Adam Baker

St John's College
Centre for Doctoral Training in Healthcare Innovation
Oxford Centre for Human Brain Activity

DPhil Thesis

Department of Engineering Science

Supervisors

Dr Mark Woolrich
Dr Penny Probert Smith

Temporal dynamics of resting state brain connectivity as revealed by magnetoencephalography

This thesis is submitted to the Department of Engineering Science, University of Oxford, in partial fulfilment of the requirements for the degree of Doctor of Philosophy
Submitted March 2014

Temporal dynamics of resting state brain connectivity as revealed by magnetoencephalography

Adam Baker
St John's College

Abstract

Explorations into the organisation of spontaneous activity within the brain have demonstrated the existence of networks of temporally correlated activity, consisting of brain areas that share similar cognitive or sensory functions. These so-called resting state networks (RSNs) emerge spontaneously during rest and disappear in response to overt stimuli or cognitive demands. In recent years, the study of RSNs has emerged as a valuable tool for probing brain function, both in the healthy brain and in disorders such as schizophrenia, Alzheimer's disease and Parkinson's disease. However, analyses of these networks have so far been limited, in part due to assumptions that the patterns of neuronal activity that underlie these networks remain constant over time. Moreover, the majority of RSN studies have used functional magnetic resonance imaging (fMRI), in which slow fluctuations in the level of oxygen in the blood are used as a proxy for the activity within a given brain region.

In this thesis we develop the use of magnetoencephalography (MEG) to study resting state functional connectivity. Unlike fMRI, MEG provides a direct measure of neuronal activity and can provide novel insights into the temporal dynamics that underlie resting state activity. In particular, we focus on the application of non-stationary analysis methods, which are able to capture fast temporal changes in activity. We first develop a framework for preprocessing MEG data and measuring interactions within different RSNs (**Chapter 3**). We then extend this framework to assess temporal variability in resting state functional connectivity by applying time-varying measures of interactions and show that within-network functional connectivity is underpinned by non-stationary temporal dynamics (**Chapter 4**). Finally we develop a data driven approach based on a hidden Markov model for inferring short lived connectivity states from resting state and task data (**Chapter 5**). By applying this approach to data from multiple subjects we reveal transient states that capture short lived patterns of neuronal activity (**Chapter 6**).

This thesis is submitted for the degree of Doctor of Philosophy at the University of Oxford.

Hilary Term 2014

Acknowledgements

This thesis would not have been possible without help, support, and data from many people and organisations. In particular, I would like to thank:

Dr Mark Woolrich for his outstanding supervision and for being an enthusiastic, insightful and affable mentor throughout my DPhil.

Dr Penny Probert Smith for providing excellent co-supervison and for maintaining a keen interest in my work.

Dr Matthew Brookes for his excellent collaboration and provision of datasets used in this thesis.

Professors Steve Smith and Tim Behrens for their valuable insights and discussions.

My colleagues at OHBA for creating an enjoyable working environment throughout my DPhil, in particular, Henry Luckhoo, Louisa Stacey, Tiina Parviainen, Hamid Mohseni, Sam Harrison, Diego Vidaurre and Giles Colclough.

The academic and support staff at OHBA, including Prof. Kia Nobre, Dr Sven Braeutigam, Dr Jennifer Swettenham, Dr Anling Rao and Judith Ponsford.

The Centre for Doctoral Training in Healthcare Innovation for funding this DPhil under the RCUK Digital Economy Programme and for providing training and support during the first year of my DPhil and beyond.

St John's College for providing me with generous funding, accommodation and a community throughout my undergraduate and postgraduate degrees.

The Department of Engineering Science and the Department of Psychiatry for providing administrative and financial support.

The Organisation for Human Brain Mapping for generous travel awards.

Rebecca Findlay for her unending support, positive outlook, and proof-reading skills.

My family, in particular my parents Amanda and Keith, for their love and support and everything else they have given me.

Contents

List of Figures

CHAPTER

1

Introduction

1.1 Brain activity during rest

In the absence of any explicit physical or cognitive task, an awake person is said to be resting. In this *resting state*, the brain is far from idle, but instead appears to exhibit spatio-temporally organised fluctuations in neural activity. These fluctuations emerge spontaneously during the resting state and disappear in response to overt stimuli or cognitive demands.

The idea that the brain is constantly active is by no means a new one. On introducing the electroencephalogram (EEG) in 1929, Hans Berger rhetorically asked “Is it possible to demonstrate the influence of intellectual work upon the human electroencephalogram, insofar as it has been reported here?”, to which he concluded “Of course, one should not at first entertain too high hopes with regard to this, because mental work... adds only a small increment to the cortical work which is going on continuously and not only in the waking state.” (Berger, 1929). Despite these early observations, for a long time the biological significance of spontaneous activity went largely unexplored. This was in part due to a tendency within the neuroscience community to consider the brain as primarily reflexive, responding to incoming stimuli and processing them in a bottom-up manner (Raichle, 2010). Under this hypothesis, spontaneous activity was assumed to correspond to noise resulting from random fluctuations that have no computational relevance. This classical view of the brain as a passive analyser of environmental stimuli is evident from decades of "activation" studies in which the response of the brain to a particular cognitive demand is compared to a control condition or baseline. In these studies the assumption is that ongoing activity is sufficiently random that it may be averaged out leaving only the activity evoked by the task. As a result a large proportion of published work has focussed on identifying brain areas in which experimental manipulation results in increased activity relative to the baseline condition. While such studies have proved hugely successful in elucidating the neural mechanisms associated with controllable cognitive and sensory tasks, one must acknowledge that this approach will only be able to provide a limited perspective on brain function. The limitation of this bias towards evoked activity is perhaps most readily apparent when considering that the increase in energy required by the brain to perform a task amounts to only 5% relative to that needed for ongoing activity (Raichle, 2010).

More recently, an alternative model of the brain has emerged in which the brain does not merely respond to external stimuli, but continuously generates predictions about forthcoming events (Engel et al., 2001). When viewed in this way, spontaneous activity is not considered as random noise, but rather as coordinated fluctuations that encode active representations that are modulated, rather than determined by, sensory information (Deco & Corbetta, 2011).

Elucidating the role of ongoing brain activity is particularly challenging, since its unconstrained nature makes it difficult to manipulate using traditional experimental paradigms. However, the observation from positron emission tomography (PET) studies that certain regions experience decreased activity during a range of different tasks led to the idea that ongoing activity represents a distinct cognitive state that may be deactivated by shifting the brain's focus from internal processes to the external environment (Nyberg et al., 1996; Shulman et al., 1997). This set of regions were later confirmed using functional magnetic resonance imaging (fMRI) and labelled the *default mode network* (DMN) (Greicius et al., 2003; Raichle et al., 2001). Around the same time, it was discovered that spontaneous activity as measured by the blood oxygenation level dependent (BOLD) response is organised into patterns of coherent activity (Biswal et al., 1995; Fox et al., 2005; Beckmann et al., 2005). These consistent patterns of co-activation have become known as *intrinsic connectivity networks* or *resting state networks (RSNs).* Notably, these networks strongly overlap with regions that are functionally related. For instance, the sensorimotor network comprises the set of regions that are activated during a motor task.

1.1.1 Functional significance of resting state activity

The biological significance of resting state networks remains poorly understood. Several studies have suggested a link between resting state activity and high level cognitive processes such as mind-wandering, autobiographic memory and moral reasoning (Buckner et al., 2008; Christoff et al., 2009). However, this theory has been challenged by findings that these patterns of coactivation persist during sleep (Fukunaga et al., 2006) and may be observed in monkeys and rats (Vincent et al., 2007; Lu et al., 2012). The fact that patterns of covariation appear to encompass regions that are simultaneously activated by goal-directed tasks suggests that resting state activity may serve to regulate networks that might be relevant for an impending task (Engel et al., 2013). Others have suggested that spontaneous activity might relate to ongoing internal explorations, in which the brain generates predictions about the optimal network configurations for future inputs (Deco et al., 2011).

Resting state functional connectivity strongly overlaps with the structural organisation of the brain. Studies using diffusion MRI have revealed axonal connections between regions that are functionally connected, suggesting that functional connectivity is mediated through white matter fibres (Sporns et al., 2000; Hagmann et al., 2008; Greicius et al., 2009). However, while structural connectivity appears to provide a substrate for functional connectivity, it does not fully explain it; structural connectivity appears to only account for about half of the variance in functional connectivity and some pairs of regions with no direct structural connection nonetheless show strong functional connectivity (Honey et al., 2009).

1.1.2 Resting state activity and disease

Many neurological and psychiatric diseases have been shown to be associated with disrupted brain functional and/or structural connectivity (Greicius, 2008). There is evidence to suggest that this holds for disorders including Alzheimer's disease, Parkinson's disease, schizophrenia, depression, autism, multiple sclerosis, and stroke (Engel et al., 2013).

Alzheimer's disease (AD) has been shown to be associated with abnormal functional connectivity within the default mode network (Greicius et al., 2004; Buckner et al., 2009). In particular it has been shown that AD patients have decreased correlations between the prefrontal and parietal lobes, and increased positive correlations within the prefrontal, parietal and occipital lobes, compared with healthy controls (Wang et al., 2007).

The principal pathology underlying Parkinson's disease (PD) is the dysfunctional activity of neural circuits within the basal ganglia, a region of the brain strongly associated with motor control. In particular, the symptoms of akinesia (inability to initiate movement) have been linked to excessive beta band synchrony within basal ganglia cortical loops, which prevents the normal desynchronisation required for the initiation of movement (Brown, 2003; Hammond et al., 2007; Schnitzler & Gross, 2005).

There is also some evidence to suggest that functional connectivity is impaired in schizophrenia (Buckner et al., 2008), with some studies reporting widespread functional disconnectivity in resting state fMRI (Liang et al., 2006).

In summary, functional connectivity appears to play an important role in coordinating healthy brain function and its dysfunction is closely linked with a wide range of neurological and psychiatric disorders. The ability to better characterise functional connectivity will likely help us to better understand the mechanisms that underlie these disorders.

1.2 Resting state haemodynamic fluctuations

With the advent of functional neuroimaging modalities such as fMRI and PET it became possible to study the whole brain structure of spontaneous activity. However, the first evidence that segregated brain areas are functionally related came entirely by accident. While attempting to identify responses to a finger tapping experiment, Biswal and colleagues were stymied by "noise" with a characteristic 1/f frequency profile that existed even during rest (Biswal, 2012). To characterise this low frequency signal, they correlated the time series from a voxel in the sensorimotor cortex with every other voxel in the brain and found that the strongest correlations existed between the left and right sensorimotor cortices (Biswal et al., 1995). Subsequent studies extended this correlation analysis to identify resting state functional connectivity in other functionally related brain regions including the visual cortex (Hampson et al., 2004), language areas (Hampson et al., 2002) and anterior and posterior cingulate cortex (Greicius et al., 2003).

These findings were initially disconcerting to the fMRI community. After all, most experimental paradigms sought to identify task-related signal changes or *activations* from a baseline. If spontaneous activity made up the majority of the measured signal then this baseline becomes more troublesome to determine. As a result, the study of task-independent activity quickly grew, both as a means to identify a baseline for comparing evoked responses, and in its own right in order to elucidate the functional role of this ongoing activity. Raichle et al. (2001) sought to identify this baseline of activity in terms of the oxygen extraction fraction (OEF) as measured using PET and further demonstrated the existence of a *default mode* corresponding to ongoing activity that persists during the resting state and is suspended during task-driven behaviour. In particular they revealed that certain brain regions including the posterior cingulate cortex (PCC) and anterior cingulate cortex (ACC) show greater activity during rest than in task. Greicius et al. (2003) demonstrated the existence of a default mode network (DMN) by investigating the BOLD signal correlations between these two regions and the rest of the brain. A later study by Fox et al. (2005) confirmed these findings, and demonstrated the existence of a further network involving frontal and parietal regions whose activity increases during task and anti-correlates with the DMN.

Up until this point, studies of resting state functional connectivity had been performed using seed based correlation, by computing correlations between the time course at a predetermined seed voxel and those at all other voxels in the brain.

However this approach does not provide a particularly parsimonious description of whole brain functional connectivity and the resulting correlations are strongly dependent on the chosen seed region. More recently, data driven approaches have emerged that do not require a predetermined region of interest. Most notable is the emergence of independent component analysis (ICA), which decomposes a mixture of signals into a set of components that are maximally independent in the spatial or temporal dimension (for a more complete introduction to ICA, see Section 3.2.2). Regardless of whether independence is maximised in the spatial or temporal dimension, ICA has been shown to identify brain areas whose time courses share common fluctuations. Beckmann et al. (2005) used spatial ICA to decompose whole brain BOLD recordings concatenated temporally over ten subjects, into 23 independent spatial maps. Of these maps, 8 were shown to correspond to resting state networks including the default mode network, lateralised fronto-parietal networks, sensorimotor, visual and auditory networks. Damoiseaux et al. (2006) further showed that these RSNs are highly reproducible over multiple recording sessions and subjects and correspond to BOLD signal changes comparable with those measured during task-related experiments. Smith et al. (2009) also demonstrated that the networks revealed by spatial ICA show close correspondence with those elicited during task activation by comparing the RSN spatial maps with those from the BrainMap functional imaging database. Ten spatial maps from this study are reproduced in Figure 1.1

1.3 Electrophysiological correlates of resting state networks

While fMRI has been a vital tool for characterising resting state networks, it provides an indirect measure of neuronal activity via the haemodynamic response.

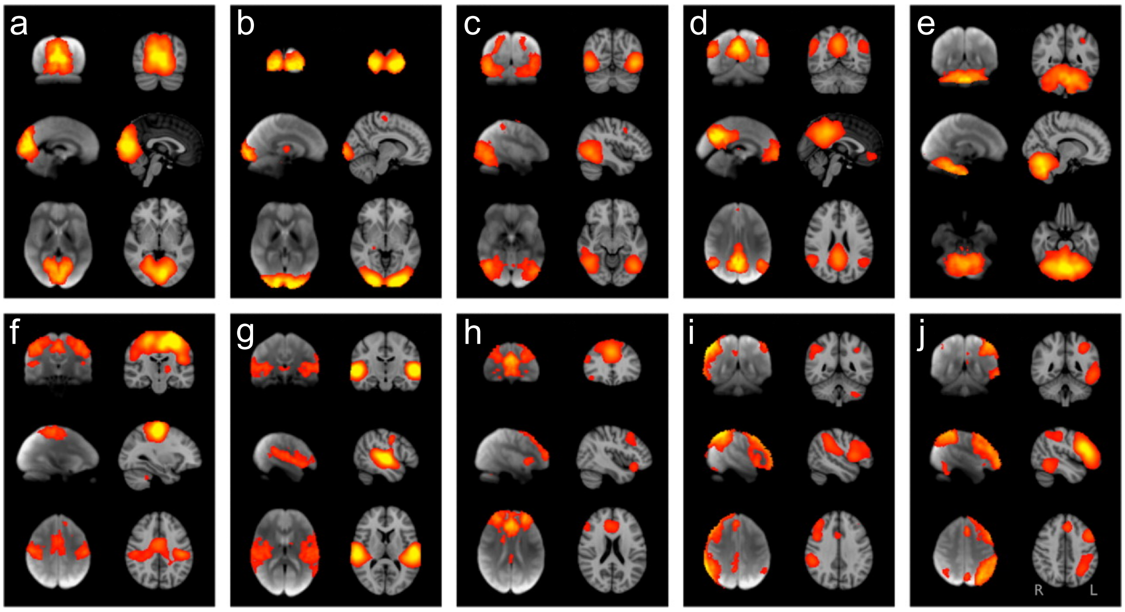

Figure 1.1: Ten RSNs identified from a spatial ICA decomposition of resting state BOLD data from a group of 36 subjects (left column in each panel) and corresponding task activations from the BrainMap database (right column in each panel). Each panel represents a different RSN: (a) medial visual areas, (b) occipital visual areas, (c) lateral visual areas, (d) default mode network, (e) cerebellum, (f) sensorimotor network, (g) auditory network, (h) executive control network, (i) right and (j) left lateralised frontoparietal networks. Figure reproduced from Smith et al. (2009).

A number of studies have sought to determine the electrical correlates of BOLD fluctuations. Logothetis et al. (2001) used simultaneously acquired fMRI and intracortical electrode recordings from anaesthetised monkeys to show that BOLD activity best correlates with local field potentials (LFPs). These signals capture the summation of all synaptic current within a local extracellular space and therefore reflect neural activity related to the input and local processing of a given area, as opposed to the spiking activity which represents the output of that area. While this result confirmed that BOLD signal changes reflect the underlying neural activity, the use of invasive electrode recordings is not suitable for measuring whole brain functional connectivity in humans (except in specific clinical settings).

Non-invasive electrophysiological recordings such as EEG and MEG have also been widely used to characterise resting state activity. In particular, combined EEG and fMRI acquisition has allowed a direct comparison of spontaneous haemodynamic and electrophysiological activity. Laufs et al. (2003) recorded simultaneous fMRI and EEG in the eyes open and eyes closed resting state. After filtering the recorded EEG into different frequency bands, they compared fluctuations in band-limited power (BLP) with the BOLD time courses at each voxel. They found that EEG power in the alpha band (8-12 Hz) correlated negatively with BOLD in lateral frontal and parietal cortices (regions typically associated with attentional processes) and that power in the 17-23 Hz range of the beta band correlated positively with BOLD in regions that make up the default mode network. Mantini et al. (2007) adopted this approach to study the relationship between band limited EEG power and BOLD RSN time courses identified using ICA from concurrently recorded fMRI. They found that each RSN was associated with variation in EEG power across multiple different frequency bands. No single frequency band was associated with the RSN time courses, suggesting that RSNs are not underpinned by one specific oscillation but rather a combination of different rhythms. In agreement with Laufs et al. (2003) they found that those networks typically associated with sensory or attentional processes correlate negatively with EEG power, while those associated with the default mode and self-referential processes correlated positively (Mantini et al., 2007).

Concurrent EEG-fMRI has proved a useful tool for investigating the electrophysiological correlates of BOLD RSNs. However, accurate reconstruction of the sources that generate a measured EEG at the scalp is challenging due to the inhomogeneous conductivity profile of the head. MEG on the other hand is better able to reconstruct the sources of activity within the brain, since the propagation of the magnetic field is not affected to the same extent by the tissue properties and geometry of the head. Resting state MEG studies have demonstrated that patterns of functional connectivity similar to those observed with fMRI exist in electrophysiological recordings. Correlation in band-limited power (or amplitude) has so far proved most fruitful in recapitulating this functional connectivity (Brookes et al., 2011b,a; de Pasquale et al., 2010; Liu et al., 2010).

As with fMRI, BLP correlations were first observed between the left and right hemispheres (Nikouline et al., 2001). This study focused on beta band power correlations. A decade later, Liu et al. (2010) extended this work to characterise BLP correlations across a range of different frequency bands in both the eyes open and eyes closed resting state, and also in sleep. They found that resting state MEG recordings exhibit fluctuations in BLP on a time scale similar to that of BOLD (<0.1 Hz) and that the BLP correlation structure across different sensors was broadly the same over different frequency bands. However, they found that interregional BLP correlations were strongest in the beta bands whereas for frequencies in the gamma band this correlation decreased with distance. This study was performed at the sensor level and thus was not able to relate patterns of BLP correlation to the underlying anatomy. However, subsequent studies have performed similar analyses at the source level, by applying source reconstruction techniques to estimate the neural activity across the brain prior to measuring functional connectivity. de Pasquale et al. (2010) computed BLP correlations between a number of seed voxels and voxels from the rest of the brain. The seed locations used corresponded to nodes of the default mode network (DMN) and dorsal attention network (DAN) defined from fMRI. They found patterns of functional connectivity that were spatially similar to those found from fMRI. Interestingly, they also showed evidence that MEG BLP is highly non-stationary, and that interhemispheric functional connectivity is increased if this non-stationarity is taken into account. In agreement with other studies, they found the strongest correlations in the theta, alpha and beta bands. Although this study was able to show that MEG RSNs exhibit similar spatial topographies to BOLD RSNs, the use of seed-based correlation necessitated that regions of interest be specific *a priori*.

Recently, Brookes et al. (2011b) have shown that resting state networks may be identified *independently* by applying temporal ICA to band-limited amplitude envelopes concatenated over multiple subjects. They found eight MEG counterparts of BOLD RSNs, either within the alpha or beta frequency bands. The spatial maps associated with these RSNs showed a strong correspondence with those from fMRI (as shown in Figure 1.2), supporting the hypothesis that resting state networks have a neuronal basis.

1.4 Temporal dynamics in resting state networks

Until very recently, resting state functional connectivity has been characterised using methods that assess temporal correlations over the duration of the recording, typically several minutes. Such approaches necessarily assume that the connectivity between regions is temporally stationary (in the sense that the mean and/or variance of interregional coupling strength does not change over time). While this assumption of temporal stationarity greatly simplifies connectivity analyses in terms of computation and visualisation, the notion that interactions between distinct brain regions are temporally stationary over several minutes represents a gross oversimplification of the complex spatio-temporal properties of spontaneous brain activity (Hutchison et al., 2013; Deco et al., 2011). Temporal changes in activity and functional connectivity are a well established phenomenon of spontaneous as evidenced by many years of theoretical (Friston, 1997), computational

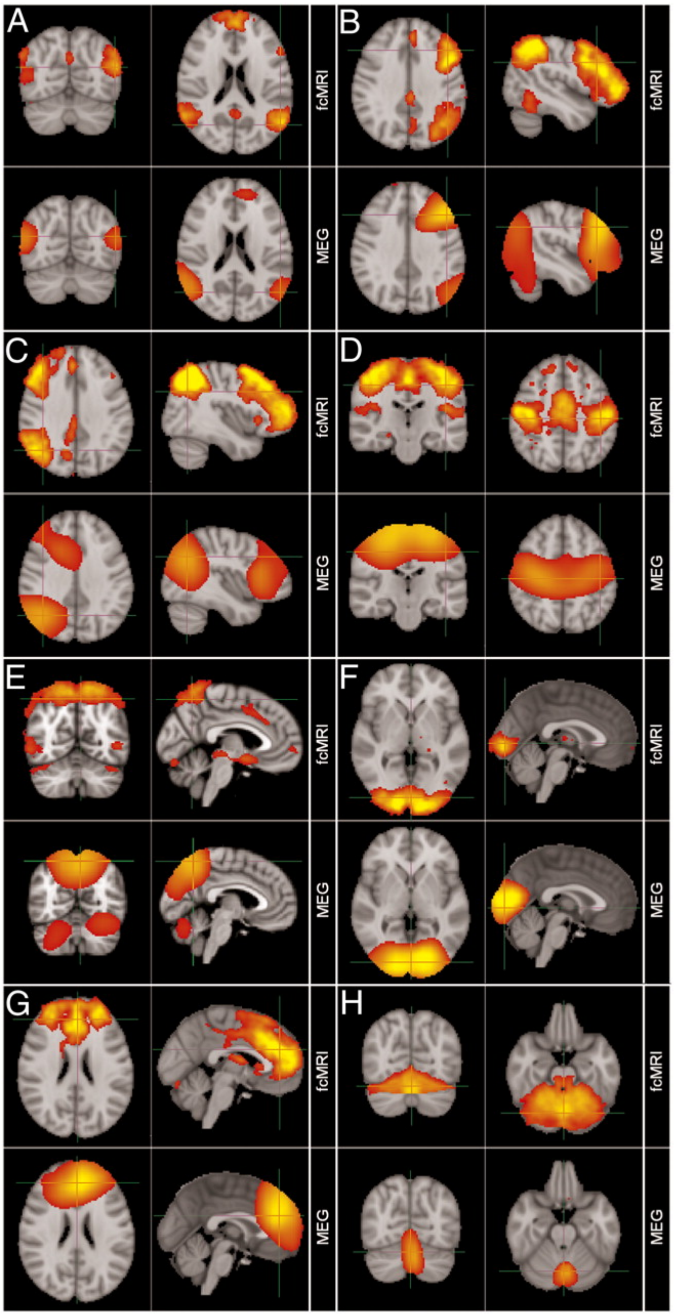

Figure 1.2: Comparison of 8 RSNs identified by applying spatial ICA to fMRI data (top row in each panel) and temporal ICA to MEG BLP (bottom row in each panel). The eight RSNs represent (A) default mode network, (B) left lateralised frontoparietal networks, (C) right lateralised frontoparietal networks, (D) sensorimotor network, (E) medial parietal regions, (F) visual network, (G) frontal lobes including anterior cingulate cortex, (H) cerebellum. Figure reproduced from Brookes et al. (2011b).

(Deco et al., 2011; Honey et al., 2007) and empirical (Lehmann et al., 1998; Bassett et al., 2006) studies. This evidence suggests that spontaneous brain activity is underpinned by complex spatiotemporal dynamics that may be better characterised using time-varying measures of interactions.

These arguments suggest that a better understanding of spontaneous brain activity may be achieved if we relax the assumption that patterns of functional connectivity are stable over time. Non-stationary analyses, where functional connectivity is able to vary throughout the duration of the scan, may provide new insights into the temporal dynamics that underlie resting state connectivity. Indeed, the last few years has seen a number of studies emerge addressing this issue using fMRI (Chang & Glover, 2010; Kiviniemi et al., 2011; Allen et al., 2012; Liu & Duyn, 2013), EEG (Britz et al., 2010; Musso et al., 2010; Yuan et al., 2012; Chang et al., 2013) and MEG (Brookes et al., 2014; de Pasquale et al., 2010).

1.5 Summary

The study of ongoing brain activity and connectivity is essential for advancing our understanding of brain function in health and disease. Central to this goal is the development of novel methodologies for unlocking the complex dynamics of neuronal interactions. MEG will likely play a key role in this, as it provides a direct measure of neuronal activity at high temporal resolution. This modality has already shown great promise for characterising the spatial and spectral composition of functional brain networks. However, to better understand the temporal dynamics that underlie functional connectivity we must develop new techniques that depart from the assumption that brain connectivity is stationary in time.

This thesis contributes to a growing line of research into characterising the temporal dynamics underlying resting state functional connectivity using MEG. In particular, a number of novel methods are developed which form the key contributions of this work. **Chapter 2** provides a background to how the magnetoencephalogram is generated, measured and used to estimate sources of neuronal activity. In **Chapter 3** we present a framework for preprocessing MEG data, with a specific focus on removing sources of interference that may confound brain connectivity measures. This preprocessing approach is then used in conjunction with a pipeline for measuring functional connectivity to identify resting state networks from MEG data. In **Chapter 4** this pipeline is further developed to characterise *non-stationary* functional connectivity between nodes of two resting state networks. In particular, we develop an approach to test if temporal variability in functional connectivity is significantly greater than that which would be expected if the coupling strength was stationary. Using this approach we show that within-network functional connectivity is significantly non-stationary and determine the minimum time scales at which significant temporal variability may be identified. In **Chapter 5** we present a novel methodology for assessing non-stationary functional connectivity based on a hidden Markov model (HMM). This approach provides a data driven decomposition of MEG data into a number of discrete states that represent transient patterns of spontaneous activity. We validate this approach, both on simulated data with known spatial and temporal characteristics, and on real MEG data for a simple visual task. Using these simple datasets we demonstrate that the method is able to identify short lived changes in activity. The HMM approach is then applied to resting state data from single subjects and identifies short-lived states with spatial patterns of coactivation that strongly resemble RSNs. In **Chapter 6**, this approach is applied to group-concatenated data in order to identify states that are consistent within and across subjects. We further characterised these transient states in terms of their temporal properties, and assess cross-network interactions.

Finally, in **Chapter 7** we conclude this thesis, and consider how this work might be developed in the future to further our understanding of spontaneous brain activity at rest and during tasks.

CHAPTER 2

Magnetoencephalography

MEG provides a direct measure of neuronal activity by recording the magnetic field generated by synchronous activity within a local population of several thousands of neurons. However, the electrophysiology underlying the generation of a detectable signal is complex and one must take care to understand exactly which aspects of neuronal activity we are able to measure. In the following sections we provide a brief overview of how the magnetoencephalogram is generated.

2.1 Electrophysiological basis of MEG signals

2.1.1 Single neuron electrophysiology

The principal mode of neural communication is via the generation of electrical impulses called action potentials. Generation of an action potential within a single neuron is mediated by the opening and closing of voltage dependent ion channels in the neuronal membrane. Neurons maintain a voltage gradient across their membrane by controlling different intracellular and extracellular concentrations of electrically charged ions by means of ion pumps. At rest, sodium, chloride and calcium ions are held at higher extracellular concentrations whereas potassium is held at a higher intracellular concentration. These different electrochemical concentrations are balanced by the membrane's resting potential. Synaptic inputs to the neuron cause the membrane potential to either rise (depolarisation) or fall (hyperpolarisation). If there is sufficient depolarisation to exceed the membrane's threshold potential, sodium channels begin to open, allowing an inward flow of sodium ions resulting in further depolarisation. This feedback process causes a rapid influx of sodium ions and a resulting action potential to be generated which propagates along the neuronal axon, terminating at the synapse. This process is followed by a closing of the sodium channels and opening of the potassium channels which causes a repolarisation of the membrane.

Whether or not an action potential is generated depends on whether depolarisation of the membrane is sufficient to reach the threshold potential. Changes in membrane potential are largely caused by postsynaptic potentials (PSPs) that occur at the dendrites. These potentials are mediated by either excitatory or inhibitory neurotransmitters that act to open or close specific ligand-gated ion channels in the membrane. Excitatory PSPs (EPSPs) are generated primarily by glutamate and inhibitory PSPs (IPSPs) are primarily generated by the neurotransmitter GABA. The changes in potential caused by neurotransmitter release at a single synapse are far too small to trigger an action potential. However, a neuron will receive inputs from hundreds to thousands of afferent neurons and the summation of these IPSPs and EPSPs act to modulate the membrane potential, thus dictating whether or not an action potential will be generated. This summation can be both spatial and temporal and is thus dictated by the number of afferents and their firing rate. This summation is illustrated in Figure 2.1.

