University of Nottingham
UK | CHINA | MALAYSIA

Magnetoencephalography for the Investigation and Diagnosis of Mild Traumatic Brain Injury

Current and next-generation analysis and instrumentation

by

Lukas Rier, M.Sci.

Faculty of Science
University of Nottingham

This thesis is submitted for the degree of
Doctor of Philosophy

School of Physics and Astronomy

June 2022

CONTENTS

ABSTRACT

Mild Traumatic Brain Injury (mTBI), (or concussion), is the most common type of brain injury. Despite this, it often goes undiagnosed and can cause long term disability—most likely caused by the disruption of axonal connections in the brain. Objective methods for diagnosis and prognosis are needed but clinically available neuroimaging modalities rarely show structural abnormalities, even when patients suffer persisting functional deficits. In the past three decades, new powerful techniques to image brain structure and function have shown promise in detecting mTBI related changes. Magnetoencephalography (MEG), which measures electrical brain activity by detecting magnetic fields outside the head generated by neural currents, is particularly sensitive and has therefore gained interest from researchers. Numerous studies are proposing abnormal low-frequency neural oscillations and functional connectivity—the statistical interdependency of signals from separate brain regions—as potential biomarkers for mTBI. However, typically small sample sizes, the lack of replication between groups, the heterogeneity of the cohorts studied, and the lack of longitudinal studies impedes the adoption of MEG as a clinical tool in mTBI management. In particular, little is known about the acute phase of mTBI.

In this thesis, some of these gaps will be addressed by analysing MEG data from individuals with mTBI, using novel as well as conventional methods. The potential future of MEG in mTBI research will also be addressed by testing the capabilities of a wearable MEG system based on optically pumped magnetometers (OPMs).

The thesis contains three main experimental studies. In study 1, we investigated the signal dynamics underlying MEG abnormalities, found in a cohort of subjects scanned within three months of an mTBI, using a Hidden Markov Model (HMM), as growing evidence suggests that neural dynamics are (in part) driven by transient bursting events. Applying the HMM to resting-state data, we show that previously reported findings of diminished intrinsic beta amplitude and connectivity in individuals with mTBI (compared to healthy controls) can be explained by a reduction in the beta-band content of pan-spectral bursts and a loss in the temporal coincidence of bursts respectively. Using machine

learning, we find the functional connections driving group differences and achieve classification accuracies of 98%. In a motor task, mTBI resulted in reduced burst amplitude, altered modulation of burst probability during movement and decreased connectivity in the motor network.

In study 2, we further test our HMM-based method in a cohort of subjects with mTBI and non-head trauma—scanned within two weeks of injury—to ensure specificity of any observed effects to mTBI and replicate our previous finding of reduced connectivity and high classification accuracy, although not the reduction in burst amplitude. Burst statistics were stable over both studies—despite data being acquired at different sites, using different scanners. In the same cohort, we applied a more conventional analysis of delta-band power. Although excess low-frequency power appears to be a promising candidate marker for persistently symptomatic mTBI, insufficient data exist to confirm this pattern in acute mTBI. We found abnormally high delta power to be a sensitive measure for discriminating mTBI subjects from healthy controls, however, similarly elevated delta amplitude was found in the cohort with non-head trauma, suggesting that excess delta may not be specific to mTBI, at least in the acute stage of injury.

Our work highlights the need for longitudinal assessment of mTBI. In addition, there appears to be a need to investigate naturalistic paradigms which can be tailored to induce activity in symptom-relevant brain networks and consequently are likely to be more sensitive biomarkers than the resting state scans used to date. Wearable OPM-MEG makes naturalistic scanning possible and may offer a cheaper and more accessible alternative to cryogenic MEG, however, before deploying OPMs clinically, or in pitch-side assessment for athletes, for example, the reliability of OPM-derived measures needs to be verified. In the third and final study, we performed a repeatability study using a novel motor task, estimating a series of common MEG measures and quantifying the reliability of both activity and connectivity derived from OPM-MEG data. These initial findings—presently limited to a small sample of healthy controls—demonstrate the utility of OPM-MEG and pave the way for this technology to be deployed on patients with mTBI.

ACKNOWLEDGEMENTS

It is no exaggeration that the work presented in this thesis would not have been possible without the support from a great number of people. First and foremost, I want to thank my supervisors Prof. Matthew Brookes and Dr Nikos Evangelou for the opportunity to carry out this project and for the guidance, inspiration and confidence they gave me in these past few years.

Secondly, I would like to thank the members, past and present, of the MEG group and the SPMIC in general for providing an environment of comradery which made working here a joy. In particular, I would like to thank Dr Lauren Gascoyne who was an integral part of Team mTBI and an excellent source of advice and laughs, Dr George O'Neill for being a fountain of MEG knowledge, and Dr Zelekha Seedat who's work on HMM's is a great inspiration to me. I want to thank Team OPM, especially Dr Elena Boto, Dr Ryan Hill, and Dr Niall Holmes for their support whenever I ventured into the world of "hot" MEG.

Additionally, I wish to thank Dr Chris Allen, for his efforts in all things paperwork and recruitment, for being a great source of motivation, and for his work on the systematic review that informed so many of our decisions. Here, I would also like to thank the DREEAM Team who made data collection possible while facing a tumultuous couple of years in the NHS.

My gratitude also goes out to Dr Ben Dunkley and his group in Toronto, who made this work possible with their generous sharing of data and expertise, and to Dr John Scadding for his enthusiasm and asking the most enlightening, awkward questions.

Finally, I would like to thank Aisling Macdonald, whose love and support is keeping me (mostly) sane, Mary and Angus, who shared their home (and dog) with me during a global pandemic, Maeve Macdonald, for the craic and, last but certainly not least, my parents Karin and Helmut, Oma Paula, and Renate for their unwavering support even though I am a thousand miles away: Ohne Enk wars nia gongen!

I would like to further express my gratitude to Dr Karen Mullinger and Dr Caroline Witton for their thorough reading of this document and the constructive, enjoyable conversation during the examination of this work.

Chapter 1 INTRODUCTION

Awareness of the vulnerability of the human brain has a long history: perhaps the most ancient medical records in existence today—known as the *Edwin Smith Surgical Papyrus*—contain case reports of injuries to the head and face (Breasted, 1930) and are dated around 1600 BCE.

In the UK today, a person with a head injury is admitted to hospital every three minutes (Headway, 2016)[1]. While the debilitating symptoms of severe injury can be more reliably ascribed to macroscopic damage to the brain, for milder cases of injury, the mechanisms underlying any long-term sequelae are less well understood. This provides the main motivation for this thesis, which is aimed to explore mild Traumatic Brain Injury (mTBI) via the measurement of the magnetic signatures of brain activity (magnetoencephalography).

MTBI is recognised as a 'silent epidemic' and a serious public health concern (National Center for Injury Prevention and Control, 2003). A minority, but substantial number, of patients—given the high incidence of mTBI—suffer from persistent post-concussive symptoms and can experience chronic, lifelong deficits, for example, impaired attention, altered balance, slowness in thinking etc., that result in reduced quality of life (Raskin et al., 2014). The insidious long term health effects of mTBI are of great concern in the world of contact sports, such as rugby and American Football, where increasing evidence suggests repeated mTBI and sub-concussive blows can precipitate cognitive decline and even neurodegenerative disease, such as early-onset dementia.

In this introductory chapter, a primer on the definition, diagnosis, and mechanism of injury of mTBI will be provided. A range of neuroimaging methods used to study mTBI will be described, with a focus on the current state

[1]www.headway.org.uk/about-brain-injury/further-information/statistics/ (accessed 1st June 2022)

of the literature on Magnetoencephalography (MEG) and mTBI. Finally, the overarching aims of this thesis will be outlined alongside a brief overview of the chapters to come.

1.1 MILD TRAUMATIC BRAIN INJURY

The sophisticated reports of head injuries in Egyptian papyri in the third millennium BCE point to a long history of human interest in brain injury. Brain injuries that occur *without* fracturing of the skull were first recognised and described as *concussions* by Arabic physician Rhazes around 900 CE (Levin et al., 2014). Despite this long history of the study of brain injury, the mild end of the severity spectrum is yet to be uniquely defined. However, the most commonly used definitions of mTBI overlap considerably, with the score on the Glasgow Coma Scale (GCS, see Table 1.1), the duration (if applicable) of any loss of consciousness (LOC) and post-traumatic amnesia (PTA) as the common features used for diagnosis. For example, a common definition of mTBI would be loss of consciousness for no more than 30 minutes, GCS of no less than 13, and post-traumatic amnesia for no more than 24 hours. Additionally, some definitions include the absence of abnormalities on MRI or CT scanning.

RESPONSE	SCORE CRITERIA
Ocular	1 No eye opening
	2 Eye opening to pain
	3 Eye opening to sound
	4 Eyes open spontaneously
Verbal	1 No verbal response
	2 Incomprehensible sounds
	3 Inappropriate words
	4 Confused
	5 Orientated
Motor	1 No movements
	2 Abnormal extension to pain
	3 Abnormal flexion to pain
	4 Withdrawal from pain
	5 Localising pain
	6 Obeys commands/ moves on purpose

Table 1.1 Score parameters and criteria for the Glasgow Coma Scale. To obtain the GCS score, the sum of the highest applicable score for each response parameter is recorded. This gives a score from 3 to 15. Commonly, a score of 13-15 is classed as "Mild", 9-13 as "Moderate" and 3-8 as "Severe" TBI.

While most individuals with a **single** mTBI are said to recover within approximately 3 months of injury, several studies reported that up to 50% of individuals continue to suffer the sequelae of mTBI (McInnes et al., 2017; Nelson et al., 2019; Wilson et al., 2021) even in the absence of abnormalities in clinical imaging.

The true number of people globally suffering an mTBI is estimated at 42 million per year (Gardner and Yaffe, 2015). Although less than half of these cases likely result in hospital visits (Cassidy et al., 2004), mTBI nonetheless poses a significant burden on health services (Bazarian et al., 2005; James et al., 2019).

One of the difficulties with estimating the true rate of mTBI may be the disparity in definitions used across research studies (Carroll et al., 2004). Furthermore, concerns about high false-negative rates in mTBI diagnoses have been raised (Powell et al., 2008). Despite the high incidence, relatively little is known about how neural mechanisms underlie the sequelae experienced by mTBI patients. However, recent years have seen increasing interest. In particular, recent media attention on athletes with Chronic Traumatic Encephalopathy (CTE; Omalu et al., 2006, 2005)—likely caused by repeated exposure to concussive (or sub-concussive) blows to the head—and an increase in military personnel suffering consequences of blast-induced mTBI, have raised public awareness about the long-term consequences of mTBI and increased research interest in those populations. For this reason, there has been renewed interest in developing methods to understand, diagnose and manage mTBI.

1.1.1 Injury Mechanism

MTBI is caused by events in which a mechanical force is exerted on the brain. This could be an impact causing rapid deceleration of the moving or rotating head (from a fall for example), or acceleration of the head (e.g. from a punch). The resulting rotational, linear, and tensile forces acting on the brain rapidly deform the tissue which is thought to be the primary cause of injury. The elasticity and structure of the brain mean that a single impact event results in 'shaking', exposing the tissue to a series of acceleration/deceleration events until the motion is complete (Bigler and Maxwell, 2012). Forces which cause rotations about the brain's centre of mass put particular strain on central brain structures, e.g. the corpus callosum, (Shenton et al., 2012) which contain axons—long nerve fibres in the white matter which form the important connections between brain regions (or hemispheres in the case of the corpus callosum). In severe TBI, abrupt shearing of axons can occur, leading to irreversible cell death. This is not believed to be the case in mTBI. Instead, the initial deformation of axons triggers a neurometabolic cascade which may cause axonal damage (Giza and Hovda,

2014)—the likely cause of various long-term symptoms experienced by a substantial subset of mTBI patients.

1.1.2 Symptoms and long-term effects of mTBI

MTBI can cause a variety of symptoms. These include—but are not limited to—cognitive deficit (e.g. attention, concentration), headaches, balance problems, and emotional changes such as increased irritability, anxiety, and depression (Belanger et al., 2005; Raskin et al., 2014; Vanderploeg et al., 2015). If such symptoms are caused by mTBI, or significantly worsen after a mild brain injury, and do not resolve within the expected time of around one month, they are often referred to as Post-concussion Syndrome (PCS). However, several definitions for PCS exist and do not always align with the definitions of mTBI (Raskin et al., 2014). In most cases, mTBI is limited to a single event. However, in some cases, multiple injuries are common (e.g. in contact sports). Recent evidence suggests that repeated mTBI and sub-concussive blows can indeed precipitate cognitive decline and even neurodegenerative disease, such as early-onset dementia (Costanza et al., 2011; Hume et al., 2017; Pearce et al., 2018).

1.2 NEUROIMAGING METHODS

As we have seen above, the clinical means of diagnosing mTBI rely on coarse measures—such as the GCS—as well as information that may only be available from witness accounts—such as LOC and PTA durations. Objective biomarkers which reliably diagnose mTBI, or predict the outcomes of recovery, are lacking. However, a plethora of imaging modalities have been developed in the past few decades that are increasingly used to investigate mTBI (and other neurological diseases) and may offer insights into the mechanisms behind the subtle injuries suffered by the brain. What follows below is a brief overview of some of the most prominent structural and functional imaging techniques and their relevance to mTBI.

1.2.1 Structure

Computed Tomography (CT) (Hounsfield, 1973) uses the variability of X-ray absorption by different tissues to construct 3D images. This is done by exposing the scanned subject to a narrow beam of X-rays from multiple angles. CT is the primary clinical tool used in assessing mTBI as it is highly sensitive to skull fractures or other pathologies that may require medical intervention. However, abnormalities are not common (indeed, some definitions of mTBI specifically require a normal clinical CT scan). In the case of "uncomplicated" mTBI—defined as mTBI without imaging findings—a lack of CT findings does not conclusively predict a lack of symptoms post-injury (Lange et al., 2009). Additionally, the necessary—albeit small—dose of ionising radiation imparted by CT scans may make alternative imaging methods more desirable.

Magnetic Resonance Imaging (MRI) (Lauterbur, 1973; Mansfield and Grannell, 1973) exploits the magnetic properties of Hydrogen nuclei which are abundant in the water contained in brain tissues. The magnetic moments of the protons exposed to the strong, uniform magnetic field of an MRI scanner become aligned, inducing a bulk magnetisation. This is perturbed using an electromagnetic pulse at Radio Frequency (RF) leading to a decaying precession of the nuclear magnetisation. In returning to their net alignment (a process called relaxation), the protons emit RF energy which is measured and used to generate intensity maps showing the distribution of water in the head. Contrasts between tissues are the consequence of their relaxation times, parameters that determine the time scales at which the perturbed nuclei return to the net magnetisation state. The longitudinal relaxation time T_1 determines the time taken for the magnetisation to return to equilibrium, the transverse relaxation time T_2 determines the rate at which the precessing nuclei lose coherence in the plane transverse to the main field. Altering the sequences of perturbation pulses and times between measurements can be used to generate a multitude of different contrasts which can be optimised to suit specific diagnostic needs. **T1 weighted**, **T2 weighted** images as well as **Fluid Attenuated Inversion**

Recovery (FLAIR; similar to T2 but with nulled signal in the cerebrospinal fluid or CSF) are common in clinical contexts. MRI is more often used in moderate and severe cases of TBI and is useful in detecting contusions or swellings for example, however—similar to CT—clinical MRI shows limited findings in mild cases.

Even though clinical imaging does not reveal abnormalities in the majority of cases, microscopic injuries—which may underlie the reported long-term symptoms—cannot be ruled out. More advanced MRI sequences that may be more sensitive to these "invisible" changes have piqued the interest of researchers:

Diffusion Tensor Imaging (DTI) (Basser et al., 1994) is an MRI technique sensitive to the diffusion of water molecules and can be used to infer how constrained this diffusion is in a given direction. In a neuronal fibre, for example, diffusion is most constrained in the directions perpendicular to the fibre. This means DTI can provide information about the structural integrity of brain tissue and may be sensitive to the types of diffuse axonal injury predicted in mTBI. DTI has, therefore, shown some promise in the study of mTBI, however, methodological heterogeneity and contradictory findings mean that further research using this modality is needed (Asken et al., 2018; Shenton et al., 2012).

Susceptibility Weighted Imaging (SWI) (Haacke et al., 2004; Wharton et al., 2010) is sensitive to field inhomogeneity caused by variation in the magnetic properties of tissues. It provides excellent contrast between brain tissue and blood products—especially at high field strengths—which may stem from microscopic haemorrhages. These are present in some cases of mTBI but not visible on standard clinical imaging (CT or T1/T2 weighted MRI), usually conducted at field strengths up to 3T (Park et al., 2009; Wang et al., 2014).

1.2.2 Function

Even with advanced techniques like DTI and SWI, there is a lack of detectable macroscopic damage in mTBI. Nevertheless, functional impairments are common and so, functional imaging techniques have played an important part in recent advances in mTBI research. Any cognitive impairment following mTBI should—in principle—be the consequence of a change in the patient's neural activity. The goal of functional imaging is to infer neural activity, either directly—by measuring the electrical currents flowing in the brain—or via metabolic processes that support neural function:

Positron Emission Tomography (PET) (Raichle et al., 1983) generates functional images by tracking the location of tracer molecules tagged with radioactive isotopes administered to a patient. The unstable nuclei of Oxygen-15 or Fluorine-18 for example, decay by positron emission. When these positrons collide with an electron a short distance away, the annihilation of both particles emits two gamma rays in opposite directions. Using a ring of gamma detectors around the subject, the location of annihilation can be estimated. The anatomical origin of the signal can be determined by overlaying the generated maps onto structural scans (most commonly CT). Function can be inferred by measuring blood flow in the brain via Oxygen-15, but other isotopes can be used to probe various metabolic processes. PET has had some success in detecting abnormal glucose metabolism in TBI and may be sensitive to injury severity and progress of recovery (for a review see Byrnes et al., 2014). Some evidence for compensatory mechanisms post-mTBI has been found using PET (Chen et al., 2003), where resting metabolism did not differ in mTBI patients, while cerebral blood flow was abnormal under a task condition.

Functional Magnetic Resonance Imaging (fMRI) (Ogawa et al., 1990) provides a measure of function by exploiting the differing magnetic properties of oxygenated and deoxygenated haemoglobin. Increased neural activity raises metabolic demand which is associated with changes in blood flow, volume and

oxygenation in the active region—the hemodynamic response. Approaches based on the oxygenation of blood are also referred to as Blood Oxygen Level Dependent (BOLD) fMRI. This technique can be used to determine active regions during tasks as well as provide measures of functional connectivity by evaluating statistical relationships between signals from distal brain regions as a proxy for how well these regions are connected. Although fMRI provides a measure of brain activity with excellent spatial resolution, the BOLD response occurs on time scales of seconds, therefore, temporal resolution is low.

FMRI is currently not used in clinical assessments of mTBI, however, this modality may be sensitive to mTBI and further research is warranted (see McDonald et al. (2012) for a review). For example, the temporal dynamics in the BOLD signal have been shown to vary longitudinally with recovery from mTBI (Churchill et al., 2020), and fMRI connectivity may be related to symptom severity (Churchill et al., 2018) and neuropsychological outcomes after injury (Shi et al., 2021). fMRI has also been shown to yield classification rates (>80%) when trying to distinguish mTBI from healthy controls (Vergara et al., 2018, 2017).

Cerebral Blood Flow (CBF) and Cerebrovascular Reactivity (CVR) are two measures of vascular function which can be found via MR imaging. CBF, the perfusion of blood within brain tissue, commonly measured in mL per $100g$ of tissue per minute, can be quantitatively estimated without the use of contrast agents using Arterial Spin Labelling (ASL) (Williams et al., 1992). This subtraction technique relies on the contrast between pairs of acquisitions, one of which involves the excitation of water protons in the blood on their way to the brain. Tracking the magnetisation of labelled blood then allows for the estimation of perfusion. CVR refers to the ability of the cerebral vasculature to respond to challenges such as changes in CO_2 levels in the body by constriction or dilation of vessels. CVR can be measured by temporarily exposing subjects to elevated levels of CO_2 and estimating the ratio between the change in the

chosen measure of vascular activity (perfusion using ASL or BOLD activity) and the change in CO_2 concentration .

The roles of CBF and CVR abnormalities as biomarkers for mTBI are not yet clearly established, however, associations between symptoms and cerebrovascular abnormalities have been reported (e.g. (Churchill et al., 2017)) suggesting that perfusion and CVR measures could be crucial in furthering our understanding of recovery from mTBI (see (Andre, 2015; Ellis et al., 2016; Lunkova et al., 2021; Wang et al., 2020) for reviews). Further research is needed to determine whether vascular MRI should be included as part of a diagnostic battery of tests for mTBI. Given that these techniques are generally well tolerated, non-invasive and sources of complementary information to other functional imaging techniques, their further exploration is warranted.

Functional Near-Infrared Spectroscopy (fNIRS) (Chance et al., 1993; Hoshi and Tamura, 1993; Kato et al., 1993) is another technique to measure brain function via haemodynamic changes. In this case, the absorption properties of haemoglobin are exploited. fNIRS consists of sources of near-infrared light which penetrates the scalp and skull, allowing it to scatter inside the brain. Photodetectors on the scalp yield a measurement of the absorption spectra that can be used to measure changes in blood oxygenation. The use of fNIRS in mTBI research has so far been rare, although some groups have reported abnormal functional connectivity in mTBI populations (Hocke et al., 2018; Sharma et al., 2020; Urban et al., 2015). Low spatial and temporal resolutions remain the primary limitations of this technique, however, fNIRS is wearable, portable and cheaper than the techniques mentioned above and hence could potentially be deployed for pitch-side assessments in sports if a biomarker of sufficient sensitivity and specificity could be found.

Magnetoencephalography (MEG) (Cohen, 1972, 1968) measures the magnetic signatures of neuroelectrical activity using highly sensitive magnetic field detectors. The source of the measured magnetic fields are electrical

currents generated by synchronously firing neurons in the brain. This makes MEG a direct way of assessing neural function. Magnetic fields are not perturbed by the conductivity profile of the head, pass through the skull unhindered and can therefore be used to reconstruct the activity in the brain with high spatial and excellent temporal resolution. Conventionally, the magnetic signals are measured using cryogenically cooled Superconducting Quantum Interference Devices (SQUIDS). Similar to MRI and PET, the rigid nature of the instrumentation makes MEG sensitive to movement artefacts. However, more recent developments of Optically Pumped Magnetometers (OPMs) have enabled wearable MEG which also has the significant advantages of higher sensitivity and spatial resolution, as well as lifespan compliance. The current state of the MEG literature was expounded in a recent systematic review by Allen et al. (2021) and will be summarised in Section 1.3. Briefly, recent work has suggested that MEG offers ~90% classification accuracy in identifying patients with mTBI making it arguably the most promising modality for investigation of mTBI.

Electroencephalography (EEG) (Berger, 1929) measures electric potentials on the surface of the head—caused by the same neuroelectric activity responsible for the MEG signal—by using an array of electrodes which make electrical contact with the scalp. EEG has an excellent temporal resolution, however, the low conductivity of the skull causes smearing in the field patterns produced, limiting spatial resolution considerably. One advantage of EEG, similar to fNIRS, is its portability. Although EEG is commonly used in hospitals, the value of EEG in the clinical assessment of mTBI is unclear. EEG studies from the second half of the 20th century frequently report slowing of neural oscillations and epileptiform activity. However, methodological concerns, like the requirement for visual interpretation of the EEG signal and the retrospective nature of much of the available research, cast doubt on the usefulness of EEG (Nuwer et al., 2005). Quantitative EEG, which involves data analysis techniques (such as frequency decomposition, coherence and statistical analyses based on

normative data) may alleviate issues of inter-rater-reliability, however, these techniques are less common in clinical contexts, require additional expertise and further research is needed to determine their utility for mTBI assessment (Amyot et al., 2015; Arciniegas, 2011; Conley et al., 2018; Gaetz and Bernstein, 2001; Haneef et al., 2013; Nuwer et al., 2005).

1.3 EXISTING MTBI RESEARCH USING MEG

Imaging modalities based on electrophysiology, such as EEG and MEG give us the most direct, non-invasive means to probe neural function. This makes them potentially useful tools to detect the subtle changes in brain function caused by mTBI. The combination of spatial and temporal sensitivity of MEG is particularly promising and a wide range of research has been carried out to develop biomarkers using this modality. Two common abnormalities differentiate patients from controls: abnormal neural oscillations—with a particular focus on enhanced slow waves or cortical slowing (e.g. Huang et al., 2014; Lewine et al., 1999; Proskovec et al., 2020)—and abnormal functional connectivity (e.g. Antonakakis et al., 2017, 2016; Pang et al., 2016; Zhang et al., 2020). Excess slow-wave activity is a non-specific marker which has generally been associated with poor brain health (e.g. in dementia). Functional connectivity is perhaps a more logical candidate biomarker since altered communication between regions could be readily explained by diffuse white matter damage. These findings were collated in the (at the time of writing) only systematic review on mTBI and MEG (Allen et al., 2021). While its outlook is generally positive, the review highlights some general features of the populations studied:

> *"Twenty-five papers* [out of 37] *examined a civilian population with mixed mechanisms of injury, in five papers the population recruited from was unclear. Five papers examined a military population with two of these specifically focussed on blast injury. Two papers include both military personnel and civilians. Ten of the papers*

recruited only patients with mTBI and persisting PCS. The study sizes ranged from six to 84 participants with mTBI. Mean time between injury and MEG assessment ranged from six days to 13 years but was unreported in nine papers. There was a male bias in the mTBI population of all included papers, with 17 reporting exclusively male participants. The mean mTBI sample age ranged from 25 to 42 years. Year of publication spanned 1999 to 2020."

Here, the heterogeneity in time post-injury of the scans, the bias towards entirely male cohorts, and the inclusion of exclusively symptomatic subjects are noteworthy. In particular, the acute stage of injury has been explored to a lesser degree compared to chronic mTBI populations. In addition, the majority of control groups only consisted of healthy individuals. This is a methodological shortcoming of much of mTBI research in general (Satz et al., 1999). The inclusion of control groups exposed to trauma (excluding the brain) is important if more specific symptom profiles and biomarkers of mTBI are to be found. In addition to cohorts, the reported data analysis techniques varied widely across articles (with little to no replication between groups). Thus, while MEG is promising for the assessment of mTBI, the lack of cohesion across findings means it is not yet ready to be used as a clinical tool.

1.4 THESIS AIMS

This thesis has, broadly, three aims:

I. Based on recent findings of the role of transient bursts in functional connectivity and the repeated findings of abnormal functional connectivity in mTBI we aimed to explore the application of Hidden Markov Models to MEG data from mTBI subjects in the subacute stage of injury (several weeks to months post mTBI).

II. To address the lack of research in the acute stage of injury, we aimed to replicate previous MEG findings in a cohort of subjects in the acute stage of injury and compare this mTBI cohort with healthy as well as orthopaedic trauma controls who had experienced a non-head injury.

III. Finally, we aimed to further the field of wearable MEG by testing the feasibility of using this new technology for future longitudinal assessment of mTBI via a repeatability study.

1.5 THESIS OUTLINE

Chapter 2 gives an introduction to brain anatomy and describes the origin of the electrical activity in the brain. The types of macroscopic brain activity detectable by MEG are described and an overview of the means of measuring those signals is given. This includes the operating principles and instrumentation of conventional SQUID-based MEG systems as well as those using optically pumped magnetometers (OPMs).

Chapter 3 contains the theoretical principles of the techniques used to infer activity inside the brain from the acquired MEG signals, as well as secondary methods used to analyse MEG data. The MEG forward and inverse problems, the estimation of functional connectivity, and the application of the Hidden Markov Model (HMM) to MEG data are covered.

Chapter 4 describes the application of HMMs to MEG data collected at the Hospital for Sick Children, Toronto, Canada, in the subacute phase of mTBI (a few weeks to several months). Analysing data collected at rest and during a task and, utilising the wealth of information provided by the HMM, we gain insight into abnormal neural dynamics found in subjects recovering from mTBI. We apply machine learning techniques to evaluate the diagnostic utility of features found using the HMM.

Chapter 5 outlines the MEGAbIT study, an observational case-control study conducted at the University of Nottingham, which investigated mTBI in the rarely studied acute phase of injury (within 14 days). We present two analyses of resting-state data. Firstly, we apply the HMM-based methodology developed in Chapter 4, secondly, we calculate the oft-cited low-frequency power, to compare abnormalities found in the acute phase with those in the more commonly reported chronic phase. In addition to contrasting mTBI subjects with healthy controls, we also include a cohort of orthopaedic trauma controls to explore whether any abnormalities could be ascribed to a general trauma effect.

Chapter 6 contains a proof-of-principle study exploring the potential of an OPM-based MEG system for future longitudinal studies of mTBI. We introduce a novel task paradigm designed to challenge subjects in domains such as visuomotor control and attention and assess the intrasubject repeatability of brain measures of neural activation and functional connectivity.

Chapter 7 will summarise our findings and offer concluding remarks.

1.6 COVID-19 IMPACT STATEMENT

Before the COVID-19 pandemic, the MEGAbIT study (presented in Chapter 5) was intended to constitute the bulk of this thesis, with planned analyses of task as well as resting-state MEG data, and the presentation of results from structural

MR imaging at 7T (DTI and SWI). Pandemic-related restrictions, however, significantly hindered subject recruitment leading to lower sample sizes than specified during the planning stage in the study protocol. The lack of an adequately sized control group for the task MEG and MRI data, where legacy datasets from healthy controls were not available, led to the exclusion of these data from this thesis. Analysis of the resting-state MEG data presented in Chapter 4—which were provided by external collaborators—motivated further exploration of our Hidden Markov Model (HMM) approach (see 3.4) resulting in the work presented in 5.3, in addition to the more "classical" analysis of low-frequency oscillations (see 5.4) which was planned prior to data-collection.

The lack of access to the OPM-MEG facilities during the early stages of the pandemic impeded further development of the wearable MEG system, preventing any collection of OPM-MEG data from volunteers participating in our mTBI study. Instead, the work presented in Chapter 6 was conducted using healthy volunteers, with a view to testing the feasibility of using OPMs in future longitudinal studies of mTBI.

Chapter 2 MAGNETOENCEPHALOGRAPHY

Since the advent of scientific investigation of the brain, students attributed distinct functions to separate parts. While this was often subject to much speculation, the compartmentalisation and assignment of functions to specific structures in the brain is still the essence of modern neuroscience. This chapter will outline the macro and microscopic makeup of the brain and describe the basis of electrophysiological function and the aspects relevant to the MEG signal, before giving an overview of the types of activity which can be measured using MEG. The second half of this chapter will contain a description of the theory and instrumentation which underlies the acquisition of MEG data, including an overview of methods that are used to reduce various sources of interference in MEG recordings.

2.1 THE ORIGIN OF THE MEG SIGNAL

2.1.1 The macroscopic structure of the brain

When describing the anatomy of the human brain, the coarsest subdivision one can make is perhaps the distinction between grey and white matter. Grey matter mostly includes the cortex—the corrugated outermost surface which is thought to be responsible for the bulk of high-level processing—while the white matter mainly consists of the fibres connecting cortical regions. The longitudinal fissure divides the brain into two hemispheres which are approximately symmetric in structure (but not always in function), and connected via the white matter fibres in the corpus callosum. As shown in Figure 2.1 each hemisphere is further subdivided into four major lobes (frontal, temporal, parietal and occipital), which are named after the corresponding cranial bones. Each lobe of cortex contains a multitude of grooves—the sulci—surrounding the gyri—the cortical folds—which in turn can contain multiple functionally distinct regions.

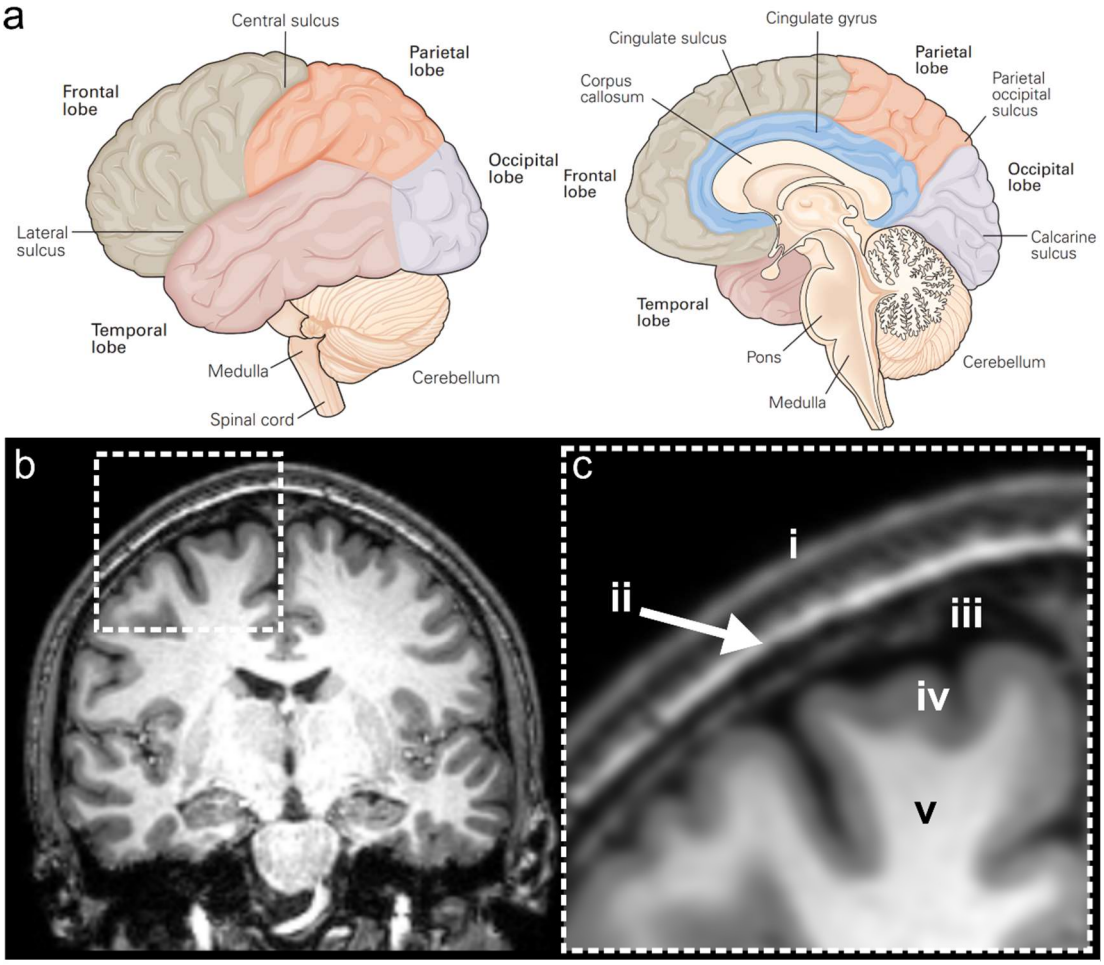

Figure 2.1 a) Macroscopic organisation of the brain showing the 4 main lobes (from (Schwartz J et al., 2012)*). b) Coronal slice of a T1 weighted MRI with inset in c. c) MRI image showing the scalp (i), skull (ii), cerebrospinal fluid (iii), grey matter (iv), white matter (v).*

2.1.2 The microscopic structure of the brain

At the cellular level, the brain consists of two major classes of cells: neurons and glia. Glia, which are named after the Greek word for 'glue' greatly outnumber neurons (or nerve cells) in the central nervous system. While they are *not* responsible for binding neurons together, they generally surround neuronal cells, with some types providing insulation between the electrically active neurons, others fulfilling roles in the immune system. Glia do not directly contribute to electrical signalling in the brain although they play an important

role in ensuring the efficient passing of electrical currents between neurons and are responsible for metabolic support.

2.1.2.1 *The Neuron*

The human brain contains around 10^{11} neurons (Schwartz J et al., 2012), which can be categorised into hundreds of types. In aid of simplicity, neurons have four constituent parts:

- The soma, which comprises the main part of the cell body and contains the cell nucleus.
- The dendrites, which are short protrusions extending from the soma and receive synaptic inputs from other neurons.
- The axon, a long tubular process that carries electrical impulses to other neurons.
- The presynaptic terminals, which connect the axon to the dendrites of other neurons.

The soma and neuronal dendrites predominantly sit in the grey matter while the axons can span long distances and form the white matter fibres carrying signals to and from other cortical regions and the brain stem. The cortex itself can be further subdivided into 6 distinct layers or laminae (Purves et al., 2001a) which can be seen in Figure 2.2. Each layer is largely defined by its characteristic cellular makeup, with layer IV containing many stellate neurons with axons connecting to nearby cells, for example, while layers V and VI are richer in pyramidal neurons with axons leaving the cortex to connect to deeper brain structures such as the thalamus. Layers II and III contain small pyramidal neurons with axons projecting to other cortical regions enabling communication between parts of the cortex.

The different configurations of neurons in the cortical laminae produce a variety of characteristic electrophysiological behaviours which are responsible for

common features of the MEG signal. For example, layer V has been shown to generate oscillations with a frequency in the range of 4-12 *Hz* in invasive electrophysiological recordings in animals for example (Silva et al., 1991), which corresponds to the theta and alpha band signals that are regularly observed in MEG recordings. Circuits in more superficial layers (e.g. II/III) are thought to give rise to higher frequency signals (Bonaiuto et al., 2018; Buffalo et al., 2011).

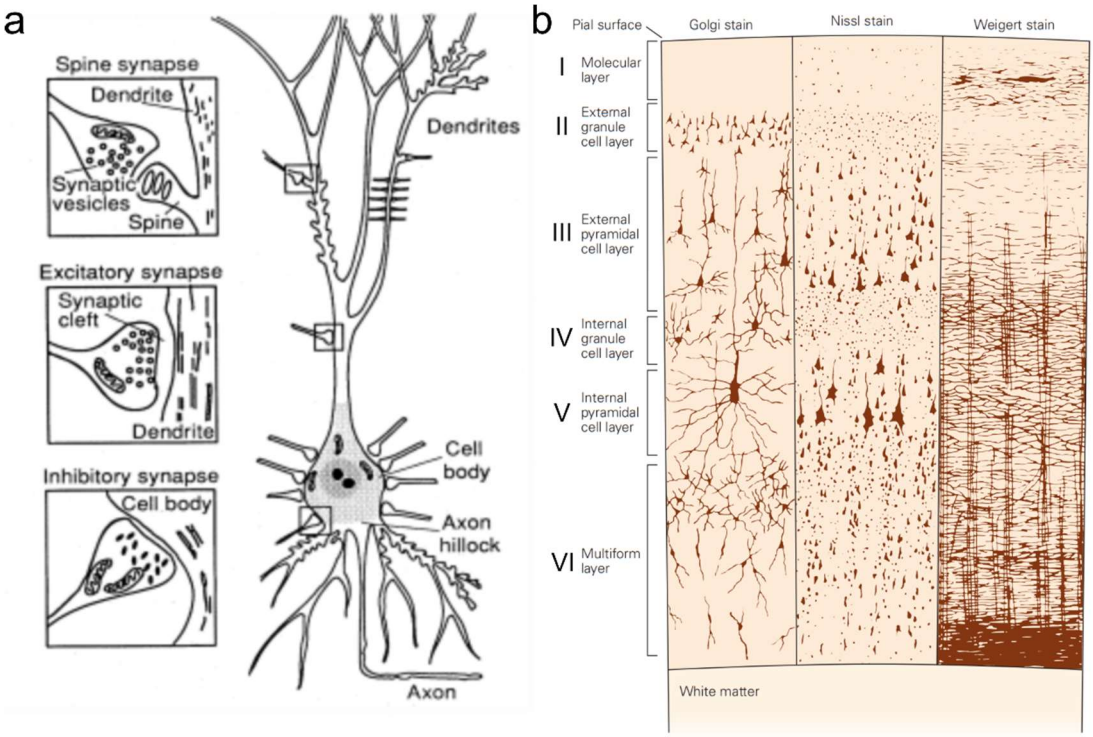

Figure 2.2: a) Schematic representaion of a pyramidal neuron adapted from (Hamalainen et al., 1993). b) Cortical stains showing the layer structure of the neocortex (from (Kandel, 2013))

The dendritic structure around the neuronal soma is an important factor in differentiating cell types in the cortex. Cortical stains reveal two distinct morphologies: stellate neurons with short dendrites radiating symmetrically from the soma, and pyramidal neurons with triangular soma (giving them the name 'pyramidal') and asymmetrically arranged dendrites which are arranged perpendicular to the cortex.

As will be described in more detail in the following sections, the currents generating the magnetic fields that are responsible for the MEG signal are carried by the dendrites of cortical neurons. Given that magnetic fields from a large number of synchronised cells need to superimpose to generate a measurable extra cranial field one can assume that the MEG signal is driven by pyramidal neurons, which constrain the direction of dendritic currents to be roughly parallel with their neighbouring pyramidal cells, allowing for constructive interference. The fields produced from currents in the dendrites of stellate neurons—on the other hand—can be assumed to destructively interfere.

2.1.2.1.1 Neurotransmission

The basis of all (electrical) signalling in the human nervous system is the action potential. First measured in the axon of a squid by Hodgkin and Huxley (1939), the action potential is a spike in intracellular potential that travels along the axon from the axon hillock towards the output synapses of the firing neuron. These spikes are transient deviations from the resting state of a neuron which is defined as a potential of −70 mV between the intracellular and extracellular environment, across the cell membrane. This resting potential is achieved via a balance of the intra- and extra-cellular concentrations of ionic species of potassium (K^+), sodium (Na^+), calcium (Ca^{2+}) and chlorine (Cl^-). This equilibrium of concentrations is maintained by ion channels and pumps the former selectively allowing specific ions to permeate the membrane, the latter actively transporting ions across the membrane. Ion channels contain gating mechanisms that can be closed or open, depending on the potential—in the case of voltage-gated channels—or on the attachment of ligands such as neurotransmitters, Ca^{2+} ions or cyclic nucleotides, which are organic messenger molecules (Purves et al., 2001b).

2.1.2.1.2 The action potential

While a potential of -70 mV is maintained by neuronal ion channels and pumps at rest, changes in the cellular environment such as an influx of neurotransmitters from an input synapse can initiate the opening of ion channels causing alterations in the polarisation of the cell. In the case of an excitatory synapse, a positive change in potential is caused while input at inhibitory synapses is followed by a decrease in potential. If the integration of these changes leads to a depolarisation that reaches a threshold of approximately -55 mV, voltage-gated sodium and potassium channels open leading to an influx of Na^+ and efflux of K^+. The threshold voltage of -55 mV presents a point of no return, with a rapid influx of Na^+ increasing the potential to $+40$ mV; the peak of the neuronal *spike*. At this potential, Na^+ channels are closed and K^+ ions flow out of the cell, eventually achieving a state of hyperpolarisation (a potential of < -70 mV) before the resting potential is recovered. Action potentials are usually initiated at the axon hillock, which contains a higher concentration of Na^+ channels and is, therefore, most susceptible to depolarisation. Once depolarisation is initiated at the axon hillock, the change in potential causes a depolarisation event at an adjacent section of the membrane. The depolarisation edge (Figure 2.3), which is followed by a repolarisation edge, propagates along the axon without a loss in amplitude.

2.1.2.1.3 The post-synaptic potential

Once an action potential from a pre-synaptic cell reaches an output synapse, neurotransmitters are released which initiate the opening of ion channels at the dendrite of the postsynaptic cell. The resulting change in polarisation—the post-synaptic potential—leads to an intracellular current of flowing ions along the dendrite, also known as the primary current. This post-synaptic current decays exponentially in strength with distance (r) from the input synapse,

$$J(r) \propto e^{-\frac{r}{\lambda}}, \tag{2.1}$$

where the characteristic decay length $\lambda = \sqrt{(r_m/r_i)}$ is dependent on the resistivities r_m and r_i of the membrane and intracellular fluid respectively. λ is typically given as 0.1 to 0.2 mm (Hari and Puce, 2017). From afar, this current can be modelled as a current dipole

$$Q = \lambda I, \tag{2.2}$$

which can be rewritten using Ohm's law to give

$$Q = \frac{\Delta V \lambda}{R}, \tag{2.3}$$

where R is the cells' internal resistance ($R = r_i \lambda$). Given that r_i, depends inversely on the cross-sectional area of the dendrite in question, equation 2.3 can be expressed in terms of the dendritic diameter d and the intracellular conductivity σ_i:

$$Q = \frac{\Delta V \sigma_i \pi d^2}{4}. \tag{2.4}$$

Typical values for the terms in equation 2.4 have been determined experimentally ($d = 1\ \mu m,\ \sigma_i = 1\ \Omega^{-1}m^{-1}, \Delta V = 25\ mV$), and yield $Q \sim 20\ fAm$ for a single post-synaptic current (Hamalainen et al., 1993). Neuromagnetic fields are commonly on the order of $10fT$, which, for a typical sensor geometry and source depth, would require a $\sim 1\ nAm$ source. This means that the MEG signal must be produced by around 10^4 to 10^5 synchronised dendrites.

To restore the resting potential, ions are pumped against the concentration gradient leading to a restorative current between the ion pump and the synaptic junction termed the volume current, which flows outside the post-synaptic cell and in the opposite direction to the primary current (see Figure 2.3).

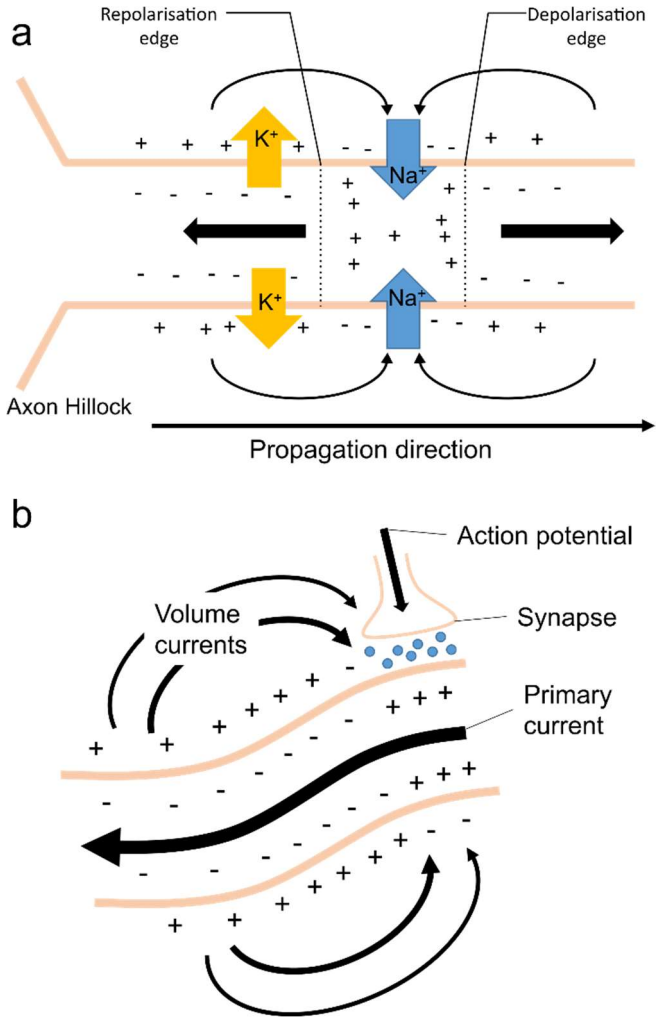

Figure 2.3 a) Action potential. b) Post synaptic and associated volume current.

2.1.3 Measurable and non-measurable effects

As we have seen three candidate currents may produce the fields measured using MEG. The current due to the action potential can be modelled as two opposing current dipoles pointing away from the depolarisation edge as it travels along the axon (Figure 2.3), lasting approximately 1 ms. The primary post-synaptic current can be modelled as a single current dipole, typically lasting on the order of tens of milliseconds (Hamalainen et al., 1993). Given the typical current dipole strengths assumed for the action, and post-synaptic potentials, it is necessary for a large number of neurons to fire in synchrony. This means that the integration of the longer-lasting postsynaptic currents is

much more likely. Additionally, the magnetic fields produced by a current dipole and quadrupole decay with distance, r, as $\frac{1}{r^2}$ and $\frac{1}{r^3}$ respectively. Hence it is likely that the primary currents are the more important contributor to the MEG signal.

The volume currents also produce a magnetic field, however, in the case of a spherical approximation of the head, the field due to this type of current can be ignored(Hamalainen et al., 1993). This approximation is made to simplify the analytical solution of the forward problem (discussed in chapter 3), however more realistic head models can be used to include the contribution of volume currents to the MEG signal.

In summary, the MEG signal originates in the dendrites of a large number of synchronised pyramidal neurons in the cortex and some subcortical grey matter structures, with the current dipoles due to the primary post-synaptic potentials providing the bulk of the signal.

2.2 TYPES OF ELECTRICAL BRAIN ACTIVITY

Given the excellent temporal resolution of MEG, a great wealth of information can be extracted from the evolution of the measured neuromagnetic fields. The brain can be exposed to a plethora of situations and tasks, or be studied at rest, in health and disease, with the MEG signal providing a window into its function (or dysfunction). In this section, I will outline a series of types of activity that can be measured using MEG.

2.2.1 Spontaneous activity

Early electroencephalography experiments conducted by Hans Berger (Berger, 1929) lead to the discovery of spontaneous or intrinsic brain activity which is present at all times without the requirement of specific stimuli. Berger observed a rhythmically changing signal with a frequency of approximately 10 Hz which

originated in the occipital lobe and increased in amplitude when the subject closed their eyes. This characteristic signal has since been defined as a rhythm in the alpha band which spans the 8-13 Hz frequency range. Other spontaneously occurring oscillatory signals have been observed and categorised since then (see Table 2.1).

Label	**Canonical Frequency Band**
Delta ($\boldsymbol{\delta}$)	$< 4\ Hz$
Theta ($\boldsymbol{\theta}$)	4-8 Hz
Alpha ($\boldsymbol{\alpha}$)	8-13 Hz
Beta ($\boldsymbol{\beta}$)	13-30 Hz
Gamma ($\boldsymbol{\gamma}$)	$> 30\ \ Hz$

Table 2.1 Canonical frequency bands of neural oscillations

A commonly studied feature of intrinsic brain activity is the oscillatory power in the signal at frequencies of interest which can be estimated using a discrete Fourier transform or adjacent data analysis techniques.

2.2.2 Evoked activity

When a subject is exposed to an external, transient stimulus, brain regions relevant to the stimulus will often produce a time- and phase-locked signal. In the event of a brief tone being played to a subject, for example, the auditory cortex will produce a stereotypically shaped signal with a fixed latency relative to the stimulus. Averaging the signals from a large number of trials will result in the cancellation of any non-phase locked signals leading to a reduction of random noise by a factor of $\sqrt{\mathrm{N}}$, N being the number of trials. This allows for an accurate estimation of the stereotypical signal shape which can then be used to make inferences about any abnormal or altered functioning of the relevant brain region. When used in conjunction with source modelling, the measurement of

evoked activity can be of use when mapping the cortical organisation of sensory input.

2.2.3 Induced activity

Not all stimulus-related signals are precisely phase-locked. Induced responses are commonly observed as modulations of spontaneous activity which occur locked in time to a stimulus but with variable phases. Some often reproduced examples are decreases in beta band power in the motor cortex during movement (Pfurtscheller et al., 1996; Pfurtscheller and Lopes Da Silva, 1999) or modulation of occipital alpha power in response to a visual stimulus (Pfurtscheller et al., 1994). Experimental designs used to measure induced activity still involve repetitive stimulation over several trials, however, simple averaging would result in the cancellation of the non-phase-locked signals. A common approach to circumvent these issues is to filter the signal to a frequency band of interest and to estimate a measure of signal power before averaging across trials. This can be done by simply squaring the filtered signal, however, the use of the Hilbert transform (see Chapter 3) to estimate the amplitude envelope is more common. This approach can be extended to produce time-frequency spectra (TFS) in order to visualise the evolution of induced responses during a trial. Figure 2.4 shows an example TFS where a subject performs a motor task (see the maze paradigm in Chapter 6). One can clearly observe the modulation of amplitude in the beta band: beta-band power is reduced during movement and is followed by an increase above baseline after movement cessation before the signal returns to baseline.

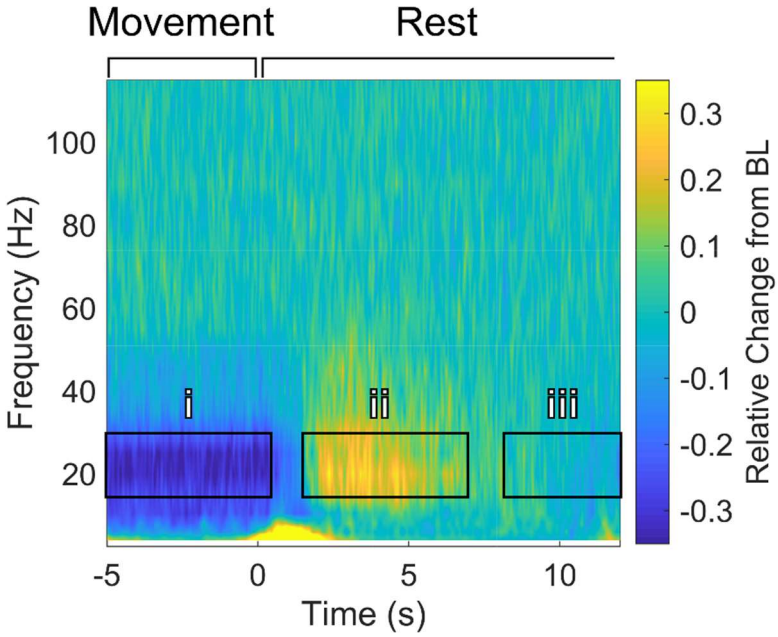

Figure 2.4 Example TFS from a single optically pumped magnetometer placed above the left motor cortex. Data was recorded from a single subject while performing a motor task with their right arm and hand. i) shows the beta-band desynchronization during movement, ii) the rebound above baseline when the movement stops at 0s, and iii) the time window used as baseline to calculate relative change.

2.2.4 Functional connectivity

The view that the brain can be compartmentalised into functionally distinct regions is pervasive at least since the postulations by Gall in the early 19th century (Gall, 1835). While there exists ample evidence of specific cortical regions performing specialised functions, the vast range and complexity of human (and animal) behaviours are only possible with large parts of the brain acting in concert. Advances in neuroscientific research at the cellular and whole-brain level over the past few decades have produced a picture of the human brain as a highly adaptable and flexible interconnected network not only in structure but also in function.

Functional connectivity can be defined as a statistical dependence between the signals produced in distinct parts of the cortex and is thought of as representing

information flow or communication between separate regions. Fries proposed the *communication via coherence hypothesis* (Fries, 2005), which suggests that the synchronisation of neural oscillations in spatially distinct parts of the cortex provides a mechanism for information transfer between regions. This model is based on the principle that the oscillatory behaviour of a neuronal population represents repeating windows of increased excitability and hence an increased likelihood of communication. Effective communication between two neuronal populations can, therefore, be achieved when these windows of optimal communication coincide (see the illustration in Figure 2.5). This implies a synchronisation of the oscillatory patterns of the two oscillations.

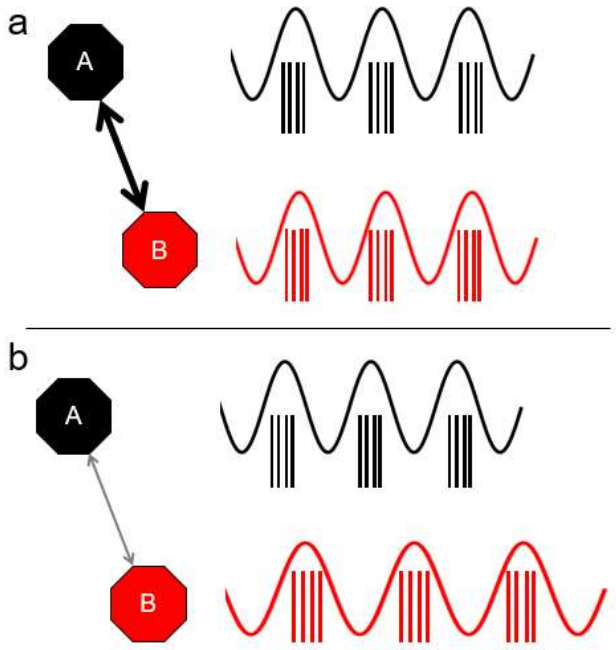

Figure 2.5 Schematic representation of two neuronal populations exhibiting oscillatory behaviour. a) Populations A and B are firing at the same frequency and approximately in phase which represents strong connectivity. b) Populations A and B are still exhibiting oscillatory behaviour. They are, however not synchronised, which would imply weak or no connectivity between A and B.

Functional connectivity has been measured using several neuroimaging modalities such as PET and fMRI, however, given its excellent temporal resolution, MEG is a particularly useful modality to investigate the dynamic communication between brain regions at rest but also during tasks. Importantly, it not only provides us with a better understanding of the mechanisms behind cognition but also a way of assessing clinical populations in which effective communication across the brain is disrupted. Abnormalities in functional connectivity, as measured by MEG, have been reported in a number of diseases and are actively researched in the field of mTBI (see section 1.3). Methodological details about how functional connectivity is estimated from MEG data are discussed in Chapter 3.

2.2.5 Transient bursts

Measures based on neural oscillations are ubiquitous in the study of cognition in health and disease, in the investigation of task responses, and at rest. Effects like the movement-related desynchronization and post-movement rebound in the beta band for example are well described in the literature and can be aberrant in clinical populations (e.g. Parkinson's disease (Heinrichs-Graham et al., 2014; McColgan et al., 2020)) However, recent developments in the field have put into question whether the 'oscillatory' picture of these types of activity is valid in all cases (van Ede et al., 2018). Sustained modulation of neural activity within a frequency band is conventionally measured via trial averaged time-frequency analyses. Similarly, spectral power at rest can be assessed by estimating an average power spectrum over an entire MEG recording. While such measures are appropriate and of great use in the study of brain function, an inspection of unaveraged time-frequency representations of MEG data reveals that the signals are, in fact, composed of stochastically occurring, transient events. Only when averaged do they appear to follow a smoothly varying oscillation.

While this way of reconceptualising neural oscillations is increasingly studied, it is yet to be determined what implications this 'burst' picture of oscillatory activity will have on our understanding of brain activity. However, this approach brings with it an additional wealth of measures, such as burst rate, duration and timing relative to stimuli, which may provide useful features in studying disease. Chapter 3 will describe a study of subjects with mTBI, where a burst approach was used to generate potential biomarkers for mTBI.

2.3 MEASURING THE MEG SIGNAL

We have so far established the principles behind the electrical signals generated in the brain. Based on our understanding of the action potential, the post-synaptic currents, and the organisation of cortical neurons we can conclude that synchronised, post-synaptic currents in tens of thousands of dendrites produce a weak but measurable magnetic field outside the head. In the following sections, I will cover two technologies that have enabled us to perform these measurements. Firstly, I will describe the physical principles of magnetic field sensing using cryogenically cooled Superconducting Quantum Interference Devices, or SQUIDs, which are used in the current state-of-the-art magnetoencephalography equipment and were used in the data collection for the experiments presented in Chapter 4 and Chapter 5. Secondly, I will cover the Optically Pumped Magnetometer (OPM). This type of magnetic field sensor has seen increasing excitement in the MEG field as it allows for room temperature measurements of the MEG signal using wearable sensor arrays, but crucially, at very short distances from the scalp, promising higher sensitivity and spatial resolution compared to conventional MEG. OPMs were used to produce the results discussed in Chapter 6. Finally, a range of techniques used to reduce interference from external sources is discussed.

2.3.1 Superconductor-based MEG

When designing a sensor to measure bio-magnetic fields, a perhaps obvious approach would be to exploit the electromagnetic phenomenon of currents being induced in loops of electrically conductive materials when exposed to varying magnetic fields. Given that the fields produced by the primary post-synaptic currents (described in section 2.1.3) are only a fraction of a pT in strength, however, means that resistive coils would lack the necessary sensitivity unless constructed using an excessively large number of turns. Cohen demonstrated that cerebral magnetic fields could be measured this way in 1968 (Cohen, 1968), using a 1 million turn sensing coil. Cohen also showed, shortly after, that the magnetic fields produced by the brain could be measured with greater SNR using a SQUID magnetometer (Cohen, 1972).

2.3.1.1 *Superconductivity*

With the discovery that the electrical resistance of mercury vanishes at liquid helium temperatures in 1911 (Onnes, 1911), Dutch physicist Heike Kammerlingh Onnes first described the phenomenon of superconductivity and showed that it occurred at a critical temperature T_c. Investigating the magnetic properties of materials in the superconductive state, Meissner and Ochsenfeld showed that superconductors behaved like ideal diamagnets, expelling any magnetic fields within their bulk, but that this effect could be destroyed above a critical field B_c. It was also found that some materials appeared to undergo two separate phase transitions at fields B_{c1} and B_{c2} (where $B_{c1} \ll B_{c2}$), between which magnetic flux can penetrate the material and increases with the strength of the external field. Above the critical field B_{c2} superconductivity no longer occurs (Rjabinin and Shubnikow, 1935). This second class of materials are called type-II superconductors (in contrast to type-I with only one critical field B_c), and form the basis of SQUID magnetometers.

It was not until long after Onnes' first descriptions of the effect that a theoretical explanation of superconductivity was successfully developed. Bardeen, Cooper and Schrieffer proposed such a theory (BCS theory hereinafter) in 1957 (Bardeen et al., 1957). BCS theory proposes the formation of electron pairs, or Cooper pairs—pseudo-particles that act like bosons and can therefore occupy the same quantum state. Bardeen et al. showed that this phenomenon favours a state where the motion of all Cooper pairs must occur such that their momentum and direction of motion is equivalent which gives rise to the observed supercurrents.

The combined quantum mechanical wave function describing the ensemble of Cooper pairs can be described using an order parameter of the form

$$\psi(\mathbf{r}) = \psi_0 e^{i\boldsymbol{\theta}} \tag{2.5}$$

where ψ_0 is the ground-state wave function of the Cooper pair and θ is a phase term that depends on the centre of mass $\boldsymbol{r}$ common to all cooper pairs and their momentum $\boldsymbol{q}$ ($\theta = \boldsymbol{q} \cdot \boldsymbol{r}$). The current density $\boldsymbol{j}$, as a function of position, can be written as

$$\mathbf{j}(\mathbf{r}) = -\frac{e}{m}|\psi(\mathbf{r})|^2(\hbar\boldsymbol{\nabla}\theta + 2e\mathbf{A}) \tag{2.6}$$

Equation 2.6 shows how the current density within a superconductor relates to the magnetic vector potential $\boldsymbol{A}$, the electron charge e and mass m, the reduced Planck constant $\hbar$ and the phase θ of the wave function. Given that superconductors expel all magnetic fields, the field in the bulk of the material is zero which means that the current density must be zero according to Ampere's law. We can therefore rewrite 2.6 to obtain

$$\hbar\nabla\theta = -2e\mathbf{A} \tag{2.7}$$

and integrate around a closed path C within the superconductor. The left-hand term in 2.7 will yield the phase difference of the wave function along C as shown in 2.8

$$\hbar\Delta\theta = -2e\oint \mathbf{A}.\,d\mathbf{l}\,. \tag{2.8}$$

If we consider Stokes theorem

$$\int_S \nabla \times \mathbf{F} \cdot d\mathbf{S} = \oint_C \mathbf{F} \cdot d\mathbf{r} \tag{2.9}$$

and the fact that $\nabla \times \mathbf{A} = \mathbf{B}$ (the magnetic field), the path integral term in 2.8 can be expressed as an integral over the surface enclosed by C:

$$\hbar\Delta\theta = -2e\oint \mathbf{A}.\,d\mathbf{l} = -2e\iint_S \mathbf{B}.\,d\mathbf{S}\,. \tag{2.10}$$

The surface integral of **B** can be seen as the magnetic flux through the loop defined by C. Additionally, we can use the property of the wave function that the phase change around the loop must be a multiple of 2π to rewrite 2.10 as

$$2\pi n\hbar = -2e\boldsymbol{\Phi}, \tag{2.11}$$

showing that magnetic flux in a superconducting loop is quantised!

2.3.1.2 *The Josephson Effect*

In order to allow for indirect measurement of the magnetic flux through a superconducting loop, the Josephson Effect can be utilised. It was demonstrated that superconductivity was not lost when introducing an insulating barrier—a Josephson junction—between two superconductors (Josephson, 1962). Josephson demonstrated that, up to a critical current, Cooper pairs would be able to tunnel across the barrier without resistance. However, once the critical current is reached (by applying a bias current across the junction for example), a potential difference can be measured across the junction.

The SQUIDs used in currently available, commercial systems utilise two Josephson junctions in a *DC SQUID* configuration which is shown in Figure 2.6.

For such a system it can be shown that the potential difference across the SQUID varies periodically with the flux through the superconducting loop as seen in equation 2.12

$$V \propto A \cos\left(\frac{e}{\hbar}\boldsymbol{\Phi}\right) \tag{2.12}$$

where A is the effective area of the SQUID loop and the period is the flux quantum $\boldsymbol{\Phi_0}$. Given that $\boldsymbol{\Phi_0} = 2.7 \times 10^{-1}$ Wb or Tm^2 SQUIDs are highly sensitive to the minute magnetic fields produced by the brain.

Due to the periodicity of the function in 2.12, any potential measured does not correspond to a unique flux. A feedback circuit with a lock point as shown in Figure 2.6 can, therefore, be used to measure the applied flux in an approximately linear part of the transfer function. If the flux is varied to deviate from this lock point, the feedback circuit will apply an equal and opposite potential to maintain a constant flux. This compensatory signal is used to provide the flux measurement.

In order to shield the SQUIDs from interference, and to ensure optimal capturing of flux produced by the brain, cryogenic MEG systems use flux transformers made of superconducting coils (see diagram in Figure 2.6), which are coupled to a SQUID further away from the head surface.

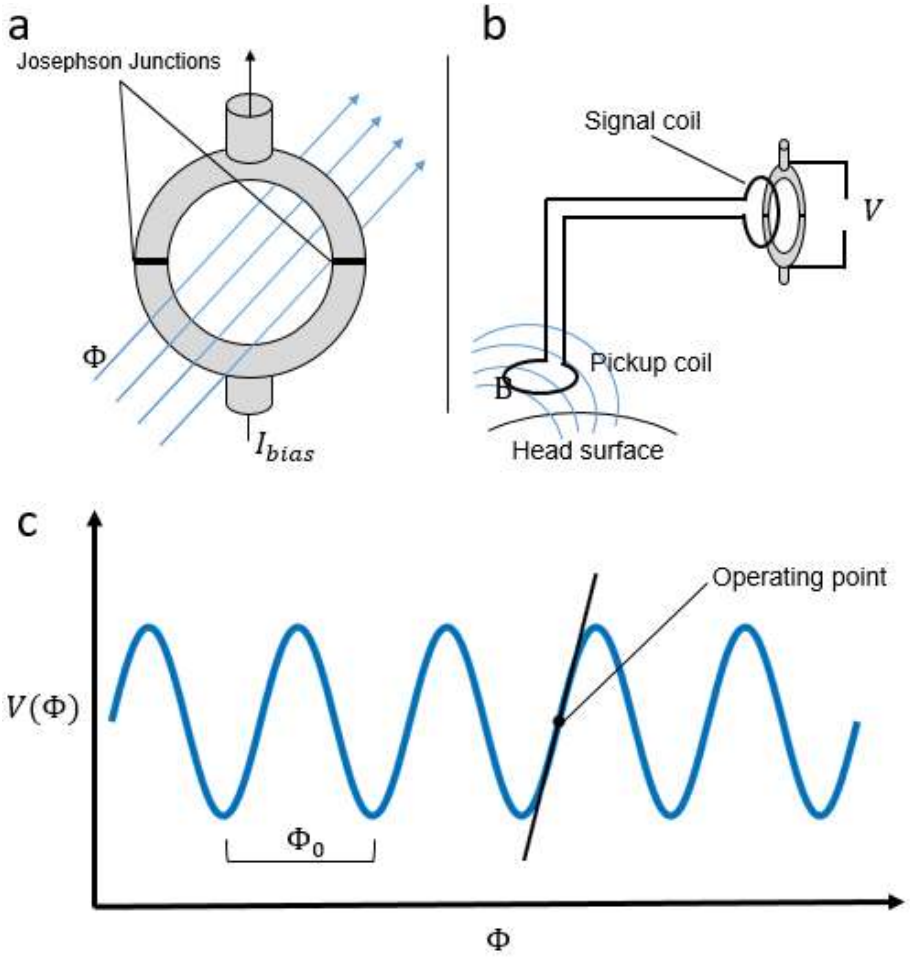

Figure 2.6 a) Schematic showing a SQUID. Two Josephson junctions interrupt a ring of superconducting material. If a bias current above the critical current flows across both junctions and an external magnetic field is applied, a voltage across the junctions can be measured. b) Flux transformer setup. Neuromagnetic fields induce a current in a superconducting coil which is inductively coupled via a signal coil to the SQUID. c) Relationship between the voltage across a SQUID magnetometer given an input flux Φ. In order to keep the SQUID within the linear regime around the operating point, a negative feedback circuit applies an offset current to counteract any changes in flux from the lock point. This feedback signal provides an indirect measure of the changes in flux through the SQUID.

2.3.2 Optically pumped magnetometers

For more than four decades, SQUIDs were the only viable technology for MEG sensors. Advances in quantum sensing technology in recent years, however, enabled the manufacturing of cryogen-free magnetic field sensors with the sensitivity required for magnetoencephalographic measurements. Optical pumping is the controlled manipulation of the spin angular momentum of a particle (Gerlach and Stern, 1922) by irradiation with light photons (Kastler, 1973). Through the absorption and emission of laser light photons, it is possible to bring an ensemble of atoms into a state that is sensitive to changes in an applied magnetic field, allowing for sensitive magnetometry. Early work on optically pumped magnetometers (OPMs) demonstrated that vapours of alkali metals could be brought into a state suitable for the measurement of weak magnetic fields (Bell and Bloom, 1957; Bloom, 1962). While initial experimental setups did not reach the sensitivity or form factor required for use in MEG sensors, the application of modulation fields (Dupont-Roc et al., 1969; Kastler, 1973; Slocum and Marton, 1973) and the utilisation of high-density vapours (Happer and Tang, 1973) allowed for miniaturisation and sensitivities comparable to SQUID based sensors (Budker et al., 2000).

In recent years, progress in microfabrication techniques enabled the production of OPM sensors suitable for neuromagnetic measurements (Schwindt et al., 2005, 2004) and early demonstrations of MEG measurements included the recording of evoked responses to median nerve and auditory stimulation (Johnson et al., 2010, 2013) as well as visual alpha oscillations (Kamada et al., 2015). Simulations also demonstrated that on-scalp sensors could deliver improvements in SNR and spatial resolution over conventional MEG (Boto et al., 2016; Iivanainen et al., 2017). As well as the commercialisation and large scale manufacturing of OPMs, the construction of the first, multi-sensor MEG systems generated much excitement in the field. At the time of writing, three companies produce MEG-suitable OPM sensors: QuSpin Inc., FiledLine Inc. and

Kernel. Using a single QuSpin sensor, Boto et al. successfully measured evoked and induced fields by repeatedly placing the sensor at varying locations mimicking early single-SQUID MEG experiments (Boto et al., 2017). Shortly after, a multichannel array of 13 OPMs was used to localise motor-related responses (Boto et al., 2018). In the following years, much progress was made with regards to the form factor of the sensors themselves as well as the methods to mount the sensors on the head. While early attempts relied on fixed arrays (Borna et al., 2017) or heavy, personalised head casts (Boto et al., 2018), more user-friendly, general-purpose mounting solutions were soon adopted. By adapting commercially available cycling helmets Hill et al. demonstrated the adaptability of OPM sensor arrays to study human brain activity at all ages (Hill et al., 2019). Iivanainen et al. used an additively manufactured helmet with sensor slots covering the entire head allowing for adaptable positioning of sensors (Iivanainen et al., 2019). An additively manufactured helmet was also used by Boto et al. (2021) and Hill et al. (2020). Hill et al. also compared the rigid helmet approach with a lightweight flexible cap (Hill et al., 2020).

