surface recordings, such as EEG or MEG, because they record the average activity of a large population of neurons. This property has allowed researchers to come up with further evidence of the correlation between different rhythms to different mental states and cognitive functions. The amplitude of any oscillatory activity in the brain correlates directly to the degree of underlying synchrony [7, 31].

The brain coordinates its activity by synchronizing the firing patterns of groups of neurons. Macroscopically, this mechanism is observed as neural synchrony. Thus, it is logical to extract two conclusions: first, that specific brain functions and different brain areas produce different rhythmic activity and, second, that abnormal rhythmic behavior is indicative of some form of brain damage or disease [8, 31]. An interesting explanation for the superposition of different frequency bands in a neural signal at any given moment is that the brain performs multiple functions by frequency multiplexing of these oscillations [32].

The neural recordings that were analyzed for the following parts of this project were obtained from the Prefrontal Cortex (or neocortex) of mice (see Fig. B.1). The PFC plays a very important role in sensorimotor integration, managing representation, planning and execution of actions at the highest cognitive level [31]. Neural signal oscillations in the PFC are thought to be produced by the communication of the PFC with other brain areas, such as the visual cortex (producing waves in the beta and gamma bands) and the hippocampus (producing theta waves). Gamma and beta waves are related to attention, while theta waves are largely attributed to hippocampal activity and memory management. Reward signals are also mediated between the prefrontal cortex and deep brain structures, during learning processes [33].

The detection and study of neural synchrony is also very useful from a diagnostic standpoint, as a number of CNS dysfunctions display an abnormal oscillatory footprint. Erroneous function of neuronal groups generates abnormal synchrony patterns in diseases characterized by cognitive deficits, such as schizophrenia, epilepsy, autism and Alzheimer's, as well as other neurodegenerative disorders, for instance Parkinson's disease [8].

A summary of the most common neurological imbalances that cause abnormal oscillations in neural record-

ings is shown in Table 2.2.

Table 2.2: Typical neurological disorders creating abnormal neural synchrony. (Modified from [8])

Disorder	Neural Synchrony	Dysfunctions
Schizophrenia	reduction of local- and long-range synchronization	perception, executive process, memory attention, social cognition
Epilepsy	increase in local synchrony, reduction in long-range synchronization	specific cognitive deficits in relationship to seizure focus
Autism	reduced functional connectivity and evidence of impaired neural synchrony	perception, executive functions, social cognition, attention, memory
Alzheimer's	reduced neural synchrony during resting state and evidence of reduced functional connectivity	working memory, perception, attention, executive processes
Parkinson's	increase of neural synchrony in the basal ganglia and between subcortical-cortical structures	motor functioning, perception, working memory, attention, executive functions

As has been mentioned above, neural tissue displays low-pass characteristics. This can explain why high frequency oscillations are much more localized than lower-frequency events that can travel along larger neuronal networks [30].

All the above goes to show that the correct detection and identification of neural synchrony events are of paramount importance both for research and diagnostic purposes.

3

NEURAL SIGNAL ANALYSIS AND STATE OF THE ART

In this chapter, the characteristics of neural signals will be analyzed and some standard methods of biomedical signal analysis will be presented. The methods in this Chapter are the ones that will be used in the analysis (see Chapter 5), while additional methods can be found in Appendix C.

3.1. NEURAL SIGNALS

The ability to record the electrical activity in the Nervous System is not an end in and on itself. It would be counterproductive to simply harvest signals without being able to interpret them in some manner. Research into neural recordings has concluded that there is an overabundance of clinically significant information hidden within the features of these signals. In order to extract this information, the signals must be suitably cleaned and processed, using one or more methods of signal processing, as will be analyzed below.

3.1.1. THE NATURE OF NEURAL SIGNALS

Having acquired the neural signals, the logical next step is to categorize them, based on their characteristics, which will, obviously, impact the selection of the methods used for their analysis. Since it is almost impossible to predict the behavior of these signals over time, it is safe to assume that neural signals are *stochastic*. However, it has also been proposed that neural signals result from a nonlinear dynamic system, which displays deterministic, albeit chaotic behavior, in order to achieve the level of randomness that these signals display. That being said, the deterministic approach yields models so complex that they indeed approach a true stochastic process, so it is much preferred to simply regard neural signals as stochastic as treat them as such.

Since the stochastic approach has been selected, the analysis will make use of the signal's statistics (or *moments*), the most basic of which are the mean, the variance and the standard deviation. Higher order statistics, such as kurtosis and skewness, which, in simple terms, display how different a process is from a simple Gaussian, can also be used. Their use is, however, uncommon, since the signals are usually treated as *wide-sense stationary*, so only their first two moments are considered sufficient to satisfactorily describe the process.

Regarding these signals as stationary is very convenient. That being said, it is easy to observe that they display transient behavior over longer time periods, which actually renders them *non-stationary*. Given this inconvenience, it is usually wise not to study the entire signal as a whole. This means that the signal has to be *segmented* (See C.3) into parts of small duration, during which the signal's statistics remain largely unchanged, in other words, into stationary (or at least quasi-stationary [29]) segments [34]. This not only makes the analysis simpler, but can be used as a powerful tool to discover features in the waveform that could not be identified with visual scrutiny or simple temporal or spectral methods [7].

3.1.2. NOISE AND ARTIFACTS

Neural signals are characterized by extremely low amplitude. It is to be expected that any recording of such signals will suffer from a substantial amount of noise or, to use the more general term, *artifacts*. There is a plethora of artifact sources in neural recordings. A very broad categorization is the distinction into *technical* and *physiological* artifacts [7].

Technical artifacts are caused by the recording equipment itself, coming in the form of internal or thermal noise in the amplification and recording devices, noise in the digitization process or even in the form of powerline interference, also known as *mains hum*, at 50 or $60\,Hz$. Physiological artifacts originate in the biological structures of the subject being recorded. In the case of neural recordings, the most important artifacts are muscle activity, including respiration, blinking and eye movements (which constitute the Electrooculogram), cardiac activity as well as background EEG in ERP recordings.

Measures must be taken in order to reduce the levels of noise prior to the analysis of the signals, since in most cases the signal of interest is buried under an avalanche of noise and unwanted features, making it practically impossible to isolate [29].

3.2. POPULAR METHODS

After harvesting, neural signals can be subjected to any signal analysis techniques the researcher deems suitable for the application at hand. First and foremost, noise removal techniques must be applied on the signal, that will remove both random noise and physiological artifacts (see 3.1.2) produced by areas of the brain that are irrelevant to the analysis. Once a clearer signal has been acquired, different methods can be used to extract any characteristics the signal may display.

3.2.1. ARTIFACT AND NOISE REMOVAL

The first measure against noise in the recordings is the selection of suitable equipment, which meets specifications that support low-noise applications, or even modified to meet low-noise criteria [35]. Also, the experiment must be conducted in a relatively noise-free and quiet environment. In many cases, electromagnetic interference is reduced by encasing the recording devices and the subject of the experiment in a Faraday cage [29].

RANDOM NOISE AND FILTERING

In [7], both the need for artifact removal and various guidelines are mentioned. One of the obvious tools for high frequency noise reduction is, of course, low pass filtering. However, this must be used in moderation, so that clinically significant brain activity is not lost. In the case of neural signals, information that could be obtained from spikes, which are spread on a wide bandwidth, can become irretrievable due to the careless application of a low pass filter. High pass filters are routinely used [36–40], for the elimination of any slow transient changes in the signal, as well as the elimination of possible DC offsets in the recording.

All the above amount to a process with band pass properties. Besides DC factor removal, the high pass action at very low frequencies also removes any slow or infraslow artifacts that are caused by the activity of non-neuronal glial cells [15]. A high cutoff frequency for the low pass filter might be necessary for the optimal representation of neuronal spikes, which are characterized by a large bandwidth. Various band pass filters used in literature are shown in Table 3.1. It is important to note that it is not unusual for a wider band pass filter to be used during the actual recording of the signal, while a more focused one is used during the processing of the signal.

In [29], the author states that in essence, the methods suitable for noise and interference removal are averaging (see C.2), a moving average filter in the time domain, filtering in the frequency domain, the design of an optimal Wiener filter and adaptive filtering.

MAINS HUM

The interference from the electrical network introduces a 50 or $60\,Hz$ oscillation in the recordings. Measures of electromagnetic shielding, like Faraday cages, can be used to prevent this, but usually a notch filter at either of these frequencies is eventually applied on the signals [7, 17] during or after harvesting. Removal of the mains hum is, arguably, much easier than that of physiological artifacts, as the signal to be removed

Table 3.1: Band pass filters used in literature

Reference	Low Cutoff	High Cutoff	Signal
Foti et al. [39]	$0.1Hz$	$30Hz$	
Santesso et al. [37]	$1Hz$	$30Hz$	
Holroyd et al. [41]	$1Hz$	$10Hz$	EEG
Holroyd et al. [42]	$0.1Hz$	$20Hz$	
Bokura et al. [43]	$0.1Hz$	$50Hz$	
Lehmkuhle et al. [36]	$0.3Hz$	$300Hz$	ECoG
Charlety et al. [44]	$0.3Hz$	$10kHz$	Single Neuron
Destexhe et al. [23]	$300Hz$	$5kHz$	LFP
Donoghue et al. [22]	$10Hz$	$100Hz$	LFP
	$0.3Hz$	$3kHz$	Neural Discharge
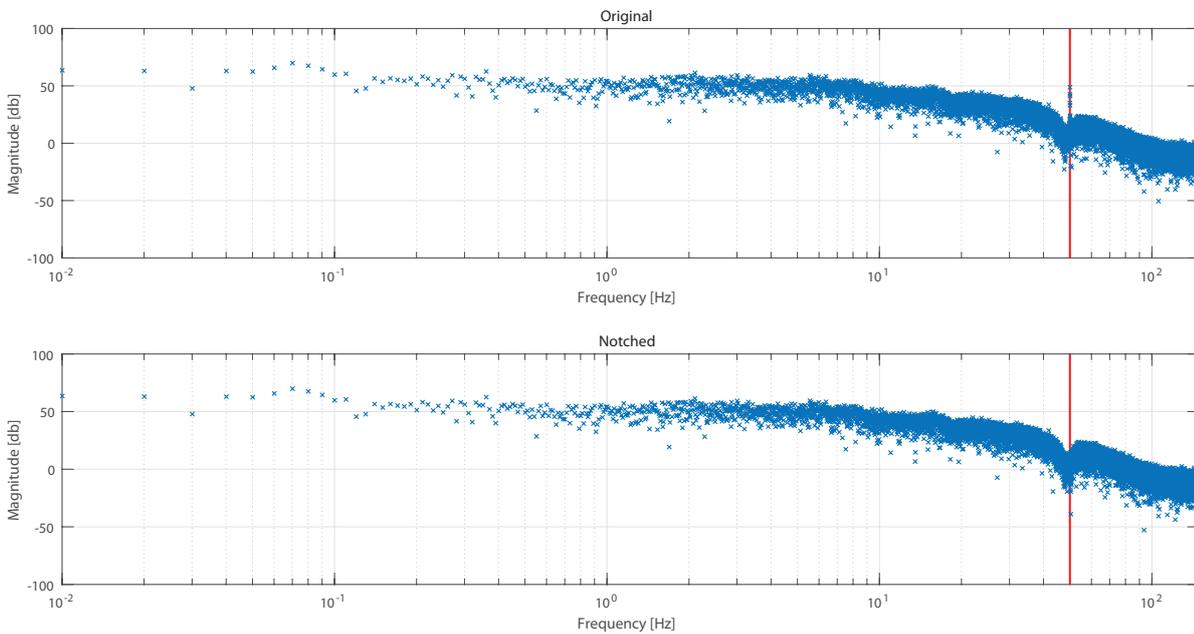
Figure 3.1: The effect of a notch filter to remove a $50Hz$ interference, highlighted with a red line. The spectrum of the original signal is shown in the top graph, the notched in the bottom graph

is perfectly known a priori, allowing for the optimal design of a notch filter, which will sufficiently suppress these specific frequencies without damaging the rest of the valuable signal. The effects of a notch filter on the power spectrum contaminated by a $50Hz$ hum are displayed in Figure 3.1. Note that in this recording, a notch filter has already been applied, possibly during the harvesting of the signals, however, this notch was not correctly centered on $50Hz$, allowing the existence of a peak in the spectrum.

OCULAR ACTIVITY

Another artifact in neural recordings is ocular activity, wherein eye movements (constituting the EOG) and blinks create electrical signals infecting the neural recordings. These are both caused by charge separation due to friction of the retina and cornea with surrounding tissues. Eye blinks are detected in the waveforms as positive peaks lasting some tenths of a second, typically containing frequencies up to 8 to $13Hz$ [7] but of negligible amplitude, that does not significantly interfere with the alpha band in the same frequencies [17]. The EOG consists of lower frequencies, typically around $6-7Hz$, and affects the recordings to a higher degree.

Clearly, the EOG artifact has to be dealt with before the analysis is initiated. Many methodologies exist for EOG removal. Sörnmo in [7] states that removal requires a reference signal, in other words, the EOG must be recorded simultaneously with the neural signals, so that it can later be identified and removed from the

neural recording, a method which is also used in [43]. Also in [7], it is suggested that an adaptive algorithm for EOG rejection would be more appropriate, proposing a method based on the Least Means Squares Algorithm, while Independent Component Analysis (see C.5) is used for artifact removal in [37] and [45].

However, despite the need for additional equipment, it seems to be much more useful to record EOG when possible. In a very popular publication by Gratton et al. [46], a method for EOG removal is proposed, which is based on actual EOG signals harvested by the subject during the experiment. This algorithm has been used and referred to extensively in literature (including [7, 38, 39, 39, 41, 42]).

OUTLIER REMOVAL
Another useful action in the preprocessing of the signal is the removal of the so-called *outlier values*, wherein by visual or other scrutiny, values that are deemed exorbitantly out of trend are manually removed from the recording and not taken into consideration in further analysis. An extremely high or low amplitude seems to be the only suitable criterion for the removal of such values in the case of neural recordings, as a sudden change in frequency content could be of significant clinical importance and should be analyzed in greater detail.

Similar to the selection of the band pass filter range (See 3.2.1.1), the selection of the threshold for the removal of outlier values is up to the researcher's judgment, but, also, depends on the nature of the signal itself as well as the recording method. In [47], the threshold is set at $250mV$, while any values out of the $\pm100mV$ are excluded in [43]. Also, a semi-automated procedure based on custom criteria is mentioned in [39].

In some cases, for instance in single cell recordings, the reverse is true. It is usually assumed that a fast and large deflection in the signal is useful and labeled as an actual neuronal spike if it has an amplitude of at least 3 times that of the surrounding noise [44].

It is important to stress that outlier rejection such as this can reduce the amount of usable data dramatically if the latter is limited or the artifacts too frequent [45].

3.2.2. SAMPLING CONSIDERATIONS
Signals harvested from the brain carry a great amount of information. It is natural to assume that a high sampling frequency is necessary to make sure that high frequency features are correctly sampled and no information is lost or corrupted due to aliasing. Indeed, the signals are routinely harvested at relatively high sampling rates, usually from over $1kHz$ up to even 32. However, the amount of data produced by these experiments, especially when so high sampling rates are used, is enormous, which can make storing, parsing and analyzing the data very computationally challenging.

Used correctly, downsampling (or subsampling) of the signals can be very beneficial as far as algorithm complexity and speed are concerned. Aliasing can be avoided by downsampling in moderation, so that the resulting sampling frequency maintains a value larger than the Nyquist frequency corresponding to the fastest feature expected to be found in the signal. Also, low-pass filtering prior to downsampling (decimation) ensures that the Nyquist criterion is met at the new sampling frequency, so that no aliasing is, again, avoided.

The use of downsampling is very common in literature, while the level of downsampling varies greatly, depending also on the nature of the signal at hand (see Table 3.2).

The researchers in the table above use simple downsampling, which, in essence, entails keeping only every Mth sample of the signal, where M is a downsampling factor. Additionally, Ding et al. in [34] use bootstrap resampling to reduce the size of their data.

An example of downsampling in data can be seen in Figure 3.2. The original signal was harvested at a rate of $30kHz$. It is clear from the bottom graph that even after resampling at $1kHz$, the time domain signal retains most of its characteristics. Also, looking at the power spectra on the right of the Figure, it is clear that there is no valuable information lost and the spectrum also retains its characteristics. The only difference is in the overall magnitude, which is decreased as the downsampling factor increases.

Table 3.2: Downsampling examples in literature

Reference	Original f_s	Downsampling f_s
Shimamoto et al. [48]	$1.5 - 3kHz$	$1kHz$
Ray et al. [21]	$4kHz$	$1kHz$
Zhang et al. [49]	$2048Hz$	$512Hz$
Manning et al. [50]	$32kHz$	$2kHz$
Sanders et al. [51]	$1kHz$	$250Hz$
Le van Quen et al. [52]	$10kHz$	$2kHz$
Kühn et al. [20]	$25kHz$	$1250Hz$
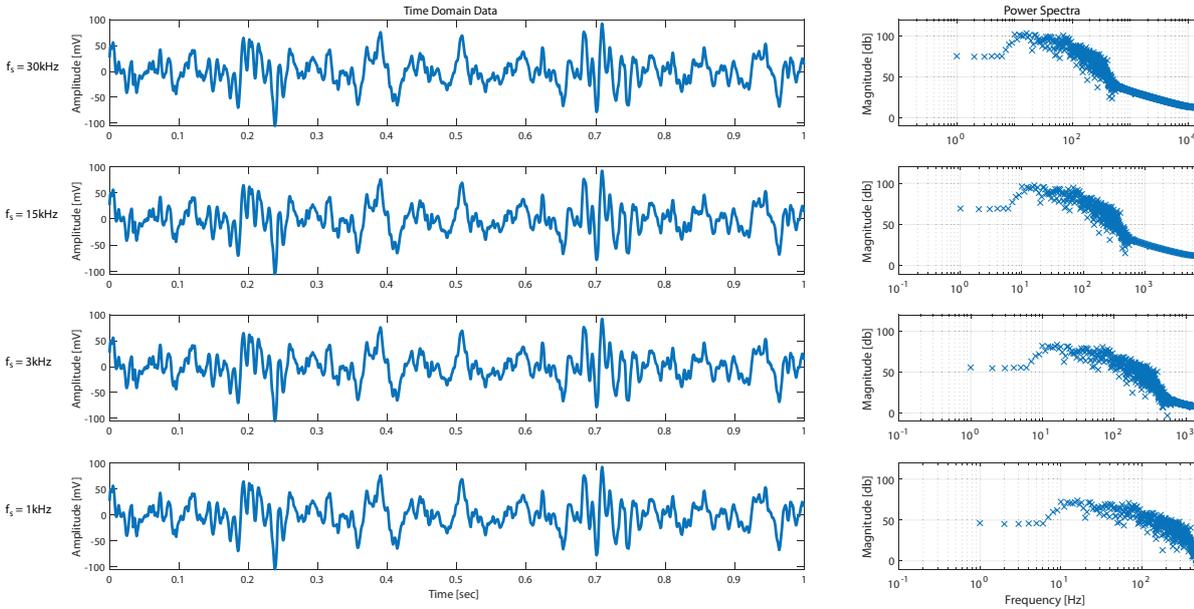
Figure 3.2: Downsampling of an LFP recording harvested from the PFC. Each row of figures displays a time domain signal at different sampling frequencies and its corresponding power spectrum

3.2.3. PARAMETRIC ESTIMATION AND SIGNAL MODELING

Non-parametric analysis is a very common and rapid way of studying stochastic data (see C.1). Despite its relative simplicity, however, nonparametric estimation of the signal is not a panacea. For instance, the Fourier Transform cannot function properly in very short data segments, which are used in many cases of neural signal processing, yielding highly biased spectral estimators [34].

To overcome this disadvantage, parametric stochastic modeling of the recordings can be applied, during which the signal is assumed to be the output of a linear system driven by Gaussian white noise. The parameters of the system are calculated by fitting the model to the real neural recordings, based on a suitable error criterion (most usually the Mean Square Error). Arguably, a well-designed parametric approach can be more accurate in estimating the signal properties than a nonparametric one for short-time recordings, while the parameter estimation itself can be viewed as a feature extraction process.

The model architecture that has gained the most traction in the field of neural signal processing is that of linear *autoregressive modeling* [7]. The general form of an AR model is:

$$x(n) = -\alpha_1 x(n-1) - \alpha_2 x(n-2) - \ldots - \alpha_p x(n-p) + v(n) \tag{3.1}$$

where p is the model order, x is the output signal, v is the input noise driving the system and α_i are the model parameters. The basic version of the AR model in Equation 3.1 can be used to satisfactory results, however it is preferable to use variations of it, such as the time-varying version:

$$x(n) = -\alpha_1(n)x(n-1) - \alpha_2(n)x(n-2) - \ldots - \alpha_p(n)x(n-p) + v(n) \tag{3.2}$$

which is better at handling nonstationary effects, such as slow changes in the EEG rhythms. In case the recording is conducted from a number of channels, a *multivariate* approach is also available:

$$\boldsymbol{x}(n) = -\boldsymbol{A}_1\boldsymbol{x}(n-1) - \boldsymbol{A}_2\boldsymbol{x}(n-2) - \ldots - \boldsymbol{A}_p\boldsymbol{x}(n-p) + \boldsymbol{v}(n) \tag{3.3}$$

where $\boldsymbol{x}$ is a $M \times 1$ vector and $\boldsymbol{A}_i$ are $M \times M$ matrices of temporal and spatial parameters of the model. It is interesting to note that the elements of the $\boldsymbol{A}$ matrices away from the main diagonal represent the degree of correlation between different channels.