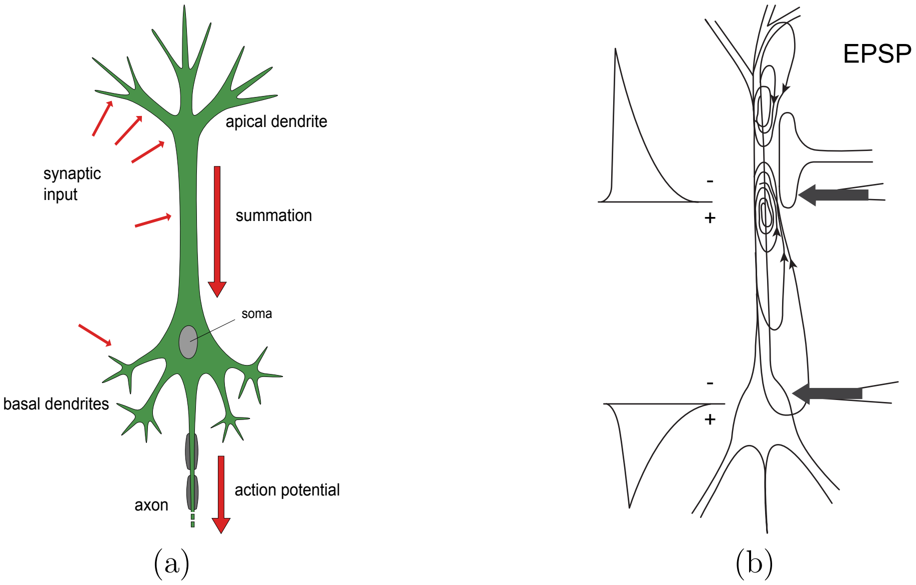

(a) (b)

Figure 2.1: (a) Structure of a single pyramidal neuron. The neuron receives synaptic input at its apical and basal dendrites. Changes in membrane potential caused by excitatory and inhibitory postsynaptic potentials sum as they propagate towards the soma. If the change in membrane potential reaches the threshold potential, an action potential will be initiated which travels along the neuron's axon. (b) An excitatory postsynaptic potential caused by an excitatory synapse at an apical dendrite. Depolarisation of the membrane is caused by the inward flow of positive ions resulting in an active sink at the level of the synapse and distributed passive sources along the soma-dendritic membrane (reproduced from Hansen et al., 2010).

2.1.2 Generation of macroscopic magnetic fields

The intracellular currents that arise from changes in the membrane potential underlie the magnetic field that is measured at the scalp. The generation of currents within the dendrites may be modelled as a current dipole (shown in Figure 2.1) which, in accordance with the right hand rule of electromagnetism, gives rise to a magnetic field around its axis. However, the field produced by a single PSP is far too small to be detectable. It has been estimated that the current dipole moment of a single neuron in the cortex is in the order of 0.2 pAm (Okada, 1981). Since the dipole moment required to produce a measurable magnetic field is of the order of 10 nAm (Hämäläinen et al., 1993), the number of neurons that must be synchronously active may be calculated as approximately 50,000 (Hansen et al., 2010). Thus the measured field measured by MEG corresponds to a current-dipole that represents a population of active neurons. It is worth noting that PSPs, as opposed to action potentials, are the primary generators of the MEG signal. This is in part due to the fact that the latter are best described by quadrupolar sources, whose field strength decays more rapidly with distance (Hämäläinen et al., 1993). Furthermore, action potentials have very short duration, of the order of 1 ms, and there would need to be a very high level of synchrony to generate a measurable field.

However, it is not sufficient for this population just to be synchronously active; these neurons must also be spatially aligned such that their individual dipole moments reinforce one another. This requires that the dendrites be aligned in an *open field* configuration, in which their dipoles have similar orientations (Lorente de No, 1947). This is in contrast to a *closed field* configuration in which orientations are randomly distributed such that their dipoles cancel each other out (Figure 2.2). Thus only those neurons that are arranged in such a structure will generate a macroscopic field that may be measured by MEG. Fortunately, within the cortex, pyramidal neurons are arranged in palisades with their apical dendrites extending from layers III and V towards the cortical surface (Figure 2.2).

The architecture of the cortex is such that a measurable field will be generated if there is sufficient synchronous activity. However, in order to determine whether this field will be recorded outside the skull, one must also take into account the macroscopic structure of the cortex. Since the folding of the cortical sheet results in the formation of gyri and sulci, some neuronal populations will have dendrites that are perpendicular to the skull while others will have dendrites that are parallel to the skull. These configurations will result in current dipoles that are oriented either radially or tangentially to the surface. It can be shown that, for the case

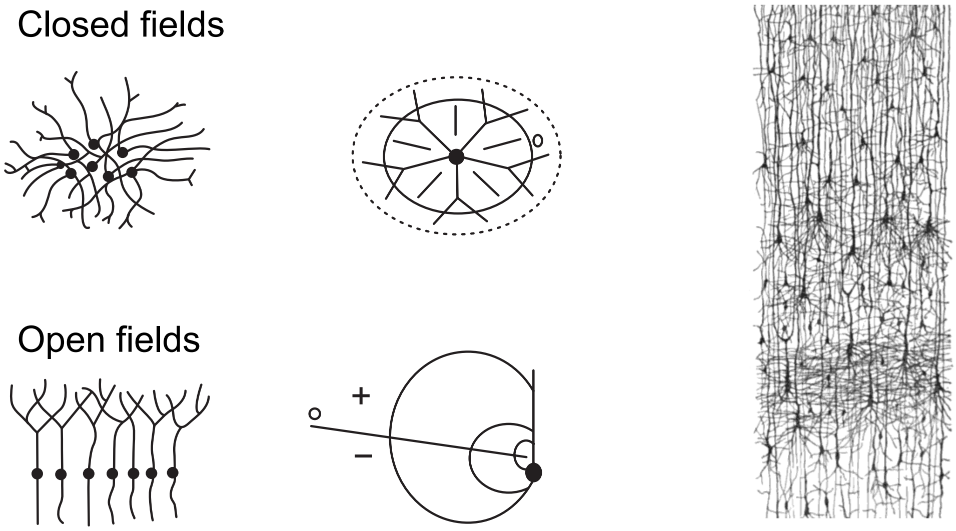

(a) (b)

Figure 2.2: (a) Examples of neuronal configurations that result in open and closed fields (adapted from Hansen et al., 2010). The dipole moments of neurons in a closed field arrangement act to cancel one another out, whereas those of neurons in an open field arrangement reinforce one another. (b) Section through the cortex drawn by Ramón y Cajal. Pyramidal neurons are arranged in palisades, resulting in an open field configuration.

of a spherical conductor, radial sources result in no field external to the surface (Sarvas, 1987). To the extent that the head may be approximated to a sphere, the implication is that only those current dipoles that have a tangential component will be measurable. In reality, the head is not perfectly symmetrical, so in some regions radial sources will generate a measurable field (Ahlfors et al., 2010). However, for most areas of the brain MEG will be sensitive primarily to sources that lie at the sulcal bank, and relatively blind to sources at the gyral crown and sulcal fundi. By comparison, the electrical fields measured by EEG are visible for both radial and tangential dipoles. These orientations are shown in Figure 2.3.

It is worth noting that MEG (and EEG) is relatively insensitive to subcortical sources. This is in part due to the fact that magnetic field decays rapidly with distance. Furthermore, the more complex cytoarchitecture of deep brain structures such as the thalamus means that the organisation of dendrites does not necessarily result in open fields (Attal & Schwartz, 2013).

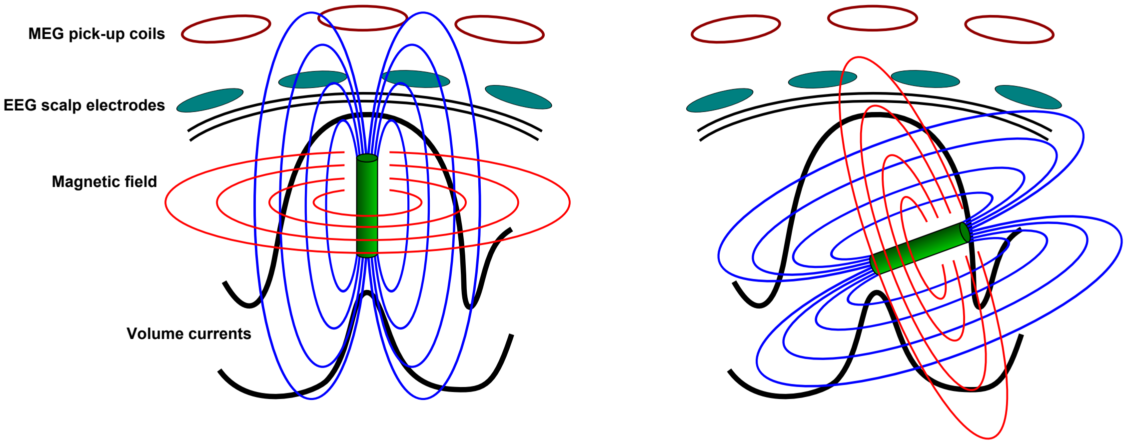

Figure 2.3: Schematic showing the generation of magnetic fields (red) and volume currents (blue) arising from a dipole in two different orientations. Only those orientations with a tangential component (right) will result in a measurable magnetic field outside of the head.

2.2 Measurement of the brain's magnetic field

The magnetic fields produced by the brain are tiny, with field strength of approximately 10 femtotesla. To put this into perspective, the strength of the Earth's magnetic field is around 30 microtesla, 8-9 orders of magnitude greater. Measurement of these fields requires highly sensitive detectors and magnetic shielding capable of reducing the impact of electromagnetic interference.

The first recording of the brain's magnetic field was performed by Cohen (1968), who used room temperature coils to measure the human alpha rhythm. However, it was the application of the recently invented superconducting quantum interference device (SQUID) that allowed detection of the magnetoencephalogram and demonstrated the validity of MEG as a viable neuroscientific tool (Cohen, 1972). These highly sensitive devices comprise a ring of superconducting material broken by a very thin layer of electrical insulation known as a Josephson junction. The superconducting ring acts as a perfect diamagnet. When exposed to an external magnetic field, a shielding current is induced, which in turn generates a magnetic field that exactly cancels the external field. The induced current depends on the external field and would thus provide a measurement of the field strength. However, this current is not observable, since any measurement would disrupt the superconducting loop. The introduction of the Josephson junction does not disrupt the flow of current since electrons can pass through the insulator by means of quantum tunnelling, but results in a flux-dependent resistance across the SQUID. This resistance allows the flux to be determined, by applying a constant bias current and measuring the voltage. SQUIDs are typically designed to be very small (less than 1 mm in diameter), since this improves their sensitivity. However, this means they have poor coupling to the magnetic field. For the purposes of MEG, this coupling is first enhanced using flux transformers. These consist of pickup coils that act to capture magnetic flux across a larger area and deliver it to the SQUID.

Multiple pickup coils may be used in one sensor, and various arrangements of these coils can be used to allow different aspects of the field to be measured. Magnetometers are the most basic configuration and comprise a single pickup coil which measures the field strength perpendicular to the coil. These sensors are sensitive to both nearby and distant sources and are thus able to detect fields from deeper in the brain at the expense of increased susceptibility to external interference (Plonsey & Robert, 1995). Gradiometers consist of a pickup coil and a compensation coil which are arranged in parallel to each other, but are wound in opposing directions. This design measures the spatial gradient of the field and it is thus less sensitive to distant sources, which manifest as a uniform field across both coils (therefore with a small gradient). The two coils that comprise a gradiometer may be arranged either to lie side-by-side in the same plane (planar gradiometer) or stacked one above the other (axial gradiometer). Planar gradiometers are maximally sensitive to sources directly beneath them whereas axial gradiometers are sensitive to deeper sources (though not as deep as magnetometers) as shown in Figure 2.4.

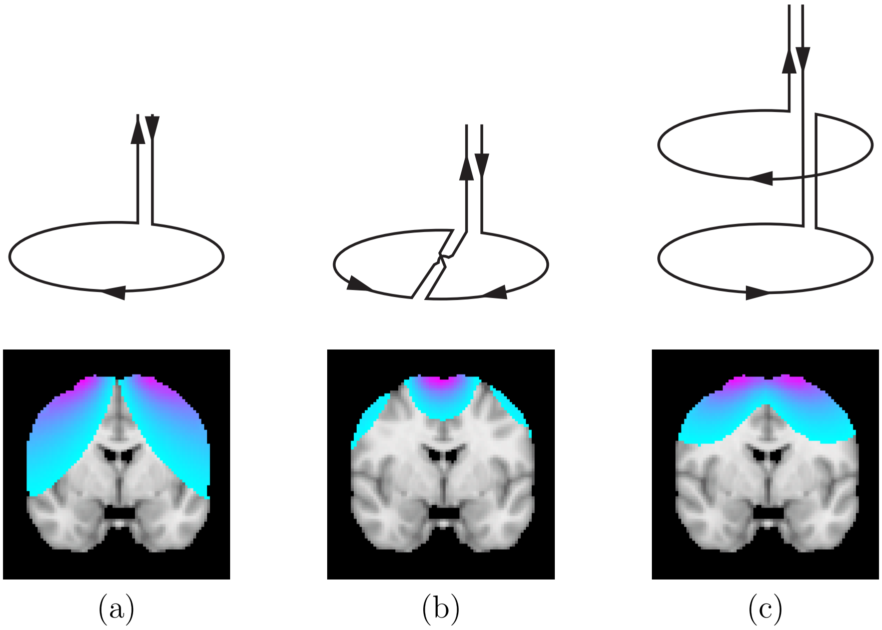

Figure 2.4: Examples of (a) magnetometer, (b) planar gradiometer, (c) axial gradiometer and their spatial sensitivity (shown as the magnitude of their leadfields at each location within the head). Planar gradiometers are maximally sensitive to sources directly beneath the coil, whereas axial gradiometers and magnetometers have similar sensitivity patterns and are sensitive to off-centre sources, but magnetometers are sensitive to deeper sources.

The first MEG measurements were single-channel recordings using only one SQUID sensor. In order to map activity across the whole head, the location of the sensor would be moved to different positions over the scalp and the same experiment repeated for each new location. This approach was time consuming and error prone, and precluded the study of neural activity that is not stimulus-locked. The first whole head device, containing arrays of multiple sensors arranged in a helmet surrounding the head, was introduced in 1992 (Ahonen et al., 1993). Modern devices now contain several hundred sensors and, in some cases, contain combinations of magnetometers and gradiometers allowing for simultaneous measurement of deep and superficial sources.

2.3 Sources of interference

As previously discussed, the magnetic field generated by the brain are several orders of magnitude weaker than that of the earth. Furthermore, the fields associated with sources of ambient electromagnetic noise such as power lines, traffic and cell phones, are all much stronger than the MEG signal. The field strength of some of these sources compared to the MEG signal are shown in Figure 2.5. In order for the MEG sensors to record neural activity in the presence of these sources of interference, advanced magnetic shielding is required. The primary means of protection comes from the use of a magnetically shielded room (MSR) which provides a low-reluctance path along the walls of the room for external fields to travel. Although the MSR is able to significantly minimise sources of external noise, it does not provide any protection from sources of interference from within the room. For this reason, equipment such as display projectors are placed outside the MSR, or else are made of non-magnetic materials. However, a final important class of interference is physiological noise originating from the subject themselves. Typical sources of physiological noise include the magnetocardiogram due to the heartbeat, muscle contractions, eye movements and blinks. These sources of interference are not readily excluded from the measured signal and must be removed offline using signal processing techniques. Such techniques are discussed later in **Chapter 3**.

2.4 MEG source reconstruction

The recorded MEG signal consists of a measurement of the magnetic field (or, in the case of gradiometers, the gradient of this field) at each sensor location. These measurements represent a weighted summation of all active sources within the brain. In many cases, this sensor level description of neuronal activity may be sufficient to answer a particular neuroscientific question, for instance by comparing evoked

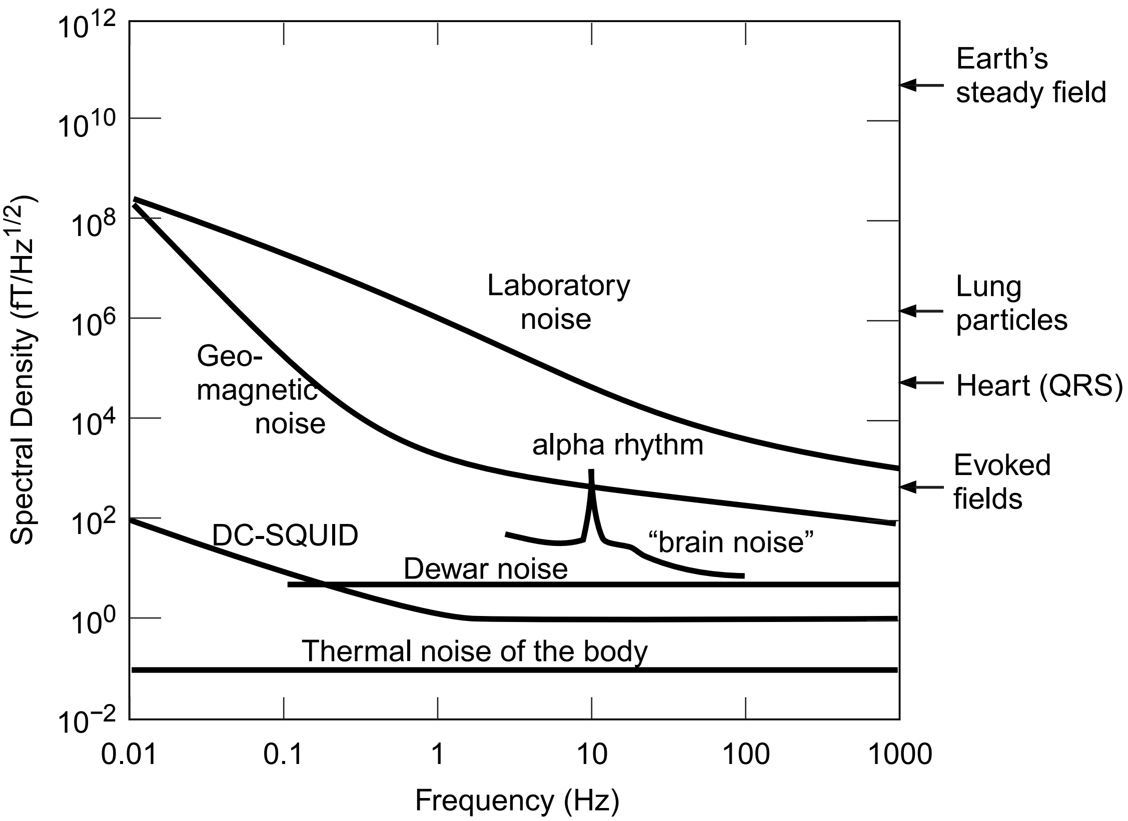

Figure 2.5: Magnetic field strength of various sources demonstrating that the MEG is a much weaker source than most types of interference (reproduced from Hansen et al., 2010).

responses between different experimental conditions or populations. Indeed, a large number of MEG (and EEG) studies operate exclusively at this level. However, if one is concerned with relating these measurements to the underlying anatomy, then a means of estimating the sources of neuronal activity is required. This is a two-fold problem. First, one requires a *forward model* that provides a complete description of how the magnetic field generated by a known neuronal current source propagates through the head and is measured by the sensor array. Such a model is governed by Maxwell's equations which may be solved analytically or numerically given realistic assumptions about the geometry and magnetic properties of the head, and the positions of the head and sensor array. However, for the purpose of estimating the underlying sources of activity one is faced with the reciprocal problem, that is, obtaining an estimate of the neuronal activity given the measured field. This is known as the *inverse problem* and is inherently ill-posed; any given distribution at the sensors could result from an infinite number of possible neuronal current distributions. In the following sections, both aspects of source reconstruction will be discussed.

2.4.1 Forward modelling

As discussed in Section 2.1.2, synchronous activity within a population of similarly aligned neurons results in a measurable magnetic field that can be approximated as a current dipole. Modelling patches of active cortex as a dipolar source in this way provides a convenient representation of cortical activity, since a current source of arbitrary extent is readily broken down into smaller subcomponents each represented by a current dipole (Baillet et al., 2001). The forward model describes how a current dipole at a particular location within the brain is related to the magnetic field at the sensor array. Formally, for a unit dipole at location in three dimensional space $\mathbf{q}$ the potentials $\mathbf{x_q}$ recorded at the N sensors are related to the (3×1) dipole moment $\mathbf{m(q)}$ by the following equation:

$$\mathbf{x_q} = \mathbf{H(q)m(q)} \tag{2.1}$$

where $\mathbf{H}$ is a $N \times 3$ lead field matrix that represents the physical properties of the medium through which the signal travels. Since this medium is linear, the contributions from all L active dipoles may be summed to give the $N \times 1$ potentials $\mathbf{x}$ recorded at the sensors resulting from all the neuronal activity:

$$\mathbf{x} = \sum_{i=1}^{L} \mathbf{H}(\mathbf{q}_i)\mathbf{m}(\mathbf{q}_i) + \mathbf{n} \tag{2.2}$$

where **n** is the measurement noise.

2.4.2 Lead field estimation

Estimating the lead fields for a particular location of interest requires a solution to the magnetic field propagation from a current dipole at that location to the sensor array. The relationship between electrical current and magnetic field is described by Maxwell's equations. Since the frequencies associated with electrophysiological signals (typically less than 1 kHz) are much slower than the speed at which electromagnetic waves propagate, a quasi-static approximation may be adopted. The relationship between a the current $\mathbf{J}(\mathbf{r}')$ at a location $\mathbf{r}'$ and the magnetic field $\mathbf{B}(\mathbf{r})$ at another location $\mathbf{r}$ is then given by the Biot-Savart law:

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi} \int \mathbf{J}(\mathbf{r}') \times \frac{\mathbf{r} - \mathbf{r}'}{\parallel \mathbf{r} - \mathbf{r}' \parallel^3} dv' \tag{2.3}$$

where dv' is the differential element of volume and μ_0 is the permitivity of free space (Baillet et al., 2001; Hämäläinen et al., 1993). Solutions to this equation exist in analytical form for simple head geometries such as a sphere. However, the problem is further complicated by the effects of volume conduction. In addition to the magnetic field induced by the current dipole, fields are also generated by *return currents* that arise due to movement of ionic species throughout the volume of the head. The distribution of these currents depends on the conductivity profile of the head, which must therefore be modelled in order to correctly estimate the measured field. The primary inhomogeneities in conductivity are between the different tissue layers of the white matter, cortex, skull and scalp. One solution is to model these different layers as concentric spheres. It may be shown that a set of concentric spheres results in the same MEG fields as a single homogeneous sphere, regardless of the radii or conductivity profile of the spheres (Sarvas, 1987). This attractive property, in addition to the fact that analytic solutions exist for this geometry, has meant that spherical head models have been widely used for MEG forward modelling.

A number of techniques have been developed that make use of these properties, while attempting to model more realistic head shapes. Huang et al. (1999) proposed using multiple overlapping spheres by fitting a sphere to each sensor individually. Nolte (2003) developed a model that approximates the head shape using a single shell built from a spherical harmonic basis set. Both these methods provide a more realistic fit to the head shape by accounting for local variations, while preserving the spherical assumption. These approximations have been shown to work reasonably well despite their assumptions of homogeneous conductivity and isotropy (Leahy et al., 1998). A number of more realistic models such as the boundary element method (BEM) may be applied if an accurate representation of the head's geometry is obtained (e.g. from structural imaging). Even more realistic models may be derived using the finite element method (FEM), which allows modelling of anisotropy (for instance along white matter tracts) and inhomogeneities in the conductivity of different layers (Baillet et al., 2001). Solving for the lead fields using these models require numerical solutions and long computation times.

2.4.3 Coregistration

Computation of the lead fields requires a model of the propagation of the magnetic field from the neuronal sources to the sensors. An important requirement is therefore that the relationship between the position of the head and the sensor array is known. Furthermore, in order to accurately model the head geometry, and to localise neural currents in the brain, one also requires a mapping between the head position within the scanner and the brain anatomy, which is typically obtained with structural MRI. Thus MEG source reconstruction involves determining a mapping between three coordinate systems: the MEG device coordinate system, the head coordinate system (within the MEG device) and the MRI device coordinate system.

Localisation of the head within the MEG helmet may be achieved using a number of electromagnetic head position indicator (HPI) coils that are attached to the subject's head prior to entering the scanner. By periodically energising these coils, their position within the MEG sensor array may be identified. This allows a rigid-body registration between the head and the MEG device coordinate systems, such that differences in the position of a subject are accounted for. The HPI coil locations may be recorded outside the scanner using a three-dimensional digitiser (e.g. the Polhemus Isotrack; Polhemus Corporation, Colchester, Vermont, USA). In addition to the positions, a number of anatomical landmarks, or fiducials, are recorded which may be independently identified on both the MRI and the actual head. The fiducials typically used are the nasion, and the left and right preauricular points. These points allow a rigid body registration to be made between the head coordinate system and the MRI device coordinate system. Additionally, a number of head shape points may be recorded which allows for a more accurate registration to the structural image. Since the HPI and fiducial points are defined by the digitiser, the head coordinate system provides a link between the MEG device coordinate system and the MRI space. An additional registration (generally affine or non-linear) is sometimes performed, between the subject's MRI and a standard brain coordinate system, such as the Montreal Neurological Institute (MNI) coordinate system. This allows results from different subjects to be defined in a common coordinate system such that averages or other group statistics may be performed.

2.5 The inverse problem

In Section 2.4.2 we discussed how MEG forward modelling allows the lead fields to be estimated. Armed with a model of how neuronal currents manifest at the sensor level, it is now possible to consider how the inverse problem might be solved. As discussed in Section 2.4, the inverse problem is ill-posed, since the observations at the sensor array may be explained by an infinite number of underlying source distributions. Solving the inverse problem therefore requires additional constraints to be imposed so as to identify those source distributions that are most biologically plausible. Typical constraints include assuming that only a few sources are simultaneously active, that sources are distributed smoothly and/or sparsely, or by constraining sources to be located only on the cortical surface. A number of different approaches to source reconstruction have been proposed and the most popular of these methods will now be discussed in turn.

2.5.1 Dipole fitting

One of the most basic source localisation techniques is to assume that the measured activity is produced from a small number of equivalent current dipoles. A single dipole may be modelled by six parameters, accounting for its location in three-dimensional space, and its magnitude along these three dimensions. In the case of a single modelled dipole, solving for these parameters becomes a well-posed problem that may be solved using optimisation techniques, since the number of observations is now much greater than the number of unknowns. Indeed, given that there are typically of the order of a few hundred sensors, the problem may be solved for several simultaneously active dipoles, although in practice this number is typically limited to no more that five (Hansen et al., 2010). Dipole fitting relies on two assumptions: firstly, that the sources are focal and secondly that the number of dipoles to be included in the model is known. Thus the technique is well suited to situations in which one expects only a few sources to be active, for instance in the case of sensory evoked potentials. However, the approach is not well suited for modelling more distributed cortical activity, or for whole brain connectivity analyses.

2.5.2 Distributed source imaging

Whereas dipole fitting assumes that sources are focal, a family of inverse solutions exist that model neuronal activity by a large number of distributed sources. In these approaches the brain is divided up, either into a regular three dimensional grid of voxels, or a set of vertices constrained to the cortical surface. In both cases, reasonable spatial resolution results in several thousand possible source locations and the model is therefore severely undetermined. Solving for the magnitudes of each dipole source therefore requires additional constraints to be imposed. Two approaches that are commonly used are the minimum current estimate (MCE) and minimum norm estimate (MNE), which minimise the total current (L1-norm) or the sum of squared currents (L2-norm), respectively (Hämäläinen & Ilmoniemi, 1994; Matsuura & Okabe, 1995). One drawback to these techniques is that they favour superficial sources over deeper sources. This feature arises due to the fact that sources farther away from the sensors must have greater magnitude in order to produce the same measured field as sources close to the surface, and will therefore be penalised more by the minimisation constraint. Furthermore, the L2-norm used in MNE produces smooth distributions, and as such is less well suited to reconstructing focal sources. Other distributed source imaging techniques formulate the inverse problem in a Bayesian framework. Solutions may then be found by assuming a particular prior on the source distribution (Wipf & Nagarajan, 2009; Friston et al., 2008). Due to the under-determined nature of distributed source models, further constraints are often used, such as constraining dipoles to lie on the cortical surface, and/or with dipole orientations perpendicular to this surface. These approaches help to reduce the number of degrees of freedom, but at the potential expense of excluding sources that do not conform to these assumptions (such as those from within subcortical regions) (Baillet et al., 2001).

2.5.3 Beamforming

Another widely adopted approach to source reconstruction is to use beamforming. Rather than formulating the inverse problem as a generative model which is then solved by assuming a set of constraints, beamformers estimate the activity independently for each location of interest by constructing a spatial filter that is optimally tuned for that location. In the same way that frequency filters pass components of a signal with the desired temporal features, spatial filters are designed to pass components of a signal from a specified spatial location, while attenuating those components from other locations (Van Veen & Buckley, 1988). This is achieved by selectively weighting the contribution of each sensor to the output of the beamformer (Hillebrand et al., 2005). Hence the mapping between the recorded MEG data and the estimated source activity is determined by a set of beamformer weights. Several methods exist for estimating of these weights, the most widely accepted being the linearly constrained minimum variance (LCMV) approach developed by Van Veen et al. (1997). In this approach, a weights matrix is found that minimises the variance at the filter output while passing the signal at the desired location with unity gain. A distributed representation of activity over the entirety of source space may be obtained by dividing the brain into a grid and estimating a set of weights and associated current estimates for each location in turn.

Beamforming offers a number of advantages over other source reconstruction techniques. Firstly, there is no requirement to assume a particular number of sources *a priori*, making them well suited for imaging distributed sources. Secondly, due to the fact that the spatial filters adaptively suppress activity from outside the location of interest, beamformers are able to reduce interference from other locations or non-neuronal artefacts. Thirdly, in contrast to minimum norm solutions which favour smooth current distributions, beamformers are able to reconstruct focal sources accurately (Barnes & Hillebrand, 2003), and are sensitive to both deep and superficial sources. However, beamformers also suffer from two major limitations. Firstly, there is an implicit assumption that all sources of activity are mutually uncorrelated, and violation of this assumption can result in severe errors in the source estimate (Van Veen et al., 1997). Secondly, it can be shown that sensitivity to noise increases with depth into the head (Van Veen et al., 1997). In practical use this limitation is easily overcome by applying one of a number of normalisation techniques (Huang et al., 2004). However, great care must be taken if one is to compute group statistics based on multiple separately beamformed sessions, since this normalisation provides a biased estimate of the true dipole variance (Luckhoo et al., 2014).

2.6 LCMV beamforming

Beamforming has been extensively used throughout this thesis, due to its advantages in reconstructing whole-brain oscillatory activity. In this section we provide a mathematical formulation of the LCMV beamformer and a discussion of relevant implementation decisions.

