

# Signal Processing for Magnetoencephalography

Rupert Benjamin Clarke

*Submitted for the degree of Doctor of Philosophy*

*University of York  
Department of Electronics*

September 2010

## **Abstract**

Magnetoencephalography (MEG) is a non-invasive technology for imaging human brain function. Contemporary methods of analysing MEG data include dipole fitting, minimum norm estimation (MNE) and beamforming. These are concerned with localising brain activity, but in isolation they do not provide concrete evidence of interaction among brain regions. Since cognitive neuroscience demands answers to this type of question, a novel signal processing framework has been developed consisting of three stages. The first stage uses conventional MNE to separate a small number of underlying source signals from a large data set. The second stage is a novel time-frequency analysis consisting of a recursive filter bank. Finally, the filtered outputs from different brain regions are compared using a unique partial cross-correlation analysis that accounts for propagation time. The output from this final stage could be used to construct conditional independence graphs depicting the internal networks of the brain.

In the second processing stage, a complementary pair of high- and low-pass filters is iteratively applied to a discrete time series. The low-pass output is critically sampled at each stage, which both removes redundant information and effectively scales the filter coefficients in time. The approach is similar to the Fast Wavelet Transform (FWT), but features a more sophisticated resampling step. This, in combination with the filter design procedure leads to a finer frequency resolution than the FWT.

The subsequent correlation analysis is unusual in that a latency estimation procedure is included to establish the probable transmission delays between regions of interest. This test statistic does not follow the same distribution as a conventional correlation measures, so an empirical model has been developed to facilitate hypothesis testing.

# Contents

|   |           |
|---|-----------|
| <b>List of Figures</b>                        | <b>7</b>  |
| <b>Acknowledgements</b>                       | <b>9</b>  |
| <b>1 Introduction</b>                         | <b>10</b> |
| 1.1 Chapter Overview . . . . .                | 12        |
| <b>2 Background</b>                           | <b>13</b> |
| 2.1 The Human Brain . . . . .                 | 13        |
| 2.1.1 Macroscopic View . . . . .              | 13        |
| 2.1.2 Microscopic View . . . . .              | 17        |
| 2.2 Magnetoencephalography . . . . .          | 21        |
| 2.2.1 History . . . . .                       | 21        |
| 2.2.2 Instrumentation and Operation . . . . . | 23        |
| 2.2.3 Alternatives . . . . .                  | 27        |
| 2.2.4 Comparison . . . . .                    | 32        |
| 2.3 MEG Analysis . . . . .                    | 34        |
| 2.3.1 Classic Model . . . . .                 | 35        |
| 2.3.2 Minimum Norm Estimation . . . . .       | 37        |
| 2.3.3 Beamforming . . . . .                   | 39        |
| 2.3.4 Model Improvements . . . . .            | 40        |
| 2.3.5 Discussion . . . . .                    | 42        |
| 2.4 Summary . . . . .                         | 43        |

|          |  |           |
|----------|--|-----------|
| <b>3</b> | <b>Concepts</b>                                | <b>45</b> |
| 3.1      | Realistic Brain Models . . . . .               | 45        |
| 3.2      | Connectivity . . . . .                         | 46        |
| 3.3      | Combined Signal Processing Framework . . . . . | 48        |
| 3.3.1    | Summary . . . . .                              | 51        |
| <b>4</b> | <b>Minimum Norm Estimation</b>                 | <b>52</b> |
| 4.1      | Definition . . . . .                           | 52        |
| 4.2      | Solution . . . . .                             | 55        |
| 4.3      | Regularisation . . . . .                       | 56        |
| 4.3.1    | Regularisation Parameter Selection . . . . .   | 58        |
| 4.4      | MNE In Practice . . . . .                      | 62        |
| 4.4.1    | Source Space . . . . .                         | 63        |
| 4.4.2    | Depth Weighting . . . . .                      | 64        |
| 4.5      | Examples . . . . .                             | 65        |
| 4.5.1    | Introduction . . . . .                         | 66        |
| 4.5.2    | Method . . . . .                               | 66        |
| 4.5.3    | Results . . . . .                              | 68        |
| 4.5.4    | Conclusions . . . . .                          | 68        |
| <b>5</b> | <b>Time-frequency Analysis</b>                 | <b>71</b> |
| 5.1      | Relevance to MEG . . . . .                     | 71        |
| 5.2      | Short-time Fourier Transform . . . . .         | 75        |
| 5.2.1    | Discrete Fourier Transform . . . . .           | 75        |
| 5.2.2    | Definition of STFT . . . . .                   | 79        |
| 5.2.3    | Discussion . . . . .                           | 81        |
| 5.3      | Wavelet Analysis . . . . .                     | 83        |
| 5.3.1    | Continuous Wavelet Transform . . . . .         | 84        |
| 5.3.2    | Discrete Wavelet Transform . . . . .           | 86        |

|          |  |            |
|----------|--|------------|
| 5.3.3    | Multiresolution analysis . . . . .           | 88         |
| 5.3.4    | Discussion . . . . .                         | 89         |
| 5.4      | A Novel Filter-bank Analysis . . . . .       | 90         |
| 5.4.1    | Time-domain View . . . . .                   | 91         |
| 5.4.2    | Frequency-domain View . . . . .              | 98         |
| 5.5      | Summary . . . . .                            | 116        |
| <b>6</b> | <b>Statistical Framework</b>                 | <b>119</b> |
| 6.1      | Theory . . . . .                             | 119        |
| 6.1.1    | Correlation . . . . .                        | 120        |
| 6.1.2    | Partial Correlation . . . . .                | 121        |
| 6.1.3    | Cross Correlation . . . . .                  | 122        |
| 6.2      | A Test Statistic . . . . .                   | 123        |
| 6.2.1    | Latency Estimate . . . . .                   | 125        |
| 6.2.2    | Null Model . . . . .                         | 125        |
| 6.3      | Higher order statistics . . . . .            | 129        |
| 6.4      | Application . . . . .                        | 135        |
| 6.4.1    | Introduction . . . . .                       | 135        |
| 6.4.2    | Methods . . . . .                            | 136        |
| 6.4.3    | Results . . . . .                            | 146        |
| 6.4.4    | Discussion . . . . .                         | 146        |
| 6.5      | Summary . . . . .                            | 150        |
| <b>7</b> | <b>Summary and Further Work</b>              | <b>152</b> |
| 7.1      | Summary . . . . .                            | 152        |
| 7.1.1    | Aims . . . . .                               | 152        |
| 7.1.2    | MEG in the context of Neuroimaging . . . . . | 152        |
| 7.1.3    | MEG Analysis . . . . .                       | 153        |
| 7.1.4    | Time-Frequency Decomposition . . . . .       | 153        |

|          |  |            |
|----------|--|------------|
| 7.1.5    | Connectivity Analysis . . . . .        | 154        |
| 7.2      | Further Work . . . . .                 | 155        |
| 7.2.1    | Source Signal Estimation . . . . .     | 155        |
| 7.2.2    | Time-frequency Decomposition . . . . . | 155        |
| 7.2.3    | Connectivity Analysis . . . . .        | 156        |
| <b>A</b> | <b>Spherical Head Model</b>            | <b>157</b> |
|          | <b>Abbreviations</b>                   | <b>162</b> |
|          | <b>References</b>                      | <b>165</b> |

## List of Figures

|      |  |    |
|------|--|----|
| 2.1  | Gross anatomy of the brain . . . . .                                     | 15 |
| 2.2  | Lateral view of cerebrum . . . . .                                       | 16 |
| 2.3  | General morphology of a neuron . . . . .                                 | 19 |
| 2.4  | The first MEG measurement, made in 1971 . . . . .                        | 23 |
| 2.5  | MEG System at York Neuroimaging Centre . . . . .                         | 24 |
| 2.6  | Experimental current dipole . . . . .                                    | 36 |
| 3.1  | Data flow diagram of MEG signal processing framework . . . . .           | 50 |
| 3.2  | Conditional independence graph . . . . .                                 | 51 |
| 4.1  | Representative diagram of an L-curve . . . . .                           | 60 |
| 4.2  | MNE solutions from simulated dipoles . . . . .                           | 69 |
| 5.1  | Time and frequency domain representations of a MEG signal . . . . .      | 73 |
| 5.2  | Effect of windowing on DFT . . . . .                                     | 78 |
| 5.3  | Spectral leakage in DFT due to rectangular windowing . . . . .           | 79 |
| 5.4  | Sidebands in Fourier spectrum introduced by rectangular window . . . . . | 80 |
| 5.5  | Spectrogram of linear chirp . . . . .                                    | 82 |
| 5.6  | Two common mother wavelets . . . . .                                     | 85 |
| 5.7  | Multiresolution Analysis . . . . .                                       | 89 |
| 5.8  | Hann window . . . . .  | 93 |
| 5.9  | Sequence of raised cosine wavelets . . . . .                             | 93 |
| 5.10 | Zero D.C. by symmetry . . . . .  | 96 |

|  |     |
|--|-----|
| 5.11 Fourier transform of raised cosine wavelets . . . . .                         | 98  |
| 5.12 Transfer function of raised cosine wavelet . . . . .                          | 99  |
| 5.13 Asymmetrical slopes of adjacent wavelet scales . . . . .                      | 99  |
| 5.14 Impulse responses of ideal and windowed filters . . . . .                     | 103 |
| 5.15 Illustration of filter design and action upon MEG signal . . . . .            | 107 |
| 5.16 Impulse response of filter bank . . . . .                                     | 112 |
| 5.17 Examples of filter bank output . . . . .                                      | 113 |
| 5.18 Reconstruction of input signal . . . . .                                      | 114 |
| 5.19 Short-time Fourier transform . . . . .  | 115 |
| 5.20 Change in power with respect to baselines . . . . .                           | 116 |
|  |     |
| 6.1 Bootstrapped test statistic with noise input . . . . .                         | 128 |
| 6.2 Trends in null distribution of test statistic vs. sample size . . . . .        | 130 |
| 6.3 Trends in null distribution of test statistic vs. sequence length . . . . .    | 131 |
| 6.4 Cross correlation with simulated data . . . . .                                | 133 |
| 6.5 Removal of confounding influence using PCA . . . . .                           | 134 |
| 6.6 Approximate locations of auditory and speech areas on inflated brain . . . . . | 136 |
| 6.7 Location of chosen regions of interest . . . . .                               | 138 |
| 6.8 Minimum norm estimates from averaged data . . . . .                            | 139 |
| 6.9 Epoch average of gamma activity in regions of interest . . . . .               | 141 |
| 6.10 Left hemisphere amplitude changes (A & B) . . . . .                           | 142 |
| 6.11 Left hemisphere amplitude changes ( $C_1$ & $C_2$ ) . . . . .                 | 143 |
| 6.12 Right hemisphere amplitude changes (D & E) . . . . .                          | 144 |
| 6.13 Right hemisphere amplitude changes (F) . . . . .                              | 145 |
| 6.14 Interaction model during auditory stimulation (35.3 – 44.1 Hz) . . . . .      | 147 |
| 6.15 Interaction model during auditory stimulation (44.1 – 55.1 Hz) . . . . .      | 147 |

## **Acknowledgements**

I owe a considerable debt of gratitude to the many people without whose support this thesis could not have been produced. It would be impossible for me to name everyone who has contributed in one way or another, but the following deserve special mention:

Gary Green

David Halliday

Aziz Asghar

Pádraig Kitterick

Andre Gouws

Will Woods

Mark Hymers

Ash Jansari

and all the staff at YNiC.

## **Financial Assistance**

This PhD was generously funded by the York Neuroimaging Centre (YNiC) in partnership with the University of York. I am grateful for their support.

## **Declaration**

The work presented in this thesis is entirely that of the author except where otherwise indicated. No part of this document has been submitted for previous examination.

# 1 Introduction

Magnetoencephalography (MEG) is a non-invasive functional neuroimaging modality that is closely related to the long-established practice of electroencephalography (EEG). When neurons are activated in living organisms, ionic currents flow, producing extremely low-intensity magnetic fields. This phenomenon is called *neuroelectromagnetism*. The technology is available to record these fields externally and produce graphical representations of them. The primary subject of measurement is the human brain.

There are two broad areas of application for MEG: clinical assessment and cognitive neuroscience [1]. In the former application, epilepsy patients with intractable seizures that do not respond to drug therapy are examined using MEG prior to neurosurgery [2]. Here, the clinician is interested in locating the focus of epilepsy. In this respect, MEG offers heightened accuracy over EEG, albeit at greater expense. In the second application, cognitive psychologists make use of MEG to study and characterise dynamic brain activity in both healthy and diseased states.

Commercial MEG systems have been available for over 20 years, and compared with other functional neuroimaging methods, MEG is minimally invasive while achieving comparable spatial resolution and significantly superior time resolution [1, 3]. In view of this, it may seem surprising that MEG systems are not more commonplace than is actually so. One explanation is that the field of MEG analysis is impoverished of suitable techniques for answering the questions posed by neuroscience. In particular, while the problem of locating neuromagnetic sources within the brain has been addressed in many ways, relatively little has been done to assist the neuroscientist with making inferences from that information. One might wish to infer, for example, that certain stimuli cause increased

activity in a particular brain region. In the search for potential areas of novel contribution in MEG analysis, the present author concluded that a combination of selected digital signal processing algorithms and statistical analysis might offer additional insight in this respect. This idea formed the basis for the project described herein.

In this thesis, an engineered approach to connectivity analysis is proposed, which consists of using an existing source estimation process to extract signal estimates from a small number of regions of interest in the cerebral cortex. This is to be followed by a time-frequency analysis that subdivides neurophysiological frequency bands. Finally, the spectrotemporal representation shall be subjected to statistical analysis to determine the associations between regions of interest. In the course of developing the proposed framework, a novel time-frequency decomposition was also developed consisting of a filter-bank made up of cascaded low-pass finite impulse response (FIR) filters. It is possible to tune the characteristics of the filter-bank, but a particular tuning was suggested which subdivides traditional neurophysiological frequency bands by a factor of approximately 3. The final development of the approach was a connectivity analysis that was based upon partial correlation methods.

The earlier chapters of this thesis will provide a fundamental background in neuroscience, MEG recording and contemporary MEG analysis in order to contextualise the subsequent content. A discussion of some concepts for novel approaches to MEG analysis is then presented. Finally, the full implementation of one particular concept is described in detail over several chapters, with examples of its application to typical experimental data. The tools developed are designed to contribute to furthering understanding of human brain function.

## 1.1 Chapter Overview

The remaining chapters of the thesis are described below.

**Background** In this chapter, the underpinning knowledge required to understand the techniques discussed in the thesis is presented. Established forms of MEG analysis are also considered.

**Concepts** This chapter describes some research topics that were considered, forming a philosophical argument for developing a new MEG analysis framework, which is presented last.

**Minimum Norm Estimation** This chapter comprises a detailed study on the theory and practice of minimum norm estimation, which is an important component of the proposed framework. Some examples of its use are included, based on simulated data.

**Time-frequency analysis** An argument for using time-frequency representations in the interpretation of MEG data is presented in this chapter, followed by a review of existing techniques. A novel approach to time frequency analysis is then developed, with examples of its application to MEG data. This constitutes the second stage of the framework.

**Statistical Framework** This chapter charts the development of a suite of tools for analysing MEG data preprocessed using the techniques described in the previous two chapters. The integration of all of the techniques making up the framework is discussed, with a preliminary example of its application to MEG data.

**Summary and Further Work** A summary of the work presented throughout the thesis is given. Additional work that could be carried out to continue and enhance the concepts that have been developed is then suggested.

## 2 Background

Magnetoencephalography is one of a number of functional neuroimaging modalities, which otherwise include Electroencephalography (EEG), functional Magnetic Resonance Imaging (fMRI) and Positron Emission Tomography (PET). Functional neuroimaging is concerned with the detection of time-dependant activity within the signaling systems of the nervous system. Such measurements are intended to elucidate the role of the brain in cognition and behaviour, motor function, endocrinology and autonomic functions such as breathing, circulation and digestion. This chapter provides the appropriate theoretical and historical background to the work described later in the thesis.

### 2.1 The Human Brain

For the purposes of understanding the requirements of an MEG analysis method, it is useful to first consider the anatomy of the brain and its general function. The first part of this section studies the gross anatomy of the brain. Proceeding this is a section that examines the microscopic elements that make up the cerebral cortex, and how these are thought to contribute to MEG signals.

#### 2.1.1 Macroscopic View

The brain is categorised as part of the *central nervous system* (CNS), which also includes the spinal cord. Sensory information from receptors around the body arrives at the spinal cord via *afferent* nerves. Motor function is brought about by *efferent* nerves that join the spinal cord to remote effectors. Together, these nerves make up the *peripheral nervous*

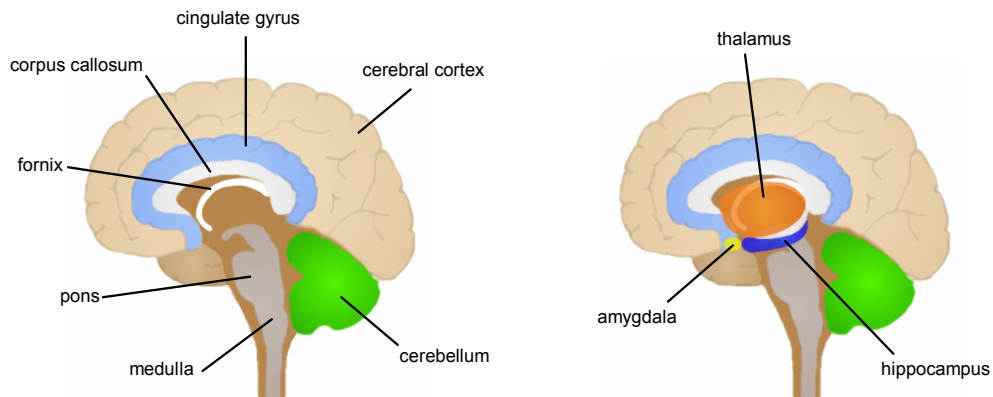
*system*. From the point where the peripheral nerves join the spinal cord, further connections ascend to the brain stem. Here, they mostly decussate, i.e. they cross from left to right and vice versa. Consequently, sensory and motor function from one side of the body is associated with the brain hemisphere on the opposite side. Descending (motor) nerves pass through the *internal capsule* within the midbrain, having originated in the motor area of the cortex, the *precentral gyrus*. Ascending nerves terminate in the *thalamus*, where further connections are made to the *postcentral gyrus*, which processes somatosensory information — sensations of pressure, temperature, pain etc [4].

## Brain

From an embryological point of view, the brain develops in three sections: the *forebrain*, *midbrain* and *hindbrain* [4]. The forebrain splits into two halves that curve over to form the *cerebral hemispheres*. These envelope the deeper structures of the forebrain (the limbic system). The midbrain forms part of the brainstem, joining the *basal ganglia* of the cerebral hemispheres to the hindbrain. The hindbrain forms the medulla, pons and cerebellum. Figure 2.1 shows the major structures in the brain.

The cerebral hemispheres consist of an outer layer of grey matter with many folds. This is known as the cerebral cortex. Grey matter is so-called because it consists mostly of cell bodies, which appear grey in a preserved brain [4]. Underneath is a core of white matter, made up of myelinated nerve fibres, which are bundles of axons. The myelin sheaths are electrically insulating, and serve to isolate nerves from each other [4]. Some nerves, known as *association* fibres, connect different parts of the cortex within the same hemisphere. *Projection* fibres make connections between the cerebral cortex and either subcortical structures, the brainstem or spinal cord. Lastly, *commissural* fibres connect equivalent regions within the two hemispheres [4]. A large body of commissural fibres exists, referred to as the *corpus callosum*.

The inferior (deeper) parts of the cerebral hemispheres which surround the brainstem, are known together as the limbic system. This includes the *cingulate gyrus*, which is involved



(a) Medial sagittal section

(b) Medial sagittal section with selected inferior lateral structures overlaid

Figure 2.1: Gross anatomy of the brain

with autonomic functions such as heart rate, as well as many other functions including attention [5]. Also included are the *hippocampus* and *parahippocampal gyrus*. These are involved with memory [5]. The *hypothalamus* connects the brain with the endocrine system, thus affecting the production of hormones. The *thalamus* provides an interface between the cortex and other parts of the CNS. The *fornix* is composed of white matter, and connects the hippocampus to both the *mamillary bodies* and part of the cortex. Finally, the *amygdala* is located in the temporal region and is implicated in emotional function. In general, the limbic system as a whole is associated with emotion, learning and memory [5].

The cerebellum is a large neuronal structure immediately below the cerebral hemispheres to the rear of the brainstem. Thick fibre tracts attach it to the pons. Its function is predominantly related to movement and proprioception (the sense of movement and location) [5].

## Cerebral Cortex

In MEG studies, the cerebral cortex is particularly important because MEG is primarily sensitive to superficial sources due to their proximity to the sensors. It consists of considerable numbers of neurons of two major types: pyramidal cells and stellate cells [4]. The densities of these cells varies throughout the cortex, contributing to functional specialisation. Pyramidal cells are elongate structures possessing an apical dendrite. This is a long dendrite from which several shorter branches emerge, receiving very many synaptic connections from other neurons (see §2.1.2). The apical dendrite usually extends towards the cortical surface. Pyramidal cells vary in height from about  $10\ \mu\text{m}$  up to  $100\ \mu\text{m}$  for some neurons in the motor cortex. The axons of pyramidal cells connect to other brain regions. In comparison, stellate cells are more rounded with no apical dendrite. They mainly serve to make short connections between neurons, each with relatively few synapses.

The many folds that appear in the cerebral cortex are called *sulci* [4]. The convex regions between them are *gyri*. Although the cortex has very many small sulci, there are a few much larger ones, sometimes called *fissures*, which divide the cortex into distinct lobes. The four lobes, shown in figure 2.2, are the *frontal*, *parietal*, *occipital* and *temporal* lobes. As well as being physically separated, these lobes are functionally distinct.




Figure 2.2: Lateral view of cerebrum

taken from: Gray, H. *Anatomy of the Human Body* (Philadelphia: Lea & Fabiger, 1918)

### 2.1.2 Microscopic View

The microanatomy of the CNS generally concerns two types of cell from which it is composed. Those are neurons and glial cells (neuroglia) [4, 5]. Neurons are the decision-making units in the CNS. Neuroglia do not directly take part in information transmission or processing. However, they are vital to the operation of the CNS. This section describes the morphology of neurons and neuroglia, as well as the mechanism of communication between neurons.

#### Morphology

Neurons feature various narrow projections (*cell processes*) leading away from the cell body (*soma*) [4, 5]. Axons are the processes that carry signals to other neurons. They are frequently sheathed in myelin, a fatty substance that is an electrical insulator. They terminate at a number of synapses, which are the connections to other neurons. Those connections are made onto *dendrites*, further neuronal processes that receive signals. Dendrites have many branches, the most extreme of which have the highest concentration of protruding spines that are the postsynaptic processes (synaptic receptacles). Unlike axons, they are never myelinated. Within the soma is a nucleus and various organelles, which take part in the development and metabolism of the cell. The cell, including its processes, is otherwise filled with a continuous body of cytoplasm. Figure 2.3 depicts the structure of a neuron.

The neurons fall into four categories depending on their form [4, 5]. Unipolar neurons have a single process extending from their soma, which branches to form both dendritic and axonal terminals. Bipolar neurons have a separate axon and dendrite extending in different directions. Pseudounipolar neurons are similar to unipolar neurons in appearance, but are in fact bipolar neurons whose dendrite and axon are fused together. Finally, multipolar neurons are those with a considerable branching dendrite structure (the dendritic field), which immediately surrounds the soma. They also possess an axon that extends away from the soma and can potentially synapse with distant neurons. This axon branches rarely, if at

all, over most of its length. Multipolar neurons are the type usually occurring in the brain.

Although neurons form the signaling network of the CNS, there are several times as many neuroglia present. These fall into three categories [4, 5]. Astrocytes are approximately rounded cells with numerous branches extending from the soma. They surround neurons, serving to provide nutrients to them. Notably, they form the *blood-brain barrier*. This transports ions between the blood and extracellular fluid of the brain, while rejecting blood-borne agents that may disrupt the function of the neurons. Oligodendrocytes are myelin-producing cells, with each one myelinating several axons in spiraling concentric laminations. Microglia are phagocytes that remove the debris of damaged tissue.

In addition to the myelin produced by oligodendrocytes, another type of myelination is performed by Schwann cells [4, 5]. Each Schwann cell only myelinates part of a single axon. Whether myelinated by Schwann cells or oligodendrocytes, the myelinated regions are periodically interrupted, exposing part of the axonal membrane. This exposed part is called a node of Ranvier, and serves to enhance the transmission of action potentials as described in the next section.

## **Signaling**

Signaling between neurons takes place by way of the synapses. A single synapse consists of three parts. This first component is the *presynaptic terminal*, which is one of the synaptic boutons at the extremities of an axon. The second is the *synaptic cleft*, a gap of approximately 20 nm that occurs between the presynaptic terminal and the remaining component of the synapse, the *postsynaptic process*. As previously stated, this is usually a dendritic spine occurring on another neuron. If this is the case then it is called an axodendritic synapse. Synapses can also connect directly onto the soma of the receiving neuron, in which case they are classed as axosomatic. Other types of synapse do exist (e.g. dendrodendritic), but in far smaller numbers [4, 5].

The cytoplasm within a neuron is at a lower electrical potential than the extracellular fluid [5]. If the extracellular fluid is defined as being at 0V, the voltage of the cytoplasm

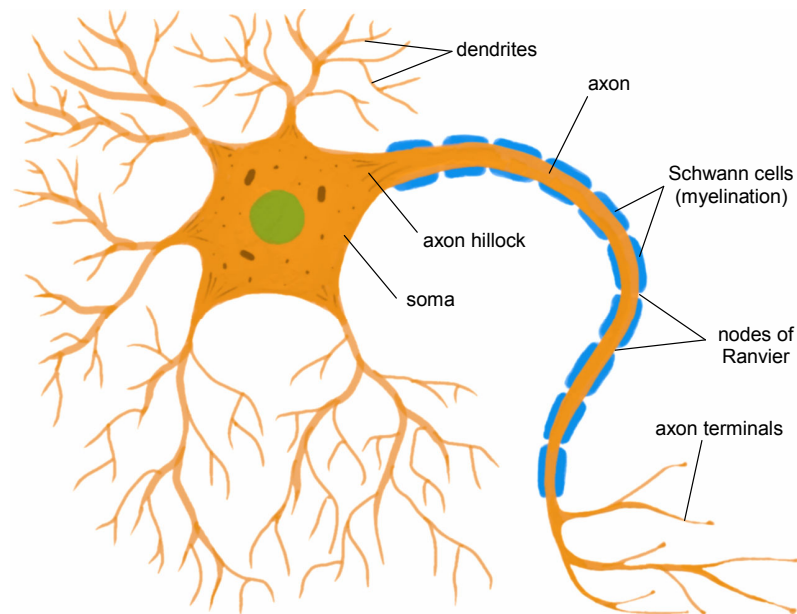


Figure 2.3: General morphology of a neuron

is negative. This negative voltage is known as the *membrane potential* as it occurs across the cell membrane. In the neuron's resting state, the membrane potential is maintained at about  $-70$  mV by active ion pumps in the neuronal membrane [5]. These pump sodium ( $\text{Na}^+$ ) ions out of the cell and potassium ( $\text{K}^+$ ) ions into it, in a 3:2 ratio. Thus the intracellular concentration of  $\text{K}^+$  increases while that of  $\text{Na}^+$  decreases. The 3:2 imbalance causes the neuron to become negatively charged. An equilibrium is maintained because the ion pumps act against diffusion and passive electrical currents, which cause a limited opposing transport of ions through non-gated ion channels. A voltage is developed across the membrane due to the diffusion of  $\text{K}^+$  down the concentration gradient. The cell is said to be polarised.

Changes in the postsynaptic membrane potential occur when chemical neurotransmitters are released from the presynaptic terminal into the synaptic cleft [5]. Such releases are caused by the firing of the presynaptic neuron. The subsequent change in membrane potential is brought about by the action of gated ion channels. These respond to the

neurotransmitter by allowing movement of ions across the postsynaptic membrane. This action is “graded” — the ionic current (and its resultant effect on membrane potential) varies in magnitude according to the neurotransmitter concentration [5]. Also, the direction of the current varies between synapses, depending on whether they are excitatory or inhibitory. Excitatory synapses raise the membrane potential (decreasing its magnitude) while inhibitory synapses lower it. It should be noted that the membrane potential is not constant throughout the cell due to the resistivity of the cytoplasm. As such, changes in membrane potential will diminish with distance from the ionic current source that caused them.

In the axon hillock (where the axon emerges from the soma) and the axon itself, voltage-dependent ion channels exist in the cell membrane [5]. If the membrane potential exceeds a particular threshold, these channels allow  $\text{Na}^+$  ions to enter the cell. The ions will do so due to the potential difference and the concentration gradient created by the ion pumps. The influx of  $\text{Na}^+$  ions further depolarises the neuron, increasing the membrane potential and opening yet more  $\text{Na}^+$  channels. The resulting avalanche in depolarising ionic currents causes the membrane potential to spike, generating an *action potential* that continues down the axon [5]. Once the cell is depolarised,  $\text{K}^+$  channels open, repolarising the neuron. In fact, the repolarisation overshoots briefly, during which time a second spike is less likely to occur (the membrane potential is further than usual from the threshold). Also, the  $\text{Na}^+$  channels are briefly inactivated. Both factors limit the rate at which action potentials can be generated.

The magnitude of the action potential decays exponentially along the axon as the volume currents diverge. This limits both the distance that the signal can travel, and the speed. At successive gated ion channels, the rising edge of the action potential takes longer to exceed the threshold because of the attenuation. However, in myelinated axons, this effect is mitigated by the presence of the myelin [5]. The fatty coating insulates sections of the axon such that the voltage dropped along the axon is reduced, through the elimination of volume currents flowing out of the cell. The nodes of Ranvier provide points at which the action

potential is regenerated by voltage-dependent, gated ion channels [5]. The transmission speed is increased because the signal effectively jumps from one node to the next via the low-loss pathway.

As well as synapses mediated by neurotransmitters, a further type of synapse exists that uses electrical instead of chemical transmission [5]. In these synapses, the presynaptic and postsynaptic membranes are in very close contact. Electrical transmission is facilitated by *gap junctions* [5]. These are groups of protein molecules situated in the pre- and postsynaptic membranes, which form pores that result in the continuity of cytoplasm from one neuron to the next. Electrical synapses permit much quicker communication between neurons than chemically mediated synapses. Most are also bidirectional.

MEG signals are thought to be primarily due to dendritic currents [6, 7]. Axonal current sources consist of a *depolarisation front* that represents the commencement of an action potential and travels down the length of the axon, followed closely by a *repolarisation front*. The separation of the two depends on the duration of the action potential, but is consistently very small. This can be modelled by two oppositely-polarised current dipoles (see §2.3), forming a quadrupole. Quadrupolar sources produce higher order spatial fields that drop off rapidly with distance, thus contributing very little to the signal measured by MEG [7].

## 2.2 Magnetoencephalography

This section is intended as a primer on MEG. It includes a brief history of neuroelectromagnetic measurement, followed by a description of the instrumentation used in MEG and its use. Rival modalities are also considered and the place of MEG in functional neuroimaging is then established.

### 2.2.1 History

The field of medicine has been aware of the role of electricity in nerve function since the 1700s, owing to pioneers such as Luigi Galvani, whose stimulation of frog muscles in 1781 firmly established the connection between electrical stimulation and muscle

contraction [8]. The subsequent invention of the galvanometer the following century facilitated the measurement of endogenous currents in frog muscles, firstly by Carlo Matteucci in 1838. Then Emil du Bois-Reymond characterised the nerve impulse or *action potential* in 1841 [9].

In the second half of the 19th century, theories on the structure and organisation of the central nervous system began to emerge. John Hughlings-Jackson proposed that brain regions might have different functions after noting the ordered progression of muscle contractions during epileptic seizures [5]. Marc Dax, Pierre Paul Broca and Carl Wernicke all used autopsy to identify regions of the cerebral cortex that, when damaged by stroke, impaired language ability [10, 5]. These discoveries lent credence to the notion of functional specialisation in the brain.

The development of histological staining of tissue samples by Camillo Golgi in 1873 allowed Santiago Ramón y Cajal to produce highly-renowned illustrations of neuronal structures [9]. Using similar techniques, Korbinian Brodmann categorised 52 distinct regions of the cerebral cortex based on differences in microanatomy in 1909 [5].

Meanwhile, Richard Caton had conducted invasive intra-cranial EEG on animal cortices, having begun investigations in 1875. Many other investigators followed suit, but not until 1924 did Hans Berger carry out the initial extra-cranial EEG measurements on human beings. Berger succeeded in identifying alpha and beta oscillations with his technique, which was brought to the attention of the medical fraternity when these results were replicated by Adrian and Matthews in 1934 [9].