In this thesis, an array of QuSpin sensors, mounted in an additively manufactured helmet were used to perform the experiments discussed in Chapter 6. In the following section, the QuSpin OPM will be described in more detail as well as the theoretical principles of optical pumping.

2.3.2.1 *QuSpin OPM Sensors*

To date, QuSpin has produced 3 generations of OPM sensors (Gen-1 to 3), a combination of which were used in chapter 6. Gen-1 sensors were released in 2016 (Figure 2.7). Miniaturisation and weight reduction in the connectors made Gen-2 and 3 sensors more useful for head-mounted multi-sensor MEG arrays. Most recently, triaxial sensors, which can measure three orthogonal components of a magnetic field simultaneously were introduced. While the 3 sensor types differ slightly in the layout of internal components, the physical principles behind their operation are identical.

The simplest OPM setup includes a glass cell containing high-pressure vapour, a laser and a photodiode as a detector (Figure 2.7). For the vapour, usually, alkali metals are used given the simple electronic structure with a single electron in the outer shell. QuSpin sensors use Rubidium-87 in a 3×3×3 mm^3 glass cell, and a diode laser with a wavelength of 795 nm. The high-density vapour necessary for stable sensing (as discussed in 2.3.2.3) is generated by increasing the cell temperature to 150℃ using resistive heating coils. Note that the cells are not directly in contact with the outer casing of the sensor which is not heated beyond body temperature.

Electromagnetic coils around the cell are used to offset the background field (up to 50nT (Shah and Hughes, 2015) as well as the modulation fields which will be described later. A schematic diagram of the optical path can be seen in Figure 2.7. The laser light is circularly polarised, collimated and guided through the cell before hitting a photodetector. The latter provides a measurement of the light intensity which is modulated by any changes in the external field.

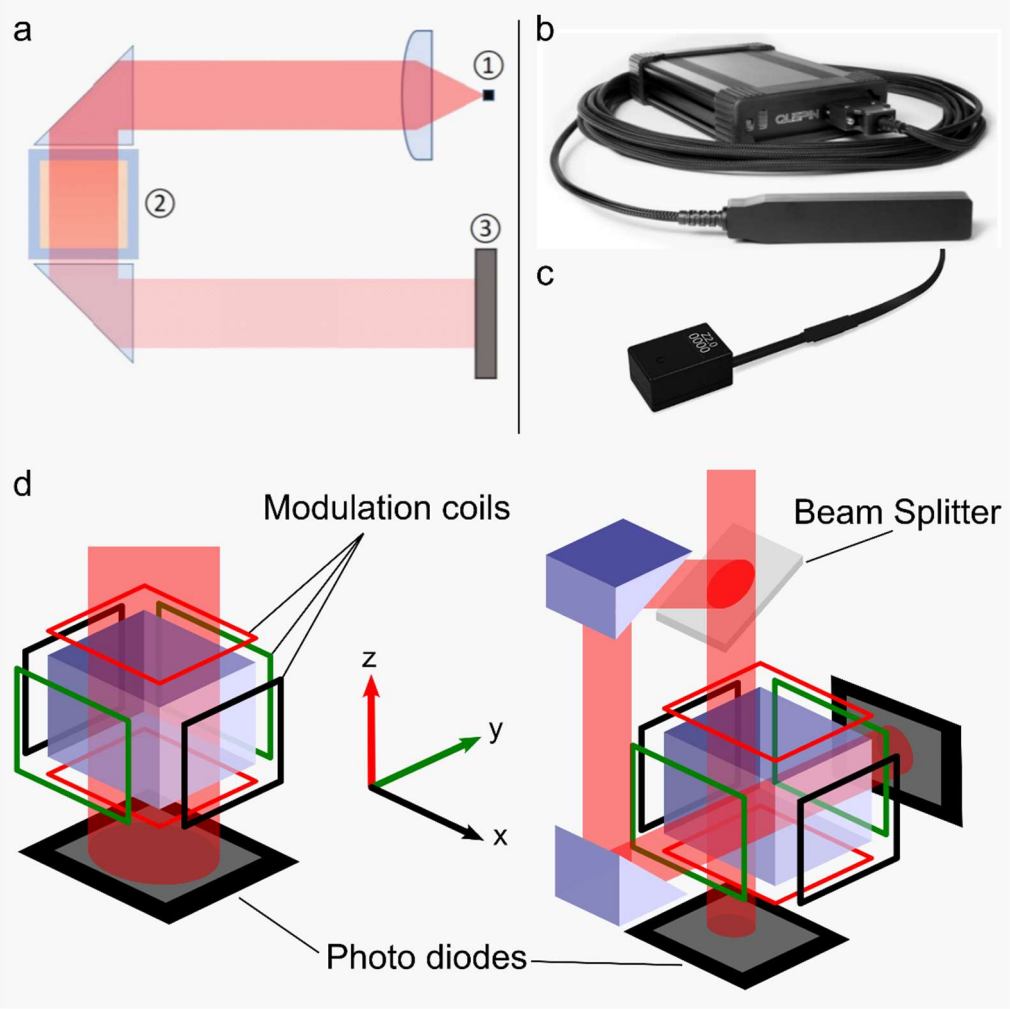

Figure 2.7 a) Essential components of OPM sensors. LASER diode and collimating lens (1). Glass cell containing ^{87}Rb vapour (2). Photodiode (3). b) QuSpin Gen-1 sensor and control electronics. c) QuSpin Gen-2 sensor. d) Field zeroing and modulation coils for dual (left) and triaxial sensors (right).

2.3.2.2 *Optical Pumping*

Optical pumping refers to the use of a light source to bring an entire population of atoms (or the great majority) into the same quantum state with a specific energy level. This is achieved by tuning the LASER light to a frequency that, when absorbed by ^{87}Rb atoms will induce a specific change in energy and angular momentum. This quantum analogue of classical angular momentum

and its behaviour in a magnetic field gives rise to the effect that allows for optical pumping to occur and hence forms the basis of OPM sensing.

The total angular momentum of an atom $\boldsymbol{F}$ is a combination of its nuclear ($\boldsymbol{I}$) and electron ($\boldsymbol{J}$) angular momenta,

$$\boldsymbol{F} = \boldsymbol{J} + \boldsymbol{I} \tag{3.13}$$

where **J** is itself a combination of the electron's orbital angular momentum **L** and spin angular momentum **S**, hence

$$\boldsymbol{F} = (\boldsymbol{L} + \boldsymbol{S}) + \boldsymbol{I}. \tag{2.14}$$

F is associated with an atom's magnetic moment as shown in equation 2.15,

$$\boldsymbol{\mu_F} = \gamma \mathbf{F}, \tag{2.15}$$

where γ is the gyromagnetic ratio of the atom; γ has a value of approximately $7HznT^{-1}$ for ^{87}Rb. The magnetic moment of the atom in the presence of a magnetic field **B** is associated with a potential energy

$$\mathrm{E} = -\boldsymbol{\mu_F} \cdot \mathbf{B}, \tag{2.16}$$

which means that different values of angular momentum will result in different energy levels that can be occupied by the atom. When exposed to laser light, atoms can absorb any energy and angular momentum associated with the incoming photons and therefore change their quantum state. Given that we are in the quantum realm, the outer shell electron of a ^{87}Rb atom can exist in a series of discrete energy states. The absorption or release of a photon can therefore only occur if the resulting state transition corresponds to a change between two of these possible quantum states.

The set of possible states is determined by the interaction of the electron's orbital and spin angular momenta which causes a phenomenon termed fine

structure splitting. For the vector quantities **L** and **S**, dimensionless quantum numbers L and S can be defined, respectively, which are defined as follows:

$$|\mathbf{S}| = \sqrt{\mathrm{S}(\mathrm{S}+1)}\hbar, \text{where } \mathrm{S} \in \left\{\frac{1}{2}\right\}; \tag{2.17}$$

$$|\mathbf{L}| = \sqrt{\mathrm{L}(\mathrm{L}+1)}\hbar, \text{where } \mathrm{L} \in \{0,1,2,3,4\}. \tag{2.18}$$

The value of the momenta is reliant on $\hbar$, the reduced Planck's constant.

All possible values for the dimensionless quantum number J, associated with the electron angular momentum can be found by taking integer steps in the range between $|\mathrm{L}-\mathrm{S}|$ and $|\mathrm{L}+\mathrm{S}|$. In the ground state, where $\mathrm{L}=0$, this means that $J = \frac{1}{2}$, for the first excited state $\mathrm{L} = 1$, J can take the values $\frac{1}{2}$ and $\frac{3}{2}$. Transitions between the $L = 0$ state and the two $\mathrm{L} = 1$ states are known as the D1 (for the $\mathrm{J} = \frac{1}{2}$ state) and D2 (for the $\mathrm{J} = \frac{3}{2}$ state) transitions. The wavelengths of the photons that need to be absorbed for these transitions to occur are 795 and 780 nm respectively. In QuSpin OPMs only D1 transitions are selected due to the 795 nm laser.

Further splitting of the fine structure states occurs due to the interaction between the total angular momenta of the electron (**J**) and the nucleus (**I**). This hyperfine structure splitting is defined by the quantum number F which can take values from $|\mathrm{I}-\mathrm{J}$ to $|\mathrm{I}+\mathrm{J}|$, again in integer steps. The nuclear quantum number I for ^{87}Rb is $3/2$, F can therefore only take the values 1 or 2. A final splitting of the hyperfine structure is called Zeeman splitting (Zeeman, 1897) and occurs due to the interaction of the total angular momentum **F** with the applied magnetic field. The relevant component of the angular momentum falls along the laser beam axis and can be defined as $\mathrm{F}_{\mathrm{laser}} = \hbar \mathrm{m}_{\mathrm{f}}$, where m_{f} can take any value (in integer steps) from – F to F. The full energy spectrum containing these levels can be seen in Figure 2.8.

The σ^+ circular polarisation of the laser light carries an angular momentum of +1. The absorption of 795 nm photons produced by the laser can therefore induce any D1 transition with a change in m_f of +1. Once an electron de-excites, a photon of the same wavelength is emitted, however with a random angular momentum m_f which can take the values −1, 0 or 1. Over time, this leads to an accumulation of atoms in the $L = 0$, $F = 2$ state. Given the selection rule $\Delta m_f = +1$, no further D1 excitation is possible as there are no $F = 3$ levels in the $L = 1$ state. This means that no further photons can be absorbed, rendering the gas transparent and leading to the maximum signal at the photo-diode. At this stage, the gas is highly polarised and produces a strong net magnetisation which experiences a torque due to any magnetic field components perpendicular to the laser beam axis and is therefore magnetically sensitive.

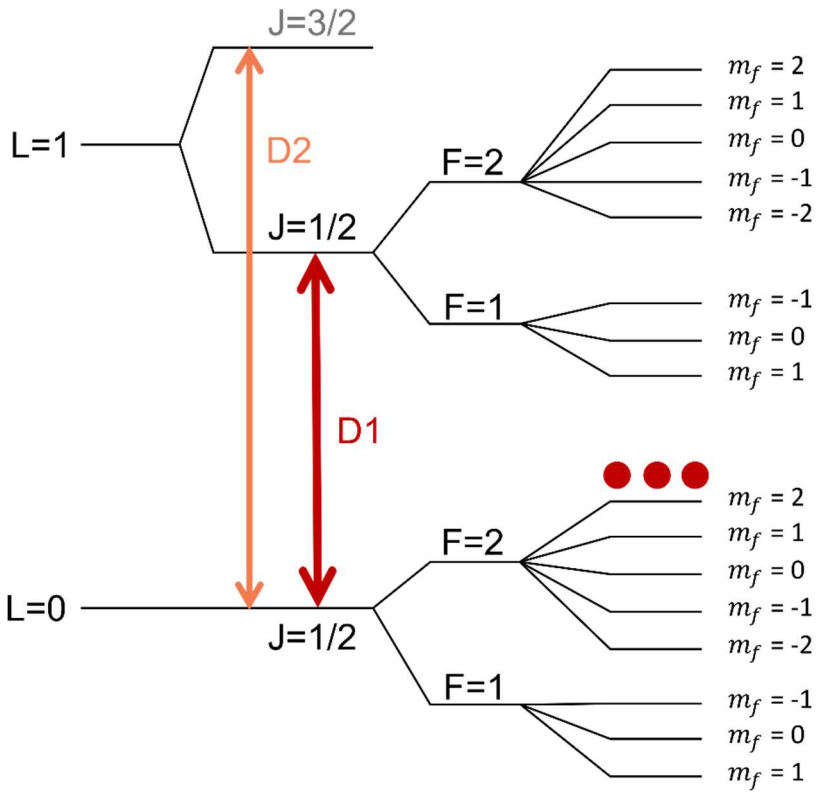

Figure 2.8 Energy level diagram for ^{87}Rb. A 795 nm circularly polarised photon will cause a D1 transition with a change in m_f of +1. Over time, all atoms will accumulate in the $L = 0$, $F = 2, m_f = 2$, state.

2.3.2.3 *Spin exchange relaxation*

Even with continuous pumping, the net polarisation of the gas is not permanent due to relaxation processes which cause a loss of sensitivity. Collisions with other atoms, the cell walls and field inhomogeneity (Happer, 1972) are examples of such processes that can cause relaxation. Spin-exchange collisions are considered to be the dominant form of relaxation in high-density vapours, however. In such events, two Rubidium atoms collide leading to a potential exchange of spins. If one of the atoms involved in the collision changes its F-level, the sign of its spin will flip, causing, in turn, a change in the gyromagnetic ratio γ which leads to rapid decoherence in the vapour as the atoms now precess out of phase. The relaxation rate due to spin-exchange R is given as $R = (T_{SE})^{-1}$ where T_{SE} is the average time between collisions. Perhaps counterintuitively, this effect can be mitigated by increasing the vapour pressure (through heating for example), which increases the number of collisions and therefore R. If $R \gg \gamma B$, the spin-exchange collisions occur much faster than the natural precession frequency and a net precession is maintained (Happer and Tam, 1977)—albeit at a slightly slower frequency. This average precession frequency in the spin-exchange relaxation free (SERF) regime is given by

$$\Omega' = \frac{\gamma B}{Q}, \tag{2.19}$$

where Q is a slowing factor dependent on the fraction of polarisation in the system. The precession of the atoms in this regime is subject to statistical variation which depends on two factors, the external field B and the spin-exchange rate R. To achieve the highest possible polarisation, B must be kept as low as possible, which can be achieved using the on-board coils and external shielding (see section 2.3.3.1-3). Additionally, the spin-exchange rate must be kept high by ensuring a high vapour density. In the Rubidium-based OPMs used in Chapter 6 this is done by heating the vapour-cells to $\sim 150°C$ while in operation (Shah et al., 2007). As mentioned in 2.3.2.1, the outer surfaces of the

QuSpin OPM are not heated beyond body temperature, however, a large number of densely mounted OPMs without adequate dissipation of heat through the sensor helmet—for example—could cause discomfort. Appropriate design of the sensor helmet or mount is, therefore, crucial. Note that the temperature of operation cannot be decreased by increasing the cell's internal pressure or by adding a larger quantity of Rubidium as this could result in the formation of Rb droplets, effectively decreasing the number of atom-photon interactions and potentially disrupting the optical path through the transparent cell walls. OPMs based on different alkali metals such as Caesium (Petrenko et al., 2021) or He-4 OPMs (Labyt et al., 2019), which can operate at significantly lower temperatures ($\sim 100°C$ and room-temperature respectively), have shown some promise, however, further optimisation of the sensitivity and size of such devices is needed before their use for OPM-MEG is viable.

2.3.2.4 *The OPM signal*

The net polarisation **P** achieved in the glass cell through optical pumping is sensitive to the magnetic field and its behaviour can be described by the Bloch equations (Bloch, 1946):

$$\frac{\mathrm{d}\mathbf{P}}{\mathrm{dt}} = \frac{1}{\mathrm{Q}}\left[\mathbf{P} \times \gamma\mathbf{B} - (\mathbf{P} - \mathrm{P}_0\hat{\mathrm{z}})\frac{1}{\mathrm{T}}\right], \qquad (2.20)$$

where P_0 is the equilibrium polarisation achieved by pumping when $\mathrm{B} = 0$ and T is a time constant characterising the combination of relaxation and pumping rate of the system.

In the steady-state, where $\mathrm{d}\mathbf{P}/\mathrm{dt} = 0$, we can solve equation 2.20 which yields

$$
\begin{aligned}
P_x &= \frac{\frac{B_y}{\gamma T}+B_x B_z}{\left(\frac{1}{\gamma T}\right)^2+B_x^2+B_y^2+B_z^2} P_0; \\
P_y &= \frac{-\frac{B_x}{\gamma T}+B_y B_z}{\left(\frac{1}{\gamma T}\right)^2+B_x^2+B_y^2+B_z^2} P_0; \\
P_z &= \frac{\left(\frac{1}{\gamma T}\right)^2+B_z^2}{\left(\frac{1}{\gamma T}\right)^2+B_x^2+B_y^2+B_z^2} P_0.
\end{aligned}
\tag{2.21}
$$

If we measure the field in the x-direction, for example, the sensor coils minimise the field in the y and z directions, allowing a simplification of equation 2.21 by setting $B_y = B_z = 0$:

$$
\begin{aligned}
P_x &= 0, \\
P_y &= \frac{-\gamma T B_x}{1+(\gamma T B_x)^2} P_0, \\
P_z &= \frac{1}{1+(\gamma T B_x)^2}.
\end{aligned}
\tag{2.22}
$$

Figure 2.9 shows the polarisation in y and z due to changes in the field in x. Unfortunately, when measuring the polarisation along the laser axis (z), the sensing curve follows a Lorentzian shape (also referred to as zero-field or Hanle resonance (Hanle, 1925)) which does not give unique polarisation values for positive and negative fields of the same magnitude. If an external oscillating field is applied along the sensitive direction, this can be alleviated, however.

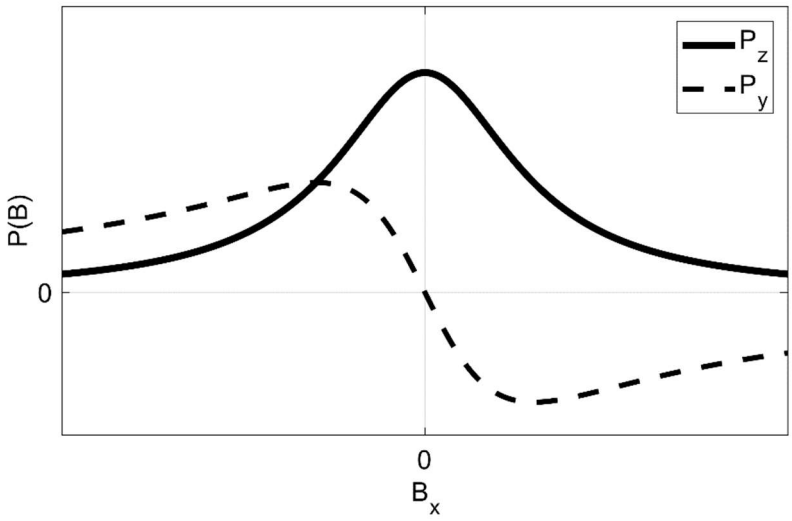

Figure 2.9 Polarisation of an optically pumped vapour with magnetic field strength applied in x. The sensitive direction along the laser axis displays a Lorentzian relationship while the polarisation along y shows a dispersion curve.

Assuming a modulation field along the x-axis of the form

$$B_x' = b_x \cos\omega_m t, \tag{2.23}$$

where b_x is the modulation field magnitude and ω_m is the modulation frequency (which is much larger than the Larmor frequency, in our case 923 rad. Hz^{-1}), the polarisation can be described by

$$\frac{d\mathbf{P}}{dt} = \frac{1}{Q}\left[\mathbf{P} \times \gamma[\mathbf{B} + (b_x \cos\omega_m t)\hat{\mathbf{x}}] - (\mathbf{P} - P_0\hat{\mathbf{z}})\frac{1}{T}\right]. \tag{2.24}$$

While the steady-state assumption cannot be applied here, a solution for P_z was found by Cohen-Tannoudji et al. (1970). In the zero-field case where B_y and B_z are nulled, the polarisation along the laser axis can be described by equation 2.25:

$$P_z = \frac{\gamma T B_x}{1 + (\gamma T B_x)^2} J_0\left(\frac{\gamma b_x}{\omega_m}\right) J_1\left(\frac{\gamma b_x}{\omega_m}\right) P_0 \sin\omega_m t. \tag{2.25}$$

The terms J_n are Bessel functions of the first kind, and evaluate to a constant. Using a lock-in amplifier, the demodulated signal V can therefore be described as

$$V(B_x) \propto \frac{\gamma T B_x}{1 + (\gamma T B_x)^2} \tag{2.26}$$

This dispersion relationship means that positive and negative deflections from the zero-field point can be distinguished. Additionally, the response function is approximately linear in the near-zero field regime as is shown in Figure 2.10.

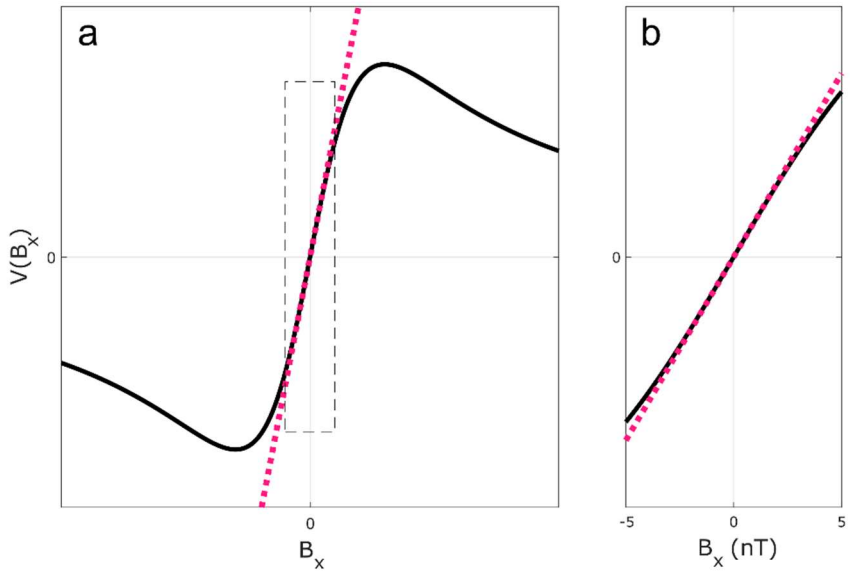

Figure 2.10 a) Plot of the dispersion relationship relating magnetic field with the voltage output for the QuSpin OPM. The magenta dotted line shows an approximate linear relationship in the limit of small B. b) inset of a showing the good correspondence of the dispersion curve with the linear approximation around $B = 0$.

2.3.3 Interference reduction in MEG recordings

We have so far established that neuromagnetic signals are extremely weak but measurable by using highly sensitive devices such as a SQUID or OPM. The signals of interest are, however, dwarfed by fields due to environmental or

biological sources of magnetic noise. As shown in Figure 2.11, typical brain signals are drastically smaller than the Earth's magnetic field, magnetic noise due to electronic equipment and signals generated by the heart.

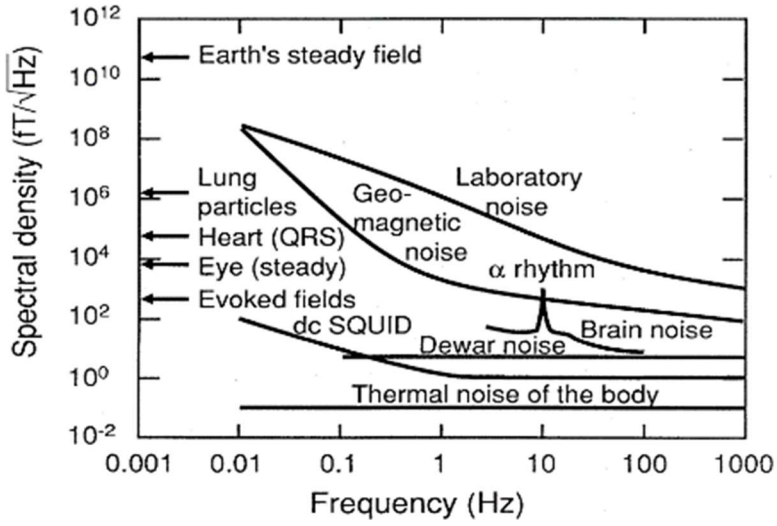

Figure 2.11 Typical amplitudes and spectral densities of environmental and biomagnetic noise (from Hamalainen et al., 1993).

This presents a considerable challenge, as SQUIDS, as well as OPMs, need to operate at very low fields. In the case of SQUIDs, the environmental noise could fluctuate so far beyond the lock point as to make measurements of flux ambiguous[2]. The zero-field OPMs used in the work presented here also require a magnetically quiet environment which cannot be achieved using only a static compensation field from the onboard coils. In this section, I will describe equipment and methodologies used to remove or reduce environmental and unwanted biomagnetic noise that is common in SQUID and OPM-MEG recordings, such as passive and active shielding, gradiometers, and independent component analysis (ICA) for artefact removal in software.

2 For commercial SQUID MEG systems this is an oversimplification. CTF MEG systems such as the ones used in the experimental work for this thesis employ electronic means to extend the working range of the SQUID sensors. This does however not remove the need for the interference techniques described in this section.

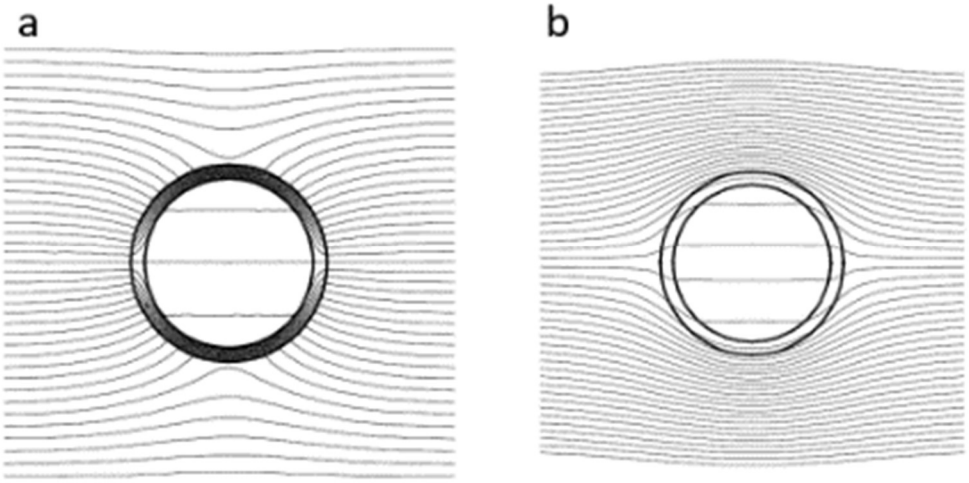

Figure 2.12 Shielding mechanisms. a) Flux shunting in high permeability materials to shield from DC fields (usually <10 Hz) b) Shielding due to eddy currents for high frequency fields. Both fields represent cross-sections from a cylindrical field. Adapted from (Celozzi et al., 2008)

2.3.3.1 *Magnetic shielding*

Maintaining a magnetically quiet environment has been crucial for the measurement of neuromagnetic signals since the very earliest MEG experiments (Cohen, 1968; Hamalainen et al., 1993). This can be achieved by exploiting the behaviour of magnetic fields at air-metal-boundaries. According to Ampere's and Gauss' laws, a static magnetic field must behave as shown in Figure 2.12a (Celozzi et al., 2008). Field lines must be perpendicular to the surface when entering into a material of higher relative permeability. For commonly used shielding materials with very high permeability such as mu-metal, fields can drastically change direction when entering the material and are 'guided' around the shielded area. This effect is only effective for quasi-constant fields, often referred to as DC fields which vary at frequencies of < 10 Hz. For high-frequency magnetic interference (AC fields) above 10 Hz, we can utilise the effect described by Faraday's law: in a conductor exposed to a time-varying magnetic field, an electric field is induced. The induced current due to this field will generate flux opposing the applied field leading to a shielding effect as shown

in Figure 2.12b. Generally, high conductivity materials such as aluminium are used to achieve effective shielding of AC fields. Magnetically shielded rooms (or MSRs) for MEG are usually large cuboids, consisting of several layers of mu-metal and aluminium, accessible via a door. However, some OPM-MEG setups have been reported that used human-size cylindrical shields (Borna et al., 2017).

2.3.3.1.1 Degaussing

The mu-metal used in the construction of MSRs is a ferromagnetic alloy and can therefore be magnetised, leading to a background field inside the MSR which is undesirable, especially when using zero-field OPMs. Degaussing, a process that can be used to reduce such remnant fields, was originally developed to demagnetise the hulls of ships that were vulnerable to magnetically activated sea mines (Kelly, 1946). A large, decaying sinusoidal oscillating current is applied to wires wrapped around the magnetised object, which gradually reduces the remnant magnetisation. In the work presented here, degaussing was only available in the OPM-based experiments. Ideally, degaussing should occur before each experiment or each time the door of the MSR is opened and closed.

2.3.3.2 *Gradiometry*

In addition to the shielding provided by the MSR, the inherent geometry of the sensor arrangement can be used to reduce sources of interference. As established in section 2.1.3, sources of neuromagnetic fields can be modelled as current dipoles, meaning that their magnetic field decays with distance (R) as $\frac{1}{R^2}$. We can deduce that their n-th order gradients, therefore, decay as $\frac{1}{R^{2+n}}$. If we assume a distant source of interference that can be modelled as a dipolar field, we can show that the gradient due to the nearby neural field is much larger than the gradient due to a distant source. Measurements of the field gradients instead of the fields themselves can, therefore, greatly reduce the influence of distant sources of interference on the captured signal. Gradients can be measured directly by adding counter-wound coils to the flux transformer to form

gradiometer pickups. Several such coil configurations exist, the most common being first-order axial and planar gradiometers as shown in Figure 2.13.

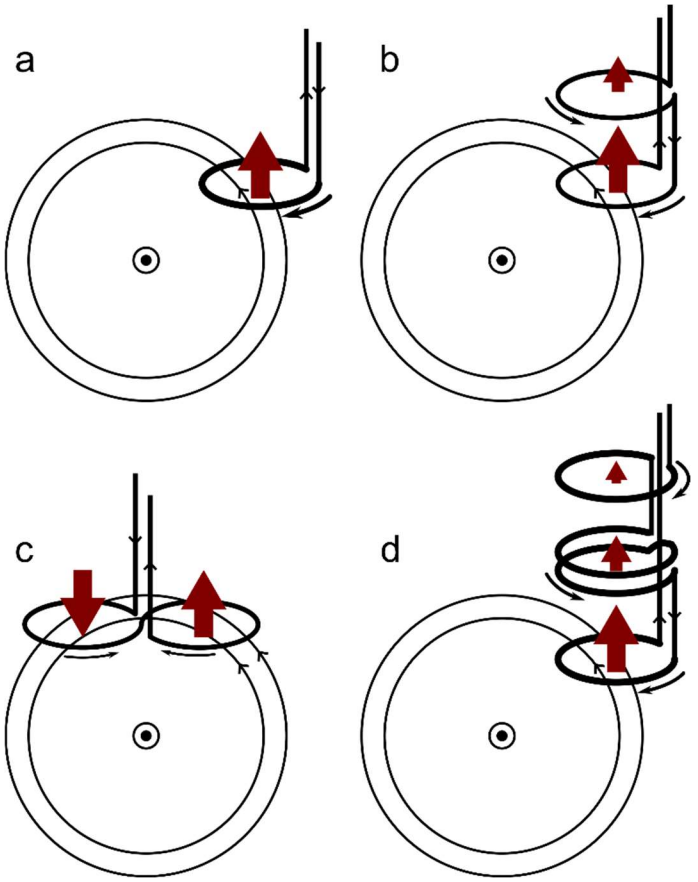

Figure 2.13 Pickup coil configurations. a) Magnetometer. b) First order axial gradiometer. c) Planar gradiometer d) Second order axial gradiometer. The red arrows depict the relative contributions of the dipolar field to the flux through each coil.

The CTF systems used in chapters 4 and 5 employ first-order axial gradiometers. Second-order gradiometers can be produced using arrangements such as the one shown in Figure 2.13d, however, their larger size reduces practicality. CTF systems, therefore, use an array of reference sensors distal to the pickup coils which measure three orthogonal field components and can be used to construct synthetic, 3rd order gradiometers in software.

Gradiometers can also be constructed in OPM arrays, by subtracting the signals from adjacent sensors, however, this was not done in any of the experimental work presented here.

2.3.3.3 *Active shielding for OPM arrays*

While passive shielding in combination with degaussing procedures and on-board nulling coils can produce an adequate environment for the operation of OPM sensors, problems arise when scanned subjects are allowed to move through any remnant field. Even small translations and rotations can lead to large artefacts in OPM-MEG data (Boto et al., 2018) or even prevent data collection altogether[3]. To reduce any remnant fields and field gradients, active shielding can be used. Holmes et al. showed that a set of bi-planar coils driven with fixed currents can be used to significantly reduce the background field strength and inhomogeneity (Holmes et al., 2019, 2018). This is achieved by measuring the field magnitude and gradients near the subject's head before applying currents to the coils, minimising any unwanted fields at the reference sensors. In the experiment discussed in Chapter 6, a set of 'fingerprint coils' (Holmes et al., 2019) was used. The generated fields superpose to minimise the remnant field near a set of reference sensors (see Figure 2.17) which are placed by the subject's head, again leading to a more favourable environment for OPM-MEG.

2.3.3.4 *Interference reduction using independent component analysis*

Environmental noise can be effectively reduced using the shielding techniques described above, however, signals from unwanted biomagnetic fields can still be present in the data. Signals due to eye movements and blinks (referred to as EOG artefacts after the electrooculogram) or the heartbeat (ECG artefacts after

3 In a conventional shielded room with background fields of tens of nT, even small rotations of less than 5 degrees would lead to measured field changes larger than the operational range of the OPMs used in the work presented here ($\pm$1.5nT). Additionally, any field inhomogeneity or gradients in the background field would lead to artefacts from sensor translations.

the electro-cardiogram) are not as effectively reduced as more distant external interference due to their relative proximity to the gradiometers. It can therefore be advantageous to remove these unwanted signals in software before proceeding with further analysis of MEG data. A commonly used method in MEG, as well as EEG data analysis, is temporal Independent Component Analysis (tICA) (Barbati et al., 2004; Jung et al., 2000; Rong and Contreras-Vidal, 2006; Vigário, 1997). ICA is a solution to the "cocktail party problem" which refers to the task of listening to a specific person in a crowded and noisy room. In the context of MEG, the goal of ICA is to find a set of independent time courses, weighted sums of which contribute to the total measured signal. Signals due to the heart and eyes can be assumed to have unique spatiotemporal characteristics and are represented by a small subset of components that can be removed. The data can be represented mathematically by an n×t matrix $\mathbf{X}$, where n is the number of sensors and t is the number of samples in the recording, where

$$\mathbf{X} = \mathbf{AS}. \tag{2.27}$$

Here, $\mathbf{S}$ is a $n_c \times t$ matrix, where $n_c \leq n$, the rows of which are the linearly independent components. $\mathbf{A}$ is an $n \times n_c$ mixing matrix, the columns of which represent the field patterns of the independent components. Inspection of these field patterns alongside the time courses of activation can be useful in determining which components correspond to which type of artefact. Figure 2.14 shows two example components extracted from MEG data recorded using a 275 channel CTF MEG scanner at the University of Nottingham. Artefacts can be categorised by inspecting the topographies (columns of $\mathbf{A}$) in conjunction with their corresponding time courses (rows of $\mathbf{S}$). We can see a typical artefact due to eye blinks in Figure 2.14a, with a characteristic topography showing large weights in sensors close to the eyes. Setting any row of $\boldsymbol{S}$ to zero before multiplying with the mixing matrix will lead to the removal of the

corresponding component. An example of this is shown in Figure 2.15, where a short segment of data is shown before and after removing an EOG artefact.

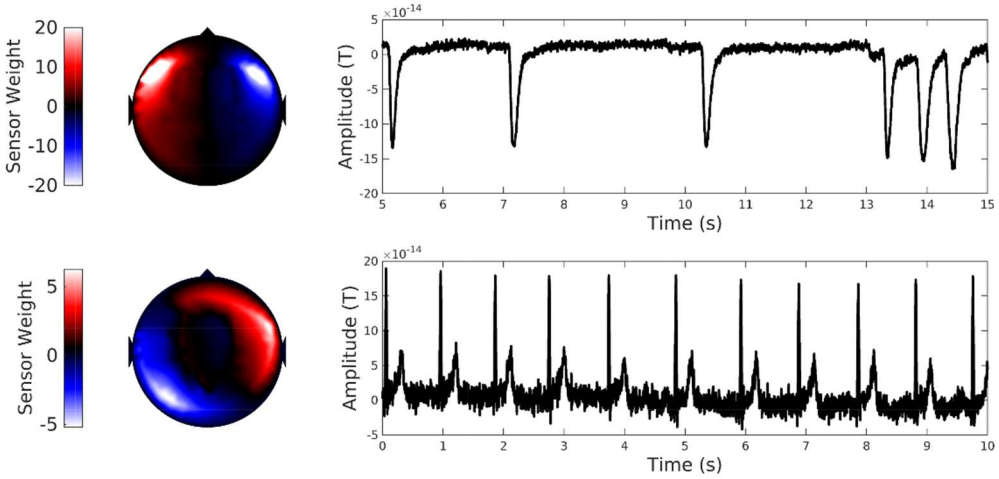

Figure 2.14 Example topographies and component time courses found using ICA on MEG data. The top row shows an EOG artefact. The component topography shows large weights near the eyes and the component time course contains short spiking events whenever a blink occurs. The bottom row shows an ECG artefact. The time course contains a typical ECG pattern with narrow spikes (QRS complex) followed by a broader peak (T wave).

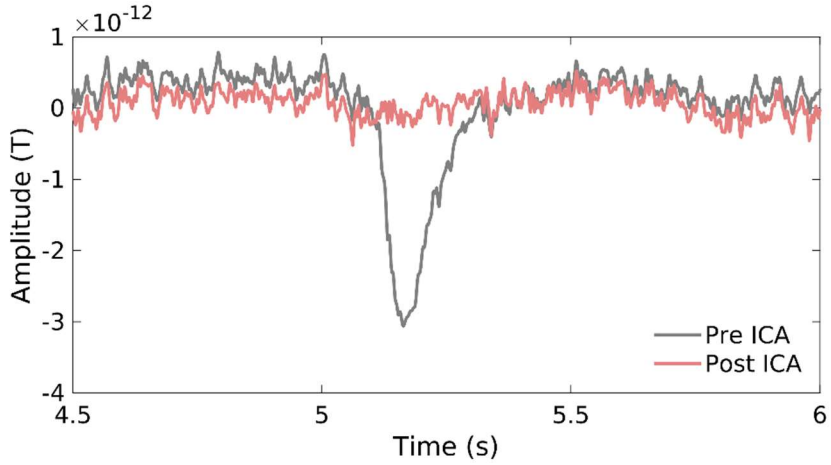

Figure 2.15 Segment of raw MEG data from a frontal sensor containing artefacts due to eye blinks (grey line) and clean data from the same sensor after removing EOG artefacts using ICA decomposition.

2.4 PERFORMING MEG EXPERIMENTS

We have so far discussed the origin of the MEG signal and the sensor and interference reduction technology needed to measure it. In this final section, I will outline how these elements are incorporated into the full set of instrumentation needed to perform MEG experiments and describe the general procedure followed when collecting data. Lastly, I will describe two methods of coregistration—or alignment—of the MEG sensors with the anatomy of the scanned subject.

Before a MEG experiment takes place, potential subjects are provided with all necessary information concerning the study so informed consent can be obtained. Subjects need to undergo safety screening to assess suitability for a MEG and MRI scan and are to remove all metal from their person. After providing any instructions or opportunities for practising a task performed during the acquisition, the subject can be prepared for the scan and positioned inside the scanner.

2.4.1 Instrumentation and acquisition

MEG scanner instrumentation contains a multitude of elements. As seen in section 2.3.3, MEG sensors need to be housed in the magnetically quiet environment provided by a magnetically shielded room. Waveguides in the room's walls are used to feed connections from the control room to the interior of the MSR to enable data acquisition and control of the sensors as well as stimulus equipment. Typically, a back-projection screen is used in conjunction with a projector outside the MSR which throws an image through a larger waveguide. Other types of peripheral equipment include speakers or MEG compatible earphones for auditory stimulation, eye tracking cameras and button boxes to allow patient responses during tasks. Figure 2.16 and Figure 2.17 show simplified schematics of SQUID-based and OPM-based MEG systems respectively.

2.4.1.1 *SQUID MEG setup*

The subject is positioned on the adjustable support in a sitting or supine position with their head inside a helmet-shaped cavity in the gantry (see Figure 2.16a). To limit head movement and ensure subject comfort, padding is placed around the head. To allow observation of the subject and communication with the experimenter, an intercom and video system connects the MSR interior with the experimenter in the control room. During the acquisition, a dedicated electronics rack handles the SQUID measurements and analogue to digital conversion (ADC) before transferring the sensor data to the acquisition PC for storage and real-time visualisation. The acquisition PC also receives input from peripherals such as button boxes and eye-tracking equipment. A separate stimulus PC controls the presentation and timing of stimuli and sends additional signals to the acquisition PC. The timing of stimuli can therefore be matched with corresponding MEG signals.

As shown in Figure 2.16b, the SQUID sensors are housed in a cryogenic Dewar filled with liquid helium. The use of Helium at a temperature of $\sim 4K$ bears the potential risk of cryogenic burns and asphyxiation due to the displacement of Oxygen by Helium gas in the unlikely case of a leak. Safety precautions include the continuous possibility of communication with the MEG operator via the intercom, Oxygen sensors with alarms measuring the 0_2 levels in the MSR and flow rate meters which would show abnormal evaporation of Helium from the Dewar.

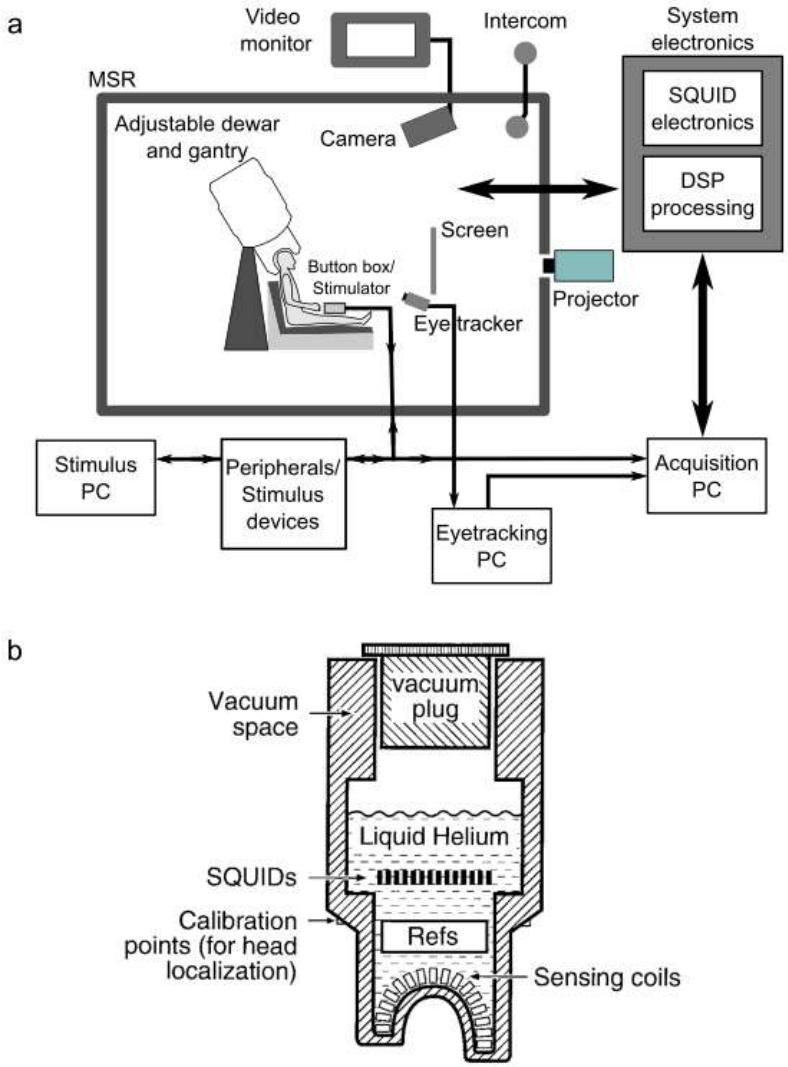

Figure 2.16 Instrumentation for SQUID-based MEG. a) Schematic showing an MSR housing a cryogenic MEG system. b) Interior layout of a MEG Dewar. Adapted from (Vrba and Robinson, 2001).

2.4.1.2 *OPM-MEG setup*

The subject is positioned in the centre of the MSR (sitting), between two nulling coils as shown in Figure 2.17. An additively manufactured helmet holding an array of OPM sensors is worn by the participant and a chin strap is used to hold the helmet in place. Padding can be used to maximise comfort and prevent relative motion between the head and helmet. In the system used for the study in Chapter 6, each OPM sensor is attached to the OPM control electronics via a

separate cable. To prevent the weight of these cables from exerting forces on the connections to the sensor, the cables are bundled and attached to shoulder straps. This allows the participant to carry the cable weight like a backpack and increases freedom of motion.

As mentioned in section 2.3.3.3, active magnetic field zeroing is used to reduce motion-related artefacts in OPM-MEG data. To achieve this, a set of reference sensors placed near the subject's head are used to sample the background field and field gradients. A nulling algorithm is then used to apply currents to a set of nulling coils that generate equal and opposite fields, destructively interfering with unwanted background fields. This, as well as data acquisition, is controlled by the acquisition PC. A separate stimulus PC is responsible for controlling the presentation of audio or visual stimuli, as well as the acquisition of optional motion tracking data. A set of infrared cameras mounted at 6 locations in the room can be used to record the positions of infrared-reflective markers placed on the helmet, or the subject's hand for example. This can be used to quantify head motion and to record subject interaction with the task, for example. The stimulus PC is also used to send trigger signals to the acquisition PC to provide temporal alignment of the stimuli with the recorded MEG data.

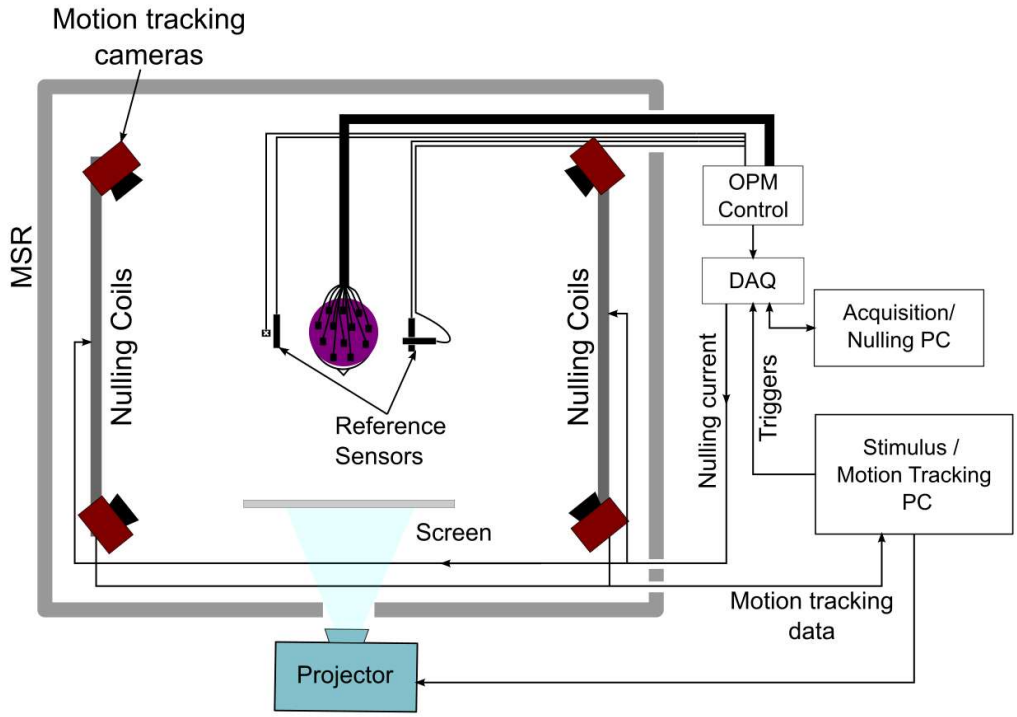

Figure 2.17 Schematic of the OPM-MEG system at the University of Nottingham.

2.4.1.3 *Quality control*

Several steps can be taken to ensure the highest possible MEG data. Before running an experiment, it is advisable to record at least one minute of data without any subjects in the room. Inspecting the frequency distribution of such "empty-room-noise" data can reveal any new or known sources of interference and ensures that all sensors are working. A series of tests can be conducted (Eye blinks, deep breath, jaw clench) by the scanned subject to ensure that artefacts are not excessively large. After recording, data should be inspected in small segments; segments containing large artefacts should be marked and removed. For the CTF-MEG used in this work, one common artefact is the SQUID reset[4] which presents as a large spike in the data. Segments, where head position exceeds a predetermined threshold of a few millimetres, should also be discarded. Any sensors containing excessive noise or any signals indicating faults should be excluded from analysis in both MEG modalities.

2.4.2 Coregistration

While the analysis of MEG signals at the sensor level can provide useful information about the timing and approximate location of neurological processes, it is preferable to reconstruct the signal sources within the brain. As will become clear in Chapter 3, such *source space* analysis requires knowledge of the positions of the MEG sensors relative to the head. In most cases, this is done by the coregistration of an MRI image of the head with the geometry of

4 The periodic sensitivity function of SQUIDs (equation 2.12 in section 2.3.1.2) inherently constrains the dynamic range of the sensors. If fluctuations in the signal exceed peak-to-peak amplitudes greater than half of the flux quantum, uniqueness of the measured flux values cannot be guaranteed. In CTF MEG scanners this limitation is circumvented via high frequency modulation of the flux, extending the range to $\pm\Phi_0$ relative to the lock point (Forgacs and Warnick, 1967). In the case where the signal drift exceeds magnitudes greater than Φ_0 a counter is incremented, the locking circuit is disengaged, then reengaged at a new operating point, thus extending the dynamic range further. Unfortunately, this resetting of the lock point causes a large transient artefact in the MEG signal.

the sensor array. In the work presented in this thesis, two distinct processes were used for coregistration, to accommodate the differing requirements of the cryogenic and wearable OPM-MEG systems.

2.4.2.1 *CTF MEG systems*

For data collected using the CTF MEG scanner, electromagnetic fiducial coils were attached to the nasion and the preauricular points of the scanned subject. During the scan, the coils are continuously energised which allows for the localisation of the coils within the scanner. To establish the location of the fiducial coils relative to an individual's anatomy, the subject's head/face shape including the relative position of the fiducial coils was obtained using either a Polhemus FASTRAK 3D digitiser or a 3-dimensional optical imaging system (Structure IO camera; Occipital Inc., San Francisco, CA, USA) in conjunction with the SKANECT software, which uses a combination of infrared depth sensing and visible light image acquisition (the latter was used during the COVID-19 pandemic to ensure a safe distance between subject and experimenter). The head shape was aligned to an MR image by extracting the scalp using edge detection techniques and using an automatic least-square-distance registration algorithm. This provided information about the relative position of the fiducial coils and the brain. In cases where an anatomical MRI is not available, a template MRI can be aligned to the scalp digitisation with adequate source localisation results (Gohel et al., 2017). Figure 2.18 shows how the digitisation of the head shape and fiducial coil locations (panels a and b) are aligned with the scalp extracted from the MRI (panel c), yielding the fiducial positions relative to the subject's anatomy (panel d). Note that head shapes generated using the Polhemus method are generally coarser than the one depicted in Figure 2.18a, usually consisting of 500-600 points.

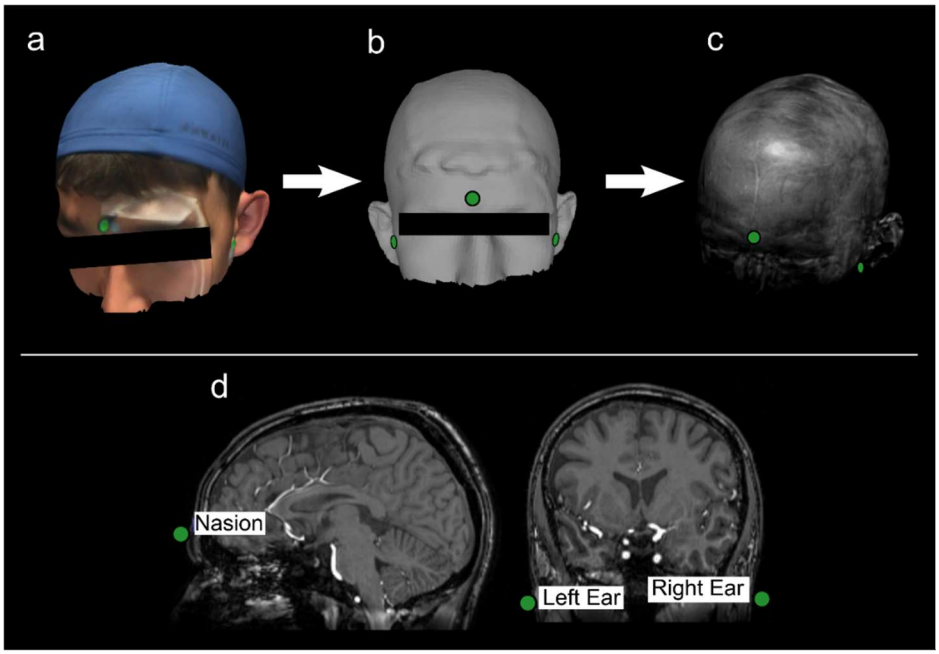

Figure 2.18 Coregistration procedure for conventional MEG. Three fiducial coils (green dots), which are enegised during the recording and can thus be localised by the MEG scanner, are attached to the nasion and preauricular points. a) A scan using a 3D optical imaging system, or head digitisation using a POLHEMUS digitiser is used to create a 3D representation of the head shape and the relative positions of the fiducial coils (b). This shape is coregistered to the scalp (c)—extracted from an anatomical MRI—using an iterative closest point algorithm, yielding the positions of the fiducial markers relative to the brain anatomy as shown in (d).

2.4.2.2 *OPM-MEG systems*

Wearable MEG has additional requirements to enable source localisation. While sensor positions and orientations are fixed and known within the conventional system from the manufacturing of the scanner, the modular nature of OPM-MEG means that the locations and orientations of each sensor need to be determined for each new configuration. To achieve this, an additively manufactured helmet was used to hold sensors at known relative positions and orientations (Hill et al., 2020). The coregistration of the sensors with the scanned individual's anatomy is done in several stages. First, digitisation of the head and face is achieved using the Structure IO camera while the subject is wearing a swimming cap (approximating the scalp surface). This "head-only"

digitisation is then aligned with the extracted scalp shape from an anatomical MR image using the MeshLab software (Cignoni et al., 2008). A second digitisation is generated while the subject is wearing the helmet containing the sensors. By matching prominent features of the subject's face (and/or markers placed on the face) in the "head-only" digitisation with this second "helmet-and-face" model, the helmet location relative to the anatomy can be established. A 3D model of the helmet can then be aligned with the "helmet-and-face" digitisation to obtain a final transformation between the anatomical MRI and the helmet coordinate systems which can be applied to the sensor locations and orientations which are known from the 3D printing process of the helmet. Figure 2.19 shows an example of the coregistration procedure for OPM-MEG performed for the investigation presented in Chapter 6.

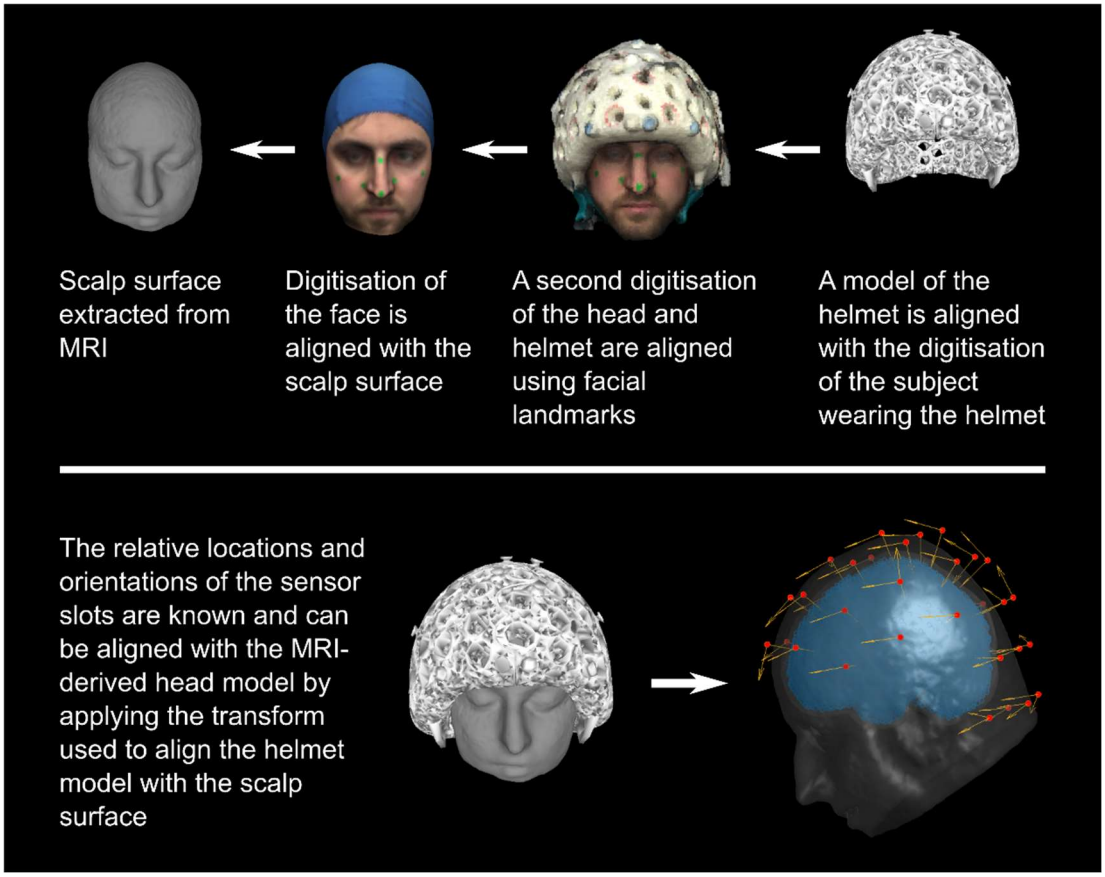

Figure 2.19 Coregistration procedure for OPM-MEG.

2.5 SUMMARY

In this chapter, I have outlined the origin of the MEG signal and how we can measure it even in the presence of interference that dwarfs neuromagnetic signals. In the next chapter, I will first outline the theoretical principles of source reconstruction, which allow us to generate three-dimensional maps of the neural currents underlying the MEG signal and subsequently introduce the mathematical framework of secondary analysis methods such as functional connectivity estimation and Hidden Markov Models.

Chapter 3 MEG SIGNAL ANALYSIS

The arrays of magnetic field sensors described in Chapter 2 allow us to sample and map the neuromagnetic fields near the surface of the head. While the resulting *sensor space* MEG data can yield useful information about the spectrotemporal makeup of electrical brain activity and the approximate area of activation (i.e. which hemisphere or lobe may contain the source of activity), a large number of sources may be producing fields at any point in time, a mixture of which will be measured by the sensors. Thus, it is preferable to estimate the current distribution inside the brain producing these fields—commonly referred to as source space reconstruction. Because we can assume sources to be dipolar, the approximate location of an active cortical source may be inferred from the observed field pattern. Given a dipolar source of known strength and orientation, the field pattern can be predicted by solving the *forward problem*. Unfortunately, the number, strength and orientations of sources can generally not be known. Additionally, constructive and destructive interference of magnetic fields means that any observed pattern could be generated by an infinite amount of magnetic source configurations. The task of estimating the current distribution in source space—termed the *inverse problem*—is, therefore, ill-posed. This chapter will outline the mathematical framework underlying the source space analysis of MEG data, describing solutions to the forward and inverse problems for neuromagnetism. Additionally, two secondary analysis methods applied to source space data will be discussed: functional connectivity analysis and Hidden Markov Model (HMM) analysis.

3.1 THE FORWARD PROBLEM

We can calculate the expected magnetic field from a cortical dipolar source by following an approach described by (Sarvas, 1987), where the head is modelled as a conductive volume G (bounded by the surface S) containing a source with a continuous current density **J**. The electromagnetic field generated by such a source is governed by Maxwell's equations:

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon}, \tag{3.1}$$

$$\nabla \cdot \mathbf{B} = 0, \tag{3.2}$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \tag{3.3}$$

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

where μ_0 and ϵ are the magnetic permeability and electric permittivity, **E** and **B** are the electric and magnetic fields respectively and **ρ** is the charge density. A quasi-static approximation can be made as the fields generated by the brain typically vary at frequencies less than 100 Hz. (Hamalainen et al., 1993) This means that the time derivative terms vanish, yielding

$$\nabla \times \mathbf{E} = 0; \tag{3.5}$$

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J}. \tag{3.6}$$

The field **B** at a location **r** outside a conductor, resulting from the current density **J** at a position $\mathbf{r}'$ inside the conductor, can be calculated using the Biot-Savart law:

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

where dv is a volume element of the volume G. As shown in Chapter 2, post-synaptic currents consist of the primary and volume currents; we can therefore write

$$\begin{aligned}\mathbf{J}(\mathbf{r}') &= \mathbf{J_P}(\mathbf{r}') + \mathbf{J_V}(\mathbf{r}') \\ &= \mathbf{J_P}(\mathbf{r}') - \sigma\boldsymbol{\nabla}\mathrm{V}(\mathbf{r}'),\end{aligned} \tag{3.8}$$

where $\mathbf{J_P}$ and $\mathbf{J_V}$ are the primary and volume current densities. We also see that the volume current can be expressed via Ohm's law in terms of the conductivity of the extracellular medium σ—which we assume to be constant over G—and the electric potential V. Substituting 3.8 into 3.7 and defining $\mathbf{a} = \mathbf{r} - \mathbf{r}'$, gives

$$\mathbf{B}(\mathbf{r}) = \frac{\boldsymbol{\mu_0}}{4\pi}\left[\int_{\mathrm{G}} \mathbf{J_P}(\mathbf{r}') \times \frac{\mathbf{a}}{|\mathbf{a}|^3}\,dv \; - \sigma \int_{\mathrm{G}} \boldsymbol{\nabla}\mathrm{V}(\mathbf{r}') \times \frac{\mathbf{a}}{|\mathbf{a}|^3}\,dv\right]. \tag{3.9}$$

Using the vector calculus identity $\nabla f \times \nabla g = \nabla \times (f\nabla g)$, and setting $f = \mathrm{V}(\mathbf{r}')$ and $\nabla g = \mathbf{a}/\,|\mathbf{a}|^3$ the integral term due to the volume current can be written as

$$\mathbf{B_V}(\mathbf{r}) = -\frac{\boldsymbol{\mu_0}}{4\pi}\left[\sigma \int_{\mathrm{G}} \boldsymbol{\nabla} \times \mathrm{V}(\mathbf{r}')\frac{\mathbf{a}}{|\mathbf{a}|^3}\,dv\right]. \tag{3.10}$$

From Stoke's theorem, $\mathbf{B_V}(\mathbf{r})$ can be expressed in terms of a surface integral:

$$\mathbf{B_V}(\mathbf{r}) = -\frac{\boldsymbol{\mu_0}}{4\pi}\left[\sigma \int_{\mathrm{S}} \hat{\mathbf{n}} \times \mathrm{V}(\mathbf{r}')\frac{\mathbf{a}}{|\mathbf{a}|^3}\,ds\right] \tag{3.11}$$

where $\hat{\mathbf{n}}$ is an outward-facing unit vector normal to the surface S at the surface element ds. We can therefore rewrite equation 3.9 using 3.11:

$$\mathbf{B}(\mathbf{r}) = \frac{\boldsymbol{\mu_0}}{4\pi}\left[\int_{\mathrm{G}} \mathbf{J_P}(\mathbf{r}') \times \frac{\mathbf{a}}{|\mathbf{a}|^3}\,dv \; - \sigma \int_{\mathrm{S}} \hat{\mathbf{n}}(\mathbf{r}') \times \mathrm{V}(\mathbf{r}')\frac{\mathbf{a}}{|\mathbf{a}|^3}\,ds\right]. \tag{3.12}$$

Equation 3.12 is a simplified form of the Geselowitz formula (Geselowitz, 1970). Multi-compartment models that take varying conductivities of the brain, skull, scalp etc. into account also exist. However, the true conductivities of the various types of tissue in vivo are generally not well known and a series of geometric simplifications can be made to yield a forward solution that is invariant to conductivity. The most basic approximation is the single-sphere model which is illustrated in Figure 3.1.

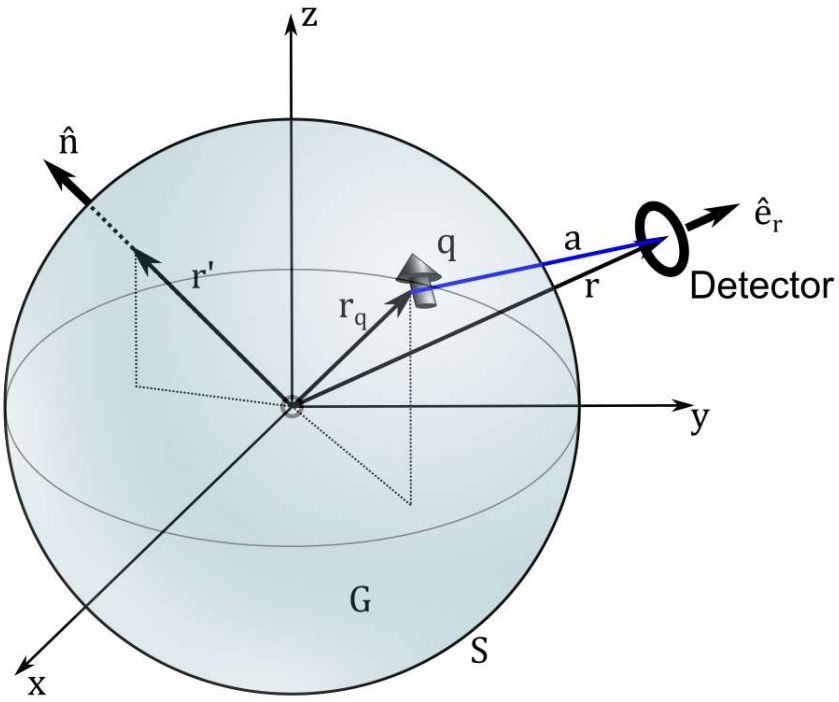

Figure 3.1 Single sphere model for the biomagnetic forward solution. The head is approximated as a spherical volume G bound by the surface S. A current dipole $\boldsymbol{q}$ *is found at the location* $\boldsymbol{r_q}$*. We consider fictitious currents at all points on the surface* $\boldsymbol{r'}$ *generating the magnetic signal in the detector at r.* $\hat{\boldsymbol{e}}_{\boldsymbol{r}}$ *is the radial unit vector and* $\hat{\boldsymbol{n}}$ *is the vector normal to the surface at the surface element ds.*

If we assume the head to be a spherical volume of uniform conductance, and that the detector only measures the radial component of the field, we can consider the radial component of $\mathbf{B}(\mathbf{r})$:

$$B_r(\boldsymbol{r}) = \boldsymbol{B}(\boldsymbol{r}) \cdot \hat{\boldsymbol{e}}_{\boldsymbol{r}}; \tag{3.13}$$

in full:

$$\mathrm{B_r}(\mathbf{r}) = \frac{\mu_0}{4\pi}\left[\int_{\mathrm{G}} \mathbf{J_P}(\mathbf{r}') \times \frac{\mathbf{a}}{|\mathbf{a}|^3} \cdot \hat{\boldsymbol{e}}_{\boldsymbol{r}} dv \; - \sigma \int_{\mathrm{S}} \hat{\mathbf{n}}(\mathbf{r}') \times \mathrm{V}(\mathbf{r}') \frac{\mathbf{a}}{|\mathbf{a}|^3} \cdot \hat{\boldsymbol{e}}_{\boldsymbol{r}} ds \right], \tag{3.14}$$

where $\hat{\boldsymbol{e}}_{\boldsymbol{r}}$ is the radial unit vector, in this case, normal to the sensitive plane of the MEG sensor. Expanding **a** and considering only the vector terms in the surface integral we can write

$$\begin{aligned}\hat{\mathbf{n}}(\mathbf{r}') \times \mathrm{V}(\mathbf{r}') \frac{\mathbf{a}}{|\mathbf{a}|^3} \cdot \hat{\mathbf{e}}_{\mathbf{r}} &\propto \left(\hat{\mathbf{n}}(\boldsymbol{r}') \times \boldsymbol{r} - \underbrace{\hat{\mathbf{n}}(\boldsymbol{r}') \times \boldsymbol{r}'}_{=\mathbf{0}\ as\ \hat{\mathbf{n}} \parallel \boldsymbol{r}'} \right) \cdot \hat{\mathbf{e}}_{\mathbf{r}} \\ &= \underbrace{\hat{\mathbf{n}}(\boldsymbol{r}') \times \boldsymbol{r}}_{\perp \hat{\mathbf{e}}_{\mathbf{r}}} \cdot \hat{\mathbf{e}}_{\mathbf{r}} \\ &= 0.\end{aligned} \tag{3.15}$$

This means that—for a spherical geometry—the volume current term vanishes, yielding

$$\mathrm{B_r}(\mathbf{r}) = \frac{\mu_0}{4\pi} \int_{\mathrm{G}} \mathbf{J_P}(\mathbf{r}') \times \frac{\mathbf{a}}{|\mathbf{a}|^3} \cdot \hat{\boldsymbol{e}}_{\boldsymbol{r}} dv\,, \tag{3.16}$$

which, crucially, does not depend on the conductivity of the medium.