2.6.1 Mathematical formulation

From Section 2.4.1, the signals measured at the sensors are related to the underlying dipole moment by the forward model in equation Section 2.2. As described previously, the spatial filter relates the length T MEG data recorded at the N sensors $\mathbf{x}$ to the output by the $(N\times 3)$ weights matrix $\mathbf{W}$. Hence, the $(3\times T)$ output $\mathbf{y}$ from a spatial filter centred at location $\mathbf{q}_0$ is given by:

$$\mathbf{y}(\mathbf{q}_0) = \mathbf{W}^T(\mathbf{q}_0)\mathbf{x} \tag{2.4}$$

An ideal narrowband spatial filter has unit response in the pass band and complete attenuation in the stop band. In reality, perfect attenuation in the stop band is not possible. Instead, the LCMV approach minimises the contribution from signals in the stop band by minimising the variance at the filter output, subject to the constraint of unity gain in the pass band. Mathematically, the LCMV approach is expressed as:

$$\min_{\mathbf{W}(\mathbf{q}_0)} tr\left\{\mathbf{C}(\mathbf{y})\right\} \text{ subject to } \mathbf{W}^T(\mathbf{q}_0)\mathbf{H}(\mathbf{q}_0) = \mathbf{I} \tag{2.5}$$

where $\mathbf{I}$ is the identity matrix. The cost function in equation 2.5 requires the covariance matrix of the filter output $\mathbf{C}(\mathbf{y})$. However, since $\mathbf{y} = \mathbf{W}^T(\mathbf{q}_0)\mathbf{x}$, it can be shown that this is related to the data covariance matrix $\mathbf{C}(\mathbf{x})$ by the expression $\mathbf{C}(\mathbf{y}) = \mathbf{W}^T(\mathbf{q}_0)\mathbf{C}(\mathbf{x})\mathbf{W}(\mathbf{q}_0)$. Equation 2.5 can therefore by expressed as:

$$\min_{\mathbf{W}(\mathbf{q}_0)} tr\left\{\mathbf{W}^T(\mathbf{q}_0)\mathbf{C}(\mathbf{x})\mathbf{W}(\mathbf{q}_0)\right\} \text{ subject to } \mathbf{W}^T(\mathbf{q}_0)\mathbf{H}(\mathbf{q}_0) = \mathbf{I} \tag{2.6}$$

The above equation may be solved using the Lagrange multiplier method, resulting in the following solution for the weights matrix $\mathbf{W}(\mathbf{q}_0)$:

$$\mathbf{W}(\mathbf{q}_0) = \left[\mathbf{H}^T(\mathbf{q}_0)\mathbf{C}^{-1}(\mathbf{x})\mathbf{H}(\mathbf{q}_0)\right]^{-1}\mathbf{H}^T(\mathbf{q}_0)\mathbf{C}^{-1}(\mathbf{x}) \tag{2.7}$$

It is now possible to calculate the output from the beamformer $\mathbf{y}$ using 2.4 and 2.7. This represents an estimate of the moment at location $\mathbf{q}_0$. The estimate of the source strength or power from the activity, is the value of the cost function $tr\left\{\mathbf{C}(\mathbf{y})\right\}$ at the minimum. The estimated power projected by the beamformer may therefore be shown to be:

$$\mathbf{P}(\mathbf{q}_0) = tr\left\{\left[\mathbf{H}^T(\mathbf{q}_0)\mathbf{C}^{-1}(\mathbf{x})\mathbf{H}(\mathbf{q}_0)\right]^{-1}\right\} \tag{2.8}$$

The sensitivity of the beamformer varies with distance from the sensors, such that deep sources will result in a much larger value than superficial sources. To account for this the expression in 2.8 may be normalised (often known as weights normalisation) by an estimate for the projected noise, based on the noise covariance matrix $\mathbf{C}_n$:

$$\hat{\mathbf{P}}(\mathbf{q}_0) = \frac{tr\left\{\left[\mathbf{H}^T(\mathbf{q}_0)\mathbf{C}^{-1}(\mathbf{x})\mathbf{H}(\mathbf{q}_0)\right]^{-1}\right\}}{tr\left\{\left[\mathbf{H}^T(\mathbf{q}_0)\mathbf{C}_n^{-1}\mathbf{H}(\mathbf{q}_0)\right]^{-1}\right\}} \tag{2.9}$$

This expression thus represents the power of the neural activity and is commonly known as the neural activity index or the pseudo Z statistic (Van Veen et al., 1997; Vrba & Robinson, 2001).

2.6.2 Dipole orientation

In (2.4) the estimated source $\mathbf{y}(\mathbf{q}_0)$ is assumed to be a $(3{\times}T)$ matrix, corresponding to the estimated activity in each of the three spatial dimensions. This formulation is commonly referred to as a *vector beamformer*. An alternative formulation is that of the *scalar beamformer* in which a single dipole orientation is estimated (under the assumptions of the equivalent current dipole model). Estimation of the source orientation is typically determined either by searching through all possible angles to find the maximum source power (Robinson & Vrba, 1999) or by performing an eigendecomposition of the source power and retaining the eigenvalue that yields maximum source power (Sekihara et al., 2001). Alternatively, given an accurate representation of the subject's cortical folding, dipole orientations may be constrained to lie perpendicular to the cortical surface (Baillet et al., 2001).

An advantage of scalar beamformers is that source estimates are represented by a single one-dimensional time course. However, in cases where the source of interest is weak, the presence of other strong sources may leak into the region of interest resulting in a poor estimate of the genuine source orientation. For the sake of simplicity, a scalar beamformer has been used throughout this thesis.

2.7 Summary

This chapter has provided a brief background to magnetoencephalography, with a particular focus on source reconstruction approaches. The methods developed in this thesis for estimating functional connectivity depend on a robust estimate of neuronal activity. It is therefore important to bear in mind the limitations and assumptions of MEG source reconstruction in order to arrive at meaningful inferences about the nature of long range communication in the brain. In the next chapter we develop a pipeline for identifying MEG resting state networks by combining beamforming with measures of functional connectivity.

CHAPTER 3

MEG functional connectivity

3.1 Chapter Abstract

In **Chapter 2** we described how source-space estimates of neuronal activity may be reconstructed from MEG data recorded at the sensors. These estimates describe the time varying electrical potential at each voxel within the brain and allow the characterisation of activity within specific anatomical regions. Although such approaches are well motivated by the notion of functional segregation (Friston, 1994; Tononi et al., 1994), the brain does not work simply within isolated regions, but as a network. In order to gain a better understanding of brain function, techniques to characterise the interactions between different neuronal assemblies are required. In this chapter we introduce a number of techniques for assessing resting state functional connectivity with MEG. We also outline a number of confounds that make interpretation of these measures challenging, and describe approaches for mitigating these confounds. Finally, we describe an analysis pipeline for detecting resting state functional connectivity in MEG, and demonstrate its application to mapping

resting state networks.

3.2 Assessing brain functional connectivity

The aim of functional connectivity (FC) analysis is to characterise the relationship between signals recorded from distinct neuronal populations as a means of assessing how these populations interact. Functional connectivity describes the statistical dependency between two or more signals and includes measures such as correlation and coherence. An important distinction is the difference between *functional* and *effective* connectivity. The former is an information theoretic measure that describes observable phenomena, whereas the latter is a parametric description of how these phenomena are generated (Friston, 2011). In other words, effective connectivity corresponds to a model of the underlying neuroanatomy and can therefore elucidate the causal nature of these interactions. As a result, effective connectivity has the potential to provide a much more informative description of brain function. However, assessing effective connectivity in the resting-state is extremely challenging due to the complexity of generating realistic whole-brain models and the unconstrained nature of resting state activity. For these reasons, this thesis focuses on measuring functional connectivity. It should be noted however that productive steps are being made in developing whole-brain models of spontaneous activity (for a review see Deco et al. (2011)) and a more complete understanding of resting state activity will likely come from a combination of functional and effective connectivity.

Measures of functional connectivity describe the statistical dependencies between two or more neuronal signals. While most measures may readily be extended to an arbitrary number of signals, the most basic approach is to assess pairwise functional connectivity, describing the relationship between two signals, for instance in terms of their correlation or coherence.

3.2.1 Seed-based functional connectivity

The most straightforward applications of functional connectivity are seed-based approaches. In these approaches, a predetermined location of interest is specified as the *seed voxel* and statistical measures of functional connectivity are computed between the time course at this voxel, and the time course at another *target voxel*. In many cases all voxels in the brain will be sequentially used as the target voxel resulting in a whole brain statistical map of functional connectivity with the seed location. An obvious limitation of seed-based approaches is the requirement for a predetermined region of interest (ROI) to use as the seed, the choice of which will affect the measured functional connectivity. Approaches for defining regions of interest are discussed in Section 3.3.3.

Seed-based functional connectivity has been used extensively in resting state fMRI studies, where functional connectivity is estimated using the correlation between the measured BOLD activity. For example, in the pioneering study of Biswal et al. (1995), the correlation between a seed within the motor cortex and all other voxels revealed regions of high correlation in the motor cortex contralateral to the seed. Correlation is a measure of the instantaneous dependencies between two time series and is well suited to fMRI due to the temporal smoothness of the BOLD signal (since the haemodynamic response causes neuronal activity to be temporally blurred, neuronal activity that is simultaneous but not phase-locked will still result in correlated BOLD activity).

In MEG and EEG however, straightforward correlation between the measured time courses does not generally provide a meaningful measure of functional connectivity. A significantly non-zero correlation would require the activity between neuronal populations to be time-locked with zero phase-lag, and that this relationship be maintained over a sufficiently large number of time points. In practice these conditions are rarely met. A more appropriate measure is coherence, which requires only a fixed phase relationship between the two signals, and therefore allows for the possibility that neuronal populations have a fixed non-zero phase lag. Such a relationship is more realistic and may arise out of axon conduction and synaptic delays (Buzsáki & Draguhn, 2004). For stimulus locked data, coherence and other measures of phase locking have proved to be useful tools in elucidating functional connectivity (see e.g. Siegel et al. (2012) for a review). In task-free paradigms, such measures have generally been less successful. This may be, in part, due to the unconstrained nature of the resting state, such that a stationary fixed-phase relationship is not sufficiently maintained over the duration of the recording. Recent work has suggested that phase interactions may be identified from resting state MEG recordings (Marzetti et al., 2013). However, it is unclear to what extent regions in which strong phase-locking is observed are simply those with high source power (Hillebrand et al., 2011).

3.2.1.1 Functional connectivity from band-limited envelope correlation

A more successful measure for assessing resting state MEG functional connectivity is band-limited power (BLP) correlation. Rather than computing the correlation between the raw time courses, a measure of time varying power (or alternatively amplitude) is first computed from the band-limited (i.e. frequency filtered) MEG data. This may be computed using one of a number of techniques, for instance using time-frequency decompositions such as the Hilbert or wavelet transforms. This time course corresponds to the envelope of oscillatory power or amplitude within a particular frequency band and is shown schematically in Figure 3.1. Functional connectivity may then be computed as the correlation between these envelope time courses. Since the envelope captures only changes in signal amplitude, this measure does not require a phase-locked relationship.

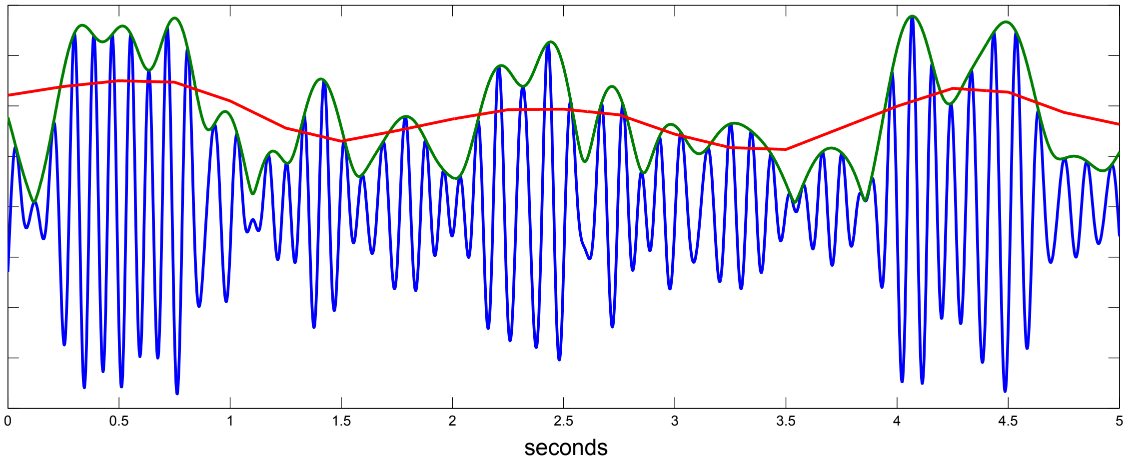

Figure 3.1: Raw time course (blue), amplitude envelope (green) and amplitude envelope temporally averaged within a 1 second sliding window (red) for a simulated alpha oscillation.

Envelope correlation has proved successful in identifying resting state functional connectivity in MEG. For instance, Brookes et al. (2011a) demonstrated significant envelope correlation between the left and right sensorimotor cortices. Furthermore, by computing the correlation between these two regions across a range of frequency bands, they showed that this functional connectivity has a frequency dependency that peaks in the beta band. An important aspect of this study was an investigation of the time scales at which significant envelope correlation could be detected. The authors used two measures of functional connectivity based on windowing the envelopes and computing either the average *within-window correlation*, or the correlation between *averaged envelopes* (by first computing the average amplitude of the envelope within each window; red in Figure 3.1). The within-window correlation measure was significant for window sizes of 1 second and longer. However, the correlation of averaged-envelope provided the most robust measure of functional connectivity, with the most significant correlation for window sizes of 0.5-2 seconds. These findings were further validated for regions other than the sensorimotor network by Luckhoo et al. (2012).

3.2.2 Independent component analysis

An alternative to seed-based methods is to use data driven approaches to decompose whole brain measurements into a subset of components that describe unique spatiotemporal characteristics of the data. By far the most widely adopted of these approaches is independent component analysis (ICA), which decomposes the data into additive subcomponents that are statistically independent. ICA is a method for *blind source separation*, i.e. it aims to separate a set of mixed signals without knowledge of the mixing process or the underlying signals. Provided the number of observations is equal to or greater than the number of sources (i.e. the problem is not under-determined), this may be achieved by maximising the statistical independence of the estimated components. One way of achieving this is to maximise the non-Gaussianity of the estimated signals. This approach is motivated by the central limit theorem that states that a linear mixture of signals will have a more Gaussian distribution than any of the mixed signals.

Application of ICA to an $N \times T$ set of mixed signals $\mathbf{X}$ (where N is the number of signals and T is the number of samples) results in an $M \times T$ set of independent components $\mathbf{S}$ and an $N \times M$ mixing matrix $\mathbf{A}$ that describes the contribution of each component to each measurement:

$$\mathbf{X} = \mathbf{A}\mathbf{S} \tag{3.1}$$

Additionally, one can estimate an unmixing matrix $\mathbf{W}$ as the (pseudo) inverse of **A** that describes the contribution of each measurement to each component. In the case where the measurements are the time course at each voxel, **W** corresponds to a spatial map describing the weighting of each voxel to that component. Voxels that are strongly represented by a particular component may therefore be said to be functionally connected, since they share the same time course to some extent.

The matrix **X** represents the ($voxels \times time\,points$) spatiotemporal data. The orientation of this matrix (i.e. voxels represented by rows or columns) determines whether independence is maximised along the spatial or temporal dimension. In fMRI, spatial ICA has been most widely used, since robust application of ICA requires a large number of samples in the independent dimension, and there are orders of magnitude more voxels that time points in fMRI. On the other hand, the high temporal resolution of MEG means that temporal ICA is more appropriate. The dimension along which independence is maximised will result in different decompositions of the data. In the case of spatial ICA, this decomposition will result in a set of uncorrelated spatial components, and therefore this assumption promotes sparse (i.e. non-overlapping) spatial maps. Temporal ICA on the other hand does not necessarily promote sparsity in space, and therefore can identify spatial maps which overlap but have different temporal “modes” (Smith et al., 2012). Interestingly, when applied to MEG data, temporal ICA reveals spatial maps that more strongly resemble those identified from applying spatial ICA to fMRI data, despite their different assumptions of independence (see Figure 1.2). This is potentially due to the increased spatial smoothness of MEG data (Smith et al., 2012).

3.2.2.1 Group ICA via temporal concatenation

In both fMRI and MEG, ICA has largely been applied to group-concatenated data for two reasons. Firstly, to ensure enough samples for a robust decomposition, and secondly, because one is generally interested in networks that are represented across a group rather than in individuals. It is of course tempting to consider applying ICA to each subject separately and then deriving a group average *post-hoc*. However this approach proves challenging in practice, since each decomposition will result in a different set of independent time courses and mixing matrices, and there will be no clear one-to-one correspondence between these decompositions. Instead, the data from all subjects are concatenated along the temporal dimension, resulting in a single $voxels \times (time\,points \times subjects)$ data matrix $\mathbf{X}_{group}$ that represents the group.

3.2.2.2 Application of ICA to MEG functional connectivity

Motivated by the observation that interhemispheric functional connectivity could be identified from correlations in the amplitude envelope of source reconstructed MEG activity, Brookes et al. (2011b) assessed whole brain functional connectivity by applying temporal ICA to the envelopes at all voxels. In this study a model order of 25 was used (i.e. the dimensionality of the data was reduced to 25 prior to performing ICA). Of these 25 components, 8 had spatial maps (i.e. the unmixing matrix) that corresponded to RSNs that had been previously reported from fMRI studies. The best spatial correspondence with fMRI RSNs was achieved when the beta band envelopes were used, with the exception of the default mode network, which was best found in the alpha band. As with seed-based estimates of functional connectivity, the amplitude envelopes were first temporally averaged and down-sampled as shown in Figure 3.1. Thus the identified networks correspond to regions whose amplitude envelopes covary on a slow temporal scale.

3.2.3 Effective connectivity

So far we have considered measures of functional connectivity. These measures quantify statistical dependencies between neuronal signals, i.e. they are nonparametric. Another branch of connectivity measures is *effective connectivity*, which provide a parametric description of the causal relationships between neuronal signals. This loose definition captures a large number of measures. The most straightforward are autoregressive (AR) models. An autoregressive model represents a time series as a weighted sum of its p previous values (where p defines the "order" of the model), and additive noise. These models are readily extended to multivariate autoregressive models (MVAR) that model a signal based on its previous p values *and* the previous p values of a set of other signals. This model provides a common basis for a range of coupling measures including partial directed coherence (Baccalá & Sameshima, 2001), direct transfer function (Kaminski & Blinowska, 1991), and Granger causality (Granger, 1969). Since MVAR models use previous values of other time series to construct the model, these coupling measures can be used to estimate not only the strength of connections but also their directionality, by assessing causality (in the Granger sense). Put simply, a time series x may be said to "Granger cause" another time series y, if knowledge of past values of x may be used to better predict future values of y. Multivariate autoregressive models have been applied to EEG and MEG task data (Schlögl & Supp, 2006; Supp et al., 2007; Michalareas et al., 2013) but have so far been relatively unexplored for resting state connectivity.

Other measures of effective connectivity come in the form of generative models. In these approaches, the underlying neuronal sources and their connections are modelled, for instance using neural mass models or coupled oscillators. The data generated by such a model may then be projected to sensor level using a realistic head model and compared with the recorded data to assess the performance of the model. In some approaches, including dynamic causal modelling (DCM), this is achieved within the framework of Bayesian model comparison, by which different models of network structure are compared in terms of their model evidence (i.e. the probability of the observed data, given the particular model in question) (Friston, 2011). This allows hypotheses about the underlying neuronal dynamics and network structure to be tested. However, it should be noted that a particular model may only be considered "best" (as judged by its Bayesian model evidence) in the context of the other models it is compared to, which form only a tiny fraction of the entire search space (Lohmann et al., 2012). Due to exponential growth of this search space with increased model complexity, DCM has so far shown greatest success for networks with only a small (<10) number of nodes (Smith et al., 2013b). However, a recent extension of DCM to model intrinsic fluctuations has been shown to enable the application of DCM to resting state data (Friston et al., 2011).

Other work has focused on simulating resting state activity using large scale network models whose structure is defined by anatomical connectivity. By combining measurements of structural connectivity with realistic assumptions about the neuronal dynamics, interactions and delays, these models have demonstrated simulated functional connectivity that closely matches measured functional connectivity in fMRI (Honey et al., 2009; Deco et al., 2013; Cabral et al., 2011) and MEG (Cabral et al., 2013). Notably, Cabral et al. (2013) used a set of 90 weakly coupled oscillators within a brain network model defined from diffusion tensor imaging (DTI) to reveal emergent band-limited amplitude correlations that matched those observed in real MEG data.

3.3 Challenges in MEG functional connectivity

In this section we review a number of challenges that can confound source space estimates of MEG functional connectivity, and how these confounds may be mitigated. This summary will provide the motivation for a number of techniques used in the pipeline for assessing resting state functional connectivity.

3.3.1 Effect of field spread

Field spread refers to the propagation of magnetic fields generated by neuronal sources such that they may be measured at distant locations. While this property makes recording neuronal activity outside of the head possible, the fact that sources are simultaneously recorded by multiple sensors makes the interpretation of functional connectivity at the sensor level particularly challenging, since an observed level of correlation or coherence between pairs of sensors may be caused entirely by a single neuronal source, rather than by distinct coherent sources. Source reconstruction techniques such as beamforming and minimum-norm estimates have been shown to limit the effect of field spread (Schoffelen & Gross, 2009). However due to the ill-posed nature of the inverse problem it is not completely abolished, meaning that the weights used to project the data into source space will be correlated (Palva & Palva, 2012). This weights correlation means that there will be *signal leakage* or *cross-talk* between the source space estimates of neural activity at spatially distinct voxels, resulting in spurious functional connectivity between voxels that may be entirely driven by this leakage. Furthermore, this weights correlation is inhomogeneous across the head meaning that simply ensuring a minimum separation between seed and target voxels may not be sufficient to ensure connectivity measures are not spurious.

3.3.1.1 Signal leakage correction

Due to the challenges involved in interpreting seed-based connectivity in the context of signal leakage, a number of functional connectivity measures have been proposed to mitigate these effects. Since the propagation of magnetic fields to the sensors is much faster than the sampling rate (i.e. the speed of light versus 1 ms), there will be no measurable delay, or phase difference, between the field measured at different sensors, meaning that spurious interactions resulting from field spread (or signal leakage between voxels, since source projection involves a linear weighting of the sensors) will be at zero phase-lag. Conversely, genuine neuronal interactions will be expected to be at non-zero phase-lag due to axon conduction and synaptic delays (Buzsáki & Draguhn, 2004). Thus functional connectivity measures that rely only on lagged interactions provide a means to separate spurious from genuine connectivity. Such measures include the imaginary part of coherency (Nolte et al., 2004) and the phase lag index (Stam et al., 2007).

An alternative strategy is to remove linear dependencies from the two signals of interest prior to computing pairwise functional connectivity. This may be achieved by orthogonalising the two signals with respect to one another, by subtracting from one signal $\mathbf{x}$ the univariate projection on $\mathbf{x}$ of the second signal $\mathbf{y}$:

$$\mathbf{x} \perp \mathbf{y} = \mathbf{x} - \mathbf{y}\left(\mathbf{y}^{+}\mathbf{x}\right) \tag{3.2}$$

where $\mathbf{y}^{+}$ is the pseudo-inverse of $\mathbf{y}$. The orthogonalised signal $\mathbf{x} \perp \mathbf{y}$ is then used in place of $\mathbf{x}$ when computing pairwise functional connectivity. This technique has been successfully applied in conjunction with envelope correlation, to ensure that interactions between the amplitude of two signals is not driven purely by weights correlation (Brookes et al., 2012b; Hipp et al., 2012). It should be noted that the orthogonalisation may be performed in one of two directions ($\mathbf{x}$ to $\mathbf{y}$ or $\mathbf{y}$ to $\mathbf{x}$). This choice is somewhat arbitrary but may result in slightly different functional connectivity scores. One approach is to simply perform the orthogonalisation in each direction independently and average the two resulting functional connectivity scores.

3.3.2 Interference and artefacts

As described in **Chapter 2**, the low signal-to-noise ratio (SNR) of MEG means that the measured signal will contain not only neuronal signals, but also interference from environmental or physiological noise sources. These sources of interference can have a severe impact on functional connectivity metrics, since many of these artefacts will simultaneously affect multiple channels (or voxels after source projection) and will therefore result in statistical dependencies between signals that may be incorrectly interpreted as genuine neuronal interactions. For this reason, the ability to reliably detect and remove sources of interference via signal preprocessing represents a key step in the analysis of functional connectivity.

3.3.3 Choice of ROI

When assessing pairwise functional connectivity it is common to restrict the analysis to a subset of ROIs. This is because computing the full connectivity matrix between all pairs of voxels would result in excess of 10 million computations for each subject, which can quickly render computation prohibitive. Interpreting this dense representation of the data is also challenging in itself.

A number of approaches exist for defining regions exist. Firstly, the brain may be segmented into a number of *parcels* that correspond to anatomically or functionally distinct subregions. Anatomical parcellations may be defined based on cytoarchitectonics (Brodmann, 1909), cortical curvature (Van Essen, 2005) or structural connectivity (Johansen-Berg et al., 2004). While such approaches have been employed for studies of MEG functional connectivity (e.g. Hillebrand et al., 2011), the boundaries of structurally derived parcels may not correspond to functional connectivity boundaries, resulting in suboptimal parcels, or a mixing of time series across multiple parcels. For these reasons, it has been argued that structural parcellations should be avoided for functional connectivity analysis in favour of data driven functional parcellations (Smith et al., 2011). These approaches are typically defined by clustering the brain into functionally homogeneous subregions, for instance using hierarchical clustering or spatial ICA. The key parameter in such approaches is the number of components to estimate: low model orders (less than 30) will typically result in RSNs, whereas higher model orders will split these networks into subcomponents corresponding to individual network nodes (Smith, 2012).

Another option is to apply graph theory to compute summary measures that describe the global connectivity structure. Hipp et al. (2012) used such an approach to characterise the correlation structure of MEG BLP in terms of each voxel's number of connections and "hubness".

3.4 Pipeline for assessing resting state functional connectivity

In this section we demonstrate the application of band-limited envelope correlation to identify functional connectivity from resting state MEG data. This analysis represents a reproduction of findings from the existing literature that have applied temporally stationary measures of functional connectivity to MEG data (Brookes et al., 2011a; Hipp et al., 2012). Reproducing the findings from these studies is a prerequisite step for testing a pipeline for assessing MEG resting state functional connectivity, and will lay the foundations for the more novel *non-stationary* methodologies described later in this thesis.

3.4.1 Overview of pairwise functional connectivity pipeline

The pipeline used to identify pairwise (seed-based) functional connectivity may be summarised as follows:

1. **Artefact rejection:** Artefacts due to environmental and physiological interference removed through a combination of (i) manual identification of bad channels and time epochs, (ii) ICA decomposition and removal of artefact components, (iii) semi-automated identification and rejection of muscle artefacts.
2. **Frequency band specification:** The sensor space data is filtered into one or more frequency bands for subsequent source space analysis.
3. **Source reconstruction:** Projection of the data into source space, either at specific regions of interest or across the whole brain by defining a grid of voxels spanning the entire volume of the head.
4. **Signal leakage correction:** The signals from the seed and target voxels are orthogonalised with respect to one another as in (3.2).
5. **Envelope computation:** The envelopes of the orthogonalised signals are computed as the magnitude of the analytic signal derived from the Hilbert transform.
6. **Sliding window averaging:** The envelopes of both the seed and target voxels are temporally averaged within a sliding window of 1 s duration, with 75% overlap between adjacent windows, resulting in a low-pass filtered envelope.
7. **Envelope correlation:** The Pearson correlation coefficient is computed between each pair of averaged envelopes.

3.4.2 Data and preprocessing

Resting state MEG data were acquired from ten healthy subjects (mean age 27 ± 8 years, three females). The subjects were asked to lie in the scanner with their eyes open while ten minutes of data were recorded. The dataset used in this chapter are described in full in Section 5.7. The data were converted to SPM8 `http://www.fil.ion.ucl.ac.uk/spm` and down-sampled to 200 Hz.

3.4.3 Artefact rejection

As described in Section 3.3.2, environmental and physiological artefacts can potentially confound functional connectivity measurement. We used a number of techniques to remove these artefacts as much as possible, without removing neuronal activity.

3.4.3.1 Manual artefact rejection

Before proceeding with automated artefact rejection techniques, the most severe artefacts were identified and removed by visual inspection of the data. This was achieved by flagging channels and time epochs as bad, such that they may be excluded from subsequent analyses. It is important to perform this process prior to applying more sophisticated preprocessing tools, since channels and/or epochs with very high variance can degrade the performance of these tools (for instance, the correlation between ICA components and externally measured signals such as the electrocardiogram (ECG) can be confounded if both are contaminated by a high variance artefact). In contrast to trial-wise data where individual trials may be discarded, resting state data has no such natural partitioning into epochs. We adopted an approach based on defining "bad epochs" of arbitrary length (as shown in Figure 3.4.3.1b). By recording the time points that fall within these epochs, these data may be excluded from later stages of the pipeline (while allowing the temporal continuity of the data to be maintained where necessary).

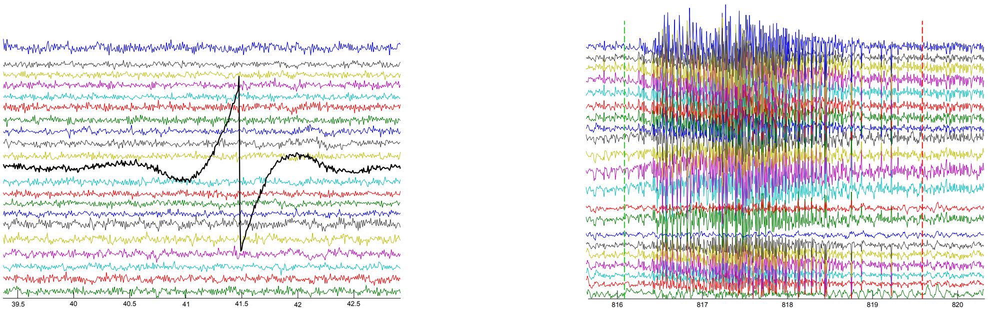

(a) (b)

Figure 3.2: Example of flagging (a) a bad channel (shown in black) and (b) a bad epoch (defined as those time points that lie between the green and red dashes). These channels and time points were excluded from subsequent stages of the pipeline.