Going back to the study of Eq. 3.1, the output signal x can be viewed as the output of a linear filter of order p, to which white Gaussian noise is fed. The transfer function of this filter is:

$$H_p(z) = \frac{1}{A_p(z)} = \frac{1}{1 + a_1 z^{-1} + a_2 z^{-2} + \ldots + a_p z^{-p}} \tag{3.4}$$

Clearly, the function above is defined only by its poles, which is why AR models are called *all-pole* models. The factor $A_p(z)$ is a linear *prediction error filter*. This relation becomes apparent by noticing that the AR model in essence implements an FIR predictor of a signal's value from a linear combination of p past samples:

$$\hat{x}_p(n) = -\alpha_1 x(n-1) - \alpha_2 x(n-2) - \ldots - \alpha_p x(n-p) \tag{3.5}$$

The noise v can be omitted in Eq. 3.5, because it is assumed to be white and stationary, so it would not affect the quality of the prediction in any way. The prediction error in this case would be:

$$\begin{aligned}
\varepsilon_p(n) &= x(n) - \hat{x}_p(n) = \\
&= x(n) + \sum_{k=1}^{p} a_k x(n-k) \overset{3.1}{=\!=} \\
&= -\sum_{k=1}^{p} a_k x(n-k) + v(n) + \sum_{k=1}^{p} a_k x(n-k) \\
&= v(n)
\end{aligned} \tag{3.6}$$

which further supports the omission in Eq. 3.5. Eq. 3.6 can be viewed as the output of a linear time-invariant FIR filter with a transfer function of:

$$A_p(z) = 1 + \alpha_1 z^{-1} + \alpha_2 z^{-2} + \ldots + \alpha_p z^{-p} \tag{3.7}$$

which, in turn, is the transfer function of the linear prediction filter formulated in 3.5. The parameters α_i can be estimated by minimizing the variance of the error ε:

$$\sigma_\varepsilon^2 = E[\varepsilon_p^2(n)] \tag{3.8}$$

By setting $z = e^{j\omega}$ in 3.7, the power spectrum of the AR model is obtained:

$$S_x(e^{j\omega}) = \frac{\sigma_v{}^2}{|A_p(e^{j\omega})|^2} = \frac{\sigma_v{}^2}{\left|1 + \sum_{k=1}^{p} a_k e^{-j\omega k}\right|^2} \tag{3.9}$$

Returning to Equation 3.4, the unit circle on the complex z-plane is a useful tool for the visualization of the system poles [29]. An example is shown in Figure 3.3. Poles can be either real (the red X on the real axis in the Figure), or come in complex pairs of two (the blue Xs in the Figure).

In order for the poles to be stable, they have to lie within or on the unit circle ($|z| \le 1$, highlighted in Figure 3.3). The unit circle itself is a different representation of the frequency domain. The frequency increases in a counterclockwise manner. The DC factor corresponds to $z = 1$, while half the sampling frequency lies at $z = -1$. So the angle of each rhythm corresponds to the frequency that it models. The distance from the unit circle is also important, as the closer a pole lies to the unit circle, the more pronounced that particular frequency is in the time domain signal.
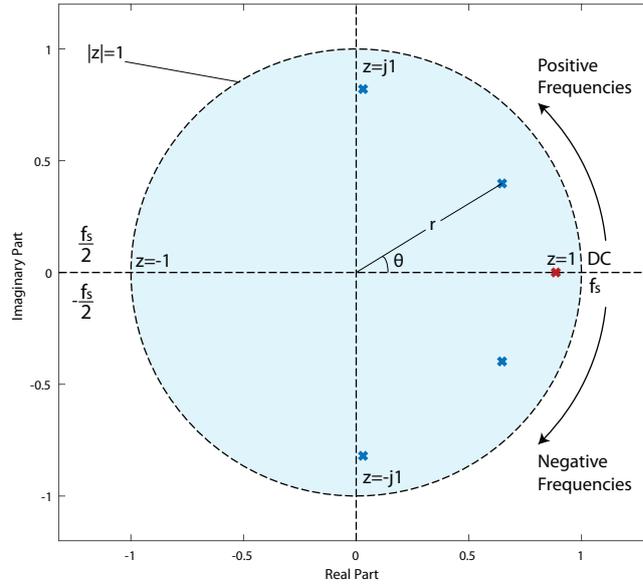
Figure 3.3: Poles of an AR model in the unit circle

Another variation of the AR model is the autoregressive moving average model, which defines the output as the linear combination of p past output samples and q past input samples plus the current input like so:

$$x(n) = -\alpha_1 x(n-1) - \alpha_2 x(n-2) - \ldots - \alpha_p x(n-p) + b_0 v(n) + b_1 v(n-1) + b_q v(n-q) \tag{3.10}$$

where α_i and b_j are fixed parameters and v is the input Gaussian white noise with variance σ_v^2. The corresponding transfer function can be obtained by applying the z-transform on 3.10:

$$H(z) = \frac{B(z)}{A(z)} = \frac{b_0 + b_1 z^{-1} + \ldots b_q z^{-q}}{1 + a_1 z^{-1} + \ldots a_p z^{-p}} \tag{3.11}$$

It is obvious that this is no longer an all-pole model, since the b_n coefficients are nonzero. Its power spectrum is:

$$S_x(e^{j\omega}) = \left| \frac{b_0 + b_1 e^{j\omega} + \ldots b_q e^{j\omega q}}{1 + a_1 e^{j\omega} + \ldots a_p e^{j\omega p}} \right|^2 \sigma_v^2 \tag{3.12}$$

It can be argued that this is the more general formula, since if q is set to zero, it reverts back to a simple AR model, but if p is set to zero, then the transfer function takes the following form:

$$H(z) = \sum_{k=0}^{q} b_q(k) z^{-k} \tag{3.13}$$

which is the transfer function of a *moving average* (all zero) process, the spectrum of which is given by:

$$S_x(e^{j\omega}) = |B_q(e^{j\omega})|^2 \sigma_v^2 \tag{3.14}$$

One of the pivotal decisions in signal modeling is the selection of the model order, which depends entirely upon the goal of the analysis. More detailed analysis (in this context, analysis that includes the modeling of higher frequency features) requires a higher model order. However, a lower model order will be more robust to high frequency noise, while providing satisfactory spectral detail for detection and classification of significant events in the signal. To that end, models of order as low as 8 or 5 have been successfully used in literature [7, 53].

Modeling the expected behavior of the signal is a useful first step in artifact removal, since these features will be derived from the model and stand out in the analysis, making them easy to detect and remove [54]. This procedure can make use of both linear and nonlinear models [55]. The original or altered versions of the methods described above are also used for different reasons across literature, such as modeling ERPs with

the use of Adaptive Multivariate AR Modeling [34] and implementing a *directed transfer function* to model the interaction between EEG signals with MVAR models [49].

As will be made clear further on, it is also possible to extract valuable information from the signal without going through the trouble of training models, which can be very computationally expensive. For example, the ICA method (see C.5) can perform blind source separation and feature extraction on the signal without the necessity for a well defined parametric model [45].

3.2.4. MATCHED FILTERS

It is quite common that a very specific waveform, with its features known *a priori*, is going to be detected in the signal. In this case, it is possible to use *matched filtering* for the detection of known waveforms surrounded by unwanted features and noise. By passing the signal through a matched filter, the SNR of the output is maximized, displaying the desired waveforms in greater relief, as the amplitude of the output increases at the time instances that these events occur.

One can argue that the design of these filters subverts the classic filter design procedure, insofar as, usually, the specifications of filters are expressed in the frequency domain, while in the case of matched filters the starting point is the observation of repeated features in the time domain. So, in very broad terms, the wanted time domain feature, or a similar signal template, is used as the impulse response of the filter. This, of course, means that the output in the time domain is given by the convolution of the original signal with the template used for the matched filter. However, in this context, the same operation can be expressed as the calculation of the cross correlation between the two signals, so, matched filters are also called *correlation filters* [29].

In different fields, such as telecommunication applications, the signal can be, for instance, a predefined series of pulses, facilitating the definition of the suitable matched filter for the detection of this time series. However, in a biomedical context, signals are, as it has been made clear above, non-stationary and are regarded as stochastic. As a result, it is almost impossible to create a template without first examining the signal itself. So, in this case, the template used for the design of the matched filter has to emerge from the signal itself [7], rendering the analysis inherently data-based.

It emerges that the optimal template for the matched filter is a reversed and scaled version of a signal snippet. The proof of this is presented in Appendix D [29].

Matched filters are commonly used for the detection of specific biological signal features, such as brain rhythms, EEG spike-and-wave complexes as well as identification of individual cardiac rhythm waves [29].

4

THE EXPERIMENT AND DATA COLLECTION

4.1. PREPARATION OF ANIMALS

The animals used in the experiments were wild type b-6 mice. The mice were anaesthetized using isoflu-orane gas and their heads were immobilized using a stereotaxic device. Their temperature was kept stable at 37°C. Pre-operative analgaesia included the subcutaneous administration of $0.05ml$ of Rimadyl solution and $0.05ml$ of Temgesic solution. A small amount of Lidocaine was dripped into the open wound before the exposure of the skull and the drilling process.
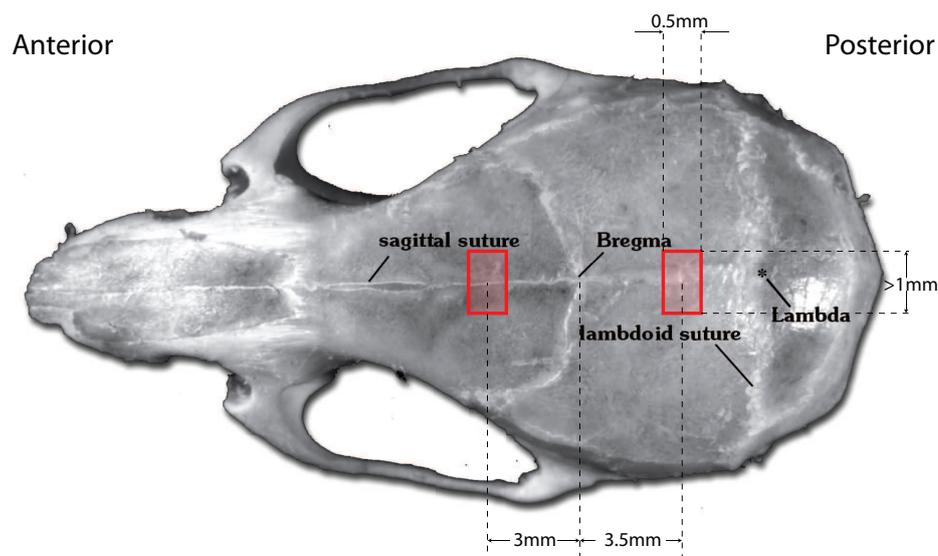
Figure 4.1: The sites of the cranial windows in relation to the bregma. Image modified from [6]

Two cranial windows were opened (see Figure 4.1), one for the introduction of the VTA stimulation electrodes and one for the introduction of the carbon fiber and 32 channel probes. The first is located $3.5mm$ posterior to the bregma and the second $3mm$ anterior. After the craniotomies were complete, a pedestal, used for the fixation of the mouse onto the experimental setup, and a pin for electrical grounding were fixed onto the skull using charisma, creating two watertight compartments. Excess skin was stitched and the mice were removed from the surgery table and were left to resuscitate. Post operative analgesia included the subcutaneous ad-ministration of Rimadyl solution every 24 hours after the surgery.

4.2. EXPERIMENTAL PROCEDURE

The animals were placed on a roller, which provided freedom of movement, however their heads were fixed onto the setup using the pedestal. A reference electrode was connected to the grounding pin. The watertight
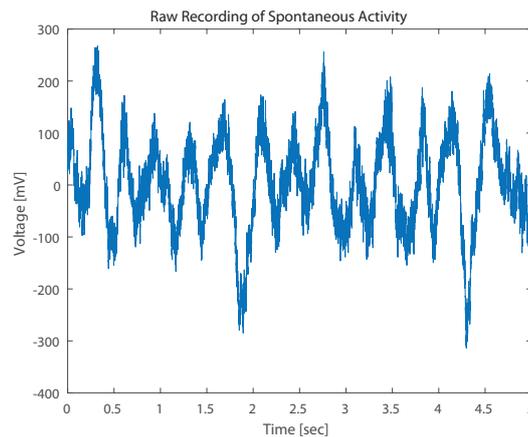
Figure 4.2: Five seconds of one channel of spontaneous activity recording

compartments were filled with saline (NaCl isotonic solution), to promote conductivity to the grounding pin. The electrodes were then introduced into the brain through the cranial windows using an electronic and a hydraulic manipulator. The electrodes were connected to an OpenEPhys FPGA board, which interfaced them with its software on a computer responsible for monitoring and storing the data, which were digitized at $30kHz$.

Unfortunately, the carbon fiber electrode system, which was responsible for the monitoring of dopamine levels in the PFC did not function as expected, so the neurotransmitter marker was not taken into consideration. That being said, it is important to mention that the presence of this particular marker does not affect the analysis itself, but would only be used for the extraction of results related to the correlation of neural activity and dopamine concentration.

SPONTANEOUS ACTIVITY
During the recording of the baseline activity, the stimulation electrodes were kept inactive, allowing for the spontaneous firing of cortical neurons, naturally occurring in the brain.

Spontaneous cortical activity was recorded so that is can be studied independently as well as in comparison to the neural activity during stimulation of the VTA. As will be analyzed later on, significant neural synchrony events were observed in these recordings and a method was developed for their detection.

An example of the raw signal is seen in Figure 4.2. The recording is clearly contaminated by a low-frequency artifact, consistent with respiratory muscle activity [56]. The respiratory EMG could not be monitored in this particular experiment, so measures have to be taken to remove this and other artifacts prior to further analysis of the recordings.

STIMULATION
Various VTA stimulation paradigms were considered for the elicitation of responses in the PFC. It has been shown [57] that visible responses are more likely to appear when the stimulation pulses mimic the firing behavior of the dopaminergic neurons themselves. Specifically, three types of pulses were used:

- biphasic pulses of $0.5ms$ duration and amplitude of $0.1 - 1mA$ every 10 seconds

- trains of 5 biphasic pulses at $20Hz$

- trains of 10 biphasic pulses at $5Hz$

The stimulation electrode comprises of two very thin electrodes, each responsible for delivering pulses to the left and right VTA.

The detection algorithm developed in Chapter 5 will be applied to these stimulation recordings, in order to specify if there is a change in the behavior of the signal upon stimulation.

<div align="right">

5

</div>

ANALYSIS AND RESULTS

5.1. DESCRIPTION OF SIGNALS

As has been stated before, the signals analyzed in this project were Local Field Potentials harvested from neurons located in the Prefrontal Cortex, more specifically in the mesocortical dopaminergic pathway. Two kinds of signals were recorded. The first is what is assumed to be a baseline signal, where an idle brain was being monitored, without the presence of external stimuli. The second is the recording of the responses of the PFC during and after the electrical stimulation of the VTA. Both these recordings are rich in diagnostic information, and the analysis of both makes use of different approaches and methods, as will be made clear in this chapter.

The signals harvested with the OpenEPhys software were 32-channel LFPs, amplified with an Intan Technologies amplifier chip digitized at $30kHz$. When electrical stimulation was introduced into the brain, the timing, shape and amplitude of the pulses was simultaneously recorded in an additional analog channel. The latter signal was, obviously, not subjected to the same analysis as the LFP recordings, but is used as a marker for the verification and quality of the detection results.

In total, three LFP signals were isolated from different recordings and were used in the analysis that follows. These were:

- **Signal 1:** A baseline signal, where the deep brain was not stimulated and displays the spontaneous activity of the PFC

- **Signal 2:** A signal displaying the activity of the PFC while the VTA is being stimulated with large amplitude pulses. 11 pulses are contained in this recording

- **Signal 3:** A signal displaying the activity of the PFC while the VTA is being stimulated with small amplitude pulses. 14 pulses are contained in this recording

As has been stated above, the stimulation pulses follow the paradigm put forward in [57].

Oscillations were marked in Signal 1 and 2 manually, to assist in the design and validation of their own detection scheme, as will be made clear later on. Signal 1 is used for the AR modeling and parametric analysis steps and Signal 2 is used for validation. Signals 2 and 3 are used for the evaluation of the stimulation effect on the electrical activity of the PFC.

5.2. ANALYSIS OF SPONTANEOUS ACTIVITY

The fact that the brain receives no regulated input during the recording does not mean it remains inactive. On the contrary, the brain is very much active, with its neurons firing irregularly. An example can be seen in Figure 5.1. However, every so often, the group of neurons close to the recording electrodes receive input from different parts of the brain and fire in a coordinated manner. These events are deemed, in a number of cases, diagnostically significant and potentially contain information on how a part of the neural circuitry operates, so being able to detect such events in a neural signal is a very useful tool.
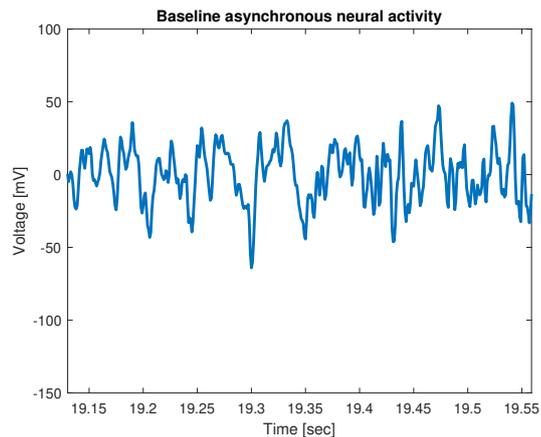
Figure 5.1: Asynchronous neural activity

5.2.1. NEURAL SYNCHRONY

This coordinated firing of a group of neurons is what we call *neural synchrony*, a phenomenon best exemplified by the study of delta brain rhythms (See Chapter 2.3). Visual scrutiny of an event like this will reveal short parts of the signal that are characterized by lower frequency content and higher amplitude.

A simple explanation of this behavior is that irregular activity is more spread out in frequency and has a lower amplitude because it results from the superposition of the electrical effects of action potentials in the extracellular space which, however, fire at different times, and most of which cancel each other out to a certain extent. However, when a group of neurons fires synchronously, their action potentials are aligned and are superimposed into a larger, slower wave.

Thus defined, neural synchrony can be a very useful marker of various diagnostic information concerning the CNS, including the identification of delta rhythms, the presence of which in awake patients is indicative of brain damage, and also the detection of certain types of epileptic seizures [58]. Following a similar line of reasoning, it is safe to assume that the detection of neural synchrony events in the recordings analyzed are indicative of the momentary activation of the mesocortical dopaminergic pathway, which can be considered the *spontaneous activity* of this particular brain area, insofar as it does not constitute a response to artificial and controlled external stimuli. In this section, the methods used for the detection of these oscillatory events will be elaborated on.

5.2.2. ANALYSIS AND DETECTION SCHEMES

At this point, it is useful to make an important distinction. The analysis of the neural recordings will provide valuable information about the features of the signals and potentially useful diagnostic information that might be contained within them. When these features are specified and quantified, a detection scheme will be formulated for their identification.

From a computational point of view, this means that the detection scheme will, hopefully, be able to function with much less demanding processes, using standardized information about the signals extracted and selected from a very detailed analysis.

5.2.3. PERFORMANCE METRICS

The quality of any detection algorithm is quantified with the use of specific performance metrics. The calculation of the performance metrics relies on the number of correct and incorrect detections. Specifically:

- **True Positives** are correct detections

- **False Positives** are cases of overdetection, meaning that the algorithm marked something that was not there

- **False Negatives** are produced when the algorithm fails to detect a desired pattern

- **True Negatives** are cases wherein the algorithm does not react because there is indeed nothing to detect

These counters are usually visualized in the *Confusion Matrix*. The performance metrics are calculated as different ratios among these four. The most commonly used are *accuracy, precision* and *recall*, defined as:

$$Accuracy = \frac{TP}{TP + FP + TN + FN} \tag{5.1}$$

$$Precision = \frac{TP}{TP + FP} \tag{5.2}$$

$$Recall = \frac{TP}{TP + FN} \tag{5.3}$$

$$Specificity = \frac{TN}{TN + FP} \tag{5.4}$$

Accuracy is a measure of how many correct predictions were made overall. Precision expresses the percentage of positives that were actually correct, while recall is the percentage of the patterns that were correctly identified as positives. Specificity can be viewed as the opposite of recall, defining the percentage of negatives that the algorithm correctly identified as such.