The first biomagnetic measurements were magnetocardiographic (MCG) measurements carried out by Baule and McFee in the 1960s [11]. For this, they used two similar coils on a ferrite core, each with a considerable number of turns. These were wired in series with opposite polarity, to form a type of gradiometer that would detect the nearby biomagnetic field but not the distant environmental noise fields. The measurements were made outdoors, as far as possible from noise sources. This approach was successful, but produced noisy results. Not long afterwards, David Cohen began making similar measurements, this

time using a low-noise amplifier and a magnetically shielded environment, with improved results [12]. Using this set-up, Cohen also attempted the first MEG measurements in 1968, but the signal-to-noise ratio was very poor. Then in 1969, a device known as the SQUID (Superconducting QUantum Interference Device, §2.2.2), was invented by Jim Zimmerman. The SQUID is a very low noise detector, and when Cohen learnt of this development, a collaboration was organised with Zimmerman on the biomagnetism research [1]. They succeeded in producing a very clean MEG trace using one of Zimmerman's experimental devices. By 1971, commercially manufactured SQUIDs had become available, which Cohen employed to make the first MEG recording [13] (Fig. 2.4). The outstanding results prompted widespread interest in magnetoencephalography.

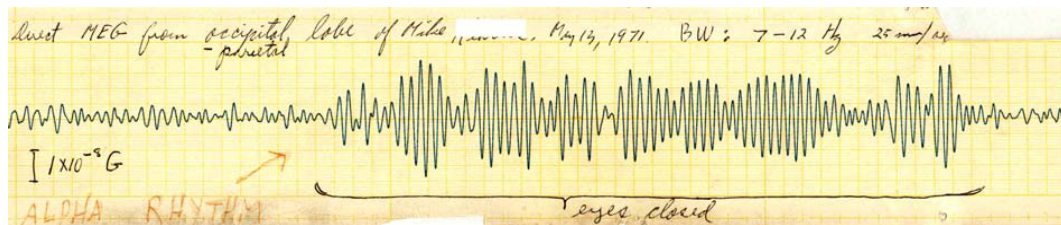



Figure 2.4: The first MEG measurement, made in 1971 (taken from [12])

### 2.2.2 Instrumentation and Operation

Modern MEG equipment measures neuromagnetic fields using several hundred magnetometers, closely arranged around the subject's scalp. The centrepiece of a whole-head MEG system is usually a large housing incorporating a helmet-shaped cavity. The subject's head rests inside the cavity while recording takes place. Internal magnetometers are located immediately around the helmet so as to minimise their distance from the neuromagnetic source, yielding the strongest possible signal. Figure 2.5 shows the MEG system used to capture the results presented in the later chapters.

Neuromagnetic fields arise mainly from postsynaptic currents that flow within the dendrites of neurons, as discussed in §2.1.2. The magnetic field of individual action potentials in the brain is not measurable using MEG. The smallest extra-cranial magnetic



This MEG system (4-D Neuroimaging Magnes 3600WH) contains 248 primary measurement coils, with 28 additional sensors for measuring environmental noise fields. The angle of the Dewar can be adjusted to accommodate measurements with the subject in seated or supine positions. In the photograph, the Dewar is in an intermediate position that reduces the boil-off of liquid helium when not in use. *Photograph: R.B. Clarke*

Figure 2.5: MEG System at York Neuroimaging Centre

fields that can be detected are the combined effect of more than  $10^4$  near-simultaneous neuron activations [14]. This signal is representative of neuronal activation in general because the highly-connected neurons are inclined to become active together in localised groups. Even so, the fields are only on the order of tens of femtoTesla [14], which is so minute that cryogenically-cooled SQUID magnetometers are required to achieve adequate signal-to-noise ratios. The majority of the MEG housing encloses a Dewar containing the liquid helium cryogen that supports the operation of these highly sensitive devices.

A SQUID is a Superconducting Quantum Interference Device, a specialised magnetometer having much lower output noise than conventional designs. MEG uses d.c. SQUIDs, which consist of a superconducting ring with two Josephson junctions [3, 1]. This is simply a ring made from a particular alloy, interrupted by miniscule gaps that are the Josephson junctions. Electrons can cross the junctions due to quantum tunnelling effects, with a phase shift developing across the junction. The superconducting property of the alloy is only achieved at a very low temperature. A cryogenic environment at 4.2K surrounds all components of the sensors. Neuromagnetic fields induce current in a pick-up coil, which is connected across a complementary signal coil, forming a flux transformer. The signal coil is then inductively coupled to the SQUID, which is connected across the input of a low noise amplifier. The SQUID is biased with a direct current. Then, due to quantum phase effects, the voltage across the SQUID becomes a periodic function whose amplitude depends on the magnetic flux coupled to the SQUID. It is this varying amplitude (not the periodic signal) that becomes the MEG signal.

The neuromagnetic fields being measured are many orders of magnitude smaller than typical ambient magnetic fields, such as the geomagnetic field of the planet Earth (which is about  $50\mu\text{T}$ ). MEG equipment must be sensitive enough to measure the relevant neurological fields, but insensitive to much larger noise signals. Problematic noise signals can originate from several sources. These include the ambient magnetic fields of the physical environment, the intrinsic noise of the sensors, and the input noise of the amplifiers receiving the sensor signals. MEG systems employ specific design features to combat each

of these potential sources of noise. Firstly, environmental noise fields are greatly attenuated by housing the complete measuring instrument in a magnetically-shielded room. Secondly, SQUID magnetometers are utilised for their superior noise performance. Lastly, the output from the SQUIDs is amplified using low-noise instrumentation amplifiers. These features can lead to spectral noise densities of less than  $5\text{fT}/\sqrt{\text{Hz}}$  [15]. Many systems implement analogue and digital noise reduction to further improve the signal-to-noise ratio.

During MEG investigations, the output from every magnetometer (or occasionally a subset of them) is recorded over time. Stimuli such as sounds (auditory stimuli) or images (visual stimuli) are often presented to the subject during the recording to investigate neurological responses. Spontaneous brain activity can also be measured. Recent MEG systems digitise the signals for storage and manipulation. After the recording, the data may be analysed using various processes to establish the location and time course of neuronal activity. The spatial resolution can be better than 3 mm [3, 14]. However, this depends on several factors including the analysis method, number of sensors, the location of the activity and the presence of noise. MEG analysis methods will be addressed in §2.3.

Temporal resolutions in MEG range down to 1 ms. The technology might allow for a greater bandwidth, but there is a practical limitation on the highest frequencies that can be detected in the presence of noise. The power spectrum of most noise signals is white, but the neuromagnetic signals diminish with increasing frequency [16], eventually disappearing below the noise floor. Noise is a constant consideration in experimental design for MEG. Signal averaging over multiple trials is commonly used to improve noise rejection [7]. The data are captured in short time segments, known as *epochs*. Each epoch captures data from a single stimulus presentation, and often for brief periods before and afterwards. In subsequent offline analysis, the average is taken across sets of similar epochs that were gathered under a particular experimental condition. The time variable is usually referred to the stimulus onset, so signals that are phase-locked to the stimulus are coherent between trials. Averaging attenuates incoherent signals, such as noise, by the square root of the number of trials [7].

### 2.2.3 Alternatives

This section briefly considers the main non-invasive functional neuroimaging modalities that complement MEG. These are positron emission tomography (PET), functional magnetic resonance (fMRI), and electroencephalography (EEG). These modalities are compared with MEG in terms of suitability for particular applications. Some of the practical and safety considerations associated with their use will also be discussed.

#### PET

A metabolically active compound that has been labelled with a radioactive tracer is administered to the subject intravenously or sometimes by inhalation. The tracer is an isotope with a short half-life that undergoes  $\beta^+$  decay, such as oxygen-15. When a positron is emitted, it travels a very short distance before it collides with a nearby electron. The two are annihilated and a pair of photons are emitted at nearly  $180^\circ$  to one another. Some of these photon pairs arrive at a ring of detectors in the scanning equipment. The arrivals of photon pairs are correlated in order to reconstruct an image showing how the molecule or its metabolites are distributed in the plane of the detector ring. This gives an indication of localised metabolic rate, and hence the intensity of neural activation [17]. In neuropsychological applications the rate of metabolism is a comparative measurement against a baseline rate, i.e. changes rather than absolute levels are of interest. Some clinical applications do study the absolute levels since abnormal rates of metabolism may indicate pathology [14].

The spatial and temporal resolution of PET depend on the tracer that is employed, but resolutions on the order of a few millimetres and response times of approximately one minute can be achieved. A distinct advantage of this modality is the ability to target a particular neurotransmitter system by labelling an appropriate precursor compound [14]. Because it involves ionising radiation, the use of PET must be strictly limited in order to minimise the radiation dose received by the subject. Radiation exposure is more acceptable in clinical diagnosis where there is foreseeable benefit to the patient. Conversely, PET

for research purposes is highly restricted. Compared with fMRI, EEG and MEG, PET is regarded as being more invasive due to the requirement to administer radiopharmaceuticals.

### **fMRI**

Functional MRI is a specialised form of magnetic resonance imaging, a modality normally used to provide structural images of the anatomy. In this specialised application of MRI, the technology is adapted to produce tomographic images of changes in haemodynamic responses. These are believed to be the markers of temporal changes in neural activity. MRI in general relies on the phenomenon of nuclear magnetic resonance, which is briefly explained in the following paragraphs. A more detailed discussion of the relevant theory can be found in [18], from which the following information was taken. After the general treatment of MRI, there follows a dedicated paragraph on fMRI.

A subatomic particle has a spin angular momentum associated with it, even if it is stationary. This means it is effectively rotating about its centre of mass. If the particle is charged, the movement of charge results in a circular current flow, which produces a magnetic dipole. A magnetic dipole can be described in terms of the dipole moment, which is a vector having a direction and magnitude. When subjected to a strong uniform magnetostatic field, the dipole moment will precess about the direction of the field, in accordance with Larmor's Theorem. It will also tend to align with the field<sup>1</sup>. The rate of precession is proportional to the strength of the field. More importantly, it is also proportional to the ratio between the charge on the particle and its mass, or equivalently, the ratio between the dipole moment and the angular momentum of the particle. This ratio is known as the *magnetogyric ratio*, and the corresponding frequency is the *Larmor frequency*.

Once a particle is precessing at the Larmor frequency, then if it is further subjected to a second, much smaller, orthogonal magnetic field that rotates about the static field with equal frequency, then the dipole moment will also precess about that field as it rotates.

---

<sup>1</sup>To be entirely accurate, the particle's dipole will either align with the magnetic field or counter to it. However, in a large group of similar particles, a slight majority will be in alignment with the field, so this is the net effect.

Thus, the dipole moment is deflected away from the direction of the static field at a rate proportional to the smaller field's strength and the particle's magnetogyric ratio. Clearly, the deflection is much slower than the original precession due to the smaller field strength. Other types of particle that are precessing about the static field at a different frequency are not significantly affected. Thus a means of selectively deflecting particles with a particular magnetogyric ratio is provided. In practice, if the strong static field is being produced by a powerful solenoid, the rotating field can be produced by a single perpendicular transmitter coil. Only one coil is required to produce the rotating field by recognising that a sinusoidal field is equal to the sum of two fields rotating in opposite directions, such that one rotates at a frequency which is the negative of the other. Therefore, a radio-frequency (r.f.) sinusoidal signal is used to drive the coil. The field rotating at the negative frequency differs from the Larmor frequency by twice its value, so has no significant effect on dipole precession.

If the rotating field is applied continuously, the dipole continues to be deflected until it is counter-aligned with the static field, before returning to alignment. In this way, it oscillates in and out of alignment with the static field at the frequency determined by the magnetogyric ratio of the particle and the magnitude of the rotating field. The electromagnetic field produced by this particle spin is maximal when the deflection is  $90^\circ$ . If considerable numbers of similar particles undergo this process, a signal is detectable. In essence, this can be used to measure the abundance of a particular particle within a sample.

By applying an r.f. pulse exactly long enough to deflect the dipole moments through  $90^\circ$ , the resulting signal from the dipoles continues to be emitted once the pulse is removed. However, the signal does not last indefinitely. The additional energy given to the particle dissipates due to thermal conduction, and the dipole moment returns to alignment with the static field. The signal decays exponentially, with a time constant  $T_1$ , the *longitudinal relaxation time*. A second form of exponential decay occurs because slight inhomogeneities in the static field cause the dipole moments to precess about the static field at slightly different rates. They point in the same direction (perpendicular to both fields), immediately after the  $90^\circ$  pulse is applied, but as time goes on, they spread out in the plane perpendicular

to the static field. The resulting decay is described by the *transverse relaxation time*,  $T_2$ .

If a second r.f. pulse twice as long as the first is applied after a chosen interval,  $\tau$ , the dipole moments are deflected through a further  $180^\circ$ . The relative offsets of the dipole moments due to the transverse spreading are reversed, so they begin to converge as they continue to precess about the static field. After a further duration of  $\tau$  has elapsed, the dipoles become exactly aligned as they were at the end of the  $90^\circ$  pulse. Therefore, the emitted signal reaches a maximum at this time, called the *spin echo*. It is much easier to measure this signal because it occurs in the absence of any excitation pulses, and at a known time.

In order to make structural images using nuclear magnetic resonance, a gradient is introduced in the static field. The r.f. pulses are tuned to the Larmor frequency of protons located at a particular displacement along the field gradient. The magnetic dipole moments of protons in this narrow slice will be selectively deflected, so the density of protons in that slice can be determined. Three-dimensional images can be resolved by taking numerous different slices and repeating the process with the gradient occurring in the three orthogonal directions.

Functional MRI uses the same technology to produce time-varying images of functional activation in the brain. In a process called the haemodynamic response, oxygenation of blood in the brain responds to changes in local neuron activation by overcompensating for the increase in metabolism that it causes in the region of the neurons. This fact is employed by fMRI in that oxyhaemoglobin is diamagnetic but deoxyhaemoglobin is paramagnetic and possesses a much larger magnetic dipole moment. The local magnetic field is affected so as to reduce the relaxation time of nearby protons. The effect on the spin echo can be detected and used to determine changes in the oxygenation level of the blood in different regions of the brain, and thus the functional activity.

It takes a matter of seconds for the haemodynamic response to reflect an increase in neural activity, which limits the useful temporal resolution of fMRI. Consequently, improvements are generally sought in the spatial resolution, where  $3 \times 3 \times 4$  mm elements

are typical [14]. Due to the superior temporal resolution, the technology is more suitable than PET for detecting the subtleties of functional mapping in the brain. The reliance of both on metabolic processes rather than direct electrical activity is a limitation not present in either EEG or MEG. Because strong magnetic fields and r.f. electromagnetic radiation are used in fMRI, electronic and metallic implants are contraindications for undergoing a scan. Otherwise, the procedure is considered harmless because it has no influence at the molecular level.

### EEG

In the most common form of EEG, a number of electrodes are applied to the subject's scalp, where they are used to measure surface electrical potentials at several locations simultaneously. Typically, a cap is placed over the scalp that has apertures at standardised electrode locations. A conductive gel is injected through each hole prior to insertion of a needle electrode, forming a good electrical contact with the scalp. An international standard exists for the location and name of EEG electrodes, called the *10-20 system* [8]. This defines 21 electrodes based on the proportion of the distance around the skull at which they are placed. These distances are measured from the nasion point above the nose and between the eyes to the inion point, which is central to the back of the head where there is a bony protrusion at the bottom of the skull. The electrodes are located at the intersections of a grid of lines that are separated by 10% or 20% of the distance around the skull.

In many respects the process of recording and analysing the signals from the electrodes is similar to MEG; multiple signals are recorded, and the source models used for analysis are often identical [8] (source modelling is introduced in §2.3). However, there are some crucial differences. First, potentials must be measured relative to a reference. This could be obtained from another of the scalp electrodes, some combination of all of the scalp electrodes, or a reference electrode in a “neutral” location, such as the earlobe. The chosen reference clearly has a bearing on the outcome of the experiment. The second important difference is the need to take into account the conductive properties of head tissue. Analysis

of EEG data relies on conduction models, which aim to predict how internal current sources give rise to surface potentials. This behaviour can be very complicated and difficult to predict, particularly since the skull is highly insulative, so the path of volume currents to the electrodes is convoluted and indirect. It also attenuates and filters the measured signals, particularly where the source is distant from the electrode. This impairs the signal-to-noise ratio. As a result, the accuracy with which sources can be localised is limited to a few centimetres [14].

The cost of the equipment required for EEG is particularly low in comparison to all of the other modalities reviewed herein. The temporal resolution is also excellent — better than 1 ms. The measured data are directly influenced by neural spiking, providing a clear insight into neuronal processes. It is particularly ideal for clinical applications where diagnoses can be made by identifying EEGs characteristic of certain pathologies. The ambiguity and limited accuracy of source localisation curbs the use of EEG for investigating functional mapping and connectivity. The procedure is more invasive than MEG due to the application of electrodes, which is time-consuming and intrusive.

### 2.2.4 Comparison

The sensitivity profile of magnetometer coils describes how greatly they respond to sources at different angles and distances from the sensor. It diminishes with distance, especially in the case of more conventional gradiometer-type sensors. Gradiometers are deliberately employed to measure the 1st-order field component so as to reject distant environmental noise. As a result, MEG is more suited to measuring superficial activity, since the signal-to-noise ratio of deeper sources is smaller. This contrasts with PET and MRI, which exhibit consistent performance over the entire experimental volume.

The data collected in an MEG experiment do not allow the signal from a specific region of interest to be isolated with ease. There are various ways of attempting this, which are discussed in the following section. However, external magnetic field measurements contain ambiguities that prevent the absolute separation of regional sources in the absence of prior

information about the statistical relationships between them. Consequently, MEG cannot entirely replace invasive procedures such as electrocorticography (ECoG — the placement of microelectrodes onto the cortical surface during surgery). PET and fMRI have no such limitations, but the nature of the resulting information leaves ample territory for MEG investigations, which reveal far more about transient activity owing to their freedom from reliance on haemodynamic response.

The initial equipment costs for MEG are particularly high when compared with other modalities. The running costs are not insubstantial either, due to the consumption of liquid helium as a cryogen, which must be replenished on a regular basis [1].

With the exception of EEG, MEG is the only non-invasive functional neuroimaging modality capable of studying detailed electrical activity in the brain [7]. It is often stated that MEG has superior temporal resolution to PET and fMRI, but this does not fully describe the level of informational detail achieved. In particular, phase relationships between different sources can be studied due to the preservation of signal polarity. The transient characteristics of the signal are faithfully reproduced, something that no other non-invasive neuroimaging modality facilitates. It is as much the nature as the magnitude of the measured activity that is of interest in an MEG study; the accurate reproduction of transients in the MEG signal translates into an enhanced ability to draw comparisons between activations in different brain regions. The temporal resolution is only limited by the sample rate, which is typically similar to EEG. However, the low pass filtering effect of the brain tissue limits the potential for sample rate increases in EEG. It is theoretically possible to measure higher frequency activity using MEG, with the caveat that brain electrical activity has lower energy at such frequencies. The system signal-to-noise ratio then becomes the limiting factor.

Unlike the surface potentials measured by EEG, which are due to currents travelling via circuitous routes, the extra-cranial magnetic fields recorded in MEG pass directly through the tissue of the head. This is a unique advantage that improves the accuracy of source localisation. Under favourable conditions sources may be localised to within about 3 mm [14].

MEG excels in terms of being both safe and non-invasive, being an entirely passive procedure, although a structural MRI scan is usually required to provide anatomical data with which the results of the MEG experiment are co-registered. That is, the alignment between the MEG sensor coils and the subject's brain is determined for the purpose of modelling the source space, as discussed in the following section.

## 2.3 MEG Analysis

The fundamental purpose of functional neuroimaging is to identify which parts of a brain are active and at what times. In the case of MEG, this entails calculating neuronal activity from magnetic fields observed outside the head. In physical terms, the head can be viewed as a volume conductor, and neuronal activity as a time-varying current density. Helmholtz proved in 1853 that infinitely many internal current densities can produce a given magnetic field outside a volume conductor. Therefore, knowing the magnetic field is not sufficient to know for certain which current density produced it. Furthermore, MEG recordings do not describe a magnetic field completely. They merely sample the field in many places, leaving out further information that could distinguish between solutions. As a result of this uncertainty, there is no universally correct estimation procedure, and several are in contemporary use.

Simple methods take the form of *source localisation*, which attempts to identify the origin of an MEG signal to within a small brain region. Algorithms that perform source localisation generally start with a model for the magnetic field produced by neuronal currents, because this problem is more tractable. The *forward model* is a function that maps current densities, described by a set of parameters, to vectors of magnetic field measurements. The domain of the function is the *source space*, which includes all possible parameter combinations. The function's range is the *sensor space*, a vector space with the same dimension as the number of sensors. Each orthogonal direction in the sensor space represents the output of a different sensor. To solve the original problem, an inverse solution is found by manipulating the parameters until the predicted and actual measurements are

closest.

More advanced forms of source estimation allow for simultaneous activity taking place over prescribed surfaces or volumes, leading to tomographic maps of activity. The same underlying forward models are used nonetheless. Three methods are presented below that represent the state-of-the-art in MEG analysis. The model underpinning all of these methods has undergone various refinements since its introduction, which are also considered. Some comparisons will then be drawn between the methods.


### 2.3.1 Classic Model

The classic approach to MEG analysis is to model the neuronal activity as a single current dipole within a homogeneous, spherical conductor. A current dipole consists of a current source and sink of equal magnitude, separated by a small distance. As well as a location, the dipole has a magnitude and direction, known as the dipole moment, which is orientated from the sink to the source. The dipole moment is the product of the (scalar) current and the displacement vector between the source and sink. In the theoretical dipole, the displacement is infinitesimal and the current infinite, with the magnitude of the dipole moment being finite. Of course, a practical dipole has both finite displacement and current, but is closely approximated by a theoretical dipole when regarded from a sufficient distance [3, 8, 19]. The spherical conductor geometry facilitates analytical solution of the forward problem; i.e. what a dipole with given parameters would generate at the sensors. The following result is derived in Appendix A.

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi F^2} (F\mathbf{Q} \times \mathbf{r}_0 - \mathbf{Q} \times \mathbf{r}_0 \cdot \mathbf{r} \nabla F) \quad (2.1)$$

$$\begin{aligned} \text{where } F &= a(ra + r^2 - \mathbf{r}_0 \cdot \mathbf{r}), \\ a &= |\mathbf{r} - \mathbf{r}_0|, r = |\mathbf{r}| \end{aligned}$$

$\mathbf{B}(\mathbf{r})$  is the magnetic field as a function of position, which is represented by the position vector  $\mathbf{r}$ .  $\mathbf{Q}$  is the dipole moment and  $\mathbf{r}_0$  is the dipole location. The constant  $\mu_0$  is the permeability of free space. Figure 2.6 shows an example of a magnetic field produced



This is a 3-dimensional image of the magnetic field pattern from an experimental current dipole, recorded using an MEG instrument (4-D Neuroimaging Magnes 3600 WH). This is a whole-head device, so the sensor array is in the shape of a helmet. Blue/red represent fields orientated in/out of the head.

Figure 2.6: Experimental current dipole

by a current dipole. A notable property of this spherical conductor model is that a dipole orientated along any radius of the sphere does not produce a field at the sensors.

Interpreting MEG recordings with the help of this model calls for the dipole parameters to be estimated from a set of instantaneous field measurements as produced by the sensors. This inverse problem has no analytical solution, and must instead be solved numerically. A cost function is defined as the sum, over all sensors, of squared differences between each sensor measurement and the field predicted by the dipole model at the same location and orientation as that sensor. The cost function is then minimised for the dipole parameters using an iterative algorithm, eventually leading to an optimal least-squares fit. While this approach facilitates a solution, it unfortunately embodies some major assumptions that are difficult to justify.

Unlike a current dipole, practical neuronal currents are widely distributed around the brain, not infinitesimal in size. Although multiple trial averaging may be used to isolate a focal source, the spherical conductor remains a poor model for the brain volume. Sometimes a few dipoles are fitted to the data, and more sophisticated conduction models such as finite

element or boundary element models can be used. However, the number of dipoles used dramatically affects their estimated locations; confidence in the results is not good without prior knowledge of the quantity of focal sources.

### 2.3.2 Minimum Norm Estimation

The dipole model introduced above can be greatly improved upon by modelling the current density as a dense grid of dipoles covering a putative source volume or surface. Instead of trying to estimate the location, direction and magnitude of a single dipole, a large complement of dipoles with fixed locations and directions is used. In this way, only the dipole amplitudes need to be estimated.

A predefined grid, chosen by the analyst, sets the dipole locations. Up to three dipoles are placed at each point. The use of three mutually orthogonal dipoles per grid point allows the direction of the estimated current density at a particular point to vary arbitrarily. The three dipoles each account for a different scalar component of the 3-dimensional current density vector. Alternatively, only one or two dipoles are placed at each point. In these cases, the dipole directions are based on assumptions about the signal source. For instance, the component of the current density radial to the spherical volume used in the forward model may be omitted due to its lack of influence on the measurements. Each location then has a perpendicular pair of dipoles in the tangential plane. The typical orientation of pyramidal cell dendrites within the cerebral cortex (detailed in §2.1.1) is commonly cited to justify the use of only one dipole, normal to cortex. Similarly, the grid locations themselves are distributed over relevant anatomical structures that are thought to contribute to the MEG signal.

The dipole amplitudes, here denoted  $\mathbf{j}(t)$ , are linearly related to the sensor signals by the *lead field matrix*,  $\mathbf{L}$ . A row of  $\mathbf{L}$  represents the sensitivity of an individual sensor to each dipole in the grid. The recorded signals,  $\mathbf{b}(t)$ , are a vector function of time, while  $\mathbf{e}(t)$  represents additive measurement noise.

$$\mathbf{b}(t) = \mathbf{L}\mathbf{j}(t) + \mathbf{e}(t) \quad (2.2)$$

$\mathbf{L}$  is calculated using the forward solution presented in the single dipole model. Each column,  $\mathbf{l}_i$ , corresponds to a unique combination of dipole location ( $\mathbf{r}_0$ ) and orientation ( $\mathbf{Q}$ , which is a moment with unit magnitude). The elements of  $\mathbf{l}_i$  are calculated as follows:

$$\mathbf{l}_i = \begin{pmatrix} \mathbf{B}(\mathbf{r}_1) \cdot \mathbf{o}_1 \\ \mathbf{B}(\mathbf{r}_2) \cdot \mathbf{o}_2 \\ \vdots \\ \mathbf{B}(\mathbf{r}_m) \cdot \mathbf{o}_m \end{pmatrix}, \quad (2.3)$$

where  $\mathbf{r}_{1\dots m}$  are the positions and  $\mathbf{o}_{1\dots m}$  the orientations of the sensors.  $\mathbf{B}(\mathbf{r})$  is the three-dimensional vector-valued function given by Eq. 2.1. It is used here to predict the magnetic field at a sensor location  $\mathbf{r}$ . A sensor's output due to a source at  $\mathbf{r}_0$  with direction and strength  $\mathbf{Q}$  is evaluated with the dot product  $\mathbf{B}(\mathbf{r}) \cdot \mathbf{o}$ . Subsequently Eq. 2.2 yields a set of linear equations to be solved for the dipole amplitudes.

Unfortunately, the number of amplitudes being estimated usually exceeds the number of measurement channels. The problem is underdetermined and has no unique solution ( $\mathbf{L}$  is not invertible). The Moore-Penrose pseudoinverse can be used instead [20]:

$$\hat{\mathbf{j}}(t) = \mathbf{L}^+ \mathbf{b}(t) \quad (2.4)$$

$$\text{where } \mathbf{L}^+ = \mathbf{L}^T (\mathbf{L}\mathbf{L}^T)^{-1}$$

This constrains the problem by also minimising the Euclidean norm of the solution, which is the square root of the sum of squared source amplitude estimates.

$$\hat{\mathbf{j}}_{MN}(t) = \operatorname{argmin} \|\hat{\mathbf{j}}(t)\| \quad (2.5)$$

Owing to the choice of constraint, this method is known as *Minimum Norm Estimation* (MNE). The minimum norm estimate is optimal in the sense that it contains no energy that is not represented in the measurements [3, 20]. All other solutions consist of the minimum-norm solution plus a component that is invisible from the sensors' point of view. Chapter 4 will provide a more comprehensive treatment of MNE.

### 2.3.3 Beamforming

Like MNE, beamforming uses linear transformation to estimate sources of MEG signals. It also uses fixed dipoles — in this case a volumetric grid with three orthogonal components at each grid node. Beamforming differs from MNE in that spatial filters constructed from the signal covariance matrix are used to estimate the dipole amplitudes [21]. These filters can be derived in various ways, but a common type is the *minimum-variance beamformer* [22], described below.

First, a weighting matrix  $\mathbf{W}$  is defined, which relates the dipole amplitudes to the source signals.

$$\mathbf{j}(t) = \mathbf{W}^T \mathbf{b}(t) \quad (2.6)$$

The columns of the weight matrix,  $\mathbf{w}_i$ , represent individual spatial filters to be calculated from the signal covariances, which are estimated over a chosen time window on the data. For discrete-time data, the estimated covariance matrix,  $\mathbf{S}$  can be calculated as follows. The elements of  $\mathbf{b} = (b_1, b_2, \dots, b_m)^T$  are scalar functions of time. If the data window contains measurements at time points  $t_{1\dots n}$ , an  $m \times n$  matrix of mean-centred measurements can be defined:

$$\mathbf{M} = \begin{bmatrix} b_1(t_1) - \bar{b}_1 & b_1(t_2) - \bar{b}_1 & \cdots & b_1(t_n) - \bar{b}_1 \\ b_2(t_1) - \bar{b}_2 & b_2(t_2) - \bar{b}_2 & \cdots & b_2(t_n) - \bar{b}_2 \\ \vdots & \vdots & \ddots & \vdots \\ b_m(t_1) - \bar{b}_m & b_m(t_2) - \bar{b}_m & \cdots & b_m(t_n) - \bar{b}_m \end{bmatrix} \quad (2.7)$$

$$\text{where } \bar{b}_i = \frac{1}{n} \sum_{j=1}^n b_i(t_j)$$

Now the data covariance matrix may be defined in terms of  $\mathbf{M}$

$$\mathbf{S} = \frac{1}{n} \mathbf{M} \mathbf{M}^T \quad (2.8)$$

An expression for the power,  $P_i$ , in each time course to be estimated can be written in terms of the data covariance matrix and the weights.

$$P_i = \mathbf{w}_i^T \mathbf{S} \mathbf{w}_i \quad (2.9)$$

The weights are then determined by minimising  $P_i$  subject to the constraint that  $\mathbf{w}_i^T \mathbf{l}_i = 1$ ,  $\forall i$ , where  $\mathbf{l}_i$  are the forward-predicted field patterns created by hypothetical sources at each grid point. These vectors are also the columns of the lead field matrix defined in the previous section. The constraint sets the filter gain to unity at the region of interest while the gain is minimised elsewhere, creating a “virtual electrode” at the region of interest. The solution is given by

$$\mathbf{w}_i = \frac{\mathbf{S}^{-1} \mathbf{l}_i}{\mathbf{l}_i^T \mathbf{S}^{-1} \mathbf{l}_i} \quad (2.10)$$

Since the weights are dependent on the data, the method evidently relies on a statistical model to achieve signal separation.

### 2.3.4 Model Improvements

The three analysis methods described above rely on a forward model for predicting sensor signals from putative current densities. Conventionally, a conductive sphere is used to represent the head, primarily because this enables analytic solution of the forward problem (as in Appendix A). However, as the head is somewhat aspherical, the suitability of such a model is questionable, and it is easy to conceive of more realistic models that may require numerical solution but offer closer predictions of real-world MEG measurements. A better forward model corresponds to improved solutions to the inverse problem. To this endeavour, both boundary element (BEM) and finite element models (FEM) have been created which more accurately simulate the propagation of currents through the cephalic tissue. These inhomogeneous models are based on the anatomy of the subject, obtained from CT or MRI data. A realistic model might simulate different conductivities for the skull, brain, cerebrospinal fluid etc. These enhancements have been demonstrated to improve the accuracy of source localisation [23]. The improvement comes at the expense of considerably increased computation; solution of the most detailed FEM models requires the use of a supercomputer [24].

Analytic solutions are possible with certain conductor geometries because internal ohmic currents generate no magnetic field outside the conductor [25]. If the geometry is inaccurate,

this result does not hold true. Experimental data suggests that volume currents do indeed produce a measureable field [23]. An alternative model uses a piecewise homogeneous conductor, partitioned into arbitrarily shaped volumes of different conductivity [24]. In this case, the effect of the volume currents is given by surface integrals over the interfaces between homogeneous regions. The boundary element method discretises these surfaces into numerous, usually triangular, elements. The integral over each discrete element can then be evaluated by a suitable approximation (linear, quadratic etc.). The accuracy is determined by the size of the surface elements and the chosen approximation.

Finite element models discretise volume rather than surfaces, permitting the simulation of inhomogeneous and anisotropic media [24]. Inhomogeneity could be modelled using BEM by the use of large numbers of homogeneous subvolumes with intricate surface detail, but the balance of computational efficiency is in favour of FEM in such instances. The results produced by FEM also tend to be more accurate than those of BEM. However, BEM is most efficient in simpler models [26].

### Multiple Sphere Model

One means of improving the simpler spherical model is by substituting one of several overlapping spheres into the model for every sensor in the forward calculation [27]. The sphere that best approximates the head depends on the location and orientation of the sensor in question, so it is beneficial to adapt the parameters accordingly.