3.4.3.2 ICA artefact rejection

The measured magnetic fields comprise a mixture of neuronal signals, non-neuronal physiological signals (such as the magnetocardiogram), and environmental signals (such as mains noise). As described in Section 3.2.2, ICA may be used to decompose a mixture of signals into its constituent components and is widely used for denoising MEG and EEG recordings (Mantini et al., 2011). Here we used ICA to decompose the sensor data into 150 temporally independent components (tICs), which were classified into brain and artefact components. From (3.1), the sensor data $\mathbf{X}$ may be decomposed using ICA into a set of unmixed sources $\mathbf{S}$ and a mixing matrix $\mathbf{A}$. By assuming that these sources may be perfectly separated into brain and artefact components, (3.1) may be rewritten as:

$$\mathbf{X} = \begin{bmatrix} \mathbf{A}_{brain} & \mathbf{A}_{artefact} \end{bmatrix} \begin{bmatrix} \mathbf{S}_{brain} \\ \mathbf{S}_{artefact} \end{bmatrix} \tag{3.3}$$

From this decomposition, an estimate of the denoised sensor data $\mathbf{X}_{brain}$ may be recovered in one of two ways. Firstly, the brain components $\mathbf{S}_{brain}$ may be multiplied by the portion of the mixing matrix corresponding to these components $\mathbf{A}_{artefact}$ to exclude the artefact components:

$$\mathbf{X}_{brain} = \mathbf{A}_{brain}\mathbf{S}_{brain} \tag{3.4}$$

Alternatively, the denoised sensor data may be estimated by subtracting the artefact data $\mathbf{X}_{artefact}$ from the original sensor data $\mathbf{X}$. It is typical for the dimensionality of the data to be reduced prior to performing the ICA decomposition. Estimating the denoised sensor data by subtracting the artefact components therefore has the advantage that neuronal activity that was excluded by the dimensionality reduction stage may be retained, by subtracting the artefact data from the full rank sensor data (Mantini et al., 2011). This subtraction may be expressed as follows:

$$\mathbf{X}_{brain} = \mathbf{X} - \mathbf{A}_{artefact}\mathbf{S}_{artefact} \tag{3.5}$$

For datasets in which external signals (such as the electrocardiogram or electrooculogram (EOG)) were simultaneously acquired, this was achieved by computing the temporal correlation between each tIC and these external signals. For datasets without these external signals, classification was performed manually by visual inspection of the tICs and their associated spatial topographies. Three types of artefacts were rejected at this stage:

- **Cardiac components**: For datasets with simultaneous ECG recordings, components were classified as cardiac related if the correlation between the ECG and the tICs exceeded 0.3. For datasets without ECG recordings, cardiac components were identified by visual inspection, by looking for components that contained regular (∼1 Hz) deflections resembling the QRS complex. An example cardiac tIC and its associated spatial topographies for both sensor types is shown in Figure 3.3a.

- **Eye blink components**: For datasets with simultaneous EOG and/or eye-tracker recordings, components were classified as blink related if the correlation between these signals and the tICs exceeded 0.3. For datasets without these recordings, blink components were identified by visual inspection, by looking for components with high kurtosis (>20), a regular blink structure, and characteristic spatial topographies as shown in Figure 3.3b.
- **Mains components**: Components were classified as mains-related if the power spectrum of the tIC had a global maximum at 50 Hz. An example mains component is shown in Figure 3.3c.

3.4.3.3 Muscle artefact rejection

As a final preprocessing step, the data were compared with electromyogram (EMG) signals recorded from the jaw (where available) in order to remove artefacts due to muscle activity. Contractions of the facial muscles as measured by EMG typically show a power spectrum that peaks between 30 and 80 Hz (Muthukumaraswamy, 2013); the frequency range of these artefacts therefore overlaps with the gamma frequency band. Although subjects were requested to remain as still as possible, some instances of muscles activity are inevitable during a 10-15 minute recording session (typically as the result of swallowing or jaw clenching). We identified these brief instances of muscle activity by first high-pass filtering the EMG and MEG signals above 60 Hz (such that the bandwidth of the filtered data was between 60 Hz and the Nyquist limit of 125 Hz). A time course of BLP was computed by applying

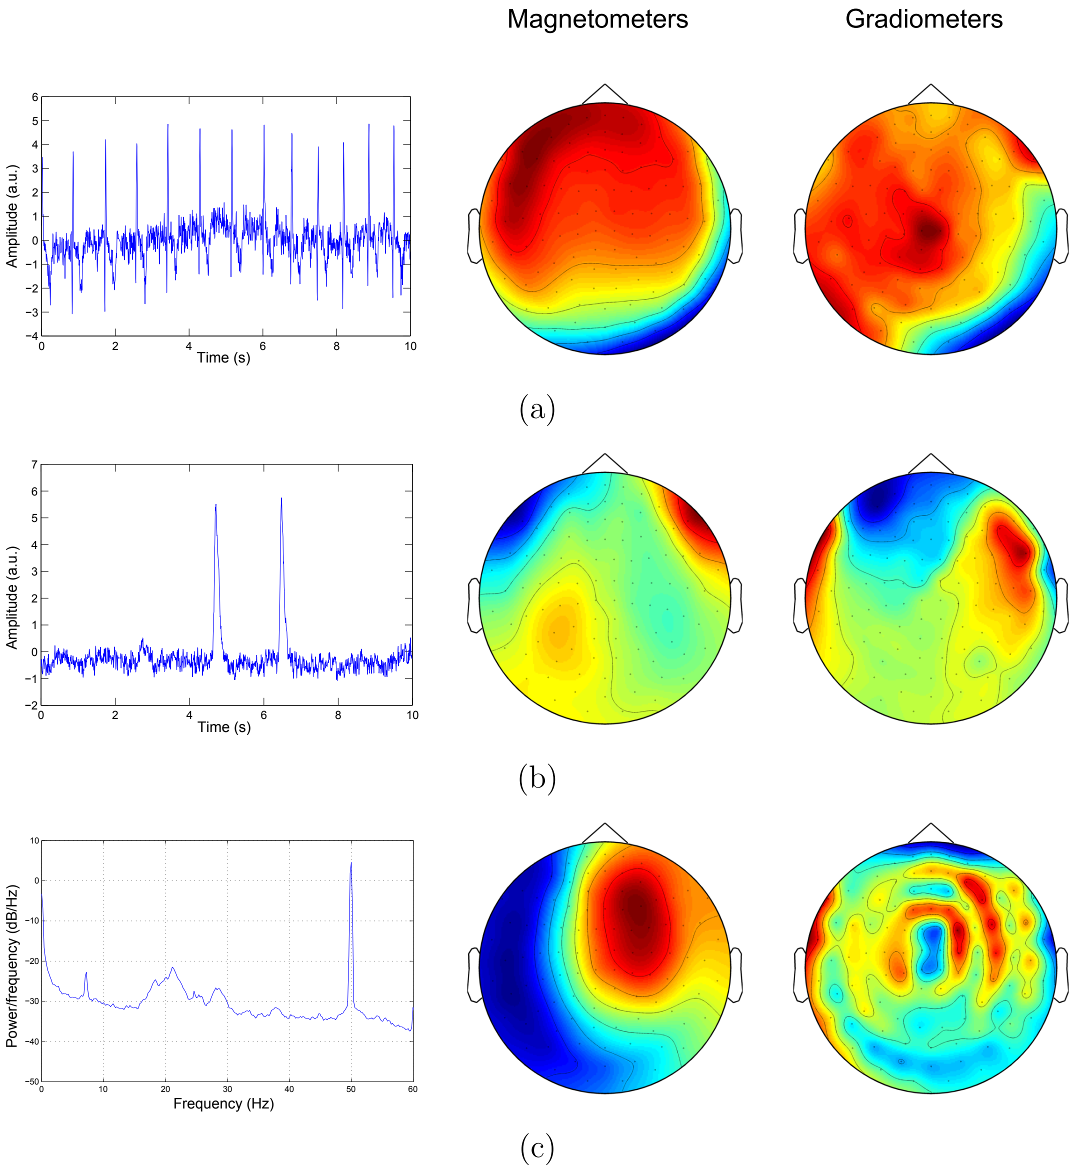

Figure 3.3: ICA time courses/power spectra and sensor topographies for (a) a cardiac component (b) an eye blink component (c) a 50 Hz mains component. The sensor topographies show the weighting of each sensor to the artefact component (i.e. a single column of the mixing matrix). These topographies are shown for the magnetometers and gradiometers separately.

a 400 ms moving average filter to the square of the signal amplitude. Figure 3.4a shows an example of band-limited EMG power recorded from the masseter muscle of the jaw. There are a number of transient periods of high EMG power, indicating increased muscle activity within these time points. Figure 3.4b shows the BLP averaged across all MEG channels for this same recording. Periods of high MEG power in the 60 - 125 Hz band occur at the same time points, demonstrating that this muscle activity is also measured by the MEG sensors. To identify and remove these artefacts, we applied the following heuristic method:

1. High-pass filter each EMG channel above 60 Hz.
2. Compute BLP by applying 400 ms moving average filter to the square of the filtered signal.
3. Remove baseline drifts in power by high-pass filtering the BLP above 0.1 Hz.
4. Identify outliers as all time points exceeding a threshold of 10 times the interquartile range.
5. Define events by the start and end time points of contiguous periods of outliers.
6. Merge events that are within 1 s of one another.
7. Remove any events that are of less than 200 ms in duration.
8. Dilate events by 500 ms.

Two detected muscle artefacts are shown in Figure 3.4c. The data within all muscle artefact events were excluded from later stages of the pipeline in the same way as the epochs manually marked as bad. We adopted the more conservative approach of removing periods of data rather than using the EMG time course to identify ICA components, since it was found that these artefacts were typically correlated with a large number of tICs. This is possibly because periods of muscle activity are very transient and present in a small proportion of the recording compared with more sustained artefacts such as the ECG or EOG. Consequently we chose not to risk removing genuine neuronal data by rejecting multiple tICs, and instead removed those brief periods in time in which muscle activity was high.

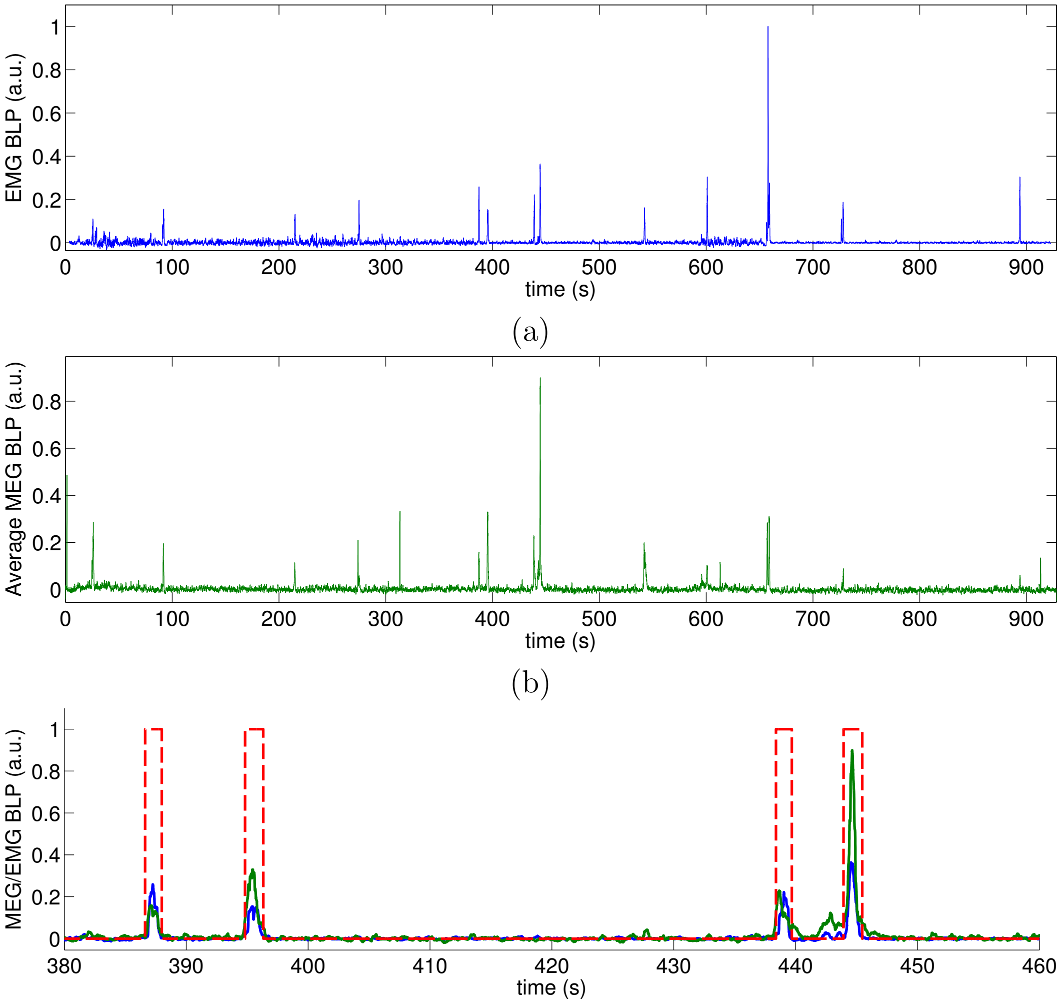

(a)

(b)

(c)

Figure 3.4: Top: Band limited EMG power in the 60-80 Hz band recorded from the jaw. Middle: Band limited MEG power in the 60-80 Hz band, averaged over all channels. Bottom: identification of two instances of increased muscle activity. The EMG and MEG BLP are shown in blue and green respectively and the dashed red lines show the time points marked as bad.

3.4.4 Results

We demonstrate the use of the pipeline for estimating source-space functional connectivity by reproducing two key results that have been shown in previous studies.

Firstly, we perform a whole-brain seed-based connectivity analysis of three resting state networks to show that we can identify the same spatial distribution of envelope correlation shown in previous studies (Brookes et al., 2011a; Hipp et al., 2012). Secondly, we reproduce findings that demonstrate the frequency dependency of this functional connectivity.

3.4.4.1 Seed-based functional connectivity

Following artefact rejection, the data for each subject were band-pass filtered into the 4 to 30 Hz band and source-projected using the LCMV beamformer as described in Section 2.6. Source time courses were estimated at each voxel on a regular 8 mm grid spanning the entire brain. Regions of interest to be used as the seed locations were defined in the right hemisphere for nodes of the sensorimotor network (right primary motor cortex), visual network (right V1) and default mode network (right inferior parietal lobule). These ROIs were defined from previously published fMRI RSNs as the MNI coordinates corresponding to the peak voxel within each RSN's spatial map (Smith et al., 2009). The fMRI RSNs and the nodes used are shown in Figure 3.5.

For each ROI, seed-based functional connectivity was computed between the time course at the seed voxel and all other voxels in the brain. Functional connectivity was estimated using a combination of signal leakage correction and amplitude envelope correlation, following the procedure described in Section 3.4.1. These correlations were mapped across the brain, resulting in a spatial map for each seed and each subject. The spatial maps were averaged over all subjects resulting in a single spatial map for each seed. Figure 3.6 shows the spatial map of envelope correlation with the three seeds rendered onto the MNI cortical surface (thresholded at 50% of the maximum correlation in each map). In each map, there is strong functional connectivity between the right hemisphere seed and homologous areas

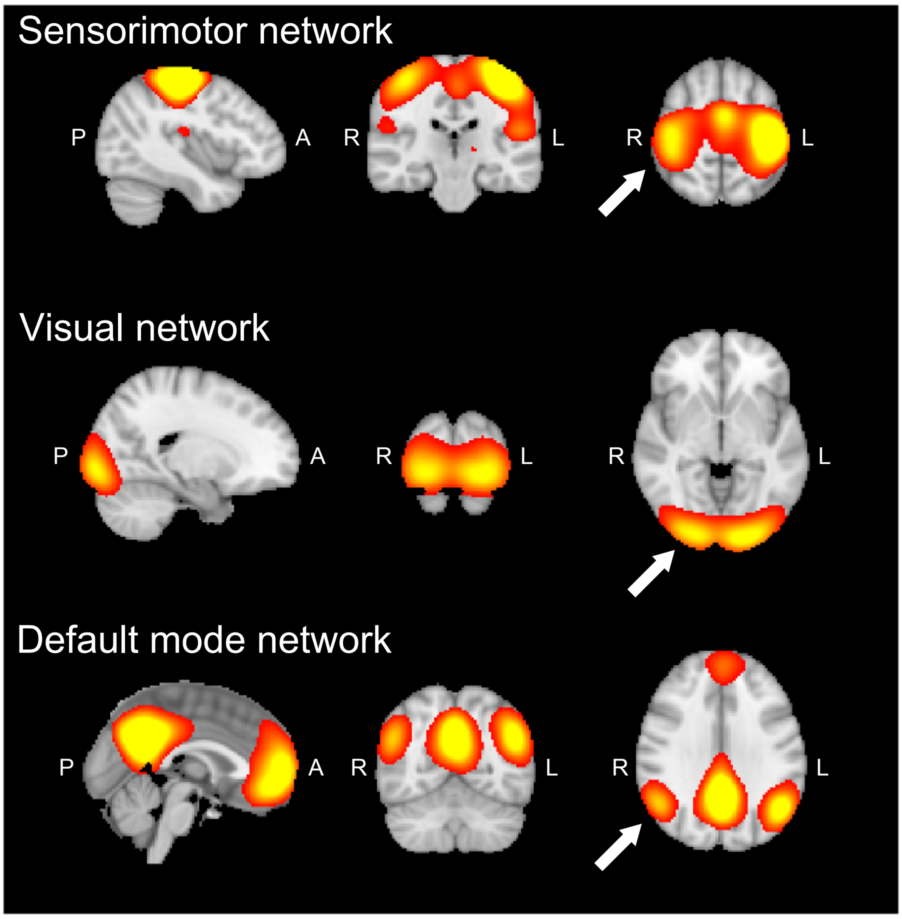

Figure 3.5: Z-statistical spatial maps thresholded at Z = 3 for the sensorimotor, visual and default mode networks taken from Smith et al. (2009). The white arrows indicate the right hemisphere network nodes used as ROIs in the seed-based connectivity analysis. For each node, the MNI coordinates corresponding to the peak voxel were used as the seed region.

in the left hemisphere, in strong agreement with previous studies (Brookes et al., 2011a; Hipp et al., 2012). Seeds placed in the right hemisphere nodes of the visual and sensorimotor networks revealed functional connectivity in homologous regions in the left hemisphere. The seed placed within the right inferior parietal lobule (rIPL) of the DMN revealed functional connectivity with the lIPL but also with ipsilateral nodes of the DMN including the right temporal lobe, and right medial frontal gyrus.

3.4.4.2 Frequency dependency of MEG functional connectivity

In the previous section we demonstrated the application of seed-based functional connectivity to identify MEG resting state networks. These results represent a reproduction of existing findings shown by Brookes et al. (2011a) and Hipp et al. (2012). An important finding from these studies was that functional connectivity is strongly dependent on frequency, with increased BLP correlation within the alpha and beta bands. To assess the frequency dependency of inter-hemispheric functional connectivity, a single target voxel was defined for each network as the voxel in the left hemisphere that was most strongly correlated with the right hemisphere seed. The approximate locations of these seed and target voxels are shown in Figure 3.6a as white stars and circles respectively. The same pipeline as before was applied to estimate envelope correlation between these three seed-target pairs after frequency filtering the data into 25 overlapping 2 Hz frequency bands spanning the entire 4-30 Hz band of the beamformed data. Figure 3.6b shows the group averaged envelope correlation for each network over all frequency bands. The shaded region represents the standard error of the mean over all subjects. For the visual network, functional connectivity peaks in the alpha band, whereas the sensorimotor and default mode networks have peaks in both alpha and low beta bands. These frequency dependencies show a strong agreement with the findings of Hipp et al. (2012).

3.5 Summary

In this chapter we have introduced the concept of functional connectivity and discussed some of the issues in applying these measures to resting state MEG data. Understanding the advantages and limitations of these techniques is fundamental if one is to apply functional connectivity in a way that results in meaningful results.

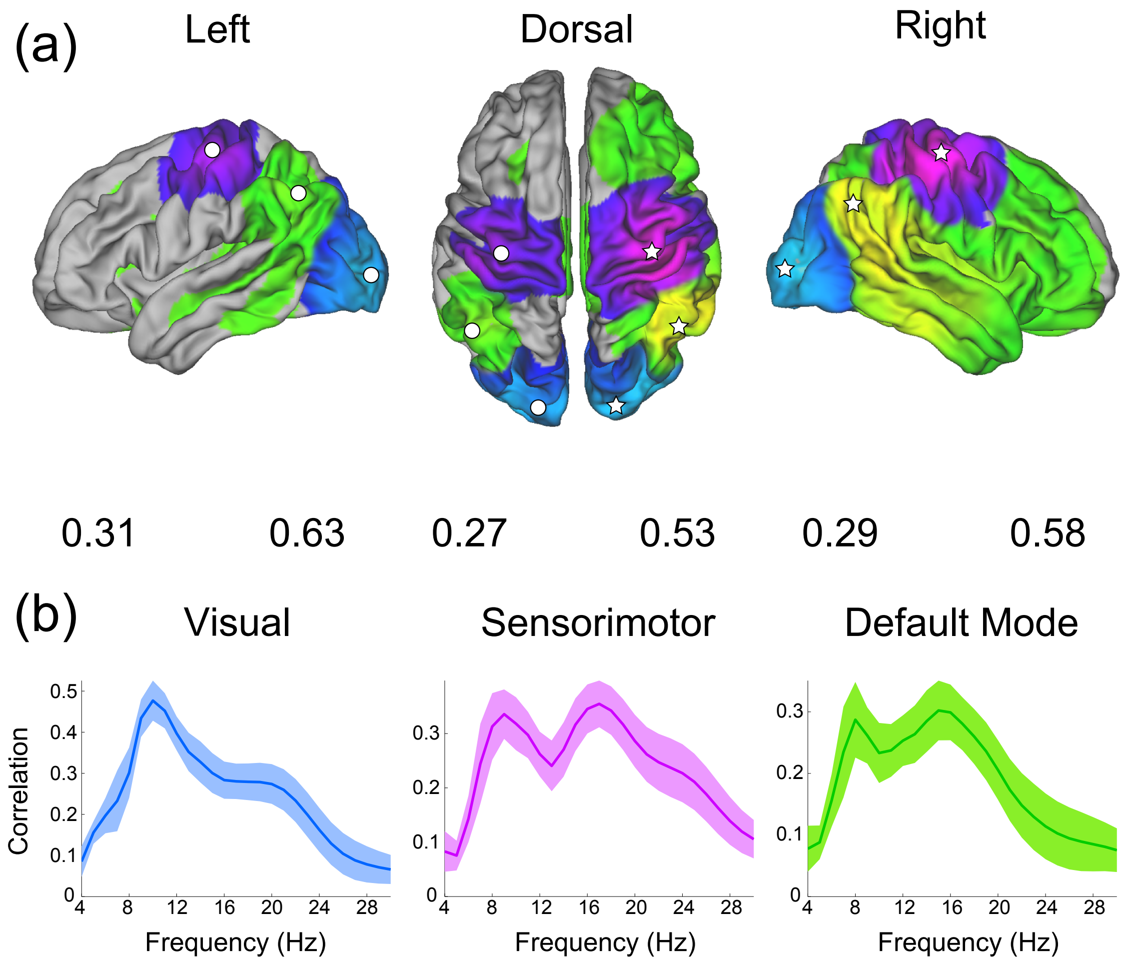

Figure 3.6: (a) Surface rendered spatial maps of seed-based resting state functional connectivity with each of the three RSN nodes (sensorimotor, visual and default mode networks). The maps represent the average correlation score over nine subjects. Each map has been thresholded at 50% of the maximum correlation, and lighter colours represent higher correlation scores. In cases where the maps of more than one network overlap, the map with highest correlation score is shown. (b) Frequency dependency of functional connectivity, computed between each right hemisphere seed and the mirrored location in the left hemisphere (shown by white stars and circles in (a) respectively).

Based on this discussion and previous literature, we have presented a pipeline for assessing resting state functional connectivity. This pipeline includes techniques for preprocessing MEG data such that spurious correlations due to artefacts are mitigated. In addition, signal leakage correction is employed in order to prevent weights correlation from confounding functional connectivity measures. We validated the pipeline by assessing interhemispheric functional connectivity between three resting state networks, using regions of interest derived from fMRI RSNs. Specifically, we demonstrated that the spatial extent and frequency dependency of MEG functional connectivity is consistent with results from previous studies. The results presented in this chapter lay the foundations for novel work in the following three chapters where we develop *non-stationary* measures of functional connectivity. In **Chapter 4** we extend the seed-based approach to study the time-varying nature of resting state functional connectivity. In **Chapter 5** and **Chapter 6** we depart from seed-based approaches and develop a more sophisticated *data driven* approach to assessing non-stationary whole brain functional connectivity based on a hidden Markov model.

CHAPTER 4

Investigating the temporal dynamics of resting state functional connectivity with MEG

4.1 Chapter Abstract

This chapter focuses on the application of time varying functional connectivity measures to investigate the temporal dynamics of resting state connectivity. We focus on connectivity within the default mode and the sensorimotor networks, which can be robustly identified independently across multiple subjects in MEG data. Regions of interest (ROIs) corresponding to nodes of these networks were first identified using ICA, and the estimated neuronal activity at these ROIs were then subject to further time-varying functional connectivity analysis using sliding window correlations on the band-limited power (BLP) time courses for pairs of regions at a time. By comparing these measures to those computed from simulations of realistic null data, we demonstrate that within-network envelope correlations are

significantly temporally non-stationary, and determine the minimum time scales at which significant temporal variability may be identified. Furthermore, we show that this temporal variability occurs at those frequencies that have previously been shown to reveal strong temporally stationary functional connectivity. We further computed measures of bimodality to provide evidence that temporal variability in MEG functional connectivity is associated with bi-state behaviour, in which there is dynamic switching between periods of low functional connectivity and transient periods of high functional connectivity.

4.2 Introduction

Until recently, studies of resting-state functional connectivity (FC) have typically assessed statistical dependencies between different brain regions over the duration of a recording session (usually of several minutes duration), resulting in a single time-averaged value of functional connectivity. While this approach greatly simplifies connectivity analyses, the assumption that interactions between distinct brain regions are temporally stationary over several minutes represents a gross oversimplification of the complex spatio-temporal properties of spontaneous brain activity (Hutchison et al., 2013; Deco et al., 2011). Temporal changes in activity and functional connectivity are a well established phenomenon of electrophysiological recordings. For example, fluctuations in spontaneous activity have been shown to underlie trial-by-trial variability of evoked responses (Arieli et al., 1996; Makeig et al., 2004). Different task demands may also influence whole brain functional connectivity. For instance it has been shown that the patterns and strength of whole brain networks may be modulated by cognitive load during working memory tasks (Luckhoo et al., 2012) or by task relevance (Brookes et al., 2012b). Since the resting state may be thought of as a condition of undirected wakefulness in which the mind occupies a number of mental states, one would similarly expect there to exist dynamic changes in functional connectivity. For example, attempts to modulate internal activity by requesting subjects to recall song lyrics or events from their day have been shown to produce changes in whole brain BOLD functional connectivity (Shirer et al., 2012). Finally, straightforward changes in state, such as eyes open versus eyes closed, affect the spectral properties of spontaneous activity and patterns of functional connectivity (Wu et al., 2010).

These arguments suggest that a better understanding of resting state functional connectivity may be achieved if we are able to relax the assumption that patterns of coactivation are stable over time. Non-stationary analyses of functional connectivity, where functional connectivity is able to vary throughout the duration of the scan, may provide new insights into the temporal dynamics that underlie resting state connectivity. Electrophysiological measures such as MEG and EEG will likely play a key role in characterising these dynamics, since their temporal resolution allows them to capture interactions at time scales that are inaccessible to fMRI.

In this chapter, we provide an overview of recent efforts to develop time-varying analyses of functional connectivity and describe the new challenges associated with assessing non-stationary functional connectivity. With these challenges in mind, we present an analysis of time-varying functional connectivity within a number of known resting state networks. The focus of this work is to determine the time scales and frequency bands at which significant temporal non-stationarity may be observed using MEG.

4.3 Time varying functional connectivity

In recent years there has been a growing interest in characterising temporal changes in resting state functional connectivity. In this section we will review a number of approaches that have been developed and discuss some general limitations that must be considered when assessing non-stationary functional connectivity.

4.3.1 Measures of time varying functional connectivity

4.3.1.1 Sliding window functional connectivity

The most widely used method for assessing time varying connectivity is simply to compute measures of functional connectivity within a sliding window. In these approaches, a time window with a particular duration is defined and those data points within that time window are used to compute a particular functional connectivity metric. The time window is then shifted by a number of samples and another measure of functional connectivity is computed. In this way, a sequence of functional connectivity measures is generated that describes the temporal evolution of functional connectivity over the duration of the recording. The most important consideration is the duration of the window used. Short windows with relatively few samples will result in a less robust estimation of functional connectivity, whereas longer windows will be less able to resolve transient changes in functional connectivity. Furthermore, short time scales will be unable to resolve connectivity between signals below a certain frequency. Choosing an appropriate time window therefore involves consideration of the SNR of the data, the frequencies of interest, and the time scales at which temporal variability is believed to exist.

A number of recent studies have applied sliding window functional connectivity to study resting state connectivity. Using fMRI, Chang & Glover (2010) computed sliding window correlation and coherence between the PCC and nodes of the default mode and dorsal attention network and showed evidence for considerable temporal variability in functional connectivity. de Pasquale et al. (2010) studied the same two RSNs using MEG by applying sliding window correlation to the BLP of nodes of these networks. They demonstrated that periods of high within-network correlation alternate with periods of low correlation on a time scale of approximately ten seconds.

Due to the increase in information yielded by sliding window functional connectivity measures, some studies have summarised the degree of temporal variability by computing further statistics on the functional connectivity time series. For instance, Chang & Glover (2010) computed the temporal standard deviation over the sequence of sliding window correlation coefficients to identify those brain regions with variable PCC connectivity. Allen et al. (2012) defined a "zone of instability" by using an iterative partitioning algorithm to separate brain regions into groups with more or less variable functional connectivity. They found that dorsal attention, default mode and superior occipital regions had more variable connectivity with other regions. The same group further interrogated their sliding window correlations by searching for reproducible transient patterns of whole brain correlation that repeat at different points in time. By applying k-means clustering, they identified several "connectivity states" that showed strong departures from average connectivity patterns. In particular, they showed that DMN connectivity is highly variable over time, with some nodes exhibiting state-dependent interactions with other networks. Using MEG, de Pasquale et al. (2010) developed an algorithm to identify windows of time in which the nodes of a particular network are maximally correlated (with respect to a control region). By comparing cross-network interactions within and outside of these windows, they showed that the DMN interacts most strongly with other networks when its internal correlation is high (de Pasquale et al., 2012).

It should be noted that, given sufficient time points, any conventional functional connectivity analysis may be applied in a time varying fashion. For instance, Kiviniemi et al. (2011) performed spatial ICA on a sliding window basis using time windows of around 100 s.

4.3.1.2 Co-activation patterns

Traditional analyses of resting state activity are typically based on assessing signal covariation over all available time points. It has been shown however, that the pattern of activity corresponding to a particular RSN may be identified using only a few data points that represent periods of time at which the signal amplitude within the network is high (Tagliazucchi et al., 2012). Based on this observation Liu & Duyn (2013) proposed a method for identifying *co-activation patterns (CAPs)* by clustering single BOLD volumes based on their spatial similarity. A key preprocessing step was to first reduce the dataset by retaining only the 15% of frames with the highest amplitude in a predetermined seed region. Using the PCC and left intraparietal sulcus (IPS) as seed regions they identified several CAPs, some of which covered distinct regions within the DMN and DAN and some that corresponded to functionally relevant regions outside of these networks including the visual and motor cortices. Based on these findings, the authors suggested that "interregional BOLD correlations result from instantaneous coactivations (or codeactivations) of multiple brain regions at some critical time points rather than from continuous, sustained interregional neuronal interactions" (Liu & Duyn, 2013).