In some cases, the use of only one of these metrics, however, results in a biased evaluation of the algorithm quality. Different metrics can be defined that provide a more objective evaluation, such as the F-measure, or F-score, defined as the harmonic mean of the precision and recall values:

$$F-score = 2 \times \frac{Precision \times Recall}{Precision + Recall} \tag{5.5}$$

What is more, another metric is routinely used [59] for the evaluation of detection algorithms like the one developed in this project. This metric is the Average Detection Rate, defined as the mean of specificity and recall:

$$ADR = \frac{Specificity + Recall}{2} \tag{5.6}$$

Finally, a visual representation of the algorithm's performance is the usage of the ROC (Receiver Operating Characteristic) space. For the creation of the ROC graph, the True and False Positive Rates are used. Respectively, these correspond to recall and 1-specificity. The basic ROC space is given in Figure 5.2.
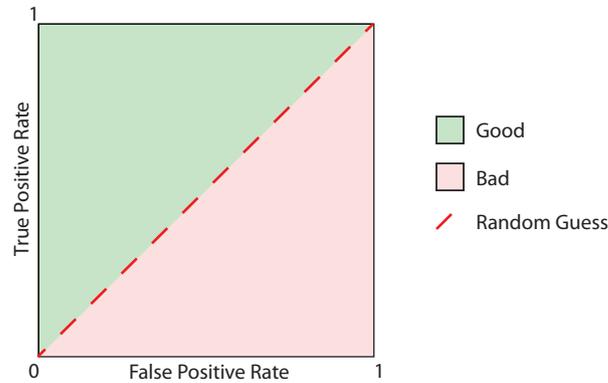
Figure 5.2: The ROC space

As can be seen in the Figure, the upper and lower parts of the ROC space are separated by the *line of no-discrimination*, which represents, theoretically, the TPR and FPR values when the detection is random. The points in the upper (green) part correspond to a good detection scheme, while the points in the lower (red) part, to a bad one. The $(0,1)$ point, which corresponds to $TPR = 1$ and $FPR = 0$ is considered *perfect classification*. The goal is that an algorithm scores as close to this point as possible on the ROC space.

The constituents of the parametric analysis, described later on, will be evaluated based on the aforementioned metrics. The best set of parameters will, eventually, be used on a different recording for validation purposes (see 5.3).

5.2.4. PREPROCESSING

As always, the first step in the analysis of these signals preprocessing. Signal 1 was used in this step. In this case, the problems the signal presents are the presence of a certain number of artifacts and the huge amount of samples due to the high sampling frequency, slowing the analysis down. Downsampling is definitely going to take place, however, in order to avoid additional artifacts due to aliasing, it will be performed last. In Figure 5.3 a part of the raw recording and its power spectrum are displayed. In this case, the power spectrum has been calculated by windowing the waveform and averaging the resulting spectra of the windows. The windows were 0.5 seconds long and non-overlapping.
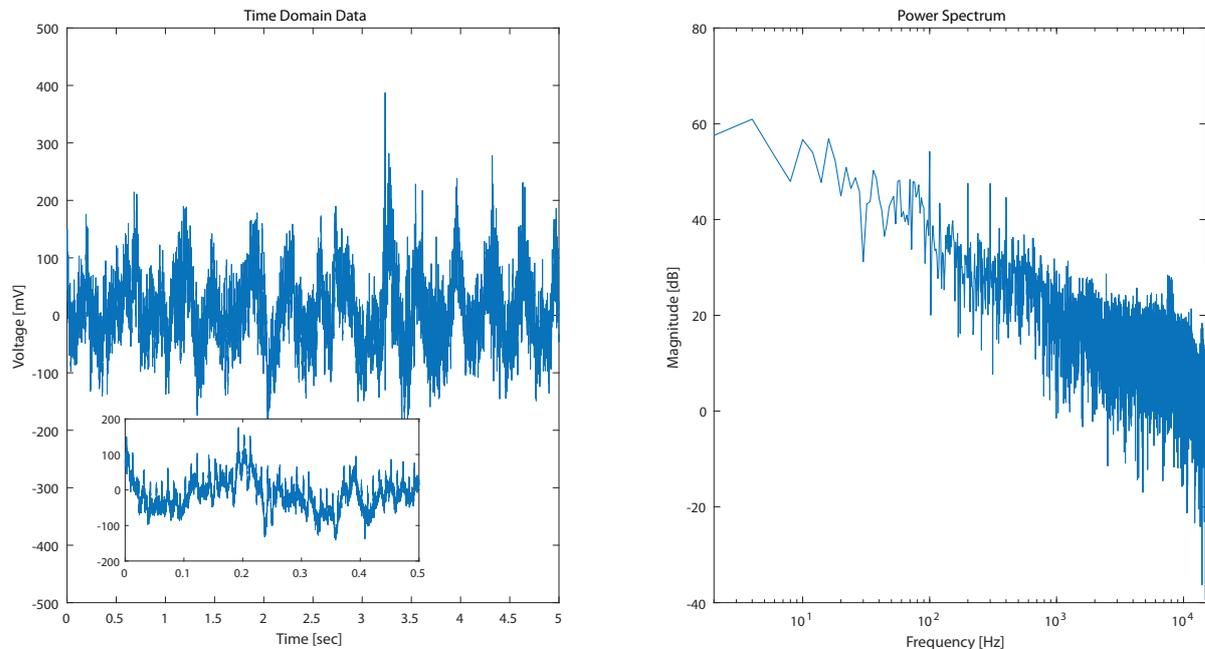
Figure 5.3: Raw signal in the time domain and its power spectrum

By looking at the figure on the left and also by taking into consideration its zoomed-in version in the smaller box, it is easy to deduce that the recording is dominated by a low-frequency and very high-amplitude wave. The frequency of this wave fluctuates around the value of $2Hz$ (judging both from the time domain signal and the peak in the power spectrum), so it can identified this as an EMG artifact caused by respiration muscles, within a limited margin of error [56].

What is more, the power spectrum reveals the presence of harmonics of $100Hz$ in the recording, caused most likely by the rectification of the $50Hz$ mains hum on one of the devices in the experimental setup.

The focus of this experiment is not the identification of single-neuron firing, but the change in the behavior of LFP signals during synchronous events. This means that not the entire spectrum of the recording has to be taken into account. Usually, the upper limit for the LFP spectrum is thought to be around $100Hz$, due to the low-pass characteristics of brain tissue [15]. With this in mind, the time series is low-pass filtered at $300Hz$, as a compromise between allowing potentially useful frequencies higher than $100Hz$ in and, at the same time, reducing the effect of high-frequency noise.

A high-pass filter is applied with a $10Hz$ cutoff frequency, eliminating the respiration artifact and most of the alpha brain rhythms. The mouse was placed in a darkened room during the experiment, so it is assumed that beta activity ($14-30Hz$) is also largely absent.

Visual examination of the spectrum reveals the presence of $100Hz$ harmonics that can easily be suppressed using a comb filter. However, this is only a small part of the signal, which means that the energy contained in the unwanted frequencies is, perhaps, too small for these frequencies to be easily identified and dealt with.

Spectrum displaying all unwanted features
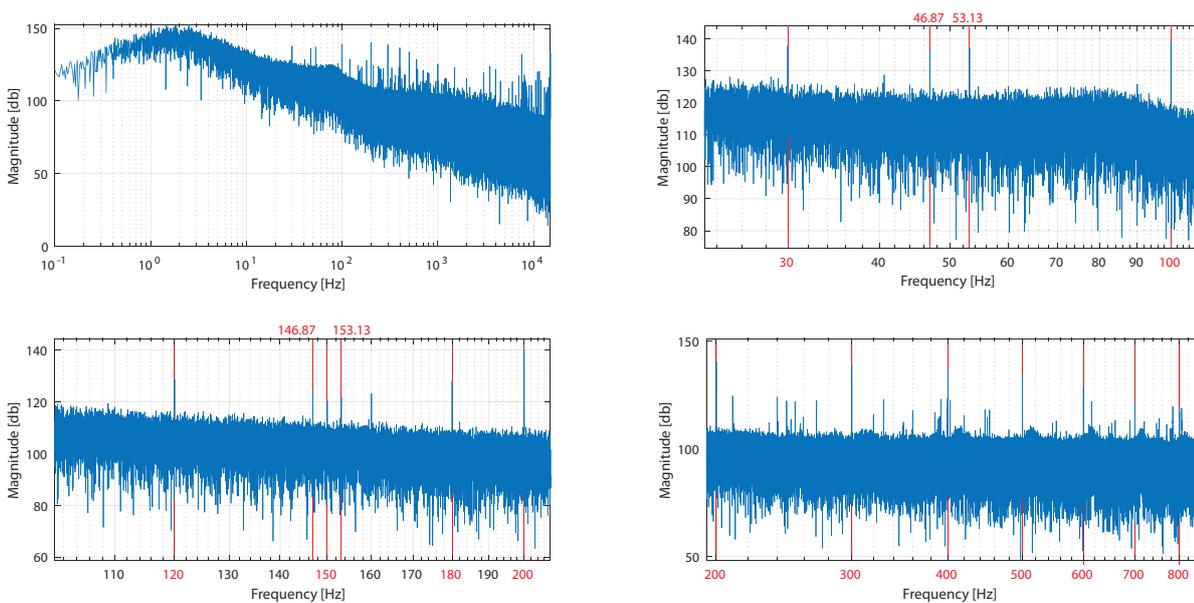
Figure 5.4: Spectrum displaying a larger number of artifacts. The red lines indicate unwanted frequencies present in the spectrum

The spectrum of a longer part of the signal will reveal more unwanted features, highlighted in Figure 5.4.

Besides the $100Hz$ harmonics, peaks at multiples of $30Hz$ are also visible, as well as peaks at about $3.1Hz$ on either side of 50, 150, 250Hz and so on. A comb filter is, then, applied to these frequencies to suppress their peaks in the spectrum. Finally, the filtered signal is downsampled to a sampling frequency of $1kHz$, which is low enough for the analyses not to be cumbersome and slow, and high enough to maintain satisfactory levels of detail in the signal and the analysis results. The product of this preprocessing paradigm is seen in Figure 5.5.

By zooming into the filtered version of the signal, the synchronous events are clearly visible, with some examples given in Figure 5.6.

These events fulfill the aforementioned criteria for neural synchrony, so it was deemed useful to establish a procedure to analyze and detect them within the time series of the neural recording. A number of methods and tools were used to this end, namely Singular Value Decomposition, Wavelet Transform, AR Modeling, and Time-Frequency Analysis. The final goal is to choose the one that provides with the best marker for the detection of the synchronous events. The analysis was initially executed on the short 10-second signal snippet shown in Figure 5.5, to rapidly test the effectiveness of the methods locally, but their robustness was also tested by analyzing larger parts of the signal eventually.

5.2.5. ANALYSIS AND DETECTION SCHEME

Various methods were applied on the recordings to specify which one displayed the best behavior, in this case, which one could correctly isolate the oscillatory activity. Initially, AR modeling was selected for the first part of the analysis, since it both offered satisfactory identification of the areas of interest and also specific insight into their characteristics. Matched filters were developed further along the analysis, in order to function as detection tools for synchronous activity.

ANALYSIS WITH AR MODELING

As has been stated before, neural signals are characterized by high nonstationarity. This means that there is little point in creating an AR model based on the entire signal. Instead, a sliding window approach was selected, wherein a number of samples were isolated in each iteration, based on which an all-pole model was created. At every step, the poles of the model were stored, so that their behavior over time could be studied

Preprocessing Results
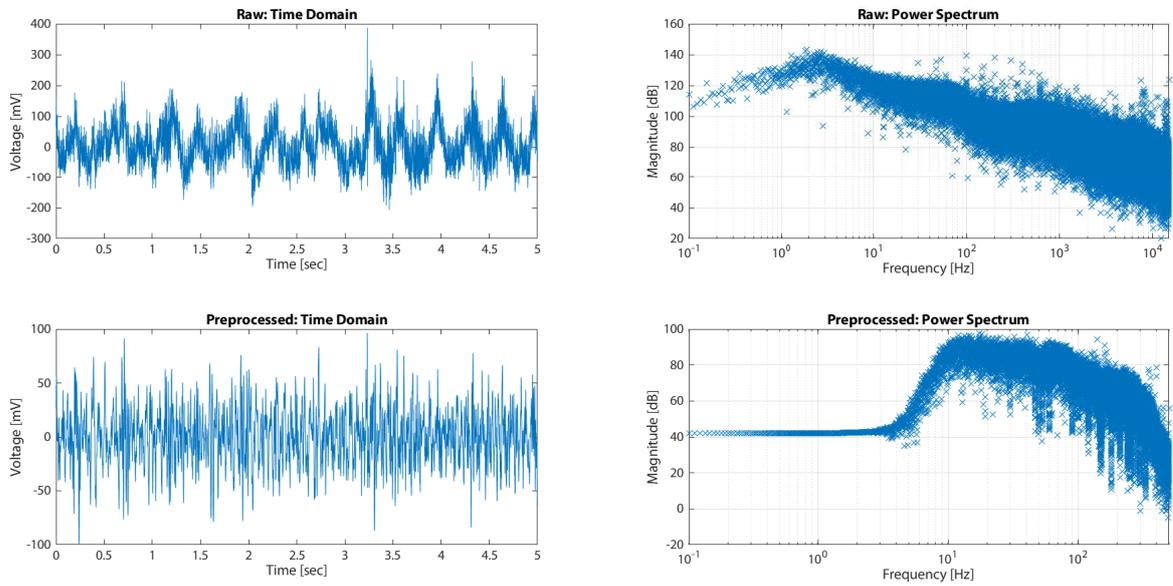
Figure 5.5: Raw and preprocessed signal and corresponding spectra

Synchronous neural activity events
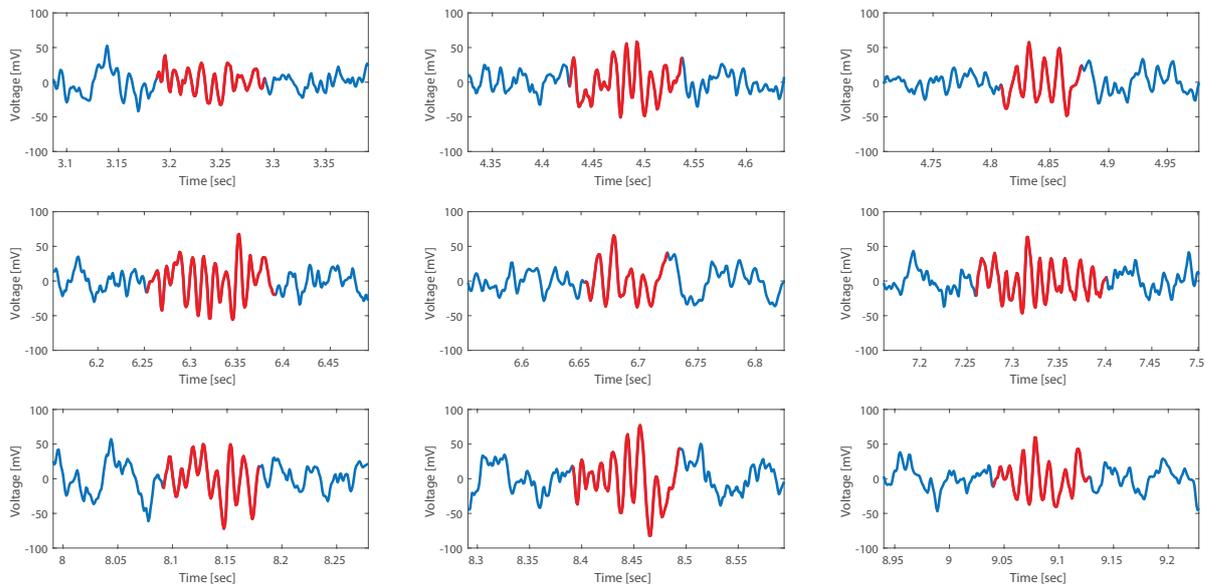
Figure 5.6: Examples of neural synchrony in the signal, indicated in red

in detail. A flowchart of this procedure is displayed in Figure 5.8.

It is important to define the windowing process that will be used, namely with what kind of window the data will be multiplied before modeling. The selection of a window whose amplitude diminishes close to its borders will decrease the undesirable effects of spectral leakage, present when a simple rectangular window is used. Left untreated, leakage introduces extra unwanted frequencies in the spectrum due to discontinuities at the beginning and end of the data window. A good candidate is a Hanning window, seen in Figure 5.7. After its application, the values close to the edges of the data tend to zero, effectively eliminating the aforementioned discontinuities and keeping the signal clean of additional noise.
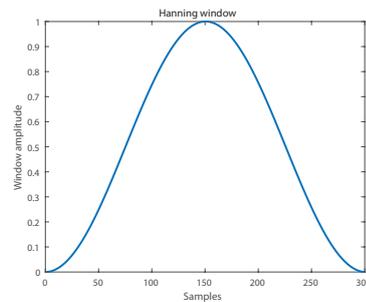
Figure 5.7: A 301-sample wide Hanning window

What is more, a compromise must be made about the length of the window itself. The main idea behind analyzing the signal in parts is that it is assumed that these parts are indeed stationary (see 3.1.1). Following this line of reasoning, selecting a large window runs the risk of the signal contained within it ceasing to be stationary, which defeats the whole purpose of windowing. On the other hand, the oscillatory events under scrutiny might not fit inside a very small window, in which case, it is possible that not all of their features are taken into account while modeling. A middle solution was reached after visual examination of the signals, which revealed that most of the oscillations have an average length of less than 150 samples (or 0.15 seconds), which was, eventually, the window length selected for the algorithm.

Another important decision in this case is the selection of the model order. The time series is complicated enough to discourage the use of a low-order model, however, a very high order will result in the modeling of high-frequency features and, possibly, noise, which is not desired, leading to a certain compromise between the guarantee of satisfactory modeling and loss of detail. Typically, the model orders used range from 5 to 8 [7, 53].

The parameters that can be monitored in this kind of analysis are the positions of the rhythms of the model, in this case, the positions of the poles, since AR models are all-pole models. In general, poles are complex numbers, usually denoted as $\alpha + \beta i$, however, their polar representation ($r\angle\theta$) provides more valuable information. In this form, θ is the normalized frequency that a particular pole models, while r is indicative of how dominant this frequency is. The closer to the unit
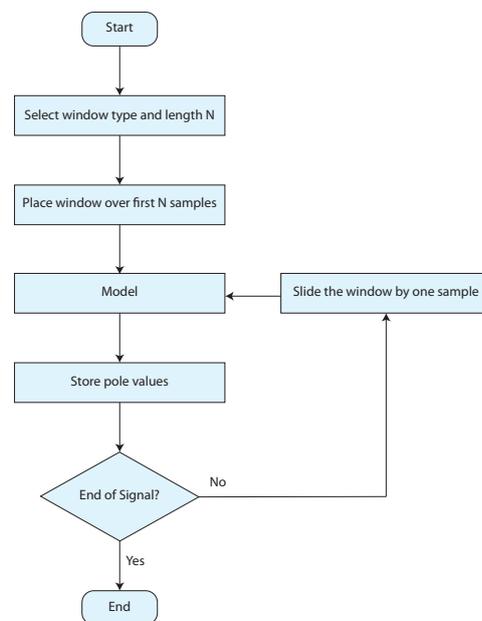
Figure 5.8: The AR Modeling scheme

circle, the more energy is contained in a particular frequency.

The basic idea behind this procedure is that the behavior of some or all of the poles will change in a similar way during the oscillation events. The first approach is the monitoring of all the individual parameters and assessing whether one of them actually behaves in a way consistent with the hypothesis. Besides that, it is important that the application of a threshold on this behavior is possible, because this will facilitate the detection process. The norm of the poles was considered as a suitable monitor.

Once the norms were calculated for all time instances, a thresholding scheme had to be implemented. Whenever the magnitude of the norm exceeded this threshold value, that duration would be marked as a significant event. In order to manage any outlier values that surpassed the threshold for a very short period of time, a second threshold, in this case in time, was also applied.

It was observed that the norm of one specific pair of poles displayed a behavior consistent with the appearance of oscillations in the signal, so the thresholding scheme was applied on this norm. In Figure 5.9 the signal and the norm are displayed. The areas of the signal highlighted red are events that were singled out by the algorithm. The parameters used in this analysis were:

- Order = 5

- Window Size = 150 samples, equivalent to 0.15 seconds

- Threshold = 0.91

- Minimum Duration = 0.06 seconds
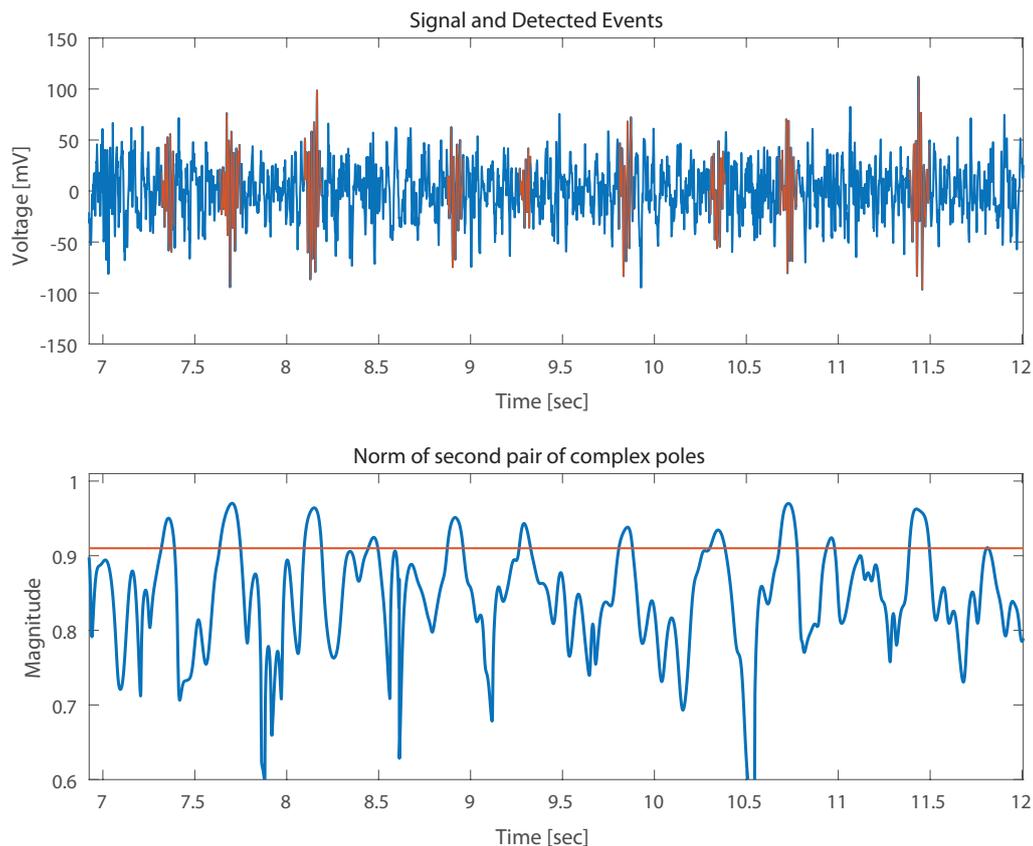
Figure 5.9: The pole norms over time in a small part of the analyzed signal. The detected events are highlighted in red in the top graph, while the threshold value is the red line in the bottom graph

Note that in Figure 5.9 the norm exceeds the threshold more than once, but only the ranges larger than the minimum duration set were identified as significant events and marked as such.