So far, we have assumed that the primary current is distributed across the whole volume (i.e. at any position **r**′). However, we can approximate it as an equivalent current dipole **q** at a single point in the volume ($\mathbf{r_q}$). Using the Dirac delta function

$$\delta(x - x') = \begin{cases} +\infty, & x = x' \\ 0, & x \neq x' \end{cases} \tag{3.17}$$

we can simplify equation 3.16 to

$$\begin{aligned}\mathrm{B_r}(\mathbf{r}) &= \frac{\mu_0}{4\pi} \int_{\mathrm{G}} \boldsymbol{q} \delta\left(\mathbf{r}' - \mathbf{r_q}\right) \times \frac{\mathbf{a}}{|\mathbf{a}|^3} \cdot \hat{\boldsymbol{e}}_{\boldsymbol{r}} dv \\ &= \frac{\mu_0}{4\pi} \boldsymbol{q}(\mathbf{r_q}) \times \frac{\mathbf{a}}{|\mathbf{a}|^3} \cdot \hat{\boldsymbol{e}}_{\boldsymbol{r}}.\end{aligned} \tag{3.18}$$

To allow for arbitrary (i.e. non-radial) detector orientations we need to estimate the full magnetic field vector at **r**. Using the assumption that there are no currents outside the head ($\mathbf{J} = 0$ outside G) and $\boldsymbol{\nabla} \times \mathbf{B} = 0$, from equation 3.6 we can express $\mathbf{B}(\mathbf{r})$ in terms of the magnetic scalar potential U:

$$\mathbf{B}(\mathbf{r}) = -\mu_0 \nabla \mathrm{U}(\mathbf{r}). \tag{3.19}$$

For a conservative field[5] $\mathbf{B}$, the fundamental theorem of calculus for line integrals states that

$$-\frac{1}{\mu_0}\int_a^b \mathbf{B}\cdot dl = U(b) - U(a), \tag{3.20}$$

which means that the integral along any path between two points in the field is given by the difference in scalar potential. If we consider a line integral from $\mathbf{r}$ to infinity along a radial path we can substitute 3.18 into 3.20:

$$U(l=\infty) - U(l=0) = -\frac{1}{\mu_0}\int_0^\infty \mathbf{B}(l)\cdot\hat{\boldsymbol{e}}_r\cdot \mathrm{d}l$$

$$\underbrace{U(l=\infty)}_{=0} - U(l=0) = -\frac{1}{\mu_0}\int_0^\infty B_r(\boldsymbol{r}+l\hat{\boldsymbol{e}}_r)\cdot \mathrm{d}l$$

$$U(l=0) = U(\boldsymbol{r}) = \frac{1}{4\pi}\int_0^\infty \frac{\mathbf{q}\times\left(\boldsymbol{r}+l\hat{\boldsymbol{e}}_r-\boldsymbol{r}_q\right)\cdot\hat{\boldsymbol{e}}_r\cdot \mathrm{d}l}{\left|\boldsymbol{r}+l\hat{\boldsymbol{e}}_r-\boldsymbol{r}_q\right|^3}, \tag{3.21}$$

where dl is the line segment at position l relative to $\boldsymbol{r}$, and U is assumed to vanish at infinity. Given that $(\boldsymbol{q}\times l\hat{\boldsymbol{e}}_r)\cdot\hat{\boldsymbol{e}}_r = 0$ by definition, equation 3.21 can be simplified to

$$U(\boldsymbol{r}) = \tfrac{1}{4\pi}\mathbf{q}\times\left(\boldsymbol{r}-\boldsymbol{r}_q\right)\cdot\hat{\boldsymbol{e}}_r\int_0^\infty \frac{\mathrm{d}l}{\left|\boldsymbol{r}+l\hat{\boldsymbol{e}}_r-\boldsymbol{r}_q\right|^3}\,, \tag{3.22}$$

which yields

$$U(\boldsymbol{r}) = -\tfrac{1}{4\pi}\frac{\mathbf{q}\times\boldsymbol{r}_q\cdot\boldsymbol{r}}{F}\,, \tag{3.23}$$

with $F = |\mathbf{a}|(|\mathbf{r}||\mathbf{a}| + |\mathbf{r}|^2 - (\mathbf{r}_\mathrm{q}\cdot\mathbf{r}))$ and $\mathbf{a} = \mathbf{r}-\mathbf{r}_\mathrm{q}$. If we recall equation 3.19 and evaluate the gradient of U we get the solution to the single sphere model from (Sarvas, 1987):

5 This only applies outside the conductor volume where there are no charges and the curl of **B** is zero.

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi F^2}\left(F\mathbf{q} \times \mathbf{r}_q - \mathbf{q} \times \mathbf{r}_q \cdot \mathbf{r}\nabla F\right), \qquad (3.24)$$

with $\nabla F = \left(\frac{|\mathbf{a}|^2}{|\mathbf{r}|} + \frac{\mathbf{a}}{|\mathbf{a}|} \cdot \mathbf{r} + 2|\mathbf{a}| + 2|\mathbf{r}|\right)\mathbf{r} - \left(|\mathbf{a}| + 2|\mathbf{r}| + \frac{\mathbf{a}}{|\mathbf{a}|} \cdot \mathbf{r}\right)\mathbf{r}_q$.

While equation 3.24 is not adequate at modelling activity from brain regions that are not represented by spherical geometry (Hamalainen et al., 1993; Sarvas, 1987), this expression presents an easily computed, analytic solution to the forward problem that allows us to understand the principles of forward-modelling.

More advanced solutions exist that address some of the shortcomings of the single sphere model:

Boundary element methods (BEMs) estimate the forward field by discretising a subject's anatomy and usually treat the head as a set of nested compartments with different conductivities (usually 3 shells: brain, skull and scalp). BEM-based forward solutions most accurately take into account individual brain geometry and can therefore be seen as the most accurate forward models. In practice, however, accurate segmentation of the head anatomy and the numerical forward solution comes at a higher computational cost.

The multiple spheres model is an extension of the single sphere model developed by Huang et al. (1999). This head model is generated by defining a set of overlapping spheres which approximate the individual skull geometry local to each MEG sensor—usually extracted from and MRI. Applying the single sphere solution separately for each sphere and sensor can then be used to construct a full forward solution. Huang et al. showed that this method greatly improved the accuracy of the forward solution compared to a single sphere approach, with comparable performance to a BEM but greatly improved computational efficiency. For this reason, a multiple spheres model was used to produce the results in Chapter 4 and Chapter 5.

The single-shell model proposed by Nolte uses a correction term to the single sphere solution based on spherical harmonics to account for areas that deviate from spheroidal geometries. (Nolte, 2003) The coefficients for a fixed number of harmonics in the basis set are found by minimising the error in the boundary condition that the forward field is tangential to the surface of the conductor. Increasing the model order and, therefore, the number of harmonics in the correction increases the model accuracy but at the cost of increased computational complexity. Still, a commonly used implementation of the single-shell model (Oostenveld et al., 2011) has been shown to yield comparable results to a three shell BEM with a significantly reduced computation time (Stenroos et al., 2014).

3.2 THE INVERSE PROBLEM

The MEG signal measured at any point in time results from a summation of all the current dipole sources in the brain. We can formulate this statement mathematically via an integral over the brain volume V as shown in equation 3.25

$$\mathbf{b}(\mathrm{t}) = \int_V \mathbf{l}_{\boldsymbol{\theta}} \mathrm{q}_{\boldsymbol{\theta}}(\mathrm{t}) dV \,. \tag{3.25}$$

Here, $\mathbf{b}(\mathrm{t})$ is an $n\times1$ vector of the magnetic field measurements at the n sensors, $\mathbf{l}_{\boldsymbol{\theta}}$ is the $n\times1$ vector of fields that would be measured given a unit-strength current dipole at $\boldsymbol{\theta}$, which denotes the position and orientation vector (i.e $\left(\boldsymbol{r}_{\boldsymbol{q}}, \frac{\boldsymbol{q}}{|\boldsymbol{q}|}\right)$ in Figure 3.1) for each location in V, assuming that the dipole orientation at each position is fixed. $q_{\boldsymbol{\theta}}$ is the current dipole strength at $\boldsymbol{\theta}$ and time t.

In general, the brain volume is divided into a set of M discrete locations or voxels with the assumption that each voxel contains one dipolar source[6]. Equation 3.25 can be expressed as a sum over all voxels:

$$\mathbf{b}(t) = \sum_{m=1}^{M} \mathbf{l_\theta q_\theta}(t). \tag{3.26}$$

We can then write down the generative model of all sources in matrix form:

$$\mathbf{b}(t) = \mathbf{LQ}(t), \tag{3.27}$$

where **L** is an n×M matrix containing the forward solutions for all M locations and **Q** is a vector containing the M dipole strengths for all voxels. The inverse problem can now be formulated using the terms in equation 3.27:

*How can we estimate the current distribution **Q** in the brain given a set of magnetic field measurements **b**?*

A naïve approach to the inverse problem might be the solving of the matrix equation in 3.27, however, this would require an inversion of **L**. Given that the number of possible sources M greatly exceeds the number of sensors n, **L** is non-square. Moreover, field spread can lead to linear dependencies between the columns of **L**, making the inversion of **L** a non-trivial problem. There exist a variety of methods to circumvent this issue by introducing additional assumptions. The method used in the experimental work presented in chapters 4-6 is known as beamforming.

3.2.1 Beamforming

Initially developed for applications in communications, RADAR/SONAR, Geo- and Astrophysical exploration among others (Van Veen and Buckley, 1988),

[6] This simplified case can be expanded by assuming three orthogonal dipolar sources at each location such that arbitrary dipole orientations can be modelled.

beamforming is a popular spatial filtering technique for EEG and MEG source localisation (Robinson and Vrba, 1999; Van Veen et al., 1997). The goal of spatial filtering is to reconstruct a signal at a pre-specified location and orientation $\boldsymbol{\theta}$ while rejecting signals that originate elsewhere. This is illustrated in Figure 3.2: the brain volume of a subject is discretised into voxels, one of which is chosen as the source location of interest. A set of weights w_n is chosen such that a weighted sum of the sensor signals gives an estimate $\widehat{\mathbf{q}}_{\boldsymbol{\theta}}$ of the true signal at $\boldsymbol{\theta}$. Mathematically,

$$\begin{aligned}\widehat{\boldsymbol{q}}_{\boldsymbol{\theta}}(t) &= w_{\boldsymbol{\theta}1}b_1(t) + w_{\boldsymbol{\theta}2}b_2(t) + \cdots + w_{\boldsymbol{\theta}N}b_N(t) \\ &= \boldsymbol{w}_{\boldsymbol{\theta}}^{\boldsymbol{T}}\boldsymbol{b}(\boldsymbol{t}),\end{aligned} \tag{3.28}$$

which is the general formulation of a spatial filter. Beamforming uses variance (or power) minimisation to determine the filter $\boldsymbol{w}_{\boldsymbol{\theta}}$, with the constraint that the power at $\boldsymbol{\theta}$ is unchanged. Each filter, therefore, attenuates all signals from unwanted locations while preserving the signal of interest. This constrained minimisation problem is formulated in equation 3.29:

$$\min_{\mathbf{w}_{\boldsymbol{\theta}}}(\widehat{\sigma}^2), \text{ subject to } \mathbf{w}_{\boldsymbol{\theta}}^{\mathrm{T}}\mathbf{l}_{\boldsymbol{\theta}} = 1 \tag{3.29}$$

where

$$\widehat{\sigma}^2 = \langle \widehat{\boldsymbol{q}}_{\boldsymbol{\theta}}^2(t) \rangle. \tag{3.30}$$

Here, $\widehat{\sigma}^2$ is the variance of the estimated source signal $\widehat{\boldsymbol{q}}_{\boldsymbol{\theta}}$. Substituting equation 3.28 into 3.30 we obtain

$$\begin{aligned}\widehat{\sigma}^2 &= \left\langle \left(\mathbf{w}_{\boldsymbol{\theta}}^{\mathrm{T}}\mathbf{b}(\mathrm{t})\right)\left(\mathbf{w}_{\boldsymbol{\theta}}^{\mathrm{T}}\mathbf{b}(\mathrm{t})\right)^{\mathrm{T}} \right\rangle \\ &= \mathbf{w}_{\boldsymbol{\theta}}^{\mathrm{T}}\langle \mathbf{b}(\mathrm{t})\mathbf{b}(\mathrm{t})^{\mathrm{T}} \rangle \mathbf{w}_{\boldsymbol{\theta}} \\ &= \mathbf{w}_{\boldsymbol{\theta}}^{\mathrm{T}}\mathbf{C}\mathbf{w}_{\boldsymbol{\theta}}\end{aligned} \tag{3.31}$$

where $\mathbf{C}$ is the covariance matrix of the sensor data $\mathbf{b}$.

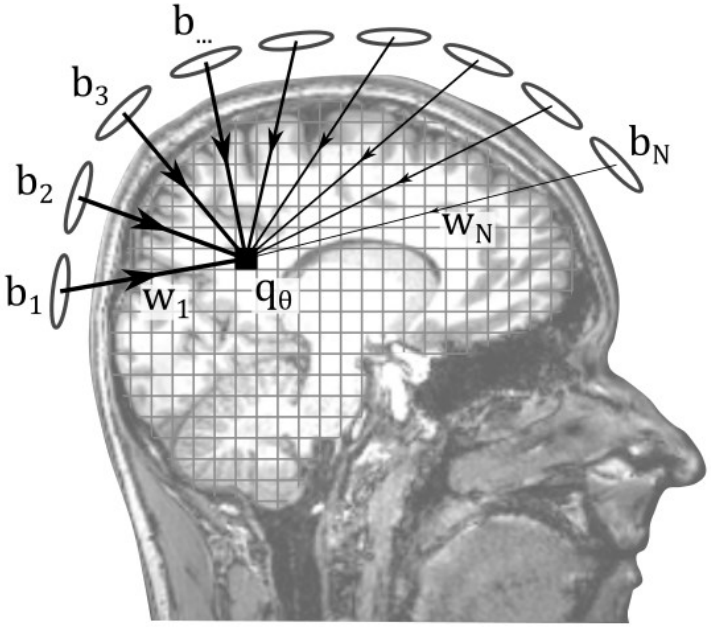

Figure 3.2 Spatial Filtering. For each location q_θ , a set of weights $[w_1, w_2, \ldots w_N]$ are chosen such that a weighted sum of the signals measured by the N sensors gives an estimate of the signal at a pre-selected location. In this schematic, weights are represented by the arrows, the thickness of which represents the size of the weight. As an example, sensors close to the chosen source location have larger weights.

The constrained minimisation problem in equation 3.29 can be solved using a Lagrange multiplier. To find the minimum of a function $f(x)$ with an additional constraint $g(x) = 0$, we can define the Lagrangian

$$\mathcal{L}(x, \lambda) = f(x) + \lambda g(x), \tag{3.32}$$

where λ is the Lagrange multiplier. If we let $f(x) = \hat{\sigma}^2$, and

$$g(x) = 0 = \mathbf{w}_\theta^{\mathrm{T}} \mathbf{l}_\theta - 1 \tag{3.33}$$

given the constraint from 3.29, we get

$$\mathcal{L}(\mathbf{w}_\theta, \lambda) = \mathbf{w}_\theta^{\mathrm{T}} \mathbf{C} \mathbf{w}_\theta + \lambda\left(\mathbf{w}_\theta^{\mathrm{T}} \mathbf{l}_\theta - 1\right). \tag{3.34}$$

The minimum can then be found by solving

$$\frac{\partial \mathcal{L}(\mathbf{w_\theta}, \lambda)}{\partial \mathbf{w_\theta}} = 0$$

$$\frac{\partial \mathcal{L}(\mathbf{w_\theta}, \lambda)}{\partial \mathbf{w_\theta}} = 2\mathbf{C}\mathbf{w_\theta} + \lambda \mathbf{l_\theta} = 0 \tag{3.35}$$

Equating 3.35 with the constraint in 3.33 we obtain the Lagrange multiplier

$$\lambda = -\frac{2}{\mathbf{l_\theta^T} \mathbf{C}^{-1} \mathbf{l_\theta}}. \tag{3.36}$$

Rearranging 3.35 to

$$\mathbf{w_\theta} = -\frac{\lambda \mathbf{C}^{-1} \mathbf{l_\theta}}{2} \tag{3.37}$$

and substituting 3.36 we obtain the solution for the beamformer weights:

$$\mathbf{w_\theta^T} = \frac{\mathbf{l_\theta^T} \mathbf{C}^{-1}}{\mathbf{l_\theta^T} \mathbf{C}^{-1} \mathbf{l_\theta}}. \tag{3.38}$$

Issues arise if the sensor covariance matrix $\mathbf{C}$ is singular or close to singular. Singular matrices have a zero determinant leading to instabilities in the inverse. In MEG, a common method to alleviate this is Tikhonov regularisation (e.g. Brookes et al., 2008). The regularised covariance matrix $\mathbf{C_r}$ is defined as

$$\mathbf{C}_r = \mathbf{C} + \mu \mathbf{I}, \tag{3.39}$$

where $\mathbf{I}$ is the identity matrix and μ is a positive, scalar regularisation parameter. Tikhonov regularisation adds μ to all elements of the diagonal, ensuring that none of the eigenvalues is close to zero. Typically, μ is chosen to be a small fraction of the maximum singular value of $\mathbf{C}$.

3.2.1.1 *Estimating source orientation*

In practice, the assumption that both the source location and orientations are known is not valid. While the possible source locations are pre-determined by the discretisation of the brain volume, the orientations are usually unknown.

Based on the principle that the ideal source orientation is the one yielding the highest SNR, there are two approaches to determining source orientation:

3.2.1.1.1 Exhaustive search

Forward modelling has shown that MEG signals likely arise from sources tangential to the surface of the brain with little to no contribution from radial sources. We can thus constrain the possible source orientation to the tangential plane. By modelling multiple dipoles with a range of azimuthal angles ϕ between 0 and π, we can estimate the corresponding spatial filters $\mathbf{w}_\phi$ and calculate the SNR using a pseudo-Z statistic $\mathcal{Z}_\phi$:

$$\mathcal{Z}_{\boldsymbol{\phi}} = \frac{\mathbf{w}_{\boldsymbol{\phi}}^{\mathrm{T}} \mathbf{C} \mathbf{w}_{\boldsymbol{\phi}}}{\mathbf{w}_{\boldsymbol{\phi}}^{\mathrm{T}} \sigma_{\mathrm{n}}^2 \mathbf{I} \mathbf{w}_{\boldsymbol{\phi}}}, \tag{3.40}$$

where σ_{n}^2 is the variance of the uncorrelated noise, assumed to be the same for each sensor. The final set of chosen weights are those corresponding to the maximum value of $\mathcal{Z}_{\boldsymbol{\phi}}$.

3.2.1.1.2 Singular value decomposition

A second method to determine optimal source orientation was proposed by Sekihara et al. (Sekihara et al., 2004). Beamformer weights and the corresponding source time courses are calculated for three orthogonally oriented dipoles. Singular value decomposition of the $3 \times N_t$ matrix containing the three time courses with N_t time points yields a set of three orthogonal time courses or principal components. The component corresponding to the largest singular value and therefore the direction of maximum variance is equivalent to $\hat{\mathbf{q}}$ at the optimum orientation.

3.2.1.2 *Beamforming in practice*

Recalling Section 2.2, the various types of activity observed require different types of source statistics. For intrinsic signals, information about source power and the locations of maximum activation can be found by generating a map of power values $\hat{\sigma}^2$ using equation 3.31 ($\hat{\sigma}^2 = \mathbf{w}^T\mathbf{C}\mathbf{w}$). One issue with this approach is a bias towards the centre of the brain introduced by the beamformer. As SNR decreases with distance to the sensor, weights for deeper sources tend to be increased to compensate for this reduction in signal leading to an increased source variance (Hall et al., 2013). Bias correction techniques usually involve the scaling of weights by their norm:

$$\mathbf{w}_{\text{corrected}} = \frac{\mathbf{w}}{\|\mathbf{w}\|}. \tag{3.41}$$

The resulting source maps are related to the pseudo-Z statistic (equation 3.40)

$$\mathcal{Z} = \frac{\mathbf{w}^{\mathrm{T}}\mathbf{C}\mathbf{w}}{\mathbf{w}^{\mathrm{T}}\mathbf{\Sigma}\mathbf{w}}, \tag{3.42}$$

where $\mathbf{\Sigma}$ is a diagonal matrix that determines the type of scaling and can be set to the identity matrix in the most basic case. Beamformer images can reveal patterns of activation for a given condition, however, care must be given when comparing scaled source maps from multiple recordings. Luckhoo et al. showed that the scaling of weights can introduce strong confounds that may reverse the sign of an effect when comparing multi-session recordings in different conditions (Luckhoo et al., 2014). This is of great importance when comparing the intrinsic activity of two different subject cohorts like a patient and control group for example.

When studying stimulus-related, induced activity, it is useful to compare an active condition with a rest or control condition to obtain the brain regions

which are relevant to the stimulus. A common contrast used to achieve this is the pseudo-T statistic, defined as

$$\mathcal{T}_{\theta} = \frac{\mathbf{w}_{\theta}^{\mathrm{T}}\mathbf{C}_{\mathbf{a}}\mathbf{w}_{\theta} - \mathbf{w}_{\theta}^{\mathrm{T}}\mathbf{C}_{\mathbf{c}}\mathbf{w}_{\theta}}{\mathbf{2}(\mathbf{w}_{\theta}^{\mathrm{T}}\mathbf{C}_{\mathbf{c}}\mathbf{w}_{\theta})}, \tag{3.43}$$

where $\mathbf{C_a}$ and $\mathbf{C_c}$ are the data covariance matrices for the active and control periods within a single scan.

Source maps, generated using a Z- or T-statistic, provide spatial information about active regions in the brain but do not fully exploit the temporal resolution of MEG. To extract the full wealth of spectro-temporal information of the MEG signal virtual electrodes, or VEs can be constructed for any location of interest $\boldsymbol{\theta}$:

$$\mathrm{VE} = \hat{\mathbf{q}}\,(\mathrm{t}) = \mathbf{w}^{\mathrm{T}}\mathbf{b}(\mathrm{t}). \tag{3.44}$$

Fourier and time-frequency analyses of VEs give insight into the spectral make-up of the MEG signal and allow us to assess task-based changes in the whole spectrum of brain activity. In the following sections, I will outline the basic principles of two additional classes of analysis techniques that can be applied to source reconstructed time courses and that were used in the research presented in this thesis: Functional connectivity analysis and Hidden Markov Modelling.

3.3 FUNCTIONAL CONNECTIVITY

As discussed in Chapter 2 (section 2.2.4), brain regions that are separated in space but exhibit a statistical dependency between their activation patterns are said to be functionally connected. Much of the early literature concerning functional connectivity was not based on electrophysiology but relied on functional MRI. Here, FC became synonymous with the correlation between the BOLD signal time courses extracted from separate brain regions. Biswal and colleagues reported that low-frequency fluctuations in the BOLD signal from the motor cortices were correlated within and between the hemispheres in the

absence of a task, providing initial evidence that the formation of functional networks is a phenomenon intrinsic to brain activity (Biswal et al., 1995). Since then, several networks related to sensory processing and cognitive functioning have been described and shown to be highly consistent across subjects (e.g. visual, motor and the default mode networks; see van den Heuvel and Hulshoff Pol (2010) for a review). The spatial congruity of these fMRI-derived functional networks with the distinctive patterns observed when mapping the power of band-limited electrophysiological oscillations (e.g. (Brookes et al., 2011a, 2011b); see the review by Hall et al. (2014)) inspired further research into the statistical relationships between MEG signals from distal brain regions.

Given that the BOLD signal is an indirect measure of neuronal function and is limited by the time scale of the BOLD effect which is on the order of seconds, the high sampling rates of electrophysiological measurements can offer additional insights into the dynamics of functional integration between brain regions. The wealth of information contained in the MEG signal, however, does not limit us to a definition of a single measure of FC.

3.3.1 Overview of connectivity measures

Unlike for BOLD fMRI, no simple definition of FC exists for MEG-derived signals. Instead, oscillations limited to the *classical* frequency bands are commonly considered, allowing for within-band (see Bastos and Schoffelen (2016) for a review) and between-band FC measures (for example coupling between theta and gamma oscillations (Jensen and Colgin, 2007)). Most FC measures quantify either the amplitude to amplitude or phase to phase coupling of bandlimited signals (see Figure 3.1 for an illustration). However, phase-amplitude coupling measures are available (e.g. Seymour et al., 2017).

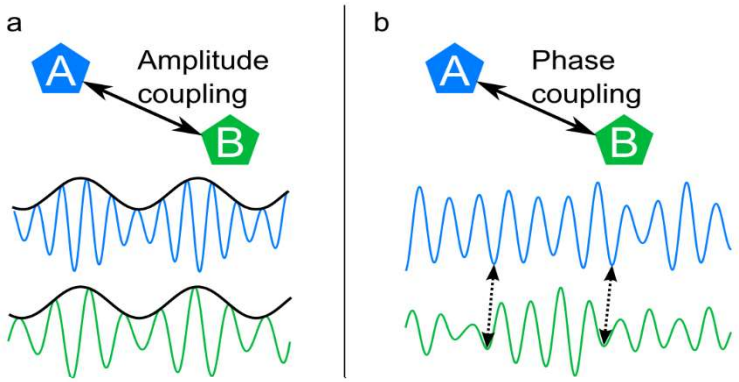

Figure 3.1 Illustration of connectivity measures. a) Two signals from sources A and B exhibit amplitude coupling if their envelopes (black lines) are correlated over time. Note that the signals need not be of the same frequency or frequency band. b) Two signals from sources A and B are phase coupled if their phase difference is constant. In this case there is a phase difference of π/2 between the signals.

In most cases, source reconstructed signals are filtered to a frequency band of choice and the amplitude and/or phase information are extracted (see Figure 3.2). Given the methods used in the experimental part of this thesis, I will focus on within-band FC measures.

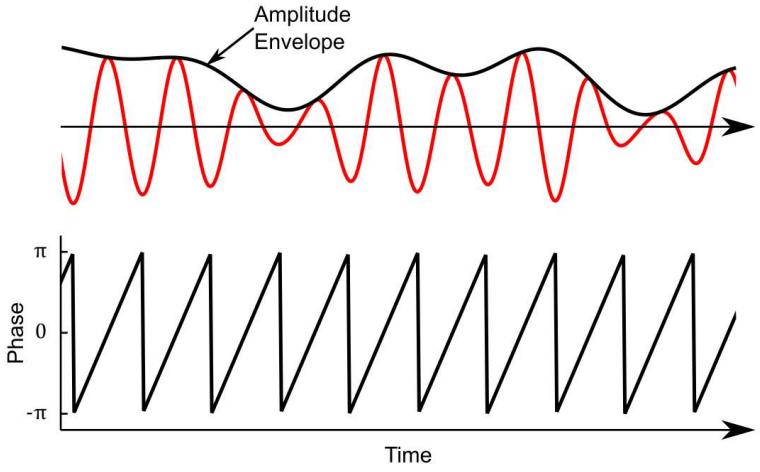

Figure 3.2 Illustration of the amplitude (top) and phase (bottom) of a signal (red) evolving in time.

Amplitude coupling measures can be calculated using the Pearson correlation between the amplitudes of two signals—commonly referred to as amplitude envelope correlation (AEC; O'Neill et al., 2015b). For AEC, spurious connectivity

can arise due to spatial leakage. It is advised to apply corrections to reduce this effect (Sadaghiani et al., 2021); a popular technique to achieve this is described in section 3.3.2.

Phase coupling techniques generally quantify the stability of the difference between the phases of two signals over time. Widely used measures are coherence, the phase difference derivative and the phase-locking value (Lachaux et al., 1999), phase lag index (Stam et al., 2007) and imaginary coherence (Nolte et al., 2004). While leakage also should be taken into account for phase-based measures, phase lag index and imaginary coherence are less affected by leakage as they remove instantaneous interactions between signals by default.

Various methods exist to estimate the amplitude and phase evolution of a signal (e.g. Hilbert transform and wavelet transform). In the work presented here, a method based on the Hilbert transform was used for its popularity and ease of implementation.

Equation 3.45 is an expression for the complex analytical signal $\hat{z}(t)$ derived from the signal vector $\hat{q}(t)$, where H is the Hilbert transform:

$$\hat{z}(t) = \hat{q}(t) + iH[\hat{q}(t)]. \tag{3.45}$$

H is a convolution operation on the signal which can be expressed as shown in equation 3.46:

$$H[\hat{q}(t)] = P\left[\frac{1}{\pi}\int_{-\infty}^{-\infty}\frac{\hat{q}(\tau)}{\mathrm{t}-\tau}d\tau\right], \tag{3.46}$$

where $P[\,]$ denotes the Cauchy principal value of the integral. This avoids any issues caused by the singularity which arises when $t = \tau$.

The complex analytical signal contains the amplitude and phase information at each time point which can be calculated as shown in equations 3.47 and 3.48 respectively:

$$\text{Envelope} = ||\hat{z}|| = \sqrt{\left(\hat{q}(t)\right)^2 + (H[\hat{q}(t)])^2} \tag{3.47}$$

$$\phi(t) = \tan^{-1}\left(\frac{H[\hat{q}(t)]}{\hat{q}(t)}\right) \tag{3.48}$$

3.3.2 Leakage correction

The ill-posed nature of MEG source reconstruction has the undesirable consequence that time-courses in source space are not independent, leading to spatial leakage and spurious connectivity for a variety of FC measures (Schoffelen and Gross, 2009). Several methods have been developed to correct for this effect before connectivity measures are calculated (Brookes et al., 2012; Colclough et al., 2015; Hipp et al., 2012; O'Neill et al., 2015a). In general, most methods aim to remove any leakage confound from the data by assuming that it presents as interaction with no time lag between sources; true connectivity is assumed to involve a time delay (although we know from invasive recordings that this is not true in all cases (Vicente et al., 2008)). For the work reported in Chapters 4 and 5, symmetric orthogonalisation (Löwdin, 1950) was used to remove the effect of leakage. Introduced by Colclough et al. (2015) for application to MEG connectivity analyses, this technique aims to adjust the full set of reconstructed time-courses, usually corresponding to a set of regions in a brain atlas parcellation, to achieve orthogonality between all of them. Figure 3.3 shows a 2D analogue where two vectors (black) are rotated by the same amount to achieve a right angle between them (blue).

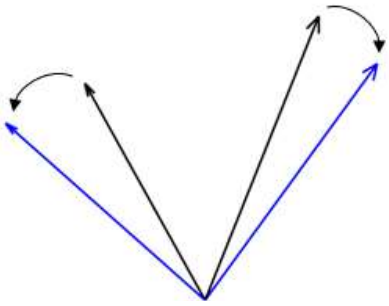

Figure 3.3 Illustration of leakage correction via symmetric orthogonalisation in 2D. Two vectors (black) representing the time-courses to be corrected are adjusted by the same angle to result in orthogonal vectors(blue).

Mathematically, this leakage correction method is based on singular value decomposition. An $s \times T$ matrix $\boldsymbol{D}$ of virtual sensor time courses, where s is the number of virtual sensors and T is the number of time points can be factorized to obtain

$$\boldsymbol{D} = \boldsymbol{U}\boldsymbol{S}\boldsymbol{V}^T, \tag{3.49}$$

where $\boldsymbol{U}$ and $\boldsymbol{V}$ are matrices containing the eigenvectors of $\boldsymbol{D}\boldsymbol{D}^T$ and $\boldsymbol{D}^T\boldsymbol{D}$ respectively, while $\mathbf{S}$ contains the square root of the eigenvalues shared between $\boldsymbol{U}$ and $\boldsymbol{V}$. The matrix of leakage corrected data $\widehat{\boldsymbol{D}}$ can be calculated using

$$\widehat{\boldsymbol{D}} = \boldsymbol{U}\boldsymbol{V}^T. \tag{3.50}$$

One limitation of this method is that the number of simultaneous time courses that can be orthogonalised must be less than or equal to the rank of the data. For MEG data this depends on the number of sensors and whether any dimensionality reduction techniques were used during pre-processing. As an example, raw MEG data will generally have a rank equal to the number of sensors, however, removing components using ICA could reduce the rank. The brain parcellations used in this work had a region count (78) much smaller than the number of sensors in the conventional MEG systems (~270). The relatively smaller array of OPM sensors used in Chapter 6, however, would not provide enough information to correct for leakage using symmetric orthogonalisation

whilst using the same parcellation. To overcome this, we opted to use a pairwise orthogonalisation (Brookes et al., 2012). Briefly, if two time courses $\boldsymbol{y}$ and $\boldsymbol{x}$, reconstructed for a seed ($\boldsymbol{y}$) and test ($\boldsymbol{x}$) location in the brain, are used to calculate a measure of functional connectivity, any leakage from the test location towards the seed location will be reflected as a zero-lag contribution to the seed signal. This contribution can be expressed as a linear projection of the signal in $\boldsymbol{x}$ onto the signal $\boldsymbol{y}$:

$$\beta = \boldsymbol{x}^{\dagger}\boldsymbol{y}, \tag{3.51}$$

where $\boldsymbol{x}^{\dagger}$ denotes the pseudo-inverse of $\boldsymbol{x}$. The leakage corrected seed time course $\boldsymbol{y}'$ can then be estimated using

$$\boldsymbol{y}' = \boldsymbol{y} - \boldsymbol{x}\beta. \tag{3.52}$$

Connectivity measures can then be calculated using $\boldsymbol{x}$ and $\boldsymbol{y}'$.

3.4 HIDDEN MARKOV MODEL ANALYSIS

3.4.1 From oscillations to transient events

Conventional analyses of the spectral properties of MEG signals typically involve Fourier-based estimation of frequency content—e.g. short-time Fourier or wavelet transforms—using the assumption of stationarity over a time window of interest, and averaging over time or trials. Recent work exploring non-averaged signals, however, found that this approach might be flawed as it may misrepresent the true dynamics underlying the estimated spectra (Jones, 2016; Little et al., 2019; Quinn et al., 2019; Sherman et al., 2016; van Ede et al., 2018; Zich et al., 2020). An example is shown in Figure 3.4, where smoothly varying beta activity appears to be the result of short bursts of high amplitude at the single-trial level.

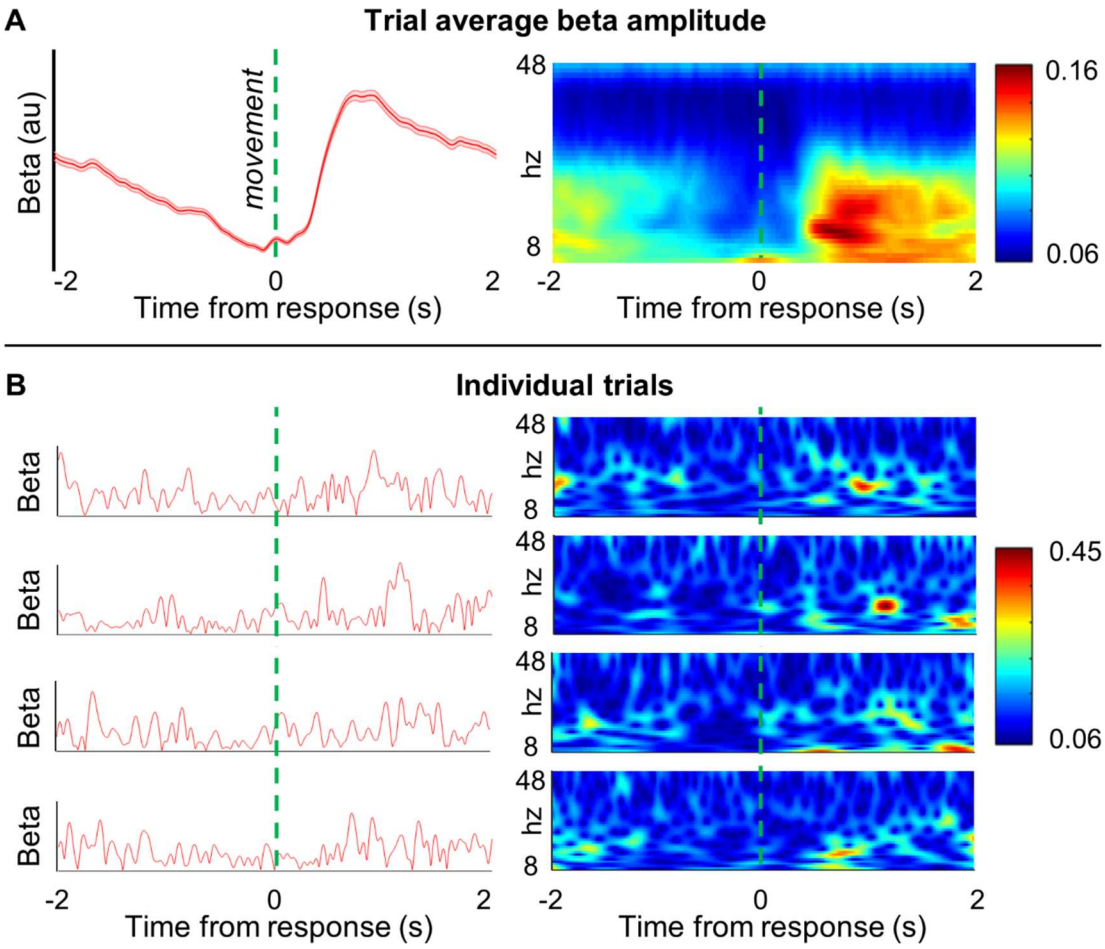

Figure 3.4 Evidence for burst activity (adapted from (Little et al., 2019)). Panel A shows the average time-frequency spectrogram (right) and evolution of beta-band activity (left) throughout the trials of a motor task. When inspecting the unaveraged signals for individual trials (panel B), we see short periods of high-amplitude activity instead of the continuous band observed in A.

Figure 3.5 shows an example of two signals with equivalent spectra even though the underlying dynamics are very different. Using an approach adapted from (Quinn et al., 2019) two possible models of brain activity were employed to illustrate this. First, characteristic brain 'noise' was simulated by generating pseudo-random Gaussian noise and filtering to obtain the characteristic $\frac{1}{f}$ spectrum. For the *burst* time-course, short sinusoids of constant amplitude and a frequency of 22.5 Hz (centre of the beta band) were generated and windowed using a tapered cosine function to remove any sharp edges. This burst signal was added to the 1/f-noise. The *sustained beta* time-course was generated by

amplitude modulation. Gaussian noise was limited to 3% of its bandwidth by setting all higher frequencies to 0 in the frequency domain and performing an inverse Fourier transform to yield a modulation time-course. This time-course was used to modulate a 22.5 Hz sinusoid which was added to a second vector containing 1/f-noise. Power spectra for both signals were estimated using Welch's method, averaging over all non-overlapping windows of length 0.5*s*. We can see that the averaged spectra are virtually equivalent, showing the expected peak in the beta band (at the chosen frequency of 22.5 Hz).

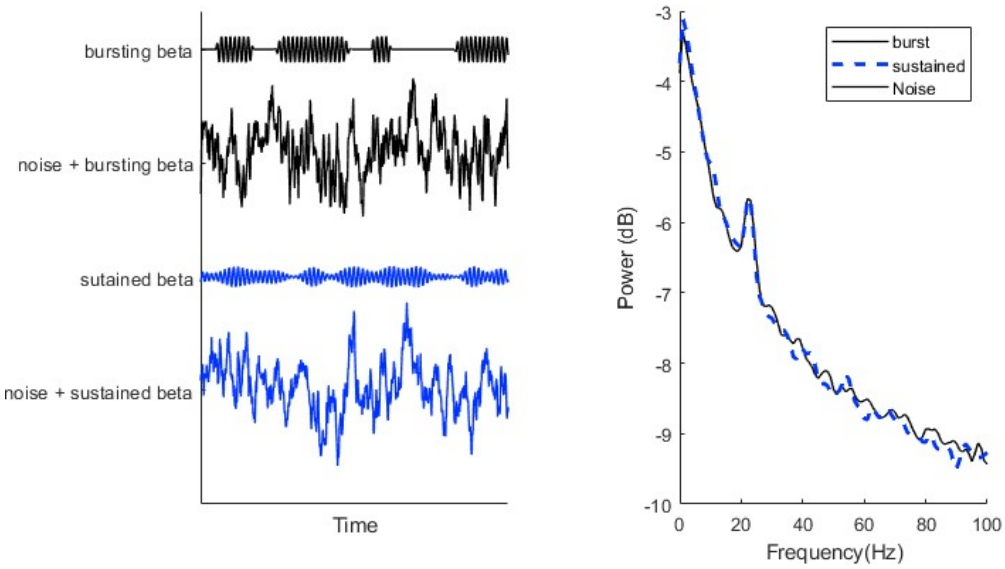

Figure 3.5: Differing generative models of MEG time-courses can result in equivalent spectral profiles. Note that the noise for each signal was generated separately

As demonstrated in Figure 3.5, time-averaging of MEG signals could deprive us of important information regarding the underlying dynamics of brain data. The common use of differences in spectral power when distinguishing experimental conditions or patient groups from controls might yield an incomplete picture. As an example, one might ask whether any difference in beta band power in a patient group stems from a change in the number of bursts per unit time, the duration of bursts, the amplitude of the bursts, or indeed a combination of these factors. For task data, burst measures may provide a baseline-free method of investigating patterns of activation. Recent work by Seedat et al. (2020) first

demonstrated that transient bursts also played an important role in functional connectivity. Functional network patterns reconstructed using AEC also emerged when assessing the temporal coincidence of burst events in separate brain regions.

In chapters 4 and 5, we used a Hidden Markov model or HMM to achieve the segmentation of MEG time courses. In the following sections, I will introduce the HMM and outline the principles of variational inference which are used to estimate the model parameters.

3.4.1.1 *Hidden Markov models for bursts extraction*

3.4.1.1.1 Markov chains

Systems that evolve in time by transitioning probabilistically between a set of reoccurring states can be mathematically described using a Markov chain. A graphical model of such a system with two states is shown in Figure 3.6. Given a starting state, a chain of states can be sampled by choosing the next state only using the outgoing probabilities of the previous state. This concept is known as the Markov property and can be mathematically formulated as

$$P(X_t) = P(X_t|X_{t-1}) \tag{3.53}$$

where t indicates the current time-point.

The graphical Markov model in Figure 3.6 can be described using a transition matrix

$$T = \begin{bmatrix} 0.7 & 0.3 \\ 0.2 & 0.8 \end{bmatrix} \tag{3.54}$$

where each row indicates the transition probabilities from one state. Each element T_{ij} (where $i \neq j$), in this matrix gives the probability of changing from state i to state j, while the diagonal elements denote the likelihood of staying in

the same state. Note that the values contained in the rows are probability distributions and sum to 1.

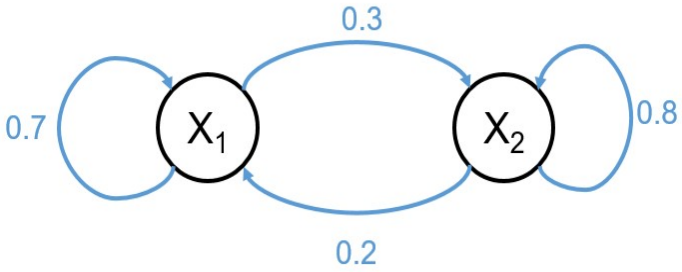

Figure 3.6 Markov Model with two states X_1 and X_2. The arrows and numbers indicate the possible transitions and the corresponding probabilities.

Markov models can be used when the investigated system has pre-defined, discrete and knowable states. In many cases, however, the exact states underlying a sequence of measurements might not be known and could be corrupted by noise or other nuisance signals (Rabiner, 1989). For our desired application of finding burst states (versus non-burst states) in noisy neuroimaging data, both of these factors are true. Thus an extension to the simple Markov Model is needed to account for this added signal corruption.

3.4.1.1.2 The Hidden Markov Model

The graphical signal model for an HMM with two hidden states and two possible observations is shown in Figure 3.7. Here the transitions between the now hidden states X_i are described by the same matrix T as in section 3.4.1.1.1. The model observations Y_i are related to the underlying hidden states via a set of emission probabilities, shown in red in Figure 3.7.

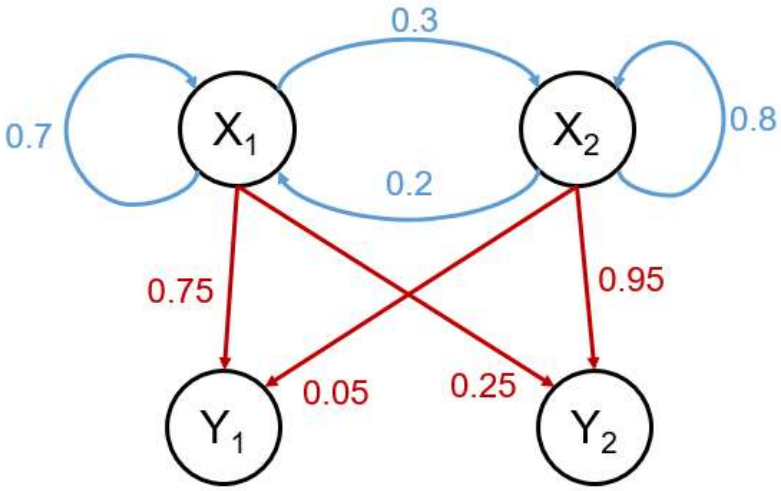

Figure 3.7 Graphical representation of a Hidden Markov Model. Blue arrows and values denote the transition probabilities between the hidden states X_i; red arrows denote the emission probabilities of observing states Y_i given each hidden state.

The emission probabilities can be mathematically written in matrix form

$$\boldsymbol{E} = \begin{bmatrix} 0.75 & 0.25 \\ 0.05 & 0.95 \end{bmatrix} \tag{3.55}$$

where each element E_{ij} describes the likelihood of making a measurement Y_i if the currently active state is X_j. A toy example of such a system could be described as follows:

Let's consider our friend Mark, who is prone to weather-related mood swings. On any given day Mark is either happy (Y_1) or sad (Y_2). If the weather is sunny (X_1) Mark is more likely to be happy (75% chance), while a cloudy day is highly likely to put him in a bad mood (95% chance). Where Mark lives, the weather is relatively stable, meaning that it is more likely to stay the same than change from one day to the next, i.e. the transition probabilities on the diagonals of T are larger than the off-diagonal transitions which indicate a change of state.

If we contact Mark on three consecutive days to ask him about his mood, we can attempt to infer the weather on those days using information contained in the transmission and emission matrices. We can do this by considering the probability of a sequence $X = X_{t0}, X_{t1}, X_{t2}$ given the sequence of observations $Y = Y_{t0}, Y_{t1}, Y_{t2}$ —or $P(X|Y)$—and taking the sequence X which maximises $P(X|Y)$. A possible sequence is shown in Figure 3.8.

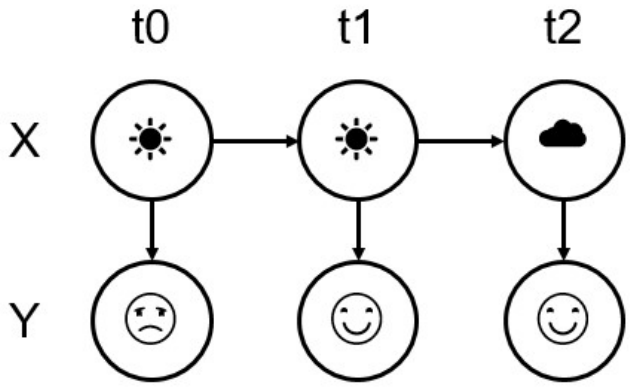

Figure 3.8 One possible sequence of weather states given a measured sequence of moods.

To estimate $P(X|Y)$, we need to find an alternative expression using Bayes' theorem:

$$P(X|Y) = \frac{P(Y|X)P(X)}{P(Y)}. \tag{3.56}$$

The Markov property (equation 3.53) applies to the sequence of hidden states which only depend on the respective previous time-point. $P(X)$ can therefore be written as the product of the transition probabilities of each step from t_i to t_{i+1} and the probability of the first hidden state visited in the sequence (as we don't know past states before the observed sequence). Mathematically

$$\begin{aligned} P(X) &= P(X_{t0})P(X_{t1}|X_{t0})P(X_{t2}|X_{t1}) \\ &= P(X_{t0}) \prod_{i=t1}^{tn-1} P(X_{ti}|X_{ti-1}), \end{aligned} \tag{3.57}$$

where the terms after the product operator are the elements of the transition matrix and n is the number of time steps. The first term is not explicitly stated but can be estimated from the transition matrix. The probability of the initial state is the likelihood of the state occurring given an infinite sequence, which is defined by the stationary distribution. In our example sequence shown in Figure 3.8, it is the overall chance of a sunny day.

It is possible to estimate the stationary distribution—often written as $\boldsymbol{\pi_0}$—by calculating the left eigenvector of the transition matrix, or solving

$$\boldsymbol{\pi_0 T} = \boldsymbol{\pi_0}. \tag{3.58}$$

In this example $\boldsymbol{\pi_0} = [0.4, 0.6]$, meaning that we can expect a sunny day 40% of the time. Finally, we can write down an expression for $P(Y|X)$. The probability of observing any Y_i given the underlying state X_i only depends on the emission probabilities. Hence, for any sequence

$$P(Y|X) = \prod_i P(Y_i|X_i)\,. \tag{3.59}$$

Substituting equations 3.57 and 3.59 into the numerator[7] of 3.56 we can get the expression for the probability of a given sequence

$$\begin{aligned} P(X,Y) &= P(Y|X)P(X) \\ &= P(X_{t0}) \prod_{i=t1}^{tn-1} P(X_{ti}|X_{ti-1}) \prod_{i=t0}^{tn-1} P(Y_i|X_i)\,. \end{aligned} \tag{3.60}$$

The likelihood of our example sequence in Figure 3.8 is, therefore:

$$\begin{aligned} P([☀, ☀, ☁], [☹, ☺, ☺]) &= 0.4 \times (0.7 \times 0.3) \times (0.75 \times 0.25 \times 0.05) \\ &\cong 0.00079 \end{aligned} \tag{3.61}$$

[7] Given that we want to maximise the fraction with respect to X, which does not feature in the denominator, we can focus on the numerator of P(Y|X). This quantity is often referred to as the joint posterior probability: a distribution over all hidden parameters and observed variables combined.

Maximising equation 3.60 with respect to the sequence of hidden states will yield the most likely sequence X given our observations Y. A naïve approach to finding this sequence would be to calculate $P(X, Y)$ for every possible sequence X, this, however, becomes intractable for long sequences like the ones we might expect when analysing neuromagnetic signals. Additionally, the parameters governing our underlying signal generator—the transition matrix—as well as the variables governing our observations—the observation model or emission matrix in our example—are usually unknown. Indeed, the number of possible hidden states K, which determines the size of the $K \times K$ transition matrix might not be known.

For this thesis, our goal is to segment MEG time-series into short-lived states, but importantly without having access to precise knowledge about any underlying brain states or the exact processes of signal generation. We aim to learn all these parameters—state transition probabilities and the variables governing our observations—from the data itself and some assumptions about the properties of our data. Following a similar notation to Woolrich et al. (2013), we can express the distribution over these parameters, for an HMM of length in time T, given our data as

$$P(s, \theta, \pi | y), \tag{3.62}$$

where $s = \{s_1, \dots s_T\}$ is the set of hidden state variables, θ is the set of parameters which define our observation model, π is the set of transition probabilities, a $K \times K$ matrix for a model with K states, and $y = \{y_1, \dots y_T\}$ is our dataset. Applying Bayes' rule (equation 3.56), the Markov property (equation 3.53), and the relation between observations and hidden states defined by the HMM we can write down the joint probability distribution over all parameters and variables:

$$P(s,\theta,\pi|y) = \frac{P(s_0|\pi_0)\prod_t P(s_t|s_{t-1},\pi)P(y_t|s_t,\theta)P(\theta)P(\pi)}{P(Y)}$$

$$\therefore P(y,s,\theta,\pi) = \underbrace{P(s_0|\pi_0)\prod_t P(s_t|s_{t-1},\pi)}_{Markov\ Chain}\underbrace{P(y_t|s_t,\theta)}_{\substack{Observation\\Model}}\underbrace{P(\theta)P(\pi)}_{Priors}. \quad (3.63)$$

Here, $P(\theta)$ and $P(\pi)$ are non-informative priors for the distributions over the transition probabilities and observation model parameters (Woolrich et al., 2013). Having obtained an expression for $P(y,s,\theta,\pi)$ we are now tasked with finding the set of hidden variables $\{s,\theta,\pi\}$—we will call H—which optimally fits our data and describes our system. Optimising the vast number of variables needed for such a system is intractable, especially for large datasets such as MEG recordings. To alleviate this issue we can use variational inference.

3.4.1.1.3 Variational Bayesian Inference

In the previous section, we concluded that estimating all of our unknown model parameters and variables for the joint distribution over all hidden variables H and the observed data y is likely to be intractable. In such situations, it is often possible to approximate our distribution via a simpler surrogate distribution. An example of this is illustrated in Figure 3.9, where a complex distribution P with multiple modes is approximated by a Gaussian distribution Q over the same space H. Our goal is to find $Q(H)$ that closely approximates $P(H|y)$. To solve this problem we can employ the principles of variational inference which allow us to define a cost function based on the Kullback-Leibler divergence or KL divergence, the optimisation of which will yield an approximation of our desired distribution (Bishop, 2006).

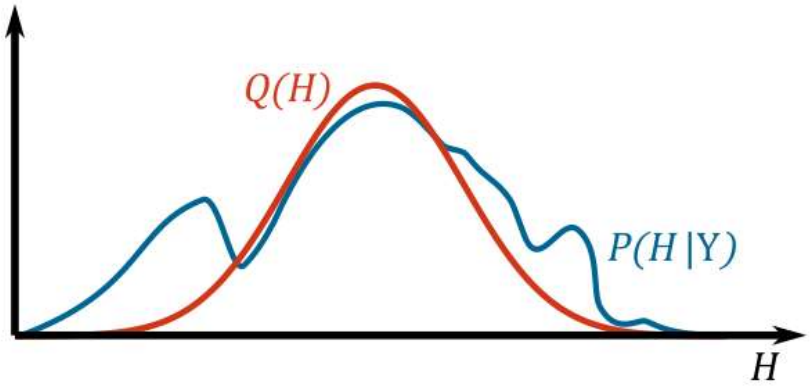

Figure 3.9 A complicated distribution over a set of hidden variables H given a set of observations Y might be intractable to estimate. Variational inference can be used to estimate a surrogate distribution Q which approximates P and can be used to find parameters H given the observed data.

The KL divergence (D_{KL}) is defined as the expected value of the logarithmic difference between two distributions which we can rewrite in terms of our known joint distribution in equation 3.63:

$$D_{KL}\big(Q(H) \parallel P(H|y)\big) = \int Q(H)\log\left(\frac{Q(H)}{P(H|y)}\right)dH$$

$$= \int Q(H)\log\left(\frac{Q(H)P(y)}{P(y,H)}\right)dH$$

$$= \int Q(H)\log\left(\frac{Q(H)}{P(y,H)}\right)dH + \log\big(P(y)\big). \qquad (3.64)$$

Note that the second integral vanishes as the distribution over the already observed data $P(y)$ does not depend on the hidden variables and the integral over the probability distribution $Q(H)$ is by definition 1. Inspecting 3.64 we see that minimising D_{KL} amounts to maximising the integral term, as the logarithm of $P(y)$ is a negative constant. It is common to refer to this as free energy minimisation, where the free energy $\mathcal{F}$ is defined as

$$\mathcal{F} = -\int Q(H)\log\left(\frac{P(y,H)}{Q(H)}\right)dH\,. \qquad (3.65)$$

I will now outline additional assumptions that are made to allow for the estimation of our hidden variables.

An important assumption of the approximate distribution Q is that the hidden variables are independent of each other (mean-field assumption (Haft et al., 1999)), which means that Q can be factorised into separate distributions Q_i for each variable H_i:

$$Q(H) = \prod_i Q_i(H_i)\,. \tag{3.66}$$

Note that each Q_i is a probability distribution, hence the additional constraint

$$\int Q_i(H_i) dH_i = 1. \tag{3.67}$$

Crucially, the set of all H contains the hidden state variables and the model parameters, which means that minimising $\mathcal{F}$ separately, with respect to the distributions $Q_i(H_i)$ gives the update formulae for these parameters, allowing us to compute the observation model parameters, state variables and transition probabilities (Haft et al., 1999). Sequentially improving the estimates of $Q_i(H_i)$, optimises the overall estimate of Q. As described by (Haft et al., 1999; Rezek and Roberts, 2005), the update equations for Q_i can be derived in a general form, regardless of the type of observation model:

$$\log\big(Q_i(H_i)\big) = \int Q_{j\setminus i}\big(H_{j\setminus i}\big)\log\big(P(y,H)\big)\, dH_{j\setminus i}. \tag{3.68}$$

Here, '$j \setminus i$' indicates the set of all j except i, which can be interpreted as fixing all variables except H_i. We can therefore improve the estimate of Q_i, which in turn optimises Q. The entire optimisation procedure involves repeating this process, cycling through all $Q_i(H_i)$, until a convergence criterion is reached.

The general applicability of equation 3.68 provides great flexibility when applying HMMs, as the observation model can be tailored to the problem and

data at hand. The next section will outline the choice of observation model used throughout this thesis.

3.4.1.1.4 Choice of observation model

In the basic case described in section 3.4.1.1.2, the observation model was referred to as the emission matrix. For applications such as modelling electrophysiological signals, more complex distributions need to be chosen. A relatively simple example would be a Gaussian observation model, where each state is assumed to generate observations drawn from a normal distribution with the mean and variance as parameters unique to the state. One drawback of such a model is that the Markov property constrains each observation to only be dependent on the previous time-point. This is likely not the case for signals originating in the brain. An elegant solution to include more temporal context is the time-delay-embedded HMM, or TDE-HMM, first introduced by Vidaurre et al. (2018). Here, the observation model is constructed using the autocovariance of the modelled signal over a set of embedded lags, or time-points around the time point of interest. The TDE-HMM effectively models each state based on the spectral content in a pre-defined window around each time point. The experimental work in Chapters 4 and 5 was conducted by using a TDE-HMM.

3.4.1.1.5 Practical considerations

When implementing an HMM, a series of choices ought to be made to ensure successful and consistent inference. This includes defining the complexity of the model, hyper-parameters for the observation model, and the structure of the input data. As is the case with Bayesian inference, appropriate priors will have to be chosen. This section will cover these concerns as well as briefly describe the model outputs provided by the HMM implementation used in subsequent chapters.

3.4.1.1.5.1 Model size

A primary concern with HMM analysis of MEG signals is the lack of knowledge about the number of underlying states generating the observed MEG signal. It is necessary for us to pre-specify the number of Hidden states, however, an excessively large number of states may lead to overfitting. We relied on two factors when choosing the number of states:

1. What can be considered a useful number of states for the experiment?
2. Are any states redundant?

Firstly—reminding ourselves of the transient burst model of brain activity from Figure 3.5—we should note that our primary interest was to detect reoccurring, transient bursting activity in MEG time courses. Assuming that this activity is present in *one* of the model states, the number of additional states becomes secondary. We, therefore, base our chosen state dimensionality on previous work (Seedat et al., 2020):

Briefly, Seedat and colleagues showed that choosing three, six or ten states still results in a single state with a high positive correlation to the beta-band envelope and—crucially—with spectral content invariant to the number of states. This suggests that our state of interest is not a result of under-fitting and cannot be consistently split into two or more states by simply increasing the model size. We therefore consider choosing more than three states unnecessary.

Secondly, we utilised a functionality of the model implementation which allows for states to be automatically removed at the training stage, provided they make up only a small proportion of the modelled signal, allowing for those time-points to be absorbed into the remaining states. Note that this never occurred during the analyses presented in Chapters 4 and 5 suggesting that three states are adequate.

For TDE-HMMs an additional parameter is the range of time-lags that are used to define the observation model. Wider time windows will allow for better estimation of low-frequency components in the signal, while short windows will cause the model to focus on higher frequency information.

3.4.1.1.5.2 Data structure and model output

HMMs can be applied to time series of any form. This means that investigators must make decisions with regards to any pre-processing of the MEG data passed to the HMM; narrow or broadband filtering, channel or source-level data, raw data or envelope time-courses are all aspects to be considered. Depending on the research question, the HMM can be applied in a multivariate manner where all channels or brain locations are modelled simultaneously. In the work presented here, however, univariate modelling was used exclusively. Given that electrophysiological data have a high temporal resolution, downsampling should be used to reduce the space of hidden states and improve computational efficiency.

For our purposes, the most important result produced by a TDE-HMM is the set of time courses describing the likelihood of each state being occupied at every point in time. Other features like the transition matrix and the autocovariance matrices of each state were not used for the experimental work presented here.

3.5 SUMMARY

In this chapter, the theoretical framework behind source reconstruction—beamforming in particular—was discussed. I described analysis methods that allow us to investigate different aspects of brain function such as synchrony and information exchange between separate brain regions via functional connectivity, or interesting dynamic activity when applying HMMs. Equipped with this toolbox of techniques and the instrumentation described in Chapter

2, we set out to use MEG to capture signals characteristic of mTBI and develop new experimental paradigms to be used in future mTBI research.

Chapter 4 TRANSIENT BURSTS IN MTBI

The excellent sensitivity of MEG in probing neural function makes it a promising tool for assessing mTBI with numerous studies proposing abnormal neural oscillations as a potential biomarker (for a systematic review see (Allen et al., 2021) and Chapter 1 for a summary). However, growing evidence suggests that neural dynamics are (in part) driven by transient, pan-spectral bursting events. Using this model to investigate mTBI, we applied an HMM to resting-state and motor-task MEG data, to show that previously reported findings of diminished intrinsic beta amplitude and connectivity in individuals with mTBI can be explained by changes in the beta band spectral content of bursts and a loss in the temporal coincidence of bursts respectively. In the motor task, mTBI resulted in reduced burst amplitude, altered modulation of burst probability during movement and decreased connectivity in the motor network. The work presented in this chapter suggests that mTBI impairs the coordination of neural activity across the brain and proposes a potent new method for understanding mTBI.

The work presented in this chapter was previously published in *NeuroImage: Clinical* (Rier et al., 2021).

4.1 INTRODUCTION

As laid out in Chapter 1, mTBI contributes substantially to the burden on health services (Cassidy et al., 2004; James et al., 2019), yet objective diagnoses and predictors of outcome after mTBI are lacking. Whilst most individuals with mTBI are expected to recover quickly and spontaneously, some studies report that up to 50% of patients continue to experience persistent and debilitating symptoms which—despite the 'mild' description—can severely impact quality of life (McInnes et al., 2017; Nelson et al., 2019; Wilson et al., 2021). By some definitions, mTBI is not associated with demonstrable abnormalities on routine clinical imaging with CT or MRI. Yet neurocognitive sequelae and symptoms suggest that neural dysfunction can persist after injury. Consequently, new, objective means to understand the neuropathology of mTBI, prognosticate, and inform intervention, are required.

Electrophysiological imaging provides a potentially potent approach to understanding the functional consequence and impairment after mTBI. MEG allows for sensitive imaging of neural activity and dynamics, allowing direct access to brain electrophysiology. Further, as we have seen, inverse modelling yields 3D images of brain electrophysiology, with unmatched non-invasive insight into micro-, meso- and macroscopic neural circuits that dynamically form and dissolve to underpin cognition. As described in Chapter 2, the MEG signal is dominated by 'neural oscillations' (rhythmic electrical activity generated by neural assemblies) which are traditionally divided into canonical frequency bands (delta (1-4 Hz), theta (4-8 Hz), alpha (8-13 Hz) beta (13-30 Hz) and gamma (30 + Hz). Abnormalities in these signals are an indication of pathology, and several putative atypical oscillatory signatures of mTBI have been reported.

Much of the literature has focused on the "pathological slowing" of rhythmic brain activity, including elevated low-frequency power, in particular the delta

band (Huang et al., 2014; Lewine et al., 1999; Proskovec et al., 2020). These observations are consistent with animal studies where increased slow-wave activity was observed following the controlled induction of white-matter lesions (Gloor et al., 1977), suggesting abnormal delta is a consequence of white matter damage. More recent work suggests high-frequency dysfunction: for example, Huang and colleagues reported abnormal resting-state gamma activity in subjects with combat-related mTBI (Huang et al., 2020) and Zhang et al. demonstrated a decrease in beta amplitude in mTBI patients (Zhang et al., 2020). These high-frequency signals are driven by distinct neurophysiological processes and open new routes to understanding the pathology of mTBI.

The heterogeneity of sequelae of mTBI is thought to be driven by diffuse yet subtle white matter damage, particularly around the corpus callosum (Aoki et al., 2012), affecting communication between distal, functionally specific regions. The degree of communication between regions can be tested directly using functional connectivity assessment. For MEG, this tends to mean assessment of either coherence, phase synchrony, or amplitude envelope correlation, between oscillations in frequency bands of interest (O'Neill et al., 2015a). The beta band is of particular interest since it appears crucial in the establishment of canonical resting-state networks (Brookes et al., 2011a; Hipp et al., 2012) as well as the dynamic orchestration of neural activity (Little et al., 2019; Sherman et al., 2016). Given the presumed diffuse nature of the disruptions caused by mTBI, the beta band is, therefore, a good candidate to assess this brain injury. In support of this, Zhang et al. showed that machine learning can classify injured subjects with high accuracy and that beta connectivity was an important feature for classification, showing that dysfunctional beta connectivity is an important marker of mTBI (Zhang et al., 2020).

All of the above assumes the "classical" view that oscillatory brain function is largely smoothly modulating rhythmic activity. However, recent findings suggest that this view is impoverished (as laid out in Chapter 3). Paradigm-

shifting work now shows that transient, dynamic burst states are a fundamental mode of neural functioning (van Ede et al., 2018). "Oscillatory" power does not stem from oscillations *per se*, but rather it results, at least in part, from transient events of high amplitude, the spectral content of which intersects canonical frequency ranges, particularly the beta band (Jones, 2016; Little et al., 2019; Sherman et al., 2016; Zich et al., 2020). When averaging across trials, these events sum to give the impression of slowly modulating oscillations—when, in fact, apparently sustained beta activity around motor events results from pan-spectral bursts whose probability of occurrence changes throughout the task, for example (Little et al., 2019; Seedat et al., 2020).

This evolution in the understanding of neural dynamics shows that classical measures of oscillatory amplitude should also be supplemented with nascent metrics like "burst likelihood" and "burst duration", as they may be driving the 'classic' amplitude metrics. The burst model also offers a novel means of calculating communication and connectivity. Seedat et al. used a Hidden Markov Model (HMM) to identify bursts and showed that the beta band connectome can be explained by the temporal coincidence of bursting (Seedat et al., 2020). Using a similar technique, Gascoyne et al. (2021) showed that previously reported beta abnormalities in schizophrenia can be explained using the burst model. The neurobiological origins of the beta burst have been studied through neural modelling (Jones, 2016; Sherman et al., 2016), which could provide a deeper understanding of the origin of signal abnormalities within the cortex and suggests that important features of these transient neural signals can be lost through averaging data from many epochs or trials. The use of these methods in mTBI would reveal novel information about the neural impairment underlying mTBI sequelae.

In this chapter, we analysed MEG recordings in 52 subjects (29 mTBI subjects and 23 healthy controls)—a subset (50 subjects) of which was previously presented by Zhang et al.—using a burst framework. Zhang and colleagues

showed that there were intrinsic deficiencies in both spontaneous beta power and connectivity (Zhang et al., 2020). Here, using an HMM approach, we tested the hypothesis that *the previous observation of abnormal beta activity can be explained by deficits in burst parameters (specifically, low beta burst amplitude)*. In the same resting-state data we hypothesised that *abnormalities in connectivity can be explained by the lack of temporal coordination between bursts in spatially separate brain regions*. We further used a machine learning (ML) approach to determine the characteristic connections that are disrupted in mTBI and to evaluate the utility of burst connectivity measures in ML-based mTBI diagnosis. We also present a novel analysis of (previously unpublished) motor task data acquired in the same subjects to test a hypothesis of abnormal connectivity during the well-known post-movement beta rebound (PMBR)—a marker for interhemispheric recalibration of the motor system, and by extension, an index of corpus callosum integrity (Tewarie et al., 2019). Assuming that mTBI disrupts white matter integrity, particularly around the corpus callosum, we hypothesised that *the coincident bursts in the motor cortex which drive the measurable PMBR and connectivity will be in relative deficit in mTBI subjects*.

4.2 METHODS

In total 52 subjects were recruited to the study; 29 mTBI patients and 23 healthy controls. All subjects underwent two MEG recordings: a resting state and a motor task. A flow chart showing subject inclusion is given in Figure 4.1. Groups were matched for age, sex and handedness. All subjects gave written informed consent to take part in the study, which had been approved by the Research Ethics board of the Hospital for Sick Children, Toronto, Canada. All subjects in the mTBI group were scanned within 90 days of their injury (e.g. acute-subacute phase of injury; 39 $\pm$ 21 days (mean $\pm$ standard deviation (SD)), range: 7-88 days); mTBI subjects had undergone investigation by MRI (T1 weighted anatomical; for acquisition parameters, see Appendix A) and no

positive clinical indications were found following review by a neuro-radiologist. In addition to MEG data acquisition, subjects were assessed for symptom severity using the Sport Concussion Assessment Tool 2 (SCAT2) on the day of scanning. The symptom evaluation contained in the SCAT2 battery of assessments asked subjects to subjectively rate the severity of 22 symptoms which span a range of clinical domains such as somatic (e.g. headache, neck pain), cognitive (e.g. feeling "in a fog", difficulty concentrating), emotional and behavioural changes (e.g. more emotional, irritability), and sleep disturbance (e.g. drowsiness, Trouble falling asleep) (McCrory et al., 2009).

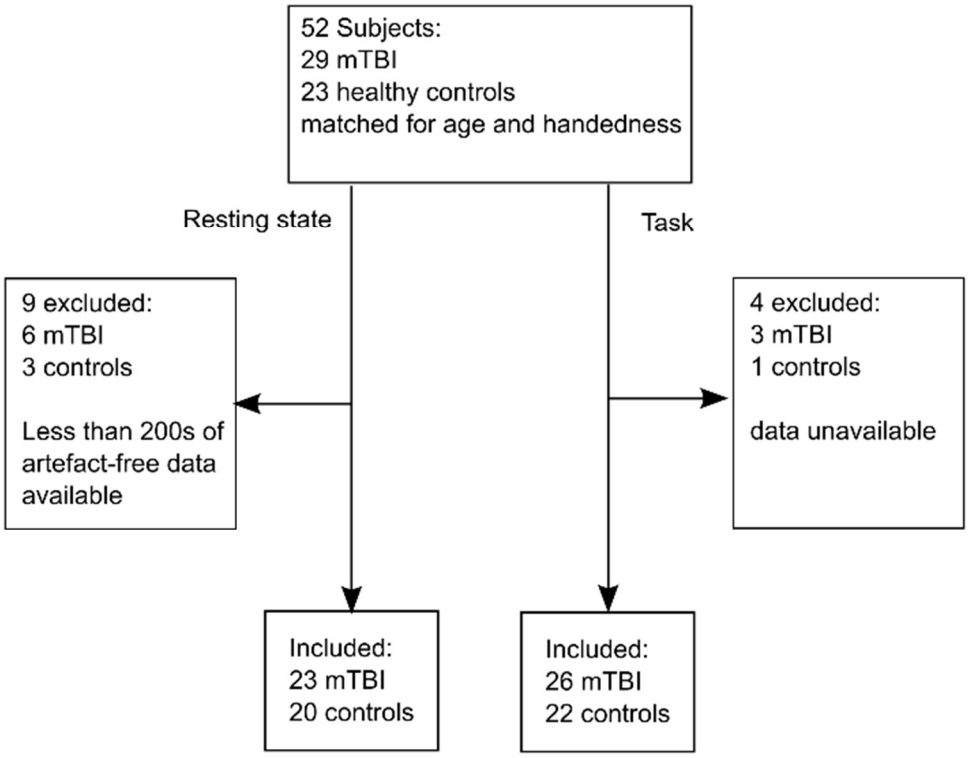

Figure 4.1 Patient enrolment flow chart

4.2.1 Paradigms

Data were recorded during two separate paradigms:

Resting-state: Subjects were asked to lie still, with their head in the MEG helmet, whilst 300 *s* of resting-state data were collected. Participants were in the supine position and were instructed to keep their eyes open. 23 individuals with mTBI (all male, aged 30 $\pm$ 7 (mean ± SD)) were included in the study alongside 20 healthy controls (all male, aged 28 $\pm$ 5 (mean ± SD)). (These data were previously presented by (Zhang et al., 2020))

Motor task: We used a stimulus matching task with a 2-alternate forced-choice design that requires bilateral motor responses, previously described by Oh et al. (2014). Subjects were simultaneously presented with three images—two stimuli on the upper left and right of a central fixation point and a third target image immediately below the fixation point. The subjects were asked to match one of the two stimuli presented in the upper visual field with the third target stimulus in the lower field, on either the colour or shape dimension, and indicate which of the stimuli matched the target by pressing a button with the corresponding left or right hand. Stimuli presentations lasted until a response was given (up to a maximum of 4*s*) with a jittered inter-stimulus interval of 800 - 1200 milliseconds. Recordings lasted for 370 trials. 26 individuals with mTBI (all male, aged 30 $\pm$ 7 years (mean ± SD)) were included in the study alongside 22 healthy controls (all male, aged 28 $\pm$ 5 (mean ± SD)). A diagram depicting example stimuli is shown in Figure 4.2.