Two methods have been proposed for determining the sphere used for each sensor [27]. The first requires the sensitivity characteristics of the sensors (the *lead fields*) to be calculated on a predetermined grid using a boundary element method. For each sensor, the sphere centre is then adjusted so as to maximise the agreement between lead fields calculated using the multiple sphere model and the BEM model. The principal benefit in this case is the ability to estimate an ECD with an accuracy approaching that of the BEM model, without interpolating between the discrete points of the BEM lead fields. To achieve this using a boundary element model would require re-evaluation of the computationally intensive

model upon every iteration of the minimisation procedure.

The second method obviates the need to construct another model at all. Instead, a geometric approach is used to fit the spheres to the curvature of the skull local to the relevant sensor [27]. To achieve this, the interior surface of the skull and the surface of the sphere are converted to meshes containing equal numbers of points. The contribution to the relevant sensor from the volume currents at each point on the two meshes is calculated. A cost function is then defined as the sum of squared differences between the contributions from the corresponding points on the two meshes [27]. The sphere parameters can be adjusted using the downhill simplex algorithm to minimise the cost function.

Compared with the single sphere model, the accuracy of source localisation performed using the overlapping sphere model is much closer to that of the boundary element model [27]. As an enhancement to the classical dipole fitting method, the multiple sphere model is particularly attractive since the computational cost is almost negligible. A multiple sphere model was used for the later examples in this thesis.

### **2.3.5 Discussion**

Three classes of MEG source estimation have been considered: dipole fitting, MNE and beamforming. Dipole fitting approaches are the most basic. The model is very rudimentary and cannot be used to resolve arbitrarily-shaped current densities. In practice it is good at estimating the location of a single focal source, such as the epileptic focus in the brains of patients with certain types of epilepsy. The required numerical minimisation procedure presents an implementation challenge because the cost function has several local minima; care must be taken to promote convergence on the global minimum. Due to the possibility of “silent sources” that do not contribute to the measurements, numerically determined dipole parameters may include spurious components.

Unlike dipole fitting, MNE facilitates resolution of arbitrary distributions of current density. By fixing the dipole locations the problem is linearized, which enables simple analytical solution. Due to the underdetermined nature of the problem there is not

enough information in the data for the estimated dipole amplitudes to be independent, so estimates inherently appear smooth. Also, the pseudoinverse is ill-conditioned, leading to gross amplification of measurement noise in the solution. This can be mitigated using *regularisation*, whereby the solution is allowed to deviate slightly from the data in order to improve numerical stability. The penalty is reduced sharpness in the spatial resolution.

Like MNE, beamforming can be used to resolve tomographic maps of activity. However, the spatial filtering approach achieves much sharper distinction between nearby sources than MNE [21]. This is achieved by taking into account temporal as well as spatial information. The statistical model underpinning the spatial filters requires an assumption of linear independence between sources [22]. If this condition is not met, signal leakage or migration between regions can occur.

It is apparent that none of the three methods considered above is universally superior. Ultimately, the application will determine the preferable approach, based on alignment with the requirements. More infrequently, MEG analysis also includes general blind source separation methods such as *principal component analysis* or *independent component analysis*[28]. Otherwise, the raw recordings may simply be viewed in the time or frequency domain.

## 2.4 Summary

In this chapter, the anatomy of the central nervous system has been considered on different scales. The most important structures in this context are the two cerebral hemispheres of the brain, which are thought to contribute most to the MEG signal. The cerebral cortex is densely packed with neurons, which are interlinked by cell processes, forming neural networks. The dendrites, which synapse onto other neurons, carry the impulsive currents that are responsible for the majority of the neuromagnetic signal.

Three classes of MEG source estimation have also been considered: dipole fitting, MNE and beamforming. The classical method of dipole fitting is unable to resolve spatially extended sources, but is good at localising highly focal ones. MNE is a linear transformation

that maps sensor signals onto a putative source-space. It tends to blur the current density estimates because there is a far greater number of sources to be estimated than there are sensor signals. Beamforming is also a linear transformation, but it adapts to the data by constructing spatial filters from the signal covariance matrix. This allows beamforming to achieve sharper spatial resolution than MNE.

## 3 Concepts

The underlying objective of the project described herein was to provide a novel contribution to the field of MEG analysis, and a number of broad areas were considered for research. The sections in this chapter discuss those candidate topics, which should provide some context for the concept that was chosen for development.

### 3.1 Realistic Brain Models

The smallest signals measured by MEG are due to the ensemble effect of at least a million simultaneous synaptic currents [3]. As previously stated in §2.1.2, the currents are due to the interaction between neurons, which form a large network of computational elements. Even if these elements are stochastic, it is nonetheless possible to predict certain trends in their behaviour. As a result, the measurable signals arising from the neuronal currents are liable to take on predictable characteristics in terms of spatial distribution and progression over time. One way of making the biomagnetic inverse problem introduced in §2.3 more tractable would be to take into account the probable spatial distribution and time course of neuronal sources, based on a Bayesian model for the logical behaviour of the brain. Rather than arbitrary current distributions, the model parameters would be estimated from the measurements, leading to a complete description of the neural sources.

Despite the apparent complexity of the human brain, the relevant models need not necessarily be incomprehensibly large. Bearing in mind the limited amount of genetic information that codes for the development of the human brain, one would expect to find a large degree of self-similarity. It is therefore conceivable that a model consisting of

a manageable number of abstract elements would provide an adequate substitute for real networks of inordinate numbers of neurons.

Aspects of brain electrical activity have been simulated on various different scales. These begin with individual neuron models, such as the Hodgkin-Huxley model [29], extending to vast artificial neural networks made up of semi-realistic elements exhibiting many of the behaviours of their biological counterparts [30]. Much of the work in this area has focused on producing networks with intelligent characteristics, without being concerned with the adequacy of the network as a model for predicting activity in biological networks.

A potential problem with using approximate or abstract elements in brain models is that the emergent properties of networks comprising such elements could be rather different from those of biological neural networks. Even the most approximate models that exhibit brain-like characteristics are extremely complicated and require enormous computing resources. Repeating a simulation as part of an iterative process would see these demands increase exponentially. Considerable efforts have already been invested in producing models which explain characteristic empirical MEG/EEG signals [31, 32, 30]. So-called *neural mass* models attempt to describe the mean behaviour of large populations of neurons with small numbers of state variables.

A form of MEG analysis that is informed by models of the brain has been developed recently, entitled *dynamic causal modelling* (DCM) [33]. Specifically, DCM employs neural mass models to investigate dynamic connectivity among populations of neurons, based on information obtained from functional neuroimaging experiments (including MEG). Given the extent of these developments, it was felt that there was limited scope for offering novel contributions in this area within the time frame of the project. The next section investigates connectivity as an alternative area of research.

## **3.2 Connectivity**

One of the most pertinent questions in neuroscience is how different regions of the brain influence each other. Many of the traditional analytical tools applied to MEG data provide

little in the way of determining this information, leading to conjecture and uncertainty surrounding the meaning of experimental results. Subjective observations of similarity between activations at different locations are not sufficient to draw firm conclusions about their relationship. There is clearly a need for more rigorous means of establishing interdependence between observed patterns of activity. One area of particular interest is dynamic connectivity, which refers to the hypothesis of logical reorganisation of the brain according to task.

Connectivity measures in MEG are typically based on linear correlation between variables. At the most rudimentary, these variables may simply be different sensor channels. In this case, observed correlation should be treated with caution, because the signals at each sensor are known to be combinations of a common group of underlying signals. Also, one cannot assume that different sensors are selective of particular sources. While it is true that the sensitivity profile of MEG sensors favours sources in the locality of the sensor, the sources themselves do not usually produce the greatest magnetic field intensity at the most proximal sensor. In solution to these problems, correlation may instead be determined between signals in “source space”. This means that the random variables on which the measures are calculated are estimates of source signals from chosen points within the source volume, derived using some external method. Potential choices of method for obtaining these source signals are discussed in §3.3.

When interpreting measures of correlation in multivariate data, it is important to realise that observations of two variables which have no influence over each other may nonetheless be correlated. This readily occurs if the two variables in question have common dependency on a third variable; the first two variables may be *conditionally independent*, given the third variable. In these circumstances, measures of *partial correlation* are useful. Partial correlation coefficients quantify the association between pairs of variables, conditional upon the other variables in the system. Partial correlation alone cannot be used to draw conclusions about causality since it reveals nothing about the direction of influence. Also, these measures only exclude the possibility of high correlation being due to the influence of

a third variable if all variables in the system are known. Partial correlation coefficients form the basis for conditional independence graphs — graphical models that represent the connectivity structure of the variables in a system. Evidently, this has applications in representing connectivity, including dynamic connectivity, between signal sources in the brain.

Dynamic causal models, mentioned briefly in the previous section, take a different approach to establishing connectivity. Although a variety of models are amenable to the principal of DCM, the system to be studied is normally represented by a vector of state variables, with a set of parameters governing the system dynamics. Then a set of differential equations expresses the rate of change of each state variable as a function of the state variables, the inputs to the system (if any) and the system parameters [33]. Established connectivity models in this form are used to simulate the behaviour of small cortical regions. From this, a forward model such as those described in §2.3 is used to generate predicted MEG signals. The parameters of the connectivity model are then optimised to fit the data. Typically, this type of analysis is applied to experimental data to provide evidence of changes in connectivity according to the task being carried out by the subject.

DCM is only one answer to the limitations of statistical methods. The use of correlation measures for studying connectivity is developed further in the following section, which describes how they can be incorporated in a multi-stage processing framework that is tailored to the analysis of MEG data.

### 3.3 Combined Signal Processing Framework


The output signals from the MEG channels consist of various weighted sums of intracranial sources, plus noise. The first requirement of a system that can be used to explore connectivity is the ability to isolate individual sources as far as possible. It is insufficient to simply select individual sensor channels near regions of interest, even though sensitivity diminishes as distance squared. Firstly, the regions in question may not be known *a priori*. Secondly, depending on the direction of current flow, there may be a null in the

magnetic field pattern at the most proximal sensors. Lastly, the signals on different sensor channels are highly correlated, with considerable leakage from distant regions. Statistical associations between channels do not necessarily indicate connections in the brain.

In the previous chapter, two methods for projecting the raw data into a putative “source space” were introduced. MNE and beamforming both have the ability to estimate a time course from a given source location, almost as though an electrode were inserted directly into the brain. By producing tomographic maps of activation, both methods also provide the means to identify a few regions of interest where changes in activity are observed. This could be performed by a human being or a computer algorithm. The result of this selection process is a manageable number of time series to be used as inputs to a statistical analysis. This provides real benefits in terms of data reduction. MEG data sets are typically very large and incomprehensible. Distilling the salient features by estimating the activity in a few important locations reduces the volume of data. For these reasons, a source-space projection was chosen as a pre-processing stage prior to statistical analysis, shown as Stage 1 in Figure 3.1.

The spectrum of MEG signals covers several decades, while the power falls off with increasing frequency. Interaction at different frequencies is one of the postulates to be investigated. One could calculate statistical measures directly on the time series produced by source-space projection, but the dominant low-frequency components would hide any higher frequency interactions. Simple high pass filters or whitening of the data are not the solution. There is no clear choice of cutoff frequency, and the wide-band signal would still contain a *mêlée* of signals produced by different neural mechanisms. An intermediate processing stage consisting of a time-frequency analysis has been incorporated into the analysis framework. This allows frequency-specific effects to be analysed separately, or even cross-frequency effects where the amplitude envelope in one band is seen to influence that of another band at a distant site. It appears in Figure 3.1 as Stage 2.

The final stage of the framework is the statistical analysis, shown as Stage 3 in Figure 3.1. Ordinary correlation coefficients are uninformative when comparing observations from




**A** Raw MEG signals **B** Estimated source signals  
**C** Time-frequency representation of sources **D** Correlation coefficients

Figure 3.1: Data flow diagram of MEG signal processing framework

distant brain regions because there is inevitably a finite propagation time. These propagation times can be estimated using cross-correlation and adjusted for. However, in a system with many variables it is necessary to estimate partial correlations that compare pairs of variables, conditional upon the remaining ones. Otherwise, two variables may appear to be correlated because they are both driven by a common third variable. Once delay-adjusted partial correlation coefficients have been determined, the system can be represented graphically using conditional independence graphs. A conditional independence graph is a useful way to visualise a model of a multivariate system. Random processes are represented by nodes on the graph. Edges between the nodes represent statistically significant partial correlations. The term “conditional independence graph” reflects the fact that the absence of an edge between two processes is due to those processes being independent, conditional upon the other processes shown in the graph. An example of these is shown in Figure 3.2.

A data flow diagram of the three-stage process is shown in Figure 3.1. Using this approach it should become possible to draw conclusions about functional connectivity that are rooted in the data and hence supported by experimental evidence. For the source space projection, the MNE method has been chosen because it permits current estimates that provide a good fit to the data without assuming statistical independence between sources.



In this example of a conditional independence graph, processes 1, 3 and 4 are conditionally independent of each other, meaning that any correlation between them is attributable to their common relationship with process 5. Likewise, process 2 is conditionally independent of process 5, but could be related via process 3.

Figure 3.2: Conditional independence graph

#### 3.3.1 Summary

A novel signal processing framework for the analysis of MEG data has been devised, combining a source model inversion, a time-frequency transformation and a partial correlation analysis for the identification of neural connectivity. By assuming as little as possible about the source distribution and the nature of sources, many of the pitfalls encountered with other MEG analyses should be avoided. At the same time, the framework has the potential to provide a much-needed objective view of interactions between cortical regions. Full details on the theory and implementation of the three constituent processes of the framework are presented in the following chapters.

## 4 Minimum Norm Estimation

Minimum norm estimation (MNE) is a particular approach to reconstructing magnetic sources from MEG recordings. Source reconstruction is the principal objective of many forms of MEG analysis. However, MNE is distinct from other methods due to the type of source model used and the means of estimating the model parameters. As we shall see, MNE can be used to estimate arbitrary spatial distributions of neuronal activity, and assumes nothing about statistical associations between activity in different regions of the source space. A frequently cited disadvantage of MNE is weak spatial resolution [34, 35]. However, precise source localisation is not a key requirement of the framework proposed herein. In this case, the aim of MNE is to find cohesive assemblies of neurons and estimate their time course. It is not crucial to know the exact focus of such assemblies, and consequently the spatial resolution is unimportant.

Unlike the chapters concerning the other modules in the framework, this chapter serves mainly as a literature survey, with a view to further justifying the inclusion of MNE in the framework. The technique as it is described here is supposed to represent the state of the art without offering novel contributions. Discussions regarding the detailed implementation of MNE will not be exhaustive. Instead they will focus on the algorithms used to obtain the experimental results presented in Chapter 6.

### 4.1 Definition

It is perhaps unfortunate that, in the context of MEG analysis, the word “source” describes a number of different concepts. In classical dipole fitting methods a putative source is

represented by an individual dipole, the *equivalent current dipole* (ECD) [3]. Either a single source is assumed or a small number of postulated sources are modelled by an equal number of dipoles. In both cases, the source dimensions are held to be vanishingly small. Moreover, the location and direction are permitted to vary with time (unless this has been constrained a priori). This definition of a source is potentially ambiguous since multiple dipoles might be required to represent a spatially extended neuronal population acting in concert [36]. Such a population could equally be termed a “source”. Furthermore, the apparent migration of one dipole source could represent, for example, cessation of activity in one region accompanied by onset of activity in another [37]. Should this really be regarded as one source that has moved or two sources whose major activity took place in succession?

An alternative definition for a source may be more helpful: a static current density that is modulated by a scalar function of time, called the amplitude. With this definition, the temporal progression of each source can be described by an unambiguous amplitude-versus-time signal. Any variation in the topography or directional characteristic of the total current density comes about as a result of the net contribution of multiple sources. The magnetic fields due to individual sources are summed at the sensors by the principal of superposition. If a model exists that is able to predict how each static source contributes to the field measurements, this definition makes it convenient to write down the linear relationship between MEG signals,  $\mathbf{b}(t)$ , and source signals,  $\mathbf{j}(t)$ . Noise is represented by  $\mathbf{e}(t)$

$$\mathbf{b}(t) = \mathbf{L}\mathbf{j}(t) + \mathbf{e}(t) \quad (4.1)$$

This formula simply states that the field at each sensor is a weighted sum of the source signals. By contrast, the source-to-signal relationship is much less straightforward if the current densities underpinning the source model are permitted to vary, as in ECD models. Using those models, the measurements have a nonlinear relationship to the source parameters [38, 39]. By using fixed current densities and estimating only their amplitudes, the problem is simplified, yet topographical changes in the total current density can still be accounted for by variations in the relative amplitudes of the underlying current densities.

Up to this point, nothing has been stated about the nature of the static current densities,

nor their cardinality. Without a priori knowledge of the source distributions, the current densities cannot readily be fixed. However, given a putative source-space geometry, such as the brain volume, then a uniform grid of current dipoles may be used to approximate arbitrary current densities [40]. The resulting matrix of weightings for each sensor is then known as a *lead field matrix*. Each row in the matrix is a lead field vector, which describes the sensitivity of one particular MEG channel towards different regions of the source volume (and different directions of current flow within those regions). It would be desirable to invert the lead field matrix so that the topography of the activations can be found:

$$\hat{\mathbf{j}}(t) = \mathbf{L}^{-1}\mathbf{b}(t) \quad (4.2)$$

Unfortunately, the problem is almost always underdetermined; there are fewer MEG channels than source amplitudes to be estimated, hence more unknowns than equations. The lead field matrix has no inverse, therefore no unique solution exists to the inverse problem. It is nonetheless possible to find a unique solution provided there are additional constraints. The minimum norm estimate is thus named because it results in a solution vector,  $\hat{\mathbf{j}}$ , that has minimal Euclidean length, or norm; viz. the sum of squared source amplitudes given as the MNE is less than that of any other solution that satisfies the lead field relation in Equation 4.1. Calculating the minimum norm estimate is a case of finding  $\hat{\mathbf{j}}$  such that

$$\|\hat{\mathbf{j}}\|^2 = \sum_{\forall i} \hat{j}_i^2 \quad \text{is minimised, with } \mathbf{L}\hat{\mathbf{j}} = \mathbf{b} .$$

The MNE is favoured over other solutions because the estimated current density only has components that are represented in the measurements. Any other solution will consist of the MNE plus a current density that produces no MEG signal [3, 20]. More formally, if  $F$  is the infinite function space containing all possible current densities, the MNE is the projection of the true current density in the subspace of  $F$  that is spanned by the lead field vectors [41, 3].

## 4.2 Solution

Previously, we defined the MNE as the source amplitude vector with the shortest Euclidean length that reproduces the measured data upon application of the forward model. In order to find this, we must substitute  $\mathbf{L}^{-1}$  in Equation 4.2 (which does not exist) with a suitable inverse operator. The inverse operator that gives the minimum norm solution is the Moore-Penrose pseudoinverse [42], denoted  $\mathbf{L}^+$ . This may be found using some factorisation of  $\mathbf{L}$ . In general, a least-squares problem such as this might be solved using singular value decomposition (SVD) [40]:

$$\mathbf{L} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T \quad (4.3)$$

Here, if  $\mathbf{L}$  is  $m \times n$ ,  $\mathbf{U}$  and  $\mathbf{V}$  are  $m \times m$  and  $n \times n$  orthogonal matrices respectively.  $\mathbf{\Sigma}$  is an  $m \times n$  matrix with nonzero values only on the leading diagonal:

$$\Sigma_{ij} = \begin{cases} \sigma_i & i = j \\ 0 & \text{otherwise} \end{cases}$$

The *singular values*,  $\sigma_i$  are positive or zero, and decrease monotonically, i.e.  $\sigma_1 \geq \dots \geq \sigma_m \geq 0$ .

The above factorisation enables computation of the pseudoinverse using the formula,

$$\mathbf{L}^+ = \mathbf{V}\mathbf{\Sigma}^+\mathbf{U}^T \quad (4.4)$$

where  $\mathbf{\Sigma}^+$  is an  $n \times m$  diagonal matrix whose diagonal elements are given by:

$$\Sigma_{ii}^+ = \begin{cases} \frac{1}{\Sigma_{ii}} & \Sigma_{ii} \neq 0 \\ 0 & \text{otherwise} \end{cases}$$

This approach requires prior computation of the SVD. In practice, this is not necessary; an  $m \times n$  matrix of rank  $r$  can also be factorised into a product between an  $m \times r$  matrix and an  $r \times n$  matrix. We shall denote the multiplicands  $\mathbf{C}$  and  $\mathbf{D}$  respectively, so that  $\mathbf{L} = \mathbf{CD}$ . Subsequently, the pseudoinverse of a real-valued lead field matrix is given by [43]

$$\mathbf{L}^+ = \mathbf{D}^T (\mathbf{DD}^T)^{-1} (\mathbf{C}^T \mathbf{C})^{-1} \mathbf{C}^T. \quad (4.5)$$

Since the various MEG sensors have different perspectives on the magnetic field, the lead fields are linearly independent; the lead field matrix is of full rank. The rank is consequently equal to whichever is the smaller of  $m$  and  $n$ . As there are normally fewer sensor channels than estimated sources ( $m < n$ ), the rank is equal to  $m$ , or the number of sensors. In this case, Equation 4.5 can be simplified by choosing  $\mathbf{C}$  to be identity, implying that  $\mathbf{D} = \mathbf{L}$ . The inverse operator can then be calculated by

$$\mathbf{L}^+ = \mathbf{L}^T (\mathbf{L}\mathbf{L}^T)^{-1} . \quad (4.6)$$

In place of Equation 4.2, we now have a calculable formula for the minimum norm estimate in its fundamental form:

$$\hat{\mathbf{j}}(t) = \mathbf{L}^T (\mathbf{L}\mathbf{L}^T)^{-1} \mathbf{b}(t) \quad (4.7)$$

Notably, the inverse operator depends only on the forward model and not on the measurements themselves, showing that MNE does not rely on any statistical model for the data. Usefully, the inverse operator only needs to be calculated once for a whole data set.

The above solution pays no regard to measurement error, which may be problematic in applications like MEG that typically suffer from poor signal-to-noise ratio (SNR). Incorporating a model for the noise into the solution can be achieved through regularisation, which is discussed in the next section.

### 4.3 Regularisation

MEG sensors exhibit very low sensitivity towards sources in deep regions of the brain, and to those that are predominantly radial in their orientation [25, 7, 20]. For these sources, the error in the measurements (i.e. the noise) can become massively amplified when they are projected in the source space. This noise amplification takes place because the problem is ill-conditioned, thus a small change at the input causes large changes at the output. The practical consequence of this numerical instability is that MNE will sometimes resolve physiologically implausible source distributions because they happen to minimise the power in the solution. For example, the minimum norm solution might contain an extraordinarily

high-amplitude peak that is sharply localised in a low-sensitivity region, when in fact a more credible solution might have a smoother distribution and modest amplitude. A widely-used remedy to this phenomenon is regularisation.

Regularisation is the practice of finding an approximate solution that improves the numerical stability. By approximating the inverse operator, the process itself introduces error, but this is to be traded off against the attenuation of measurement error. Multiple regularisation approaches have been developed [44], but they all have the effect of suppressing sources that are not well represented in the measurements. Often, the degree of regularisation can be controlled by a single parameter which is adjusted to some optimal point where the effect of measurement noise is reduced without introducing significant approximation error.

One form of single-parameter regularisation that can very simply be applied to discrete data has become popular in this application. Tikhonov regularisation [44, 39, 20] is performed by adding a scaled identity matrix,  $\lambda\mathbf{I}$ , to the reciprocal part of the inverse operator.

$$\mathbf{L}_R^+ = \mathbf{L}^T (\mathbf{L}\mathbf{L}^T + \lambda\mathbf{I})^{-1} \quad (4.8)$$

Here, the regularisation parameter is  $\lambda$ . The non-regularised solution in Equation 4.6 is given by setting  $\lambda = 0$ . Greater values of  $\lambda$  increase the severity of regularisation. To account for the effect of measurement noise on the transform, a noise covariance matrix,  $\mathbf{S}$ , can be substituted for the identity matrix. In addition, it is sometimes desirable to incorporate *a priori* information about the statistical relationships between sources. This may be obtained from another functional neuroimaging modality such as PET or fMRI. In such cases the source covariance matrix,  $\mathbf{R}$ , is applied as follows:

$$\mathbf{L}_R^+ = \mathbf{R}\mathbf{L}^T (\mathbf{L}\mathbf{R}\mathbf{L}^T + \lambda\mathbf{S})^{-1} \quad (4.9)$$

Evidently, if  $\mathbf{R}$  and  $\mathbf{S}$  are identity matrices, the solution is unchanged from the previous case.  $\mathbf{R} = \mathbf{I}$  implies that sources are independent and of equal variance;  $\mathbf{S} = \mathbf{I}$  makes similar assumptions about the noise components on the sensor channels [39].

Another regularisation technique that has been applied to MEG is *truncated SVD* [41]. Referring to the SVD method for finding the inverse operator (Equation 4.4), it is notable that the columns of  $\mathbf{V}$  (the *right singular vectors*) represent abstract source distributions. These form an orthogonal basis from which other source distributions may be constructed. The corresponding singular values then determine the effective gain acting upon each component as it is transmitted to the sensor array. As these values approach zero, the compensating gain required in the inverse operator becomes very large, leading to ill-conditioning of the matrix and amplification of measurement error. Regularisation can therefore be achieved by setting a proportion of the smallest singular values to zero (which remain zero in  $\Sigma^+$ ). The corresponding components are thereby assumed to contribute nothing to the data. The number of components to reject is determined by the regularisation parameter. This parameter can either set a threshold as a fraction of the largest singular value, or simply specify the number of values to be removed.

Application of a regularised inverse results in an estimate that does not precisely fit the data. If the discrepancy lies within the bounds of error then the regularised solution can be justified. Accordingly, there is a limit on the extent to which regularisation should be applied. A concrete procedure is therefore needed to establish an appropriate regularisation parameter.

### 4.3.1 Regularisation Parameter Selection

It was previously stated that regularisation is a compromise between approximation error and suppression of noise, governed by a regularisation parameter,  $\lambda$ . It is not possible to determine the magnitude of the approximation error introduced by regularisation, as this requires knowledge of the true source current density. If the current density is concentrated away from regions of low sensitivity then higher values of  $\lambda$  are more acceptable. Consequently there is no definitive way to choose the value of  $\lambda$ , and a diversity of approaches can be found in the literature. Three contemporary methods are described here.

### L-curve

The L-curve is a graphical plot of the norm of the regularised solution,  $\|\mathbf{L}_R^+ \mathbf{b}\|$ , versus the norm of the residual,  $\|\mathbf{L}\hat{\mathbf{j}}_R - \mathbf{b}\|$ , on a log-log scale [45]. It provides a means of visualising the balance between improvement of the solution afforded by regularisation and the consequent discrepancy with the data. The name derives from the characteristically L-shaped curve, which is parameterised by  $\lambda$ . Increasing  $\lambda$  moves down the curve. Figure 4.1 shows the typical form of an L-curve. The steep section of the L-curve occurs when additional regularisation is of most benefit since the residual increases slowly compared with the decrease in the solution norm. This decrease is indicative of noise suppression. In the horizontal section beyond the “knee”, the residual norm increases in magnitude more rapidly. This implies that too much regularisation has been applied and the solution is rapidly deviating from the data. It is therefore desirable to choose a value of  $\lambda$  close to the knee [46].

The concept of an L-curve has existed for several decades, but it is only more recently that a systematic procedure for locating the knee has been published [46]. Since then, the method has been applied to inverse problems across many fields, including MEG analysis [47, 48]. The reader is referred to [46] for full details, but in essence  $\lambda$  is calculated as the point of maximum curvature on the curve  $(\hat{\rho}, \hat{\eta})$ , where  $\hat{\rho} = \log \|\mathbf{L}_R^+ \mathbf{b}\|$ , and  $\hat{\eta} = \log \|\mathbf{L}\hat{\mathbf{j}}_R - \mathbf{b}\|$ . The curvature is calculated by

$$\kappa(\lambda) = \frac{\hat{\rho}'\hat{\eta}'' - \hat{\rho}''\hat{\eta}'}{((\hat{\rho}')^2 + (\hat{\eta}')^2)^{\frac{3}{2}}}.$$

The situation is complicated if, as in truncated SVD, the regularisation parameter is discrete, because the L-curve consists of piecewise sections and is not continuously differentiable. In this case, cubic splines can be fitted to the points and differentiated instead to determine which point is closest to the knee. However, since points are often clustered away from the gross features of the curve, the high-frequency fluctuations must first be smoothed out to avoid converging on an irrelevant part of the curve.

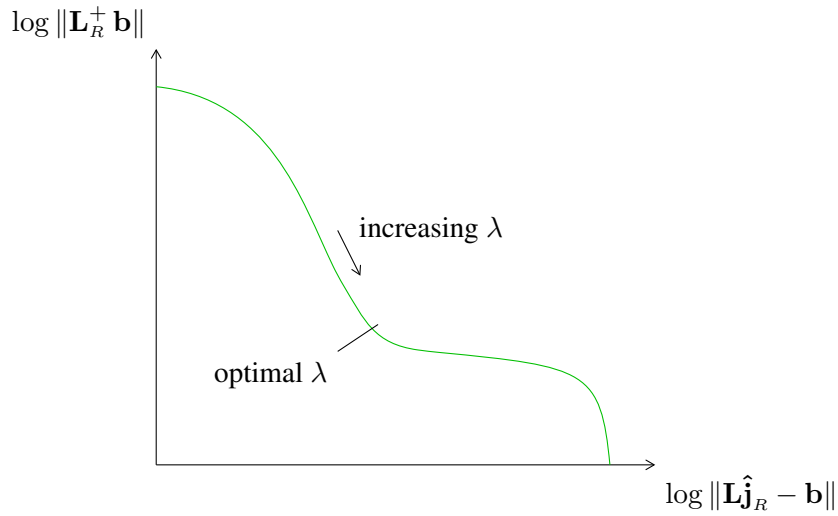


Figure 4.1: Representative diagram of an L-curve

### Equal SNR

As discussed earlier in this section, the need for regularisation arises because errors are amplified by projecting the data in the source space. For data that progress over time, a corresponding degradation can be observed in the SNR. Consequently, one way to account for the effect of noise on MNE is to regularise to the extent that the SNR of the transformed data approximates that of the sensor-space data [35]. To achieve this, comparable measures of the data and solution SNR must be defined. The following is equivalent to the definition given in [35].

The columns,  $\mathbf{l}_i$ , of the  $m \times n$  lead field matrix represent field topographies at the sensors arising from hypothetical sources in different regions of the source space. It is helpful to also define  $\hat{\mathbf{l}}_i$  that have been normalised to have unit power (or Euclidean length  $\sqrt{m}$ ). Then the SNR of the data is given by the ratio between the mean of the leadfield topographies' root mean square (RMS) values and that of the source topographies (note that the RMS of a vector  $\mathbf{a} \in \mathbb{R}^n$  is equal to  $|\mathbf{a}|/\sqrt{n}$ ):

$$\text{SNR}_D = \frac{\sum_{i=1}^n |\hat{\mathbf{l}}_i|}{n\sqrt{m}} \div \frac{\sum_{t=1}^T |\mathbf{b}(t)|}{T\sqrt{m}} \quad (4.10)$$

Due to the leadfield topography normalisation, the dividend equals one. The SNR is then defined simply as the reciprocal of the time average of the source topography RMS.

$$\text{SNR}_D = \frac{T\sqrt{m}}{\sum_{t=1}^T |\mathbf{b}(t)|} \quad (4.11)$$

For the SNR of the solution, the normalised leadfield topographies are transformed using the regularised inverse operator,  $\mathbf{L}_R^+$ . The resulting  $n$ -dimensional source-space vectors represent the spatial blurring properties of the regularised inverse. The solution SNR is the mean RMS of each transformed leadfield topography divided by the time average of the solution RMS:

$$\text{SNR}_S = \frac{\sum_{i=1}^n |\mathbf{L}_R^+ \hat{\mathbf{i}}_i|}{n\sqrt{n}} \div \frac{\sum_{t=1}^T |\mathbf{L}_R^+ \mathbf{b}(t)|}{T\sqrt{n}} \quad (4.12)$$

$$= \frac{T \sum_{i=1}^n |\mathbf{L}_R^+ \hat{\mathbf{i}}_i|}{n \sum_{t=1}^T |\mathbf{L}_R^+ \mathbf{b}(t)|} \quad (4.13)$$

Using these measures, it is now possible to find a value of  $\lambda$  such that  $\text{SNR}_S \approx \text{SNR}_D$ . The original authors do not suggest how to find this value, but the relation could be rewritten as a minimisation problem.

$$\min_{\lambda} \left| \frac{\sum_{i=1}^n |\mathbf{L}_R^+ \hat{\mathbf{i}}_i|}{\sum_{t=1}^T |\mathbf{L}_R^+ \mathbf{b}(t)|} - C \right| \quad (4.14)$$

$$\text{where } C = \frac{n\sqrt{m}}{\sum_{t=1}^T |\mathbf{b}(t)|}$$

The constant  $C$  is independent of  $\lambda$ , requiring calculation once per data set. The expression can then be minimised numerically, with the approximation precision determined by the convergence criteria.