In total, 168 oscillatory events were isolated by the AR modeling scheme. Now it is possible to visualize the changes occurring in the model more clearly. By creating models from only these parts of the signal and from random parts that do not contain interesting features and plotting their poles on the complex z-plane, we get Figure 5.10. The poles in blue are the poles of the events detected with the thresholding scheme, while the red ones are taken from windows positioned at random along the signal.
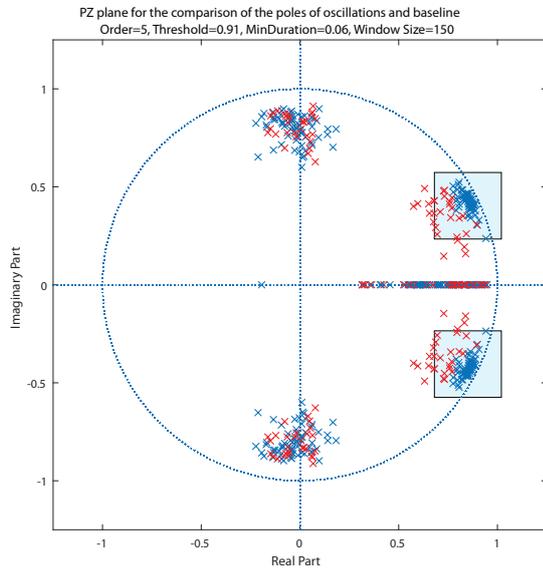
Figure 5.10: The poles of the system at different time windows, on the complex plane. Poles of the oscillations are shown in blue, while poles that are produced from random epochs of the signal are shown in red
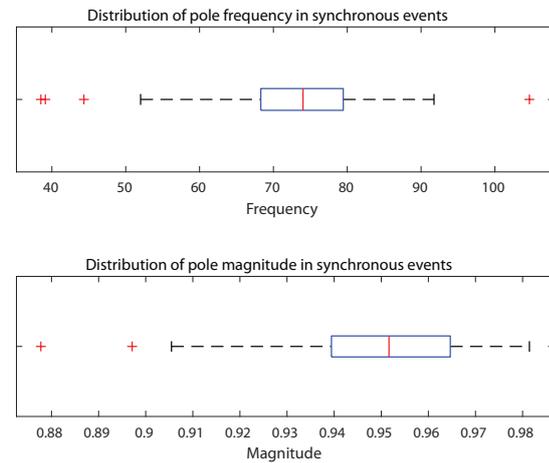
Figure 5.11: The distributions of pole frequency and magnitude in the detected events

Evidently, the model poles during the oscillations are concentrated in a specific area in the unit circle, denoted with a blue square in Figure 5.10, while poles resulting from areas of asynchronous activity do not follow any specific pattern. By using the properties of the unit circle in the complex z-plane (see 3.2.3), it is possible to specify the frequencies these poles represent. Their distribution is visualized in the box plot in Figure 5.11. From the AR model analysis, it can be surmised that the prefrontal cortex displays synchronous events of frequencies ranging from 70 to $80\,Hz$. Also, despite the AR modeling scheme behaving quite well, its application is somewhat impractical, due to high computational cost of the algorithm. Thus, although this approach can provide useful information about the signal and the characteristics of the synchronous events, it cannot be easily implemented in a real application of simple detection of these events.

DETECTION WITH MATCHED FILTERS
In order to detect synchronous events faster, without using the slow modeling approach described above, it is possible to design a different scheme, purely for detection, based on the knowledge gained from the AR analysis itself. A logical step forward would be the design of bandpass filters centered around the frequencies identified by the boxplots in Figure 5.11. However, since it was observed that the oscillations displayed similar behavior throughout the length of the time series, the design of Matched Filters was also deemed a very good candidate for the backbone of the detection scheme.

Thus, the oscillations detected by the AR scheme were used as templates for the matched filters. As a first step, the output of the matched filter was assessed visually, at which point, an important observation was made. The output of the matched filter yields clear oscillatory activity during significant events and asynchronous behavior everywhere else. Hence, in order to suppress the latter activity, while keeping the oscillations intact, the signal was also filtered with a differentiated version of the templates. In this manner, the sinusoidal-resembling oscillations would remain largely unchanged, excluding a phase shift, while the filter

output during asynchronous activity would have a smaller amplitude. An isolated oscillation, the resulting filter templates and their power spectra are shown in Figure 5.12
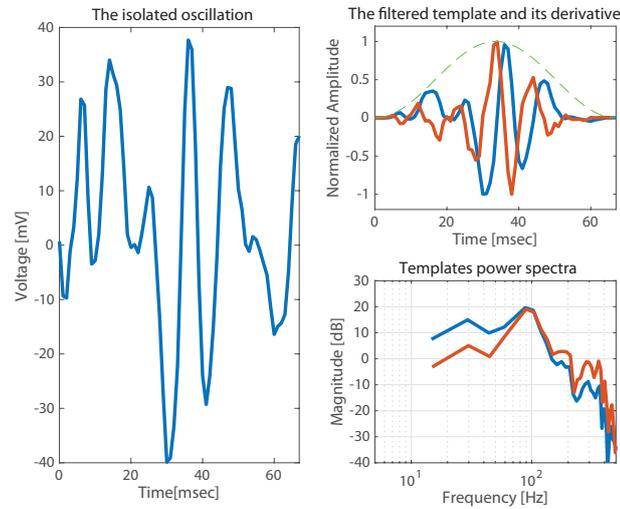
Figure 5.12: An oscillation isolated by the AR scheme. The original is shown in blue and the differentiated version in red

This differentiation scheme facilitates thresholding, since, provided that a suitable template is used, the output of the matched filter is much more pronounced and easily discernible during oscillatory events. This output is also normalized, so that the same thresholding scheme can be used for all recordings, regardless of their amplitude. This way, the threshold takes values between 0 and 1.

Eventually, the marker that undergoes thresholding is the upper envelope of the filter output that is created by using the differentiated version of the template. This behavior can be seen in Figure 5.13.
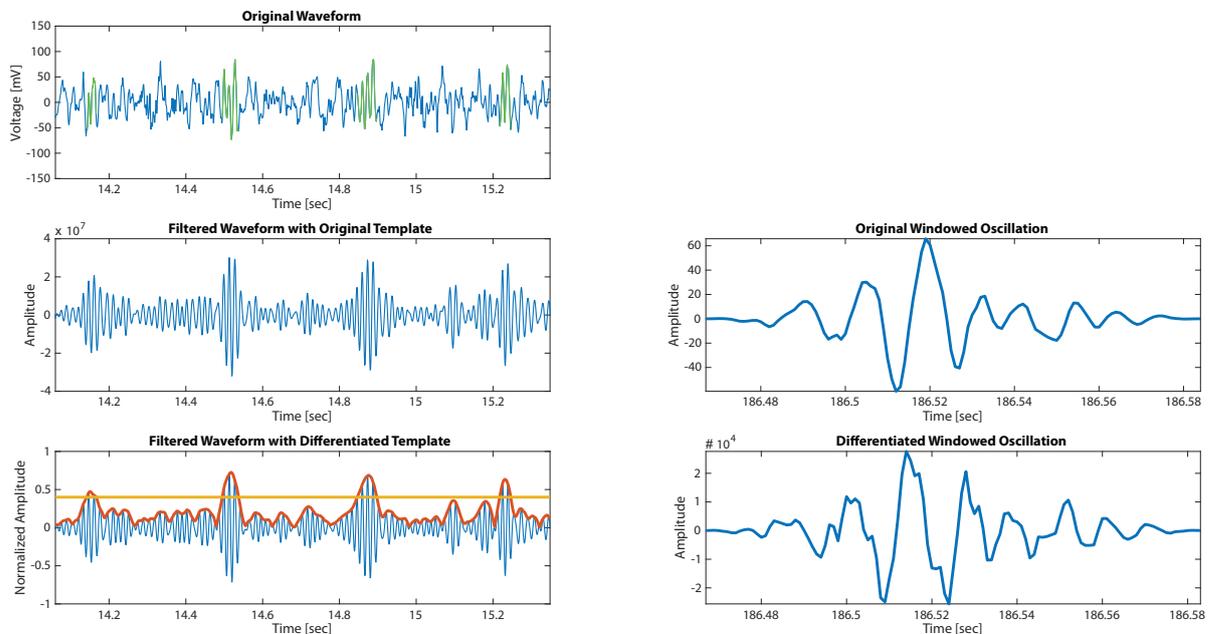
Figure 5.13: Comparison between matched filtering with the original and the differentiated template. The green areas on the signal are identified as oscillations by the detection scheme
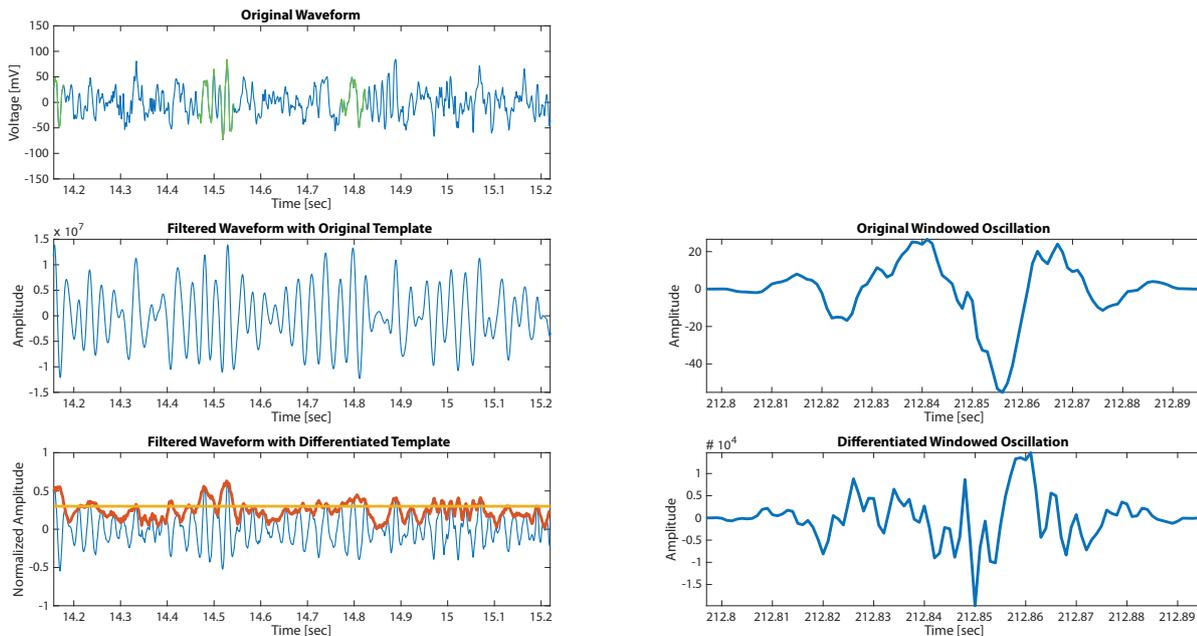
Figure 5.14: Comparison between matched filtering with the original and the differentiated template.
In this instance the template is unsuitable for filtering.

Using the template shown in Figure 5.13, the envelope of the output displays clear peaks that can be effectively thresholded based on a user-defined scheme. In other words, this particular template is a good representation of the behavior under scrutiny. In Figure 5.14 there is an example of how the output looks like if an unsuitable template is used. In this case, the output cannot be thresholded, while its amplitude actually drops at around 14.85 sec, where a clear synchronous event can be seen in the original waveform.

Clearly, the quality of the detection scheme is a function of three parameters; namely the filter template used, the threshold set for the amplitude of the envelope of the filtered signal and the minimum accepted duration that it can take values over the threshold. In order to specify the parameter set that optimizes the algorithm, a parametric analysis was designed, shown in Figure 5.15.
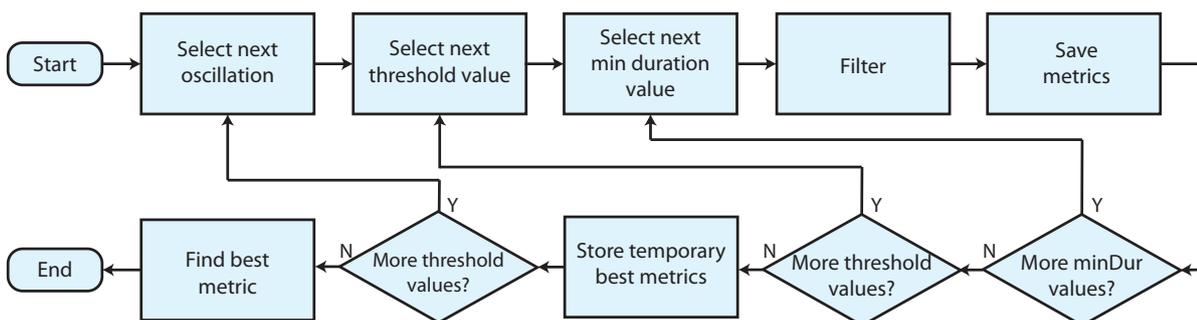
Figure 5.15: The parametric analysis for the specification of the best parameter set for the matched
filtering detection scheme

In every iteration of the algorithm, a different set of parameters is used and evaluated. The performance metrics are stored for each of these sets, to be evaluated at a later time. For each of the templates, the following parameters were used:

- Threshold: 0.2 to 0.8 with a step of 0.1, a total of 7 values

• Minimum Duration: 0.01 to 0.06 with a step of 0.01, a total of 6 values

so, eventually, for the 168 oscillations that were isolated by the AR scheme, the total amount of iterations sums up to $168 \times 6 \times 7 = 7056$.

The performance of each iteration is also visualized on the ROC space, shown in Figure 5.16, which will also be used for the selection of the best parameter set, based on the distance of each point to $(0,1)$ (see Figure 5.2).
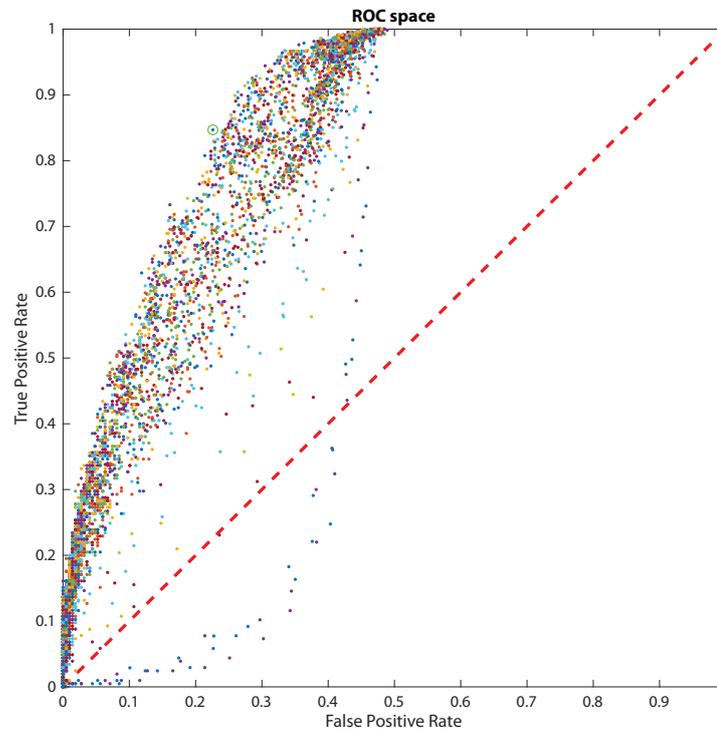
Figure 5.16: The ROC space for the parametric analysis

Each point in the Figure represents a different parametrization, while each color represents a different filter template. The point closest to $(0,1)$ is highlighted with a green circle. Besides the minimum distance on the ROC space, the parameters yielding the highest F-score and ADR were also taken into account. In summary, this is shown in Table 5.1. Note that the oscillation index is the number of the detected event out of the 168

	Best Value	Oscillation Index	Threshold	Minimum Duration
F-score	0.7735	98	0.3	0.03 sec
ADR	0.83786	131	0.4	0.02 sec
ROC distance	0.2523	98	0.3	0.03 sec

Table 5.1: The parameters corresponding to the best metrics

detected in total. Finally, the two templates that will be used for the validation of the scheme are shown in Figures 5.17 and 5.18.

These parameter sets will be tested on the validation dataset (Signal 2), to evaluate the performance of the filtering and thresholding scheme on a different waveform.

5.3. VALIDATION OF THE DETECTION SCHEME

Signal 2 was used for the validation of the detection scheme. The noise levels in this signal were higher than the ones found in Signal 1, used for the training of the scheme. Low-pass filtering at a lower frequency, namely
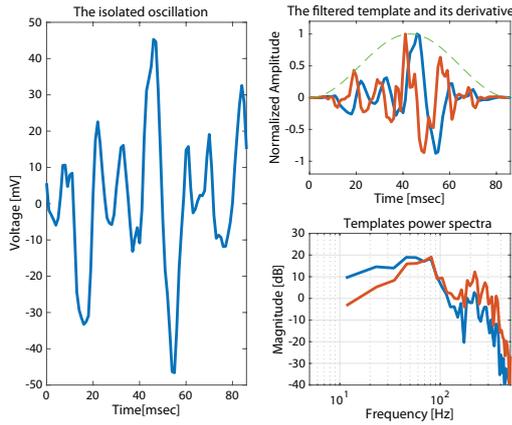
Figure 5.17: Oscillation 98
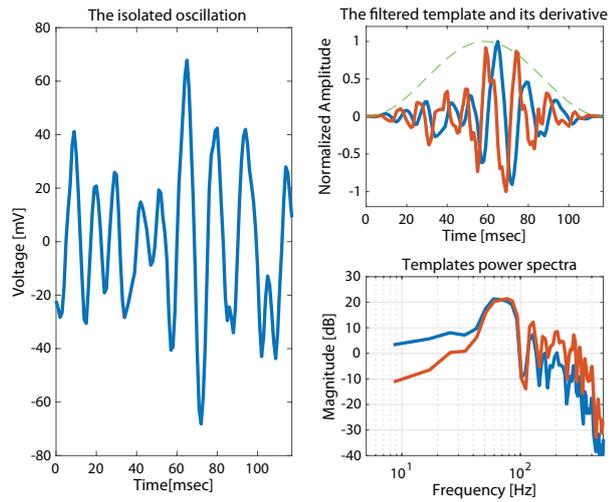
Figure 5.18: Oscillation 131

150 instead of $300Hz$ facilitated the visual identification of the oscillations and the labeling of the data. Once the recording was labeled, it was filtered using each of the two parameter sets identified in the previous section, yielding the results shown in Figures 5.19 and 5.20 and Table 5.2.
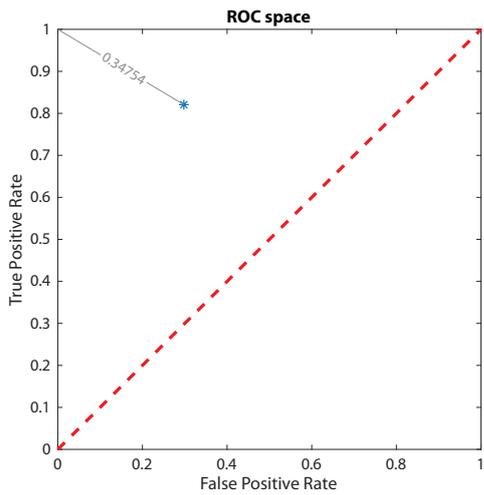
Figure 5.19: Oscillation 98
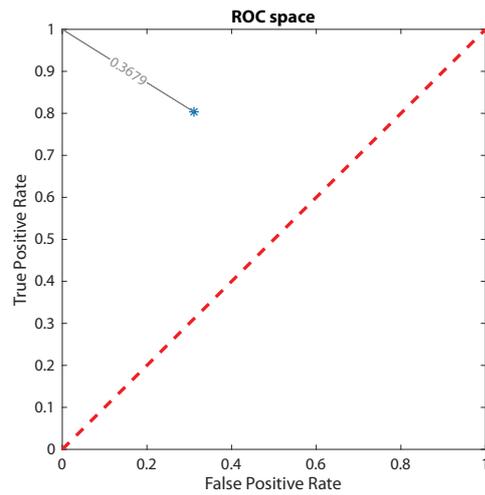
Figure 5.20: Oscillation 131

	Parameter Set 1	Parameter Set 2
F-score	0.61339	0.5794
ADR	0.7631	0.74625
ROC distance	0.34754	0.3679
TP	142	135
FP	148	163
FN	31	33
TN	349	361

Table 5.2: Validation metrics

where:

- Paremeter Set 1: Oscillation 98, Threshold = 0.3, Min Duration = $0.03 sec$

• Paremeter Set 2: Oscillation 131, Threshold = 0.4, Min Duration = 0.02sec

It can be surmised from the Figures and the Table that using the first parameter set yields slightly better results. That being said, the number of false positives detected by the scheme is very high. Examples of true and false positives can be seen in Figures 5.21 and 5.22. In the top graphs in these Figures, black denotes the labeling of the signal, while green denotes a detection by the filter.
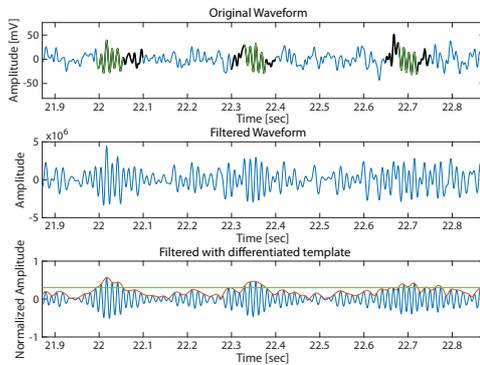
Figure 5.21: True Positives in the Validation stage. The manual annotations in the data are displayed in black, the events detected by the algorithm are shown in green
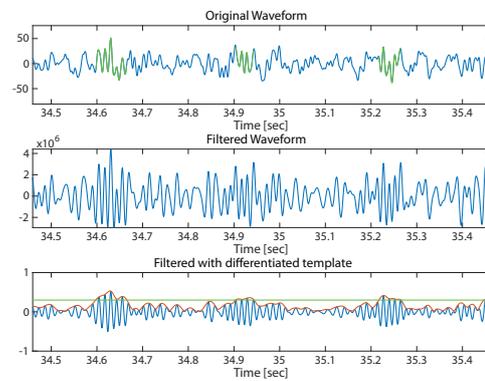
Figure 5.22: False Positives in the Validation stage

5.4. RESPONSE TO STIMULATION

The two signals were used to examine the response of the PFC to stimulation of the VTA are Signals 2 and 3, which were subjected to the detection scheme described above.

