



University of
Nottingham
UK | CHINA | MALAYSIA

Revealing Distinct Neural Signatures in Magnetoencephalography with Hidden Markov Models

Zelekha Abid Seedat

Thesis submitted to the University of Nottingham
for the degree of Doctor of Philosophy

March 2022

Contents

Abstract	6
Acknowledgements	9
Chapter 1	
Introduction	11
1.1 Functional Neuroimaging	11
1.1.1 Invasive measures of brain function.....	12
1.1.2 Non-invasive measures of brain function.....	12
1.1.3 Indirect measures of brain function	14
1.2 Thesis Aim.....	15
1.2.1 Data Interpretation.....	16
1.3 Thesis Layout	19
Chapter 2	
Magnetoencephalography	21
2.1 MEG Signal Generation.....	21
2.1.1 Neurons	21
2.1.2 The Synapse	22
2.1.3 Neurotransmission	24
2.1.4 Types of Neurons.....	26
2.2 Measurable Effects	27
2.2.1 Spontaneous Rhythms.....	28
2.2.2 Evoked Effects.....	29
2.2.3 Induced Responses	30
2.2.4 Functional Connectivity	31
2.3 MEG Data Acquisition	34
2.3.1 SQUIDS	35
2.3.2 Noise Reduction.....	45
2.3.3 Experimental Procedure	48
2.3.4 Coregistration with a structural MRI	51
Summary.....	52
Chapter 3	
MEG Data Analysis	54
3.1 The Forward Problem	54
3.1.1 Estimating the Lead Fields.....	54
3.1.2 Single Sphere Model	57

3.1.3 Alternative Head Models	60
3.2 The Inverse Problem	62
3.2.1 A Generative Model of Sources.....	62
3.2.2 Beamforming	63
3.2.3 Source Orientation	65
3.2.4 Depth Correction	66
3.2.5 Beamformer Artefact Rejection	66
3.2.6 Parcellation	67
3.2.7 Leakage correction	71
3.3 Hidden Markov Modelling.....	75
3.3.1 Model Concept.....	76
3.3.2 Model Inference.....	80
3.3.3 Model Application	85
Summary.....	88
Chapter 4	
The Role of Transient Spectral Bursts in Functional Connectivity	90
4.1 Introduction	91
4.2 Methods	94
4.2.1 Paradigms and data acquisition.....	94
4.2.2 Source localisation	95
4.2.3 Hidden Markov model implementation	96
4.2.4 Burst parameters and comparison to established methods	97
4.2.5 Functional connectivity.....	99
4.3 Results	102
4.3.1 Identification of bursts in resting state data using the HMM	102
4.3.2 Burst dynamics during a visuomotor task	107
4.3.3 Transient spectral bursts and their role in functional connectivity.....	109
4.4 Discussion	112
4.5 Conclusion	118
Chapter 5	
The Post-Stimulus Response Across the Cortex	120
5.1 Introduction.....	121
5.2 Methods	124
5.2.1 Mitigating the Impact of the COVID-19 Pandemic	124
5.2.2 Grip-Force Experiment	126

5.2.3 Visual Experiment	128
5.2.4 Data Analysis.....	130
5.3 Results	137
5.3.1 Grip-Force Stimulus Results.....	137
5.3.2 Visual Stimulus Results	142
5.4 Discussion	147
5.4.1 Classical analysis	147
5.4.2 HMM analysis	148
5.4.3 HMM state interpretation	149
5.4.4 The impact of state durations on connectivity.....	150
5.4.5 Future work	153
5.5 Conclusion	154
Chapter 6	
Hidden Markov Modelling of the Interictal Brain.....	155
6.1 Introduction.....	156
6.2 Methods	160
6.2.1 Patient identification and data collection	160
6.2.2 Data Processing	162
6.2.3 Hidden Markov Modelling.....	164
6.2.4 Beamforming	165
6.2.5 Comparison to existing methods.....	166
6.3 Results	168
6.3.1 Focal spike and wave case	168
6.3.2 Focal epilepsy with polymorphic activity.....	170
6.3.3 Multifocal epilepsy with polymorphic activity	171
6.3.4 Group results	174
6.4 Discussion	175
6.5 Conclusion	178
Chapter 7	
Conclusion	180
7.1 Overview.....	180
7.2 Future Work.....	182
7.3 Closing remarks	183
Appendix A.....	185
Appendix B.....	189

Appendix C.....	192
References.....	198

Abstract

Magnetoencephalography (MEG) is a functional neuroimaging method which measures the magnetic fields produced by neural communication in the brain. Specifically, the fields induced by dendritic current flow in assemblies of pyramidal neurons. Because these magnetic fields are generated directly by brain electrophysiology, and are mostly unperturbed by the skull, MEG data are rich in spatial and temporal information. This thesis is chiefly concerned with interpreting these data in a way that produces useful results whilst minimising bias.

Hidden Markov modelling (HMM) is a robust statistical method which has been applied to fields as diverse as speech recognition and financial market prediction. It parses data into a number of 'hidden states', each with their own unique characteristics, in an unsupervised way. Because it is data-driven, it can create a model unique to each participant's brain activity and specific to each task. In addition, the HMM framework itself is flexible so it can be applied to both sensor and source-space data and can be applied to multiple channels (multivariate) or to a single time course (univariate). Choice of an observation model allows states to be characterised by amplitude, spatial, or spectral content depending on the research question.

The aim of this thesis is to apply hidden Markov modelling (HMM) to whole-head MEG data to identify repeated patterns of transient neural activity occurring throughout the brain. Once these patterns were identified, the interaction between these short 'bursts' of activity across the cortex was established which provided a unique measure of functional connectivity.

Three studies were undertaken:

The role of transient spectral bursts in MEG functional connectivity: In recent years, the smoothly varying neural oscillations often studied in MEG (such as those trial-averaged responses in the traditional neurophysiological (such as alpha/beta) frequency bands) have been shown to be made up of single-trial high-amplitude 'bursts' of activity. These bursts can be observed in the beta

frequency band and are therefore often referred to as *beta bursts*. In this study, a novel time-delay embedded HMM was used to identify bursts in broadband data based on their spectral content for MEG data from 66 healthy adult participants. The burst amplitude, duration and frequency of occurrence were characterised across the cortex in resting state data, and in a motor task the classic movement-related beta desynchronisation and post movement beta rebound were shown to be made up of changes in burst occurrence. A novel functional connectivity metric was then introduced based on the coincidence of bursts from distal brain regions, allowing the known beta band functional connectome to be reproduced. Bursts coincident across spatially separate brain regions were also shown to correspond to periods of heightened coherence, lending evidence to the communication by coherence (Fries 2005, 2015) hypothesis.

Post-stimulus responses across the cortex: During a motor task, both primary (during stimulation) and post stimulus responses (PSR) can be observed. These are well characterised in the literature, but little is known about their functional significance. The PSR in particular is modified in a range of seemingly unrelated neurological conditions with variable symptoms, such as schizophrenia (Robson et al. 2016), autism spectrum disorder (Gaetz et al. 2020) and multiple sclerosis (Barratt et al. 2017), indicating that the PSR is a fundamental neurophysiological process, the disturbance of which has implications on both healthy and pathological brain function. This work therefore tested the hypothesis that the PSR is present across the cortex. MEG data were acquired and analysed from two experiments with 15 healthy adult volunteers each – the first was a right-hand grip task with visual feedback, the second involved passive left visual field stimulation. Both experiments varied stimulus duration (2s, 5s and 10s) with a 30s rest-period between trials to allow characterisation of the full PSR. A univariate 3-state time-delay-embedded hidden Markov model (HMM) was used to characterise the spatial distributions of the primary and PSR across the cortex for both tasks. Results showed that for both tasks, the primary response state was more bilateral over

the sensorimotor or visual areas (depending on task) where the PSR state was more unilateral and confined to the contralateral sensorimotor or visual areas (again, dependant on task). A state coincidence metric was then used to investigate the integration of the primary and PSR states across brain regions as a measure of task-related functional connectivity.

Hidden Markov modelling of the interictal brain: Epilepsy is a highly heterogeneous disease with variations in the temporal morphology and localisation of epileptiform activity across patients. Unsupervised machine learning techniques like the HMM allow us to take into account this variability and ensure that every model is tailored to each individual. In this work, a multivariate time-delay embedded HMM was used to identify brain states based on their spatial and spectral properties in sensor-level MEG data acquired as part of standard clinical care for patients at the Children’s Hospital of Philadelphia. State allocations were used together with a linearly constrained minimum variance (LCMV) beamformer to produce a 3D map of state variance, hence localising probable epileptogenic foci. Clinical MEG epilepsy data are routinely analysed by excess kurtosis mapping (EKM) and so the performance of the HMM was assessed against this for three patient groups, each with increasingly complex epilepsy manifestation (10 patients in total). The difference in localization of epileptogenic foci for the two methods was $7 \pm 2\text{mm}$ (mean \pm SD over all 10 patients); and $94 \pm 13\%$ of EKM temporal markers were matched by an HMM state visit. It is therefore clear that this method localizes epileptogenic areas in agreement with EKM and in patients with more than one focus the HMM provides additional information about the relationship between them.

Acknowledgements

This thesis exists because of all the help and support I have received from everyone over the last four and a half years. All the thank yous in this page will never be able to express how grateful I am for that.

First, to my supervisors, Matt and Karen, thank you for all of your help and guidance throughout my master's and my PhD, and for encouraging me to pursue my own ideas and love of research. And Matt, thank you for helping me find my dream job at Young Epilepsy and looking out for me while I've settled in.

Thank you to my colleagues in the MEG team at Nottingham, especially George, Lauren, Daisy and Elena who helped me a lot in my first year. Thank you, Lukas, for helping me get through the HMM theory section!

To my collaborators, Mark and Andrew in Oxford, your expert help and knowledge with using the HMM has been invaluable. And a huge thank you to Bill at CHOP, for giving me oodles of data and helping me get into the field of epilepsy research which I love. Thank you for being a brilliant mentor and for reminding me I have "true grit".

And thank you to my friends, for your spontaneous bunches of flowers and for making me laugh. Especially Ya Chun, Lizzie, Felipe, Nicola, Clara, and Ryan.

To my parents, thank you for your boundless love and support, and for encouraging my love of learning from the moment I could talk. And to Nani who has taught me to live my life every day and who loves hearing my science updates, whether they're interesting or not! I found those hidden Markovs!

To Hassan, thank you for being with me every step of the way. You encouraged me to apply in the first place and believed in me when I couldn't. Thank you for being with me on countless hospital visits, for sharing lockdown with me, and for moving halfway across the country for this new job. I can't wait to see what adventures the next chapter brings us.

Chapter 1

Introduction

In-vivo brain imaging is an invaluable tool for understanding both healthy and pathological brain function. Broadly, these techniques can be divided into two categories: structural neuroimaging methods such as magnetic resonance imaging (MRI) or X-ray computed tomography (CT), and functional neuroimaging methods such as magneto/electroencephalography (M/EEG) and functional MRI (fMRI). It is the second category to which this thesis belongs. Structural imaging helps us to establish *where* different structures reside in the brain. In contrast, functional neuroimaging helps us to establish *what* each structure does and its role in relation to the rest of the brain and body. Indeed, in the case of medical diagnostics, there are some diseases where there are no obvious structural abnormalities but there is significant divergence from healthy brain function.

1.1 Functional Neuroimaging

The fundamental units of brain function are neurons – single cells which transmit information between one another and to the rest of the body in the form of electrical potentials and chemical synapses. The human brain comprises 86 billion neurons (Azevedo et al. 2009) which must coordinate in such a way that regions of the cortex can become specialised for a particular function (for example, the precentral cortex specialises in motor control), and those specialised cortical regions must also coordinate with one another (for example, the visual and motor cortices must communicate for effective hand-eye coordination). It is the messages that these neurons transmit to one another through electrical potentials along dendrites and axons, or via the movement of chemical neurotransmitters across synapses (the junction between two neurons) which we wish to capture with functional neuroimaging. There are a number of imaging modalities, each with their own pros and cons, which will be discussed in the rest of this section.

1.1.1 Invasive measures of brain function

The functional neuroimaging method with the greatest spatial (millimetres) and temporal (milliseconds) resolution is intracranial EEG (iEEG), including both electrocorticography (ECoG) where a large part of the skull is removed (craniotomy) and an electrode grid is placed in the subdural space, and stereotaxic EEG (sEEG) which uses depth electrodes placed in the brain through small boreholes in the skull. In the case of sEEG, each depth electrode has approximately 10-14 recording contacts, so by placing electrodes deep into the cortex, it is possible to directly intercept the electrical currents produced by populations of neurons surrounding each contact. However, there are a number of limitations to human iEEG recordings, the first is that these methods are highly invasive and are therefore reserved only for patients suitable for brain surgery or experimental animals. This in itself presents more difficulties: because data are only from pathological brains the comparison with healthy brain function (which is of interest in this thesis) is likely to be compromised. The placement of electrodes will also be decided by clinical need, which often favours placement in the temporal lobe compared with other brain locations (Parvizi and Kastner 2018). Animal models prove exceptionally useful; however, it is questionable whether the complex aetiology of a human mental health condition can be faithfully reproduced in an animal. Another limitation of iEEG is the sparse sampling of brain activity - placing enough iEEG electrodes to cover the entire cortex would be clinically impossible.

It is therefore beneficial to consider neuroimaging methods which are non-invasive (to allow study of healthy human brains) and which operate on a macroscopic scale so that whole-brain activity can be recorded.

1.1.2 Non-invasive measures of brain function

The oldest and perhaps most widespread method is scalp EEG, first demonstrated by Hans Berger in 1929 (Berger 1929). This is a direct measure of brain function, where the electric potentials generated by the synchronous firing of many thousands of neurons in the brain are detectable on the scalp by measuring the potential difference between scalp-mounted electrodes. To

make electrical contact, the scalp must be cleaned, and gently abraded before conductive gel is used to stick the electrodes down. The technology is relatively inexpensive and is used internationally for diagnosis of neurological conditions such as epilepsy. Despite its high temporal resolution (similar to iEEG), the spatial resolution of EEG is greatly hampered because the spatial signature of electrical potentials is distorted by the poor electrical conductivity of the skull, making it difficult to localise the source of measured signals.

In comparison to EEG, MEG (Cohen 1972) measures the magnetic fields generated by the same neuronal activity, but has much greater spatial resolution because the magnetic fields induced by current flow in neuronal assemblies are relatively unimpeded by the presence of the skull. This makes it much easier to reconstruct the source of MEG signals. Because MEG is also a direct measure of cerebral activity, the temporal resolution is equivalent to iEEG. Most commercial MEG systems are based on magnetic field sensors which operate at cryogenic temperatures. This makes them expensive to buy and maintain as sensors must be kept in a dewar of liquid helium. However, a new generation of MEG systems are being developed with sensors which operate at room temperature (Boto et al. 2017), potentially bringing the maintenance costs down and allowing sensors to be placed on the scalp in a wearable helmet (this is particularly important for children or patients who struggle to sit still).

The three imaging modalities introduced so far directly measure the electrophysiological component of brain function, rather than the neurochemical element. Neurotransmitters, such as gamma aminobutyric acid (GABA) and glutamate regulate the firing of neurons through inhibition and excitation respectively – in other words, GABA reduces the likelihood of a neuron firing an action potential and glutamate increases the likelihood. Nuclear magnetic resonance (NMR) spectroscopy can be used to image the concentration levels of these neurotransmitters in voxels of the brain, thus serving as another direct measure of brain function. Proton functional magnetic resonance spectroscopy (^1H -fMRS) (Prichard et al. 1991) allows the

time-evolution of the concentration levels of these neurotransmitters to be measured in relation to a task with temporal resolution of less than a minute (Stanley and Raz 2018). Although the temporal and spatial resolution of ^1H -fMRS is coarser than other neuroimaging methods, its unique contribution to understanding the role of key metabolites facilitating brain function cannot be overstated.

1.1.3 Indirect measures of brain function

In studies where temporal resolution is not critical and access to direct electrophysiology and/or neurochemistry is not required, other non-invasive neuroimaging methods can be used. This section introduces indirect functional brain imaging methods which detect the downstream effects of neural activity (like blood flow) and have lower temporal resolution than M/EEG but generally excellent spatial resolution.

The neuroimaging methods in this section all exploit the fact that metabolic demand increases in parts of the brain which are active, meaning that glucose and oxygen must be delivered to support the cellular processes that mediate brain function (and which produce the bioelectromagnetic fields measurable with EEG and MEG). This changing metabolism is tightly coupled, spatially, to active assemblies of cells, allowing the precise mapping of active cortex. However, the increase in blood flow facilitating the increased metabolic demand occurs after the neurons have begun firing, resulting in low temporal resolution.

There is a difference in the magnetic susceptibility of oxygenated (diamagnetic) and deoxygenated blood (paramagnetic) which can be measured with MRI. Functional MRI (Ogawa et al. 1990) exploits this, using the oxygen level of the blood as an endogenous contrast agent – by acquiring multiple images over short timescales the evolution of the blood oxygen level dependant (BOLD) response can be measured with high spatial resolution.

Positron emission tomography (PET) (Ter-Pogossian et al. 1975) is another common neuroimaging method whereby a participant is injected with a

radioactive tracer which travels in the blood to the brain, highlighting locations with increased blood flow or energy uptake (which increase following activation of neurons). The radioactive tracer then decays into a more stable isotope through emission of a positron which annihilates on contact with an electron in the surrounding tissue. This annihilation produces two gamma photons, travelling in opposite directions (with 180° between them) which are then detected by gamma cameras surrounding the patient. Analysis of the two photons using back projection reconstruction allows the location of the annihilation event to be calculated, thus highlighting regions of the brain which are active. PET is a core tool for cancer management but the exposure to radiation from the tracer means that this method is rarely used for healthy research participants.

Functional near infrared spectroscopy (fNIRS) (Jöbsis 1977) is an optical imaging method that exploits the fact that near infrared light can be detected through the scalp using scalp-mounted optodes. There is a difference in the characteristic absorption spectra of oxygenated and deoxygenated haemoglobin which means that shining near infrared light through the scalp and detecting the reflected light as it reaches the optodes allows us to map the blood oxygenation, and therefore blood flow across the cortex. However, this technique has poor spatial resolution due to the scattering of light in the media between cerebral blood vessels and the scalp-mounted optodes as well as poor temporal resolution.

1.2 Thesis Aim

The aim of this thesis is to develop and apply new methods which capitalise on the spatial and temporal dynamics of functional neuroimaging data to allow better understanding of whole-brain function on a millisecond time scale.

It is therefore important that the imaging modality of choice operates on a scale that covers the whole brain and is non-invasive to allow investigation of healthy adult research participants. The spatial and temporal resolution must also be great enough to distinguish the neuronal activity arising in different

cortical areas and across different frequency bands. Of the numerous functional imaging methods available, MEG covers all bases by being both non-invasive and a direct measure of neuronal activity, as well as combining good temporal resolution with good spatial resolution. For these reasons MEG is the imaging modality of choice for this thesis.

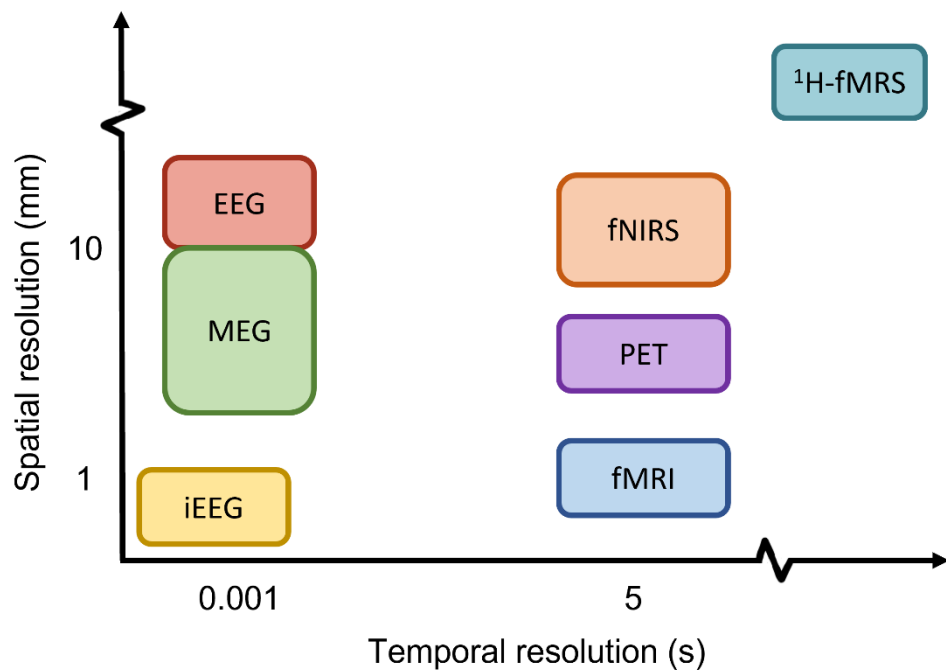


Figure 1.1: The approximate spatial and temporal resolutions of each of the functional neuroimaging methods introduced in this chapter.

1.2.1 Data Interpretation

Once MEG has been identified as the most appropriate non-invasive functional neuroimaging method for understanding whole-brain function on a millisecond timescale, and the immediate challenge of building a system sensitive enough to detect the minute magnetic fields produced by the brain has been overcome, the challenge then becomes interpreting these data which are so rich in spatial and spectral content.

Classical MEG analysis techniques often involve projecting data into source space by means of a spatial filter, followed by frequency filtering, Hilbert transforming and trial-averaging. This helps us to identify neural responses induced by a stimulus in a known frequency band of interest. However established this technique may be, there is inherent bias in the way the data

are analysed which can cause us to miss key information. For example, analysis of single-trial data shows that the smoothly varying oscillatory responses we are used to observing are in fact made up of punctate bursts of activity (Shin et al. 2017; Sherman et al. 2016; Jones 2016; Little et al. 2019). See Figure 1.2.

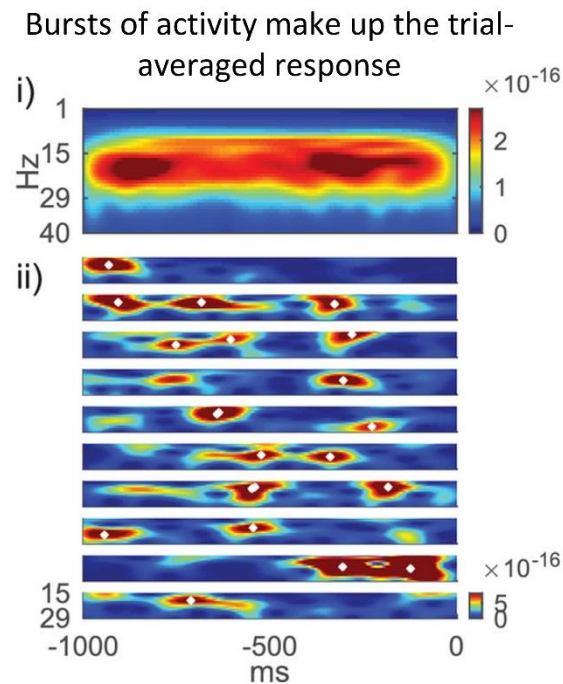


Figure 1.2: By examining single-trial data (ii), it is clear that the classical trial-averaged beta band (13-30Hz) response (i) is made up of transient bursts in the beta frequency band. The colour plots here are time-frequency spectrograms with time on the x-axis and frequency on the y-axis. Figure adapted from (Shin et al. 2017).

These bursts of activity are altered in neurological diseases (Tinkhauser et al. 2017; Rier et al. 2021; Gascoyne et al. 2021) and have been shown to mediate long-range resting state functional connectivity (Seedat et al. 2020). They therefore play an important role in understanding both healthy and pathological brain function and were previously unstudied because data were not often considered at a single-trial level. A similar bias exists when investigating the frequency content of a signal - the data are often filtered into one of the common neurophysiologic frequency bands despite the fact that waveform shapes at a single trial level have complex spectral content.

A possible solution to this problem could be to employ a data-driven machine learning algorithm which deals with raw data and identifies relevant patterns

and similarities without the need to prespecify a frequency band of interest. There are a multitude of techniques to choose from, and the most appropriate must be sought. First, the method should deal with timeseries data because such is the nature of MEG data. The next decision is whether to use a supervised or unsupervised machine learning approach.

Supervised machine learning requires a training dataset with labels assigned to it. The algorithm then learns patterns associated with each of the labels so that it can allocate a label to completely new data. This is useful in classification cases where there is a single sought-after result. For example, classifying patients with or without Alzheimer's disease based purely on a structural MRI of their brain. Unsupervised machine learning approaches are useful for exploratory investigations, where the exact outcome is unknown (it cannot necessarily be labelled), and we only wish to glean information from the patterns which the algorithm identifies. The latter is the approach that this thesis takes.

Hidden Markov modelling (HMM) is one such unsupervised machine learning approach – timeseries data are used as input, and the output is a series of *hidden states* where each state represents a different pattern of activity in the MEG data. One of the key benefits of using an HMM is that it is a flexible framework – by changing the observation model the state definitions can be changed from one which differentiates states based on their amplitude, to one that is based on their spectral content for example. And the input data can be both univariate (such as a single time course from a single MEG virtual electrode) or multivariate (characterising states based on multiple MEG channels which can be useful for mapping the spatial pattern of state activity).

It is for these reasons that HMMs have been applied to MEG data in recent years with some success. For example, identification of the timepoints associated with individual HMM states allowed Woolrich et al. to develop a temporally adaptive beamformer to aid in MEG source localisation (Woolrich et al. 2013). This was followed by investigations into functional connectivity (Baker et al. 2014; Vidaurre et al. 2018) and has also helped to disentangle the

complex MEG signals surrounding neurological diseases such as Parkinson's (Heideman et al. 2020) and schizophrenia (Gascoyne et al. 2021). This thesis uses variations of the HMM to identify transient bursts of neuronal activity in both healthy adult MEG data, and in data from paediatric patients with epilepsy.

1.3 Thesis Layout

First, a detailed theory of MEG is introduced in Chapter 2. This begins with a description of the biological origins of the MEG signal, including the physics underpinning neurotransmission and the types of neurons involved. The effects measurable with MEG are then introduced: evoked and induced responses, followed by functional connectivity. The final part of this chapter covers the hardware required for MEG data acquisition and the way in which the different components of a MEG system integrate.

The theory behind MEG data analysis is introduced in Chapter 3. In the first part of this chapter, the mathematics behind source localisation is covered including solutions to the forward and inverse problems. Then a thorough introduction to HMMs and their application to MEG research is addressed.

Chapter 4 is the first experimental chapter and is about the detection of beta bursts in regions across the cortex. First a method for detecting beta bursts using an HMM is introduced which is compared with existing techniques. This method is then applied to both resting state and visuomotor task data. It is then shown that the well-characterised resting state functional connectivity matrices can be derived using a measure of burst overlap.

The same burst detection algorithm is applied in Chapter 5 to two experiments with the aim of characterising the post-stimulus response and its distribution across the cortex. This is done for both a passive visual task and a motor task with visual feedback. Data acquisition is described, as well as the impact of the COVID-19 pandemic on this lab-based work.

In Chapter 6 the HMM framework is applied to sensor-space MEG data from 10 paediatric epilepsy patients with the aim of localising interictal activity.

There are three patient groups, each with increasingly complex epilepsy and the performance of the HMM is compared with excess kurtosis mapping (a technique used in some clinical MEG settings around the world).

A reflection and conclusion of the work presented in this thesis is given in Chapter 7.

Chapter 2

Magnetoencephalography

Magnetoencephalography (MEG) is the measurement of extra cranial magnetic fields induced by current flow in neuronal assemblies in the brain. By applying mathematical models to these measured fields, we are able to infer changes in neuronal current flow throughout cognition. This offers a direct but non-invasive measurement of neuronal activity and provides insight into the functional significance of individual brain regions as well as the way the brain coordinates activity between regions.

This chapter is split into three parts. Section 2.1 will describe the neurophysiology underpinning the generation of MEG signals. The resulting macroscopic measurable effects will then be introduced in Section 2.2, and the final part of this chapter (Section 2.3) will describe the technology used to acquire these data.

2.1 MEG Signal Generation

The human brain is made up of billions of neurons, and together with the spinal cord it forms the central nervous system. A single neuron can be thought of as a fundamental unit of the nervous system – a type of highly specialised biological wire, allowing information to be carried from every part of the body to the brain and back again, integrating the central and peripheral nervous systems. It is the current flow, generated by the creation and transmission of potential differences in the neurons of the brain, which induces the magnetic fields measurable outside the head in MEG.

This section will describe the biology underpinning MEG signal generation and the physics of neurotransmission.

2.1.1 Neurons

Neurons are a fascinating sub-group of animal cells with the ability to conduct electrical activity from one part of the body to another. There are many types of neurons and each comprise many parts, but there are four key components:

the soma (or cell body), the dendrites, axon, and myelin sheath. A diagram of a single neuron is shown in Figure 2.1.

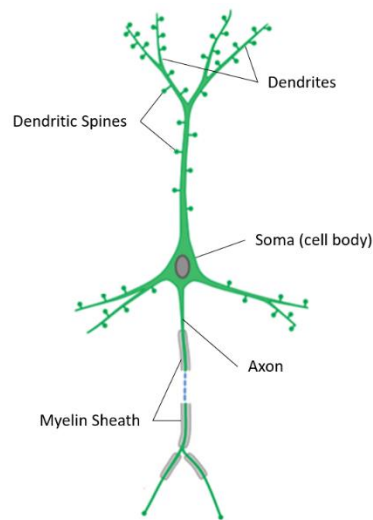


Figure 2.1: A simplified diagram of a neuron showing the cell body and dendrites which largely make up the grey matter of the cerebral cortex, and the axon covered in a myelin sheath which largely makes up the white matter in the central nervous system. Figure adapted from (Heidelberg et al. 2004).

Information is passed from cell to cell via the synapse. The synapse is the gap between cells across which neurotransmitters (chemicals which regulate cell membrane permeability) relay information. The incoming signal from the axon of the presynaptic cell terminates at the synaptic cleft (the gap between the cells). In pyramidal neurons most of the excitatory inputs are received at the synapse on dendritic spines which are present in large numbers (Heidelberg et al. 2004).

2.1.2 The Synapse

When a nerve impulse reaches the presynaptic terminal of a neuron it is accompanied by an influx of calcium ions into the intracellular fluid. This triggers the migration of synaptic vesicles towards the presynaptic membrane where they fuse with the membrane and release a chemical neurotransmitter into the synaptic cleft. This neurotransmitter then binds to receptor molecules on the postsynaptic membrane, allowing ion-specific membrane channels to open. Most synaptic inputs are received by the dendrites of the recipient cell and the flow of sodium ions into the cytoplasm produces an excitatory post-

synaptic potential which travels towards the soma. Both excitatory and inhibitory inputs (dependant on the flow of ions into and out of the cell at the synapse) can arrive at the soma and the sum of these will determine whether an action potential¹ 'fires' along the axon of the neuron.




Figure 2.2: A nerve impulse arrives at a synapse, triggering the release of neurotransmitter into the synaptic cleft. The neurotransmitter binds to receptor molecules on the postsynaptic membrane allowing sodium ions to enter the cytoplasm and triggering a post-synaptic potential.

The post-synaptic potential comprises an intracellular primary current and an extracellular volume current. The primary current lasts approximately 10ms and travels towards the soma. The ions in the extracellular space flow towards the area of repolarisation after the primary current has passed, resulting in current flow in the opposite direction to the primary current, towards the synapse. This is referred to as the volume current and is dependent on the conductivity of the extracellular tissue.

¹ Action potentials are described in detail in section 2.1.3




Figure 2.3: The postsynaptic potential in the dendrites of neurons is made up of an intracellular primary current and an extracellular volume current. Figure adapted from (Vrba 2000).

The contributions of both the primary and volume currents to the measurable MEG signals are considered in Chapter 3, where it is shown that the impact of the volume current is negligible.

2.1.3 Neurotransmission

The flow of neurotransmitters across a synapse allows information to be shared *between* neurons, but the transmission of information *within* a single neuron (the longest of these in the human body can be up to one meter (Fletcher and Theriot 2004)) must be fast and efficient. This section describes the action potential, sometimes referred to as a nerve impulse.

Every cell has a 'resting potential' characterised by the presence of positively and negatively charged ions in the intracellular and extracellular fluid as well as the relative membrane permeabilities of these ions. This potential difference across cell membranes forms the basis of electrical conduction in neurons. When the post-synaptic currents described in the previous section flow along the dendrites of the neuron and arrive at the soma, they combine to produce an overall intracellular potential. It is only when the potential at the axon hillock (the portion of the cell body adjacent to the axon) reaches a firing threshold of approximately -55mV that an action potential is initiated.

An action potential is a series of rapid changes in the local membrane potential. The **resting potential** of a neuron is -70mV and is largely caused by the passive diffusion of positive potassium ions (K^+) out of the cell along a concentration gradient. When a nerve impulse arrives at a particular location on the axon, the local membrane permeability changes to allow positive sodium ions (Na^+) to flow into the cell. This process is called **depolarisation** and results in a positive membrane potential of $+35\text{mV}$. Once this potential is reached, the sodium channels in the cell membrane close and an increased number of potassium channels open. This allows more potassium to flow out of the cell and, after a brief period of hyperpolarisation, returns the membrane potential to -70mV . Part of this **repolarisation** phase involves sodium and potassium ions being actively pumped, against their concentration gradients, out of and into the cell respectively. This process is demonstrated in Figure 2.4. The adjacent part of the axon then undergoes the same process thus allowing the action potential to travel along the length of the axon. It is worth noting that the brief period of hyperpolarisation prevents the action potential from travelling in both directions along the axon, allowing directionality of nerve impulses.

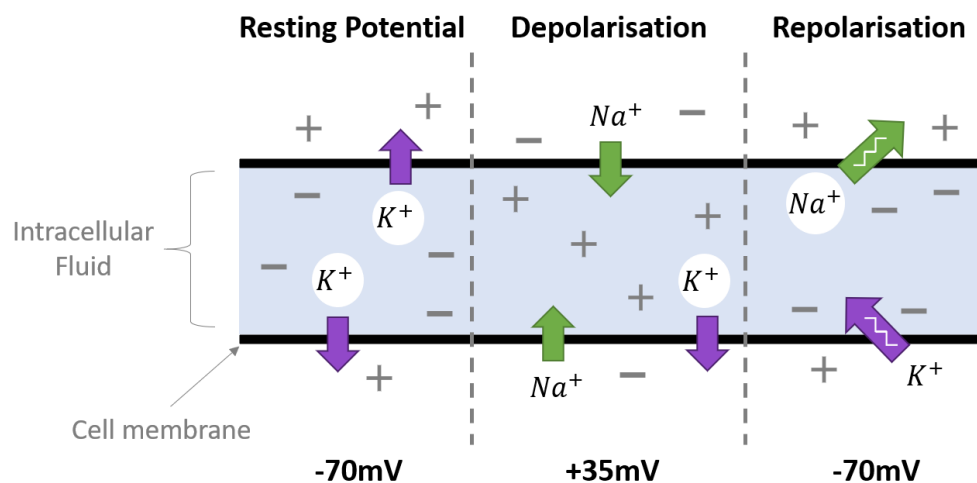


Figure 2.4: The membrane potential and ion permeabilities during the resting, depolarisation, and repolarisation phases of an action potential.

We therefore have two possible generators of the magnetic fields which are measurable outside of the head: action potentials and post-synaptic potentials. Action potentials last around 1ms and can be approximated as two current dipoles pointing in opposite directions. The depolarisation edge of the action potential is modelled as a single dipole in the direction of propagation along the axon and the repolarisation edge as a dipole oriented towards the soma. This is the same as approximating the action potential as a quadrupolar current. In comparison, the post-synaptic potential lasts approximately 10ms and can be modelled as a dipole. The magnetic field produced by these currents is dependent on distance and drops off more quickly for the quadrupolar field ($\frac{1}{r^3}$ compared with the $\frac{1}{r^2}$ of a dipolar field). This, combined with the fact that post-synaptic potentials last longer than action potentials, leads us to believe that the signals measurable outside of the head by MEG are a superposition of the magnetic fields induced by many thousands of individual postsynaptic potentials occurring in the dendrites of neurons.

2.1.4 Types of Neurons

The neurons in the brain are highly specialised and can be difficult to separate into distinct types; they have different functions, arise in specific locations, use various neurotransmitters, and have particular dendritic geometries. It is this latter characteristic that we are most interested in, given that the signals we measure in MEG arise in the dendrites of neurons. We can therefore categorise neurons into two types; these are stellate cells and pyramidal cells, shown alongside one another in Figure 2.5. In stellate cells the dendrites radiate from the soma in all directions while the dendrites of a pyramidal neuron are oriented in a common direction.




Figure 2.5: Reconstructions of a pyramidal cell (A) and a stellate cell (B). The dendrites of a pyramidal cell largely point in the same direction whereas the dendrites of a stellate cell have a roughly symmetric geometry, pointing in every direction. Figure adapted from (Churchill et al. 2004).

When many stellate neurons fire in synchrony, the variable orientation of the dendrites leads to field cancellation and hence a loss in the total measurable signal. Conversely, when many pyramidal neurons fire in synchrony, the signal is amplified because the dendrites are oriented in a common direction. The microscopic origin of the MEG signal is therefore likely to be the post-synaptic potentials in the dendrites of pyramidal neurons and approximately 10,000-50,000 pyramidal cells must be simultaneously active before the macroscopic signal can be detected (Murakami and Okada 2006).

2.2 Measurable Effects

The signals produced by these large populations of neurons, which are detected outside of the head, involve complex layers of information which must be untangled if we are to understand the mechanisms underlying both healthy and pathological brain function. Because of the millisecond temporal resolution of EEG and MEG, these data are rich in spectral information and have been studied over the last century with a variety of methods, allowing us to extract different types of brain activity. Measurable neuroelectric effects fall roughly into four categories: spontaneous rhythms, evoked effects, induced responses, and functional connectivity.

2.2.1 Spontaneous Rhythms

Following work in animals conducted by Richard Caton (Caton 1875), in 1929 Hans Berger used EEG to discover the first of a number of 'spontaneous rhythms' in the human brain, that is rhythms produced by the brain not related to any external stimulus. Berger referred to this oscillatory signal in the frequency range 8-13Hz as alpha activity and found that its amplitude is modulated by opening and closing one's eyes (Berger 1929). Further classifications of frequency bands of interest occurred organically over the following decades and we now refer to them as delta (0.2 to 3.5Hz), theta (4 to 7.5Hz), alpha (8 to 13Hz), beta (14 to 30Hz) and gamma (30 to 90Hz) (Lopes da Silva 2013). The neural activity falling into each of these frequency bands is often strongest in certain brain regions and related to a specific brain function. For example, beta oscillations have the highest amplitude over the sensorimotor cortex and are modulated by motor tasks. An increase in beta activity is also associated with the disruption of these motor networks in diseases such as Parkinson's where symptom severity correlates with beta amplitude (Tinkhauser et al. 2017).

In fact, oscillatory activity is observed on all neurophysiological scales – from the microscopic activity of a single neuron, through the mesoscopic synchronised activity of neuronal ensembles, and the macroscopic oscillations correlated between distal brain regions. But it can be confusing to refer to these spontaneous rhythms as oscillations because pure alpha or beta oscillations do not exist in isolation, instead the neural activity contains spectral peaks in any of the above-mentioned frequency ranges. It is therefore important to keep in mind the underlying waveform shape when describing this oscillatory activity, because although parsing the data into frequency bands of interest is useful for narrowing the scope of the problem and has shed light on the way the brain works, very different waveform shapes can yield near-identical spectra as shown by Jones in 2016 (Figure 2.6).

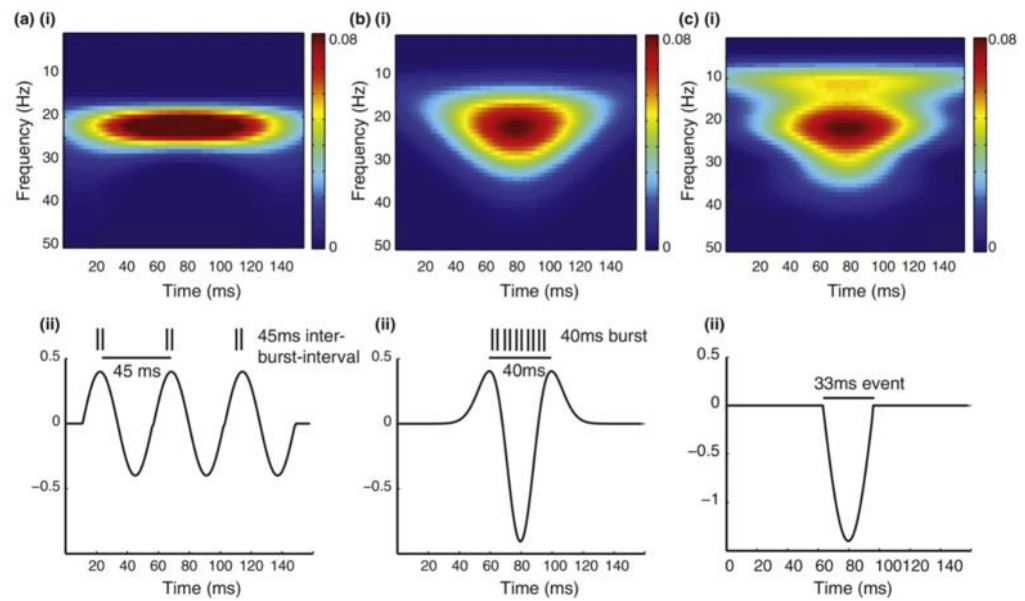


Figure 2.6: Different underlying waveforms (ii) can yield very similar spectra (i), in this case all peaking at a frequency of 22Hz. The waveforms are (a) three cycles of a 22Hz sine wave, (b) an inverted Ricker wavelet, and (c) a single deflection. Their corresponding time-frequency spectrograms are shown in (ai), (bi) and (ci) respectively. The implication on the functional significance of these different underlying waveform shapes must be considered when analysing data in the frequency domain. Figure from (Jones 2016).

In much the same way, the interaction *between* distinct frequency bands of neuronal activity, so-called cross-frequency coupling, especially between theta and gamma oscillations, may simply be a result of fast transitions in waveforms (Kramer, Tort, and Kopell 2008). If this is the case, then such cross-frequency coupling does not tell us something about the interaction between functionally distinct neural processes, but rather about the temporal morphology of the brain activity. This is a subtlety which must be considered when interpreting the functional significance of spontaneous rhythms.

2.2.2 Evoked Effects

Evoked responses occur when an external stimulus is introduced, and the activity is both time and phase-locked. This means that any positive or negative peaks in the neural activity occur at the same time (relative to stimulus presentation) over multiple trials, so that averaging over trials yields a clear waveform shape (having averaged out any spurious non-time locked effects).

Following Hans Berger's 1929 publication, Adrian and Matthews further studied alpha oscillations (termed the 'Berger Rhythm') and not only repeated the results of Hans Berger but also identified a characteristic evoked potential in the occipital cortex produced by a flickering light source (Adrian and Matthews 1934). The visual evoked response is now clearly understood and variations in its morphology can be indicative of a number of diseases affecting visual pathways, visual cortex or retina (Creel 2019). Chapter 6 of this thesis is concerned with data from a visual experiment; however, we will not be studying the evoked effect, but rather the induced response to a visual stimulus.

2.2.3 Induced Responses

A change in the spontaneous rhythms caused by an external stimulus is known as an induced response. These are time-locked but not phase-locked to the stimulus presentation which means that averaging over many trials will attenuate the signal, rather than emphasise it (as would be the case for an evoked effect); this is because the peaks and troughs of the waveform induced by the stimulus do not line up across trials. To study an induced response, the data are first frequency-filtered into the band of interest. The amplitude envelope of this frequency-specific signal is then calculated, and it is this that is averaged over many trials to reveal the change in oscillatory power relative to stimulus presentation.

A well characterised induced response is that of the post-movement beta rebound (PMBR). Following the cessation of a motor task (such as a button-press or finger abduction), there is a short increase in the amplitude of beta activity in the contralateral motor cortex which gradually returns to baseline. This increase in beta activity is also known as an event related synchronisation (ERS). In contrast during the movement an event related desynchronisation (ERD) of the beta activity occurs. The functional significance of this PMBR in healthy people is under debate, but it is altered in a number of neurological disorders such as schizophrenia (Robson et al. 2016; Gascoyne et al. 2021), autism spectrum disorder (Gaetz et al. 2020), mild traumatic brain injury (Rier

et al. 2021), and multiple sclerosis (Barratt et al. 2017), where the severity of illness correlates with the deviation of the PMBR from normal. It has also been shown that these relationships to neurological disorders are unique to the PMBR and aren't seen with respect to the ERD. This idea that the PMBR is a unique brain response is investigated more in Chapter 5.

In recent years, these induced responses are being increasingly analysed on a single trial level, rather than averaging over many trials. This is because what would appear to be a sustained oscillation in trial-averaged data is actually made up of transient 'bursts' of activity (Shin et al. 2017) - where the timings, durations, frequencies of occurrence and amplitudes of these bursts provide further information about the fundamental neural processes involved in cognition. See Chapters 4 and 6 for experimental work in understanding the role of burst activity in resting state data and in response to simple visual, motor, and combined visuomotor tasks.

Many of the neuroelectric effects considered in this thesis will be induced responses to a stimulus (either motor or visual). However, a large portion of this work also studies resting state functional connectivity (relationships between spontaneous neural activity originating at different cortical locations).

2.2.4 Functional Connectivity

The brain constitutes many different cortical areas, each with their own functional significance, and each made up of large populations of neurons firing in synchrony. These brain regions not only coordinate activity between the many thousands of neurons they comprise, but also activity between disparate brain locations, so that an efficient and coordinated response to external stimuli might be affected.

Coordinated activity between different areas of the brain – functional connectivity - can be captured using a variety of techniques. One such method posits that the optimal window for transfer of information between brain regions is when their electrical potentials peak at the same time, or in other

words their oscillations are in phase. This is the ‘communication by coherence’ hypothesis (Fries 2005); the oscillations from different cortical locations share windows for input and output of information when they are coherent. There are a number of phase-based connectivity metrics including phase coherence (Mormann et al. 2000), imaginary coherence (Nolte et al. 2004), phase-lag index (Stam, Nolte, and Daffertshofer 2007) and phase difference derivative (PDD) (Breakspear and Williams 2003).

Phase Difference Derivative (PDD)

One of the benefits of PDD is that it produces a time course of phase coherence allowing us to identify the moments in time when connectivity is maximal. It involves first calculating the instantaneous phase of a ‘seed’ and ‘test’ signal, θ_1 and θ_2 respectively, and then subtracting them to get the phase difference,

$$\varphi_{1,2}(t) = \theta_1(t) - \theta_2(t), \quad (2.1)$$

where $\varphi_{1,2}$ represents phase difference. The derivative of this phase difference then allows us to assess the stability of the phase relationship: where the PDD values are close to zero, the phase difference is constant, and the coherence is high

$$\frac{d\varphi_{1,2}(t)}{dt} \approx 0. \quad (2.2)$$

Similarly, where the PDD values are high, the phase difference between the two signals is rapidly changing and the coherence is low. These values can then be negatively transformed using

$$x_{new} = e^{-|x|}, \quad (2.3)$$

so that a high value is consistent with high coherence.

Amplitude Envelope Correlation (AEC)

Another key measure of functional connectivity is amplitude envelope correlation (AEC). In this case, it is the amplitude, rather than the phase of the signal, which implies a functional connection between distal regions of the brain. In fact, network structures calculated via AEC in resting state MEG match

frequently observed functional magnetic resonance imaging (fMRI) resting state networks (Brookes, Woolrich, et al. 2011; Hipp et al. 2012).

For every brain location, MEG data are first filtered into a frequency band of interest to give, $\hat{q}(t)$. The analytic signal, $\hat{z}(t)$, is then derived as

$$\hat{z}(t) = \hat{q}(t) + iH[\hat{q}(t)], \quad (2.4)$$

where H is the Hilbert transform. The Hilbert transform is the convolution of $\hat{q}(t)$ with the function $h(t) = \frac{1}{\pi t}$ so that

$$H[\hat{q}(t)] = \frac{1}{\pi} P \int_{-\infty}^{\infty} \frac{\hat{q}(\tau)}{t-\tau} d\tau, \quad (2.5)$$

where P is the Cauchy principal value which accounts for the singularity when $t = \tau$.

The amplitude envelope of this signal is then given by

$$E[\hat{q}(t)] = \sqrt{(\hat{q}(t))^2 + (H[\hat{q}(t)])^2}. \quad (2.6)$$

The resulting data are often down sampled to ~ 1 Hz yielding a slow amplitude time course for each brain region. For every pair of regions, the Pearson's correlation coefficient, $r_{E_1 E_2}$, is computed as a measure of the connectivity strength between brain areas. AEC is typically represented in matrix format.

$$r_{E_1 E_2} = \frac{\sum_{i=1}^n (E_{1i} - \bar{E}_1)(E_{2i} - \bar{E}_2)}{\sqrt{\sum_{i=1}^n (E_{1i} - \bar{E}_1)^2} \sqrt{\sum_{i=1}^n (E_{2i} - \bar{E}_2)^2}} \quad (2.7)$$

where E_1 and E_2 are the slow amplitude time courses for any two distinct regions and \bar{E}_1 and \bar{E}_2 are their corresponding mean values.

In this thesis, I will use both AEC and PDD to validate my own novel connectivity metric based on beta bursting events, see Chapter 4. I will show that it is in fact the 'bursts' of activity mentioned in Section 2.2.3 above, which provide optimal windows of both high amplitude and high phase coherence for transfer of information between brain regions, thus unifying these two established methods.

Throughout this section we have assumed that functional connectivity is derived from true brain activity (with no signal leakage). However, due to the nature of the inverse problem, it is impossible to infer the activity at one location in the brain without it being contaminated by activity from another brain region. This can manifest as a zero-phase-lag correlation between regions which is not due to coordinated brain activity, rather it is a limitation of our imaging capabilities. This problem is addressed in detail in Section 3.2.7.

2.3 MEG Data Acquisition

The first MEG recordings using a single channel SQUID² system were made by David Cohen in 1972, and like Berger, his predecessor, he measured spontaneous alpha oscillations which were modulated by opening and closing of the eyes (Cohen 1972). Given that MEG signals originating from the brain are on the order of tens to hundreds of femtoTesla (fT) which is vanishingly small when compared to the Earth's magnetic field (tens of microTesla), this was no mean feat.




Figure 2.7: The first MEG signals recorded using a single channel SQUID system mounted inside a magnetically shielded room. Figure adapted from (Cohen 1972).

² See Section 2.3.1

The most common commercially available MEG systems today boast arrays of hundreds of sensors, but the hardware used for field measurement has changed very little. This section will describe the theory underpinning signal detection.

Another key component of MEG data acquisition is effective shielding and interference reduction from both external and internal sources of magnetic fields not originating from the brain. These include the Earth's static magnetic field, as mentioned above, as well as the fields produced by electronic equipment, cars and even other biological magnetic fields (with the human heart producing fields at hundreds of picoTesla and lung particles producing nanoTesla fields) (Vrba 2000). This will be covered in Section 2.3.2.

2.3.1 SQUIDS

Until recently³, most commercially available MEG systems used arrays of superconducting quantum interference devices (SQUIDS) to measure the magnetic fields produced by the brain. This section describes the physics underpinning SQUIDS and how they are incorporated into a large multichannel array.

Superconductivity

To understand how a SQUID works, we must first understand superconductivity. As a metal is cooled, its resistance falls gradually until it reaches a non-zero value at 0K. However, in the case of a superconductor, the resistance of the metal will abruptly drop to 0Ω at a critical temperature, T_c , forming a perfect conductor (Onnes 1911).