4.3.1.3 EEG Microstates

Temporal variability in spontaneous activity has long been appreciated in electrophysiological studies using EEG. By visualising spontaneous EEG activity as a series of instantaneous scalp electric field maps, early researchers noted that different topographies would dominate for periods of 100 ms (Lehmann, 1971). These topographies appeared to be quasi-stable, recurring at different points in time throughout the recording. Data driven segmentation of these momentary scalp maps led to the identification of a set of characteristic EEG topographies or *microstates*. Interestingly, relatively few of these microstates were required to describe the observed data; for instance Koenig et al. (2002) found that only 4 maps were required to explain 79% of the variance in the data. EEG microstates have been shown to influence cognition and perception (Britz et al., 2009; Mohr et al., 2005) as well as characterising aspects of spontaneous thoughts (Lehmann et al., 1998). As a result, it has been suggested that microstates may represent the fundamental “building blocks of cognition” or the “atoms of thought” that underlie spontaneous activity.

Recently it has been suggested that EEG microstates may represent the electrophysiological signatures of resting state networks. A number of studies have sought to investigate this by correlating fMRI RSNs with simultaneously acquired EEG recordings. Britz et al. (2010) and Musso et al. (2010) independently demonstrated that BOLD activity associated with a particular microstate identified from the EEG data can reveal spatial patterns with similar structure to a number of ICA derived RSNs. Yuan et al. (2012) further showed evidence of temporal correlation between the fast dynamics of EEG microstates and the slow fluctuations of fMRI RSNs. However, directly relating the EEG topologies of microstates to the underlying anatomy is challenging due to the inhomogeneous conductivity profile of the head, although emerging work suggests that the brain sources associated with microstates do correspond to the same anatomical regions as a number of BOLD RSNs (Yuan et al., 2013).

4.3.2 Issues and limitations of time varying functional connectivity

Since time varying measures of functional connectivity are based on only a subset of the available data, estimates will be more strongly affected by noise and artefacts than measures that are computed over the entire duration of the recording. In MEG, sources of interference such as head motion, eye blinks, muscle contractions and scanner artefacts can result in spatially structured functional connectivity patterns, since they will be recorded to different extents by multiple sensors (as demonstrated in Figure 3.3). Projection into source space will preserve some of this spatial structure. Moreover, since many of these artefacts are transient in nature, techniques such as sliding window correlation will result in temporal changes in connectivity that may be entirely artefactual. Great care must therefore be taken to remove or otherwise account for these sources of interference so as not to incorrectly interpret these fluctuations in functional connectivity as representing the underlying dynamics of neuronal activity.

A further consideration is that estimates of functional connectivity will depend on the properties of the measured signals themselves. For instance, if the amplitude of the MEG signal in one or both regions momentarily decreases, but the noise level stays constant, then the measured functional connectivity during this time window may vary due to the resulting decrease in SNR. Similar changes in functional connectivity can also result from changes in noise (Friston, 2011). Thus accounting for temporal non-stationarity in the power spectra and autocorrelation of the analysed signals is also important in terms of appropriately interpreting dynamic changes in functional connectivity.

Approaches that rely on sliding windows must also choose an appropriate window size. Here there is a trade-off between the ability to detect transient interactions and sensitivity to noise. An important consideration is to ensure that the window size used is large enough to resolve the lowest frequencies of interest. However, due to the approximate 1/f power spectrum of MEG (and fMRI) the signal will be dominated by the lowest resolvable frequency. Moreover, for a given window size, interactions between signals at higher frequencies will be more likely to exhibit temporal non-stationarity within a window (Hutchison et al., 2013). For these reasons, multi-scale analyses (for instance using time-frequency decompositions such as the wavelet transform) may provide a more meaningful characterisation of functional connectivity dynamics.

4.4 Materials and methods

In Section 4.3 we reviewed a number of approaches for assessing time-varying functional connectivity and highlighted some limitations that must be considered when applying these methods. We now proceed to describe our pipeline for assessing time-varying functional connectivity using MEG.

4.4.1 Data acquisition

MEG data were acquired for ten healthy subjects. For each subject, four sessions comprising fifteen minutes of data were acquired giving 10 hours of resting MEG data in total. During each session, the subjects were asked to sit still and loosely fixate on a fixation cross.

MEG data were acquired using a 306 channel Elekta Neuromag system (Helsinki, Finland) comprising 102 magnetometers and 204 planar gradiometers. The data were acquired at a sampling frequency of 1000 Hz with a 0.1 Hz high-pass filter. Localisation of the head within the MEG helmet was achieved using three electromagnetic head position indicator (HPI) coils. By periodically energising these coils their position within the MEG sensor array was identified. Prior to data acquisition, the HPI coil locations, the position of three fiducial points (the nasion, and left and right pre-auricular points), and the head shape were recorded using a three-dimensional digitiser (Polhemus Fastrack). The location of the MEG sensors were coregistered to each individual subject's structural MRI by matching the digitised head surface to the head surface extracted from the anatomical image. Electrodes were also placed on the wrist to record the electrocardiogram (ECG), above and below the eye to record the electrooculogram (EOG), and on the neck and jaw to record the electromyogram (EMG) associated with neck and facial movements. Blinks and saccades were recorded using an Eyelink-1000 infrared eye tracker (SR Research, Osgoode, Canada).

4.4.2 Data preprocessing

The data were visually inspected and any channels containing scanner artefacts were noted. MaxFilter software (Elekta Neuromag) was used to remove external noise using the Signal Space Separation (SSS) method and to down-sample the data to 250 Hz (Taulu et al., 2005). Channels previously noted as bad were explicitly excluded from the SSS algorithm to prevent scanner artefacts from propagating over other channels. The filtered and down-sampled data were converted to SPM8 format (`http://www.fil.ion.ucl.ac.uk/spm`) and visually inspected to identify periods of data corresponding to obvious artefacts or with abnormally high variance. The channels and time samples flagged as bad were excluded from all subsequent analysis steps. Independent component analysis (ICA) (`http://research.ics.aalto.fi/ica/fastica/`) was used to decompose the sensor data into temporally independent components (tICs), which were classified into brain and artefact components. Classification was based on looking for characteristics in the tICs that were indicative of the cardiac cycle, eye blinks or saccades, or mains noise, as described in Section 3.4.3.2. As a final preprocessing step, muscle artefacts were rejected following the procedure in Section 3.4.3.3.

4.4.3 Identification of functional networks

The aim of this study was to investigate how functional connectivity within resting state networks varies over time. A prerequisite was therefore to identify a set of ROIs that correspond to functionally connected locations. One option would be to follow the approach of de Pasquale et al. (2010) (and employed in Section 3.4.4.1) by using regions corresponding to fMRI RSNs. However, since it has been shown that equivalent RSNs may be found from MEG, we opted to use ROIs based on MEG RSNs (Brookes et al., 2011b). To this end, we first performed a whole brain connectivity analysis using group temporal ICA (Hyvärinen & Oja, 2000), following the method of (Brookes et al., 2011b).

Beamforming was used to provide an estimate of the neural activity at every vertex of a regular 6 mm grid spanning the entire brain. Prior to beamforming, the data were first band-pass filtered into the beta frequency band (8 - 30 Hz), since the alpha and beta bands have been previously shown to yield network maps with good correspondence to those found with fMRI (Brookes et al., 2011b). Following beamforming, the oscillatory envelope at each voxel was derived using the Hilbert transform and temporally averaged within sliding windows with a width of 1 s and 50% overlap between consecutive windows. This process is equivalent to applying a low-pass filter and down-sampling. Averaging the envelope within 1 s windows has been shown to optimally enhance estimates of resting state functional connectivity and down-sampling reduces the amount of data that must be input to the ICA algorithm (Luckhoo et al., 2012). The envelopes were temporally concatenated across all sessions and subjects. These concatenated data were reduced to 25 dimensions using principal component analysis (PCA) and then separated into 25 independent time courses using temporal ICA, each with a corresponding spatial map which is a spatial representation of one column of the ICA mixing matrix.

Each of the 25 spatial maps were visually inspected and components representing the default mode network and the sensorimotor network were selected based on their correspondence with previously reported fMRI and MEG topographies. The spatial maps corresponding to these networks are shown in Figure 4.1. The default mode network includes right and left inferior parietal lobule (IPL), right and left medial frontal gyrus (MFG), and the posterior cingulate cortex (PCC).

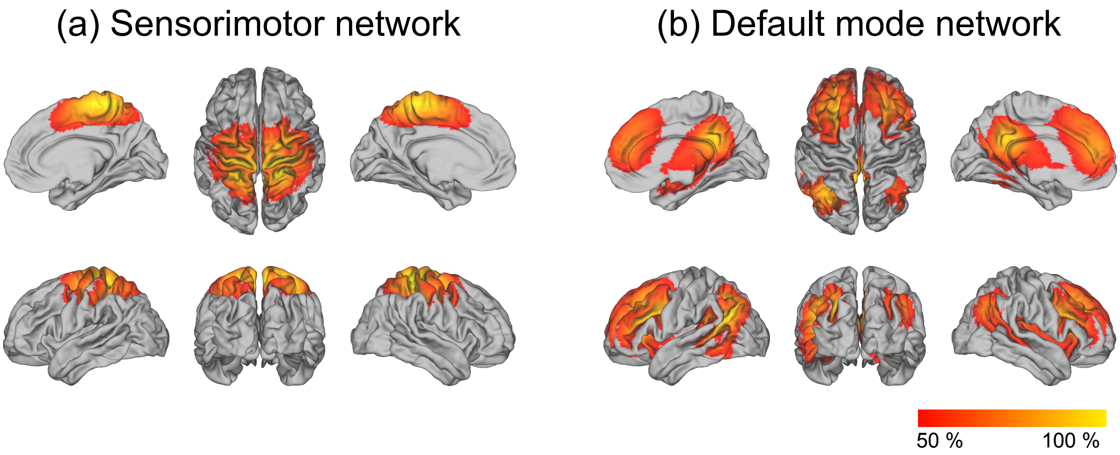

Figure 4.1: Group average ICA spatial maps corresponding to the sensorimotor network (a) and default mode network (b) derived by applying ICA to amplitude envelopes of 8-30 Hz beamformed data. Each map been thresholded to show only values between 50% and 100% of the value of the maximum voxel.

4.4.3.1 Defining session-specific ROIs from the group RSNs

Concatenation of source activity estimates across all subjects allows for a more robust estimation of functionally meaningful networks, as it ensures that the components found using ICA are consistent across multiple subjects. However, the spatial map for a given component will correspond to the group average of the component and will not necessarily be a true spatial representation of the network for a particular subject and session. To obtain session-specific network maps the portion of the tIC time courses corresponding to that particular session was identified and these truncated time courses were used in a general linear modelling (GLM) analysis (Friston et al., 1996; Brookes et al., 2004; Woolrich et al., 2009). Specifically, we performed a multiple linear regression at each voxel, with the down-sampled envelope time course for the n^{th} session $\mathbf{Y}_n$ as the dependent variable, and the truncated ICs $\mathbf{X_n}$ as the independent variables.

$$\mathbf{Y}_n = \mathbf{X}_n \boldsymbol{\beta}_n + \boldsymbol{\epsilon} \tag{4.1}$$

This yields a set of regression parameter estimates $\boldsymbol{\beta}_n$ that explain the contribution of the envelope time course at that voxel to each group level ICA component. Hence, the spatial map over all voxels of the regression parameter corresponding to a network of interest provides a session-specific network map.

Each spatial map represents a distributed network. Since we were interested in within-network functional connectivity we manually split these networks into their constituent network nodes (a node is defined here as a spatially distinct cluster of voxels). These nodes represent the ROIs used for subsequent time-varying connectivity analysis. A schematic showing the method for identifying these session-specific networks is shown in Figure 4.2. The mean and standard error of the MNI coordinates corresponding to each network node are provided in Table 4.1.

4.4.4 Reconstruction of ROI time series

Having identified a set of ROIs that correspond to functional subregions the next step is to assign to each a time series that best represent the activity within that ROI. We took the set of voxels corresponding to each ROI and reconstructed the

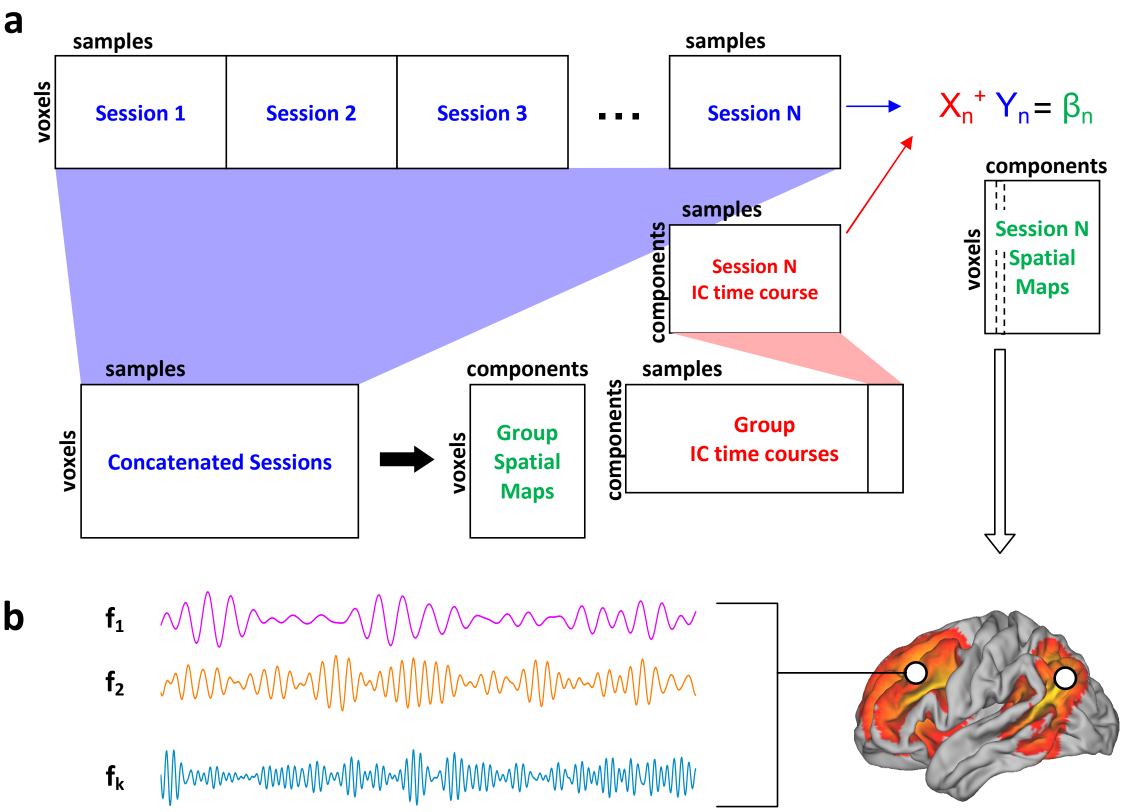

Figure 4.2: (a) Identification of session-specific ROIs. Low-pass filtered and down-sampled amplitude envelopes are temporally concatenated across all sessions. The concatenated time series at each voxel are decomposed using ICA into 25 temporally independent component time courses and spatial maps (mixing matrix). The independent time courses for a particular session n are included as independent variables $\mathbf{X}$ in a GLM with the low-pass filtered and down-sampled envelope for that session as the dependent variable $\mathbf{Y}$. Session-specific spatial maps are given by the regression parameter estimates $\boldsymbol{\beta}_n$. (b) Example session-specific spatial map for the default mode network. Band limited activity at nodes of the network are reconstructed using beamforming for each k frequency bands at the MNI coordinates corresponding to the centre of each node (white circles).

activity for each voxel using a beamformer. This set of time courses then provides a description of activity within the ROI. Since the subsequent analysis is based on pairwise measures of functional connectivity, we opted to summarise each ROI by a single time course (but see Brookes et al., 2014, for a multivariate extension that allow functional connectivity to be assessed between ROIs spanning multiple voxels). There are a number of possible ways to define a single ROI time course.

Network	Region	ROI centre: mean ± std (mm)		
		x	y	z
Default Mode	Right inferior parietal lobule	46.8 ± 6.3	-60.7 ± 8.2	31.0 ± 8.7
	Right medial frontal gyrus	32.5 ± 6.2	25.8 ± 9.1	35.1 ± 7.5
Default Mode	Left inferior parietal lobule	-41.9 ± 6.0	-60.1 ± 5.0	29.6 ± 8.1
	Left medial frontal gyrus	-30.0 ± 3.9	28.1 ± 12.8	29.9 ± 12.9
Sensorimotor	Right sensorimotor cortex	29.0 ± 6.6	-25.8 ± 9.3	59.0 ± 5.6
	Left sensorimotor cortex	-29.6 ± 6.9	-28.4 ± 7.7	58.5 ± 3.6

Table 4.1: MNI coordinates of the three node pairs studied. The mean and standard deviation are across all sessions.

Activity across the entire ROI may be captured by simply averaging the time series for all voxels together. However, in the case of source reconstructed MEG data, straightforward averaging across an ROI should be avoided due to ambiguity in the sign at each voxel (this ambiguity arises out of the fact that a dipolar current source will give rise to equivalent magnetic field patterns regardless of its polarity). Thus the source estimates at each voxel within an ROI will not necessarily have a consistent sign and averaging across these time series will be meaningless (Hillebrand et al., 2011). An alternative approach is to perform a principal component analysis across the ROI and retain only the first component (Supp et al., 2007). However in this work we took a more straightforward approach of assuming that the activity within an ROI may be adequately summarised by the time course at a single voxel in the centre of the ROI. For functionally derived

ROIs this assumption is likely to be reasonable, since each ROI will correspond to a functionally distinct subregion. Moreover, given the large point spread function of MEG source estimates, a single voxel's time course will capture much of the variance across the ROI, provided the regions are relatively small.

4.4.5 Estimating time varying functional connectivity

In order to assess the temporal dynamics of functional connectivity, the activity at nodes of the RSNs defined previously was reconstructed across multiple frequency bands before computing time varying functional connectivity between pairs of these nodes. This process may be considered as comprising 3 steps:

1. Source reconstruction of virtual electrode signals by reapplying the beamformer across multiple frequency bands for nodes of the network.
2. Computation of the signal amplitude at each node as a function of time and frequency using the wavelet transform.
3. Computation of time varying functional connectivity over all frequencies.

Each of these steps will now be discussed in detail.

4.4.5.1 Source reconstruction of ROI time series

In order to investigate changes in time varying connectivity, the source activity was estimated at ROIs corresponding to the RSN nodes, again using beamforming. As before, the LCMV beamformer was used, and the covariance matrix was computed over the whole duration of the scan. However, whereas the beamformer used to define the RSNs was based only on the alpha and beta frequency bands, here we were interested in connectivity at multiple frequency bands over a wider range. The ability of the beamformer to accurately reconstruct sources of activity is dependent on the data covariance matrix used. As such the data used to define the covariance matrix should be matched as well as possible to the data being analysed. On the other hand, the estimation of the data covariance matrix can be improved by pooling over data from multiple frequency bands (Brookes et al., 2008). Hence there is a choice of whether to use one wideband covariance matrix that includes the entire frequency range, or to use multiple narrower bands. However, Brookes et al. (2008) showed that severe errors on the estimated covariance matrix were only found for below around 30 s of data (for a bandwidth of 5 Hz), so this should not be considered a problem for the fifteen minutes of resting data used here. As a result, we chose to tune the beamformer to each frequency band analysed by using a separate covariance matrix to compute the weights for each frequency band. The centre frequency of the band was the frequency corresponding to the wavelet scale, and the bandwidth was set to capture the frequency resolution of the wavelet transform. The wavelet transform naturally trades off temporal and spectral resolution such that higher frequencies have better temporal resolution and poorer spectral resolution. Therefore the bandwidth of the data used to construct the covariance matrix was greater for higher frequencies. Fifteen frequency bands were used spanning the range 4 to 48 Hz with logarithmic spacing. The width of the bands used to beamform the data were set such that each captured the frequency resolution of the wavelet transform at that frequency. Since the frequency resolution is inversely proportional to the centre frequency, wider bands were used for higher frequencies. The frequency bands used are shown in Figure 4.3a.

4.4.5.2 The wavelet transform

After source reconstruction, the beamformed time courses were decomposed into a time-frequency representation using the continuous Morlet wavelet transform. The

Morlet wavelet may be considered as a sinusoid, localised in time by a Gaussian, where the scale parameter S controls the time length of the wavelet, and the Morlet factor ω_0 governs the number of oscillations of the sinusoid within the wavelet. Thus for a given ω_0 the scale is inversely proportional to frequency and from this point on we will use these terms interchangeably. The parameter ω_0 controls the trade off between time and frequency resolution, and was here chosen as $\omega_0 = 6$ which provides a good compromise between temporal and spectral resolution (Grinsted et al., 2004). The wavelet transform was chosen over other spectral estimation approaches as it is well suited to non-stationary data and combines spectral decomposition with power envelope estimation, although other time-frequency decompositions such as the Hilbert transform may work equally well with the appropriate settings.

4.4.5.3 Time varying functional connectivity

Time varying functional connectivity was assessed by computing sliding window band-limited power correlations, i.e. the correlation between the envelopes of the two signals of interest, $\mathbf{x}$ and $\mathbf{y}$, within a sliding time window:

$$\mathbf{C_{xy}}\left(\mathbf{t},\mathbf{f}\right) = \frac{\sum_{\mathbf{t}-\frac{T_f}{2}}^{\mathbf{t}+\frac{T_f}{2}} \left(\mathbf{P_x}\left(\mathbf{t},\mathbf{f}\right) - \overline{\mathbf{P_x}}\left(\mathbf{t},\mathbf{f}\right)\right)\left(\mathbf{P_y}\left(\mathbf{t},\mathbf{f}\right) - \overline{\mathbf{P_y}}\left(\mathbf{t},\mathbf{f}\right)\right)}{\sum_{\mathbf{t}-\frac{T_f}{2}}^{\mathbf{t}+\frac{T_f}{2}} \left(\mathbf{P_x}\left(\mathbf{t},\mathbf{f}\right) - \overline{\mathbf{P_x}}\left(\mathbf{t},\mathbf{f}\right)\right)^2 \sum_{\mathbf{t}-\frac{T_f}{2}}^{\mathbf{t}+\frac{T_f}{2}} \left(\mathbf{P_y}\left(\mathbf{t},\mathbf{f}\right) - \overline{\mathbf{P_y}}\left(\mathbf{t},\mathbf{f}\right)\right)^2} \tag{4.2}$$

where $\mathbf{P_x}\left(\mathbf{t},\mathbf{f}\right)$ is the magnitude of the wavelet transform of the signal $\mathbf{x}$ in the time window $\mathbf{t} - \frac{T_f}{2} < \mathbf{t} \leq \mathbf{t} + \frac{T_f}{2}$ and the over bars denote the mean of this quantity over that interval. The size of the window T_f is a function of frequency such that the window size in seconds is given by $T_f = \frac{\Delta}{f}$ where Δ is a constant that corresponds to a fixed number of periods at any frequency. The rationale for using a frequency dependent window size, rather than a window of fixed duration, is that the effective number of samples in the power envelope of the signal is inversely proportional to the period of the underlying carrier wave. The uncertainty in the value of the correlation coefficient depends on the number of samples used in the estimate. Using a frequency dependent window length therefore calibrates this uncertainty such that it is the same across all frequency bands. The relationship between frequency and window size is shown in Figure 4.3b.

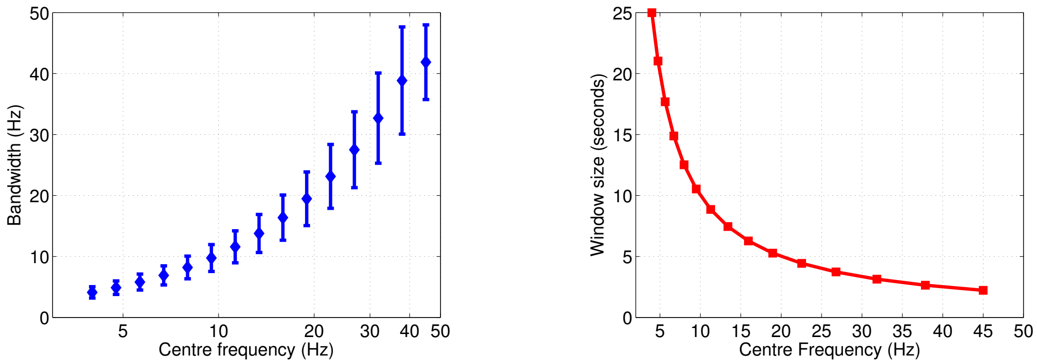

(a) (b)

Figure 4.3: (a) Frequency bands used for beamforming. Centre frequencies are logarithmically spaced from 4 to 48 Hz. The width of the bands used to beamform the data were set such that each captured the frequency resolution of the wavelet transform at that frequency. Since the frequency resolution is inversely proportional to the centre frequency, wider bands are used for higher frequencies. (b) Size of the sliding window used for estimating time-varying functional connectivity versus centre frequency for $\Delta = 100$. Since $\Delta = 100$ corresponds to a fixed number of cycles for any frequency, the window size in seconds is inversely proportional to frequency.

4.4.6 Functional connectivity summary statistics

The sliding window correlation method described above was applied separately to the pair of nodes for each network and recording session, resulting in a time-frequency representation of the power correlation between each pair of brain regions. In order to characterise the functional connectivity (including the extent of its temporal non-stationarity) in different networks and frequency bands we defined a set of summary metrics that are calculated separately for each pair of brain areas and for each frequency band. With the exception of the first measure, the *Average stationary FC*, these are computed by collapsing the sliding window BLP correlations over the time dimension.

4.4.6.1 Average stationary FC

Temporally stationary functional connectivity was computed for each frequency band as the correlation between the two signal power envelopes over all time points. The power envelopes were first temporally down-sampled within one second windows as per (Brookes et al., 2011b). The two signals are equivalent to those input to ICA in order to define the network ROIs, apart from the frequency band. This measure differs from the *average within-window FC* measure in that the envelopes are averaged within time windows and then correlated over all windows, whereas the *average within-window FC* measure computes the correlation within each time window and then averages.

4.4.6.2 Average within-window FC

The average functional connectivity was assessed by computing the mean of the time varying envelope correlation computed over all time points for a particular frequency. This measure captures the average level of non-stationary functional connectivity over the duration of the scan.

4.4.6.3 Within-window FC variability

Rather than exploring the average level of functional connectivity, we were primarily interested in how the functional connectivity varies over time. To this end, we have used the standard deviation of the time varying envelope correlation as a summary measure of how variable this coupling is.

4.4.6.4 Bimodality

The standard deviation of the time varying connectivity is a measure that will be sensitive to the amount of temporal non-stationarity, but it is also of interest to determine the nature of the underlying distribution of correlation values. Increases in standard deviation could be due to a continuum of dynamic states of varying amounts of functional connectivity, which would cause increased spreading (and therefore increase standard deviation) of a unimodal distribution of functional connectivity. However, an alternative hypothesis is that there are a limited number of states, the most parsimonious description being that each pair of brain areas has a bi-state relationship; where one state is characterised by low functional connectivity, and the second is characterised by high functional connectivity. Under this hypothesis, we would expect a mixture of distributions, one with a mean close to zero, representing low functional connectivity, and a second with a positive mean, representing the average correlation during periods of high functional connectivity. To explore this bimodality, we used two different measures:

Skewness Since the networks under consideration are positively correlated, we are specifically looking for an increase in the positive tail. To this end we have used the skewness (the third moment) of the distribution as a measure of bimodality.

BIC ratio This is the ratio between the approximate Bayesian model evidences (Bayes Information Criterion) when inferring a one-component versus a two-component Gaussian mixture model; this is a more explicit measure of bimodality than the skewness.

4.4.7 Significance testing of functional connectivity statistics

As previously discussed, spurious functional connectivity may arise due to signal leakage between nearby voxels, or differences in power across different frequency bands. Furthermore, apparent temporal variability in neuronal coupling, as revealed by time-varying functional connectivity measures may simply be the result of random fluctuations in the two time series. In order to appropriately interpret the measures described in Section 4.4.6, it is necessary to perform significance testing against an appropriate null hypothesis.

Significance testing of the statistics in Section 4.4.6 was performed using a Monte Carlo procedure based on simulating surrogate data under a particular null hypothesis. For temporally stationary functional connectivity, the null hypothesis is that the observed correlation is due entirely to signal leakage and/or increased SNR within a particular frequency band. As in Section 3.3.1.1, signal leakage was accounted for by first orthogonalising one signal with respect to the other prior to computing their envelopes. It was therefore unnecessary to model this in the surrogate data. To generate surrogate amplitude envelopes with the same spectral characteristics as the real envelopes, a procedure based on phase randomisation was used. This results in data with the same power spectra as the real data, but with random phase such that any genuine correlation is destroyed. For each pair of ROIs, multiple realisations of surrogate envelope pairs were generated using phase randomisation of the original envelopes. From these simulated envelope data, the same stationary functional connectivity measures were computed as for the real data, resulting in a null distribution of correlation values.

For measures of temporal non-stationarity in the functional connectivity (e.g. the *within-window FC variability* measure), it is also necessary to test whether the observed variability is significantly greater than that which would be expected if the coupling between the two signals was stationary about a given mean i.e. that the variation cannot be solely explained by the uncertainty in the estimated correlation value. This was achieved by augmenting the phase randomisation with a linear mixing procedure that ensured that the stationary correlation between each pair of surrogate envelopes matched that of the real data. Measures of functional connectivity variability were then computed from these simulated envelopes.

4.4.7.1 Phase randomisation

For each ROI pair and frequency band, 500 realisations of surrogate envelope pairs were generated using phase randomisation of the original envelopes following the method of Prichard (1994).