In order to obtain a more compact picture of the behavior of the signal, the latter was segmented into epochs of 20 seconds, each centered around a stimulation pulse. Each epoch was, in turn, divided into 20 zones of 1 second each. Then, the numbers of oscillations initiating in each zone was calculated for each of the epochs and were, eventually, summed. This is visualized in the peri-stimulus histograms, seen in Figures 5.23 and 5.24.

It should be noted that the fact that the number of detections before and after the stimulation pulse is the same in Figure 5.24 is purely coincidental. The same graph, taking 10 out of 14 epochs into consideration, showing that the oscillations on either side of the pulse is not always the same, is displayed in Figure 5.25.

It is clear that the results extracted from these two signals are conflicting. In Figure 5.23, the stimulation pulse appears to momentarily silence the oscillatory activity in the PFC, before it reverts to normal. In Signal 3, however, there is an accumulation of oscillations on either side of the pulse.
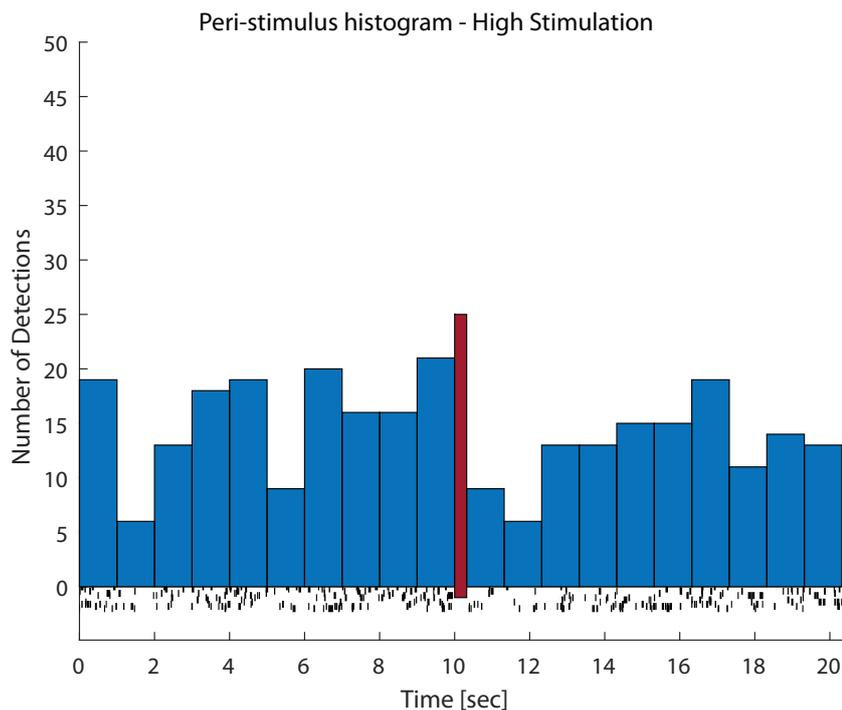
Figure 5.23: Peri-Stimulus Histogram for Signal 2. The black lines at the bottom represent the oscillations detected by the scheme. The blue bars represent the total number of detections in the 1-second time windows and the red bar represents the stimulation pulse.
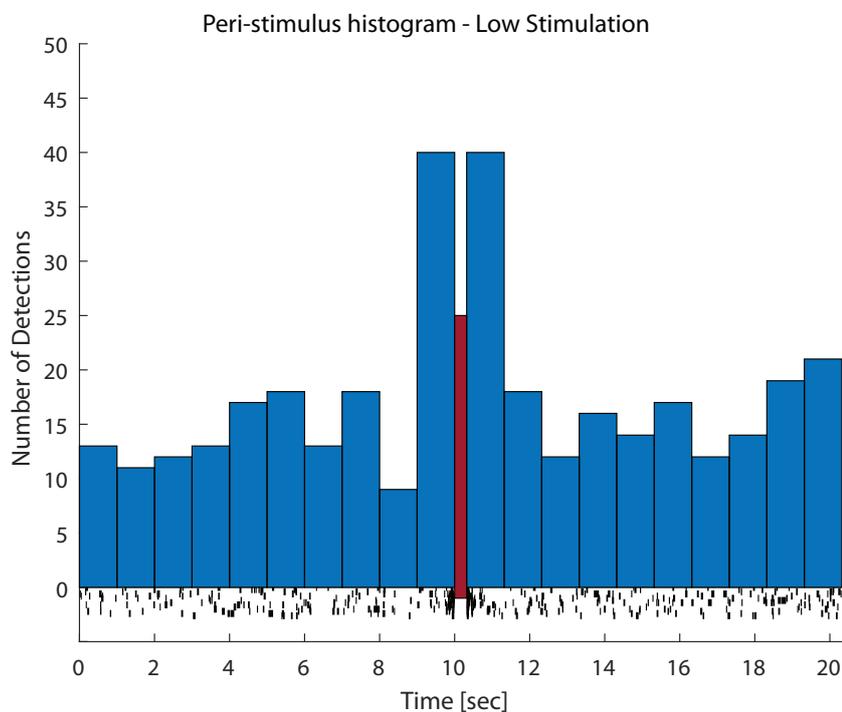
Figure 5.24: Peri-Stimulus Histogram for Signal 3. The black lines at the bottom represent the oscillations detected by the scheme. The blue bars represent the total number of detections in the 1-second time windows and the red bar represents the stimulation pulse.
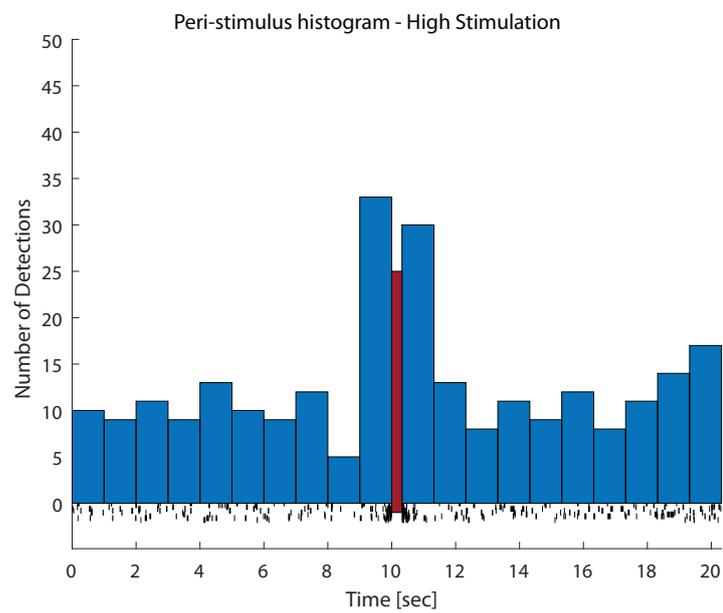
Figure 5.25: Peri-Stimulus Histogram for Signal 3 for 10 out of 14 stimulation pulses. The black lines at the bottom represent the oscillations detected by the scheme. The blue bars represent the total number of detections in the 1-second time windows and the red bar represents the stimulation pulse.

6

DISCUSSION AND CONCLUSION

The distillation of the procedure described in the previous chapter can be seen in Figure 6.1.
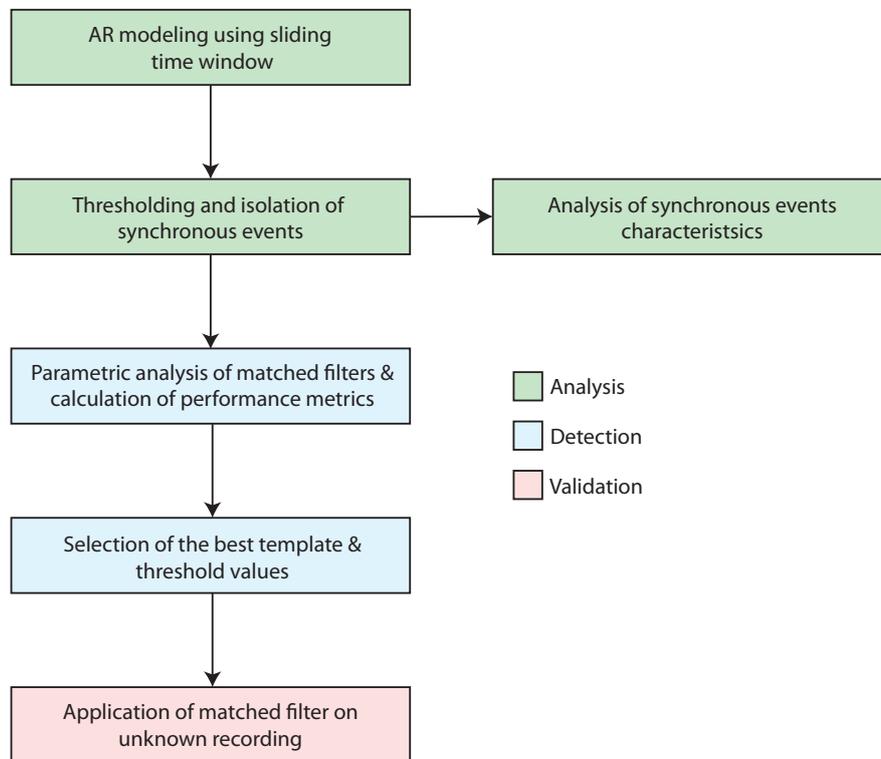
Figure 6.1: The analysis and detection scheme

6.1. ANALYSIS – AR MODELS

The selection of the model order was not made based on the observed characteristics of the oscillations; it was a compromise between lacking in detail using a lower order or running the danger of modeling noise using a higher model order. Indeed, the fifth order model effectively isolates gamma-range oscillations in the $70-80Hz$ range, which are modeled by one of the complex pole pairs of the model.

What is more, the selection of a moderate model order for the AR scheme was made so that the algorithm was not extremely computationally intensive. Possibly, a higher model order can isolate oscillations in different frequencies as well, paving the road for the study of different neural synchrony events using this method.

The performance of the thresholding scheme in the analysis phase using AR models was evaluated visually, as the labeling of the data was not available at the time. If the data had been labeled, a parametric analysis similar to the one performed in the optimization of the matched filter parameters could be implemented, keeping in mind that each iteration would last much longer. Based on data labels, some of the oscillations isolated by the AR modeling scheme could have been discarded as insignificant, which would also reduce the computation time of the parametric analysis in the detection phase. The data annotations were not used to corroborate the results of the AR analysis, as the isolation of the oscillations was deemed satisfactory and the isolated events were used for the detection in the following step.

6.2. DETECTION – MATCHED FILTERS

By looking at Figure 5.12, the bandpass characteristics of the filter templates and, by extension, the isolated oscillations is apparent. The frequency band that can be regarded as the passband of the spectrum overlaps with the distribution of frequencies displayed in the boxplots in Figure 5.11, ensuring a first level of consistency in the results.

The spectrum also showcases some frequency content below the passband, which is steeply attenuated after the differentiation, while the passband remains largely unchanged. This fact, coupled with the observation that, indeed, the differentiated version of the template yields better results than the original one, reinforces the argument that the main feature of these oscillations is indeed accumulated in the passband, since the system works better when the latter is more pronounced than other features of minor importance.

Even though their passband is similar, in other words, their main characteristics can be found in the same frequencies, oscillation 98 differs from oscillation 131 in the sense that it appears to display some minor characteristics away from the passband that are shared by more oscillatory events, making its function as a filter template more effective.

It would be a gross oversight not to address the issue of the large number of false positives. At this point, it has to be emphasized that the labeling of the data was carried out manually. One of the criteria for the selection of the labeled parts of the waveform is duration, quantified as the number of peaks in each oscillatory event. Oscillations of less than 3 peaks were either not considered as significant or overlooked entirely. However, even these very short events appear to share characteristics with the templates used, which increases their correlation beyond the detection thresholds, creating these false positives.

In any case, the algorithm appears to be oversensitive even to the shortest occurrences of synchronous activity in the band of interest, creating a large number of false positives. The criteria for the best parameter selection were the best F-score, ADR and distance from $(0, 1)$ on the ROC space. If the goal is to minimize the false positives, different or less strict criteria have to be used for the parameter selection. For example, a different point on the ROC space can be selected, closer to the y-axis, displaying a reduced FPR. Of course, this could be at the expense of the other metrics and the performance of the scheme as a whole.

Seeing as the labeling of the data is inherently flawed, any conclusion regarding the oscillatory behavior of the PFC in Signals 1 and 2 (as defined in 5.1) has to be taken with a grain of salt, as the level of credibility of the results as well as the validation of the detection scheme are only as good as the manual labeling of the data. At its present state, the detection scheme can be used as a guide for researchers to isolate this particular kind of neural synchrony and select which of the detected oscillations are significant.

Since a specific kind of neural synchrony displays specific characteristics in time and frequency, the use of matched filters is a crude, yet effective way of implementing an *ad-hoc* scheme for their detection. What is more, since synchronous activity varies based on brain area and, possibly, medical condition, the analysis and detection scheme put forward in this project can be used to isolate and detect different kinds of oscillatory activity in the brain, such as epileptic seizures.

6.3. RESPONSE TO STIMULATION

Signals 2 and 3 were used for the study of the effect of VTA stimulation on the PFC activity and they yield completely different results regarding the PFC responses.

In any case, the amount of data that was analyzed was limited, so there can be no definitive conclusion extracted from this analysis. That being said, an informed assumption can be made for this behavior nonetheless, in the hope that it motivates a more comprehensive data collection and analysis.

There appears to be a direct relation between the amplitude of the stimulation pulses and the response elicited in the PFC. High amplitude pulses momentarily suppress the occurrence of oscillations in the $70-80$ Hz band, while oscillations seem to appear at a much higher rate around pulses of lower amplitude.

These two behaviors are inherently different, in the sense that in the first case the brain appears to respond to each individual stimulation, which is manifested in the suppression of neural synchrony. However, in Signal 3, the level of synchrony, and by extension, the number of oscillations, is high even before the onset of the stimulation pulse, as if the brain is *expecting* a stimulation at that specific time.

Considering the fact that stimulation pulses are delivered at fixed time intervals, an assumption would be that the brain has been *conditioned* and responds accordingly in expectation of the incoming stimulations [60, 61]. If this was the case, the number of oscillations around the stimulation would increase with every incoming pulse. Unfortunately, the part of the signal containing the first pulses administered to the mouse was not contained in the waveform that was analyzed, so, again, no safe conclusion can be reached, however, this is a valid indication of conditioning.

The higher amplitude pulse in Signal 2 might actually be affecting a structure called the tail of the Ventral Tegmental Area (tVTA) or the Rostromedial Tegmental Nucleus (RMTg), located directly posterior to the VTA, which has been shown to act as a major inhibitor of dopamine networks [62]. This could explain why the neurons in the mesocortical dopaminergic pathway, in which the recording electrode is introduced, are silenced immediately following the stimulation pulse.

7

FUTURE WORK

The original goal of this project was to correlate the activity of the PFC with the concentration of the neuro-transmitter dopamine within it. As has been stated previously, the electrode probe was introduced into the mesocortical dopaminergic pathway of the mouse. A valid assumption would be that the oscillatory activity that was detected was really the coordinated activity of dopaminergic neurons, relaying messages from the basal ganglia towards the cortex.

This cannot be stated with absolute certainty, though, unless markers like the local concentration of dopamine around the electrode is measured, as well as the activity of deep brain structures related to dopamine, such as the VTA or the substantia nigra.

Even then, however, it is doubtful if the oscillations were a result of the activity of dopaminergic neurons. In order to specify if dopaminergic neurons create the oscillations, the only fool-proof way is to introduce a dopamine blocker into the brain, thus canceling out the effects of dopamine. If the oscillations are no longer detected in the signal, it can only mean that they had been caused by the synchrony of dopaminergic neurons. Despite the interesting insight this procedure would provide, this was deemed to be outside the scope of this project.

A more detailed approach would be to create a matched filter bank, using more oscillations as templates, as an endeavor to compensate for the true negatives produced by the detection algorithm. Synchrony at different frequencies could also be detected and analyzed by utilizing a higher-order AR model, at the risk of introducing noise into the analysis.

What is more, further enhancement of the matched filter could be achieved by the creation of a completely customized template, based on the spectrotemporal and other characteristics of the detected oscillations. This would entail a more detailed study of this type of neural synchrony, providing additional insight into its characteristics and nature. Also, this approach could completely uncouple the analysis, isolation and detection procedures.

If it is desired that neural synchrony, such as the one presented in this project, is detected in real time, the speed and computational cost of the filtering and thresholding algorithms must be addressed. This is one of the reasons why the matched filter approach is much more preferable to AR modeling.

Matched filtering and the thresholding of the output is fast enough to be used in real time, during the actual harvesting of the data. A user-defined filter template and thresholding scheme can be used in harvesting software, such as the OpenEPhys suite, for on-line detection and evaluation of oscillatory activity.

The methods used in this project are purely signal-processing-oriented and the quantifiable characteristics of neural synchrony are based on an Autoregressive Modeling scheme and the Fourier Transform. It can be argued that with modern techniques, such as deep learning, this procedure could be automated and would possibly reveal even more interesting features hidden in the signal. While this is true, it is always useful to re-

turn to the manual basic analysis methods, like the ones presented in this report, which can provide valuable insight on the nature of the signals. Information such as the one gathered by simple signal analysis methods could be invaluable in the study of neuronal firing, for instance in the identification and study of different neuronal species.

For consistency in the conclusions drawn that relate to the brain responses, more experiments must take place and recordings gathered. Histological analysis of the stimulated areas will provide consistency along the recorded data, while they might also shed some light on the hypothesis of the potential activation of the RMTg.