In 1957, Bardeen, Cooper and Schrieffer put forward a theory to explain this counterintuitive phenomenon; today it is referred to as BCS theory. As an electron moves through a conductor it will temporarily distort the crystal lattice structure, forming a "wake" of positive charge close to the electron. If a

³ Large arrays of optically pumped magnetometer (OPM) systems are now being used in a research environment and are expected to supersede conventional SQUID based systems because sensors can be placed much closer to the head for improved signal to noise ratio and because they have a lower maintenance cost as liquid helium is not required for cooling.

neighbouring electron is close by, it will be attracted towards this area of positive charge and the two electrons will bind together to form a 'Cooper Pair'. It seems unlikely that two like-charged particles would be bound together in this way but the attraction to the region of positive charge is greater than the repellent Coulomb force, which is reduced by the effect of nearby electrons (sometimes referred to as the screened Coulomb interaction). The distance between electrons in a pair is known as the coherence length and is dependent on the conductor material but is on the order of 100nm (Bardeen, Cooper, and Schrieffer 1957).

Since these Cooper pairs are bound pairs of electrons with opposite spin, they exhibit bosonic properties – namely, many pairs can exist in the same quantum state. As more pairs fall into the same state, the binding energy within the pairs increases, making superconductivity a cooperative phenomenon. It is the condensation of these Cooper pairs into a ground state at temperatures below the critical temperature which produces the effect of superconductivity (Bardeen, Cooper, and Schrieffer 1957). The thermal energy required to break a Cooper pair in this condensate state is $3.5k_B T_C$ where k_B is the Boltzmann constant and T_C is the critical temperature of the conductor.

The individual wavefunctions of the Cooper pairs constructively interfere so that we can describe the whole ensemble with a single canonical wavefunction,

$$\psi(\mathbf{r}) = \psi_0 e^{i\mathbf{k}\cdot\mathbf{r}}, \quad (2.8)$$

where ψ_0 is the ground state of the Cooper pairs, \mathbf{r} is the position (or centre of mass) of the two electrons and \mathbf{k} is the wave vector of the condensate.

Relating Superconductivity with Magnetic Field Measurements

Superconductors can be classified into two types by their interaction with applied magnetic fields. Screening currents flowing on the surface of a superconductor produce an equal and opposite magnetic field to an applied one, which means that the field inside the superconductor is zero; this is known as the Meissner effect (Meissner and Ochsenfeld 1933). This effect is

seen below a critical applied field, B_C , which is dependent on temperature. In type I superconductors, this transition is abrupt at B_C . In type II superconductors, there is a transition period where magnetic flux partially penetrates the material at a lower critical field B_{C1} , and the superconducting nature (zero resistance) persists until an upper critical field B_{C2} is reached – at which point there is complete flux penetration. Because of the ability to retain superconducting properties at higher fields, type II superconductors are generally used in SQUID magnetometers.




Figure 2.8: Below a critical field, B_C , the Meissner effect means that magnetic flux is expelled by the superconductor. This is due to screening currents flowing on the superconductor surface. Figure adapted from (Hook and Hall 1974).

Now we will consider a superconducting ring and the mathematics relating the behaviour of its macroscopic wavefunction with the magnetic flux passing through the loop (Hook and Hall 1974; Feynman, Leighton, and Sands 1963; London 1948):




Figure 2.9: A closed loop Γ in the body of a superconducting ring. Figure adapted from (Feynman, Leighton, and Sands 1963).

The local current density, $\mathbf{j}(\mathbf{r})$, associated with the canonical wavefunction is

$$\mathbf{j}(\mathbf{r}) = \frac{i\hbar e}{2m}(\psi^*\nabla\psi - \psi\nabla\psi^*) - \frac{2e^2}{m}\psi^*\psi\mathbf{A}, \quad (2.9)$$

where \hbar is the reduced Planck constant, e is the electron charge, m is the electron mass and \mathbf{A} is the magnetic vector potential. This equation satisfies the conservation of probability when ψ satisfies the time-dependent Schrodinger equation, and can be derived by considering the relationship between the momentum and velocity vectors for a particle (Appendix C (Hook and Hall 1974)).

By taking the most general form of Equation 2.8, $\psi(\mathbf{r}) = |\psi_0|e^{i\theta(\mathbf{r})}$, and inserting this into Equation 2.9, we get

$$\begin{aligned} \mathbf{j}(\mathbf{r}) &= \frac{i\hbar e}{2m} \left(\psi_0 e^{-i\theta} (e^{i\theta} \nabla \psi_0 + i\psi_0 e^{i\theta} \nabla \theta) - \psi_0 e^{i\theta} (e^{-i\theta} \nabla \psi_0 - i\psi_0 e^{-i\theta} \nabla \theta) \right) \\ &\quad - \frac{2e^2}{m} |\psi(\mathbf{r})|^2 \mathbf{A} \\ &= \frac{i\hbar e}{2m} (2i|\psi(\mathbf{r})|^2 \nabla \theta) - \frac{2e^2}{m} |\psi(\mathbf{r})|^2 \mathbf{A} \\ &= \frac{-e}{m} |\psi(\mathbf{r})|^2 (\hbar \nabla \theta + 2e\mathbf{A}). \end{aligned} \quad (2.10)$$

We can then assume that the current density far from the surface (within the body of the superconductor) is zero, $\mathbf{j} = 0$, due to the Meissner effect⁴ and Equation 2.10 simplifies to

$$\hbar \nabla \theta = -2e\mathbf{A}. \quad (2.11)$$

Now, integrating around the closed loop, Γ , of the superconducting ring produces

$$\hbar \oint \nabla \theta \cdot d\mathbf{l} = -2e \oint \mathbf{A} \cdot d\mathbf{l} \quad (2.12)$$

and since the line integral of \mathbf{A} around a loop is equal to the flux, Φ , through the loop,

$$\hbar \oint \nabla \theta \cdot d\mathbf{l} = -2e\Phi. \quad (2.13)$$

⁴ The Meissner effect can be described by the London equation (London 1948).

Considering the superconducting ring again, we can solve the left-hand side of Equation 2.13. At any point on the ring, there can only be one value of the wavefunction, $\psi(\mathbf{r})$, so the change in θ from the start to the end of the loop, $\Delta\theta$, must be $2\pi n$. Therefore

$$2\pi n\hbar = -2e\Phi \quad (2.14)$$

which shows that the flux through a closed loop is quantised. The flux quantum, $\Phi_0 = h/2e$, has a value of $2.07 \times 10^{-15} \text{Wb}$.

The Josephson Effect

In 1962, Josephson measured a current caused by the tunnelling of Cooper pairs across a junction (created by an oxide layer) in a superconductor (Josephson 1962). This is similar to the tunnelling of electrons into a potential barrier. On either side of the barrier, the superconductors have wavefunctions $\psi_1 = |\psi_1|e^{i\theta_1}$ and $\psi_2 = |\psi_2|e^{i\theta_2}$. Provided that both superconductors are made from the same material and are at the same temperature, $|\psi_1|$ will equal $|\psi_2|$, but because they are not in contact their phases will be different, $\theta_1 \neq \theta_2$. However, if the oxide layer is thin enough, these superconductors are said to be *weakly coupled*. This means that they are in enough contact so that the lowest energy state has $\theta_1 = \theta_2$, but by applying a small current through the junction (or a potential difference across it), we can generate a phase difference.

Inside the oxide layer, the canonical wavefunction becomes a superposition of the two decaying wavefunctions,

$$\psi = \sqrt{\frac{n_s}{2}} (e^{i\theta_1 - K(x+d/2)} + e^{i\theta_2 + K(x-d/2)}), \quad (2.15)$$

as shown in Figure 2.10. This assumes that the contributions from both ψ_1 and ψ_2 are negligible by the time they reach the other side of the barrier. Note that K^{-1} is the characteristic decay length of the barrier, and n_s is the

number of electrons per unit volume (this is divided by 2 to account for there being 2 electrons per Cooper pair).

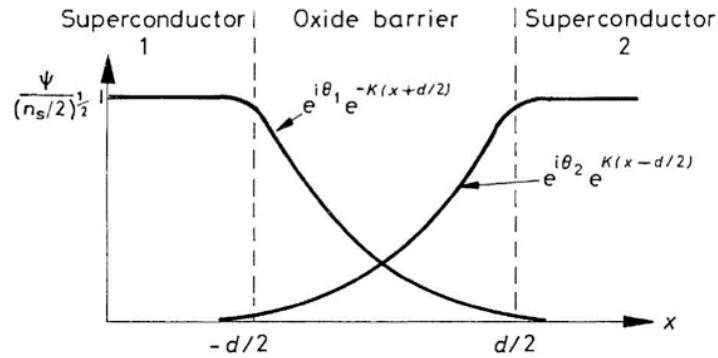


Figure 2.10: A Josephson junction comprises two superconductors separated by an oxide layer. Cooper pairs can tunnel into this barrier as shown by the exponential decay of the wavefunctions from both superconductors 1 and 2. Figure from (Hook and Hall 1974).

We can then substitute the wavefunction from Equation 2.15 into Equation 2.9 where $A = 0$ as before, to find the Cooper pair current density through the barrier,

$$j = \frac{ie\hbar n_s}{2m} K e^{-Kd} (-e^{i(\theta_1 - \theta_2)} + e^{i(\theta_2 - \theta_1)}) = j_0 \sin \delta. \quad (2.16)$$

Here, δ is the difference between θ_1 and θ_2 across the barrier and the maximum current density in the oxide layer is $j_0 = e\hbar n_s K e^{-Kd} / m$, corresponding to a phase difference of $\pi/2$. This relationship between current flow and phase is known as the Josephson equation, and the flow of Cooper pairs through the junction is called the DC Josephson effect (Josephson 1962).

Quantum Interference

By connecting two Josephson junctions in parallel, quantum interference effects between the currents flowing through each of the junctions can be measured (Jaklevic et al. 1964). This is directly analogous to the interference patterns measured in double slit electron beam interference experiments.




Figure 2.11: Two Josephson junctions connected in parallel produce quantum interference effects relating the current, J_{total} , through the system to the flux passing through the loop. This superconducting quantum interference device (SQUID) allows us to make highly accurate measurements of minute neuromagnetic fields. Figure from (Feynman, Leighton, and Sands 1963).

Let us consider a superconducting ring as before, but this time it is broken in two places by a Josephson junction, as shown in Figure 2.11. The total current, J_{total} , flowing from point P to Q will be the sum of the currents flowing in both the upper and lower branches of the loop, with junctions a and b respectively. Because these currents begin and end at the same points, their phase differences from P to Q must be the same. The total phase difference will therefore be a sum of the phase differences across the junction, either δ_a or δ_b , and that over the path. We can use Equation 2.11 to write down the phase difference over the two paths first separately,

$$\Delta\theta_{PQ} = \delta_a - \frac{2e}{\hbar} \int_P^{Q, via a} \mathbf{A} \cdot d\mathbf{l}$$

$$\Delta\theta_{PQ} = \delta_b - \frac{2e}{\hbar} \int_P^{Q, via b} \mathbf{A} \cdot d\mathbf{l},$$

and then together since $\Delta\theta_{PQ}$ must be the same,

$$\delta_b - \delta_a = \frac{-2e}{\hbar} \oint \mathbf{A} \cdot d\mathbf{l} = \frac{-2e}{\hbar} \Phi. \quad (2.17)$$

Adding the two line integrals together produces the closed loop integral around Γ as before, so now the phase difference depends on the magnetic flux through the loop. We can write the phase differences relative to a δ_0 ,

$$\delta_a = \delta_0 + \frac{e}{\hbar} \Phi \text{ and } \delta_b = \delta_0 - \frac{e}{\hbar} \Phi,$$

and then the total current is

$$\begin{aligned} J_{total} = J_a + J_b &= J_0 \left(\sin \left(\delta_0 - \frac{e\Phi}{\hbar} \right) + \sin \left(\delta_0 + \frac{e\Phi}{\hbar} \right) \right) \\ &= 2J_0 \sin(\delta_0) \cos \left(\frac{e\Phi}{\hbar} \right). \end{aligned} \quad (2.18)$$

Irrespective of the value of δ_0 , $\sin(\delta_0)$ cannot be larger than 1. Therefore, the maximum current through the system is

$$J_{max} = 2J_0 \left| \cos \left(\frac{e\Phi}{\hbar} \right) \right|, \quad (2.19)$$

varying periodically as a function of the magnetic flux through the loop. Since Φ is an integer multiple of the tiny flux quantum, the output current is sensitive to minute changes in magnetic field. A loop with two Josephson junctions, and an area of 1cm^2 will have a maximum and minimum critical current separated by just 10^{-11}T (Hook and Hall 1974).

SQUID Operation

This quantum interference, where current through the system is modulated by an applied magnetic field is the basis for using SQUIDS as magnetometers. In most available MEG systems today, these SQUIDS operate at low temperatures (4.2K) and consist of a superconducting ring split by two Josephson junctions. A magnetic flux flowing through this ring results in a measurable voltage drop across the Josephson junctions (Zimmerman, Thiene, and Harding 1970).

To reduce the effects of noise on the SQUID measurements, the SQUIDS are placed further away from the head, and are inductively linked to a pick-up coil close to the head surface, see Figure 2.12A. A constant bias current is applied to the SQUID which is inductively linked to both the pick-up coil and a feedback

system. When a neuromagnetic field passes through the pick-up coil it produces a current which in turn creates a secondary magnetic field. This secondary field creates a magnetic flux through the SQUID ring and a voltage drop is detected.

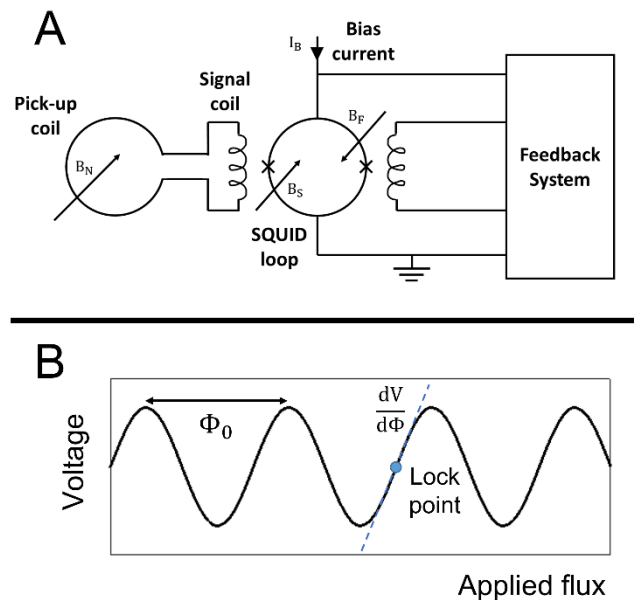


Figure 2.12: SQUID operation, with a schematic diagram of the sensor circuitry (A). A pick-up coil detects the neuromagnetic field (B_N) and is inductively linked (B_S) to the SQUID loop. The magnetic flux through this loop results in a voltage drop which is balanced by a feedback system (B_F). The measured variable is the voltage output of the feedback system. (B) When the constant bias current, I_B , exceeds a critical value, the flux-voltage transfer function is sinusoidal.

When the constant bias current exceeds a critical value, the voltage is a sinusoidal function of magnetic flux which means that many values of magnetic flux correspond to the same voltage, see Figure 2.12B. This ambiguity can be solved using a negative feedback system - an inductively linked circuit which produces an equal and opposite field to the signal coil so that the total flux through the SQUID ring is kept at a known constant value. This constant value is known as the 'lock point' and is at the steepest part of the function, when $dV/d\Phi$ is at its maximum value. It is therefore the voltage produced by the feedback system, needed to balance the flux through the SQUID loop, which is used as a measure of neuromagnetic field. SQUIDs operating in this way have a high sensitivity with a noise floor $< 10 \text{ fT}/\sqrt{\text{Hz}}$, a bandwidth of several kHz and a dynamic range on the order of nT.

The low operating temperature of SQUIDs require them to be kept submerged in a dewar of liquid helium. An insulating vacuum is therefore used to protect the scanning participant and operator, this places the coils several centimetres from the head surface (Figure 2.13B).

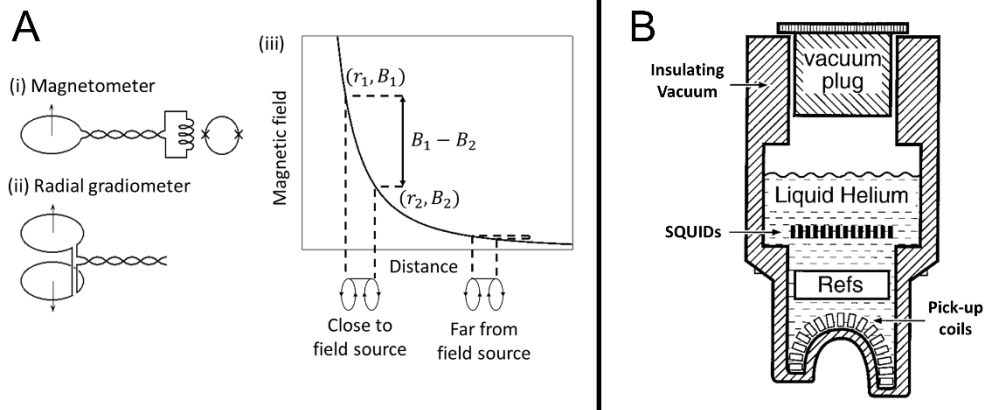


Figure 2.13: MEG hardware. (A) Variations in pick-up coil design. The axial gradiometer measures a net current proportional to the field gradient (A iii) and helps to reduce interference from field sources at some distance from the sensor. (B) The relative positions of the MEG hardware inside the dewar. Figures adapted from (Vrba and Robinson 2001).

2.3.2 Noise Reduction

Both hardware and software components of the MEG system can be used to effectively reduce interference from sources of magnetic field other than the brain.

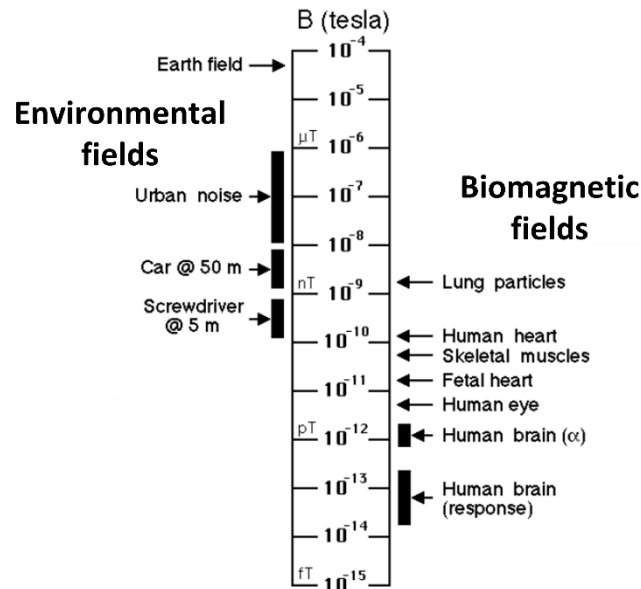


Figure 2.14: Magnetic fields which occlude MEG signals. These include both environmental field sources such as those produced by common lab equipment or nearby traffic, and biomagnetic fields produced by the research participant's other organs like their heart and eyes. Figure adapted from (Vrba 2000).

Hardwired Gradiometers

The design of the pick-up coil plays a large role in reducing interference from magnetic field sources far from the sensor (Figure 2.13A). The simplest coil design is a single loop magnetometer. A magnetic field creates a flux through the loop which in turn induces a current in the coupling coil, inductively coupling the pick-up coil with the SQUID. However, the pick-up coil will inevitably detect a combination of neuromagnetic field and noise from the environment. Given that the pick-up coil will only be several centimetres from the head, and that sources of magnetic interference will be much further away, a pick-up coil which is sensitive to near fields is desirable. This can be achieved with a hardwired gradiometer, and in fact the SQUID MEG systems used to acquire data in this thesis employ axial gradiometers which consist of two oppositely wound coils separated by a short distance. It is the difference in current induced in these two coils which is consequently measured. Since magnetic field drops off with distance, r , as $1/r^2$, the difference in field will be

much greater for close field sources than for distant interference sources. This is demonstrated in Figure 2.13A(iii).

Magnetically Shielded Rooms

Placing the MEG system in the centre of a magnetically shielded room (MSR) creates a magnetically quiet environment in which to conduct experiments. These rooms are generally built from layers of copper and mu-metal, where the mu-metal shields low frequency fields (0 to 10Hz) and the Eddy currents in the highly conductive copper layers provide attenuation of higher frequency fields (greater than 10Hz) (Vrba 2000).

The effectiveness of a magnetic shield can be calculated as the shielding factor

$$SF = \frac{|B_0|}{|B_S|}$$

where B_0 is the magnitude of the field at a given point, and B_S is the magnitude of the field at the same point after the shielding has been installed. Therefore, larger values for the shielding factor correspond to greater shielding, whereas values close to one correspond to little or no shielding.

Low frequency magnetic fields, less than 10Hz, are often referred to as DC fields and are largely shielded through a process called flux shunting. At a boundary between two materials with different relative permeabilities, such as mu-metal (with $\mu_r > 80,000$ (Jiles 1998)) and air ($\mu_r \approx 1$), both Ampère's law and Gauss's law must be obeyed. In other words, the tangential component of the magnetic field, \mathbf{H} , must be continuous across the interface, and the normal component of the field, \mathbf{B} , must also be continuous where $\mathbf{B} = \mu_0\mu_r\mathbf{H}$ inside the material and $\mathbf{B} = \mu_0\mathbf{H}$ in the air (Hoburg 1995). This results in the field on the air side of the interface being pulled towards the material almost perpendicular to the surface, then shunted through the material almost parallel to its surface, and released on the other side (Celozzi, Araneo, and Lovat 2008) as shown in Figure 2.15B.




Figure 2.15: Passive magnetic shielding mechanisms include Eddy current cancellation (A) and flux shunting (B). Figure adapted from (Celozzi, Araneo, and Lovat 2008).

High frequency magnetic fields, greater than 10Hz and often referred to as AC fields, are largely shielded by Eddy current cancellation which is why a highly conductive material is chosen (such as copper). Provided the magnetic field is time-varying, $\mathbf{B}(t)$, and the conductivity of the material, σ , is high, there is an induced electric field $\mathbf{E}(t)$, according to Faraday's law

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

The induced electric current density creates a magnetic field opposing the incident field, forcing it to be repelled and to run parallel to the surface of the shielding material (Celozzi, Araneo, and Lovat 2008). The induced current density decays exponentially into the material, with a characteristic decay length equal to the skin depth, δ . For an AC source with angular frequency $\omega = 2\pi f$, the skin depth can be calculated as in (Hoburg 1995),

$$\delta = \sqrt{\frac{2}{\omega \mu_0 \mu_r \sigma}}.$$

In copper, for example, where $\mu_r \approx 1$, the skin depth at 60Hz frequency is 8.5mm. Intuitively, one might expect to have shielding thickness greater than the skin depth for maximal shielding. However, this is not the case because using a shielding thickness much larger than the skin depth allows current loops to circulate through the shield, inducing flux densities inside the shielded area. For this reason, a good shielding factor can be achieved with shield thicknesses less than the skin depth.

Synthetic Gradiometers

Synthetic gradiometry follows this same principle of measuring field gradient to remove the contributions from interference sources far from the brain. In this case the difference in field is measured between a set of reference sensors (labelled 'Refs' in Figure 2.13B) and the pick-up coils surrounding the head.

The hardwired axial gradiometers in the SQUID MEG systems used for this thesis, as shown in Figure 2.13A, are first order gradiometers. Second order gradiometers can be built from two hardware first order gradiometers by calculating the displacement between individual gradiometer centres. This process of creating synthetic gradiometers can be expanded to higher orders and the CTF MEG system in Nottingham allows one to apply up to third order gradiometers.

These third order synthetic gradiometers may be applied during data acquisition or later during analysis.

Active Noise Cancellation

Reference sensors placed close to the head measure the background fields at this location (these are different to the reference sensors inside the dewar of liquid helium in a SQUID MEG system). An array of large biplanar coils surrounding the subject can then actively cancel out the measured field by applying an equal and opposite field. When used in conjunction with a magnetically shielded room, this field nulling technique can reduce background fields to less than 0.5nT (Holmes et al. 2018). Although this technique is not often used in a conventional SQUID MEG laboratory, it has become crucial in the development of new OPM-based MEG systems where the operation of the sensors requires near-zero background field (Holmes et al. 2018).

2.3.3 Experimental Procedure

Prior to scanning any participant, both safety and consent forms must be completed. Despite MEG being a passive scan, the safety form is important because the functional MEG data are typically mapped onto a structural MRI of the individual's brain. Due to the large magnetic field participants are

exposed to in an MRI scanner ($\geq 3T$), any metal inside the participant and within the vicinity of their head would make them unsuitable for scanning. Similarly, pacemakers and other implants are a contraindication of MRI.

The MEG laboratory consists of a control room, containing the SQUID electronics and the acquisition and stimulus presentation computers, as well as the MEG system housed inside a magnetically shielded room, see Figure 2.16. A participant is positioned with their head in the helmet of the MEG dewar and can be seated or supine. Fiducial coils are placed on the participant at the nasion and right and left preauricular points to allow for head localisation (so the position of the coils relative to the sensors is known) and coregistration of the MEG functional data with MRI structural data at a later stage. These coils can be energised throughout the experiment to track head motion relative to sensors, or at the beginning and end of an experimental run. The participant must remain completely still throughout the experiment and a motion tolerance limit can be set to ensure this is the case. Because the boiloff of the liquid helium in the dewar carries a small risk of asphyxiation, an oxygen alarm is placed inside the shielded room with the participant. The door to the MSR is then closed and the participant can be communicated with via an intercom and video monitoring system.

Data are processed by the SQUID electronics before being sent to the acquisition computer. The stimulus computer is used to interact with the participant and can project visual, auditory and tactile stimuli, as well as acquire behavioural measures from the participant (using, for example, button-press responses or an eye tracker). The exact timings of stimulus presentation and participant responses can be passed to the acquisition computer to place temporal markers within the MEG data.




Figure 2.16: The MEG lab. The participant is scanned seated or supine in the MSR with their head inside the helmet of the dewar. The scanner operator can communicate with them via an intercom and a range of stimuli can be presented. This set-up allows both neurophysiological and behavioural data to be acquired during an experiment. Figure adapted from (Vrba 2000).

Some of the data in this thesis were acquired during the COVID-19 pandemic and as such the experimental procedure was altered to reduce the risk of a participant or scanner operator contracting the virus. These precautions are described in greater detail in Chapter 5.

The data analysed in this thesis were acquired in two different settings: the CTF-MEG lab in the Sir Peter Mansfield Imaging Centre at Nottingham (Chapters 4 and 5), and the CTF-MEG lab in the Children’s Hospital of Philadelphia USA (Chapter 6).

2.3.4 Coregistration with a structural MRI

It is possible to map functional MEG data onto a structural MRI scan so that the source of the MEG signal can be localised (this process will be described in detail in Chapter 3). For the greatest spatial accuracy of source localisation, the position of the participant's brain relative to the MEG sensors should be known.

The first step is to track the location of the head inside a CTF-MEG system using three coils, each placed at a fiducial location (the nasion and left and right preauricular points). These coils are energised at non-physiologic frequencies so their positions inside the MEG helmet can be found. The location of those same coils must then be found relative to the participant's MRI. There are various ways of doing this and owing to the fact that the data used in this thesis were acquired in different labs and at different times, three different coregistration methods were employed.

Polhemus head digitisation

For the data used in Chapter 4 a Polhemus FASTRACK digitiser (Polhemus, Colchester, VT, USA) was used to create a 3D representation of the surface of each participant's head along with the relative fiducial coil locations. This involves seating the participant in front of a transmitter box and placing a receiver on top of their head. Then a stylus, whose position is known relative to the receiver, is used to draw the surface of the head and face. An anatomical T_1 -weighted MRI is acquired using an MPRAGE sequence at an isotropic resolution of 1mm^3 . Note that this MRI must be acquired after the MEG scan because of long-lived tissue magnetisation effects at high fields. A surface matching algorithm then maps this 3D head digitisation to the scalp surface extracted from the participant's MRI using an iterative closest point algorithm (Besl and McKay 1992). The position of the fiducial coils is therefore known relative to the MRI, allowing complete coregistration of the MEG sensors with the brain.

Skanect head digitisation

During the COVID-19 pandemic, the decision was made to switch from using the Polhemus digitiser (which requires close contact with a participant for roughly 10 minutes) to an optical imaging system which could be used at a distance to reduce the risk of virus transmission. A Structure Sensor camera (Occipital Inc., San Francisco, CA, USA) was used, together with an iPad operating Structure Sensor Skanect software to produce a 3D model of each participant's head. It does this by using an infrared laser projector to place dots on the surfaces in its field of view which are then detected by an infrared camera to determine the shape and distance of any objects. A visible light camera is also used so that colour can be mapped onto the 3D shape. The acquisition takes approximately 3 minutes.

First, green stickers are placed on the fiducial coils so that a colour threshold can be used to determine the coil positions relative to the head surface. A swimming cap is then used to flatten down any hair so that a true head shape is obtained, and the Skanect image is taken. An MRI is acquired exactly as described in the previous section, and the head surface from the Skanect is matched with the scalp surface extracted from the MRI using the same surface matching algorithm.

MRI contrast markers

Because the data in Chapter 6 were acquired in a hospital setting, the MEG scan and MRI were done in the same session. The coregistration procedure is then greatly simplified: first a pen is used to mark the location of the MEG fiducial coils on the skin so that after the MEG scan, patients can make their way to the MRI where contrast markers are placed at those same fiducial points. This completes the coregistration of the fiducial markers from the MRI and MEG enabling spatial mapping of the MEG sensors to brain anatomy.

Summary

The first part of this chapter described the microscopic origins of the MEG signal, from the basic biology of a neuron to the chemical and electrical exchanges underpinning neurotransmission. We have learned that it is the

post-synaptic potentials in the dendrites of tens of thousands of pyramidal neurons in the cortex which result in measurable MEG activity. From this signal, depending on the type of experiment and the way the data are analysed, several measurable effects can be derived. These include spontaneous rhythms, evoked responses, induced effects, and functional connectivity. The latter part of the chapter detailed MEG data acquisition using conventional SQUID sensors. The experimental procedure was described as well as coregistration with an MRI and the noise reduction techniques allowing us to measure the magnetic fields induced by current flow in the brain despite surrounding sources of magnetic noise being many orders of magnitude larger.

The next chapter will cover the analysis of MEG data. This includes mathematical models for transforming sensor data into a 3D source space so that we can estimate the location of MEG sources inside the brain, as well as the theory underpinning hidden Markov modelling.

Chapter 3

MEG Data Analysis

Large sensor arrays and high temporal resolution means that MEG data are rich in spatial, temporal and spectral information. In the first instance, the neuromagnetic fields from all brain regions superpose meaning that the same sensors can detect fields from multiple regions. These sensor-level data are extremely useful, yielding topographic plots of neural activity in a similar fashion to those used in most EEG systems. However, by exploiting the large number of sensors surrounding the head, and with the use of mathematical models, it is possible to reconstruct the source signal at each location in the brain. The first two sections of this chapter are concerned with source reconstruction.

In the first section, a model of the field measured at the sensors for a known source location and signal amplitude is proposed. Section 3.2 uses this ‘forward’ model, together with the measured MEG data to solve the ‘inverse’ problem using a technique called beamforming. This produces an estimate of the neuromagnetic signal arising at each brain location.

The final part of this chapter, Section 3.3, describes the theory behind hidden Markov modelling - a statistical timeseries analysis which uncovers recurrent patterns of activity in data and is applied to MEG throughout this thesis.

3.1 The Forward Problem

Before we can localise the source of the magnetic fields measured at the MEG sensors, we must first understand what the magnetic field would look like (at the sensors) from a known current distribution in the brain. This is known as the MEG forward problem.

3.1.1 Estimating the Lead Fields

Any consideration of electromagnetism must begin with Maxwell’s equations. These equations form the basis of classical electromagnetism and by bringing

them together, Maxwell described the occurrence of magnetism, electricity, and light in a single unified theory (Maxwell 1873).

Gauss' law for electricity

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\varepsilon} \quad (3.1)$$

Gauss' law for magnetism

$$\nabla \cdot \mathbf{B} = 0 \quad (3.2)$$

Faraday's law of induction

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \quad (3.3)$$

Ampère's law

$$\nabla \times \mathbf{B} = \mu_0 \left(\mathbf{J} + \frac{\partial \mathbf{E}}{\partial t} \right) \quad (3.4)$$

Here, \mathbf{E} and \mathbf{B} are the electric and magnetic fields respectively, ρ is the charge density, t is time, and ε and μ are the permittivity and permeability of a medium respectively. Note that the permeability of tissue in the head can be assumed to be that of free space, μ_0 . \mathbf{J} is the total current density. Since neuroelectromagnetic fields vary in time at low frequencies (<100Hz) they can be described as quasi-static (i.e. $\partial x/\partial t = 0$) (Hämäläinen et al. 1993) and the last two Maxwell equations can be expressed thus

$$\nabla \times \mathbf{E} = 0 \quad (3.5)$$

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J}. \quad (3.6)$$

In 1987, Sarvas applied basic mathematical concepts to model both the biomagnetic forward and inverse problems. The following maths is based on his calculations (Sarvas 1987). Using the Biot-Savart law, the magnetic field $\mathbf{B}(\mathbf{r})$, measured at a location \mathbf{r} is

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi} \int_G \mathbf{J}(\mathbf{r}') \times \frac{\mathbf{r}-\mathbf{r}'}{|\mathbf{r}-\mathbf{r}'|^3} dv', \quad (3.7)$$

where \mathbf{r}' is a location inside volume G , and dv' is a volume element. Recall from Chapter 2 that the total current density is made up of both a primary current, \mathbf{J}_p , (inside the dendrites of pyramidal neurons) and a volume current, \mathbf{J}_v , made up of the displaced ions outside of the neuron in the extracellular

space flowing towards the area of repolarisation. Here we will consider the contributions of both current sources

$$\mathbf{J} = \mathbf{J}_p + \mathbf{J}_v. \quad (3.8)$$

Since the volume current is an ohmic current (charge carriers in a conducting medium responding to changes in the local electric field), it can be written as $\mathbf{J}_v = \sigma \mathbf{E}$ where σ is the conductivity of tissue which we will assume to be uniform inside G . Since the electric field and electric potentials are related by $\mathbf{E} = -\nabla V$, we can substitute the total current density into Equation 3.7 to give

$$\begin{aligned} \mathbf{B}(\mathbf{r}) &= \frac{\mu_0}{4\pi} \int_G (\mathbf{J}_p(\mathbf{r}') - \sigma \nabla V(\mathbf{r}')) \times \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} dv' \\ &= \frac{\mu_0}{4\pi} \int_G \mathbf{J}_p(\mathbf{r}') \times \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} dv' - \frac{\mu_0 \sigma}{4\pi} \int_G \nabla V(\mathbf{r}') \times \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} dv', \end{aligned} \quad (3.9)$$

where V is the electric potential and we now have a separate term for the primary and volume current contributions.

Now, by focussing on the volume current term specifically,

$$\mathbf{B}_v(\mathbf{r}) = -\frac{\mu_0 \sigma}{4\pi} \int_G \nabla V(\mathbf{r}') \times \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} dv',$$

and by using the vector identity $\nabla a \times \nabla b = \nabla \times (a \nabla b)$, where $\nabla a = \nabla V$ and $\nabla b = \mathbf{r} - \mathbf{r}' / |\mathbf{r} - \mathbf{r}'|^3$, we can express the volume current thus

$$\mathbf{B}_v(\mathbf{r}) = -\frac{\mu_0 \sigma}{4\pi} \int_G \nabla \times V(\mathbf{r}') \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} dv'. \quad (3.10)$$

Next, we recall the generalised Stokes' theorem, linking the volume integral to a surface integral:

$$\int_G \nabla \times \mathbf{A} dv' = \int_S \mathbf{n} \times \mathbf{A} ds \quad (3.11)$$

where \mathbf{n} is a vector normal to the surface element, ds , pointing away from the volume centre. This allows us to express the volume currents mathematically

as if they were flowing on a conductive surface. This bears no semblance to reality but is a useful mathematical formalism,

$$\mathbf{B}_v(\mathbf{r}) = \frac{-\mu_0}{4\pi} \int_S V(\mathbf{r}') \mathbf{n}(\mathbf{r}') \times \frac{\mathbf{r}-\mathbf{r}'}{|\mathbf{r}-\mathbf{r}'|^3} ds. \quad (3.12)$$

Now including the primary current term as \mathbf{B}_0 , the magnetic field induced by current flow within the dendrites of the neurons, the total magnetic field produced by this current density is

$$\mathbf{B}(\mathbf{r}) = \mathbf{B}_0(\mathbf{r}) - \frac{\mu_0\sigma}{4\pi} \int_S V(\mathbf{r}') \mathbf{n}(\mathbf{r}') \times \frac{\mathbf{r}-\mathbf{r}'}{|\mathbf{r}-\mathbf{r}'|^3} ds. \quad (3.13)$$

This useful interpretation of volume currents as surface currents flowing on the boundary between regions with different conductivities was introduced by Geselowitz in 1970 when he addressed the interpretation of magnetocardiogram results (Geselowitz 1970). It is for this reason that Equation 3.13 is referred to as the Geselowitz formula.

3.1.2 Single Sphere Model

It is at this point that we introduce a head model to address the forward problem. The simplest model of the head is as a spherical conductor with volume G , surface S and uniform conductivity σ as shown in Figure 3.1. We also assume that the MEG system only detects radial components of the magnetic field⁵.

⁵ In most conventional cryogenic MEG systems, the sensors are arranged around the head and measure along one axis, but they are not radial to a sphere and so, even with this assumption, a vector expression for field is needed. New OPM-MEG systems have the capability of making triaxial measurements which should improve spatial filtering accuracy (Brookes, Boto et al. 2021).

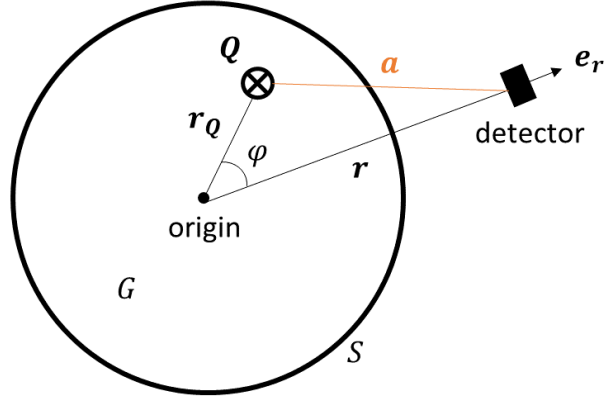


Figure 3.1: A diagram of the single sphere head model. The head is modelled as a homogeneous conductor with volume G , bounded by spherical surface S . The fields produced by a single dipole with magnitude and direction \mathbf{Q} , and detected at a sensor with position \mathbf{r} from the origin of the sphere, are described by the Sarvas equation (Equation 3.28).

Taking this into consideration, Equation 3.13 becomes

$$B_r(\mathbf{r}) = \mathbf{B}(\mathbf{r}) \cdot \mathbf{e}_r$$

$$B_r(\mathbf{r}) = \frac{\mu_0}{4\pi} \int_G \mathbf{J}_p(\mathbf{r}') \times \frac{\mathbf{r}-\mathbf{r}'}{|\mathbf{r}-\mathbf{r}'|^3} \cdot \mathbf{e}_r dV' - \frac{\mu_0\sigma}{4\pi} \int_S V(\mathbf{r}') \mathbf{n}(\mathbf{r}') \times \frac{\mathbf{r}-\mathbf{r}'}{|\mathbf{r}-\mathbf{r}'|^3} \cdot \mathbf{e}_r dS. \quad (3.14)$$

Now, we can simplify this equation by analysing just the volume current term which is

$$\mathbf{n}(\mathbf{r}') \times (\mathbf{r} - \mathbf{r}') \cdot \mathbf{e}_r = (\mathbf{n}(\mathbf{r}') \times \mathbf{r} - \mathbf{n}(\mathbf{r}') \times \mathbf{r}') \cdot \mathbf{e}_r. \quad (3.15)$$

Since the cross product of parallel vectors is zero, $\mathbf{n}(\mathbf{r}') \times \mathbf{r}' = 0$, and the dot product of perpendicular vectors is zero, $(\mathbf{n}(\mathbf{r}') \times \mathbf{r}) \cdot \mathbf{e}_r = 0$. This term therefore vanishes, removing the contribution from the volume currents altogether! We see that

$$B_r(\mathbf{r}) = \frac{\mu_0}{4\pi} \int_G \mathbf{J}_p(\mathbf{r}') \times \frac{\mathbf{r}-\mathbf{r}'}{|\mathbf{r}-\mathbf{r}'|^3} \cdot \mathbf{e}_r dV'. \quad (3.16)$$

Importantly, this means that B_r is not dependent on tissue conductivity.

So far, we have assumed that the current is evenly distributed throughout the volume G , now we will consider the current from a dipole which exists at a single point, \mathbf{r}_Q , within the volume. We can therefore write

$$\mathbf{J}_p(\mathbf{r}') = \mathbf{Q}\delta(\mathbf{r}' - \mathbf{r}_Q), \quad (3.17)$$

where \mathbf{Q} is the magnitude and direction of the current dipole and $\delta(\mathbf{r})$ is the Dirac delta function. Equation 3.16 then becomes

$$B_r(\mathbf{r}) = \frac{\mu_0}{4\pi} \int_G \mathbf{Q}\delta(\mathbf{r}' - \mathbf{r}_Q) \times \frac{\mathbf{r} - \mathbf{r}'}{|\mathbf{r} - \mathbf{r}'|^3} \cdot \mathbf{e}_r dv' \quad (3.18)$$

and the remaining integral can be simplified since

$$\int_{-\infty}^{\infty} \delta(x) dx = \begin{cases} 0, & x \neq 0 \\ 1, & x = 0 \end{cases}$$

so that Equation 3.18 becomes

$$B_r(\mathbf{r}) = \frac{\mu_0}{4\pi} \mathbf{Q} \times \frac{\mathbf{r} - \mathbf{r}_Q}{|\mathbf{r} - \mathbf{r}_Q|^3} \cdot \mathbf{e}_r dv' \quad (3.19)$$

when $\mathbf{r}' = \mathbf{r}_Q$. Because the sensors are not radially oriented to a sphere we must also calculate the non-radial components of the field outside of the spherical conductor using Maxwell's equations (Equation 3.6). Since current density \mathbf{J} is zero outside of volume G , $\nabla \times \mathbf{B} = 0$. This means that we can express the magnetic field in terms of a magnetic scalar potential (which we will call U),

$$\mathbf{B}(\mathbf{r}) = -\mu_0 \nabla U(\mathbf{r}). \quad (3.20)$$

To find an expression for $U(\mathbf{r})$ we fix \mathbf{r} outside of G and consider a line integral along the radius of G , $\mathbf{r} + t\mathbf{e}_r$, where $0 \leq t \leq \infty$. Since U vanishes at infinity,

$$U(\mathbf{r}) = - \int_0^{\infty} \nabla U(\mathbf{r} + t\mathbf{e}_r) \cdot \mathbf{e}_r dt. \quad (3.21)$$

Using Equation 3.20

$$U(\mathbf{r}) = \frac{1}{\mu_0} \int_0^{\infty} B_r(\mathbf{r} + t\mathbf{e}_r) \cdot \mathbf{e}_r dt, \quad (3.22)$$

and substituting in the expression for B_r from Equation 3.19 produces

$$U(\mathbf{r}) = \frac{1}{4\pi} \mathbf{Q} \times (\mathbf{r} - \mathbf{r}_Q) \cdot \mathbf{e}_r \int_0^{\infty} \frac{dt}{|\mathbf{r} + t\mathbf{e}_r - \mathbf{r}_Q|^3}. \quad (3.23)$$

Solving this integral produces the result

$$U(\mathbf{r}) = -\frac{1}{4\pi} \frac{\mathbf{Q} \times \mathbf{r}_Q \cdot \mathbf{r}}{F} \quad (3.24)$$

where $F = |\mathbf{a}|(|\mathbf{r}||\mathbf{a}| + |\mathbf{r}|^2 - \mathbf{r}_Q \cdot \mathbf{r})$ and $\mathbf{a} = \mathbf{r} - \mathbf{r}_Q$. Using Equations 3.20 and 3.24 we obtain the Sarvas equation (Sarvas 1987) which is the general solution for the single sphere model:

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi F^2} (F\mathbf{Q} \times \mathbf{r}_Q - \mathbf{Q} \times \mathbf{r}_Q \cdot \mathbf{r} \nabla F) \quad (3.25)$$

where $\nabla F = \left(\frac{|\mathbf{a}|^2}{|\mathbf{r}|} + \frac{\mathbf{a}}{|\mathbf{a}|} \cdot \mathbf{r} + 2|\mathbf{a}| + 2|\mathbf{r}| \right) \mathbf{r} - \left(|\mathbf{a}| + 2|\mathbf{r}| + \frac{\mathbf{a}}{|\mathbf{a}|} \cdot \mathbf{r} \right) \mathbf{r}_Q$. This equation tells us that radially oriented sources in a spherical conductor cannot be detected because their component of the magnetic field vanishes. However, a study completed in 2002 showed that source depth rather than orientation is the main factor compromising MEG sensitivity and most neuronal ensembles have detectable tangential components (Hillebrand and Barnes 2002).

3.1.3 Alternative Head Models

The approximation of the head as a single sphere produces a useful estimate for the forward model. However, it is an oversimplification of head geometry and is therefore inadequate for regions of the brain which deviate the most from a spherical shape (for example, the frontal and fronto-temporal areas) (Hämäläinen and Sarvas 1989). To overcome this problem, much work has been done to extract realistic head models from the anatomical MRIs of research participants with multiple layers of homogeneous conductivity separated by boundaries (Hämäläinen and Sarvas 1989; Meijs et al. 1987), with the forward problem being solved numerically using the Geselowitz formula. The most accurate model to-date is a three shell boundary element model (BEM), where the brain, skull and scalp make up each layer (Stenroos, Hunold, and Haueisen 2014). However, these models are computationally expensive so an alternative head model based on multiple overlapping spheres was introduced (Huang, Mosher, and Leahy 1999) with a fraction of the computational cost and improved accuracy compared with a single sphere.

In this method, for each individual MEG sensor a local sphere is fitted to the curvature of the skull closest to the sensor. The head shape is then modelled as a set of overlapping spheres, see Figure 3.2, which is more realistic than a single sphere.




Figure 3.2: A diagram explaining the multiple local spheres head model from (Huang, Mosher, and Leahy 1999). A local sphere is fitted for each sensor and the complete set of spheres forms a more geometrically accurate model of the head than a single sphere approximation.

Although this makes sense intuitively, the bigger problem is defining which parts of the skull are 'local' to each sensor. We could define a tiny part of the skull immediately beneath the sensor as local, or at the other extreme the whole skull could be defined as 'local' which would result in the same single-sphere model depicted in Figure 3.1. Huang et al. solved this problem by first choosing an initial estimate for the sphere centre and radius, and then minimising the cost function based on the difference between the secondary currents arising from a realistic mesh of the innermost skull surface and the local sphere.

A multiple local spheres model was used throughout this thesis because of its improved accuracy compared with the single sphere model and improved computational speed compared with a BEM. However, it is worth noting that increased computing power is now more readily available, so boundary element methods are becoming more popular, especially as software such as FieldTrip mean it can be readily implemented. This makes it likely that future studies will utilise these.

3.2 The Inverse Problem

Now that we have a good mathematical foundation for the forward model, we can use this to estimate the spatial and temporal signatures of the neuronal sources which produce the magnetic fields measured by our MEG sensors, often referred to as the biomagnetic inverse problem.

3.2.1 A Generative Model of Sources

A generative model of sources posits that the magnetic fields measured at the sensors are the superposition of all the fields produced by all locations in the brain. This can be written mathematically as

$$\mathbf{b}(t) = \int_V \mathbf{l}_\theta q_\theta(t) dV \quad (3.26)$$

where \mathbf{b} is an N-element vector, each element of which corresponds to a MEG sensor measuring the superposition of all the fields generated by dipoles across the whole brain volume, V . Mapping these dipoles to the measured fields are the forward fields, \mathbf{l}_θ for a dipole with unit strength and with position and orientation determined by θ (i.e., $\theta = [\mathbf{r}', \delta]$ where \mathbf{r}' is a vector location within V and δ relates to orientation). Note that the lead fields and the source orientation are assumed to be time independent. It is only the dipole strength $q_\theta(t)$ which varies with time, t .

By parsing the brain into a discrete number of voxels (M), and by assuming that there is only one dipole per voxel, we can write this as a sum

$$\mathbf{b}(t) = \sum_{m=1}^M \mathbf{l}_{\theta_m} q_{\theta_m}(t), \quad (3.27)$$

or in matrix formalism

$$\mathbf{b}(t) = \mathbf{L} \cdot \mathbf{q}(t), \quad (3.28)$$

where \mathbf{b} is still an N-element column vector (N being the total number of MEG sensors), \mathbf{L} is an $N \times M$ lead field matrix, and \mathbf{q} is an $M \times 1$ matrix containing the instantaneous dipole strengths, so the columns of \mathbf{L} give the contribution of each dipole to all sensors and the rows of \mathbf{L} give the dipoles that contribute to the signal at a single sensor.

Intuitively, the inverse problem can be solved by inverting our generative model to obtain $\mathbf{q}_{\theta_m}(t)$. This would seem simple given that we know our measured fields, $\mathbf{b}(t)$, and our lead field matrix, \mathbf{L} , based on the forward model. However, the inverse problem is mathematically ill-posed. This means that it is theoretically impossible to solve. From a simple physics point of view, this is because an infinite number of current distributions could result in the same measured fields. This is, in part, due to field cancellation from overlapping source fields, as well as the presence of ‘silent sources’ such as radially oriented dipoles or closed loops. Mathematically speaking, there are many more potential source locations than MEG sensors - since $M \gg N$, \mathbf{L} is not square, and even if it was the linear dependence between the columns of \mathbf{L} mean that \mathbf{L}^{-1} is impossible to compute.

It is, however, possible to *estimate* a solution to the biomagnetic inverse problem under a specific set of assumptions which construct the inverse solution. This can be done a number of ways, methods include minimum norm estimation (Hämäläinen and Ilmoniemi 1994), equivalent current dipole models (Brenner et al. 1978; Tuomisto et al. 1983) and linearly constrained minimum variance (LCMV) beamformers (Van Veen et al. 1997). Each of these methods enlist their own assumptions which can be explored in more detail in (Hämäläinen et al. 1993; Mosher, Baillet, and Leahy 2003; Hillebrand and Barnes 2002). Because of its proficiency at reconstructing spontaneous (oscillatory) activity, and its ability to minimise biomagnetic and environmental interference sources, the method used in this thesis is LCMV beamforming.

3.2.2 Beamforming

The beamformer is a type of spatial filter which assigns a weighting to each MEG sensor in such a way that the combination of all sensors and their associated weights produces an estimate of source strength at every location in the brain. This can be represented as

$$\hat{\mathbf{q}}_{\theta}(t) = \mathbf{W}_{\theta}^T \mathbf{b}(t) \quad (3.29)$$

where $\hat{\mathbf{q}}_{\theta}$ is the source strength estimate at location and orientation θ , and the weights, \mathbf{W}_{θ} , are an $N \times 1$ matrix applied to the sensor measurements $\mathbf{b}(t)$ from all N MEG sensors.