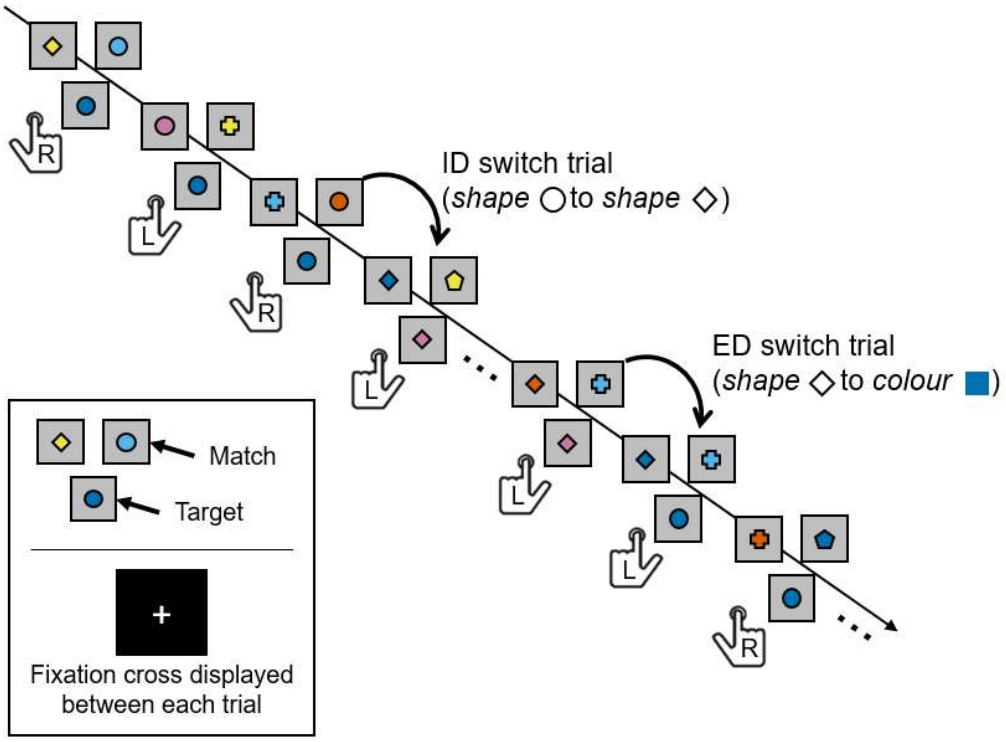

Figure 4.2 Stimulus matching task. In each trial, subjects were presented with three shapes displayed within grey squares on a black background (see inset): a centrally placed target shape, and two stimulus shapes on the left and right of the target. Subjects were asked to press a button with their left or right index finger to indicate a match between the target and the left or right stimulus shape respectively; the hand symbols below each of the examples indicate the correct response for the examples given. The target and stimuli could either match in shape or in colour and were pseudo-randomly selected to prevent ambiguous trials where shape and colour could be matched. After at least three consecutive trials with the same target stimulus and the same matching rule ("shape" or "colour") a switch trial occurred which would entail either an intradimensional (ID) shift (same matching dimension i.e. shape matching but to a different shape) or an extradimensional (ED) shift (different matching dimension i.e. shape to colour). A white fixation cross was displayed between each trial.

4.2.2 Data collection

All MEG data were collected using a 151-channel whole-head MEG system (CTF, Coquitlam, Vancouver, Canada) operating in a third-order gradiometer configuration, at a sampling frequency of 600 Hz. Before entering the scanner, three fiducial markers were placed on the subject's head (left and right

preauricular and nasion). These coils were energised to enable continuous tracking of the head position throughout the scan, and consequently, enabled estimation of head movement.

Note that data acquisition, pre-processing and source localisation were not performed by the author.

4.2.3 Pre-processing

The MEG data were filtered using a 4th order Butterworth band-pass filter between 1-150 Hz and notch filters at 60 Hz (mains interference) and 120 Hz (harmonic). The continuous head motion information was used to define epochs in which the head position was maintained within 5 mm of its starting location, and head velocity was less than $5mm/s$; other epoch exclusion criteria were: SQUID resets in the MEG signal (see section 2.4.1.3, footnote 4) and/or signal discontinuities exceeding $\pm 2pT$. Each epoch was visually assessed for adequate quality by a trained MEG expert. Epochs not meeting these criteria were rejected. Independent component analysis (ICA) was applied to attenuate ocular (EOG) and cardiac (ECG) related artefacts. Identified components were removed manually following visual inspection.

For the Resting-state data, the post-ICA artefact-free recordings were segmented into 10s epochs and only subjects with $> 200s$ (20 epochs) of clean data were included in the final analyses. The mean and standard deviation of available epochs for the resting state paradigm was 24.1 ± 0.6 for the mTBI group and 23.9 ± 0.6 for the control group.

The task data were epoched into $6s$ segments; $3s$ before and $3s$ after a button press. Trials containing artefacts, or incorrect responses were excluded and only set-shift trials were used. For the task, an average (mean ± SD) of 175 ± 49 and 172 ± 33 epochs was included for the mTBI and control groups respectively. For the mTBI group, this included 88 ± 29 left button presses and $88 \pm$

23 right button presses. For the control group, there were 87 $\pm$ 21 left button presses and 85 $\pm$ 19 right button presses. There were no significant differences in trial count between groups as assessed by a 2-sided T-test. Note also that reaction time did not differ significantly between groups. For left button presses, reaction times were 0.50 $\pm$ 0.12 s for controls and 0.55 $\pm$ 0.08 for mTBI (p = 0.11 using a two-sided T-test). For right button presses, reaction times were 0.54 $\pm$ 0.10 s for controls and 0.53 $\pm$ 0.07 for mTBI (p = 0.82 using a 2-sided T-test).

4.2.4 Data analysis

Following collection and pre-processing, the data analysis pipeline was similar to that introduced by Seedat et al. (2020).

4.2.4.1 *Source localisation:*

We characterised brain activity in 78 cortical regions (Gong et al., 2009) defined according to the Automated Anatomical Labelling Atlas (AAL) (Tzourio-Mazoyer et al., 2002). To this end, we used a linearly constrained minimum variance (LCMV) beamformer (Van Veen et al., 1997) implemented in FieldTrip (Oostenveld et al., 2011), which estimates electrophysiological activity at a predefined location/orientation in the brain, whilst minimising contributions from all other sources—both in the brain and environmental interference (see section 3.2.1). These estimates are known as virtual sensor time series. A single virtual sensor was placed at the centre of mass of each of the 78 AAL regions. A single shell head model (Nolte, 2003) was used to construct the forward solution for each participant by co-registering the MEG sensor locations onto an anatomical T1-weighted, age-appropriate MRI template (MNI152) using SPM12 through FieldTrip. A common spatial filter was computed for each region, using all artefact-free trials to generate a covariance matrix. The covariance matrix was regularised with 5% Tikhonov regularisation. The beamformer was applied to reconstruct the broadband time series for the

centroid of each of the AAL regions. Sources were projected to the dominant orientation by taking the eigenvector of the source covariance with the largest eigenvalue (Sekihara et al., 2004). The reconstructed virtual time series were filtered between 1 and 48 Hz. The resulting data comprised 78 regional electrophysiological time courses, per subject.

As mentioned in Chapter 3 one significant challenge in MEG connectivity estimation is that the ill-posed nature of the MEG inverse problem causes "leakage" of signal between virtual sensors at separate locations (Brookes et al., 2012). This leakage manifests as a zero-time lag linear summation of activity from other regions and, for this reason, orthogonalisation of virtual sensors before connectivity calculation results in a marked reduction of leakage, albeit at the cost of true zero-phase-lag connectivity (O'Neill et al., 2015a). Here, we employed multivariate symmetric orthogonalisation as described in Chapter 3 (Colclough et al., 2015): To recap briefly, a set of orthonormal time-courses that are closest to the virtual sensor data, and for which there is a simple analytic solution, is found. Secondly, the solution is finessed by iteratively adjusting the lengths of the orthogonal vectors until the solution is as close as possible to the uncorrected time courses. The resulting data contain the 78 orthogonalised virtual sensor time series. Following the application of this procedure, the time courses were downsampled to a 100 Hz sampling rate, mean centred, and variance normalised.

4.2.4.2 *Hidden Markov Model:*

All 78 regional time courses for each subject were processed independently using a univariate time-delay embedded HMM. The details of the HMM have been described extensively in section 3.4 and will not be repeated here. Briefly, the HMM assumes that a series of mutually exclusive hidden "states" governs each electrophysiological time course. This means that for every brain region, each time point is associated with a single state. Here, we used an HMM with time-delay embedding (Vidaurre et al., 2018) where each state is characterised

by a different autocovariance pattern defined over a specified time window (of duration 230ms). These state autocovariance patterns contain the spectral information of the signal when a particular state is active and consequently states are derived based upon specific repeating spectro-temporal patterns of activity. Previous work has shown that using the HMM in this way enables accurate identification of the pattern associated with the pan-spectral bursts that underlie the beta oscillatory signal (Seedat et al., 2020).

The same HMM was applied as described by Seedat et al. (2020). Specifically, model training was undertaken using variational Bayesian inference as described in Chapter 3. For each of the 78 time courses, we assumed 3 states, and so the model output was a set of 3 time courses representing the probability of each state being active over all time. To identify which of the three states corresponded to the pan-spectral bursts of interest (henceforth termed the "burst state"), we measured the correlation between the state probability and the amplitude of beta oscillations (defined by the application of a (Morlet-based) continuous wavelet transform to the regional time course and extracting those values corresponding to the 13-30Hz frequency band). The state whose probability time course correlated highest with the beta envelope was taken as the burst state while the remaining 2 states were defined to be non-burst states. The probability time courses for each state were subsequently binarised by assuming that if the probability exceeded two thirds, then the given state had been entered. These binary time courses enabled the identification of bursts and post-hoc analysis allowed measurements such as burst (and non-burst) state duration, burst amplitude, and coincidence. This method was applied to each of the 78 time courses independently, for all subjects and tasks.

4.2.4.3 *Summary metrics for resting-state data:*

For the resting-state data, we aimed to use the HMM output to investigate whether the previously observed (Zhang et al., 2020) beta deficits in power and

connectivity could be explained in terms of our burst framework. We calculated the following features from the HMM output:

- **Burst amplitude**: The maximum value of the beta envelope during each visit to the burst state.
- **Total burst time**: The proportion of time spent in the burst state throughout the resting state recording. (Calculated as the total time burst state was active, divided by the total experimental duration.)
- **Burst duration**: The average time spent in the burst state, on each visit.
- **Burst count**: The number of visits to the burst state per second, calculated by counting the total number of visits in a recording and dividing by the total length of the recording.
- **Functional connectivity:** Here we define functional connectivity as the temporal coincidence between bursts in two regions. Burst coincidence between regions i and j was measured using the Jaccard index, $J^{i,j}$—a ratio of the intersection over the union of two binary state time-courses. This yields a value between 0 and 1; 1 means perfect coincidence; 0 means no coincidence. Mathematically,

$$J^{i,j} = \sum_{t}^{N_t} \left(B_t^i \wedge B_t^j\right) \Big/ \sum_{t}^{N_t} \left(B_t^i \vee B_t^j\right), \qquad (4.1)$$

where B_t^i indicates whether the burst state was occupied in region i at time point t (1 indicating burst, 0 indicating non-burst). N_t is the number of samples, and $\wedge$ and $\vee$ are the logical AND and OR operators respectively. Figure 4.3 shows a schematic representation of this measure.

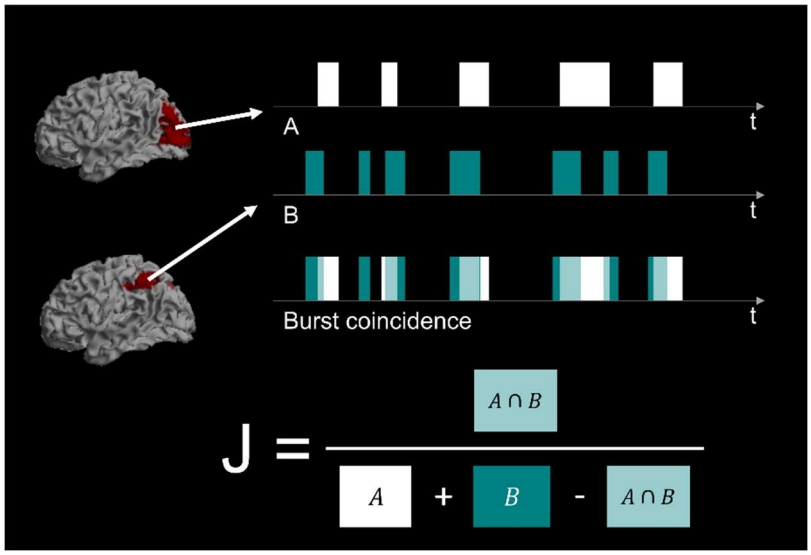

Figure 4.3 Illustration of the burst connectivity metric. The intersection and union of the binarised time-courses from two brain regions are used to calculate the Jaccard Index (J)

4.2.4.4 *Machine learning analysis of burst coincidence connectomes*

The burst coincidence connectomes derived from the resting state recordings were further analysed using a machine learning pipeline with the aim of determining their utility in distinguishing subjects with mTBI from healthy subjects. The pipeline included recursive Random Forest Feature Selection (rRF-FS) and binary classification using a support vector machine (SVM) model. All possible, unique connections between the 78 regions of the AAL atlas ($(78\times77)/2 = 3003$ connections) were used as input features. Using 10-fold cross-validation to evaluate classification performance, the rRF-FS was used to select the most important connections via a variable importance threshold, and consensus voting procedure (Zhang et al., 2020, 2016) (see Appendix B for more details), before SVM classification (Appendix C). Model performance was measured via the mean area under the receiver operating characteristic curve (ROC-AUC) after each cross-validation iteration. The statistical significance of the SVM model was tested by repeating the SVM training and classification 100 times, randomly permuting the sample group labels. Making use of the excellent interpretability of Random Forest Feature Selection, we assessed whether the

chosen connections are among those driving the global reduction in burst connectivity.

4.2.4.5 *Summary metrics for task data:*

For the task data, we used a similar approach, again measuring burst characteristics and connectivity, but this time in the context of task timing and the well-known beta band features (the movement-related beta decrease (MRBD) and the post-movement beta rebound (PMBR)). We calculated the following:

- **Burst probability time courses**: For a single region, binary burst time courses were reshaped into a matrix of time (within a trial) by the number of trials. For each trial timepoint, we assessed the probability (across trials) that a burst occurred. This was calculated as simply the sum of the total number of trials showing a burst at time t, divided by the total number of trials. These probability time courses were calculated within each subject and then averaged across subjects. We expected to see a decrease in burst probability during movement (corresponding to the MRBD) and an increase upon movement cessation (delineating the PMBR).
- **Burst statistics during PMBR window**: The PMBR window was defined from $\mathbf{0.45}\boldsymbol{s}$ to $\mathbf{0.85}\boldsymbol{s}$ relative to the button press (Pfurtscheller et al., 1996). For each trial in every subject, burst amplitude and duration were calculated for each burst which fell within, began, or ended during the PMBR window. These values were averaged within each trial, and subsequently across trials and subjects.
- **Burst coincidence time courses**: The time evolution of the burst coincidence was generated by expanding on the Jaccard index method. For every pair of regions, the binary burst time courses for both regions were reshaped into a matrix of time within a trial, by the number of trials. The Jaccard index was then calculated for each time point, t, within a trial (i.e.

rather than calculating the Jaccard index over all time as was done above in the functional connectivity section, we calculated it over all trials for each time point within a trial). This enabled us to define a time course of burst coincidence probability.

4.2.5 Statistical testing:

4.2.5.1 *Resting-state data*

The resting-state HMM yielded 78 values of burst amplitude, non-burst amplitude, and total burst time. To avoid a-priori assumptions on the brain regions or connections implicated in mTBI (which may differ between subjects) we collapsed these metrics across all regions/connections. This left a single global mean value for each metric. We then computed the difference between groups (mTBI and controls) and assessed statistical significance using a non-parametric Wilcoxon rank-sum test, corrected for multiple comparisons across the three measures using the Benjamini-Hochberg procedure (Benjamini and Hochberg, 1995). This enabled direct testing of the hypothesis that previous findings of abnormal beta activity can be explained by abnormalities in burst amplitude.

Independently, our connectivity analyses generated 3003 values of connectivity (i.e. one value for each connection, between all possible pairs of the 78 AAL regions). Again, to avoid a-priori assumptions on the most prominent connections affected by mTBI, values were collapsed across all connections, and the global mean connectivity was calculated. This was done for all subjects and a Wilcoxon rank-sum test was used to test significance. This allowed us to test the hypothesis that abnormalities in connectivity can be explained by a lack of coordination between bursts in spatially separate brain regions.

4.2.5.2 *Motor task data*

For the task paradigm, we first focused on the PMBR as this has been shown to be a sensitive marker of network communication and connectivity. We set out to test the hypothesis that interhemispheric connectivity would be disrupted in mTBI (i.e. by damage to white matter tracts in the corpus callosum). Specifically, we expected burst structure to be deficient during the PMBR, manifesting in changes in the burst features as well as synchrony between bursts.

We calculated the mean burst probability and burst amplitude during the PMBR. This was done separately for the left and right motor cortices, and left button press and right button press—yielding a total of 8 measurements. For completeness, we also measured the overall modulation of burst probability (i.e. the difference in burst probability between the PMBR and the MRBD windows). In all cases, a Wilcoxon rank-sum test was used to test for significant differences between groups ($p < 0.05$) and FDR correction was used to correct for multiple comparisons across the 12 separate tests (left/right cortex; left/right button press; 3 separate metrics). This allowed for testing of the hypothesis that the bursts which drive the measurable PMBR will be altered in individuals with mTBI.

To test our hypothesis that connectivity between motor regions during the PMBR will be diminished in the patient cohort, the group difference in the mean connectivity between the left and right motor cortices during the PMBR window was calculated, and statistical significance was assigned using a Wilcoxon rank-sum test. Here, an average was calculated over both trial conditions.

4.2.5.3 *Relationship with symptoms*

Finally, the relationships between symptom severity—as measured via the SCAT2—and MEG derived metrics (global burst amplitude, and burst connectivity) in the resting state, were assessed using Spearman correlation. We

reasoned that burst amplitude and connectivity might have a monotonically decreasing relationship with symptom severity (i.e., those with more severe symptoms would have diminished burst amplitude and connectivity). To test this, within the mTBI group only, we calculated the correlation between symptom severity and MEG measures. $P < 0.05$ and FDR correction was used to determine statistical significance. The association between symptom severity and the mean connectivity of the connections selected via the rRF-FS procedure was also measured using Spearman correlation, to determine whether the subset of features would show a stronger relationship than the whole brain measure of connectivity. In addition, we also tested the same correlation across the combined group of subjects (i.e., patients and controls). Note that this combined measure doesn't suggest a relationship between symptoms and MEG measures at an individual level. Rather, a significant correlation would be likely to be driven solely by a group difference. For this reason, the combined Spearman correlations were only used to support the group observations described above.

4.3 RESULTS

4.3.1 Spontaneous beta bursts are abnormal in mTBI

Figure 4.4 shows the burst statistics captured during the resting state recording. In Figure 4.4a, the upper panel shows the spatial distribution of beta amplitude during the bursts identified by the HMM. In agreement with previous work (Seedat et al., 2020), in control subjects, the amplitude is maximal over the posterior frontal and parietal lobe, in particular the sensorimotor cortices. This amplitude appears diminished in the mTBI group although the spatial pattern is similar. The lower panel of Figure 4.4a shows beta amplitude during the non-burst states. Here, the overall amplitude is much lower (as expected) and the spatial pattern is no longer apparent. Figure 4.4b shows the corresponding

spatial distributions of the differences in the burst (upper plot) and non-burst (lower plot) beta amplitude between the two groups.

In Figure 4.4c, the violin plots show the distribution (across subjects) of whole-brain burst and non-burst amplitudes; the plot on the left shows burst amplitude; the plot on the right shows the amplitude during the non-burst periods. Note that, in agreement with our hypothesis, during bursts, there is a significant ($p = 0.0165$; Wilcoxon sum rank test) drop in amplitude for patients relative to controls. However, here there is no such measurable difference in the non-burst windows. This demonstrates an important point: in (Zhang et al., 2020) the authors showed diminished resting-state beta amplitude; here we extend this finding to show that this reduction is driven largely by the spectral content of the bursts, and specifically that the amplitude of the beta component is reduced. Figure 4.4d contains violin plots showing the average time spent in the burst state. Note that there is no significant difference between patients and controls; further suggesting that the beta deficit in patients stems from a reduction in burst amplitude and not a reduction in the number or density of occurrence of the bursts. Finally, Figure 4.4e shows a scatter plot of global burst amplitude versus SCAT-2 symptom severity. Spearman correlation for the combined data points (patients and controls) gave $R = -0.399$; $p = 0.0088$, which is likely driven by the difference between groups. However, for the patients only, we found no evidence ($R = -0.189$; $p = 0.387$, Spearman Correlation) of a direct correlation between global burst amplitude and symptom severity. Note that we also tested the correlation between our burst metric and a more traditional approach to measuring beta amplitude. In the latter, a Morlet-wavelet transform was applied to the beamformer reconstructed data and values corresponding to the canonical beta band ($13 - 30Hz$) were extracted to obtain the amplitude envelope. The mean envelope value was then computed (over the entire experimental duration) and averaged across brain regions. We found a significant correlation ($R = 0.93$; $p = 5x10 - 20$; Pearson correlation)

between burst amplitude and beta amplitude (over all time), demonstrating clearly that the traditional beta metric (collapsed over all time) is at least partly driven by the beta amplitude during the burst windows (which only account for ~30% of the overall experimental duration).

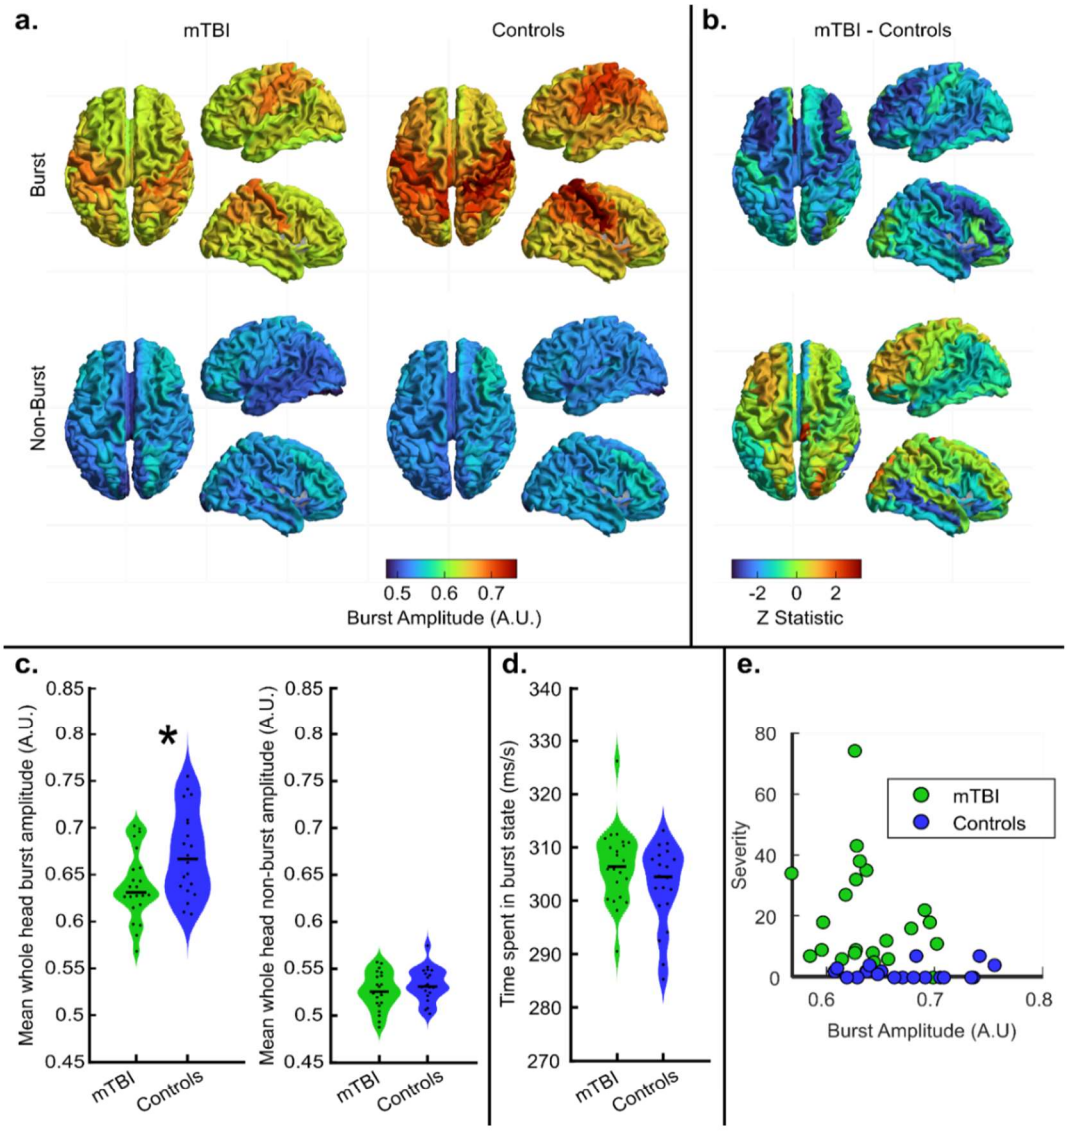

Figure 4.4 Resting-state burst statistics. a) The upper panel shows the spatial distribution of beta amplitude during bursts. The lower panel shows beta amplitude in the non-burst windows. In both cases, patients are shown on the left and controls on the right. b) Spatial signature of the differences in beta amplitude during bursts (top) and non-burst periods (bottom) (i.e. the difference between patients and controls, left and right in a) respectively) c) The left panel shows average burst amplitude, the right panel shows beta amplitude in the non-burst states, demonstrating no significant difference between patients and controls. In both cases, results are collapsed across 78 brain regions and each data point represents an individual subject. Note that burst beta amplitude is significantly ($p = 0.016$; Wilcoxon sum rank test) diminished in patients. d) Violin plot showing the average time spent in the burst state (no significant difference between groups) e) Scatter plot of global burst amplitude versus SCAT-2 symptom severity. Spearman correlation for the combined data points (patients and controls) gave $R = -0.399$; $p = 0.0088$. For the patients only, we found $R = -0.189$; $p = 0.387$.

4.3 RESULTS

Figure 4.5 shows the results of functional connectivity assessment in resting-state data. Figure 4.5a shows connectome matrices calculated via assessment of burst coincidence in patients (upper left panel) and controls (lower left panel). The corresponding upper and lower right panels show the spatial distribution of the 5% of connections (i.e. the $3003 \times 0.05 \approx 150$ connections) with the highest Jaccard index, plotted on a glass brain. The result for controls mirrors previous findings (both calculated using burst coincidence (Seedat et al., 2020) and the more widely used amplitude correlation metrics (e.g. Hunt et al., 2016); the largest connectivity tends to be found between homologous regions of the occipital, motor, sensory, and posterior parietal cortices. Interestingly, the overall connectivity pattern was maintained in patients, but the absolute values of connectivity are diminished.

Figure 4.5b shows the spatial signature of the differences in connectivity between patients and controls. Here we plot the 2% of connections with the highest differences between groups. The spatial signature shows that the largest resting-state connectivity differences are between posterior parietal regions. Note again that these findings extend previous work; Zhang et al. demonstrated using 'classical' methods that beta connectivity is diminished in mTBI. Here, using the same data, we show that these differences can be explained through an assessment of burst coincidence between brain regions.

This effect is formalised in the violin plot in Figure 4.5c. Here, each data point represents the mean strength of all connections in a single subject, and we note that connectivity is diminished significantly ($p = 0.031$; Wilcoxon rank sum test) in the mTBI group relative to controls. A significant correlation between symptom severity and global burst connectivity was also found (correlation derived using data from the mTBI group and controls combined) further verifying this measure (Spearman $R = -0.39; p = 0.01$). However, when taking into account individuals with mTBI only, the correlation with symptomology

was not significant (Spearman $R = -0.327; p = 0.128$). See the Appendix D for a scatter plot of global burst connectivity against severity.

The glass brain plot in Figure 4.5d shows the connections which most accurately separate patients and controls, as selected using the data-driven rRF-FS approach. The violin plot shows the mean connectivity over the selected features (connections) for each subject. Note that the separation between the groups is greatly improved compared to the global measure, as would be expected from this type of ML approach. Interestingly, many of the connections which distinguish patients are interhemispheric, potentially implicating damage in the corpus callosum. The average ROC-AUC classification accuracy across the ten cross-validation folds was 0.98 with a SD of 0.08. Finally, we measured a significant correlation between symptom severity and the average burst coincidence of the rRF-FS features; for the combined mTBI and control groups we found ($R = -0.72;\ p = 8 \times 10^{-8}$; Spearman Correlation). However, we note that this is influenced by the ML approach which yields features that best differentiate the groups and might therefore inflate the value of the calculated Spearman correlation. More interestingly, we found that for the mTBI group only, there was a significant correlation between connectivity and symptom severity ($R = -0.45; p = 0.03$; Spearman Correlation) which could not have been driven by the ML approach. A scatter plot of symptom severity against rRF-FS selected burst connectivity is shown in Figure 4.5e.

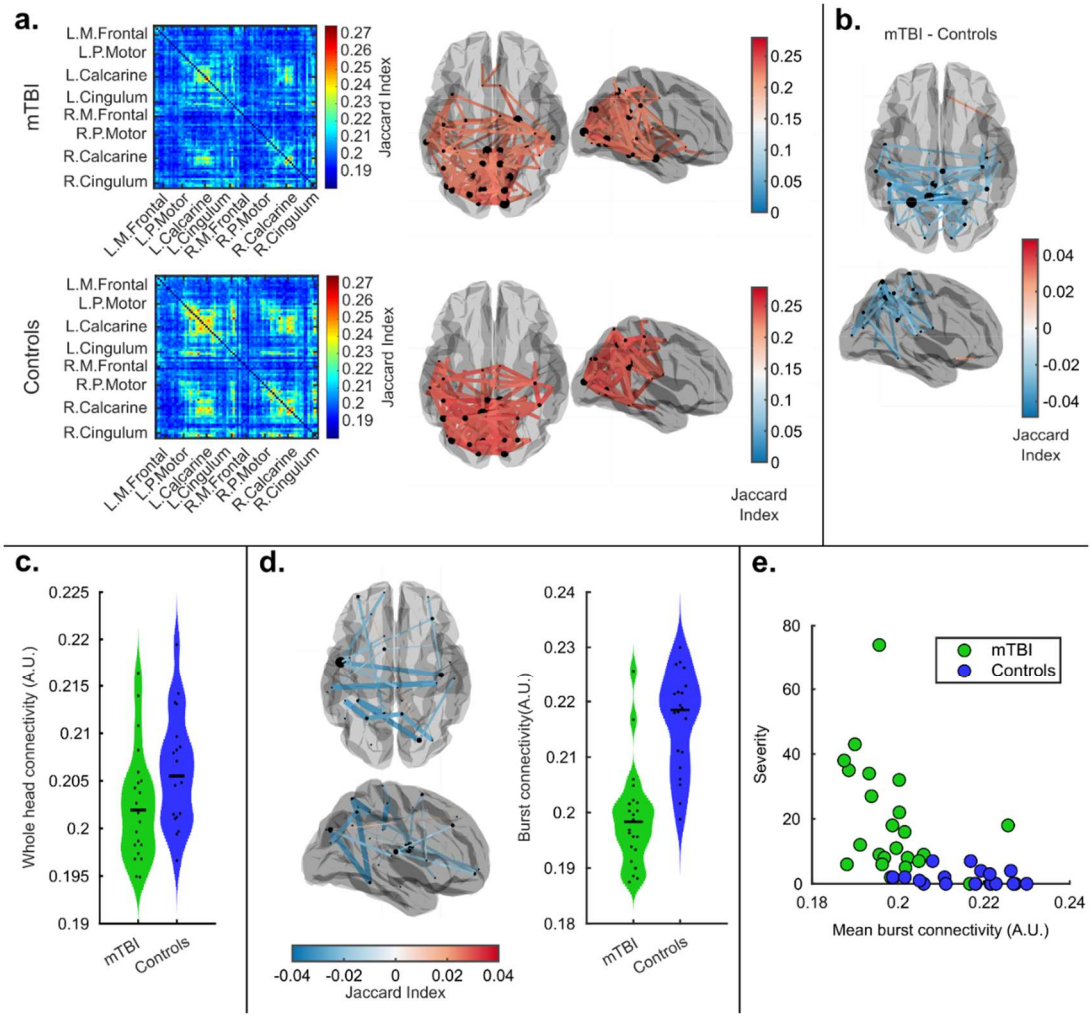

Figure 4.5 Resting-state functional connectivity. a) Resting-state connectome matrices for patients (upper panel) and controls (lower panel). In both cases, connectivity is calculated via the assessment of burst coincidence between regions. The glass brains show the spatial structure of the information in the corresponding connectome matrices. The red lines show 5% of connections with the highest functional connectivity value. b) The glass brain plot shows the 2% of connections with the largest difference between patients and controls. c) Whole-brain functional connectivity assessment in patients and controls. Each data point represents the average connectivity across all connections, for a single individual (i.e. the overall sum of all of the matrix elements in the connectomes shown in (a) divided by the number of connections). Whole head connectivity – computed via assessment of burst coincidence – is significantly ($p = 0.031$; Wilcoxon sum rank test) diminished in patients. d) Glass brain showing the connections selected using recursive Random Forest Feature selection and violin plot showing the average connectivity for those connections in both groups. e) Scatter plot showing the relationship between symptom severity and burst connectivity in the rRF-FS selected connections. Spearman correlation for the combined data points gave $R = -0.72$; $p = 8x10 - 8$. Spearman correlation for the mTBI group only gave $R = -0.45$; $p = 0.03$.

4.3.2 MTBI disrupts the dynamic neural repertoire of the motor system

Figure 4.6 shows burst statistics during the motor task. Note that the task contained both left and right finger movements and these have been analysed separately. We also examine effects in both contra- and ipsilateral cortices. (Statistical analyses employed FDR correction to account for multiple comparisons). Figure 4.6a shows the temporal evolution of burst probability throughout the task. The upper two plots show burst probability in the left motor cortex, for left (left) and right (right) button presses. The lower two plots show burst probability in the right motor cortex, for right (left) and left (right) button presses. In all cases, the green trace shows the patients and the blue trace shows healthy controls. The solid lines show the mean across subjects and the shaded regions show the standard error across subjects. Notice that a characteristic response is seen in both regions and all conditions, whereby the burst probability is diminished during the movement itself (i.e. around time $t = 0$) and is enhanced immediately following the movement. Previous work (Little et al., 2019) has shown that this change represents the basis of the MRBD and PMBR. There does appear to be a systematic effect in the contralateral cortex whereby the modulation of the burst probability is lower in patients.

In testing the burst probability during the rebound window, we found a significantly diminished likelihood of bursts in contralateral cortices (Right motor, left button $p = 0.0099$ *; Left motor, right button $p = 0.0248$ *; Wilcoxon Rank Sum test) but not ipsilateral cortex (* indicates significance following FDR correction). The overall modulation of burst probability—as measured by the difference in probability in the rebound and desynchronization periods—was also significantly different in both conditions in the contralateral cortex; (left motor cortex, right press: $p = 0.024$*; right motor cortex, left press: $p = 0.0047$*; Wilcoxon Rank Sum test). In the ipsilateral cortex, the modulation of burst probability was not significant following multiple comparison correction.

Figure 4.6b shows the amplitude during the HMM identified bursts during the rebound window. Here, we find significantly reduced beta burst amplitude in both left and right motor cortex during left button presses: (left-hand movement; left motor cortex, $p = 0.007$*. Left-hand movement; right motor cortex, $p = 0.005$*; Wilcoxon Rank Sum test). Similar trends were also observed for right-hand button presses but this was only significant in the right motor cortex ($p = 0.0003$*; Wilcoxon Rank Sum test).

Overall, these results support our previous findings shown in Figure 4.4, demonstrating significantly reduced burst amplitude in the beta frequency band. Correlation analysis of symptom severity with burst amplitude in the contralateral and ipsilateral motor cortices found no significant relationship after correcting for multiple comparisons.

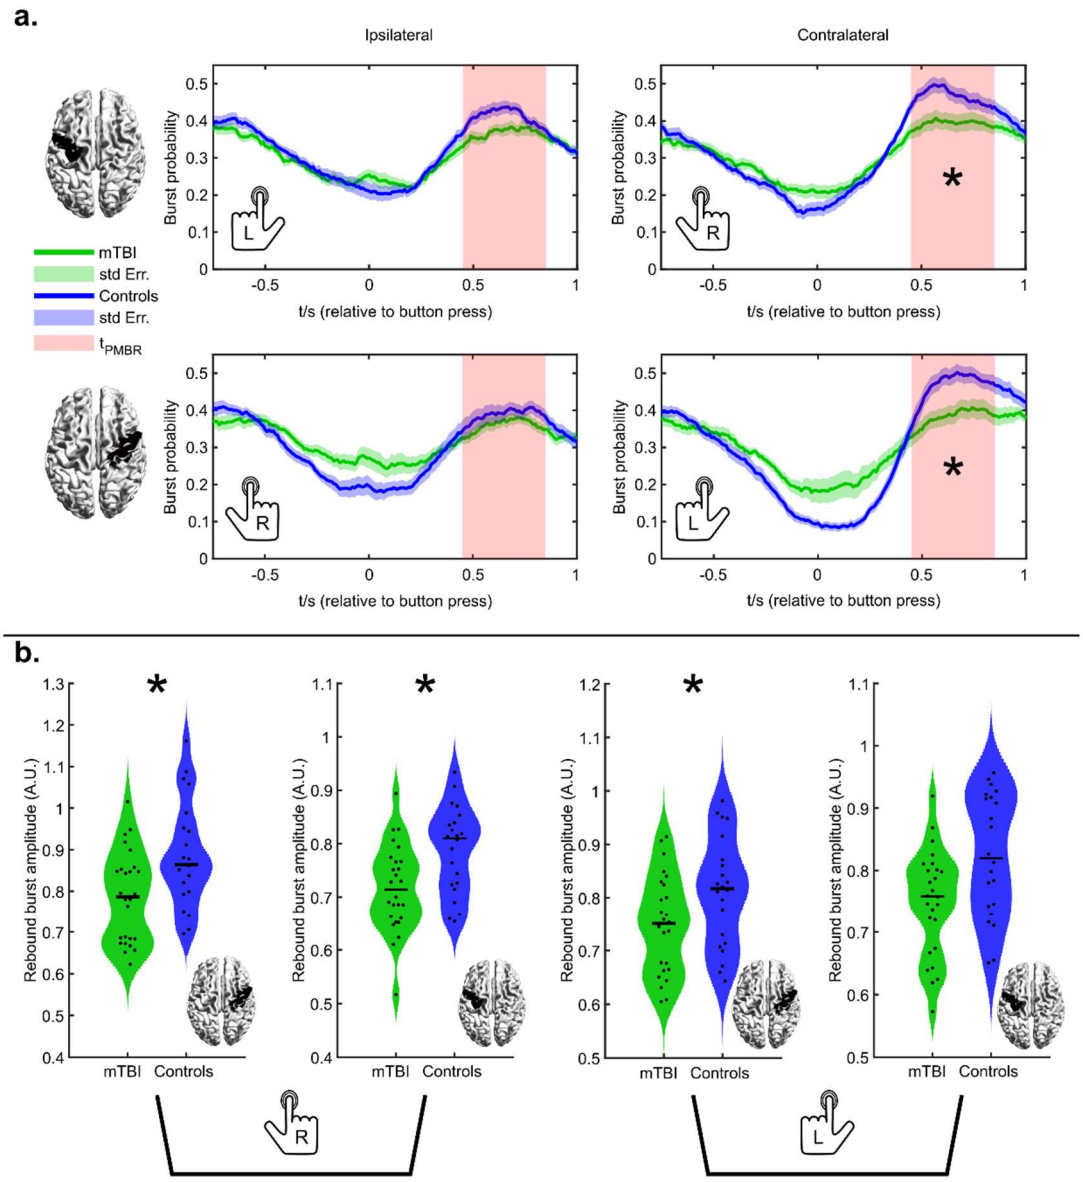

*Figure 4.6 Burst statistics during a motor task. a) Time courses showing the probability of bursts throughout a button press. The upper panel shows the case for the left primary motor cortex during a left hand (left) and right hand (right) movement. The lower panel shows the case for the right primary motor cortex during a right hand (left) and left hand (right) movement. b) Burst amplitude in the beta band during the post-movement beta rebound. Left-hand plots show contralateral and ipsilateral cortices during a left-hand button press. Right-hand plots show ipsilateral and contralateral cortices during a right-hand button press. * indicates statistical significance ($p < 0.05$, Wilcoxon Rank Sum test) following FDR correction for multiple comparisons.*

Figure 4.7 shows transient functional network patterns during finger movement. Figure 4.7a shows the temporal evolution of burst coincidence between the left and right primary motor cortex, during the task. The green trace shows the patients and the blue trace shows healthy controls. In agreement with previous work (O'Neill et al., 2015b), we observe lower connectivity during unilateral movement, followed by an increase on movement cessation; this supports the theory that the PMBR is a time of elevated connectivity between primary cortices and other regions (Tewarie et al., 2019), and carries a top-down inhibitory influence. The largest difference between patients and controls, in terms of connectivity, occurs during the post-movement rebound period. To test this statistically, the violin plot on the right-hand side shows connectivity estimated during the rebound window, demonstrating a significant ($p = 0.021$; Wilcoxon rank sum test) reduction.

Finally, Figure 4.7b shows the spatial signature of the dominant 2% of functional connections at 4 selected time points during the task. It is interesting to note how this spatial pattern changes considerably during the task with posterior parietal and sensorimotor connections, apparent at the start, giving way to frontal connectivity, which then, in turn, gives way to dominant interhemispheric motor network connectivity during the rebound window. The upper set of images in Figure 4.7b show the case for patients, the middle set for controls, and the lower set show the highest 2% of differences between the two groups. Note that these differences map out a clear motor network around the time of the post-movement rebound, showing a clear and significant reduction in burst coordination within this time window. No significant correlation between symptom severity and PMBR connectivity (between the motor cortices) was found (combined groups: $R = -0.296$; $p = 0.043$, patients only: $R = -0.005$; $p = 0.98$; Spearman Correlation. See Appendix D for a scatter plot).

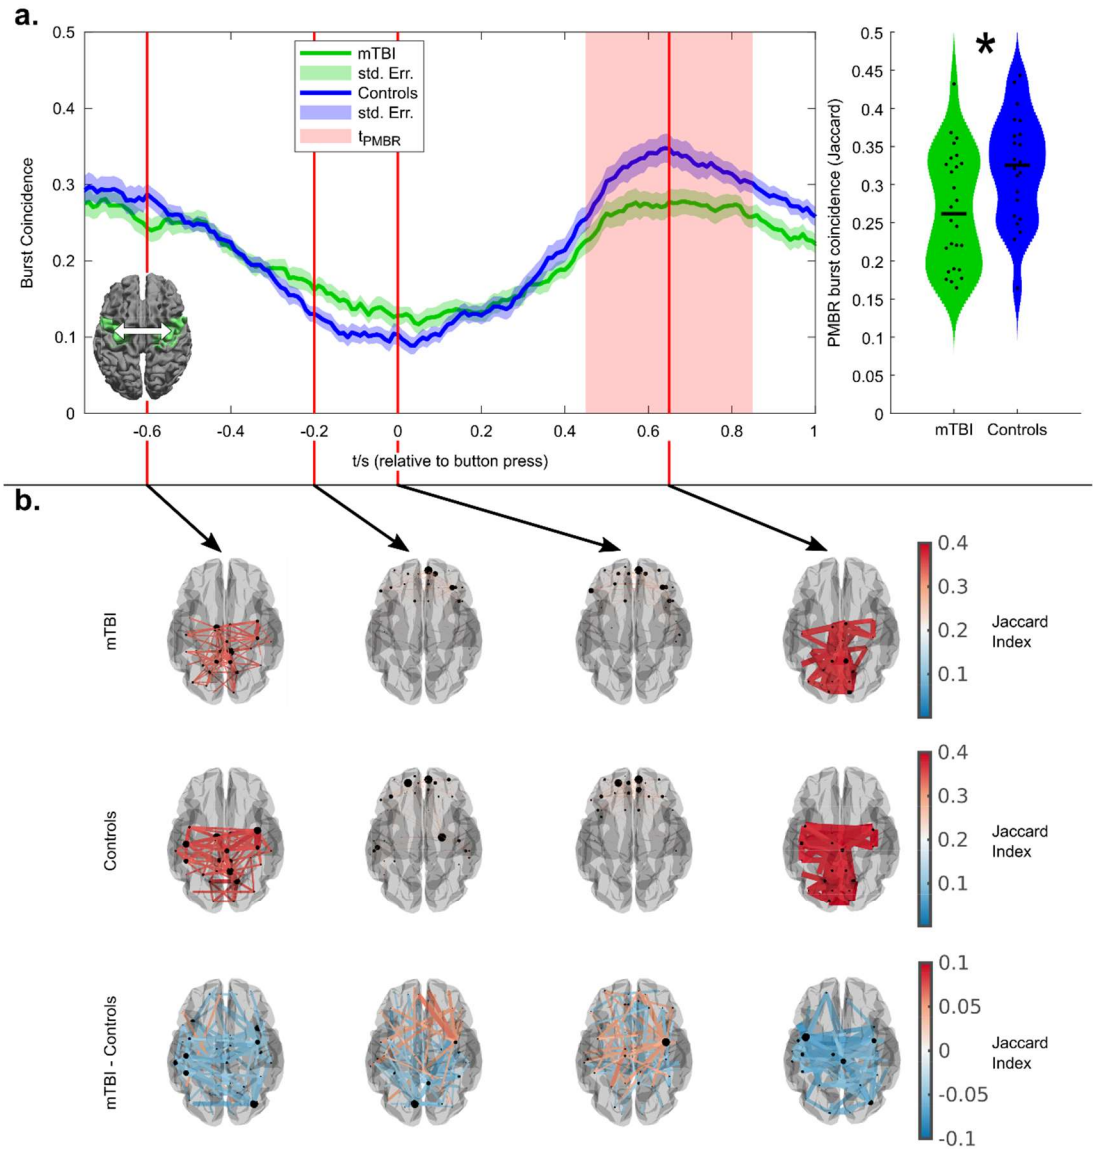

Figure 4.7 Functional connectivity during a motor task. a) Time course showing the temporal evolution of burst coincidence throughout finger movement. Note that burst coincidence is significantly ($p = 0.021$; Wilcoxon rank sum test) less likely during the beta rebound window, in patients relative to controls. b) The spatial distribution of dominant connections at 4 time points during the task. The upper set shows the case for patients. The centre set shows the case for controls, and the lower set shows the dominant differences. In all cases, the 2% of connections with the highest values are shown. Notice that during the rebound window, the dominant differences are between bilateral motor regions. See the online supplementary materials for (Rier et al., 2021) for a video showing the whole time evolution of the functional connections.

4.4 DISCUSSION

4.4.1 A summary of findings:

Our resting-state findings demonstrated abnormalities in both beta burst amplitude and the coincident bursting that mediates connectivity. In Figure 4.4, results show clearly that the diminishment of beta amplitude in patients is not general over all time; rather it is mostly a property of bursts. Indeed, when the burst state was active, significant differences between patients and controls could be found which were not mirrored during non-burst windows; this means that lower beta amplitude is occurring, approximately one-third of the time. It is important to note that the null result in the non-burst windows does not necessarily imply that there is no information of interest outside the bursts; indeed, a more subtle analysis, perhaps looking at individual brain regions, may well demonstrate significant differences between patients and controls. Nevertheless, at least at the whole-brain level, our results suggest that the largest difference in beta amplitude between patients and controls occurs during the active burst state.

Again, using resting-state data, we were able to show that the previous findings of reduced connectivity (Zhang et al., 2020)—as measured by amplitude envelope correlation—are driven by a drop in burst coordination across the brain (see Figure 4.5). Whilst we showed that connectivity over the whole brain was diminished in mTBI, post hoc analyses suggested that the most affected connections were interhemispheric, between regions of polymodal parietal cortices. These regions are well-known foci of sensorimotor integration and contain dense connections with numerous brain areas (Whitlock, 2017). We saw a significant correlation between connectivity and symptom severity in patients, following the selection of a subset of connections using the rRF approach. This demonstrates the utility of this data-driven method to determine which connections are most affected. This result should be used to inform the hypotheses of future studies, potentially leading to stronger results.

In our task data, we probed more specific hypotheses relating to the MRBD and PMBR. These effects have been observed in electrophysiological imaging data for many years and are one of the most robust findings in neuroscience. Recent work (Tewarie et al., 2019) has linked the PMBR to a period of enhanced connectivity within a wider network encompassing primary motor and motor planning regions, with the implication that this network provides top-down inhibitory influence to "shut down" and recalibrate the motor cortex post-movement. Given the global and inter-hemispheric nature of this network, we hypothesised that diffuse white matter damage might make the PMBR abnormal in mTBI patients. Results supported these hypotheses, showing diminished modulation of burst probability, lower amplitude, and a deficit in coincident bursts, during the PMBR in patients. Whilst it remains a possibility that this effect was due to the way in which the task was performed (e.g., it's conceivable that patients and controls may move differently), there was no measurable difference in reaction times between groups and so it's likely that the observed deficit is a fundamental feature of mTBI. Interestingly these motor-specific findings did not correlate significantly with overall symptom severity, however, this may be because our symptom scores are global, and we did not correlate them with an objective assessment of motor function specifically. Previously, mTBI sequelae have been demonstrated to include motor deficits (Pearce et al., 2018), and these findings represent a putative mechanism.

4.4.2 Mechanistic interpretations: the importance of beta-band phenomena

Despite the unmet clinical need for rapid and reliable biomarkers for mTBI, imaging signatures for concussive injuries are difficult to establish given the heterogeneity of mechanism, injury, symptoms, and long-term outcome. Neurophysiological indices derived from MEG have shown significant promise in their sensitivity and specificity in the assessment of mTBI. Previous work (Huang et al., 2014) has demonstrated that abnormal neural oscillations,

particularly in the low-frequency delta band, can be used to successfully distinguish mTBI patients from healthy controls. More recent works (Huang et al., 2020; Zhang et al., 2020) show that higher frequency oscillations (beta and gamma) can also be reliable markers; as well as emerging measures of electrophysiological dynamics including functional connectivity. However, these previous measures have assumed the 'classical' picture that oscillations apparent in time-averaged brain signals represent ongoing rhythmic processes. In this chapter, we used the emerging theory of pan-spectral bursts—transient, temporally discrete brain dynamics—elucidated using a Hidden Markov Model, to further this field and show the multiscale, multiterminal impact of a single mTBI on dynamic neural activity. Our results, which show diminished intrinsic and spontaneous bursting dynamics in the beta range, suggest that the neuronal assemblies responsible for generating transient beta events are affected by mTBI. This finding mirrors, in part, recent findings in fMRI which report altered dynamics and signal complexity after mTBI (Churchill et al., 2020). A decrease in burst amplitude might be caused by a reduction in excitatory drive from the thalamus (Sherman et al., 2016) which would support the notion that long-range connections to the cortex are affected by concussive events. Our finding that diminished beta band connectivity in mTBI is due to a lack of burst coordination in long-range networks (at rest and during a task), particularly within inter-hemispheric connections in the parietal lobes, could be explained by a variety of mechanisms. These could include axonal damage with possible subsequent Wallerian degeneration, reduced myelination and other forms of neurodegeneration affecting neural synchrony.

Changes in beta band power and disruption of beta band connectivity at rest, as well as beta abnormalities during motor tasks, have been reported in patients with neurodegenerative disease such as Parkinson's disease (Bosboom et al., 2006; Heinrichs-Graham et al., 2014; McColgan et al., 2020; Stoffers et al., 2007). Recent reports on a cohort of subjects with repeated sports-related mTBI provide evidence of increased inflammation and accumulation of Tau protein

(which is thought to contribute to the development of neurodegenerative diseases) (Marklund et al., 2021). Given the likely connection between mTBI and chronic neurodegenerative disease (Gardner and Yaffe, 2015) the beta band is of particular relevance to mTBI. This relevance is further implied by studies (e.g. Brookes et al., 2011; Hipp et al., 2012) demonstrating that canonical brain networks—including those associated with sensory/motor function, and the higher-order attentional networks—can be elucidated via an analysis of beta-band oscillations. These findings are in broad agreement with other emerging theories that beta oscillations carry top-down inhibitory influence on primary cortices. They also agree with a predictive coding model which suggests that beta oscillations carry internal forward models (whilst gamma oscillations reflect prediction errors for example (see e.g. Bastos et al., 2015; Buschman and Miller, 2007; Kopell et al., 2000). Taken together, we propose that beta-band processes support, and offer a sensitive measure of, long-range communication and connectivity. In mTBI, it is hypothesised that diffuse axonal injury (DAI) (Gazdzinski et al., 2020) and in particular disruption of the white matter around the corpus callosum, is one of the main drivers of neuropathology. This being the case it is likely that the long-range networks (particularly interhemispheric) mediated by beta oscillations would be disrupted by injury. It is therefore likely that assessment of beta-band phenomena would provide a marker of mTBI. Indeed, such a marker would reflect the disruption of the connectome (either localised coup- and countercoup injury or diffuse injury) which is likely to exist in mTBI (Browne et al., 2011; Kirov et al., 2013) but is too subtle to be measured reliably using clinically available structural imaging.

Additionally, mTBI can cause a plethora of neurochemical changes, impact neurotransmission, and cause long-lasting metabolic impairment (see MacFarlane and Glenn (2015) for a review), which could disturb the effective generation of burst activity. Again, this is measurable using functional imaging but will escape detection in structural imaging. An important consideration here is time post-injury; neurochemical and neurometabolic changes, and their

recovery to normal levels, vary in time scales, with some hyperacute effects resolving within minutes (e.g. K^+ ion efflux and altered glutamate levels) and others remaining altered for several days (e.g. Ca^{2+} accumulation; Churchill et al., 2020; Giza and Hovda, 2014; Jamjoom et al., 2021; MacFarlane and Glenn, 2015). In the long term, inflammation and glial activation in the subcortical white matter but also changes in myelination, Wallerian degeneration, or oedema, might disrupt functional connectivity. Longitudinal, multimodal studies covering time points soon after injury are required to produce a clear picture of how electrophysiological signals are altered during recovery.

The burst model of beta band activity has gained traction in the neuroscientific community in recent years, not only as a conceptual model but moreover because it has spawned some of the first mathematical models suggesting how beta oscillations might be generated physiologically (Sherman et al., 2016). This, therefore, affords the possibility that, in time, we can link the putative beta band markers of mTBI to a neurophysiological model. One consequence of the bursting hypothesis is that beta effects, whilst typically thought to be ongoing over all time, only occur for a small percentage of the overall duration of a MEG scan. It is with this in mind that we might expect a bursting framework to offer a more sensitive measure than classical analyses, where the effects sought might become averaged out over time.

It is worth noting that while the selection process used to define our burst state depends on beta band activity, the rich spectral information offered by our HMM shows that lower frequencies are dominant in the burst state spectrum (see Appendix E for example spectra), although group differences are most prominent in the beta range.

4.4.3 Clinical perspective

Previous papers have tended to focus on differentiation between individuals with mTBI and controls, with recent findings offering between 80 and 100%

classification accuracy using ML. In the current work, we have also employed an ML approach—applied only to the connectivity data—which shows a reasonably high classification accuracy. This shows that the putative biomarkers of mTBI outlined here might also offer diagnostic capability. However, mTBI remains a diagnosis based on clinical history, patient experience and self-reported symptoms, and is ultimately delivered by medically qualified clinicians. A brain scan, however accurate, that adds nothing other than to confirm the diagnosis, is of little practical help when determining a treatment pathway. Instead, we would argue that 1) the ability to understand the sequelae of mTBI using hypothesis-driven, objective assessment and 2) the ability of an imaging modality to be able to predict patient outcome or rehabilitation needs are significantly more important. MEG holds some promise in this area, but to realise its potential, we must employ the most sensitive markers of illness in a multivariate framework. While it is difficult to prevent overfitting of classifiers trained on relatively small datasets, employing feature selection techniques, such as the random forest approach outlined in this study, will be crucial in selecting candidate markers and will also aid hypothesis generation for future research.

4.4.4 Limitations

There are several limitations to the current study which should be expounded. In terms of methodology, source localisation is impacted by the ill-posed nature of the MEG inverse problem. Brain regions, particularly those in close proximity, can exhibit signal leakage which will artificially inflate connectivity measures. To mitigate this, we applied orthogonalisation which eliminates zero phase lag effects; this is known to reduce artefactual connectivity but at the expense of real zero-phase-lag interactions (Brookes et al., 2012). This means that, whilst the long-range connections that we observe to differ between groups are real, there may be other (particularly short-range) connections that have been missed in this study due to this methodological limitation. A second technical

limitation relates to the identification of the beta burst state. The method uses a univariate HMM in which three states are identified for each brain region. Following this, the state whose temporal occurrence most strongly correlates with high beta power is designated the burst state. This means that, hypothetically, if a brain region was dominated by noise, a burst state would still be selected (i.e., as the one most strongly correlated with the beta envelope, even if that correlation was weak) and could be meaningless. Related, two regions exhibiting a coincident burst may not necessarily be exhibiting a functional connection. This is a fundamental limitation of the univariate nature of the method, however, we note that this same methodology has been applied successfully in previous papers (Gascoyne et al., 2021; Seedat et al., 2020). It is known to reliably extract connectome metrics (based on burst coincidence), and coincident bursts have been demonstrated to occur during periods of increased phase locking. For these reasons, it remains likely that the method offers a true picture of oscillatory dynamics and connectomics.

Additionally, two experimental limitations should be mentioned. Firstly, our motor task employed relatively short inter-stimulus intervals. The beta rebound is known to last for more than 7s, and so the rebound following a trial will not be given a chance to fully relax before the following trial begins (Fry et al., 2016; Pakenham et al., 2020). We note that to gain a sufficient signal to noise ratio, a large trial count is desirable, however, maintaining a large trial count whilst also allowing the rebound to fully relax between trials would lead to an experiment that was too long for patients to tolerate. Consequently, the experimental design used, whilst potentially cofounded, is practical. For future work, optimisation of this paradigm or a more suitable task for the investigation of motor function should be considered.

A further limitation was the sample size used in our ML analysis. Using large numbers of features and small sample size, it is conceivable that overfitting might occur when applying ML techniques such as the ones we have presented

here. However, high classification accuracy does not imply overfitting in every case. As demonstrated in previous work which used a similar ML pipeline and sample size (Zhang et al., 2016), the combination of feature reduction and 10-fold cross-validation allowed us to train SVMs that correctly classified the testing samples. Importantly, there was no information leakage between training and testing samples during 10-fold cross-validation as feature selection and reduction were conducted without including the testing samples during each of the folds. We are confident that the high classification accuracy is representative of what could be achieved in a larger study. However, we acknowledge that the relatively small sample size studied here remains a limitation; further samples are needed to ultimately validate this, and indeed all mTBI MEG findings. Our study was also limited by the behavioural measures that were acquired. Specifically, the SCAT2 is a generic questionnaire-based assessment that covers many symptom domains. It remains a useful measure of symptomology but is not specific. This is likely why within-group correlations between MEG derived metrics were weak (in the case of connectivity), or absent (in all other cases). In the task, it is very unlikely that such a broad and general symptom score would correlate with the subtle MEG metrics that only relate to a single brain network. In future studies, behaviour should be more accurately assessed, potentially using motor specific tasks/questionnaires to better assess the relationships between neural measures and symptoms.

Finally, the choice of study participants should be considered. As laid out in Chapter 1, the majority of studies investigating the neural activity of subjects with exposure to mTBI focus on chronically symptomatic cases, with MEG data acquired several months after the initial injury. While the present study did not exclude subjects based on the presence or absence of symptoms, the majority of MEG data for the mTBI group were collected in the subacute phase of injury (several weeks to a few months post mTBI) from subjects with *persistently* more severe symptoms than the healthy controls. Given the moderate sample size and the wide range of time post-injury, it is difficult to generalise these results to

other mTBI populations. It is unclear at which phase these abnormalities develop. Given the dearth of high-quality longitudinal research, studies at time points in the acute regime are needed to elucidate the evolution of pathology post mTBI. The study in Chapter 5 will address some of these weak points.

4.5 CONCLUSION

We show that fundamental and mechanistically plausible neural bursting phenomena in the brain are disrupted by a single mTBI in the sub-acute phase following injury when structural MRI is normal. This dysregulated neural repertoire exists intrinsically and during dynamic recalibration of the motor system after behavioural output. These results point towards a mechanism whereby white matter damage disrupts network function; whilst the damage itself may be too subtle for structural imaging to see, its functional consequences are accessible using MEG.

Chapter 5 THE MEGABIT STUDY

As laid out in Chapter 4, the transient burst framework offers an exciting new way to investigate potential signatures of mTBI in the MEG signal. However, the choice and size of the study cohort leave common shortfalls in the mTBI-MEG literature identified by Allen et al. (2021) unaddressed. Firstly, subjects are typically scanned a long time after injury and, secondly, by contrasting the patient cohort with healthy controls, most studies do not account for a potential non-specific trauma effect. Here, we present "The role of MEG in Assessment and diagnosis in mTBI" (MEGAbIT), an observational study of subjects either exposed to mTBI or non-specific orthopaedic trauma in the acute phase—less than 14 days post-injury—as well as a group of healthy controls from an existing normative dataset. Using resting-state MEG data, we applied our HMM-based analyses to explore whether our previous findings of reduced burst amplitude and burst coincidence connectivity would also hold in the acute regime, and be specific to mTBI rather than non-specific orthopaedic trauma. We found that burst statistics are stable for both studies—even though data were acquired at different sites, using different scanners—and replicated our finding of reduced burst connectivity. Our analyses did not reveal the previously reported abnormalities in burst amplitude. In addition, we also employed a more conventional analysis of delta power. Although excess low-frequency power appears to be a promising candidate marker for persistently symptomatic mTBI, insufficient data exist to confirm this pattern in acute mTBI. We found abnormally high delta power to be a sensitive measure for discriminating mTBI subjects from healthy controls, however, similarly elevated delta amplitude was found in a cohort with orthopaedic trauma, suggesting that excess delta may not be specific to head trauma, at least in the acute stage of injury.

5.1 INTRODUCTION

The current literature concerning potential MEG-based markers for mTBI can be separated into two overarching themes: Abnormal low-frequency oscillations, and abnormal functional connectivity (Allen et al., 2021), with the majority of studies focusing on resting-state data. Low-frequency abnormalities such as increased delta-band power appear to be the most commonly studied, however, the methods used to reconstruct source-space signals and/or to extract power, as well as the choice of inclusion criteria vary widely between research groups. A similar phenomenon is apparent in the connectivity literature where a variety of functional connectivity metrics and graph theoretical measures have been reported. Allen et al. also show that many studies utilised small sample sizes (see Chapter 1) and were subject to a moderate to high risk of bias. The majority of published results were collected long after exposure to mTBI and—given the lack of longitudinal data—leave the full evolution of pathology after mTBI unexplored.

As discussed in Chapter 3 (section 3.4) and Chapter 4, HMMs can give us additional insights into the temporal dynamics of the MEG signal that are ignored in conventional analyses. Not only do we gain information about the amplitudes of punctate events but changes in their probability of occurrence and duration can also be measured; their temporal coincidence can be used as a measure of functional connectivity. In Chapter 4 we showed that burst measures derived from our HMM analysis reveal abnormalities in a cohort of mTBI patients in the subacute stage of injury (up to three months post-injury), adding a further functional connectivity metric to the plethora of existing measures in need of replication.

From (Allen et al., 2021) we can conclude that while low-frequency power and functional connectivity appear to be the most commonly studied potential MEG markers of mTBI, two important questions are left unanswered:

1. Are there measurable brain changes in the acute to subacute stage (less than 14 days post-injury), before long-term symptoms can be reported?
2. Given that most studies use healthy individuals as controls, are these changes due to non-specific trauma, i.e. the secondary effects of any injury, not just mTBI?

To address these questions, here we aimed to reproduce our previous findings of abnormal burst connectivity (Chapter 4) and replicate previously reported abnormal delta power in a new dataset collected from subjects scanned in the acute phase of injury—within two weeks of presentation to the emergency department with mTBI.

Using resting-state MEG data recorded as part of the MEGAbIT study, as well as a normative dataset of age-matched healthy controls from previous resting-state studies, we performed two main types of analysis:

- Hidden Markov modelling—as performed in Chapter 4—to assess whether abnormal resting-state bursting metrics can be also observed in the acute stage.
- "Conventional" measurement of delta-band power in source reconstructed MEG data and analysis using a maximum Z-score method described by Huang et al. (2014).

5.1.1 Chapter outline

The following section will give methodological details on the design of the MEGAbIT study, data acquisition and preprocessing steps that were common to both the HMM and delta-power analysis pipelines. This will be followed by a description of the HMM analysis pipeline and discussion of results (section 5.3) before section 5.4 will outline the methods, and discuss the results, of the delta-power analysis.

5.2 METHOD

5.2.1 The Role of MEG in Assessment and Diagnosis In mTBI (MEGAbIT)

MEGAbIT is a single site, case-control observational study, aimed to characterize the clinical utility of functional and structural brain imaging measures, including task and resting-state MEG and several MR imaging sequences.

5.2.1.1 *Study design*

Two groups of participants were recruited from the Queens Medical Centre Emergency Department in Nottingham:

mTBI:

Adult subjects who have suffered an mTBI (40 participants planned)

Trauma Controls (TC):

Age and sex-matched non-head trauma controls who have suffered injuries requiring hospitalisation for less than 24 hrs (20 participants planned).

All volunteers attended a single session at the Sir Peter Mansfield Imaging Centre in the acute stage of injury (< 14 Days post-injury). To minimise potential incidental MR findings and to ensure homogeneity of the MEG data, an upper age limit of 35 years was chosen. Table 3 contains the full inclusion and exclusion criteria for recruitment.

Inclusion Criteria	The volunteer is willing and able to give informed consent to participate in the study.
	Age 18-35.

	Diagnosed in ED with mTBI(no abnormality on clinical structural imaging, Loss of consciousness (LOC) $< 30min$, amnesia for $< 24\ hrs$, Glasgow Coma Scale (GCS) ≥ 13 at all times and 15 within 24 hrs. OR Diagnosed in the Emergency Department (ED) with non-head trauma, matched for age and sex with mTBI group.
Exclusion Criteria	The patient requires $> 24\ hrs$ hospitalisation.
	The participant has a contraindication to undergo 7T MRI.
	The participant is unable to read text on a PC screen at $\sim 1m$ without glasses.
	Pregnancy.
	Any neurological, developmental or psychiatric disorders e.g. brain tumour, stroke, epilepsy, Alzheimer's disease, schizophrenia, post-traumatic stress disorder, major depressive disorder, bipolar disorder or history of learning disability.
	Previous hospital attendance with a traumatic brain injury.
	Substance or alcohol abuse within six months of enrolment.
	Taking of medications that could alter MEG signals: opioids and synthetic opioids (excluding codeine and dihydrocodeine), anti-epileptic drugs, sedatives, neuroleptics, and hypnotics.
	Extensive metal dental hardware e.g. braces and large metal dentures (excluding fillings), implanted medical devices or other metal objects in the head, neck, or face areas that although they

	hold no risk to participants during a MEG recording may cause non-removable artefacts in the MEG data.
	Participants who have participated in another research study involving an investigational product in the past 12 weeks.
	Any other significant disease or disorder, which, in the opinion of the Investigator, may put the participants at risk because of participation in the study, or may influence the result of the study, or the participant's ability to participate in the study.
Additional exclusion criteria due to the COVID-19 pandemic	Diagnosed or suspected COVID-19 infection (any symptoms of COVID-19 infection prior to the SPMIC visit).
	Participants requiring to self-isolate due to COVID-19 related government regulations

Table 3: Inclusion and exclusion criteria for the MEGAbIT study

5.2.1.2 *MEGAbIT Study Protocol*

Informed consent to study participation was ensured for all volunteers. Participants were screened for any contraindications to MEG or MRI scanning. Prior to any neuroimaging, participants were asked to complete a set of questionnaires to assess symptoms. The neurobehavioural symptom inventory (NSI) was used to get an estimate of the severity of impairment following the injury (Vanderploeg et al., 2015). This measure consists of 22 questions assessing affective, somatosensory, cognitive and vestibular symptoms (Meterko et al., 2012), each scored from 0 to 4. The sum of scores yields a severity score ranging from 0 (no symptoms) to 88 (maximum severity in all symptoms). Other questionnaires completed but not reported in this thesis include a depression checklist (PHQ-9, Kroenke et al., 2001), an assessment for

anxiety (GAD-7, Spitzer et al., 2006) and a checklist for post-traumatic stress (PCL-C, Ruggiero et al., 2003).

Cognitive impairment was assessed using a battery of neuropsychological measures which are laid out in Table 4. Results from these assessments are not included in this thesis.

	Assessment	Measured variable
(Rey, 1958)	RAVLT (Rey Auditory Verbal Learning Test)	Verbal episodic memory
(Wechsler, 2009, 1944)	WAIS-IV (Wechsler Adult Intelligence Scale IV)	Processing speed
(Reitan, 1971)	Trail Making Test	Executive functioning

Table 4: Battery of tests to assess cognitive impairment.

We also assessed recovery to normal life using the Glasgow Outcome Scale – Extended (GOSE, Wilson et al., 1998) and a questionnaire on healthcare utilisation.

Subjects performed three tasks while MEG data were acquired using a 275-channel CTF MEG. These included a 10-minute resting-state scan (eyes open), a 2-back working memory task (Huang et al., 2019), and a covert attention

switching task (Bauer et al., 2014; Hoogenboom et al., 2006). Only resting-state data were used for this thesis.

MR imaging was performed after the MEG recording using a 7T Phillips Achieva. Sequences acquired included a T1-weighted MPRAGE (1*mm* isotropic resolution) which was used for MEG source reconstruction, susceptibility-weighted imaging (SWI), diffusion tensor imaging (DTI), and fluid-attenuated inverse recovery imaging (FLAIR) for a subset of participants. For this thesis, only MPRAGE anatomicals were used.

Secondary outcome measures to assess recovery included a 3 and 6-month online assessment of self-reported symptoms and 2-back performance in an online experiment.

5.2.2 Data availability

At the time of writing, 41 participants have been recruited into the MEGAbIT study (31 mTBI and 10 orthopaedic trauma controls, henceforth denoted TC). No scans were acquired for 1 mTBI subject (who could not remove all jewellery); a further 3 participants with available data were excluded (see Figure 5.1). Recruitment was significantly disrupted during the COVID-19 pandemic, leading to reduced subject counts compared to the original goal set out in the study protocol. Recruitment of mTBI patients posed a considerable challenge even after the partial lifting of pandemic-related restrictions.

To produce an example of the challenges in recruiting this demographic: during three months of recruitment, 821 patients attending the Queens Medical Centre Emergency Department (Nottingham) were screened for possible mTBI. Out of 205 mTBI cases, 60 patients expressed interest in participating in the study, however, only seven subjects were successfully recruited during this period.

Given the availability of normative resting-state data from healthy controls (HC), data were analysed despite the small number of TC subjects available. The available HC datasets were sourced from two previous studies conducted using the CTF MEG scanner at the University of Nottingham:

The UK MEG partnership (UKMP) dataset: Sixty healthy subjects were scanned as part of an MRC funded initiative to amass a large normative database of resting-state MEG data (Hunt et al., 2016) (33 female, age 42 ± 12 years (mean $\pm$ std. deviation)). **The Neurodevelopment dataset:** A further 36 datasets from healthy volunteers were collected in a previous study conducted by (Brookes et al., 2018) (17 Female, age 22 ± 2 years).