For a univariate (envelope) time series $\mathbf{p_x}(\mathbf{t})$ (where $\mathbf{p_x}(\mathbf{t})$ is a vector of time points representing the amplitude envelope for a single frequency of $\mathbf{P_x}(\mathbf{t}, \mathbf{f})$) of length N samples, the Fourier transform $\boldsymbol{P}_{\mathbf{x}}(\boldsymbol{\omega})$ may be obtained by applying the discrete Fourier transform operator $\mathcal{F}$:

$$\boldsymbol{P}_{\mathbf{x}}(\boldsymbol{\omega}) = \mathcal{F}\left\{\mathbf{p_x}(\mathbf{t})\right\} = \sum_{n=0}^{N-1} \mathbf{p_x}(t_n) e^{i\boldsymbol{\omega} n \Delta t} \tag{4.3}$$

This complex valued Fourier transform may be expressed in terms of its amplitude $\mathbf{A}(\boldsymbol{\omega})$ and phase $\boldsymbol{\phi}(\boldsymbol{\omega})$ as:

$$\mathbf{P_x}(\boldsymbol{\omega}) = \mathbf{A}(\boldsymbol{\omega}) e^{i\boldsymbol{\phi}(\boldsymbol{\omega})} \tag{4.4}$$

A phase randomised Fourier transform may be obtained by rotating the phase $\boldsymbol{\phi}(\boldsymbol{\omega})$ of each frequency $\boldsymbol{\omega}$ by a random angle $\boldsymbol{\xi}(\boldsymbol{\omega})$, which is chosen uniformly in the range $0 < \boldsymbol{\xi}(\boldsymbol{\omega}) < 2\pi$ radians. The phase randomised Fourier transform is then given by:

$$\widetilde{\boldsymbol{P}_{\mathrm{x}}}(\boldsymbol{\omega}) = \mathbf{A}(\boldsymbol{\omega})e^{i(\boldsymbol{\phi}(\boldsymbol{\omega})+\boldsymbol{\xi}\boldsymbol{\omega})} \tag{4.5}$$

The phase randomised envelope $\widetilde{\mathbf{p}_{\mathbf{x}}}(\mathbf{t})$ may then be obtained by applying the inverse Fourier transform:

$$\widetilde{\mathbf{p}_{\mathbf{x}}}(\mathbf{t}) = \mathcal{F}^{-1}\left\{\boldsymbol{P}_{\mathrm{x}}(\boldsymbol{\omega})e^{i\boldsymbol{\xi}(\boldsymbol{\omega})}\right\} \tag{4.6}$$

and similarly for a second envelope $\widetilde{\mathbf{p}_{\mathbf{y}}}(\mathbf{t})$ by rotating the phase by a random angle $\boldsymbol{\zeta}(\boldsymbol{\omega})$:

$$\widetilde{\mathbf{p}_{\mathbf{y}}}(\mathbf{t}) = \mathcal{F}^{-1}\left\{\boldsymbol{P}_{\mathrm{y}}(\boldsymbol{\omega})e^{i\boldsymbol{\zeta}(\boldsymbol{\omega})}\right\} \tag{4.7}$$

From these expressions it is evident that the magnitude of each Fourier component is the same as for the original data. Furthermore, by the Wiener-Khintchine theorem, the autocorrelation of the surrogate data is preserved (Prichard, 1994).

4.4.7.2 Preserving the distribution of the surrogate envelopes

One limitation of the phase randomisation approach is that it does not preserve the distribution of the original data since the surrogate data will be Gaussian distributed. For original data that is approximately Gaussian distributed this would not present an issue. However, since we are using the amplitude envelopes, our data is strongly non-Gaussian and therefore the surrogate data would not accurately model our data. To ensure that the distribution of the data was preserved we applied a power transform to the original envelopes $\mathbf{p}_{\mathbf{x}}(\mathbf{t})$ and $\mathbf{p}_{\mathbf{y}}(\mathbf{t})$ prior to performing phase randomisation. The power transform used was the Box-Cox transformation:

$$\mathbf{p}'_{\mathbf{x}}(\mathbf{t})_\lambda = \begin{cases} \dfrac{\mathbf{p}_{\mathbf{x}}(\mathbf{t})^\lambda - 1}{\lambda}, & \lambda \neq 0 \\ \log \mathbf{p}_{\mathbf{x}}(\mathbf{t}), & \lambda = 0 \end{cases} \tag{4.8}$$

and equivalently for $\mathbf{p}_{\mathbf{y}}(\mathbf{t})$. The parameter λ is the power to which the data must be raised to ensure a normal distribution. Surrogate envelopes $\widetilde{\mathbf{p}'_{\mathbf{x}}}(\mathbf{t})$ and $\widetilde{\mathbf{p}'_{\mathbf{y}}}(\mathbf{t})$ were then generated by randomising the phase of these transformed envelopes. The surrogate envelopes retain the power spectrum and autocorrelation of the power-transformed envelopes, but have an approximately Gaussian distribution. To generate envelopes with the same distribution as the original envelopes, the surrogate envelopes were raised to the inverse of λ:

$$\widetilde{\mathbf{p}_{\mathbf{x}}}(\mathbf{t}) = \widetilde{\mathbf{p}'_{\mathbf{x}}}(\mathbf{t})^{\frac{1}{\lambda}} \tag{4.9}$$

This ensures the distribution of the surrogate envelopes is approximately equal to that of the original data.

4.4.7.3 Preserving stationary functional connectivity

For the purpose of testing the significance of temporal variability in functional connectivity, we simulated pairs of surrogate envelopes which maintained the same linear correlation (i.e. stationary functional connectivity) as the real data. This may be achieved by setting $\boldsymbol{\zeta}(\boldsymbol{\omega}) = \boldsymbol{\xi}(\boldsymbol{\omega})$ for all $\boldsymbol{\omega}$ such that that the phase differences (and hence cross-correlations) between the two signals $\widetilde{\mathbf{p}_{\mathbf{x}}}(\mathbf{t})$ and $\widetilde{\mathbf{p}_{\mathbf{y}}}(\mathbf{t})$ are equal to those of the original data. However, the power transformation procedure described above destroys the linear relationship between $\widetilde{\mathbf{p}_{\mathbf{x}}}(\mathbf{t})$ and $\widetilde{\mathbf{p}_{\mathbf{y}}}(\mathbf{t})$. Instead, to ensure the correlation between the envelopes matched that of the real data, these two signals were linearly mixed:

$$\widetilde{\mathbf{p}_{\mathbf{x}}}(\mathbf{t}) = (1-\alpha)\widetilde{\mathbf{p}_{\mathbf{x}}}(\mathbf{t}) + \alpha\widetilde{\mathbf{p}_{\mathbf{y}}}(\mathbf{t}) \tag{4.10}$$

where α is a mixing coefficient that was set using an iterative scheme by minimising the magnitude of the error between $corr(\mathbf{p}_{\mathbf{x}}(\mathbf{t}), \mathbf{p}_{\mathbf{y}}(\mathbf{t}))$ and $corr(\widetilde{\mathbf{p}_{\mathbf{x}}}(\mathbf{t}), \widetilde{\mathbf{p}_{\mathbf{y}}}(\mathbf{t}))$. Under the assumption that the distributions of $\mathbf{p}_{\mathbf{x}}(\mathbf{t})$ and $\mathbf{p}_{\mathbf{y}}(\mathbf{t})$ are approximately the same, this procedure ensures that the surrogate envelopes maintain the same correlation, power spectrum and distribution as the real envelopes.

For each node pair, 500 realisations of simulated data were generated. From these surrogate envelopes, the same measures of connectivity were computed as for the real data, resulting in a null distribution of connectivity measures.

4.4.7.4 Significance testing

To assess significance at a particular frequency, the functional connectivity summary statistics were compared with those computed from the Monte Carlo simulations. Under the null hypothesis, one would expect the measured functional connectivity not to differ significantly from that which would be expected from random fluctuations in the connectivity estimate. To test this, a Z-statistic was computed for each summary statistic by subtracting the mean and dividing by the standard deviation (over realisations) of the null distribution corresponding to that particular summary measure. These Z-statistics were averaged all sessions and subjects resulting in a group averaged Z-statistic. The threshold for significance at the group level is then given by the Z-statistic corresponding to a p-value of 0.05 (Z = 1.645). This approach to assessing group level significance is equivalent to that used in Brookes et al. (2011a) and corresponds to a fixed-effects analysis.

4.5 Results

4.5.1 Single subject example

Figure 4.4a shows an example of the wavelet power correlation ($\Delta = 100$) computed between two nodes of the default mode network (rIPL and rMFG) for a single session. Qualitatively transient periods of high correlation are revealed, indicated by red patches in the time-frequency plane. The standard deviation of the wavelet power correlation computed over all time points for each frequency bin is shown in Figure 4.4b against the significance threshold, which corresponds to the 95th percentile of the null distribution. Values exceeding this significance threshold correspond to a standard deviation over time that exceeds that which could be expected from random fluctuations about a given temporally stationary average correlation. Figure 4.4c shows the distributions of correlation values over all time points for the 6 Hz, 14 Hz and 39 Hz bands and the individual components of a two component Gaussian mixture model fit. Frequencies with highly variable time-varying correlation (e.g. 14 Hz) have a distribution that is more bimodal and positively skewed than those frequencies at which non-significant variability is observed.

4.5.2 Stationary vs. non-stationary estimates of functional connectivity

Figure 4.5 compares group averaged estimates of functional connectivity based on temporally stationary envelope correlation (green) with two measures derived from time varying measures of connectivity (light and dark blue) as a function of frequency. The window parameter Δ was set as 100 (corresponding to 100 oscillatory cycles). For the temporally stationary envelope correlation, maximum functional

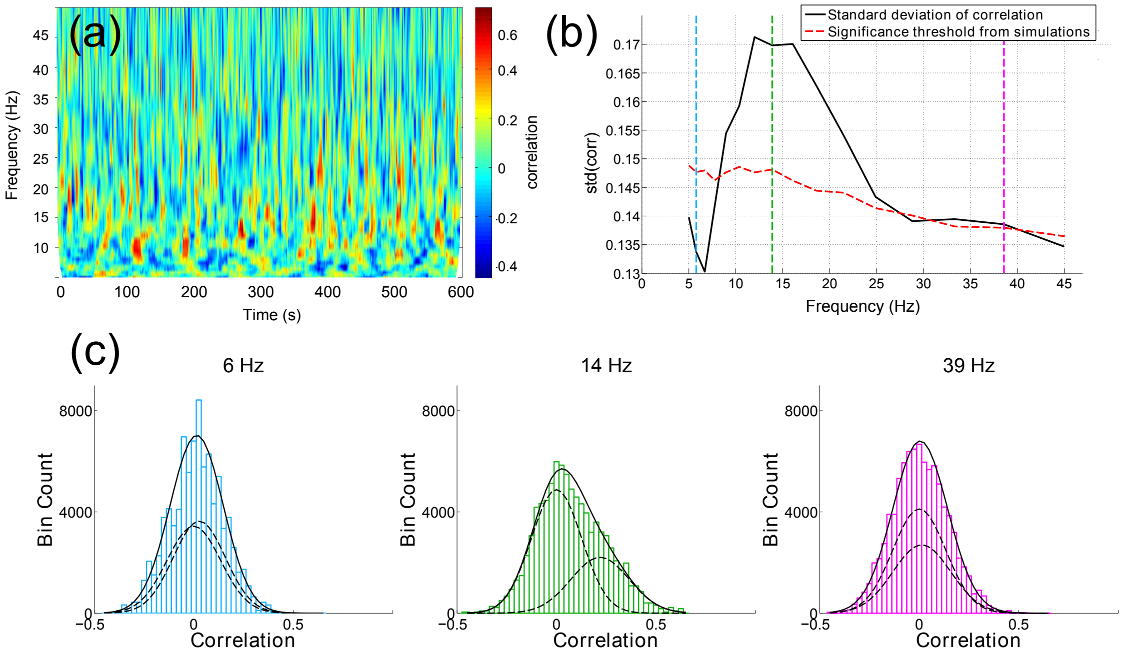

Figure 4.4: Example of the wavelet power correlation between nodes of the default mode network (rIPL and rMFG) over 10 minutes for a single subject. Transient periods of high correlation are revealed, indicated by red patches in the time-frequency plane. (b) Plot showing the standard deviation of the wavelet power correlation computed over all time points for each frequency bin (black) against the significance threshold which corresponds to the 95th percentile of the null distribution. Values exceeding this significance threshold correspond to a standard deviation over time that exceeds that which could be expected from random fluctuations about a given average correlation. (c) Histograms of correlation values over all time points at 6 Hz (blue), 14 Hz (green) and 39 Hz (magenta) overlaid with a two component Gaussian mixture model fit (overall fit in bold and individual components in dashed). Frequencies with highly variable time-varying correlation show a bimodal distribution.

connectivity is observed in the alpha and beta bands, consistent with previous findings. A similar relationship between frequency and functional connectivity is also revealed by taking the average value of the time varying estimate of envelope correlation *average within-window FC*; light blue), indicating that this functional connectivity may also be detected at shorter time scales than the low frequency fluctuations that the *average stationary FC* measure will be sensitive to.

More interesting is that the standard deviation of the time varying envelope correlation also displays a prominent peak at the same frequencies, demonstrating that the functional connectivity measured using this method is significantly non-stationary. Since this significance threshold is based on simulations that control for an increase in the average envelope correlation, this observed increase in variability is above and beyond that which could be expected from uncertainty in the estimate alone. This result suggests that correlation observed between the two power time courses are not stationary over time, but are rather driven by transient periods of high correlation.

4.5.3 Increased variability in functional connectivity is driven by transient correlations

The threshold for significant variability in functional connectivity is such that it cannot be explained simply by an increase in the average connectivity. However, the possibility remains that the observed variance in the correlation time course results from a broadening of the distribution rather than from transient synchronisation events. To discriminate these scenarios we used measures of bimodality, namely the *skewness* and the *BIC ratio*. Figure 4.6 is equivalent to Figure 4.5 but now shows the *within-window FC variability* (standard deviation) alongside these two bimodality measures. There is a strong correspondence between frequencies associated with high *within-window FC variability* and those associated with an increase in *skewness* and *BIC ratio*, although the latter is not significant. This result shows that the distribution of correlation values does not just broaden when high standard deviation is seen but becomes increasingly bimodal. This observation is consistent with the existence of transient periods of high correlation.

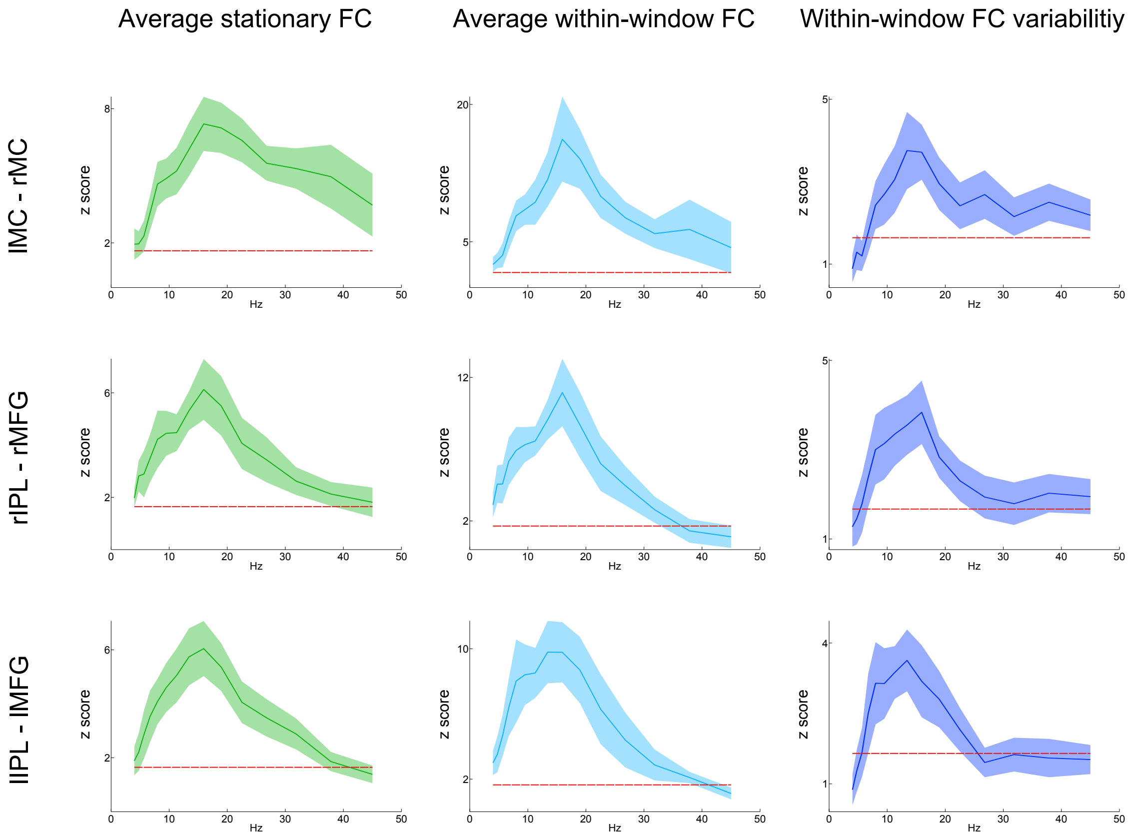

Figure 4.5: Group averaged Z-statistics for the three functional connectivity measures computed at each frequency between nodes of the default mode and sensorimotor neworks. In each case the line and shaded region represents the mean and standard error respectively across all subjects. The dashed red line shows the 95% confidence limit derived from simulations. For the *average stationary FC* and *average within-window FC* measures, the Z-statistics correspond to a test of the null hypothesis that there is no significant power correlation. For the *within-window FC variability* measure, the z-statistics correspond to testing the null hypothesis that variability in the time course of power correlation may be explained by random fluctuations in the estimated correlation alone. Z-statistics above the threshold therefore indicate that there exists significant variability in functional connectivity within the network at a given frequency band.

4.5.4 Effect of time window Δ

Figure 4.7 shows the effect of varying the time window Δ on the *within-window FC variability* Z-statistic for each of the three networks. The parameter Δ corresponds to the number of oscillatory cycles over which the power correlation is computed. For each network, the window size that best captures an increase in standard

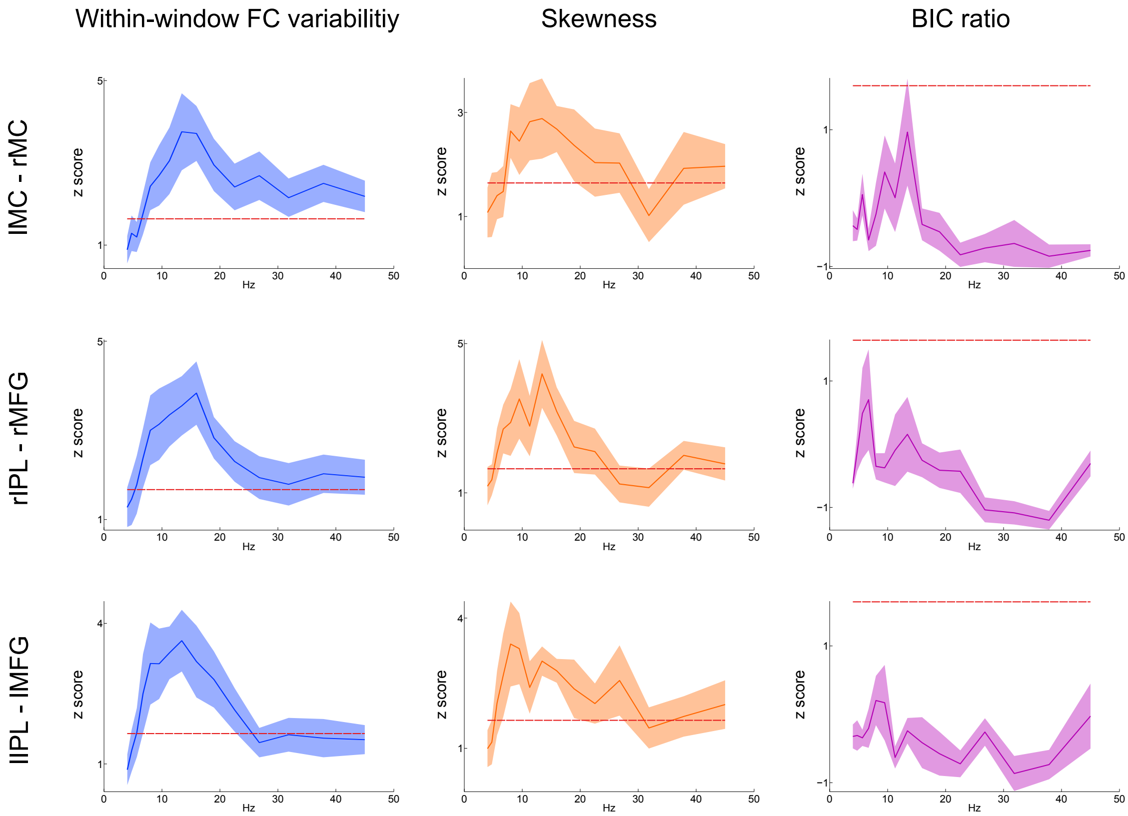

Figure 4.6: Group averaged Z-statistics for the two bimodality measures compared to the *within-window FC variability* measure, computed at each frequency between nodes of the default mode and sensorimotor neworks. In each case the line and shaded region represents the mean and standard error respectively across all subjects. The red line shows the 95% confidence limit derived from simulations. The Z-statistics correspond to testing the null hypothesis that variability in the time course of power correlation may be explained by random fluctuations in the estimated correlation alone. There is a strong correspondence between frequencies exhibiting high standard deviation in time varying connectivity and the *skewness* and *BIC* ratio measures (however, only the *skewness* measure is significant). This correspondence demonstrates that variability in the functional connectivity arises due to an increase in positive correlation values, rather than from a broadening of the distribution.

deviation is in the range 50 - 250 cycles. Below 50 cycles, it is not possible to identify an increase in variability above that which is expected from random fluctuations in the power of both signals.

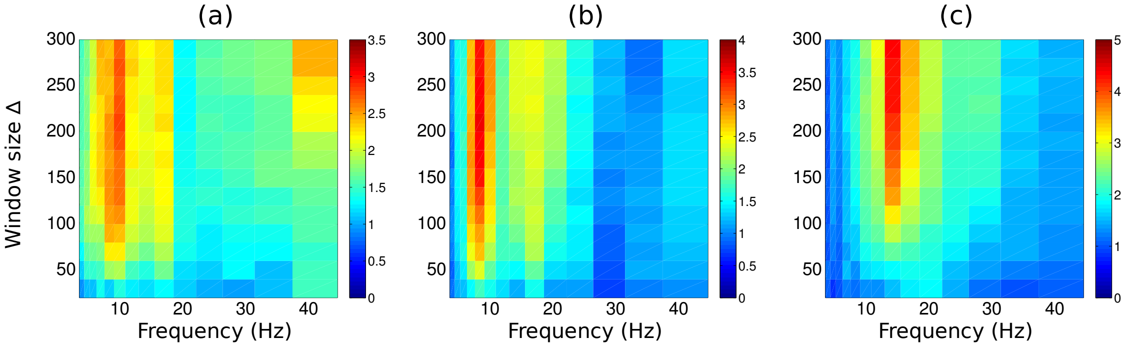

Figure 4.7: Effect of window size Δ on the group averaged *within-window FC variability* Z-statistic computed between (a) rIPL and rMFG, (b) lIPL and lMFG, (c) rMC and lMC. The parameter Δ corresponds to the number of oscillatory cycles at any frequency. For each network the window size that best captures an increase in *within-window FC variability* is in the range 50 to 250 cycles. Below 50 cycles, it is not possible to identify an increase in variability above that which is expected from random fluctuations in the power of both signals.

4.6 Discussion

This chapter has presented a method based on the wavelet transform to investigate variations in power correlations across time and frequency. We demonstrate the use of this method for assessing the temporal dynamics of resting state connectivity between nodes of the default mode and the sensorimotor networks. By computing summary measures that capture the temporal variability of these correlations, we show that the frequency bands that have previously been shown to have the strongest band-limited power correlations have significantly non-stationary temporal dynamics, when compared with measures computed from simulations of realistic null data with imposed stationarity. These results confirm the findings reported qualitatively elsewhere that MEG RSNs based on signal envelope correlations are markedly non-stationary (de Pasquale et al., 2010, 2012). By varying the size of the window used to compute time-varying correlations, we identify the minimum time scales at which this non-stationarity can be observed using our current method and provide a means by which the non-stationarity can be tested for significance.

Further, we found evidence that this non-stationarity may be a result of bimodal network dynamics. The following discussion considers the implications of these results and highlights the methodological strengths and weaknesses of this work.

4.6.1 Frequency specificity of MEG RSNs

In this chapter we computed time varying envelope correlation across multiple frequencies spanning 4 to 48 Hz. The highest correlations were found in the alpha and beta bands for all three ROI pairs, both when using stationary and non-stationary measures. These results are in line with those reported in other studies of electrophysiological resting state connectivity (Laufs et al., 2003; Mantini et al., 2007; Liu et al., 2010; Brookes et al., 2011b; de Pasquale et al., 2012). Here we show that this same frequency dependency is also observed when looking at the temporal variability in this correlation.

4.6.2 Bimodality of resting state connectivity

Alongside the standard deviation, we also computed measures of bimodality to further investigate the notion that resting state functional connectivity connectivity is a transient phenomenon. Our working hypothesis was that, if resting state connectivity is driven by transient events in which the network becomes fully synchronised, then this might manifest as a bimodal distribution of correlation values, where one mode represents the transient events and the other mode represents the noise in the correlation measure. Results from both measures of bimodality, the *skewness* and the *BIC ratio*, were similar to those from the *within-window FC variability* measure in the frequencies that exceed the group level significance threshold (although the *skewness* proved to be a more robust measure to compute as indicated by the non-significance of the *BIC ratio*). This result shows that the increase in temporal variability may be attributed to an increase in transiently strong correlations; rather than a general multiplicative scaling of the functional connectivity, which would only result in a broadening of the distribution.

These findings support the idea that resting state connectivity is a non-stationary phenomenon. Complementary results from the study by de Pasquale et al. (2012) suggest that temporal non-stationarities may play a role in mediating interactions between different networks whereby periods of weaker within-network synchronisation allow some nodes of one network to interact with other networks. A related explanation for these non-stationarities that has been put forward is that the networks seen in the resting state are manifestations of multiple overlapping functional modes, or transiently synchronising networks, that, when viewed as an average over time, result in the well-known RSNs (Smith et al., 2012). Under this theory, the non-stationarities in correlation values investigated in this work may arise from the node-pair in question synchronising as part of different transiently synchronising networks at different points in time.

4.6.3 Time scale of non-stationary functional connectivity

By varying the size of the window over which time varying connectivity was computed, the time scale on which non-stationarity may be observed was investigated. The results in Figure 4.7 suggest that there is an optimum range that best captures transient synchronisation in RSNs between 50 and 250 cycles. For instance in the beta band at 20 Hz, this corresponds to a time window of 2.5 to 12.5 seconds. This is an interesting result as it confirms previous work by de Pasquale *et al.* which showed that BLP correlations in the default mode and dorsal attention networks are non-stationary on a time scale of about 10 seconds (de Pasquale et al., 2010). However, it remains unclear whether this window length genuinely reflects the underlying time scale at which network synchronisation occurs, or if it just represents the amount of information needed to robustly detect a correlation between the two signals when using a sliding window approach. If the latter is the case then in order to probe connectivity on a finer temporal resolution than the minimum time scale shown here, it will be necessary to depart from temporal correlations and apply more adaptive methods that can intelligently pool data from disjoint periods in time. Finally, the mechanism by which envelope correlations are related to the rapid dynamics of the underlying neuronal signal remain unclear.

4.6.4 Methodological considerations

In this study we focussed our attention on two resting state networks, the default mode network and the sensorimotor network. These networks were selected on the basis that they are robustly found across multiple subjects and sessions and have been repeatedly shown in previous MEG studies (Brookes et al., 2011b; Luckhoo et al., 2012).

The networks used in this analysis were first identified using ICA following the method described in (Brookes et al., 2011b). This involves first temporally averaging the signal envelopes within 1 s wide time windows which equates to applying a low-pass filter. Selecting networks identified on the basis of slow fluctuations in amplitude may seem incongruous in the context of this analysis where we are interested in the temporal non-stationarities of MEG functional connectivity. However, while the temporal resolution is indeed significantly reduced through this process, it should be emphasised that the signal envelope is governed by, and can be associated with, the underlying neuronal dynamics of the carrier frequency, and so the high temporal resolution information is being used. Furthermore, the aim of this work was to determine the extent to which resting state connectivity is driven by transiently high correlations, so selecting known RSNs was a prerequisite for this study. Nonetheless, selecting ROIs using this temporally stationary approach may not be optimal if the analysis described here were to be applied more generally to multiple brain regions, since it will only find networks that are stable over the duration of the scan, and with amplitude modulations that may be detected after temporal averaging is applied. It is possible that transient networks may exist that will not be detected unless this non-stationarity is accounted for in the same manner as (de Pasquale et al., 2010).

Having identified RSNs of interest, time series representing the neuronal activity at nodes of these networks were reconstructed using a beamformer. Variations in the position of these ROIs across different subjects and sessions were accounted for by first identifying a session-specific spatial map for each network. From these maps, single MNI coordinates corresponding to the centre of each node of the network were used to reconstruct a representative time signal for this node. This method corresponds to selecting a set of ROIs for each network that maximises the within network connectivity. Since the aim of this study was to investigate variability in resting state connectivity rather than looking for strong connectivity *per se*, this method should not be considered circular. An alternative approach to reconstructing the activity at each node would be to use a wider ROI defined by regions where the ICA derived spatial map exceeds a particular threshold. The beamformed time series across this ROI could then be averaged, resulting in a representative signal for that RSN node. Since we only considered power correlations in this study this approach may have been more appropriate, however if the method were extended to other characteristics of the underlying time series such as the phase, then one would need to overcome ambiguities in the sign for the orientation of the source at each voxel (Hillebrand et al., 2011).

A limitation of the beamforming approach used here was that a single set of weights were used for all time points, therefore assuming that the data covariance structure is stationary across the duration of the scan. A more appropriate method might be to allow for sources of activity to change position and orientation over time by using a time varying beamformer (Dalal et al., 2008), although this approach still requires a sliding window that is large enough to allow sufficiently good estimation of the covariance matrix for use in beamforming (e.g. >30 s for the beta band (Brookes et al., 2008).

Significance testing was achieved through the use of Monte Carlo simulations of realistic null data. It is relevant to draw a distinction between the method employed here and that used in the study by de Pasquale et al. (2010) in which transient periods of high correlation were defined by identifying maximally correlated windows. The latter method iteratively searches for periods in time where the correlation between a seed region and multiple network nodes is high compared to a location outside the network. This has the advantage that artefacts that might cause an increase in the within network correlation would not affect the measure if it also induced a high correlation with a node external to the network. However, the method relies on selecting an appropriate control node *a priori* which our method does not require.

4.7 Summary

We have proposed a method that allows the dynamics of resting state connectivity to be investigated across time and frequency on the basis of correlations in the oscillatory envelope. By computing summary measures that describe the degree of transient synchronisation and comparing these results with simulations of realistic null data we have shown that the interactions between nodes of well-known resting state networks is a non-stationary phenomenon, and that this non-stationarity is bistate in nature. We further showed that there is an optimum time period in which this non-stationarity may be observed. These results further demonstrate the ability of MEG to explore the dynamics that underlie resting state networks.