### Direct Calculation

Numerical approaches to regularisation parameter estimation such as those detailed above require initial estimates on the first iteration. Some investigators prefer to avoid iterative methods altogether. When using Tikhonov regularisation, a simple formula may be used to calculate a first estimate for  $\lambda$  [49].

$$\lambda = \frac{\text{Tr}(\mathbf{LRL}^T)}{\text{Tr}(\mathbf{S})} \quad (4.15)$$

where the trace of a matrix is defined as the sum of the elements on the leading diagonal:

$$\text{Tr}(\mathbf{A}) = \sum_{\forall i} A_{ii} . \quad (4.16)$$

Since the trace of a matrix is also equal to the sum of its singular values, it is evident from Equation 4.9 that the formula above is designed to make the regularising term,  $\lambda\mathbf{S}$ , commensurate with the unregularised inverse operator. It would only be reasonable to take this as a final value if the magnitude of the noise approached that of the signal. Otherwise, the outcome would be excessive regularisation. For MEG data, one would usually hope to achieve a SNR greater than unity. A more reasonable direct estimate based on Equation 4.15 further divides the quotient by the square of the SNR (expressed as a power ratio) [50]. Lin and colleagues suggest a value of 5 as an appropriate substitute if the SNR is not determined empirically.

$$\lambda = \frac{\text{Tr}(\mathbf{LRL}^T)}{\text{Tr}(\mathbf{S}) \text{SNR}^2} \quad (4.17)$$

In contrast with the previous two methods, this approach does not directly depend on the data. As a result, the regularised inverse can be calculated once and reapplied to several data sets, provided that the source prior matrix, noise covariance and signal-to-noise estimate remain constant.

## 4.4 MNE In Practice

The mathematical principles of MNE have already been described in this chapter. However, there are several practical considerations involved in calculating MEG source estimates

using the technique. The most important of these are detailed in this section.

#### 4.4.1 Source Space

In §4.1 the idea of using a grid of dipole sources to represent arbitrary current densities was introduced. One could define the source space as a uniform grid covering the brain volume (or the conductor volume employed in the source model). However, a significant proportion of these points would lie in deeper brain regions to which MEG is known to be insensitive. As discussed in the previous section, this would generally lead to spurious current density estimates unless these sources were suppressed using regularisation. At the same time, if a more appropriate source space was selected, the degree of regularisation could potentially be reduced, providing a better fit to the data. A total number of estimated sources greater than one or two orders of magnitude above the number of sensor channels is not worthwhile. The information recorded contains too few orthogonal components to resolve further meaningful detail. Given this limitation, putative source locations should ideally reflect the probable regions from which measured signals arise. For these reasons it is common to restrict the source space to the cerebral cortex, the geometry of which may be obtained from structural imaging such as Magnetic Resonance Imaging (MRI) or Computed Tomography (CT) [39, 40].

Because current density is a 3-dimensional vector quantity, three orthogonal components are required at each source location to represent arbitrary directions of current flow. Therefore there may be up to three times as many columns in the lead field matrix as there are source locations. In the case of a spherical head model, this may be reduced to two orthogonal tangential sources since the radial component is invisible at the sensors. More realistic head models facilitate estimates with three dimensions, so this simplification is not always available. There is evidence that MEG signals are predominantly due to dendritic currents normal to the cerebral cortex [39, 40]. This is cited as justification for assuming a single component in the normal direction at each point on the cortex [51].

### 4.4.2 Depth Weighting

The sensitivity profile of MEG magnetometers follows an inverse-square law, i.e. they are decreasingly sensitive in proportion to the source distance squared. Consequently, activations in more superficial regions of the brain appear more amplified in the measurements than those in deeper regions. Because this variation is naturally present in the lead fields, the MNE will inherently bias the estimated current density towards superficial sources. In particular, it can be demonstrated that deeper sources tend to be incorrectly localised such that they appear to be located too superficially [51].

SNR typically worsens as source depth increases as a result of distance-related signal attenuation versus constant noise power. Impaired SNR will inevitably compromise localisation accuracy. However, the technique of *depth weighting* can reduce the systematic bias caused by the non-uniformity of the lead fields [51, 52]. Depth-weighting is performed by modifying the *a priori* source covariance matrix so that sources appearing weaker in the lead fields are assumed to have greater variance. This has the effect of coercing the power in the solution towards deeper regions, thereby counteracting the inherent bias towards superficial ones. The modified matrix can be determined in one of two following ways.

#### Linear Scaling

The simplest form of depth weighting in common use is a linear scaling that requires a parameter [51]. A sensitivity value is calculated for each distinct source location, which is raised to a power specified by a depth weighting parameter,  $p$ . The diagonal elements of the source covariance matrix,  $\mathbf{R}$ , are divided by the resultant value. In other words, the prior variances for a given location are adjusted in inverse proportion to the sensitivity in that location. The sensitivity value is the sum of the power in the sensor topographies arising from the relevant source components. For example, where there are three components for each source location, the value is given by [51]

$$(\mathbf{l}_{3i}^T \mathbf{l}_{3i} + \mathbf{l}_{3i-1}^T \mathbf{l}_{3i-1} + \mathbf{l}_{3i-2}^T \mathbf{l}_{3i-2})^p \quad (4.18)$$

Here,  $3i$ ,  $3i - 1$  and  $3i - 2$  are the indices of three adjacent columns of the lead field matrix that correspond to an orthogonal set of co-located dipoles. Note that the same adjustment is applied to the three components.

This depth weighting method leaves the choice of parameter to the user. Based on established experimental protocols with human subjects, a parameter value around  $p = 0.7$  has been found to most accurately localise the expected activity [51]. Others have suggested a more modest adjustment of  $p = 0.5$  [53].

### Largest Singular Value

A more sophisticated approach to depth weighting is applicable to models with three components at every location. This involves forming a submatrix from each triplet of columns in the lead field matrix. Subsequently, the singular values of the submatrices are calculated by singular value decomposition (described in §4.2). The square-root of the largest singular value is taken to be the sensitivity in each region [52]. As with the parametric method, the diagonal elements of the source covariance matrix are divided by the sensitivity to bias the solution appropriately. It is interesting to note that the largest singular value is also the Euclidean length of the first principal component.

A comparison of the relative localisation accuracy afforded by the two depth weighting methods does not appear to have been published at the time of writing. Whichever method is chosen, it may be prudent to avoid heavily-weighted solutions, especially if they counteract the properties of regularisation, which simultaneously suppresses the deepest sources.

## 4.5 Examples

As a practical illustration of MNE, some examples have been produced using simulated data. The techniques described in this section are also replicated in all subsequent uses of MNE described herein.

### 4.5.1 Introduction

As discussed over the previous sections, MNE is susceptible to spatial blurring because it is a solution to an underdetermined system of linear equations, and the resolved source signals are not independent. In addition, regularisation, which is applied to ameliorate the effects of noise, can further distort and defocus the spatial distribution of sources. In order to demonstrate these effects in action and assess their impact, some simulated data were produced by inserting individual dipole moments into a cortical surface mesh, applying the forward model, and adding noise to the result. This produces a dataset equivalent to a single time-slice from a practical MEG recording, which was then “solved” by MNE to discover the extent of the described weaknesses. The procedure is detailed below.

### 4.5.2 Method

Experimental current dipoles were simulated in three regions of an anatomical model derived from the structural MRI scan of a human subject. The subject had participated in a previous MEG study, and the relative positions of the 248 MEG sensors in the simulation were taken from the associated records in order to accurately reflect a typical recording scenario. The forward model used to simulate the measurements was a multiple overlapping sphere head model (see §2.3).

To define a source space and to facilitate visualisation of solutions, a 3-dimensional cortical surface mesh was extracted from the MRI volume data using the software *FreeSurfer*, published by the Martinos Center for Biomedical Imaging. First the skull and outer tissue are stripped from  $256^3$ -element volume data using an enhanced watershed algorithm [54]. The cerebral hemispheres are then separated by a dividing plane through the corpus callosum, and structures inferior to the pons are removed. For each hemisphere, grey and white matter are segmented and triangular meshes are formed over both the grey/white matter boundary and the interface between grey matter and cerebro-spinal fluid (the pial surface) [55]. The software then uses surface deformation to eliminate the jagged edges of the MRI voxels and produce smooth representations of the white matter boundary and pial

surface.

The surfaces produced by the method described above contain over  $10^5$  vertices. It is impossible to estimate more independent sources than sensor channels, so this many sources are not required and merely serve to increase computational requirements. Also, the meshes produced do not aim to have a consistent density of points over the cortical surface, which would bias solutions towards regions where the density is highest. To remedy this problem, the white matter surface from each hemisphere was decimated using the MNE tools also available from the Martinos Center for Biomedical Imaging. Note that the minimum norm estimates themselves were not calculated with this suite. The decimation that was applied selects a subset of the original vertices by initially inflating the surface onto a sphere. The sphere is then approximated by a geodesic grid formed by repeated subdivision of an icosahedron. The number of subdivisions determines the eventual number of sources. In this case, five subdivisions were chosen, leading to 10242 sources in each hemisphere. The resulting grid contains an approximately uniform distribution of points over the cortical surface.

Simulated dipoles of equal strength were placed in the superior temporal gyrus (STG), the superior frontal gyrus (SFG) and the cuneus (BA17). The first dipole moment was orientated in a lateral (X) direction, the second in a vertical (Z) direction and the final at  $45^\circ$  in the coronal (X-Z) plane. For simplicity, all of the simulated dipoles were placed in the left hemisphere. Separate magnetic field patterns were simulated for each individual dipole, as well as the three dipoles together. Gaussian distributed white noise was added to the sensor space data to mimic the effects of measurement error. The noise amplitude was such that the standard deviation was a fifth of the channel standard deviation, averaged over the four field patterns. A regularised, depth weighted MNE solution was then calculated using software originally developed by P.T. Kitterick at YNiC, and adapted by the present author. Tikhonov regularisation was used, with a regularisation parameter equal to 1% of the trace of the inverse operator. Depth weighting was applied using the non-parametric largest singular value method.

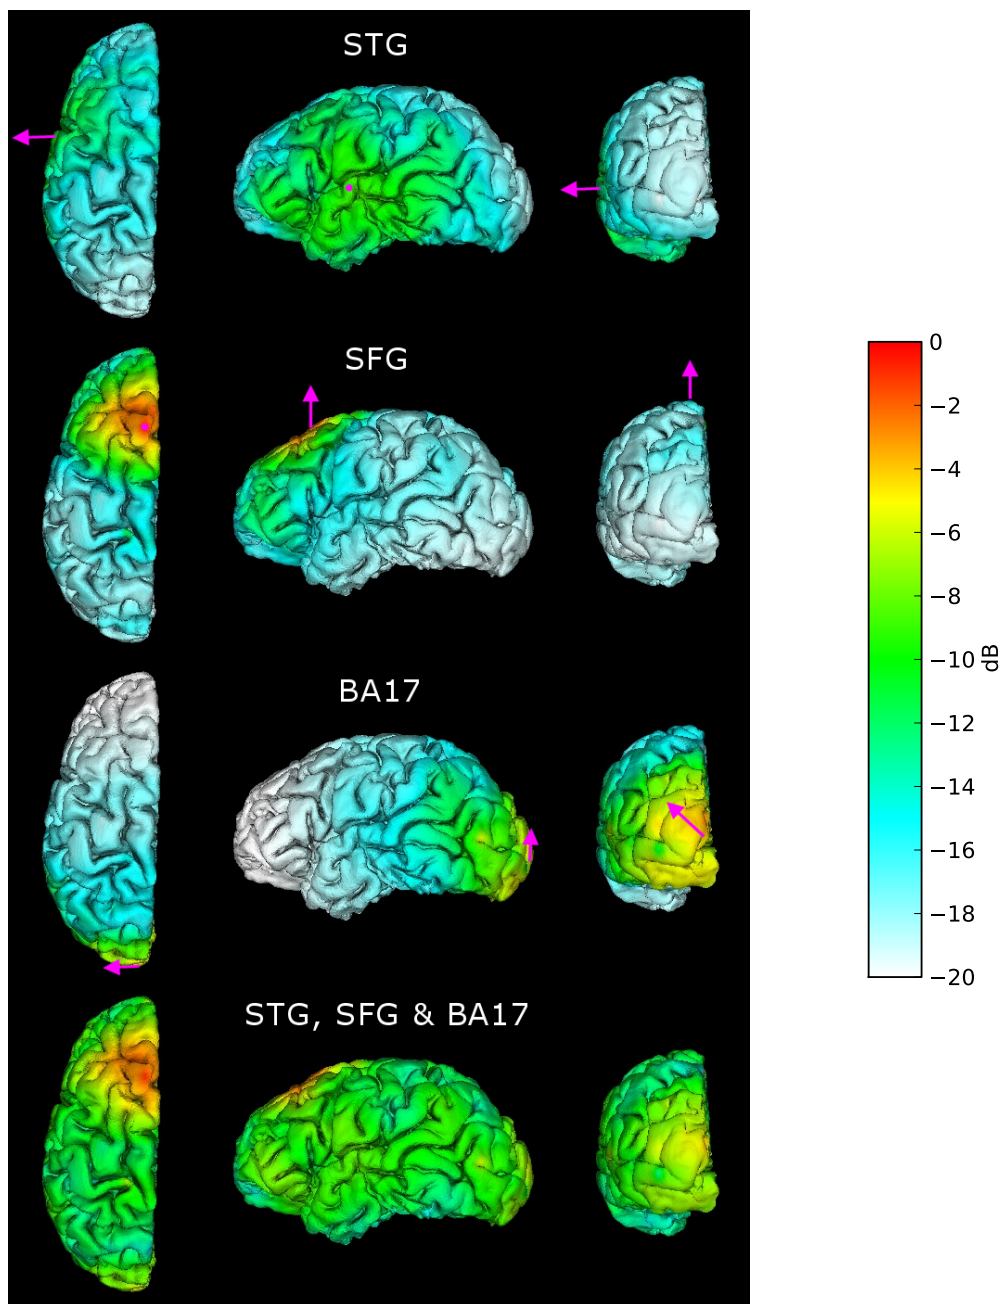
### 4.5.3 Results

The current density estimates that were produced contained three orthogonal components at each MNE grid point. For visualisation purposes, the Euclidean length was calculated at each point by summing the squares of the three components and taking the square root. The result is an unsigned magnitude, which is useful in this case for showing the localisation ability of the approach. If a signed time course estimate is required, as in later examples in this thesis, principal component analysis (PCA) can be used to determine the dominant direction of current flow in a given location during an experimental epoch. The component in this direction can then be taken as an estimated time course for the activity in that region.

Solutions on the coarse MNE grid only contain data for a subset of the vertices in the cortical surface model. Also, while the white matter boundary was used as the source space, the activity is somewhat clearer when displayed on the pial surface. To assign colour values to the dense pial surface mesh, the coarse grid was interpolated using Shepard interpolation [56]. This method was chosen because it facilitates the interpolation of data that have been sampled on an irregular grid such as the MNE grid. The *Visualization Tool Kit* (VTK) implementation of Shepard interpolation was used, within the *DataViewer3D* (DV3D) application developed by A. Gouws at YNiC. Figure 4.2 shows a selection of views on the data produced with DV3D. Although the source space for the MNE solutions included both hemispheres, only the activity in the left hemisphere has been shown due to the absence of simulated dipoles in the right hemisphere.

### 4.5.4 Conclusions

Even in the presence of a modest level of noise, it is evident that focal sources have been identified in regions of the cerebral cortex corresponding well to the locations of the experimental dipoles. However, one can observe considerable leakage of these highly-localised sources into nearby regions. In an application where this was unacceptable, another source estimation procedure might be preferred. However, there is no requirement in this project for high spatial resolution because we are more interested in the time course



From top to bottom: MNE solutions resulting from simulated dipoles in the temporal lobe (STG), frontal lobe (SFG), occipital lobe (BA17) and all three dipoles simultaneously.

Dipole moments are shown in magenta.

Figure 4.2: MNE solutions from simulated dipoles

of the activations than precisely localising them. Consequently, the localisation ability is satisfactory considering that significant leakage ( $> -10\text{dB}$ ) is not seen in regions far removed from the true source. If accurate source localisation is needed, the data could also be analysed using an alternative method such as beamforming.

With the three dipoles active simultaneously, the temporal lobe activity is not very distinct. This would be considered a poor quality result even with real MEG data. Since the head shape and position were obtained from a real experiment, this could well have been caused by the position of the subject relative to the sensors. With more realistic data, underlying sources are unlikely to reach maximal amplitude at the same time, which should aid with visual identification of them. In any case, if it was suspected that a source was present, a time course could be estimated in that region prospectively without risk to the validity of other estimates. Taking into account the benefits of MNE compared with beamforming in terms of low computational requirements and the lack of reliance on statistical assumptions, these results broadly support the use of MNE within the proposed framework.

## 5 Time-frequency Analysis

MEG signals contain significant energy over a frequency range that spans several decades. In addition, these signals commonly oscillate to some degree, whether this is an ongoing wave or just the ringing of a transient event. It is widely considered that neuronal oscillations at different frequencies embody different neurological functions [57, 16, 58], perhaps acting as independent channels of communication. To investigate these phenomena, it is advantageous to analyse the overall signal into a time-frequency representation. With this approach, underlying elements of the signal at different frequencies can be considered independently, without forgoing a sense of their temporal progression. This chapter explores these types of data representation, beginning with an argument for using them in MEG analysis. Subsequently, some common time-frequency decompositions are described. First, the Short-Time Fourier Transform (STFT) is presented, with an introduction on the Discrete Fourier Transform (DFT) and windowing. Wavelet transforms are then introduced, covering continuous and discrete types, the latter including multiresolution analysis. The chapter concludes with a description of some novel approaches to time-frequency decomposition that attempt to redress the shortcomings of more conventional methods.

### 5.1 Relevance to MEG

MEG signals are recorded as a measure of amplitude versus time, so this is often how they are presented, e.g. Figure 5.1(a). However, interpreting this unprocessed *time domain* representation is not unlike trying to comprehend an entire spectrum of radio broadcasts simultaneously. In the radio spectrum there are numerous conflicting signals — far too

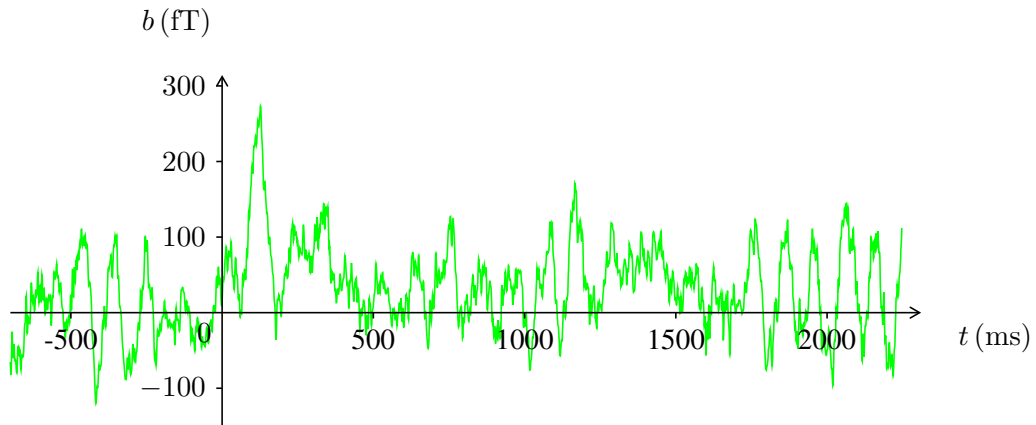
many to resolve individual transmissions. Some form of tuner or filter must be used to isolate a particular broadcast. The same could be said of neuromagnetic sources within the brain. It is commonplace to use an array of sensors in MEG, which provides the means to separate sources that are physically removed from each other. But what if two concurrent signals produced by different mechanisms originate from the same brain region? In this case, interactions between disparate regions may only become evident if sources are further distinguished by approximate frequency of oscillation.

It is well established that the spectral density of MEG signals falls off with increasing frequency [16, 3]. The dominant low-frequency components can easily overshadow high-frequency events in a wide-band representation. Many experimentalists apply filters with broad passbands to isolate frequency ranges considered relevant to their investigation. The frequency of the passband edges is usually directed by empirical knowledge of neurological frequency bands (i.e. alpha, beta, gamma, delta and theta bands [16]), combined with a postulated mechanism for the neural functions being studied. Such bands are not consistently defined in the literature. Also, the type and order of the filter are not directed by the experimental paradigm, and may ultimately be chosen arbitrarily. The absence of reasoning from this process jeopardises the integrity of the results due to selective rejection of potentially important features and the possible introduction of time-domain artefacts.

A more rigorous approach to studying frequency-specific effects might be to calculate the Fourier transform of the recorded signal. However, this *frequency domain* representation, as in Figure 5.1(b), simply reveals how energy is concentrated spectrally, obscuring features such as modulation, which is subtly manifested as sidebands around a spectral peak. In addition, bursts of neuronal oscillation may only last a few cycles. Even when these bursts reach relatively large amplitudes, their energy is not significant over an entire measurement epoch if the wavelength is much shorter than the epoch. Thus these informative events are poorly represented when a complete epoch is transformed into the Fourier domain.


From the arguments above, it could be construed that neither frequency- nor time-domain representations are ideal for visualising or analysing MEG. By creating a hybrid of both

The plots below show different representations of the same MEG data. The auditory stimuli used in this experiment consisted of various spoken words and non-words (meaningless word-like utterances), of which 200 were presented to a normal subject. Of these, 22 trials were rejected due to artefacts (interference from sources other than the brain).



This is the trial average, with the axis scaled so the onset of each stimulus occurs at time  $t = 0$ . The stimulus durations are in the range 526–557 ms due to variations in word length. To reduce noise, which is incoherent across trials, the point-wise arithmetic mean has been plotted. Although common practice, this also rejects neuromagnetic signals that are not phase-locked to the stimulus. Alpha oscillations are evident during the pre-stimulus period and towards the end of the measurement epoch. A classic evoked auditory response (M100) is visible as a spike at 100ms.

(a) Time domain



These power spectral density (PSD) estimates are derived from the time domain data by applying the DFT. The lower curve is a direct Fourier transform of the averaged data, appropriately scaled to provide a rough PSD estimate. The upper curve is the average of the PSD estimated from each individual trial, yielding a more stable estimate that also retains incoherent signal energy. The alpha oscillations are still discernible as a peak at around 10Hz, but the M100 response is not clearly manifest in the power spectrum. The other peaks at 50Hz, 100Hz and 150Hz are interference from electrical supplies in the vicinity of the instrumentation. The horizontal section at the upper end of the frequency scale is the noise floor.

(b) Frequency domain

Figure 5.1: Time and frequency domain representations of a MEG signal

approaches, a balance can be struck between a temporal view of the signal and the ability to easily distinguish different frequency components, giving rise to *time-frequency analysis*. This family of techniques regards amplitude or energy as a function of both time and frequency. Such representations accommodate the study of oscillatory activity, without being prejudiced against unforeseen outcomes.

There are a great many ways to decompose a signal in time and frequency; the chosen method must suit the nature of the signals being interrogated. The best known time-frequency analysis method is the *short-time Fourier transform* (STFT), which adds time-dependency to a conventional frequency domain view by dividing the original data into multiple time windows. However, the application of a window to a signal can introduce spurious features in the frequency domain. The choice of window width further imposes a compromise between time resolution (which is fixed across all frequencies) and frequency resolution.

An increasingly widespread family of time-frequency transforms known as *wavelet transforms* allow time resolution to vary in proportion to frequency, creating a logarithmic frequency scale [59]. These approaches are well suited to the log-spaced nature of neuromagnetic signals. They include the *continuous wavelet transform* (CWT), which offers high resolution at the expense of considerable redundancy and computational demands. Additional redundancy is highly undesirable in MEG analysis, which concerns data sets that are already vast from the outset. The *discrete wavelet transform* (DWT) can represent information without added redundancy, but is associated with a fixed octave-wide frequency resolution. It also exhibits a phenomenon known as *shift variance*, whereby the transformed output depends on the time offset of the input.

The *fast wavelet transform* (FWT) is a recursive, computationally efficient implementation of the DWT. By building on the concept of the FWT, a novel filter bank has been created that abandons the fixed frequency resolution of the existing approach without introducing excessive redundancy. This parametric transformation provides a choice as to the extent of additional redundancy, the amount of spectral smoothing and the length of ringing in the

time domain.

In the remainder of the chapter, detailed assessments of the established time-frequency transforms are made with regard to MEG analysis. The conclusions from these assessments support the subsequent development of the novel, MEG-orientated filter bank method, which is presented last.

## 5.2 Short-time Fourier Transform

The short time Fourier transform (STFT) is based on the discrete Fourier transform (DFT), so it will be instructive to consider this first. The STFT will then be defined and the application of it discussed, with an emphasis on MEG.

### 5.2.1 Discrete Fourier Transform

The DFT is a well-known analysis that transforms discrete-time input into the discrete frequency domain. Any periodic waveform,  $x(n)$ , with a period of  $N$  samples, can be generated from the sum of a discrete Fourier series. This is a harmonic series of complex waves, that can be multiplied by complex coefficients,  $X(k)$ , and summed to reproduce the waveform.

$$x(n) = \sum_{k=-\infty}^{\infty} X(k)e^{j(2\pi/N)kn} \quad (5.1)$$

The coefficients describe the amplitude and phase of each frequency component. The only frequencies present in the series are those that satisfy  $\omega_k = 2\pi k/N$ ,  $k \in \mathbb{Z}$ . These discrete frequencies are such that complete cycles necessarily occur within the  $N$ -sample sequence. Although the Fourier series is infinite, just  $N$  consecutive values of  $X(k)$  are sufficient to construct any conceivable waveform because  $e^{j(2\pi/N)kn}$  is periodic in  $k$ . Therefore having coefficients at other values of  $k$  contributes no additional information. The sum can be bounded as follows:

$$x(n) = \sum_{k=0}^{N-1} X(k)e^{j(2\pi/N)kn} \quad (5.2)$$

Notably, the number of samples in the discrete-time representation matches the number of discrete frequency components required. Normalising by a factor of  $1/N$  keeps the signal amplitude constant with respect to sequence length:

$$x(n) = \frac{1}{N} \sum_{k=0}^{N-1} X(k) e^{j(2\pi/N)kn} \quad (5.3)$$

Now an expression for  $X(k)$  can be obtained in terms of  $x(n)$  by multiplying both sides by  $e^{-j(2\pi/N)nm}$  and summing for all  $n$ :

$$\sum_{n=0}^{N-1} x(n) e^{-j(2\pi/N)nm} = \frac{1}{N} \sum_{n=0}^{N-1} \sum_{k=0}^{N-1} X(k) e^{j(2\pi/N)n(k-m)} \quad (5.4)$$

Due to the symmetry of the complex wave, the sum over  $n$  on the right-hand side evaluates to zero except when  $k = m$ . That is,

$$\sum_{n=0}^{N-1} e^{-j(2\pi/N)n(k-m)} = \begin{cases} N & \text{if } k = m \\ 0 & \text{otherwise} \end{cases}$$

Consequently, the appearance of the above expression in the right hand side of Equation 5.4 can be replaced by  $Nu_0(k - m)$ , where  $u_0$  is the digital impulse function. This gives

$$\sum_{n=0}^{N-1} x(n) e^{-j(2\pi/N)nm} = \sum_{k=0}^{N-1} X(k) u_0(k - m) \quad (5.5)$$

Finally, putting  $m = k$  results in the general form of the DFT, which is

$$X(k) = \sum_{n=0}^{N-1} x(n) e^{-j(2\pi/N)nk} \quad (5.6)$$

Although this result is derived for the case when  $x(n)$  is periodic, the DFT also uniquely represents finite sequences that are defined as zero outside the range  $0 \leq n \leq N - 1$ .

In practice, the DFT is almost always calculated using a fast Fourier transform (FFT) algorithm. This significantly reduces the computation time as the sequence length increases by exploiting the periodicity of  $e^{j\theta}$  to reduce the number of complex multiplications involved. Without modification, the classic Cooley-Tukey FFT algorithm can only be performed on sequence lengths of  $N = 2^p$ ,  $p \in \mathbb{N}$  [60]. Variants which are applicable to


arbitrary-length sequences exist, but the computational efficiency is degraded for sequences with lengths that are not powers of two.

Frequently, the input to a DFT is obtained from a continuous-time waveform that has been uniformly sampled in time, as in the case of an MEG signal. Practical signals such as these are neither periodic nor finite in length. The time series to be analysed is produced by truncating a signal that has infinite duration in reality (the magnetic fields do not actually cease when the measurements end). Mathematically speaking, a rectangular window has been applied to the signal. This has important implications for the output of the transform.


To demonstrate the effects of windowing, consider a time series of length  $N$  samples that consists of a sampled sinusoid having a frequency  $\omega_0 \neq 2\pi k/N$ . An incomplete cycle of the sinusoid occurs within the time window bounding the time series (Figure 5.2). This leads to a discontinuity at the edges of the window, which is manifested in the transformed signal as spectral leakage (Figure 5.3). Some of the energy in the signal appears to be in frequencies far removed from  $\omega_0$ . This does not represent any genuine aspect of the original signal that was sampled. Instead it is due to the choice of window length.

In fact, the effect of windowing any signal in this way gives rise to sidebands in the amplitude spectrum, regardless of the relationship between the window length and the signal wavelength (Figure 5.4). However, when a stationary oscillation with frequency  $\omega_0 = 2\pi k/N$  is transformed, the samples in the discrete frequency domain coincide with the nulls between each sideband, so the sidebands are not apparent. If the oscillation were truly stationary, there would be no other frequencies in the original signal, so the representation is complete for the original signal. Otherwise, sidebands are present but not visible, or indeed discernible by any method, from the contents of the window without additional assumptions (Figure 5.4 assumes the signal is zero outside the window).


The introduction of sidebands can be mitigated to some extent by the use of a suitable windowing function, such as a Hamming window [61]. Such windows taper to zero towards the edges, eliminating the discontinuities at the window edges. The reduction in amplitude of the sidebands comes at the expense of a reduction in the effective frequency resolution,



(a) Original sinusoidal signal




(b) Rectangular window applied at arbitrary time offset



(c) Effective signal under analysis

Figure 5.2: Effect of windowing on DFT




Power spectra from 16384-point DFTs of sinusoids having  $\omega_0 = 2\pi \times 82/16384$  (integrally related to window length) and  $\omega_0 = 2\pi \times 82.1/16384$  (not integrally related to window length, causing leakage)

Figure 5.3: Spectral leakage in DFT due to rectangular windowing of a periodic signal

or widening of the spectral peaks. In time series analysis, another way of minimising the artefacts caused by windowing is to choose a section of the signal that naturally has minimal power at the edges of the window. In MEG this tactic may be suitable for evoked responses averaged over several trials, but is less useful for induced or spontaneous signals which do not have convenient quiescent periods at the beginning and end. In any case, the STFT involves dividing the sequence into several shorter windows, so a non-rectangular window function is vital if significant sideband energy is not acceptable.

### 5.2.2 Definition of STFT

Now that the DFT has been considered, the short-time Fourier transform can be introduced. The STFT transforms time-dependent data into a domain where they are a function of both time and frequency. The frequency components at each time point are obtained by calculating the DFT upon a sliding window. The time argument to the function corresponds



Part of the continuous amplitude spectrum of a sinusoid with angular frequency  $\omega_0$  after a rectangular window of duration  $10 \times 2\pi/\omega_0$  has been applied. If the DFT was calculated on the same signal, just a single peak would be visible at  $\omega/\omega_0 = 1$ , due to the discrete frequency samples coinciding with the nulls when  $\omega \neq \omega_0$

Figure 5.4: Sidebands in Fourier spectrum introduced by rectangular window

to the offset of the sliding window. For a finite sequence,  $x(n)$ , of length  $N$ , the STFT may be expressed algebraically as

$$X(k, m) = \sum_{n=mR}^{mR+M-1} x(n)h(n - mR)e^{-j(2\pi/M)nk} \quad (5.7)$$

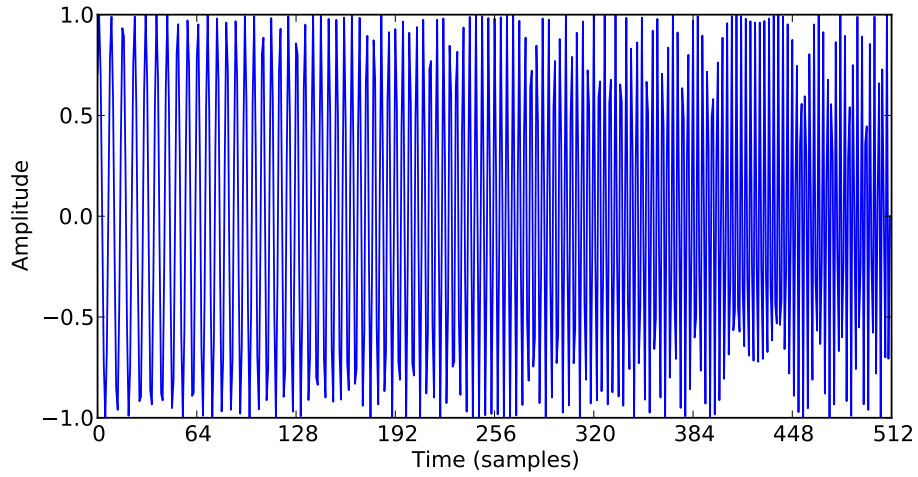
where  $k$  and  $m$  are the frequency and time indices respectively. The windowing function  $h(n)$  is defined as zero outside the range  $0 \leq n \leq M - 1$ , with  $M$  typically equal to a power of two to allow calculation using an FFT algorithm. The interval between successive time windows is determined by the step size  $R$ . This has been distinguished from  $M$  to permit the overlapping of adjacent windows by setting  $R < M$ . Overlapped windows may be desirable in the case where a non-rectangular window is being applied, as the windowing function introduces bias towards aspects of the data that are central to the time window. A particular overlap largely or fully eliminates bias when using typical window functions such as Gaussian or raised cosine (Hann). When a three-dimensional plot is produced of the squared magnitude of the transformed signal (or energy density) against time and frequency, the graph is known as a spectrogram. Spectrograms often use colour or lightness to represent power, as in the example shown in Figure 5.5.