Given the relatively wide age range present in the normative data (i.e. 9 to 67 years), a subset of these data was selected to provide an age- and sex-matched sample for the mTBI cohort included in the MEGAbIT study. The following procedure was used to choose an optimal set of controls:

First, any control subjects falling outside the age range of 16-40 were excluded. A cost matrix C was then constructed (number of mTBI subjects × number of normative datasets), setting each element ij to a cost value C_{ij} associated with assigning a control j to mTBI subject i. The cost values were chosen as follows:

$$C_{ij} = \begin{cases} \left|\Delta Age_{ij}\right|, & Sex_i = Sex_j \\ \infty, & \left|\Delta Age_{ij}\right| > 5 \\ \left|\Delta Age_{ij}\right| + 1000, & Sex_i \neq Sex_j \end{cases} \quad (5.1)$$

Here, ∞ was used to prevent large age differences between mTBI subjects and their assigned controls. A penalty of 1000 in addition to the age difference was chosen in cases where a potential control was not sex-matched. This was done to ensure sex matching for the majority of subjects while limiting age differences more strictly. When assigning a value of ∞ where there was no sex matching, no complete set of controls could be found within the chosen age constraints.

Using the cost matrix C, an optimal selection of healthy controls (HC) was found using the *Hungarian algorithm* (Kuhn, 1955) which minimises the total cost of assignment, or

$$\min_{R,L}(Trace(RCL)), \tag{5.2}$$

where R and L are permutation matrices for a particular assignment. Figure 5.2 shows the age and sex distributions for all three cohorts. The final number of subjects retained in the studies and details on age, sex and exclusions are shown as a flow diagram in Figure 5.1.

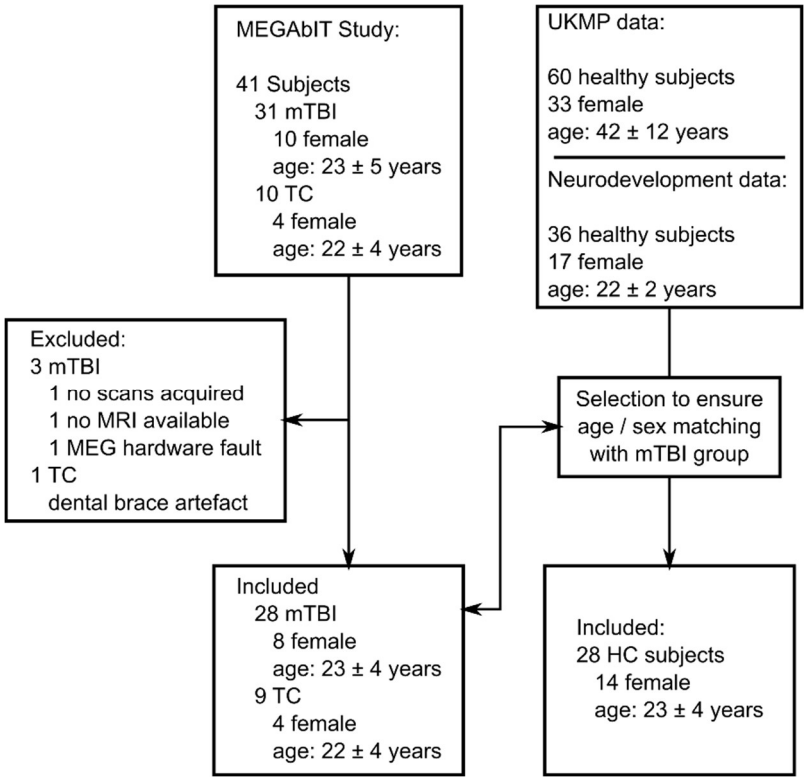

Figure 5.1 Flow chart showing group compositon and details about age and sex information as well as exclusions. Ages are given as mean ± standard deviation.

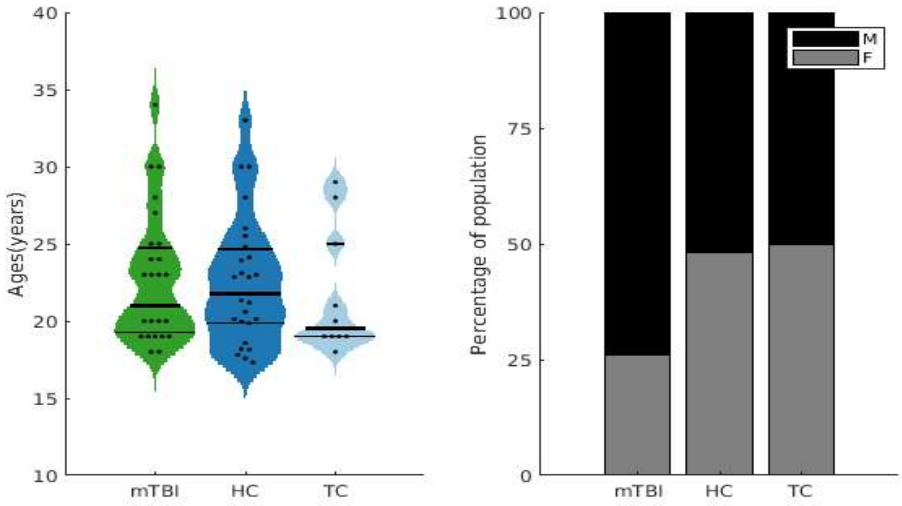

Figure 5.2 Age (left) and sex (right) distributions for each cohort included in the study. HC and TC denote the healthy controls from the normative dataset and TC the orthopaedic trauma controls recruited as part of the MEGAbIT study.

5.2.3 Scanning protocols

The resting-state scans for the **MEGAbIT** cohort lasted for 10 minutes. Subjects were sitting upright and with their eyes open. A fixation point in form of a red dot on a grey background was shown on a back-projection screen. Recordings were sampled at 600 Hz and a 150 Hz anti-aliasing filter was applied. For the **UKMP dataset**, subjects underwent a 5-minute resting-state recording, also sitting upright and with their eyes open, fixating on a red fixation dot on a grey background. The recordings were sampled at 1200 Hz with an anti-aliasing filter applied at 300 Hz. For the **Neurodevelopment dataset** subjects were scanned for 5 minutes (eyes open, supine) fixating on a marker, at a sampling rate of 600 Hz. Anti-aliasing at 150 Hz was applied here. Continuous head localisation data was available for all datasets. Note that all data was acquired on the same scanner, a 275-channel whole-head MEG system (CTF, Coquitlam, Vancouver, Canada).

5.2.4 Preprocessing

Where necessary, datasets were resampled to 600 Hz. The raw MEG data were pre-processed by applying the 3rd order synthetic gradiometer configuration (see section 2.3.3.2), mean correction over the entire recording, and a 4^{th} order Butterworth band-pass filter between 1 and 150 Hz. After splitting the data into 10s-long epochs, each recording was inspected and epochs with large artefacts like squid resets (section 2.4.1.3) were removed. Automatic head-movement rejection followed, removing trials containing movements $> 7mm$ from the average head position. The mean ($\pm$std. deviation) length of data retained was $524 \pm 76s$ for the mTBI group, $552 \pm 46s$ for the TC group, and $286 \pm 22s$ for the HC group.

5.2.5 Coregistration

In order to allow for source reconstruction, the position of each subject's head relative to the sensor array was found by recording the location of three fiducial coils—placed on the nasion and two preauricular points—which were energised continuously throughout the MEG recording. To find the positions of the fiducial coils relative to the individual subject's anatomy a Polhemus ISOTRAK was used to create a 3D digitisation of the head and face shape including the fiducial location. A T1-weighted anatomical MRI (MPRAGE sequence, 1mm isotropic resolution) was available for each subject from which the scalp was extracted and aligned with the head digitisation (see Chapter 2, section 2.4.2). This resulted in the positions of the fiducial points in anatomical space which could be used for source localisation.

During the COVID-19 pandemic, physical distancing and general hygiene guidelines made head digitisation using the Polhemus ISOTRAK impractical. For this reason, an alternative method was used for a subset of the individuals scanned during the MEGAbIT study. As described in Chapter 2 (section 2.4.2) a Structure IO camera was used to acquire a head digitisation at a distance while

the subject was wearing a swimming cap to obtain an approximate shape of the scalp surface. Coloured markers were attached to the fiducial coils to obtain a head digitisation including the fiducial positions. By extracting the points in the digitisation containing the coloured markers the fiducial locations were extracted and a digitisation file was generated to match the output from Polhemus-based digitisation using an in-house MATLAB script. Coregistration with the anatomical MR images proceeded as above with an automatic alignment of the digitisation with the scalp shape extracted from the anatomical MRIs.

5.2.6 Head motion assessment

Given the methodological difference between some HC subjects, which were scanned in a supine position, and the MEGAbIT cohort, which were sitting upright, we assessed whether HC displayed less head movement. This was especially relevant for the analysis of delta-band power in section 5.4, as head motion is likely to fall within or intersect with this frequency band.

For each subject, position data was extracted for each head-localisation coil and the magnitude of their displacement was calculated via the Euclidian distance from the position recorded at the beginning of the scan. Each displacement time-course was mean centred, linear trends were removed and a bandpass filter for the delta band was applied. We calculated the variance remaining in the filtered time-courses as a proxy for the delta power component in the motion data and averaged over the three head localisation coils.

5.3 HIDDEN MARKOV MODELLING

5.3.1 Analysis

5.3.1.1 *Pre-processing and source reconstruction*

Our HMM-based set of analyses was conducted on the resting-state data described above. Preprocessing, artefact rejection and removal of data segments with excessive motion were done as described in section 5.2.4. Artefact free data were notch filtered to suppress mains line noise and two harmonics using a hamming windowed sinc FIR filter (bandwidth $2Hz$) implemented in EEGLAB (version 2019_1 (Delorme and Makeig, 2004)). Source reconstruction was done using a multiple spheres head model (Huang et al., 1999) and the previously described beamformer with source orientation estimation via exhaustive search (see Chapter 3 section 3.2.1.1). As described in Chapter 4, 78 source locations were chosen based on the AAL Atlas (Gong et al., 2009; Tzourio-Mazoyer et al., 2002). Brains were extracted from each anatomical MRI using FSL BET (Smith, 2002), and aligned to the MNI152 template brain (Evans et al., 2012) using FSL FLIRT (Jenkinson and Smith, 2001). The inverse transformation matrices from this registration were used to transform the coordinates of the centroids for each region in the atlas (in MNI152 space) into the corresponding locations in each subject's brain. The beamformer was used to reconstruct data at each of these 78 locations. Wide-band data ($1\text{-}150Hz$)—all concatenated, clean epochs—were used to construct the covariance matrices needed for the beamformer which was implemented without regularisation. The resulting beamformer weights were used to calculate virtual time-courses at each atlas location using equation (3.28):

$$\widehat{\boldsymbol{q}}(t) = \boldsymbol{w}^T \boldsymbol{B}(t),$$

where $\widehat{\boldsymbol{q}}(t)$ is the reconstructed source time-course, $\boldsymbol{w}$ contains the beamformer weights and $\boldsymbol{B}(t)$ is the data matrix. In line with Chapter 4—and (Rier et al.,

2021; Seedat et al., 2020)—source time-courses were band-pass filtered to 1-48Hz and symmetric orthogonalisation was employed to correct for source leakage (Colclough et al., 2015).

5.3.1.2 *HMM and secondary analyses*

The HMM analysis pipeline described in Chapter 4 was applied. Briefly, using a time-delay embedded (TDE) observation model, a 3-state HMM was constructed separately for each subject and atlas region. Parameters were found using variational Bayesian inference. We classify one *burst state* by measuring the correlation between the time-courses of state probability and beta band amplitude (found using a continuous Morlet-wavelet transform). The remaining two states were classed as non-burst states. Binary time-courses were generated for each state by thresholding the state probability time-courses at 2/3 such that a value of 1 in a particular time-course coded for that state being active.

5.3.1.3 *Burst summary metrics*

From the HMM output, we generated the set of features described in Chapter 4, briefly:

- **Burst amplitude:** The maximum of the beta envelope during each visit to the state
- **Total burst time:** Time spent in the burst state per unit of time
- **Burst coincidence connectivity:** Functional connectivity was measured via the Jaccard index.

We examined whether the topographies of burst (and non-burst) amplitude and patterns of connectivity were qualitatively comparable with the results shown in Chapter 4 (and (Rier et al., 2021)), and measured group differences, aiming to establish whether the previously found deficit in amplitude and connectivity in sub-acute individuals with mTBI still applies in the acute stage, and to a cohort not restricted to exclusively male subjects. Additionally, we assessed

whether any abnormalities relative to healthy controls (HC) were specific to individuals with mTBI by comparing with the summary metrics of our orthopaedic trauma control group (TC).

5.3.1.4 *Feature selection and classification*

The connectivity matrices produced using our burst coincidence metric were analysed using the same machine learning pipeline described in Chapter 4. First, Recursive Random Forest Feature Selection (rRF-FS (Zhang et al., 2020), Appendix B) was used to determine the subset of functional connections which best distinguish subjects with mTBI from healthy controls. Second, a support vector machine (SVM) was trained for binary classification between groups to evaluate the utility of the features chosen by the rRF-FS algorithm to distinguish mTBI and HC subjects (see Appendix C). We repeated the approach from Chapter 4 and trained our ML models only using mTBI and HC subjects, excluding the TC group. As before we assessed whether the selected features improve differentiability between mTBI and HC by comparing the mean connectivity over the selected features. Additionally, we explored whether this measure would also better distinguish mTBI from the unseen TC subjects, aiming to gain information about the generalisability of the feature selection algorithm to unseen data.

5.3.1.5 *Statistical testing*

As described in Chapter 4, no a-priori assumptions were made about which brain regions or connections may be implicated in the injury sustained by each individual as sufficient reliable data about the site of impact was rarely available. We, therefore, averaged the burst amplitude, non-burst amplitude, and time spent in burst metrics across all 78 atlas regions; connectivity values were averaged across all unique pairs of regions. Group differences were calculated (mTBI vs TC, mTBI vs HC, TC vs HC) and assessed for statistical significance using the non-parametric Wilcoxon rank-sum test, corrected for multiple

comparisons using the Benjamini-Hochberg procedure (Benjamini and Hochberg, 1995). The amplitude measures were corrected for six measures (three pairs of groups over two measurements—burst and non-burst amplitude). The connectivity measures were also corrected for 6 comparisons (three pairs of groups comparing global connectivity and average connectivity over RF-FS connections only).

5.3.2 Results

5.3.2.1 *Burst statistics*

Figure 5.3a shows the spatial topography of the beta amplitude for the burst state and non-burst states in the upper and lower row of panels, respectively. The burst state plots show the familiar pattern of increased amplitude adjacent to the central sulcus in all three groups. Note that the range of values observed is comparable to the amplitudes measured in Chapter 4. As expected, the average beta amplitude was comparatively lower in the non-burst states. The violin plots in Figure 5.3b show the proportion of time spent in the burst state for all three groups which appear comparable between groups as well as with the previous results in Chapter 4. Statistical testing between each pair of groups via Wilcoxon rank-sum tests yielded no significant differences.

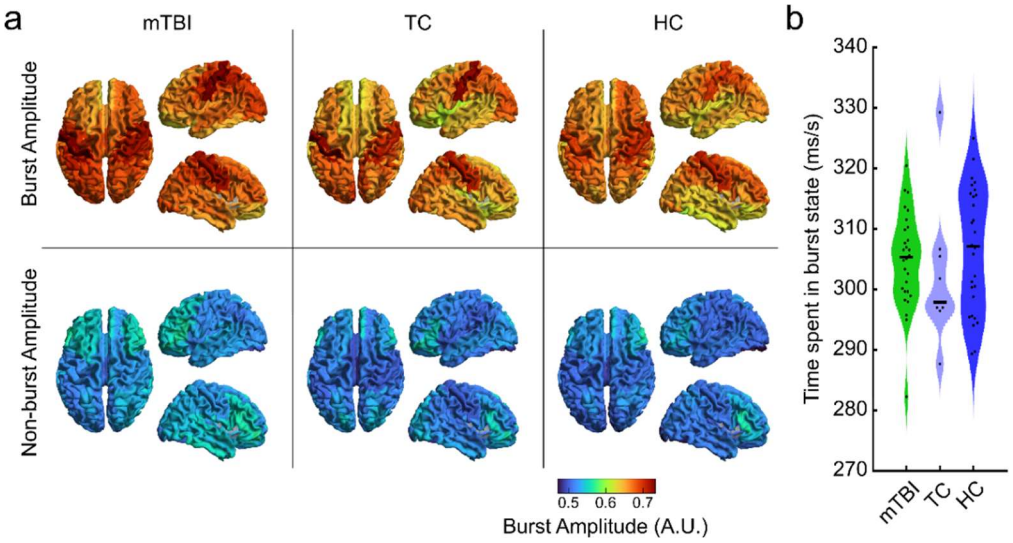

Figure 5.3 Burst statistics. a) Brain plots showing the spatial distribution of the burst (top) and non-burst (bottom) amplitude—the maximum amplitude of the beta band envelope averaged over the burst and non-burst state visits. b) Violin plot showing the proportion of time spent in the burst state for all three groups. No significant differences were found after performing a Wilcoxon rank-sum test between each pair of groups.

Figure 5.4 shows group comparisons of burst and non-burst amplitude. The violin plots in panel a show the global mean amplitude for burst and non-burst states. No significant difference between groups was found for the burst state (Wilcoxon rank-sum test), contrary to the hypothesis of finding reduced burst amplitude as seen in the subacute group (Chapter 4). After correcting for multiple comparisons, the mTBI group showed increased non-burst amplitude compared to healthy controls (Wilcoxon rank-sum test, $Z = 3.35$; $p = 0.0008$). Panel b of Figure 5.4 shows the distribution of differences between mTBI and TC (left column) and mTBI and HC (right column) for the burst state (top row) and non-burst states (bottom row).

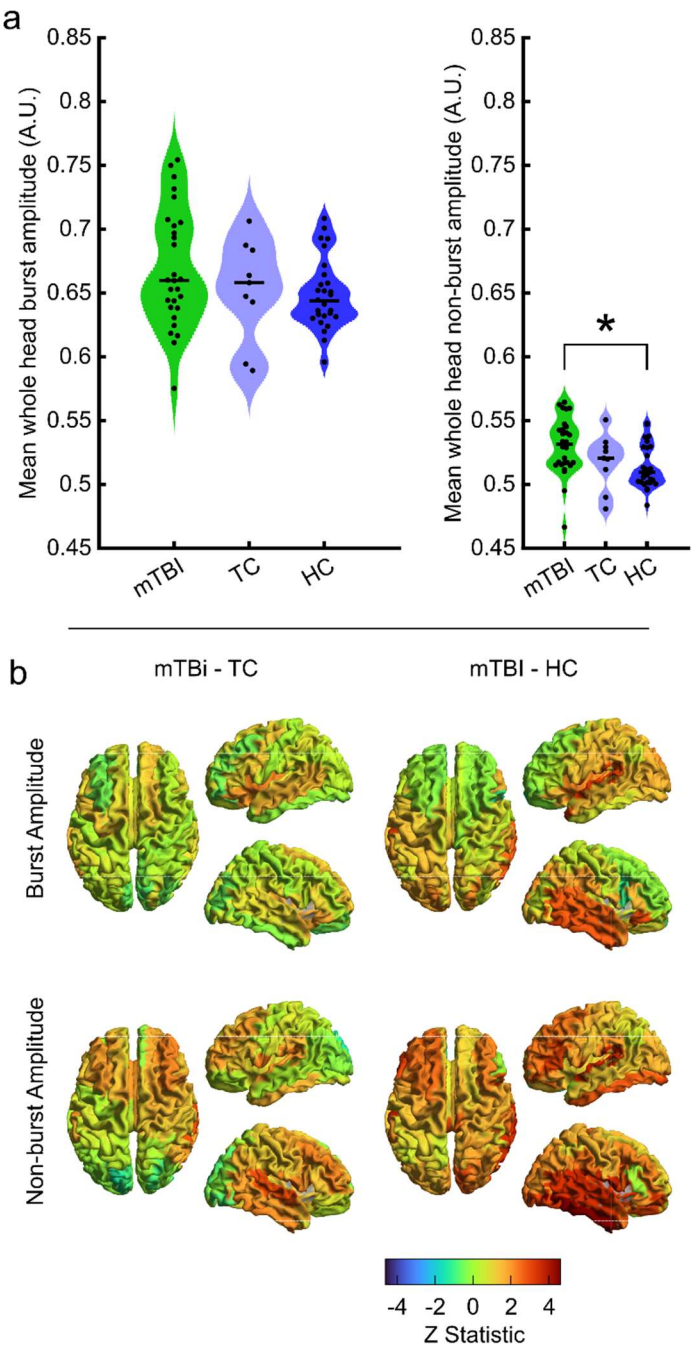

Figure 5.4 Group comparisons of burst and non-burst amplitudes. a) Violin plots showing global mean burst amplitude (left) and non-burst amplitude (right). Wilcoxon rank-sum testing showed a significant difference between non-burst amplitude in mTBI vs HC (Wilcoxon rank sum test, $Z = 3.35$; $p = 0.0008$). b) Distribution of differences between mTBI and TC (left column) and mTBI and HC (right column) for the burst state (top row) and non-burst states (bottom row).*

5.3.2.2 *Burst Coincidence Connectivity*

Figure 5.5a shows the connectivity matrices obtained using the burst coincidence Jaccard index, as well as the corresponding glass brain plots showing the top 5% of strongest connections. The network patterns in all groups match the ones previously described in Chapter 4 and published work by (Seedat et al., 2020), with the strongest connections being concentrated in the occipital, parietal and motor-sensory cortices. Similar to the observations made in the subacute mTBI cohort investigated in Chapter 4, connectivity strength is reduced in mTBI compared to controls. Figure 5.5b shows the 2% of connections with the largest absolute difference between mTBI and the two control groups. Figure 5.5c demonstrates this via a violin plot of the global mean connectivity for all subjects. We observe a statistically significant reduction in connectivity in mTBI compared to healthy controls (Wilcoxon rank-sum test, $Z = -2.34; p = 0.020$*).

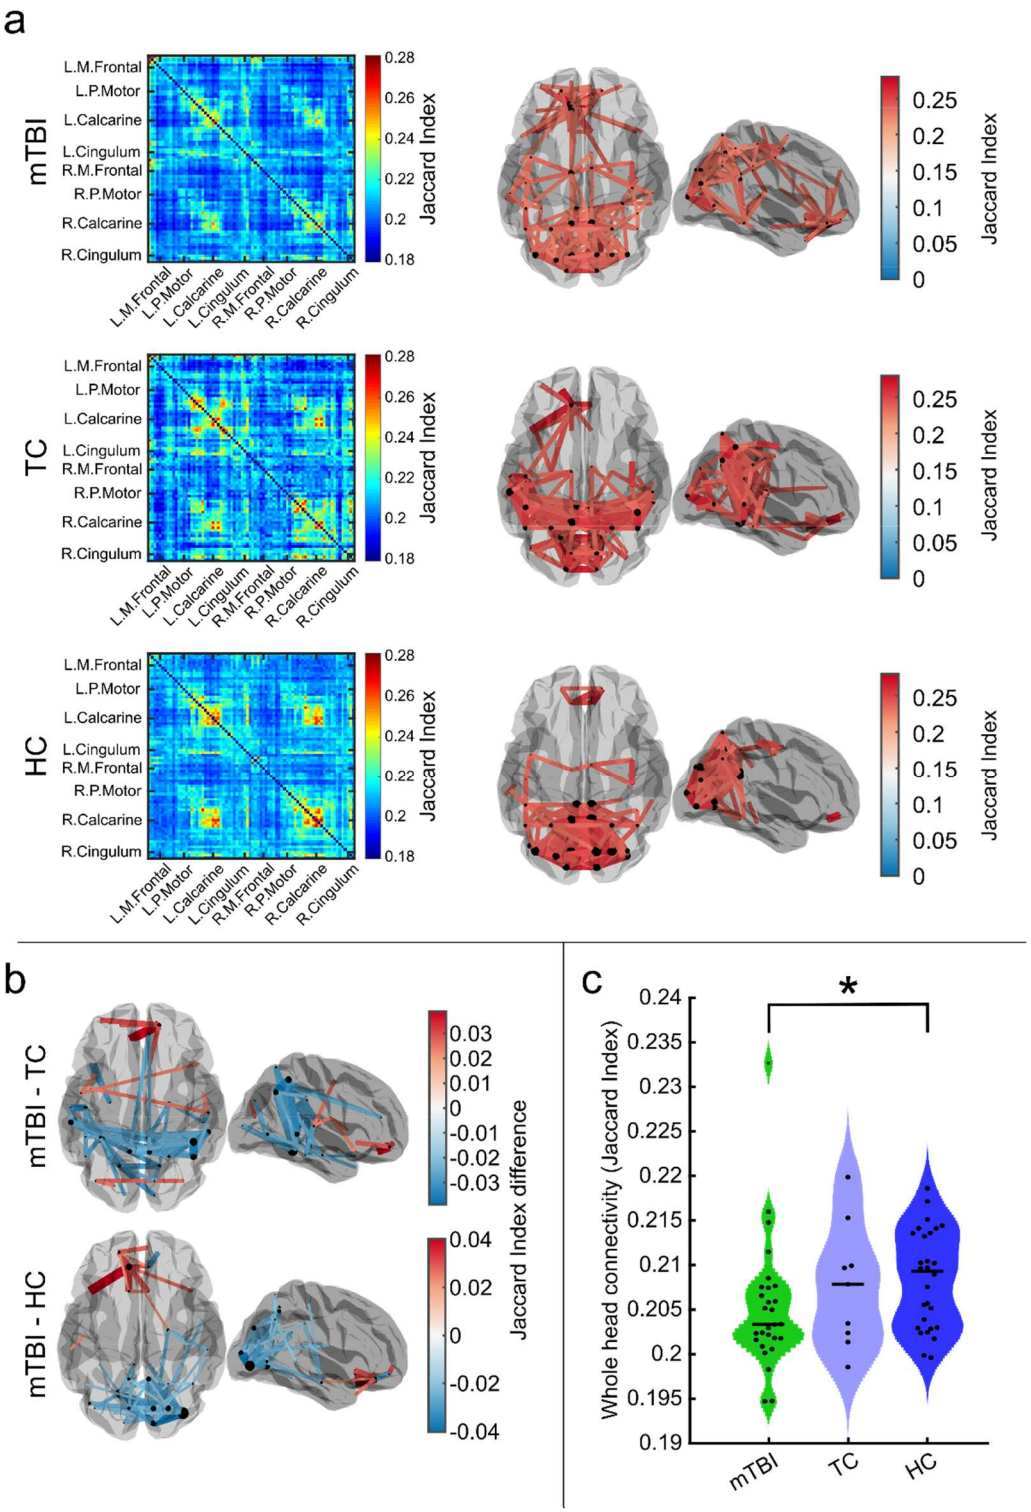

Figure 5.5 Burst coincidence connectivity. a) Average connectivity matrices for all three groups and corresponding glass brain plots showing the largest 5% of connections. b) Glass brain plots showing the top 2% of differences in connection strength when comparing mTBI subjects with each control group. c) Violin plots showing global burst coincidence connectivity for all three groups. Wilcoxon rank-sum tests between all pairs of groups showed statistically significant reduction in connectivity in mTBI compared to HC

Figure 5.6a contains a glass brain plot with differences in burst connectivity between mTBI and HC for those connections chosen as relevant features by the rRF-FS algorithm. The majority of chosen edges show reduced connection strengths in mTBI as expected and connect nodes in the posterior parietal, occipital and temporal lobes. Figure 5.6b shows the mean connection strengths over the chosen features for all three groups. Mirroring the findings from Chapter 4, mTBI subjects display significantly reduced connectivity compared to healthy controls (Wilcoxon rank-sum test, $Z = -5.07$; $p = 3.9 \times 10^{-7}$*) showing improved differentiability between the groups. Comparing mTBI to the TC group, using the same set of features, also yielded a significant difference (Wilcoxon rank-sum test, $Z = -2.7$; $p = 0.0068$*). TC subjects did not differ significantly from the HC group (Wilcoxon rank-sum test, $Z = -2.07$; $p = 0.038$). Finally we assessed whether the symptom severity metric derived from the NSI (see section 5.2.1.2) was related to the burst connectivity measure similar to the result reported in Chapter 4[8] via Spearman correlation. No significant relationship was found—neither for the combined mTBI and TC group nor for the mTBI group only. Figure 5.6c shows scatterplots of symptom severity—at baseline and 3 months post-injury—against mean burst connectivity (over rRF-FS selected connections). As HC data were not collected as part of the MEGAbIT study, severity values were not available.

The average ROC-AUC classification accuracy (mean $\pm$ standard deviation) for the SVMs trained across the ten cross-validation folds was 0.88 ± 0.18. Note here that training and classification were only done using mTBI and HC subjects.

8 It is worth noting that the severity measures acquired in MEGAbIT are not equivalent to the ones reported in Chapter 4 or Rier et al. (2021). The questionnaires used do not completely overlap in terms of assessed symptoms and are scored on different Likert scales.

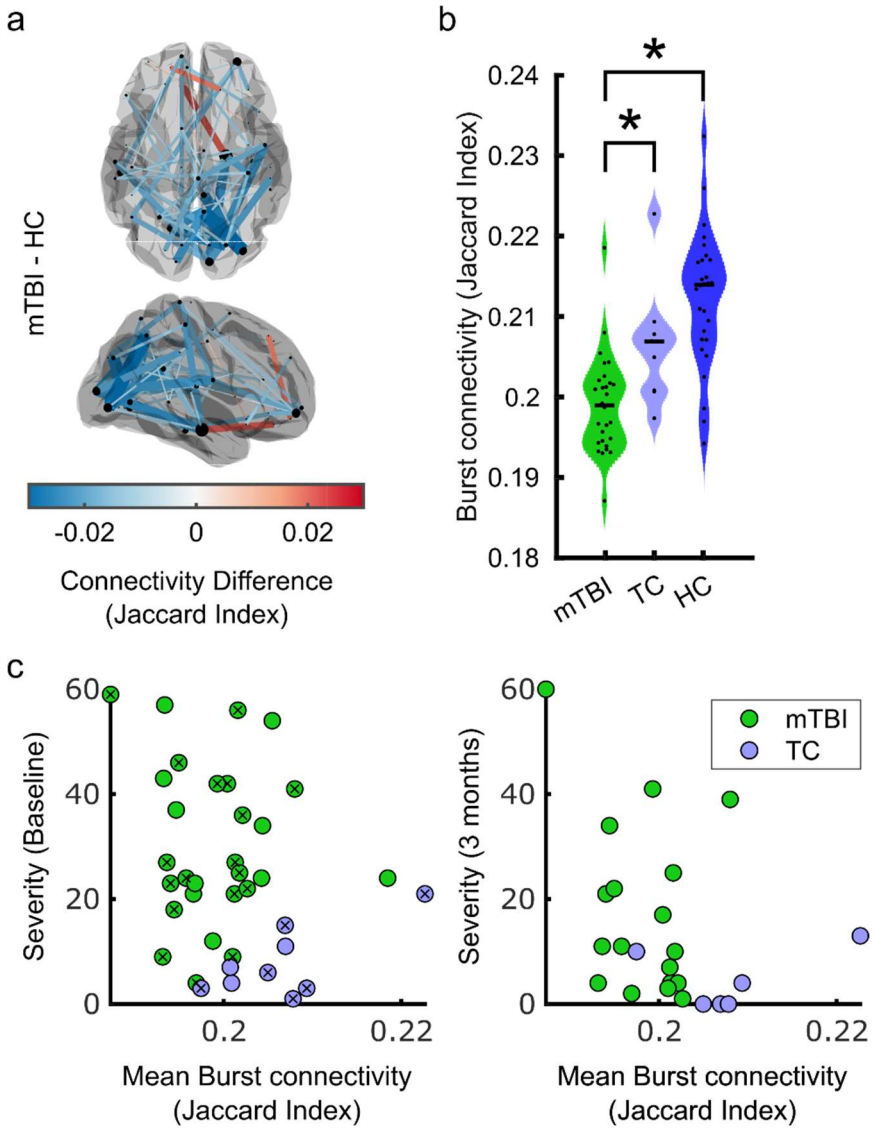

Figure 5.6 Burst connectivity for selected features. a) Connections chosen using random forest feature selection. The edge strengths are derived from the difference between mTBI and HC subjects. b) Violin plot showing the mean connection strength over the selected features shown in a) for all subjects. mTBI subjects show significantly reduced connectivity compared to healthy controls (Wilcoxon rank-sum test, $Z = -5.07$; $p = 3.94 \times 10^{-7}$). The same set of features yielded a significant difference between mTBI and TC subjects (Wilcoxon rank-sum test, $Z = -2.71$; $p = 0.0067$*). TC subjects did not differ significantly from the HC group (Wilcoxon rank-sum test, $Z = -2.07$; $p = 0.038$). c) Scatter plots of symptom severity (derived from the Neurobehavioural Symptom Inventory) at baseline and at 3 months post injury, against mean connectivity over the selected features. Crosses indicate data from subjects retained in the "3 months" plot.*

5.3.3 Discussion

The primary aim of MEGAbIT was to address the sparsity of research done at the acute stage of injury and the use of potentially inadequate controls by recruiting a control group of subjects exposed to non-specific orthopaedic trauma in addition to the mTBI cohort—both scanned within 14 days of injury. Additionally, we aimed to assess whether the resting-state findings described in Chapter 4 and Rier et al. (2021) could be replicated in our cohort while retaining sufficient specificity to distinguish mTBI from orthopaedic trauma as well as healthy controls.

5.3.3.1 *A summary of findings*

5.3.3.1.1 Burst and non-burst amplitude

Qualitatively, the measurement of the burst amplitude was comparable with those described in chapter 4 and previous work by (Seedat et al., 2020). As seen in Figure 5.3 (and for comparison Figure 4.4) the characteristic pattern of increased burst amplitude around the central sulcus is preserved in all three groups with the amplitude values falling within ranges expected from previous work (Rier et al., 2021; Seedat et al., 2020). Similarly, the time spent in the burst state was comparable, with the median values for the mTBI and HC groups falling between $300ms$ and $310ms$ —as seen previously in Chapter 4. The median for TC subjects was measured to be slightly smaller than this range although no significant difference was found between any of the groups. The stability of these measures is of great importance. Given that most research groups are limited to using small sample sizes and often use a variety of equipment, MEG-derived measures need to be comparable across sites to maximise the potential for data aggregation, improve hypothesis generation for future studies, and enable high quality multi-centre studies.

Repeating the previous measurements of global mean burst and non-burst amplitudes, we did not observe the expected reduction in mTBI relative to

healthy or trauma controls (see Figure 5.4). Inspection of the state spectra (see Appendix E) reveals more drastic between-group variability, especially in the lower frequencies (<10Hz) and the beta band compared to the data presented in Chapter 4, with no significant reduction in the beta band of the mTBI group.

Additionally, both the mTBI and TC groups appeared to show an increase in median burst amplitude (non-significant) as well as a larger variance in amplitudes compared to controls. A similar effect was seen in the non-burst amplitude, where mTBI subjects showed a statistically significant increase compared to HC. This was in contrast to our observations in the population studied in Chapter 4 where no significant effects emerged in the non-burst amplitudes. Given that the non-burst states are defined as any remaining states left after classifying our burst state, it is not possible to ascribe mechanistic meaning to the non-"bursts". It is likely that non-burst states are more strongly affected by noise, however:

HMMs can be interpreted as a form of blind source separation, where each observed state corresponds to a putative source emitting a reoccurring signal characterised by its observation model. Our burst state mostly contains periods of high SNR in the beta band which is the source of our amplitude metric. On the other hand, non-burst states may be more likely to contain reoccurring patterns of noise—signals stemming from heartbeats or periodic head movements due to heartbeats and breathing for example. These are less likely to produce intermittent peaks of beta activity but would lead to temporary increases in the noise floor, hence increasing the non-burst beta-amplitude during those periods. While our mass univariate approach to the HMM—analysing each subject and region separately—allows for greater flexibility when finding individual burst states, it is challenging to systematically investigate whether any of the non-burst states are related to the aforementioned nuisance signals, as the non-burst states are not directly comparable between subjects or regions. It is however plausible that the non-burst amplitudes diverge due to a

difference in head motion between subjects from the MEGAbIT study and our normative data (see section 5.2.6). This will be expounded in section 5.4.3.

5.3.3.1.2 Burst coincidence connectivity

Measurements of burst connectivity appeared to be robust across Chapter 4 and the current finding. The most prominent cluster of connections and nodes was found in the occipital and parietal lobes (see Figure 5.5). The previously observed trend of reduced burst connectivity in mTBI was reproduced for our acute cohort, although a statistically significant difference was only found comparing the global mean burst connectivity of mTBI and HC subjects. While statistical comparisons with the TC cohort are hampered by the small available sample size, burst connectivity values appear numerically stable between scanning sites and different experimental cohorts. Figure 5.7 repeats violin plots of global burst connectivity from Chapter 4 (Figure 4.5c) alongside the results displayed in Figure 5.5c, demonstrating their comparability. This suggests that burst coincidence connectivity could be a good candidate for longitudinal studies tracking recovery, or comparing pre- and post-mTBI MEG recordings in athletes. Future work should include test-retest characterisation of this measure to accurately assess the intra-subject repeatability and estimate the minimum effect sizes needed to show abnormal connectivity in single subjects.

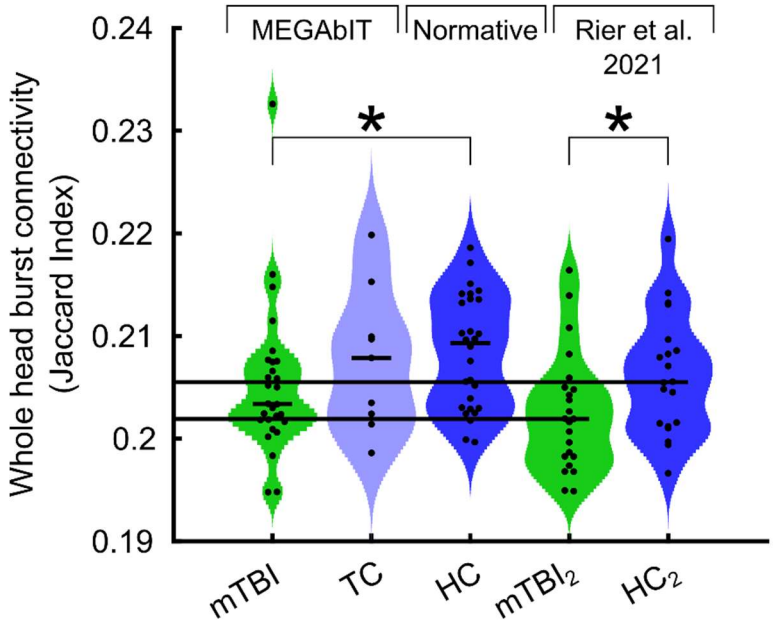

Figure 5.7 Comparison of global burst connectivity with results from chapter 4 (indicated by subscript 2). The horizontal black lines indicate sample medians for each group. Despite differing scanner types, values appear numerically stable.

The random-forest-based feature selection pipeline yielded a subset of functional connections which best differentiate mTBI subjects from healthy controls. We calculated the mean burst coincidence connectivity across the selected connections for all three groups to investigate whether they would display a more pronounced reduction in mTBI compared to HC subjects, similar to the results presented in Chapter 4 (Figure 4.5). An important detail here is that feature selection and classification were only performed using mTBI and HC subjects. In part, this was due to the pipeline implementation which was designed as a binary classifier (Zhang et al., 2020, 2016). In addition, the small number of available samples in the TC group was not sufficient to warrant training and validation of a classifier with three output classes or a second binary classifier between mTBI and TC. However, excluding TC subjects from the training process is advantageous when assessing the generalisability of the ML algorithm. By keeping the rRF-FS process naïve towards TC data, we showed that the connectivity measure derived from selected features would not only

improve the distinction between mTBI and HC but also between mTBI and TC subjects. This provides evidence that the selected features were less likely to be a result of overfitting.

Only two functional connections selected here were among the features found in Chapter 4 (left supramarginal to right postcentral gyrus and right Heschl's gyrus to the orbital part of the right middle frontal gyrus). This highlights several important points. Firstly, the patient populations studied here and in (Rier et al., 2021; Zhang et al., 2020) are not equivalent. Not only were data acquired at different times after injury—mechanisms of injury and presentation of symptoms can vary widely. Additionally, patient demographics differed, with a male-only population studied in Chapter 4. We can, therefore, not assume that the same functional connections would be disrupted and that features selected from small populations can be directly applied to unseen data. The lack of generalisability of the selected features may also raise concerns about overfitting, especially given the small sample sizes and large number of features. However, while sample size is one of the main limitations of this study, the use of 'bagging' (see Appendix B) greatly reduces the likelihood of overfitting the random forest classifier. Briefly, each decision tree in the ensemble is trained using a small, randomly chosen subset of features and subjects. This random sampling with replacement—which yields the feature importance measure via internal cross-validation—increases the stability of the random forest classifier and reduces the likelihood of overfitting in the final classification step by offering a systematic means of reducing the feature set.

SVM classification yielded good ROC-AUC accuracy (similar to Chapter 4), showing that our burst connectivity measure allows for adequate separation between mTBI and HC subjects. This result should, however, be interpreted with caution. Overfitting can arise when training SVMs on high-dimensional data in combination with small sample sizes. To combat this, we used the RF-FS process to systematically extract the most important features and 10-fold

cross-validation to assess the robustness of the SVM classifier via the standard deviation of the classification accuracy over the 10 folds. While we found the classifier to be reliable across the cross-validation steps, it does not follow that this method could be used as a diagnostic tool with similar classification accuracy. The classification accuracy does, however, provide a useful summary measure to confirm the utility of burst coincidence connectivity in studying mTBI. Future work should either combine data from several studies for a more robust and generalizable classifier or use demographic data, time post-injury and similar details to assess the utility of classifiers trained on homogeneous samples.

Given that the RF-FS procedure mainly selected connections with a connectivity deficit, mirroring the results presented in Chapter 4, and that these connections differ significantly between the TC and mTBI groups, our burst connectivity measure shows great promise.

5.3.3.2 *Mechanistic interpretations and clinical relevance—an update to Chapter 4*

By making use of the burst model of neural activity, we can go beyond the classical assumption of continuously and smoothly varying oscillations and gain further insight into those aspects of neural dynamics that drive potential abnormalities in mTBI. As laid out in Chapter 4 (section 4.4.2), there exist a number of plausible mechanisms that may be responsible for those abnormalities. In this chapter, we aimed to test whether the same abnormalities could also be found in the acute phase of injury, and included a non-head trauma control group to assess whether abnormalities were specific to mTBI.

In the subacute mTBI group, we observed altered spontaneous bursting in the form of reduced beta amplitudes during the bursts as well as disruption of burst coincidence connectivity. We postulated that the amplitude reduction could arise from a disruption of thalamocortical connections (Sherman et al., 2016) which are likely candidates for deformation in a concussive event. The absence

of this amplitude reduction, in this cohort, as well as the subtle differences in the state spectra of the investigated groups (see Appendix E), raises two questions: Are our two mTBI groups comparable enough to suggest that the previously observed significant reduction in amplitude was a statistical anomaly? Or are we observing truly different stages of recovery, which manifest in different signals? While it is possible that the groups we have thus far referred to as *acute* and *subacute* are—to a degree—interchangeable, further research is necessary to expand our knowledge about the evolution of the MEG signal during recovery from mTBI.

Current understanding of pathology after TBI (predominantly based on non-human research) suggests a variety of secondary processes that occur after injury that display different profiles of evolution. The initial deformation of tissue causes a cascade of neurochemical and metabolic changes (MacFarlane and Glenn, 2015); the disruption of the blood-brain-barrier is followed by a variety of neuroinflammatory processes such as activation of microglia and astrocytes and other imbalances can last for more than a week—albeit with varying recovery timescales (see Simon et al. (2017) for a review). Excitotoxicity has been observed in TBI in the first week post-injury (Jamjoom et al., 2021; Ruppel et al., 2001; Vespa et al., 1998). The majority of these processes are thought to resolve on the timescale of our acute window. However, pathological Tau protein accumulation and neuroinflammation have been reported at later stages of injury (Jamjoom et al., 2021). Based on the timescales of these processes the decrease in burst amplitude in subacute mTBI may be due to long-term pathology that is not yet present in the acute phase or offset by other effects such as excitotoxic hyperexcitability. Longitudinal studies combining MEG with non-imaging markers of neurometabolism and inflammation may be needed to confirm this.

Neuropathology following an mTBI is widely attributed to diffuse axonal injury (DAI) with the white matter around the corpus callosum being particularly

vulnerable (Gazdzinski et al., 2020). Given the potential disruption of neural connections (Browne et al., 2011; Kirov et al., 2013) we can plausibly expect a reduction in the synchrony between brain regions. Both acute and subacute mTBI groups show reduced burst coincidence connectivity. While this result validates our previous findings, the inclusion of our orthopaedic trauma controls calls the specificity of our measure in identifying *head* injury into question. Our observation of reduced connectivity in TC subjects compared to healthy subjects may point towards a systemic trauma-related effect. However, another possibility is that TC subjects were not subject to purely peripheral trauma. While all TC subjects attended the emergency department without mTBI-related loss of consciousness, the mechanism of injury for a number of subjects may have included acceleration/deceleration of the head, and hence a potential subconcussive event (see Appendix F for more details about the recorded injury mechanisms). While the effects of subconcussive events are poorly understood, several (non-MEG) neuroimaging studies have shown abnormalities resulting from subconcussive events, especially when subjects were exposed to reoccurring impacts (e.g. over a season of playing contact sports). While MEG studies focusing on subconcussive events are lacking, our findings may be of interest here. It is important to note that our limited sample size for the TC group impedes the formulation of strong conclusions. Our results, nonetheless, highlight the importance of adequate control groups in mTBI studies.

5.3.4 Conclusion

In this study, we replicate our previous finding of reduced resting-state burst coincidence connectivity in mTBI and show that this abnormality can be observed in the acute phase of injury. By including a group of controls exposed to non-head orthopaedic trauma, we explored whether any observed abnormalities would be due to a non-specific trauma effect. While our connectivity metric significantly differed between mTBI and trauma controls—

mirroring previous findings—burst amplitude did not differ between groups. This may point towards an evolution of pathology post injury that should be assessed in future longitudinal studies of mTBI with MEG. By investigating mTBI subjects soon after injury, our findings contribute to a more complete picture of brain changes following head trauma.

5.4 DELTA POWER

5.4.1 Analysis

5.4.1.1 *Source Reconstruction*

The coregistered, T_1-weighted anatomical MRIs (1mm isotropic resolution) for each subject were downsampled to 4mm isotropic resolution and brain extraction using FSL BET (Smith, 2002), yielded the anatomical space used for source reconstruction. A multiple spheres head model (Huang et al., 1999) was used for source reconstruction via an LCMV beamformer. For each voxel location in the downsampled brain, the source orientation was found via an exhaustive search (see Chapter 3) yielding weights corresponding to the orientation with the highest SNR. Using these beamformer weights $\boldsymbol{w}$, the source power p was estimated for each voxel using equation 3.31

$$p = \hat{\sigma}^2 = \boldsymbol{w}^T \boldsymbol{C} \boldsymbol{w},$$

$\boldsymbol{w}$ contains the beamformer weights and $\boldsymbol{C}$ is the covariance matrix generated from the band-limited data (band-pass filtered to the delta band; $1 - 4\ Hz$). No regularisation was applied to minimise leakage and maximise the interference reduction properties of the beamformer. In order to prevent the possible confound introduced via weights normalisation (Luckhoo et al., 2014), non-normalised beamformer weights were used. We assumed that the depth bias particular to beamforming will be accounted for when producing statistical images from the different subject groups. All maps of delta-power were transformed into a common space using FSL FLIRT (Jenkinson and Smith, 2001): each anatomical was aligned to the MNI152 template brain (Evans et al., 2012), yielding the transformation matrices needed to align each functional map to the template.

5.4.1.2 *Metrics derived from source reconstructed delta power*

The delta power maps were used to assess the following points:

I) *Can we replicate the commonly reported delta power increase in our mTBI cohort?*

Using un-normalised power maps to prevent bias we looked at **global delta power** as the locus of injury can vary widely. This approach also aids in avoiding the multiple comparisons problem. Additionally, we mimicked Huang et al. (Huang et al., 2014) by using a normative sample of healthy controls to calculate voxel-wise Z-scores and comparing the distribution of individual **maximum Z-scores** (Z_{max}). Mathematically,

$$Z_{i,j} = \frac{p_{i,j} - \overline{p_{i,HC}}}{s_{i,HC}} \tag{5.3}$$

where $Z_{i,j}$ and $p_{i,j}$ are the Z-score and power in voxel i of subject j, $\overline{p_{i,HC}}$ is the mean power in voxel i across the HC group, and $s_{i,HC}$ is the standard deviation of power values in voxel i across the HC group. Z_{max} is the maximum value across all $Z_{i,j}$ for a particular subject j, or

$$Z_{max,j} = \max_{i}\left(Z_{i,j}\right) \tag{5.4}$$

We assessed whether this approach can distinguish subjects with mTBI from the healthy controls used to generate the normative database.

II) *Are there voxels that consistently contain abnormal delta band signal*? We generated T-statistic maps comparing mTBI/TC group means with the HC group. This was done by calculating a two-sample T statistic at each brain voxel i in MNI152 space

$$T_i = \frac{\overline{p_{i,G}} - \overline{p_{i,HC}}}{\sqrt{\frac{s_{i,G}^2}{N_G} + \frac{s_{i,HC}^2}{N_{HC}}}} \quad (5.5)$$

where $\overline{p_i}$ are the mean power values in voxel i, s_i^2 are the variances in power across the group and subscripts G and HC indicate the compared group and healthy control group respectively. N are the respective sample sizes. A cluster-based permutation test was used to control for multiple comparisons.

III) *Is the measure of excess delta power specific to head trauma?*
We assessed whether global delta power / the maximum Z-score are sufficiently specific measures to classify orthopaedic trauma controls differently from the mTBI group when compared to the normative database. If excess delta power is a sign of mTBI, as suggested by the literature, we should expect orthopaedic trauma controls to present with normal delta power.

IV) *Is there a relationship between self-reported symptom severity and delta power?*
We calculated the Spearman correlation between mean global delta power and symptom severity as measured using the sum of scores in the Neurobehavioural Symptom Inventory (NSI) questionnaire. This was done with severity measured at baseline and also at 3 months to see whether delta power at baseline has prognostic utility.

5.4.2 Results

Figure 5.8 shows the global mean delta power for the three groups. No significant difference between the mTBI and either control group was found when performing a Wilcoxon rank-sum test: for mTBI vs HC, $Z = 1.16$ ($p = 0.25$); for mTBI vs TC, $Z = -1.79$ ($p = 0.074$). Interestingly, the TC group yielded the highest median delta power. However, no significant difference

between TC and HC was found after correction for multiple comparisons using the Benjamini-Hochberg procedure (Benjamini and Hochberg, 1995) (Wilcoxon rank-sum test; $Z = 2.35$; $p = 0.019$).

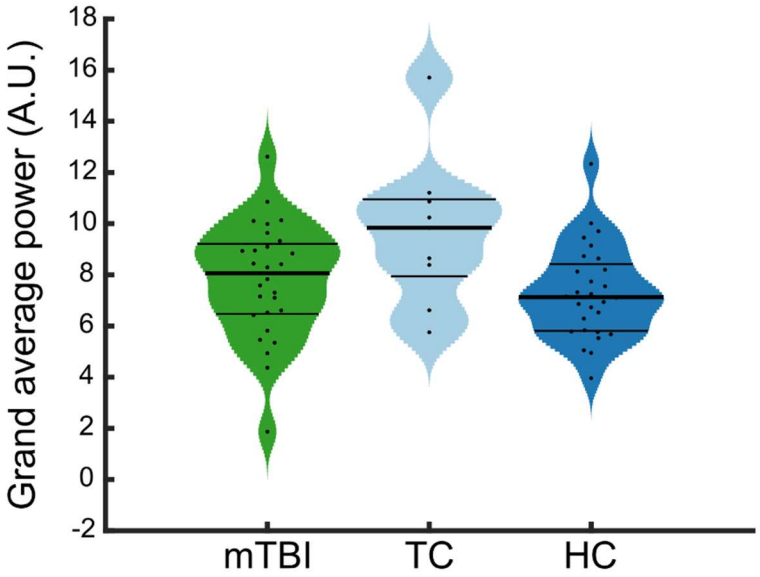

Figure 5.8 Global average delta power. Violin plots showing the mean delta power over all brain voxels for all three groups: mTBI, orthopaedic trauma controls (TC) and healthy controls.

Figure 5.9a and b respectively show the voxel-wise mappings of the T-statistic and corresponding voxels where $p < 0.05$ for mTBI vs HC. Frontal and occipital lobes appear to contain most voxels with increased delta, with a cluster in the left hemisphere showing reduced delta in mTBI compared to healthy controls. Note that when correcting for multiple comparisons using a cluster permutation test, no significantly different clusters were identified. Figure 5.9 c) and d) show the same statistics for mTBI vs TC. While cluster permutation testing yielded no significantly different clusters between groups, we observed an overall tendency of decreased delta power in mTBI compared to TC subjects.

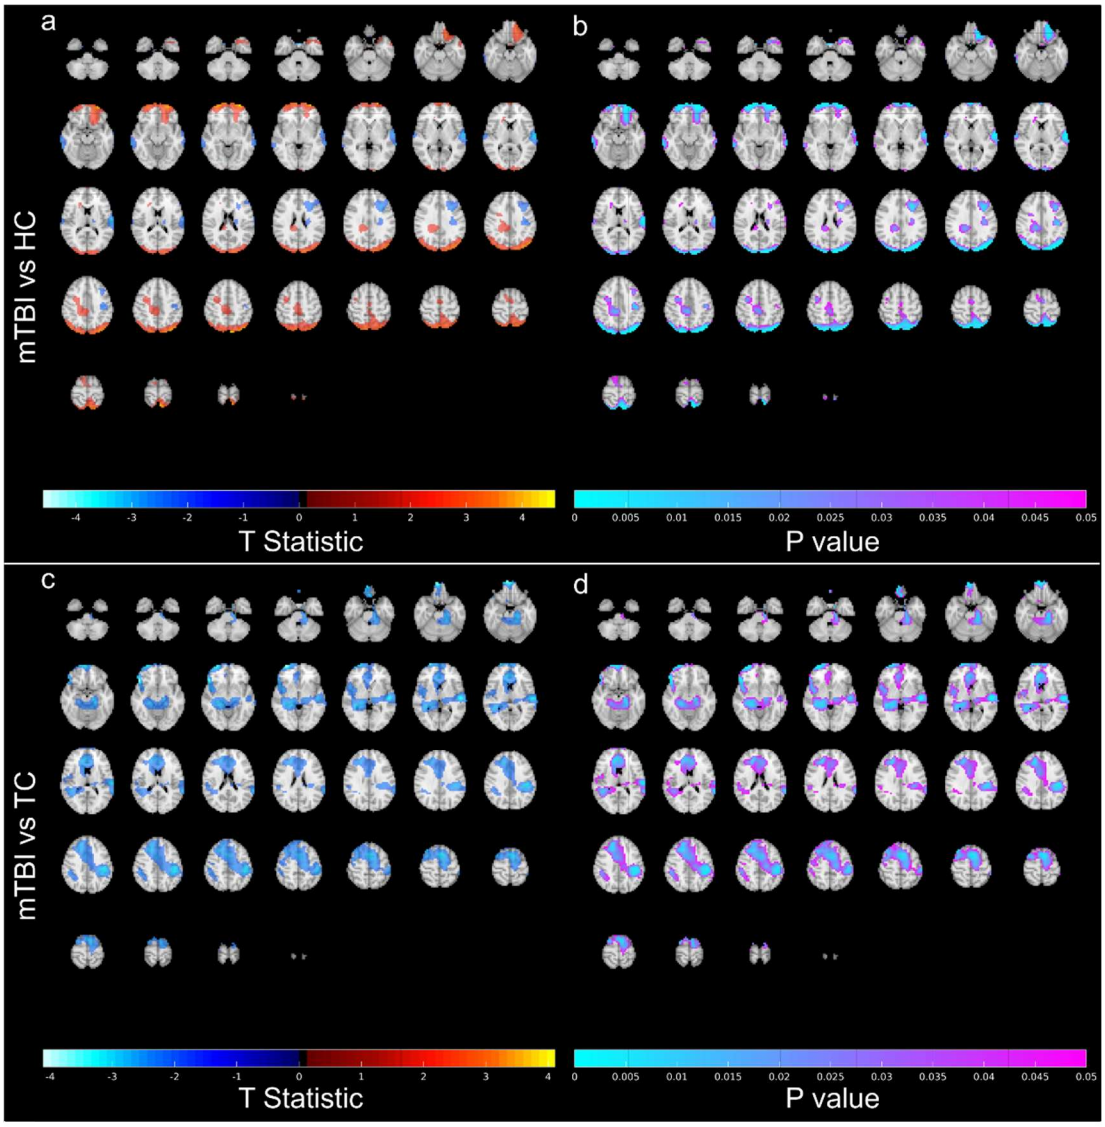

Figure 5.9 T statistical maps comparing mTBI subjects with healthy controls (top row) and orthopaedic trauma controls (bottom row). a) T-statistic – mTBI vs HC. b) Uncorrected p value – mTBI vs HC. c) T-statistic – mTBI vs TC. d) Uncorrected p-value – mTBI vs TC. T-statistical maps and P-value maps are thresholded to only include voxels with $p < 0.05$*. All maps are shown in radiological orientation (anatomical left on the right of the images).*

The approach inspired by Huang et al. (2014) yielded the distributions of maximum Z-scores shown in Figure 5.10. When only considering mTBI and HC subjects and aiming for maximal sensitivity by using the highest Z_{max} for the HC group as a classification threshold (see dashed line in Figure 5.10), 20/28 mTBI subjects—having a Z_{max} greater than the highest Z_{max} found in the HC

group—are correctly classified. However, only one TC would correctly be identified as mTBI-negative, indicating a lack of specificity for this approach.

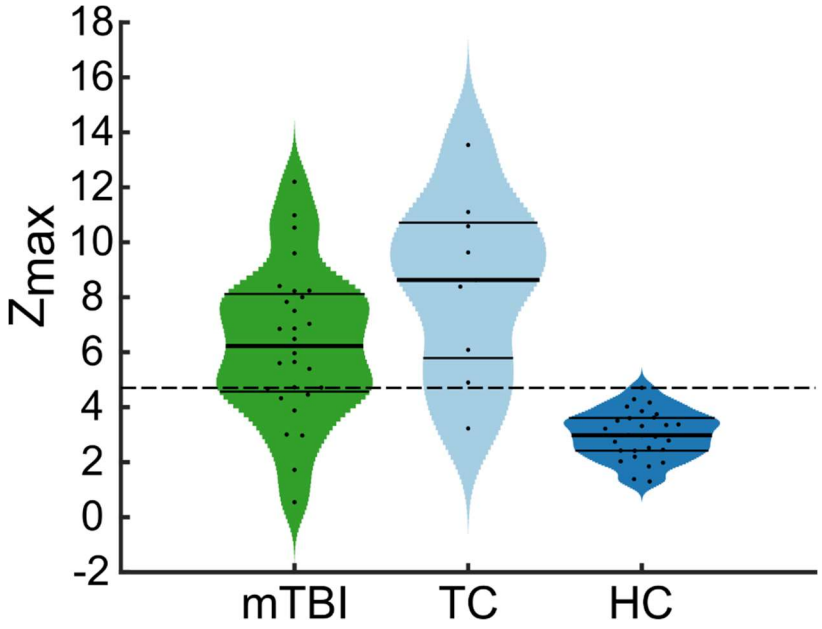

Figure 5.10 Maximum Z-scores for delta power. The group of healthy controls was used to generate maps of the voxel-wise mean and standard deviation for a normative sample, which were used to calculate individual Z-score maps for each subject. Maximum Z-scores were extracted from each subject and used to generate the violin plots. Solid black lines indicate the quartiles for each distribution. The horizontal dashed line indicates the maximum value found in the HC group—the classification threshold with zero false positives for the HC cohort.

We investigated whether there was a monotonic relationship between delta power and symptom severity. Scatter plots of symptom severity against global delta power for the mTBI and TC groups are shown in Figure 5.11. Using the severity measure recorded on the scanning day, combining the mTBI and TC groups yielded ($R = -0.43; p = 0.0080$*); the mTBI group excluding TC yielded ($R = -0.39; p = 0.042$*). For the 3 month assessment, Spearman correlation for the combined groups yielded ($R = -0.47; p = 0.021$*). The mTBI group only yielded ($R = -0.49; p = 0.037$*). Not all subjects volunteered to provide symptom information three months post-injury leading to the reduced sample size.

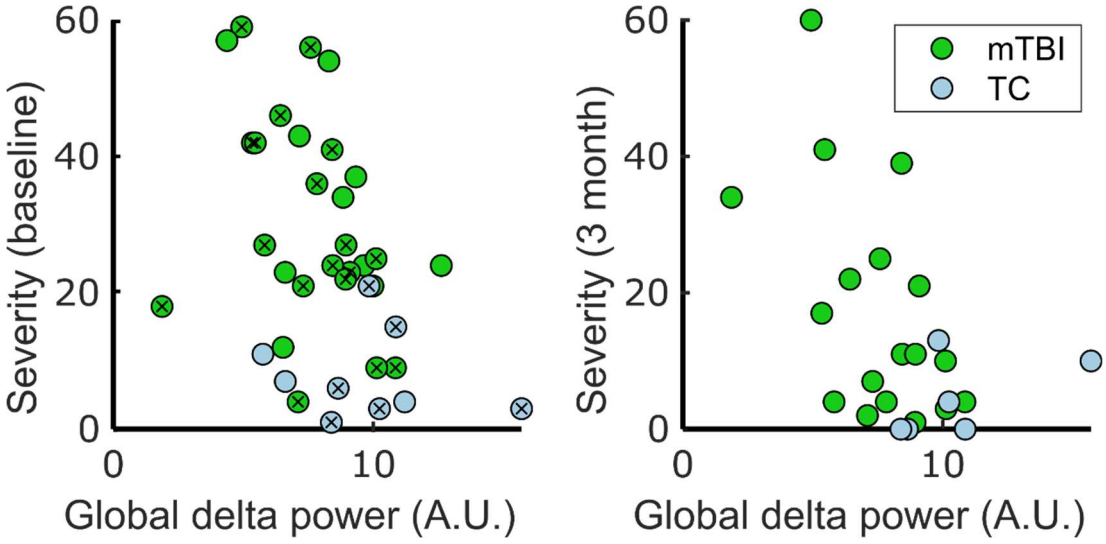

Figure 5.11 Relationship between symptom severity (sum of neurobehavioural symptom inventory scores) and global mean delta power for mTBI subjects and orthopaedic trauma controls. The plot on the left was generated using symptom scores recorded on the day of scanning. Spearman correlation between severity and power for the combined group yielded ($R = -0.43; p = 0.0080$), the mTBI group excluding TC yielded ($R = -0.39; p = 0.042$*). The right-hand plot shows a scatter plot using severity scores recorded 3 months post-scan. Note that not all subjects volunteered to complete the 3-month assessment resulting in smaller sample sizes. Crosses indicate those subjects present in both plots. For the 3 month assessment, Spearman correlation for the combined groups yielded ($R = -0.47; p = 0.021$*). The mTBI group only yielded ($R = -0.49; p = 0.042$*).*

5.4.3 Discussion

With the study presented above, we aimed to investigate the utility of excess delta power measures in distinguishing subjects with mTBI from healthy individuals and orthopaedic trauma controls in the acute stage of injury. This addresses two of the methodological issues identified in several mTBI studies synthesized in Allen et al. (2021).

Our measurements of global delta power showed no statistically significant group differences between the mTBI and either control group with TC subjects showing the highest average delta power overall. While this latter result is

inconsistent with the expectation of increased delta power in mTBI, several factors might contribute to reduced sensitivity in this global measure. The mechanisms of injury and/or location of impact to the head vary widely, are seldom known precisely and the diffuse nature of injury to the brain following mTBI doesn't enable predictions of specific regions of interest even if such knowledge was available. A global measure circumvents this issue while also avoiding the multiple comparisons problem resulting from the large number of voxels used. As can be seen in Figure 5.9, group differences between mTBI and HCs are not significant in any particular region and inconsistent in sign across the brain, which means any localised abnormalities are likely to be averaged out by our global measure.

The approach inspired by Huang et al. (2014) attempts to increase sensitivity by focusing on regions with maximal delta power, the rationale being that mTBI subjects should show strong peaks of delta activity in the most affected region. Our results mirror the findings of (Huang et al., 2014) when using this metric, who equally used a normative database of healthy controls to produce Z-score maps which in turn enabled the extraction of peaks of a maximum Z-score from each individual. Huang et al. (2014) showed good separation between mTBI subjects at the chronic stage and healthy controls (~86% positive detection rate). Our acute mTBI cohort could be distinguished from HC with reasonable sensitivity (~71% positive detection rate), however, this method resulted in all but one TC subject being classified as mTBI. We propose two speculative interpretations of this lack of specificity:

Firstly, the hypothesis of excess delta power being a specific indicator for mTBI to MEG scans performed shortly after the injury might be unfounded given the dominance of research conducted in subacute or chronic mTBI populations. The similarity between our mTBI and TC groups could, therefore, indicate that excess delta power is indicative of an effect due to non-specific trauma, at least at the acute stage.

Secondly, the differences between data collected as part of the MEGAbIT study and the HC group could be due to a systematic error caused by methodological discrepancies at the acquisition stage. While we aimed to eliminate the known confound of age (Babiloni et al., 2006; Hunt et al., 2019) by following the optimal matching procedure described in section 5.2.2, one potential source of artificially increased delta power in the MEGAbIT groups could be related to differences in head motion inside the scanner helmet. While MEGAbIT subjects were scanned in a sitting position, several HC subjects were recorded while supine, which may have reduced movement. Our posthoc assessment of subject motion (section 5.2.6) was carried out to quantify the potential influence of head movement on the delta measure. Panel a) in Figure 5.12 contains a violin plot of our motion metric obtained for each subject. Head motion was comparable for both groups scanned as part of the MEGAbIT study; a clear reduction in the motion metric can be seen in the HC group, however. Spearman correlation of the motion metric with the global delta power (Spearman $R = 0.00$; $p = 0.99$) and maximum Z-score metrics (Spearman $R = 0.26$; $p = 0.03$) revealed no statistically significant monotonic relationship after correcting for multiple correlations, however, a larger effect size was observed for the maximum Z-score. This suggests that the sensitivity of Z-score-based classification might have been lower than the calculated 0.71, had we used a sample of healthy subjects with more comparable head motion. The global delta power measure appears not to be affected by head motion.

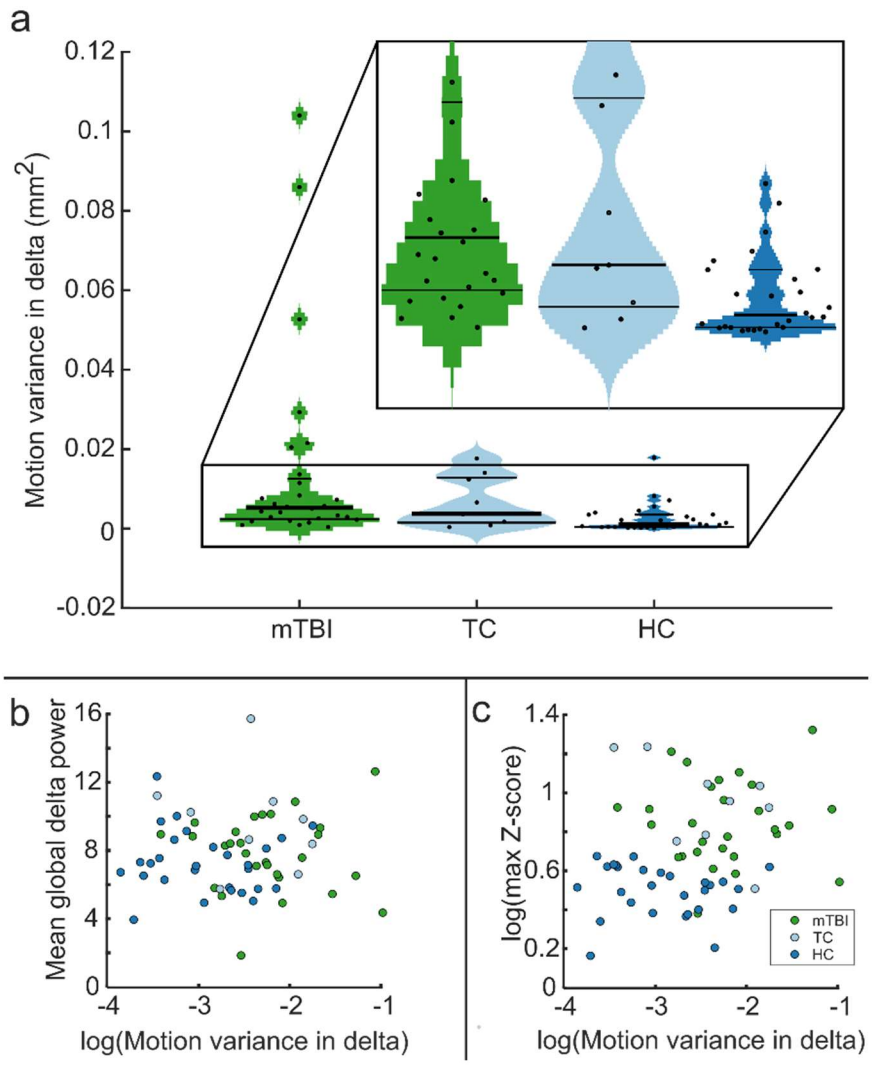

Figure 5.12 The effect of head motion on delta power. a) Variance of the displacement time courses filtered to the delta band and averaged over the three head localisation coils. Note the logarithmic scale. The amount of motion present in the mTBI and TC groups was comparable, whilst being significantly larger than for the HC cohort. A potential explanation for this could be the supine positioning of subjects used for the majority of HC subjects while volunteers for the MEGAbIT study were sitting upright during the recording. b) Scatter plots of the motion variance against the global delta power metric (top) and the maximum Z-score (bottom). Spearman correlation of the motion metric vs Mean global delta power yielded (Spearman $R = 0.00$; $p = 0.99$ *); correlation with the maximum Z-score yielded (Spearman* $R = 0.26$; $p = 0.03$*).*

When only considering data from our mTBI and TC groups, we would conclude that excess delta power does not differentiate mTBI from appropriately matched orthopaedic trauma controls. On the contrary, a reduction in delta power might

be more indicative of mTBI compared to TC. While we did not measure statistically significant differences between mTBI and TC, a significant negative association between symptom severity and delta power was found for the combined mTBI and TC groups as well as for the mTBI group only. A similar trend was found when assessing the symptom severity 3 months post-injury. This surprising result contradicts our expectation (based on data overwhelmingly collected at the chronic stage) that excess delta power is related to more severe symptoms. Longitudinal studies which would assess any changes in symptomology and MEG signal throughout recovery would be needed to establish whether individuals maintaining abnormal delta are indeed more likely to experience symptoms. Especially for studies where scans are performed in the acute stage, subject retention for longitudinal data points is crucial. A larger sample of mTBI subjects would therefore be advantageous, as up to half of all patients may experience long-term symptoms (McInnes et al., 2017; Nelson et al., 2019).

Another important limitation that will need to be addressed in future work is the small sample size in the TC group, especially given the methodological discrepancies between the HC group and our sample of mTBI subjects. Subject positioning may have influenced the amount of head motion present in the different groups. Additionally, the use of historical data for a normative dataset was not conducive to screening participants for the same exclusion criteria. The HC group, although classed as healthy at the time of the scan, was not screened for a history of brain trauma, for example.

5.4.4 Conclusion

Excess delta power has been reported in a substantial number of studies investigating mTBI using MEG, however, the heterogeneity of methods, patient cohorts, and control groups described in the literature does not allow for definitive statements about the efficacy of this measure as a biomarker for mTBI. In this study, we measured source-reconstructed delta power in a group of

subjects with mTBI in the acute stage of injury. Excess delta appeared to differentiate mTBI from healthy controls, however, similarly increased delta activity was observed in a group of subjects suffering from orthopaedic trauma. Additionally, delta power was shown to be inversely related to symptom severity, counter to our prior expectations. Limited by the small number of orthopaedic trauma controls, mTBI could not be distinguished from non-head trauma via excess delta power.