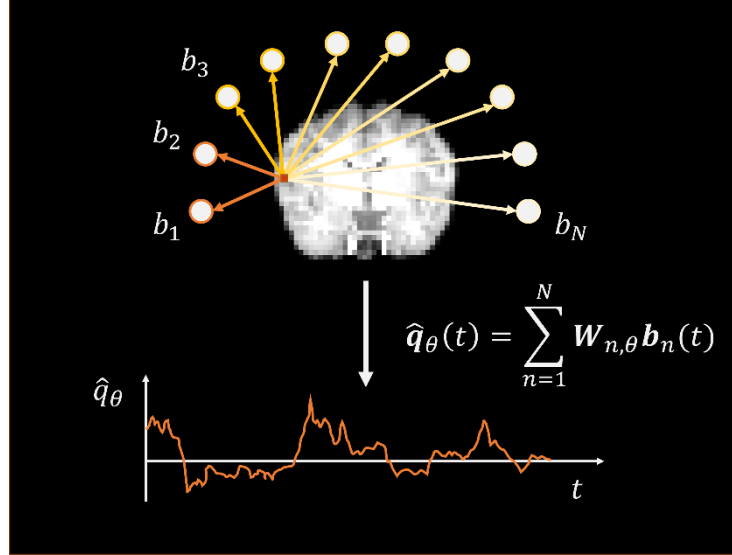


Figure 3.3: A schematic diagram describing how beamformer source reconstruction works. The brain is split into M voxels, and the signal at each location \hat{q}_{θ} is the sum over N MEG sensor measurements. The contribution from each sensor is determined by a set of weights. The different weightings for each MEG sensor are represented here by varying shades of colour and an example reconstructed time course is presented (bottom).

The weights are chosen so that the signal at a particular location is retained, but the contribution of signals from the rest of the brain (and surrounding environment) are minimised. This is done using variance minimisation

$$\min_{\mathbf{W}_{\theta}} |\langle \hat{\mathbf{q}}_{\theta}^2(t) \rangle| \quad \text{subject to } \mathbf{W}_{\theta}^T \mathbf{l}_{\theta} = 1, \quad (3.30)$$

where the angular brackets represent the expectation value. The variance is minimised subject to the linear constraint that at θ , the weights (transposed) multiplied by the lead fields specific to that location and orientation, \mathbf{l}_{θ} , remains at 1, which makes sense given that the lead fields are fitted to a unit dipole. Equation 3.30 can be solved thus (Van Veen et al. 1997)

$$\mathbf{W}_{\theta}^T = \frac{\mathbf{l}_{\theta}^T \mathbf{C}^{-1}}{\mathbf{l}_{\theta}^T \mathbf{C}^{-1} \mathbf{l}_{\theta}} \quad (3.31)$$

where $\mathbf{C} = \langle \mathbf{b}(t) \mathbf{b}(t)^T \rangle$ and the brackets $\langle x \rangle$ indicate the temporal expectation value; in practice this is approximated by the data covariance

between channels (producing an $N \times N$ matrix). A single element of the covariance matrix is calculated as

$$C_{i,j} = \frac{1}{P} \sum_{p=1}^P b_i(t_p) b_j(t_p), \quad (3.32)$$

which means that the i, j^{th} element of the matrix contains the covariance between data measured at channels i and j . P is the total number of points in the time window in which covariance is calculated. To calculate the beamformer weights we must invert the covariance matrix. However, in practice, when \mathbf{C} is close to singular, the matrix must first be regularised. This is typically achieved using the Tikhonov method so that

$$\mathbf{C}_r = \mathbf{C} + \mu \mathbf{\Sigma}, \quad (3.33)$$

where \mathbf{C}_r is the regularised matrix, μ is the regularisation parameter and $\mathbf{\Sigma}$ is approximated as the identity matrix, \mathbf{I} .

An implicit assumption of the beamformer is that spatially separate sources are not temporally correlated. This is because two distal sources which are temporally correlated would produce two dipolar fields outside the head, which would then be measurable at the MEG sensors. Since the lead fields are modelled on a single unit dipole, those spatially separate but correlated sources would be minimised according to equation 3.30.

3.2.3 Source Orientation

Both the location and orientation of the dipolar source will greatly impact on the measured signals. We have represented the source location and orientation throughout these calculations in the vector $\boldsymbol{\theta}$, but we have not yet calculated it. This is done by finding the orientation which produces the maximum signal-to-noise ratio (SNR) using an exhaustive search. This involves calculating the SNR for every possible dipole orientation and then selecting the orientation with the maximum value. Intuitively this seems like a computationally intensive step, but because our MEG sensors are unable to measure radially-oriented dipoles (see Section 3.1.2) we can restrict our search

to angles azimuthal to the brain surface. A pseudo-Z-statistic, psZ , is calculated for each location, r , and azimuthal angle, φ , so that

$$psZ_{r,\varphi} = \frac{W_{r,\varphi}^T C W_{r,\varphi}}{W_{r,\varphi}^T \Sigma W_{r,\varphi}}. \quad (3.34)$$

The orientation φ is then extracted from the maximum value of $psZ_{r,\varphi}$.

3.2.4 Depth Correction

A peculiarity of beamforming is that the source amplitudes are greatest at the centre of the brain. This is because the L2 norm of the lead fields falls towards the centre of the brain (because of the increased distance from MEG sensors), therefore increasing the norm of the beamformer weights for those deeper sources (Mosher, Baillet, and Leahy 2003). This artificial skewing of the data towards deeper sources can be counteracted by normalising the result to produce a pseudo-Z-statistic (psZ). Introducing a normalisation factor to Equation 3.29 gives us

$$\hat{q}_{psZ} = \frac{W^T}{\sqrt{W^T W}} \mathbf{b}(t), \quad (3.35)$$

where \hat{q}_{psZ} is the new estimate of the data with arbitrary units (rather than units of Am) because it has been normalised. By substituting Equation 3.31 for W^T into Equation 3.35, the final source estimate can be given by

$$\hat{q}_{psZ} = \frac{l_{\theta}^T C^{-1}}{\sqrt{l_{\theta}^T C^{-2} l_{\theta}}} \mathbf{b}(t). \quad (3.36)$$

This pseudo-Z-statistic can then be applied to each voxel sequentially, resulting in a statistical 3D image of source strength over the whole brain.

3.2.5 Beamformer Artefact Rejection

The beamformer acts as a statistical ‘lens’ focussing on one location and minimising the contribution from all other locations (both inside the brain or outside of it). This means the beamformer acts as an effective artefact rejection method (Brookes, Hale, et al. 2011). This is especially useful for sources inside or close to the head because the contribution from fields produced far from the head are already minimised by using gradiometers (see

Section 2.3.2). Ocular and muscle artefacts are reduced, and the contribution from other brain areas is also minimised (Sekihara, Hild, and Nagarajan 2006). Activity related to the participant's heartbeat is also effectively removed, whether that is the magnetocardiogram (O'Neill, Barratt, et al. 2015) or the ballistocardiogram (Brookes, Mullinger, et al. 2008).

3.2.6 Parcellation

Specific anatomical regions of the brain are responsible for different functions, for example the occipital lobe is largely associated with visual processing and the postcentral gyrus is where sensory information is registered. In some studies, we are more interested in how these larger brain regions interact with a task or with one another than we are in each voxel of the brain individually. This is of particular importance when calculating functional connectivity – dividing the cortex into larger regions requires significantly less time and computational power than voxel to voxel connectivity. It is therefore important to have a consistent way of splitting the brain into functionally relevant parcellations.

A Standard Brain

When we are comparing the localisation results from multiple participants in a study, it is important to place the results in a standard space before group analyses can be carried out. There are different standard templates to compare one's results to, but arguably the most well-known and utilised standard is the MNI152 which is the average of structural MRI images from 152 healthy young adults, composed by the Montreal Neurological Institute as part of the International Consortium for Brain Mapping (ICBM) (Mazziotta et al. 2001). Aligning any brain to this standard provides an objective way of defining coordinates which does not rely on the subjective description of anatomical location (Jenkinson and Chappell 2018). Ideally a standard brain would be used to define our parcellations, however, the natural variation in cortical features

across 152 people means that the average brain lacks the clear spatial definition needed to split the brain into useful regions (see Figure 3.4).




Figure 3.4: The MNI152 standard brain was composed by the Montreal Neurological institute and is the average over 152 MRIs from young healthy adults. It provides a standard space for comparing functional activity localisations. The process of averaging has removed fine detail from the standard, so it is not suitable for delineating parcellations.

To get around this problem, a high-resolution single-subject brain can be used to define the parcellations which can then be transformed into standard space so that comparison between regions from different subjects is possible. This method is used to define the Automated Anatomical Labelling (AAL) atlas (Tzourio-Mazoyer et al. 2002) which was the parcellation method of choice for this thesis.

AAL Atlas

This atlas is based on the MNI single subject brain from a healthy young adult male whose brain was scanned 27 times using a T1-weighted gradient echo sequence at 1mm^3 resolution. Tzourio-Mazoyer and colleagues tracked the sulci using a 3D surface rendering and manually drew them onto the render (see Figure 3.5A). Regions of interest in 2D space were then identified manually followed by a nearest-neighbours connectivity algorithm which joined the regions of interest across slices to produce 90 3D anatomical volumes of interest (45 in each hemisphere). The 2D regions of interest are shown in Figure 3.5B. The MNI single subject brain was then linearly transformed into the MNI152 space, thus providing an atlas mapping out functionally useful regions in a standard space (Tzourio-Mazoyer et al. 2002).




Figure 3.5: The AAL atlas was created by manual segmentation of the sulci (A) and regions of interest (B) on the MNI single subject brain. The lateral and parasagittal medial view of the sulci are shown in A(i) and A(ii) respectively. The 2D regions of interest for 4 different axial slices is shown in (B), the number in the bottom right corner indicates the stereotaxic z-coordinate in mm. This atlas can then be aligned with the MNI152 brain so it can be used in standard space to parcellate other single-subject brains. Figure adapted from (Tzourio-Mazoyer et al. 2002).

This method of parcellation has several advantages over the Talairach atlas (Talairach and Tournoux 1988), which was previously the atlas of choice. First, the Talairach atlas is based on the post-mortem examination of a sixty-year-old female brain. Because of its age, the brain showed some cortical atrophy. Only the left hemisphere was examined so the right hemisphere is assumed to be a mirror image. There is also no MRI counterpart to this brain making it difficult to align to as a standard. There will also be clear differences between an ex-vivo brain and a healthy in-vivo one (like the AAL atlas).

The AAL atlas is typically used in MEG for defining cortical regions for functional connectivity analyses (Tewarie et al. 2014; Brookes et al. 2016; Dauwan et al. 2019) where it has yielded useful results. However, there are limitations to this atlas. For example, it is based on the manual segmentation of a single subject's brain and does not consider the inevitable differences between healthy brains. It is also segmented based solely on anatomical landmarks rather than including information about function and connectivity like the Human Connectome Project Multimodal Parcellation atlas does

(Glasser et al. 2016) which yields greater functional separation of regions (Tait et al. 2020).

Application

In Chapters 4 and 5, the source space data are parcellated according to the AAL atlas. This was done by first taking the MNI152 brain and transforming it to a single subject's brain using a linear image registration tool (FLIRT) (Jenkinson and Smith 2001; Jenkinson et al. 2002). Using the same transform, the AAL atlas was then overlaid onto the subject's brain and a subject-specific mask was created for each region. The voxel location of the centre of mass of each region was calculated and it was this coordinate that was passed to the beamformer. A subset of the 90 available AAL regions was used, excluding deep regions which MEG does not capture well (78 regions in total, see Table 3.1). Beamformer weights were estimated for each centroid yielding 78 source time courses per person. By completing this process for each subject, the results could be compared across subjects for a single brain region.

Region no.	Description	Region no.	Description
1	Rectus_L	40	Rectus_R
2	Olfactory_L	41	Olfactory_R
3	Frontal_Sup_Orb_L	42	Frontal_Sup_Orb_R
4	Frontal_Med_Orb_L	43	Frontal_Med_Orb_R
5	Frontal_Mid_Orb_L	44	Frontal_Mid_Orb_R
6	Frontal_Inf_Orb_L	45	Frontal_Inf_Orb_R
7	Frontal_Sup_L	46	Frontal_Sup_R
8	Frontal_Mid_L	47	Frontal_Mid_R
9	Frontal_Inf_Oper_L	48	Frontal_Inf_Oper_R
10	Frontal_Inf_Tri_L	49	Frontal_Inf_Tri_R
11	Frontal_Sup_Medial_L	50	Frontal_Sup_Medial_R
12	Supp_Motor_Area_L	51	Supp_Motor_Area_R
13	Paracentral_Lobule_L	52	Paracentral_Lobule_R
14	Precentral_L	53	Precentral_R

15	Rolandic_Oper_L	54	Rolandic_Oper_R
16	Postcentral_L	55	Postcentral_P
17	Parietal_Sup_L	56	Parietal_Sup_R
18	Parietal_Inf_L	57	Parietal_Inf_R
19	SupraMarginal_L	58	SupraMarginal_R
20	Angular_L	59	Angular_R
21	Precuneous_L	60	Precuneous_R
22	Occipital_Sup_L	61	Occipital_Sup_R
23	Occipital_Mid_L	62	Occipital_Mid_R
24	Occipital_Inf_L	63	Occipital_Inf_R
25	Calcarine_L	64	Calcarine_R
26	Cuneus_L	65	Cuneus_R
27	Lingual_L	66	Lingual_R
28	Fusiform_L	67	Fusiform_R
29	Heschl_L	68	Heschl_R
30	Temporal_Sup_L	69	Temporal_Sup_R
31	Temporal_Mid_L	70	Temporal_Mid_R
32	Temporal_Inf_L	71	Temporal_Inf_R
33	Temporal_Pole_Sup_L	72	Temporal_Pole_Sup_R
34	Temporal_Pole_Mid_L	73	Temporal_Pole_Mid_R
35	ParaHippocampal_L	74	ParaHippocampal_R
36	Cingulum_Ant_L	75	Cingulum_Ant_R
37	Cingulum_Mid_L	76	Cingulum_Mid_R
38	Cingulum_Post_L	77	Cingulum_Post_R
39	Insula_L	78	Insula_R

Table 3.1: The subset of AAL regions used throughout this thesis. Table adapted from (Hunt et al. 2016).

3.2.7 Leakage correction

Because of the nature of the inverse problem, it is challenging to estimate a source at one location in the brain without it being contaminated by signals arising at a spatially separate brain location. This is often referred to as signal ‘leakage’ and manifests as a zero-phase-lag correlation between distal brain

regions. It is particularly important to reduce the effects of leakage between brain regions in functional connectivity analyses (see Section 2.2.4) where the correlation between estimated time courses implies coordination – we need to be sure that any correlation is due to true connectivity and not an artefact caused by signal leakage.

Mathematically, the signal leakage problem is described succinctly in (O'Neill, Barratt, et al. 2015), and is outlined here. First, consider a true source \mathbf{q}_1 at position \mathbf{r}_1 which is completely independent of source \mathbf{q}_2 at position \mathbf{r}_2 . Because they are independent, the covariance between the sources is zero:

$$\frac{1}{P} \mathbf{q}_1 \mathbf{q}_2^T = 0, \quad (3.37)$$

where P is the total number of time points and both \mathbf{q}_1 and \mathbf{q}_2 are vectors of size $1 \times P$. Assuming that there are no other sources inside the brain, the MEG measurements \mathbf{b} would be given by

$$\mathbf{b} = \mathbf{l}_1 \mathbf{q}_1 + \mathbf{l}_2 \mathbf{q}_2 + \mathbf{e} \quad (3.38)$$

Where \mathbf{l}_1 and \mathbf{l}_2 are the lead field vectors for brain locations \mathbf{r}_1 and \mathbf{r}_2 respectively and \mathbf{e} is the sensor noise. Beamforming gives an estimate $\hat{\mathbf{q}}_1 = \mathbf{w}_1^T \mathbf{b}$ where \mathbf{w}_1 are the beamformer weights for position \mathbf{r}_1 . Substituting in the expression for \mathbf{b} from Equation 3.38 gives

$$\hat{\mathbf{q}}_1 = \mathbf{w}_1^T \mathbf{l}_1 \mathbf{q}_1 + \mathbf{w}_1^T \mathbf{l}_2 \mathbf{q}_2 = \mathbf{q}_1 + \mathbf{w}_1^T \mathbf{l}_2 \mathbf{q}_2, \quad (3.39)$$

because of the linear constraint $\mathbf{w}_1^T \mathbf{l}_1 = 1$. Similarly, $\hat{\mathbf{q}}_2 = \mathbf{q}_2 + \mathbf{w}_2^T \mathbf{l}_1 \mathbf{q}_1$.

The covariance can then be recomputed with the estimated time courses,

$$s = \frac{1}{P} \hat{\mathbf{q}}_1 \hat{\mathbf{q}}_2^T, \quad (3.40)$$

and values for $\hat{\mathbf{q}}_1$ and $\hat{\mathbf{q}}_2$ can be substituted in from Equation 3.39. This results in

$$s = \frac{1}{P} (\mathbf{q}_1 \mathbf{q}_2^T + \mathbf{w}_2^T \mathbf{l}_1 \mathbf{q}_1 \mathbf{q}_1^T + \mathbf{w}_1^T \mathbf{l}_2 \mathbf{q}_2 \mathbf{q}_2^T + \mathbf{w}_1^T \mathbf{l}_2 \mathbf{q}_2 \mathbf{w}_2^T \mathbf{l}_1 \mathbf{q}_1^T),$$

but because the sources are independent (Equation 3.37), the first and final terms equal zero and the remaining covariance between the beamformer estimated time courses are caused by signal leakage. Given that $\text{cov}(X, X) = \text{var}(X)$, this can be rewritten as

$$s = \mathbf{w}_2^T \mathbf{l}_1 v_1 + \mathbf{w}_1^T \mathbf{l}_2 v_2, \quad (3.41)$$

where v_1 and v_2 are the variances of \mathbf{q}_1 and \mathbf{q}_2 respectively. This equation means that the signal leakage is only zero when both $\mathbf{w}_2^T \mathbf{l}_1$ and $\mathbf{w}_1^T \mathbf{l}_2$ equal zero, which is the case when the weights from source 1 and forward field vector from source 2 are orthogonal (and vice versa).

There are a number of leakage correction methods currently in use, but all of them seek to remove zero-lag correlations between regions of interest (Brookes, Woolrich, and Barnes 2012; Colclough et al. 2015; Hipp et al. 2012; Nolte et al. 2004; Stam, Nolte, and Daffertshofer 2007). This is because signal leakage is a linear effect, only introducing false correlations at zero phase lag. Because this thesis investigates the effects across all brain regions (and their connections with each other), symmetric orthogonalisation (Colclough et al. 2015) was used for leakage correction. This multivariate approach is not biased by the order in which leakage correction is done, nor is there any need for a seed region.

Symmetric orthogonalisation works in two steps: first the closest set of orthonormal timeseries are found, then the amplitude and orientation of these vectors are altered iteratively until they converge on a solution closest to the original (uncorrected vectors). This is depicted graphically in Figure 3.6. The mathematics underpinning this method can be found in (Everson 1997; Löwdin 1950; Colclough et al. 2015) but briefly, for a set of n regions of interest, containing P time samples, a matrix of the time courses, \mathbf{Z} , can be given by

$$\mathbf{Z} = \{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_n\} \in \mathbb{R}^{P \times n}, P \geq n. \quad (3.42)$$

We seek the corrected time courses, \mathbf{X} , such that the Frobenius norm is minimised

$$\epsilon = \|\mathbf{Z} - \mathbf{X}\|_F^2 = \text{trace}[(\mathbf{Z} - \mathbf{X})^T(\mathbf{Z} - \mathbf{X})]. \quad (3.43)$$

By constraining \mathbf{X} to contain only orthogonal vectors, it can be expressed as $\mathbf{X} = \mathbf{O}\mathbf{D}$ where $\mathbf{O}^T\mathbf{O} = \mathbf{I}_n$ and \mathbf{I}_n is the $n \times n$ identity matrix. \mathbf{D} is diagonal and $\mathbf{X}^T\mathbf{X} = \mathbf{D}^2$. This can be substituted into Equation 3.43 so that

$$\epsilon = \text{trace}[\mathbf{Z}^T\mathbf{Z}] - 2\text{trace}[\mathbf{Z}^T\mathbf{O}\mathbf{D}] + \text{trace}[\mathbf{D}^2]. \quad (3.44)$$

If \mathbf{D} is known and \mathbf{Z} is full-rank, then ϵ is minimised by maximising $\text{trace}[\mathbf{Z}^T\mathbf{O}\mathbf{D}]$. This can be done using the singular value decomposition of $\mathbf{Z}\mathbf{D}$. The vector magnitudes are then iterated until convergence.

Symmetric, multivariate orthogonalisation




Figure 3.6: A graphical representation of the symmetric orthogonalisation process. The original vectors represent \mathbf{Z} , the uncorrected regional time courses. The nearest set of orthonormal vectors is found, and the magnitude and direction of the orthogonal vectors is iterated until convergence when the set of orthogonal vectors closest to the original data are found. Figure adapted from (Colclough et al. 2015).

3.3 Hidden Markov Modelling

Hidden Markov models (HMMs) are a useful mathematical framework, and have been used to model data from different fields of study including finance and market prediction (Hassan and Nath 2005; Mamon and Elliott 2014), speech recognition (Rabiner 1989), brain computer interfacing (Obermaier et al. 2001; McCormick, Ma, and Coleman 2010) and biomedical imaging (Rezek and Roberts 2005). This is, in part, due to their simplicity and flexibility. Variations in model parameters such as the number of states, the type of observation model and the pre-processing of the data prior to model inference, all allow adaptations of the model to specific scientific questions whilst providing a single framework to link, for example, different imaging modalities together.

In recent years, HMMs have gained traction as a technique for modelling the electrophysiological data produced by neurons in the brain. In 2005, Rezek and Roberts produced a mathematically rigorous framework for application of HMMs to biological signals including an EEG sleep study (Rezek and Roberts 2005). This was followed by Woolrich et al. who applied the model directly to MEG data and was able to segment the data into a number of states. Interestingly, the identified state time windows were on the order of ~ 100 ms, equivalent to those observed in transient EEG microstates. These were then used to create a temporally-adaptive beamformer whose covariance matrix was calculated using only time windows when a given state was active (Woolrich et al. 2013). The model was soon extended to the study of functional connectivity in MEG data, this time identifying the default mode network (a network which has long been studied using fMRI) and observing its interaction with other transient brain states (Baker et al. 2014). More complex HMMs with observation models based on multivariate autoregression and time-delay embedding have identified states based on their spectral content and used them to investigate whole-brain dynamic networks (Quinn et al. 2018; Vidaurre et al. 2018; Vidaurre et al. 2016).

It should be said that the HMM is not proposed as a genuine neurobiological model, but rather a useful way to interpret the data. It would be illogical to assume that the brain obeys the Markov property – that the state of the whole brain (or even a single anatomical region of the brain) at time t is dependant only on the state it was in at time $t-1$. A cumulative effect over many preceding time points seems much more likely (a matter which is explored in more detail in Section 3.3.3 below). However, in much the same way that ICA can be used to extract useful information from brain data, the HMM is able to segment the data in an unsupervised manner, linking together spatial and temporal patterns of activity which recur throughout a given dataset.

The various HMMs implemented in this thesis are based on the HMM toolbox created by the team at the Oxford centre for Human Brain Activity (OHBA), and the following mathematics describes the theoretical basis of the model implementation. Useful reading includes (Rabiner 1989), (Rezek and Roberts 2005) and [OHBA's GitHub HMM-MAR wiki](#).

3.3.1 Model Concept

Markov Chains

Markov chains are a useful way to describe, mathematically, the relationship between variables. In the simplest case we would have just two variables, for example, let us assume that an awake research participant can be in two states – they either have their eyes open, or they have their eyes closed. These states are mutually exclusive and there is no intermediate state. This can be represented by a diagram, as shown in Figure 3.7. The probability of transitioning from one state to the other (and of staying in the same state at the next time point) is given by the number next to each transition arrow. The sum of the transition probabilities from any given state is equal to one. The so-called Markov property of this chain is that the state of the participant at timepoint t depends only on the state they were in at the timepoint $t-1$. Mathematically this is

$$P(x_t | x_1 x_2 x_3 \dots x_{t-1}) = P(x_t | x_{t-1}), \quad (3.45)$$

which means that the sequence of states leading up to time t are irrelevant save the point at $t-1$. As Markov chains become more complex with a greater number of states, it is helpful to express the transition probabilities in a matrix,

$$A = \{a_{ij}\} = \begin{bmatrix} 0.6 & 0.4 \\ 0.3 & 0.7 \end{bmatrix}, \quad (3.46)$$

where a single element, a_{ij} , describes the probability of transitioning from state i to state j .




Figure 3.7: A graphical representation of a simple Markov chain. There are two states, ‘eyes open’ and ‘eyes closed’, with a fixed probability of transitioning between states (or remaining in the same state) shown by the arrows. For example, the probability of transitioning to an ‘eyes closed’ state from an ‘eyes open’ one is 0.4. This can also be seen in the transition probability matrix on the right. The Markov property means that the state which the model is in at time point t depends only on the state at time point $t - 1$.

Hidden Markov Models

The Markov chain can be extended to include a set of ‘hidden’ states which are not directly discernible, but they govern the sequence of observable data. For example, if we take the research participant from our Markov chain example and place them in another room so that we can’t see if their eyes are open or closed, we can infer their state by observing their MEG data; we are most likely to see an increase in alpha activity if their eyes are closed, and we are likely to see a reduction in alpha if their eyes are open. The likelihood of observing an increase or decrease in alpha activity, given that the participant is in one of the

hidden states, can be captured in an emission probability matrix. This is shown in Figure 3.8.

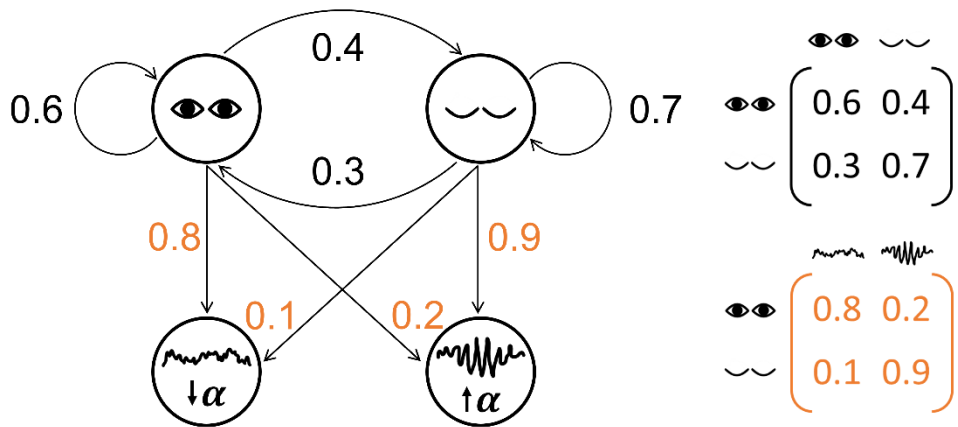


Figure 3.8: A graphical representation of a hidden Markov model. This is an extension of the Markov chain where the hidden states (eyes open, eyes closed) govern a sequence of observable data (an increase or decrease in alpha activity). The transition probabilities remain as before, but now we must also consider emission probabilities (orange). These describe the probability of observing a result from a particular state. For example, if the participant is in an ‘eyes closed’ state, there is a 0.9 chance of observing an increase in alpha activity.

The Markov chain has remained the same (the sequence of states is still governed by the Markov property (Equation 3.45) and a transition probability matrix) but the model has been extended to include the probabilities of observing a result. These emission probabilities can also be written in matrix form,

$$B = \{b_{i,j}\} = \begin{bmatrix} 0.8 & 0.2 \\ 0.1 & 0.9 \end{bmatrix}, \quad (3.47)$$

where each element of the vector, $b_{i,j}$, corresponds to the likelihood of observing result j when in state i .

This model can be used to calculate the probability of any sequence of events occurring. An example sequence is given in Figure 3.9.




Figure 3.9: An example sequence of events. The probability of transitioning between states (x) is given by matrix A . The probability of observing either an increase or decrease in alpha activity (y), given that we are in a particular state, is given by matrix B .

The probability of these events occurring, $P_{sequence}$, is the joint probability of the observation and state sequences. This is just a product of six terms,

$$P_{sequence} = P(x_0 = open)P(y_0 = \alpha_{down} | x_0 = open)P(x_1 = closed | x_0 = open)P(y_1 = \alpha_{up} | x_1 = closed)P(x_2 = closed | x_1 = closed)P(y_2 = \alpha_{up} | x_2 = closed), \quad (3.48)$$

where the probabilities highlighted in blue can be taken directly from the transition probability matrix and those highlighted in orange can be taken directly from the emission probability matrix. However, we don't yet have a value for the initial probability – that of being in an 'eyes open' state at time point 0.

The initial state probabilities are an inherent property of the model and are sometimes referred to as the stationary distribution. The stationary distribution vector, π_0 , can be calculated as the normalised left eigenvector of the transition probability matrix,

$$\pi_0 A = \pi_0, \quad (3.49)$$

and element π_{0_i} represents the proportion of time the model spends in state i . Note that the elements of π_0 must sum to 1 because they represent a probability distribution. For our example the stationary distribution turns out to be

$$\pi_0 = [0.43 \quad 0.57]. \quad (3.50)$$

In keeping with the literature, the full transition probability matrix will now be referred to as π , rather than A , and the initial state probabilities as π_0 . Using these values, the probability of the sequence occurring can be calculated as

$$P_{sequence} = 0.43 \times 0.8 \times 0.4 \times 0.9 \times 0.7 \times 0.9 = 0.078.$$

It is therefore mathematically simple to calculate the probability of any given sequence of events occurring, with the general form of Equation 3.48 being

$$P(X, Y) = P(x_0) \prod_{t=1}^T P(x_t | x_{t-1}) P(y_t | x_t), \quad (3.51)$$

as described by (Rezek and Roberts 2005). Here X and Y are the state and observation sequences respectively, with total duration T time points. Equation 3.51 requires us to know the sequence of hidden states, but by their very nature they are hidden, and so we must infer the most likely underlying state sequence using only the observed brain data. The simplest (and most computationally intensive) method for working this out would be to compute the probability for every permutation of state sequence, given the observed brain data (alpha activity) sequence, and find the most likely. Mathematically this is just

$$\operatorname{argmax}_{X=x_1, x_2, \dots, x_n} P(X = x_1, x_2, \dots, x_n \mid Y = y_1, y_2, \dots, y_n). \quad (3.52)$$

3.3.2 Model Inference

The simple but complete model presented in the previous section (Figure 3.8) has allowed us to explore the concept of a hidden Markov model, with a full description of the transition and emission probabilities.

However, the purpose of this thesis is to use a hidden Markov model to uncover patterns in functional brain data which tell us something new about

the way in which the brain works. In this case, the model is unknown – there is no predetermined set of states, transition, or emission probabilities. The only observables we have are the MEG data which we must use to infer the entirety of the model. This section describes how variational Bayes inference is used to compute the most likely model of brain activity.

Bayes' Theorem

Statistics do not exist in a vacuum; there is always some context in which we observe the likelihood of something happening. Thomas Bayes sought to incorporate this context, or prior information, into the calculation of probability (Bayes 1763). The following maths is an intuitive derivation of Bayes' theorem.

We can consider the probability of observing two events, A or B. The probability of observing A, $P(A)$, is just the subset of outcomes in which A is true, divided by the set of all outcomes. The same can be said of $P(B)$. These events can also both occur at the same time, $P(A \text{ and } B)$, shown by the shaded intersection ($A \cap B$) in Figure 3.10A.

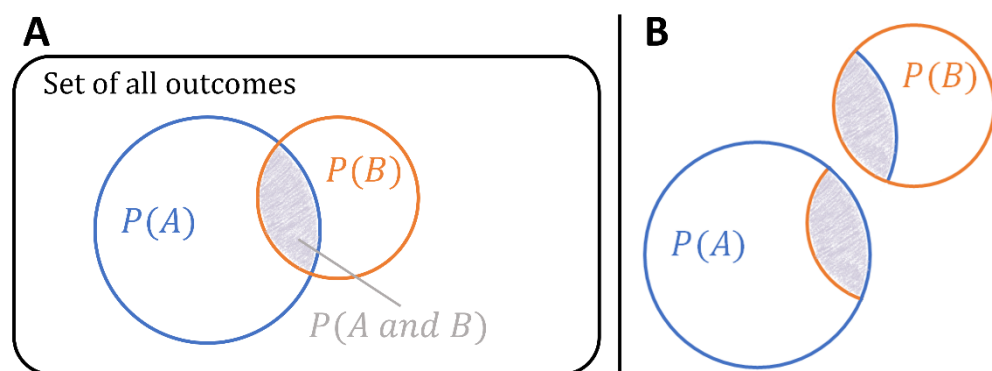


Figure 3.10: There is a probability of observing events A or B. These probabilities can be represented as a subset of all possible outcomes. The blue circle represents the subset of all outcomes in which A is observed, $P(A)$, likewise the orange circle represents the probability of observing B, $P(B)$. However, these events do not occur in isolation, and it is possible to calculate the probability of both events occurring ($A \cap B$, shaded in grey).

There are several ways to calculate $P(A \text{ and } B)$. For example, it is the subset of all outcomes where both A and B are true (shaded area), divided by the set of all outcomes. But it is also a subset of the outcomes where A is true (Figure 3.10B lower left) and a subset of the outcomes where B is true (Figure 3.10B

upper right). This allows us to write two equations: the probability of observing A given that B is true,

$$P(A|B) = \frac{P(A \text{ and } B)}{P(B)}, \quad (3.53)$$

and the probability of observing B given that A is true,

$$P(B|A) = \frac{P(A \text{ and } B)}{P(A)}. \quad (3.54)$$

By substituting Equation 3.54 into 3.53,

$$P(A|B)P(B) = P(B|A)P(A), \quad (3.55)$$

which can be rearranged to produce Bayes' theorem

$$P(A|B) = \frac{P(A)P(B|A)}{P(B)}. \quad (3.56)$$

All we have done is to use our prior knowledge about an event occurring to shrink the set of possible outcomes, yielding a probability value which is more relevant to the context in which the event occurs.

This theorem can be used to systematically update our beliefs. In an experiment it is useful to know the probability of a hypothesis (H) being true, given the observed evidence (E). Using Bayes theorem this is just

$$P(H|E) = \frac{P(H)P(E|H)}{P(E)}, \quad (3.57)$$

where the *posterior* distribution, $P(H|E)$, is the belief about the hypothesis after seeing the evidence. $P(H)$ is often referred to as the *prior* which is the belief or context in which our probabilities are calculated, and $P(E|H)$ is the *likelihood* of observing the evidence given that the hypothesis is true. Each time new evidence is introduced, the hypothesis is updated. In fact, the core of Bayes' theorem is that the evidence should not determine our beliefs about a system, but it should update them.

In the context of optimising the parameters of our HMM, the hypothesis is an initial estimate of the model, and the evidence is the observed MEG data which inform and update the posterior probability of the model.

Variational Learning

It is often the case that the probability distribution we are attempting to characterise is complex and mathematically intractable. This is certainly the case when calculating the most likely posterior distribution of a model of the human brain. The solution is to use a variational framework to approximate the Bayesian posterior densities. The following mathematics for using variational Bayes inference in the field of biosignal analysis are described in detail in (Rezek and Roberts 2005).

In variational learning, the distance between the actual (intractable) distribution P , and an approximate (tractable) distribution Q , is minimised.




Figure 3.11: An illustration of variational inference which aims to minimise the distance between an intractable distribution, P , and a simpler, approximate distribution Q , which in this case is Gaussian.

The distance metric used is the Kullback-Leibler divergence, sometimes referred to as the variational free energy in the field of statistical physics, which is given by the integral,

$$\mathcal{F} = \int Q(H) \log \frac{Q(H)}{P(H|V)} dX + \log P(V). \quad (3.58)$$

Distributions P and Q are functions of all hidden variables, H (model parameters and hidden states) conditioned on the observations V (MEG data).

Assuming that the HMM has K states with hidden state variables $X = \{x_0, \dots, x_T\}$ for T time points, and observations $Y = \{y_1, \dots, y_T\}$ (MEG data), the full true posterior probability of the model is

$$P(Y, X, \theta, \pi) = P(x_0|\pi_0) \prod_t P(x_t|x_{t-1}, \pi) P(y_t|x_t, \theta) P(\theta) P(\pi) \quad (3.59)$$

where π are the transition probabilities, θ are the model parameters, and $P(\theta)$ and $P(\pi)$ are non-informative priors (Woolrich et al. 2013; Rezek and Roberts 2005). Since this is a fully probabilistic model, the observation model linking the hidden states with the observed data is a continuous distribution. For example, this could be a Gaussian distribution as shown in Figure 3.12. The choice of observation model characterises the relationship between the hidden states and the data and so different models will be appropriate for different purposes.




Figure 3.12: A schematic diagram of a hidden Markov model where the hidden states, x_t are linked to the observed values y_t via an observation model. In this case the observation model is Gaussian, and each state is characterised by the model parameters (i.e., a specific mean and standard deviation value per state).

A sensible choice of observation model for application to MEG data is a multivariate normal distribution (MVN), as was used in (Woolrich et al. 2013; Baker et al. 2014). For state k this is

$$P(y_t | x_t = k, \theta) \sim MVN(\mu_k, \Sigma_k), \quad (3.60)$$

with the parameter set $\theta_k = \{\mu_k, \Sigma_k\}$, where μ_k is a $K \times 1$ mean vector and Σ_k is a $K \times K$ covariance matrix.

Now that we have the full true posterior distribution of the model, a sensible approximation (Q -distribution) must be produced; first we use the mean field assumption

$$Q(H) = \prod_{i=1}^T Q(H_i), \quad \int Q(H_i) dH_i = 1 \quad (3.61)$$

which means that the Q -distribution factorises over all hidden variables, with the constraint that the distributions integrate to unity (as they represent probabilities). It is then possible to partition the subset of hidden state

variables further to include both the model parameters ($\theta = \{\theta_1, \dots, \theta_M\}$) and hidden state variables ($S = \{S_0, \dots, S_T\}$). This allows us to write the Q -distribution as a simple chain structure

$$Q(H) \triangleq Q(S)Q(\theta) = Q(S_0) \prod_{t=1}^T Q(S_t|S_{t-1}) \prod_{j=1}^M Q(\theta_j) \quad (3.62)$$

which is mathematically tractable, and for which exact update equations can be derived (Rezek and Roberts 2005; Haft, Hofmann, and Tresp 1999). With each update, the distance between the two distributions is minimised, and this is continued until convergence is reached.

Note that the priors are chosen to be conjugate distributions which means that a prior Gaussian density will produce a posterior Gaussian density (Woolrich et al. 2013). The conjugate distributions chosen for the observation models used in this thesis are outlined below, but for the initial and transition state probabilities, a Dirichlet density was used.

3.3.3 Model Application

The model was applied using the MATLAB toolbox created by the Oxford Centre for Human Brain Activity, which can be found here <https://github.com/OHBA-analysis/HMM-MAR>.

The HMM is a powerful tool for unsupervised segmentation of data into specific states which are revisited many times throughout the course of a single MEG scan. It can also be used to select patterns of task-related brain activity without the need for prespecifying a baseline or rest period. However, the vast number of model parameters which can be selected prior to model inference can produce its own challenge: the user must understand how to tailor the model to each scientific problem to produce any meaningful results. This section outlines some of the most important parameters to consider.

Observation models

The choice of observation model is arguably the most important model consideration; it determines the relationship between the hidden states and the observed data, and therefore establishes the whole structure of the model. In the preceding chapter, a Gaussian model was described, and this is indeed

the simplest observation distribution. One can increase the level of complexity by introducing a multivariate normal distribution. The parameter set used to define this observation model is simply the mean and variance of the data, which means that each state characterises the power content of the signal. This makes the model better suited to amplitude envelope (Hilbert envelope) data in a prescribed frequency band where it is the amplitude modulation which is most interesting; for example, when studying data from a motor task which has been filtered into the beta band.

The observation model used throughout this thesis was the time-delay embedded (TDE) model. This was first introduced by (Vidaurre et al. 2018) and allows one to characterise each hidden state by the frequency content of the raw data (not the amplitude envelope of the signal). It builds the model by computing the autocovariance of the data over several ‘embedded lags’ or time-steps surrounding time point t . The size of the time window (the number of lags) will act almost as a lens – shorter time windows will naturally bias the model towards higher frequency components whereas longer time windows will focus the state characteristics on lower frequencies so there is still a choice to be made about the broader frequency band of interest. However, each state has a unique spectrum associated with it allowing us to distinguish between, for example, an alpha and a beta state. One of the key advantages of the time-delay embedding is that it builds context into the model. The Markov property means that the state in which the model is found depends only on the very last time point, which is an unlikely mode of operation for the brain as a whole. It is much more likely that the brain integrates previous information to update its forward model, or beliefs about the world. Using the TDE-HMM means that each state is defined by the interaction of time point t with the surrounding time points – this builds up an autocovariance matrix for each state. The model still obeys the Markov property because the current state only depends on the preceding state, but the state observation models now include information about the timepoints surrounding each state visit. This model is appropriate when the exact frequency of interest is unknown, when you would like to

model the interaction between states characterised by different frequencies, or when identifying transient states expected to have a specific neural spectral profile (such as beta bursts, see Chapter 4).

Number of states

Pre-specification of the number of states is required to compute the HMM, this is because the model will almost always favour more states as a way of minimising the free energy (Kullback-Leibler divergence) of the model. However, an infinite number of states would lead to severe overfitting and so an appropriate number should be selected by the user.

Rather than assuming that there is a correct number of states, a suitable way of framing the question would be “What is an informative number of states?”. If increasing the number of states does not yield any more information regarding your scientific question, then the extra states are redundant. It is good practice at the beginning of each project to run the model inference several times with different numbers of states to determine the most informative model output.

Model Input

Another component of the HMM which requires careful consideration is the data used for model inference - any timeseries data can be fed into the HMM inference function and the best-fitting model to describe those data will be produced.

First, the model inference can be done at sensor level or in source space (beamformed) data. Then the number of channels must be considered. A univariate (single channel) model would be appropriate when there is only one location of interest, for example from a peak virtual electrode time course, or from a single brain parcel. When the research question involves the whole brain, a multivariate model may be more appropriate. In this case, the spatial properties of the states are modelled as well as those features of the observation model (such as power or spectral content). However, because the temporal and spatial characteristics of the state are combined in a single

model, it becomes impossible to separate out those two elements at a later stage (this is important for the work reported in Chapters 4 and 5).

The data should also be pre-processed in manner which makes the model more informative. First the data should be frequency filtered into a band of interest – this could be a specific neurophysiological frequency range such as 8-13Hz alpha activity, or it could be a broader band such as 1-48Hz which includes most of the interesting neurophysiologic frequencies (except high gamma activity) but stops short of including powerline noise (50Hz, UK). The data should then be down sampled so that the frequencies of interest cover the bulk of the Nyquist range (for example, data filtered 1-48Hz should be down sampled to ~100Hz) for the most efficient model inference. If the data are sensor-level, then removing obvious artefacts (ECG, ocular) from the data prior to model inference is sensible. Having said that, the HMM acts as a useful artefact rejection method in itself, often identifying a ‘heartbeat state’ in the data which can simply be discarded at a further analysis stage.

Model Output

The model output (specific to the HMM-MAR toolbox) includes a structure containing information about the most probable state parameters. For a Gaussian observation model this would simply be $\theta_k = \{\mu_k, \Sigma_k\}$ from Equation 3.60, and for the TDE model this would be an autocovariance matrix per state. The model also outputs state probability time courses. These are the instantaneous probabilities of being in each state at any given time point. It is also possible to decode the model to produce the most likely sequence of hidden states over the whole timeseries (Viterbi path (Rabiner 1989)), however this was not useful for the projects in this thesis. It was more informative to know the likelihood of being in any given state at a single given timepoint.

Summary

The first two sections of this chapter were chiefly concerned with source reconstruction – that is transforming sensor space data into source space through the use of a spatial filter. This involves first creating a ‘forward model’

describing the magnetic fields which would be observable at the MEG sensors given a known current distribution inside the head. This model is then used, together with a beamformer to estimate a solution to the inverse problem (that is recreating the source of the magnetic fields measured at the sensors). The theory behind brain parcellation and leakage correction was also described. The final part of this chapter introduced hidden Markov modelling; the statistical model used throughout the rest of this thesis for detecting recurrent transient spectral events in MEG data. This was accompanied by a discussion of the various model parameters and when it is appropriate to use them for answering a certain research question. The hope is that this will prove to be useful for future students wishing to utilise this model in their research.

Chapter 4

The Role of Transient Spectral Bursts in Functional Connectivity

This first experimental chapter investigates transient bursts of activity, observable with MEG, in both resting state and task data for healthy adult volunteers. The data were acquired by a previous student and all analyses were completed by me. This work has been published in Neuroimage (Seedat et al. 2020) and has been presented at national (MEG-UK 2019) and international conferences (OHBM 2019).

Neural oscillations dominate electrophysiological measures of macroscopic brain activity, and fluctuations in these rhythms offer an insightful window on cortical excitation, inhibition, and connectivity. However, in recent years the ‘classical’ picture of smoothly varying oscillations has been challenged by the idea that many ‘oscillations’ may actually be formed from the recurrence of punctate high-amplitude bursts in activity, whose spectral composition intersects the traditionally defined frequency ranges (such as alpha or beta band). This finding offers a new interpretation of measurable brain activity, however neither the methodological means to detect bursts, nor their link to other findings (such as connectivity) have been settled. Here, a new approach to detecting bursts in MEG data is implemented; a time-delay embedded Hidden Markov Model (HMM) is used to delineate single-region bursts which are in agreement with existing techniques. However, unlike existing techniques, the HMM looks for specific spectral patterns in time course data. Burst characteristics are also investigated – namely, the distribution of burst duration, frequency of occurrence and amplitude across the cortex in resting state MEG data. During a motor task, the movement related beta decrease (MRBD) and post movement beta rebound (PMBR) are shown to be driven by changes in burst occurrence. Finally, this work shows that the beta band functional connectome can be derived using a simple measure of burst overlap, and that coincident bursts in separate regions correspond to a period

of heightened coherence. In summary, this work introduces a new methodology for burst identification and connectivity analysis which will be important for future investigations of neural oscillations.

4.1 Introduction

Neural oscillations comprise rhythmic fluctuations in electrical potential observed across neuronal assemblies, as described in Chapter 2.2. A vast body of work suggests that an increase in the amplitude of 'low' frequency (such as alpha (8-13 Hz) and beta (13-30 Hz) band) oscillations is a marker of increased functional inhibition whereas increased high frequency (gamma (30+ Hz) band) amplitude is a marker of excitation (Pfurtscheller and Lopes da Silva 1999). Modulation of neural oscillations during both simple sensory and cognitive tasks are some of the most robust measurements in brain imaging – for example, ballistic finger movements generate a drop in beta amplitude during movement (the movement related beta decrease (MRBD)), followed by a transient increase in amplitude (above baseline) upon movement cessation (the post movement beta rebound (PMBR)) (Jurkiewicz et al. 2006). Whilst the functional significance of these effects is unknown, their importance is underlined by several demonstrations of abnormalities across a number of disorders including developmental conditions (e.g. autism (Buard et al. 2018)), severe psychoses (e.g. schizophrenia (Robson et al. 2016)), and neurodegenerative disorders (e.g. Parkinson's disease (Gross et al. 2001)). Some light on the role of beta oscillations has been shed by pharmacological manipulation; for example, several studies have shown that alteration of GABA levels results in changes in beta modulation (Muthukumaraswamy et al. 2013). Further studies have shown a direct link between GABA concentration and beta amplitude (Gaetz et al. 2011). Such studies support a hypothesis that beta oscillations are related to inhibition. However, other studies have suggested that they are related to long range connectivity between distal brain regions (Engel et al. 2013), with demonstrations of large-scale spatio-temporal correlations in oscillatory envelopes measured across the cortex (Brookes, Woolrich, et al. 2011; Hipp et al. 2012). Precisely how the inhibitory

modulation of oscillatory amplitude helps to drive connectivity remains unknown.

Recent work has begun to change the way that the research community thinks about oscillations, particularly in the beta band (van Ede et al. 2018). The 'classical' picture is that the brain generates an ongoing oscillation whose amplitude varies over time depending on the task being undertaken. The MRBD was thought of as a drop in oscillatory amplitude, whilst the PMBR represents a smooth increase in amplitude above the resting level. However, this picture was largely a result of time and/or trial averaging. More recent studies (Jones 2016; Sherman et al. 2016; Shin et al. 2017) looking at unaveraged beta oscillations, both at rest and during tasks, have shown that rather than a smooth oscillation, the beta rhythm is actually formed from the recurrence of discrete and punctate events; each event can be thought of as a very short (a few hundred milliseconds) burst of activity. These 'beta bursts' occur with a characteristic probability, which is altered by a task. For example, during movement execution the probability of a beta bursts becomes lower; during the PMBR that probability becomes higher (Little et al. 2019). This means that, when summed over large numbers of trials, bursts combine to give the impression of a smooth decrease, followed by an increase in oscillatory amplitude (the MRBD and PMBR). Interestingly, (Little et al. 2019) have shown behavioural relevance of bursts by demonstrating that the timing of the last burst, prior to movement, predicts movement onset time. The classical view of event related synchronisation and desynchronization is, therefore, likely inadequate. This, in turn, has significant implications across a range of nascent neuroscientific findings, including the application to a clinical setting (Tinkhauser et al. 2017), the inhibition hypothesis, and the interpretation of electrophysiological functional connectivity (Engel et al. 2013). This means that much work needs to be done to understand mechanisms and implications. For example, the methodological means to detect single-region bursts is not settled, with most studies choosing an empirically-derived thresholding of beta band limited data. However, it's likely that bursts are not limited to the beta

band, and a model driven, broad band approach which, in addition to identifying bursts, can also characterise their spectral content would provide an important step forward.

The role of neural oscillations in functional connectivity is based primarily on the concept of 'communication by coherence' (Fries 2005, 2015). The premise is that if neural oscillations in two separate brain regions are in phase, then this provides periods of mutually high electrical potential, which offer optimal windows for the transfer of action potentials, and hence information. In resting state data, brain regions that are highly connected would therefore be hypothesised to exhibit high coherence, and this has proven to be the case with a number of studies showing phase-locking between regions (Engel et al. 2013; Vidaurre et al. 2018). However, other studies have suggested that temporal correlation between the amplitudes of neural oscillations also offers a means to measure connectivity. For example, Brookes et al. (Brookes, Hale, et al. 2011) and Hipp et al. (Hipp et al. 2012) independently showed that some resting state networks commonly observed in functional magnetic resonance imaging (fMRI), could also be observed in MEG by correlating the envelopes of beta band oscillations. More recent studies (Colclough et al. 2016; Liuzzi et al. 2017) have shown that whilst network measures made using coherence and amplitude envelope correlation can be similar, in individual subjects amplitude metrics are typically more repeatable. Moreover, there is now growing evidence that within-network functional connectivity is underpinned by coordinated neuronal dynamics that fluctuate on a very rapid (in the order of 100ms) timescale (Baker et al. 2014). This implies that burst-like activity might be driving network coordination. However, no-one has specifically probed the extent to which bursts drive functional connectivity, and assessment of how these bursts temporally coincide across regions might offer a novel means to measure and interpret functional connectivity.