A

SYMBOLS AND ABBREVIATIONS

- ADR: Average Detection Rate
- AP: Action Potential
- AR: Autoregressive
- ARMA: Autoregressive Moving Average
- BPF: Band Pass Filter
- CNS: Central Nervous System
- DA: Dopamine
- DTFT: Discrete-Time Fourier Transform
- ECoG: Electrocorticogram
- EEG: Electroencephalogram
- EMG: Electromyogram
- ERP: Event Related Potential
- fMRI: Functional Magnetic Resonance Imaging
- FN: False Negative
- FP: False Positive
- GABA: Gamma Aminobutyric Acid
- HPF: High Pass Filter
- ICA: Independent Component Analysis
- ISI: Inter Spike Interval

- LMS: Least Mean Squares
- LFP: Local Field Potential
- LPF: Low Pass Filter
- MA: Moving Average
- MEG: Magnetoencephalogram
- MSE: Mean Squared Error
- NCC: Neural Correlates of Consciousness
- PCA: Principal Component Analysis
- PFC: Prefrontal Cortex
- RMTg: Rostromedial Tegmental Nucleus
- ROC: Receiver Operating Characteristic
- SN: Substantia Nigra
- SNR: Signal-to-Noise Ratio
- SVD: Singular Value Decomposition
- TN: True Negative
- TP: True Positive
- tVTA: Tail of the Ventral Tegmental Area
- VTA: Ventral Tegmental Area

B

MOUSE BRAIN ATLAS

The structures in the mouse brain that were mentioned in this report are shown in Figure B.1
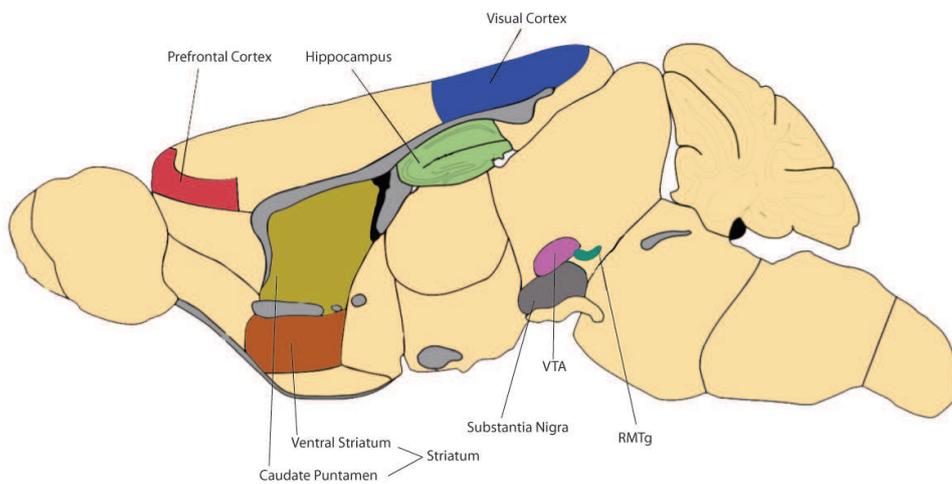
Figure B.1: A mouse brain section displaying the approximate locations of the structures mentioned in this report. Image modified from GENSTAT at
http://www.gensat.org/imagenavigator.jsp?imageID=36597.

C

ADDITIONAL METHODS

The following signal processing methods came up during the literature research for this project, but were not, eventually, used in the analysis. However, for the sake of completion they are presented in detail below.

C.1. NONPARAMETRIC ESTIMATION OF THE SIGNAL PROPERTIES

It has been established that neural signals are largely stochastic and non-deterministic. This fact means that since the signal cannot be known a priori, its properties have to be estimated.

The transient properties of the signal are unknown, so an approach in the frequency domain is necessary for nonparametric estimation. The analysis begins with the estimation of the autocorrelation function $r_x(k)$ of the signal $x(n)$:

$$\hat{r}_x(k) = \frac{1}{N} \sum_{n=0}^{N-1-k} x(n+k)x(n), \quad k = 0, \ldots, N-1 \tag{C.1}$$

and the knowledge that, in general, the power spectrum of a signal is given by:

$$S_x(e^{j\omega}) = \sum_{k=-\infty}^{\infty} r_x(k)e^{-j\omega k} \tag{C.2}$$

Combining C.1 and C.2, yields the estimate of the power spectrum of the signal, called the *periodogram*:

$$\hat{S}_x(e^{j\omega}) = \sum_{k=-N+1}^{N-1} \hat{r}_x(k)e^{-j\omega k} \tag{C.3}$$

Also, by remembering that the correlation function can be written as:

$$\hat{r}_k = \frac{1}{N} x(k) * x(-k) \tag{C.4}$$

and obtaining its Fourier Transform to obtain the spectral estimate, it emerges that the latter can be calculated by calculating the values of the square of the magnitude of the Discrete-Time Fourier Transform in all N points.

$$\hat{S}_x(e^{j\omega}) = \frac{1}{N} \left| X(e^{j\omega}) \right|^2 = \frac{1}{N} \left| \sum_{n=0}^{N-1} x(n)e^{-j\omega n} \right|^2 \tag{C.5}$$

Despite its elegant computation, the periodogram in and on itself is not a consistent estimate of the power spectrum, due to the fact that its variance does not approach zero as the number of samples increases. Also, the formula in Equation C.5 implies that the DTFT is computed in fixed-length rectangular windows, something that introduces leakage into the estimated power spectrum. These disadvantages are not beyond a measure of control, however, since they can be reduced by the use of *averaging* (see C.2) and *windowing*. Windowing refers to the use of a different window than the simple rectangular for the computation of the Fourier Transform, which displays less pronounced side lobes, thus reducing the leakage created by the discontinuities at the window borders. Among the most popular windows are the Hanning, Blackman, Hamming and Bartlett windows [7, 29].

C.2. AVERAGING

In its core, averaging aims for the reduction of noise in a signal displaying *repetitive patterns*. In simple terms, it exploits the random nature of the noise by aligning and adding a number of instances of a signal pattern and dividing by their number, thus reducing the effect of the noise and yielding a clearer signal. Note that, of course, this relies on the assumption that the noise contained in the signal is uncorrelated over time.

What is more, another proviso for the use of averaging is the guarantee that the process under scrutiny can be *replicated*. In neural signal terms, it is assumed that a repeated stimulus elicits a similar response in the brain every time, which is recorded with one of the methods mentioned above. The recorded noisy responses are then averaged to yield a single waveform containing the response itself, clear of noise.
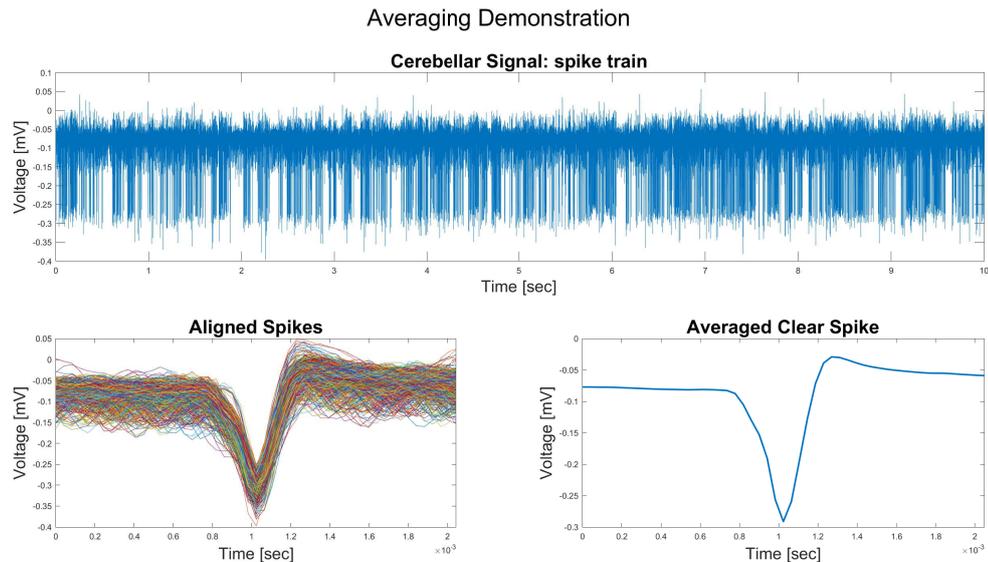
Figure C.1: Demonstration of averaging on signal harvested from the Cerebellum

The application of this method is displayed in Figure C.1. The spikes from the signal are detected and aligned (top and bottom left graphs). They are then averaged over time, producing the clear spike (bottom right). It is also apparent from the Figure that averaging mimics the effect of a linear time-invariant filter applied on the spikes. In the case of neural recordings, which are characterized by relatively low SNR, hundreds or thousands of repetitions of the desired pattern have to be elicited by a suitable stimulus for the averaged signal to contain acceptable levels of noise.

The version of averaging described above is known as *ensemble* averaging and is commonly used for noise reduction. However, an important disadvantage of this method is its inability to keep up with dynamic transient changes in the recording. In order to tackle this, a modified recursive version of the averager, taking past values into consideration, is used instead, consulting in *exponential averaging* [7].

Averaging can be ineffective if there are extreme values scattered across the recording. It is wise to remove any outlier values (see 3.2.1.4) before averaging, to achieve enhanced robustness in the averaging algorithm. Care must be taken when the signal is infected with artifacts. Clearing the signal of these artifacts should take place prior to averaging. Also, if the desired signals are elicited with the help of an external stimulus, the latter can be strategically timed so as not to coincide with time-locked artifacts, such as the Electrocardiogram.

If artifacts persist, *weighted* averaging can be used. By computing a model for the ensemble to be averaged and selecting a performance criterion, one can create a weight vector, which when multiplied with the ensemble yields the weighted average. Weighted averaging is also recommended when the statistical properties of the noise as well as its levels vary across the recording. Herein lies a fundamental difference between

weighted and ensemble averaging, in the fact that in the former, *a priori* knowledge of the nature of the noise variance is necessary.

Besides averaging in the time domain, the tool can be also used in the frequency domain, aiming at the reduction of noise variance contained in the power spectra. A very typical example is the Welch method [63], itself an extension of the Bartlett method [64]. Here, the signal is first segmented (see C.3) and the power spectrum for each segment is calculated. Then, these spectra are averaged, yielding the *averaged periodogram*. A disadvantage of this method is the reduction of the frequency resolution of the signal by a factor of K, where K is the number of segments. As stated before, usually a window is applied on the segments, in order to avoid *leakage*. The difference between the Welch and Bartlett methods is that in the former, the windows are allowed to overlap. This process can also used on estimated spectra.

Averaging is being used extensively in research. Besides its random noise cancelling properties, averaging can be used to determine the baseline of a signal, and then subtracting it from each of the signal's epochs, a process which yields the *non-event-related activity* contained in the signal [46, 47]. Holroyd et al. [38, 41, 42] average neural signal epochs to detect deflections in the signal related to Error-Related Negativity and Positivity and reward prediction. Similar deflections of neural signals are also detected by averaging in [17]. Connoly et al. [19] provide an example of frequency domain averaging, which aims in the detection and removal of artifacts. Ding et al. [34] use averaging for spectral estimation in order to derive a model for cortical ERPs. Spectral estimation by the use of averaging is also used in [20]. An application of the Welch method with the use of a Hanning window is presented in [36]. Averaging is also regarded as an important tool for signal preprocessing when other methods of analysis, such as wavelet analysis (see C.4), are applied on the recording [65].

Obviously, averaging is only as efficient as the way it has been set up. If, for instance, there are neural signal deflections in the ensemble that are not properly aligned, either due to variations in timing or due to errors in the alignment procedure itself, they will definitely be lost in the averaging, potentially leading to loss of valuable information [18]. Loss of information can also be a problem when the response itself varies from trial to trial in amplitude (and, in the case of EEG, scalp distribution), in which case the use of ensemble averaging is not recommended [45].

C.3. SEGMENTATION

As has been stated above, neural signals display high nonstationarity, with their spectrotemporal and statistical characteristics changing over time. However, parts of the signal that are sufficiently short can be assumed to be at least quasi-stationary or, at any rate, display the same kind of behavior throughout their duration. Signal segmentation cuts the signal into these smaller pieces, based on a number of user-defined parameters.

Segmentation can be *fixed*, in which case all segments have the same size, or *adaptive*, where the segments are of different sizes that depend on the transient behavior of the signal. Since changes in the statistics of the signal can be detected by segmentation algorithms, it can also be used as a tool for feature extraction, providing useful clinically significant information contained in the signal that would be normally overlooked during visual examination of the signal [7].

For the segmentation of the signal, two time windows are defined: the *reference* and *test* windows. The part of the signal contained in the reference window, which can be of fixed or increasing size (see Figure C.2), is analyzed and its statistics are compared to those of the part contained in the sliding test window. A *dissimilarity measure* $\Delta(n)$ is also defined, based on which the statistics of the two windows are compared. A segment boundary is defined by the algorithm when the dissimilarity measure surpasses a predefined threshold η. It is important that the change is sufficiently abrupt to be detected by the algorithm. A slow change can be regarded as transient behavior, which this algorithm is not sensitive to.

A very appealing choice for a dissimilarity measure is the comparison of the spectral content of the two windows, quantified by the *Spectral Error Measure*. Two methods are analyzed in [7] for the computation of the SEM: the periodogram and the whitening approaches.
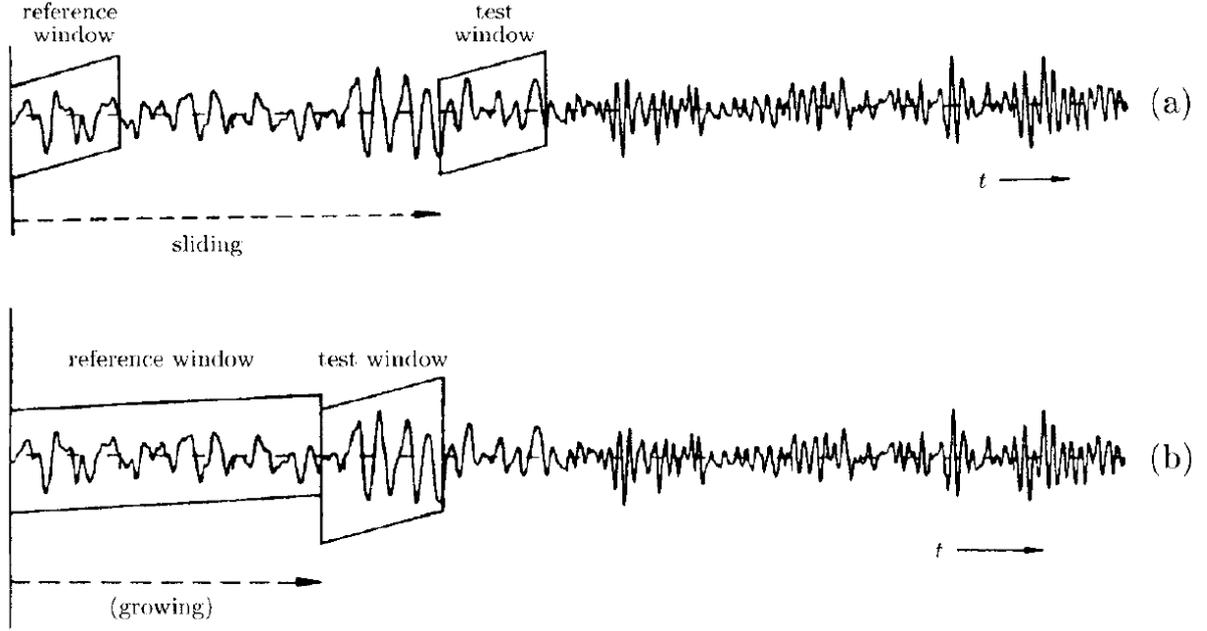
Figure C.2: Segmentation of an EEG signal with a fixed (a) and an increasing-size (b) reference window. Courtesy of [7]

The periodogram approach, as the name suggests makes use of the periodograms of the reference and test windows. The *running periodogram* is defined as:

$$S_x(e^{j\omega}, n) = \sum_{k=-\infty}^{\infty} r_x(k, n) e^{-j\omega k} \tag{C.6}$$

where r_x is the correlation matrix of the signal. The periodogram of the reference window is defined as $S_x(e^{j\omega}, 0)$. Eventually, the normalized SEM is given by:

$$\Delta(n) = \frac{\frac{1}{2\pi} \int_{-\pi}^{\pi} (S_x(e^{j\omega}, n) - S_x(e^{j\omega}, 0))^2 d\omega}{\frac{1}{4\pi^2} \int_{-\pi}^{\pi} (S_x(e^{j\omega}, n) d\omega \int_{-\pi}^{\pi} (S_x(e^{j\omega}, 0) d\omega} \tag{C.7}$$

The normalization is used so that decreases in signal power are detected as easily as increases. Normalization guarantees that the definition in Equation C.7 deals with changes in spectral power in a symmetric way.

The time domain version of the definition is:

$$\Delta(n) = \frac{\sum_{k=-\infty}^{\infty} (r_x(k, n) - r_x(k, 0))^2}{r_x(0, n) r_x(0, 0)} \tag{C.8}$$

The whitening approach makes use of an p-th order AR model to describe the signal, which is decorrelated or whitened using a linear prediction error filter $A_p(z)$. When a spectral change occurs, the error process $e(n)$ deviates from whiteness and this is where the segment boundary is set. Being a white process, the spectral power of the reference window is constant and equal to its variance:

$$S_e(e^{j\omega}, 0) = r_e(0, 0) = \sigma_e^2 \tag{C.9}$$

so the SEM is formulated anew as:

$$\Delta(n) = \frac{\frac{1}{2\pi} \int_{-\pi}^{\pi} (S_x(e^{j\omega}, n) - \sigma_e^2)^2 d\omega}{\left[\frac{1}{2\pi} \int_{-\pi}^{\pi} (S_x(e^{j\omega}, n) d\omega\right]^2} \tag{C.10}$$

and its time domain representation is:

$$\Delta(n) = \frac{r_e(0, n)}{r_e(0, 0)} - 1 \tag{C.11}$$

This type of signal segmentation must abide by certain specifications, namely few *false alarms*, few *missed detections* and low *detection delay*. A false alarm is the erroneous detection of a changes in the signal and the creation of smaller-than-necessary segments, a condition called *oversegmentation*. Missed detections are the exact opposite, namely, the failure to detect changes, leading to large, possibly nonstationary segments. The detection delay must be kept low so that the segment limit is as close as possible to the onset of a change in the signal properties. At any rate, segmentation is a very useful tool for the fist step of feature recognition in a signal and its parameters and properties also depend on the tools and methods that will be used an applied on the signal at a later time [66].

Despite its feature recognition abilities, segmentation is more commonly used in its fixed version. Most usually, signals are segmented around the onset of a stimulus or other significant time domain features of the signal (see Table C.1). The duration of the segment in this case is colloquially known as *epoch*, which is, then, subjected to further analysis. It has already been stated that the Welch and Bartlett methods are making use of segmentation to provide with smaller parts of the signal to undergo averaging. Rafiee et al. [67] use segmentation before analyzing the signal using wavelet analysis, which will be discussed in C.4.

Table C.1: Epochs duration in literature

Reference	Duration Before Reference	Duration After Reference
Santesso et al. [37]	$200ms$	$800ms$
Holroyd et al. [42]	$200ms$	$800ms$
Foti et al. [39]	$200ms$	$1000ms$
Foti et al. [40]	$200ms$	$800ms$
Bokura et al. [43]	$200ms$	$1000ms$
Zhang et al. [49]	$1s$	$1s$

C.4. TIME-FREQUENCY ANALYSIS AND WAVELETS

It is common practice to analyze signals by expanding them using a series of basis functions, such as in the case of the Fourier and Karhunen-Loève transforms. The former uses sines and cosines as basis functions, while the latter uses the largest eigenvalues of the correlation matrix of the signal to produce its truncated series expansion. However, these basis functions are not flexible enough to achieve efficient tracking of temporal changes in neural recordings [7]. In fact, a simple power spectrum gives no temporal information about the features of the signal. The necessity for a different approach is apparent. The solution comes in the form of spectrotemporal analysis (or time-frequency analysis), in the simplest form of which, the spectral characteristics of the signal are observed over time.

A staple of spectrotemporal analysis and a logical next step is the Short Time Fourier Transform, in which the signal is segmented in windows of short length and the Fourier Transform of each window is then computed, thus effectively creating a spectrum as a function of time, the magnitude of which is visualized in the *spectrogram*, an example of which is shown in Figure C.3.

Nevertheless, the STFT is still bound by the limitations set by the trigonometric basis functions of the Fourier Transform. It would be convenient if more flexible basis functions could be used, better suited for fitting in any kind of signal better than simple sines and cosines. This is where *wavelets* offer an elegant and versatile solution.

Wavelets are functions whose energy is concentrated in a short time window. Two very important criteria must be met for a function to qualify as a wavelet: those of finite energy and admissibility, the latter of which guarantees the band-pass properties of the wavelet function [29]. What makes wavelets so versatile is the fact that they can be *scaled* and *translated* in time, based on user-assigned parameters. With this in mind, it can be surmised that with these two degrees of freedom wavelets achieve good localization in time and frequency, so that they can cover both coarse and finer details of the signal. This quality makes them excellent candidates for use in *multiresolution analysis* [7].

Spectrogram of epileptic neural signal
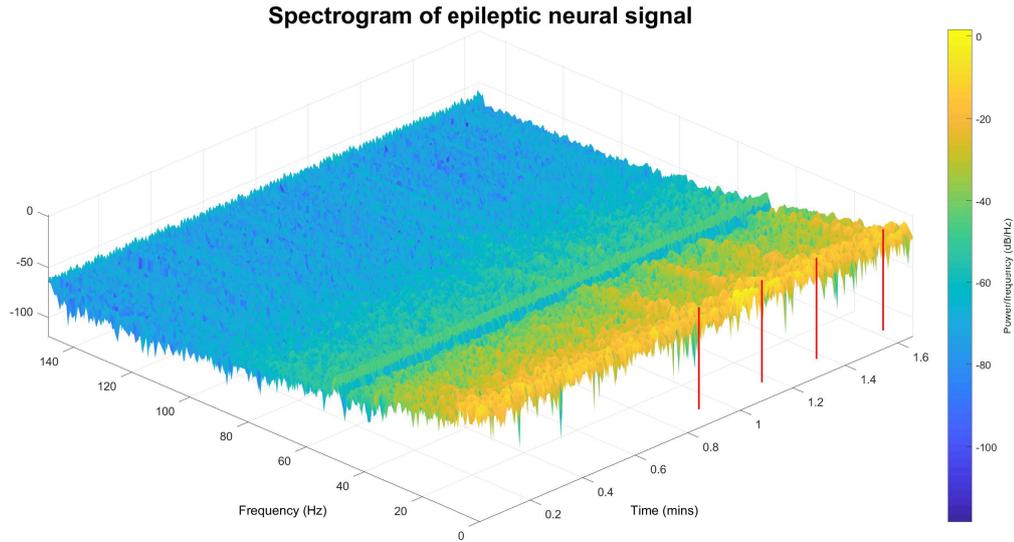
Figure C.3: The spectrogram of an LFP contaminated by epileptic seizures. The onset and duration of the seizures is visible in the spectrogram, as the frequency content of the signal increases for short periods of time and it is denoted with red vertical lines. What is more, this particular signal is contaminated by mains hum at $50Hz$, which is visible in the entire duration of the signal as a persistent peak in the spectrogram.