CHAPTER 5

Hidden Markov modelling of transient MEG functional connectivity

5.1 Chapter Abstract

In this chapter we develop a data driven approach for characterising transient changes in functional connectivity using a hidden Markov model (HMM). This approach is used to infer a discrete set of states, which capture patterns of source space covariance that repeat in a predictable way. The chapter starts with an overview of the HMM and its mathematical formulation. The model is then applied to simulated data with known covariance structure to demonstrate that the HMM can identify transient changes in signal variance and covariance. We then apply the HMM to real MEG data acquired during a simple visual stimulus task. By comparing the inferred HMM states to results from a traditional trial-wise GLM analysis, we show that the HMM is able to capture the spatial and temporal properties of the P100 evoked response, demonstrating the ability of this model to capture transient changes in MEG activity. Finally, the approach is used to infer HMM states from resting state data acquired from individual subjects. By mapping changes in activity specific to each inferred state, we identified repeating patterns of activity whose spatial topographies strongly resembled RSNs.

5.2 Introduction

In **Chapter 4** we developed a methodology for assessing the temporal dynamics of MEG functional connectivity using sliding window correlation. Using this method, we demonstrated that functional connectivity within resting state networks exhibits significant temporal variability on a time scale of seconds to tens of seconds. These results add to a growing body of evidence across multiple neuroimaging modalities that resting state functional connectivity is a highly dynamic phenomenon that may be better characterised by non-stationary measures.

While sliding window approaches have the potential to enable a better understanding of the dynamics that underlie resting state connectivity, the ability of these methods to identify rapid changes in functional connectivity is fundamentally limited by the need for a sufficient number of continuous time points for a robust estimate of the covariance structure. In Section 4.5.4 we demonstrated that it was not possible to identify increases in temporal variability using a sliding window shorter than around 50 cycles. This means that this technique will be blind to changes in connectivity that occur at time scales faster than several seconds. However, in order to provide an effective substrate for cognitive processes, functional networks should be able to rapidly reorganise and coordinate on a sub-second temporal scale (Bressler & Tognoli, 2006).

In this chapter we develop an approach for characterising time varying functional connectivity at sub-second temporal resolution. This is based on the hidden Markov model (HMM) which can identify the points in time at which particular brain states re-occur, providing an intelligent pooling over multiple short-lived events. We demonstrate that this approach provides a meaningful decomposition of spontaneous MEG activity and can provide new insights into the role of fast transient distributed brain activity.

5.3 Hidden Markov model

Hidden Markov models are well established tools for modelling time series data which have been widely applied to biomedical signals. Examples of their application may be found in speech recognition, EEG sleep staging, biomedical signal segmentation and brain computer interfacing. The popularity of these models may be attributed to their ability to segment data into a sequence of discrete states in an unsupervised fashion and to their relative simplicity. Another factor is their ability to model highly non-stationary data that are typical in many biomedical applications.

5.3.1 Introduction to HMMs

Hidden Markov models get their name from two defining properties. First, they assume that the data is generated by a process whose state is unobserved or *hidden*. While the states themselves are hidden, each is associated with an observation model that provides a probabilistic mapping of the state to the observed data. Secondly, HMMs assume that these hidden states satisfy the *Markov* property. This property refers to the fact that the system has no memory beyond the current state, i.e. the next state in the sequence depends only on the current state the system is in, and not on the sequence of states that preceded it. This property may be described within the Bayesian framework as a conditional independence between state s_t and all states prior to s_{t-1}

$$P(s_t|s_1...s_{t-1}) = P(s_t|s_{t-1}) \tag{5.1}$$

Here we assume the observations $\mathbf{y}_t$ are conditional only on the state at time t:

$$P(\mathbf{y}_t|s_1...s_t) = P(\mathbf{y}_t|s_t) \tag{5.2}$$

These independence properties means that the joint distribution of the state sequence and observations may be factored as follows:

$$P(\mathbf{y}, \mathbf{s}) = P(s_0) \prod_t^T P(s_t|s_{t-1}) P(\mathbf{y}_t|s_t) \tag{5.3}$$

Where $P(s_0)$ is the initial state probability. This factorisation may be represented graphically as a dynamic Bayesian network as shown in Figure 5.1, where the arrows denote conditional dependencies. From this factorisation it is evident that the HMM contains two conditional probability distributions: $P(s_t|s_{t-1})$ and $P(\mathbf{y}_t|s_t)$. The *state transition model* $P(s_t|s_{t-1})$ defines the probability of the current state given the value of the previous state. Most HMMs typically assume a discrete state space, such that s_t can take on one of K integer values. The state transition model $P(s_t|s_{t-1})$ is then defined by a $K \times K$ *state transition matrix* containing the probabilities of transitioning between each pair of states. The *observation model* $P(\mathbf{y}_t|s_t)$ provides a mapping of the hidden states to the observed data, and can take one of a number of forms depending on the observed variables. For instance, the observation model may represent a multinomial distribution for discrete observations, a Gaussian distribution for continuous observations, or a Poisson distribution for count processes.

Fitting an HMM to the data involves inferring both the transition model and the observation model for each state. Since these models are effectively coupled by the conditional dependence relationships, this inference is performed simultaneously for both $P(s_t|s_{t-1})$ and $P(\mathbf{y}_t|s_t)$. Herein lies a strength of the model over simple mixture models, since the temporal dynamics of the system will provide some regularisation of the inferred observation models by favouring distributions that repeat in a predictable way. In addition to inferring the observation model for each state and the state transition matrix, it is also possible to infer the probability of being in each state at any point in time. This state sequence may then be used to characterise the temporal properties of the model.

The hidden Markov model described above has previously been applied to MEG data in order to derive time-varying beamformer weights (Woolrich et al., 2013). In this approach, states were inferred from the sensor level recordings and the resulting HMM state time courses were used to derive state-specific spatial filters by computing covariance matrices from data pooled over each state. This allows improved spatial localisation of neuronal activity by tuning the filter properties to that which is required at different points in time.

In this chapter we use an HMM with a multivariate normal distribution observation model to model transient resting state functional connectivity in source projected MEG data. Our working hypothesis is that transient interactions between different brain regions should manifest as unique patterns of activity that recur at different points in time. These patterns should be modelled by the HMM as a number of discrete states each with a multivariate normal distribution over the source-space data that describes these interactions. Under the hypothesis that these patterns are short lived, these states may be unobservable with traditional approaches to assessing functional connectivity that average over either the entire duration of the recording, or within windows several seconds long. Inferring an HMM however may allow us to recover these putative transient states in an unsupervised fashion without any *a priori* assumptions about the time scales or brain areas involved.

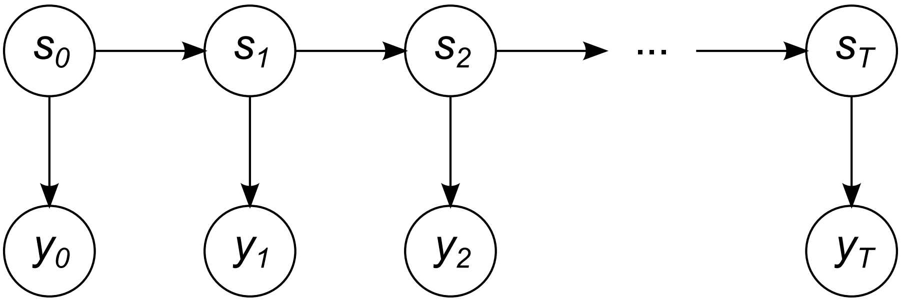

Figure 5.1: A Bayesian network showing the conditional independence relations for a hidden Markov model.

5.3.2 HMM Inference

We assume an HMM of length T, state space dimension K, hidden state variables $\mathbf{s} = \{s_1 ... s_T\}$ and observed data $\mathbf{y} = \{\mathbf{y}_1 ... \mathbf{y}_t\}$, where $\mathbf{y}_t$ correspond to the $(M \times 1)$ observations at time t (e.g. sensors or voxels). The full true posterior probability of the model is then given by:

$$P(\mathbf{y}, \mathbf{s}, \boldsymbol{\Theta}) = P(s_0|\boldsymbol{\pi}_0) \prod_{t}^{T} P(s_t|s_{t-1}, \boldsymbol{\pi}_t) P(\mathbf{y}_t|s_t, \boldsymbol{\theta}) P(\boldsymbol{\pi}_t) P(\boldsymbol{\pi}_0) P(\boldsymbol{\theta}) \tag{5.4}$$

where $P(\boldsymbol{\pi}_t)$, $P(\boldsymbol{\pi}_0)$ and $P(\boldsymbol{\theta})$ are chosen to be non-informative priors. The HMM parameters $\boldsymbol{\Theta} = \{\boldsymbol{\pi}_0, \boldsymbol{\pi}_t, \boldsymbol{\theta}\}$ consist of $\boldsymbol{\pi}_0$ which parameterises the initial state probability $P(s_0)$, $\boldsymbol{\pi}_t$ which is the $K \times K$ transition probability matrix that determines the state transition probabilities $P(s_t|s_{t-1})$, and $\boldsymbol{\theta}$ which parameterises the observation probabilities $P(\mathbf{y}_t|s_t)$. The term $P(\mathbf{y}_t|s_t, \boldsymbol{\theta})$, is the observation model. In this work we assume that the observation model for state k is a multivariate normal distribution with $\boldsymbol{\theta}_k = \{\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k\}$, where $\boldsymbol{\mu}_k$ is the $(M \times 1)$ mean vector, and $\boldsymbol{\Sigma}_k$ is the $(M \times M)$ covariance matrix:

$$P(\boldsymbol{y}_t|s_t = k, \boldsymbol{\theta}) \sim \mathcal{N}(\boldsymbol{\mu_k}, \boldsymbol{\Sigma}_k) \tag{5.5}$$

The prior distributions over the HMM parameters $\boldsymbol{\Theta} = \boldsymbol{\pi}_0, \boldsymbol{\pi}_t, \boldsymbol{\theta}$ are chosen to be non-informative conjugate distributions. The approximate posterior distributions will then be functionally identical to the prior distributions (i.e. a Gaussian prior density is mapped to a Gaussian posterior density), making the model tractable to certain kinds of inference (see Rezek & Roberts (2005) for details).

In this work we use variational Bayes (VB) inference on the HMM (Rezek & Roberts, 2005). This is fully probabilistic and furnishes us with full posterior distributions on the model parameters $P(\boldsymbol{\Theta}, \mathbf{s}, \mathbf{y})$. Aside from the inferred posterior distribution over the observation model $P(\mathbf{y}_t|s_t, \boldsymbol{\theta})$, we are also interested in determining those points in time at which particular states are active. The relevant output is the marginal posterior inference on the state variables $P(s_t|\mathbf{y})$. This is obtained using Viterbi decoding. For the purpose of computing summary statistics of state life time and occupancy, we have chosen to hard classify the states as being on or off by choosing the most probable *a posteriori* state u_t at each time point:

$$u_t = \arg\max_k P(s_t = k|\mathbf{y}) \tag{5.6}$$

5.3.3 Summary statistics

As in (Woolrich et al., 2013) we defined a number of summary statistics to describe the temporal characteristics of the inferred states. These statistics were computed from the state time courses as follows.

Fractional occupancy is defined as the fraction of time spent in each state:

$$\text{fractional occupancy } (k) = \frac{1}{T}\sum_t (u_t == k) \tag{5.7}$$

where $u_t == k$ is one if $u_t = k$ and is zero otherwise, and T is the length of the state sequence in samples. The mean life time is defined as the average amount of time spent in each state before transitioning out of that state:

$$\text{mean life time } (k) = \frac{\sum_t (u_t == k)}{\text{number of occurences } (k)} \tag{5.8}$$

The mean interval length is similarly defined as the average amount of time spent between consecutive visits to a particular state:

$$\text{mean interval length } (k) = \frac{T - \sum_t (u_t == k)}{\text{number of occurences } (k)} \tag{5.9}$$

where the number of occurrences is given by:

$$\text{number of occurences } (k) = \sum_t (((u_t == k) - (u_{t-1} == k)) == 1) \tag{5.10}$$

5.4 HMM inference of simulated data

In this section we provide a brief demonstration of HMM inference from trivial simulated data. The simulated data comprised three simulated alpha oscillations which were independently generated by band-pass filtering Gaussian noise. Correlation was introduced between the independent signals by multiplying the data with the Cholesky decomposition of a known covariance matrix. Three states with zero mean and known covariance were simulated within the data by enforcing different covariance structures within the first, second and third tertiles of the data. Two sets of simulated data were generated. In the first simulation, each state corresponded to a doubling in variance in one of the three simulated alpha oscillations with respect to the variance in the other two states, but with zero covariance for all states. In the second simulation, the variance was held constant and each state corresponded to an increase in covariance between one pair of simulated alpha oscillations (where the covariance was equal to half of the variance). The simulated covariance matrices and signals are shown in Figure 5.2.

From these simulated data an HMM with three states was inferred. Because the signals had been band-pass filtered into the 8 - 12 Hz frequency band, the mean of the data was zero. Accordingly, the inferred observation models were constrained to have zero mean. This step was taken to ensure that the model did not fit to the peaks of single oscillations, but rather captured changes in signal amplitude by variance changes.

In both simulations, inference of a 3 state HMM was able to recover the covariance structure and the time points at which each state was active as shown in Figure 5.2.

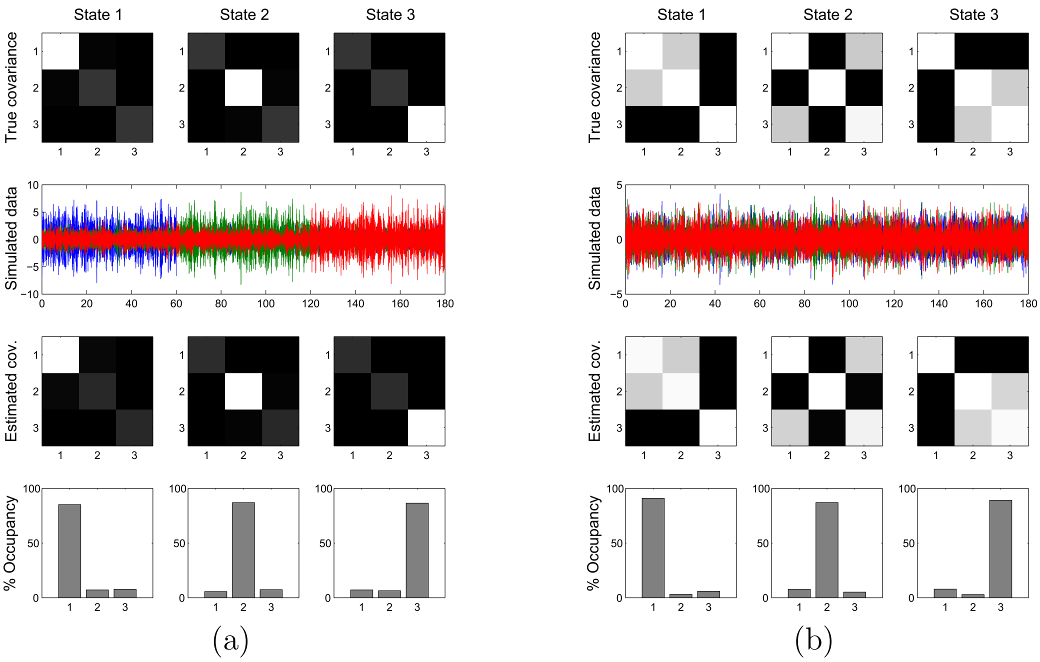

Figure 5.2: HMMs inferred from 3 simulated time courses (2nd row) with three states with known covariance structure (1st row). Three states were inferred using an HMM and the estimated state covariance matrices (3rd row) and fractional occupancy (4th row) are shown for each state. In simulation (a) the states were defined based only on changes in variance of the three simulated signals. In simulation (b) the states were defined only by changes in covariance. For both simulations, the HMM was able to recover the covariance structure and the times at which each state was active.

5.5 Application of HMM to MEG data

Having demonstrated the ability of the hidden Markov model to recover known states from simulated data, we applied the technique to real MEG data with the aim of identifying transient patterns of spontaneous activity. In Section 5.6, we first apply the HMM to task data where the timing and brain regions associated with a visual evoked response are known. In Section 5.7, we apply the technique to resting state data for single subjects. First however, it is necessary to address some additional methodological considerations for applying the technique to MEG data.

5.5.1 Dimensionality reduction

In Section 5.4 the HMM was applied to simulated data comprising three simulated alpha oscillations. In reality, the dimensionality of MEG data is much higher, at approximately 300 for sensor data, and several thousand for source space data (depending on the grid size). Prior to inferring the HMM it was first necessary to reduce the dimensionality of the data to a more manageable size for computational efficiency (the HMM toolbox used in this work was limited to a data dimensionality of approximately 40 due to computational stability). As described in Section 3.3.3, a number of techniques exist for defining ROIs. One option would be to use network nodes that are known to be functionally connected as in **Chapter 4**. However, restricting the analysis to those regions that are identified as functionally connected using temporally stationary approaches would bias the results to interactions that are maintained over longer periods of time and may exclude regions that only engage transiently. For example, the PCC has been previously demonstrated to engage transiently with a number of other regions rather than maintaining strong membership of one particular network (de Pasquale et al., 2012). Based on this reasoning we adopted a more data driven approach to dimensionality reduction, by applying principal component analysis (PCA) and retaining only the 40 components with highest variance. This approach makes no explicit assumptions on the spatial distribution of the observations, other than that signals of interest are likely to have high variance. After performing PCA dimensionality reduction, the data were whitened by normalising the variance of each component. It is worth noting that this approach is exactly equivalent to the first stage of FastICA, such that the lower dimensional subspace may be projected back to the original voxel-wise space by an unmixing matrix.

5.5.2 Mapping state-specific activity

By performing a principal component analysis prior to inferring the HMM, the dimensionality of the data is reduced to a computationally manageable amount. However, this means that the multivariate normal distributions that define each state span this reduced subspace, and are therefore not readily interpretable in terms of the underlying anatomy. To interrogate those brain areas associated with each state, the obvious solution is to simply project the state means and covariance matrices back to the original voxel space using the eigenvectors derived during the PCA.

We are primarily interested in characterising changes in activity that are specific to each state. However the multivariate normal distributions will also capture activity that is common across states. Moreover, the covariance matrix for each state will contain not only covariance due to genuine neuronal interactions, but also covariance due to signal leakage between voxels. As a result of this leakage term, the covariance matrices associated with different states appear very similar. Since the covariance due to signal leakage is consistent over states (as the beamformer weights do not change over the duration of the recording), simply viewing the difference between states provides an acceptable means to mitigate the effects of signal leakage.

Here we use the inferred HMM state time courses and regress them back onto the data to obtain maps representing the difference between states, which will therefore have signal leakage subtracted out. This is achieved by fitting a general linear model (GLM) to the amplitude envelope at each voxel, with the state time courses as regressors (Friston et al., 1996; Brookes et al., 2004; Woolrich et al., 2009). This yields a set of K spatial maps of the parameter estimates which provide a description of how well each regressor (state time course) explains the amplitude fluctuations at each voxel. These parameter estimates were expressed in terms of partial correlation coefficients by normalising both the data and the regressors to have zero mean and unit variance prior to fitting the GLM.

5.6 Validation on single session MEG task data

We previously demonstrated the ability of the hidden Markov model to recover transient states from simulated data. In this section the HMM is applied to MEG data recorded during a simple task to validate the method on real data. We used a single session from an existing dataset from a face recognition study in which images of faces and motorbikes were presented (data courtesy of Susie Murphy and Catherine Harmer, University of Oxford). For the purpose of validating the HMM, we focussed only on the response to the visual stimulus for a single subject, and did not examine differences between the faces or motorbikes conditions. The aim of this analysis was to determine whether the HMM could identify states that were time-locked to the visual stimulus. Crucially, by using task data, we can obtain a form of "ground-truth" by also performing a traditional trial-wise GLM analysis of the evoked response. This enables a characterisation of the brain areas and timing of stimulus-locked activity. Under the hypothesis that the states inferred by the HMM capture transient changes in oscillatory dynamics, we would predict that the model would encode this stimulus-locked activity in the form of states that repeat at the same time points across multiple trials, and with spatial maps that represent brain areas associated with visual processing.

5.6.1 Experimental paradigm

The paradigm consisted of 360 trials split evenly into eight blocks. Within each trial, following the presentation of a fixation cross, participants were presented with

400 to 600 ms of a blank grey screen followed immediately by the 100 ms presentation of an image of either a face or a motorcycle. Within each block, 45 trials took place in a pseudo-random order. Thirty of the trials in each block involved the presentation of a face, while the remaining fifteen trials involved the presentation of a motorcycle. During image presentation, participants were instructed to maintain their gaze where the fixation cross formerly was (i.e. at the centre of the image).

5.6.2 Data processing

External noise was removed using the signal space separation (SSS) method (Taulu et al., 2005) and the data was down-sampled to 200 Hz, using the MaxFilter software (Elekta-Neuromag). After converting to SPM8 format (`http://www.fil.ion.ucl.ac.uk/spm`), the data were band-pass filtered between 1 and 48 Hz and sorted into trials by epoching the continuous data between 200 ms prestimulus and 350 ms post-stimulus onset. The LCMV beamformer was used to project the data onto a regular 8 mm grid spanning the entire brain.

5.6.3 GLM analysis

As a comparison for the HMM, we identified event related activity associated with the presentation of the visual stimulus by fitting a trial-wise general linear model (GLM) to the source reconstructed data at each voxel and time point. The design matrix contained two regressors corresponding to each of the trial conditions (motorbikes and faces). By contrasting the regression parameters for each condition, it would be possible to perform a test for evoked activity that is significantly higher in one condition versus the other. However, for the purpose of validating the HMM we opted to simply identify evoked activity that was locked to the presentation of a visual stimulus (regardless of whether this stimulus was a face or motorbike). To this end, the regression parameters for both conditions were averaged at each time point, resulting in a visual event-related field (ERF) for each time point and voxel.

5.6.4 HMM analysis

The trial-wise data were temporally concatenated across trials such that the dimensions were voxels × (time points × trials). Principal component analysis was used to reduce the dimensionality of the data to 40 components which were subsequently whitened. From these 40 principal components, an HMM with five states was inferred. This low model order was selected due to the paucity of data for the single session, and due to the expectation that relatively few distinct patterns of activity would exist within a simple task environment. Each inferred HMM state is associated with a unique multivariate normal distribution over the observations (principal components), defined by an $(M \times 1)$ mean vector and an $(M \times M)$ covariance matrix where $M = 40$ (the number of principal components). As before, the mean of the multivariate normal distribution was forced to be zero. To account for variations in the inference due to different initialisations, ten realisations were performed for each inference and the model with the lowest free energy (an approximation to the Bayesian model evidence) was chosen. The state time courses corresponding to these states were identified using Viterbi decoding as described in Section 5.3.2.

5.6.4.1 Identifying task dependent states

To determine if any of the inferred states were stimulus-locked, a GLM analysis was performed independently for each inferred HMM state time course using the same trial-wise regressors as in Section 5.6.3, i.e. instead of fitting the GLM to the voxel-wise trial-wise data as described previously, the model was instead fitted to the HMM state time courses after epoching into trials. This results in a time course similar to that of an ERF that characterises the task dependence of each state over time, but here expresses a time course of the fractional occupancy for each state.

5.6.4.2 Identifying state-specific activity

The method proposed above enables the identification of task-locked states. However, it does not provide any information on the spatial topography of the inferred states. To map activity that was locked to the states (rather than the task) we again performed a GLM analysis but this time fit the data at each voxel to the HMM state time courses as described in Section 5.5. Prior to fitting the GLM, the trial-wise data at each voxel were concatenated across trials and the Hilbert envelope of the data was computed, such that the GLM would identify state-specific fluctuations in amplitude. Five regressors corresponding to each of the state time courses were fitted to the concatenated data at each voxel resulting in a spatial map of regression parameters for each state. By variance normalising the data prior to fitting the GLM, these parameters were expressed as partial correlations.

5.6.5 Results

Figure 5.3 shows the results from the trial-wise GLM. The maximum t-statistic for the average of the faces and motorbikes condition was found within the visual cortex at t = 100 ms. The spatial map at this time point is shown in Figure 5.3a. The event related field at the location within the visual cortex with the highest t-statistic is shown in Figure 5.3b. The P100 evoked response is clearly identifiable as the prominent peak at 100 ms post stimulus.

Figure 5.4 shows results from the GLM analysis applied to the HMM state time

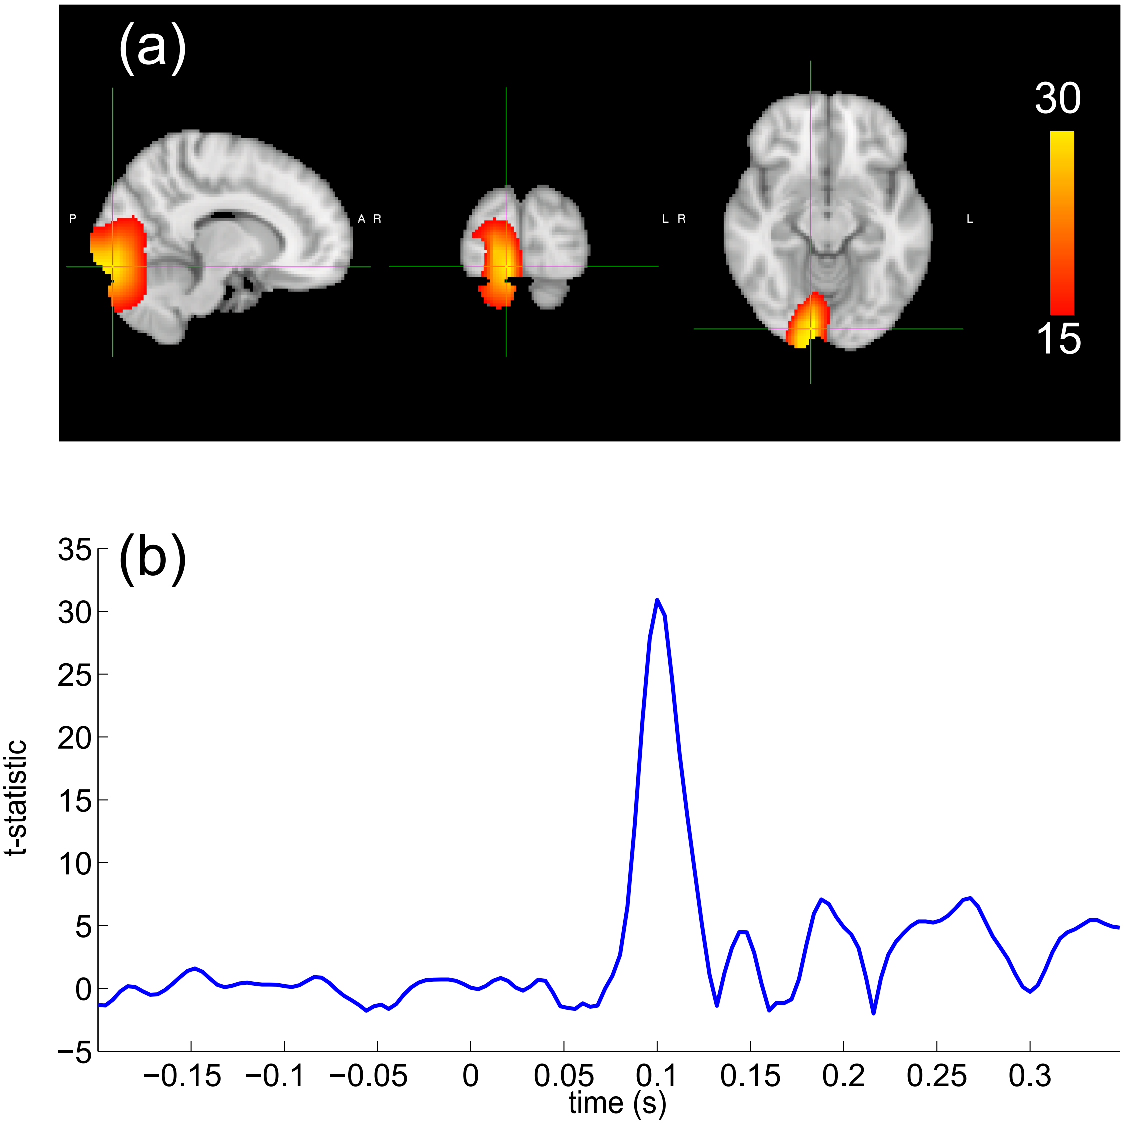

Figure 5.3: (a) Spatial map of uncorrected t-statistics for the average of the faces and motorbikes conditions at t = 100 ms post-stimulus showing stimulus locked activity within the visual cortex. The map is thresholded between 50 and 100 percent of the maximum t-statistic. (b) time course of event related activity from 200 ms pre-stimulus to 350 ms post-stimulus for the voxel with the highest t-statistic (crosshairs in (a)). Easily identifiable is the prominent peak at 100 ms post-stimulus, corresponds to the P100 evoked response.

courses. Each plot shows the t-statistic for the average of the faces and motorbikes condition derived from fitting the GLM to a particular state time course. An increase in the t-statistic corresponds to a state being more frequently visited during a particular point in time within the trial. State 5 shows a strong task-dependency, with a prominent peak at t = 100 ms. States 1 and 2 are negatively correlated with stimulus presentation, denoting that these states were less frequently occupied at this time point.

Fitting the state time courses to the data at each voxel enables the identification of brain regions whose amplitude fluctuations are locked to a particular state. Figure 5.5 shows spatial maps of the partial correlations for each of the five states. Red/yellow and blue colours represent state-specific increases and decreases in amplitude respectively. State 5, which was found to be task-dependent in Figure 5.4, corresponds to increased activity in the visual cortex. Figure 5.5 also shows the fractional occupancy time courses for each state across trials. These occupancies are expressed as percentages, where 100 percent occupancy would indicate that a certain state is always visited at a particular point within the trial (these time courses represent the same information as shown by the t-statistics in Figure 5.4, but expressed in absolute terms rather than relative to the pre-stimulus baseline).