Chapter 6 NOVEL TASK PARADIGMS USING OPM-MEG

The experimental results presented so far in this thesis have added to a catalogue of findings that demonstrate the ability of MEG to detect abnormalities in brain function after mTBI. However, to fully understand the variation an injured brain undergoes, we need to develop an understanding of the changes in MEG signal over time (i.e. longitudinal measures). In addition, populations at high risk of mTBI, such as athletes, may benefit from baseline assessments which can be used to reveal post-injury abnormalities and track recovery. There also appears to be a need to investigate naturalistic paradigms which can be tailored to induce activity in symptom-relevant brain networks and consequently are likely to be more sensitive biomarkers than the resting state scans used to date. The introduction of OPMs for MEG makes naturalistic scanning possible. Moreover, the cryogen-free nature of OPMMEG may offer a means for simpler, cheaper and more accessible scanners to be deployed for longitudinal studies. However, before they can be used for this purpose, we need to quantify the reliability of OPM-based measures of brain activity. In this final experimental chapter, we introduce a novel naturalistic task paradigm which showcases the flexibility of OPM-MEG and induces motor activity shown to be abnormal in subjects with mTBI in Chapter 4 (albeit using a task which was not optimised to assess motor function). We estimate a series of common MEG measures and combine neural fingerprinting and the Bland-Altman method to quantify the reliability of both activity and connectivity derived from OPM-MEG data. These initial findings—presently limited to a small sample of healthy controls—demonstrate the utility of OPM-MEG and pave the way for this technology to be deployed on patients with mTBI.

6.1 INTRODUCTION

As discussed in previous chapters, MEG has shown great promise in the study of mild traumatic brain injury. However, the introduction of MEG-based assessments of mTBI to the clinic cannot yet be justified. In particular, within-subject, longitudinal changes in the MEG signal have been largely unexplored with (the previous chapter's study notwithstanding) most studies concentrating on chronic cases, sometimes scanned many months, or even years, post-injury. It is therefore not clear how observed abnormalities evolve over time. Additionally, the vast majority of studies only collect resting-state data. This means that the networks affected by mTBI may not be specifically probed. For example, a balance task would be ideal to investigate vestibular deficits. The use of MEG to probe the formation and dissolution of networks supporting affected function, especially in a naturalistic way, would offer a significant step forward.

While head injuries are prevalent in the wider population, recent media attention (e.g. (Reuters, 2021, 2020)) has rightly been given to athletes who are at high risk of concussive blows to the head (e.g. Rugby/Football/American Football players). Consequently, these groups are important stakeholders for mTBI research. Those with an elevated risk of a head injury may offer a unique opportunity for longitudinal studies to contrast baseline neural activity (i.e. before a season) with post-injury signals, or to explore the cumulative effects of exposure to subconcussive blows which may also have a long-term impact on brain health (Choe, 2016; Davenport et al., 2016). Additionally, MEG could be used to longitudinally track recovery to establish objective criteria for return-to-play, replacing or supplementing current tools such as the SCAT battery of tests (Echemendia et al., 2017). Such neuropsychological tests may be more vulnerable to subjective administration and learning effects—as well as purposeful underperformance at baseline in the case of professional athletes (Bailey et al., 2006; Echemendia and Julian, 2001)—than more objective imaging-based assessments.

Unfortunately, a major drawback of conventional MEG is the prevalence (and cost) of available scanners, which are mostly located in universities. This may complicate the logistics of performing a large number of baseline scans and makes pitch-side assessments impossible.

Recent advances in the manufacturing of OPMs have removed the need for cryogenic cooling to perform MEG. Moreover, the small form factor of OPM has enabled wearable MEG with sensors being mounted closer to the scalp, hence achieving increases in signal to noise ratio and spatial resolution (Boto et al., 2016; Iivanainen et al., 2017), and allowing for more naturalistic task paradigms without restricting head movements. The customisability of OPM-MEG avoids the one-size-fits-all nature of conventional MEG, meaning uniform coverage is easy to achieve. OPM-MEG is also (at the time of writing) less costly than conventional MEG. All these factors point to the study of mTBI becoming more feasible.

Indeed, in the longer term, a "MEG-scanner-in-a-van" system positioned near sports grounds may even become possible.

As shown in Chapter 4, mTBI may manifest in motor network abnormalities measurable by MEG. However, one major limitation of the task used was that the paradigm was not optimised to assess motor function (Rier et al., 2021). The short trial duration, for example, may not be suitable to assess the entire range of motor responses. Among the most commonly observed neurophysiological effects in the motor cortex are the movement-related beta desynchronization (MRBD)—a decrease of power relative to intrinsic beta-band activity—and the post-movement beta rebound (PMBR)—a relative increase occurring after movement cessation—which then returns to baseline levels and is most prominently localised in the motor and sensory areas around the central sulcus (for reviews see (Cheyne, 2013; Kilavik et al., 2013)). Another effect observed is altered functional connectivity in the bilateral sensorimotor network,

coinciding with the PMBR (O'Neill et al., 2017; Tewarie et al., 2019). Recent work has shown that PMBR can be relatively long-lasting compared to the trial lengths used in Chapter 4 (Fry et al., 2016; Pakenham et al., 2020), which should be taken into consideration when designing motor-specific tasks.

To induce these motor-related effects and to tailor our paradigm to domains potentially affected in mTBI, we aimed to design a challenging motor task which would also require a level of cognitive effort and attention. We took inspiration from The Trail Making Test (TMT), a common neuropsychological test, which requires subjects to trace a path connecting a series of numbers or alternating numbers and letters on a page in sequence and as quickly as possible (Bowie and Harvey, 2006). Instead of performing repeats of the TMT, which may be subject to learning effects, our volunteers were required to solve a series of randomly generated mazes in the shortest time possible. This task is, therefore, close to the TMT sometimes used in mTBI assessment. However, it requires a degree of subject movement that would preclude the use of conventional MEG. We, therefore, utilised our wearable MEG system based on optically pumped magnetometers (OPMs).

Although OPM-MEG has shown great promise, wearable MEG brings with it several challenges. While the sensor positions are fixed relative to the head to allow for source reconstruction, movement of the sensor array through any magnetic field or field gradients present inside the magnetically shielded room can induce artefacts in the signal or even prevent sensors from operating if their dynamic range is exceeded.

While recent (and ongoing) work on active and passive shielding has been instrumental in addressing this issue, the sensitivity of OPM-MEG to individual differences in the MEG signal and the stability and repeatability of OPM-MEG based measures of neural activity need to be assessed. This is especially important for longitudinally acquired data—to track recovery for example—

where any observed differences due to plasticity in recovery would need to be contrasted with the natural within-subject variability between scans.

6.1.1 Summary of Aims

We aimed to use our task paradigm, inspired by the Trail Making Test, to challenge subjects in the domains of visuomotor control, attention and concentration while allowing free movement of the head via wearable OPM-MEG. We expected to observe common brain responses to motor tasks such as the MRBD and PMBR (localised to the sensorimotor areas) and the formation of a motor network revealed via amplitude envelope correlation. To evaluate the feasibility of using OPM-MEG (and our novel task) for longitudinal assessment, we aimed to assess the reliability and test-retest agreement of each functional measure.

6.2 METHODS

We performed a repeatability study by scanning 6 individuals (all right-handed, 2 Female) twice (on consecutive days). In order to exploit the flexibility of OPM-MEG, we devised the maze task which required participants to freely move their arm and allowed them to direct their gaze without any restriction of head motion.

6.2.1 System Setup and Data Acquisition

The OPM-MEG array used for this study was previously described in (Hill et al., 2020) and followed the principles laid out in Chapter 2. A rigid, additively manufactured helmet containing 100 possible sensor slots was used to hold 37 2^{nd}-generation OPM sensors (QuSpin Inc. Colorado, USA). These included 33 dual-axis and 4 triaxial sensors giving a total of 78 channels. The entire setup is shown schematically in Figure 6.1.

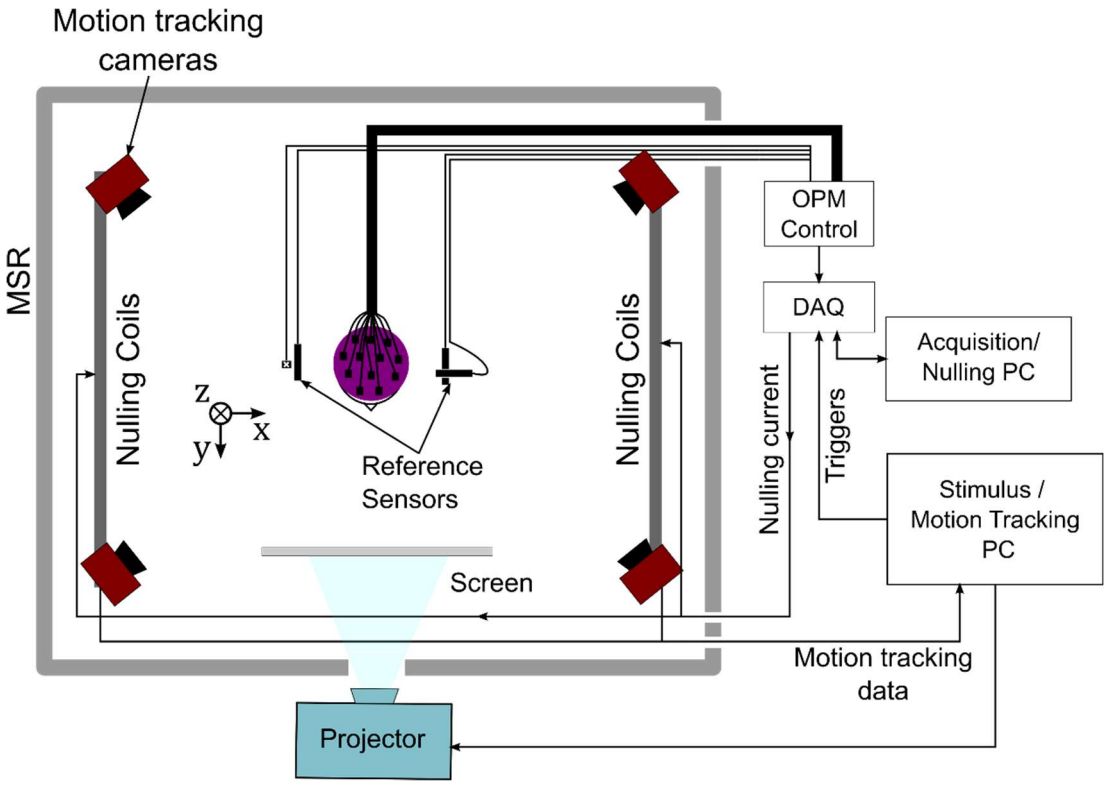

Figure 6.1 Schematic of the OPM-MEG system at the University of Nottingham (Repeated from Chapter 2).

Sensors were preferentially placed in regions near the bilateral motor cortices, parietal and occipital regions with approximately uniform spacing between sensors (see Figure 6.4 for an example of sensor coverage).

Sensors were connected to the control electronics via cables fixed onto a backpack worn by the subjects. The first 60 *cm* of cable connected to each sensor consisted of a lightweight flex cable to reduce helmet weight.

As was introduced in Chapter 2, several means of interference reduction can be combined to ensure a high-quality signal and to keep background fields within the limited operating range of the OPM sensors: Firstly, the sensor array is housed in a magnetically shielded room (MSR) consisting of 4 layers of mu-metal and a single layer of copper (MuRoom, Magnetic Shields Limited, Kent, UK). The mu-metal layers are equipped with degaussing coils (Altarev et al., 2015). A sinusoidal current with exponentially decaying amplitude is applied to

the degaussing coils which periodically magnetises and demagnetises the inner mu-metal layers of the MSR until a pre-defined but small magnetisation is achieved. This reduces the magnetic field at the centre of the MSR to ~2-5 nT. Additionally, the magnetic field around the helmet was controlled using a set of biplanar field nulling coils ((Holmes et al., 2019), Cerca Magnetics Limited, Kent, UK) wound on $1.6m^2$ square planes ($1.5m$ separation) as shown in Figure 6.1. This system of coils contains 8 individual elements capable of generating the three uniform field components and 5 independent linear gradients. Using a reference array consisting of 2 pairs of first generation OPM sensors placed on either side of the seated participant's head (see reference sensors in Figure 6.1), the background field components and first-order field gradients (in x) across the participant's head are sampled and currents are applied to the nulling coils to cancel these background fields.

Once subjects were equipped with the sensor helmet and pointer (see Figure 6.3a), the MSR was degaussed and the remaining background fields and gradients were cancelled as described above. OPM sensors were then initialised: the field nulling coils internal to the sensors were energised to cancel any remaining field and bring the magnetometers into the zero-field regime. Subsequently, a calibration step was performed to produce a sensitivity of 2.7 $V(nT)^{-1}$. All sensor and trigger channels were sampled at 1200 Hz by a DAQ which was connected to a computer used for data storage.

Before beginning the experiment, subjects were given a chance to familiarise themselves with the task, including several practice trials, to minimise any learning effects between runs.

6.2.2 Task: Trail Making Paradigm

A set of mazes was pseudo-randomly generated (e.g. Figure 6.2). Briefly, the procedure for this is as follows:

The screen was divided into a grid of 11×11 cells. Starting from the cell in the top left corner, one of its unvisited neighbours is chosen at random and added to the stack[9] *S* which keeps track of the steps taken on the current path. Once the current path reaches a "dead end" where no unvisited neighbours are available, we "backtrack" by sequentially removing elements from the stack until we find a cell with available neighbours or return to the starting cell. Removing the starting cell, the final element from the stack (step 4), indicates that a complete maze was generated and halts the loop (step 3). Note that mazes generated using recursive backtracking contain no cycles and hence only one simple path between each pair of cells.

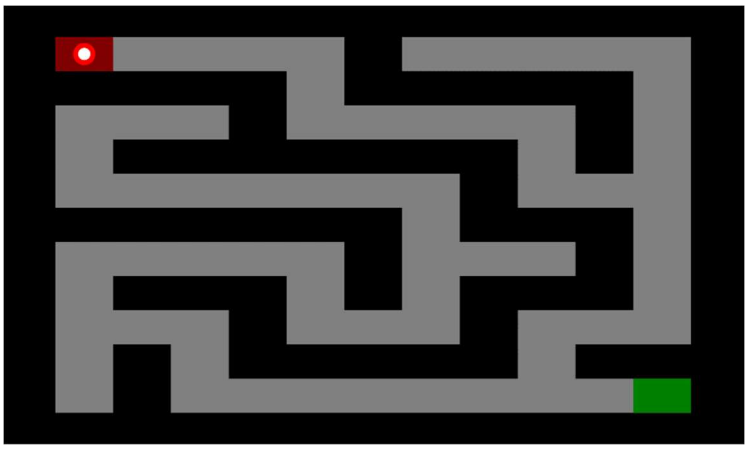

Figure 6.2 Example maze used during the trail making task.

9 In computing, a stack is an abstract data structure. Elements can be added to the top of the stack using the "push" operation. At each point, only the top-most element can be accessed and removed from the stack using the "pop" operation.

A more formal description of the algorithm is given below.

Maze generation: "Recursive Backtracker" Algorithm

```
1 : Generate a square grid of cells separated by walls
2 : Push a cell to stack (S) and mark it as visited (the top-left cell was chosen here)
3 : While S not empty do:
4 :     Pop cell from S; current_cell ← cell
5 :     If current_cell has any unvisited_neighbour then:
6 :         Push current_cell to S
7 :         next_cell ← Randomly select one unvisited_neighbour
8 :         Remove wall between current_cell and next_cell
9 :         Push next_cell to S and mark it as visited
```

The series of generated mazes were displayed on the back-projection screen and volunteers were asked to solve each maze as quickly as possible by pointing at the screen with a pointer and tracing the path from the upper-left corner (red tile) to the lower right corner (green tile) by moving their hand.

After 11*s* of maze traversal, in which the maze may or may not have been completed, subjects were instructed to place their pointing hand on their lap and rest while gazing at a fixation point at the centre of the projected image. Trials lasted for a total of 22*s*. A new maze would then be shown and the next trial would start once a subject pointed to the red starting tile. Each run of the experiment consisted of 40 trials, with the same sequence of pseudo-random mazes.

6.2.2.1 *Subject interaction with the paradigm*

The interactive element of the task required a real-time estimation of where the participant was pointing and a means of ensuring that each maze was solved

correctly. To achieve this, subjects were given a *pointer*—a small rigid body holding 5 asymmetrically arranged infrared reflectors (see Figure 6.3a). The pointer was tracked using six Flex 13 infrared, motion-tracking cameras (*OptiTrack*, NaturalPoint Inc., Corvallis, Oregon, U.S.A.) positioned around the subject. The motion tracking cameras were calibrated and controlled using the NaturalPoint *Motive* software platform. The NaturalPoint *NatNet SDK* software was used to allow sampling and recording of motion data via MATLAB. The fixed and asymmetric nature of the pointer facilitated unambiguous tracking of its location and orientation.

As shown in Figure 6.3b, knowledge of the position and orientation of the pointer could be used to construct a ray in the pointing direction relative to the room coordinates. To estimate the point of intersection of this ray with the screen, we estimated the position and extent of the projected image by placing a reflective marker in the lower-left corner and by measuring the width and height of the image. The coordinate position of the reflector was sampled during the initialisation of the experiment. The orientation of the screen plane was assumed to be parallel to the *xz*-plane (Figure 6.3b), which was defined during the calibration of the motion tracking cameras.

Once the intersection of the pointing ray and the screen plane was estimated, two circular dots were overlaid on the maze image to provide feedback to the participant. A red dot henceforth termed the tracer, and a white dot termed the cursor. The cursor dot was used to display the current pointing position on the screen (similar to a LASER pointer) and vanished if the subject was pointing to a location off-screen. The tracer was initialised in the starting location of the maze (red tile; Figure 6.2) and was restricted to locations between the maze walls. The goal of the task was to use the cursor to "collect" the tracer from the starting position and to "drag" the tracer through the maze into the green target tile to complete the maze. While the cursor was allowed to cross walls, doing so would leave the tracer on the maze path requiring the participant to return to

the tracer. This ensured that subjects used controlled movements and accurately traced the maze path. Completion of the maze was defined as the moment in which the tracer entered the green target tile.

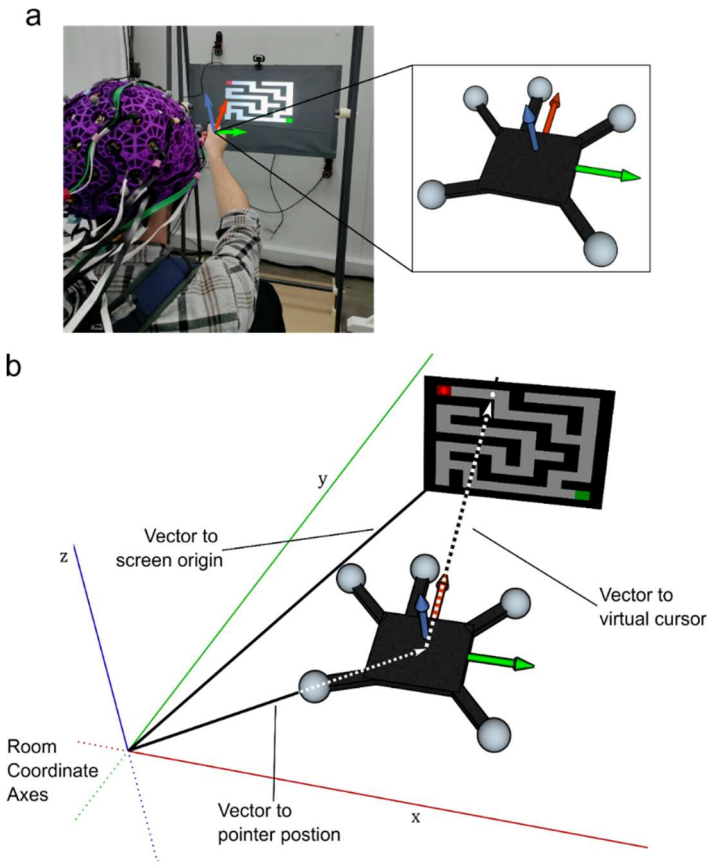

Figure 6.3 Subject interaction with the paradigm. a) Subject pictured wearing the OPM-MEG array during the active part of the task. A maze projected on the back projection screen is shown. The subject is steering a virtual cursor through the maze using the pointer *shown in the box on the right. The pointer is a small rigid body with* 5 *asymmetrically arranged infrared reflectors which were detected using an OptiTrack motion capture system. b) Schematic of the system geometry used to determine the pointing direction of the subject's hand. The room coordinates are defined during the calibration of the OptiTrack cameras. Using this coordinate system, the position of the screen was determined by placing a reflective marker in the lower left corner of the projected image and by measuring its extent using a ruler. The plane of the projection screen was defined to be parallel to the xz-plane. By sampling the position and orientation of the rigid body a ray could be traced from the pointer onto the plane of the projection screen. The virtual cursor (white dot) controlled by the subject was displayed at the intersection of this ray and the projected image.*

6.2.2.2 *Behavioural measures*

We extracted two behavioural measures to assess task performance and the likelihood of behavioural differences or a learning effect between runs:

- **Time to completion:** Mean time from starting the maze—when the cursor dot enters the red starting tile (see Figure 6.2 or Figure 6.3b)—to reaching the green target tile averaged over trials completed in less than $\mathbf{11}\boldsymbol{s}$.
- **Completed Trials:** Percentage of trials completed within a maximum of $\mathbf{11}\boldsymbol{s}$.

6.2.3 Coregistration

To enable source localisation, an optical coregistration procedure (Zetter et al., 2019) shown in Figure 6.4 was used to estimate the position of the sensor array relative to subject anatomy:

A T1 weighted anatomical MRI (1mm isotropic resolution) was available for each subject and scalp meshes were extracted using SPM12 as part of FieldTrip (Oostenveld et al., 2011). The positions and orientations of sensors relative to the helmet were known a-priori from the computer-aided design (CAD) file of the helmet. Alignment of the 3D helmet model and the scalp mesh required two optical scans performed using a 3-dimensional optical imaging system (Structure IO camera—Occipital Inc., San Francisco, CA, USA) in conjunction with the SKANECT software (see Figure 6.4): one digitisation of the subject's face while wearing the sensor array and one of the subject's head and face without the helmet.

After each recording and before the digitisation, subjects were asked to stick coloured, adhesive markers to facial landmarks (nose, cheekbones) to aid the optical coregistration procedure. Using MeshLab, the mesh of the subject's head was aligned to the scalp mesh by selecting at least 4 landmarks such as the tip

of the nose, the corners of the eyes and the centre of the brow. The point cloud containing the participant's face and helmet was aligned to the head-only point cloud using prominent facial features and the coloured adhesive markers.

Finally, the CAD model of the helmet was aligned with the point cloud of the subject wearing the helmet, yielding a transformation matrix which aligns the helmet with the anatomical structures. This transformation matrix was then applied to the sensor locations—which were known from the 3D helmet model—yielding sensor locations in the coordinate system of the anatomical.

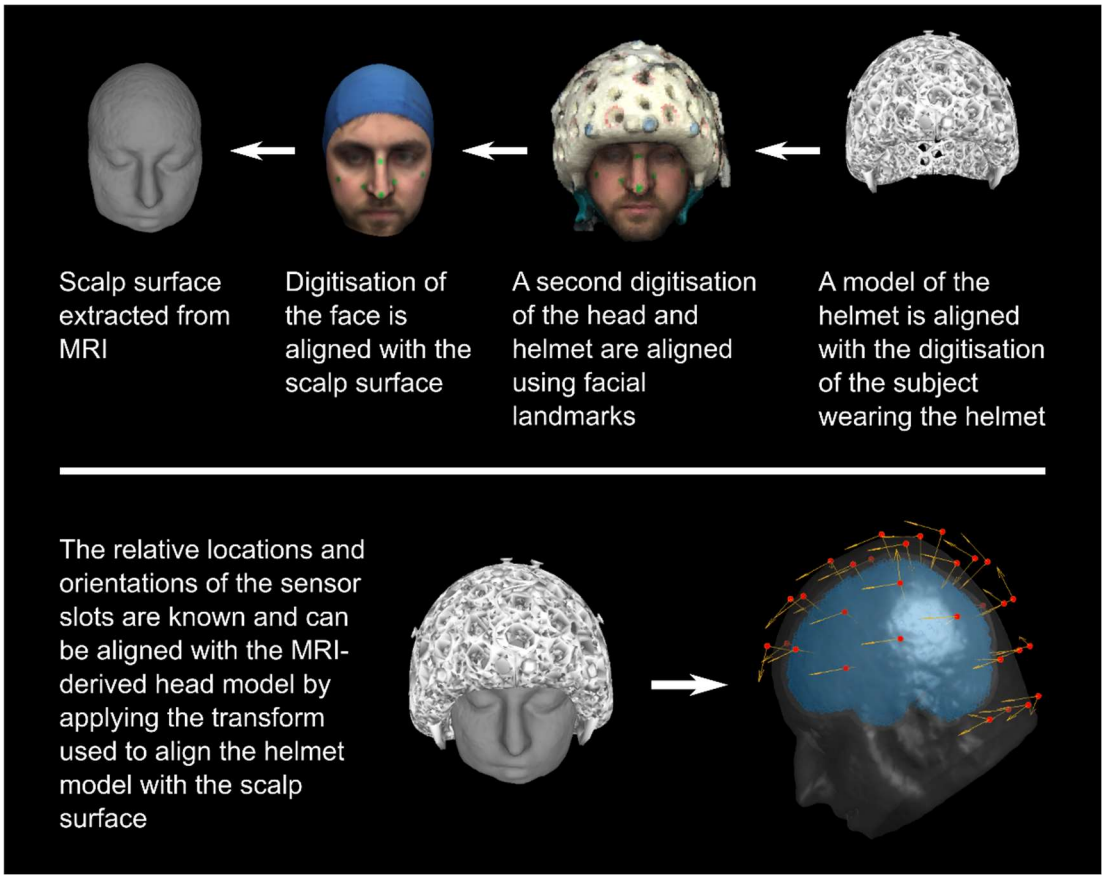

Figure 6.4 Coregistration procedure. Computer aided design (CAD) of the sensor helmet provides us with accurate knowledge of the relative positions and orientations of the sensors. To achieve full coregistration with the subject, we perform two 3D digitisations of the subject's head—one while wearing the sensor array and one while wearing a swimming cap to flatten any hair as much as possible. The head-only scan is aligned with the scalp surface extracted from the anatomical MRI and facial features are used to align the scan containing the helmet and participant's face. Aligning the CAD model with this latter digitisation yields coregistration between the subject anatomy and the sensor locations as shown in the image in the bottom right. Red dots depict sensor positions and yellow arrows indicate all sensitive axes of each sensor: one radial to the scalp surface for each sensor and an additional one or two tangential directions for dual and triaxial sensors respectively.

6.2.4 Preprocessing

Before data inspection, a notch filter (second-order IIR filter at 50 Hz) was applied to reduce mains interference. For inspection, data were filtered using a 1-150 Hz bandpass filter and divided into epochs lasting 22 s and beginning 5 s

before maze completion. All epochs were inspected and any noisy channels or trials containing artefacts were identified. Channels containing excessive noise or showing signs of malfunction were removed and excluded from subsequent analyses. The remaining channels were subject to homogeneous field correction (Hill et al., 2022; Tierney et al., 2021), further reducing interference. To avoid edge artefacts due to filtering, data were filtered to the bands of interest before unwanted trials were removed. For time-frequency analyses, the field corrected data were filtered between 1 and 150 Hz; for source-localisation of the MRBD and to estimate our functional connectivity measure, data were filtered to the beta band (13-30 Hz). Each time a 4^{th}-order Butterworth filter was used.

6.2.5 Analysis

After pre-processing and data cleaning, we constructed a series of common measures of brain activity to assess whether the expected beta band responses would be induced by our paradigm. First, TFS analysis of the sensor-level data was conducted to visualise whether movement-related desynchronization and a post-movement rebound could be observed in the vicinity of the left motor cortex:

Mean-field corrected data were band-pass filtered to 26, narrow, overlapping frequency bands between 1 and 120Hz. The amplitude envelope for each band was constructed by calculating the absolute value of the analytic signal, which was calculated by applying the Hilbert transform. These envelopes were averaged over trials and normalised relative to a baseline period (7-12 s after maze completion at 0 s).

6.2.5.1 *Source Reconstruction*

A scalar beamformer (see Chapter 3, (Robinson and Vrba, 1999)) was used to generate source-space maps of activation and signals for subsequent analyses. To generate the source model, the brain was extracted from the anatomical

MRIs using FieldTrip (Oostenveld et al., 2011) and downsampled to voxels of $4mm$ isotropic resolution. A single sphere model (Sarvas, 1987) was used to calculate the lead fields (see Chapter 3) for each voxel. Source orientations at each voxel position were estimated by calculating lead fields for two orthogonal directions in a plane tangential to the head-surface before determining the optimal source orientation using a singular value decomposition approach (see Chapter 3, section 3.2.1.1.2, (Sekihara et al., 2004)). To estimate source maps of task-induced activity, data were filtered to the beta band (13-30 Hz), and 'bad' trials removed. Clean, filtered data were concatenated and covariance matrices were calculated to generate the beamformer weights after applying Tikhonov regularisation with 5% of the maximum singular value of the covariance matrix (Chapter 3, (Brookes et al., 2008)). Using an active window set during maze traversal ($t = -5\text{-}0\ s$) and a control window during the resting period ($t =$ $7\text{-}12\ s$) a pseudo-T-statistical (T) image was generated to show regions exhibiting MRBD with negative T-values.

To allow for the estimation of a group average, facilitate comparisons between subjects and perform further analyses, source maps were aligned to a common space—the MNI152 template anatomy (Evans et al., 2012)—using the FLIRT tool as part of the FSL suite (Jenkinson et al., 2002; Jenkinson and Smith, 2001).

Using broad-band (1-150 Hz) data, virtual electrodes were constructed at the peak MRBD location found in the group average, transformed to each subject's anatomy and run to estimate a trial averaged TFS using the same methods as described for the sensor-level analysis above. The envelope of the beta-band signal derived from the TFS analysis was used for subsequent analyses to assess repeatability.

6.2.5.2 *Functional Connectivity*

As previously described in Chapter 4 and Chapter 5, functional connectivity was estimated using 78 source locations defined by the AAL Atlas (Gong et al.,

2009; Tzourio-Mazoyer et al., 2002). The extracted brains from each MRI were coregistered to a template brain—again using FLIRT—and the inverse transformation of this alignment was used to determine coordinates of the atlas region centroids at the corresponding locations in each subject's brain. To estimate beta-band amplitude envelope correlation, data filtered to 13-30 Hz, and the whole duration of concatenated clean trials, were used to estimate the covariance matrix. As described in section 3.3, virtual electrodes were constructed at the atlas locations, Hilbert envelopes were estimated and pairwise leakage correction was applied (see Chapter 3, section 3.3.2 (Brookes et al., 2012)). The corrected time courses were down-sampled to a sampling rate of 5Hz and used to calculate the Pearson correlation coefficient between each seed and test envelope.

6.2.5.3 *Note on Reproducibility measures*

Showing that a choice of paradigm or instrumentation is suitable for future application in within-subject diagnostic MEG scans requires information regarding the reproducibility of MEG-derived brain measures. Here, two key concepts can be defined: *Reliability*—the sensitivity of a method to between-subject differences—and *Agreement*—the numerical similarity between repeated measures (Haghayegh et al., 2020).

While between-subject variability has been seen as a nuisance in the past (and indeed it is treated as such in the previous chapters of this thesis), more recent work on neural fingerprinting has shown that interindividual differences in brain activity, especially functional connectomes in MEG (Sareen et al., 2021) and fMRI (Amico and Goñi, 2018; Finn et al., 2015), can be used to reliably identify subjects based on their patterns of activation and may also be related to behavioural differences (Finn et al., 2015). We used a fingerprinting measure to assess this ability to differentiate subjects despite noise and measurement errors to characterise the intrasubject *reliability* of OPM scans.

Image consistency for source-localised activity has previously been assessed using Pearson correlation between repeated measures of the same subject for OPM-MEG, and SQUID MEG systems (Hill et al., 2020). Similarly, correlations between the MEG signals acquired using OPMs at two different sites (in the same subject) have been reported (Hill et al., 2022). Likewise, the consistency of reconstructed functional connectivity matrices was assessed using correlation coefficients (Boto et al., 2021). While these results indicate a degree of linearity between the shapes of activation patterns, time courses or connectivity matrices from repeated measures, they fall short of characterising *agreement*. The correlation coefficient only measures the strength of the linear relationship and is scale-independent (i.e. the slope of the line of best fit is irrelevant). In the case of separate but related quantities, a significant correlation may indicate a relationship of interest. For repeated measures of the same quantity, however, it would be surprising if there was no relationship to be found. Additionally, highly correlated repeated measures, which are of different orders of magnitude, for example, do not agree; a fact obscured by the correlation coefficient.

These misconceptions were collated by Bland and Altman (1986), who described a simple method of measuring agreement using a plot of the difference between measurements against their average. Here, we used these 'Bland-Altman' plots to numerically estimate the agreement between repeated measures of brain activity and give limits of agreement which may be used in future works to generate hypotheses or estimate necessary sample sizes to reliably measure effects between patient and control populations.

6.2.5.4 *Fingerprinting*

To assess whether our OPM-MEG array was sensitive to repeatable, interindividual differences in brain activity, we performed spatial, spectrotemporal, and connectivity-based neural fingerprinting analyses as illustrated in Figure 6.5. The chosen imaging measure was vectorised, e.g. the pseudo-T statistics measured at each voxel (in MNI 152 space) were rearranged

into a 1*D* vector for each subject and run; for connectivity matrices, the upper triangle of the symmetric matrix was used. The first runs were used as a "matching database"; each second run was compared to each vector in the matching database and Pearson correlation coefficients were calculated to populate the identifiability matrix—a 6×6, asymmetric matrix. For each row, the value on the diagonal of this matrix gives the within-subject correlation, while off-diagonal values are between-subject correlations of each run in the matching database with the other subject's second runs. A successful fingerprint means row maxima are found on the diagonal.

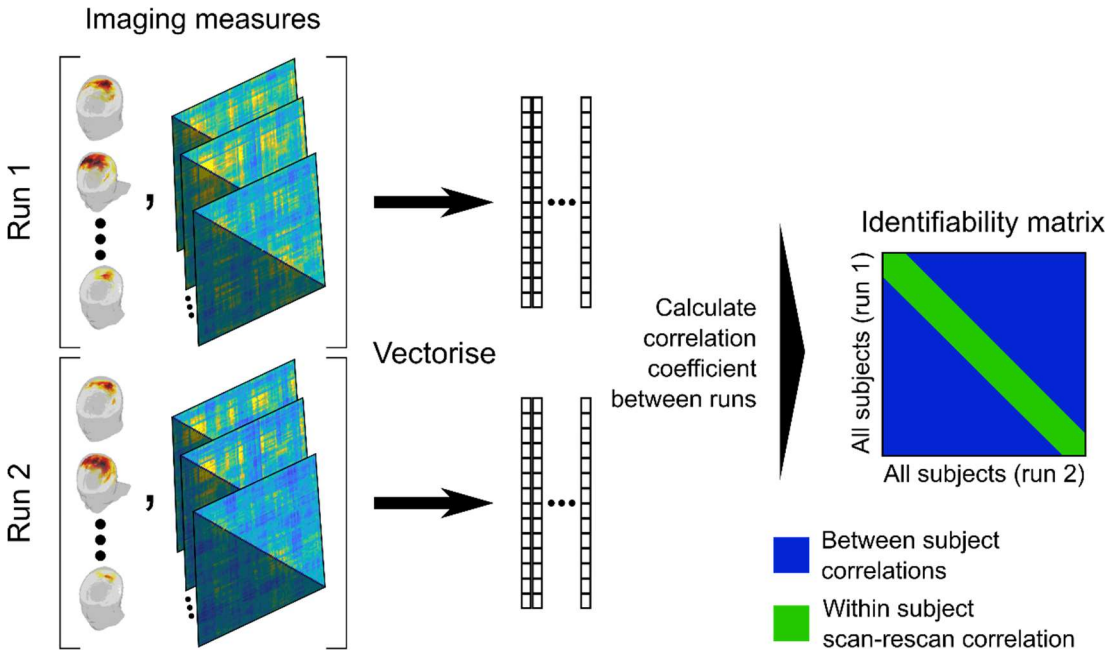

Figure 6.5 Fingerprinting analysis. The imaging measure of choice is vectorised for both runs in each subject. For the source maps of induced activity, this includes coregistration with a template brain and extraction of relevant values using a cortical mask (union of AAL atlas regions) which are flattened into a vector. For the symmetric connectivity matrices, only unique edge strengths (upper or lower triangle of the matrix) are extracted and flattened. The Pearson correlation coefficients between the first and second runs are calculated for all pairs of subjects yielding the identifiability matrix. Values on the diagonal contain the within subject, scan-rescan correlations. Off-diagonal elements contain correlations between subjects. Note that this matrix will not be symmetric. For an ideal fingerprint measure, row maxima should lie on the diagonal, meaning all subjects can be correctly identified from the second scan based on their first scan.

6.2.5.5 *Bland-Altman Analysis*

We estimated repeatability measures based on the Bland-Altman plot (Martin Bland and Altman, 1986), which consists of a scatter plot of the differences between pairs of measurements against their means (see Figure 6.6c for an example). Bland-Altman plots were used to better visualise the agreement between runs of the experiment. Additionally, two measures were extracted to give a numerical estimate of the agreement between measures: The *bias*, which we define as the mean of differences between paired measurements (Blue line in Figure 6.6c), and the *range,* which is 1.96 times the standard deviation of the differences[10] (absolute difference between the black horizontal lines and the blue line in Figure 6.6c). Bland and Altman refer to the 'limits of agreement' (bias $\pm$ range)—the confidence interval of the differences—or the confidence interval (CI) within which 95% of differences lie. A bias close to zero is desirable as it indicates that the overall magnitude of the measurement is stable between experiments on average; a small range (relative to the magnitude of the measurement of interest) is desirable as it gives an estimate of the smallest true change that can be reliably measured.

10 This stems from the fact that 95% of the area under the standard normal distribution corresponds to the area between -1.96σ and 1.96σ (the 95% confidence interval), where σ is the standard deviation.

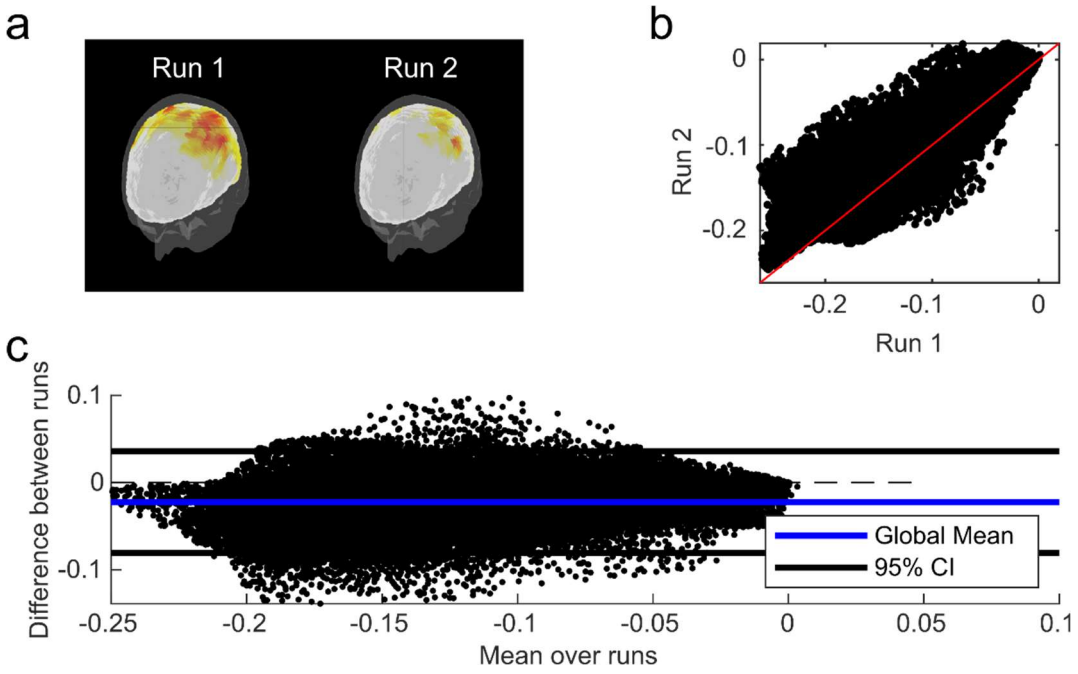

Figure 6.6 Repeatability analysis using Pearson Correlation and Bland-Altman analysis. a) Example of MRBD source maps from two separate scans of a single subject. b) Scatter plot of the vectorised MRBD maps shown in panel a. The red line shows the line of equality between both measures (not a linear fit!). The Pearson correlation coefficient of this example is 0.84. *c) Bland-Altman plot of the vectorised source maps shown in a. The blue horizontal line shows the mean of differences between runs or* bias (here, bias = −0.022). *The black horizontal lines indicate the* 95% *confidence interval (*95% *CI) of differences, or* ±1.96 *standard deviations which we will call* range *of differences* (here, range = 0.03).

6.3 RESULTS

6.3.1 Outline of MEG Measures

The main MEG results presented here will focus on three main aspects:

- A '**sanity check**', qualitative examination of group averages to show whether the novel task paradigm would elicit expected patterns of activity (MRBD, PMBR, connectivity in the motor network)
- **Neural fingerprinting** matrices
- **Bland-Altman plots** to show repeatability between runs and estimates of *bias* and *range*—as defined in 6.2.5.3—for within-subject repeats.

We will cover each of these aspects for three common measures of brain activity: Source maps of task-induced beta desynchronization, the spectrotemporal evolution of neural oscillations with a focus on the beta band, and beta band functional connectivity measured using amplitude envelope correlation.

6.3.2 Data availability and Behavioural Results

After data cleaning, an average of 7 ± 2 channels and 1 ± 2 trials (mean $\pm$ std. deviation) were excluded from MEG signal analysis. Task performance was good with all subjects completing a majority of trials within the given time as shown in Figure 6.7. Maze completion time and the proportion of completed trials did not vary significantly between runs. The average difference between runs was 0.2 ± 0.2 *s* for completion time, and $-4.5 \pm 3.8\%$ for the proportion of completed trials (mean $\pm$ std. error).

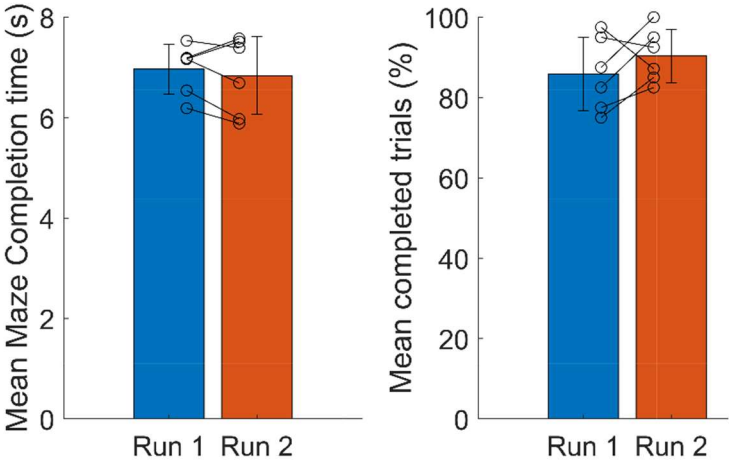

Figure 6.7 Performance measures. Average completion time (left) and percentage of completed trials (right) for both runs. Error bars indicate the standard deviation. Black circles connected with lines indicate values from the same subject.

6.3.3 Sensor-level TFS and MRBD source maps

Figure 6.8a shows a TFS averaged over subjects for each run for a single sensor located approximately over the left motor regions (see Figure 6.8b for an example of sensor coverage including the sensor for which the TFS plots were generated. As expected, both plots show a reduction in beta-band power in the final 5*s* of the active portion of the trial, with a rebound following maze completion at 0*s*. Figure 6.8c shows a pseudo-T statistical image (in MNI152 space) averaged over all twelve recordings. As expected, beta-desynchronization is most prominent in the motor cortices with the largest reduction found in the left hemisphere, contralateral to the hand used during the task.

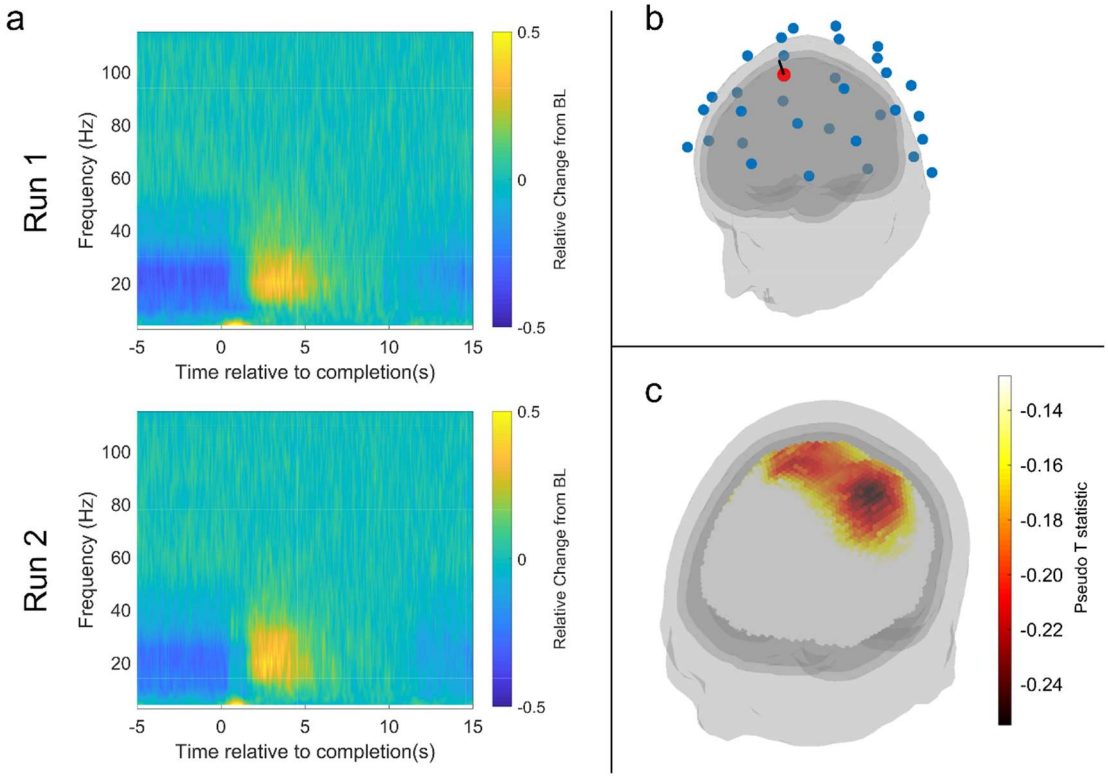

Figure 6.8 Sensor and source level analysis. a) Mean TFS for a single sensor located approximately above the left motor cortex (radial sensitive axis shown as a black line in panel b), averaged over trials and subjects. Movement related desynchronization of beta band activity is clearly visible in the final 5s of the active period, followed by a rebound after maze completion at 0s. The baseline (BL) period was chosen to be 7-12s after maze completion—the final 5 seconds of the resting period prior to starting the next trial. b) Example of the arrangement of sensors for a single subject (blue dots) and the sensor chosen to produce the TFS in a). c) Mean pseudo-T statistic across subjects showing beta desynchronization (active window: −5-0s; control window: 7-12s). Peak desynchronization was found in the left sensorimotor region.

Figure 6.9 shows all MRBD source maps. While the majority of images show activation in the expected regions, between-subject variability (left to right) appears high, with patterns displaying more consistency when making within-subject (top to bottom) comparisons.

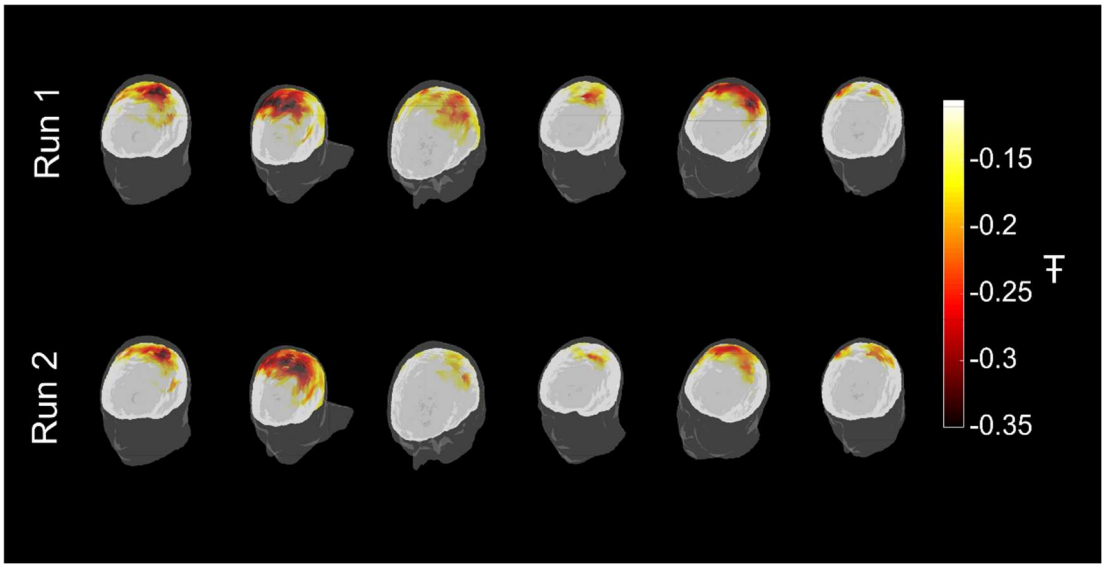

Figure 6.9 Pseudo T statistic images of beta desynchronization during the active window for all subjects and both runs. Most runs show peaks in the vicinity of the left motor cortex as expected.

In support of this, quantitative analysis via spatial fingerprinting using the MRBD source maps showed excellent sensitivity to interindividual differences with all subjects' second runs being correctly identified based on their first run (see Figure 6.10a). Mean within-subject and between-subject correlations were found to be 0.84 and 0.50 respectively. Figure 6.10b shows a Bland-Altman plot for all subjects combined. The global mean of differences is close to zero suggesting low bias between runs. Figure 6.10c and d show within-subject bias and range respectively. Range values are comparable between subjects.

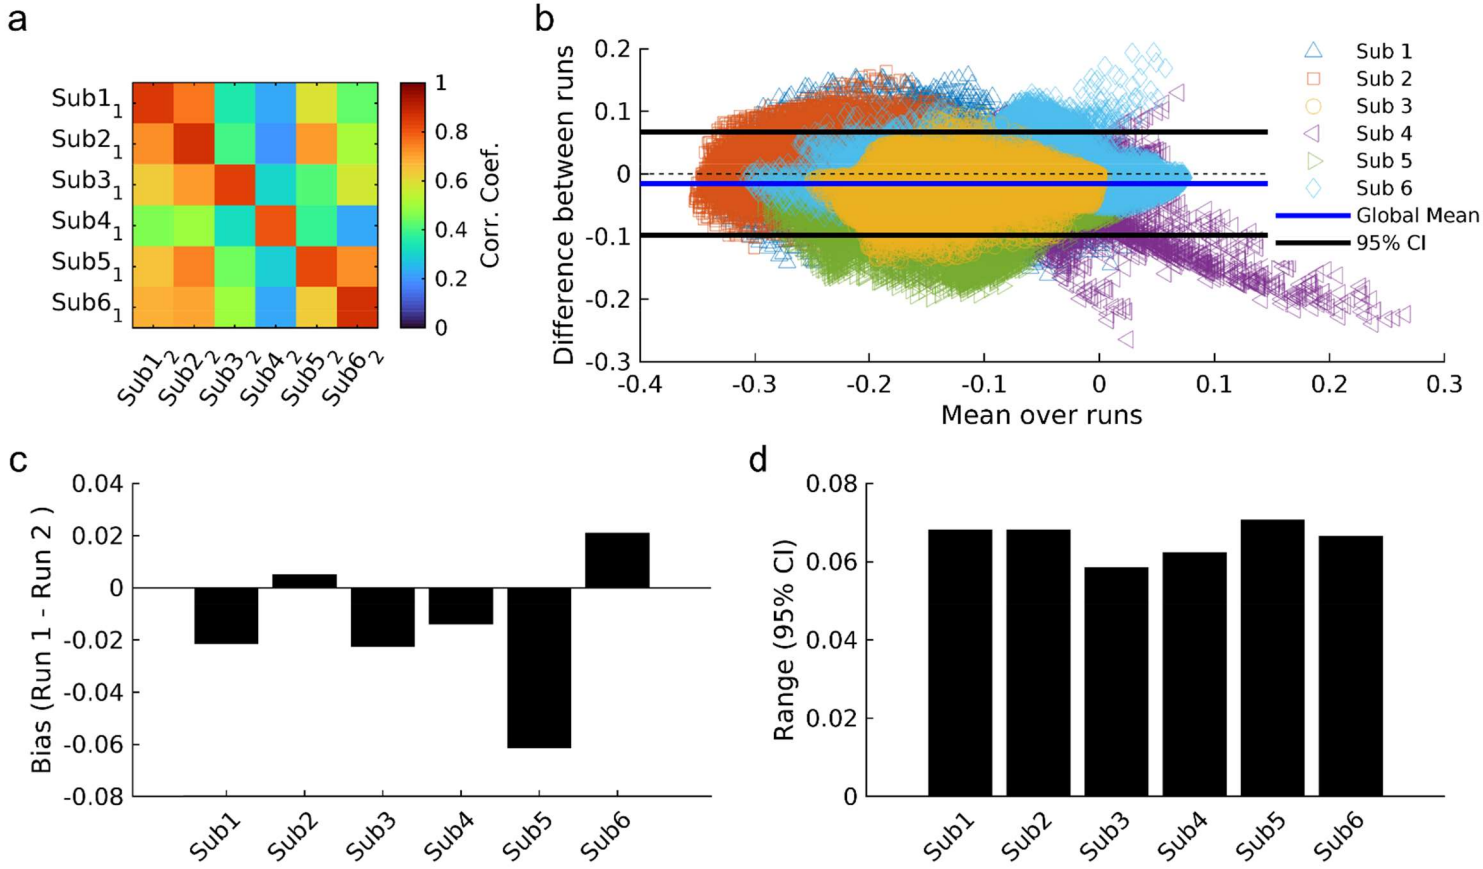

Figure 6.10 Repeatability of MRBD assessed using spatial fingerprinting in source space. a) Fingerprinting matrix. All row maxima are found on the diagonal, meaning all subjects were correctly identified. b) Bland-Altman plot of MRBD (all source space voxels) over all subjects. The global mean is close to zero suggesting low bias between measurements. c) Within-subject bias (mean difference between runs over MRBD source maps). d) Within-subject range (95% confidence interval of differences over MRBD source maps).

6.3.4 Source-space beta envelope

Figure 6.11a shows group average time-frequency spectra generated from virtual electrodes reconstructed at the peak MRBD location. A reduction in beta-band power can be seen during the active portion of the task in both runs. Figure 6.11b focuses on the evolution of the beta amplitude throughout the trials, showing the mean beta time-course over all runs with the standard deviation being represented by the shaded area, again revealing the MRBD. Compared to the sensor-level TFS, the PMBR is less prominent in the average time course. Inter-subject variability in the amplitude of the rebound appears to be elevated.

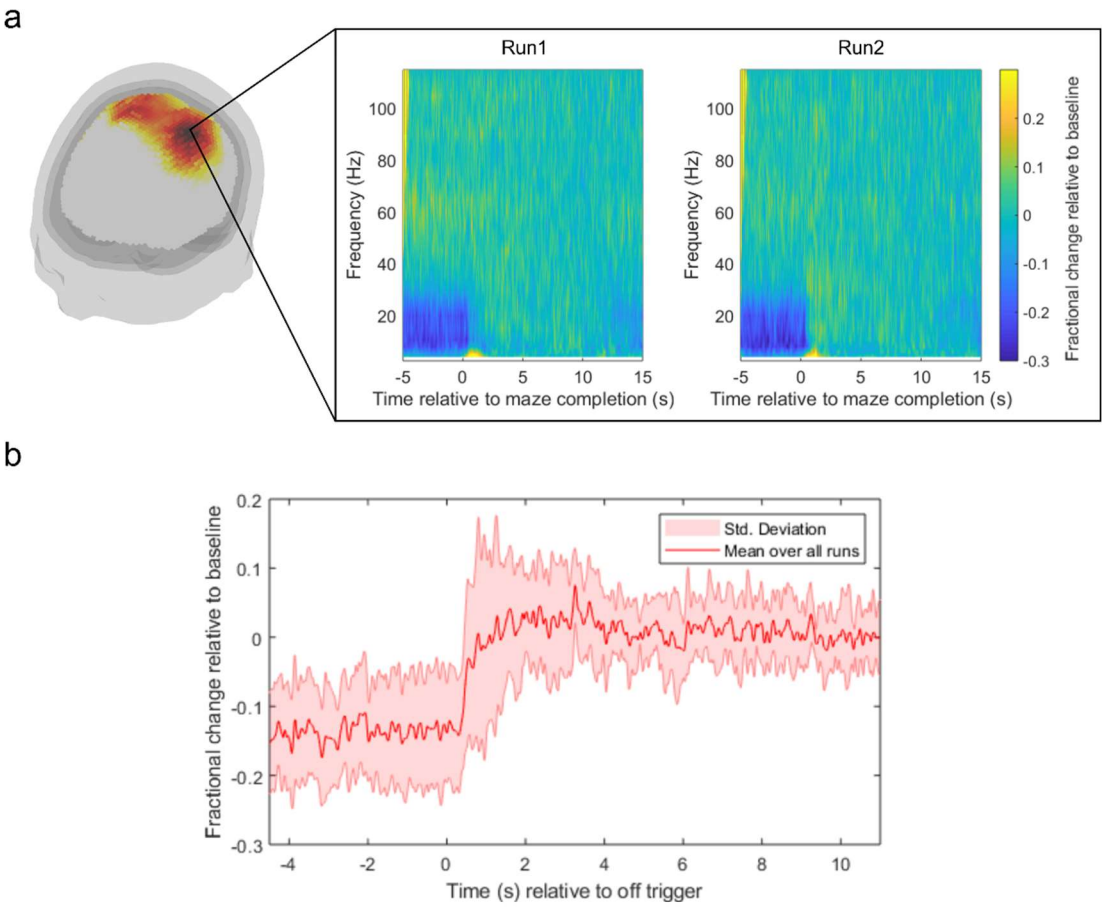

Figure 6.11 a) Mean TFS of the time-courses reconstructed at the peak MRBD location. b) Mean of the trace of the beta band envelope over all subjects and runs, extracted from the TFS in a.

Figure 6.12 shows fingerprinting and Bland-Altman analysis results for the beta envelope. Fingerprinting (Figure 6.12a) was poor with only one subject being identified correctly. Mean within-subject and between-subject correlations were found to be 0.63 and 0.54 respectively. Bland-Altman analysis (Figure 6.12b-d) reveals low bias relative to the range values found.

Figure 6.12 Repeatability of the beta envelope reconstructed at the peak MRBD location. a) Fingerprinting matrix. Only one of 6 subjects was correctly identified. b) Bland-Altman plot for the source beta-band envelope over all subjects. The global mean is very close to zero suggesting low bias between measurements. c) Within-subject bias (mean difference between runs). d) Within-subject range (95% confidence interval of differences over the beta envelopes).

6.3.5 Amplitude envelope correlation

Figure 6.13 depicts the mean whole trial connectivity matrix averaged over all scans and corresponding glass brains showing the strongest 5% of connections (the strongest 150 of a total of 3003 unique connections). As expected, the dominant connections form a canonical motor network, with the strongest nodes found in the vicinity of the bilateral motor cortices.

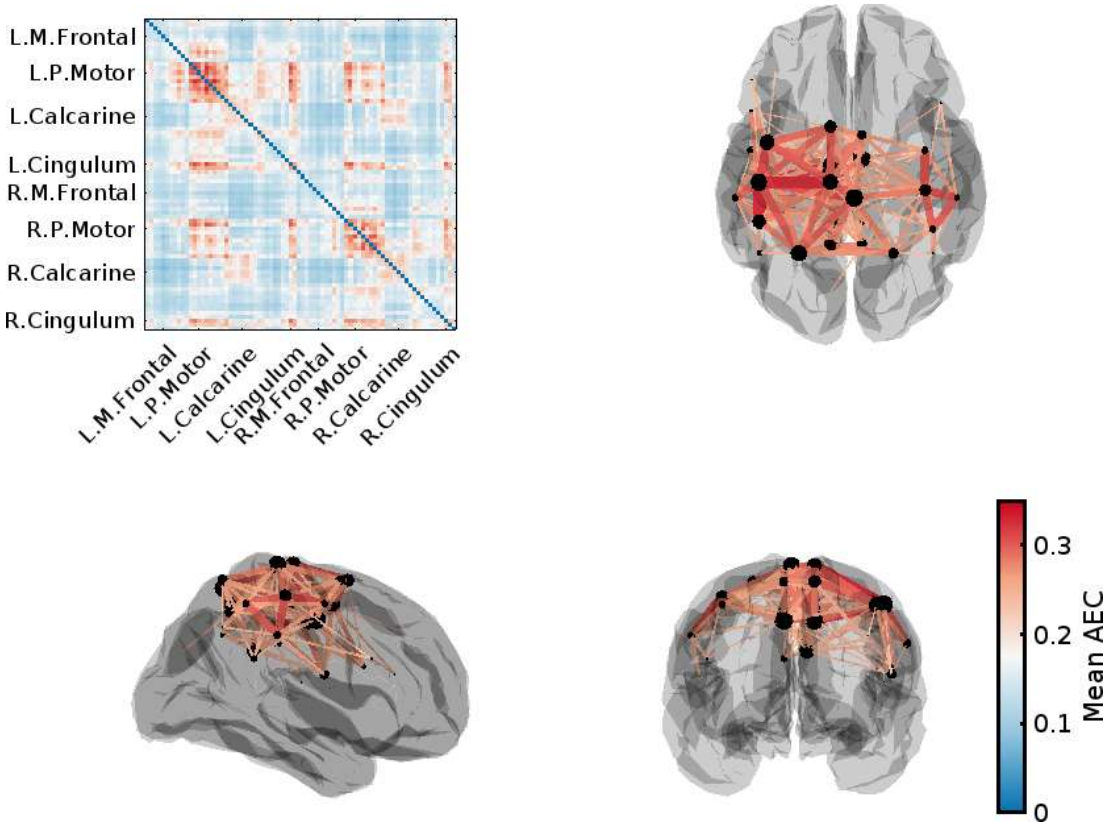

Figure 6.13 Beta band functional connectivity as measured by amplitude envelope correlation. The connectivity matrix and glass brains show the average connectivity collapsed over trials and subjects. The strongest 5% of connections in the glass brains clearly show the bilateral motor network.

Figure 6.14 contains glass brain plots for each subject and run displaying the top 5% of connections as measured by beta band AEC. The expected motor network patterns appear in most runs; in a similar vein to the MRBD maps, much inter-subject variability is apparent. Five out of six subjects were correctly identified using beta AEC connectome fingerprinting (see Figure 6.15a). Bland-Altman analysis (Figure 6.15b-d) revealed low bias.

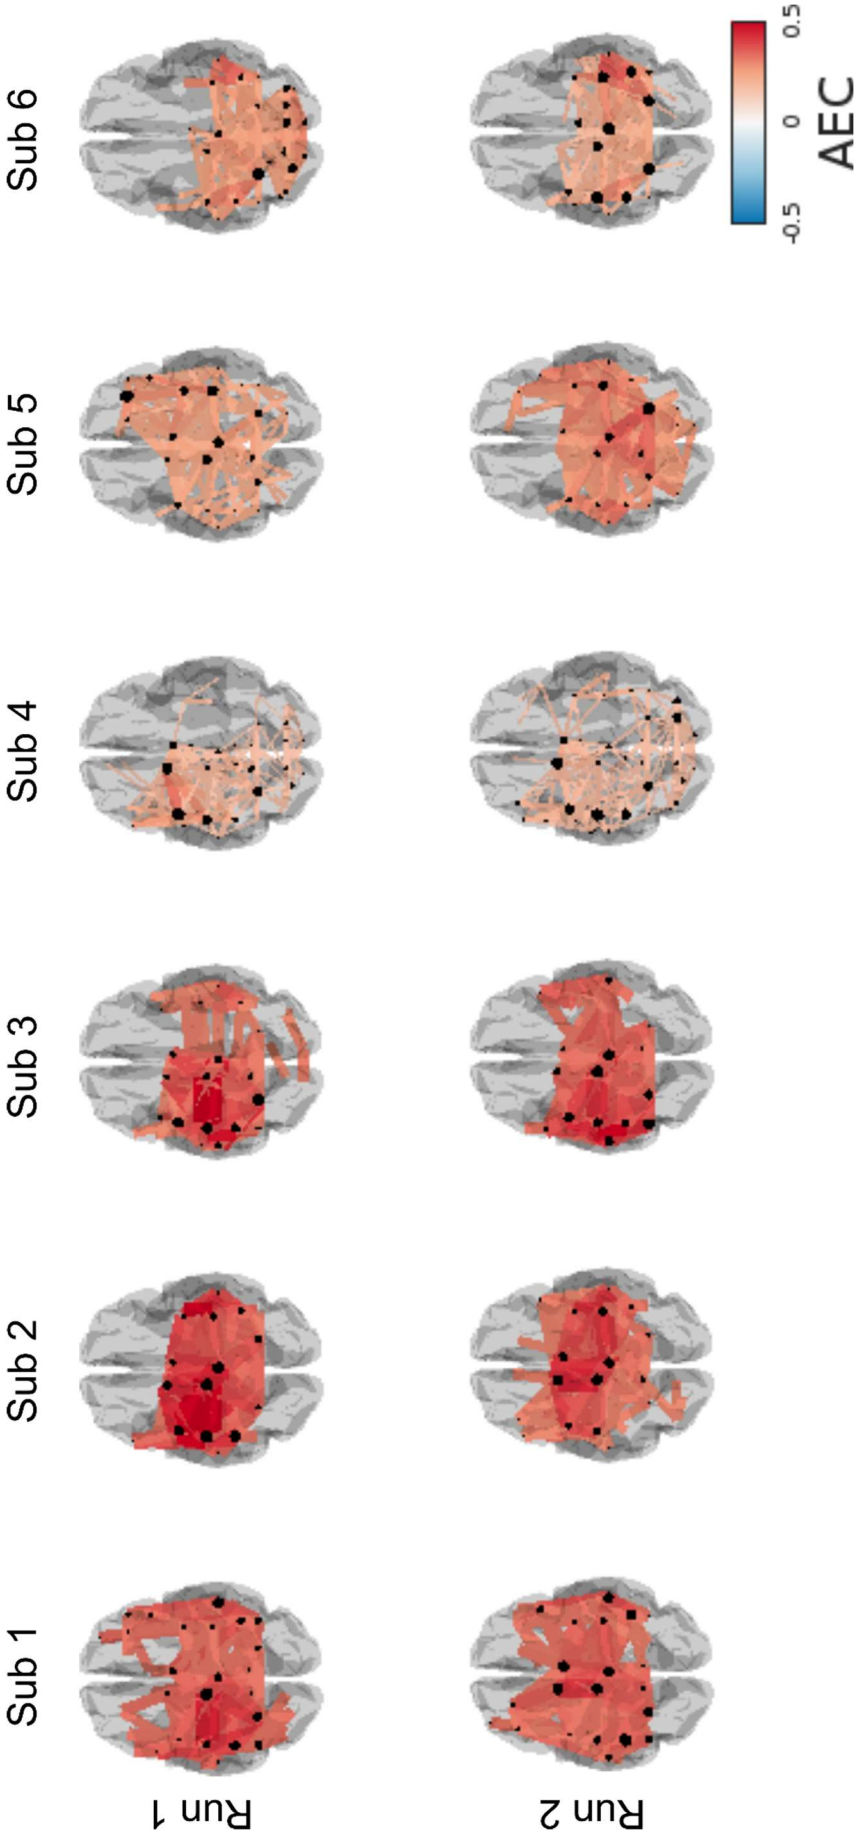

Figure 6.14 Glass brain plots showing beta band functional connectivity for each subject and run.

Figure 6.15 Repeatability of average beta-band connectivity. a) Fingerprinting matrix. Five out of six subjects were correctly identified. b) Bland-Altman plot for beta band connectivity over all subjects. The global mean is very close to zero suggesting low bias between measurements. c) Within-subject bias (mean difference between runs). d) Within-subject range (95% confidence interval of differences over the beta envelopes).

6.4 DISCUSSION

In this study, we have assessed the within-subject consistency of OPM-MEG by repeated scanning of subjects completing a novel task paradigm and reconstructing a series of brain activity measures related to motor function, which was previously found to be abnormal in mTBI (see Rier et al. (2021); Chapter 4). This was done to quantitatively estimate the range of measurement variability, which may inform the optimisation of both system design and task paradigms in longitudinal studies of mTBI using OPM-MEG.

6.4.1 A summary of Findings

6.4.1.1 *Behavioural measures*

Two behavioural measures were collected to ensure homogeneous task performance across subjects and to assess whether any learning effect may have occurred between repeated runs: we measured the mean time taken to complete each trial as well as the number of trials which were completed within the time limit of $11s$. Although the small number of subjects scanned did not warrant statistical comparison, the mean differences between runs of completion time and completed trials were close to 0, indicating that performance was (qualitatively) stable between runs. This can be attributed to the random sequence of mazes used. The conventional trail making test, in contrast, uses a single standardised pattern which may be memorised more easily, resulting in faster completion times after each repeat.

6.4.1.2 *Qualitative examination of activation patterns*

To assess whether the novel task employed in this study would elicit the expected patterns of brain activity we inspected time-frequency spectra of sensors located above the left motor regions contralateral to the hand used during the task. Averaging over trials, and stimulus locking at the point of maze completion, the final seconds of movement show a clear reduction in beta power

(MRBD), followed by a rebound (PMBR) immediately after movement cessation in both runs. Source localisation of the MRBD (contrasted with a resting baseline period) successfully showed the expected pattern of activation across the bilateral motor regions with peak beta-power reduction in the left motor cortex when averaging over all runs. We also assessed spectrotemporal dynamics in source space, by reconstructing a virtual electrode at the peak MRBD location and producing time-frequency spectra as well as time courses of the beta envelope. Again the MRBD could be seen clearly. The grand average over subjects and runs did not show a PMBR. This can be expected as the peak location of PMBR and MRBD have been shown to be anatomically separated (Fry et al., 2016; Jurkiewicz et al., 2006). Functional connectivity measurement using amplitude envelope correlation revealed the expected topology, with the strongest connections forming a network between the frontal and parietal regions associated with the motor network in and between both hemispheres.

6.4.1.3 *Fingerprinting and Bland-Altman Analysis*

We used a fingerprinting measure in conjunction with Bland-Altman plots to assess intra-subject repeatability for three measures based on our OPM-MEG recordings: source-space maps of the MRBD, the trial-averaged envelope of beta-band activity, and the beta-band connectome averaged over the whole experiment. This was done to test the sensitivity of OPM-MEG to interindividual differences as well as quantitatively assess the agreement between repeated measures of brain activity when using wearable OPM-MEG. Pseudo-T statistical images of the MRBD revealed substantial variability between subjects, however, individual topographies remained relatively constant between runs (see Figure 6.9). This was confirmed by our fingerprinting measure which showed the mean within-subject correlation of 0.84 dropping to 0.50 between subjects. This resulted in 100%, successful identification. In addition to high identifiability—meaning the within-subject correlations are higher than between-subject correlations—Bland-Altman plots

revealed low bias and range compared to the peak MRBD values, suggesting good test-retest reliability.