### 5.2.3 Discussion


Depending on the requirements of the analysis, different parameters should be chosen for the STFT. The longer the window length, the greater the frequency resolution, but the poorer the time resolution. This restriction is essentially a consequence of the uncertainty principle, which can be interpreted to mean that the frequency and the duration of an oscillation cannot both be measured to arbitrary precision. More specifically, the measurements must satisfy the following inequality [62]:

$$\Delta t \Delta f \geq \frac{1}{4\pi} \quad (5.8)$$

where  $\Delta t$  and  $\Delta f$  represent the standard deviations of the time and frequency windows defining the elements in the time-frequency plane. The uncertainty principle *does not* imply that the time and frequency resolutions must be constant everywhere in the time-frequency



(a) Time-domain plot of 512-sample chirp



(b) Energy density of chirp estimated by STFT using rectangular 32-sample window

Figure 5.5: Spectrogram of linear chirp

plane, but this *is* a property of the STFT. Arguably, the implications of this property make the STFT unsuited to the analysis of MEG signals, as the next paragraph explains.

Empirical investigations have established natural frequency bands for neurological signals [16]. The divisions between these bands are uniformly spaced on a logarithmic scale. This contrasts with the frequency scale of the DFT (and consequently the STFT), which is linear. A pertinent aspect of this issue is the increase in the lowest resolved frequency when using shorter windows. Making the window long enough to accommodate low frequencies blurs the temporal detail at high frequencies. Energy in the original signal that occurs at wavelengths significantly longer than the window length will only be represented by changes in the DC component of successive windows. This component may not even be reproduced, for example, in a graphical representation where a logarithmic frequency scale has been used.

Undesirable artefacts are introduced by the use of windowing. As with the DFT in general, various windowing functions may be applied that provide different balances between frequency resolution and sideband rejection. This is particularly important to the STFT because window edges will inevitably occur where local signal power is high and the resulting sideband energy equally so.

The overriding limitation of the STFT is the fixed time resolution across the frequency scale. Higher-frequency oscillations are capable of carrying more information, so it would be desirable to increase the time resolution with respect to frequency. The increasingly popular *wavelet transform* is a time-frequency transform that does facilitate frequency-dependent time resolution. It is described in the next section. A more comprehensive mathematical treatment of the STFT can be found in [63].

## 5.3 Wavelet Analysis

A wavelet transform provides a means of resolving a signal into some combination of basis functions that are known as *wavelets*. It will become evident in this section that a wavelet approach to time-frequency analysis overcomes many of the problems associated

with the STFT. Nonetheless, certain limitations inevitably remain. Wavelet methods have occasionally been applied to MEG data [64, 50, 65]. The following information on wavelets was obtained from [59], except where indicated.

### 5.3.1 Continuous Wavelet Transform

A wavelet is a function that typically describes a short-lasting oscillatory fluctuation. For the analysis, a particular *mother wavelet*,  $\psi(t)$ , is chosen, which is then dilated and translated by various amounts and compared with the signal being interrogated. This wavelet function must have finite energy and zero mean:

$$\text{zero mean} \quad \int_{-\infty}^{\infty} \psi(t) dt = 0 \quad (5.9)$$

$$\text{finite energy} \quad \int_{-\infty}^{\infty} |\psi(t)|^2 dt < \infty \quad (5.10)$$

The dilation and translation of the wavelet are governed by a dilation parameter  $a$ , and a translation parameter  $b$ , so the modified wavelet is

$$\psi\left(\frac{t-b}{a}\right)$$

Usually, the wavelet function is defined to include a normalisation factor of  $1/\sqrt{a}$  so as to give the wavelet equal energy at all scales:


$$\psi_{a,b}(t) = \frac{1}{\sqrt{a}} \psi\left(\frac{t-b}{a}\right) \quad (5.11)$$

The *continuous wavelet transform* of a signal,  $x(t)$ , is then given by the following convolution integral

$$T(a,b) = \int_{-\infty}^{\infty} x(t) \psi_{a,b}^*(t) dt \quad (5.12)$$

where  $\psi_{a,b}^*(t)$  is the complex conjugate of the wavelet function. Many different mother wavelets having specific properties have been applied. These include both real-valued functions such as the Mexican hat wavelet (Figure 5.6(a)), and complex-valued functions like the Morlet wavelet (Figure 5.6(b)). If a complex wavelet is used, then a real-valued

time domain signal will be transformed into real and imaginary wavelet components, which can be represented as a magnitude and phase. This provides a good means of comparing time-varying signals that may be correlated in amplitude but phase-shifted with respect to each other.



(a) The Mexican hat wavelet

(b) The Morlet wavelet

The Mexican hat wavelet is given by  $\psi(t) = (1 - t^2)e^{-t^2/2}$ , which is the second derivative of a Gaussian function. The Morlet wavelet, given approximately by  $\psi(t) = \pi^{-1/4} e^{j\omega_0 t} e^{-t^2/2}$ , for  $\omega_0 \gg 0$ , is a complex wave shaped by a Gaussian envelope with standard deviation  $\sigma = a$ . The parameter  $\omega_0$  determines the frequency of the complex wave and thus the centre frequency of the wavelet, which is independent of the envelope size. In the diagram above, the dotted line represents the imaginary part

Figure 5.6: Two common mother wavelets

The wavelet transform may be thought of as a bank of bandpass filters [66], with the wavelet function representing the impulse response of the filter. The wavelet scale is then inversely proportional to both the centre frequency and the bandwidth of the filter. If the magnitude squared is plotted against displacement and scale, the graph is called a *scalogram* and displays the relative power in the signal at different time offsets and wavelet scales. Alternatively, a change of variables allows the transform to be expressed in terms of frequency. If a Morlet wavelet is also used, a graphical plot of the magnitude squared is akin to the STFT spectrogram using a Gaussian window, and likewise displays energy density.

The fundamental difference between the spectrogram calculated using the STFT and the Morlet wavelet is that the latter has a time resolution proportional to frequency. In contrast, the STFT imposes the same window length on all frequencies. With the Morlet wavelet, the effective window length varies according to scale, and is proportional to wavelength. The frequency resolution can be made constant when viewed on a logarithmic scale. This sort of decomposition is most applicable to signals with energy distributed over a bandwidth spanning several orders of magnitude.

Because Equation 5.12 represents a convolution, the convolution theorem can be applied, which allows the wavelet transform to be calculated in the frequency domain. The calculation of each wavelet scale becomes a pointwise multiplication between the Fourier transform of the wavelet and that of the signal. The inverse Fourier transform then yields the desired result. This is advantageous in that the Fourier transform of the wavelet is usually known analytically, while the FFT algorithm can be used to perform the transformation of the signal, significantly decreasing calculation time.

In principal, the CWT is defined as a continuous function of the scale and translation variables, but in practice, values will be calculated at discrete points for graphical representation or further processing. A notable problem with this approach is that the wavelet representation is necessarily sampled at a very high resolution to avoid loss of information, leading to a smooth appearance but large redundancy in the output. An alternative approach is to use the *discrete wavelet transform*, described below.

### 5.3.2 Discrete Wavelet Transform

The discrete wavelet transform (DWT) is a type of wavelet analysis in which the wavelet is restricted to discrete translations and scales as follows:

$$T_{m,n} = \int_{-\infty}^{\infty} x(t)\psi_{m,n}(t) dt \quad (5.13)$$

$$\text{where } \psi_{m,n} = \frac{1}{\sqrt{a_0^m}} \psi\left(\frac{t - nb_0a_0^m}{a_0^m}\right) \quad (5.14)$$

This is related to the CWT (Equation 5.12) via the following substitutions:

$$a = a_0^m$$

$$b = nb_0a^m$$

Thus the DWT is defined in the same way as the CWT, but instead of being a continuous function of translation and scale, the transform now produces a series of coefficients. The translations and scales are defined in terms of the discrete integer variables  $m$  and  $n$ . The parameter  $a_0$  determines the common ratio (or dilation step) between successive scales, while  $b_0$  sets the size of the discrete translation steps. The latter are scaled in the same manner as the wavelet function itself, such that longer scales have proportionally larger intervals between translations. Commonly,  $a_0 = 2$  and  $b_0 = 1$ , producing a wavelet discretisation known as a *dyadic grid*.

The above definition of discrete wavelets permits the transform to discard information, so that a 1:1 mapping does not exist between the original signal and its transform. Additionally, the transform may also have redundancy. However, an *orthonormal basis* provides a non-redundant representation from which the signal may be exactly reconstructed. The requirements for an orthonormal wavelet basis are that the product between wavelets at different translations and scales is zero, and that the wavelets have unit energy:

$$\int_{-\infty}^{\infty} \psi_{m,n}(t)\psi_{p,q}(t) dt = \begin{cases} 1 & \text{if } m = p \text{ and } n = q \\ 0 & \text{otherwise} \end{cases} \quad (5.15)$$

Under these conditions, the inverse discrete wavelet transform perfectly reconstructs the original signal:

$$x(t) = \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} T_{m,n}\psi_{m,n}(t) \quad (5.16)$$

It is particularly simple to construct an orthonormal wavelet basis by choosing a dyadic grid, which has resulted in the DWT becoming synonymous with dyadic scaling in the literature.

### 5.3.3 Multiresolution analysis

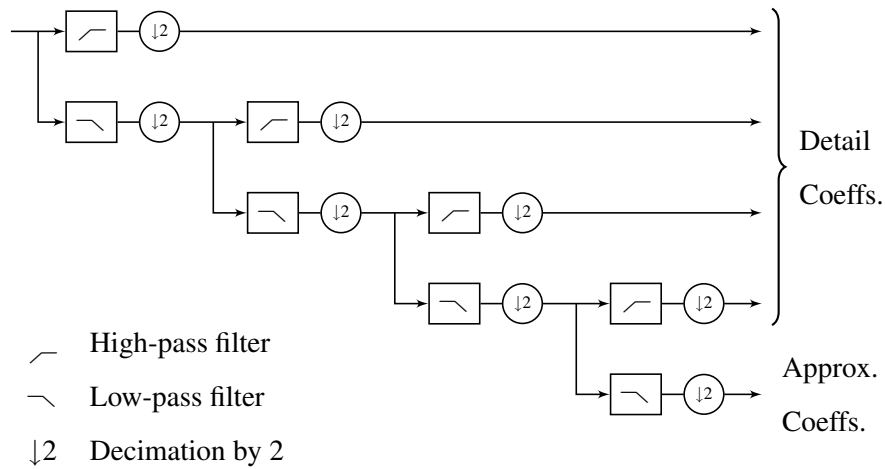
For a continuous-time signal, the discrete wavelet coefficients form an infinite series. Because wavelets represent bandpass filters, they each have a finite low frequency cutoff. An infinite number of filters of increasing scale is therefore required to represent an arbitrary signal. This is an inconvenient situation, but it can be remedied by introducing a *scaling function*,  $\phi(t)$  [67]. The scaling function is related to the specific mother wavelet being used, and acts as a low-pass filter. It provides an approximation of the signal down to a given scale,  $m_0$ . The remaining discrete wavelet coefficients provide the detail that is absent from the approximation. This gives rise to the *multiresolution representation*, wherein the signal is represented as a weighted sum of the scaling function at all translations and the wavelet function at all translations with scales up to  $m_0$ :

$$x(t) = \sum_{n=-\infty}^{\infty} S_{m_0,n} \phi_{m_0,n}(t) + \sum_{m=-\infty}^{m_0} \sum_{n=-\infty}^{\infty} T_{m,n} \phi_{m,n}(t) \quad (5.17)$$

In this representation,  $S_{m_0,n}$  are known as the *approximation coefficients*, while  $T_{m,n}$  are the *detail coefficients*. If the signal to be transformed is a discrete-time signal of finite length, then a finite number of coefficients are needed. In effect, a discrete-time signal is already a low-pass approximation of a continuous-time signal, which may be resolved into a further approximation with accompanying detail coefficients at a chosen number of scales. A finite-length sequence demands only a finite number of possible wavelet translations, leading to a finite set of coefficients that completely represent the signal.

Typically, the DWT is implemented as a succession of quadrature mirror filter pairs [67], with each high-pass section providing a new set of detail coefficients and the low-pass section providing a new approximation, which is fed into the next pair of filters. The signal is decimated at each stage, providing only the necessary sampling interval with increasing scale. This scheme is often known as the fast wavelet transform (FWT), by analogy with the FFT. It is depicted in Figure 5.7 for the case when there are four wavelet scales.

An orthonormal DWT does not exhibit shift invariance; i.e. the DWT of a time-shifted version of a signal is not necessarily the same as the similarly shifted transform of the



This data flow diagram shows the sequence of processes required to perform multiresolution analysis for four wavelet scales with dyadic scaling.

Figure 5.7: Multiresolution Analysis

original signal. This is sometimes an undesirable characteristic, for instance, when cross-correlation measures are to be calculated (as they will be in the latter part of the signal processing framework). A simple way to eliminate the shift variance is to remove the decimation step from the filtering process, producing intermediate points between the discrete samples of the dyadic grid. The filter coefficients must also be stretched to take into account the new sample rate at each stage. This is achieved by inserting zero-valued filter coefficients between the existing ones. The number of wavelet coefficients generated at each scale is equal to the number of time points in the original signal, so this is a highly redundant representation. Discrete wavelet transforms that are approximately shift invariant and modestly redundant have also been devised.

### 5.3.4 Discussion

Wavelet methods are attractive because they can accommodate logarithmic frequency scales, and do not suffer from windowing artefacts in the same way that the STFT does (in the sense that any spectral spreading is an explicit characteristic of the wavelet). The

CWT gives total freedom with respect to translation and scaling, which avoids problems with shift variance but generates excessively redundant data sets. Furthermore, infinitely many wavelet scales are required to fully represent a signal, unless the bandwidth of the signal is known to be constrained. The redundancy in particular is not acceptable for the intended application for at least two reasons. Firstly, it will reduce calculation efficiency. More importantly, it will have the effect of greatly expanding the already-voluminous MEG data.

By comparison with the CWT, discrete wavelet methods provide the additional benefit of compact representation. This comes at the expense of introducing shift-variance in the transform, but this can be overcome if necessary. The paramount objection to using the discrete wavelets is the implicit requirement for dyadic grid scalings, which yields a very poor frequency resolution of one octave. In fact, it should be possible to use other scalings, a concept that is explored in the following section.

## **5.4 A Novel Filter-bank Analysis**

As previously discussed in this chapter, common time-frequency analyses have serious limitations that hinder their effectiveness in many applications like magnetoencephalography. To reiterate,

- The STFT relies on windowing of the data, which can introduce artefacts towards the window edges. Also, the time resolution is fixed across all frequencies and the frequency resolution linear.
- The CWT is slow to calculate and produces vast quantities of highly redundant data that is smooth in appearance. Energy at frequencies above and below a finite pass-band is ignored.
- The DWT remedies the problems of the CWT, but introduces shift-variance and conventionally enforces coarse frequency resolution.

Due to the perceived failings of conventional time-frequency analyses, a novel method was sought. The discrete wavelet transformation was taken as the starting point because of its efficient representation and fast calculation using the FWT.

Most of the literature concerning discrete wavelet decomposition seems to assume that a dyadic grid scaling is always required [66]. Such scaling implies a dilation step of two, so the frequency resolution obtained by the process is a whole octave. In applications such as image compression, the factor of two is useful because it simplifies the construction of an orthogonal basis, which allows for minimal redundancy in the transformed data. In contrast, the coarse resolution is often objectionable in the analysis of time series data. For example, a frequency-modulated oscillation will mostly be confined to a single wavelet scale if the frequency deviation is less than the carrier frequency. As a result, the modulating signal is not overtly manifested in the output because the energy in the carrier cannot be seen to move between scales over time. It is not trivial to define a discrete wavelet basis that is efficient yet has a dilation step in the interval  $(1, 2)$ .

The challenge is to define a wavelet basis that is as closely orthonormal as possible, in order to avoid redundancy, but without being dyadic in nature. The transform should have a more-or-less uniform response across time and frequency. Without this constraint, a subsequent correlation analysis could suggest spurious relationships between unrelated signals that have both gained certain features from the transformation. Two approaches to meeting these criteria were considered. Only the second has been shown to be of virtue, but they are both described here for completeness.

#### 5.4.1 Time-domain View

The first approach to non-dyadic scaling that was considered takes inspiration from the Morlet wavelet, which is usually associated with continuous wavelet transforms. For information, the Morlet wavelet is defined as follows [59]:

$$\psi = \pi^{-1/4} (e^{j\omega_0 t} - e^{-\omega_0^2/2}) e^{-t^2/2} \quad (5.18)$$

In essence, this is a complex wave shaped by a Gaussian envelope. Simply applying such an envelope to a complex wave results in a function that has a DC component, which is not admissible as a wavelet function. To compensate, the  $-e^{-\omega_0^2/2}$  term in the brackets is included to adjust for zero DC. A simplified form omits this term because it is insignificant for large values of  $w_0$ . The parameter  $w_0$  determines the number of cycles of the complex wave that occur within the standard deviation of the Gaussian envelope. It governs the centre frequency of the wavelet function's passband. Indirectly,  $\omega_0$  also affects the bandwidth of the wavelet, with higher numbers resulting in greater frequency selectivity. It is this control that makes the Morlet wavelet an interesting model for a novel wavelet function.


The idea inspired by the above was to create a discrete wavelet function from a sinusoid or complex wave by applying a raised cosine amplitude envelope. By setting the discrete translation steps to half a wavelength at the frequency of the raised cosine, the combined envelope of wavelets at different translations sums flat, meeting the criterion of no bias in the time axis. The raised cosine envelope can be expressed as follows:

$$u(t) = \begin{cases} \frac{1}{2} (1 + \cos \frac{2\pi t}{T}) & -\frac{T}{2} \leq t \leq \frac{T}{2} \\ 0 & \text{otherwise} \end{cases} \quad (5.19)$$

This is often called a Hann window, defined here with the peak at zero and width  $T$  (see Figure 5.8). Developing the concept further, the frequency of the enveloped wave must be an even multiple of the frequency of the cosine used to create the window. This causes the different translations to behave as different windows onto the same wave (see Figure 5.9). The ratio between the envelope width and the period of the oscillation determines the bandwidth of the wavelet in the frequency domain, as with the Morlet wavelet. The wavelet function is thus defined as


$$\psi(t) = \begin{cases} \frac{1}{2} (1 + \cos \frac{2\pi t}{T}) \cos \frac{2\pi n t}{T} & -\frac{T}{2} \leq t \leq \frac{T}{2} \\ 0 & \text{otherwise} \end{cases}, \quad n \text{ even} \quad (5.20)$$

Unlike the simplified Morlet wavelet, this function inherently contains no DC. Since  $\psi$  in this case is zero outside the interval  $[-\frac{T}{2}, \frac{T}{2}]$ , the DC component or mean of the function



The Hann window is formed by taking one cycle of a cosine and offsetting it so that the minima touch the time axis. In this case, the function has been scaled so that the central maximum reaches an amplitude of one.

Figure 5.8: Hann window



Five consecutive translations of a particular wavelet scale are shown. The wavelets are in-phase where they overlap, so they are not orthogonal. However, their product is small so redundancy is minimal.

Figure 5.9: Sequence of raised cosine wavelets

is given by a definite integral:

$$\bar{\psi} = \frac{1}{T} \int_{-T/2}^{T/2} \frac{1}{2} \left( 1 + \cos \frac{2\pi t}{T} \right) \cos \frac{2\pi n t}{T} dt$$

We can show that this evaluates to zero for all  $n \in \mathbb{N}$ ,  $n \neq 1$ . For simplicity, the constant multiplying factors  $\frac{1}{T}$  and  $\frac{1}{2}$  can be stripped. So can the factor that determines the window length, by a change of variables,  $\theta = 2\pi t/T$ . We just need to prove that

$$\int_{-\pi}^{\pi} (1 + \cos \theta) \cos n\theta d\theta = 0 \quad (5.21)$$

First, the integral will be split into separate terms for which antiderivatives will be found individually:

$$\int (1 + \cos \theta) \cos n\theta d\theta = \int \cos n\theta d\theta + \int \cos \theta \cos n\theta d\theta \quad (5.22)$$

The first term on the right hand side is trivial. By inspection,

$$\int \cos n\theta d\theta = \frac{\sin n\theta}{n} \quad (5.23)$$

The second term is more complicated, but can be integrated by parts:

$$\begin{aligned} \text{Let } X &= \int \cos n\theta \cos \theta d\theta \\ &= \int u_x \frac{dv_x}{d\theta} d\theta \end{aligned}$$

$$\text{where } u_x = \cos n\theta, \quad \frac{dv_x}{d\theta} = \cos \theta$$

$$\frac{du_x}{d\theta} = -n \sin n\theta, \quad v_x = \sin \theta$$

Integrating by parts,

$$X = \cos n\theta \sin \theta + n \int \sin n\theta \sin \theta d\theta \quad (5.24)$$

Working on the resultant integral in the same way,

$$\begin{aligned} \text{Let } Y &= \int \sin n\theta \sin \theta d\theta \\ &= \int u_y \frac{dv_y}{d\theta} d\theta \end{aligned}$$

$$\text{where } u_Y = \sin n\theta, \quad \frac{dv_Y}{d\theta} = \sin \theta$$

$$\frac{du_Y}{d\theta} = n \cos n\theta, \quad v_Y = -\cos \theta$$

Hence,

$$Y = -\sin n\theta \cos \theta + n \int \cos n\theta \cos \theta d\theta$$

$$= -\sin n\theta \cos \theta + nX \quad (5.25)$$

Substituting (5.25) into (5.24) produces an equation to be solved for  $X$ :

$$X = \cos n\theta \sin \theta - n \sin n\theta \cos \theta + n^2 X \quad (5.26)$$

$$X = \frac{n \sin n\theta \cos \theta - \cos n\theta \sin \theta}{n^2 - 1} \quad (5.27)$$


Note that this function is not defined for  $n = 1$ . From (5.23) and (5.27), the definite integral is given by

$$\int_{-\pi}^{\pi} (1 + \cos \theta) \cos n\theta d\theta = \left[ \frac{\sin n\theta}{n} + \frac{n \sin n\theta \cos \theta - \cos n\theta \sin \theta}{n^2 - 1} \right]_{-\pi}^{\pi} \quad (5.28)$$

It is evident that this will evaluate to zero for integer  $n$  (other than one) and limits of  $\pm\pi$  because the sine of any multiple of  $\pi$  is equal to zero, and this appears as a factor in all three terms at both limits. Extending the concept to a complex wave, the imaginary part also has no DC. This can be determined by inspection since it is a product between the windowing function, which has even symmetry, and a sine stretched about the origin, which has odd symmetry. The resulting function is therefore odd as well. Consequently, integrating symmetrically about  $\theta = 0$  will always evaluate to zero (since the function is continuous). This is graphically represented in Figure 5.10.

The Fourier transform of the complex version can be found analytically by redefining the wavelet function in terms of some common functions for which standard results exist, then applying the Convolution Theorem. Here is the alternative definition, where  $u(t)$  is the Heaviside step function:

$$\psi(t) = \left( u\left(t + \frac{T}{2}\right) - u\left(t - \frac{T}{2}\right) \right) \times \frac{1}{2} \left( 1 + \cos \frac{2\pi t}{T} \right) e^{j2\pi nt/T} \quad (5.29)$$



The function  $\text{Im}\{\psi(-t)\} = -\text{Im}\{\psi(t)\}$ . Thus, the negative region (shaded blue) cancels with the positive region (red) when integrating about the origin.

Figure 5.10: Zero D.C. by symmetry

This definition can be written in terms of three functions that have standard Fourier transforms.

$$\text{Let } A = u\left(t + \frac{T}{2}\right) - u\left(t - \frac{T}{2}\right)$$

$$\text{Let } B = e^{j2\pi nt/T}$$

$$\text{Let } C = \cos \frac{2\pi t}{T}$$

$$\psi(t) = \frac{1}{2}(AB + ABC) \quad (5.30)$$

The Fourier transforms of  $A$ ,  $B$  and  $C$  are

$$\mathcal{F}\{A\} = \frac{2 \sin \frac{\omega t}{2}}{\omega} = T \text{sinc} \frac{\omega T}{2}$$

$$\mathcal{F}\{B\} = 2\pi \delta\left(\omega - \frac{2\pi n}{T}\right)$$

$$\mathcal{F}\{C\} = \pi \left( \delta\left(\omega - \frac{2\pi}{T}\right) + \delta\left(\omega + \frac{2\pi}{T}\right) \right)$$

By applying the Convolution Theorem,

$$\mathcal{F}\{AB\} = T \text{sinc} \left( \frac{\omega T}{2} - \pi n \right) \quad (5.31)$$

$$\mathcal{F}\{ABC\} = \frac{T}{2} \left( \text{sinc} \left( \frac{\omega T}{2} - \pi(1+n) \right) + \text{sinc} \left( \frac{\omega T}{2} + \pi(1-n) \right) \right) \quad (5.32)$$

Hence, the Fourier transform of  $\psi$  is given by

$$\Psi(\omega) = \frac{T}{2} \left( \text{sinc} \left( \frac{\omega T}{2} - \pi n \right) + \frac{\text{sinc} \left( \frac{\omega T}{2} - \pi(1+n) \right) + \text{sinc} \left( \frac{\omega T}{2} + \pi(1-n) \right)}{2} \right) \quad (5.33)$$

This function is plotted in Figure 5.11 for several values of  $n$ . The plot notably illustrates sidebands occurring either side of the main passband. Their presence is an inevitable consequence of confining the impulse response to a finite interval within the time domain. The first null on each side is located  $4\pi/T$  radians from the centre frequency of the passband, which is given by


$$\omega_c = \frac{2\pi n}{T}$$

Beyond the main lobe, the nulls occur at intervals of  $2\pi/T$ . It then becomes clear that odd values of  $n$  may also be used (with the exception of  $n = 1$ ) and the wavelet function will remain admissible. However, a sinusoid at a frequency of  $\omega_c$  will alternate from being in-phase to being out-of-phase with each translation step.

The frequency transfer function of  $\psi$  is plotted in Figure 5.12 with logarithmic scaling in both axes. What this graph highlights is the significance of the sidebands, and the asymmetry of the main lobe when viewed on a logarithmic frequency scale. On the first of these matters, the sideband rejection seems inadequate at less than 32dB referred to the maximum gain. This is problematic in that it is obstructive towards efficient sampling of the signal after convolution with the wavelet. The upper sidebands that exceed the Nyquist frequency (half the sample rate) will be aliased into the main passband. To minimise this effect, the Nyquist frequency should correspond with the upper null beyond which the sidebands are deemed to be insignificant. For example, keeping aliasing noise below -60dB requires a sample rate in Hz of

$$\frac{2(n+7)}{T} .$$

Other than this, the presence of the sidebands is not that inconvenient, perhaps even beneficial. They would aid the intuitive representation of transient events in the signal,



Varying the value of  $n$  has the effect of translating the raised-cosine wavelet function in the frequency domain. Only discrete integer values can be used or else the Fourier transform would not pass through the origin, indicating a d.c. component. This would not be admissible as a wavelet function.


Figure 5.11: Fourier transform of raised cosine wavelets

suppressing ringing and “premonition” type effects in the resulting analysis.

Considering now the shape of the main lobe, its apparent asymmetry may not seem important to begin with, as a pure consequence of the axis scaling, but it becomes critical in light of the geometric series relationship between wavelet scales. The overlapping slopes between adjacent wavelet scales are asymmetrical, independent of axis scaling (ref. Figure 5.13). This feature prevents the sum of scales from producing a flat response, regardless of scale factor or normalisation. As a result, this design of wavelet function does not meet the criterion of being unbiased towards particular frequencies or time offsets. In view of that fact, a more suitable wavelet approach was sought.

### 5.4.2 Frequency-domain View

For efficiency reasons, the DWT is typically implemented using a cascade of filters, as described in §5.3.3. When using a dyadic grid, which is the norm, the sequence is decimated



The sidebands and central lobe appear skewed on a logarithmic frequency scale, such as this octave scale, portending a discrepancy between the overlapping slopes of adjacent bands when incorporated into a filter bank.

Figure 5.12: Transfer function of raised cosine wavelet




Figure 5.13: Asymmetrical slopes of adjacent wavelet scales

by a factor of two after each filter stage so that it is critically sampled. It also allows the same set of filter coefficients to be used at each stage. In effect, the filters are stretched by a factor of two, which halves the cutoff frequency. The filters used must therefore be half-band filters. The low pass filter in each stage serves to prevent the upper portion of the signal's bandwidth from being aliased into the lower portion when the sample rate is halved.

It is quite simple to extend this concept of a fast wavelet transform to the non-dyadic case by realising that it is effectively a cascade of resampling operations. The only difference compared with the usual resampling process is the addition of a second, high pass output, which gives the residual from the antialiasing filter. Notably, the high pass output can be decimated as well as the low pass output. Even though it contains frequencies that would exceed the Nyquist frequency there will be no ambiguity with lower frequencies because they have been rejected by the filter.

In order to achieve a non-dyadic grid, one can apply a common method that is used to resample a discrete-time signal by a non-integer factor. This involves upsampling the signal by an integer factor first, then applying the antialiasing filter, before decimating by a second integer factor. Now the final sample rate is fractionally related to the initial one. By using this method at each stage of a FWT, the factor relating adjacent wavelet scales can be changed from two to any rational number, especially one between 1 and 2, e.g.  $\frac{3}{2}$  or  $\frac{5}{4}$ . There is a corresponding improvement in the frequency resolution. The process will be described in more detail at the end of the section.

Using the approach described above presents a filter design problem. The criteria set down at the beginning of §5.4 now have to be interpreted as constraints on the filter characteristics, as follows:

- To maximise orthogonality, the overlap between high- and low-pass sections should be minimal (subject to other constraints).
- Each filter stage must have complimentary high- and low-pass sections whose outputs sum to produce an approximately flat amplitude response. This will avoid bias towards particular frequencies.

- The filter-pairs should be in phase where they overlap to prevent over-representation of energy in these regions, i.e. the individual amplitude responses also sum to unity.
- The filters should ideally be zero-phase to allow comparison of the relative phase between bands. The consequent symmetry of the impulse response also prevents skewing of the signal's amplitude envelope.

There are several methods commonly used to design digital filters that could potentially be used here. The Bilinear Transform is a method used to design recursive digital filters that are equivalent to analogue filters. Alternatively, the poles and zeros of the filters may be placed manually in the  $z$ -plane to achieve a target response. Infinite impulse response filters produced using these methods are suboptimal for the intended application. They possess inherent phase distortion, and produce an infinite sequence at the output even when applied to a finite input sequence. This output may not have the intended properties once it has been truncated. In the frequency domain, the truncation convolves the signal spectrum with the Fourier transform of the bounding window (typically a sinc function), which introduces energy outside the passband of the filter. Finite impulse response (FIR) filters are preferable. These can be designed to be linear phase. Because the filter length is restricted to a given number of taps, the output of the filtered data is finite in length, albeit longer than the input sequence. It is not necessary to truncate the output, in which case the action of the filter will be exactly as expected. The design of the filter can be optimised to meet a given specification using a variety of methods.


One way to design a linear phase FIR filter is to start with an ideal brick wall filter, which has an infinite impulse response. For a brick wall low pass filter, the transfer function is a rectangular function, which becomes a sinc function in the time domain when the inverse Fourier transform is applied. Such a filter is non-realisable because it is not causal [63]. The impulse response extends infinitely in both the positive and negative directions on the time axis. Future input samples would be required to evaluate the output (as well as an infinite record of the signal). However, the sinc function tends to zero as time approaches  $\pm\infty$ , so the contribution to the filter output from samples that are far removed from the

present time point is insignificant. By applying a symmetrical window to the sinc function, a zero-phase FIR filter can be created that approximates a brick wall filter. Although this is still not causal, the filter can be realised if sufficient delay is introduced because the extent of the “future” input required is finite. Now it is linear-phase rather than zero-phase (a delay can be seen as a phase shift that is proportional to frequency), but if the delay through all paths in the system is compensated for then the distinction is academic. Figure 5.14 depicts the steps described above.