In this study, in a subject cohort of resting state and task positive MEG data, the role of bursts in mediating connectivity is assessed. In the first part of this chapter, a method is introduced to detect single-region transient bursts in

source localised MEG data. This method, which is based upon the premise of a Hidden Markov Model (Baker et al. 2014; Vidaurre et al. 2016; Woolrich et al. 2013; Quinn et al. 2018; Vidaurre et al. 2018) offers an unsupervised and objective means to identify bursts in broadband (univariate) data. To ensure that the full spectral profile of these bursts is captured, the model is not constrained to the beta band. Instead, transient spectral bursts are detected within a broad frequency range (1 to 48Hz) and findings are compared to a 'classical' beta burst framework. Following this, the role of these bursts in functional connectivity is assessed. Specifically, this hypothesis is tested: *the beta band electrophysiological connectome can be derived based upon an analysis of coincident bursts.*

4.2 Methods

Data were acquired by a previous student as part of the United Kingdom MEG partnership programme and have been published previously (Hunt et al. 2019). 75 subjects took part in the study, which was approved by the University of Nottingham Medical School Research Ethics Committee.

4.2.1 Paradigms and data acquisition

All MEG data were acquired using a 275-channel CTF MEG system, at a sampling rate of 1,200 Hz. The system was operated inside a three layer magnetically shielded room (MSR) and in third order synthetic gradiometer formulation to reduce the effects of external interference. All subjects were seated. Data were acquired during two separate paradigms:

Resting state: Five minutes of resting state MEG data were acquired; the subject was asked to sit with their eyes-open, and 'think of nothing in particular'. A fixation cross was displayed (by projection through a waveguide in the MSR) on a back projection screen which was placed approximately 40 cm in front of the subject to give them something to focus on during data acquisition.

Visuo-motor task: The task comprised presentation of a visual stimulus – a high contrast vertical square wave grating – for a jittered interval of 1.5s to 2s.

Upon cessation of visual stimulation, the participant was asked to make a single right index finger abduction. Fifty trials employed a short inter trial interval (ITI) of 4 seconds, and a further 50 trials employed a longer ITI of 8 seconds. As the PMBR has been shown to last in excess of 6 seconds (Fry et al. 2016; Pfurtscheller and Lopes da Silva 1999; Pakenham et al. 2020) only the long ITI trials are analysed in the current work.

All subjects also underwent an anatomical MRI scan, using a Philips 7T Achieva MRI scanner (Philips) running a phase sensitive inversion recovery (PSIR) sequence (field of view: $240 \times 216 \times 16\text{mm}^3$, 0.8mm isotropic resolution). To co-register MEG functional data to MRI structural data, a 3D head digitisation was obtained. Prior to MEG acquisition, three coils were placed at fiducial locations on the head (the nasion and pre-auricular points). The locations of these coils, relative to the subjects scalp and face surface were digitised (Polhemus Inc.). The equivalent head surface was extracted from the anatomical MRI scan and a surface matching algorithm was employed to compute the locations of the fiducial markers relative to the MRI (hence the brain anatomy). The coils were energised during MEG recording to localise their position inside the MEG helmet; this was done continuously throughout data acquisition. Knowledge of the location of the fiducial markers inside the MEG helmet enabled complete coregistration of the MEG sensor geometry to brain anatomy. It also allowed motion tracking of the subject's head.

Following acquisition, data were inspected visually; 9 subjects were removed from the resting state paradigm and 12 were removed from the visuo-motor paradigm either due to movement (>5 mm) or artifacts in the data. This left a total of 66 (age 38 ± 12 ; 35 female) and 63 (age 38 ± 12 ; 34 female) subjects for the resting and visuomotor paradigms respectively.

4.2.2 Source localisation

The cortex was parcellated into 78 regions according to the automated anatomical labelling (AAL) atlas (Tzourio-Mazoyer et al. 2002), as described in Chapter 3.2.6. The brain was divided into voxels (4mm cubic grid) and the voxel location at the centre of mass of each parcellation was then passed to a scalar

linearly constrained minimum variance (LCMV) beamformer. The data covariance was computed within a 1–150Hz frequency window and a time window spanning the whole experiment (Brookes, Vrba, et al. 2008). Regularisation was applied to the covariance matrix using the Tikhonov method with a regularisation parameter equal to 5% of the maximum eigenvalue of the unregularised covariance matrix. The forward model was based upon a dipole approximation (Sarvas 1987) and a multiple local sphere head model (Huang, Mosher, and Leahy 1999) . Dipole orientation was determined by rotating the dipole in the tangential plane and finding the orientation which gave the highest SNR. This process, which has been used previously (Brookes et al. 2018), resulted in a single regional time course of activity for each of the 78 parcels in the AAL atlas. These time courses were then frequency filtered between 1 and 48Hz, and temporally down-sampled to a sampling rate of 100Hz. Symmetric orthogonalisation was used to reduce the effect of spatial leakage (Colclough et al. 2015).

4.2.3 Hidden Markov model implementation

A hidden Markov model was employed for detection of single-region beta bursts. Each regional time course was treated independently and a univariate 3-state Time-Delay-Embedded (TDE) HMM (Vidaurre et al. 2018) was inferred. This choice of observation model means that each state is characterised by a different autocovariance pattern defined over a specified time window (duration 230ms). These state autocovariance patterns contain the spectral information of the signal when that state is active. Given that most of the literature suggests bursts exist in the beta neurophysiological frequency band, it makes sense to characterise states by their spectral content. However, the full 1-48Hz data were passed to the HMM so the bursts were not constrained to the beta band and any cross-frequency bursting could be detected.

Model inference used a variational Bayesian method which seeks to minimise the free energy of the system (as described in Section 3.3.2). For each separate regional time course, 3 states were assumed (though if insufficient evidence for a third state existed, the model was able to collapse to 2 states). Although

the rest of this paper relates to this 3-state model, results from a 6-state and a 10-state model can be found in Appendix A. The output of the model was therefore a set of 3 state time courses, each representing the probability of a region's time course existing in a specified state, at any given time point. The correlation was then measured between these state probability time courses and the amplitude envelope of beta oscillations in the same brain region (estimated by applying a (Morlet) wavelet transform to the regional time course and extracting those values corresponding to the 13-30Hz frequency band to derive the instantaneous envelope). The state whose probability time course correlated highest with the beta amplitude envelope was taken to represent the 'transient spectral burst state'. This burst probability time course (which reflects the inherent uncertainty in the model) was then binarised; where the instantaneous probability of being in the burst state was greater than two thirds, it was assumed the state had been entered and that a burst had begun. (In post-hoc analyses a threshold probability of $\frac{1}{2}$ was also tested.) This is an improvement over a traditional, heuristic approach (thresholding the beta envelope) in two ways. First, the burst identification is not based purely on the amplitude time course of an a-priori specified frequency band, but on a generative model of bursting across a broader frequency range. This allowed the full spectral profile of bursts to be considered without assuming that the only interesting part of their spectrum falls within the beta band. Second, the thresholding is more principled, being based on the *probability* of being in a burst state.

4.2.4 Burst parameters and comparison to established methods

Having applied the HMM, the characteristics of the state visits (i.e., the bursts), and how they vary across brain regions was examined. To this end the binary time courses associated with the burst state, alongside the regional time course data, were used to determine four burst 'features':

Burst duration: The time that a region spent in the burst state, on each visit.

Burst amplitude: The maximum value of the beta envelope during each visit to the burst state.

Frequency of occurrence: The number of visits to the burst state normalised by time.

Burst interval: The length of time between visits to the burst state.

Each of these values was measured independently for each region and subject, then the values were averaged over subjects and plotted as a function of cortical location (AAL region). For resting state data, these values were recorded as a time average over the entire 5-minute resting state recording. For the visuo-motor task data, they were measured in three windows; during the MRBD, the PMBR, and a 'rest' period at the end of each trial. These windows were taken, approximately, to be 0 - 1 s (MRBD), 1 - 3 s (PMBR) and 4 - 7 s (rest) relative to the offset of the grating, however they were also allowed to modulate, independently for each trial, according to when the bursts actually occurred; when a burst fell at the edge of a time-window, the window was extended to include that burst. Similarly, the MRBD window was extended to capture the time between the edge of the window and the occurrence of a subsequent (or previous) burst. These burst-modulated time windows are shown alongside the results for completeness. For each subject, this procedure resulted in 16 measurements (4 burst parameters, each measured in the resting state and in 3 separate task windows). A Wilcoxon non-parametric rank sum test was used to statistically test for a difference in burst parameters between the resting state and the task. To correct for multiple comparisons, the Bonferroni correction was applied: because three tests were computed (one for each window - MRBD, PMBR and rest) for each of the four burst parameters, the significance threshold was divided by twelve and reduced from 0.05 to a corrected value of 0.0042.

The results from the HMM derived broad-band (1-48Hz) bursts were compared to a more conventional beta band thresholding approach, to test the extent to which the HMM was identifying the same bursts as the established method. To this end, using the mean corrected beta (13-30Hz) envelope of a single

region (left sensory cortex) from the resting state data, three different thresholds were applied (1.5, 2.5 and 3 times the standard deviation (measured independently for each subject)) to produce a binary time course showing periods of high beta amplitude (the thresholded beta bursts). The percentage of overlapping bursts picked up by both the HMM and thresholding techniques were then quantified; the percentage picked up by the HMM only, and the percentage picked up by thresholding only. These values were measured independently for each subject and averaged over subjects. The broad-band nature of this HMM implementation also enabled characterisation of the spectral properties of the bursts in both the resting state and visuo-motor task data. This was done using a state-specific multitaper analysis (Vidaurre et al. 2016), resulting in a single spectrum for each state and each subject, in each location in the brain. Spectra were averaged and plotted alongside the standard error over subjects.

4.2.5 Functional connectivity

To measure functional connectivity, relationships between all possible pairs of regional time courses from the 78 AAL regions were assessed; this method results in 3003 measures of inter-regional functional connectivity. Two different methods were used:

Amplitude envelope correlation (AEC): AEC (Brookes, Woolrich, et al. 2011; Hipp et al. 2012; de Pasquale et al. 2010) is a well characterised connectivity measure in which the Pearson correlation between oscillatory envelopes is computed, see Chapter 2.2.4. Here, AEC was calculated in the theta (4-8 Hz), alpha (8-13 Hz), beta (13-30 Hz), low gamma (30-48 Hz) and high gamma (48-70 Hz) bands. Following frequency filtering, the envelope was calculated using a Hilbert transform and temporally down sampled to 1s. AEC was measured between all pairs of brain regions to yield a (78 x 78 element) matrix. These were generated independently for each subject.

Burst coincidence: In accordance with our hypothesis, it was crucial to establish whether the AEC connectivity matrix could be reproduced by simply considering the likelihood of transient spectral bursts co-occurring in spatially

separate brain locations. To calculate this, for any one pair of AAL regions, the binary time courses were extracted (the outputs of the HMM) for the burst state in each region and the Jaccard index was calculated. This index measures the intersection over the union (and hence the similarity) between two binary time courses; larger values indicate greater burst coincidence. This was computed for all possible region pairs and averaged across subjects. Jaccard indices were then transformed to pseudo-z-statistics by subtracting the mean and dividing by the standard deviation (across the whole matrix).

It should be pointed out that AEC and burst coincidence are, to a degree, related. Specifically, if we assume bursts are periods of high amplitude beta oscillations then the burst time courses should mirror, to an extent, the amplitude envelope. However, a significant amount of information from the envelope itself has been removed by generating the binary time courses. Therefore, we might assume that if the burst coincidence and the AEC derived connectomes are in strong agreement, then AEC is mainly driven by coincident bursts as distinct from lower amplitude coherent fluctuations.

To assess the relationship between AEC and burst coincidence, an element-by-element scatter plot was produced and a linear correlation was measured between the burst coincidence matrix (which recall is measured in the broad band (1-48Hz)) and the AEC matrices, derived in the theta, alpha, beta, low and high gamma bands. To assess how these correlation values compared to a null hypothesis, a pseudo-matrix approach was used (Tewarie et al. 2016; Hunt et al. 2016). First, an eigenvalue decomposition of the burst coincidence matrix was performed. Each eigenvector was then randomised using a phase-based technique. Post-randomisation reconstruction yielded a pseudo-matrix, similar in structure to the genuine matrix, but not reflecting genuine functional connectivity. Comparison of the correlation between real and pseudo-matrices revealed the statistical significance of the relationship. This was done using a non-parametric Wilcoxon rank-sum test between real and pseudo-correlation values obtained for each subject.

To further investigate the contribution of bursts to AEC, the AEC matrix (in the beta band) was recalculated having removed the periods of time during which the coincident bursts occurred – if bursts drive the AEC measure, then the AEC values should be significantly diminished when the bursts were removed. The change in AEC values was quantified (average across the whole matrix) as well as the correlation with the coincident burst matrix.

Finally, to test whether coincident bursts were also representative of time windows of high coherence (phase-locking) between regions, 6 exemplar (interhemispheric) connections were selected across the left and right precentral, postcentral, occipital, parietal and frontal cortices (see Figure 4.5 for exact locations). These were denoted the ‘seed’ and ‘test’ regions. For the 6 connections, all the coincident bursts in both regions were found and, within a time window centred on the midpoint of each of the overlapping bursts, the phase difference derivative (PDD) was measured (Breakspear and Williams 2003; Tewarie et al. 2019). PDD is a measure that captures the stability of phase relationships between time courses: first the instantaneous phase of the seed and test signals is computed and the difference between them calculated, then the phase difference is differentiated. This method is described in detail in Chapter 2.2.4. The result is a time course of coherence between regions. Here, PDD was measured in the $\pm 0.5s$ time window (relative to burst centre) for all coincident bursts, and averaged over bursts. This produced an average time course of changing coherence throughout a coincident burst. Note that PDD was used, as distinct from simpler (and more direct) measures of phase locking (such as coherence) because it yields a time course measure throughout the burst duration. To contrast this with a control condition, PDD was also calculated for ‘non-coincident’ bursts by using the same time window ($\pm 0.5s$), surrounding an equivalent number of bursts that occurred in the seed region, but without a coincident burst in the test region. Higher coherence was expected for coincident bursts compared to non-coincident bursts. These measurements were generated for all subjects, and our hypothesis was tested using a non-parametric Wilcoxon rank-sum test.

4.3 Results

The results are spread across three sections. The first describes the detection of bursts in resting state data, this is followed by an analysis of the burst metrics during a visuomotor task. The final section demonstrates the role of bursts in resting state connectivity.

4.3.1 Identification of bursts in resting state data using the HMM

Figure 4.1 shows, qualitatively, a comparison between a simple thresholding technique (Figure 4.1A), and the HMM (Figure 4.1B) for identifying bursts. The probability time course (blue), and the binarised state time course (red), for the burst state, are shown (overlaid) in Figure 4.1C. A snapshot of 20s of data extracted from the left somatosensory cortex from a single subject is shown. Figure 4.1D shows a time-frequency spectrogram (derived from a continuous (Morlet) wavelet transform) of the same data segment, for comparison. Note that despite being applied to broad-band data, the HMM identifies bursts of high amplitude oscillations which are plainly visible in the time-frequency decomposition and also identified using the thresholding technique. However rather than simply using the amplitude envelope, the HMM makes use of the richer information contained in the waveform shape to identify bursts. Figure 4.1E shows the spectrum of oscillatory frequencies associated with each of the three identified states for this single subject. The burst state was identified as state 2, which in this case peaks in the beta band, but is accompanied by a large alpha band component.

The HMM and thresholding techniques are further compared, across all 66 subjects, in Figures 4.1F and G. Figure 4.1F shows the percentage of HMM bursts which were matched by thresholded bursts, for three different values of threshold. As would be expected, this percentage drops with increasing threshold, but for a low threshold (1.5 times the standard deviation of the envelope) the vast majority (>90%) of HMM-identified bursts were accompanied by a period of above threshold beta amplitude, suggesting that the HMM is a viable way to find punctate periods of high amplitude beta

envelope. Similarly, Figure 4.1G shows the percentage of threshold bursts, which were matched by the HMM. Here, again as would be expected, the percentage increases with threshold, and for a high threshold (3 times the standard deviation) >70% of threshold bursts were also found by the HMM. These statistics suggest that there is agreement between the two methods. However as noted above, the HMM identification of bursts is based on a model of richer spectral content across a broader frequency range, rather than amplitude in a predefined frequency band; and further, the thresholding in the HMM is more principled, by being based on the probability of being in a burst state (which has a marked tendency to have values either very close to 0.0 or very close to 1.0). This is demonstrated by a comparison between two different probability threshold values as shown in Figures 4.1F and G. Changing this probability value from two thirds to a half makes no significant impact.

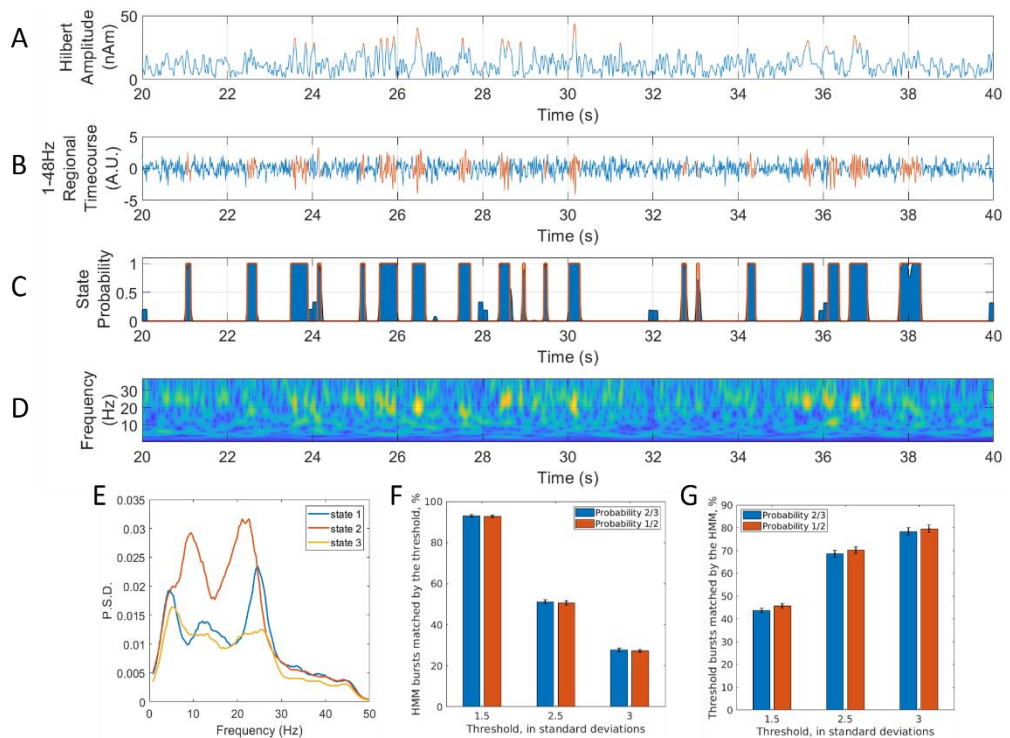


Figure 4.1: Burst detection using the TDE-HMM in resting state data. A) Beta amplitude envelope with bursts (in red) identified using a simple threshold technique (2.5 times the standard deviation). B) Broadband (1-48 Hz) data with bursts (again in red) identified using the HMM. C) Probability time course for the burst state (blue) with binarisation overlaid (red). D) Time frequency decomposition of the data, generated using a wavelet transform. Note that A-D show 20 s of data from a single representative subject. E) Spectra, showing the component frequencies of the three states identified by the HMM from the same single subject. The state most closely correlated with the beta envelope (and hence identified as the beta burst state) is state 2 (red). F) The percentage of HMM identified bursts that were matched by the threshold technique, plotted as a function of threshold itself. Two different probability thresholds on the HMM (two thirds and a half) are shown in blue and red respectively. G) The percentage of bursts identified by the threshold technique and matched by the HMM. The error bars in F) and G) show the standard error over participants.

Figure 4.2 shows the properties (duration, amplitude, frequency of occurrence, and interval time) of the HMM identified bursts as a function of cortical location, in the resting state. The upper (brain) plots show the mean burst parameters across subjects. The lower (line) plots show the same thing but include error bars to enable visualisation of the variation in burst parameters across subjects (note that in general, the variation across subjects is small compared to the regional variation). Generally, the parietal, occipital and temporal regions generate fewer bursts of longer duration compared to the frontal regions, which tend to generate more bursts of shorter duration. Consistent with the known role of beta in the sensorimotor system, the highest

amplitude bursts appear in the primary sensorimotor cortices with the lowest amplitudes in the frontal regions. Interestingly, the time between bursts is also longest in the sensorimotor cortices. On average, the burst durations were found to be of the order 300ms. However, this is also a function of model parameters (see Appendix A).

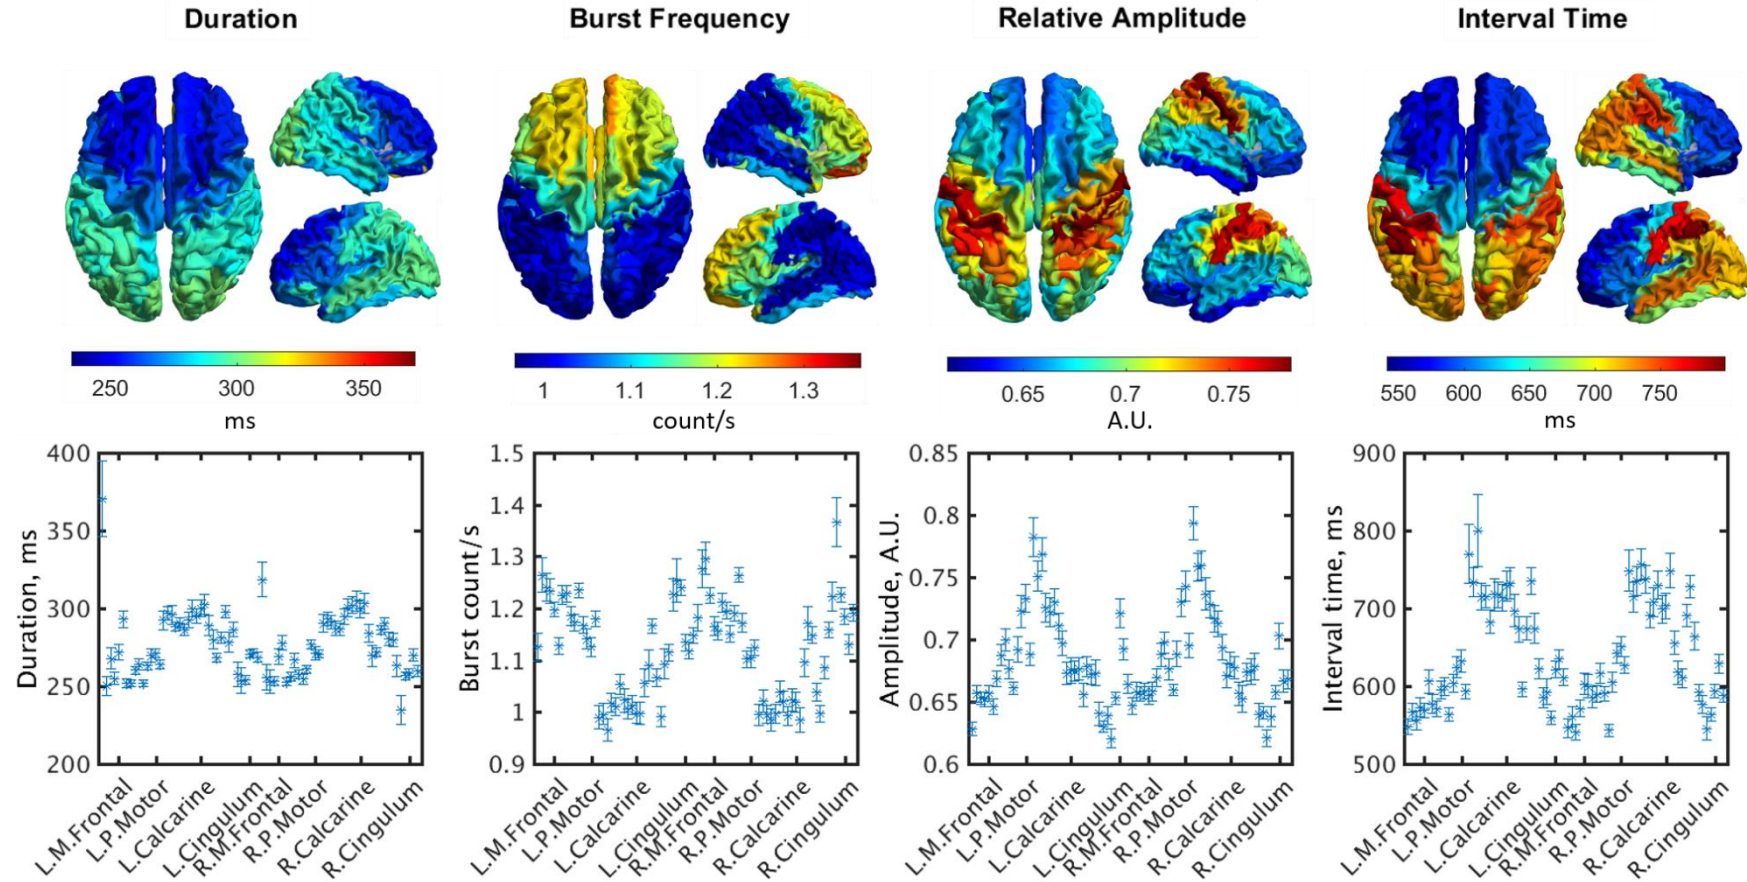


Figure 4.2: Resting state burst statistics. Averaged burst duration, frequency of occurrence, amplitude and interval time across the cortex. The brain plots show the mean values as a function of AAL region. The line plots show the same information, but the error bars allow visualisation of standard error in each measurement across subjects.

4.3.2 Burst dynamics during a visuomotor task

Figure 4.3 shows the results of the HMM burst analysis of the visuomotor task data. Figure 4.3A shows the difference in burst frequency between a window encapsulating the beta rebound (1s to 3s relative to visual stimulus offset (modulated, on a trial-by-trial basis, by neighbouring bursts)), and a window capturing the MRBD (0s to 1s (also modulated by neighbouring bursts)). The spatial signature is clear with a peak centred on the contralateral primary sensorimotor area as would be expected. Plots 4.3B to 4.3E relate to this peak region. Figure 4.3B shows a raster plot of the binary time course for the burst state, extracted from the left primary sensory cortex. The x-axis represents time, and the y-axis represents trials. All subjects have been concatenated (vertically) and time zero corresponds to the offset of visual stimulation. As shown, the burst state displays a greater likelihood of occurrence during the PMBR window compared to rest, whilst in the MRBD window, the burst state is less likely to occur. These relationships are formalised in Figure 4.3C which show the burst parameters in the three windows associated with the task, as well as in the resting state data. Note here that the distribution represents all bursts in all subjects; the black line shows the mean value, and the red line denotes the median. In agreement with the raster plot and previous results (Little et al. 2019), a marked drop in burst frequency is seen during the MRBD compared to the PMBR (or indeed rest and resting state). As would be expected, decreased burst count is accompanied by decreased duration and increased interval between bursts during the MRBD. Interestingly, there is relatively little change in the burst amplitude in the three different windows of the task. There are also substantial differences between the bursts occurring in resting state and the bursts occurring in the visuomotor task data. The bursts detected in the rest window of the visuomotor task are significantly longer than bursts occurring in the resting state; there are also fewer of them, with a shorter interval time. Figure 4.3D shows the average frequency spectrum of the burst state across subjects in both the resting state and task data – note the pan-spectral nature of these bursts. Note also the difference between the two spectra; the spectral content which characterises the bursts appears to

change with task. Figure 4.3E shows the burst-modulated time windows for all trials over all participants.

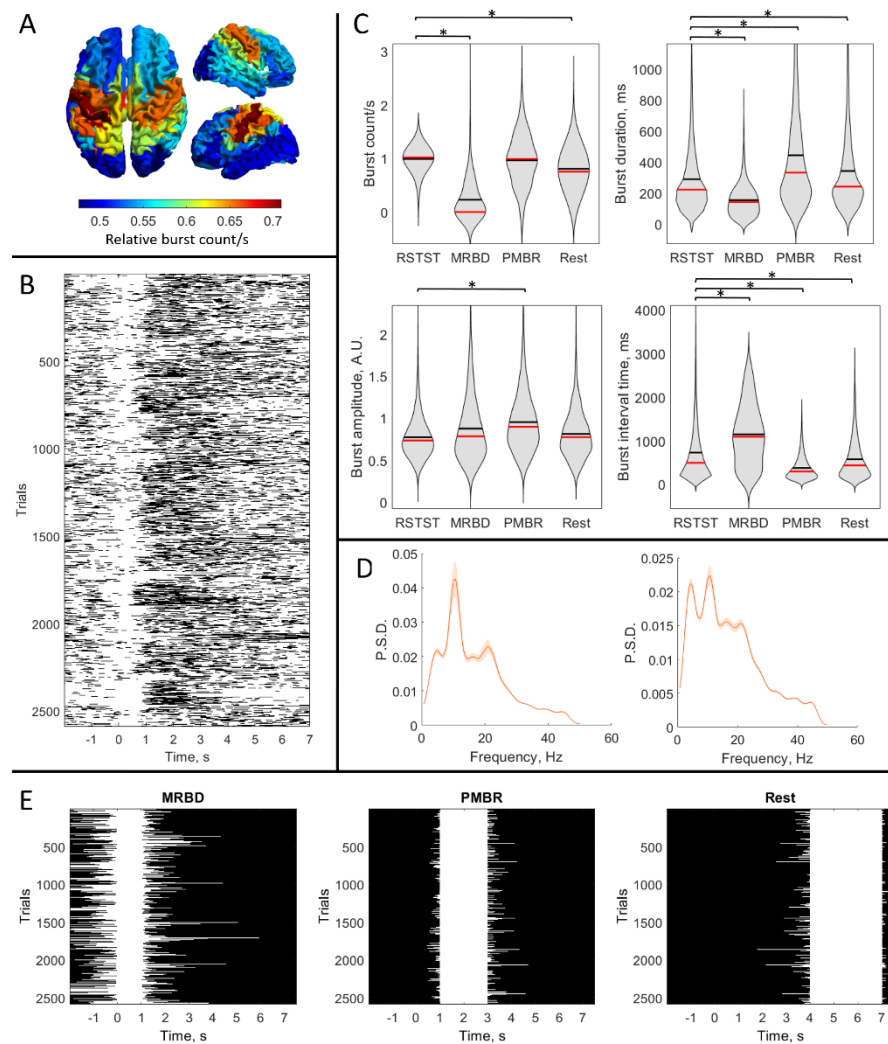


Figure 4.3: HMM-identified bursts during a visuomotor task. A) Image showing the difference in the frequency of burst occurrence between the PMBR and MRBD windows. Notice that, as expected, the greatest difference in burst frequency is in the left sensorimotor cortex. B) Raster-plot showing the occurrences of the burst state for all trials in all subjects. C) Violin plots showing differences in burst parameters (duration, frequency, amplitude, and interval time) in the three separate task windows, and in resting state (RSTST) data. The stars above the violin plots indicate a significant difference (with a corrected p -value < 0.05) in burst parameters when compared with resting state. D) mean spectrum of the burst state, averaged over subjects, for both resting state (left) and task data (right). The shaded area indicates the standard error over subjects. E) Burst-modulated time windows for the 3 task conditions, extended to include all bursts which fall on the edge of the specified windows. These windows are shown in white and were used in the assessment of burst parameters. Panels B, C, D and E relate to the left sensory cortex.

4.3.3 Transient spectral bursts and their role in functional connectivity

In Figure 4.4 burst coincidence is compared with AEC as a means of investigating the whole brain resting state functional connectome. Figure 4.4A shows the result of AEC analysis applied to beta band data. The matrix shows all-to-all functional connectivity (averaged across subjects), whilst the brain shows the strongest 20% of connections; the colour and width of the lines represent the magnitude of individual connections (that is, the envelope correlation values between regions) whilst the radius of the black circles represents connectivity strength between that region and the rest of the brain. As would be expected from previous studies (for example, (Hunt et al. 2016)), connectivity measured in the beta band is strongest in the sensorimotor, posterior parietal and visual areas. Figure 4.4B shows the results of the coincident burst analysis. The elements of the matrix represent a measure of burst coincidence (estimated using the (pseudo z-transformed) Jaccard index) between AAL regions. The brain plot shows the top 17% of inter-regional coincidence values. There is a clear similarity between the beta AEC and broadband coincident burst results, and this is shown quantitatively in Figure 4.4C. The bar plot shows correlation between connectivity methods (burst coincidence and AEC) for all frequency bands (error bar shows standard error across subjects). An example of this correlation for the averaged matrix is also shown. The dotted line shows correlation between the burst matrix and pseudo-matrices derived from eigenvector phase randomisation; all bands demonstrated correlation beyond chance. The result shows clearly that similar information can be derived from AEC and coincident bursts, with the strongest similarity in the beta band. Figure 4.4D shows AEC calculated in envelope data with bursts removed. Note that the colour scale has reduced by a factor of ten when compared with the original AEC matrix (Figure 4.4A). The left panel of Figure 4.4D shows the beta band AEC matrix and the scatter plot on the right shows AEC values, calculated with and without bursts removed, plotted against each other. Note that whilst some structure remains, the amplitude of the connectivity values is markedly diminished compared to Figure 4.4A,

highlighting the fact that much of the connectivity information is contained within the bursts. Quantitatively, the AEC values were reduced from 0.146 ± 0.009 to 0.0036 ± 0.0004 as an average over the whole matrix – a reduction of 97.5%. The correlation with burst coincidence was altered from 0.33 ± 0.02 to 0.076 ± 0.009 . Given that burst coincidence makes up a fraction of the data (for example, 10.52% of the time course in sensorimotor cortex at resting state) but explains more than 90% of the correlation in the AEC matrix, it is clear that bursts play an important role in driving functional connectivity.

Finally, Figure 4.5 shows resting state PDD, measured as a function of time in a -0.5s to +0.5s time window relative to the centre of coincident (blue) and non-coincident (orange) bursts. The figure shows 6 example connections; left-right motor cortex (A), left-right sensory cortex (B), left-right superior parietal cortex (C), left-right visual cortex (D), left-right inferior frontal cortex (E) and left-right orbitofrontal cortex (F). As shown, in the motor, sensory, visual and parietal cortices a coincident burst reveals an increase in broadband coherence between regions, which is not mirrored by non-coincident bursts, with a significant difference ($p < 0.05$) measured in the visual and parietal cortices. This effect was not seen in the frontal cortex. These results support the hypothesis that transient spectral bursts, coincident between regions, offer a means to support network connectivity via a brief period of coherent oscillatory activity.

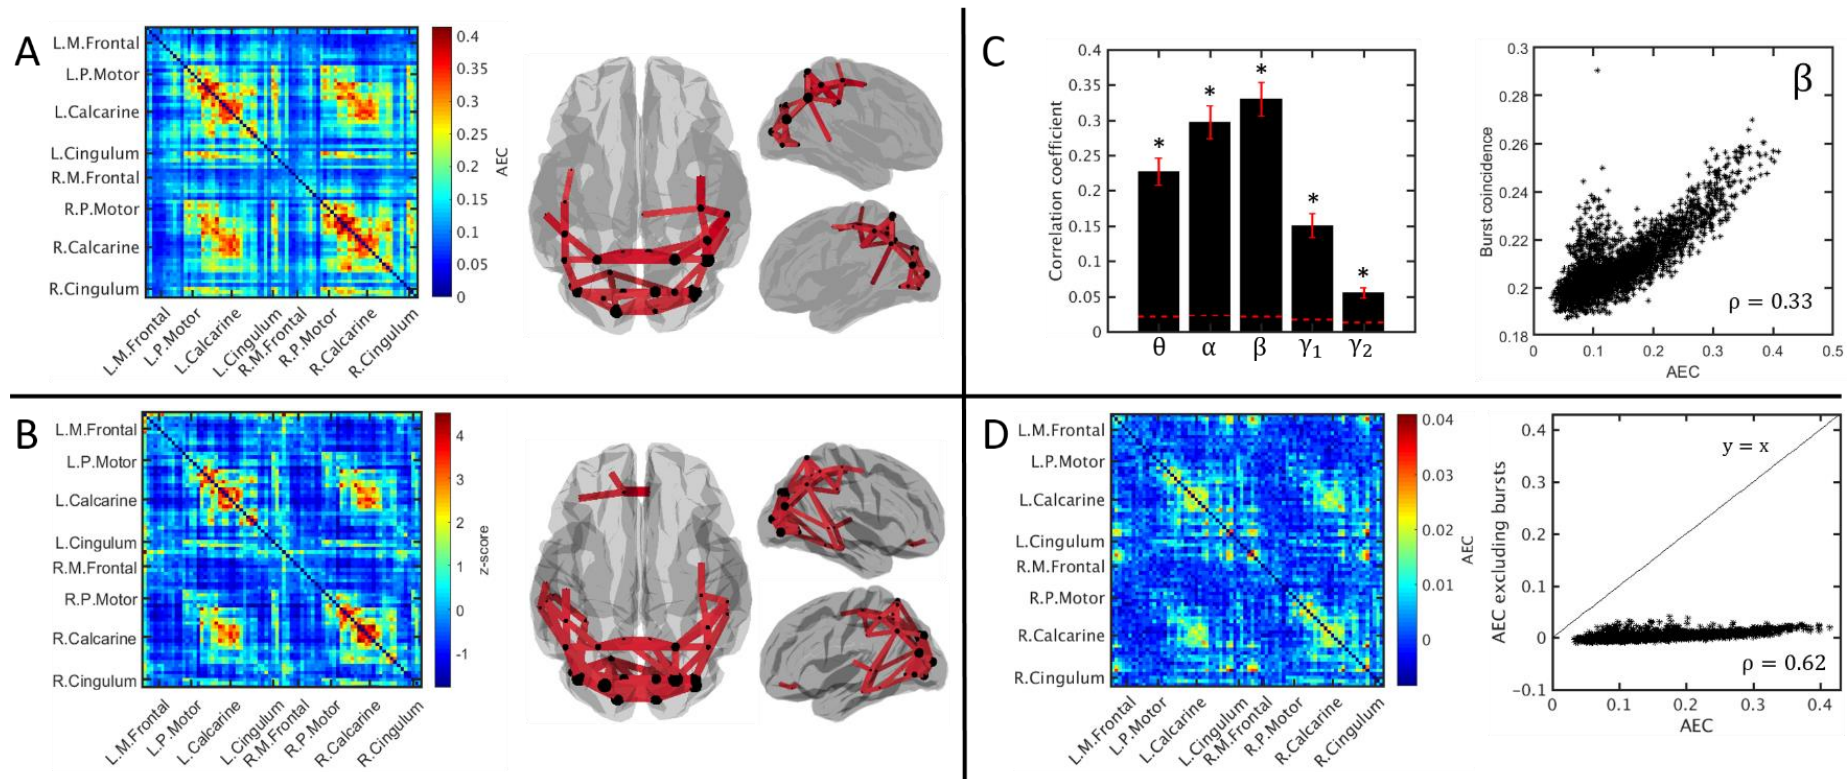


Figure 4.4: Comparison of AEC and coincident bursts in resting state functional connectivity. A) Beta band AEC averaged over 66 subjects. The matrix shows all-to-all AEC measurements whereas the brain plot shows the 20 % of strongest connections. The thickness and colour of the lines denotes the strength of connectivity, and the circles denote connectivity strength between that region and the rest of the brain (the sum in one direction of the matrix). B) Burst coincidence (estimated using the Jaccard index in the broad (1-48Hz) band and transformed to a pseudo z-statistic). Again, the matrix shows all-to-all connections, and the brain plot shows the highest 17% of connections measured. Note similarity between (A) and (B). C) Correlation between the connectome measured using coincident bursts, and AEC measured in all frequency bands. The bar chart shows average values, and the error bars show standard error over subjects. The dotted red lines indicate the ‘null correlation’ values that would be obtained if the structure of the coincident burst matrix were in fact random. The stars indicate where there is a significant difference ($p < 0.05$) between the ‘null correlation’ and actual correlation values. The scatter plot shows the correlation for the beta band case. D) The matrix shows AEC calculated in beta band data with all bursts removed. Note the different colour scale. The scatter plot shows elements of the matrix in (A) plotted against elements of the matrix in (D).

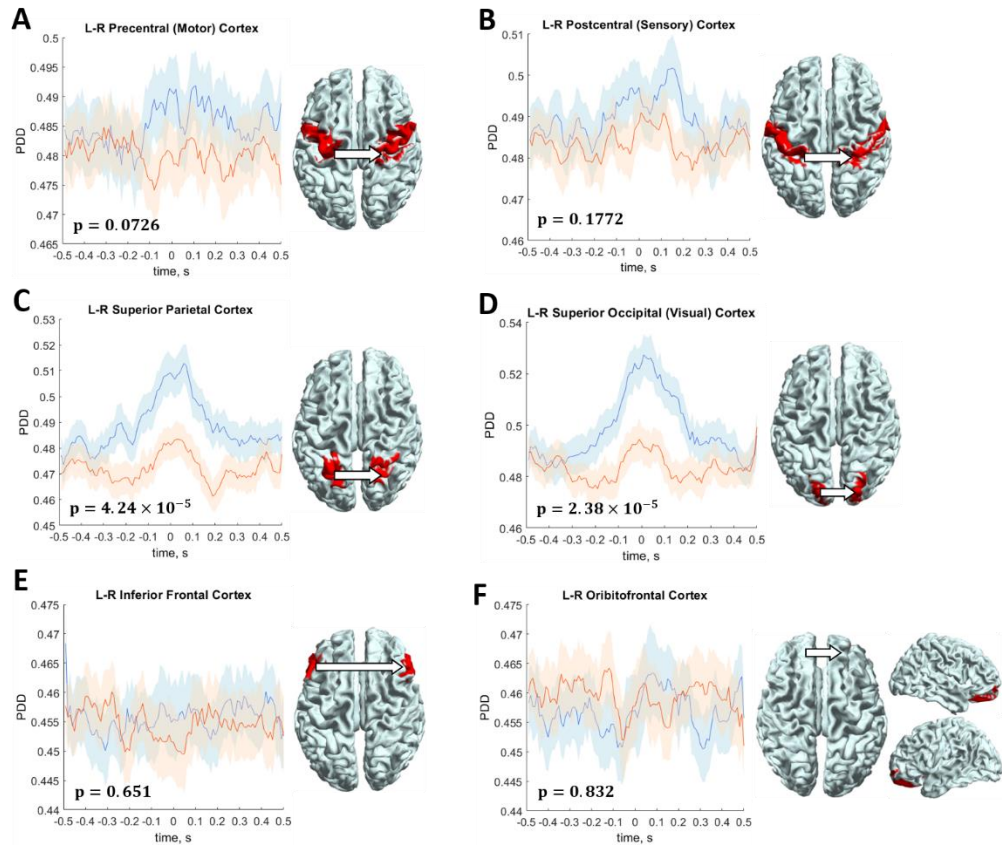


Figure 4.5: Coherence underlying coincident bursts in the resting state. In all 6 plots, phase difference derivative between two spatially separate brain regions is plotted as a function of time in the -0.5 s to $+0.5$ s window relative to the centre of a transient spectral burst. The blue lines represent the case where bursts were coincident between regions. The orange line represents the equivalent case for non-coincident bursts. The shaded area represents standard error across subjects. 6 example connections are shown; left right motor (A) sensory (B) parietal (C) visual (D) and frontal (E-F) cortices. Brain regions are shown inset.

4.4 Discussion

Recent work has shown that neural “oscillations” are generated, at least in part, via a summation of a set of punctate high amplitude “bursts” whose fundamental frequencies intersect with the traditional neurophysiological frequency bands. This new finding has paved the way for a new mechanistic interpretation of the role played by “oscillations” in cognition. However, this work is in its infancy, and even the methodological means to detect bursts is not settled. To date, most studies have employed a simple thresholding of band limited data, but this treats individual bands in isolation without reference to a broader pan-spectral picture of bursts. In addition, there has been no consensus reached on the threshold used. For example recent studies

have used: 2 standard deviations on the mean (Lundqvist et al. 2016); 6 times the median (Shin et al. 2017); and the 98th percentile (Sherman et al. 2016). Because of this, a more principled and broad band approach to detect bursts would be useful.

In previous work, the HMM has been used on multi-region (i.e. multivariate) brain data to reveal visits to short-lived transient brain states with spatially and spectrally distinct patterns (Quinn et al. 2018; Vidaurre et al. 2018). Here, we showed that the HMM also offers a useful means to detect spectral bursts in single-region (i.e. univariate) data, as an alternative to the simple thresholding approach. Importantly, our use of a time-delay embedded HMM, where state characteristics are based upon data autocovariance, means that bursts are identified by the temporal morphology of the signal over a broader frequency range. This is distinct from a simple thresholding technique where only the amplitude of the signal within a prespecified frequency range is taken into account. The new method therefore allows for the complete spectral content of bursts to be elucidated. In fact, the bursts identified by the HMM were not confined to the beta band, but demonstrate a spectrally-specific shape that includes a prominent alpha peak. This was the case even when the HMM was asked to derive a higher number of states (see Appendix A).

Despite the differences there was still agreement between the HMM and the simple thresholding method, with >90% of the HMM identified bursts being mirrored by a period of high beta amplitude (>1.5 standard deviations from the mean). Additionally, when choosing a suitably high threshold (3 standard deviations from the mean), >70% of the bursts identified by the threshold technique were also found by the HMM. This shows that the transient spectral bursts identified by our HMM are (in the vast majority of cases) equivalent to the traditional beta bursts. This is of significant importance for future studies where, we argue, “beta bursts” should be treated as spectral events that span multiple frequency bands.

In the resting state we see marked variation of burst parameters with cortical region. Generally, the highest burst frequency, and shortest duration was in

the frontal regions; the lowest frequencies occurred across the parietal, temporal and occipital regions. This finding could suggest a fundamentally different functional role of bursts in the frontal areas, but could equivalently suggest that the MEG signal in these regions simply has lower signal to noise ratio (likely resulting from the distance between the MEG sensors and the brain). Indeed, the latter hypothesis is supported by a number of findings that beta band power is highest in parietal, occipital and temporal regions. The highest amplitude bursts were found in the sensorimotor regions; this is consistent with the known role of beta in the sensorimotor system. Interestingly the longest duration bursts are observed in the occipital lobe. Whilst the reason for this is unclear, it could relate to the broad band nature of our HMM implementation – recall here that the HMM was applied in the 1-48Hz frequency band, and the burst state spectra clearly covers both the alpha and beta bands. This is because bursts in the occipital cortex are weighted towards an alpha dominance whereas bursts in the sensorimotor cortex are weighted towards beta (see Appendix A).

Our task positive data, shown in Figure 4.3, were in agreement with previous findings (Little et al. 2019). During movement, we observe a drop in the frequency of bursts; their duration is decreased, and the interval between bursts increased, compared to both the resting phase of the visuomotor task, and the true resting state. Conversely, during the beta rebound, as expected we see a marked increase in burst frequency, and also an increase in duration, with the time between bursts decreasing relative to rest, and resting state. Figures 4.3C and D show significant differences between the bursts detected in resting state and those in the rest period of the task data. It is likely that this is due to the relatively short inter-trial interval of 8s since recent work has shown that it takes an average of 11s for activity to return to baseline (Pakenham et al. 2020). Interestingly, although the mean burst amplitude was greater during the beta rebound, this increase was relatively small compared to the large standard deviation of burst amplitudes potentially suggesting that burst amplitude is relatively stable, and that it is burst frequency and duration

that are most affected by task performance. This finding agrees with previous work based upon a thresholding approach, and so provides further validation of the HMM as a means to detect bursting in MEG data.

Our connectivity findings show that the electrophysiological connectome, derived using amplitude envelope correlation, can also be generated by using a simple measure of the temporal coincidence of bursts. Whilst similarity between burst coincidence (calculated in the broad band) and AEC was apparent for a number of frequency bands, it was strongest for the beta band. The beta band connectome has garnered significant interest in recent years due to i) the resemblance of beta networks to the canonical resting state networks observed in fMRI (Brookes, Woolrich, et al. 2011; Hipp et al. 2012), ii) the ability to estimate dynamic spectro-temporal changes in network connectivity (O'Neill, Bauer, et al. 2015; O'Neill et al. 2018), and iii) synergies between electrophysiological and structural data (Hunt et al. 2016). However, an implicit assumption in much of this work was that the oscillatory amplitude varied smoothly over time and correlations in this smooth variation drives connectivity. While there is evidence that the amplitude envelope of oscillations in the resting state is related to transient, bursting amplitude changes from previous HMM work (Baker et al. 2014), the extent to which single-region transient spectral bursts dominate functional connectivity has not previously been investigated.

Here we have shown that, by only taking into account the periods when the single-region bursts from two spatially separate regions overlap, the same static functional connectome emerges. In other words, the known beta connectome can be considered to be driven by coincident bursts. This is further supported by the fact that when periods of identified bursts are removed from an AEC calculation we see a reduction in average AEC values of 97.5%. There is a link between coincident bursts and AEC; high values of AEC are generated by coherent fluctuations in the beta envelope. Given that the HMM nominally pulls out bursts of high beta amplitude, the application of the HMM effectively binarises the beta envelope, with periods of high amplitude

set to one and low amplitude set to zero. Finding the burst overlap effectively amounts to correlating the binarised beta envelopes and therefore it is perhaps not surprising that there is a link between burst overlap and AEC. However, the binarisation removes much of the information content of the envelope - the burst state only accounts for a fraction of the total recording (e.g. 10.52% of the time course in sensorimotor cortex in the resting state accounts for overlapping bursts). The fact that we can get practically the same information from a simple burst analysis suggests that AEC is predominantly driven by bursts, rather than for example, coherent but low amplitude fluctuations which have been lost by the binarisation. However, this said, given that some structure remains in the AEC derived matrix even when bursts have been largely removed suggests that bursts are not solely responsible for the observed connectome.

The final figure shows that when coincident bursts occur, they provide a short period during which the underlying 'oscillations' from the two regions involved are likely to be more coherent. This is consistent with the findings of (Engel et al. 2013), in which resting networks in MEG were shown to be well characterised by visits to short-lived transient brain states, each with spatially distinct patterns of oscillatory power and coherence in specific frequency bands. This obviously lends itself to the communication by coherence hypothesis – that the brain organises its oscillations in such a way as to provide temporally coincident windows of mutually high electrical potential within which we might expect passage of action potentials, and hence information (Fries 2015). Bringing together this finding with other work it is tempting to speculate that coincident bursts offer a means for a broad brain network to provide some inhibitory influence on a specific region, such as the motor cortex; the more frequent the bursts, the more inhibition that region receives. This finding of short periods of coherence also offers some explanation of previous methodological findings. In general it has been shown that amplitude envelope correlation offers a more robust way to measure functional connectivity compared to coherence – at least in individual subjects (Colclough

et al. 2016; Liuzzi et al. 2017). This study shows that bursts represent punctate events, occurring on average once per second in the resting state, which are coherent between regions. This means that over all time (for example, in a 300s resting state recording), we will only see coherence for a very small fraction of that time. It follows that estimated coherence over an entire recording would therefore represent a less reliable measure. Conversely in the amplitude envelope correlation case, it's the beta bursting compared to rest which drives the connectivity measure, and therefore it benefits from both the periods of coherence (with a high beta envelope) and periods of no coherence with a lower envelope.