Fingerprinting analysis of the trial-averaged beta-band envelopes was less successful. This does, however, reveal the importance of assessing differences between repeated measures instead of (or in addition to) correlation analysis. We observed very low bias compared to the signal of interest (the MRBD in this case) indicating no systematic difference between repeated measures. On the other hand, range values were large, likely reflecting the amount of noise in the envelopes. Even with perfect agreement of the underlying signal, an increase in noise is monotonously associated with a decrease in the correlation coefficient. Additionally, the shape of the MRBD response and return to baseline can be expected to be very similar across subjects. The high similarity of the underlying signal but with added noise may explain why between-subject correlation was relatively high and comparable to within-subject correlation, in turn leading to poor identifiability. In the context of potential longitudinal assessments of mTBI using this measure, effect sizes indicating an abnormality would need to be very large to be reliably measured. Here, analysis of the range provided by Bland-Altman plots may be a means of optimising the number of trials needed to reduce the remaining noise to the desired level by trail-averaging.

Connectivity matrices showed moderate identifiability. Intra-subject correlations were high (0.63), however, there was also a strong correlation between subjects (0.54) which—similarly to the beta envelope—can be attributed to the specificity of task-induced signals. Bland-Altman plots showed promisingly low bias, and range values, indicating good intra-subject repeatability.

While it is possible to make judgements about the repeatability of these measures based on bias and range values compared to the average observation, Bland-Altman analysis—by design—does not give a definite dichotomous test

of repeatability. Instead, it provides a quantitative limit on the effect size that can be reliably measured. Whether measured differences of such a size would be clinically relevant will have to be determined before designing a diagnostic tool based on OPM-MEG.

6.4.2 Limitations

There are several limitations to note in this study. Firstly, the small number of subjects reduces our precision in estimating the identifiability measure derived from our fingerprinting matrices. Additionally, it was not feasible to produce useful Bland-Altman plots for single measures per subject—like global mean connectivity for example.

Our choice of producing Bland-Altman plots across voxels, time points or connectome edges allowed us to estimate the average within-subject range over all samples. There are, however, two implicit assumptions: First, that samples within a measure are independent, and secondly that the data are homoscedastic, i.e. the spread of differences is constant regardless of the magnitude of the measurement. Given that all of our measures have a degree of structure, independence between each sample cannot be fully guaranteed. In order to apply the method as intended by Bland and Altman, separate plots and estimates of bias and range would have to be produced for each voxel, time point etc. Inspection of the Bland-Altman plots produced in Figure 6.10, Figure 6.12, and Figure 6.15, indicates that the differences are mostly homoscedastic, however, some outliers are visible. In Figure 6.10 for example, the data for subject 4 show a number of samples that deviate from the otherwise normal distribution, however, given that their mean values are positive they likely do not fall within the region of interest for the MRBD measured here.

Because the assumptions laid out here are only loosely met, care needs to be taken when interpreting our results. The range values produced should not be seen as a true estimate of the minimum effect size that can be reliably measured.

Rather, it provides a worst-case estimate (that may still be of use for paradigm optimisation).

Future work, with a larger sample of subjects, could pre-define a more specific measure such as mean connectivity, mean MRBD/PMBR amplitude etc., which would yield a single measure per subject and repeat, and allow for Bland-Altman analysis over subjects, which would give a better estimate of the expected bias and range.

Some limitations related to the experimental method also need to be mentioned. Firstly, the single-sphere model used for source reconstruction is likely suboptimal; in work by Boto et al. (2016), OPM-MEG was shown to be sensitive to errors in the forward model. For this study, the simplicity of implementation and compatibility with existing analysis pipelines justified the use of the single sphere model.

The operation of the sensors in the SERF regime (see Chapter 2, section 2.3.2) requires heating of the glass cell central to the operation of OPMs. This means that sensors may be subject to overheating after prolonged continuous use—especially if air circulation is impeded by the helmet holding the sensors. This can limit the duration of experiments and hence the number of trials.

Additionally, a heavy helmet can reduce the time subjects can comfortably spend wearing the array. Here, the combined weight of the helmet and sensors was approximately $1.7kg$ (Hill et al., 2020). At the time of writing, the helmet prototype used in the work presented here has been succeeded by a more lightweight iteration with improved ventilation (Hill et al., 2022).

A further limitation is posed by the number of sensors available at the time of scanning and the resulting limited coverage. Given the task employed, sensors were placed with a focus on parietal and occipital regions as well as more posterior parts of the frontal lobe to ensure good sensitivity to the primary

motor areas. While this was kept constant between subjects and repeats, the removal of sensor channels due to faults or noise could not be guaranteed to be the same for each run. While beamforming is based on ensuring unit gain at the source location, changes in coverage will influence the SNR of the reconstructed signal which in turn affects the fingerprinting performance and reproducibility. For future whole-brain analyses, an array with full coverage of the head with triaxial sensors should yield improved results (Boto et al., 2022). It is worth noting that the development of OPM sensors is still progressing and increased stability of the experimental setup in the near future can be expected.

Finally, our methodology could be improved by recording hand and head movements in tandem with the MEG, data. This would allow for more accurate characterisation of how subjects choose to execute the task. Knowledge about head movement, in particular, could elucidate whether part of the inter-subject variability is simply due to subtle differences in the amount of motion through the remnant fields in the MSR or actual behavioural or strategic differences in navigating the maze.

6.4.3 Implications for future mTBI research

The need for longitudinal studies of mTBI is clear. Especially in populations such as athletes with a high risk of head injury, regular baseline MEG scans could be compared with post-injury data to elucidate whether MEG can be used to track recovery. The results presented in this chapter show that OPM-MEG can reliably reconstruct functional measures which appear to remain stable within subjects. The Bland-Altman approach was used to quantify the consistency between scans, which will be crucial for future longitudinal studies. It also offers a numerical means to optimise experimental parameters (such as the number of trials, stimulus choice, number and distribution of OPM sensors), by clearly revealing the amount of bias and variability between repeated recordings.

Our task shows that experiments which mimic pitch-side assessments will be possible—especially with further improvements in the active control of the magnetic environment of OPM-MEG systems (Holmes et al., 2021; Rea et al., 2021)—paving the way for new tasks, such as balance assessments, which would be impossible when using conventional MEG.

6.5 CONCLUSION

We have demonstrated the utility of OPM-MEG for naturalistic paradigms that involve more complex movements than would be practical in cryogenic MEG. Additionally, we have shown that OPM-MEG derived brain activity measures have good intrasubject reliability as indicated via the construction of 'neural fingerprints' when noise is sufficiently controlled. Bland-Altman analysis of repeated measures allowed us to give an estimate of the variation of imaging measures between sessions which should provide a basis for more specific hypothesis generation and paradigm optimisation. Accurate knowledge about scan-to-scan variability will be crucial if longitudinal, within-subject assessments are to be introduced to the clinic or as part of concussion test batteries in sports medicine or armed forces contexts.

Chapter 7 CONCLUDING REMARKS

The work presented in this thesis straddles the *old* and *new* of mTBI research using MEG. In Chapter 4, we analysed existing MEG data from mTBI subjects in the subacute stage of mTBI with known abnormalities in functional connectivity. Using a Hidden Markov Model, we gained new insights into the mechanism behind the observed connectivity deficits, namely a lack of coincidence between transient bursts of activity. Chapter 5 described the collection of new data, drawn from subjects with mTBI in the acute stage as well as non-head injuries, which was contrasted with healthy controls. Here, we validated our HMM-based analysis pipeline—which revealed the same connectivity deficit seen in Chapter 4. We also probed for excess delta power using more tried and tested methods. In Chapter 6 we ventured into the potential future of mTBI (and generally MEG) research, demonstrating the capability of OPMs to enable MEG during naturalistic tasks and showcasing a method to quantify the repeatability of the functional measures derived from OPM-MEG.

7.1 CHOICE OF STUDY COHORTS

The existing MEG literature concerning mTBI does not provide sufficient evidence for a generalizable biomarker. This is due to the wide range of analysis techniques and choices of functional measures reported and, in particular, the choice of subjects recruited for previous studies. Several studies include only chronically symptomatic individuals in their mTBI group, scanning subjects long after the initial injury, while others (as was done in Chapters 4 and 5) do not exclude asymptomatic individuals or those who will not suffer chronic symptoms, by default. Here, an argument could be made that this choice may lead to the inadvertent investigation of two separate subgroups on the spectrum of severity i.e. chronically symptomatic individuals have suffered a more severe

injury that should be classed separately from those who recover soon after experiencing it.

Similarly, the time post-injury at which individuals are recruited varies widely, with very few subjects being scanned in the acute stage. As seen in Chapter 4 and Chapter 5, similar functional connectivity abnormalities observed in subjects scanned several weeks to months after injury can already be seen in the first two weeks. However, deficits in burst amplitude measures seen in Chapter 4 were not replicated in Chapter 5. In sum, this work suggests that if we are to understand the mechanisms behind mTBI and—crucially—the processes driving an unsuccessful recovery, longitudinal studies (or cross-sectional studies at a wider range of time points) are needed.

Furthermore, in the review by Allen et al. (2021), we saw that the choice of control groups in observational studies of mTBI may be suboptimal in the majority of publications (including in our study from Chapter 4), with healthy volunteers being chosen over controls from a population with exposure to non-head trauma. Controlling for the effect of trauma exposure could be paramount in revealing abnormalities that are truly due to mTBI rather than the secondary physiological or psychological effects of injuries in general. Thus we included an orthopaedic trauma control (TC) group in addition to healthy controls in our study from Chapter 5. Despite being limited by the small sample size in the TC group—recruitment of which was significantly disrupted by the COVID-19 pandemic—results from our analyses of burst coincidence connectivity revealed a consistent deficit in mTBI when compared to both TC and healthy control subjects. However, when analysing slow waves, TCs appeared closer to the mTBI patients than HCs. Whether this is because slow waves are a general feature of brain activity following trauma or if some TCs could have experienced a subconcussive injury during their trauma (and thus have an undiagnosed head injury) is unknown. Either way, this finding highlights the critical need for careful choice of controls.

A final point pertaining to the choice of studied subjects is gender bias. The cohort investigated in Chapter 4, for example, was exclusively male, mirroring much of the existing literature. In Chapter 5, subjects were not excluded based on sex, however, the majority of subjects recruited were male. One justification for this is that epidemiological studies of mTBI frequently report an increased risk of mTBI for young males compared to females. Recent evidence, however, suggested that mTBI in females may lead to worse outcomes, highlighting the importance of a more balanced approach for future studies (Bazarian et al., 2010; Levin et al., 2021).

7.2 BURSTS GO OUT OF SYNC IN MTBI

In chapters 4 and 5 we have shown that subjects with mTBI display a deficit in functional connectivity manifesting as a reduction in the temporal overlap of transient bursts of activity in pairs of brain regions. Our HMM-based burst measures revealed consistent abnormalities—despite the groups from Chapters 4 and 5 being recruited at different sites, scanned using different scanners, and representing different phases of injury. This suggests that they could be useful candidate measures for future multicentre studies.

Recent work on bursting activity suggests that transient events in the MEG signal may offer a more precise window into the underlying neural activity compared to ongoing neural oscillations (Jones, 2016). Abnormal bursts could, therefore, not only provide a marker of mTBI but allow for a better understanding of neurophysiology. We may, for example, ask how bursts are synchronised in a healthy brain. This pioneering work produced the first neurobiological models of beta-bursts, including their origins in deep structures and the neurochemical processes that support them. Bringing imaging observations together with these new models will enable us to question e.g. if burst synchronisation is mediated by deeper structures in the brain. The lack of synchrony in mTBI may point to the disruption of certain white matter

connections or a deficit in specific neurotransmitter concentrations. Additionally, multimodal studies of burst coincidence and advanced imaging (e.g. DTI) to characterise the structural integrity of the white matter could be crucial in furthering our understanding.

7.3 SENSITIVITY OF GLOBAL MEASURES AND 'N=1' ASSESSMENTS

The mechanisms of injury (which vary widely in mTBI) and the complex biomechanics of brain trauma (i.e. the forces acting on the brain while the injury is acquired) make it near impossible to predict which regions in the brain are likely to be affected—even when information such as the location and direction of impact is known. It is for this reason that global measures were chosen for our analyses in Chapters 4 and 5. However, the loss of sensitivity due to averaging over brain regions or functional connections may severely limit the utility of such measures.

One of the goals in the search for MEG biomarkers of mTBI should be to determine whether clinical use of MEG is warranted. A clinician should be able to perform an 'N=1' study, scanning a single individual with possible mTBI and be able to conclusively give a diagnosis. The overlapping distributions of our global measures appear to be inadequate for this purpose.

If the goal is to reliably and quickly diagnose, we need to find the most important brain regions to look for functional abnormalities and/or clearer hypotheses about the mechanisms underlying mTBI symptoms. We have shown with the approach in Chapters 4 and 5, that machine-learning-based feature selection can significantly improve upon the global measure of burst connectivity—and lead to high classification accuracies—but is somewhat sensitive to the population studied. Test subjects need to be from a similar enough population to those used to train the models or training data needs to be abundant and as diverse and broad as possible. Whether this will mean more extensive research of mTBI in smaller subgroups (like groups of athletes with

common injury mechanisms, age ranges, blast-injury cohorts etc.) or pooling of existing data for more generalizable models remains to be seen. An objective diagnosis of mTBI using MEG would be useful. However, the predictive power to determine which patients will make a full recovery and which will suffer chronic symptoms may be of even greater clinical utility. The extent to which we can do this is currently unknown. What is known, and shown in Chapter 6, is that MEG is stable over time. Hence, longitudinal tracking of patients through recovery may offer important insights.

7.4 DELTA POWER

Counter to our expectations, the experimental results in Chapter 5 (Section 5.4) did not reveal a clear increase in delta band power in the mTBI cohort. As mentioned in the previous section, however, this may be attributable to the loss of sensitivity when using global measures. The approach of deriving a maximum Z-score for each individual based on a normative database of healthy controls (Huang et al., 2014) was designed to identify the most abnormal brain region for each subject and to yield a single measure with which individual diagnosis could be achieved using a simple threshold. While this approach appeared to identify the majority of mTBI subjects, trauma controls were also largely classified as being part of the mTBI group. This raises several important points: Firstly, the choice of an appropriate control group can—once again—not be overstated. Our group of healthy controls was chosen from a large set of normative data collected for previous studies for example. While the same analysis procedures were used to analyse each subject's data, possible changes in the approach to data acquisition, such as upgrades to the MEG system, and other experimental details may change intractably over time. Caution is, therefore, advised when using historically acquired data.

Secondly, the oft-reported excess low-frequency power may not generalise to our mTBI cohort. Given the scarcity of data on the acute phase, it is plausible

that excess delta waves are a feature of chronically symptomatic mTBI which may emerge at later stages. This further highlights the need for longitudinal studies with several time points post-injury.

A third—albeit speculative—explanation for the observed similarity of our mTBI and TC groups is that a subset of TC subjects were exposed to a sub-concussive blow which went undiagnosed in light of the more severe primary orthopaedic injury. Researchers considering future mTBI studies may want to specify stricter exclusion criteria (see 5.2.1.1) for any TC cohort, focusing on a category of injury which does not involve acceleration/deceleration of the head such as ankle sprains for example.

7.5 OPMs: THE FUTURE OF MEG?

While observational studies such as those presented in Chapters 4 and 5 will progressively add to our understanding of mTBI, the need for longitudinal measures is clear. Especially in the context of concussions in sports, unique opportunities to not only assess recovery but also measure pre and post-injury differences may arise. This may bring the added benefit that athletes could improve estimates of an adequate return-to-play interval by tracing recovery using individual baseline measures. Our work in Chapter 6 showed that OPM-MEG is sensitive to individual differences in the brain's response to naturalistic tasks, and our repeatability assessments demonstrate that OPM-MEG is a viable technique for longitudinal measurement of brain function. The field of OPM-MEG is rapidly evolving; wearable MEG systems will become more accessible in the coming years and improved interference reduction techniques (Holmes et al., 2021) will bring OPM-MEG closer to a pitch-side assessment. Moreover the introduction of new, naturalistic task-paradigms—such as balance tasks—could better target specific groups of symptoms. OPM-MEG represents an extremely promising technique for future study and—potentially—management of mTBI.

7.6 FINAL THOUGHTS

In the search for a measure sensitive enough to detect the subtle brain changes caused by mTBI—and throughout the empirical work in this thesis—we have explored a smorgasbord of techniques which are aimed to aid diagnosis but also to non-invasively learn about the unknown processes that result in the symptoms of mTBI. While consistent findings of deficits in functional connectivity—for example—are encouraging, it is important to note that all of our work is limited by the clinical definition and diagnosis. In my opinion, one cause for the heterogeneity of subjects recruited into MEG studies of mTBI may be that a variety of (currently undefined) clinical phenotypes[11] fall into the category of mTBI and may display uniquely different abnormalities in the MEG signal. It is important that—as evidence from MEG and other imaging modalities mounts—clinical phenotypes based on imaging measures are identified. In particular, more focus should be placed on differences between individuals who recover quickly, compared to those who continue to suffer long-term symptoms. This would be a crucial step towards MEG as a tool for prognosis in mTBI.

In its current form, MEG is not yet suitable for routine clinical assessment of mTBI, however, given the high classification accuracies reported here and in the wider literature, we can conclude that the MEG signal can reveal useful information about mTBI. In particular, burst-based measures (Chapter 4 and Chapter 5) appear to be suitable for diagnostic classification and may offer a means to develop models of mTBI at the cellular scale.

As it is beyond the scope of this thesis to prove the superiority of MEG in classifying mTBI compared to other modalities, further work to test this is needed. However, the combination of functional and structural measures from

[11] These may include groups with common symptoms, mechanisms of injury, previous exposure to sub-concussive blows, etc.

different modalities may be necessary to improve not only classification but stratification of patients with mTBI into clinically relevant phenotypes.

Despite millennia of observations and descriptions of brain injuries, the uncertainty about how mTBI leads to the debilitating long-term symptoms experienced by many remains vast. Nevertheless, I believe we have made an important step towards understanding mTBI. The idea that subtle changes in brain tissue integrity, too small to be observable with structural scanning, generate a functional deficit which can be detected using MEG is compelling. Whether used for diagnosis, triage, management or simply as a research tool to better understand mTBI, it seems the future of MEG in mTBI is important. The ultimate goal will be to improve the lives of those suffering from the complex sequelae of brain injury. The findings presented in this thesis, whilst not conclusive, demonstrate one thing clearly: MEG has a significant role to play if we are to realise this ambition.

APPENDICES

A MRI ACQUISITION PARAMETERS FOR CHAPTER 4

Anatomical T1-weighted images were acquired on a 3T Magnetom Tim Trio (Siemens AG, Erlangen, Germany) located at The Hospital for Sick Children in Toronto, using a 12-channel head coil running an MPRAGE pulse sequence (TR=2,300 ms; TE=2.9ms; Flip angle=9°; Field of view=240 × 256 × 192 mm; slice thickness=1mm).

B RANDOM FOREST FEATURE SELECTION AND CROSS-VALIDATION

Training complex machine learning classifiers using high-dimensional data and small sample sizes can be difficult. In addition, the learnt features of machine learning models are often difficult to interpret. The analysis developed by Zhang et al. (Zhang et al., 2020, 2016), which was briefly described in section 4.2.4.4, includes two main techniques to overcome this: Firstly, we use several nested stages of cross-validation to avoid overfitting despite limited sample sizes, secondly, we use a decision tree-based feature selection algorithm called "random forest" to select important functional connections which were used for further analyses, reducing model complexity while preserving the excellent interpretability of decision trees. This appendix is intended as a primer on decision trees and random forests and contains a more in-depth description of the feature selection pipeline used to produce the relevant results in section 4.3.1.

B.1.1 DECISION TREE MODELS

Decision trees can be used to model multivariate data by splitting the feature space into regions which are assigned a simple model—the region average or the most prevalent category for example. (Hastie et al., 2001). Tree-based

models can be used for classification as well as regression applications, with the aptly named CART (Classification and Regression Tree) being a popular method (Breiman et al., 1984).

CART trees are built recursively following the following steps:

1. Create the root node
2. For each feature/variable, split the dataset between each pair of neighbouring data points and calculate a "split-quality" measure of choice for all possible splits and features
3. Choose the best split for each feature and select the feature with the best split value
4. Create two child nodes using the chosen split feature/value
5. For each child node repeat steps 2-4 until the stopping criterion

At each node, the goal is to find such a split that the resulting nodes are as homogeneous as possible. For classification trees, this means that each child node should mostly contain data points from one class only, for regression trees, this means choosing a split such that the mean of the data points on each side of the split provides an accurate model.

For our binary classification trees, we used the commonly implemented Gini impurity, which is defined as

$$GI = 1 - (p_{G1})^2 - (p_{G2})^2 \,, \qquad (A.1)$$

where p_{Gn} are the probabilities of a data point from class or group n occurring at the current node. GI has a maximum of 0.5 where an equal number of points from both groups are present. An optimal split will minimise the weighted mean of GI across all child nodes (weighted by the number of observations per node). Completely pure nodes will yield $GI = 0$, which can be used as a stopping

criterion. In order to reduce model complexity, other stopping criteria such as the minimum node size—the minimum number of observations needed to create an additional split—or a maximum number of splits can be used. Figure A.1 shows an example of a classification tree. The toy dataset consists of 40 observations of two variables in two classes. A decision tree was trained to partition the feature space spanned by the two variables into sections coloured to represent the output class chosen for points within the decision boundaries.

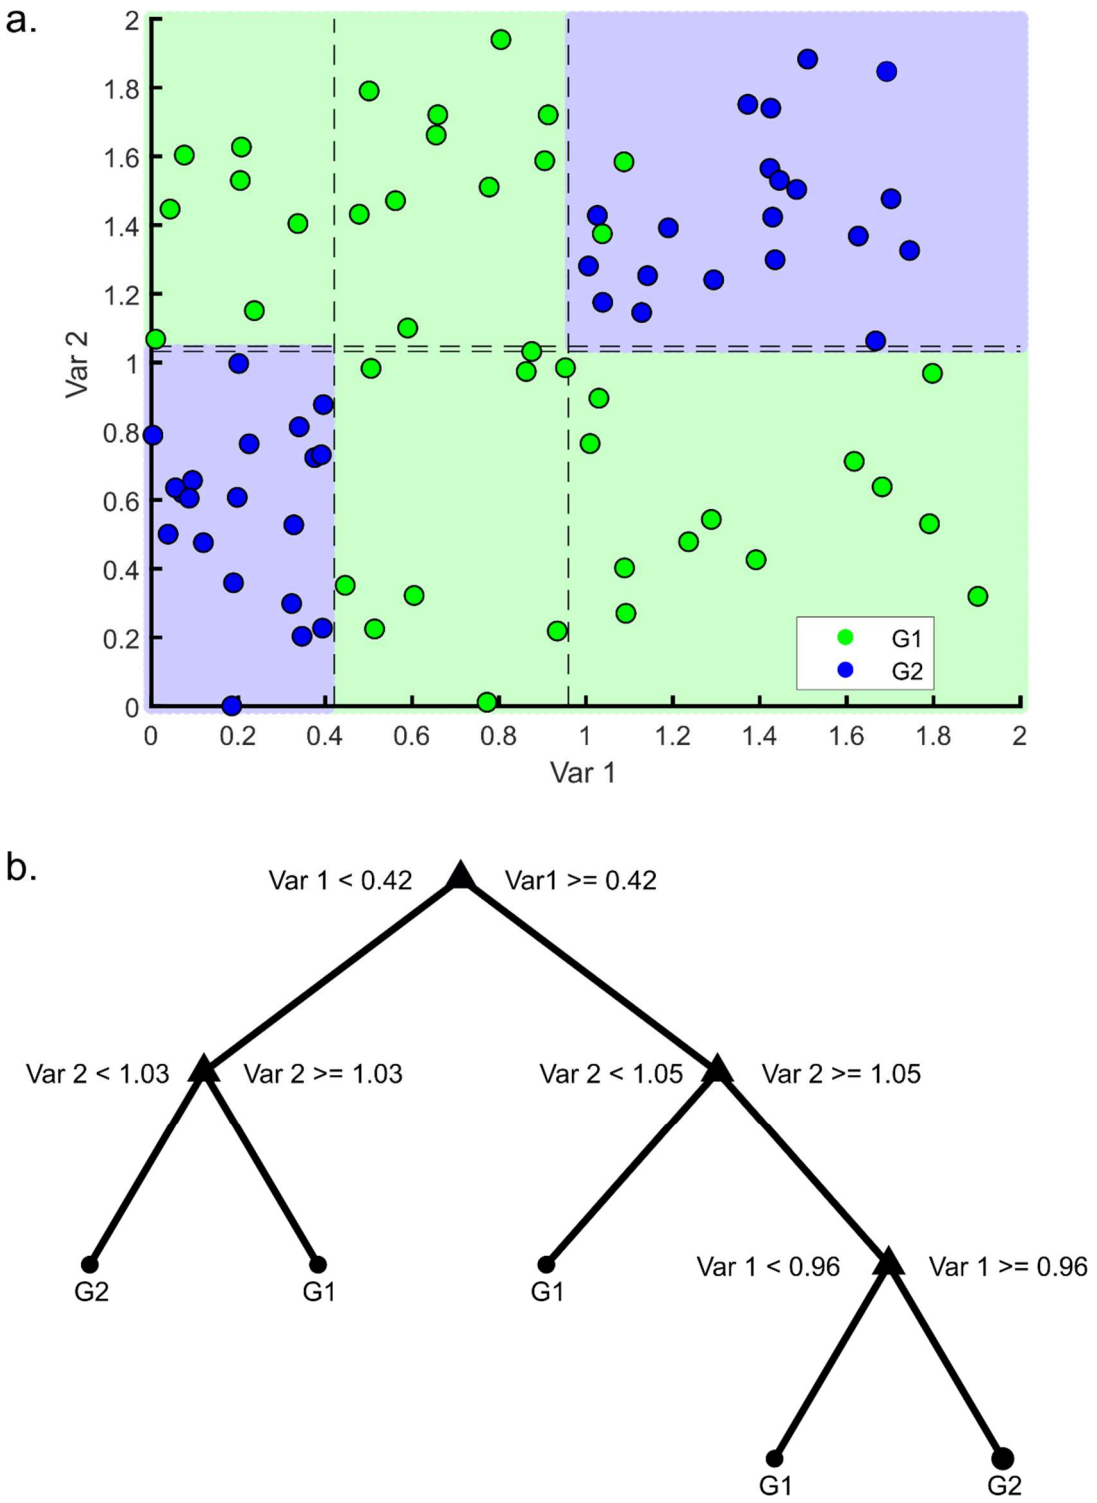

Figure A.1 Classification tree example. a. example data with two classes (G1 and G2) and two features or variables (Var1 and Var2). The points represent observations, and are coloured to reflect the group membership. Dashed lines indicate the splits made by the CART model and the partitions are coloured to the predicted class by the tree shown in b. b. Classification tree diagram with split values at each node. Leaf nodes indicate the chosen class (either G1 or G2).

While classification trees can be used to generate class labels for observations, regression trees partition the feature space and yield a simple model of the regressed quantity in each partition. Usually, the means of the observations on each side of the decision boundary are chosen as the *predictions* for each partition. The split values are optimised as follows: For each feature, and all possible splits, the residuals between each data point and the predictions in each partition are calculated. The split with the lowest mean of the squared residuals is chosen. An example of a regression tree model can be seen in Figure A.2b-c.

Decision trees can model complex data consisting of many variables of different types (e.g. continuous, categorical, etc.) while remaining interpretable to humans, who can simply follow the decision steps. However, the chosen splits are highly sensitive to training data, meaning a small change in the training data can result in very different trees (Bishop, 2006). We remedy this by employing an ensemble learning technique called a random forest.

B.1.2 RANDOM FORESTS

Random forest classifiers consist of a large number—or ensemble—of decision trees, each of which is trained on a randomly chosen subset of the available data (Breiman, 2001). Mathematically,

$$\{h(\mathbf{x}, \mathbf{\Theta}_k), k = 1, \dots, N_{trees}\}, \qquad (A.2)$$

where h is a decision tree trained on a randomly selected subset of the training data $\mathbf{\Theta}_k$ and $\mathbf{x}$ is the input vector. A new data point will be classified by each tree and—in the case of a classification model—majority voting over the ensemble will decide the outcome. Training individual trees on only a subset of the training data removes bias towards the training set resulting in favourable outcomes. The selection of training samples is commonly done using random sampling with replacement, termed Bootstrap Aggregating or *Bagging* (Breiman, 1996). Introducing randomness into the training process prevents

correlation between trees and reduces bias towards outliers and makes random forests robust when only small sample sizes are available. This applies especially when bagging is used in conjunction with random feature selection, where only a random subset of features is made available when choosing a split value for each node. Bagging has the additional advantage of providing an internal means of cross-validation at the tree level. The bootstrap sample—or *bag*—used to construct each tree, leaves the rest of the observations as validation data—the *out-of-bag* (OOB) sample. The classification error for the OOB sample ($errOOB_k$) for the kth tree gives a measure of generalisation error and can also be used to determine feature importance.

B.1.3 RANDOM FOREST FEATURE SELECTION

For the purposes of this thesis, the main application of random forests was not to produce diagnostic classifiers but to robustly select the most important functional connections. Combining bagging and random permutation of features can be used to estimate a measure of variable importance (VI) from a random forest (Breiman, 2001; Genuer et al., 2010). As described above, $errOOB_k$ is an estimate of the generalisation error of each tree. To estimate variable importance, we repeat the calculation of this error after the random permutation of feature m in the OOB sample, $\widetilde{errOOB}_k$, and compute

$$VI_m = \frac{1}{N_{tree}} \sum_{k}^{N_{tree}} \left(\widetilde{errOOB}_k - errOOB_k \right). \qquad (A.3)$$

An increase in error after permuting the values for a given feature indicates that the structure of its data is important for a successful classifier; a larger increase in error means higher variable importance.

Ranking features using variable importance is a crucial part of feature selection, however, this approach introduces the additional problem of including all the most important features while creating the smallest possible feature set. We

follow a two-step procedure introduced by (Genuer et al., 2010) to achieve this in a data-driven manner, without the need for a hypothesis on the number of selected features.

Firstly, variable importance is calculated using equation $A.3$. This is repeated several times (50 times in our pipeline) and means and standard deviations of VI are calculated. Features are sorted in descending order of VI as shown in Figure A.2a. Here an initial selection step is used to reduce the feature space while still retaining an excess number of features for further reduction.; keeping the ranking of the variables the same as in Figure A.2a, a CART regression tree (Figure A.2c) is used to model the standard deviations of the variable importance (σ_{VI}). We select the feature with the minimum CART prediction value (Figure A.2b) and use its VI value as a threshold (Figure A.2a). In the example shown in Figure A.2, 27 features remain.

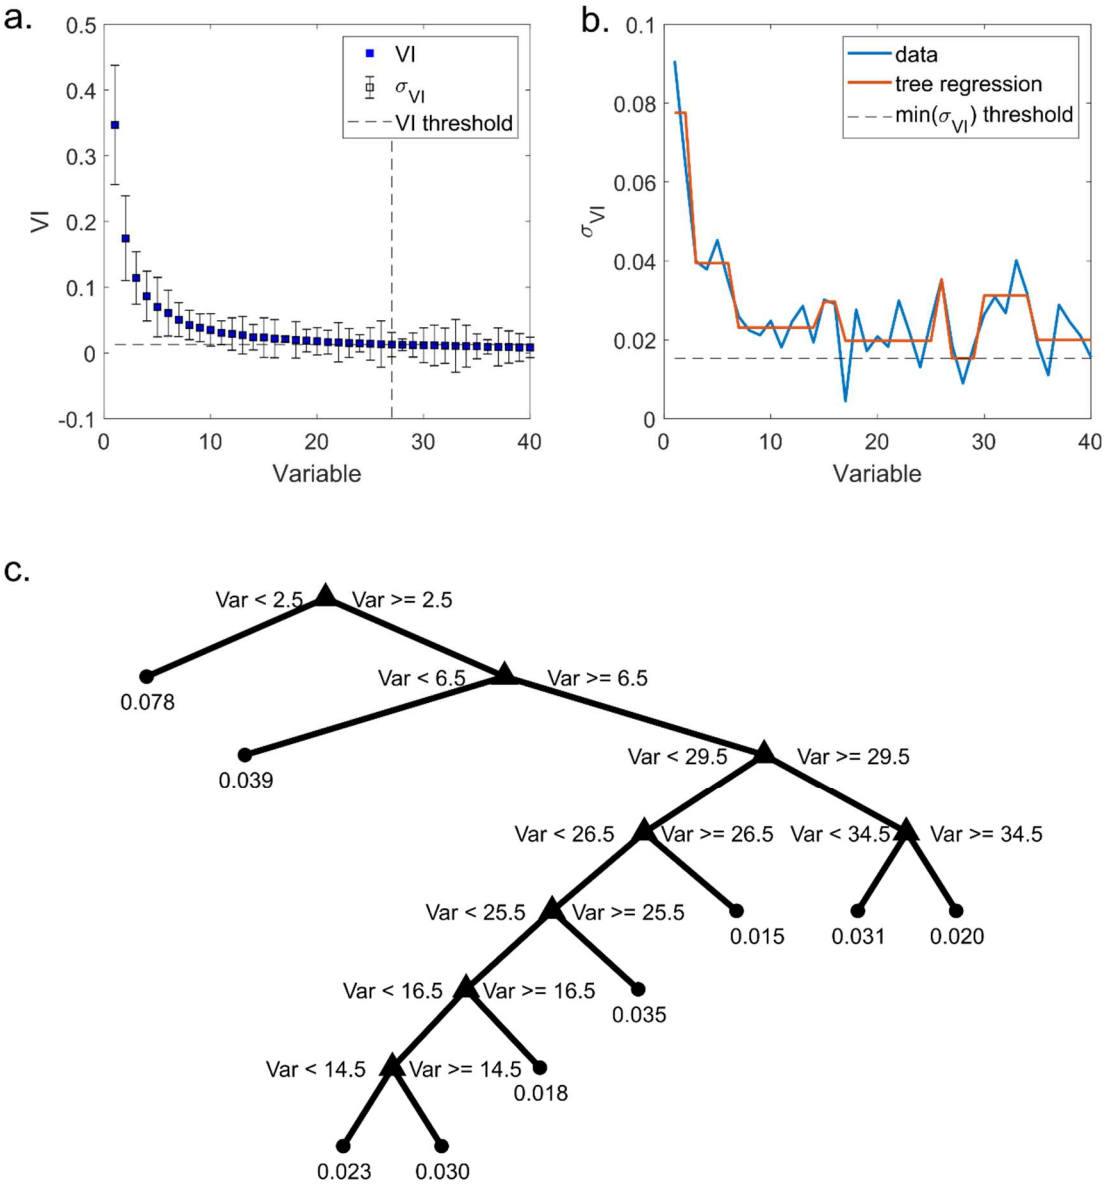

Figure A.2 Initial feature selection step-toy example. a. Mean variable importance (VI) sorted in descending order. The error bars represent the standard deviation of the variable importance (σ_{VI}) over the bootstrapping repeats. The dashed lines indicate the threshold used to produce the reduced feature set (in this case variables 1-27*). b. Plot of the standard deviations of VI, keeping the order of variables the same as in panel a. Using the regression tree shown in panel c, the minimum σ_{VI} and corresponding variable was found and chosen as the threshold (dashed line in panel a). c. regression tree used to model the regression line in panel b.*

After the initial selection step, a series of increasingly complex models are created by sequentially adding the remaining features in descending order of importance. Average OOB errors are calculated for each model after 50

bootstrap repeats. Finally the distribution of average OOB errors (σ_{errOOB}) is used to select the final model as shown in Figure A.3. Out of those models with $errOOB < (\min(errOOB) + \sigma_{errOOB})$, the one with the fewest features is selected.

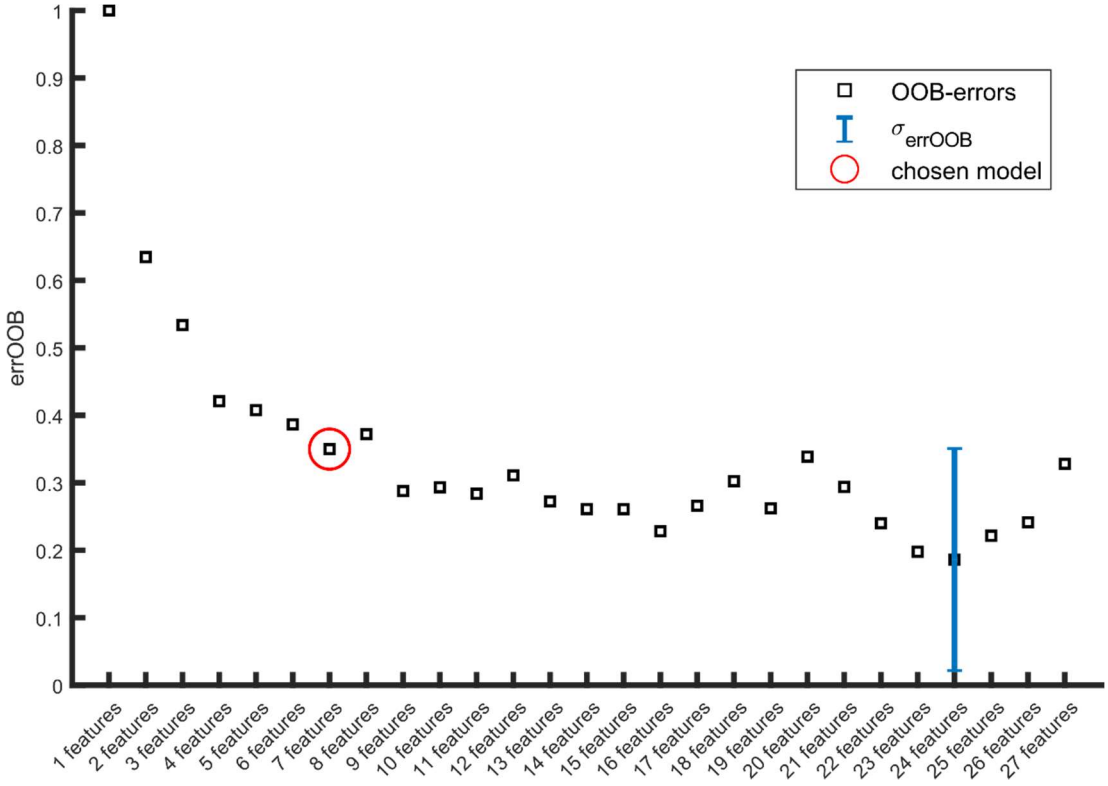

Figure A.3 Model choice after sequential feature selection-toy example. Starting from the feature with highest variable importance, features are sequentially added and out-of-bag errors are calculated. Error rates tend to decrease with increasing model complexity but with diminishing returns for very complex models. We calculate the standard deviation of the error over all models (σ_{errOOB}, blue bar) and consider those models with errors within $\pm\sigma_{errOOB}$ of the model with the minimum error. The model with the fewest features in this range is chosen for the current cross-validation step.

B.1.4 FULL FEATURE SELECTION PIPELINE

The final implementation of our pipeline combines the aforementioned initial and sequential feature selection steps as designed by (Zhang et al., 2016) with nested 10-fold cross-validation (CV) and support vector machine (SVM)

classification as described in (Zhang et al., 2021, 2020). Our approach, however, differed from Zhang et al in two ways. Firstly, in (Zhang et al., 2020) the dataset was split into 80% training and 20% testing data before performing a k-fold CV (on the training data only). Secondly, Zhang et al. (2020) use the final feature list to train an additional SVM—tested on entirely naïve data—to yield a final classification accuracy which was used to compare the utility of connectivity and power spectral density measures. Here, we used the entire dataset for 10-fold CV and calculated the mean and standard deviation of the internal SVM classification accuracy to obtain an estimate of the utility of the chosen connections. Figure A.4 depicts the steps relevant to the work in chapters 4 and 5. Available data are randomly split into 10 sets. For each of the 10 CV steps, 9 sets are used as training data. Using random forests with 501 trees[12], variable importance is estimated using bagging, OOB error estimation and feature permutation. This is repeated 50 times with random bootstrap samples. Feature ranking according to mean VI and CART modelling of the standard deviation of VI are used for initial feature reduction.

At each sequential feature selection (SFS), step random forests of 501 trees are trained as above and OOB errors are estimated 50 times. Using the thresholding technique outlined in Figure A.3, a list of features is produced and used to train an SVM (see Appendix C). After 10-fold CV—yielding ten feature lists—all those features being selected 2 or more times are included in the final list.

12 An odd number of trees is chosen to avoid ties in the majority voting procedure which is used to generate the classification output from the random forest.

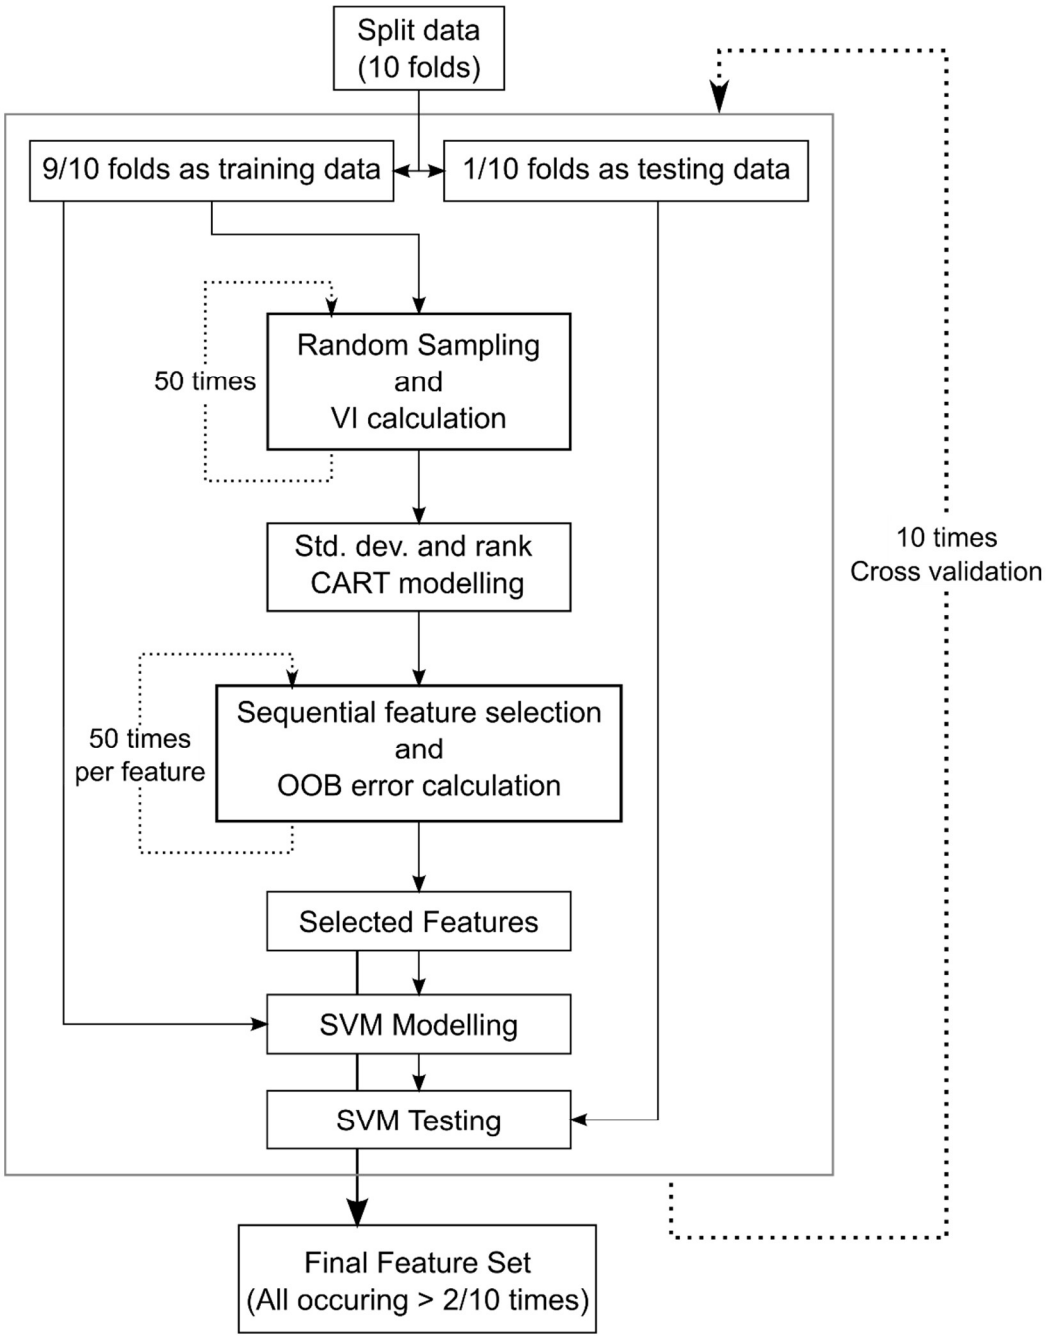

Figure A.4 Recursive random forest feature selection pipeline. The complete dataset is split into ten folds, nine of which are used as training data in each cross-validation (CV) step. After initial estimation of variable importance using 501 trees (repeating 50 times) and modelling of the standard deviation of variable importance against variable rank using a CART model, initial feature reduction is performed. Sequential feature selection yields a list of features which are used to train a Support vector machine (SVM) which includes an internal 10-fold CV to optimise the SVM hyper-parameters. The trained SVM is tested using naïve data and the area under the receiver operating characteristic is recorded as classification accuracy. After performing all 10 repeats of the outer 10-fold CV, the final feature set is chosen by including those features which appear at least twice throughout the CV.

C SUPPORT VECTOR MACHINES FOR CLASSIFICATION

SVMs are a popular solution to classification problems differentiating between two groups of data which attempts to find a decision boundary that minimises generalisation error when classifying unseen data. Being a supervised machine learning technique, SVMs use training data to estimate a decision boundary. While many such boundaries may exist—especially for high-dimensional data—SVMs aim to find an optimal solution by choosing the one which maximises the *margin* (Bishop, 2006). Maximum margin classifiers find a decision boundary such its distance to any of the training data points is maximised. While many applications can make use of hyperplanes which linearly separate the training samples (e.g. Figure A.5, top), this does not always apply. To allow for non-linear decision boundaries, a non-linear kernel can be used (Boser et al., 1992). Conceptually, this involves non-linear transforms of the data which are added to the feature space before finding the maximum margin hyperplane in the now extended feature space. This will then map to a non-linear decision boundary in the native feature space of the training data. In this thesis, all SVM classifications were performed using the radial basis function (RBF) kernel as implemented in the RBioFS package by (Zhang et al., 2016). An example use of a radial basis kernel for SVM classification can be seen in Figure A.5 (bottom). Using the RBF kernel introduces an additional hyperparameter which needs to be specified in the implementation of the SVM classifier. Instead of opting for default values, Zhang et al. used an internal 10-fold CV step when training the SVM to optimise hyperparameters via a grid search before training the final classification model (Zhang et al., 2016). Here, we follow the same procedure.

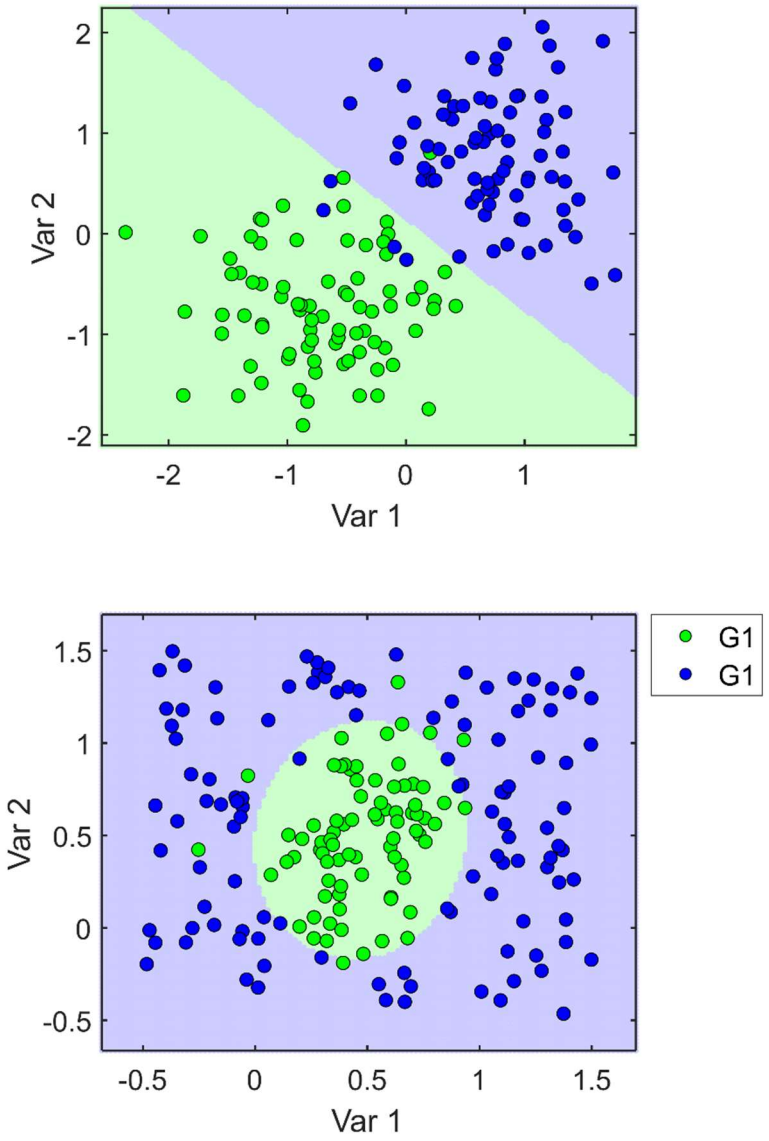

Figure A.5 Examples of SVM classification. Each plot contains scatter plots of the two groups of data used to train separate SVMs. The shaded areas correspond to the class label given to new observations within them. The top axis shows the training data and result of a basic SVM classifier. As the data is defined on two dimensions, the boundary dividing the regions assigned to the two classes is a straight line which maximises classification accuracy and the perpendicular distance to all training data points. In the bottom axis, a linear split of the plane would not result in adequate separation of the two groups. The SVM trained to classify these data used a radial basis function kernel to extend the feature space and generate the non-linear decision boundary.

D SCATTER PLOTS OF BURST CONNECTIVITY MEASURES WITH SYMPTOM SEVERITY

Figure D.1 shows supplementary scatter plots for results presented in Chapter 4. Note that one control subject reported much more severe symptoms than the average. Given that the exclusion criteria for controls did not require presentation without symptoms, we did not exclude this subject from analysis.

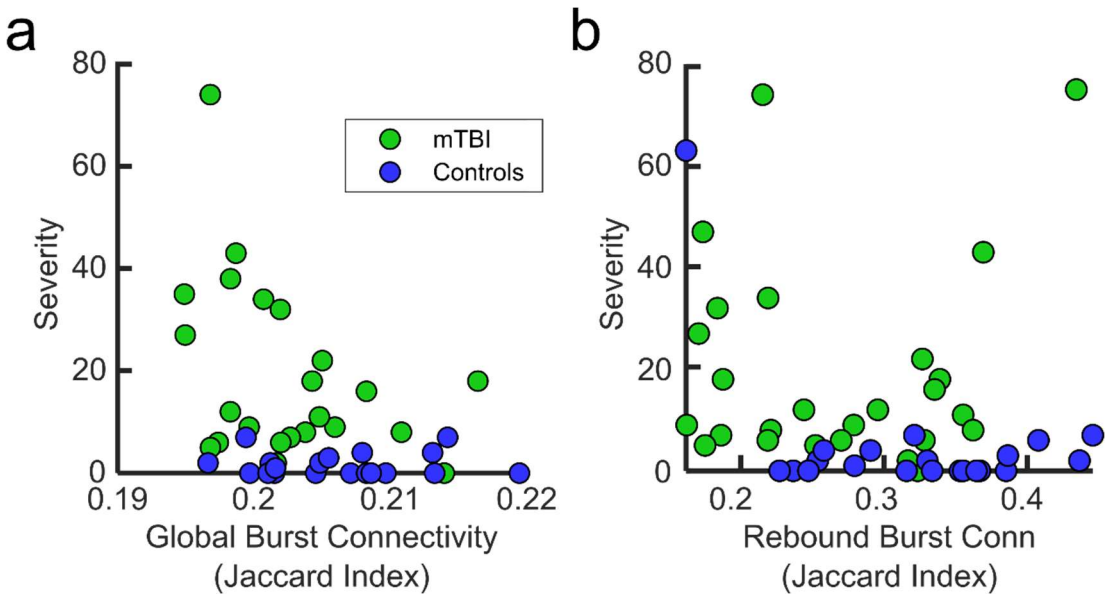

Figure D.1 a) Scatter plot of symptom severity against global burst connectivity b) Scatter plot of symptom severity against global mean connectivity during the post-movement rebound.

E SPECTRAL CONTENT OF THE BURST STATE (CHAPTERS 4 AND 5)

The Hidden Markov Model used in Chapter 4 and 5.3 allows us to characterise bursts of brain activity across a broad frequency range (1-48 Hz in this case), without prior assumption of a band of interest such as the beta band for example. A state-specific multitaper analysis (Vidaurre et al., 2016) was used to generate spectra for each state in each region and subject. Spectra for the "burst state" used for further analyses in Chapters 4 and 5 are shown in Figure E. 1. Note that while the bursts are chosen to coincide with high-amplitude events in the beta band, their spectral profiles are much richer with prominent low-frequency components, particularly in the alpha band. While group comparisons of these spectra are beyond the scope of this thesis, it is worth noting that in the subacute cohort, a reduction of power in the beta range appears dominant, while variability between the acute mTBI, HC and TC cohorts appears across a range of frequencies. This warrants further investigation in the future.

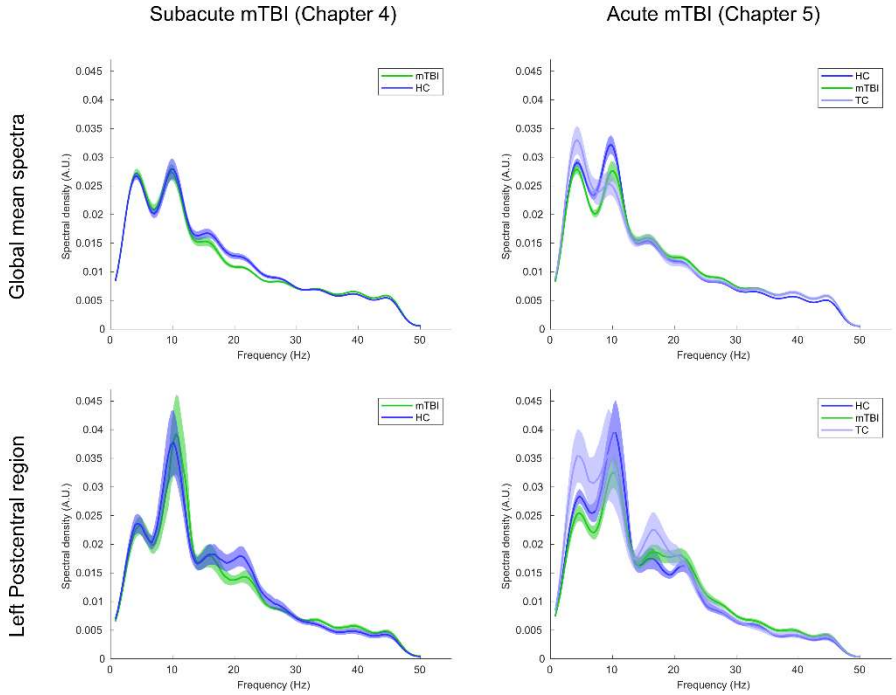

Figure E. 1 Global mean spectra (top row) for the burst state averaged over all 78 AAL regions, and example of single-region spectra for the left postcentral AAL region (bottom row). Lines represent the mean spectra for each group defined in Chapters 4 and 5 and shaded regions the standard error across groups.

F TYPES OF INJURY RECORDED IN THE MEGABIT STUDY

Figure F. 1 shows the mechanisms of injury recorded for our mTBI and TC cohorts. Falls were common in both groups and one of the nine TC subjects sustained an injury in a collision while playing contact sports.

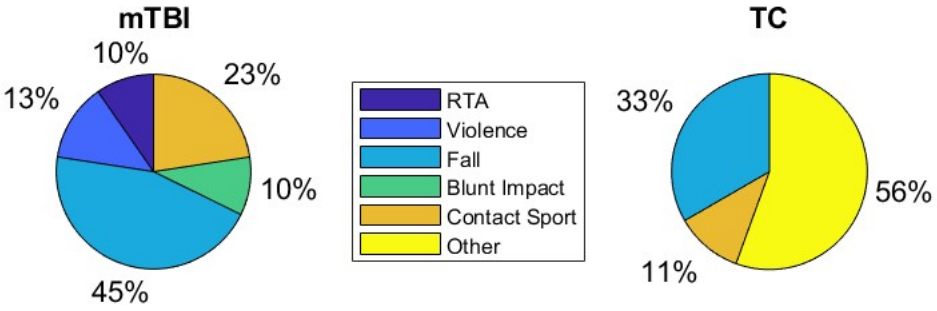

Figure F. 1 Injury Mechanisms for mTBI and trauma control groups recorded in the MEGAbIT study.

Figure F. 2 shows the distribution for LOC duration. No data was available for one subject. The majority of subjects lost consciousness for less than one minute.

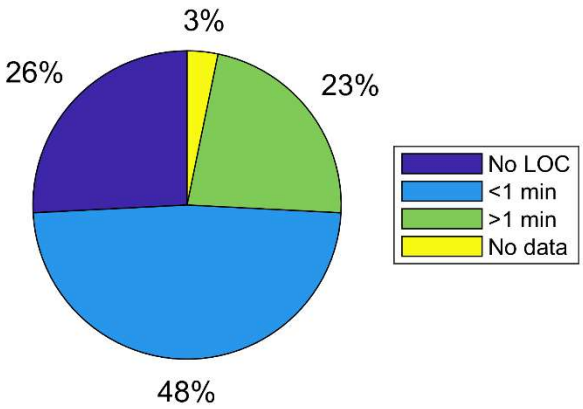

Figure F. 2 Distribution of Loss of Consciousness duration for the mTBI group studied in Chapter 5

REFERENCES

Allen, C.M., Halsey, L., Topcu, G., Rier, L., Gascoyne, L.E., Scadding, J.W., Furlong, P.L., Dunkley, B.T., das Nair, R., Brookes, M.J., Evangelou, N., 2021. Magnetoencephalography abnormalities in adult mild traumatic brain injury: A systematic review. NeuroImage Clin. 31, 102697. https://doi.org/10.1016/j.nicl.2021.102697

Altarev, I., Fierlinger, P., Lins, T., Marino, M.G., Nießen, B., Petzoldt, G., Reisner, M., Stuiber, S., Sturm, M., Taggart Singh, J., Taubenheim, B., Rohrer, H.K., Schläpfer, U., 2015. Minimizing magnetic fields for precision experiments. J. Appl. Phys. 117, 233903. https://doi.org/10.1063/1.4922671

Amico, E., Goñi, J., 2018. The quest for identifiability in human functional connectomes. Sci. Rep. 8, 1–14. https://doi.org/10.1038/s41598-018-25089-1

Amyot, F., Arciniegas, D.B., Brazaitis, M.P., Curley, K.C., Diaz-Arrastia, R., Gandjbakhche, A., Herscovitch, P., Hinds, S.R., Manley, G.T., Pacifico, A., Razumovsky, A., Riley, J., Salzer, W., Shih, R., Smirniotopoulos, J.G., Stocker, D., 2015. A Review of the Effectiveness of Neuroimaging Modalities for the Detection of Traumatic Brain Injury. J. Neurotrauma. https://doi.org/10.1089/neu.2013.3306

Andre, J.B., 2015. Arterial Spin Labeling Magnetic Resonance Perfusion for Traumatic Brain Injury. Top. Magn. Reson. Imaging 24, 275–287. https://doi.org/10.1097/RMR.0000000000000065

Antonakakis, M., Dimitriadis, S.I., Zervakis, M., Micheloyannis, S., Rezaie, R., Babajani-Feremi, A., Zouridakis, G., Papanicolaou, A.C., 2016. Altered cross-frequency coupling in resting-state MEG after mild traumatic brain injury. Int. J. Psychophysiol. https://doi.org/10.1016/j.ijpsycho.2016.02.002

Antonakakis, M., Dimitriadis, S.I., Zervakis, M., Papanicolaou, A.C., Zouridakis, G., 2017. Altered rich-club and frequency-dependent subnetwork organization in mild traumatic brain injury: A MEG resting-state study. Front. Hum. Neurosci. https://doi.org/10.3389/fnhum.2017.00416

Aoki, Y., Inokuchi, R., Gunshin, M., Yahagi, N., Suwa, H., 2012. Diffusion tensor imaging studies of mild traumatic brain injury: A meta-analysis. J. Neurol. Neurosurg. Psychiatry 83, 870–876. https://doi.org/10.1136/jnnp-2012-302742

Arciniegas, D.B., 2011. Clinical electrophysiologic assessments and mild traumatic brain injury: State-of-the-science and implications for clinical practice. Int. J. Psychophysiol. 82, 41–52. https://doi.org/10.1016/J.IJPSYCHO.2011.03.004

Asken, B.M., DeKosky, S.T., Clugston, J.R., Jaffee, M.S., Bauer, R.M., 2018. Diffusion tensor imaging (DTI) findings in adult civilian, military, and sport-related mild traumatic brain injury (mTBI): a systematic critical review. Brain Imaging Behav. 12, 585–612. https://doi.org/10.1007/s11682-

017-9708-9

Babiloni, C., Binetti, G., Cassarino, A., Dal Forno, G., Del Percio, C., Ferreri, F., Ferri, R., Frisoni, G., Galderisi, S., Hirata, K., Lanuzza, B., Miniussi, C., Mucci, A., Nobili, F., Rodriguez, G., Luca Romani, G., Rossini, P.M., 2006. Sources of cortical rhythms in adults during physiological aging: A multicentric EEG study. Hum. Brain Mapp. 27, 162–172. https://doi.org/10.1002/hbm.20175

Bailey, C.M., Echemendia, R.J., Arnett, P.A., 2006. The impact of motivation on neuropsychological performance in sports-related mild traumatic brain injury. J. Int. Neuropsychol. Soc. 12, 475–484. https://doi.org/10.1017/S1355617706060619

Barbati, G., Porcaro, C., Zappasodi, F., Rossini, P.M., Tecchio, F., 2004. Optimization of an independent component analysis approach for artifact identification and removal in magnetoencephalographic signals. Clin. Neurophysiol. 115, 1220–1232. https://doi.org/10.1016/j.clinph.2003.12.015

Bardeen, J., Cooper, L.N., Schrieffer, J.R., 1957. Theory of Superconductivity. Phys. Rev. 108, 1175. https://doi.org/10.1103/PhysRev.108.1175

Basser, P.J., Mattiello, J., LeBihan, D., 1994. MR diffusion tensor spectroscopy and imaging. Biophys. J. 66, 259–267. https://doi.org/10.1016/S0006-3495(94)80775-1

Bastos, A.M., Schoffelen, J.M., 2016. A tutorial review of functional connectivity analysis methods and their interpretational pitfalls. Front. Syst. Neurosci. 9, 175. https://doi.org/10.3389/fnsys.2015.00175

Bastos, A.M., Vezoli, J., Bosman, C.A., Schoffelen, J.M., Oostenveld, R., Dowdall, J.R., DeWeerd, P., Kennedy, H., Fries, P., 2015. Visual areas exert feedforward and feedback influences through distinct frequency channels. Neuron 85, 390–401. https://doi.org/10.1016/j.neuron.2014.12.018

Bauer, M., Stenner, M.-P., Friston, K.J., Dolan, R.J., 2014. Attentional Modulation of Alpha/Beta and Gamma Oscillations Reflect Functionally Distinct Processes. J. Neurosci. https://doi.org/10.1523/JNEUROSCI.3474-13.2014

Bazarian, J.J., Blyth, B., Mookerjee, S., He, H., McDermott, M.P., 2010. Sex Differences in Outcome after Mild Traumatic Brain Injury. J. Neurotrauma 27, 527. https://doi.org/10.1089/NEU.2009.1068

Bazarian, J.J., McClung, J., Cheng, Y.T., Flesher, W., Schneider, S.M., 2005. Emergency department management of mild traumatic brain injury in the USA. Emerg. Med. J. 22, 473–477. https://doi.org/10.1136/EMJ.2004.019273

Belanger, H.G., Curtiss, G., Demery, J.A., Lebowitz, B.K., Vanderploeg, R.D., 2005. Factors moderating neuropsychological outcomes following mild traumatic brain injury: A meta-analysis. J. Int. Neuropsychol. Soc. https://doi.org/10.1017/S1355617705050277

Bell, W.E., Bloom, A.L., 1957. Optical Detection of Magnetic Resonance in Alkali Metal Vapor. Phys. Rev. 107, 1559. https://doi.org/10.1103/PhysRev.107.1559

Benjamini, Y., Hochberg, Y., 1995. Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. J. R. Stat. Soc. Ser. B 57, 289–300. https://doi.org/10.1111/j.2517-6161.1995.tb02031.x

Berger, H., 1929. Über das Elektrenkephalogramm des Menschen. Arch. Psychiatr. Nervenkr. 87, 527–570. https://doi.org/10.1007/BF01797193

Bigler, E.D., Maxwell, W.L., 2012. Neuropathology of mild traumatic brain injury: Relationship to neuroimaging findings. Brain Imaging Behav. 6, 108–136. https://doi.org/10.1007/s11682-011-9145-0

Bishop, C.M., 2006. Pattern recognition and machine learning, Information science and statistics. Springer, New York.

Biswal, B., Zerrin Yetkin, F., Haughton, V.M., Hyde, J.S., 1995. Functional connectivity in the motor cortex of resting human brain using echo-planar mri. Magn. Reson. Med. 34, 537–541. https://doi.org/10.1002/mrm.1910340409

Bloch, F., 1946. Nuclear Induction. Phys. Rev. 70, 460. https://doi.org/10.1103/PhysRev.70.460

Bloom, A.L., 1962. Principles of Operation of the Rubidium VaporMagnetometer. Appl. Opt. Vol. 1, Issue 1, pp. 61-68 1, 61–68. https://doi.org/10.1364/AO.1.000061

Bonaiuto, J.J., Meyer, S.S., Little, S., Rossiter, H., Callaghan, M.F., Dick, F., Barnes, G.R., Bestmann, S., 2018. Lamina-specific cortical dynamics in human visual and sensorimotor cortices. Elife 7. https://doi.org/10.7554/ELIFE.33977

Borna, A., Carter, T.R., Goldberg, J.D., Colombo, A.P., Jau, Y.Y., Berry, C., McKay, J., Stephen, J., Weisend, M., Schwindt, P.D.D., 2017. A 20-channel magnetoencephalography system based on optically pumped magnetometers. Phys. Med. Biol. 62, 8909–8923. https://doi.org/10.1088/1361-6560/aa93d1

Bosboom, J.L.W., Stoffers, D., Stam, C.J., van Dijk, B.W., Verbunt, J., Berendse, H.W., Wolters, E.C., 2006. Resting state oscillatory brain dynamics in Parkinson's disease: An MEG study. Clin. Neurophysiol. 117, 2521–2531. https://doi.org/10.1016/j.clinph.2006.06.720

Boser, B.E., Guyon, I.M., Vapnik, V.N., 1992. Training algorithm for optimal margin classifiers, in: Proceedings of the Fifth Annual ACM Workshop on Computational Learning Theory. Publ by ACM, pp. 144–152. https://doi.org/10.1145/130385.130401

Boto, E., Bowtell, R., Krüger, P., Fromhold, T.M., Morris, P.G., Meyer, S.S., Barnes, G.R., Brookes, M.J., 2016. On the Potential of a New Generation of Magnetometers for MEG: A Beamformer Simulation Study. PLoS One 11,

e0157655. https://doi.org/10.1371/JOURNAL.PONE.0157655

Boto, E., Hill, R.M., Rea, M., Holmes, N., Seedat, Z.A., Leggett, J., Shah, V., Osborne, J., Bowtell, R., Brookes, M.J., 2021. Measuring functional connectivity with wearable MEG. Neuroimage 230, 117815. https://doi.org/10.1016/j.neuroimage.2021.117815

Boto, E., Holmes, N., Leggett, J., Roberts, G., Shah, V., Meyer, S.S., Muñoz, L.D., Mullinger, K.J., Tierney, T.M., Bestmann, S., Barnes, G.R., Bowtell, R., Brookes, M.J., 2018. Moving magnetoencephalography towards real-world applications with a wearable system. Nature. https://doi.org/10.1038/nature26147

Boto, E., Meyer, S.S., Shah, V., Alem, O., Knappe, S., Kruger, P., Fromhold, T.M., Lim, M., Glover, P.M., Morris, P.G., Bowtell, R., Barnes, G.R., Brookes, M.J., 2017. A new generation of magnetoencephalography: Room temperature measurements using optically-pumped magnetometers. Neuroimage 149, 404–414. https://doi.org/10.1016/j.neuroimage.2017.01.034

Boto, E., Shah, V., Hill, R.M., Rhodes, N., Osborne, J., Doyle, C., Holmes, N., Rea, M., Leggett, J., Bowtell, R., Brookes, M.J., 2022. Triaxial detection of the neuromagnetic field using optically-pumped magnetometry: feasibility and application in children. Neuroimage 252, 119027. https://doi.org/10.1016/j.neuroimage.2022.119027

Bowie, C.R., Harvey, P.D., 2006. Administration and interpretation of the Trail Making Test. Nat. Protoc. 1, 2277–2281. https://doi.org/10.1038/nprot.2006.390

Breasted, J.H., 1930. The Edwin Smith Surgical Papyrus: published in facsimile and hieroglyphic transliteration with translation and commentary in two volumes.

Breiman, L., 2001. Random Forests. Mach. Learn. 2001 451 45, 5–32. https://doi.org/10.1023/A:1010933404324

Breiman, L., 1996. Bagging predictors. Mach. Learn. 24, 123–140. https://doi.org/10.1007/bf00058655

Breiman, L., Friedman, J.H., Olshen, R.A., Stone, C.J., 1984. Classification And Regression Trees, 1st ed. Routledge, New York. https://doi.org/10.1201/9781315139470

Brookes, M.J., Groom, M.J., Liuzzi, L., Hill, R.M., Smith, H.J.F., Briley, P.M., Hall, E.L., Hunt, B.A.E., Gascoyne, L.E., Taylor, M.J., Liddle, P.F., Morris, P.G., Woolrich, M.W., Liddle, E.B., 2018. Altered temporal stability in dynamic neural networks underlies connectivity changes in neurodevelopment. Neuroimage 174, 563–575. https://doi.org/10.1016/j.neuroimage.2018.03.008

Brookes, M.J., Hale, J.R., Zumer, J.M., Stevenson, C.M., Francis, S.T., Barnes, G.R., Owen, J.P., Morris, P.G., Nagarajan, S.S., 2011a. Measuring functional

connectivity using MEG: Methodology and comparison with fcMRI. Neuroimage 56, 1082–1104. https://doi.org/10.1016/j.neuroimage.2011.02.054

Brookes, M.J., Vrba, J., Robinson, S.E., Stevenson, C.M., Peters, A.M., Barnes, G.R., Hillebrand, A., Morris, P.G., 2008. Optimising experimental design for MEG beamformer imaging. Neuroimage 39, 1788–1802. https://doi.org/10.1016/J.NEUROIMAGE.2007.09.050

Brookes, M.J., Woolrich, M., Luckhoo, H., Price, D., Hale, J.R., Stephenson, M.C., Barnes, G.R., Smith, S.M., Morris, P.G., 2011b. Investigating the electrophysiological basis of resting state networks using magnetoencephalography. Proc. Natl. Acad. Sci. U. S. A. 108, 16783–16788. https://doi.org/10.1073/pnas.1112685108

Brookes, M.J., Woolrich, M.W., Barnes, G.R., 2012. Measuring functional connectivity in MEG: A multivariate approach insensitive to linear source leakage. Neuroimage 63, 910–920.