Windowing the impulse response does affect the transfer function of the filter, which will be convolved with the Fourier transform of the window. In general, the longer the window that is used, the steeper the roll-off and the greater the stop-band attenuation achieved by the filter. The exact relationship between these parameters is a property of the chosen window function. Consequently, it is only possible to directly optimise the filter for one of the two parameters (by adjusting the number of taps). Meeting both specifications tends to cause one of them to be exceeded. It is not trivial to generate a window that produces a desired ratio between the two parameters, so filters produced using this method are inefficient on the number of taps used.

Another way to design nonrecursive, linear phase digital filters is by optimising the filter coefficients using numerical methods. One such approach, known as the Parks-McClellan method (after the authors of the original Fortran implementation), uses the Remez exchange algorithm to solve a minimax problem [68]. This allows the designer to specify the filter in terms of a target frequency response. The maximum deviation from the target response will be minimised for a given number of taps. The target response must consist of series of bands where the desired gain is constant or that of a differentiator (a first-order high-pass response). Specified bands are separated by “don’t care” regions of finite width in which the filter gain cannot be specified. The bands may each be weighted so that the allowable deviation in one band is different from that of another.

In the simplest application of the Parks-McClellan method, a low- or high-pass filter can be designed where the transition bandwidth and filter length are specified by the



These plots demonstrate the steps used to produce a low pass filter using the windowing method. The top axes show the impulse response of an ideal filter, which extends *ad infinitum* in both directions. In the middle axes, a window (shown in gray) has been applied so it has 31 taps, but 15 of these occur in negative time, meaning that the filter is still not realisable. The problem is overcome by introducing a 15-sample delay, as shown in the bottom axes.

Figure 5.14: Impulse responses of ideal and windowed filters

designer. The weights can be manipulated to vary the balance between the resulting passband ripple and the stopband attenuation on a continuous scale. However, these two parameters are not specified directly, since they are also a function of the number of taps and the transition bandwidth. Fewer taps equates to more ripple in the passband and poorer attenuation in the stopband. Reduction of the transition bandwidth also comes at the expense of more ripple and worse attenuation. The principal advantage over the windowing method is the precise control over the trade-off between parameters. This has given rise to the development of an empirical formula that effectively gives the designer control over all but one of the following parameters: passband ripple, transition bandwidth, stopband attenuation and filter length [69]. Kaiser's formula is as follows:

$$\begin{aligned}\delta_1 &= \frac{1}{10^{r/20} - 1} \\ \delta_2 &= 10^{a/20} \\ N &= -\frac{10 \log_{10} \delta_1 \delta_2 - 13}{14.6 \Delta F}\end{aligned}\tag{5.34}$$

where  $r$  is the passband ripple in dB,  $a$  is the stopband attenuation in dB,  $N$  is the number of taps and  $\Delta F$  is the transition bandwidth. All of the proposed constraints detailed earlier can be met by employing this method. Although the transition band is considered to be a “don't care” region, it is evident (without proof) from the consistency in the shape of the responses produced, that a smooth transition occurs between the passband and stopband. Intuitively, this makes sense; any other response would increase the slope of the transfer function in that region. In turn, this would unnecessarily increase the ripple in the surrounding bands, so it does not occur. A procedure for designing the filter and performing the time-frequency analysis is outlined below.

The steps required to design and apply the pair of filters used at each stage of the time-frequency decomposition are:

1. Choose a ratio between two integers  $\left(\frac{\zeta}{\eta}\right)$  to be the common factor between the Nyquist frequency  $\left(\frac{f_s}{2}\right)$  of successive bands. For example,  $\frac{4}{5}$  (upsampling by 4, decimation by 5).
2. Choose a low frequency passband edge somewhat lower than this to allow for the transition bandwidth. For example,  $\frac{3}{4} \times \frac{f_s}{2}$ . This is also the stopband edge of the high pass filter. Smaller transition bands require more taps, extending the ringing and delay time.
3. Design the high pass filter with the passband edge at  $\frac{\zeta}{\eta} \times \frac{f_s}{2}$  and the transition bandwidth as determined above.
4. Design the low pass filter at the upsampled rate  $(\zeta f_s)$ . This requires the normalised cutoff frequency and transition bandwidth to be divided by the upsampling factor,  $\zeta$ . It will need  $\zeta$  times as many taps to keep the other filter specifications the same.
5. Apply the high pass filter to the input data to produce the high pass output. A decimation may be applied after the filter to reduce redundancy.
6. Upsample the input data by inserting  $(\zeta - 1)$  zeros between each sample. Apply the low pass filter, followed by a decimation by  $\eta$ .

The same filters are used at every stage in the analysis, so they need only be designed once. The last two steps are repeated as necessary to produce the desired number of bands. The high pass outputs from each stage are the outputs of the filter bank, while the low pass outputs feed into the subsequent stage. The final low pass output gives the residual from the input that is not represented in any other output.

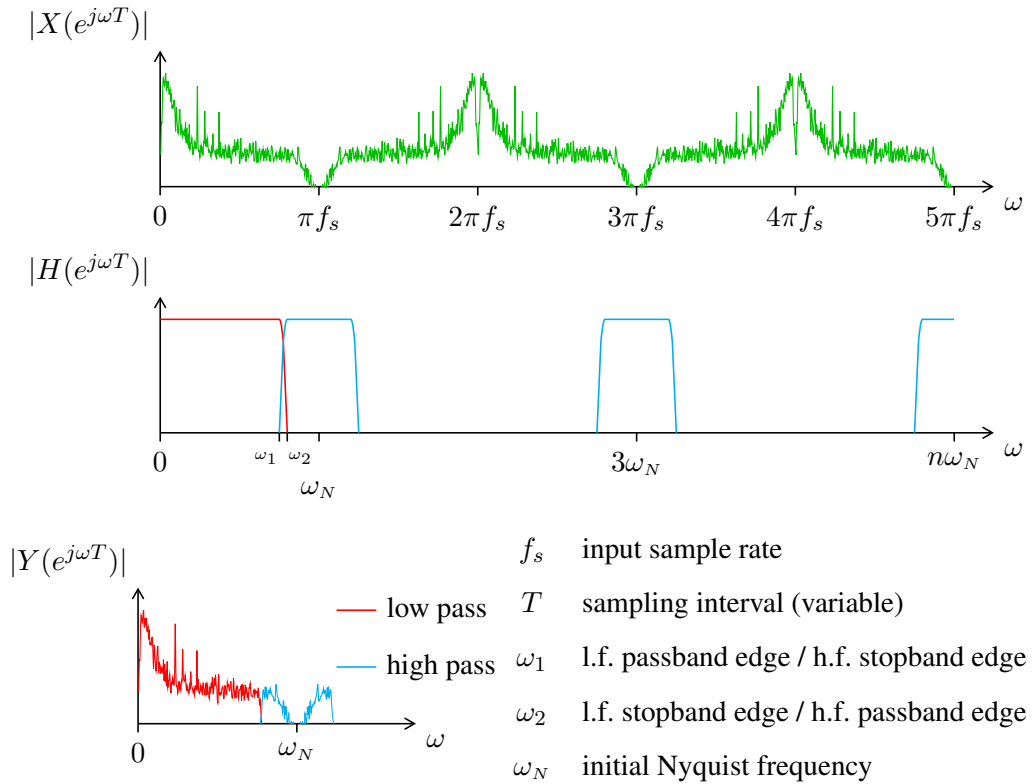
In order to understand why this process works it is necessary to appreciate that each sample of the input signal represents a scaled impulse in the continuous time domain. This is the origin of the term *Pulse Code Modulation* (PCM) to describe discrete-time sampling.

PCM sampling produces modulation products either side of the sampling frequency and its multiples. Usually, the signal to be sampled is limited to half the sampling rate (the Nyquist frequency) so that the lower sideband of the first harmonic does not encroach upon the baseband signal. The low pass filter used to limit the bandwidth is called an *antialiasing filter* because it prevents “aliasing,” or ambiguity between signal components occurring above and below the Nyquist frequency. With this constraint, it is possible to reconstruct the baseband signal from the discrete representation by using a suitable low pass filter.

The appropriate way to upsample a PCM signal is to insert zeros between the data points. A digital reconstruction filter should be applied afterwards because the modulation products from sampling the signal at the original rate now extend below the increased Nyquist frequency. In this application, the same filter also serves as the antialiasing filter for the subsequent decimation, which resamples the signal again at a lower frequency. Furthermore, the filter’s impulse response fulfils the same purpose as the wavelet scaling function in the DWT, determining the ultimate characteristics of the analysis. In a practical filter, the transition bandwidth must be finite, which leads to a degree of overlap between the high- and low-pass filters in each stage. This is analogous with the crossover network in a 2-way loudspeaker; in both applications the bandwidth is divided into two outputs that sum to recreate the input. Permitting a reasonable amount of overlap has the advantage that the ringing of the filter in the time domain is reduced, which will otherwise be visible in the output of the analysis. The filter cutoff must be lower than the final Nyquist frequency by some margin so that the transition band can be accommodated. The normalised cutoff frequency  $\left(\frac{2f_c}{f_s}\right)$  should be greater than the square of the common ratio  $\left(\frac{\zeta}{\eta}\right)$  so that the low pass cutoff does not fall below the high pass cutoff of the subsequent stage.

### Implementation

It is desirable for both filters to have an odd number of taps, as the delay will otherwise be fractional. The half-sample offset would prevent any alignment between samples in



The Fourier spectrum of the input signal is plotted on the top axes, with angular frequency on the horizontal scale. Because the signal is real, there is also a negative frequency component that mirrors the positive frequency one (not shown). When  $T = f_s^{-1}$ , images of the baseband signal appear either side of each harmonic of the sample rate. After upsampling,  $T = (\zeta f_s)^{-1}$ , and these images now fall within the baseband signal. The subsequent low pass filter (middle axes, red) must reject them. The high pass filter (cyan) is designed at the original sample rate, so it has the same periodicity as the input signal in the frequency domain (the period of the low pass filter is  $\zeta$  times longer, so the first image is off the scale). The bottom axes show the filter outputs before decimation.

Figure 5.15: Illustration of filter design and action upon MEG signal

different bands, making indexing them complicated. However, when the upsampling factor  $\zeta$  is even, the low pass filter is always assigned an even number of taps. To remedy this situation, the specified number of taps for either filter is increased by one if it is even. Doing so does not affect the transition bandwidth using the Parks-McClellan method, and the other specifications are exceeded.

A practical consideration in implementing the cascade of filters and decimations is that of sequence length. When a sequence is decimated, the length of the output sequence is usually the input length divided by the decimation factor, without the remainder. Over many successive decimations, this behaviour results in a shortening of the sequence, so the lower frequency outputs will represent progressively shorter durations. Clearly, this is unwanted because genuine parts of the signal are being lost. One way to avoid the problem is to ensure that the length of the output sequence is always rounded *up* to the next sample. This actually makes the duration of the sequence grow over several decimations, which can make the transform inefficient with large numbers of bands. The optimal solution is to make the filtering and decimation procedure aware of the original sequence duration so it can ensure that it is only just exceeded at each stage.

A full implementation of the filter bank scheme has been written by the present author in Python, using the SciPy package. SciPy provides functions similar to those found in MATLAB's signal processing toolbox. The filter bank runs very efficiently by making use of FFTs and the convolution theorem to perform most of the calculations in the frequency domain. Upon initialising the filters, their impulse responses are padded with zeros and transformed into the frequency domain for a selection of FFT lengths. For optimum efficiency, these lengths are all powers of two. When a filter operation is performed, the FFT length chosen is the smallest power of two greater than or equal to the complete convolution of the sequence with the filter. The output is truncated back to the original length. This behaviour circumvents the problem of wraparound otherwise caused by the assumption of stationarity implicit in using an FFT.

The transform is not shift invariant, i.e. when the input is shifted in time the output does

not simply get shifted by the same offset. Nor can it be shift invariant, considering that the sample rate of the outputs is different from that of the input. The smallest increment (in samples) that will maintain the same temporal alignment between the samples at each output is  $\zeta(B-1)$  where  $B$  is the number of bands. Offsetting the input by integer multiples of this value *will* produce a shifted version at the output, but other offsets will not. In practice, comparing or visualising the outputs from this transform requires some kind of interpolation. There are all sorts of approaches to interpolation offering varying degrees of accuracy, e.g. linear, cubic spline, FFT interpolation. The most accurate but slowest method uses an ideal filter (convolution with a sinc function) to resample the data to the display resolution.

A very useful extension to the transform is to combine it with a Hilbert transform [70], which is defined in the continuous time domain as

$$\mathcal{H}\{x(t)\} = x(t) * \frac{1}{\pi t} \quad (5.35)$$

From the following Fourier transform, it is evident that the Hilbert transform constitutes a  $90^\circ$  phase shift for all frequencies, transforming between sine (odd) and cosine (even) components.

$$\mathcal{F}\left\{\frac{1}{\pi t}\right\} = j(1 - 2u(\omega)) = \begin{cases} e^{-j\pi/2} & \omega > 0 \\ e^{j\pi/2} & \omega < 0 \end{cases} \quad (5.36)$$

where  $u(t)$  is the Heaviside step function. Hence, by the convolution theorem,

$$\mathcal{H}\{x(t)\} = \mathcal{F}^{-1}\{X(\omega)j(1 - 2u(\omega))\} \quad (5.37)$$

If this transformed signal is taken to be the imaginary part and the input signal the real part, the resulting complex signal is called the *analytic signal* [70], which can be found by removing the negative frequency components.

$$\overset{\nabla}{x}(t) = x(t) + j\mathcal{H}\{x(t)\} = \mathcal{F}^{-1}\{X(\omega)2u(\omega)\} \quad (5.38)$$

This additional process is useful in that the modulus of the analytic signal follows the envelope of the oscillations, while the argument follows the phase. One of the major benefits

afforded by using the filter bank approach and Hilbert transform, compared with other methods like spectral coherence, is the ability to perform correlation analyses that treat magnitude and phase as independent phenomena. Others have applied either the Hilbert transform or analytic wavelets to neurophysiological data for similar purposes [71], but this is a more elegant solution that integrates the two, obviating the need to preprocess data with filters prior to wavelet analysis.

Filters called *Hilbert transformers* possessing the appropriate phase shift may be designed using the Parks-McClellan algorithm. However, they have an impulse response with odd symmetry, requiring that a high pass filter has an even number of taps and a fractional delay in samples [63]. This could not simply be summed with the usual filter output to produce analytic signals. Instead, it is convenient to zero the negative frequency components in the Fourier domain, especially if one already calculates the output by FFT. In this case it is wise to increase the FFT length by some margin. While the convolution of an FIR filter with a finite sequence also has finite length, its Hilbert transform is not subject to this constraint, tending instead towards zero beyond the bounds of the real sequence. In the implementation used to produce the included examples, the FFT length was incremented to an additional power of two to accommodate this extra temporal spreading. No more than this is required since the extremes of the filters' impulse responses substantially diminish towards zero by nature with reasonable filter parameters.

With the inclusion of the Hilbert transform, the shift invariance becomes less important because the phase relationship between the sample clock and the oscillatory input no longer affects the appearance of the envelope. When visualised, the effect is comparable to a spectrogram calculated using the STFT, which shows energy density. Without the Hilbert transform, the oscillations themselves are visible and the interpolation filter is crucial to recovering them from the sparse coefficients.

### Examples

As an illustration of the properties of the filter bank, the impulse response has been graphed in Figure 5.16. This is the magnitude of the analytic output, displayed with a logarithmic amplitude scale. The filter bank parameters used in this and all subsequent examples are shown in Table 5.1. There is visible ringing in the time domain as a consequence of the relatively sharp transitions in the frequency response of the filters. With practical input signals, such ringing is not likely to be so apparent.

Figure 5.17 illustrates the effect of applying the filter bank to the averaged epoch data used in Figure 5.1 near the beginning of the chapter. Figure 5.17(a) shows the real part of the output from the process, with a linear amplitude scale. Sinc interpolation was used to recover the oscillations at the display resolution. If these oscillations are summed along the frequency axis, a close approximation to the original signal is obtained (see Figure 5.18). In Figure 5.17(b), the logarithmic amplitude of the analytic signal was taken and linearly interpolated. For comparison, a spectrogram calculated using Fourier methods is shown in Figure 5.19. A 256-point FFT window with a Hann window function was used. The windows were overlapped by 50% and the logarithm of the resulting spectral densities was linearly interpolated along the time axis. These parameters allow the best comparison with the filter bank output. With a logarithmic frequency scale, the resolution of the STFT is seen to increase with frequency while the temporal resolution remains constant. Conversely, the filter bank output has more temporal detail in higher bands and equal bandwidths on a log scale. Slight ringing is visible in the very highest bands at either end of the filter bank output, which is due to discontinuities at edges of the rectangular window encompassing the entire sequence.

Measurements from MEG experiments often consist of spontaneous activity followed by the response to a stimulus. Hence it is useful to evaluate the response in terms of increases or decreases in power, with respect to a baseline level established shortly before the stimulus was presented. Figure 5.20 depicts the change in power in each band with respect to the baseline level in that band. The edge effects in the top two bands are rather noticeable here

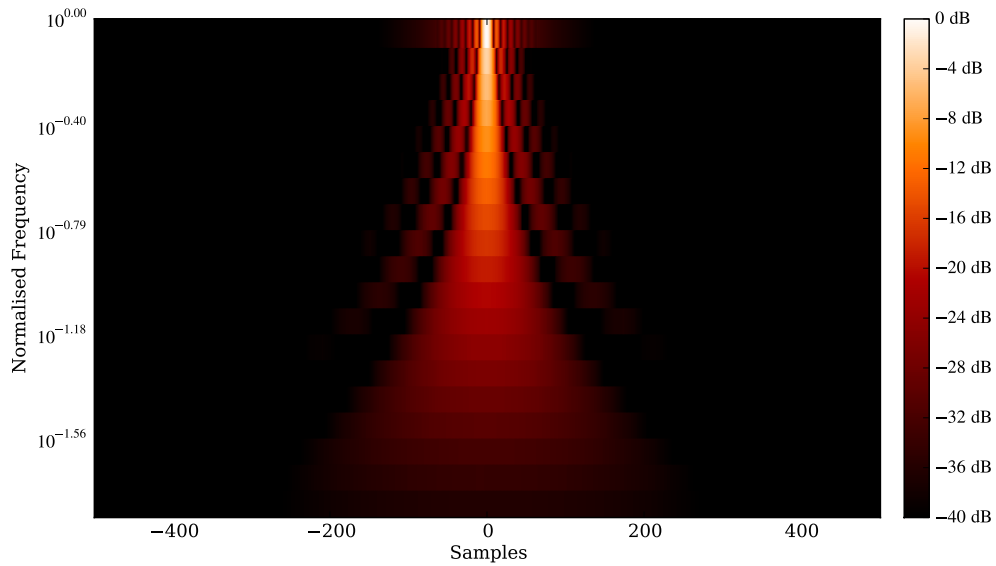




Figure 5.16: Impulse response of filter bank

because there is very little energy in those bands overall, resulting in low baselines and increased sensitivity. The most extreme time points can be dewighted so that the output may be submitted to correlation analyses without spurious results. This form of output is then particularly useful for identifying consistencies between separate measurements, i.e. whether a stimulus induces characteristic variations in power at specific frequencies and physical regions. Hence it is also this form that will be used in subsequent stages of the signal processing framework.

It is evident from these results that the broad aims of the filter bank method have been fulfilled. The essential benefits of discrete wavelet techniques have been retained, i.e. logarithmic frequency scale and efficient representation. Furthermore, a substantial improvement in the frequency resolution has been achieved. Calculation of the analytic signal at each bandpass output allows envelopes free of phase information to be compared. The fact that the input signal can be recovered from the oscillatory output to within a small error demonstrates that information is largely retained by the process.




(a) Oscillatory output




(b) Amplitude envelopes

Figure 5.17: Examples of filter bank output



(a) Comparison of input with reconstruction



(b) Reconstruction error in standard deviations

Owing to repeated truncation of the low pass outputs during the recursive filter bank process, some low-frequency energy is lost towards the very edges of the sequence. These effects aside, the absolute error in the reconstructed signal remains under 0.1 standard deviations over the majority of the sequence.

Figure 5.18: Reconstruction of input signal

| <i>Parameter</i>     | <i>Ratio</i> |
|----------------------|--------------|
| Upper Band Edge      | 0.8          |
| Lower Band Edge      | 0.75         |
| Upsampling Factor    | 4            |
| Downsampling Factor  | 5            |
| Passband Ripple      | 0.00868 dB   |
| Stopband Attenuation | 60 dB        |

Table 5.1: Default filter bank parameters




Figure 5.19: Short-time Fourier transform




Figure 5.20: Change in power with respect to baselines

## 5.5 Summary

For the intermediate time-frequency analysis stage of the proposed signal processing framework, a number of conventional methods were considered. The best known method is the *short-time Fourier transform* (STFT). The STFT is not well suited to MEG signals. Neurophysiological signals naturally fall into roughly equal bands on a logarithmic scale, in contrast with the linearly-spaced DFT bins. If the window is made long enough to accommodate the lowest frequencies present, the high frequency temporal resolution is poor. Windowing sections of data in this way causes spectral leakage because of discontinuities at the window edges and the assumption of stationarity. Non-rectangular windows mitigate this effect but impair frequency resolution by widening the spectral peaks.

As an alternative to the STFT, wavelet methods were also considered. Wavelets are short enveloped oscillations that act much like band pass filters. They can be tuned to different frequencies by stretching them in time. A range of tunings are convolved with the data to determine the frequency distribution of signal energy. Wavelet analyses are a natural choice

for a logarithmic scale because the wavelet bandwidth is proportional to its centre frequency. The *continuous wavelet transform* (CWT) is defined for any time offset or frequency. There is not usually an obvious way to discretize these scales, so a dense number of points is calculated. This leads to very smooth and highly redundant representations of the data. Large redundancy is not appropriate for an application where data reduction is important.

A *discrete wavelet transform* (DWT) exists that is only defined at discrete frequencies and time points. For this, a logarithmic frequency scale is enforced. The time intervals are spaced according to frequency, such that high frequencies have proportionally finer time resolution. The conventional discretization uses a frequency resolution of one octave, doubling the number of time points with each increment in frequency. This *dyadic grid* scaling can provide a representation of the data with no additional redundancy. The *fast wavelet transform* (FWT), an efficient implementation of the DWT, uses a cascade of half-band filters to generate the wavelet coefficients. Successive outputs are the residual from each halving of the signal resolution. The octave-wide frequency resolution is deemed too coarse for the intended application, so the DWT is also unsuitable.

The FWT is equivalent to resampling the data at half the rate, and retaining the residual part using a high pass filter that complements the antialiasing filter. Based on this observation, a novel type of filter bank has been developed that uses an established method for resampling by fractional factors. At each stage, the data is upsampled by an integer factor. This is accomplished by inserting zero-valued samples between the existing ones. At this point, the sidebands from pulse code modulation (sampling) at the original sample rate are still present. Next, a low pass filter is applied to isolate the baseband signal and reduce its bandwidth to the new resolution. Crucially, the new resolution can be any integral fraction of the initial sample rate, e.g.  $\frac{2}{3}$ ,  $\frac{4}{5}$  etc. The filtered sequence is then decimated by the denominator of the fraction to produce the next approximation. The band pass output of each stage is generated using a complementary high pass filter prior to the upsampling process.

The major benefit of this novel approach is liberation from the dyadic grid associated

with conventional DWTs. Redundancy in the data is minimal, but the design of the filters is crucial to the success of the process. It is necessary to set the filter cutoff some margin below the Nyquist frequency to allow for a transition band. Otherwise out-of-band signal energy will be aliased into the low pass output upon decimation. Finite impulse response (FIR) filters are used, designed using the widespread Parks-McClellan algorithm. Other filter types were considered, but this method allows both the stop band rejection and pass band ripple to be simultaneously constrained. Therefore the filter can be optimised to a particular error tolerance in this application.

Representations of MEG data produced by the novel filter bank reveal detail that is not clear in other time-frequency analyses. Furthermore, calculation of the analytic signal using the Hilbert transform allows detachment between amplitude variations and oscillation phase. These properties make this an ideal preprocessing step for statistical comparisons of neuronal activity among distinct regions of the cerebral cortex.

## 6 Statistical Framework

The goal of the framework proposed in Chapter 3 is to be able to draw objective conclusions concerning connectivity between brain regions. To achieve this requires the introduction of a statistical model which can be employed in hypothesis tests. Only by using these formal techniques can reliable inferences be made from the data. Generalisations about brain function made without recourse to statistical analysis are liable to be subjective and may not withstand closer scrutiny. This chapter introduces the use of inductive statistics to draw wider conclusions from the data than simple descriptive statements specific to one set of measurements. Examples of the application of these techniques within the framework are given at the end.

### 6.1 Theory

In statistics, a population is a collection of entities whose characteristics one can gather data on. A particular characteristic of a population is assumed to follow a random distribution. That is, individual observations are unpredictable, but the observations can be described collectively by a probability density function. It is frequently impossible to observe every member of a population, especially if it is infinite. However, by measuring a representative sample of the population, conclusions can often be inferred regarding the population as a whole [72].

The most important tool in statistical inference is the hypothesis test. A hypothesis in statistics is a statement about some population parameter [73]. This could be a measure of location (e.g. the population mean), a measure of spread (e.g. standard deviation), or

some other description of the population's distribution. The validity of the statement is then assessed by consulting a test statistic, which is a random variable with a known distribution.

The test statistic is an estimator for the population parameter referred to in the hypotheses [73]. Knowledge of its distribution permits a particular value to be assessed for statistical significance. If the test statistic is found to have a value that is inconsistent with its distribution under the *null hypothesis*, then the *alternative hypothesis* (the assertion to be validated) is inferred to be the cause. An “inconsistent” value is one that is highly improbable under the null hypothesis. A hypothesis test requires a critical *p-value* to be defined first, which is a probability below which the likelihood of the calculated test statistic must fall for it to be deemed significant.

In the context of the proposed MEG analysis framework, neuronal activations in regions of interest are modeled as random processes that depend on time. MEG recordings produce time series that contain observations on these random processes. In the interests of separating effects at different frequencies, these time series have previously been filtered with narrowband filters (see Chapter 5). The signal amplitudes as determined by the modulus of the analytic signal are the observations. It may seem unusual to preprocess the data so heavily before statistical analysis, but it serves to enhance particular features in the data that might signify connectivity where others do not. To establish the existence of neuronal connections, the time dependent variations in activity can be compared between brain regions. Measures of correlation are the obvious statistics to use for this purpose.

### 6.1.1 Correlation

Correlation is a measure of association between two random variables that depend on a common parameter. The degree of correlation between a pair of variables is quantified by a correlation coefficient, which is a measure of how much they covary. If the random variables are labeled  $X$  and  $Y$ , the *product-moment correlation coefficient* (PMCC) is given by [72]

$$\rho_{xy} = \frac{\sum_{\forall i} (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{\forall i} (x_i - \bar{x})^2 \sum_{\forall i} (y_i - \bar{y})^2}} \quad (6.1)$$

where  $x_i, y_i$  are the random variates, and  $\bar{x}, \bar{y}$  are the sample means. The PMCC differs from the covariance only by a normalisation factor that compensates for the scale of X and Y, leading to a coefficient that is comparable between experiments. This measure relates to a least-squares linear regression model, and the square of the coefficient is the proportion of the variation that is explained by the model [72].

### 6.1.2 Partial Correlation

In many cases, a pair of variables can be correlated due to the influence of a *confounding variable*, with which both variables are also correlated. They may have no correlation at all that cannot be explained by this third variable. To quantify the degree of correlation between X and Y that is not explained by Z, a partial correlation coefficient may be calculated [74].

$$\rho_{xy.z} = \frac{\rho_{xy} - \rho_{xz}\rho_{yz}}{\sqrt{(1 - \rho_{xz}^2)(1 - \rho_{yz}^2)}} \quad (6.2)$$

The ability to remove the influence of confounding variables is only possible when all such variables are known about and have been observed with the other correlates [74]. It is not possible to establish causality using these measures.

### Conditional Independence Models

Partial correlation coefficients gives rise to conditional independence models, which can represent a whole system of random processes and their interdependencies. These are typically represented by a graph consisting of nodes and edges. The nodes are the random variables, and edges between them represent associations between those variables. The absence of an edge between two variables signifies that they are independent, conditional upon the other variables in the system.

### 6.1.3 Cross Correlation

In their native representation, MEG measurements are a function of time. In the previous discussions, the pairs of observations on the random processes were treated as any other sample taken from a population. In the case of time series data such as MEG recordings, this assumes that statistically associated variations occur simultaneously. Substantial transmission delays have long been established to occur in neurons and across synapses [29]. Only three explanations for ordinary correlation between signals measured in different brain regions can then be plausible. The first is that the onset of response to an external stimulus occurs after approximately equal delay time in those particular regions. The second is that similar variations occur spontaneously, possibly due to some synchronisation mechanism or past event. More likely, deficiencies in the measurement process itself are to blame, especially in the case of MEG, where regional selectivity can be difficult to achieve.

Cross correlation provides a means of estimating the correlation between variables as a function of time lag ( $l$ ). This allows for the possibility of transmission delays as signals pass from one brain region to another. If an association between the pair exists at a particular time offset, a corresponding positive or negative peak is seen in the sequence. The unbiased estimate of the signal cross correlation is given by [75]

$$\hat{r}_{xy}(l) = \frac{1}{N-l} \sum_{n=0}^{N-1-l} x(n)y(n+l) \quad (6.3)$$

The normalisation  $1/(N-l)$  is used so that the sequence does not rise and fall as the proportion of overlapping points in the sum changes.

One benefit of studying cross correlations is the ability to determine the direction of influence between variables. If one concedes that influences from future to past cannot take place, then the location of the cross correlation peak reveals which variable drives the other. Cross-correlations of periodic signals tend to contain a series of maxima and minima, whose spacing relates to their periods of oscillation. The identification of the peak correlation, if there is one, can be rendered troublesome by this phenomenon. The use of

amplitude envelopes produced by the filter bank in this framework neatly avoids the issue. The oscillations are removed by taking the envelope, and the bandwidths are so narrow as to avoid excessive amplitude modulation in the case of oscillations at nearby frequencies.

## 6.2 A Test Statistic

The linear regression model underpinning statistical correlation measures is relevant only to stationary processes. This means that the statistical parameters are assumed not to change over time. One of the postulates to be investigated using this MEG analysis framework is that of dynamic connectivity. In that case, one cannot make such assumptions, so the data are not amenable to conventional correlation statistics. In particular, the MEG recordings are typically divided into short segments, called epochs. The epochs are categorised according to experimental condition, and one expects the associations between random processes to change under different conditions. In the calculation of most correlation coefficients, the sample mean is subtracted from the observations. Only the variation about the mean is predicted by a linear regression model; the d.c. component is fixed in the model parameters. In the case of amplitude envelopes, failure to take account of the d.c. component would cause overall increases or decreases in neuronal activity in response to a stimulus to be ignored.

The fact that the assumptions involved in the linear regression model are not valid for the signals in question does not mean that one cannot make use of correlation-like measures. In the field of signal processing, a cross-correlation sequence is often defined without subtracting the mean from the data. Furthermore, most MEG experiments are designed with “passive” periods immediately prior to the stimuli. This data can provide a reference level that may be used to centre the data, so that long term changes unrelated to the stimulus are discounted.

Another disadvantage of conventional cross-correlations relates to the distribution of the coefficients in the resultant sequence. These are typically normalised using fixed variance estimates based on the entirety of each input sequence. With finite length sequences,

different parts of the sequence necessarily take part in the calculation of each coefficient. Because the average power in these signal segments can vary, the coefficients are not bounded to the interval  $[-1, 1]$ , as one might expect.

For the cross correlation calculations in this framework, an unconventional definition was used in which the coefficients were normalised by the r.m.s. power in the overlapped parts of the sequence. Also, the coefficients are normalised by the number of overlapping points, as in the previous definition. Both normalisations can be performed in one calculation by evaluating the square root of the product of the sum of squares.

$$\frac{1}{\sqrt{\sum_{n=n_{x_1}}^{n_{x_2}} x^2(n) \sum_{n=n_{y_1}}^{n_{y_2}} y^2(n)}} \quad (6.4)$$

$$\begin{aligned} \text{where } n_{x_1} &= \max(0, -l) \\ n_{x_2} &= \min(N - 1 - l, N - 1) \\ n_{y_1} &= \max(l, 0) \\ n_{y_2} &= \min(N - 1 + l, N - 1). \end{aligned}$$

The signal powers must be used in the calculation rather than the usual variances because the data are not mean-adjusted. This cross correlation shall be denoted  $c_{xy}(n)$ .

$$c_{xy}(n) = \frac{\sum_{n=0}^{N-1-l} x(n)y(n+l)}{\sqrt{\sum_{n=n_{x_1}}^{n_{x_2}} x^2(n) \sum_{n=n_{y_1}}^{n_{y_2}} y^2(n)}} \quad (6.5)$$

It is important to note that this is not a statistical cross correlation, which is used to infer associations between random processes that are always present (even outside a particular measurement window). This is simply a measure of similarity between the deflections of two signals from a baseline level during a particular epoch. In common with statistical correlation coefficients, a large positive or negative value indicates greater similarity. The mean of this value over multiple trials will become the eventual test statistic for hypothesis testing.