There are a number of limitations to the current approach which should be understood. First, the HMM requires several parameter choices; these include the selection of the number of states, the length of the window used to compute the autocovariance matrices (which characterise each state), and the threshold on the probability time courses to select which state is "active". It should be noted that parameter choices are typically needed in all analysis approaches. Nonetheless, the HMM output is reasonably robust to changing the number of states (see Appendix A, Figure A1): the burst state spectrum remains constant regardless of whether the model employs 3, 6 or 10 states. Also, the spatial signatures of burst frequency, duration, amplitude, and interval remain qualitatively similar regardless of the choice. However, quantitatively the burst duration, frequency and interval times all change significantly when switching from 3 to 6 to 10 states. This is likely due to the mutual exclusivity assumption made by the model – but also means that burst properties are subjective, depending on parameter choice. The length of the window used to define the autocovariance matrices was 230ms; this value depends on the frequency resolution one wishes to capture. For example, making the window shorter focusses it on higher frequency components (Vidaurre et al. 2018). So here again, a subjective choice is required dependant on the question being asked. The value of the probability threshold is however less subjective since, as shown in Figure 4.1, the probabilities are relatively

binarised anyway. These limitations on parameter choice should be considered carefully in future studies using this technique.

Aside from parameter choice, one of the key assumptions made by the HMM is that the states are mutually exclusive. This means that the model is unable to cope with the co-occurrence of multiple states. This is not intended to represent the true physiological nature of brain activity. Rather, that this assumption and this method provides a useful description of brain activity. Despite this simplification, the binarised burst time courses enable derivation of the electrophysiological connectome which is in good agreement with that derived via AEC. Of course, the “true” connectome (an underlying ground truth) remains unknown. However, it is known that AEC has significant correlation with the resting state networks that are measured using functional magnetic resonance imaging (Hipp et al. 2012; Brookes, Woolrich, et al. 2011). Given the vast literature showing the critical role played by these networks in supporting healthy brain function, and their breakdown in a wide variety of disorders, it is reasonable to consider this as a useful description of the connectome. The fact that these network connectivities can now be understood in the context of transient spectral electrophysiological bursts, will represent an important step in our understanding of how those networks are mediated.

4.5 Conclusion

This work has demonstrated that the time-delay-embedded HMM offers a new way to interrogate MEG data, delineating oscillatory bursts which are in good agreement with the simple thresholding techniques used in the literature. However, unlike arbitrary thresholding approaches applied to oscillatory amplitudes in a single frequency band, here the HMM has been used to look for specific morphological (spectral) patterns in single-region time course data; in this way it provides an objective means to derive transient spectral bursts with a meaningful threshold in broadband data. The distribution of bursts changes across the cortex, with the lowest burst frequencies, highest amplitudes, and highest durations in the parietal, occipital and temporal lobes.

Analysis of task data was in good agreement with previous work in showing that the 'classical' MRBD corresponds to a lower burst frequency, and shorter burst duration whereas the PMBR corresponds to a higher burst frequency, and longer duration. Finally, the well-known beta connectome, which is typically calculated using amplitude envelope correlation, can also be derived by a simple measure of burst overlap. Further analysis suggested that when bursts are coincident, they facilitate a period of phase locking which likely encourages communication by coherence. In summary, this study offers a new methodology for both burst identification and connectivity analysis, which will be important for future MEG investigations of neural oscillations and their perturbation by disease.

The following chapter employs this method of burst identification in two basic neuroscience experiments – an entirely passive visual experiment and a grip force task which encouraged integration between the visual cortex and motor output via visual feedback. The HMM burst-detection framework is then applied, producing a comprehensive analysis of the role of the post-stimulus response and its distribution across the cortex.

Chapter 5

The Post-Stimulus Response Across the Cortex

Following the characterisation of transient bursts of activity in resting state data and their role in connectivity, the same model was applied to task data. This was briefly touched upon in Chapter 4 where it was shown that these same bursts make up the post-movement beta rebound in a finger abduction task. In this chapter, different tasks are introduced to identify and characterise task-related states across the whole cortex – a primary response state and a post-stimulus response state.

The primary aim of this work is to identify if distinct neuronal processes are present in the post-stimulus period compared with the primary response period using an HMM. The secondary aim is to use the resulting HMM states to reveal how the connectivity of different brain regions is altered between the primary response and the post-stimulus response (PSR) states. Data were analysed from two different experiments – the first was a ‘grip-force’ experiment which was primarily a sensorimotor task with visual feedback where participants held a bar in their right hand and exerted a grip-force on it when cued. The second was an entirely passive visual experiment. Using the HMM, the spatial distributions of the primary and post-stimulus responses were characterised across the cortex for both tasks using state probability maps.

For the grip-force task the PSR was greatest in the contralateral motor cortex and bilateral over the visual cortex, reflecting the visual feedback component of the task. In contrast, the primary response state was bilateral over sensorimotor areas, extending into the parietal cortex. A similar result was obtained for the visual experiment, the PSR was confined to the contralateral visual region and the primary response was more bilateral over visual cortex. The difference in the spatial distributions of the primary and PSR provides evidence that they are distinct neuronal processes. A classical (filtered

amplitude envelope) analysis of the visual data showed that the PSR is not modulated by stimulus duration, implying that it is primarily modulated by task difficulty. For those regions that exhibit the responses of interest, connectivity was calculated. This allowed us to gain a new understanding of the task-dependant relationship between brain locations and how it alters between primary response and post-stimulus response neuronal mechanisms. A strength of the HMM is that it allows PSRs to be identified across the cortex regardless of the frequency content of the signal (data aren't prefiltered into a classic neurophysiologic frequency band).

5.1 Introduction

The neuronal response to a basic sensory stimulus can be separated into at least two phases; the primary response (during the task) and the post-stimulus response (PSR, immediately following task cessation). The most well-characterised PSR is the post-movement beta rebound (PMBR) which is an increase in beta (13-30Hz) amplitude (above baseline) upon movement cessation and follows the movement related beta desynchronisation (MRBD) which occurs during the task. The PMBR has been studied extensively in the healthy population (e.g. (Jurkiewicz et al. 2006)) and is shown to be modified in a range of neurological diseases including schizophrenia (Robson et al. 2016), autism spectrum disorder (Gaetz et al. 2020), and multiple sclerosis (Barratt et al. 2017). The significance of this PSR in a range of seemingly unrelated neurological conditions which manifest in such variable symptoms indicates that it is a fundamental neurological process, the disturbance of which has implications on both healthy and pathological brain function. Not only does the PMBR have potential as a biomarker for disease diagnosis, it is clear that further investigation of this phenomenon is also needed to understand the dynamics of healthy brain function.

To understand the role of the PSR, we must first be able to modulate it in a controlled and predictable manner. In 2016, Fry et al. used a wrist flexion paradigm to show that a higher force output results in a PMBR with a greater amplitude, and that the duration of the PSR may depend on stimulus duration

or task difficulty (Fry et al. 2016). This was investigated further by Pakenham et al. who showed, using a grip-force experiment, that the amplitude of the PMBR changes with grip duration (Pakenham et al. 2020). As part of their study, they employed a multivariate HMM to discover a bilateral sensory-motor brain state with a probability time course exhibiting the same temporal dynamics as the PMBR (this state was more likely to be active in the period immediately after stimulus cessation). However, their work also showed that two other states displaying the same beta rebound phenomenon were distributed over other cortical areas, lending credence to the idea that the post-stimulus rebound is a more general mechanism through which regions of the brain coordinate activity. This is a hypothesis supported by (Tewarie et al. 2019) and (Hipp et al. 2012) which show that long range functional connectivity is mediated by beta band oscillations. A functional connectome can also be derived from the same 'bursts' that make up the PMBR (Seedat et al. 2020; Little et al. 2019) as described in Chapter 4.

Given that these bursts of activity at a single-trial level make up the trial-averaged PMBR in a finger abduction task (Sherman et al. 2016; Seedat et al. 2020), it is unsurprising that they can be linked to task performance (Little et al. 2019; Shin et al. 2017) and neurological diseases in the same way that the beta PSR is. For example, longer duration, higher amplitude beta bursts are correlated with symptom severity in Parkinson's disease (Tinkhauser et al. 2017) and the amplitude, duration and number of bursts is altered in established schizophrenia patients when compared with healthy control subjects (Gascoyne et al. 2021).

By contrast, the MRBD (primary response) is relatively unaffected by task parameters such as force output (Fry et al. 2016) or speed of force development (Stancák Jr. and Pfurtscheller 1995). However, the beta desynchronisation associated with motor tasks begins before the movement is undertaken, during a planning and preparation stage. It is during this time window that the MRBD can be systematically modulated by task parameters. For example, in the period immediately before movement execution, the

reduction in beta power is proportional to the uncertainty of movement direction (Tzagarakis et al. 2010). The fact that the MRBD and PMBR are altered by different task parameters, coupled with the fact that they localise to different cortical areas (Jurkiewicz et al. 2006; Fry et al. 2016; Stancák Jr. and Pfurtscheller 1995), supports the idea that they are two distinct functional processes which should be investigated each in their own right.

One theory posits that the MRBD (primary response) is a “cortical gate” which temporarily “switches off” the beta oscillations responsible for resting long-range connectivity to allow local neurons to process information during a task (Brookes et al. 2015). The following PSR is then thought to be an active inhibitory process which drives reintegration of the brain regions involved in stimulus processing – in other words, it acts as a “stop” sign once the stimulus is removed to bring excited brain regions back to baseline (Mullinger et al. 2017). Whilst focus has been on beta band responses when investigating PSRs to date, the beta band response will act in concert with neuronal responses in other frequency bands, indeed the interaction with theta, alpha and gamma band activity is well documented. For example, in 2017 Chung et al. showed that neural oscillations in the theta, alpha and beta frequency bands coordinate when executing upper limb movement (Chung et al. 2017), and in imagined motor actions frontal beta activity has been shown to couple with occipito-parietal gamma oscillations (de Lange et al. 2008). In eyes-closed resting state data, gamma power (30-70Hz) is phase-locked to posterior alpha (8-13Hz) activity (Osipova, Hermes, and Jensen 2008); in a burst-based interpretation of these results, one might reasonably posit that single-trial bursts are in fact made up of both alpha and gamma spectral components, possibly with a characteristic waveform. Furthermore, PSRs have been clearly observed in the 8-13Hz frequency band both in motor and visual cortex (Mullinger et al. 2017; Mullinger et al. 2013) suggesting that constraining analysis to specific frequency bands is limiting. If this is the case, then several brain regions recruited during the task should exhibit a pan-spectral neuronal signature during the stimulus period, which we hypothesize will be unique in

different brain regions which have different functional roles. Then, during the PSR, a different neuronal signature will be present which is hypothesized as necessary to re-establish the long-range connectivity known to be present at rest (Brookes, Woolrich, et al. 2011; Hipp et al. 2012; Tewarie et al. 2019; Seedat et al. 2020; Engel et al. 2013). Therefore, our hypothesis is that this post-stimulus activity will be functionally connected to re-establish resting state networks. Further, we hypothesize that the HMM, as implemented in Chapter 4, is a method well-suited to identifying a primary and PSR state in different cortical regions (each characterised by neural activity with region-specific spectral content) and that this can be used to investigate how connectivity alters between these states.

In this chapter grip-force data from (Pakenham et al. 2020) and an entirely passive visual experiment were used to characterise the PSR across the cortex. The same univariate hidden Markov model described in Chapter 4 was used, with new state selection criteria employed to identify the primary and post-stimulus states. This information was then used to characterise the spatial pattern of activity across the cortex and to derive the functional connectivity of these two separate phenomena. The aim of this work was to answer the questions: *where are PSRs seen across the cortex and how do these regions interact with one another?*

5.2 Methods

The data used in this chapter are from two different experiments; a grip-force experiment with data acquired by a previous student and a visual experiment with data which I acquired myself. Data acquisition was unavoidably impacted by the COVID-19 pandemic and for this reason I have included a section describing the work done to mitigate its effect.

5.2.1 Mitigating the Impact of the COVID-19 Pandemic

On March 11th, 2020, Dr Tedros Adhanom Ghebreyesus, the Director-General of the World Health Organisation, declared that the outbreak of virus COVID-19 was officially classified as a pandemic. This announcement was shortly followed by “work from home” orders by the University of Nottingham on

March 16th, and the UK government officiated a nationwide lockdown on March 26th. Access to laboratories for experimental work (other than those experiments relating to the research of COVID-19) was withdrawn for the duration of the spring-summer 2020 and, given that pilot data had just been acquired for the visual experiment with the aim of starting experiments proper, closure of the lab undoubtedly impacted this study.

To mitigate the impact of the lab closure, the analysis pipeline was built before data were acquired. This was possible partly because there were pilot data, and partly because the structure of the experiment, by design, was similar to one conducted by a previous PhD student. The resulting comprehensive data analysis pipeline enabled analysis of visual data as soon as the lab reopened, and the data were collected.

So that scanning could commence as soon as COVID-19 restrictions were relaxed, it was important to get the lab ready and the correct paperwork in place to meet new health and safety regulations. For this reason, it was necessary to work alone in the lab on several occasions during the summer to program the visual experiment on the stimulus computer and to ensure the triggers, eye tracker, projector, and other pieces of equipment worked. I also revised the lab's standard operating procedures and risk assessments ensure they complied with COVID safety regulations. This involved organising meetings with colleagues so that everyone's needs were met, and all pieces of equipment were considered. Appropriate personal protective equipment (PPE) was introduced (face masks, visors, gloves, and aprons) along with rigorous cleaning procedures and limiting contact with research participants.

One of the ways in which close contact with research participants was reduced was to introduce an optical imaging system for head digitisation (structure IO camera (Occipital Inc., San Fransisco, CA, USA) coupled to an Apple iPad operating Skanect software) to the CTF-MEG lab at Nottingham. This procedure is already used in Nottingham's OPM-MEG lab (Hill et al. 2020) and replaces the Polhemus 3D head digitisation equipment (Polhemus, Colchester,

VT, USA) which requires physical contact with the participant for approximately ten minutes.

Other COVID-19-related safety procedures included the participant completing a safety questionnaire to ensure that they had no COVID-19 symptoms in the days leading up to their scan. There was also a maximum room occupancy in place and doors and windows were opened to allow circulation of fresh air. Where more than one researcher was working in the lab (as was the case for the visual experiment), one person was designated to work with the participant (placing them correctly in the scanner and completing the 3D head digitisation) and the other was assigned to the MEG control station. This limited the number of people the participant was in close contact with and the number of people touching any given surface, reducing the risk of spreading infection.

5.2.2 Grip-Force Experiment

These data were acquired by a previous student and have been published (Pakenham et al. 2020). In total, 15 participants were recruited (10 female; age: 27 ± 3 years (mean \pm standard deviation)), all of them self-reported as being right-handed and gave written informed consent. Electromyography (EMG) and MEG were acquired simultaneously in all patients. Ethics permission was obtained from the University of Nottingham Medical School Research Ethics Committee.

The experiment involved patients holding a grip-force bar (Current Designs, Philadelphia, USA) in their right hand while they lay supine in the MEG scanner. The grip-force bar was attached to a fingerless glove so that participants could relax their hands completely between trials without dropping it (Figure 5.1A). A visual stimulus was presented to the participant which showed a target profile of force output - the force measured by the grip-force bar was overlaid on the target in real time so that the participant received constant visual feedback and could adjust their grip accordingly (Figure 5.1B). The target profile appeared 2s before the grip started and persisted for 0.5s after the grip was released, and each grip was either 2, 5 or 10s in duration. The force output

was not modulated for the trials, only the duration. To determine the target force, participants were asked to exert their maximum grip-force on the bar for 1-2s (maximum voluntary force, MVF). This was repeated twice with a 15s rest period between each grip. The average for each of these grips was found (over a 200ms epoch) and compared with baseline (400ms taken at the end of the MVF recording). The MVF was taken to be the maximum average value from any of those three grips. The target force was then set at 30% of the MVF. Each trial consisted of a grip (for either 2, 5 or 10s) with 30s rest between contractions to allow the PMBR to return to baseline. During the rest period a fixation cross was presented on the screen and participants were asked to remain still and to relax their hands completely. There were 15 trials for each grip duration in a single run (45 trials in total), and two runs were acquired during the experiment with a 15-minute break between them. At the beginning and end of each run muscle fatigue was measured by two 3s-long 100% MVF target profiles with a 30s rest between them. The layout of each run is shown in Figure 5.1C.




Figure 5.1: Experimental design for the grip-force task. Participants held a grip bar with their right hand (A) and when prompted by a visual cue, they applied a force to the bar. This allowed them to trace a square wave (lasting either 2, 5 or 10s) with their right-hand grip (B). In a single experimental run, there were multiple trials (a mixture of 2, 5 and 10s grip durations arranged in a pseudo-random order), each followed by a 30s rest period. Figure taken from (Pakenham et al. 2020).

The data were acquired on a 275-channel CTF MEG system (CTF, Coquitlam, BC, Canada) operating in 3rd order synthetic gradiometer configuration at a

sample rate of 600Hz (see Section 2.3). Before scanning, coils were placed at the participant's fiducial locations: nasion, left and right preauricular points. These coils were energised at the beginning and end of the experiment so that the position of the head relative to the MEG sensor geometry could be localised. This meant any head movement during the experiment could be noted. The head surface was also digitised using a Polhemus Digitrack (Polhemus, Colchester, VT, USA) which produced a 3D scalp surface complete with fiducial locations. This was then matched with the scalp surface extracted from a T1-weighted anatomical MRI of the head, allowing coregistration of the brain position within the MEG scanner. The MRI was acquired with a 1mm isotropic MPRAGE sequence on either a 3T or 7T Philips Achieva MRI scanner.

To monitor the onset and offset of muscle contractions in the right forearm related to the task and during the rest period, electromyography was used. This involved placing Ag/AgCl electrodes (EasyCap GmbH, Germany) in a bipolar configuration over the forearm extensor bundle (channel 1) and the forearm flexor bundle (channel 2). The read-out from these electrodes was sent through an ExG amplifier (Brain Products GmbH, Germany) to BrainVision Recorder (v1.1). A sample rate of 1000Hz was used with an operation frequency band of 0.016-250Hz. A marker was placed at the start of the experiment to synchronise the EMG recordings with the MEG data acquisition.

5.2.3 Visual Experiment

This experiment was performed by me, including paradigm design and coding, all lab preparations, data acquisition.

Data were acquired in 15 healthy volunteers (7 female) with average age 26 ± 5 years (mean \pm standard deviation). All participants self-reported as right-handed and gave written informed consent. On the day of scanning two experiments were undertaken – a visual experiment (reported in this chapter) and an n-back experiment. The n-back experiment forms part of another student's thesis and will not be reported here. Two runs of each experiment were completed, and to ensure that the order of experiments had no bearing on the results, the order of experiments was varied between subjects. Ethics

permission was granted for both experiments by the University of Nottingham Medical School Research Ethics Committee.

The experimental paradigm was similar in structure to the grip-force study, except in this case the stimulus was entirely passive so that any modulation in the PSR amplitude or duration would be due to the stimulus duration and not task difficulty. Participants sat upright in the MEG scanner and looked at a fixation point in the centre of a screen. Each trial consisted of a visual stimulus (lasting either 2, 5 or 10s) in the left visual hemifield followed by a 30s rest period to allow post-stimulus activity to return to baseline. There were 15 trials of each stimulus duration presented in a pseudo-random order, so in total there were 45 trials per run (see Figure 5.2B). The visual stimulus itself was a circular grating with concentric rings moving towards their centre (see Figure 5.2A). This stimulus is known for eliciting a strong visual gamma response but during our pilot experiments it also produced a strong alpha response which was observed by (Tan, Gross, and Uhlhaas 2016; Hoogenboom et al. 2005), making it a suitable visual stimulus for our purposes. To reduce the likelihood of an after effect a Michelson contrast of 1/3 was used. A single experimental run lasted approximately 27 minutes.




Figure 5.2: Experimental design for the visual experiment. Participants looked at a red fixation dot in the centre of a grey screen for the duration of the experiment. A circular grating with concentric rings moving towards their centre appeared in the left visual hemifield periodically (A). The stimulus was presented for a duration of either 2, 5 or 10s (in a pseudo-random order) and there was a 30s rest period between each stimulus (B).

Given that the participant was asked to sit in a darkened room and stare at a fixation point for almost half an hour at a time, it was important to make sure they did not fall asleep. An eye tracker (Eyelink, Ottawa, Canada) was used which interfaced with the MEG electronics so that it was possible to remove

any data segments where the participant appeared to fall asleep (eyes closed for longer than 3s).

MEG data were acquired using a 275-channel CTF system (CTF, Coquitlam, BC, Canada) operating in 3rd order synthetic gradiometer configuration and with a sample rate of 600Hz. Triggers were placed in the MEG data for each stimulus on/offset. Fiducial coils were placed on the participant's nasion and left and right preauricular points, these were energised at the beginning and end of the experiment so that the head could be localised within the scanner helmet and to measure any head movement during the experiment. All participants also had a T1-weighted MRI using a 1mm isotropic MPRAGE sequence on a 3T Philips Achieva scanner in a separate scanning session (never immediately before MEG data acquisition). A 3-D head digitisation was acquired with a Structure IO camera and Skanect software as described in Sections 2.3.4 and 5.2.1.

5.2.4 Data Analysis

For the visual experiment, first a more 'classical' analysis approach was used to investigate the response profile to the stimulus both in the primary and post-stimulus periods. This form of analysis had already been performed for the grip-force data (Pakenham et al. 2020). Following this, both experiments were analysed using a parcel-wise HMM to identify regions where task-related modulation occurred.

Data pre-processing

In order to map the MEG sensor geometry to each participant's brain anatomy, the 3D head digitisation (from either the Polhemus or the Skanect) was mapped onto the scalp surface formed from the MRI. Knowledge of the fiducial coil location relative to the brain (from the 3D digitisation) and relative to the MEG sensors (by energising the coils inside the MEG helmet) allowed coregistration of the brain anatomy with the MEG sensors.

The MEG data were then chopped into trials relative to stimulus offset (the exact timing of the grip offset was determined from the EMG traces, see (Pakenham et al. 2020) for further details) and all trials were bandpass filtered

1-150Hz and visually inspected for artefacts (such as those related to SQUID resets or ocular and muscle artefact). Those trials with artefacts were removed from further analyses. All data were then concatenated (across all stimulus durations) and an LCMV beamformer (Robinson and Vrba 1998; Van Veen et al. 1997; Van Veen and Buckley 1988) was used for source localisation.

Source localisation

The participant's brain was extracted from their MRI using a brain extraction tool (BET) (Smith 2002) and was divided into a 4mm cubic grid of voxels. Then a covariance matrix was computed using all of the pre-processed data (concatenated across all stimulus durations); it was regularised with a 5% regularisation value using the Tikhonov method. A multiple local sphere head model (Huang, Mosher, and Leahy 1999), with a dipole approximation (Sarvas 1987), was used for the forward model and the dipole orientation in the tangential plane which produced the highest SNR was used. The covariance matrix, together with the forward model, was used to generate a time course of neurophysiological activity for each voxel.

Initial visual data analyses

Since the response to visual stimuli is strongest in the alpha frequency band, the source localisation was carried out as described above except that the data were first filtered 8-13Hz and then the covariance matrix estimation was made (using data from all stimulus durations). This was then used, together with the forward model described above, to produce the beamformer weights. Instead of using these to generate time courses of activity at each voxel, they were used to estimate the location of the voxel with the greatest change in activity between the stimulus and post-stimulus periods.

An 'active' window was defined in the period 0-2s from stimulus onset (primary response window) and an active covariance matrix (C_a) was estimated with data from this time window (from all stimulus durations). A post-stimulus window was also defined, this time in the period 0.5-2.5s from stimulus off, and a post-stimulus covariance matrix (C_p) was estimated. For each voxel, a pseudo-T-statistic (PsT) was calculated as

$$PsT = \frac{(W_{\theta}'C_aW_{\theta}) - (W_{\theta}'C_pW_{\theta})}{2W_{\theta}'C_pW_{\theta}},$$

where W_{θ} are the beamformer weights at location and orientation θ . This resulted in a statistical image of the whole brain where the peak PsT value corresponded to the maximal difference between primary and PSRs.

The time course of activity from the peak voxel was then derived (with a covariance matrix estimated using data filtered 1-150Hz). Data were then filtered 8-13Hz and Hilbert transformed to produce an amplitude envelope time course of alpha activity for every trial. This time course was then averaged over trials to produce an induced response per participant. Pakenham et al. showed that the PSR can be incorrectly analysed using a short rest period as the neuronal responses have not returned to true baseline. This study sought to replicate that finding with the visual data. A baseline window typical of MEG experiments (5-7s post-stimulus) was used, as was a 'true' baseline of 15-23s after stimulus off, akin to the baseline used in Pakenham et al. These baseline-corrected time courses were then averaged over participants to produce a single time course of alpha activity for each stimulus duration.

A repeated measures ANOVA was then used to interpret the results to see if there was any statistical difference between the rest period and post-stimulus window, and to see if there was any difference in rebound amplitude between the stimulus durations. The difference between the primary response and the rest period was also tested using the repeated measures ANOVA, as was the difference in primary response between stimulus durations.

Hidden Markov modelling

For each participant, the down-sampled (4mm resolution) brain was parcellated according to the Automated Anatomical Labelling (AAL) atlas (Tzourio-Mazoyer et al. 2002) to produce 78 functionally-specific brain regions. First, a linear image registration tool (FLIRT (Jenkinson et al. 2002; Jenkinson and Smith 2001)) was used to calculate the transform from the 4-mm MNI152

brain⁶ to the participant's brain, and then the same transform was applied to the MNI152-aligned AAL regions (and their centroids) to produce subject-specific parcellations. The voxel time courses were then computed for the centre of mass of each parcellation, producing a single time course of activity for each of the 78 AAL regions. Symmetric orthogonalisation was then applied to reduce the effect of signal leakage as described in Section 3.2.7 (Colclough et al. 2015).

For each of the 78 brain regions a univariate, 3-state, time-delay embedded HMM was employed (directly analogous to the model used in Chapter 4) using the HMM-MAR toolbox. The parcellated data were first down-sampled to 100Hz and then frequency filtered into a broad band (1-48Hz) so that the model could characterise states based on the spectral content across a range of physiological frequency bands (theta, alpha, beta, low gamma). This was made possible by using a time-delay embedded observation model (Vidaurre et al. 2018) where each state is defined by an autocovariance pattern over a 230ms time window – this pattern contains the spectral information.

Because this model is data-driven, a more accurate model is likely to arise when more data are used for model inference. For this reason, the down-sampled and frequency filtered data for a single brain region were concatenated over all participants and then passed to the model inference - a variational Bayesian inference method which minimised the free energy of the system was used. Where there was insufficient evidence for three states, the model was able to collapse into fewer. States were assumed to be mutually exclusive. The output of the HMM was a time course for each of the three states describing the probability of the state being active at any one point in time. This model inference was completed for each brain region separately resulting in 3 state probability time courses per region.

⁶ Montreal Neurological Institute standard brain averaged over 152 healthy adults, see Section 3.2.6

State selection

Each of the three state time courses were thresholded at two thirds probability to produce a binarised time course of state activity. The binary state time courses were then chopped back into trials for each of the three stimulus durations (2, 5 or 10s) and arranged in a raster plot (see Figure 5.3A for the 2s grip-force raster plots over all subjects). By averaging over trials, the probability of entering the state at each timepoint was established (Figure 5.3B). These state time courses were then normalised (divided by their maximum probability) and baseline corrected with a 20-24s post-stimulus baseline (Figure 5.3C). The state most likely to be entered during the stimulus was selected as the primary response state, and likewise the PSR state was the one most likely to be entered into from either 1-5s (grip-force exp.) or 2-4s (visual exp.) after stimulus off.

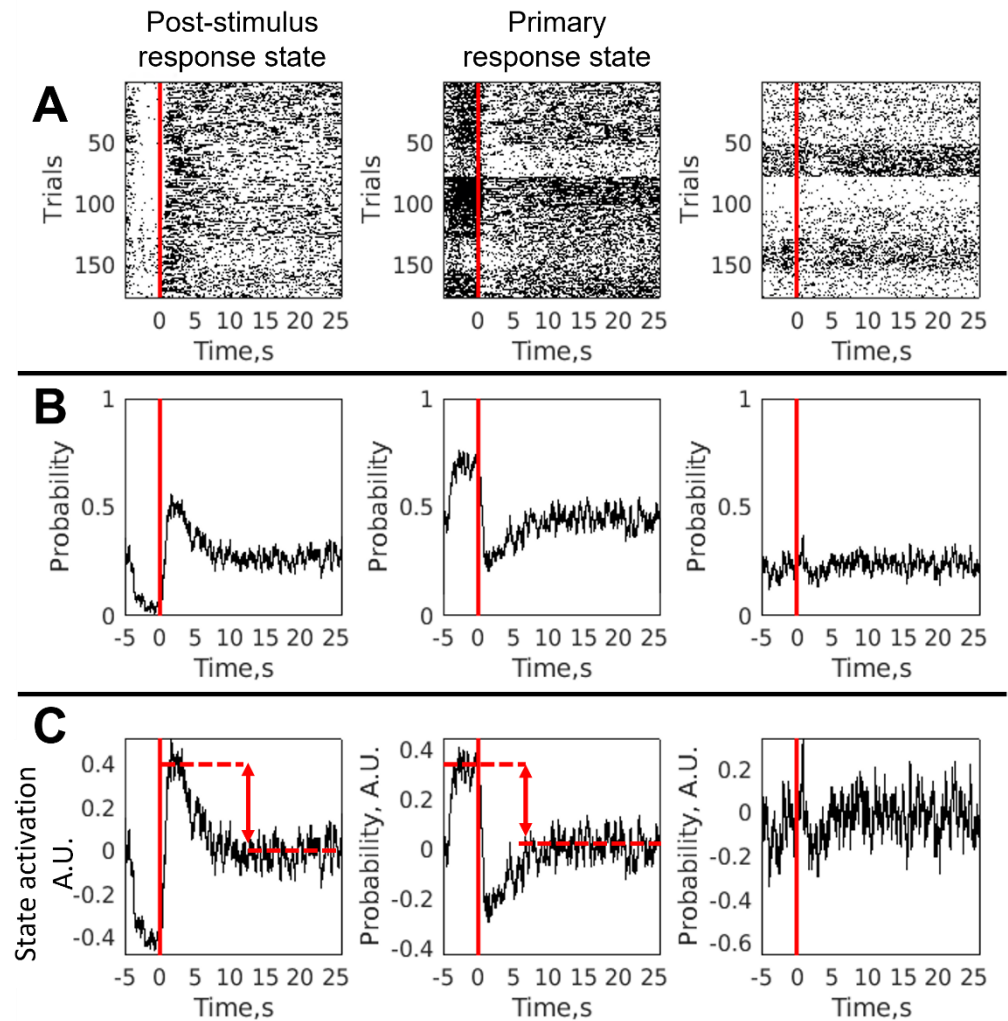


Figure 5.3: Primary and PSR state selection. For each of the three states, a raster plot of state activation (for all trials in all subjects) was obtained (panel A shows this for states 1, 2 and 3 from left to right). The average over trials created a time course of state activation probability (B), which was then normalised and baseline corrected to create a time course of state activation in panel C. For the stimulus window, the difference from baseline was calculated and the state with the greatest state activation was selected as the primary response state (state 2 in this case). Similarly, the PSR state was the one most likely to be entered in the time window immediately after stimulus cessation (state 1 in this case). The strength of the response is indicated by the arrows in panel C. In all cases the solid red line indicates stimulus offset. These plots relate to the 2s grip-force stimulus.

State activation maps

To understand the spatial distribution of these phenomena (both primary and post-stimulus) the strength of the response was calculated for each brain region. This was taken as the difference from baseline (20-24s from stimulus off) in the normalised state activation time courses from Figure 5.3C. This approach was chosen because a region with no modulation in brain activity

related to the stimulus, would also show no modulation in the state activation time courses (such as state 3 in Figure 5.3C). Therefore the difference value would be very small, reflecting the lack of task-related modulation.

State Spectra

Because the HMM state identification is data-driven and is inferred using broadband (1-48Hz) data, the spectral content of the primary and post-stimulus states was allowed to vary across brain regions – this makes sense because different frequency bands dominate in different cortical areas, and therefore task related activity across the cortex is likely to vary in its spectral content. A state-specific multitaper analysis (Vidaurre et al. 2016) was used to investigate this variation in primary and PSR state spectra, yielding a spectrum for each state, brain region and task.

State connectivity

In Chapter 4, the resting state connectome was recreated using only a coincident burst metric. Here we extended the same method to investigate coordinated activity between regions exhibiting either a primary or PSR.

First the binarised state time courses for the primary and PSR states were extracted for each brain region. These time courses were then windowed so that only the ‘bursts’ of activity during the *relevant time window* were used in the connectivity calculation⁷:

- Primary response window was from 4s to 0s preceding stimulus offset.
- PSR window was from 1-5s (grip-force expt.) or 2-4s (visual expt.) after stimulus offset.

The Jaccard index was then computed for every pair of regions as a measure of similarity between binary time courses (where each time course is concatenated over all participants). This resulted in a single connectivity matrix over all subjects for the primary and PSR states.

⁷ Connectivity was also calculated using state visits across the whole experiment. Results for this are shown in Appendix B.

To ensure that these connectivity matrices were meaningful (and not measuring the overlap between randomly-occurring bursts unrelated to the task), brain regions with no task-related modulation were excluded from the analysis. These regions were identified by first calculating the probability of entering either the primary or PSR states (those same values used to create the state activation maps) and then calculating the standard deviation in the relevant time window. Where the probability was less than the standard deviation, the region was disregarded. For the remaining regions, the state allocations were checked to see if the primary and PSR states were ever selected as the same.

5.3 Results

The results are split into two sections, the first discussing the grip-force experiment and the second discussing the visual experiment.

5.3.1 Grip-Force Stimulus Results

During the pre-processing stage one participant was removed from further analysis because of hand movements during the rest period (identified using EMG) so results are presented from the remaining 14 subjects. After bad trials were removed there were an average of 13 ± 2 , 13 ± 2 , and 12 ± 3 (mean \pm standard deviation over all subjects and runs) trials remaining for the 2, 5 and 10s stimulus durations respectively.

The state activation maps for the grip-force experiment are shown in Figure 5.4 for both the primary and PSRs. Interestingly, the two phenomena exhibit different patterns of activity across the cortex. The primary response is strongest over the sensorimotor areas, extending into the parietal cortex and is near-symmetrical over hemispheres. In contrast, the likelihood of entering the PSR state is greatest over the contralateral sensorimotor area and bilateral visual regions. This strong visual component was initially surprising but is likely to be a result of the continuous visual feedback throughout the task – participants were able to see their grip-force overlaid on the target profile (see Figure 5.1). This analysis shows that coordinating the visual feedback with their motor output produces a clear PSR in both areas (Figure 5.4B). Given that the

motor component of the task only recruited their right hand, and the visual component of the task was presented to them in both visual hemifields, it also makes sense that the visual PSR is bilateral whereas the motor PSR is more unilateral. What is perhaps surprising is that there is very little post-stimulus activity identified by the HMM analysis in the ipsilateral sensorimotor areas where one might expect to see some activity with a reduced amplitude (as previously reported for a simple right index finger abduction task (Jurkiewicz et al. 2006)).

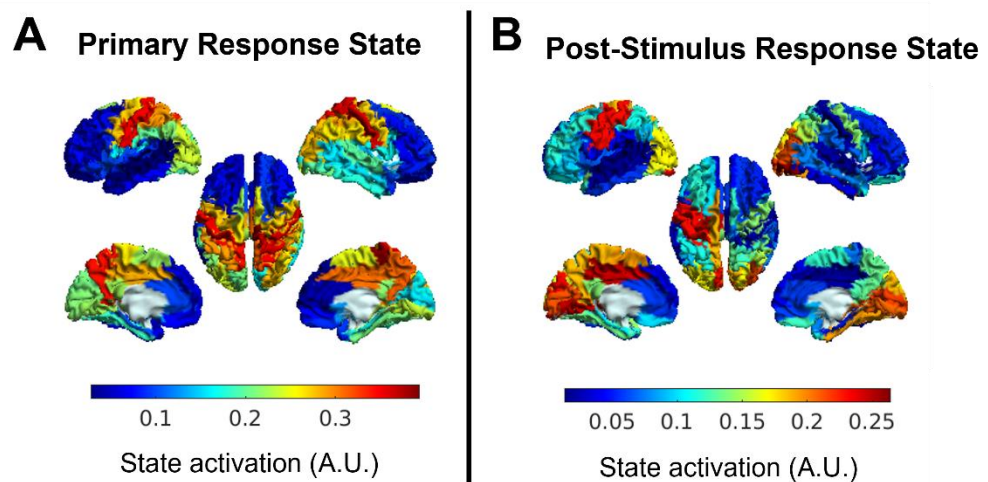


Figure 5.4: Spatial maps showing the distributions of the primary and PSR states across the cortex for the grip-force task. The likelihood of entering the primary response state is greatest over the sensorimotor areas, extending into the parietal cortex (panel A). The likelihood of entering a post-stimulus state (B) is greatest in the contralateral sensorimotor area and bilateral over the visual cortex.

The variation in state spectra is shown in Figure 5.5 for all 78 brain regions, where each region is represented by a different coloured line. The overall PSD for the PSR state (Figure 5.5B) is greater than the primary response state (Figure 5.5A). It also includes a larger alpha component, and for some regions

a beta component. These peaks are likely to reflect the visual feedback and motor components of the task respectively.

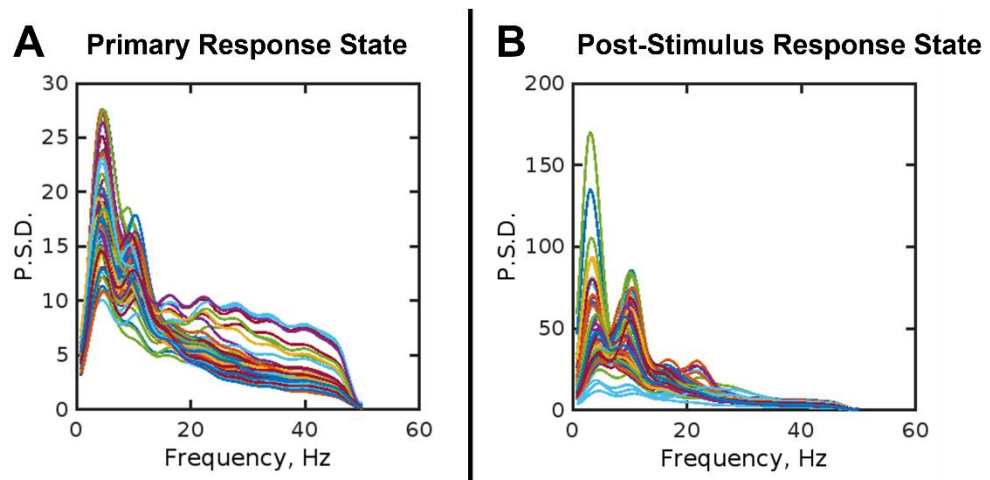


Figure 5.5: The spectral content of the primary (A) and PSR (B) states is shown here for all 78 brain regions. There is some variation across cortical locations. The primary response state has markedly lower PSD values which are dominant in the theta band with some regions incorporating an alpha component. The PSR state also has a theta peak, but most sensors also include an alpha and beta component.

For the connectivity analysis, regions in the brain which did not show a clear stimulus and/or PSR to the task were excluded (dark blue areas of the connectivity matrix (Figure 5.6, left panel) and grey areas of the brain (Figure 5.6, right panel)). Of the remaining regions, the primary and post-stimulus responses never occupied the same state. The burst-wise connectivity analyses of the primary and PSRs for those regions where a clear response was seen are shown in Figure 5.6. The bursts of activity for the primary response state are most likely to be coincident between the left and right sensorimotor regions. By averaging the connectivity matrix in one direction, the overall connectivity of each region was established, and showed that the contralateral sensorimotor area has the greatest overall primary response connectivity to the rest of the brain (Figure 5.6A, right).

In contrast, the PSR state visits are coincident between visual hemispheres and contralateral sensorimotor areas (Figure 5.6B). The region of the brain where the connectivity to the rest of the brain was the strongest was the right mid-occipital gyrus (Figure 5.6B, right). There is little evidence for a bilateral

sensorimotor network in the PSR, which is perhaps unsurprising given that the ipsilateral motor areas did not exhibit clear post-stimulus activity (Figure 5.4B).

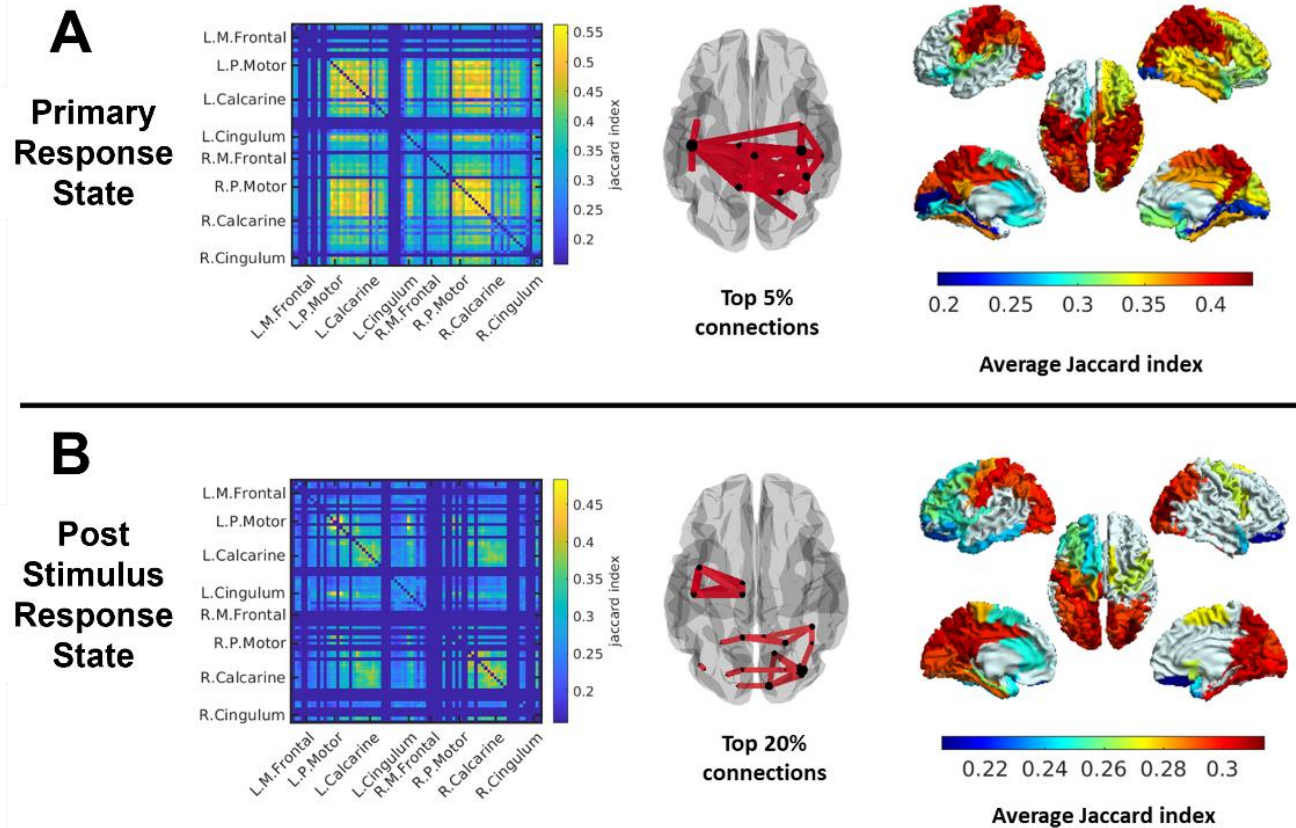


Figure 5.6: State connectivity maps for the grip-force task. Panel A shows that the primary response state exhibits clear left-right connectivity over the sensorimotor areas. The Jaccard connectivity metric for all pairs of regions is shown on the left. The top 5% of connections from this matrix is plotted in the middle, and on the right the sum over the matrix in one dimension produces a measure of the overall connectivity per region. The brain region exhibiting the highest average connectivity was the left sensory cortex with Jaccard index 0.43 ± 0.06 (mean \pm std over regions). The connectivity for the PSR state is shown in panel B – there is connectivity within the contralateral sensorimotor areas and between the left and right visual areas implying that the visual feedback component of the task evokes a rebound-type response across both hemispheres. The region with the highest average connectivity was the right mid-occipital cortex with Jaccard index 0.31 ± 0.03 (mean \pm std over regions). Note that the dark blue stripes in the connectivity matrix and grey areas of the brain in the right-hand plots are due to regions being excluded from connectivity analyses because there was no evidence of task-related activity. Note the difference in scales used between panels A and B.

5.3.2 Visual Stimulus Results

The data from two participants were removed from further analysis. For one of these participants the 3D head digitisation did not save correctly so there was no way to reliably coregister the participant's brain with the MEG sensor geometry. For the other participant less than half of the data were acquired because of technical issues in the lab during the scanning session. The results are therefore from the remaining 13 subjects.

After bad trials were identified (those with SQUID resets, ocular or muscle artefacts) and removed, there were on average 14 ± 1 (mean \pm standard deviation) trials left for each of the 2, 5 and 10s stimulus durations.

Classical analysis

The alpha time course, averaged over participants, for each of the stimulus durations is shown in Figure 5.7. There is a clear primary response (with a significant difference from rest ($p < 0.005$, RM-ANOVA)) regardless of the definition of the control window. However, there is only a significant difference between the post-stimulus and control windows when a conventional rest period of 5-7s is used ($p = 0.017$, RM-ANOVA). Interestingly, when a "true" baseline is used (15-23s from stimulus off), there is no significant difference between the two windows ($p = 0.261$, RM-ANOVA). This bias is not carried through into the HMM analysis because the model inference uses data from the entire experiment and does not require a specified baseline.

Regardless of the baseline used, there is no difference in the PSR amplitude between stimulus durations ($p = 0.324$, RM-ANOVA, for the 5-7s baseline; $p = 0.68$, RM-ANOVA, for the 15-23s baseline). Interestingly from Figure 5.6 there is also no clear modulation of post-stimulus rebound duration with stimulus duration. For the primary response, there is a significant difference between the 10s and 2s stimulus durations ($p = 0.012$, RM-ANOVA) when a 5-7s baseline is used, but there is no significant difference between stimulus durations when a 15-23s baseline is used.




Figure 5.7: The effect of changing the baseline on the visual PSR. Both A and B show the average amplitude envelope of alpha oscillations from the peak virtual electrode in the visual cortex for each of the stimulus durations. The shaded area accompanying each time course is the standard error over participants. When using the conventional baseline window of 5-7s from stimulus off there is a measurable PSR but no significant difference in amplitude between stimulus lengths (A). The post-stimulus window is not significantly different from rest when a true baseline of 15-23s post-stimulus is used (B).

HMM analysis

The spatial distribution of the primary and PSR states are shown in Figure 5.8. There is a bilateral pattern of activity across visual regions for the primary response state (Figure 5.8A), matching the bilateral response from the grip-force task. The strongest response was in the contralateral primary visual cortex. There is very little activity across the rest of the brain which makes sense given that this was a passive visual experiment and therefore not cognitively demanding. In contrast the PSR is clearly confined to the contralateral visual cortex. This mirrors the unilateral post-stimulus state result from the grip-force task. The pattern of activity across the cortex is clearly different for the primary and PSRs across two different experiments; in

both cases there is a clear bilateral/unilateral pattern of activity for the primary/post-stimulus responses, respectively.

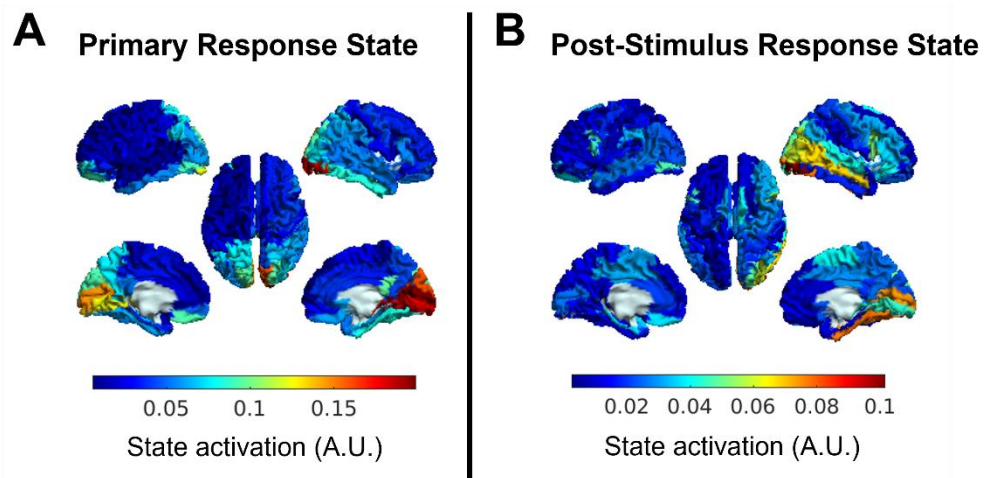


Figure 5.8: The spatial distribution of the primary and PSR states for the visual experiment. The primary response is bilateral over the visual cortex (panel A) whereas the PSR is confined to the contralateral visual regions (panel B).

Both the primary and PSR states show clear theta and alpha frequency components (Figure 5.9). Each coloured line represents the spectrum from one of the 78 AAL brain regions, and the variation in spectral content over regions is expected because only a relatively small number of brain regions exhibited task-related activity.

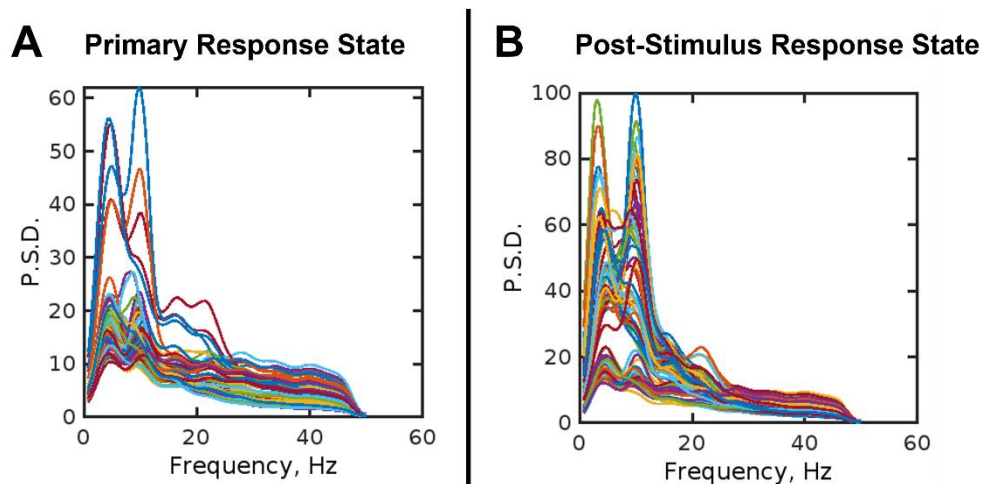


Figure 5.9: The variation in spectral content of the primary and PSR states across brain regions is shown in panels A and B respectively. Both states exhibit large variation across brain regions, with a small number of locations showing both a theta and alpha component.

For the connectivity analyses, regions with no evidence of a task-related response were removed and the results are shown in Figure 5.10. Of the

remaining regions, the primary and post-stimulus responses never occupied the same state. The primary response state coincidence across regions was greatest bilaterally across the lateral visual areas, whereas the PSR state coincidence was more focussed on the primary visual areas in the contralateral cortex. The strongest connection to the rest of the regions showing a primary response was in the right inferior temporal cortex (Jaccard index 0.35 ± 0.05 , mean \pm std over regions) whilst for the PSR it was in the right mid-occipital cortex (Jaccard index 0.35 ± 0.02 , mean \pm std over regions). It is interesting to note that whilst the PSR was generally more localised (Figure 5.8 panel B) this did not result in a weaker connectivity. Note that the strongest connectivity values do not necessarily correspond to those regions with the greatest likelihood of entering the primary or post-stimulus states.