The basic wavelet subjected to scaling and translation to yield the family of wavelets to be used in the analysis is named the *mother wavelet* (usually denoted with $\psi(t)$). The scaling of the mother wavelet reveals an inverse relation between localization in time and localization in frequency, in the sense that a smaller scale of the mother wavelet is more localized in time, its bandwidth increases, rendering it less localized in frequency and vice versa.

Much like the Fourier Transform, the *Continuous Wavelet Transform* is given by convolving the signal under scrutiny with scaled and translated versions of the mother wavelet, with the signal being exactly recoverable from its transform through the *reconstruction equation*. By discretizing the scaling and translation parameters, the *Discrete Wavelet Transform* is obtained. The signal is recovered by the inverse DWT, also known as the *wavelet series expansion*.

The general form of the convolution integral that yields the series expansion of a function using orthonormal basis functions is:

$$w_k = \int_{-\infty}^{+\infty} x(t)\phi_k(t)dt \tag{C.12}$$

In wavelet analysis, the basis functions ϕ_k are replaced by a family of wavelets:

$$\psi_{s,\tau}(t) = \frac{1}{\sqrt{s}}\psi\left(\frac{t-\tau}{s}\right) \tag{C.13}$$

where ψ is the mother wavelet and s and τ are the scaling and translation parameters respectively. Thus, inserting C.13 into C.12 yields the Continuous Wavelet Transform:

$$w(s,\tau) = \int_{-\infty}^{+\infty} x(t)\frac{1}{\sqrt{s}}\psi\left(\frac{t-\tau}{s}\right)dt \tag{C.14}$$

which is, in effect, a mapping on the time-scale domain. Equation C.14 can also be regarded as the linear filtering of the signal $x(t)$ with a filter of impulse response $\frac{\psi(-t/s)}{\sqrt{s}}$, which showcases the band-pass qualities of the CWT. However, the CWT function is highly redundant, so the next logical step is a discretization of the

parameters to reduce the amount of data provided by the CWT. The scaling and translation parameters are dyadically sampled:

$$s = 2^{-j} \qquad \tau = k2^{-j} \tag{C.15}$$

which, used in C.13, yields the discretized wavelet function:

$$\psi_{j,k} = 2^{j/2}\psi(2^j t - k) \tag{C.16}$$

By inserting C.16 into C.14, the discrete wavelet transform is obtained:

$$w_{j,k} = \int_{-\infty}^{\infty} x(t)\psi_{j,k}(t)dt \tag{C.17}$$

Multiresolution analysis is based upon the assumption that the signal can be broken down to two parts: a coarse and a fine part, representing the main features and the details of the signal respectively. Here, the concept of the *scaling function* (also known as the *father wavelet*) must be defined. Both the scaling and wavelet functions are necessary for wavelet analysis. The difference between the scaling and wavelet functions is the fact that the former approximates the main features of the signal, low frequency content in other words, while the latter deals with the high-frequency, finer details of the signal.

A scaling function related to a particular wavelet function can be used for the production of an orthonormal set of scaling functions at different translations in time.

$$\phi_{0,k}(t) = \phi(t - k) \tag{C.18}$$

These scaling functions $\phi_{0,k}$ span a subspace $\mathcal{V}_0$. Different scales of $\phi_{j,k}$ will span different subspaces $\mathcal{V}_j$, which are nested into each other:

$$\dots \mathcal{V}_{n-2} \subset \mathcal{V}_{n-1} \subset \mathcal{V}_n \subset \mathcal{V}_{n+1} \subset \mathcal{V}_{n+2} \dots \tag{C.19}$$

with the level of detail increasing with n. Consecutive scales of ϕ (designated as $\phi(t)$ and $\phi(2t)$) are related through the *refinement equation*:

$$\phi(t) = \sum_{n=-\infty}^{\infty} h_\phi(n)\phi_{1,n}(t) = \sqrt{2} \sum_{n=-\infty}^{\infty} h_\phi(n)\phi(2t - n) \tag{C.20}$$

where h_ϕ is a series of *scaling coefficients*.

The wavelet functions ψ span different subspaces $\mathcal{W}$, that account for the finer features of the signal. The relation between the subspaces is:

$$\mathcal{V}_j = \mathcal{V}_{j-1} \oplus \mathcal{W}_{j-1} \tag{C.21}$$

or, using a recursion until a j_0 value:

$$\mathcal{V}_{j+1} = \mathcal{V}_{j_0} \oplus \mathcal{W}_{j_0} \oplus \mathcal{W}_{j_0+1} \oplus \dots \oplus \mathcal{W}_j \tag{C.22}$$

In signal terms, this can be expressed as:

$$x(t) = x_{j_0}(t) + \sum_{j=j_0}^{\infty} y_j(t) \tag{C.23}$$

where x_{j_0} is the coarse approximation of the signal and y_j is the "detail" signal at a different resolution for each value of j.

In simple terms, wavelet multiresolution analysis on a signal yields a superposition of scaled and translated versions of the mother wavelet, approximating the original signal as closely as possible. Finer details of the signal, which are, obviously, characterized by a broader bandwidth, are represented by smaller scales of the mother wavelet, while slower transient events at low frequency bands by larger scales. In effect, the signal is broken down to its frequency components, which are easy to visualize, while whatever frequencies are not desired can be easily removed. Herein lie the filtering and denoising capabilities of the wavelet transform, in the sense that smaller scales of the mother wavelet can be discarded, effectively removing high frequency

Mother Wavelet Examples
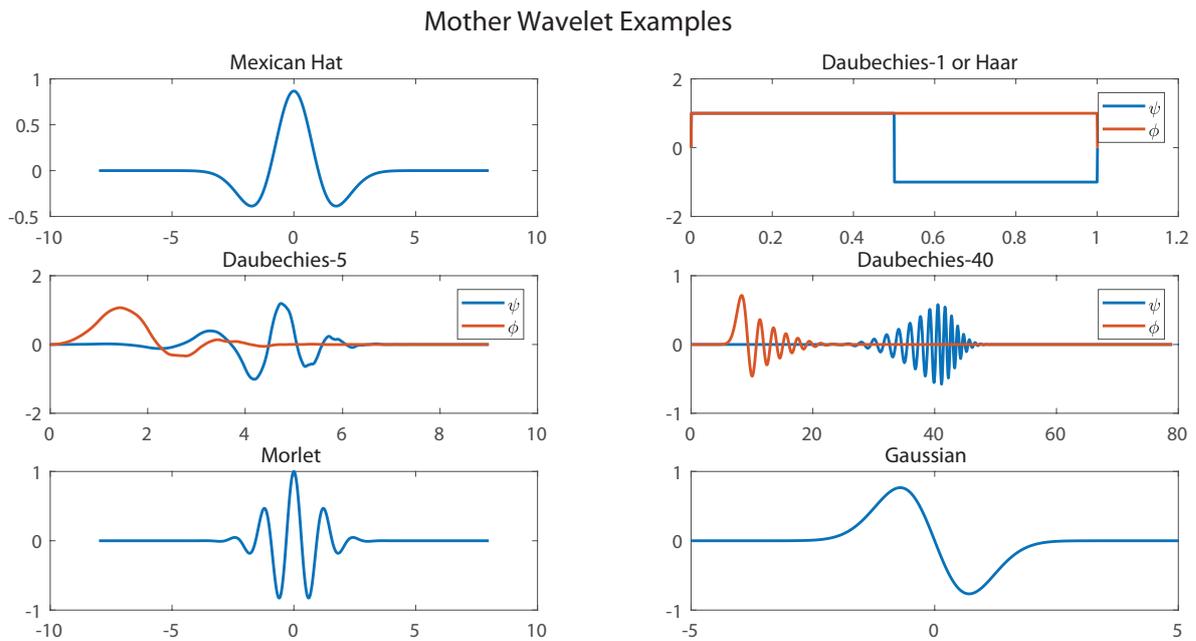
Figure C.4: Mother Wavelet Examples. The scaling function is shown in red where it can be defined

noise from the recording, without affecting any lower frequencies. The breakdown of the signal into scaled and translated versions of $\psi(t)$ facilitates the detection and extraction of features from the signal, since features of different frequency and power will be represented by different scales of the mother wavelet.

An abundance of wavelet functions has been proposed over the years. Examples include the Mexican Hat, Daubechies, Morlet, Coiflets, Gaussian and Shannon wavelets. It should be noted also that a scaling function can not be defined for all wavelet families. Some of them are displayed in Figure C.4.

Wavelets are a highly versatile tool that is being used extensively in research. Examples include their use for various types of LFP and EEG analysis [51, 68–72] as a precursor to ICA [73] and for detection of clinically significant events in the neural recordings [29, 48, 52]. It has been stated in [54] that wavelet analysis is much better suited for the examination of single trial evoked potentials than more traditional methods such as Wiener filtering or LMS. The multiresolution analysis ability of wavelets also allows for effective feature extraction from the signal, by matching scaled versions of the mother wavelet to the various components contained in the signal [29, 50].

It appears that the most popular type of wavelet is the Morlet wavelet, whose default version is used by a large number of researchers [48, 50, 52, 71], while some of them prefer its complex version [51, 69, 72]. The Daubechies, Symlets, Coiflets, Meyer and Mallat wavelets are also used in certain cases [54, 70]. Despite the popularity of the Morlet wavelet, Rafiee et al. in [67] claim that, after a comparative analysis, that the Daubechies-44 wavelet is the best suited for the analysis of various biological signals, including the EEG, because of its ability to fit into sharp asymmetric spikes, which feature prominently in neural recordings.

C.5. PRINCIPAL AND INDEPENDENT COMPONENT ANALYSIS

Most usually, neural data comprise of the simultaneous effect of a number of sources, including noise. The superposition of these effects is the raw signal recorded from the brain. There is obvious merit in reducing the dimensionality of the signal, while at the same time removing those signal components that are of no use to further analysis, such as noise. All this can be managed with the use of Principal Component Analysis.

In its core, PCA assists in the creation of a set of uncorrelated signal components which, when properly combined, satisfactorily approximate the original signal. The analysis begins by assuming that the signal under

scrutiny is represented by a $K \times 1$ vector **y**, which can be written as a deterministic linear transformation as:

$$\mathbf{y} = \mathbf{Wx} = \sum_{k=1}^{K} x_k \mathbf{W}_k \tag{C.24}$$

where

$$\mathbf{W} = [\mathbf{w}_1 \mathbf{w}_2 \dots \mathbf{w}_K] \tag{C.25}$$

with $\mathbf{w}_k$ being $K \times 1$ column vectors. Also the expression $|W| \neq 0$ must hold. The $K \times 1$ vector **x** is a reduced approximation of the signal **y**, so the matrix **W** must be formulated in such a way so that the relation in Equation C.24 holds. Primarily, the vectors $\mathbf{W}_k$ must constitute an set of orthonormal basis functions, so:

$$\mathbf{W}^T \mathbf{W} = \mathbf{I} \quad \Leftrightarrow \quad \mathbf{W}^T = \mathbf{W}^{-1} \tag{C.26}$$

Since the matrix W is orthonormal, then the transformation is reversible:

$$x = W^{-1} y = \sum_{k=1}^{K} y_k W_k^{-1} \tag{C.27}$$

If all the components in **x** are used, then the reversible transformation provides a lossless representation of **y**. The objective, however, was the reduction of the signal's dimensionality. To that end, only $L < K$ of the components can be taken into consideration. The signal **y** can be approximated as:

$$\tilde{\mathbf{y}} = \sum_{k=1}^{L} x_k \mathbf{W}_k + \sum_{k=L+1}^{K} \tilde{x}_k \mathbf{W}_k \tag{C.28}$$

where $\tilde{x}_k$ are the components that will be omitted in the reconstruction of the signal. The expression in C.28 introduces an error of:

$$\varepsilon = \mathbf{y} - \tilde{\mathbf{y}} = \sum_{k=L+1}^{K} (x_k - \tilde{x}_k) \mathbf{W}_k \tag{C.29}$$

which, in turn, yields an Mean Square Error of:

$$\bar{\varepsilon^2} = E[\bar{\varepsilon}^T \bar{\varepsilon}] = \sum_{k=L+1}^{K} (x_k - \tilde{x}_k)^2 \tag{C.30}$$

The minimum MSE can be computed by setting its derivative equal to zero:

$$\frac{\partial \bar{\varepsilon^2}}{\partial \tilde{x}_k} = -2E[(x_k - \tilde{x}_k)] = 0 \tag{C.31}$$

The relation is satisfied if:

$$\tilde{x}_k = E[x_k] \xrightarrow{C.27} W^{-1} E[\mathbf{y}] \tag{C.32}$$

which means that the omitted components can be replaced by their mean. The minimum MSE can be rewritten as:

$$\begin{aligned}
\bar{\varepsilon^2} &= \sum_{k=L+1}^{K} E[(x_k - \bar{x}_k)^2] \\
&= \sum_{k=L+1}^{K} E\left[(W_k^{-1} y_k - W_k^{-1} \tilde{y}_k)(W_k^{-1} y_k - W_k^{-1} \tilde{y}_k)^T\right] \\
&= \sum_{k=L+1}^{K} E\left[W_k^{-1}(y_k - \tilde{y}_k)(y_k - \tilde{y}_k)^T \left(W_k^{-1}\right)^T\right] \xLeftarrow{\quad W \text{ orthonormal} \quad} \\
&= \sum_{k=L+1}^{K} W_k^T \Sigma_y W_k
\end{aligned} \tag{C.33}$$

where Σ_y is the covariance matrix of the signal y, with dimensions $K \times K$. By selecting the W_k vectors as the eigenvectors of the covariance matrix:

$$\boldsymbol{\sigma}_y W_k = \lambda_k W_k, \quad where \quad \lambda_k = W_k^T \boldsymbol{\sigma}_y W_k \tag{C.34}$$

from C.33 it emerges that:

$$\bar{\varepsilon^2}_{min} = \sum_{k=L+1}^{K} \lambda_k \qquad \text{(C.35)}$$

From C.35, it can be concluded that it is useful to order the eigenvalues in descending order so that the MSE is minimized, in the sense that when the kth component x_k is replaced by $\tilde{x}_k$, the MSE will increase by λ_k. In other words, components of the signal related to larger eigenvalues play a larger role in the shaping of the signal and their omission would bring about a larger increase in the MSE. Smaller eigenvalues correspond to less important components that can be removed without causing a deterioration in the reconstructed signal. In this manner, the amount of data to be analyzed further can be greatly reduced by the correct selection for the value of L, while maintaining a satisfactory approximation of the original form of the data.

On one hand, PCA breaks the signal down to uncorrelated components and is based on second-order statistics, as it uses the covariance matrix for analyzing the signal. On the other hand, Independent Component Analysis separates the signal in uncorrelated but also *statistically* independent components. Assuming that K channels of a signal are produced by L sources and that $K \geq L$ (meaning that the recordings are redundant), the signal y can be expressed as:

$$\boldsymbol{y}(n) = \boldsymbol{M}\boldsymbol{x}(n) + \eta(n) \approx \boldsymbol{M}\boldsymbol{x}(n) \qquad \text{(C.36)}$$

where $\boldsymbol{y}$ and $\boldsymbol{x}$ are a $K \times 1$ and a $L \times 1$ vector respectively and $\boldsymbol{M}$ is a $M \times L$ matrix, while η is additive noise, independent of the signals, which for the sake of simplicity is considered to be zero. $\boldsymbol{M}$ is called a *mixing matrix*, which is unknown. The goal of ICA is the calculation of the values of the source vector $\boldsymbol{x}$ from the vector $\boldsymbol{y}$, without knowing the matrix $\boldsymbol{M}$, effectively making this procedure a problem of *Blind Source Separation*. So, the objective is the search of a matrix that will satisfy the reverse relation of C.36:

$$\tilde{\boldsymbol{x}}(n) = \boldsymbol{W}\boldsymbol{y}(n) \qquad \text{(C.37)}$$

The matrix $\boldsymbol{W}$ is called an *unmixing* or *separation matrix* and has a size of $L \times K$ [29].

Besides the above, like PCA, ICA can be used as a tool for artifact and noise removal, as seen in [37], where ICA is used for eye movement artifact removal. Also, in [45], ICA removes all kinds of artifacts from Event Related Potential EEG data, as well as implementing blind source separation. EEG artifact removal is also conducted with the use of ICA in [74] and [73], where it is additionally used for BSS, and also in [75], in which it is recommended that a large number of electrodes is used for good quality ICA. BSS is conducted with ICA also in [76].

There is a number of algorithms for the implementation of BSS. The most popular in neural signal analysis are *FastICA*, *JADE*, *InfoMax*, and *SOBI* [77]. FastICA, JADE and InfoMax exploit the signal's High Order Statistics, which means that ICA can be implemented with these methods. SOBI uses the signal's Second Order Statistics. In general, SOS methods require a much smaller number of samples to function correctly, compared to HOS methods, which means they can perform satisfactorily in shorter time intervals. This diminishes the possibility that the time intervals contain nonstationarities, which, as stated above, is something to be avoided in this kind of analysis [78].

PCA can also be used for source separation (also called *cell sorting* in literature) in LFP recordings [68, 76]. It is argued in [42] that in ERP analyses, it is unwise to use *temporal* PCA due to its inability to handle different latencies between stimulus and response, but highly recommend *spatial* PCA, between recordings of different scalp electrodes, while the authors of [39] freely use spatiotemporal PCA on EEG data for the identification of the Error Related Negativity. Also, PCA can be applied on the features of the signal, besides the signal itself, serving different goals. PCA on a covariance can be used for component identification, whereas component measurement can be achieved by applying PCA on a correlation matrix [79].

C.6. UNITS USED IN SINGLE CELL RECORDINGS

In the case of single cell recordings, the use of quantitative measures is typical. Besides all the usual signal analysis methods that can be applied on the waveforms, two more measures are important: the inter-spike interval (ISI) and the firing rate of the cell. Since this type of recording visualizes the activity of one single cell, these two units are sufficient for the monitoring of the cell's activity, while the behavior of the spikes has been

shown to be causally linked to the presence of neurotransmitters [80].

D

PROOF FOR THE DERIVATION OF MATCHED FILTER TRANSFER FUNCTION

This proof and further information on Matched Filters and their potential applications on Biomedical Signal Analysis can be found in [29].

Consider a signal $x(t)$, with a Fourier transform of $X(\omega) = \int_{-\infty}^{+\infty} x(t)exp(-j\omega t)dt$, is the input to a matched filter with a (for now unknown) transfer function $H(\omega)$. The time domain output $y(t)$ of the filtering process is given by the inverse Fourier Transform of $Y(j\omega) = X(j\omega)H(j\omega)$:

$$
\begin{aligned}
y(t) &= \int_{-\infty}^{+\infty} X(f)H(f)exp(+j2\pi ft)df \\
&= \frac{1}{2\pi} \int_{-\infty}^{+\infty} X(\omega)H(\omega)exp(+j\omega t)d\omega
\end{aligned}
\tag{D.1}
$$

The frequencies will be expressed in Hertz instead of radians per second from this point on.

Suppose that there is white noise at the input, with PSD:

$$
S_{\eta i}(f) = \frac{P_{\eta i}}{2}
\tag{D.2}
$$

where $P_{\eta i}$ is the average noise power at the input, while as an output we have:

$$
S_{\eta o}(f) = \frac{P_{\eta i}}{2} \left| H(f) \right|^2
\tag{D.3}
$$

whose power is:

$$
P_{\eta o} = \frac{P_{\eta i}}{2} \int_{-\infty}^{+\infty} \left| H(f) \right|^2 df
\tag{D.4}
$$

the square root of which is the RMS value of the power of noise in the absence of any signal.