### 6.2.1 Latency Estimate

If separate cross correlation sequences are calculated for each experimental trial, it is possible to calculate a robust estimate for the latency between interconnected regions. A latency estimate was defined based on the peak value of the point-wise  $\frac{\text{mean}}{\text{s.d.}}$  of the sequence set.

If there are  $N$  trials and  $c_{xy}[n]$  is the cross correlation sequence, the estimated latency is defined as

$$\tau_{xy} = \underset{n}{\operatorname{argmax}} \frac{\operatorname{abs}((\operatorname{mean}(c_{xy})) [n]) - \left\lfloor \frac{\operatorname{length}(c_{xy})}{2} \right\rfloor)}{\operatorname{s.d.}(c_{xy})[n]} \quad (6.6)$$

Note that the maximum of the absolute value of the sequence is found, which may be either a negative or positive correlation. By dividing by the standard deviation, peaks in the mean cross correlation that are not consistent across trials are deweighted, which stabilises the estimate.

Once the latency has been estimated, the mean of the set of sequences at the estimated time point forms the test statistic to be used in subsequent hypothesis tests.

### 6.2.2 Null Model

Conventional correlation coefficients follow theoretical distributions that allow a hypothesis test to be conducted on an individual coefficient. In the cases where these measures are applicable, the time points in the data constitute a statistical sample from a stationary random process. Because an unconventional definition of the cross correlation is used, and the data are not assumed to be stationary, such theory no longer applies. Instead, a single time point from the point-wise mean of a set of cross correlations is used to infer if the mean of this “similarity metric”, could it be calculated on the entire population, is greater than would occur with uncorrelated subjects. Also, the statistic is positively biased by incorporation of the latency estimate, which favours time points with higher correlation. The sampling distribution of the mean, including the bias introduced by latency estimation, must be reflected in the model used for the null distribution of the test statistic.

In such a complicated scenario, it is appropriate to devise an empirical model using the

technique of bootstrapping. Bootstrapping can be used to estimate the sampling distribution of a statistic by resampling the observations with replacement a very large number of times [73]. The statistic is evaluated on each of these derived samples. If this is performed a sufficient number of times, the variation seen in the statistic should reflect the distribution of the statistic when sampling from the population.

To produce a model for the null distribution of our test statistic, uncorrelated white noise was used as an input to the process. The filter bank stage of the framework was included within the bootstrapping procedure because the null distribution depends on the power spectral density (PSD) of the data being correlated. A fortunate characteristic of the filter bank is the extremely narrow bandwidth of the filters and the decimation step before each output. These properties cause the separate outputs to have almost identical PSDs as long as the data is noise-like. The power spectrum at the input does not need to be white for this to be the case. As a consequence, the same null model applies to all bandpass outputs from a particular filter bank set-up. MEG data certainly meets the criterion of being noise-like.

The null distribution also depends on the sample size (number of trials) and length of the cross correlation sequences. Shorter sequences are more likely to produce correlations by chance since they contain less information. Likewise, small sample sizes increase the variance of the test statistic. Therefore, the model must be a function of these parameters.

### **Bootstrapping Procedure**

To conduct bootstrapping, 2500 pairs of Gaussian-distributed pseudorandom sequences were created. Following the same process to be applied to MEG data, the sequences were divided into short “passive” and “active” segments, which were separately processed using the filter bank. Amplitude envelopes were produced from each band by taking the logarithm of the modulus of the analytic signal. The passive envelopes were then averaged over time to yield baseline amplitudes, which were subtracted from the active envelopes. Cross correlation sequences were subsequently calculated for each pair of active envelopes, simulating the effect of analysing uncorrelated data.

The method used to find the null distribution for  $n$ -trial data was to take 3000 random samples of size  $n$  from the pool of uncorrelated cross correlation sequences and to evaluate the test statistic on each sample. The result is a sample of 3000 values that follow the null distribution under the relevant test conditions. The process was carried out for a range of sample sizes and sequence lengths, resulting in a comprehensive family of datasets. Histograms for a selection of these datasets have been plotted in Figure 6.1. Due to the bipolar nature of the correlation coefficients, the distributions are actually symmetrical about zero, but the signs were removed in order to better illustrate the trend.

Upon inspection of the histograms, it is obvious that there is a positive skew. After a trial evaluation of several common probability distributions, it was determined that a  $\beta$  distribution provides an excellent fit to the (unsigned) bootstrapped statistic. This distribution has the appropriate skew and support on the interval  $(0, 1)$ . The red curves in Figure 6.1 are  $\beta$  distributions whose parameters  $(\alpha, \beta)$  were estimated from the moments of the sample data using the method of moments (MoM).

The moments of the  $\beta$  distribution are given by the following equations [76].

$$\bar{x} = \frac{\alpha}{\alpha + \beta} \quad (6.7)$$

$$s^2 = \frac{\alpha\beta}{(\alpha + \beta)^2(\alpha + \beta + 1)} \quad (6.8)$$

The parameters can be estimated by solving simultaneously for  $\alpha$  and  $\beta$ . First, solving Equation 6.7 for  $\beta$  in terms of  $\alpha$  and  $\bar{x}$ , we have:

$$\bar{x}\alpha + \bar{x}\beta = \alpha \quad (6.9)$$

$$\beta = \frac{(1 - \bar{x})\alpha}{\bar{x}} \quad (6.10)$$

$$= \frac{\alpha}{\bar{x}} - \alpha \quad (6.11)$$

Substituting the result into Equation 6.8 gives:

$$s^2 = \frac{\alpha^2(1 - \bar{x})}{\bar{x} \left(\frac{\alpha}{\bar{x}}\right)^2 \left(\frac{\alpha}{\bar{x}} + 1\right)} \quad (6.12)$$

$$= \frac{\bar{x}(1 - \bar{x})}{\frac{\alpha}{\bar{x}} + 1} \quad (6.13)$$




Figure 6.1: Bootstrapped test statistic with noise input

Solving for  $\alpha$  yields the MoM estimate:

$$\frac{\alpha}{\bar{x}} + 1 = \frac{\bar{x}(1 - \bar{x})}{s^2} \quad (6.14)$$

$$\alpha = \bar{x} \left( \frac{\bar{x}(1 - \bar{x})}{s^2} - 1 \right) \quad (6.15)$$

Substituting this result into Equation 6.10 yields the MoM estimator for  $\beta$

$$\beta = (1 - \bar{x}) \left( \frac{\bar{x}(1 - \bar{x})}{s^2} - 1 \right) \quad (6.16)$$

### Model Properties

Having performed various bootstraps, the parameters of a  $\beta$  distribution describing the null distribution of the test statistic are known for specific sample sizes and sequence lengths. However, to model the distribution for arbitrary test parameters requires knowledge of their underlying relationship with the distribution parameters. Further investigation revealed that

the moments of the distribution follow highly predictable trends in relation to the number of trials. Figure 6.2 shows scatter plots of the reciprocals of the mean squared,  $\mu^2$ , and the variance,  $\sigma^2$ , versus the sample size. These quotients clearly follow a linear relationship to sample size, as shown by the linear regression lines and the absence of scatter in the points. Figure 6.3 is a similar diagram for sequence lengths. While some relationship clearly exists, it is not perfectly linear, possibly due to the temporal spreading of the filters (which does not scale according to sequence length).

The previous findings informed the design of a model for the null distribution of the test statistic that would enable it to be subjected to hypothesis testing. Linear regression models were created for predicting  $1/\mu^2$  and  $1/\sigma^2$  from the sample size, for each available data point in terms of sequence length. This allows arbitrary samples sizes to be modeled, but only for specific sequence lengths. That restriction is removed by subsequent interpolation between the lines of regression. In other words, a linear model is used for trials and a piecewise linear model is used for lengths. Once the reciprocals of the moments have been found, Equations 6.15 and 6.16 are applied to find  $\alpha$  and  $\beta$ . The model was validated by comparing its predictions with additional bootstrapped samples.

## 6.3 Higher order statistics

The methods described so far allow statistically significant associations between neuronal activations in pairs of brain regions to be established. As previously discussed, such associations may depend on the influence of a driving variable, i.e. another brain region. Conventionally, partial correlation is used to examine the network structure of a system of variables such as this. However, the test statistic used here is not distributed equivalently to an ordinary correlation coefficient. Consequently, the usual partial correlation formula is not applicable.

An alternative method was devised to remove the influence of confounding variables. Consider the simple network below. There is no direct influence between X and Y, but the two will be correlated because of their common connection with Z.




Figure 6.2: Trends in null distribution of test statistic vs. sample size





Figure 6.3: Trends in null distribution of test statistic vs. sequence length



Once the latencies of the influences between  $X - Z$  and  $Y - Z$  have been established, the latency of the resulting correlation between  $X$  and  $Y$  can be predicted using the following formula:

$$\tau_{xy} = \tau_{xz} + \tau_{zy} \quad (6.17)$$

Here,  $\tau_{xy}$  denotes “the latency of the time series  $y$  with respect to  $x$ .” This relation can be used to remove the influence of  $Z$ , once it has been established, on the correlation between  $X$  and  $Y$ . First, the estimated latencies are used to temporally align the three sequences. Then the common component among the sequences is isolated using principal component analysis. Finally, the vector projections of  $X$  and  $Y$  in the direction of the common component are subtracted prior to a correlation calculation (with the original offsets having been restored).

The process has been exemplified using surrogate data similar to neurophysiological data. Figure 6.4 shows three time series ( $x$ ,  $y$  and  $z$ ) with their cross correlation sequences beneath. These sequences were produced by summing a sinusoid having randomised phase, white noise and a simulated “response”. The response signal consists of a second, frequency-distinguished sinusoid multiplied by an envelope function of the form:

$$f_{\text{env}}(t) = \begin{cases} \frac{t}{t^A + C} & t > 0 \\ 0 & \text{otherwise} \end{cases} \quad (6.18)$$

where  $A$  and  $C$  are constants that determine the shape of the simulated response. In this example,  $A = 2.3$  and  $C = 20$ .

Figure 6.5 shows the new sequences  $x'$  and  $y'$  with the common signal (shown in




Figure 6.4: Cross correlation with simulated data




Figure 6.5: Removal of confounding influence using PCA

green) removed. The cross correlation between the new sequences reveals the underlying relationship.


The concept was expanded to higher-order systems by testing for significant associations using first-order cross correlations, then identifying “driving” variables. A driving variable is characterised as one that has no prior influences. Once such a variable is found, its relationships are labelled as “safe”, and those among its partnering variables are all recalculated with the influence of the new-found driving variable removed. The known driving variable is excluded and the process is repeated until all drivers have been found. It should be acknowledged that the algorithm is unable to detect when circular relationships exist within a system.

## 6.4 Application

The techniques developed throughout this thesis were applied to practical MEG data. These experimental data were obtained from a recent study conducted by A. Asghar and A. Jansari. In their study, which has not been published at the time of writing, a passive listening paradigm was applied to examine the neural correlates of sound-colour synaesthesia. Using data from one of their control subjects, the novel three-stage analysis was applied with the aim of elucidating the cortical dynamics of auditory and speech processing in the normal brain.

### 6.4.1 Introduction

Various distinct regions in the cerebral cortex are known to be involved in processing auditory stimuli and in speech perception [5, 77]. These include temporal lobe areas such as Heschl's gyrus, parietotemporal regions such as Wernicke's area and frontal regions including Broca's area. There is not yet a consensus on the processing streams formed by interactions between the auditory cortex and multimodal integration centres in other cortical systems [78]. This analysis aims to test a hypothesis similar to recent assertions of other authors (eg. [78, 79]). The hypothesis adopted for this experiment is as follows. Auditory responses originating in the superior temporal gyrus (STG) of each hemisphere induce activity in the ipsilateral posterior superior temporal sulci (pSTS). These subsequently communicate with the anterior superior temporal sulcus (aSTS) and inferior frontal gyrus (IFG), predominantly in the left hemisphere. A direct path between the aSTS and the IFG might also exist independently of the indirect influence via the pSTS. The approximate locations of these areas are indicated in Figure 6.6 on the surface of the subject's cerebral cortex, which has been inflated using the methods of Fischl et al. [80] to reveal the interior of the sulci. As the seat of auditory perception, the primary auditory cortex (PAC) has been specifically isolated within the STG as a whole.



The points were identified by anatomical features in the non-inflated brain. Their locations are necessarily approximate as the underlying architectonics of the cortex are not evident from its convolutions due to individual variations in anatomy.

Figure 6.6: Approximate locations of auditory and speech areas on inflated brain

Wide range gamma-band activity from 30–100 Hz has been implicated in both auditory and language processing [81, 78] and recruitment of distant neural assemblies in task-related synchronisation [82, 9]. Hence the predicted pathways are expected to be expressed in this frequency range. Parts of the analysis are consequently focused on the same range, although the time-frequency decomposition is applied to the wideband MEG signal.

### 6.4.2 Methods

Auditory stimuli consisting of words and non-words were presented bilaterally to a normal, right-handed subject at 3-second intervals. During stimulation, the MEG signal was recorded at 678.17 Hz using a 248-channel whole head magnetometer. Words and non-words had durations of  $547 \pm 10$  ms and  $537 \pm 11$  ms respectively. The variance is due to differences in word length. There were 100 presentations of each type of stimulus in total, all producing epochs containing 500ms of prestimulus activity and 1000ms of activity following the stimulus onset. The experiment was a randomised block design which also included 100 sinusoidal tone bursts lasting 500ms. For the purposes of this analysis, words and non-words are treated as a single condition and the tones are ignored. Thus the epochs

under consideration are expected to include induced responses in speech perception areas that may not be activated by non-speech sounds. Of the 200 speech stimulations, 21 epochs were later rejected due to the presence of artefacts, leaving 179 epochs for analysis.


The data were subjected to an exploratory analysis in which some parameters were informed by the data. First, minimum norm estimates of regional activity were calculated from filtered and averaged data. This was used to select and fine-tune regions of interest beyond the approximate auditory areas indicated in Figure 6.6. Using the same preprocessing, time courses were estimated in the ROIs to determine a direction of interest, because current density is a vector quantity but a scalar time series is required for the second stage of analysis. Such time series were then estimated from the unfiltered, non-averaged data using the ROIs and directions already determined from the averaged data.

Next, a time-frequency decomposition was applied to each of the time series using the filter bank method described in Chapter 5. From this output, changes in amplitude were calculated with respect to the mean amplitude during the passive prestimulus period. The trial averages of these changes were visualised to determine the frequency bands in which to perform the correlation analysis.

Finally, the correlation analysis developed earlier in this chapter was used to predict patterns of connectivity between the regions of interest.

### **Stage 1: MNE**

The recorded MEG data were preprocessed by averaging across the time-aligned epochs and applying a wide gamma-band filter (30–100 Hz). Minimum norm estimates were calculated on the epoch average following the method described in §4.5.2. Tikhonov regularisation was applied, with the regularisation parameter set to 0.75% of the trace of the inverse operator. Largest singular value depth weighting was also used. Each point in the MNE grid was allocated three orthogonal components, forming a three-dimensional current density vector. The length of the vector was taken to visualise the magnitude of the activity. To clearly represent the detail within the sulci, the current density estimates are shown on the inflated



In the left hemisphere, A=aSTS, B=pSTS and C=pIFG. In the right, D=STG, E=aSTS and F=pSTS.

Figure 6.7: Location of chosen regions of interest

cortical surface models. The estimated activity is shown in Figure 6.8 at times where significant focal activity was observed close to hypothesised auditory and speech areas. In this case, MNE has resolved very modest gamma-band peaks in the PAC and pSTS of the left hemisphere. However, strong activations can be seen in the anterior STS and posterior IFG. In the right hemisphere, peaks are observed in the expected locations.

From these images, 3 regions of interest (ROIs) in each hemisphere were selected to coincide with the focus of activation, in preference to the points indicated in Figure 6.6. This adjustment makes allowances for the uncertain physical location of relevant areas and the potential for mis-localisation in MNE. The left PAC has not been selected due to the absence of resolved activity, considering also that activity in the right PAC should equally reflect the evoked response. The chosen regions are detailed in Figure 6.7. For the purpose of conducting a test of robustness upon the later correlation analysis, two alternative nodes were selected for the pIFG ROI, labelled  $C_1$  and  $C_2$ .

The direction of greatest variance over the post-stimulus period was calculated, using SVD, for the epoch average in each ROI. The component in this direction is a bipolar signal that retains valuable phase information (which will be necessary for the subsequent time-frequency decomposition). These components are shown in Figure 6.9. Two traces are

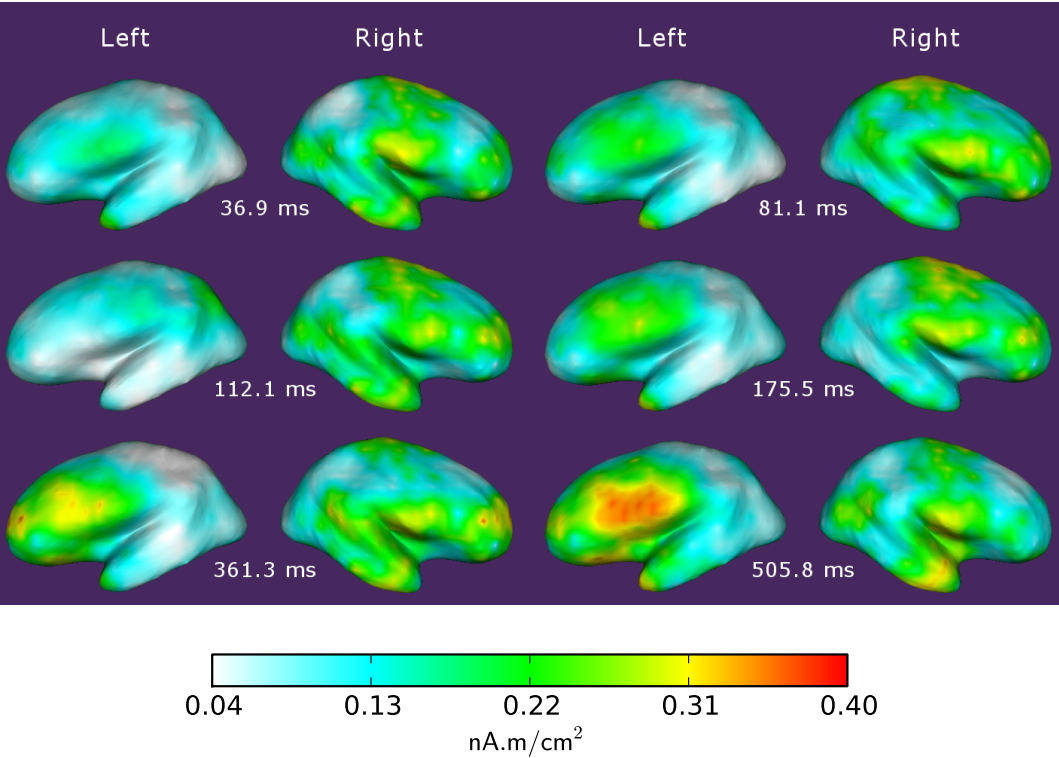


Figure 6.8: Minimum norm estimates from averaged data

shown under region C, which are the results from the alternative pIFG nodes. Oscillatory activity is evident in all of the ROIs. However, there is no obvious peak in activity in relation to the stimulus. The output from the two pIFG nodes is extremely close, as one would expect from the spatial smoothing inherent in MNE.

The correlation analysis requires individual trial data for the latency estimate and test statistic, since they are calculated across trials. Furthermore, it is not necessary to preprocess using a gamma-band filter because a subdivision of the gamma range can be selected from the filter bank output instead. In order to generate time courses for individual epochs, a submatrix of the previously calculated minimum norm inverse operator is formed by selecting only the rows corresponding to the regions of interest (three for each region). This new operator is then applied to the raw, unfiltered epochs, including the pre-stimulus period. Once again, a single component of the current density vector is taken for each region. However, the direction is now fixed so that results are comparable across trials. The component taken at each ROI is the one derived from the filtered and averaged data. To verify the results, the individual-trial time course estimates were averaged and filtered to ensuring they were consistent with those produced from the preprocessed data.

### **Stage 2: Time-Frequency Decomposition**

A filter bank having 16 bands was constructed using the default parameters from Table 5.1 (0.32-octave spacing). The filters were applied to the epoch data, then logarithmic amplitude envelopes were calculated. The mean log amplitude from each prestimulus period was subtracted from the envelope of the corresponding active period. This yields logarithmic changes in amplitude with respect to baseline levels. These time-frequency envelopes have been averaged over the epochs and visualised in Figures 6.10, 6.11, 6.12 and 6.13. Averaging after the spectrotemporal analysis accounts for signals that are not phase-locked to the stimulus and produces the most stable estimate of changes in energy density.

The time-frequency graphics reveal variations in activity taking place in numerous bands. Edge effects are visible at the extremes of the active period in the form of a cone of influence




Figure 6.9: Epoch average of gamma activity in regions of interest




Figure 6.10: Left hemisphere amplitude changes (A & B)




Figure 6.11: Left hemisphere amplitude changes (C<sub>1</sub> & C<sub>2</sub>)




Figure 6.12: Right hemisphere amplitude changes (D & E)




Figure 6.13: Right hemisphere amplitude changes (F)

spreading from the window edge. The activity in these areas appears to be below baseline levels, but this is because the signal is assumed to be zero outside the window.

### Stage 3: Correlation Analysis

As a simple illustration of the analysis that is possible with the output of the time-frequency decomposition, a correlation analysis of the type proposed earlier in the chapter was conducted separately in two adjacent frequency bands using a significance level of 0.05. The test candidates were the 10th-highest band (35.3 – 44.1 Hz) and the 9th-highest band (44.1 – 55.1 Hz). These bands are in the low-gamma range, and appear to show some consistencies between ROIs in the time-frequency plots. In both bands, the analysis was performed twice: once for the first pIFG node ( $C_1$ ), and again for the alternative node ( $C_2$ ). If the analysis is insensitive to small changes to its inputs, which could be erroneous in nature, it should produce similar results in both cases.

### 6.4.3 Results

The directed influences predicted by the correlation analysis are displayed in Figures 6.14 and 6.15. The graphs show high degrees of connectivity between the ROIs, with 8–10 edges being present out of a possible maximum of 15. The aSTS in the left hemisphere (region A) is strongly identified as driving the other ROIs. None of the other relationships is consistent among all cases. The pSTS (region F) in the right hemisphere is the only region not shown to be a driver in any of the cases.

### 6.4.4 Discussion

The results of the connectivity analysis do not support the hypothesis. Surprisingly, they imply a very highly connected network among the proposed auditory regions, with some unexpected directions of influence. For example, in the first 44.1–55.1 Hz analysis, the left hemisphere pIFG region is shown to drive the PAC of the right hemisphere. In reality, it is improbable that the networks found by the analysis accurately reflect the underlying




Figure 6.14: Interaction model during auditory stimulation (35.3 – 44.1 Hz)




Figure 6.15: Interaction model during auditory stimulation (44.1 – 55.1 Hz)

neuronal systems. There are several suspicious features of the data to suggest that the results may be spurious.

Firstly, region A (aSTC, left) is shown to drive all other regions of interest in every case, despite the fact it was only expected to drive region C (pIFG, left), and possibly region B (pSTC) if the proposed communication had a bidirectional element. At this point, it is notable that the algorithm was not designed to identify an absence of directivity. If the estimated latency is zero, it defaults to the ROI with the lower index as the “driver”, and the component is still removed from the driven variables, preserving the conditionality. The regions were indexed in alphabetical order, so it seems very likely that region A has been labelled as a driver too frequently due to this behaviour. It would be easy to establish this by reordering the regions to see if the topology of the network is affected.

Even if the algorithm preferentially awards driver status to lower-numbered ROI objects, there must also be a reason why the estimated latency was zero in the first place. It might be that the latency test is more likely to return a positive result at the point where there is maximal overlap between sequences. However, if there was no genuine connection between a variable pair, this should happen on average only 1 in 20 times if the significance level is 0.05. If there was indeed an association, it should outweigh those random events. A more plausible explanation is that the power in the input sequences is dominated by a d.c. component, which contains no latency information but is still taken into account as a general stimulus-related change in activity. To some extent this also explains the existence of correlations where they are not expected.

For a d.c. offset in the input data to have the stated effect, it would need to be consistent over many trials otherwise it would not affect the latency estimate. If the input signals used here are reconsidered, then a potential source of systematic error becomes apparent. The time-frequency decompositions assume the data to be zero outside the original sequence length. This results in a fall-off in energy towards the edges of the sequence, which encroaches further into the window in the lower frequency bands. The mechanism by which this leads to a d.c. offset in the amplitude change sequences relates to the calculation of the

baseline levels. The baselines in this example were calculated from a passive period that was half the length of the active period. Over the passive period, the proportion of energy lost to edge effects is more significant than in the active period, leading to the mean amplitude being underestimated. The effect is visible in the spectrotemporal amplitude change plots in Figures 6.10–6.13. The lowest bands all show strong overall increases in amplitude without exception. Although less apparent in the bands used in the test, the phenomenon could nonetheless be sufficient to disrupt its correct operation.

The possible incidence of false correlation at zero time lag could be further explained by signal leakage owing to the source estimation process. Because the source-space projection problem is underdetermined, the source estimates are known not to be independent. However, it was expected that the spatial discrimination ability of MNE over large cortical distances would avoid leakage between ROIs nonetheless. It is possible that this supposition was incorrect, and the resolved sources are highly correlated due to being composed of similar mixtures of underlying signals.

Another feature of the results that casts doubt over their validity is the lack of robustness in the analysis. Even though the signals from  $C_1$  and  $C_2$  appear almost identical in the averaged epoch data, significant changes were observed in the resulting network topology upon substitution.

In general, it appears that the hypothesis test is overly sensitive, producing significant test statistics too readily. In that case, it may be that the null model is inadequate, or being applied inappropriately. It has not been established that the null model continues to apply after signal removal using the PCA method. One theory is that the signal removal is too effective, creating artificial reductions in activity in the target variables, that are then regarded as correlated. It is also possible that the method of moments estimates of the distribution parameters are not sufficiently accurate. Possible tests and remedies for the problems described here are to be found in the further work section of the final chapter.

## 6.5 Summary

Statistical inference is concerned with hypothesis testing. A statistical hypothesis is a statement about a test statistic, which is a metric calculated from sample data. For the purposes of testing hypotheses about connectivity from filter-bank representations of cortical source signals, a test statistic was introduced. The statistic quantifies the similarity between two filter-bank envelopes, each representing changes in neural activity in a region of interest. It aims to do so regardless of any time lag in the relationship between them by estimating this first and making a differential time adjustment.

To conduct a hypothesis test using a test statistic, a model is needed for the distribution of the test statistic under the null hypothesis. In the absence of a theoretical model for the null distribution of our test statistic, the technique of bootstrapping was used to estimate it. The source data was filter-bank envelope changes calculated from uncorrelated white noise. A beta distribution was found to be an appropriate model for the null distribution. Beta distribution parameters were estimated for a range of sample sizes and sequence lengths using method of moments estimators on the bootstrapped data. A 2-dimensional linear model was created to facilitate interpolation between the known points in the parameter space.

To extend the hypothesis test beyond the bivariate case, a means for removing the effect of known influences from the amplitude envelope data was developed using principal component analysis. The automated removal of such an influencing signal was demonstrated. An iterative algorithm was then written to resolve all of the relationships among a group of random processes.

All stages of the proposed MEG analysis framework were combined to analyse MEG data obtained from a simple auditory experiment. First, activations were visualised using MNE, allowing the investigator to refine regions of interest. Subsequently, time course estimates were calculated for those ROIs. These time courses were decomposed into time-frequency representations, which were used to produce visualisations of neural activity in the time-frequency plane. The connectivity analysis was then applied to two bands of the

time-frequency decomposition. The results were inconclusive due to apparent failures in the connectivity analysis. Several possible modes of failure have been put forward.

## 7 Summary and Further Work

### 7.1 Summary

#### 7.1.1 Aims

The primary aim of the work presented in this thesis was to engineer a set of tools that could allow inferences to be made from magnetoencephalographic data about dynamic connectivity in the cerebral cortex.

#### 7.1.2 MEG in the context of Neuroimaging

The main contemporaries of MEG in the field of neuroimaging are fMRI, PET and EEG. EEG is low-cost, offers good temporal resolution and is widely used in clinical diagnosis, but has a weak ability to determine the spatial distribution of the signal's origins. The use of PET in cognitive neuroscience research is limited by the requirement to administer radiopharmaceuticals, which carries a risk to the subject. It is best for studying metabolic processes because specific compounds can be labelled and their passage through tissue traced. The time resolution of PET is on the order of a minute.

The most widely used modality in research is probably fMRI. This is modestly invasive and produces highly accurate and precise spatial information, but only shows haemodynamic activity, limiting time resolution to a second at best. MEG is the least invasive and offers the best time resolution while coming close to fMRI for spatial resolution of cortical sources. It is very well placed for studying cortical dynamics.

### 7.1.3 MEG Analysis

Traditional approaches to MEG analysis typically centre around source localisation. While this type of MEG analysis is relatively mature, it is only more recently that the potential for MEG to investigate cortical dynamics has begun to be realised. An engineered approach to connectivity analysis is proposed, which consists of using an existing source estimation process to extract signal estimates from a small number of regions of interest in the cerebral cortex. This is to be followed by a time-frequency analysis that subdivides neurophysiological frequency bands. Finally, the spectrotemporal representation shall be subjected to statistical analysis to determine the associations between regions of interest.

Three categories of MEG source estimation were considered for deriving the signals in regions of interest. The classical method of fitting an equivalent current dipole (ECD) to the data is a rudimentary approach capable only of resolving focal sources. It was mainly discussed as the basis for more advanced linear transformation methods such as minimum norm estimation (MNE) and beamforming. Beamforming was rejected for the proposed framework because of its reliance on the covariance matrix for signal separation. This was thought to jeopardise the ability of subsequent processes to extract meaningful information about the statistical relationships between the signal estimates. MNE was found to have poorer spatial resolution than beamforming, but this was considered to be of no importance in this application.

### 7.1.4 Time-Frequency Decomposition

The common practice of preprocessing MEG data using filters prior to analysis requires pre-supposition of relevant frequency bands. This can be eliminated, and spectral progression can be studied, by using time-frequency analysis.

To suit the requirements of MEG signal interpretation, common time-frequency decompositions were considered, including the short-time Fourier transform (STFT) and wavelet methods. The STFT was criticised for being centred around linear frequency divisions, which are unequally distributed between neurological frequency bands. It

also imposes a constant time resolution, which sacrifices temporal detail for unnecessary high-frequency spectral resolution. Wavelets were considered as a logarithmic frequency orientated alternative. However, continuous wavelet analysis was found to have problems with data redundancy and computational burden for the sizes of data sets typical of MEG. The frequency divisions used in discrete wavelet analysis were considered too wide for practical use.

Two novel wavelet-related methods were considered in place of existing methods. The first was a time-domain approach consisting of sparse convolution of the signal with Hann-windowed complex waves at different frequencies. However, the bands were found to integrate poorly, leading to spectral bias. Instead, a minimally-redundant high resolution representation was achieved by cascading identical low-pass FIR filters and sub-sampling the output at each stage. This was adopted for the proposed framework.

### 7.1.5 Connectivity Analysis

A connectivity analysis integrated with the filter-bank decomposition was developed. A test was devised for the hypothesis that two baseline-adjusted filter-bank envelopes are more closely related than can happen by chance. This was achieved by devising a test statistic that measures the similarity between envelopes, then modelling its distribution when presented with unrelated signals. The test statistic accounts for short time delays (less than the sequence length) in the relationship. It does so by estimating the latency from cross-correlation sequences calculated on large numbers of trials. The mean of the sequences at the estimated point then gives a measure of the overall similarity at the most probable time offset.

The bivariate hypothesis test was incorporated into an algorithm that attempts to test for conditionally independent relationships in a multivariate scenario by removing the influence of confounding variables using PCA.

The framework, incorporating all three stages, was applied to MEG data. However, the results were inconclusive, ostensibly due to the failure of the connectivity analysis.

## 7.2 Further Work

### 7.2.1 Source Signal Estimation

Signal leakage between sources that is inherent in MNE has been implicated in the problems encountered in the connectivity analysis. Beamforming was originally rejected due to the assumption of statistical independence between sources and the use of second-order statistics in constructing the filters. However, it has not been established empirically that this is a problem in practice. Further analysis could be carried out using beamforming to isolate the source signals.

### 7.2.2 Time-frequency Decomposition

The edge effects caused by the assumptions in the filter-bank method are postulated to have contributed to the failure of the connectivity analysis. As such, the filter-bank could be modified to resolve the problem. The effects are caused by assuming the signal to be zero outside the input sequence. At low frequencies, the symmetrical impulse response of the filter straddles the edge of the data and is not energised by any signal before or after the input sequence, leading to a fall in output approaching the window edges. Furthermore, if the signal amplitude is large at the window edges, the filter is stimulated into ringing by the discontinuity that this causes. Both effects are undesirable, but by assuming the signal to be periodic and to have even symmetry, both could be eliminated.