Browne, K.D., Chen, X.H., Meaney, D.F., Smith, D.H., 2011. Mild traumatic brain injury and diffuse axonal injury in swine. J. Neurotrauma 28, 1747–1755. https://doi.org/10.1089/neu.2011.1913

Budker, D., Kimball, D.F., Rochester, S.M., Yashchuk, V. V., Zolotorev, M., 2000. Sensitive magnetometry based on nonlinear magneto-optical rotation. Phys. Rev. A 62, 043403. https://doi.org/10.1103/PhysRevA.62.043403

Buffalo, E.A., Fries, P., Landman, R., Buschman, T.J., Desimone, R., 2011. Laminar differences in gamma and alpha coherence in the ventral stream. Proc. Natl. Acad. Sci. U. S. A. 108, 11262–11267. https://doi.org/10.1073/PNAS.1011284108/-/DCSUPPLEMENTAL

Buschman, T.J., Miller, E.K., 2007. Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. Science (80-.). 315, 1860–1864. https://doi.org/10.1126/science.1138071

Byrnes, K.R., Wilson, C.M., Brabazon, F., Von Leden, R., Jurgens, J.S., Oakes, T.R., Selwyn, R.G., 2014. FDG-PET imaging in mild traumatic brain injury: A critical review. Front. Neuroenergetics 6, 13. https://doi.org/10.3389/fnene.2013.00013

Carroll, L.J., Cassidy, J.D., Holm, L., Kraus, J., Coronado, V.G., 2004. Methodological issues and research recommendations for mild traumatic brain injury: The WHO Collaborating Centre Task Force on Mild Traumatic Brain Injury. J. Rehabil. Med. Suppl. 113–125. https://doi.org/10.1080/16501960410023877

Cassidy, J.D., Carroll, L.J., Peloso, P.M., Borg, J., von Holst, H., Holm, L., Kraus, J., Coronado, V.G., 2004. Incidence, risk factors and prevention of mild traumatic brain injury: Results of the WHO Collaborating Centre Task Force on Mild Traumatic Brain Injury. J. Rehabil. Med. Suppl. https://doi.org/10.1080/16501960410023732

Celozzi, S., Araneo, R., Lovat, G., 2008. Electromagnetic shielding., Vol. 192. ed. John Wiley & Sons, Ltd.

Chance, B., Zhuang, Z., UnAh, C., Alter, C., Lipton, L., 1993. Cognition-activated low-frequency modulation of light absorption in human brain. Proc. Natl. Acad. Sci. U. S. A. 90, 3770–3774. https://doi.org/10.1073/pnas.90.8.3770

Chen, S.H.A., Kareken, D.A., Fastenau, P.S., Trexler, L.E., Hutchins, G.D., 2003. A study of persistent post-concussion symptoms in mild head trauma using positron emission tomography. J. Neurol. Neurosurg. Psychiatry 74, 326–332. https://doi.org/10.1136/jnnp.74.3.326

Cheyne, D.O., 2013. MEG studies of sensorimotor rhythms: A review. Exp. Neurol. 245, 27–39. https://doi.org/10.1016/j.expneurol.2012.08.030

Choe, M.C., 2016. The Pathophysiology of Concussion. Curr. Pain Headache Reports 2016 206 20, 1–10. https://doi.org/10.1007/S11916-016-0573-9

Churchill, N.W., Hutchison, M.G., Graham, S.J., Schweizer, T.A., 2020. Scale-free functional brain dynamics during recovery from sport-related concussion. Hum. Brain Mapp. 41, 2567–2582. https://doi.org/10.1002/hbm.24962

Churchill, N.W., Hutchison, M.G., Graham, S.J., Schweizer, T.A., 2018. Connectomic markers of symptom severity in sport-related concussion: Whole-brain analysis of resting-state fMRI. NeuroImage Clin. 18, 518–526. https://doi.org/10.1016/j.nicl.2018.02.011

Churchill, N.W., Hutchison, M.G., Graham, S.J., Schweizer, T.A., 2017. Symptom correlates of cerebral blood flow following acute concussion. NeuroImage Clin. 16, 234–239. https://doi.org/10.1016/j.nicl.2017.07.019

Cignoni, P., Callieri, M., Corsini, M., Dellepiane, M., Ganovelli, F., Ranzuglia, G., 2008. MeshLab: An open-source mesh processing tool, in: 6th Eurographics Italian Chapter Conference 2008 - Proceedings. pp. 129–136.

Cohen-Tannoudji, C., Dupont-Roc, J., Haroche, S., Laloë, F., 1970. Diverses résonances de croisement de niveaux sur des atomes pompés optiquement en champ nul. I. Théorie. Rev. Phys. Appliquée 5, 95–101. https://doi.org/10.1051/RPHYSAP:0197000501095OO

Cohen, D., 1972. Magnetoencephalography: Detection of the brain's electrical activity with a superconducting magnetometer. Science (80-.). 175, 664–666. https://doi.org/10.1126/SCIENCE.175.4022.664

Cohen, D., 1968. Magnetoencephalography: Evidence of magnetic fields produced by alpha-rhythm currents. Science (80-.). 161, 784–786. https://doi.org/10.1126/SCIENCE.161.3843.784

Colclough, G.L., Brookes, M.J., Smith, S.M., Woolrich, M.W., 2015. A symmetric multivariate leakage correction for MEG connectomes. Neuroimage 117, 439–448. https://doi.org/10.1016/j.neuroimage.2015.03.071

Conley, A.C., Cooper, P.S., Karayanidis, F., Gardner, A.J., Levi, C.R., Stanwell, P.,

Gaetz, M.B., Iverson, G.L., 2018. Resting State Electroencephalography and Sports-Related Concussion: A Systematic Review. https://home.liebertpub.com/neu 36, 1–13. https://doi.org/10.1089/NEU.2018.5761

Costanza, A., Weber, K., Gandy, S., Bouras, C., Hof, P.R., Giannakopoulos, P., Canuto, A., 2011. Review: Contact sport-related chronic traumatic encephalopathy in the elderly: clinical expression and structural substrates. Neuropathol. Appl. Neurobiol. 37, 570–584. https://doi.org/10.1111/j.1365-2990.2011.01186.x

Davenport, E.M., Urban, J.E., Mokhtari, F., Lowther, E.L., Van Horn, J.D., Vaughan, C.G., Gioia, G.A., Whitlow, C.T., Stitzel, J.D., Maldjian, J.A., 2016. Subconcussive impacts and imaging findings over a season of contact sports. Concussion 1, CNC19. https://doi.org/10.2217/cnc-2016-0003

Delorme, A., Makeig, S., 2004. EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. J. Neurosci. Methods 134, 9–21. https://doi.org/10.1016/J.JNEUMETH.2003.10.009

Dupont-Roc, J., Haroche, S., Cohen-Tannoudji, C., 1969. Detection of very weak magnetic fields (10–9gauss) by 87Rb zero-field level crossing resonances. Phys. Lett. A 28, 638–639. https://doi.org/10.1016/0375-9601(69)90480-0

Echemendia, R.J., Julian, L.J., 2001. Mild Traumatic Brain Injury in Sports: Neuropsychology's Contribution to a Developing Field. Neuropsychol. Rev. 2001 112 11, 69–88. https://doi.org/10.1023/A:1016651217141

Echemendia, R.J., Meeuwisse, W., McCrory, P., Davis, G.A., Putukian, M., Leddy, J., Makdissi, M., Sullivan, S.J., Broglio, S.P., Raftery, M., Schneider, K., Kissick, J., McCrea, M., Dvořák, J., Sills, A.K., Aubry, M., Engebretsen, L., Loosemore, M., Fuller, G., Kutcher, J., Ellenbogen, R., Guskiewicz, K., Patricios, J., Herring, S., 2017. The Sport Concussion Assessment Tool 5th Edition (SCAT5): Background and rationale. Br. J. Sports Med. 51, 848–850. https://doi.org/10.1136/bjsports-2017-097506

Ellis, M.J., Ryner, L.N., Sobczyk, O., Fierstra, J., Mikulis, D.J., Fisher, J.A., Duffin, J., Mutch, W.A.C., 2016. Neuroimaging Assessment of Cerebrovascular Reactivity in Concussion: Current Concepts, Methodological Considerations, and Review of the Literature. Front. Neurol. 7, 61. https://doi.org/10.3389/fneur.2016.00061

Evans, A.C., Janke, A.L., Collins, D.L., Baillet, S., 2012. Brain templates and atlases. Neuroimage. https://doi.org/10.1016/j.neuroimage.2012.01.024

Finn, E.S., Shen, X., Scheinost, D., Rosenberg, M.D., Huang, J., Chun, M.M., Papademetris, X., Constable, R.T., 2015. Functional connectome fingerprinting: Identifying individuals using patterns of brain connectivity. Nat. Neurosci. 18, 1664–1671. https://doi.org/10.1038/nn.4135

Forgacs, R.L., Warnick, A., 1967. Digital-analog magnetometer utilizing

superconducting sensor. Rev. Sci. Instrum. 38, 214–220. https://doi.org/10.1063/1.1771358

Fries, P., 2005. A mechanism for cognitive dynamics: Neuronal communication through neuronal coherence. Trends Cogn. Sci. 9, 474–480. https://doi.org/10.1016/j.tics.2005.08.011

Fry, A., Mullinger, K.J., O'Neill, G.C., Barratt, E.L., Morris, P.G., Bauer, M., Folland, J.P., Brookes, M.J., 2016. Modulation of post-movement beta rebound by contraction force and rate of force development. Hum. Brain Mapp. 37, 2493–2511. https://doi.org/10.1002/HBM.23189

Gaetz, M., Bernstein, D.M., 2001. The current status of electrophysiologic procedures for the assessment of mild traumatic brain injury. J. Head Trauma Rehabil. 16, 386–405. https://doi.org/10.1097/00001199-200108000-00008

Gall, F.J., 1835. On the functions of the brain and of each of its parts: With observations on the possibility of determining the instincts, propensities, and talents, or the moral and intellectual dispositions of men and animals, by the configuration of the brain and head. Marsh, Capen & Lyon.

Gardner, R.C., Yaffe, K., 2015. Epidemiology of mild traumatic brain injury and neurodegenerative disease. Mol. Cell. Neurosci. https://doi.org/10.1016/j.mcn.2015.03.001

Gascoyne, L.E., Brookes, M.J., Rathnaiah, M., Katshu, M.Z.U.H., Koelewijn, L., Williams, G., Kumar, J., Walters, J.T.R., Seedat, Z.A., Palaniyappan, L., Deakin, J.F.W., Singh, K.D., Liddle, P.F., Morris, P.G., 2021. Motor-related oscillatory activity in schizophrenia according to phase of illness and clinical symptom severity. NeuroImage Clin. 29, 102524. https://doi.org/10.1016/j.nicl.2020.102524

Gazdzinski, L.M., Mellerup, M., Wang, T., Adel, S.A.A., Lerch, J.P., Sled, J.G., Nieman, B.J., Wheeler, A.L., 2020. White Matter Changes Caused by Mild Traumatic Brain Injury in Mice Evaluated Using Neurite Orientation Dispersion and Density Imaging. J. Neurotrauma 37, 1818–1828. https://doi.org/10.1089/neu.2020.6992

Genuer, R., Poggi, J.M., Tuleau-Malot, C., 2010. Variable selection using random forests. Pattern Recognit. Lett. 31, 2225–2236. https://doi.org/10.1016/J.PATREC.2010.03.014

Gerlach, W., Stern, O., 1922. Der experimentelle Nachweis des magnetischen Moments des Silberatoms. Zeitschrift für Phys. 8, 110–111. https://doi.org/10.1007/BF01329580

Geselowitz, D.B., 1970. On the Magnetic Field Generated Outside an Inhomogeneous Volume Conductor by Internal Current Sources. IEEE Trans. Magn. 6, 346–347. https://doi.org/10.1109/TMAG.1970.1066765

Giza, C.C., Hovda, D.A., 2014. The New Neurometabolic Cascade of Concussion.

Neurosurgery 75, S24–S33. https://doi.org/10.1227/NEU.0000000000000505

Gloor, P., Ball, G., Schaul, N., 1977. Brain lesions that produce delta waves in the EEG. Neurology 27, 326–326. https://doi.org/10.1212/wnl.27.4.326

Gohel, B., Lim, S., Kim, M.-Y., Kwon, H., Kim, K., 2017. Approximate Subject Specific Pseudo MRI from an Available MRI Dataset for MEG Source Imaging. Front. Neuroinform. 0, 50. https://doi.org/10.3389/FNINF.2017.00050

Gong, G., Rosa-Neto, P., Carbonell, F., Chen, Z.J., He, Y., Evans, A.C., 2009. Age- and gender-related differences in the cortical anatomical network. J. Neurosci. 29, 15684–15693. https://doi.org/10.1523/JNEUROSCI.2308-09.2009

Haacke, E.M., Xu, Y., Cheng, Y.C.N., Reichenbach, J.R., 2004. Susceptibility weighted imaging (SWI). Magn. Reson. Med. 52, 612–618. https://doi.org/10.1002/mrm.20198

Haft, M., Hofmann, R., Tresp, V., 1999. Model-independent mean-field theory as a local method for approximate propagation of information. Netw. Comput. Neural Syst. 10, 93–105. https://doi.org/10.1088/0954-898X_10_1_006

Haghayegh, S., Kang, H.A., Khoshnevis, S., Smolensky, M.H., Smolensky, M.H., Diller, K.R., 2020. A comprehensive guideline for Bland-Altman and intra class correlation calculations to properly compare two methods of measurement and interpret findings. Physiol. Meas. 41, 055012. https://doi.org/10.1088/1361-6579/ab86d6

Hall, E.L., Robson, S.E., Morris, P.G., Brookes, M.J., 2014. The relationship between MEG and fMRI. Neuroimage. https://doi.org/10.1016/j.neuroimage.2013.11.005

Hall, E.L., Woolrich, M.W., Thomaz, C.E., Morris, P.G., Brookes, M.J., 2013. Using variance information in magnetoencephalography measures of functional connectivity. Neuroimage 67, 203–212. https://doi.org/10.1016/J.NEUROIMAGE.2012.11.011

Hamalainen, M., Hari, R., Ilmoniemi, R.J., Knuutila, J., Lounasmaa, O. V, 1993. Magnetoencephalography theory, instrumentation, and applications to noninvasive studies of the working human brain.

Haneef, Z., Levin, H.S., Frost, J.D., Mizrahi, E.M., Mizrahi, E.M., 2013. Electroencephalography and quantitative electroencephalography in mild traumatic brain injury. J. Neurotrauma 30, 653–6. https://doi.org/10.1089/neu.2012.2585

Hanle, W., 1925. Die magnetische Beeinflussung der Resonanzfluoreszenz, in: Ergebnisse Der Exakten Naturwissenschaften. Springer, Berlin, Heidelberg, pp. 214–232. https://doi.org/10.1007/978-3-642-94259-4_7

Happer, W., 1972. Optical pumping. Rev. Mod. Phys. 44, 169–249. https://doi.org/10.1103/RevModPhys.44.169

Happer, W., Tam, A.C., 1977. Effect of rapid spin exchange on the magnetic-resonance spectrum of alkali vapors. Phys. Rev. A 16, 1877. https://doi.org/10.1103/PhysRevA.16.1877

Happer, W., Tang, H., 1973. Spin-Exchange Shift and Narrowing of Magnetic Resonance Lines in Optically Pumped Alkali Vapors. Phys. Rev. Lett. 31, 273. https://doi.org/10.1103/PhysRevLett.31.273

Hari, R., Puce, A., 2017. MEG-EEG Primer. Oxford University Press.

Hastie, T., Friedman, J., Tibshirani, R., 2001. The Elements of Statistical Learning, Springer Series in Statistics. Springer New York, New York, NY. https://doi.org/10.1007/978-0-387-21606-5

Headway, 2016. Statistics | Headway [WWW Document]. URL https://www.headway.org.uk/about-brain-injury/further-information/statistics/ (accessed 5.23.22).

Heinrichs-Graham, E., Wilson, T.W., Santamaria, P.M., Heithoff, S.K., Torres-Russotto, D., Hutter-Saunders, J.A.L., Estes, K.A., Meza, J.L., Mosley, R.L., Gendelman, H.E., 2014. Neuromagnetic evidence of abnormal movement-related beta desynchronization in Parkinson's disease. Cereb. Cortex 24, 2669–2678. https://doi.org/10.1093/cercor/bht121

Hill, R.M., Boto, E., Holmes, N., Hartley, C., Seedat, Z.A., Leggett, J., Roberts, G., Shah, V., Tierney, T.M., Woolrich, M.W., Stagg, C.J., Barnes, G.R., Bowtell, R., Slater, R., Brookes, M.J., 2019. A tool for functional brain imaging with lifespan compliance. Nat. Commun. 2019 101 10, 1–11. https://doi.org/10.1038/s41467-019-12486-x

Hill, R.M., Boto, E., Rea, M., Holmes, N., Leggett, J., Coles, L.A., Papastavrou, M., Everton, S.K., Hunt, B.A.E., Sims, D., Osborne, J., Shah, V., Bowtell, R., Brookes, M.J., 2020. Multi-channel whole-head OPM-MEG: Helmet design and a comparison with a conventional system. Neuroimage 219, 116995.

Hill, R.M., Devasagayam, J., Holmes, N., Boto, E., Shah, V., Osborne, J., Safar, K., Worcester, F., Mariani, C., Dawson, E., Woolger, D., Bowtell, R., Taylor, M.J., Brookes, M.J., 2022. Using OPM-MEG in contrasting magnetic environments. Neuroimage 253, 119084. https://doi.org/10.1016/j.neuroimage.2022.119084

Hipp, J.F., Hawellek, D.J., Corbetta, M., Siegel, M., Engel, A.K., 2012. Large-scale cortical correlation structure of spontaneous oscillatory activity. Nat. Neurosci. 15, 884–890. https://doi.org/10.1038/nn.3101

Hocke, L.M., Duszynski, C.C., Debert, C.T., Dleikan, D., Dunn, J.F., 2018. Reduced functional connectivity in adults with persistent post-concussion symptoms: A functional near-infrared spectroscopy study. J. Neurotrauma 35, 1224–1232.

https://doi.org/10.1089/NEU.2017.5365/ASSET/IMAGES/LARGE/FIGURE4.JPEG

Hodgkin, A.L., Huxley, A.F., 1939. Action potentials recorded from inside a nerve fibre [8]. Nature. https://doi.org/10.1038/144710a0

Holmes, N., Leggett, J., Boto, E., Roberts, G., Hill, R.M., Tierney, T.M., Shah, V., Barnes, G.R., Brookes, M.J., Bowtell, R., 2018. A bi-planar coil system for nulling background magnetic fields in scalp mounted magnetoencephalography. Neuroimage 181, 760–774. https://doi.org/10.1016/j.neuroimage.2018.07.028

Holmes, N., Rea, M., Hill, R.M., Boto, E., Stuart, A., Leggett, J., Edwards, L.J., Rhodes, N., Shah, V., Osborne, J., Fromhold, T.M., Glover, P., Montague, P.R., Brookes, M.J., Bowtell, R., 2021. Naturalistic hyperscanning with wearable magnetoencephalography. bioRxiv 2021.09.07.459124. https://doi.org/10.1101/2021.09.07.459124

Holmes, N., Tierney, T.M., Leggett, J., Boto, E., Mellor, S., Roberts, G., Hill, R.M., Shah, V., Barnes, G.R., Brookes, M.J., Bowtell, R., 2019. Balanced, bi-planar magnetic field and field gradient coils for field compensation in wearable magnetoencephalography. Sci. Rep. 9. https://doi.org/10.1038/s41598-019-50697-w

Hoogenboom, N., Schoffelen, J.-M., Oostenveld, R., Parkes, L.M., Fries, P., 2006. Localizing human visual gamma-band activity in frequency, time and space. Neuroimage 29, 764–773. https://doi.org/10.1016/j.neuroimage.2005.08.043

Hoshi, Y., Tamura, M., 1993. Detection of dynamic changes in cerebral oxygenation coupled to neuronal function during mental work in man. Neurosci. Lett. 150, 5–8. https://doi.org/10.1016/0304-3940(93)90094-2

Hounsfield, G.N., 1973. Computerized transverse axial scanning (tomography): I. Description of system. Br. J. Radiol. 46, 1016–1022. https://doi.org/10.1259/0007-1285-46-552-1016

Huang, M.-X., Huang, C.W., Harrington, D.L., Nichols, S., Robb-Swan, A., Angeles-Quinto, A., Le, L., Rimmele, C., Drake, A., Song, T., Huang, J.W., Clifford, R., Ji, Z., Cheng, C.-K., Lerman, I., Yurgil, K.A., Lee, R.R., Baker, D.G., 2020. Marked Increases in Resting-State MEG Gamma-Band Activity in Combat-Related Mild Traumatic Brain Injury. Cereb. Cortex 30, 283–295. https://doi.org/10.1093/cercor/bhz087

Huang, M.-X., Nichols, S., Robb-Swan, A., Angeles-Quinto, A., Harrington, D.L., Drake, A., Huang, C.W., Song, T., Diwakar, M., Risbrough, V.B., Matthews, S., Clifford, R., Cheng, C.-K., Huang, J.W., Sinha, A., Yurgil, K.A., Ji, Z., Lerman, I., Lee, R.R., Baker, D.G., 2019. MEG Working Memory N-Back Task Reveals Functional Deficits in Combat-Related Mild Traumatic Brain Injury. Cereb. Cortex 29, 1953–1968. https://doi.org/10.1093/cercor/bhy075

Huang, M.X., Mosher, J.C., Leahy, R.M., 1999. A sensor-weighted overlapping-

sphere head model and exhaustive head model comparison for MEG. Phys. Med. Biol. 44, 423–440. https://doi.org/10.1088/0031-9155/44/2/010

Huang, M.X., Nichols, S., Baker, D.G., Robb, A., Angeles, A., Yurgil, K.A., Drake, A., Levy, M., Song, T., McLay, R., Theilmann, R.J., Diwakar, M., Risbrough, V.B., Ji, Z., Huang, C.W., Chang, D.G., Harrington, D.L., Muzzatti, L., Canive, J.M., Christopher Edgar, J., Chen, Y.H., Lee, R.R., 2014. Single-subject-based whole-brain MEG slow-wave imaging approach for detecting abnormality in patients with mild traumatic brain injury. NeuroImage Clin. 5, 109–119. https://doi.org/10.1016/j.nicl.2014.06.004

Hume, P.A., Theadom, A., Lewis, G.N., Quarrie, K.L., Brown, S.R., Hill, R., Marshall, S.W., 2017. A Comparison of Cognitive Function in Former Rugby Union Players Compared with Former Non-Contact-Sport Players and the Impact of Concussion History. Sport. Med. 47, 1209–1220. https://doi.org/10.1007/s40279-016-0608-8

Hunt, B.A.E., Tewarie, P.K., Mougin, O.E., Geades, N., Jones, D.K., Singh, K.D., Morris, P.G., Gowland, P.A., Brookes, M.J., 2016. Relationships between cortical myeloarchitecture and electrophysiological networks. Proc. Natl. Acad. Sci. 113, 13510–13515. https://doi.org/10.1073/pnas.1608587113

Hunt, B.A.E., Wong, S.M., Vandewouw, M.M., Brookes, M.J., Dunkley, B.T., Taylor, M.J., 2019. Spatial and spectral trajectories in typical neurodevelopment from childhood to middle age. Netw. Neurosci. 3, 497–520. https://doi.org/10.1162/netn_a_00077

Iivanainen, J., Stenroos, M., Parkkonen, L., 2017. Measuring MEG closer to the brain: Performance of on-scalp sensor arrays. Neuroimage 147, 542–553. https://doi.org/10.1016/J.NEUROIMAGE.2016.12.048

Iivanainen, J., Zetter, R., Grön, M., Hakkarainen, K., Parkkonen, L., 2019. On-scalp MEG system utilizing an actively shielded array of optically-pumped magnetometers. Neuroimage 194, 244–258. https://doi.org/10.1016/J.NEUROIMAGE.2019.03.022

James, S.L., Bannick, M.S., Montjoy-Venning, W.C., Lucchesi, L.R., Dandona, L., Dandona, R., Hawley, C., Hay, S.I., Jakovljevic, M., Khalil, I., Krohn, K.J., Mokdad, A.H., Naghavi, M., Nichols, E., Reiner, R.C., Smith, M., Feigin, V.L., Vos, T., Murray, C.J.L., Sunshine, J.E., Yost, M.G., Ellenbogen, R.G., Kalani, R., Morrison, S.D., Crowe, C.S., Massenburg, B.B., Theadom, A., Te Ao, B.J., Balalla, S., Jones, K.M., Ofori-Asenso, R., Li, S., Sobhani, S., Hosseini, S.M., Mansournia, M.A., Yaseri, M., Anjomshoa, M., Mousavi, S.M., Haj-Mirzaian, Arv, Haj-Mirzaian, Ary, Malekzadeh, R., Poustchi, H., Roshandel, G., Sepanlou, S.G., Afarideh, M., Esteghamati, A., Esteghamati, S., Ganji, M., Kasaeian, A., Rahimi-Movaghar, A., Eskandarieh, S., Sahraian, M.A., Shams-Beyranvand, M., Abbasi, N., Farzadfar, F., Irvani, S.N., Rahimi-Movaghar, V., Salamati, P., Sharif-Alhoseini, M., Fereshtehnejad, S.M., Mate, K.K.V., Abdulkader, R., Abraha, H.N., Kassa, T.D., Nirayo, Y.L., Weldegwergs, K.G., Gezae, K.E., Zenebe, Z.M., Degefa, M.G., Kahsay, A.,

Asgedom, S.W., Gebre, A.K., Yimer, E.M., Belachew, A.B., Meles, H., Adsuar, J.C., Zodpey, S., Agrawal, S., Awasthi, A., Kumar, G.A., Ahmadi, A., Najafi, F., Rajati, F., Khazaie, H., Farzaei, M.H., Moradi, M., Rezaei, S., Soofi, M., Siabani, S., Rezaeian, S., Ahmed, M.B., Gebrehiwot, T.T., Feyissa, G.T., Hussen, M.A., Aichour, A., Aichour, I., Aichour, M.T.E., Akinyemi, R.O., Owolabi, M.O., Akseer, N., Bhutta, Z.A., Badawi, A., Alahdab, F., Kassa, G.M., Alebel, A., Wagnew, F.W.S., Belay, Y.A., Leshargie, C.T., Alghnam, S.A., Ali, B.A., Alsharif, U., Temsah, M.H., Altirkawi, K., Davitoiu, D. V., Hostiuc, S., Beuran, M., Negoi, I., Andrei, C.L., Ansari, H., Ansha, M.G., Antonio, C.A.T., Appiah, S.C.Y., Levi, M., Ariani, F., Biffino, M., Asefa, N.G., Atique, S., Rahman, M.A., Wijeratne, T., Ayala Quintanilla, B.P., Ayuk, T.B., Azzopardi, P.S., Rafiei, A., Badali, H., Mohammadi, M., Moosazadeh, M., Daryani, A., Banstola, A., Tran, K.B., Barker-Collo, S.L., Bärnighausen, T.W., Bedi, N., Tehrani-Banihashemi, A., Shabaninejad, H., Behzadifar, Ma, Kabir, A., Yousefifard, M., Moradi-Lakeh, M., Behzadifar, Me, Bekele, B., Hassen, H.Y., Henok, A., Biadgo, B., Bennett, D.A., Goulart, A.C., Bensenor, I.M., Lotufo, P.A., Wang, Y.P., Yisma, E., Berhane, A., Deribe, K., Yasin, Y.J., Demoz, G.T., Bhalla, A., Bhaumik, S., Bijani, A., Zamani, M., Bililign, N., Kumar, M., Birungi, C., Boufous, S., Brazinova, A., Brown, A.W., Car, M., Majeed, A., Rawaf, S., Rawaf, D.L., Shoman, H., Cárdenas, R., Carrero, J.J., El-Khatib, Z., Roy, N., Carvalho, F., Santos, J. V., Fernandes, E., Silva, J.P., Castañeda-Orjuela, C.A., Hoffman, H.J., Catalá-López, F., Chaiah, Y., Temsah, O., Champs, A.P., Chang, J.C., Choi, J.Y.J., Christopher, D.J., Cooper, C., Djalalinia, S., Do, H.P., Nguyen, T.H., Doku, D.T., Drake, T.M., Sheikh, A., Soyiri, I.N., Dubey, M., Santric Milicevic, M.M.M., Dubljanin, E., Faro, A., Filip, I., Radfar, A., Fischer, F., Fukumoto, T., Gankpe, F.G., Gopalkrishna, G., Haagsma, J.A., Polinder, S., Khoja, A.T., Pinilla-Monsalve, G.D., Hamadeh, R.R., Hamidi, S., Haro, J.M., Hassankhani, H., Havmoeller, R., Hegazy, M.I., Hendrie, D., Miller, T.R., Hibstu, D.T., Kassa, Z.Y., Hole, M.K., Homaie Rad, E., Hu, G., Ilesanmi, O.S., Jayaraman, S., Jha, R.P., Jonas, J.B., Moazen, B., Jorjoran Shushtari, Z., Jozwiak, J.J., Jürisson, M., Kahssay, M., Liben, M.L., Karch, A., Kengne, A.P., Khader, Y.S., Safari, H., Khafaie, M.A., Khalid, N., Khan, E.A., Khan, M.S., Usman, M.S., Siddiqi, T.J., Khang, Y.H., Khubchandani, J., Kiadaliri, A.A., Kim, D., Kim, Y.E., Yoon, S.J., Kisa, A., Winkler, A.S., Koyanagi, A., Kuate Defo, B., Kucuk Bicer, B., Lalloo, R., Moradinazar, M., Lami, F.H., Lansingh, V.C., Laryea, D.O., Latifi, A., Safiri, S., Lunevicius, R., Mahotra, N.B., Majdan, M., Manda, A.L., Mehndiratta, M.M., Mehta, V., Melese, A., Memiah, P.T.N., Mendoza, W., Mengistu, G., Shiferaw, M.S., Tekle, M.G., Meretoja, T.J., Meretoja, A., Szoeke, C.E.I., Mestrovich, T., Miazgowski, T., Mini, G.K., Mirica, A., Mirrakhimov, E.M., Molokhia, M., Monast, L., Ronfani, L., Mondello, S., Moradi, G., Moschos, M.M., Murthy, S., Musa, K.I., Mustafa, G., Naik, G., Schwebel, D.C., Nangia, V., Nascimento, B.R., Ningrum, D.N.A., Nyasulu, P.S., Renzaho, A.M.N., Ogbo, F.A., Oh, I.H., Okoro, A., Olagunju, A.T., Olagunju, T.O., Olivares, P.R., Otstavnov, S.S., Dnb, M.P.A., Pakhale, S., Pandey, A.R., Pesudovs, K., Prakash, S., Qorbani, M., Rafay, A., Rahman, M., Stokes, M.A., Shariful

Islam, S.M., Rai, R.K., Ram, U., Reis, C., Resnikoff, S., Roever, L., Sunguy, B.F., Ruhago, G.M., Saddik, B., Saldanha, R.D.F., Samy, A.M., Sanabria, J., Sartorius, B., Satpathy, M., Schneider, I.J.C., Shaikh, M.A.A., Sharif, M., She, J., Shen, J., Sheth, K.N., Shibuya, K., Shigematsu, M., Shiri, R., Shiue, I., Silveira, D.A., Soares Filho, A.M., Sinha, D.N., Soriano, J.B., Stein, D.J., Sufiyan, M.B., Sykes, B.L., Tabarés-Seisdedos, R., Tortajada-Girbés, M., Topor-Madry, R., Tran, B.X., Tudor Car, L., Ukwaja, K.N., Ullah, I., Uthman, O.A., Valdez, P.R., Vasankari, T.J., Venketasubramanian, N., Violante, F.S., Waheed, Y., Werdecker, A., Wyper, G.M.A., Yano, Y., Ye, P., Yip, P., Yonemoto, N., Younis, M.Z., Yu, C., Zaidi, Z., Zaman, S.B., 2019. Global, regional, and national burden of traumatic brain injury and spinal cord injury, 1990-2016: A systematic analysis for the Global Burden of Disease Study 2016. Lancet Neurol. https://doi.org/10.1016/S1474-4422(18)30415-0

Jamjoom, A.A.B., Rhodes, J., Andrews, P.J.D., Grant, S.G.N., 2021. The synapse in traumatic brain injury. Brain 144, 18. https://doi.org/10.1093/BRAIN/AWAA321

Jenkinson, M., Bannister, P., Brady, M., Smith, S., 2002. Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images. https://doi.org/10.1006/nimg.2002.1132

Jenkinson, M., Smith, S., 2001. A global optimisation method for robust affine registration of brain images. Med. Image Anal. 5, 143–156. https://doi.org/10.1016/S1361-8415(01)00036-6

Jensen, O., Colgin, L.L., 2007. Cross-frequency coupling between neuronal oscillations. Trends Cogn. Sci. 11, 267–269. https://doi.org/10.1016/j.tics.2007.05.003

Johnson, C., Schwindt, P.D.D., Weisend, M., 2010. Magnetoencephalography with a two-color pump-probe, fiber-coupled atomic magnetometer. Appl. Phys. Lett. 97, 243703. https://doi.org/10.1063/1.3522648

Johnson, C.N., Schwindt, P.D.D., Weisend, M., 2013. Multi-sensor magnetoencephalography with atomic magnetometers. Phys. Med. Biol. 58, 6065–6077. https://doi.org/10.1088/0031-9155/58/17/6065

Jones, S.R., 2016. When brain rhythms aren't 'rhythmic': implication for their mechanisms and meaning. Curr. Opin. Neurobiol. https://doi.org/10.1016/j.conb.2016.06.010

Josephson, B.D., 1962. Possible new effects in superconductive tunnelling. Phys. Lett. 1, 251–253. https://doi.org/10.1016/0031-9163(62)91369-0

Jung, T.-P., Makeig, S., Humphries, C., Lee, T.-W., McKeown, M.J., Iragui, V., Sejnowski, T.J., 2000. Removing electroencephalographic artifacts by blind source separation. Psychophysiology 37, 163–178. https://doi.org/10.1111/1469-8986.3720163

Jurkiewicz, M.T., Gaetz, W.C., Bostan, A.C., Cheyne, D., 2006. Post-movement beta rebound is generated in motor cortex: Evidence from neuromagnetic

recordings. Neuroimage 32, 1281–1289. https://doi.org/10.1016/j.neuroimage.2006.06.005

Kamada, K., Sato, D., Ito, Y., Natsukawa, H., Okano, K., Mizutani, N., Kobayashi, T., 2015. Human magnetoencephalogram measurements using newly developed compact module of high-sensitivity atomic magnetometer. Jpn. J. Appl. Phys. 54, 026601. https://doi.org/10.7567/JJAP.54.026601

Kandel, E.R., 2013. Principles of neural science, 5th ed. ed. McGraw-Hill Medical, New York.

Kastler, A., 1973. The Hanle effect and its use for the measurements of very small magnetic fields. Nucl. Instruments Methods 110, 259–265. https://doi.org/10.1016/0029-554X(73)90698-8

Kato, T., Kamei, A., Takashima, S., Ozaki, T., 1993. Human visual cortical function during photic stimulation monitoring by means of near-infrared spectroscopy. J. Cereb. Blood Flow Metab. 13, 516–520. https://doi.org/10.1038/jcbfm.1993.66

Kelly, H.W.K., 1946. Historical introduction to degaussing, covering pre-war work and development up to the start of the campaign against the German mine. J. Inst. Electr. Eng. - Part I Gen. 93, 430–434. https://doi.org/10.1049/JI-1.1946.0129

Kilavik, B.E., Zaepffel, M., Brovelli, A., MacKay, W.A., Riehle, A., 2013. The ups and downs of beta oscillations in sensorimotor cortex. Exp. Neurol. 245, 15–26. https://doi.org/10.1016/J.EXPNEUROL.2012.09.014

Kirov, I.I., Tal, A., Babb, J.S., Reaume, J., Bushnik, T., Ashman, T.A., Flanagan, S., Grossman, R.I., Gonen, O., 2013. Proton MR spectroscopy correlates diffuse axonal abnormalities with post-concussive symptoms in mild traumatic brain injury. J. Neurotrauma 30, 1200–1204. https://doi.org/10.1089/neu.2012.2696

Kopell, N., Ermentrout, G.B., Whittington, M.A., Traub, R.D., 2000. Gamma rhythms and beta rhythms have different synchronization properties. Proc. Natl. Acad. Sci. U. S. A. 97, 1867–1872. https://doi.org/10.1073/pnas.97.4.1867

Kroenke, K., Spitzer, R.L., Williams, J.B.W., 2001. The PHQ-9: Validity of a brief depression severity measure. J. Gen. Intern. Med. 16, 606–613. https://doi.org/10.1046/j.1525-1497.2001.016009606.x

Kuhn, H.W., 1955. The Hungarian method for the assignment problem. Nav. Res. Logist. Q. 2, 83–97. https://doi.org/10.1002/NAV.3800020109

Labyt, E., Corsi, M.-C., Fourcault, W., Palacios Laloy, A., Bertrand, F., Lenouvel, F., Cauffet, G., Le Prado, M., Berger, F., Morales, S., 2019. Magnetoencephalography With Optically Pumped 4 He Magnetometers at Ambient Temperature. IEEE Trans. Med. Imaging 38, 90–98. https://doi.org/10.1109/TMI.2018.2856367

Lachaux, J.P., Rodriguez, E., Martinerie, J., Varela, F.J., 1999. Measuring phase synchrony in brain signals. Hum. Brain Mapp. 8, 194–208. https://doi.org/10.1002/(SICI)1097-0193(1999)8:4<194::AID-HBM4>3.0.CO;2-C

Lange, R.T., Iverson, G., Franzen, M.D., 2009. Neuropsychological functioning following complicated vs. uncomplicated mild traumatic brain injury. Brain Inj. 23, 83–91. https://doi.org/10.1080/02699050802635281

Lauterbur, P.C., 1973. Image formation by induced local interactions: Examples employing nuclear magnetic resonance. Nature 242, 190–191. https://doi.org/10.1038/242190a0

Levin, H.S., Shum, D.H.K., Chan, R.C.K., 2014. Recent advances in traumatic brain injury research: Introduction. Underst. Trauma. brain Inj. Curr. Res. Futur. Dir. 1–3.

Levin, H.S., Temkin, N.R., Barber, J., Nelson, L.D., Robertson, C., Brennan, J., Stein, M.B., Yue, J.K., Giacino, J.T., McCrea, M.A., Diaz-Arrastia, R., Mukherjee, P., Okonkwo, D.O., Boase, K., Markowitz, A.J., Bodien, Y., Taylor, S., Vassar, M.J., Manley, G.T., Adeoye, O., Badjatia, N., Bullock, M.R., Chesnut, R., Corrigan, J.D., Crawford, K., Dikmen, S., Duhaime, A.C., Ellenbogen, R., Feeser, V.R., Ferguson, A.R., Foreman, B., Gardner, R., Gaudette, E., Gonzalez, L., Gopinath, S., Gullapalli, R., Hemphill, J.C., Hotz, G., Jain, S., Keene, C.D., Korley, F.K., Kramer, J., Kreitzer, N., Lindsell, C., Machamer, J., Madden, C., Martin, A., McAllister, T., Merchant, R., Nolan, A., Ngwenya, L.B., Noel, F., Palacios, E., Puccio, A., Rabinowitz, M., Rosand, J., Sander, A., Satris, G., Schnyer, D., Seabury, S., Sun, X., Toga, A., Valadka, A., Wang, K., Yuh, E., Zafonte, R., 2021. Association of Sex and Age With Mild Traumatic Brain Injury–Related Symptoms: A TRACK-TBI Study. JAMA Netw. Open 4, e213046–e213046. https://doi.org/10.1001/JAMANETWORKOPEN.2021.3046

Lewine, J.D., Davis, J.T., Sloan, J.H., Kodituwakku, P.W., Orrison, W.W., 1999. Neuromagnetic assessment of pathophysiologic brain activity induced by minor head trauma. AJNR. Am. J. Neuroradiol. 20, 857–66.

Little, S., Bonaiuto, J., Barnes, G., Bestmann, S., 2019. Human motor cortical beta bursts relate to movement planning and response errors. PLoS Biol. 17, e3000479. https://doi.org/10.1371/journal.pbio.3000479

Löwdin, P.O., 1950. On the non-orthogonality problem connected with the use of atomic wave functions in the theory of molecules and crystals. J. Chem. Phys. 18, 365–375. https://doi.org/10.1063/1.1747632

Luckhoo, H.T., Brookes, M.J., Woolrich, M.W., 2014. Multi-session statistics on beamformed MEG data. Neuroimage 95, 330–335. https://doi.org/10.1016/j.neuroimage.2013.12.026

Lunkova, E., Guberman, G.I., Ptito, A., Saluja, R.S., 2021. Noninvasive magnetic resonance imaging techniques in mild traumatic brain injury research and

diagnosis. Hum. Brain Mapp. 42, 5477–5494. https://doi.org/10.1002/hbm.25630

MacFarlane, M.P., Glenn, T.C., 2015. Neurochemical cascade of concussion. Brain Inj. https://doi.org/10.3109/02699052.2014.965208

Mansfield, P., Grannell, P.K., 1973. NMR 'diffraction' in solids? J. Phys. C Solid State Phys. 6, L422–L426. https://doi.org/10.1088/0022-3719/6/22/007

Marklund, N., Vedung, F., Lubberink, M., Tegner, Y., Johansson, J., Blennow, K., Zetterberg, H., Fahlström, M., Haller, S., Stenson, S., Larsson, E.-M., Wall, A., Antoni, G., 2021. Tau aggregation and increased neuroinflammation in athletes after sports-related concussions and in traumatic brain injury patients-a PET/MR study. NeuroImage Clin. 30, 102665. https://doi.org/10.1016/j.nicl.2021.102665

Martin Bland, J., Altman, D.G., 1986. STATISTICAL METHODS FOR ASSESSING AGREEMENT BETWEEN TWO METHODS OF CLINICAL MEASUREMENT. Lancet 327, 307–310. https://doi.org/10.1016/S0140-6736(86)90837-8

McColgan, P., Joubert, J., Tabrizi, S.J., Rees, G., 2020. The human motor cortex microcircuit: insights for neurodegenerative disease. Nat. Rev. Neurosci. https://doi.org/10.1038/s41583-020-0315-1

McCrory, P., Meeuwisse, W., Johnston, K., Dvorak, J., Aubry, M., Molloy, M., Cantu, R., 2009. Consensus statement on concussion in sport – The 3rd International Conference on concussion in sport, held in Zurich, November 2008. J. Clin. Neurosci. 16, 755–763. https://doi.org/10.1016/J.JOCN.2009.02.002

McDonald, B.C., Saykin, A.J., McAllister, T.W., 2012. Functional MRI of mild traumatic brain injury (mTBI): Progress and perspectives from the first decade of studies. Brain Imaging Behav. 6, 193–207. https://doi.org/10.1007/s11682-012-9173-4

McInnes, K., Friesen, C.L., MacKenzie, D.E., Westwood, D.A., Boe, S.G., 2017. Mild Traumatic Brain Injury (mTBI) and chronic cognitive impairment: A scoping review. PLoS One 12, e0174847. https://doi.org/10.1371/JOURNAL.PONE.0174847

Meterko, M., Baker, E., Stolzmann, K.L., Hendricks, A.M., Cicerone, K.D., Lew, H.L., 2012. Psychometric assessment of the neurobehavioral symptom inventory-22: The structure of persistent postconcussive symptoms following deployment-related mild traumatic brain injury among veterans. J. Head Trauma Rehabil. 27, 55–62. https://doi.org/10.1097/HTR.0B013E318230FB17

National Center for Injury Prevention and Control, 2003. Report to Congress on Mild Traumatic Brain Injury in the United States: Steps to Prevent a Serious Public Health Problem. Centers for Disease Control and Prevention, Atlanta, GA.

Nelson, L.D., Temkin, N.R., Dikmen, S., Barber, J., Giacino, J.T., Yuh, E., Levin, H.S., McCrea, M.A., Stein, M.B., Mukherjee, P., Okonkwo, D.O., Diaz-Arrastia, R., Manley, G.T., Adeoye, O., Badjatia, N., Boase, K., Bodien, Y., Bullock, M.R., Chesnut, R., Corrigan, J.D., Crawford, K., Duhaime, A.C., Ellenbogen, R., Feeser, V.R., Ferguson, A., Foreman, B., Gardner, R., Gaudette, E., Gonzalez, L., Gopinath, S., Gullapalli, R., Hemphill, J.C., Hotz, G., Jain, S., Korley, F., Kramer, J., Kreitzer, N., Lindsell, C., MacHamer, J., Madden, C., Martin, A., McAllister, T., Merchant, R., Noel, F., Palacios, E., Perl, D., Puccio, A., Rabinowitz, M., Robertson, C.S., Rosand, J., Sander, A., Satris, G., Schnyer, D., Seabury, S., Sherer, M., Taylor, S., Toga, A., Valadka, A., Vassar, M.J., Vespa, P., Wang, K., Yue, J.K., Zafonte, R., 2019. Recovery after Mild Traumatic Brain Injury in Patients Presenting to US Level i Trauma Centers: A Transforming Research and Clinical Knowledge in Traumatic Brain Injury (TRACK-TBI) Study. JAMA Neurol. 76, 1049–1059. https://doi.org/10.1001/jamaneurol.2019.1313

Nolte, G., 2003. The magnetic lead field theorem in the quasi-static approximation and its use for magnetoenchephalography forward calculation in realistic volume conductors. Phys. Med. Biol. 48, 3637–3652. https://doi.org/10.1088/0031-9155/48/22/002

Nolte, G., Bai, O., Wheaton, L., Mari, Z., Vorbach, S., Hallett, M., 2004. Identifying true brain interaction from EEG data using the imaginary part of coherency. Clin. Neurophysiol. 115, 2292–2307. https://doi.org/10.1016/j.clinph.2004.04.029

Nuwer, M.R., Hovda, D.A., Schrader, L.M., Vespa, P.M., 2005. Routine and quantitative EEG in mild traumatic brain injury. Clin. Neurophysiol. 116, 2001–2025. https://doi.org/10.1016/J.CLINPH.2005.05.008

O'Neill, G.C., Barratt, E.L., Hunt, B.A.E., Tewarie, P.K., Brookes, M.J., 2015a. Measuring electrophysiological connectivity by power envelope correlation: A technical review on MEG methods. Phys. Med. Biol. 60, R271–R295. https://doi.org/10.1088/0031-9155/60/21/R271

O'Neill, G.C., Bauer, M., Woolrich, M.W., Morris, P.G., Barnes, G.R., Brookes, M.J., 2015b. Dynamic recruitment of resting state sub-networks. Neuroimage 115, 85–95. https://doi.org/10.1016/j.neuroimage.2015.04.030

O'Neill, G.C., Tewarie, P.K., Colclough, G.L., Gascoyne, L.E., Hunt, B.A.E., Morris, P.G., Woolrich, M.W., Brookes, M.J., 2017. Measurement of dynamic task related functional networks using MEG. Neuroimage 146, 667–678. https://doi.org/10.1016/j.neuroimage.2016.08.061

Ogawa, S., Lee, T.M., Kay, A.R., Tank, D.W., 1990. Brain magnetic resonance imaging with contrast dependent on blood oxygenation. Proc. Natl. Acad. Sci. U. S. A. 87, 9868–9872. https://doi.org/10.1073/pnas.87.24.9868

Oh, A., Vidal, J., Taylor, M.J., Pang, E.W., 2014. Neuromagnetic correlates of intra- and extra-dimensional set-shifting. Brain Cogn. 86, 90–97. https://doi.org/10.1016/j.bandc.2014.02.006

Omalu, B.I., DeKosky, S.T., Hamilton, R.L., Minster, R.L., Kamboh, M.I., Shakir, A.M., Wecht, C.H., 2006. Chronic traumatic encephalopathy in a national football league player: part II. Neurosurgery 59, 1086–1092. https://doi.org/10.1227/01.NEU.0000245601.69451.27

Omalu, B.I., DeKosky, S.T., Minster, R.L., Kamboh, M.I., Hamilton, R.L., Wecht, C.H., 2005. Chronic traumatic encephalopathy in a National Football League player. Neurosurgery 57, 128–133. https://doi.org/10.1227/01.NEU.0000163407.92769.ED

Onnes, K., 1911. Further experiments with liquid helium. C. On the change of electric resistance of pure metals at very low temperatures etc. IV. The resistance of pure mercury at helium temperatures.

Oostenveld, R., Fries, P., Maris, E., Schoffelen, J.M., 2011. FieldTrip: Open source software for advanced analysis of MEG, EEG, and invasive electrophysiological data. Comput. Intell. Neurosci. 2011. https://doi.org/10.1155/2011/156869

Pakenham, D.O., Quinn, A.J., Fry, A., Francis, S.T., Woolrich, M.W., Brookes, M.J., Mullinger, K.J., 2020. Post-stimulus beta responses are modulated by task duration. Neuroimage 206, 116288. https://doi.org/10.1016/J.NEUROIMAGE.2019.116288

Pang, E.W., Dunkley, B.T., Doesburg, S.M., da Costa, L., Taylor, M.J., 2016. Reduced brain connectivity and mental flexibility in mild traumatic brain injury. Ann. Clin. Transl. Neurol. https://doi.org/10.1002/acn3.280

Park, J.H., Park, S.W., Kang, S.H., Nam, T.K., Min, B.K., Hwang, S.N., 2009. Detection of Traumatic Cerebral Microbleeds by Susceptibility-Weighted Image of MRI. J. Korean Neurosurg. Soc. 46, 365. https://doi.org/10.3340/JKNS.2009.46.4.365

Pearce, A.J., Rist, B., Fraser, C.L., Cohen, A., Maller, J.J., 2018. Neurophysiological and cognitive impairment following repeated sports concussion injuries in retired professional rugby league players. Brain Inj. 32, 498–505. https://doi.org/10.1080/02699052.2018.1430376

Petrenko, M. V., Dmitriev, S.P., Pazgalev, A.S., Ossadtchi, A.E., Vershovskii, A.K., 2021. Towards the Non-Zero Field Cesium Magnetic Sensor Array for Magnetoencephalography. IEEE Sens. J. 21, 18626–18632. https://doi.org/10.1109/JSEN.2021.3089455

Pfurtscheller, G., Lopes Da Silva, F.H., 1999. Event-related EEG/MEG synchronization and desynchronization: basic principles. Clin. Neurophysiol. 110, 1842–1857. https://doi.org/10.1016/S1388-2457(99)00141-8

Pfurtscheller, G., Neuper, C., Mohl, W., 1994. Event-related desynchronization (ERD) during visual processing. Int. J. Psychophysiol. 16, 147–153. https://doi.org/10.1016/0167-8760(89)90041-X

Pfurtscheller, G., Stancák, A., Neuper, C., 1996. Post-movement beta synchronization. A correlate of an idling motor area? Electroencephalogr. Clin. Neurophysiol. 98, 281–293. https://doi.org/10.1016/0013-4694(95)00258-8

Powell, J.M., Ferraro, J. V., Dikmen, S.S., Temkin, N.R., Bell, K.R., 2008. Accuracy of mild traumatic brain injury diagnosis. Arch. Phys. Med. Rehabil. 89, 1550–1555. https://doi.org/10.1016/J.APMR.2007.12.035

Proskovec, A.L., Shah, B.R., Yu, F.F., Achilleos, M., Maldjian, J.A., Davenport, E.M., 2020. Magnetoencephalography and Mild Traumatic Brain Injury. Adv. Clin. Radiol. 2, 341–350. https://doi.org/10.1016/j.yacr.2020.05.004

Purves, D., Augustine, G., Fitzpatrick, D., et al., 2001a. An Overview of Cortical Structure., in: Neuroscience. Sinauer Associates, Sunderland (MA).

Purves, D., Augustine, G., Fitzpatrick, D., et al., editors, 2001b. Ligand-Gated Ion Channels, in: Neuroscience. Sinauer Associates, Sunderland (MA).

Quinn, A.J., van Ede, F., Brookes, M.J., Heideman, S.G., Nowak, M., Seedat, Z.A., Vidaurre, D., Zich, C., Nobre, A.C., Woolrich, M.W., 2019. Unpacking Transient Event Dynamics in Electrophysiological Power Spectra. Brain Topogr. 32, 1020–1034. https://doi.org/10.1007/s10548-019-00745-5

Rabiner, L.R., 1989. A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition. Proc. IEEE 77, 257–286. https://doi.org/10.1109/5.18626

Raichle, M.E., Martin, W.R.W., Herscovitch, P., Mintun, M.A., Markham, J., 1983. Brain blood flow measured with intravenous H215O.: II. Implementation and validation. J. Nucl. Med. 24, 790–798.

Raskin, S.A., Lovejoy, D.W., Stevens, M.C., Zamroziewicz, M., Oakes, H.J., 2014. Mild Traumatic Brain Injury, in: Levin, H.S., Shum, D.H.K., Chan, R.C.K. (Eds.), Understanding Traumatic Brain Injury. Oxford University Press, USA, pp. 370–312.

Rea, M., Holmes, N., Hill, R.M., Boto, E., Leggett, J., Edwards, L.J., Woolger, D., Dawson, E., Shah, V., Osborne, J., Bowtell, R., Brookes, M.J., 2021. Precision magnetic field modelling and control for wearable magnetoencephalography. Neuroimage 241, 118401. https://doi.org/10.1016/J.NEUROIMAGE.2021.118401

Reitan, R.M., 1971. Trail Making Test Results for Normal and Brain-Damaged Children. Percept. Mot. Skills 33, 575–581. https://doi.org/10.2466/pms.1971.33.2.575

Reuters, 2021. Soccer-English football to limit high-force headers in training from 2021-22 season [WWW Document]. URL https://www.reuters.com/article/uk-soccer-england-idUKKBN2EY1GZ (accessed 6.5.22).

Reuters, 2020. Rugby-Six more former players join concussion lawsuit [WWW

Document]. URL https://www.reuters.com/article/rugby-union-concussion-idUKL4N2IX2ZY (accessed 6.5.22).

Rey, A., 1958. L'examen clinique en psychologie. [The clinical examination in psychology.], L'examen clinique en psychologie. Presses Universitaries De France, Oxford, England.

Rezek, I., Roberts, S., 2005. Ensemble Hidden Markov Models with Extended Observation Densities for Biosignal Analysis. Probabilistic Model. Bioinforma. Med. Informatics 419–450. https://doi.org/10.1007/1-84628-119-9_14

Rier, L., Zamyadi, R., Zhang, J., Emami, Z., Seedat, Z.A., Mocanu, S., Gascoyne, L.E., Allen, C.M., Scadding, J.W., Furlong, P.L., Gooding-Williams, G., Woolrich, M.W., Evangelou, N., Brookes, M.J., Dunkley, B.T., 2021. Mild traumatic brain injury impairs the coordination of intrinsic and motor-related neural dynamics. NeuroImage Clin. 32, 102841. https://doi.org/10.1016/J.NICL.2021.102841

Rjabinin, J.N., Shubnikow, L.W., 1935. Magnetic properties and critical currents of supraconducting alloys [3]. Nature 135, 581–582. https://doi.org/10.1038/135581A0

Robinson, S.E., Vrba, J., 1999. Functional neuroimaging by synthetic aperture magnetometry (SAM). Recent Adv. Biomagn.

Rong, F., Contreras-Vidal, J.L., 2006. Magnetoencephalographic artifact identification and automatic removal based on independent component analysis and categorization approaches. J. Neurosci. Methods 157, 337–354. https://doi.org/10.1016/j.jneumeth.2006.04.024

Ruggiero, K.J., Ben, K. Del, Scotti, J.R., Rabalais, A.E., 2003. Psychometric properties of the PTSD checklist—civilian version. J. Trauma. Stress 16, 495–502. https://doi.org/10.1023/A:1025714729117

Ruppel, R.A., Kochanek, P.M., Adelson, P.D., Rose, M.E., Wisniewski, S.R., Bell, M.J., Clark, R.S.B., Marion, D.W., Graham, S.H., 2001. Excitatory amino acid concentrations in ventricular cerebrospinal fluid after severe traumatic brain injury in infants and children: The role of child abuse. J. Pediatr. 138, 18–25. https://doi.org/10.1067/mpd.2001.110979

Sadaghiani, S., Brookes, M.J., Baillet, S., 2021. Connectomics of Human Electrophysiology. Neuroimage 247, 118788. https://doi.org/10.1016/j.neuroimage.2021.118788

Sareen, E., Zahar, S., Ville, D. Van De, Gupta, A., Griffa, A., Amico, E., 2021. Exploring MEG brain fingerprints: Evaluation, pitfalls, and interpretations. Neuroimage 240, 118331. https://doi.org/10.1016/j.neuroimage.2021.118331

Sarvas, J., 1987. Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem. Phys. Med. Biol. 32, 11–22. https://doi.org/10.1088/0031-9155/32/1/004

Satz, P., Alfano, M.S., Light, R., Morgenstern, H., Zaucha, K., Asarnow, R.F., Newton, S., 1999. Persistent Post-Concussive Syndrome: A proposed methodology and literature review to determine the effects, if any, of mild head and other bodily injury. J. Clin. Exp. Neuropsychol. 21, 620–628. https://doi.org/10.1076/jcen.21.5.620.870

Schoffelen, J.M., Gross, J., 2009. Source connectivity analysis with MEG and EEG. Hum. Brain Mapp. 30, 1857–1865. https://doi.org/10.1002/HBM.20745

Schwartz J, DoBaMBT, J., S, S., AJ., H., 2012. Principles of Neural Science, Fifth Edition, 5th ed. Blacklick: McGraw-Hill Publishing.

Schwindt, P., Liew, L.-A., Shah, V., Moreland, J., Hollberg, L., Schwindt, P., Liew, L.-A., Shah, V., Moreland, J., Hollberg, L., 2005. A millimeter-scale atomic magnetometer. APS 36, D6.103.

Schwindt, P.D.D., Knappe, S., Shah, V., Hollberg, L., Kitching, J., Liew, L.-A., Moreland, J., 2004. Chip-scale atomic magnetometer. Appl. Phys. Lett. 85, 6409. https://doi.org/10.1063/1.1839274

Seedat, Z.A., Quinn, A.J., Vidaurre, D., Liuzzi, L., Gascoyne, L.E., Hunt, B.A.E., O'Neill, G.C., Pakenham, D.O., Mullinger, K.J., Morris, P.G., Woolrich, M.W., Brookes, M.J., 2020. The role of transient spectral 'bursts' in functional connectivity: A magnetoencephalography study. Neuroimage 209, 116537. https://doi.org/10.1016/j.neuroimage.2020.116537

Sekihara, K., Nagarajan, S.S., Poeppel, D., Marantz, A., 2004. Asymptotic SNR of scalar and vector minimum-variance beanformers for neuromagnetic source reconstruction. IEEE Trans. Biomed. Eng. 51, 1726–1734. https://doi.org/10.1109/TBME.2004.827926

Seymour, R.A., Rippon, G., Kessler, K., 2017. The detection of phase amplitude coupling during sensory processing. Front. Neurosci. 11, 487. https://doi.org/10.3389/fnins.2017.00487

Shah, V., Hughes, J.K., 2015. Method for Detecting Zero-field Resonance. 20150212168.

Shah, V., Knappe, S., Schwindt, P.D.D., Kitching, J., 2007. Subpicotesla atomic magnetometry with a microfabricated vapour cell. https://doi.org/10.1038/nphoton.2007.201

Sharma, A., Hind, K., Hume, P., Singh, J., Neary, J.P., 2020. Neurovascular Coupling by Functional Near Infra-Red Spectroscopy and Sport-Related Concussion in Retired Rugby Players: The UK Rugby Health Project. Front. Hum. Neurosci. 14, 42. https://doi.org/10.3389/fnhum.2020.00042

Shenton, M.E., Hamoda, H.M., Schneiderman, J.S., Bouix, S., Pasternak, O., Rathi, Y., Vu, M.-A., Purohit, M.P., Helmer, K., Koerte, I., Lin, A.P., Westin, C.-F., Kikinis, R., Kubicki, M., Stern, R.A., Zafonte, R., 2012. A review of magnetic resonance imaging and diffusion tensor imaging findings in mild traumatic brain injury. Brain Imaging Behav. 6, 137–92.

https://doi.org/10.1007/s11682-012-9156-5

Sherman, M.A., Lee, S., Law, R., Haegens, S., Thorn, C.A., Hämäläinen, M.S., Moore, C.I., Jones, S.R., 2016. Neural mechanisms of transient neocortical beta rhythms: Converging evidence from humans, computational modeling, monkeys, and mice. Proc. Natl. Acad. Sci. U. S. A. 113, E4885–E4894. https://doi.org/10.1073/pnas.1604135113

Shi, J., Teng, J., Du, X., Li, N., 2021. Multi-Modal Analysis of Resting-State fMRI Data in mTBI Patients and Association With Neuropsychological Outcomes. Front. Neurol. 12, 600. https://doi.org/10.3389/fneur.2021.639760

Silva, L.R., Amitai, Y., Connors, B.W., 1991. Intrinsic oscillations of neocortex generated by layer 5 pyramidal neurons. Science (80-.). 251, 432–435. https://doi.org/10.1126/science.1824881

Simon, D.W., McGeachy, M.J., Bayır, H., Clark, R.S.B., Loane, D.J., Kochanek, P.M., 2017. The far-reaching scope of neuroinflammation after traumatic brain injury. Nat. Rev. Neurol. 13, 171–191. https://doi.org/10.1038/nrneurol.2017.13

Slocum, R.E., Marton, B.I., 1973. Measurement of Weak Magnetic Fields Using Zero-Field Parametric Resonance in Optically Pumped He4. IEEE Trans. Magn. 9, 221–226. https://doi.org/10.1109/TMAG.1973.1067647

Smith, S.M., 2002. Fast robust automated brain extraction. Hum. Brain Mapp. https://doi.org/10.1002/hbm.10062

Spitzer, R.L., Kroenke, K., Williams, J.B.W., Löwe, B., 2006. A brief measure for assessing generalized anxiety disorder: The GAD-7. Arch. Intern. Med. 166, 1092–1097. https://doi.org/10.1001/archinte.166.10.1092

Stam, C.J., Nolte, G., Daffertshofer, A., 2007. Phase lag index: Assessment of functional connectivity from multi channel EEG and MEG with diminished bias from common sources. Hum. Brain Mapp. 28, 1178–1193. https://doi.org/10.1002/hbm.20346

Stenroos, M., Hunold, A., Haueisen, J., 2014. Comparison of three-shell and simplified volume conductor models in magnetoencephalography. Neuroimage 94, 337–348. https://doi.org/10.1016/J.NEUROIMAGE.2014.01.006

Stoffers, D., Bosboom, J.L.W., Deijen, J.B., Wolters, E.C., Berendse, H.W., Stam, C.J., 2007. Slowing of oscillatory brain activity is a stable characteristic of Parkinson's disease without dementia. Brain 130, 1847–1860. https://doi.org/10.1093/brain/awm034

Tewarie, P., Hunt, B.A.E., O'Neill, G.C., Byrne, A., Aquino, K., Bauer, M., Mullinger, K.J., Coombes, S., Brookes, M.J., 2019. Relationships between Neuronal Oscillatory Amplitude and Dynamic Functional Connectivity. Cereb. Cortex 29, 2668–2681. https://doi.org/10.1093/cercor/bhy136

Tierney, T.M., Alexander, N., Mellor, S., Holmes, N., Seymour, R., O'Neill, G.C., Maguire, E.A., Barnes, G.R., 2021. Modelling optically pumped magnetometer interference in MEG as a spatially homogeneous magnetic field. Neuroimage 244, 118484. https://doi.org/10.1016/j.neuroimage.2021.118484

Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., Delcroix, N., Mazoyer, B., Joliot, M., 2002. Automated anatomical labeling of activations in SPM using a macroscopic anatomical parcellation of the MNI MRI single-subject brain. Neuroimage 15, 273–289. https://doi.org/10.1006/nimg.2001.0978

Urban, K.J., Barlow, K.M., Jimenez, J.J., Goodyear, B.G., Dunn, J.F., 2015. Functional Near-Infrared Spectroscopy Reveals Reduced Interhemispheric Cortical Communication after Pediatric Concussion. J. Neurotrauma 32, 833–840. https://doi.org/10.1089/neu.2014.3577

van den Heuvel, M.P., Hulshoff Pol, H.E., 2010. Exploring the brain network: A review on resting-state fMRI functional connectivity. Eur. Neuropsychopharmacol. https://doi.org/10.1016/j.euroneuro.2010.03.008

van Ede, F., Quinn, A.J., Woolrich, M.W., Nobre, A.C., 2018. Neural Oscillations: Sustained Rhythms or Transient Burst-Events? Trends Neurosci. https://doi.org/10.1016/j.tins.2018.04.004

Van Veen, B.D., Buckley, K.M., 1988. Beamforming: A Versatile Approach to Spatial Filtering. IEEE ASSP Mag. 5, 4–24. https://doi.org/10.1109/53.665

Van Veen, B.D., van Drongelen, W., Yuchtman, M., Suzuki, A., 1997. Localization of brain electrical activity via linearly constrained minimum variance spatial filtering. IEEE Trans. Biomed. Eng. 44, 867–80. https://doi.org/10.1109/10.623056

Vanderploeg, R.D., Silva, M.A., Soble, J.R., Curtiss, G., Belanger, H.G., Donnell, A.J., Scott, S.G., 2015. The structure of postconcussion symptoms on the neurobehavioral symptom inventory: A comparison of alternative models. J. Head Trauma Rehabil. 30, 1–11. https://doi.org/10.1097/HTR.0000000000000009

Vergara, V.M., Mayer, A.R., Damaraju, E., Kiehl, K.A., Calhoun, V., 2017. Detection of Mild Traumatic Brain Injury by Machine Learning Classification Using Resting State Functional Network Connectivity and Fractional Anisotropy. J. Neurotrauma 34, 1045–1053. https://doi.org/10.1089/neu.2016.4526

Vergara, V.M., Mayer, A.R., Kiehl, K.A., Calhoun, V.D., 2018. Dynamic functional network connectivity discriminates mild traumatic brain injury through machine learning. NeuroImage Clin. 19, 30–37. https://doi.org/10.1016/j.nicl.2018.03.017

Vespa, P., Prins, M., Ronne-Engstrom, E., Caron, M., Shalmon, E., Hovda, D.A., Martin, N.A., Becker, D.P., 1998. Increase in extracellular glutamate caused

by reduced cerebral perfusion pressure and seizures after human traumatic brain injury: A microdialysis study. J. Neurosurg. 89, 971–982. https://doi.org/10.3171/jns.1998.89.6.0971

Vicente, R., Gollo, L.L., Mirasso, C.R., Fischer, I., Pipa, G., 2008. Dynamical relaying can yield zero time lag neuronal synchrony despite long conduction delays. Proc. Natl. Acad. Sci. U. S. A. 105, 17157–17162. https://doi.org/10.1073/pnas.0809353105

Vidaurre, D., Hunt, L.T., Quinn, A.J., Hunt, B.A.E., Brookes, M.J., Nobre, A.C., Woolrich, M.W., 2018. Spontaneous cortical activity transiently organises into frequency specific phase-coupling networks. Nat. Commun. 9, 1–13. https://doi.org/10.1038/s41467-018-05316-z

Vidaurre, D., Quinn, A.J., Baker, A.P., Dupret, D., Tejero-Cantero, A., Woolrich, M.W., 2016. Spectrally resolved fast transient brain states in electrophysiological data Keywords: multivariate autoregressive model MEG Transient connectivity Bayesian modelling Spectral estimation Multitaper Coherence Partial directed coherence Sign ambiguity. Neuroimage. https://doi.org/10.1016/j.neuroimage.2015.11.047

Vigário, R.N., 1997. Extraction of ocular artefacts from EEG using independent component analysis. Electroencephalogr. Clin. Neurophysiol. 103, 395–404. https://doi.org/10.1016/S0013-4694(97)00042-8

Vrba, J., Robinson, S.E., 2001. Signal Processing in Magnetoencephalography. Methods 25, 249–271. https://doi.org/10.1006/meth.2001.1238

Wang, X., Wei, X.E., Li, M.H., Li, W. Bin, Zhou, Y.J., Zhang, B., Li, Y.H., 2014. Microbleeds on susceptibility-weighted MRI in depressive and non-depressive patients after mild traumatic brain injury. Neurol. Sci. 35, 1533–1539. https://doi.org/10.1007/S10072-014-1788-3/FIGURES/4

Wang, Y., Bartels, H.M., Nelson, L.D., 2020. A Systematic Review of ASL Perfusion MRI in Mild TBI. Neuropsychol. Rev. 1–32. https://doi.org/10.1007/s11065-020-09451-7

Wechsler, D., 2009. Subtest Administration and Scoring. WAIS--IV: Administration and Scoring Manual. San Antonio, TX Psychol. Corp.

Wechsler, D., 1944. The measurement of adult intelligence.

Wharton, S., Schäfer, A., Bowtell, R., 2010. Susceptibility mapping in the human brain using threshold-based k-space division. Magn. Reson. Med. 63, 1292–1304. https://doi.org/10.1002/mrm.22334

Whitlock, J.R., 2017. Posterior parietal cortex. Curr. Biol. https://doi.org/10.1016/j.cub.2017.06.007

Williams, D.S., Detre, J.A., Leigh, J.S., Koretsky, A.P., 1992. Magnetic resonance imaging of perfusion using spin inversion of arterial water. Proc. Natl. Acad. Sci. 89, 212–216. https://doi.org/10.1073/pnas.89.1.212

Wilson, J.T.L., Pettigrew, L.E.L., Teasdale, G.M., 1998. Structured interviews for

the glasgow outcome scale and the extended glasgow outcome scale: Guidelines for their use. J. Neurotrauma 15, 573–580. https://doi.org/10.1089/neu.1998.15.573

Wilson, L., Horton, L., Kunzmann, K., Sahakian, B.J., Newcombe, V.F.J., Stamatakis, E.A., Von Steinbuechel, N., Cunitz, K., Covic, A., Maas, A., Van Praag, D., Menon, D., 2021. Understanding the relationship between cognitive performance and function in daily life after traumatic brain injury. J. Neurol. Neurosurg. Psychiatry 92, 407–417. https://doi.org/10.1136/JNNP-2020-324492

Woolrich, M.W., Baker, A., Luckhoo, H., Mohseni, H., Barnes, G., Brookes, M., Rezek, L., 2013. Dynamic state allocation for MEG source reconstruction. Neuroimage 77, 77–92. https://doi.org/10.1016/j.neuroimage.2013.03.036

Zeeman, P., 1897. VII. Doublets and triplets in the spectrum produced by external magnetic forces. London, Edinburgh, Dublin Philos. Mag. J. Sci. 44, 55–60. https://doi.org/10.1080/14786449708621028

Zetter, R., Iivanainen, J., Parkkonen, L., 2019. Optical Co-registration of MRI and On-scalp MEG. Sci. Reports 2019 91 9, 1–9. https://doi.org/10.1038/s41598-019-41763-4

Zhang, J., Emami, Z., Safar, K., McCunn, P., Richardson, J.D., Rhind, S.G., da Costa, L., Jetly, R., Dunkley, B.T., 2021. Teasing apart trauma: neural oscillations differentiate individual cases of mild traumatic brain injury from post-traumatic stress disorder even when symptoms overlap. Transl. Psychiatry 2021 111 11, 1–14. https://doi.org/10.1038/s41398-021-01467-8

Zhang, J., Hadj-Moussa, H., Storey, K.B., 2016. Current progress of high-throughput microRNA differential expression analysis and random forest gene selection for model and non-model systems: an R implementation. J. Integr. Bioinform. 13, 306. https://doi.org/10.2390/biecoll-jib-2016-306

Zhang, J., Safar, K., Emami, Z., Ibrahim, G.M., Scratch, S.E., da Costa, L., Dunkley, B.T., 2020. Local and large-scale beta oscillatory dysfunction in males with mild traumatic brain injury. J. Neurophysiol. 124, 1948–1958. https://doi.org/10.1152/jn.00333.2020

Zich, C., Quinn, A.J., Mardell, L.C., Ward, N.S., Bestmann, S., 2020. Dissecting Transient Burst Events. Trends Cogn. Sci. https://doi.org/10.1016/j.tics.2020.07.004