At any given time instant $t = t_0$, per Eq. D.1, the magnitude of the output is:

$$
M_y = \left| y(t_0) \right| = \left| \int_{-\infty}^{+\infty} X(f)H(f)exp(+j2\pi ft_0)df \right|
\tag{D.5}
$$

Thus, the SNR can be defined as:

$$
SNR = \frac{M_y}{\sqrt{P_{\eta o}}}
\tag{D.6}
$$

itself a function of $H(f)$, so the problem of the derivation of the optimal transfer function for the matched filter is, in fact, an SNR maximization problem. The *peak-power SNR* can be defined as:

$$
\frac{M_y^2}{P_{\eta o}} = \frac{instantaneous\ power\ of\ signal}{noise\ mean\ power}
\tag{D.7}
$$

The total energy of the signal $x(t)$ is constant and equal to:

$$E_x = \int_{-\infty}^{+\infty} x^2(t)\,dt = \int_{-\infty}^{+\infty} |X(f)|^2\,dt \qquad \text{(D.8)}$$

From all the above, the following ratio can be formulated:

$$\frac{M_y^2}{E_x P_{\eta o}} = \frac{\left|\int_{-\infty}^{+\infty} H(f)X(f)exp(+j2\pi f t_0)df\right|^2}{\frac{P_{\eta i}}{2}\int_{-\infty}^{+\infty}|H(f)|^2 df \int_{-\infty}^{+\infty}|X(f)|^2 df} \qquad \text{(D.9)}$$

and since E_x is constant, maximizing this expression equates maximizing the expression in Eq. D.7. Evoking the Schwarz Inequality, it can be rewritten as:

$$\frac{M_y^2}{E_x P_{\eta o}} \leq \frac{\int_{-\infty}^{+\infty}|H(f)|^2 df \int_{-\infty}^{+\infty}|X(f)exp(j2\pi f t_0)|^2 df}{\frac{P_{\eta i}}{2}\int_{-\infty}^{+\infty}|H(f)|^2 df \int_{-\infty}^{+\infty}|X(f)|^2 df} \Leftrightarrow$$

$$\frac{M_y^2}{E_x P_{\eta o}} \leq \frac{\int_{-\infty}^{+\infty}|H(f)|^2 df \int_{-\infty}^{+\infty}|X(f)|^2 df \overset{1}{\int_{-\infty}^{+\infty}|exp(j2\pi f t_0)|^2 df}}{\frac{P_{\eta i}}{2}\int_{-\infty}^{+\infty}|H(f)|^2 df \int_{-\infty}^{+\infty}|X(f)|^2 df} \Leftrightarrow \qquad \text{(D.10)}$$

$$\frac{P_{\eta i} M_y^2}{2 E_x P_{\eta o}} \leq 1$$

The equality holds if the two functions are *collinear*, so there is a real K that satisfies $A(f) = K \times B(f)$, in other words:

$$\begin{aligned} H(f) &= K[X(f)exp(j2\pi f t_0)] \\ &= K[X(f)exp(-j2\pi f t_0)] \end{aligned} \qquad \text{(D.11)}$$

and, eventually, the inverse Fourier transform of this expression, yields the impulse response of the matched filter:

$$h(t) = Kx(-(t-t_0)) \qquad \text{(D.12)}$$

which is a scaled and reversed version of the signal, shifted in time, proving that the design of matched filters is absolutely based on the signal at hand.

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my thesis supervisor, prof.dr.ir. Alle-Jan van der Veen, for the trust that he placed in me to carry out this project and for the invaluable guidance he offered me these past months.

I owe a big thank you to my daily supervisor, ir. Farnaz Nassirinia, both for the amount of knowledge that I gained during my stay at the Neuroscience Department of Erasmus MC and for the fun we had while doing so. Soldering electrodes will never feel the same from now on.

A huge thank you to my family for their support and love and for believing in me and helping me make all this possible.

To all my friends here in Delft, Albert, María, Mauricio, Andrés, Marco, Alba, Su, Virginia, Yuun and so many others, thank you for the mutual support in this stressful time for all of us. The past two years have been unforgettable, I could not have hoped for better and I'm wishing you all the best of luck in the years to come.

My friends from back home, Panagiotis, Natassa, Makis, Pantelis, Eva, Christine, Domna, Kater, Anasta, Tony, Nick, Constantinos, Dimitris, Alex, John, Prodromos, I miss you all and I'm looking back to all the things we did with fondness and nostalgia. I'm looking forward to the next time we meet.

My labmates, Bas G, Bas W, Angelo, Rene, George, Sadaf, we all made the lab an amazing place to work. There wasn't a day I wasn't looking forward to taking the train to the lab and I have you to thank for that.

BIBLIOGRAPHY

[1] E. N. Marieb and K. Hoehn, *Human Anatomy & Physiology*, 9th ed. (Pearson Education Inc., 2013).

[2] A. C. Guyton and J. E. Hall, *Textbook of Medical Physiology*, 11th ed. (Elsevier Saunders, Philadelphia, 2006).

[3] N. K. Speed, *Journal of Chemical Information and Modeling*, Ph.D. thesis, Vanderbilt University (2013), arXiv:arXiv:1011.1669v3 .

[4] J. Malmivuo and R. Plonsey, *Bioelectromagnetism: Principles and Applications of Bioelectric and Biomagnetic Fields* (Oxford University Press, New York, 1995).

[5] T. Blakely, K. J. Miller, R. P. N. Rao, M. D. Holmes, and J. G. Ojemann, *Localization and classification of phonemes using high spatial resolution electrocorticography (ecog) grids,* 2008 30th Annual International Conference of the IEEE Engineering in Medicine and Biology Society , 4964 (2008).

[6] G. Paxinos and K. Franklin, *Paxinos and Franklin's the Mouse Brain in Stereotaxic Coordinates*, 4th ed. (Academic Press, 2012) p. 360.

[7] L. Sörnmo and P. Laguna, *Biomedical Signal Processing: Electrical Signals in Cardiac & Neurological Applications* (Elsevier Academic Press, 2004).

[8] P. Uhlhaas, G. Pipa, B. Lima, L. Melloni, S. Neuenschwander, D. Nikolic, and W. Singer, *Neural synchrony in cortical networks: history, concept and current status,* Frontiers in Integrative Neuroscience **3**, 3662 (2009).

[9] K. A. Montagu, *Catechol Compounds in Rat Tissues and in Brains of Different Animals,* Nature **180**, 244 (1957).

[10] F. M. Benes, *Carlsson and the discovery of dopamine,* Trends in Pharmacological Sciences **22**, 46 (2001).

[11] J. K. Seamans and C. R. Yang, *The principal features and mechanisms of dopamine modulation in the prefrontal cortex,* Progress in Neurobiology **74**, 1 (2004).

[12] M. Marinelli and J. McCutcheon, *Heterogeneity of dopamine neuron activity across traits and states,* Neuroscience **282**, 176 (2014), arXiv:NIHMS150003 .

[13] *American Electroencephalographic Society Guidelines for Standard Electrode Position Nomenclature,* Journal of Clinical Neurophysiology **8**, 200 (1991).

[14] M. Hämäläinen, R. Hari, R. J. Ilmoniemi, J. Knuutila, and O. V. Lounasmaa, *Magnetoencephalography—theory, instrumentation, and applications to noninvasive studies of the working human brain,* Rev. Mod. Phys. **65**, 413 (1993).

[15] G. Buzsáki, C. A. Anastassiou, and C. Koch, *The origin of extracellular fields and currents — EEG, ECoG, LFP and spikes,* Nature Reviews Neuroscience **13**, 407 (2012).

[16] S. R. Soekadar, M. Witkowski, E. G. Cossio, N. Birbaumer, S. E. Robinson, and L. G. Cohen, *In vivo assessment of human brain oscillations during application of transcranial electric currents,* Nature Communications **4** (2013), 10.1038/ncomms3032.

[17] G. Schalk and J. Mellinger, *A Practical Guide to Brain–Computer Interfacing with BCI2000* (Springer London, London, 2010).

[18] C. Bédard and A. Destexhe, *Local Field Potential Interaction with the Extracellular Medium,* in *Encyclopedia of Computational Neuroscience*, edited by D. Jaeger and R. Jung (Springer New York, New York, NY, 2015) pp. 1540–1547, arXiv:1501.01854 .

[19] A. Conolly, A. Muralidharan, C. Hendrix, L. Johnson, R. Gupta, S. Stanslaski, T. J. Denison, K. B. Baker, J. L. Vitek, and M. D. Johnson, *Local field potential recordings in a non-human primate model of Parkinsons disease using the Activa PC + S neurostimulator Allison,* Journal of Neural Engeneering **12,** 1 (2016).

[20] A. A. Kühn, T. Trottenberg, A. Kivi, A. Kupsch, G.-H. Schneider, and P. Brown, *The relationship between local field potential and neuronal discharge in the subthalamic nucleus of patients with Parkinson's disease.* Experimental neurology **194,** 212 (2005).

[21] N. J. Ray, N. Jenkinson, S. Wang, P. Holland, J. S. Brittain, C. Joint, J. F. Stein, and T. Aziz, *Local field potential beta activity in the subthalamic nucleus of patients with Parkinson's disease is associated with improvements in bradykinesia after dopamine and deep brain stimulation,* Experimental Neurology **213,** 108 (2008).

[22] J. P. Donoghue, J. N. Sanes, N. G. Hatsopoulos, and G. Gaál, *Neural Discharge and Local Field Potential Oscillations in Primate Motor Cortex During Voluntary Movements,* Journal of Neurophysiology **79,** 159 (1998).

[23] A. Destexhe, D. Contreras, and M. Steriade, *Spatiotemporal analysis of local field potentials and unit discharges in cat cerebral cortex during natural wake and sleep states.* The Journal of neuroscience : the official journal of the Society for Neuroscience **19,** 4595 (1999).

[24] M. Rasch, N. K. Logothetis, and G. Kreiman, *From Neurons to Circuits: Linear Estimation of Local Field Potentials,* Journal of Neuroscience **29,** 13785 (2009).

[25] G. T. Einevoll, C. Kayser, N. K. Logothetis, and S. Panzeri, *Modelling and analysis of local field potentials for studying the function of cortical circuits,* Nature Reviews Neuroscience **14,** 770 (2013), arXiv:NIHMS150003 .

[26] A. A. Grace and B. S. Bunney, *The control of firing pattern in nigral dopamine neurons: burst firing.* The Journal of neuroscience : the official journal of the Society for Neuroscience **4,** 2877 (1984).

[27] A. A. Grace and B. S. Bunney, *The control of firing pattern in nigral dopamine neurons: single spike firing,* The Journal of Neuroscience **4,** 2866 (1984).

[28] A. A. Grace and B. S. Bunney, *Intracellular and extracellular electrophysiology of nigral dopaminergic neurons-1. Identification and characterization,* Neuroscience **10** (1983), 10.1016/0306-4522(83)90135-5.

[29] R. M. Rangayyan, *Biomedical Signal Analysis,* 2nd ed. (John Wiley & Sons, New Jersey, 2015).

[30] G. Buzsáki and A. Draguhn, *Neuronal Oscillations in Cortical Networks,* Science **304,** 1926 (2004).

[31] K. Benchenane, P. H. Tiesinga, and F. P. Battaglia, *Oscillations in the prefrontal cortex: A gateway to memory and attention,* Current Opinion in Neurobiology **21,** 475 (2011).

[32] R. W. Friedrich, C. J. Habermann, and G. Laurent, *Multiplexing using synchrony in the zebrafish olfactory bulb,* Nature Neuroscience **7,** 862 (2004).

[33] A. K. Engel, P. Fries, and W. Singer, *Dynamic predictions: Oscillations and synchrony in top–down processing,* Nature Reviews Neuroscience **2,** 704 (2001), arXiv:NIHMS150003 .

[34] M. Ding, S. L. Bressler, W. Yang, and H. Liang, *Short-window spectral analysis of cortical event-related potentials by adaptive multivariate autoregressive modeling : data preprocessing , model validation , and variability assessment,* Biological Cybernetics **83,** 35 (2000).

[35] V. Vorobyov, N. Schibaev, G. Kovalev, and C. Alzheimer, *Effects of neurotransmitter agonists on electrocortical activity in the rat kainate model of temporal lobe epilepsy and the modulatory action of basic fibroblast growth factor,* Brain Research **1051,** 123 (2005).

[36] M. J. Lehmkuhle, S. S. Bhangoo, and D. R. Kipke, *The electrocorticogram signal can be modulated with deep brain stimulation of the subthalamic nucleus in the hemiparkinsonian rat.* Journal of neurophysiology **102,** 1811 (2009).

[37] D. L. Santesso, A. E. Evins, M. J. Frank, E. C. Schetter, R. Bogdan, and D. A. Pizzagalli, *Single dose of a dopamine agonist impairs reinforcement learning in humans: Evidence from event-related potentials and computational modeling of striatal-cortical function,* Human Brain Mapping **30**, 1963 (2009).

[38] C. B. Holroyd Nieuwenhuis, S., Yeung N. and Cohen, J.D., *Errors in reward prediction are reflected in the event-related brain potential.* Cognitive Neuroscience and Neuropsychology **14**, 4 (2003).

[39] D. Foti and G. Hajcak, *Depression and reduced sensitivity to non-rewards versus rewards: Evidence from event-related potentials,* Biological Psychology **81**, 1 (2009).

[40] D. Foti and G. Hajcak, *Genetic variation in dopamine moderates neural response during reward anticipation and delivery: Evidence from event-related potentials,* Psychophysiology **49**, 617 (2012).

[41] C. B. Holroyd and M. G. Coles, *The neural basis of human error processing: Reinforcement learning, dopamine, and the error-related negativity,* Psychological Review **109**, 679 (2002), arXiv:NIHMS150003 .

[42] C. B. Holroyd, K. L. Pakzad-Vaezi, and O. E. Krigolson, *The feedback correct-related positivity: Sensitivity of the event-related brain potential to unexpected positive feedback,* Psychophysiology **45**, 688 (2008).

[43] H. Bokura, S. Yamaguchi, and S. Kobayashi, *Event-related potentials for response inhibition in Parkinson's disease,* Neuropsychologia **43**, 967 (2005).

[44] P. J. Charlety, J. Grenhoff, K. Chergui, B. De la Chapelle, M. Buda, T. H. Svensson, and G. Chouvet, *Burst firing of mesencephalic dopamine neurons is inhibited by somatodendritic application of kynurenate,* Acta Physiologica Scandinavica **142**, 105 (1991).

[45] T.-P. Jung, S. Makeig, M. Westerfield, J. Townsend, E. Courchesne, and T. J. Sejnowski, *Analysis and visualization of single-trial event-related potentials,* Human Brain Mapping **14**, 166 (2001).

[46] G. Gratton, M. G. H. Coles, and E. Donchin, *A new method for off-line removal of ocular artifact,* Electroencephalography and Clinical Neurophysiology **55**, 468 (1983).

[47] C. L. Ehlers, T. L. Wall, and R. I. Chaplin, *Long latency event-related potentials in rats: Effects of dopaminergic and serotonergic depletions,* Pharmacology, Biochemistry and Behavior **38**, 789 (1991).

[48] S. A. Shimamoto, E. S. Ryapolova-Webb, J. L. Ostrem, N. B. Galifianakis, K. J. Miller, and P. A. Starr, *Subthalamic Nucleus Neurons Are Synchronized to Primary Motor Cortex Local Field Potentials in Parkinson's Disease,* Journal of Neuroscience **33**, 7220 (2013).

[49] H. Zhang, R. Chavarriaga, M. K. Goel, L. Gheorghe, and J. D. R. Millan, *Improved recognition of error related potentials through the use of brain connectivity features,* Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS , 6740 (2012).

[50] J. R. Manning, J. Jacobs, I. Fried, and M. J. Kahana, *Broadband Shifts in Local Field Potential Power Spectra Are Correlated with Single-Neuron Spiking in Humans,* Journal of Neuroscience **29**, 13613 (2009).

[51] T. H. Sanders, A. Devergnas, T. Wichmann, and M. A. Clements, *Canonical correlation to estimate the degree of parkinsonism from local field potential and electroencephalographic signals,* in *2013 6th International IEEE/EMBS Conference on Neural Engineering (NER)*, Vol. 344 (IEEE, 2013) pp. 158–161, arXiv:15334406 .

[52] M. Le Van Quyen, A. Bragin, R. Staba, B. Crepon, C. L. Wilson, and J. Engel, *Cell Type-Specific Firing during Ripple Oscillations in the Hippocampal Formation of Humans,* Journal of Neuroscience **28**, 6104 (2008), arXiv:NIHMS150003 .

[53] A. Iwata, N. Suzumura, and K. Ikegaya, *Pattern Classificaiton of the PCG using linear prediction analysis - Iwata et al,* Medical and Biological Engineering and Computing **15**, 407 (1977).

[54] Z. Wang, A. Maier, D. A. Leopold, N. K. Logothetis, and H. Liang, *Single-trial evoked potential estimation using wavelets,* Computers in Biology and Medicine **37**, 463 (2007).

[55] T. Milekovic, T. Ball, A. Schulze-Bonhage, A. Aertsen, and C. Mehring, *Error-related electrocorticographic activity in humans during continuous movements,* Journal of Neural Engineering **9** (2012), 10.1088/1741-2560/9/2/026007.

[56] J. Biskamp, M. Bartos, and J. F. Sauer, *Organization of prefrontal network activity by respiration-related oscillations,* Scientific Reports **7**, 1 (2017).

[57] B. L. Lewis and P. O'Donnell, *Ventral Tegmental Area Afferents to the Prefrontal Cortex Maintain Membrane Potential 'Up' States in Pyramidal Neurons via D1 Dopamine Receptors,* Cerebral Cortex **10**, 1168 (2000).

[58] S. Noachtar and J. Rémi, *The role of EEG in epilepsy: A critical review,* Epilepsy & Behavior **15**, 22 (2009).

[59] C. Strydis and R. M. Seepers, *Trading detection accuracy for battery autonomy in a wearable seizure-detection device,* 1st Workshop On Approximate Computing (WAPCO 2015) , 2 (2015).

[60] G. Koob, M. Moal, and R. Thompson, *Encyclopedia of Behavioral Neuroscience,* Science direct (Elsevier Science, 2010).

[61] I. Pavlov and G. Anrep, *Conditioned Reflexes* (Dover Publications, 2003).

[62] R. Bourdy and M. Barrot, *A new control center for dopaminergic systems: Pulling the VTA by the tail,* Trends in Neurosciences **35**, 681 (2012), arXiv:NIHMS150003 .

[63] P. Welch, *The use of fast Fourier transform for the estimation of power spectra: A method based on time averaging over short, modified periodograms,* IEEE Transactions on Audio and Electroacoustics **15**, 70 (1967).

[64] M. S. Bartlett, *Smoothing Periodograms from Time-Series with Continuous Spectra,* Nature **161**, 686 (1948).

[65] M. Akay, *Time Frequency and Wavelets in Biomedical Signal Processing,* edited by M. Akay (Institute of Electrical and Electronics Engineers, Inc., New York, 1998).

[66] M. Basseville and I. Nikiforov, *Detection of abrupt changes: Theory and application* (Prentice-Hall, New Jersey, 1993).

[67] J. Rafiee, M. A. Rafiee, N. Prause, and M. P. Schoen, *Wavelet basis functions in biomedical signal processing,* Expert Systems with Applications **38**, 6190 (2011).

[68] E. Brazhnik, A. V. Cruz, I. Avila, M. I. Wahba, N. Novikov, N. M. Ilieva, A. J. McCoy, C. Gerber, and J. R. Walters, *State-Dependent Spike and Local Field Synchronization between Motor Cortex and Substantia Nigra in Hemiparkinsonian Rats,* Journal of Neuroscience **32**, 7869 (2012), arXiv:NIHMS150003 .

[69] A. K. Sen and J. O. Dostrovsky, *Evidence of intermittency in the local field potentials recorded from patients with Parkinson's disease: A wavelet-based approach,* Computational and Mathematical Methods in Medicine **8**, 165 (2007).

[70] N. Thakor, Guo Xin-Rong, Sun Yi-Chun, and D. Hanley, *Multiresolution wavelet analysis of evoked potentials,* IEEE Transactions on Biomedical Engineering **40**, 1085 (1993).

[71] D. Rubino, K. A. Robbins, and N. G. Hatsopoulos, *Propagating waves mediate information transfer in the motor cortex,* Nature Neuroscience **9**, 1549 (2006).

[72] W. De Baene and R. Vogels, *Effects of adaptation on the stimulus selectivity of macaque inferior temporal spiking activity and local field potentials,* Cerebral Cortex **20**, 2145 (2010).

[73] N. Mammone and F. C. Morabito, *Enhanced automatic wavelet independent component analysis for electroencephalographic artifact removal,* Entropy **16**, 6553 (2014).

[74] H. Nolan, R. Whelan, and R. B. Reilly, *FASTER: Fully Automated Statistical Thresholding for EEG artifact Rejection,* Journal of Neuroscience Methods **192**, 152 (2010).

[75] P. M. R. Reis, F. Hebenstreit, F. Gabsteiger, V. von Tscharner, and M. Lochmann, *Methodological aspects of EEG and body dynamics measurements during motion,* Frontiers in Human Neuroscience **8** (2014), 10.3389/fnhum.2014.00156.

[76] M. J. Wagner, T. H. Kim, J. Savall, M. J. Schnitzer, and L. Luo, *Cerebellar granule cells encode the expectation of reward,* Nature **544**, 96 (2017).

[77] A. Kachenoura, L. Albera, L. Senhadji, and P. Comon, *ICA: A potential tool for BCI systems,* IEEE Signal Processing Magazine **25**, 57 (2008).

[78] M. Congedo, C. Gouy-Pailler, and C. Jutten, *On the blind source separation of human electroencephalogram by approximate joint diagonalization of second order statistics,* Clinical Neurophysiology **119**, 2677 (2008), arXiv:0812.0494 .

[79] C. C. Wood and G. McCarthy, *Principal Component Analysis of Event-Related Potentials: Simulation Studies Demonstrate Misallocation of Variance Across Components,* Electroencephalography and clinical Neurophysiology **59**, 249 (1984).

[80] M. F. Suaud-Chagny, K. Chergui, G. Chouvet, and F. Gonon, *Relationship between dopamine release in the rat nucleus accumbens and the discharge activity of dopaminergic neurons during local in vivo application of amino acids in the ventral tegmental area,* Neuroscience **49**, 63 (1992).