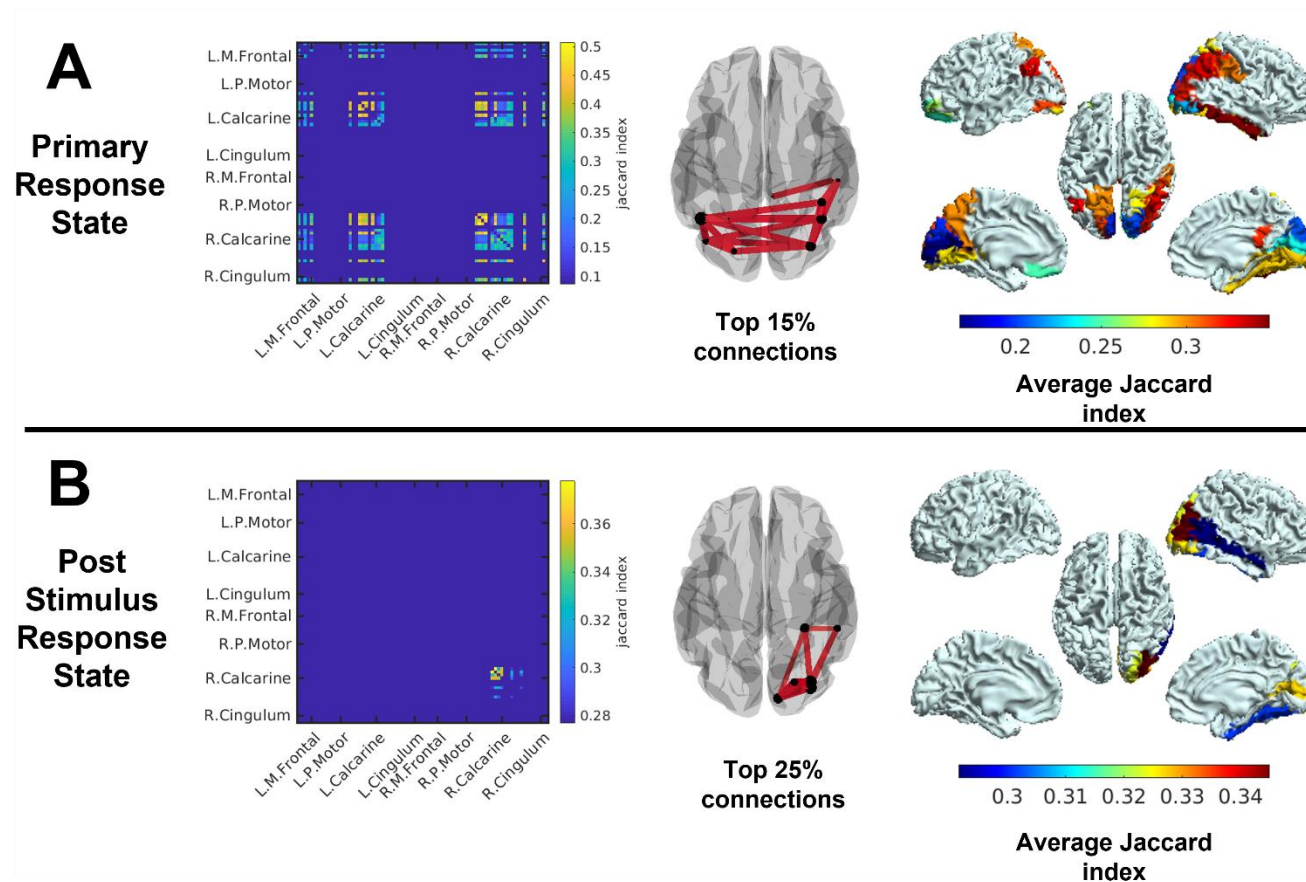


Figure 5.10: State connectivity maps for the visual experiment. Panel A shows that the primary response state exhibits clear left-right connectivity over occipital areas. The Jaccard connectivity metric for all pairs of regions is shown on the left. The top 15% of connections from this matrix is plotted in the middle, and on the right the sum over the matrix in one dimension produces a measure of the overall connectivity per region. The brain region exhibiting the highest average connectivity was the right inferior temporal cortex with Jaccard index 0.35 ± 0.05 (mean \pm std over regions). The connectivity for the PSR state is shown in panel B – there connectivity is confined to the contralateral visual areas. The region with the highest average connectivity was the right mid-occipital cortex with Jaccard index 0.35 ± 0.02 (mean \pm std over regions). Note that the dark blue stripes in the connectivity matrix and grey areas of the brain in the right-hand plots are due to regions being excluded from connectivity analyses because there was no evidence of task-related activity.

5.4 Discussion

5.4.1 Classical analysis

This work follows a number of studies attempting to understand both primary and post-stimulus responses by modulating them through experimental tasks, most notably (for example, (Pakenham et al. 2020; Fry et al. 2016)). These studies concluded that the PSR could be modulated by either the task duration or difficulty. In this work, by implementing an entirely passive visual experiment with varied stimulus durations, we were able to show that the PSR amplitude is not modulated by duration. This task also provided no evidence that the stimulus duration affected the duration of the PSR in agreement with (Pakenham et al. 2020). Therefore, these results provide more evidence supporting the hypothesis that post-stimulus amplitude is modulated by task difficulty. It is still unclear from this experiment what caused the modulation in post-stimulus duration which was observed in the rate force development task previously investigated (Fry et al. 2016).

The result of the classical analysis of the visual data replicates findings from (Pakenham et al. 2020) showing that a short inter-trial interval can alter the interpretation of the PSR. In this case the PSR could only be discerned with a 5-7s baseline as opposed to the true baseline of 15-23s. This feeds into the ongoing uncertainty of defining baseline and when the brain is “at rest”. From Figure 5.7 it appears that the brain reaches a more stable state from 10s after stimulus cessation with all traces then relatively stable with similar apparent variance. However, that then raises the question of whether the longer stimulus duration has in this case caused a bi-phasic PSR with a post-stimulus rebound and then under-swing, which warrants further investigation. What these results show is that different baseline choices can change the results of traditional analyses significantly, with baselines established in literature potentially being too close to the stimulus window. Careful consideration of the length of a rest window is therefore imperative for future studies. However, an alternative approach will be to move analysis to other data driven

methods. For example, since data from the whole experiment are used for hidden Markov model inference, there is no need to define a baseline, therefore eliminating this bias from analysis using the HMM.

5.4.2 HMM analysis

The HMM analysis used in this study follows the multivariate analysis of these data in (Pakenham et al. 2020) and the work done in Chapter 4 (Seedat et al. 2020) to create a model which allows us to tease apart the temporal relationship between regions and identify ‘bursts’ of activity at a single trial level which are related to the task. The time-delay embedded observation model allowed states to be defined by their spectral content, and the fact that the model inference was done on broadband (1-48Hz) filtered data meant that a pan-spectral picture of connectivity could emerge (with different regions exhibiting different spectral peaks). The advantage of this analysis approach over the one used in Pakenham et al. was that there was no requirement for the stimulus (or post-stimulus) state to have the same spectral content across different brain regions allowing a less constrained interrogation of the brain regions exhibiting a response. It follows that the interaction between different brain regions (with different spectral peaks) was similarly less restricted than the analysis performed by Pakenham et al. This was important to understand more about the functional role of PSRs across the cortex.

The HMM analyses of the grip-force and visual experiments show that both the primary and PSRs are observed across the cortex in a number of brain regions used to engage in the task. The spatial maps showing the likelihood of any given brain region entering the primary or post-stimulus states (Figures 5.4 and 5.8) as well as their connectivity profiles (Figures 5.6 and 5.10) provides new evidence supporting the hypothesis that the primary and PSRs are a result of different cognitive processes. This is because, in the same task, very different brain regions are recruited for each of these processes and there are two distinct spectral signatures which are found for the primary response state and post-stimulus state in regions where both are observed. It is also intriguing to see that the primary response state is bilateral across both hemispheres of

the brain with strong connectivity between them (panel A, Figs. 5.6 and 5.10) and that the PSR is very much unilateral and confined to the contralateral cortex (panel B, Figs. 5.6 and 5.10). This pattern is seen consistently across two different tasks.

The findings presented here for the connectivity between regions in different time windows are somewhat in contrast to previous work using alternative analysis strategies and to our hypotheses. Previously, connectivity between sensorimotor hemispheres, for example in the beta band, has reduced during the stimulus period, increased in the post-stimulus period, and then returned to resting levels of connectivity (O'Neill et al. 2017; Brovelli et al. 2017; Tewarie et al. 2019). However, the previous work has been generally constrained to a specific frequency band, typically the beta band. Beta band synchronisation is well known to decrease during the stimulus, but it is not suggested that the brain is not active during the stimulus. Furthermore, it is unclear when studying only the beta band if the connectivity has been reduced purely by a reduction in signal to noise cause by the decreased signal amplitude. Therefore, the lack of connectivity of the primary motor cortex with the rest of the brain when interrogating only specific frequency bands is unlikely to give the full picture of the brain dynamics at work. The HMM potentially allows a method which captures the brain entering different phases of activity in different time windows and brain regions without being constrained as to precisely what the activity is (i.e., it's spectral content) thus giving us a more holistic picture of the brain function and network connectivity. Alternatively, it is possible that the stimulus state identified with the HMM is a signature of the brain doing nothing due to pan-spectral desynchronisation (or there being insufficient signal to noise to detect what it is doing). However, even if this is the case the HMM is detecting all the regions which enter this "nothing" state and how the activity (or lack thereof) in these areas overlaps.

5.4.3 HMM state interpretation

Both a strength and limitation of this method is that it is based purely on the probability of entering a state in a given task window. This makes it incredibly

flexible as these data do not need to belong to a specific neurophysiologic frequency band, nor do they have to be confined to the desynchronisation and rebound phenomena we are used to observing. State visits simply need to occur in the time windows during and immediately after a stimulus. Inevitably this raises the question about what exactly is happening at each of these spatial locations. The HMM in this case is simply acting as a marker, highlighting regions of interest where a change in the probability of entering a state occurs within the time window of interest. Further analysis of these regions and the bursts of activity occurring in them is required to gain a fuller understanding of the biophysical meaning of the state which has been identified. This could be something as simple as extracting a time course from the regions of interest and producing an amplitude envelope of activity (as was done in Figure 5.7), or it could be to extract information about the duration and frequency of state visits and the amplitude of oscillations when the data are in a given state (analogous to the burst statistics in Chapter 4, Figure 4.2).

5.4.4 The impact of state durations on connectivity

To understand more about the connectivity patterns which were observed using the HMM in this study and what was driving this connectivity we investigated the relationship between the duration of state visits and the strength of connectivity for both the primary response and PSR states for both tasks. The duration of the state visits was longer for the primary response state (Figure 5.11) than for the post-stimulus state for both tasks. Interestingly in both states there was considerable variation in mean duration of state visits over brain regions compared with the resting state data considered in Chapter 4, as shown by Figure 5.12. This is interesting because the HMM inference process was almost identical for the two chapters, indicating fundamentally shorter bursts of activity in the absence of a task compared with either of the task-related states which were identified. The variation in state visit durations between brain regions is likely to reflect cognitive function as it is credible that some brain regions need to stay in their “task” or “post-stimulus” state longer than others to ensure the task is executed correctly. Furthermore, it is also

conceivable in a resting state (Fig 5.12) that the brain activity is fluctuating rapidly, so staying in any one state for a long period is not required.

Perhaps unsurprisingly the state durations are correlated with the connectivity measures used here. This makes sense intuitively – those regions with long state visits exhibit larger connectivity values because they are more likely to have bursts coincident with other brain regions. This may indicate that the PSR state is more meaningful with connectivity values driven by shorter targeted bursts of activity, compared with the primary response state whose connectivity may be more driven by state durations. However, for both the primary response and PSR state there is clearly a strong correlation between the state duration and the connectivity value for a given region. Further investigation into the characteristics and timings of the primary and PSR state visits is needed to draw any concrete conclusions.

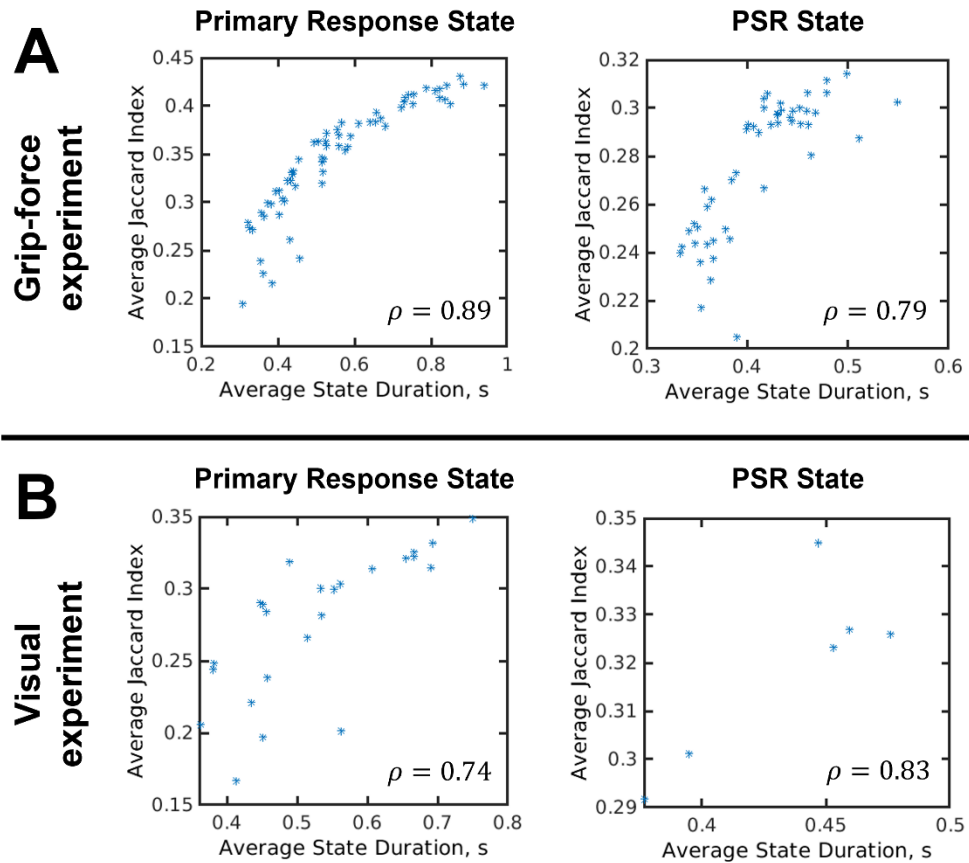


Figure 5.11: State durations correlate with Jaccard index for both the primary response state and the post-stimulus response state across both tasks. Note that the primary response state visits are generally longer than the PSR state visits. Regions with little evidence for a primary or PSR state have been removed as described in Section 5.2.4.




Figure 5.12: Burst state durations correlate with Jaccard index (connectivity value) for the resting state data presented in Chapter 4. State durations are shorter than those from the grip-force or visual experiments presented in this chapter.

5.4.5 Future work

Application of this HMM approach or similar models to a more complex task which recruits higher order brain networks may produce further evidence of a PSR across the cortex. In fact, alongside the data acquired for the visual experiment, data were collected for an n-back paradigm. It will be interesting to see how the primary and PSRs are distributed across the cortex and how they are modulated by task difficulty. Given the hypothesis that the PSR is a top-down (feedback) process and that a greater post-stimulus amplitude is expected when a mistake is made requiring a larger update to the forward model to improve performance of subsequent trials (Fry et al. 2016; Koelewijn et al. 2008), it is likely that the number or duration of post-stimulus state visits increases to reflect this. This idea of the PSR being required to update a forward model may also explain why the PSR to the visual stimulus is smaller as there was no active aspect to this task and therefore little need to update and change performance. However, further work is still required to prove this hypothesis of the functional role of the PSR.

Something that has not been explored yet is clustering, based on the temporal characteristics of when the state is active in different brain regions, or directionality of primary or PSR 'bursts' (state visits) in the connectivity results. Current analysis provides information with regards to all connections between all regions, but it is highly likely, based on other analyses (Vidaurre et al. 2018; Baker et al. 2014; Quinn et al. 2018; O'Neill et al. 2017; Tewarie et al. 2019; Brovelli et al. 2017) that during the stimulus and post-stimulus states we might be identifying more than one network. We hypothesize that some regions will be more highly connected in one network and other regions in another. These networks may then also interact hence, they may all exhibit a post-stimulus state but may not be part of the same network, even if they exhibit similar overall connectivity values in Figures 5.6 and 5.10. Directionality of connections could provide us with useful information about burst propagation and shed more light on the mechanisms (at a single-trial level) underpinning the canonical brain networks observed in both fMRI (Biswal et al. 1995; Fox

and Raichle 2007) and through amplitude envelope correlation of MEG data (Brookes, Woolrich, et al. 2011; Hipp et al. 2012). One hypothesis might be that bursts arise in the contralateral motor/visual area and propagate to connected regions during the stimulus period but that they arise in higher cognitive areas and feedback to the contralateral primary sensory areas in the post-stimulus period. Further work using this HMM analysis framework is needed to develop this concept.

5.5 Conclusion

These experiments and analyses show that the primary and PSRs are observable across the cortex in two different tasks. The PSR is not modulated by the duration of visual stimulus, making it likely that task difficulty alters the amplitude of the PSR. This agrees with the predictive coding model. The HMM analysis showed that the primary and PSRs are observed in different spatial locations, have unique spectral properties, and have different profiles of connectivity. Together, these findings strongly suggest they must reflect different cognitive processes and likely have different functional roles.

Chapter 6

Hidden Markov Modelling of the Interictal Brain

For Laylah

In Chapters 4 and 5, the hidden Markov model was used to identify transient neuronal activity in the brains of healthy adult volunteers. It was able to provide useful information relating to the timing and location of spectrally specific bursts, shedding light on the mechanisms underlying healthy brain function. In this final thesis chapter, the HMM has been extended to a paediatric epilepsy patient cohort – their unique patterns of epileptiform activity can be characterised by their location and spectral content, and in fact, trained neurophysiologists identify these distinctive patterns by eye. Therefore, the hypothesis for this chapter is that the HMM will be able to automatically identify patterns of epileptiform activity unique to each patient. This work has been submitted to Human Brain Mapping for peer review.

Epilepsy is a highly heterogeneous neurological disorder with variable aetiology, manifestation, and response to treatment. It is imperative that new models of epileptiform brain activity account for this variability, to identify individual needs and allow clinicians to curate tailored treatments. Here, we use a hidden Markov model (HMM) to create a unique statistical model of interictal brain activity for 10 paediatric patients. We use magnetoencephalography (MEG) data acquired as part of standard clinical care for patients at the Children’s Hospital of Philadelphia. These data are routinely analysed using excess kurtosis mapping (EKM), however, as cases become more complex (extreme multifocal and/or polymorphic activity), they become harder to interpret with EKM. We assessed the performance of the HMM against EKM for three patient groups, with increasingly complicated

presentation. The difference in localization of epileptogenic foci for the two methods was 7 ± 2 mm (mean \pm SD over all 10 patients); and $94 \pm 13\%$ of EKM temporal markers were matched by an HMM state visit. The HMM localizes epileptogenic areas (in agreement with EKM) and provides additional information about the relationship between those areas. A key advantage over current methods is that the HMM is a data-driven model, so the output is tuned to each individual. Finally, the model output is intuitive, allowing a user (clinician) to review the result and manually select the HMM epileptiform state, offering multiple advantages over previous methods and allowing for broader implementation of MEG epileptiform analysis in surgical decision making for patients with intractable epilepsy.

6.1 Introduction

Epilepsy is a neurological disorder affecting approximately 50 million people worldwide (Beghi, Giussani, and Collaborators 2019). It is characterized clinically by the occurrence of seizures which are generated by abnormal electrical cellular discharges in the brain. Epilepsy falls broadly into two categories: focal (originating from a single location in the brain) or generalized (originating diffusely throughout the cortex). In patients with focal epilepsy, there can be a single focus or multiple discrete foci that generate seizures. For many patients, anti-epileptic medications can control both the severity and frequency of seizures. However, approximately 30% of patients with epilepsy do not respond completely to medications (Mohan et al. 2018) and for these patients, surgical resection of affected brain tissue is a viable treatment. However, this necessitates extensive pre-surgical planning to accurately locate the affected brain area(s) prior to resection.

Current evaluation of pharmaco-resistant epilepsy is accomplished using electroencephalography (EEG) (which measures electrical activity in the brain via assessment of electrical potentials at the scalp) alongside clinical factors and structural assessment using magnetic resonance imaging (MRI). This may be augmented by the nuclear medicine techniques of positron emission tomography (PET) and especially ictal (during seizure) single photon emission

computed tomography (SPECT). If candidate epileptogenic locations are identified via imaging, a more invasive intracranial EEG (iEEG) is often initiated where electrodes are placed on the brain surface or within the grey matter. This allows electrophysiological assessment with optimal sensitivity and spatial accuracy, prior to surgery. Despite this extensive surgical planning, fewer than 50% of patients are seizure free 5 years post-surgery, with this number dropping to 38% at 10 years post-surgery (Mohan et al. 2018). It is therefore clear that a greater understanding of this heterogeneous disease, as well as improvements in clinical evaluation are required to improve patient outcome.

MEG measures the magnetic fields induced by neuronal current flow. Unlike the electric potentials measured by EEG, magnetic fields are relatively unaffected by the high resistivity of the skull, resulting in less spatial distortion of MEG compared with EEG signals, and thus improved resolution. MEG is used in a growing number of clinical settings, particularly in epilepsy. Not only does MEG provide additional information about the location of the epileptogenic zone (Stefan et al. 2011; Agirre-Arrizubieta et al. 2014; Murakami et al. 2016; Nissen et al. 2016; Gofshteyn et al. 2019), it can also be used to distinguish epileptogenic regions from eloquent cortex (Kim, Chung, and Hwang 2013). Most importantly, a recent study by Rampp et al showed, in 1000 patients, that presurgical MEG increases the chances of a patient achieving seizure freedom post-surgery (Rampp et al. 2019).

MEG therefore has significant promise for assessment of patients with epilepsy, and the recent introduction of new technologies to capture the neuromagnetic field offer even higher spatial resolution, better sensitivity, and improved practicality, at lower cost (Boto et al. 2018; Hill et al. 2020; Boto et al. 2021). This means that MEG could become more established as the technique of choice for epilepsy evaluation. However, the detection of epileptogenic activity in MEG data remains a significant challenge. Recording data during a seizure is difficult due to uncontrolled patient movement, and for this reason most MEG recordings are limited to interictal (between seizure) assessment (although see e.g. (Tang et al. 2003)). Interictal events – sharply

contoured atypical signals (known as epileptiform activity, i.e. spikes, sharps etc.) are observable in resting MEG (and EEG) data and are generally assumed to originate from seizure onset zones, meaning that spatially mapping their origin offers useful information on the location of epileptogenic cortex. However, detecting such events can be challenging for two reasons. Firstly, events are sporadic and unpredictable and, in some patients, rare. Thus, capturing an event can sometimes be a challenge without lengthy recording sessions. Even when events do occur, it can take significant time for a neurophysiologist to identify them (manually) in a MEG recording due to the high channel density relative to clinical EEG.

Second, the temporal morphology of epileptogenic activity can vary markedly between patients; some produce spikes with/or without slow wave activity, which varies dramatically in amplitude. Other patients generate polymorphic bursts of “sharp wave” activity characterized by high temporal frequency signatures. In some individuals, the epileptic pattern repeats (like a template), in others it differs on each occurrence. This makes automatic detection algorithms challenging.

There are two commonly used analysis methods for localization of interictal activity, equivalent current dipole (ECD) fitting and excess kurtosis mapping (EKM). In ECD, interictal spikes are inspected visually (Bagić et al. 2011). Once identified at the sensor level, a current dipole model is used to approximate the measured magnetic field just prior to, or at the peak of, the spike; by letting the origin of the modelled dipole vary spatially, and then selecting the point at which the model best fits the measured field, it becomes possible to localize the brain region generating the spike (Ebersole 1997; Wheless et al. 1999). This technique works reasonably well in cases where high amplitude spikes are observed in isolation but is often not useful in cases where interictal activity includes polymorphic bursts. In addition, multi-focal epilepsies are a challenge since ECD requires a-priori estimation of the number of active regions.

By contrast, EKM (Robinson et al. 2004; Schwartz et al. 2010; Gaetz et al. 2015) is an automatic method which assumes only that the epileptiform signals of

interest are sharply contoured (relative to the typically rhythmic background MEG signals). Kurtosis is a measure of the shape of a statistical distribution; in cases of abnormal activity (for example, if a dataset has large spikes), its statistical distribution includes a large “tail” and thus begins to look non-Gaussian; this is quantified by increased kurtosis (also known as the 4th moment of the distribution). By application of a kurtosis algorithm to MEG signals extracted from multiple brain locations, it becomes possible to localize areas generating abnormal activity. In addition, EKM doesn’t require a-priori estimation of the number of epileptogenic regions. Further, it is not limited to spikes, but can be used to assess any atypical activity, provided it has high kurtosis. However, excess kurtosis also has limitations; it has low sensitivity to low amplitude polymorphic activity. Also, counterintuitively, in cases with rapidly occurring high amplitude spikes, excess kurtosis has diminished sensitivity, because the kurtotic signals are so common they begin to represent the mean. For these reasons, neither ECD nor EKM are a perfect solution to analysis of MEG data in epilepsy, and other methods, which can accurately and automatically identify epileptiform activity, map its spatial origins, and (in multifocal epilepsy) characterize relationships between regions, would be useful.

Hidden Markov Modelling (HMM) has gained traction in recent years as a method to elucidate complex neural dynamics in MEG data (Quinn et al. 2018; Vidaurre et al. 2018; Higgins et al. 2021). The method works by detection of repeated patterns of activity (known as states) in MEG data; patterns can be characterized based on a number of features, including amplitude, channel covariance, and spectral properties. In the present work, we aimed to test a hypothesis that an HMM could be used to identify, in time and space, epileptiform activity. We further aimed to show that, in at least some patients, this method would offer more information than the established EKM technique. In what follows, the data processing pipeline is described, and its application to MEG data acquired in multiple different patient groups, ranging from ‘simple’ cases with focal spike-and-wave (SW) epileptiform activity for a

single locus, to more complex cases with multifocal epileptiform activity exhibiting polymorphic bursts from multiple loci.

6.2 Methods

MEG data were acquired at the Children's Hospital of Philadelphia (CHOP) as part of standard clinical practice. They were then anonymised and sent to me to apply the HMM analysis. Clinicians and scientists (most notably William Gaetz) from CHOP helped me to interpret the results.

6.2.1 Patient identification and data collection

This study was determined by the IRB to have exempt status as it constituted secondary analysis of data captured for clinical purposes under the NIH common rule (Jan 2019). All research subjects were collected as part of clinical care at the Children's Hospital of Philadelphia. Data from ten paediatric patients undergoing presurgical epilepsy evaluation between the ages of 11 months and 17 years (median 8.18 years, see Table 1), were utilized in the study. There were 2 female and 8 male patients. Nine of the 10 had focal epilepsy with or without impaired awareness, and one (Patient 6) had epilepsy with combined focal and generalized features. Six patients had right focal epileptiform discharges while 4 had left focal epileptiform discharges. Most focal discharges were in the frontal or temporal regions with one patient (Patient 8) with central localization and one patient with both frontal and posterior temporal discharges (Patient 6). When known, aetiology was primarily confirmed or suspected structural abnormality (such as focal cortical dysplasia), with one patient (Patient 1) with confirmed genetic aetiology and 2 patients with, as yet unknown aetiology.

Patient	Sex	Age at time of recordings (years)	Epilepsy type	Interictal discharge morphology	Location of EEG discharges	MEG category	MEG location(s)	Etiology
1	Male	6.73	Focal with impaired awareness	Sharp waves	Right centrottemporal (C4/T4)	Focal, spike wave	Right frontal, right precentral	Genetic
2	Male	9.62	Focal with preserved awareness	Sharp waves	Right frontal (F4/F8/F10)	Focal, spike wave	Right anterior frontal	Structural
3	Male	5.67	Focal with impaired awareness	Sharp waves	Left temporal (T3 at times shifted to F3)	Focal, spike wave	Anterolateral left parietal	Unknown
4	Female	4.93	Focal with impaired awareness and secondary generalization	Sharp waves, spike waves and focal slowing	Right frontotemporal (F8 and T4/T6)	Focal, spike wave	Right posterior frontal	Suspected structural
5	Female	17.04	Focal with preserved awareness	Sharp waves and spikes	Right frontotemporal (T4/T6, F8, T10)	Focal, polymorphic	Right superior temporal gyrus, right middle temporal gyrus, right insula	Structural
6	Male	11.55	Combined	Sharp waves, focal slowing, and irregularly generalized discharges	Left frontal (F3 sharp waves and slowing), left posterior quadrant (T5/O1)	Focal, polymorphic	Left anterior frontal and left Sylvian fissure	Suspected structural
7	Male	16.52	Focal with impaired awareness	Sharp waves	Right frontotemporal (F4/Fp2/F8)	Multifocal, polymorphic	Right anterior and mid temporal	Unknown
8	Male	0.96	Focal with preserved awareness	Sharp waves	Right central (C4/Cz)	Multifocal, polymorphic	Right perirolandic, right anteromedial occipital lobe and hippocampal tail	Structural
9	Male	6.59	Focal with preserved awareness	Sharp waves	Left frontotemporal (T3/C3, F7/T3, and F3/C3)	Multifocal, polymorphic	Left frontal	Structural
10	Male	11.88	Focal with preserved awareness	Sharp waves and focal slowing	Left frontotemporal (T3 slowing, T3/T5, Fp1/F3/F7)	Multifocal, polymorphic	Left posterior temporal and left anterior to mid frontal	Structural

Table 6.1. Patient characteristics.

Multiple 2-minute recordings were acquired using a CTF 275-channel MEG system operating in third order synthetic gradiometer configuration. (The number of acquired datasets ranged from 15 to 29; mean = 17.8). Data were acquired at a sample rate of 1200Hz. In all cases patients were scanned supine. Of the 10 patients, 6 were scanned whilst sedated using general anaesthesia.

Prior to data acquisition, three head position indicator (HPI) coils were attached to fiducial points on the head. During recording, these coils were energized (at non-physiologic AC frequencies) to allow continuous localization of their position relative to the MEG sensor array. All MEG scans were followed by an anatomical MRI during which MRI contrast markers were placed at the same fiducial points on the head. Coregistration of the MEG and MRI fiducial locations thus enabled complete spatial mapping of the MEG array relative to individual brain anatomy. This coregistration, in turn, allowed generation of functional images showing the cortical origins of epileptiform activity.

Patients were then grouped into one of three MEG categories; 4 had focal epilepsy generating interictal spike and wave activity; 2 had focal epilepsy with only polymorphic bursts, and the final 4 had multi-focal epilepsy with polymorphic bursts. These categories were determined by visual inspection of the MEG data (using EKM).

6.2.2 Data Processing

Following MEG acquisition, data were visually analysed and any 2-minute runs containing obvious interference or segments where patient motion exceeded 1 cm were removed from further analysis. Since some patients have infrequent interictal activity, all data were checked by a MEG expert to ensure they contained epileptiform activity and runs which did not were subsequently discarded. This left an average of 11 runs per subject. Our HMM-based method of mapping epileptiform activity comprised a two-step process (See Figure 1): In step one, the HMM was applied to channel level MEG data to identify time periods in which interictal epileptiform activity occurred. In Step two, a beamformer was used to localize the brain regions generating the activity identified by the HMM. These two steps are described in detail below.

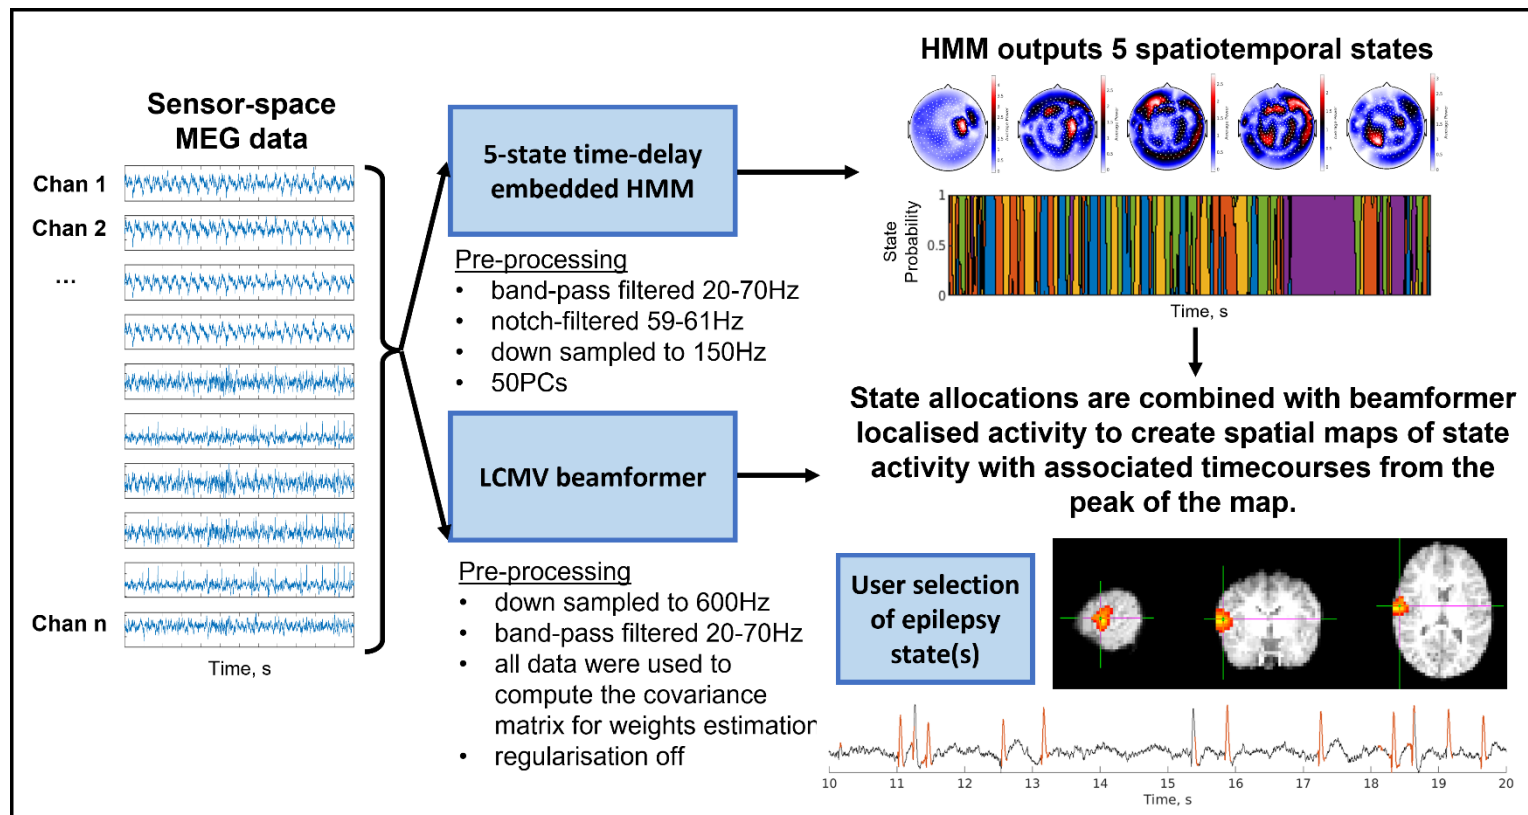


Figure 6.1: Schematic representation of the HMM-based process to identify epileptogenic activity. A multi-variate time delay embedded HMM was used to identify 5 states, each state characterized by its mean, covariance (across channels), and spectral content. A beamformer was used, along with temporal state allocations, to generate images of state activity across the cortex and a time course of activity from the peak voxel. This allowed us to identify an epilepsy state and a map of epileptogenic cortex.

6.2.3 Hidden Markov Modelling

To find spatio-temporal patterns corresponding to epileptogenic activity, we applied a multi-variate, 5-state, time delay embedded HMM, in channel space. An HMM assumes that a series of recurring mutually exclusive ‘hidden’ states governs the MEG data, such that every point in time is associated with one of the states. The sequence is assumed to be Markovian (in other words, the state active at time point, t , only depends on that active at time point $t-1$). An observation model links the HMM state to the observed values in the MEG data.

The HMM has been described extensively in previous papers (Vidaurre et al. 2018; Baker et al. 2014) and a complete mathematical description is given in Section 3.3. Briefly, in its simplest form, a HMM would describe each state using a multi-variate distribution; in other words, a mean (for all channels) and covariance (across channels). The 5 distributions that best described the data would be derived, and the probability of each data-point belonging to a specific state would be calculated. The number of states is defined a-priori and model inference would learn the sequence of states, from the observed data.

Here, we employed a more complex model which also allowed time-delay embedding (Vidaurre et al. 2018), adding information on autocovariance (defined over a specified time window (duration 70ms)). These state autocovariance patterns contain the spectral content of the signal. Consequentially, a single state is defined based upon *signal variance*, *covariance across channels*, and *spectral content*. This model had the potential to characterise both spike and wave activity and polymorphic bursts; in the former case the (typically high) amplitude of a spike, with a full width at half maximum of approximately 70ms, coupled with its distinct spectral content would characterise the state and differentiate it from ongoing “normal” activity. In the latter case, since polymorphic bursts are associated with “sharps” (high frequency activity) we again reasoned that distinct spectral content would define a state.

Prior to application of the HMM, the data used for the model inference (comprising 266-channel sensor space MEG data) were bandpass filtered between 20Hz and 70Hz (to match the standard EKM pipeline used by the Children’s Hospital of Philadelphia – see also below), notch filtered at 60Hz (to remove mains frequency artifacts), temporally down sampled to 150Hz, time-embedded using 70ms lags and a principal component analysis used to reduce the data to 50 components (to allow for faster model inference). Note that each 2-minute clinical run was considered separately, just as the EKM data were. The model inference itself was undertaken using a variational Bayesian method which seeks to minimise the free energy of the system. We computed 5-states; for each state, in addition to an observation model, we obtained a time course of the probability that the state is active. These time courses were thresholded at two thirds, thus generating a binary time course delineating when the state was “active”. From the time courses, we also obtained a *state-transition-matrix* – a 5 x 5 matrix of probabilities defining the temporal relationship between states (for example, element 2,1 in the matrix would represent the probability that state 1 followed state 2). This approach might offer useful information in cases with multi-focal epileptiform activity where one source consistently precedes the other.

6.2.4 Beamforming

Following application of the HMM (in sensor space), a linearly constrained minimum variance (LCMV) beamformer (Robinson and Vrba 1998) was used to localize the spatial signature of each state in the brain. The brain was divided into a regular 4mm grid of voxels, and each voxel time course was defined as a weighted sum of sensor measures. The beamformer weights were defined using a data covariance matrix calculated in the 20-70 Hz frequency band, and a time window spanning the entire recording. To maximize spatial resolution, no regularization was applied (Brookes, Mullinger, et al. 2008). The forward field was calculated using a current dipole approximation and a multiple local sphere volume conductor model. The beamforming parameters (frequency

filters and time windows for covariance estimation and regularization, as well as the choice of forward model) were selected to match the EKM method.

This resulted in a time course estimate of electrophysiological activity for each voxel location. To generate a functional map showing the spatial signature of each state, we frequency filtered the voxel time courses between 20 Hz and 70 Hz and then imposed the binary state time course to calculate the signal variance when the state was active, and the signal variance when the state was inactive. We then calculated the ratio, for every voxel in the brain, to delineate the brain regions which elicit changes in variance when the state switches on. In addition to the spatial maps, we also used the beamformer to derive a time course of electrophysiological activity at the peak location defined by the spatial map. Beamformer weights were defined as above but using covariance calculated in the 1-150Hz band, and a single time course showing 1-150 Hz activity extracted.

Having derived a spatial map and time course of activity for each of the 5 states, these were visually inspected by a single MEG-Epilepsy expert (W.G.). Those states whose time courses showed epileptiform activity when the state was active were identified, and the spatial localization was noted for each run. These were termed the “epileptiform state(s)”.

6.2.5 Comparison to existing methods

We compared the results of our HMM, to the more established EKM technique. We selected EKM for this comparison because of its advantages over ECD (Hall et al. 2018) and its use in large pediatric cohorts (Gofshteyn et al. 2019). To ensure that a standard EKM pipeline was followed, we used commercial software developed by CTF (Coquitlam, BC, Canada) known as SAM(g2), and the established pipeline used clinically by the epilepsy team at the Children’s Hospital of Philadelphia (Schwartz et al. 2010).

Prior to the application of EKM, all data were filtered 20-70Hz. In the SAM(g2) implementation, the brain was divided into a regular grid of 5 mm voxels and a scalar beamformer (equivalent to that described above) used to reconstruct

electrophysiological signals at each voxel. A kurtosis value was then computed for the time course at every voxel, and voxels with a kurtosis value greater than 0.5 were marked as peaks (Schwartz et al. 2010; Gofshteyn et al. 2019). Having found the spatial locations of interest, 1-150Hz time courses were analysed and temporal markers were placed at any point where the peak-to-RMS ratio exceeded 6 (Gofshteyn et al. 2019). This automated algorithm provided an estimate of both the regions and the time points which were likely generators of epileptic activity. Following this, data were inspected visually by a single expert experienced in identification of epileptic activity, to determine whether the abnormalities found by the algorithm were genuinely related to epilepsy, or were generated by sources of no interest with high kurtosis (e.g. ocular or muscle artefact). Temporal markers unrelated to epileptic activity were disregarded.

Following EKM, two separate measures were derived to quantitatively compare the output of the HMM and EKM mapping:

1. **Spatial correspondence:** Peak locations identified by the HMM (i.e. those regions whose variance increased when the HMM derived epilepsy state was entered) were compared to peak locations in kurtosis. Euclidean distance was measured between these peaks.
2. **Temporal coincidence:** All the time points identified by the EKM pipeline as containing epileptic activity were inspected to determine the number which were temporally coincident with an occurrence of the epilepsy state. Here temporal coincidence was defined as within 73ms of the EKM marker (i.e. the approximate duration of a single spike). Note that, the addition of a small amount of time around the EKM derived marker can mean that a single EKM marker is matched by more than one occurrence of the HMM epilepsy state; consequently, counterintuitively, values can exceed 100%.

All of the above (HMM, EKM, and their comparison) was applied to each 2-minute run, in each subject, separately. This meant 119 runs in total over 10 subjects.

6.3 Results

Case study results for each of the three patient groups are shown here. For detailed analysis of all patients, please see Appendix C.

6.3.1 Focal spike and wave case

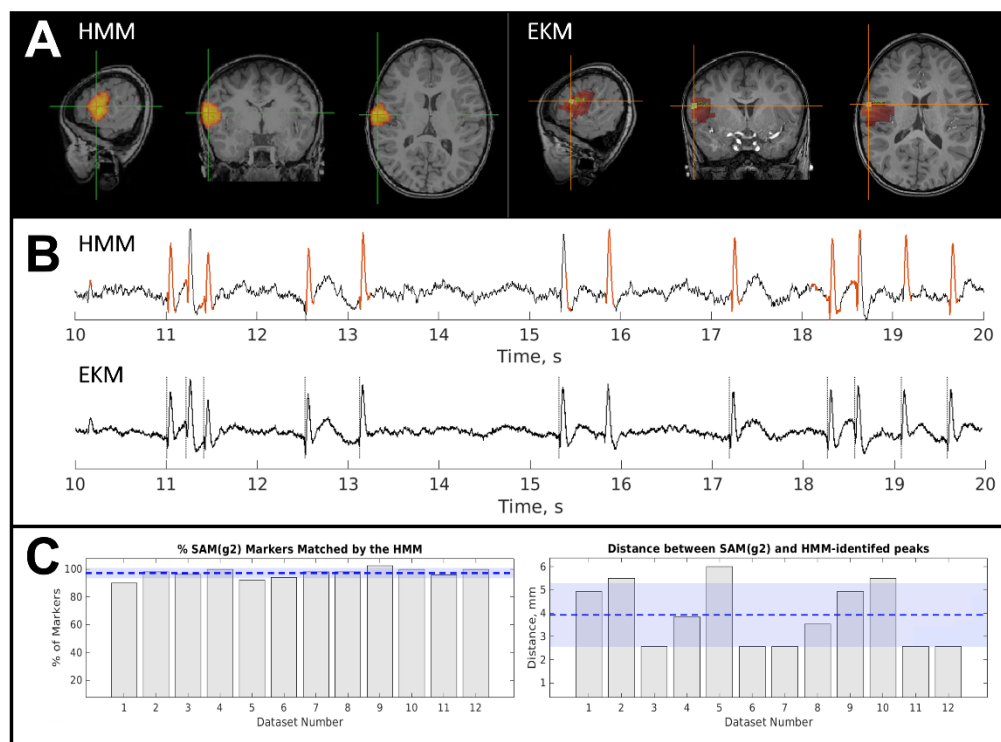


Figure 6.2: Epilepsy case 1 – patient 1; focal, spike and wave: A) The left-hand side shows the spatial signature of the epileptiform state as defined by the HMM. The right-hand side shows the spatial profile of EKM. Both maps were thresholded for visualization. B) Upper plot shows a beamformer derived time course from the peak location identified by the HMM with red regions showing occurrences of the HMM derived epilepsy state. The lower plot shows the equivalent time course from the peak in kurtosis. Dashed lines in the lower plot show time points of epileptogenic activity identified by EKM. C) Quantitative analysis over 12 runs from this subject. Left hand side shows temporal coincidence of EKM identified markers with the epilepsy state. Right-hand side shows spatial correspondence. Dashed line and blue shading shows the mean and standard deviation of the measure over all 12 2-minute runs in this subject.

Results for a single representative focal epilepsy patient are shown in Figure 6.2. Panel A shows the spatial signature of epileptogenic cortex, identified by

the HMM (left) and EKM (right). These data are taken from a single run. Panel B shows example time course segments taken from the peak voxel in the HMM map (upper trace) and the EKM map (lower trace) (note the similarity between traces due to the close spatial correspondence of the peak source locations identified by each method). In the upper (HMM) trace, the time points at which the epilepsy state was active are shown in red. In the lower (EKM) trace, the timepoints identified as containing epileptiform activity are shown by the dotted line. Note the close temporal correspondence (at least for this small segment of data). These results are quantified in Panel C; the left-hand bar chart shows the percentage of manually verified EKM markers that fell within the occurrence of the epileptiform state (this is our temporal coincidence metric). The right-hand bar chart shows the Euclidean distance between the peak from the HMM, and the peak from the EKM. In both cases, the separate bars represent different 2-minute runs in the same subject.

The patient exhibited frequent, large amplitude spikes, with some slow wave activity. Spatial correspondence was 4 ± 1 mm (mean and standard deviation over 12 2-minute runs) and $97 \pm 4\%$ of EKM-identified epileptic events were matched in time by the occurrence of the epileptiform state. Over all of the datasets, the HMM epilepsy state was active for $5.4 \pm 0.8\%$ of the time.

Three further cases of patients with focal epilepsy with spike and wave are shown in Appendix C (Figures C1, C2 and C3), all with excellent agreement between the HMM and EKM. In these cases, spatial correspondence was 6 ± 2 mm, 9 ± 5 mm, and 15 ± 3 mm, and the temporal correspondence was $102 \pm 8\%$, $100 \pm 1\%$, and $96 \pm 5\%$. In general, these support our hypothesis that the HMM performs similarly to EKM in enabling the identification of epileptiform activity in time and space, at least in patients with focal spike and wave epileptiform activity. Spatiotemporal correspondence was high in all cases and data showed good correspondence across many runs, for each subject.

6.3.2 Focal epilepsy with polymorphic activity

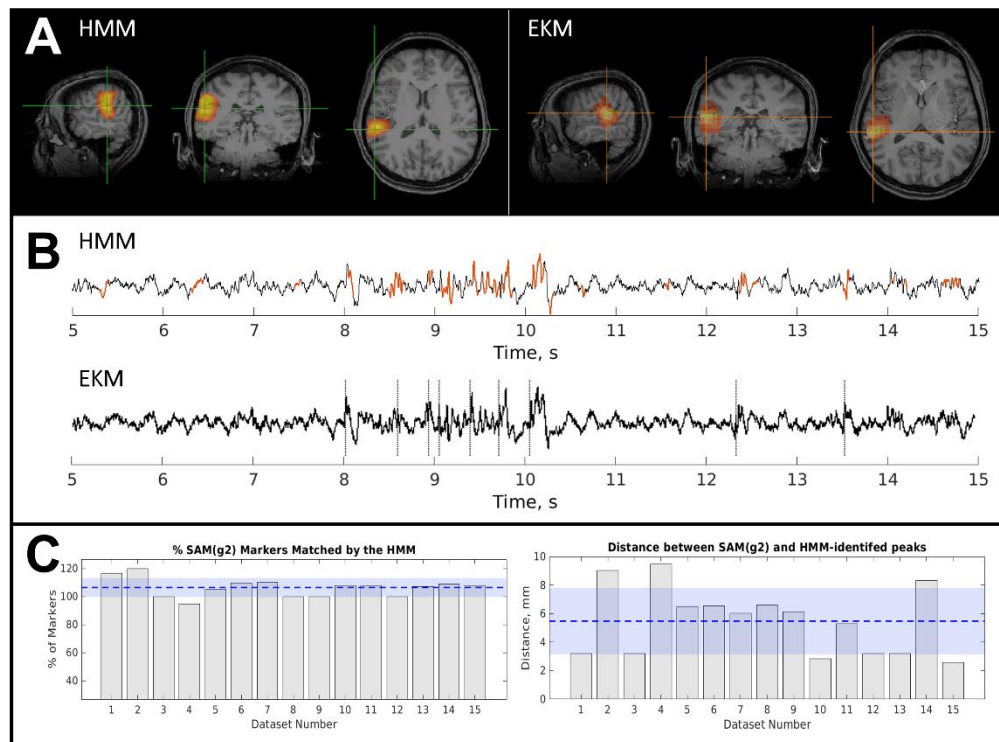


Figure 6.3: Epilepsy case 2 – patient 5; focal, polymorphic bursts: A) left: HMM derived map. Right: EKM derived map. Both thresholded for visualization. B) Upper: beamformer time course from the HMM peak; Lower: equivalent time course from the EKM peak. Dashed lines show time points of epileptogenic activity identified by EKM. Red regions show occurrences of the epilepsy state. C) Left: temporal coincidence of EKM markers with the epilepsy state. Right: spatial correspondence.

Case 2 (Figure 6.3) shows MEG data acquired in a patient with focal epilepsy, but without typical spikes in the MEG trace (Figure 6.3B). The patient's resting MEG data exhibited occasional polymorphic bursts of sharply contoured activity, which (unlike spikes) change their temporal morphology on each occurrence. Such data are not amenable to conventional ECD source analysis, however here we see that both the HMM and EKM generate a focal localization of the epileptogenic zone with excellent spatial agreement between the two methods. On average across 15 2-minute recordings, the spatial discrepancy between the HMM and EKM peak location was 5 ± 2 mm. In addition, the temporal coincidence of the epilepsy HMM state and the EKM derived markers was $106 \pm 7\%$. The percentage of time when the HMM epilepsy state was active was $6 \pm 3\%$ (average and standard deviation over all

runs). A second case of focal epilepsy with polymorphic bursts is given in Appendix C (Figure C4); results again were similar with a spatial correspondence of $5 \pm 3\text{mm}$ and temporal correspondence of $113 \pm 14\%$.