Connectivity analysis is certainly not the only application for a time-frequency decomposition of MEG data. Wavelet denoising has been applied to improving the SNR of MEG data [83]. This involves thresholding signals in the time-frequency domain so as to reject noise. The filter-bank could be used in a similar manner, and would be likely to benefit this application with its fine frequency resolution.

### 7.2.3 Connectivity Analysis

When the connectivity analysis was applied, ROIs appearing early in the list provided to the algorithm appeared to be favoured as driving variables. This could be established for certain by rearranging ROIs and reapplying the analysis. If it is found to be true, there are some modifications that might eliminate the problem. Firstly, the d.c. offset could be removed, ignored or deweighted. In any case, it contains no information about connectivity beyond grouping ROIs according to whether or not they exhibit overall changes in activity during the active period. The cause of the phenomenon is probably related to the default attribution of “driver” status under the condition of zero estimated latency. An obvious improvement would be to attribute those relationships as undirected. The signal removal could then be applied to all involved ROIs after the round of testing is complete (hypothesis tests are carried out in a round-robin fashion by the algorithm). This contrasts with the current behaviour of only removing the signal from “driven” variables.

The empirical null distribution model should be re-examined. In particular, the null model needs to be verified in the case that a correlated signal has been removed from two conditionally independent time series. If the PCA signal removal technique cannot perfectly remove the signal, the hypothesis test may produce false-positive results because the existing null model is being misapplied. Another possibility is that the model is already flawed. Bootstrapping random segments of MEG data instead of white noise may produce a different result. Also, the method of moments estimators of the beta distribution parameters are quite simplistic. Maximum likelihood estimators are widely used [84] for this purpose. These are considered to be more efficient estimators, but their calculation is more difficult.

## A Spherical Head Model

Current densities in MEG can be modelled by dipolar sources within a homogeneous conducting sphere, which represents the head. With this geometry, a closed-form analytical solution to the forward problem can be found. The following derivation is essentially reproduced from [25]. To begin with, we have Maxwell's equations, relating the electric field,  $\mathbf{E}$ ; the magnetic field,  $\mathbf{B}$ ; the current density  $\mathbf{J}$ ; the charge density,  $\rho$ ; the permittivity of free space,  $\epsilon_0$ ; and the permeability of free space,  $\mu_0$ :

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon_0} \quad (\text{A.1})$$

$$\nabla \cdot \mathbf{B} = 0 \quad (\text{A.2})$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \quad (\text{A.3})$$

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \mu_0 \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} \quad (\text{A.4})$$

However, it can be shown that the quasi-static approximations of Maxwell's equations are valid in the case of MEG measurements. That is, the derivative terms in equations A.3 and A.4 can be approximated to zero. In the magnetostatic case, the two equations become

$$\nabla \times \mathbf{E} = 0 \quad (\text{A.5})$$

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} \quad (\text{A.6})$$

Equation A.5 implies that  $\mathbf{E}$  can be represented by a scalar potential,  $V$ , such that

$$\mathbf{E} = -\nabla V \quad (\text{A.7})$$

The total current density is given by

$$\mathbf{J} = \mathbf{J}_p + \sigma \mathbf{E} \quad (\text{A.8})$$

where  $\mathbf{J}_p$  is the source current density (the so-called *primary current*), and  $\sigma\mathbf{E}$  the ohmic current in a volume conductor with conductivity  $\sigma$  (the *volume currents*). For an **unbounded** volume conductor  $G$ , the Biot-Savart law gives the magnetic field due to the current density at  $\mathbf{r}'$  as a function of  $\mathbf{r}$  (the position at which  $\mathbf{B}$  is to be determined):

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi} \int_G \mathbf{J}(\mathbf{r}') \times \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} d\nu' \quad (\text{A.9})$$

Because equation A.9 is a function of  $\mathbf{J} = \mathbf{J}_p - \sigma\nabla V$  (from equations A.8 & A.7),  $V$  must still be found in order to evaluate the integral. The divergence of a curl is always zero, i.e.  $\nabla \cdot \nabla \times \mathbf{B} \equiv 0$ , which together with equation A.6 leads to  $\nabla \cdot \mathbf{J} = 0$ . Thus,  $\nabla \cdot (\sigma\nabla V) = \nabla \cdot \mathbf{J}_p$ , and in a homogeneous conductor,

$$\nabla^2 V = \frac{\nabla \cdot \mathbf{J}_p}{\sigma} \quad (\text{A.10})$$

This is the equation to be solved for  $V$ , which requires numerical solution in the general case. However, the two identities

$$\nabla' \left( \frac{1}{|\mathbf{r} - \mathbf{r}'|} \right) \equiv \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} \quad (\text{A.11})$$

and

$$\nabla' \times \left( \frac{\mathbf{J}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} \right) = \frac{\nabla' \times \mathbf{J}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} - \mathbf{J}(\mathbf{r}') \times \nabla' \left( \frac{1}{|\mathbf{r} - \mathbf{r}'|} \right) \quad (\text{A.12})$$

can be applied to the integral in equation A.9 to give, using Stokes's Theorem,

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi} \int_G \frac{\nabla' \times \mathbf{J}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\nu' \quad (\text{A.13})$$

Now, remembering that  $\mathbf{J} = \mathbf{J}_p - \sigma\nabla V$ , and that the curl of a gradient is always zero, the term in  $V$  disappears from the numerator of the integrand if  $\sigma$  is constant:

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi} \int_G \frac{\nabla' \times \mathbf{J}_p}{|\mathbf{r} - \mathbf{r}'|} d\nu' \quad (\text{A.14})$$

Reversing the previous steps results in the expression for the magnetic field due to  $\mathbf{J}$  in homogeneous space, from now on denoted  $\mathbf{B}_0$ :

$$\mathbf{B}_0(\mathbf{r}) = \frac{\mu_0}{4\pi} \int_G \mathbf{J}_p(\mathbf{r}') \times \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} d\nu' \quad (\text{A.15})$$

Next consider an arbitrary homogeneous volume conductor,  $G_S$ , bounded by a surface  $S$ , with conductivity  $\sigma$ . Outside the conductor,  $\sigma = 0$ . From equations A.7, A.8 and A.9,

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi} \int_{G_S} (\mathbf{J}_p(\mathbf{r}') - \sigma(\mathbf{r}')\nabla V(\mathbf{r}')) \times \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} d\nu' \quad (\text{A.16})$$

This can also be written in terms of  $\mathbf{B}_0$ :

$$\mathbf{B}(\mathbf{r}) = \mathbf{B}_0(\mathbf{r}) - \frac{\mu_0\sigma}{4\pi} \int_{G_S} \nabla V(\mathbf{r}') \times \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} d\nu' \quad (\text{A.17})$$

Then, recalling equation A.11, the identity

$$\nabla V \times \nabla \left( \frac{1}{|\mathbf{r} - \mathbf{r}'|} \right) \equiv \nabla \times \left( V \nabla \frac{1}{|\mathbf{r} - \mathbf{r}'|} \right) \quad (\text{A.18})$$

allows the integral term in equation A.17 to be transformed to a surface integral by Stokes's Theorem:

$$\int_{G_S} \nabla V(\mathbf{r}') \times \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} d\nu' = \int_S V(\mathbf{r}') \mathbf{n}(\mathbf{r}') \times \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} dS$$

where  $\mathbf{n}(\mathbf{r}')$  is the unit vector normal to  $S$ , pointing out of the conductor. So the contribution to  $\mathbf{B}$  from volume currents is equivalent to the field due to a current distribution over the bounding surface.

$$\mathbf{B}(\mathbf{r}) = \mathbf{B}_0(\mathbf{r}) - \frac{\mu_0\sigma}{4\pi} \int_S V(\mathbf{r}') \mathbf{n}(\mathbf{r}') \times \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} dS \quad (\text{A.19})$$

The previous result still requires calculation of the surface potential. Consider now the case of a spherical conductor, centred at the origin. Where  $\hat{\mathbf{r}}$  is the unit vector in the direction of  $\mathbf{r}$ , the radial component of the magnetic field,  $B_r(\mathbf{r})$  is given by

$$B_r(\mathbf{r}) = \mathbf{B}_0(\mathbf{r}) \cdot \hat{\mathbf{r}} - \frac{\mu_0\sigma}{4\pi} \int_S V(\mathbf{r}') \mathbf{n}(\mathbf{r}') \times \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} \cdot \hat{\mathbf{r}} d\nu' \quad (\text{A.20})$$

But since  $\mathbf{n}(\mathbf{r}') = \hat{\mathbf{r}}'$  and  $\mathbf{r} - \mathbf{r}'$  necessarily lies in the same plane as  $\mathbf{r}$  and  $\mathbf{r}'$ , the scalar triple product above equates to zero and the volume currents disappear from the equation, leaving

$$B_r(\mathbf{r}) = \mathbf{B}_0(\mathbf{r}) \cdot \hat{\mathbf{r}} = \frac{\mu_0}{4\pi} \int_{G_S} \mathbf{J}_p(\mathbf{r}') \times \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} \cdot \hat{\mathbf{r}} d\nu' \quad (\text{A.21})$$

This result is useful because it allows an expression for  $\mathbf{B}$  outside  $G_S$  to be found. There are no volume currents outside  $G_S$  because  $\sigma = 0$ ; thus  $\mathbf{J} = 0$  also.  $\mathbf{B}$  can therefore be represented as a scalar potential,  $U$ :

$$\mathbf{B}(\mathbf{r}) = -\mu_0 \nabla U(\mathbf{r}) \quad (\text{A.22})$$

Now  $U$  can be found by taking the line integral of  $\nabla U$  along the radius, as follows:

$$\begin{aligned} U(\mathbf{r}) &= - \int_0^\infty \nabla U(\mathbf{r} + t\hat{\mathbf{r}}) \cdot \hat{\mathbf{r}} dt \\ &= \frac{1}{\mu_0} \int_0^\infty B_r(\mathbf{r} + t\hat{\mathbf{r}}) dt \end{aligned} \quad (\text{A.23})$$

For a current dipole of infinitesimal length, located at  $\mathbf{r}_0$ ,

$$\mathbf{J}_p = \delta(\mathbf{r} - \mathbf{r}_0) \mathbf{Q}$$

where  $\mathbf{Q}$  is the dipole moment and  $\delta$  is the Dirac delta. Substituting this expression into equation A.21, and calculating the integral gives

$$B_r(\mathbf{r}) = \frac{\mu_0}{4\pi} \mathbf{Q} \times \frac{\mathbf{r} - \mathbf{r}_0}{|\mathbf{r} - \mathbf{r}_0|^3} \cdot \hat{\mathbf{r}} \quad (\text{A.24})$$

Using this result and equation A.23 leads to the following expression for  $U$  due to a current dipole at  $\mathbf{r}_0$ :

$$\begin{aligned} U(\mathbf{r}) &= \frac{1}{4\pi} \mathbf{Q} \times (\mathbf{r} - \mathbf{r}_0) \cdot \hat{\mathbf{r}} \int_0^\infty \frac{1}{|\mathbf{r} + t\hat{\mathbf{r}} - \mathbf{r}_0|^3} dt \\ &= -\frac{1}{4\pi} \frac{\mathbf{Q} \times \mathbf{r}_0 \cdot \mathbf{r}}{a(ra + r^2 - \mathbf{r}_0 \cdot \mathbf{r})} \end{aligned} \quad (\text{A.25})$$

where  $\mathbf{a} = \mathbf{r} - \mathbf{r}_0$ ,  $a = |\mathbf{a}|$  and  $r = |\mathbf{r}|$ . Finally, let  $F = a(ra + r^2 - \mathbf{r}_0 \cdot \mathbf{r})$ , then using equation A.22,

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi F^2} (F \mathbf{Q} \times \mathbf{r}_0 - \mathbf{Q} \times \mathbf{r}_0 \cdot \mathbf{r} \nabla F) \quad (\text{A.26})$$

$$\text{N.B. } \nabla F = \left( \frac{a^2}{r} + \frac{\mathbf{a} \cdot \mathbf{r}}{a} + 2a + 2r \right) \mathbf{r} - \left( a + 2r + \frac{\mathbf{a} \cdot \mathbf{r}}{a} \right) \mathbf{r}_0$$

This is the forward equation for the magnetic field measured outside an isotropic, homogeneous spherical conductor due to a dipolar current source within the conductor.

As can be seen from the above results, if the dipole is radial ( $\mathbf{Q} = \lambda \hat{\mathbf{r}}$ , where  $\lambda$  is some constant), the field outside the sphere is zero (since  $\lambda \hat{\mathbf{r}} \times \mathbf{r}_0 \cdot \mathbf{r} = 0$ ).

When fitting a dipole to the data, rather than re-evaluating the expression in equation A.26 on every iteration, it makes sense to look for a way of simplifying the calculation. Lead field analysis does just this by observing that the scalar triple product in equation A.21 produces a linear relationship between the current density and the radial component of the magnetic field [43]. Thus the equation can be rewritten as a scalar product between the current density, and a function  $\mathbf{L}(\mathbf{r}, \mathbf{r}')$ , called the lead field, which is independent of the current density:

$$B_r(\mathbf{r}) = \int_{G_S} \mathbf{L}(\mathbf{r}, \mathbf{r}') \cdot \mathbf{J}_p(\mathbf{r}') d\nu' \quad (\text{A.27})$$

$$\text{where } \mathbf{L}(\mathbf{r}, \mathbf{r}') = \frac{\mu_0}{4\pi r} \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3}$$

$\mathbf{L}(\mathbf{r}, \mathbf{r}')$  represents the sensitivity of a sensor at  $\mathbf{r}$  to a source at  $\mathbf{r}'$ . In a practical MEG system, there are fixed, discrete sensor locations. A grid of possible source locations can also be defined, which allows  $\mathbf{L}(\mathbf{r}, \mathbf{r}')$  to be precalculated for every possible location. Thereafter, simple arithmetic involving the components of the dipole moment and the lead field at the dipole's location will give the predicted set of measurements. Linear interpolation can be used to estimate the contribution of sources at arbitrary locations [43].

## Abbreviations

**aSTS** Anterior superior temporal sulcus

**BA17** Brodmann area 17

**BEM** Boundary element model

**CNS** Central nervous system

**CT** Computed tomography

**CWT** Continuous wavelet transform

**d.c.** Direct current

**DCM** Dynamic causal modelling

**DFT** Discrete Fourier transform

**DV3D** DataViewer 3D (software)

**DWT** Discrete wavelet transform

**ECD** Equivalent current dipole

**ECoG** Electrocorticography

**EEG** Electroencephalography

**FEM** Finite element model

**FFT** Fast Fourier transform

**FIR** Finite impulse response

**fMRI** Functional magnetic resonance imaging

**FWT** Fast wavelet transform

**IFG** Inferior frontal gyrus

**MCG** Magnetocardiography

**MEG** Magnetoencephalography

**MNE** Minimum norm estimation

**MoM** Method of moments

**MRI** Magnetic resonance imaging

**PAC** Primary auditory cortex

**PCA** Principal component analysis

**PCM** Pulse code modulation

**PET** Positron emission tomography

**PMCC** Product-moment correlation coefficient

**PSD** Power spectral density

**pSTS** Posterior superior temporal sulcus

**r.f.** Radio frequency

**RMS** Root mean square

**ROI** Region of interest

**s.d.** Standard deviation

**SFG** Superior frontal gyrus

**SNR** Signal-to-noise ratio

**SQUID** Superconducting quantum interference device

**STFT** Short-time Fourier transform

**STG** Superior temporal gyrus

**SVD** Singular value decomposition

**VTK** Visualization Toolkit (software)

**YNiC** York Neuroimaging Centre

## References

- [1] J. Wikswo, “SQUID magnetometers for biomagnetism and nondestructive testing — important questions and initial answers,” *IEEE Transactions On Applied Superconductivity*, vol. 5, pp. 74 – 120, June 1995. 1994 Applied Superconductivity Conference, BOSTON, MA, OCT 16-21, 1994.
- [2] F. Rosenow and H. Lüders, “Presurgical evaluation of epilepsy,” *Brain*, vol. 124, no. 9, pp. 1683 – 1700, 2001.
- [3] 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,” *Reviews of Modern Physics*, vol. 65, pp. 413 – 497, April 1993.
- [4] D. Jones, *Neurons and Synapses*. No. 135 in Studies in Biology, London: Edward Arnold, 1981.
- [5] M. S. Gazzaniga, R. B. Ivry, and G. R. Mangun, *Cognitive Neuroscience: The Biology of the Mind*. New York: W.W. Norton, 2nd ed., 2002.
- [6] Y. Okada, “Toward understanding the physiological origins of neuromagnetic signals,” in *Magnetic Source Imaging of the Human Brain*, Erlbaum, 2003.
- [7] M. Hämäläinen and R. Hari, “Magnetoencephalographic (MEG) characterisation of dynamic brain activation: Basic principals of data collection and source analysis,” in *Brain Mapping: The Methods (2nd Ed.)*, Academic Press, 2002.

- [8] J. Malmivuo and R. Plonsey, *Bioelectromagnetism - Principals and Applications of Bioelectric and Biomagnetic Fields*. Oxford University Press, 1995.
- [9] E. Niedermeyer and F. Lopes Da Silva, *Electroencephalography: Basic Principals, Clinical Applications, and Related Fields*. Philadelphia: Williams & Wilkins, 1993.
- [10] H. W. Buckingham, "The Marc Dax (1770 – 1837)/Paul Broca (1824 – 1880) controversy over priority in science: Left hemisphere specificity for seat of articulate language and for lesions that cause aphemia," *Clinical Linguistics & Phonetics*, vol. 20, pp. 613 – 619, Sept – Oct 2006.
- [11] G. Baule and R. McFee, "Detection of the magnetic field of the heart," *American Heart Journal*, vol. 66, no. 1, pp. 95 – 96, 1963.
- [12] D. Cohen, "Boston and the history of biomagnetism," *Neurology and Clinical Neurophysiology*, vol. 114, pp. 1 – 4, 2004.
- [13] D. Cohen, "Magnetoencephalography: Detection of the brain's electrical activity with a superconducting magnetometer," *Science*, vol. 175, no. 4022, pp. 664 – 666, 1972.
- [14] L. Kaufman and Z.-L. Lu, "Basics of neuromagnetism and magnetic source imaging," in *Magnetic Source Imaging of the Human Brain*, Erlbaum, 2003.
- [15] 4-D Neuroimaging, *Magnes 3600WH Product Description*, 2003.
- [16] G. Buzsáki and A. Draguhn, "Neuronal oscillations in cortical networks," *Science*, vol. 304, pp. 1926 – 1929, June 2004.
- [17] M. E. Phelps, "Positron emission tomography provides molecular imaging of biological processes," *Proceedings of the National Academy of Sciences*, vol. 97, pp. 9226 – 9233, August 2000.
- [18] D. T. Edmonds, *Electricity and Magnetism in Biological Systems*. Oxford University Press, 2001.

- [19] A. Ioannides, J. Bolton, and C. Clarke, "Continuous probabilistic solutions to the biomagnetic inverse problem," *Inverse Problems*, vol. 6, pp. 523 – 542, 1990.
- [20] O. Hauk, "Keep it simple: a case for using classical minimum norm estimation in the analysis of EEG and MEG data," *NeuroImage*, vol. 21, pp. 1612 – 1621, April 2004.
- [21] M.-X. Huang, J. Shih, R. Lee, D. L. Harrington, R. Thoma, M. Weisend, F. Hanlon, K. Paulson, T. Li, K. Martin, G. Miller, and J. Canive, "Commonalities and differences among vectorized beamformers in electromagnetic source imaging," *Brain Topography*, vol. 16, no. 3, pp. 139 – 158, 2004.
- [22] B. Van Veen, W. Van Drongelen, M. Yuchtman, and A. Suzuki, "Localization of brain electrical activity via linearly constrained minimum variance spatial filtering," *IEEE Transactions on Biomedical Engineering*, vol. 44, pp. 867 – 880, September 1997.
- [23] R. Van Uiter, D. Weinstein, and C. Johnson, "Volume currents in forward and inverse magnetoencephalographic simulations using realistic head models," *Annals of Biomedical Engineering*, vol. 31, pp. 21 – 31, 2003.
- [24] P. Schimpf, J. Haueisen, C. Ramon, and H. Nowak, "Realistic computer modelling of electric and magnetic fields of human head and torso," *Parallel Computing*, vol. 24, pp. 1433 – 1460, 1998.
- [25] J. Sarvas, "Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem," *Physics in Medicine and Biology*, vol. 32, no. 1, pp. 11 – 22, 1987.
- [26] N. G. Gençer, C. E. Acar, and I. O. Tanzer, "Forward problem solution of magnetic source imaging," in *Magnetic Source Imaging of the Human Brain*, Erlbaum, 2003.
- [27] M. X. Huang, J. C. Mosher, and R. M. Leahy, "A sensor-weighted, overlapping-sphere head model and exhaustive head model comparison for MEG," *Physics in Medicine and Biology*, vol. 44, pp. 423 – 440, 1999.

- [28] R. Vigario, J. Sarela, V. Jousmaki, M. Hamalainen, and E. Oja, "Independent component approach to the analysis of EEG and MEG recordings," *IEEE Transactions on Biomedical Engineering*, vol. 47, pp. 589 – 593, May 2000.
- [29] A. Hodgkin and A. Huxley, "A quantitative description of membrane current and its application to conduction and excitation in nerve," *Journal of Physiology*, vol. 117, pp. 500 – 544, 1952.
- [30] R. Traub, D. Contreras, M. Cunningham, H. Murray, F. LeBeau, A. Roopun, A. Bibbig, W. Wilent, M. Higley, and M. Whittington, "Single-column thalamocortical network model exhibiting gamma oscillations, sleep spindles, and epileptogenic bursts," *Journal of Neurophysiology*, vol. 93, pp. 2194 – 2232, April 2005.
- [31] O. David, L. Harrison, and K. Friston, "Modelling event-related responses in the brain," *NeuroImage*, vol. 25, pp. 756 – 770, 2005.
- [32] O. David and K. Friston, "A neural mass model for MEG/EEG: coupling and neuronal dynamics," *NeuroImage*, vol. 20, pp. 1743 – 1755, 2003.
- [33] K. Friston, L. Harrison, and W. Penny, "Dynamic causal modelling," *NeuroImage*, vol. 19, pp. 1273 – 1302, 2003.
- [34] C.-H. Im, H.-K. Jung, H. Kwon, and Y.-H. Lee, "Multiresolution reconstruction of magnetoencephalography source distribution," *IEEE Transactions On Magnetics*, vol. 20, pp. 1100 – 1103, March 2004.
- [35] O. Hauk, A. Keil, T. Elbert, and M. M. Müller, "Comparison of data transformation procedures to enhance topographical accuracy in time-frequency analysis of the human EEG," *Journal of Neuroscience Methods*, vol. 113, pp. 111 – 122, 2002.
- [36] O. David and L. Garnero, "Time-coherent expansion of MEG/EEG cortical sources," *NeuroImage*, vol. 17, pp. 1277–1289, 2002.

- [37] M. Scherg and J. Ebersole, "Models of brain sources," *Brain Topography*, vol. 5, no. 4, pp. 419 – 423, 1993.
- [38] M. Huang, C. J. Aine, S. Supek, E. Best, D. Ranken, and E. R. Flynn, "Multi-start downhill simplex method for spatio-temporal source localization in magnetoencephalography," *Electroencephalography and clinical Neurophysiology*, vol. 108, pp. 32 – 44, 1998.
- [39] A. Dale and M. Sereno, "Improved localization of cortical activity by combining EEG and MEG with MRI cortical surface reconstruction — a linear-approach," *Journal of Cognitive Neuroscience*, vol. 5, pp. 162 – 176, Spring 1993.
- [40] J.-Z. Wang, "Minimum-norm least-squares estimation: magnetic source images for a spherical model head," *IEEE Transactions on Biomedical Engineering*, vol. 40, pp. 387 – 396, April 1993.
- [41] M. Hämäläinen and R. Ilmoniemi, "Interpreting magnetic fields of the brain: minimum norm estimates," *Medical and Biological Engineering and Computing*, vol. 32, pp. 35 – 42, 1994.
- [42] R. Penrose, "A generalized inverse for matrices," *Mathematical Proceedings of the Cambridge Philosophical Society*, vol. 51, no. 03, pp. 406 – 413, 1955.
- [43] J.-Z. Wang and L. Kaufman, "Magnetic source imaging: Search for inverse solutions," in *Magnetic Source Imaging of the Human Brain*, Erlbaum, 2003.
- [44] M. Bertero, C. De Mol, and E. Pike, "Linear inverse problems with discrete data: II. Stability and regularisation," *Inverse Problems*, vol. 4, pp. 573 – 594, 1988.
- [45] P. Hansen, "Regularization Tools: A Matlab package for analysis and solution of discrete ill-posed problems," *Numerical Algorithms*, vol. 6, no. 1, pp. 1 – 35, 1994.
- [46] P. Hansen and D. O'Leary, "The use of the L-curve in the regularization of discrete

- ill-posed problems,” *SIAM Journal on Scientific Computing*, vol. 14, pp. 1487 – 1503, November 1993.
- [47] S. Moratti and A. Keil, “Cortical activation during Pavlovian fear conditioning depends on heart rate response patterns: An MEG study,” *Cognitive Brain Research*, vol. 25, no. 2, pp. 459 – 471, 2005.
- [48] B. Yvert, C. Fischer, O. Bertrand, and J. Pernier, “Localization of human supratemporal auditory areas from intracerebral auditory evoked potentials using distributed source models,” *NeuroImage*, vol. 28, pp. 140–153, October 2005.
- [49] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery, *Numerical Recipes in C++: The art of scientific computing*. Cambridge University Press, 3rd ed., 2002.
- [50] F.-H. Lin, T. Witzel, M. S. Hämäläinen, A. M. Dale, J. W. Belliveau, and S. M. Stufflebeam, “Spectral spatiotemporal imaging of cortical oscillations and interactions in the human brain,” *NeuroImage*, vol. 23, pp. 582 – 595, 2004.
- [51] F.-H. Lin, T. Witzel, S. Ahlfors, S. M. Stufflebeam, J. W. Belliveau, and M. S. Hämäläinen, “Assessing and improving the spatial accuracy in MEG source localization by depth-weighted minimum-norm estimates,” *NeuroImage*, vol. 31, pp. 160 – 170, May 2006.
- [52] E. Palmero-Soler, K. Dolan, V. Hadamschek, and P. A. Tass, “swLORETA: a novel approach to robust source localization and synchronization tomography,” *Physics in Medicine and Biology*, vol. 52, no. 7, pp. 1783 – 1800, 2007.
- [53] M. Fuchs, M. Wagner, T. Kohler, and H. Wischmann, “Linear and nonlinear current density reconstructions,” *Journal of Clinical Neurophysiology*, vol. 16, pp. 267 – 295, May 1999.
- [54] F. Ségonne, A. M. Dale, E. Busa, M. Glessner, D. Salat, H. K. Hahn, and B. Fischl, “A hybrid approach to the skull stripping problem in MRI,” *NeuroImage*, vol. 22, no. 3, pp. 1060 – 1075, 2004.

- [55] A. M. Dale, B. Fischl, and M. I. Sereno, “Cortical surface-based analysis: I. segmentation and surface reconstruction,” *NeuroImage*, vol. 9, no. 2, pp. 179 – 194, 1999.
- [56] D. Shepard, “A two-dimensional interpolation function for irregularly-spaced data,” in *ACM '68: Proceedings of the 1968 23rd ACM national conference*, (New York, NY, USA), pp. 517 – 524, ACM, 1968.
- [57] M. Whittington and R. Traub, “Inhibitory interneurons and network oscillations *in vitro*,” *Trends in Neuroscience*, vol. 26, pp. 676 – 682, December 2003.
- [58] P. Fries, “A mechanism for cognitive dynamics: neuronal communication through neuronal coherence,” *Trends in Cognitive Sciences*, vol. 9, pp. 474 – 480, October 2005.
- [59] P. S. Addison, *The Illustrated Wavelet Transform Handbook*. Bristol: Institute of Physics, 2002.
- [60] J. W. Cooley and J. W. Tukey, “An algorithm for the machine calculation of complex Fourier series,” *Mathematics of Computation*, vol. 19, pp. 297 – 301, 1965.
- [61] F. J. Harris, “On the use of windows for harmonic analysis with the discrete Fourier transform,” *Proceedings of the IEEE*, vol. 66, no. 1, pp. 51 – 83, 1978.
- [62] D. Gabor, “Theory of communication,” *Journal of the IEE*, vol. 93, pp. 429 – 457, November 1946.
- [63] L. R. Rabiner and B. Gold, *Theory and Application of Digital Signal Processing*. New Jersey: Prentice-Hall, 1975.
- [64] E. Düzel, R. Habib, B. Schott, A. Schoenfeld, N. Lobaugh, A. McIntosh, M. Scholz, and H. Heinze, “A multivariate, spatiotemporal analysis of electromagnetic time-frequency data of recognition memory,” *NeuroImage*, vol. 18, pp. 185 – 197, 2003.

- [65] D. S. Bassett, A. Meyer-Lindenberg, S. Achard, T. Duke, and E. Bullmore, “Adaptive reconfiguration of fractal small-world human brain functional networks,” *Proceedings of the National Academy of Sciences*, vol. 103, pp. 19518 – 19523, December 2006.
- [66] M. Vetterli and C. Herley, “Wavelets and filter banks: Theory and design,” *IEEE Transactions on Signal Processing*, vol. 40, pp. 2207 – 2230, September 1992.
- [67] S. Mallat, “A theory for multiresolution signal decomposition: The Wavelet representation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, pp. 674 – 692, July 1989.
- [68] J. H. McClellan, T. W. Parks, and L. Rabiner, “A computer program for designing optimum FIR linear phase digital filters,” *IEEE Transactions on Audio and Electroacoustics*, vol. AU-21, pp. 506 – 526, December 1973.
- [69] J. Kaiser, “Filter design using the  $I_0 - \sinh$  window function,” in *Selected Papers in Digital Signal Processing II* (A. V. Oppenheim and L. R. Rabiner, eds.), IEEE Acoustics, Speech, and Signal Processing Society, 1976. reprinted from 1974 IEEE Int. Symp. on Circuits and Syst. Apr. 22–25, 1974, pp.20–23.
- [70] J. Bendat, “The Hilbert transform and applications to correlation measurements,” tech. rep., Brüel & Kjær, 1985.
- [71] M. Le Van Quyen, J. Foucher, J.-P. Lachaux, E. Rodriguez, A. Lutz, J. Martinerie, and F. Varela, “Comparison of Hilbert transform and wavelet methods for the analysis of neuronal synchrony,” *Journal of Neuroscience Methods*, vol. 111, pp. 83 – 98, 2001.
- [72] M. R. Spiegel, *Theory and Problems of Statistics*. New York: McGraw-Hill, 1972.
- [73] D. Moore and G. McCabe, *Introduction to the practice of statistics*. New York: W.H. Freeman, 5th ed., 2006.
- [74] J. Whittaker, *Graphical models in applied multivariate statistics*. New York: Wiley, 1990.

- [75] J. O. Smith, *Mathematics of the Discrete Fourier Transform*. W3K, 2003.
- [76] A. Gupta and S. Nadarajah, *Handbook of Beta Distribution and Its Applications*, ch. Mathematical Properties of the Beta Distribution, pp. 33–54. Marcel Dekker, 2004.
- [77] H. Cohen, *Neuroscience for rehabilitation*. Philadelphia: Lippincott, 1999.
- [78] G. Hickok and D. Poeppel, “The cortical organization of speech processing,” *Nature Reviews Neuroscience*, vol. 8, pp. 393 – 402, May 2007.
- [79] A. Leff, T. Schofield, K. Stephan, J. Crinion, K. Friston, and C. Price, “The cortical dynamics of intelligible speech,” *Journal of Neuroscience*, vol. 28, pp. 13209 – 13215, December 2008.
- [80] B. Fischl, M. I. Sereno, and A. M. Dale, “Cortical surface-based analysis: II. inflation, flattening, and a surface-based coordinate system,” *NeuroImage*, vol. 9, no. 2, pp. 195 – 207, 1999.
- [81] T. Griffiths, S. Kumar, W. Sedley, K. Nourski, H. Kawasaki, H. Oya, R. Patterson, J. Brugge, and M. Howard, “Direct recordings of pitch responses from human auditory cortex,” *Current Biology*, vol. 20, pp. 1128 – 1132, 2010.
- [82] X.-J. Wang, “Neurophysiological and computational principles of cortical rhythms in cognition,” *Physiological Reviews*, vol. 90, pp. 1195 – 1268, July 2010.
- [83] A. Ukil, “Practical denoising of MEG data using wavelet transform,” in *Proceedings of the 13th International Conference on Neural Information Processing* (I. King, J. Wang, L. Chan, and D. Wang, eds.), vol. 4233 of *Lecture Notes in Computer Science*, Springer Berlin, 2006.
- [84] T. T. Nguyen, *Handbook of Beta Distribution and Its Applications*, ch. Maximum Likelihood Estimators of the Parameters in a Beta Distribution, pp. 229 – 236. Marcel Dekker, 2004.