6.3.3 Multifocal epilepsy with polymorphic activity

The above results show relatively straightforward cases of focal epilepsy, where abnormal epileptiform activity arises from a single location in the brain. However, in more complex cases, abnormal activity can occur from more than one location (often simultaneously), and in such cases the challenge becomes determining where these multiple regions are, how they are related, and if possible, which region serves as the driver of an epileptiform network causing other areas to exhibit epileptiform activity.

Case 3 is a patient with multi-focal epilepsy; in total 15 datasets were acquired in this individual, results from a single representative run are shown in Figure 6.4. In this instance, two separate states were found to generate polymorphic epileptiform activity (shown in Figure 6.4B). The corresponding spatial maps show foci in left temporal and left frontal lobes. Interestingly, examination of state transition probabilities demonstrated that state 1 (left temporal) had an 89% likelihood of being preceded in time by state 4 (left frontal); in other words, a polymorphic burst in the frontal lobe was almost always followed by a polymorphic burst in the temporal lobe. Potentially, this suggests that the frontal lobe region is the driving signal and root cause of the epileptiform activity observed in the temporal lobe.

Across all 15 runs in this same subject, 10/15 allowed identification of 2 epilepsy states (including the run shown in Figure 4); 3/15 runs had 3 epilepsy states, and 2 runs had just 1 epilepsy state. Of the 13 datasets with more than one epilepsy state, 9/13 had a clear transition from the frontal to the temporal locations (transition probability $84 \pm 9\%$; mean and standard deviation across these 9 runs). The other 4 had no clear transitions between states. The percentage of time which the states were active was $3 \pm 1\%$ and $4 \pm 1\%$ for the temporal and frontal epilepsy states respectively.

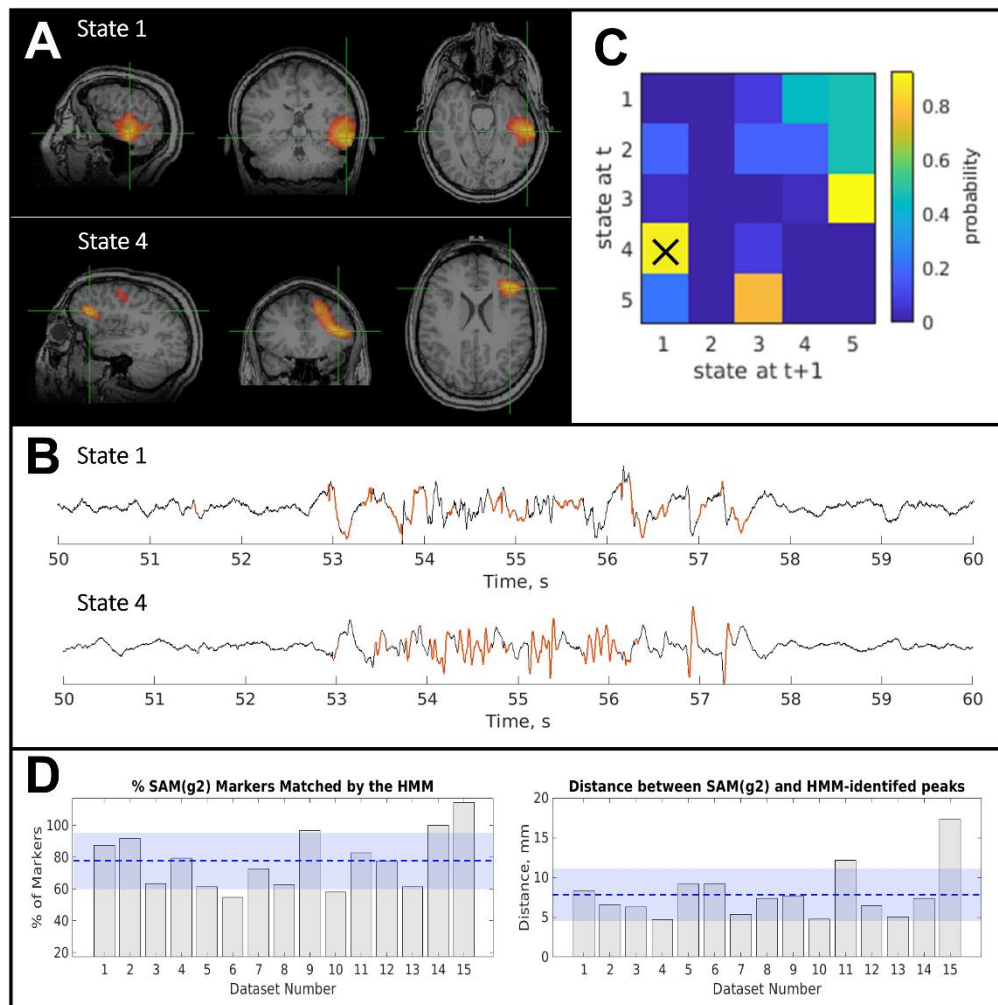


Figure 6.4: Epilepsy case 3 – patient 10; multi-focal, polymorphic bursts: A) upper: HMM derived map from state 1. lower: HMM derived map from state 4. B) Upper: beamformer time course from the HMM peak for state 1; Lower: equivalent time course from the HMM peak for state 4. Note both states show epileptiform activity. C) State transition matrix. Elements represent the probability of a transition between states; e.g. element (2, 1) would represent the probability of a transition from state 2 to state 1; element (1, 2) would represent the probability of a transition from state 1 to state 2. Panels A-C show the result from a single representative run. The average spatial and temporal correspondence between methods is shown in D with $78 \pm 18\%$ of EKM markers matched by an HMM state visit and an average Euclidian distance of $8 \pm 3\text{mm}$ between peaks. The amount of time spent in the epilepsy state was $3 \pm 1\%$ of the total time.

We also compared the results shown in Figure 6.4 to the output of our EKM algorithm. EKM generated a single map which also had peaks in frontal and temporal lobe. On average (across 15 runs) the mean spatial correspondence for the frontal lobe peak was $8 \pm 3\text{mm}$ and the equivalent distance for the temporal lobe peak was $7 \pm 4\text{mm}$, once again implying spatial correspondence between the HMM and EKM. This is impressive given the challenges posed by

such a complicated case to each of these methods. It is worth noting that although the EKM method places temporal markers in the data to help neurophysiologists assess whether a spike in one part of the brain precedes a spike in another area, the HMM provides additional information about the temporal relationship between the two brain locations using the state transition probabilities. Furthermore, the HMM uses all of the state data to estimate transitions, something which becomes particularly important in cases without clear spikes.

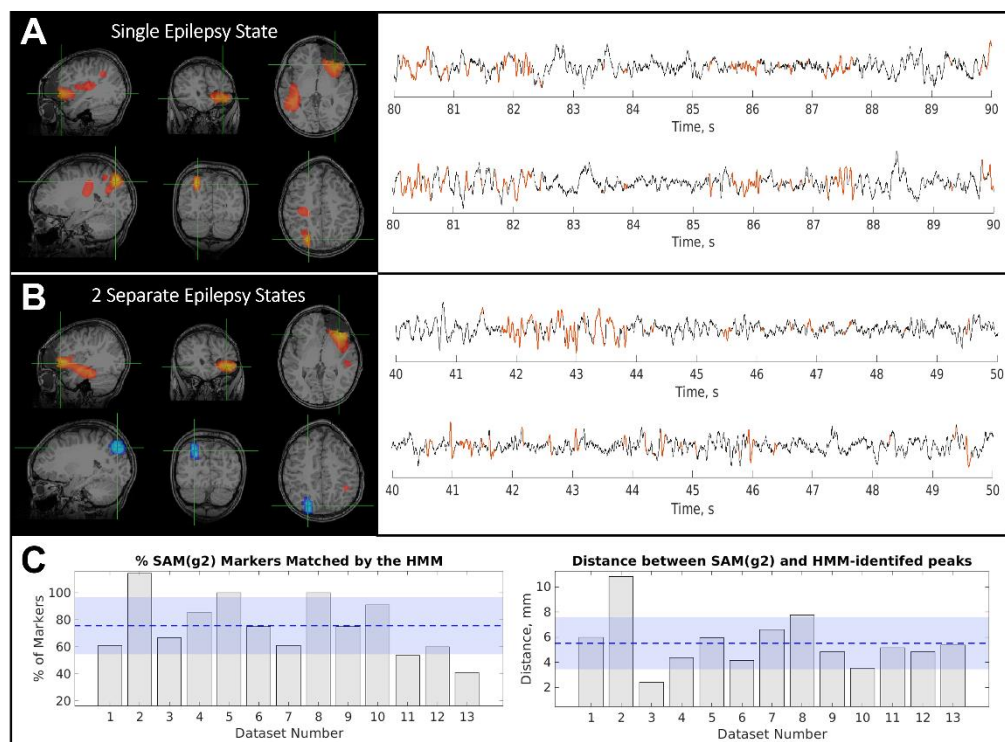


Figure 6.5: Epilepsy case 4 – patient 9; multi-focal, polymorphic bursts: A) a single run in which multiple brain locations, determined to be generating epileptiform activity, have been identified in the same state. B) a different run in the same patient where the HMM has split the candidate brain regions across states. Spatial maps were thresholded for visualization. C) the temporal (left) and spatial (right) match between the HMM and EKM. $76 \pm 21\%$ of EKM markers were matched by an HMM state visit and the average Euclidian distance between HMM peaks and the corresponding EKM peak was $6 \pm 2\text{mm}$. The total time spent in an epilepsy state was $16 \pm 6\%$ of the total time.

Case 4 (Figure 6.5) was another multi-focal epilepsy patient with a previous tissue resection in the frontal lobe. An EKM analysis showed two locations with elevated kurtosis, one close to the previous left frontal resection, and the second in right parietal cortex. Initially, the parietal activity had been dismissed as sharply contoured mu rhythm. However, the HMM showed that, in 6/13

datasets, a single HMM state described the activity in both brain locations. In other words, the state variance was high in both locations at the same time, implying that the activity (whilst of unclear clinical significance) is, in fact related. This is the case for the data shown in Figure 5A, which depicts the two separate locations in a single map and associated time courses from the two peaks. However, in a further 6/13 datasets the activity was parsed into two states (shown for dataset 2 in B as an orange and blue state). In the remaining dataset there was only the activity from the parietal location, with no apparent activity close to the resection zone. Two further cases of multi-focal epilepsy are presented in Appendix C, Figures C5 and C6. In both cases, results similar to those shown in Figures 6.4 and 6.5 were generated.

6.3.4 Group results

For each of the three patient groups, the average spatial and temporal correspondence between the two methods was found over subjects and this is shown in Table 2. There was good agreement between the two methods for all three groups (this is especially encouraging given the complexity of the multifocal patients), with less than a centimeter discrepancy between the peak locations. There was more variation in the temporal coincidence metric which is likely to be because there were much fewer EKM markers in the multifocal polymorphic data because the amplitude of the virtual electrode data rarely exceeded the threshold needed for marker placement. This meant that if the HMM missed a single marker, it resulted in a much-reduced percentage coincidence.

	Average temporal coincidence (\pm SD over subjects) (%)	Average spatial correspondence (\pm SD over subjects) (mm)
Focal spike wave	99 \pm 3	8 \pm 5
Focal polymorphic	110 \pm 5	5.45 \pm 0.03
Multifocal polymorphic	80 \pm 5	8 \pm 2

Table 6.2. Group results

6.4 Discussion

Epilepsy is a debilitating disorder in which both symptoms and treatments differ markedly across patients. In some cases where pharmacological intervention fails to control seizures, patients become candidates for surgery in which affected regions are resected. For many patients, such intervention offers seizure freedom (and thus a marked improvement in quality of life). However, success greatly depends on pre-surgical planning to accurately identify the epileptogenic region(s) and the current clinical pathway (involving EEG, MRI and iEEG) is not always successful in identifying candidate regions for resection (or indeed implantation of the iEEG). Consequently, improvements to this pathway could pave the way to enabling more patients to become eligible for surgery and might offer improved outcomes for those who do have surgery.

Recent work has demonstrated that MEG has a significant role to play in the management and treatment of epilepsy patients, offering high precision mapping of epileptogenic and eloquent cortex (Kim, Chung, and Hwang 2013; Schwartz et al. 2010). However, the current methods for analysis of MEG data are limited. ECD – still the most widely used technique – is unsuited to multi-focal epilepsies, and cases where temporal morphology of interictal events fails to include isolated high amplitude spikes. Some of the limitations of ECD are lifted via the use of kurtosis-based techniques, however these too offer limited sensitivity; they only respond to activity that is characterized by non-Gaussianity, and in multi-focal cases, EKM provides limited information about the relationship between brain regions, meaning identified brain areas can be discarded due to a lack of evidence that they are involved in an epileptic brain network. Consequently, there is room for more adaptable processing pipelines to analyse MEG data in epilepsy patients.

In this work, we have demonstrated that an HMM is a promising technique to localize epileptiform activity in both time and space. The HMM has been used increasingly in the analysis of MEG data in recent years, to elucidate the complex spatio-temporal dynamics of brain networks (Woolrich et al. 2013;

Baker et al. 2014; Vidaurre et al. 2018; Seedat et al. 2020). Here we show that, by seeking short segments of data with unique spatio-spectral characteristics, the HMM provides a useful platform for automatic detection of epileptiform activity. Importantly, the HMM doesn't rely on a single data characteristic such as high amplitude spikes, or a non-Gaussian statistical distribution. In addition, activity doesn't have to repeat in time (like a template). Rather, the HMM employs spectral characterisation so that morphological features of the data can vary on every occurrence of the epilepsy state. Consequently, as we have demonstrated, it can adapt to the heterogeneity of the disorder; the same algorithm will find spike and waves, sharps, or polymorphic bursts with equal efficiency. The flexible, data-driven approach of the HMM lends itself naturally to personalised medicine, where the output of the model will be specific to each patient.

We compared the output of our HMM to the established EKM technique (which is embedded in a number of clinical pathways for epilepsy across multiple centres). Spatial agreement was within 8 ± 3 mm (average and standard deviation across all subjects and runs in all three patient groups), showing that epileptic foci could be identified using the HMM with high spatial acuity. Temporal agreement was also reasonable, with an average of $94 \pm 13\%$ (over all subjects in all patient groups) of epileptiform events detected by the standard EKM pipeline matched by the occurrence of the epileptiform state, identified by our new HMM pipeline. Importantly, whilst the HMM looks for a specific spatio-spectral pattern, the EKM pipeline is only able to find the timing of epileptic bursts via the use of a threshold: specifically, the EKM algorithm looks for peaks whose amplitude is larger than some arbitrary threshold (6) multiplied by the RMS amplitude. Thus, epileptiform activity with lower amplitude is likely to be missed by this overly simplified technique. The HMM on the other hand relies on the occurrence of a specific multivariate spatio-spectral pattern. Consequently, it is less likely to miss epileptiform bursts.

The HMM has demonstrated advantages when looking at multi-focal epilepsy cases. In such complex cases, ECD is limited since one must make an a-priori

decision on the number of active foci. EKM is better, since it allows imaging of multiple regions, and the manual comparison of kurtosis markers can identify whether one sharply contoured event from one location consistently precedes another. Similarly, the HMM can identify multiple epileptiform foci but provides more information regarding the temporal relationship between the identified regions. Specifically, where multiple foci appear within a single state, then those regions, by definition are related temporally, in other words, the regions are “active” together. This was seen in epilepsy case 4 (Figure 6.5) in which (in 6/13 cases) an epileptogenic focus close to an area of resected tissue in the frontal lobe appeared alongside a second peak in parietal cortex. The signals from the second peak were not recognisably epileptiform and thus had initially been dismissed, however the fact that the activity generated increases in variance in synchrony with the peak close to the resection zone indicates that it is likely a part of the epileptic network causing the frontal activity. Thus, simply by virtue of the epileptic state containing two peaks, we gain more information than we would with an equivalent image acquired using EKM. Similarly in case 3 (Figure 6.4), We discovered two states containing epileptiform activity. In this case, each state contained only one region. However, the states demonstrated a temporal dependence according to the transition probability matrix; specifically, a burst of activity in temporal lobe was almost always preceded by a similar burst in frontal lobe. This again is an example of how the additional information afforded by the HMM might be used to provide valuable information about an epileptic network. In sum, the evidence presented here not only demonstrates localization of epileptiform activity, but also points the way to how the underlying mathematical framework of the HMM can offer more information beyond that gained from existing approaches.

There are a number of possible extensions to the present methodology which should be considered. First, the current version of the model operates on sensor space data, and localization is achieved using a beamformer. It is theoretically possible to apply the beamformer first, and then the HMM in

source space, and indeed this may offer additional information, for example on the connectivity between epileptic foci in multi-focal cases. However, it is also significantly more computationally demanding since the HMM must be run on many thousands of (voxel based) signals rather than a few hundred channels. This was found to be impractical, and whilst brain parcellation offers a means to reduce dimensionality, it also reduces spatial specificity which is of key importance in this application. Furthermore, the epileptic focus may not lie exactly within one anatomical region but may be split across regions posing additional challenges in using a method of parcellation for the analysis. Thus, we feel the method presented is the most practical currently, but the use of better computing may ultimately offer the opportunity to expand it. In addition, although we have applied our method to MEG data, application to EEG data or concurrent EEG/MEG data would likely also be valuable. In the former case, the limitations surrounding the spatial resolution of EEG would mean it likely that one would get low spatial specificity, but nevertheless the HMM might offer a useful means to automatically identify time points in EEG data which contain interictal events. The advent of wearable MEG systems (Hill et al. 2019; Boto et al. 2019) may also make ambulatory MEG a real possibility. In this case one would no longer have to rely on short 2-minute runs, but much longer datasets could be analysed. The increased amount of data would undoubtedly create a more accurate model which could help to overcome the variations in HMM output over multiple runs. Finally, whilst in this paper we concentrate mainly on spikes, sharps and polymorphic bursts, there are other forms of atypical signalling (for example, fast ripples (>200Hz) and slowing (1-4Hz)) which are also consistently observed in patients with epilepsy. Unlike ECD or EKM methods, HMM could readily be extended to focus on these spectral features.

6.5 Conclusion

In conclusion, the HMM offers an alternative to ECD and EKM for identifying and localizing epileptiform activity in the human brain. In 10 subjects we found good spatial agreement between methods, with the HMM able to provide

localisation which matched, within a few mm, that from EKM. Moreover, we found that the HMM ostensibly offers more information about epileptiform activity generated at multiple locations. As the use of MEG continues to grow in epilepsy, particularly given the advent of new, cheaper, and more practical MEG systems, clinical pathways should look to employ the HMM, and related techniques, as a means for spatiotemporal mapping of epileptogenic cortex.

Chapter 7

Conclusion

Throughout this thesis, hidden Markov modelling has been used to identify features in MEG data that would otherwise have been missed using traditional analysis pipelines. The HMM was employed in three different scenarios – two contributed to our understanding of the healthy adult brain, and the final chapter provided information about the epileptiform activity in the brains of paediatric patients.

7.1 Overview

In Chapter 4, the HMM identified pan-spectral bursts which were previously thought to be confined to the beta band. This highlights a key strength of the time-delay embedded HMM; it is able to characterise states based on their spectral content without prespecifying the frequency band of interest *a-priori*. Waveform shapes, in unaveraged data that comprise multiple frequency components, are captured in their entirety. The interaction of the whole waveform with, for example, the behavioural outcomes of a task, is more relevant for the understanding of brain function, than splitting that waveform into its frequency components and then correlating task outcomes with each component separately. Interestingly, the frequency content of the bursts changed with task, with the resting state bursts having a more prominent alpha component than those bursts identified in a simple visuomotor task. The latter half of Chapter 4 showed that the known resting state functional connectome could be reproduced using a coincident burst metric, implying that connectivity is driven by bursts, which make up only a fraction of the MEG signal. Bursts also offer a unification of two different functional connectivity metrics, one based on oscillatory amplitude (amplitude envelope correlation) and one based on oscillatory phase (phase difference derivative). Our work suggests that coincident bursts are in phase, offering windows of maximal coherence, when non-coincident bursts are not.

This burst detection model was then extended to two cognitive tasks in Chapter 5 with a difference in the state selection criteria – states were now characterised by the time windows in which they were most likely to occur relative to the task. This allowed a primary and a post-stimulus state to be identified independently for each region of the brain, leading to spatial maps of the probability of entering either state. One of the key strengths of this method was that the PSR in different regions of the brain did not need to conform to a specific frequency band, which is important because different frequencies are dominant in different cortical areas. In the visual task, both the primary and post-stimulus responses were confined to the visual cortex but with the distinction that the PSR was unilateral, and the primary response was bilateral. This finding was corroborated by the grip-force results – this time the responses were seen in the sensorimotor areas as well as the visual cortex, but the PSR was mostly unilateral, and the primary response was bilateral. It would be interesting to see if this pattern persists across different cognitive tasks. It is also worth noting that neither the amplitude nor duration of the PSR was modulated by the length of stimulus presentation in the (entirely passive) visual experiment. This provides evidence that the PSR is in fact modulated by task difficulty.

Chapter 6 extended the HMM analysis to clinical data from paediatric epilepsy patients. These patients were grouped according to the complexity of their epilepsy and a sensor level HMM was used to parse the data into five states. The time windows for state activation were then used together with an LCMV beamformer to create 3-D maps of state variance across the whole brain. Importantly, the epilepsy state was not selected automatically – all state maps and virtual electrode time courses were available for inspection by the user so that the most relevant state(s) could be selected manually. This is because the aim of this project was to produce a viable tool to aid clinicians who will undoubtedly want to check the output of any machine learning algorithm. The HMM identified an epileptic focus in every patient which was in agreement with the kurtosis mapping method used currently in clinical settings. In one of

the most complex patients with multifocal epilepsy exhibiting polymorphic activity, the HMM identified two epilepsy states, one for each focus. The relationship between those states was then investigated and there was a clear transition from one of the epilepsy states to the other, implying the direction of propagation of epileptiform activity. This was information which could not have been found with the established kurtosis mapping method. Crucially, in Chapter 6, the data-driven nature of the HMM allowed models of brain activity to be inferred which were *specific to each patient*.

7.2 Future Work

There is a quote, often attributed to Aristotle, which says “The more you know, the more you know you don’t know.” which is apt because this thesis has raised more questions than it has answered.

Regarding bursts and functional connectivity, it would be fascinating to know whether there is a directionality associated with those bursts which are coincident but found in spatially separate regions – do bursts from one region lag those from the other? In the case of resting state functional connectivity, in a brain region with the highest frequency of bursts, they only make up roughly 10% of the time course. The question then follows: what is the other 90% of neuronal activity if it does little to contribute towards connectivity with the rest of the brain? This question may be explored with further models, perhaps an HMM applied to the remaining non-burst data.

Post-stimulus responses are robust phenomena which are altered by pathology and therefore have the potential to form a biomarker for a number of neurological diseases. However, to use them in this way, we must first understand how they are altered by cognitive tasks. The traditional way in which the primary and post-stimulus responses are identified (by frequency filtering and envelope averaging over trials) misses key information – namely the fact that these phenomena are made up of punctate ‘bursts’ of activity on the single trial level. New analysis tools, like the HMM, allow these phenomena to be studied in greater detail, and in a way that minimises bias (because states are not constrained to a predefined frequency band). There is also much more

to be learned about the primary and post-stimulus states identified in the visual and grip-force tasks in Chapter 5. The HMM identified regions where these neural signatures exist but the primary and post-stimulus state characteristics should be further investigated. This would help us to understand how the activity in both differs, in terms of spectral content, state durations or signal amplitude.

In the case of paediatric epilepsy, the HMM was able to identify an epileptogenic focus in all ten patients. This is impressive and implies its viability as a clinical method, but it is difficult to verify such techniques against a “gold standard” when no gold standard exists. The closest we can get would be to conduct a longitudinal study on epilepsy patients who undergo surgical resection and to compare the results of the surgery with the results of the HMM to confirm the ground truth focus of the epilepsy. It would also be important to include a larger number of patients to confirm statistical significance.

7.3 Closing remarks

MEG recording is arguably the most sensitive means to measure human brain function, offering rich spatial, spectral, and temporal content which is only matched by invasive recordings. However, such cutting-edge technology requires cutting-edge analysis to interpret the data, and machine learning will allow us to do this, as this thesis has shown with hidden Markov modelling. The scope of the questions which can be answered with such a simple, yet flexible model, is impressive: from detection of neuronal signatures which underly healthy functional connectivity, to those which are modulated by cognitive tasks, and those electrophysiological signatures associated with pathology. Interpreting single-trial waveform shapes in the data used to be labour intensive but can now be done with unsupervised machine learning, helping us to interpret the data whilst minimising bias. In a clinical setting, because the models are data-driven, and therefore unique to each patient, this will surely help pave the way towards personalised medicine.

The world of brain imaging has seen multiple paradigm-shifts in recent decades; the advent of functional imaging, the realisation that patterns exist in functional data, even in the absence of a task, and the understanding of how brain networks are disrupted in disorders, to name a few. It is tempting to speculate that machine learning will inspire the next big shift in thinking, as we move away from classical analyses and towards techniques which can fully exploit the rich tapestry of information that neuroimaging provides. The human brain, with its network of approximately 86 billion neurons, remains perhaps the most complex puzzle in science, but neuroimaging, combined with techniques like the HMM, are beginning to allow us to assemble some of the pieces.

Appendix A

Panel A of Figure A1 shows the spatial signature of burst duration, frequency, amplitude, and interval time, when the HMM is required to derive 6 states, instead of 3. Figure A1B shows the derived spectrum of the burst state for 3, 6 and 10 states. Comparison of this with Figure 4.2 in the Chapter 4 shows that, qualitatively, the spatial signatures of duration, frequency, amplitude, and interval time are largely consistent regardless of the number of states selected. Similarly, the spectrum of the burst state does not change markedly. However, quantitative inspection of the burst duration, frequency and interval time shows a significant change; for example, in left sensory cortex, burst duration dropped from 300ms to 200ms when moving to six states from 3. This is likely to be a result of the mutual exclusivity of states assumption in the HMM itself; that is, the same amount of time must be split among more states, so a particular state can separate into different ones as the number of states is increased from a low number to a higher one, with only subtle differences between them. This limitation is therefore analogous to the choice of threshold in the more commonly used thresholding approaches to burst detection – in other words, the higher the threshold, the shorter the burst. Here, bursts will tend to be shorter if a larger number of states is selected. However, an important point is that whilst bursts become notionally shorter, their spectral content remains the same – meaning that states do not split

across frequencies. This important point demonstrates the robust pan-spectral nature of burst content regardless of parameter selection.

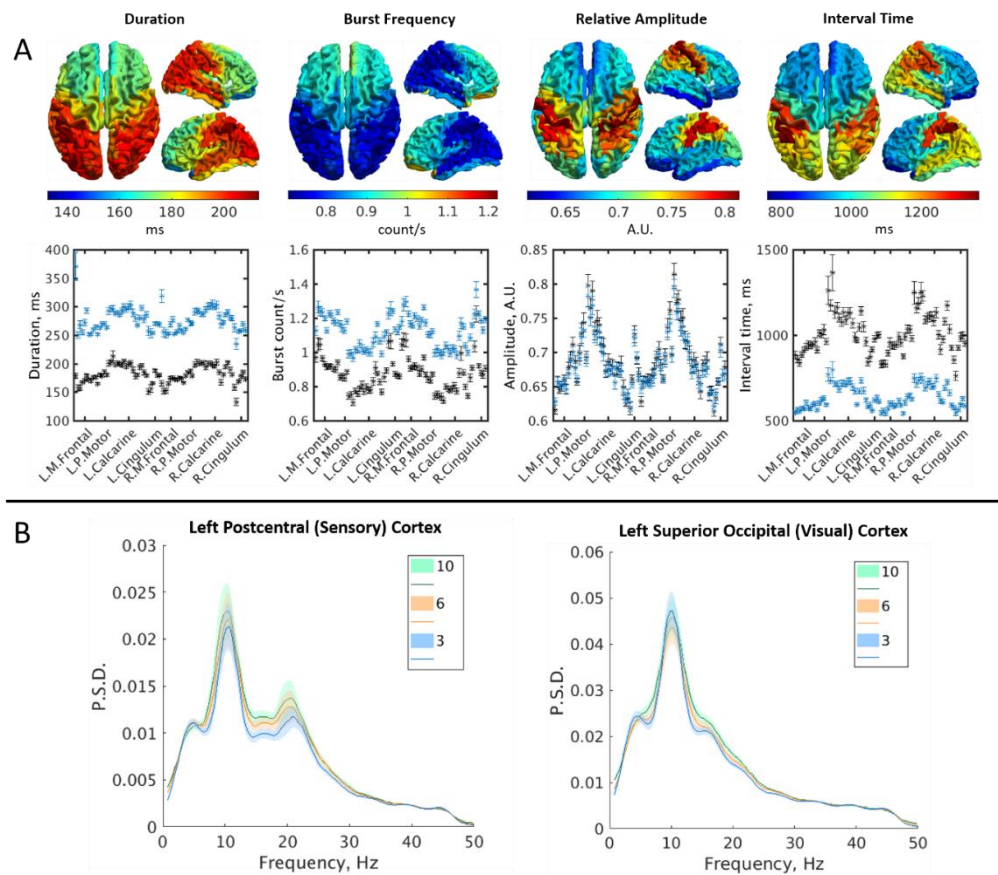


Figure A1: The effect of changing the number of states on burst characterisation. Four different burst properties are shown in (A): burst duration, frequency of occurrence, relative amplitude, and the time between bursts. The brain plots show the group-averaged results for bursts detected using 6 states in our model. These are compared directly with the 3-state burst output in the scatter plots below where the blue and black points show the 3 and 6 state outputs respectively. The error bars show the standard error over participants. B) The group-averaged burst-state spectra, computed using the multitaper, for bursts detected with a 3 (blue), 6 (orange), and 10-state (green) model. The spectra are shown for the left sensory and left visual cortices. Note that the spectra are remarkably similar, showing that the HMM consistently finds the same burst state with the same spectral content despite having changed the number of states modelled. The bursts in the sensory cortex have a more prominent beta peak than the visual cortex which is weighted more towards alpha.

Although this demonstrates a consistency in the model output, regardless of the number of states inferred, the method also relies on the selection of the burst state: this was the state with the highest correlation between its probability time course, and the amplitude of beta oscillations. Given that the increase in the number of states makes the burst durations shorter (see Figure

A1), one might reasonably expect that the burst state is splitting into two or more states. This would result in more than one state with a high correlation with the beta envelope. Therefore, more than one state might be needed to describe the full range of beta activity for a model inferring a larger number of states.

To investigate whether this was indeed the case, the correlation was calculated between the probability time course and the beta envelope for each state in the 3, 6 and 10-state models. This was done for all subjects in two brain regions – the left sensory cortex (where higher beta activity is expected) and the left orbitofrontal cortex (where little beta activity is expected). The states were sorted by their correlation strength with the amplitude of beta oscillations in that region, so state 1 was consistently the burst state, and state 2 was the next-most correlated state, and so on. The average correlation value for each state (over all subjects) is shown in Figure A2. In each case there is a clear state with a larger correlation value than the others, regardless of the number of states inferred by the model.

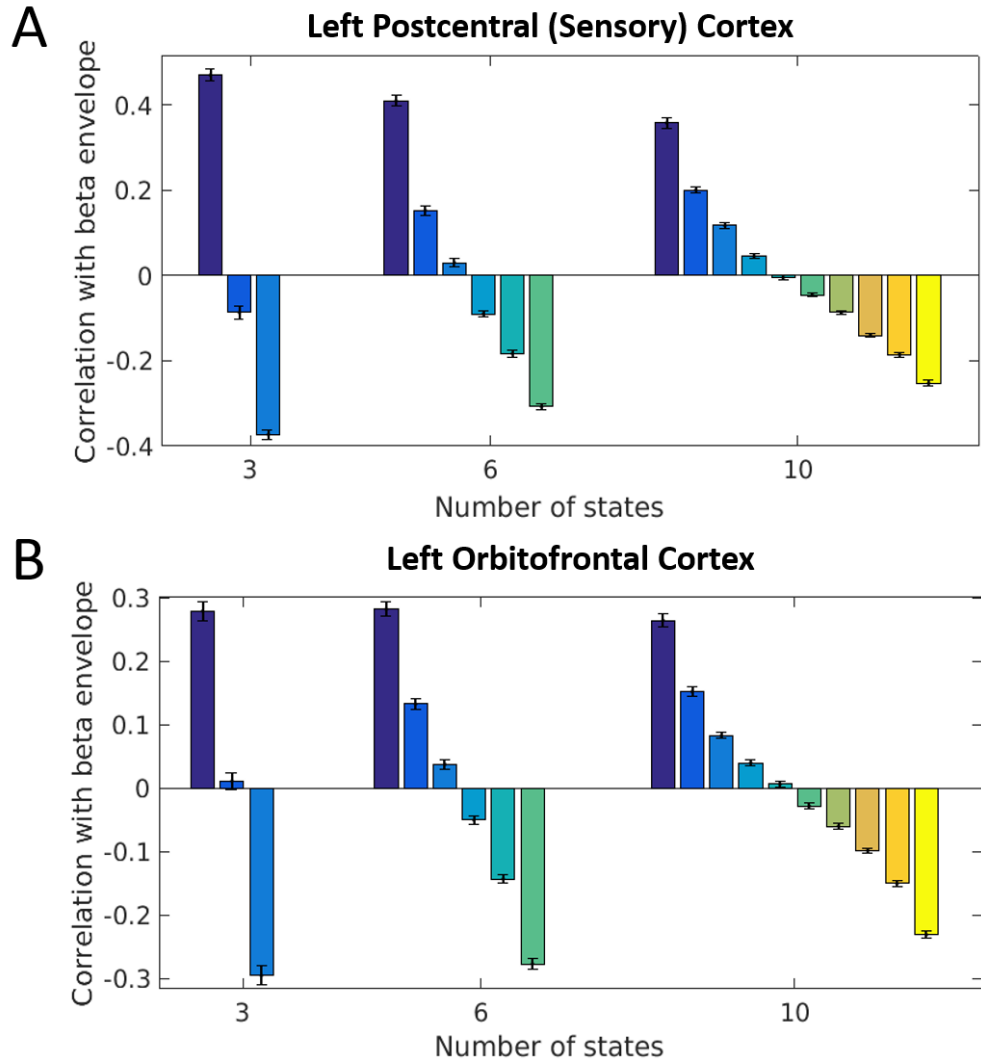


Figure A2: The correlation of each state with the beta envelope. For each of the three models (3, 6 and 10-state), the burst state was selected as the one with the largest correlation between its probability time course and the amplitude of beta oscillations. The average correlation value for each of the states (over all subjects) is shown here for two regions, the left sensory and the left orbitofrontal cortices. The first bar in each group represents the burst state; the second bar represents the next-most correlated state and so on. The error bars show the standard error over subjects. There is a clear single state which correlates best with the beta envelope in each case, indicating that the burst state is robust to changes in the number of states inferred by the model.

Appendix B

It is usual practice to investigate task-related connectivity in time windows of interest (such as the primary or post-stimulus windows or using a sliding window approach) (Brovelli et al. 2017; O'Neill et al. 2017; Tewarie et al. 2019). Given that the connectivity effects we aimed to measure were in the primary and post-stimulus time periods, those time windows were used when calculating the state-related Jaccard index connectivity in the main body of Chapter 5 (see Section 5.2.4). However, given that the timing of the state visits was crucial to their definition, the connectivity matrices should also be reproducible using all state visits (even those occurring outside of the relevant time window). The Jaccard index was therefore calculated for both the primary response and the PSR state over all time. The results for the grip-force experiment are shown in Figure B1. Although the connectivity values themselves are generally smaller than those from the windowed connectivity results, the general pattern of connectivity across the brain is maintained. This indicates that the timing of the state visits is meaningful throughout the entire experiment. Having said that, the region exhibiting the highest average connectivity for the primary response state has changed from the left sensory cortex to the right supramarginal gyrus. This is not wholly unsurprising because the supramarginal gyrus interprets tactile sensory data and the position of limbs in space. Both things will be important in the grip-force task where a visual cue is being coordinated with a motor response in the right hand. The fact that the peak location is now in the ipsilateral hemisphere is more surprising but could be explained by the strong left-right connectivity in the

primary response. The region with the highest average connectivity for the PSR remains the right mid-occipital gyrus.

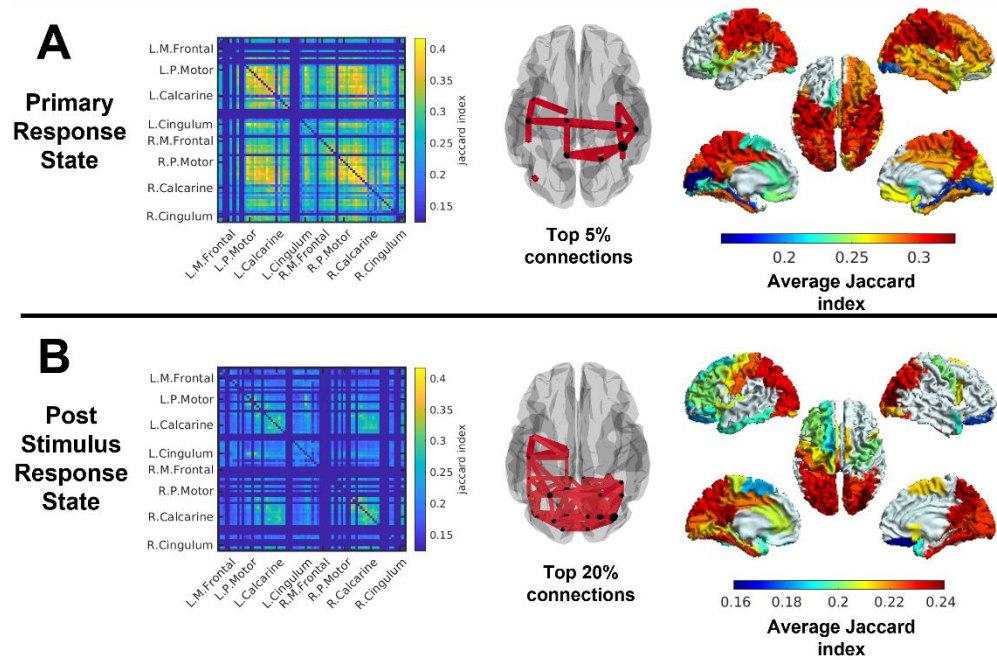


Figure B1: State connectivity maps for the grip-force experiment. The Jaccard index was calculated between every pair of regions using state visits from the whole experiment (not those state visits constrained to a particular time window). The primary response state connectivity is shown in panel A and the post-stimulus response state connectivity is shown in panel B. From left to right the figures shown the Jaccard connectivity matrix, a plot of the top 5% (primary response) and 20% of connections (PSR), and the average connectivity of each region (created by averaging over the connectivity matrix in one direction). The region with the highest average connectivity for the primary response is the right supramarginal gyrus with Jaccard index 0.33 ± 0.06 (mean \pm std over regions), for the PSR it is the right mid-occipital gyrus with Jaccard index 0.24 ± 0.05 (mean \pm std over regions).

These results are corroborated by those from the visual experiment (see Figure B2). The overall pattern of connectivity over the cortex remains unchanged but the Jaccard index values themselves have reduced. The regions with the greatest average connectivity remain the same for both the primary and post-stimulus states, right inferior temporal cortex and right mid-occipital gyrus respectively.

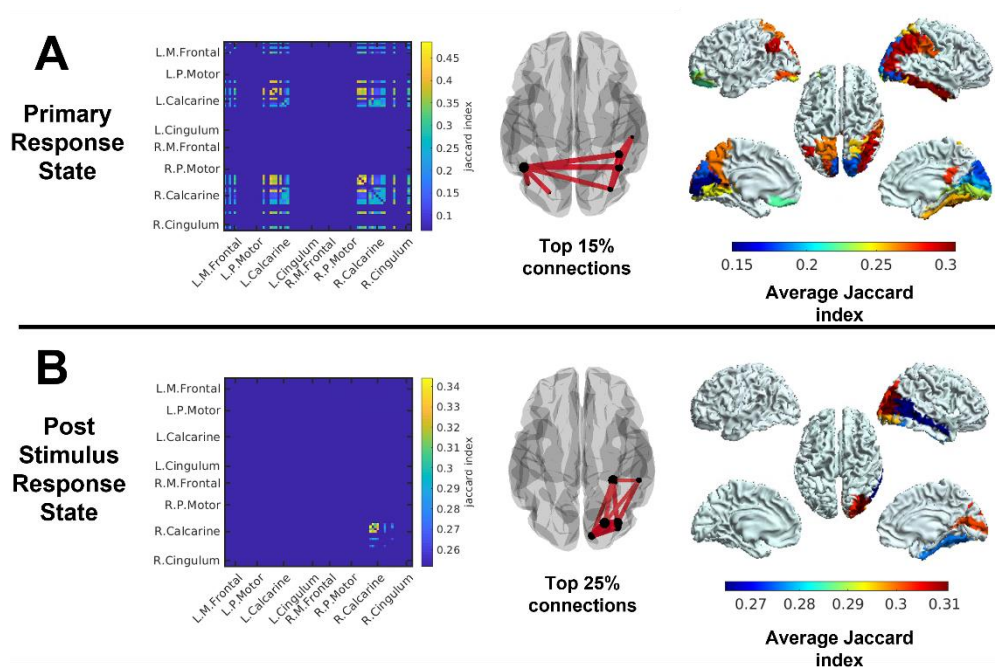


Figure B2: State connectivity maps for the visual experiment using coincidence from all state visits across the whole experiment, not just the time windows of interest. Panel A shows the results for the primary response state with the Jaccard matrix (left), top 15% of connections displayed on a brain (middle), and average connectivity for each region (right). The brain region with the highest average Jaccard index was the right inferior temporal gyrus (0.31 ± 0.07 , mean \pm std over regions). Panel B shows the result for the PSR. The most connected region was the right mid-occipital gyrus with average Jaccard index 0.31 ± 0.02 (mean \pm std over regions).

Appendix C

Results for the remaining 6 epilepsy patients are shown in this section.

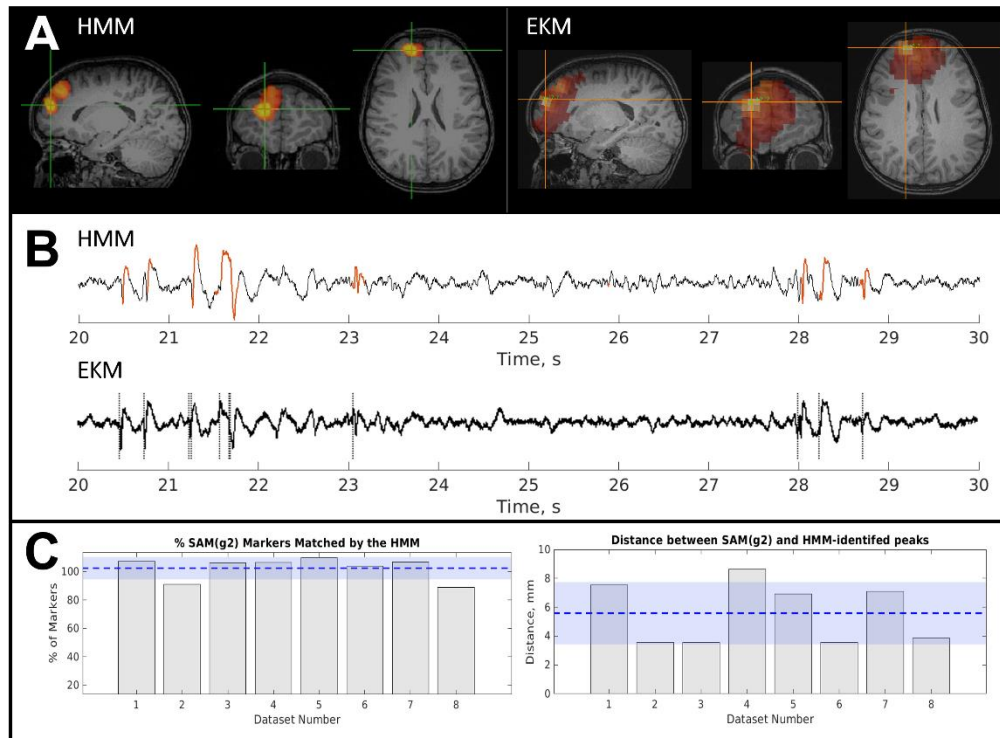


Figure C1: Patient 2 - focal, spike and wave: The HMM and EKM localize epileptogenic cortex in the right frontal lobe as shown by the spatial maps in panel A. The time course of neural activity from the peak voxel is shown in B with the HMM state activations shown in red and the EKM markers as dotted lines. The spatial and temporal correspondence between the techniques are quantified in panel C with 6 ± 2 mm between peak locations and $102 \pm 8\%$ of EKM markers matched by a HMM state visit. The average length of time spent in the epilepsy state was just $4 \pm 2\%$ (average and SD over runs) of the total time. Note that more than one state visit to a single kurtosis marker will occasionally yield values $>100\%$.

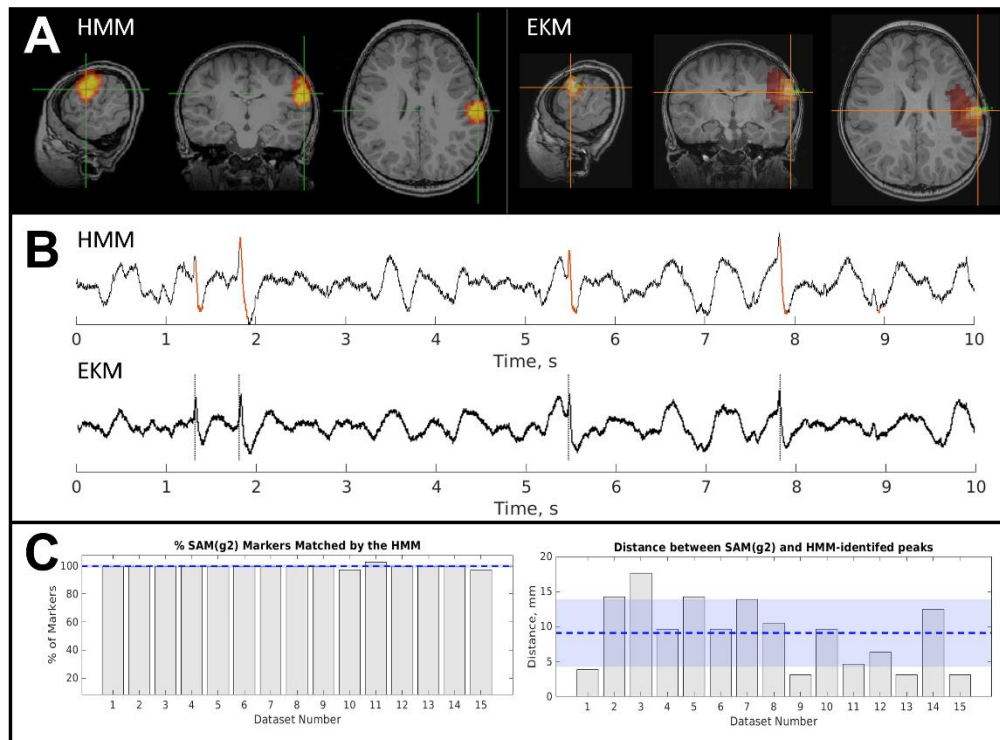


Figure C2: Patient 3 - focal, spike and wave: Panel A shows spatial maps highlighting the location of epileptogenic cortex using the HMM (left) and EKM (right). The time course of activity from the peak locations are shown in B for the HMM (upper) and EKM (lower). The HMM state visits are highlighted in red, and the EKM markers are shown as black dotted lines. There is good spatial and temporal agreement between techniques (C) with $100 \pm 1\%$ of kurtosis markers being matched by an HMM state visit and $9 \pm 5\text{mm}$ between peak locations (average and SD over runs). For this patient, the epilepsy state was active for $3.3 \pm 0.5\%$ of the total time.

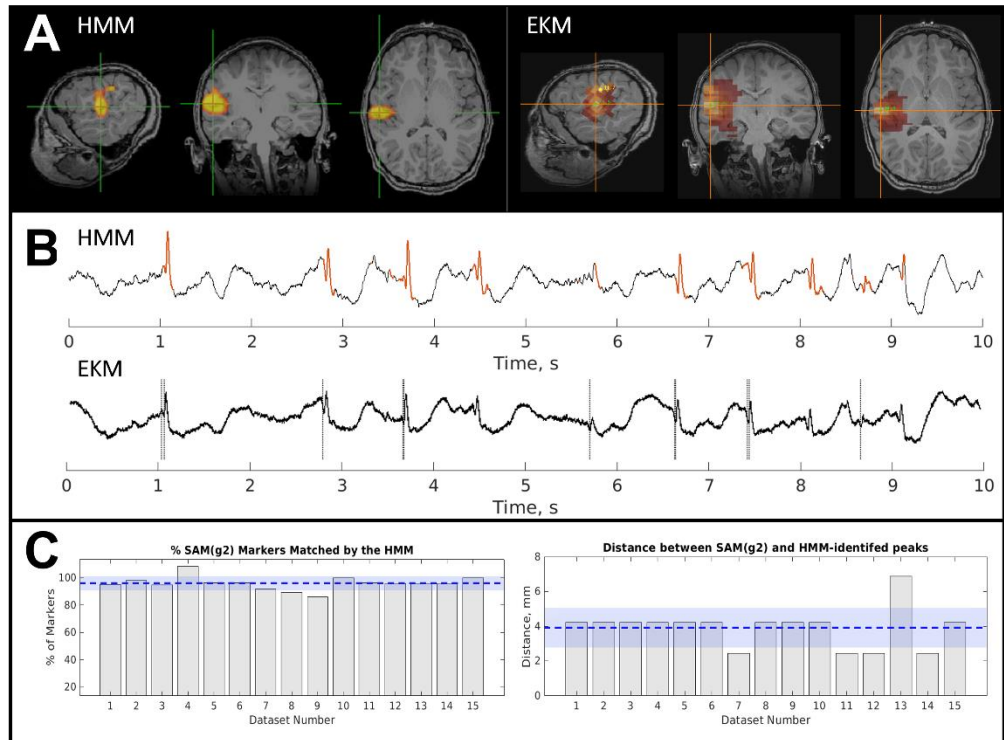


Figure C3: Patient 4 - focal, spike and wave: There is good spatial agreement between the two methods for this patient as shown in panel A (qualitatively) and quantitatively in panel C (right). There was on average $3.9 \pm 0.3\text{mm}$ between the peak locations. The time course of activity from the peak locations for both techniques is shown in panel B. HMM state visits are highlighted in red, and the dotted lines are temporal markers from the EKM method. There is also excellent temporal match between the two methods with $96 \pm 5\%$ of kurtosis markers matched by an HMM state visit (C, left). The epilepsy state was active for $10.4 \pm 0.4\%$ of the total time.

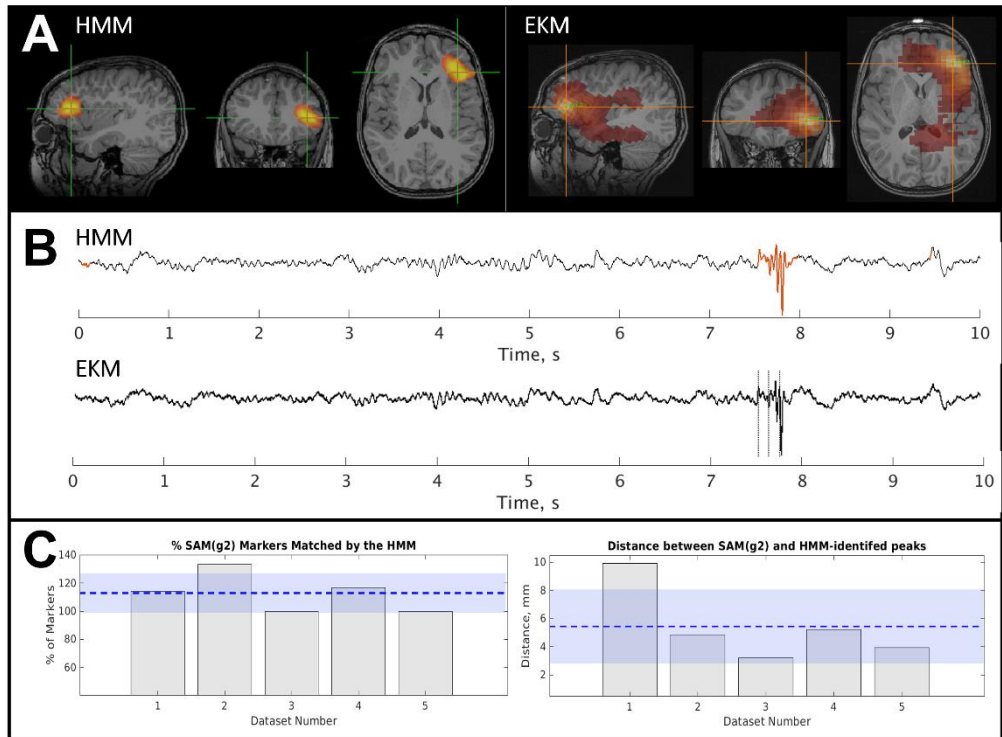


Figure C4: Patient 6 - focal, polymorphic: Interictal activity arising from the left frontal lobe was identified by both the HMM and EKM methods (A). The virtual depth electrode time course extracted from the peak voxel for both methods is shown in B with HMM state visits highlighted in red and kurtosis markers shown as dotted lines. On average over the five runs, $113 \pm 14\%$ of kurtosis markers were matched by an HMM state visit. Note that values greater than 100 are caused by multiple HMM state visits matching with a single kurtosis marker. The distance between the spatial localisations for the two methods was $5 \pm 3\text{mm}$. The HMM state was active for $1.1 \pm 0.4\%$ of the total time.

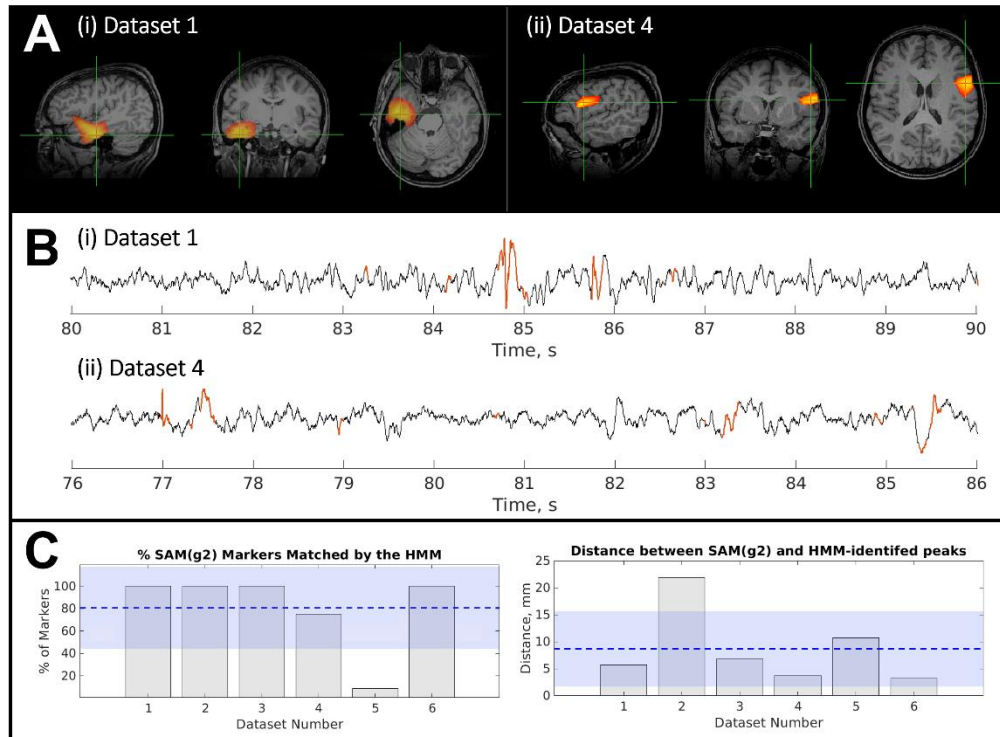


Figure C5: Figure S5: Patient 7 - multifocal, polymorphic: There was a single HMM epilepsy state associated with each dataset (run) which localized to either the right temporal lobe (A) or the left frontal area (B). Time courses of activity from the peak of those spatial maps is shown in C with the temporal lobe exhibiting some sharp activity (Ci) and the frontal areas exhibiting spikes. HMM state visits are highlighted in red. The EKM method identified both the right temporal lobe and left frontal lobe activity in datasets where both were present but the HMM did not. In the two datasets with left frontal activity the HMM failed to identify the temporal lobe activity (so is outperformed by the EKM). This is likely to be because there was a heartbeat-related state which localized to the same location, possibly occluding the epileptiform signal in this case. Using ICA to remove the heartbeat from the data prior to analyses may remedy this in the future. However, for those peak locations that were identified by the HMM, the Euclidian difference between them and the corresponding EKM peak location was $9 \pm 7\text{mm}$ (average and SD over runs). The number of EKM markers matched by an HMM state visit was $81 \pm 37\%$, with the HMM state active for $5 \pm 5\%$ of the total time. So there was still good agreement between methods.

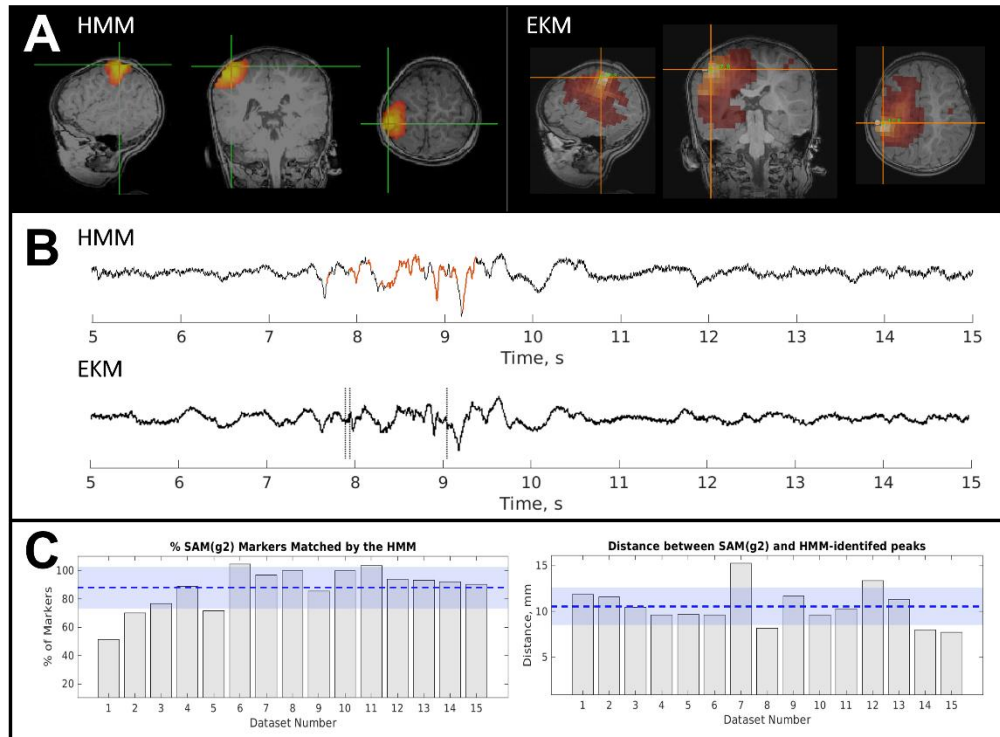


Figure C6: Patient 8 - multifocal, polymorphic: There was good agreement between methods for the spatial localization (see the maps for HMM (A, left) and EKM (A, right)) with an average distance between HMM and EKM peaks of 11 ± 2 mm. The virtual depth electrode time courses for each method are shown in B with HMM state visits highlighted in red and EKM markers as dashed lines. The temporal match was also good with $88 \pm 15\%$ of EKM markers matched by an HMM state visit. The HMM state was active for $7 \pm 3\%$ of the total time.

References

- Adrian, E. D., and B. H. C. Matthews. 1934. 'The Berger Rhythm: Potential Changes from the Occipital Lobes in Man', *Brain*, 57: 355-85.
- Agirre-Arrizubieta, Zaloa, Ngoc J. Thai, Antonio Valentin, Paul L. Furlong, Stefano Seri, Richard P. Selway, Robert D. C. Elwes, and Gonzalo Alarcon. 2014. 'The Value of Magnetoencephalography to Guide Electrode Implantation in Epilepsy', *Brain Topography*, 27: 197-207.
- Azevedo, Frederico A. C., Ludmila R. B. Carvalho, Lea T. Grinberg, José Marcelo Farfel, Renata E. L. Ferretti, Renata E. P. Leite, Wilson Jacob Filho, Roberto Lent, and Suzana Herculano-Houzel. 2009. 'Equal numbers of neuronal and nonneuronal cells make the human brain an isometrically scaled-up primate brain', *The Journal of Comparative Neurology*, 513: 532-41.
- Bagić, Anto I., Robert C. Knowlton, Douglas F. Rose, and John S. Ebersole. 2011. 'American Clinical Magnetoencephalography Society Clinical Practice Guideline 1: Recording and Analysis of Spontaneous Cerebral Activity', *Journal of Clinical Neurophysiology*, 28: 348-54.
- Baker, Adam P., Matthew J. Brookes, lead A. Rezek, Stephen M. Smith, Timothy Behrens, Penny J. Probert Smith, and Mark Woolrich. 2014. 'Fast transient networks in spontaneous human brain activity', *eLife*, 3.
- Bardeen, J., L. N. Cooper, and J. R. Schrieffer. 1957. 'Microscopic Theory of Superconductivity', *Physical Review*, 106: 162-64.
- Barratt, Eleanor L., Prejaas K. Tewarie, Margareta A. Clarke, Emma L. Hall, Penny A. Gowland, Peter G. Morris, Susan T. Francis, Nikos Evangelou, and Matthew J. Brookes. 2017. 'Abnormal task driven neural oscillations in multiple sclerosis: A visuomotor MEG study', *Human Brain Mapping*, 38: 2441-53.
- Bayes, Thomas. 1763. 'LII. An essay towards solving a problem in the doctrine of chances. By the late Rev. Mr. Bayes, F. R. S. communicated by Mr. Price, in a letter to John Canton, A. M. F. R. S', *Royal Society*, 53.
- Beghi, Ettore, Giorgia Giussani, and GBD 2016 Epilepsy Collaborators. 2019. 'Global, regional, and national burden of epilepsy, 1990–2016: a systematic analysis for the Global Burden of Disease Study 2016', *The Lancet Neurology*, 18: 357-75.
- Berger, Hans. 1929. 'Über das Elektrenkephalogramm des Menschen', *Arkiv fur Psychiatrie und Nervenkrankheiten*, 87: 527-70.
- Besl, P. J., and Neil D. McKay. 1992. 'A method for registration of 3-D shapes', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14.
- Biswal, Bharat, F. Zerrin Yetkin, Victor M. Haughton, and S. Hyde James. 1995. 'Functional connectivity in the motor cortex of resting human brain using echo-planar MRI', *Magnetic Resonance in Medicine*.
- Boto, Elena, Ryan M. Hill, Molly Rea, Niall Holmes, Zelekha A. Seedat, James Leggett, V. Shah, J. Osborne, Richard Bowtell, and Matthew J. Brookes. 2021. 'Measuring functional connectivity with wearable MEG', *NeuroImage*, 230.
- Boto, Elena, Niall Holmes, James Leggett, Gillian Roberts, Vishal Shah, Sofie S. Meyer, Leonardo Duque Munoz, Karen J. Mullinger, Tim M. Tierney, Sven Bestmann, Gareth R. Barnes, Richard Bowtell, and Matthew J. Brookes. 2018. 'Moving magnetoencephalography towards real-world applications with a wearable system', *Nature*, 555: 657-61.
- Boto, Elena, Sofie S. Meyer, V. Shah, O. Alem, Svenja Knappe, Peter Kruger, T. Mark Fromhold, Mark Lim, Paul M. Glover, Peter G. Morris, Richard Bowtell, Gareth R. Barnes, and Matthew J. Brookes. 2017. 'A new generation of magnetoencephalography: Room temperature measurements using optically-pumped magnetometers', *NeuroImage*, 149: 404-14.
- Boto, Elena, Zelekha A. Seedat, Niall Holmes, James Leggett, Ryan M. Hill, Gillian Roberts, V. Shah, T. Mark Fromhold, Karen J. Mullinger, Tim M. Tierney, Gareth R. Barnes, Richard Bowtell, and Matthew J. Brookes. 2019. 'Wearable neuroimaging: combining and contrasting magnetoencephalography and electroencephalography', *NeuroImage*, 201.

- Breakspear, Michael, and Leanne M. Williams. 2003. 'A Novel Method for the Topographic Analysis of Neural Activity Reveals Formation and Dissolution of 'Dynamic Cell Assemblies'', *Journal of Computational Neuroscience*, 16: 49-68.
- Brenner, D., J. Lipton, L. Kaufman, and S. J. Williamson. 1978. 'Somatically Evoked Magnetic Fields of the Human Brain', *Science*, 199: 81-83.
- Brookes, M.J., M.J. Groom, L. Liuzzi, R.M. Hill, H.J.F. Smith, P.M. Briley, E.L. Hall, B.A.E. Hunt, L.E. Gascoyne, M.J. Taylor, P.F. Liddle, P.G. Morris, M.W. Woolrich, and E.B. Liddle. 2018. 'Altered temporal stability in dynamic neural networks underlies connectivity changes in neurodevelopment.', *Neuroimage*, 174: 563-75.
- Brookes, M.J., J. Vrba, S.E. Robinson, C.M. Stevenson, A.P. Peters, G.R. Barnes, A. Hillebrand, and P.G. Morris. 2008. 'Optimising experimental design for MEG beamformer imaging', *Neuroimage*, 39: 1788-802.
- Brookes, Matthew J., Joanne R. Hale, Johanna M. Zumer, Claire M. Stevenson, Susan T. Francis, Gareth R. Barnes, Julia P. Owen, Peter G. Morris, and Srikantan S. Nagarajan. 2011. 'Measuring functional connectivity using MEG: Methodology and comparison with fcMRI', *Neuroimage*, 56: 1082-104.
- Brookes, Matthew J., Emma L. Hall, Sian E. Robson, Darren Price, Lena Palaniyappan, Elizabeth B. Liddle, Peter F. Liddle, S. E. Robinson, and Peter G. Morris. 2015. 'Complexity Measures in Magnetoencephalography: Measuring "Disorder" in Schizophrenia', *PLoS One*, 10.
- Brookes, Matthew J., Karen J. Mullinger, Claire M. Stevenson, Peter G. Morris, and Richard Bowtell. 2008. 'Simultaneous EEG source localisation and artifact rejection during concurrent fMRI by means of spatial filtering', *Neuroimage*, 40: 1090-104.
- Brookes, Matthew J., Prejaas K. Tewarie, Benjamin A. E. Hunt, Sian E. Robson, Lauren E. Gascoyne, Elizabeth B. Liddle, Peter F. Liddle, and Peter G. Morris. 2016. 'A multi-layer network approach to MEG connectivity analysis', *Neuroimage*, 132: 425-38.
- Brookes, Matthew J., Mark Woolrich, and Gareth R. Barnes. 2012. 'Measuring functional connectivity in MEG: A multivariate approach insensitive to linear source leakage', *Neuroimage*, 63: 910-20.
- Brookes, Matthew J., Mark Woolrich, Henry Luckhoo, Darren Price, Joanne R. Hale, Mary C. Stephenson, Gareth R. Barnes, Stephen M. Smith, and Peter G. Morris. 2011. 'Investigating the electrophysiological basis of resting state networks using magnetoencephalography', *PNAS*, 108: 16783-88.
- Brovelli, Andrea, Jean-Michel Badier, Francesca Bonini, Fabrice Bartolomei, Olivier Coulon, and Guillaume Auzias. 2017. 'Dynamic reconfiguration of visuomotor-related functional connectivity networks', *Journal of Neuroscience*, 37: 839-53.
- Buard, I., E. Kronberg, S. Steinmetz, S. Hepburn, and D. C. Rojas. 2018. 'Neuromagnetic Beta-Band Oscillations during Motor Imitation in Youth with Autism', *Autism Research and Treatment*, 2018: 9035793.
- Caton, Richard. 1875. 'Electrical currents of the brain', *British Medical Journal*, 2: 278.
- Celozzi, Salvatore, Rodolfo Araneo, and Giampiero Lovat. 2008. *Electromagnetic Shielding, Appendix B: Magnetic Shielding* (John Wiley & Sons).
- Chung, Jae W., Edward Ofori, Gaurav Misra, Christopher W. Hess, and David E. Vaillancourt. 2017. 'Beta-band activity and connectivity in sensorimotor and parietal cortex are important for accurate motor performance', *Neuroimage*, 144: 164-73.
- Churchill, James D., Jason A. Tharp, Cara L. Wellman, Dale R. Sengelaub, and Preston E. Garraghty. 2004. 'Morphological correlates of injury-induced reorganization in primate somatosensory cortex', *BMC Neuroscience*, 5.
- Cohen, David. 1972. 'Magnetoencephalography: Detection of the Brain's Electrical Activity with a Superconducting Magnetometer', *Science*, 175: 664-66.
- Colclough, G. L., Matthew J. Brookes, S. M. Smith, and Mark Woolrich. 2015. 'A symmetric multivariate leakage correction for MEG connectomes', *Neuroimage*, 117: 439-48.
- Colclough, G. L., Mark Woolrich, Prejaas K. Tewarie, Matthew J. Brookes, Andrew J. Quinn, and S. M. Smith. 2016. 'How reliable are MEG resting-state connectivity metrics?', *Neuroimage*, 138: 284-93.
- Creel, Donnell Joseph. 2019. 'Visually evoked potentials.' in Kerry H. Levin and Patrick Chauvel (eds.), *Handbook of Clinical Neurology*.

- Dauwan, M., J. I. Hoff, E. M. Vriens, Arjan Hillebrand, Cornelis J. Stam, and I. E. Sommer. 2019. 'Aberrant resting-state oscillatory brain activity in Parkinson's disease patients with visual hallucinations: An MEG source-space study', *NeuroImage: Clinical*, 22: 101752.
- de Lange, Floris P., Ole Jensen, Markus Bauer, and Ivan Toni. 2008. 'Interactions Between Posterior Gamma and Frontal Alpha/Beta Oscillations During Imagined Actions', *Frontiers in Human Neuroscience*, 2.
- de Pasquale, Francesco, Stefania Della Penna, Abraham Z. Snyder, Christopher Lewis, Dante Mantini, Laura Marzetti, Paolo Belardinelli, Luca Ciancetta, Vittorio Pizzella, Gian Luca Romani, Maurizio Corbetta, and 2010. 'Temporal dynamics of spontaneous MEG activity in brain networks', *PNAS*, 107: 6040-45.
- Ebersole, John S. 1997. 'Magnetoencephalography/Magnetic Source Imaging in the Assessment of Patients with Epilepsy', *Epilepsia*, 38: S1-S5.
- Engel, Andreas K., Christian Gerloff, Claus C. Hilgetag, and Guido Nolte. 2013. 'Intrinsic Coupling Modes: Multiscale Interactions in Ongoing Brain Activity', *Neuron*, 80: 867-86.
- Everson, Richard. 1997. 'Orthogonal, but not Orthonormal, Procrustes Problems', *Advances in Computational Mathematics*.
- Feynman, Richard P., Robert B. Leighton, and Matthew Sands. 1963. *The Feynman lectures on physics* (Reading, Mass. : Addison-Wesley Pub. Co., c1963-1965).
- Fletcher, Daniel A., and Julie A. Theriot. 2004. 'An introduction to cell motility for the physical scientist', *Physical Biology*, 1.
- Fox, Michael D., and Marcus E. Raichle. 2007. 'Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging', *Nature Reviews Neuroscience*, 8: 700-11.
- Fries, Pascal. 2005. 'A mechanism for cognitive dynamics: neuronal communication through neuronal coherence', *Trends in Cognitive Sciences*, 9: 474-80.
- . 2015. 'Rhythms for Cognition: Communication through Coherence', *Neuron*, 88: 220-35.
- Fry, Adam, Karen J. Mullinger, George C. O'Neill, Eleanor L. Barratt, Peter G. Morris, Markus Bauer, Jonathan P. Folland, and Matthew J. Brookes. 2016. 'Modulation of post-movement beta rebound by contraction force and rate of force development', *Human Brain Mapping*, 37: 2493-511.
- Gaetz, William, J. Christopher Edgar, D. J. Wang, and Timothy P. L. Roberts. 2011. 'Relating MEG measured motor cortical oscillations to resting γ -Aminobutyric acid (GABA) concentration', *NeuroImage*, 55: 616-21.
- Gaetz, William, Ronal S. Gordon, Christos Papadelis, Hisako Fujiwara, Douglas F. Rose, J. Christopher Edgar, Erin Simon Schwartz, and Timothy P. L. Roberts. 2015. 'Magnetoencephalography for Clinical Pediatrics: Recent Advances in Hardware, Methods, and Clinical Applications', *Journal of Pediatric Epilepsy*, 4: 139-55.
- Gaetz, William, Edward Rhodes, Luke Bloy, Lisa Blaskey, Carissa R. Jackel, Edward S. Brodtkin, Amy Waldman, David Embick, Stephen Hall, and Timothy P. L. Roberts. 2020. 'Evaluating motor cortical oscillations and age-related change in autism spectrum disorder', *NeuroImage*, 207.
- Gascoyne, Lauren E., Matthew J. Brookes, Mohanbabu Rathnaiah, Mohammad Z. Katshu, Loes Koelewijn, Gemma Williams, Jyothika Kumar, James T. R. Walters, Zelekh A. Seedat, Lena Palaniyappan, J. F. William Deakin, Krish D. Singh, Peter F. Liddle, and Peter G. Morris. 2021. 'Motor-related oscillatory activity in schizophrenia according to phase of illness and clinical symptom severity', *NeuroImage: Clinical*, 29.
- Geselowitz, David B. 1970. 'On the Magnetic Field Generated Outside an Inhomogeneous Volume Conductor by Internal Current Sources', *IEEE Transactions on Magnetics*, MAG-6: 346-47.
- Glasser, Matthew F., Timothy S. Coalson, Emma C. Robinson, Carl D. Hacker, John Harwell, Essa Yacoub, Kamil Ugurbil, Jesper Andersson, Christian F. Beckmann, M. Jenkinson, Stephen M. Smith, and David C. Van Essen. 2016. 'A multi-modal parcellation of human cerebral cortex', *Nature*, 536: 171-78.
- GofshTEyn, J. S., T. Le, S. Kessler, R. Kamens, C. Carr, W. Gaetz, L. Bloy, T. Roberts, E. S. Schwartz, E. D. Marsh, and 155 (2019). Synthetic aperture magnetometry and excess kurtosis

- mapping of Magnetoencephalography (MEG) is predictive of epilepsy surgical outcome in a large pediatric cohort. *Epilepsy research*, 106151. <https://doi.org/10.1016/j.eplepsyres.2019.106151>. 2019. 'Synthetic aperture magnetometry and excess kurtosis mapping of Magnetoencephalography (MEG) is predictive of epilepsy surgical outcome in a large pediatric cohort', *Epilepsy Research*, 155.
- Gross, J., J. Kujala, Matti Hämäläinen, L. Timmermann, A. Schnitzler, and R. Salmelin. 2001. 'Dynamic imaging of coherent sources: Studying neural interactions in the human brain', *PNAS*, 98: 694-99.
- Haft, M., R Hofmann, and V Tresp. 1999. 'Model-independent mean-field theory as a local method for approximate propagation of information', *Computation in neural systems*, 10: 93-105.
- Hall, M., I. A. Nissen, E. van Straaten, P. L. Furlong, C. Witton, E. Foley, S. Seri, and A. Hillebrand. 2018. 'An evaluation of kurtosis beamforming in magnetoencephalography to localize the epileptogenic zone in drug resistant epilepsy patients', *Clinical Neurophysiology*, 129: 1221-29.
- Hämäläinen, Matti, Riitta Hari, Risto J. Ilmoniemi, Jukka Knuutila, and Olli V. Lounasmaa. 1993. 'Magnetoencephalography - theory, instrumentation, and applications to noninvasive studies of the working human brain', *Reviews of Modern Physics*, 65.
- Hämäläinen, Matti, and Risto J. Ilmoniemi. 1994. 'Interpreting magnetic fields of the brain: minimum norm estimates', *Medical & Biological Engineering & Computing*, 32: 35-42.
- Hämäläinen, Matti, and J. Sarvas. 1989. 'Realistic conductivity geometry model of the human head for interpretation of neuromagnetic data', *IEEE Transactions on Biomedical Engineering*, 36: 165-71.
- Hassan, M. R., and B. Nath. 2005. 'Stock market forecasting using hidden Markov model: a new approach', *5th International Conference on Intelligent Systems Design and Applications*: 192-96.
- Heidelberg, Ruth, M. Neal Waxham, John H. Byrne, and James L. Roberts. 2004. *From Molecules to Networks : An Introduction to Cellular and Molecular Neuroscience* (Elsevier Science & Technology).
- Heideman, Simone G., Andrew J. Quinn, Mark Woolrich, F. van Ede, and Anna C. Nobre. 2020. 'Dissecting beta-state changes during timed movement preparation in Parkinson's disease', *Progress in Neurobiology*, 184.
- Higgins, Cameron, Yunzhe Liu, Diego Vidaurre, Zeb Kurth-Nelson, Ray Dolan, Timothy Behrens, and Mark Woolrich. 2021. 'Replay bursts in humans coincide with activation of the default mode and parietal alpha networks', *Neuron*, 109: 882-93.
- Hill, Ryan M., Elena Boto, Niall Holmes, Caroline Hartley, Zelekha A. Seedat, James Leggett, Gillian Roberts, Vishal Shah, Tim M. Tierney, Mark W. Woolrich, Charlotte J. Stagg, Gareth R. Barnes, Richard Bowtell, Rebecca Slater, and Matthew J. Brookes. 2019. 'A tool for functional brain imaging with lifespan compliance', *Nature Communications*, 10.
- Hill, Ryan M., Elena Boto, Molly Rea, Niall Holmes, James Leggett, Laurence A. Coles, Manolis Papastavrou, Sarah Everton, Benjamin A. E. Hunt, Dominic Sims, James Osborne, Vishal Shah, Richard Bowtell, and Matthew J. Brookes. 2020. 'Multi-channel whole-head OPM-MEG: Helmet design and a comparison with a conventional system', *NeuroImage*, 219.
- Hillebrand, Arjan, and Gareth R. Barnes. 2002. 'A Quantitative Assessment of the Sensitivity of Whole-Head MEG to Activity in the Adult Human Cortex', *NeuroImage*, 16.
- Hipp, Joerg F., David J. Hawellek, Maurizio Corbetta, Markus Siegel, and Andreas K. Engel. 2012. 'Large-scale cortical correlation structure of spontaneous oscillatory activity', *Nature Neuroscience*, 15: 884-90.
- Hoburg, J. F. 1995. 'Principles of Quasistatic Magnetic Shielding with Cylindrical and Spherical Shields', *IEEE Transactions on Electromagnetic Compatibility*, 37.
- Holmes, Niall, James Leggett, Elena Boto, Gillian Roberts, Ryan M. Hill, Tim M. Tierney, Vishal Shah, Gareth R. Barnes, Matthew J. Brookes, and Richard Bowtell. 2018. 'A bi-planar coil system for nulling background magnetic fields in scalp mounted magnetoencephalography', *NeuroImage*, 181: 760-74.

- Hoogenboom, Nienke, Jan-Mathijs Schoffelen, Robert Oostenveld, Laura M. Parkes, and Pascal Fries. 2005. 'Localizing human visual gamma-band activity in frequency, time and space', *NeuroImage*, 29: 764-73.
- Hook, J. R., and H. E. Hall. 1974. *Solid State Physics, Second Edition* (John Wiley & Sons).
- Huang, M. X., J. C. Mosher, and R. M. Leahy. 1999. 'A sensor-weighted overlapping-sphere head model and exhaustive head model comparison for MEG', *Physics in Medicine & Biology*, 44: 423-40.
- Hunt, Benjamin A. E., Elizabeth B. Liddle, Lauren E. Gascoyne, Lorenzo Magazzini, Bethany C. Routley, Krish D. Singh, Peter G. Morris, Matthew J. Brookes, and Peter F. Liddle. 2019. 'Attenuated Post-Movement Beta Rebound Associated With Schizotypal Features in Healthy People', *Schizophrenia Bulletin*, 45: 883-91.
- Hunt, Benjamin A. E., Prejaas K. Tewarie, Olivier E. Moughin, Nicolas Geades, Derek K. Jones, Krish D. Singh, Peter G. Morris, Penny A. Gowland, and Matthew J. Brookes. 2016. 'Relationships between cortical myeloarchitecture and electrophysiological networks', *PNAS*, 113: 13510-15.
- Jaklevic, R. C., John Lambe, A. H. Silver, and J. E. Mercereau. 1964. 'Quantum Interference Effects in Josephson Tunneling', *Physical Review Letters*, 12: 159-60.
- Jenkinson, M., P. Bannister, M. Brady, and S. Smith. 2002. 'Improved optimization for the robust and accurate linear registration and motion correction of brain images', *NeuroImage*, 17: 825-41.
- Jenkinson, M., and Michael Chappell. 2018. *Introduction to Neuroimaging Analysis* (Oxford University Press).
- Jenkinson, M., and S. M. Smith. 2001. 'A global optimisation method for robust affine registration of brain images', *Medical Image Analysis*, 5: 143-56.
- Jiles, David C. 1998. *Introduction to Magnetism and Magnetic Materials, Second Edition* (CRC Press).
- Jöbsis, Frans F. 1977. 'Noninvasive, infrared monitoring of cerebral and myocardial oxygen sufficiency and circulatory parameters', *Science*, 198: 1264-67.
- Jones, Stephanie R. 2016. 'When brain rhythms aren't 'rhythmic': implication for their mechanisms and meaning', *Current Opinion in Neurobiology*, 40: 72-80.
- Josephson, B. D. 1962. 'Possible New Effects in Superconductive Tunneling', *Physics Letters*, 1: 251-53.
- Jurkiewicz, Michael T., William C. Gaetz, Andreea C. Bostan, and Douglas Cheyne. 2006. 'Post-movement beta rebound is generated in motor cortex: evidence from neuromagnetic recordings', *NeuroImage*, 32: 1281-9.
- Kim, Hunmin, Chun Kee Chung, and Hee Hwang. 2013. 'Magnetoencephalography in pediatric epilepsy', *Clinical and Experimental Pediatrics*, 56.
- Koelewijn, Thomas, Hein T. van Schie, Harold Bekkering, Robert Oostenveld, and Ole Jensen. 2008. 'Motor-cortical beta oscillations are modulated by correctness of observed action', *NeuroImage*, 40: 767-75.
- Kramer, Mark A., Adriano B. L. Tort, and Nancy J. Kopell. 2008. 'Sharp edge artifacts and spurious coupling in EEG frequency comodulation measures', *Journal of Neuroscience Methods*, 170: 352-57.
- Little, Simon, James Bonaiuto, Gareth R. Barnes, and Sven Bestmann. 2019. 'Human motor cortical beta bursts relate to movement planning and response errors', *PLOS Biology*, 17.
- Liuzzi, Lucrezia, Lauren E. Gascoyne, Prejaas K. Tewarie, Eleanor L. Barratt, Elena Boto, and Matthew J. Brookes. 2017. 'Optimising experimental design for MEG resting state functional connectivity measurement', *NeuroImage*, 155: 565-76.
- London, F. 1948. 'On the Problem of the Molecular Theory of Superconductivity', *Physical Review*, 74: 562-73.
- Lopes da Silva, Fernando. 2013. 'EEG and MEG: Relevance to Neuroscience', *Neuron*, 80: 1112-28.
- Löwdin, Per-Olov. 1950. 'On the Non-Orthogonality Problem Connected with the Use of Atomic Wave Functions in the Theory of Molecules and Crystals', *The Journal of Chemical Physics*, 18.

- Lundqvist, Mikael, Jonas Rose, Pawel Herman, Scott L. Brincat, Timothy J. Buschman, and Earl K. Miller. 2016. 'Gamma and Beta Bursts Underlie Working Memory', *Neuron*, 90: 152-64.
- Mamon, Rogemar S., and Robert J. Elliott. 2014. *Hidden Markov Models in Finance* (Springer).
- Maxwell, James Clerk. 1873. *A Treatise on Electricity and Magnetism* (Clarendon Press).
- Mazziotta, J., A. Toga, A. Evans, P. Fox, J. Lancaster, K. Zilles, R. Woods, T. Paus, G. Simpson, B. Pike, C. Holmes, L. Collins, P. Thompson, D. MacDonald, M. Lacoboni, T. Schormann, K. Amunts, N. Palomero-Gallagher, S. Geyer, L. Parsons, K. Narr, G. Le Goualher, Jordan Feidler, Kenneth Smith, D. Boomsma, Hilleke Hulshoff Pol, T. Cannon, R. Kawashima, and B. Mazoyer. 2001. 'A Four-Dimensional Probabilistic Atlas of the Human Brain', *Journal of the American Medical Informatics Association*, 8: 401-30.
- McCormick, M., R. Ma, and T. P. Coleman. 2010. 'An analytic spatial filter and a hidden Markov model for enhanced information transfer rate in EEG-based brain computer interfaces', *IEEE International Conference on Acoustics, Speech and Signal Processing*: 602-205.
- Meijs, J. W., F. G. Bosch, M. J. Peters, and Fernando Lopes da Silva. 1987. 'On the magnetic field distribution generated by a dipolar current source situated in a realistically shaped compartment model of the head', *Electroencephalography and Clinical Neurophysiology*, 66.
- Meissner, W., and R. Ochsenfeld. 1933. 'A new effect upon the occurrence of superconductivity', *Naturwissenschaften*, 21: 787-88.
- Mohan, Midhun, Simon Keller, Andrew Nicolson, Shubhabrata Biswas, David Smith, Jibril Osman Farah, Paul Eldridge, and Udo Wiesmann. 2018. 'The longterm outcomes of epilepsy surgery', *PLoS ONE*, 13.
- Mormann, Florian, Klaus Lehnertz, Peter David, and Christian E. Elger. 2000. 'Mean phase coherence as a measure for phase synchronization and its application to the EEG of epilepsy patients', *Physica D: Nonlinear Phenomena*, 144: 358-69.
- Mosher, J. C., Sylvain Baillet, and R. M. Leahy. 2003. 'Equivalence of linear approaches in bioelectromagnetic inverse solutions', *IEEE Workshop on Statistical Signal Processing*: 294-97.
- Mullinger, Karen J., M. T. Cherukara, R. B. Buxton, Susan T. Francis, and S. D. Mayhew. 2017. 'Post-stimulus fMRI and EEG responses: evidence for a neuronal origin hypothesised to be inhibitory', *NeuroImage*, 157: 388-99.
- Mullinger, Karen J., Stephen D. Mayhew, Andrew P. Bagshaw, Richard Bowtell, and Susan T. Francis. 2013. 'Poststimulus undershoots in cerebral blood flow and BOLD fMRI responses are modulated by poststimulus neuronal activity', *PNAS*, 110: 13636-41.
- Murakami, H, Z. I. Wang, A. Marashly, B. Krishnan, R. A. Prayson, Y. Kakisaka, J. C. Mosher, J. Bulacio, J. A. Gonzalez-Martinez, W. E. Bingaman, I. M. Najm, R. C. Burgess, and A. V. Alexopoulos. 2016. 'Correlating magnetoencephalography to stereo-electroencephalography in patients undergoing epilepsy surgery', *Brain*, 139: 2935-47.
- Murakami, Shingo, and Yoshio Okada. 2006. 'Contributions of principal neocortical neurons to magnetoencephalography and electroencephalography signals', *The Journal of Physiology*, 575: 925-36.
- Muthukumaraswamy, S. D., J. F. M. Myers, S. J. Wilson, D. J. Nutt, A. Lingford-Hughes, Krish D. Singh, and K. Hamandi. 2013. 'The effects of elevated endogenous GABA levels on movement-related network oscillations', *NeuroImage*, 66: 36-41.
- Nissen, I. A., C. J. Stam, J. Citroen, J. C. Reijneveld, and A. Hillebrand. 2016. 'Preoperative evaluation using magnetoencephalography: Experience in 382 epilepsy patients', *Epilepsy Research*, 124: 23-33.
- Nolte, Guido, Ou Bai, Lewis Wheaton, Zoltan Mari, Sherry Vorbach, and Mark Hallett. 2004. 'Identifying true brain interaction from EEG data using the imaginary part of coherency', *Clinical Neurophysiology*, 115: 2292-307.
- O'Neill, George C., Eleanor L. Barratt, Benjamin A. E. Hunt, Prejaas K. Tewarie, and Matthew J. Brookes. 2015. 'Measuring electrophysiological connectivity by power envelope correlation: a technical review on MEG methods', *Physics in Medicine & Biology*, 60: R271-R95.

- O'Neill, George C., Markus Bauer, Mark Woolrich, Peter G. Morris, Gareth R. Barnes, and Matthew J. Brookes. 2015. 'Dynamic recruitment of resting state sub-networks', *NeuroImage*, 115: 85-95.
- O'Neill, George C., Prejaas K. Tewarie, G. L. Colclough, Lauren E. Gascoyne, Benjamin A. E. Hunt, Peter G. Morris, Mark Woolrich, and Matthew J. Brookes. 2017. 'Measurement of dynamic task related functional networks using MEG', *NeuroImage*, 146: 667-78.
- O'Neill, George C., Prejaas K. Tewarie, Diego Vidaurre, Lucrezia Liuzzi, Mark Woolrich, and Matthew J. Brookes. 2018. 'Dynamics of large-scale electrophysiological networks: A technical review', *NeuroImage*, 180: 559-76.
- Obermaier, B., C. Guger, C. Neuper, and Gert Pfurtscheller. 2001. 'Hidden Markov models for online classification of single trial EEG data', *Pattern Recognition Letters*, 22: 1299-309.
- Ogawa, S., T. M. Lee, A. R. Kay, and D. W. Tank. 1990. 'Brain magnetic resonance imaging with contrast dependent on blood oxygenation', *PNAS*, 87: 9868-72.
- Onnes, H. Kamerlingh. 1911. 'Further experiments with liquid helium d - on the change of the electrical resistance of pure metals at very low temperatures, etc v the disappearance of the resistance of mercury', *Proceedings of the Koninklijke Akademie Van Wetenschappen Te Amsterdam*, 14: 113-15.
- Osipova, Daria, Dora Hermes, and Ole Jensen. 2008. 'Gamma Power Is Phase-Locked to Posterior Alpha Activity', *PLoS One*, 3.
- Pakenham, Daisie O., Andrew J. Quinn, Adam Fry, Susan T. Francis, Mark W. Woolrich, Matthew J. Brookes, and Karen J. Mullinger. 2020. 'Post-stimulus beta responses are modulated by task duration', *NeuroImage*, 206.
- Parvizi, Josef, and Sabine Kastner. 2018. 'Human Intracranial EEG: Promises and Limitations', *Nature Neuroscience*, 21: 474-83.
- Pfurtscheller, Gert, and Fernando Lopes da Silva. 1999. 'Event-related EEG/MEG synchronization and desynchronization: basic principles', *Clinical Neurophysiology*, 110: 1842-57.
- Prichard, James, Douglas Rothman, Edward Novotny, Ognen Petroff, Takeo Kuwabara, Malcom Avison, Alistair Howseman, Christopher Hanstock, and Robert Shulman. 1991. 'Lactate rise detected by 1H-NMR in human visual cortex during physiologic stimulation', *PNAS*, 88: 5829-31.
- Quinn, Andrew J., Diego Vidaurre, Romesh Abeysuriya, Robert Becker, Anna C. Nobre, and Mark W. Woolrich. 2018. 'Task-Evoked Dynamic Network Analysis Through Hidden Markov Modeling', *Frontiers in Neuroscience*, 12: 603.
- Rabiner, Lawrence R. 1989. 'A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition', *Proceedings of the IEEE*, 77: 257-86.
- Rampf, Stefan, Hermann Stefan, Xintong Wu, Martin Kaltenhäuser, Burkhard Maess, Friedhelm C. Schmitt, Carsten H. Wolters, Hajo Hamer, Burkhard S. Kasper, Stefan Schwab, Arndt Doerfler, Ingmar Blümcke, Karl Rössler, and Michael Buchfelder. 2019. 'Magnetoencephalography for epileptic focus localization in a series of 1000 cases', *Brain*, 142: 3059-71.
- Rezek, lead A., and Stephen Roberts. 2005. 'Ensemble Hidden Markov Models with Extended Observation Densities for Biosignal Analysis.' in Dirk Husmeier, Richard Dybowski and Stephen Roberts (eds.), *Probabilistic Modeling in Bioinformatics and Medical Informatics* (Springer London: London).
- Rier, Lukas, Rouzbeh Zamyadi, Jing Zhang, Zahra Emami, Zelekha A. Seedat, Sergiu Mocanu, Lauren E. Gascoyne, Christopher M. Allen, John W. Scadding, Paul L. Furlong, Gerard Gooding-Williams, Mark Woolrich, Nikos Evangelou, Matthew J. Brookes, and Benjamin T. Dunkley. 2021. 'Mild traumatic brain injury impairs the coordination of intrinsic and motor-related neural dynamics', *NeuroImage: Clinical*, 32: 102841.
- Robinson, S. E., S. S. Nagarajan, M. Mantle, V. Gibbons, and H. Kirsch. 2004. 'Localization of interictal spikes using SAM(g2) and dipole fit', *Neurology & clinical neurophysiology*, 74.
- Robinson, S. E., and Jiri Vrba. 1998. 'Functional neuroimaging by synthetic aperture magnetometry', *Recent Advances in Biomagnetism*: 302-05.

- Robson, Sian E., Matthew J. Brookes, Emma L. Hall, Lena Palaniyappan, Jyothika Kumar, Michael Skelton, Nikolaos G. Christodoulou, Ayaz Qureshi, Fiesal Jan, Mohammad Z. Katshu, Elizabeth B. Liddle, Peter F. Liddle, and Peter G. Morris. 2016. 'Abnormal visuomotor processing in schizophrenia', *NeuroImage: Clinical*, 12: 869-78.
- Sarvas, J. 1987. 'Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem', *Physics in Medicine & Biology*, 32: 11-22.
- Schwartz, Erin Simon, J. Christopher Edgar, William C. Gaetz, and Timothy P. L. Roberts. 2010. 'Magnetoencephalography', *Pediatric Radiology*, 40: 50-58.
- Seedat, Zelekha A., Andrew J. Quinn, Diego Vidaurre, Lucrezia Liuzzi, Lauren E. Gascoyne, Benjamin A. E. Hunt, George C. O'Neill, Daisie O. Pakenham, Karen J. Mullinger, Peter G. Morris, Mark Woolrich, and Matthew J. Brookes. 2020. 'The role of transient spectral 'bursts' in functional connectivity: A magnetoencephalography study', *NeuroImage*, 209.
- Sekihara, Kunsuke, Kenneth E. Hild, and Srikantan S. Nagarajan. 2006. 'A Novel Adaptive Beamformer for MEG Source Reconstruction Effective When Large Background Brain Activities Exist', *IEEE Transactions on Biomedical Engineering*, 53: 1755-64.
- Sherman, Maxwell A., Shane Lee, Robert Law, Saskia Haegens, Catherine A. Thorn, Matti S. Hämäläinen, Christopher I. Moore, and Stephanie R. Jones. 2016. 'Neural mechanisms of transient neocortical beta rhythms: Converging evidence from humans, computational modeling, monkeys, and mice', *PNAS*, 113: E4885-E94.
- Shin, Hyeyoung, Robert Law, Shawn Tsutsui, Christopher I. Moore, and Stephanie R. Jones. 2017. 'The rate of transient beta frequency events predicts behavior across tasks and species', *eLife*.
- Smith, Stephen M. 2002. 'Fast robust automated brain extraction', *Human Brain Mapping*, 17: 143-55.
- Stam, Cornelis J., Guido Nolte, and Andreas Daffertshofer. 2007. 'Phas lag index: Assessment of functional connectivity from multi channel EEG and MEG with diminished bias from common sources', *Human Brain Mapping*, 28: 1178-93.
- Stancák Jr., Andrej, and Gert Pfurtscheller. 1995. 'Desynchronization and recovery of β rhythms during brisk and slow self-paced finger movements in man', *Neuroscience Letters*, 196: 21-24.
- Stanley, Jeffrey A., and Naftali Raz. 2018. 'Functional Magnetic Resonance Spectroscopy: The "New" MRS for Cognitive Neuroscience and Psychiatry Research', *Frontiers in Psychiatry*, 9.
- Stefan, Hermann, Xintong Wu, Michael Buchfelder, Stefan Rampp, Burkhard Kasper, Rudiger Hopfengartner, Friedhelm Schmitt, Arnd Dorfler, Ingmar Blumcke, Dong Zhou, and Daniel Weigel. 2011. 'MEG in frontal lobe epilepsies: Localization and postoperative outcome', *Epilepsia*, 52: 2233-38.
- Stenroos, Matti, Alexander Hunold, and Jens Haueisen. 2014. 'Comparison of three-shell and simplified volume conductor models in magnetoencephalography', *NeuroImage*, 94: 337-48.
- Tait, Luke, Ayşegül Özkan, Maciej J. Szul, and Jiaxiang Zhang. 2020. 'Cortical source imaging of resting-state MEG with a high resolution atlas: An evaluation of methods', *bioRxiv*.
- Talairach, J., and P. Tournoux. 1988. *Co-planar stereotaxic atlas of the human brain 3-dimensional proportional system: an approach to cerebral imaging* (G. Thieme).
- Tan, H.-R.M., J Gross, and P J Uhlhaas. 2016. 'MEG sensor and source measures of visually induced gamma-band oscillations are highly reliable', *NeuroImage*, 137: 34-44.
- Tang, Lilly, Mary Mantle, Paul Ferrari, Hagen Schiffbauer, Howard A. Rowley, Nicholas M. Barbaro, Mitchel S. Berger, and Timothy P. L. Roberts. 2003. 'Consistency of interictal and ictal onset localization using magnetoencephalography in patients with partial epilepsy', *Journal of Neurosurgery*, 98: 837-45.
- Ter-Pogossian, Michel M., Michael E. Phelps, Edward J. Hoffman, and Nizar A. Mullani. 1975. 'A Positron-Emission Transaxial Tomograph for Nuclear Imaging (PETT)', *Radiology*, 114.
- Tewarie, Prejaas K., M. G. Bright, Arjan Hillebrand, Sian E. Robson, Lauren E. Gascoyne, Peter G. Morris, J. Meier, P. Van Mieghem, and Matthew J. Brookes. 2016. 'Predicting

- haemodynamic networks using electrophysiology: The role of non-linear and cross-frequency interactions', *NeuroImage*, 130: 273-92.
- Tewarie, Prejaas K., Arjan Hillebrand, E. van Dellen, M. M. Schoonheim, F. Barkhof, C. H. Polman, C. Beaulieu, G. Gong, B. W. van Dijk, and Cornelis J. Stam. 2014. 'Structural degree predicts functional network connectivity: A multimodal resting-state fMRI and MEG study', *NeuroImage*, 97: 296-307.
- Tewarie, Prejaas K., Benjamin A. E. Hunt, George C. O'Neill, Aine Byrne, Kevin Aquino, Markus Bauer, Karen J. Mullinger, Stephen Coombes, and Matthew J. Brookes. 2019. 'Relationships Between Neuronal Oscillatory Amplitude and Dynamic Functional Connectivity', *Cerebral Cortex*, 29: 2668-81.
- Tinkhauser, Gerd, Alek Pogosyan, Simon Little, Martijn Beudel, Damian M. Herz, Huilang Tan, and Peter Brown. 2017. 'The modulatory effect of adaptive deep brain stimulation on beta bursts in Parkinson's disease', *Brain*, 140: 1053-67.
- Tuomisto, T., R. Hari, T. Katila, T. Poutanen, and T. Varpula. 1983. 'Studies of auditory evoked magnetic and electric responses: Modality specificity and modelling', *Il Nuovo Cimento D*, 2: 471-83.
- Tzagarakis, Charidimos, Nuri F. Ince, Arthur C. Leuthold, and Giuseppe Pellizzer. 2010. 'Beta-band activity during motor planning reflects response uncertainty', *Journal of Neuroscience*, 30: 11270-7.
- Tzourio-Mazoyer, N., B. Landeau, D. Papathanassiou, F. Crivello, O. Etard, N. Delcroix, and B. Mazoyer. 2002. 'Automated anatomical labelling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain', *NeuroImage*, 15: 273-89.
- van Ede, F., Andrew J. Quinn, Mark Woolrich, and Anna C. Nobre. 2018. 'Neural Oscillations: Sustained Rhythms or Transient Burst-Events?', *Trends in Neurosciences*, 41: 415-17.
- Van Veen, B. D., and K. M. Buckley. 1988. 'Beamforming: a versatile approach to spatial filtering', *IEEE ASSP Magazine*, 5: 4-24.
- Van Veen, B. D., W. Van Drongelen, M. Yuchtman, and A. Suzuki. 1997. 'Localization of brain electrical activity via linearly constrained minimum variance spatial filtering', *IEEE Transactions on Biomedical Engineering*, 44: 867-80.
- Vidaurre, Diego, Laurence T. Hunt, Andrew J. Quinn, Benjamin A. E. Hunt, Matthew J. Brookes, Anna C. Nobre, and Mark W. Woolrich. 2018. 'Spontaneous cortical activity transiently organises into frequency specific phase-coupling networks', *Nature Communications*, 9.
- Vidaurre, Diego, Andrew J. Quinn, Adam P. Baker, David Dupret, Alvaro Tejero-Cantero, and Mark W. Woolrich. 2016. 'Spectrally resolved fast transient brain states in electrophysiological data', *NeuroImage*, 126: 81-95.
- Vrba, Jiri. 2000. 'Multichannel SQUID Biomagnetic Systems.' in H. Weinstock (ed.), *Applications of Superconductivity*.
- Vrba, Jiri, and S. E. Robinson. 2001. 'Signal Processing in Magnetoencephalography', *Methods*, 25: 249-71.
- Wheless, J. W., L. J. Willmore, J. I. Breier, M. Katakai, J. R. Smith, D. W. King, K. J. Meador, Y. D. Park, D. W. Loring, G. L. Clifton, J. Baumgartner, A. B. Thomas, J. E. C. Constantinou, and A. C. Papanicolaou. 1999. 'A Comparison of Magnetoencephalography, MRI, and V-EEG in Patients Evaluated for Epilepsy Surgery', *Epilepsia*, 40: 931-41.
- Woolrich, Mark W., Adam P. Baker, Henry Luckhoo, Hamid Mohseni, Gareth R. Barnes, Matthew J. Brookes, and lead A. Rezek. 2013. 'Dynamic state allocation for MEG source reconstruction', *NeuroImage*, 77: 77-92.
- Zimmerman, J. E., Paul Thiene, and J. T. Harding. 1970. 'Design and Operation of Stable rf-Biased Superconducting Point-Contact Quantum Devices, and a Note on the Properties of Perfectly Clean Metal Contacts', *Journal of Applied Physics*, 41: 1572-80.