5.6.6 Discussion

In this section we have demonstrated the application of the hidden Markov model to real MEG data for a single subject acquired during a visual recognition task. By assessing the task-dependency of the inferred states over time, we identified a single state which was frequently visited 100 ms post-stimulus onset and which represented increased activity within the visual cortex. This state appears to encode the P100 evoked response that was identified using a traditional GLM analysis. In contrast to conventional analyses of trial-wise data, inference of the HMM was performed on the continuous data. Despite being blind to the timing of the task, the model was able to capture activity that was stimulus-locked, demonstrating the ability of the HMM to identify transient periods of activity within MEG data in an

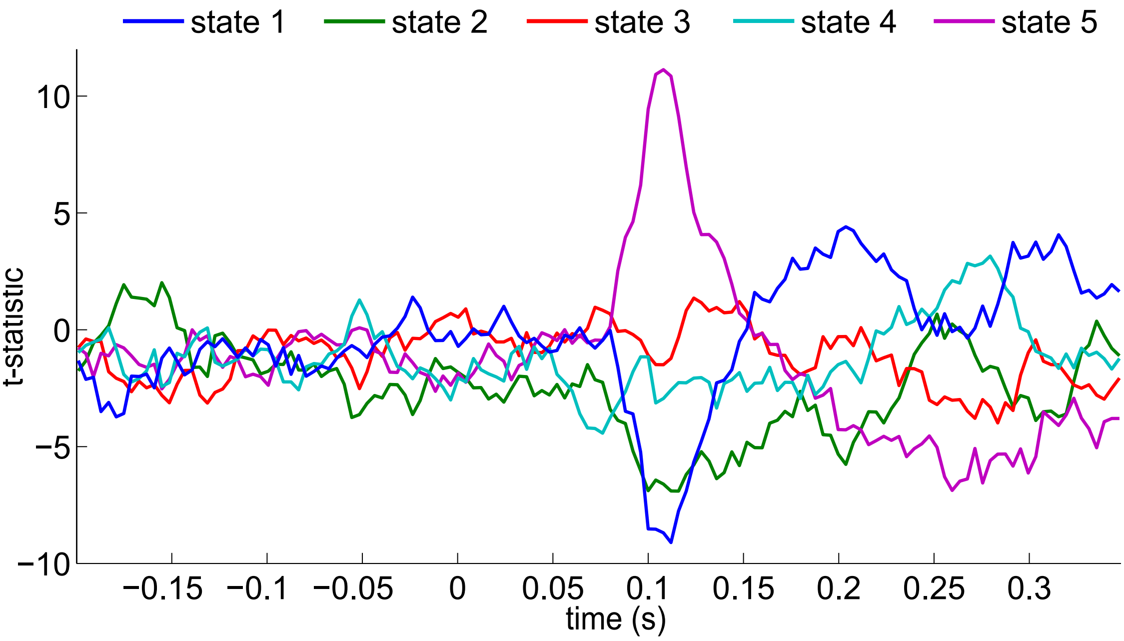

Figure 5.4: Task dependency of each of the five HMM states over all trials from 200 ms pre-stimulus to 350 ms post-stimulus. Each plot represents the t-statistic for the average of the faces and motorbikes conditions derived from a GLM fit to the state time courses. The magnitude of the t-statistic shows how often (across trials) a state is visited during at a particular point within the trial. State 5 correlates positively with the presentation of the stimulus, showing a prominent peak at t = 100 ms. States 1 and 2 are negatively correlated with stimulus presentation.

unsupervised fashion. This validation provides confidence that the HMM should be able to identify functionally meaningful states from resting state data.

5.7 Application of HMM to single session resting state data

The previous two sections have provided a validation of the HMM both on simulated data with known covariance structure and state timings, and on MEG task data where the brain regions and time scales of evoked activity were identified using a GLM analysis. In this section the HMM will be applied to resting state data for single subjects.

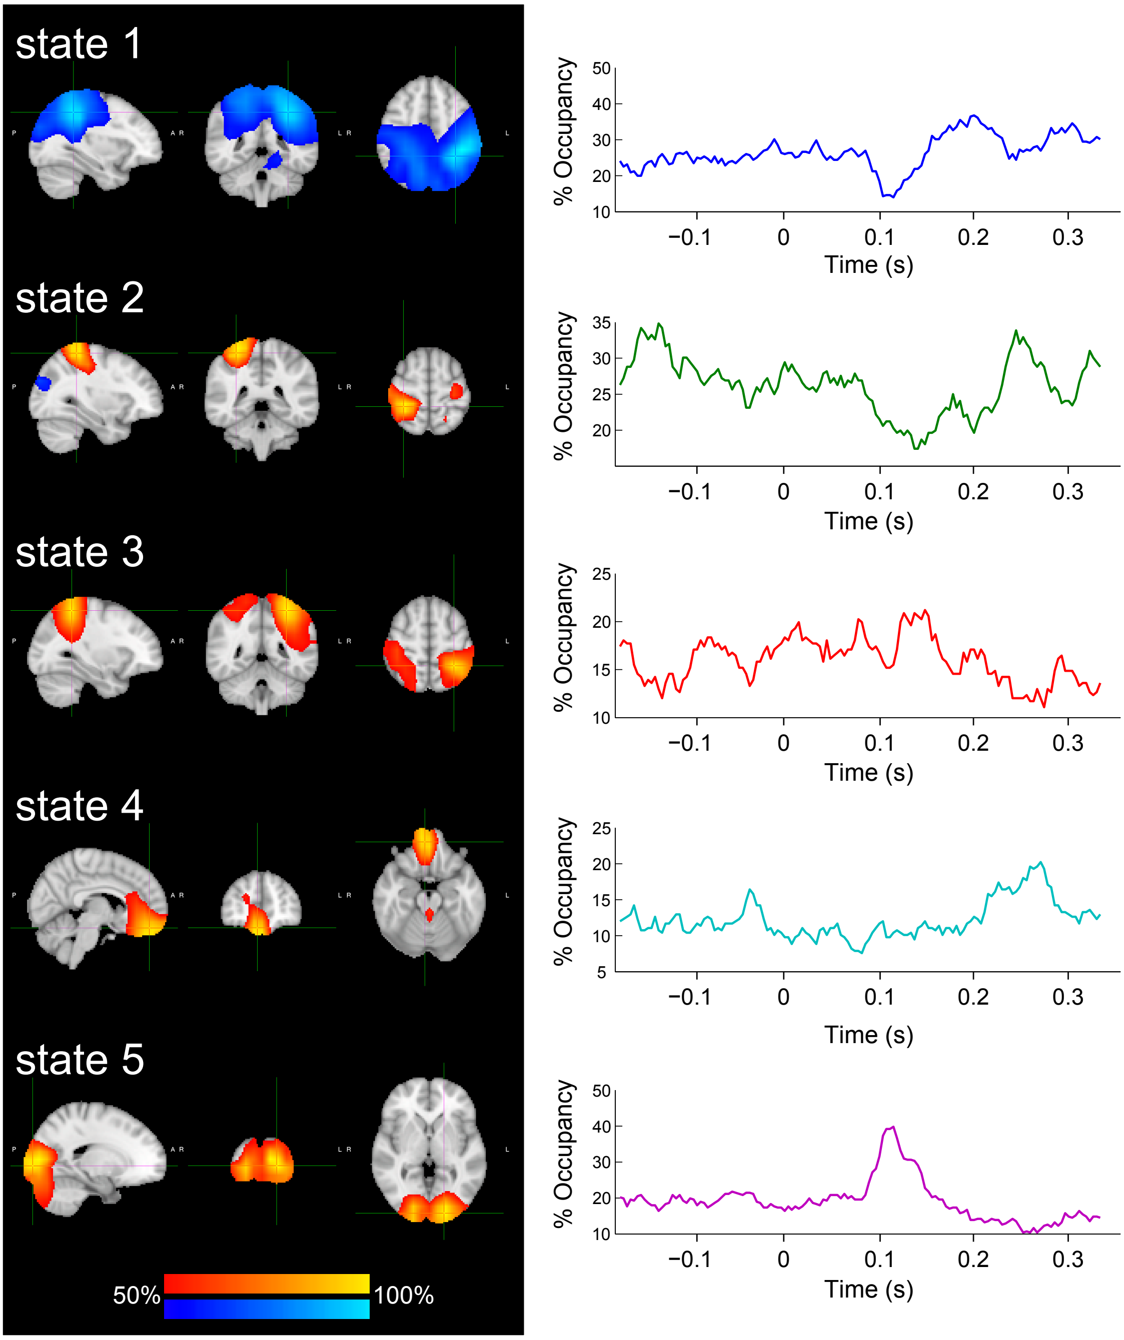

Figure 5.5: Left: Spatial maps for each of the five inferred state showing the partial correlation of the amplitude envelope with the state time courses. Each map is thresholded between 50 and 100 percent of the maximum value. Red/yellow and blue colours represent state-specific increases and decreases in amplitude respectively. Right: time courses showing the percentage occupancy of each state across trials.

5.7.1 Dataset and preprocessing

Resting state MEG data were acquired from nine healthy subjects. The subjects were asked to lie in the scanner with their eyes open while ten minutes of data were recorded. The MEG data were acquired using a 275 channel CTF whole-head system (MISL, Conquitlam, Canada) at a sampling rate of 600 Hz with a 150 Hz low-pass anti-aliasing filter. Synthetic third order gradiometer correction was applied to reduce external interference. Localisation of the head within the MEG helmet was achieved using three electromagnetic head position indicator (HPI) coils. By periodically energising these coils their position within the MEG sensor array was identified. Prior to data acquisition, the HPI coil locations, the position of three fiducial points (the nasion, and left and right preauricular points), and the head shape were recorded using a three-dimensional digitiser (Polhemus Isotrack). MR images were acquired using a 3T Phillips Achieva MR system at $1 \times 1 \times 1$ mm^3 resolution running an MPRAGE sequence. Each subject's structural MRI was registered to the MNI152 standard brain such that all subsequent source space analysis was performed in MNI space. The locations of the MEG sensors with respect to the anatomy were determined by registering the digitised head surface to the head surface extracted from the structural MRI. The data were converted to SPM8 (`http://www.fil.ion.ucl.ac.uk/spm`) and down-sampled to 200 Hz. The data for each session were preprocessed following the pipeline described in **Chapter 3** and frequency filtered into the 4-30 Hz band. This frequency band was selected on the basis that previous resting state analyses have shown strong functional connectivity within this range (Brookes et al., 2011b; de Pasquale et al., 2010).

The preprocessed MEG data were projected onto a regular 8mm grid spanning the entire brain using a scalar beamformer. Following source reconstruction, principal component analysis was used to reduce the dimensionality of the data to 40 components which were subsequently whitened. From these 40 principal components, an HMM with 8 states was inferred. As before, the mean of the multivariate normal distribution was forced to be zero, such that any changes in amplitude of the zero mean signals were modelled by changes in variance.

5.7.2 Results

5.7.2.1 Single subject HMM inference

As described in Section 5.5 spatial maps of state-specific amplitude changes were computed by fitting a GLM to the Hilbert envelope at each voxel, with each of the inferred state time courses as regressors. The spatial maps for each of the 8 inferred states are shown in Figure 5.6 for a single subject. A number of the maps look similar to well-known resting state networks. States 1 - 4 show increased activity within the default mode network (state 1), visual network (state 2), sensorimotor network (state 3) and a bilateral temporal lobe network (state 4). State 5 represents decreased activity (represented by blue colour maps) in parietal regions including the superior parietal lobule (SPL) and visual areas. States 6 - 8 show predominantly decreased activity, but this activity is less focused within particular networks.

In addition to characterising the spatial topography of each state, it is also possible to investigate their temporal properties from the state time courses. Figure 5.7a-b shows the state time courses for the same single subject for the first 5 seconds of data (a) and for the full 10 minute recording (b). It is clear that the states are short lived and are visited throughout the duration of the recording. Figure 5.7c-e describes the temporal properties of the state time courses in terms of summary statistics. The fractional occupancy Figure 5.7c describes the proportion of time spent in each state. State 1 (the putative default mode network) stands out as having lower fractional occupancy to the other states. The life times and interval lengths describe the amount of time spent within each state visit and between state visits respectively. These statistics are shown as violin plots in Figure 5.7d-e, where the shaded region corresponds to the distribution of these values and the cross shows the mean. The state life times are very short, at around 50 ms on average. The DMN (state 1) again stands out as having longer life times than other states, and less frequent visits.

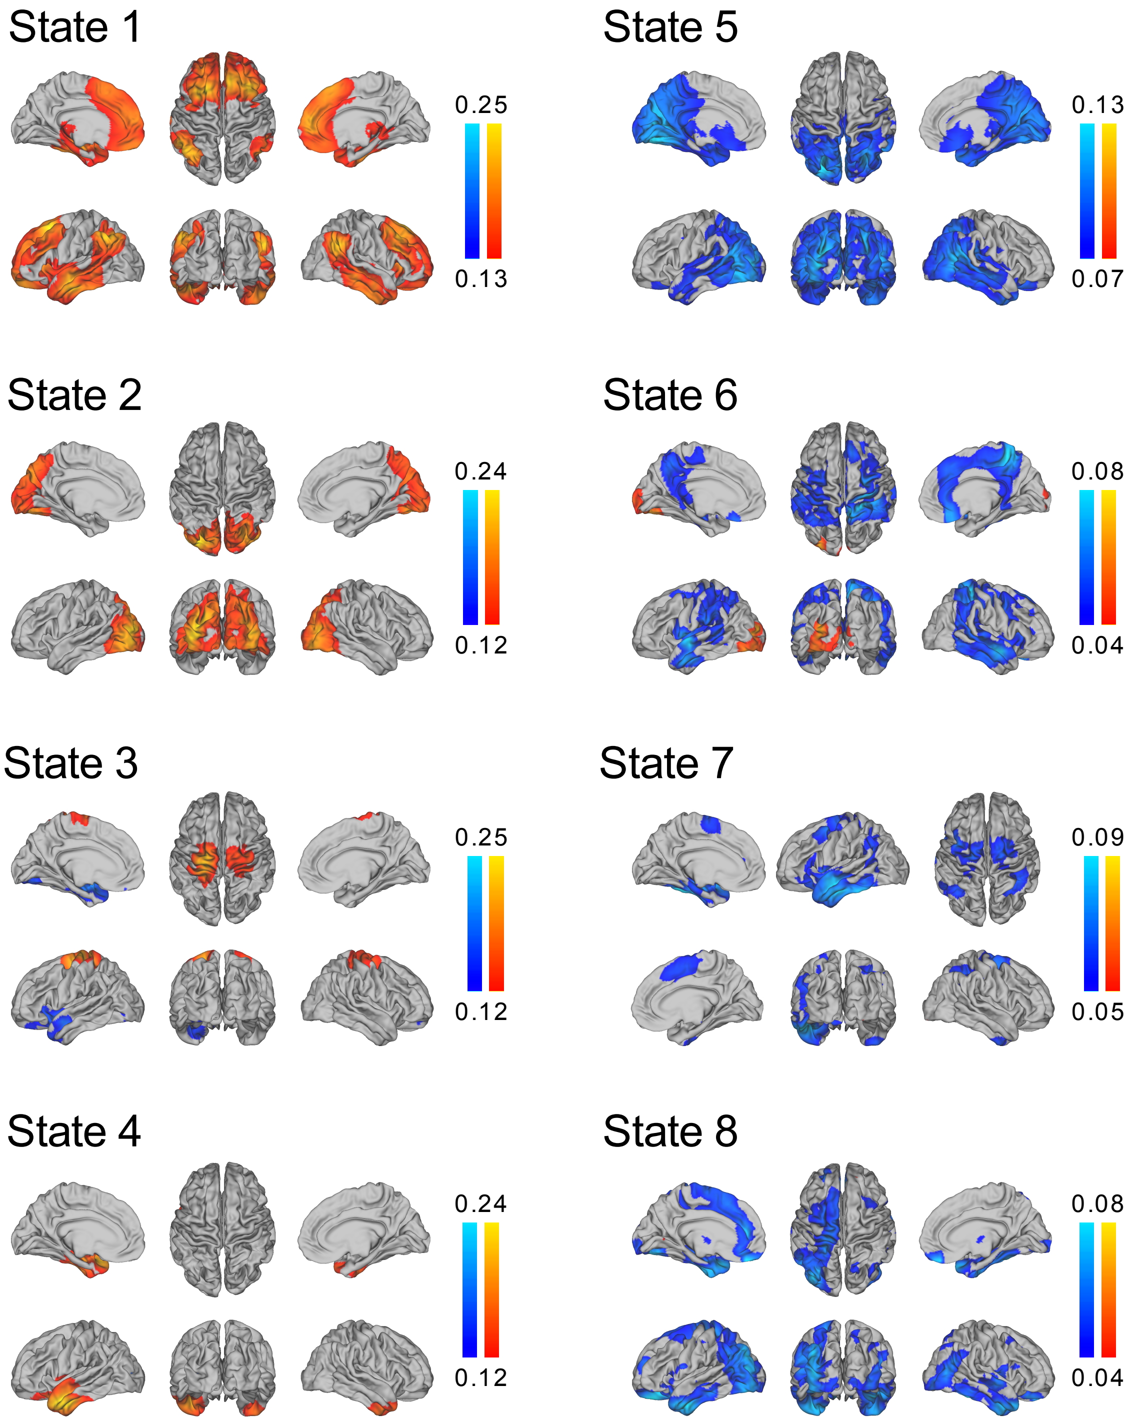

Figure 5.6: Surface rendered volumes showing the partial correlation between the state time courses and the band-limited amplitude of the source space data for each state of an 8 state HMM inferred from a single subject. The correlation values have been thresholded between 50 and 100 percent of the maximum value for each state and the colour maps represent these ranges. Red/yellow and blue colours represent positive and negative correlations respectively.

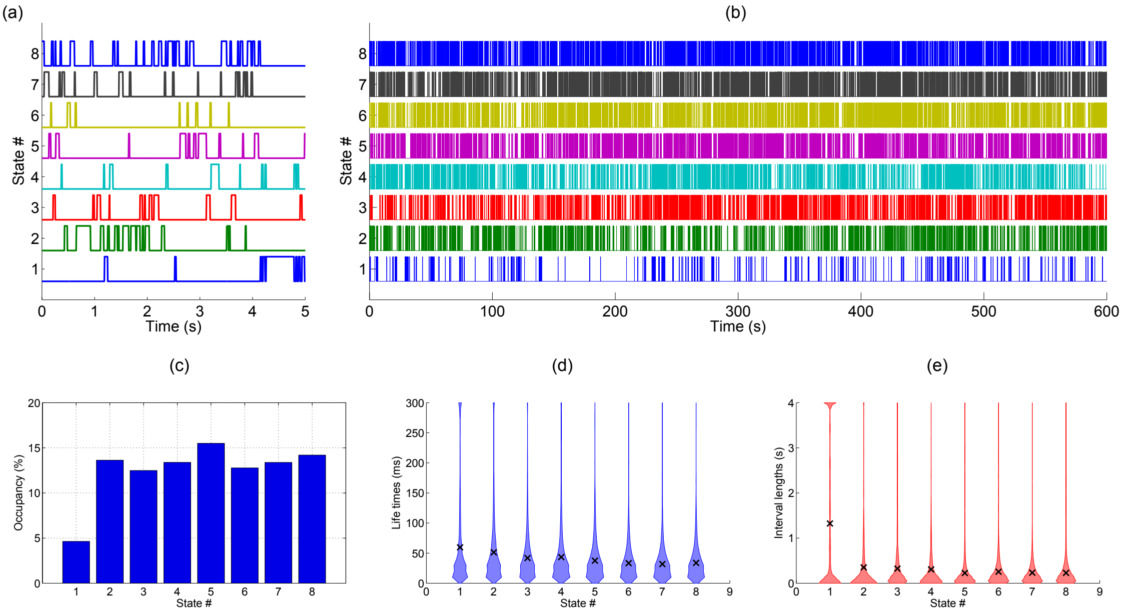

Figure 5.7: (a-b) State time courses, hard classified as on or off for each state inferred from an 8 state HMM for the first 5 seconds (a) and the entire recording (b). (c) Fractional occupancy, (d) life times, and (e) interval lengths for each inferred state. The violin plots in (d-e) shows the distribution of these values as the shaded region and the black cross shows the mean.

5.7.2.2 Matching single subject maps

Spatial maps showing the inferred states for all nine subjects are shown in Figure 5.8. To demonstrate the consistency of the states inferred across different subjects, states with similar spatial topographies were matched based on their spatial cross correlation. This was achieved by first manually selecting subject exemplars whose partial correlation maps represented unique resting state networks. We did not attempt to match all states, but only those whose spatial maps showed strong correspondence with well-known RSNs. The five exemplar maps are shown bounded in green in Figure 5.8 and correspond to the default mode, visual, sensorimotor, temporal and SPL networks. Each subject's state maps were matched to the exemplar maps iteratively by selecting the subject-exemplar pair with the greatest correlation coefficient that exceeded a threshold of 0.4. This subject map was assigned to this exemplar and the process repeated with this map removed until either all maps had been paired, or there were no remaining subject-exemplar pairs with a correlation greater than 0.4.

5.7.2.3 Equivalence of HMM states inferred from oscillatory activity and envelopes

In this section we have applied the HMM to frequency filtered estimates of neuronal activity and obtained states with a striking similarity to well-known resting state networks obtained from ICA. A key difference between our application of the HMM and the ICA approach described in (Brookes et al., 2011b) is that we used the raw band-limited data, rather than the (down-sampled) envelopes of these data. It is therefore of interest to question whether these two decompositions are capturing similar temporal dynamics, despite their different assumptions.

The HMM models the observed data within each state as from a multivariate

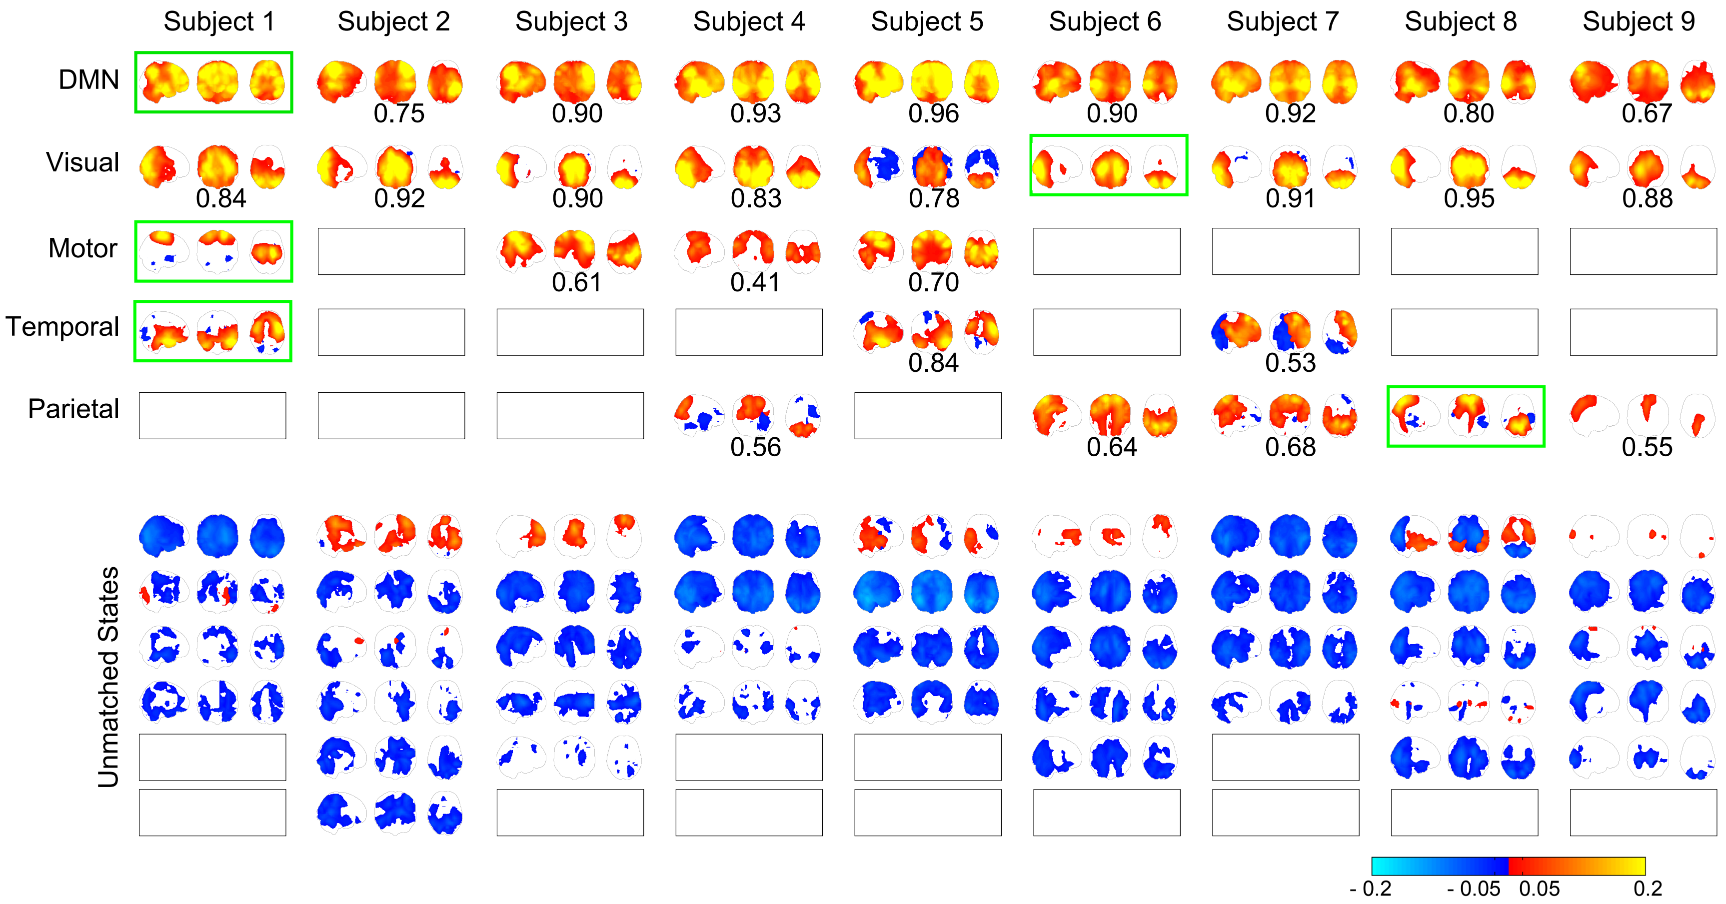

Figure 5.8: Thresholded maximum intensity projection maps for each subject and HMM state. The maps represent the partial correlation computed between each state time course and the source space band-limited (4-30 Hz) amplitude time courses. States corresponding to the DMN, motor network, visual network, temporal network and SPL were matched across subjects by computing the spatial cross correlation between an exemplar state and all states from all other subjects. A match between a state and an exemplar state was defined if the correlation coefficient was greater than 0.4 and the state was did not have a greater correlation with another exemplar. Exemplars were chosen by hand and are shown bounded by green boxes. The correlation coefficient between each matched state and the exemplar state is shown below the spatial map and the unmatched states are shown in no particular order.

normal distribution, defined by the mean and covariance of the observations. In the case of oscillatory data, the mean was constrained to be zero. As demonstrated by the simulations in Section 5.4, state transitions may be driven by changes in the variance of different observations, or changes in the covariance between observations. It is likely that the former is more probable; zero-lag correlations between distant regions does not generally provide a strong measure of functional connectivity, and numerous studies have demonstrated the importance of using band-limited power or amplitude envelopes, rather than the raw signal (Brookes et al., 2011a; Hipp et al., 2012; de Pasquale et al., 2010; Luckhoo et al., 2012). A more likely explanation is that the states reflect transient periods of increased variance in brain regions which are not necessarily phase-locked.

To investigate the changes in signal dynamics associated with state transitions we used a region of interest (ROI) analysis for the inferred states. For each state we identified from the partial correlation map the voxel most correlated with the time course for that state. We generated “pseudo evoked responses” by epoching and averaging the time course of activity at that voxel around the state onset and offset. The time series were epoched 50 milliseconds pre- and post- state onset/offset. Due to the variable life times of each state, the average was only computed for those time points at which the state was active. Figure 5.9 shows the pseudo evoked responses for (a) the raw beamformed data and (b) the envelope of these data, for the default mode, visual, and sensorimotor networks. Transitions into these states are not associated with a particular phase, as indicated by the flat response in Figure 5.9a. Epoching the envelope of the data instead of the raw beamformed data revealed an increase in amplitude around the state onset as shown in Figure 5.9b, suggesting that state transitions are driven by fast increases in signal amplitude.

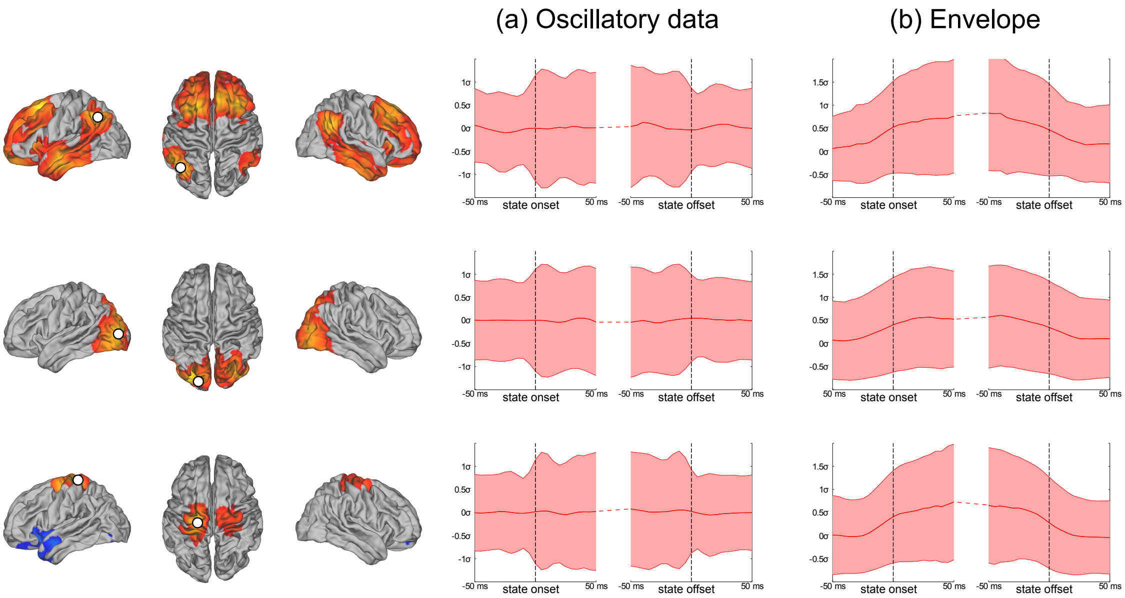

Figure 5.9: Pseudo evoked potentials generated by epoching and averaging the time course of activity at a single voxel around the onset and offset of four different states inferred for a single subject. The surface renderings represent the partial correlation maps as shown previously and the white circle indicates the location of the voxel most correlated with each state time course. (a) and (b) show the epoched average of the (a) raw beamformed data and (b) envelope of these data at the voxel of interest and the shaded region corresponds to the standard deviation over all epochs. The time series were epoched 50 ms pre and post onset/offset. Due to the variable life times of each state, the average was only computed for those time points at which the state was active.

The fact that phase information appears to be unimportant for driving the states suggests that the HMM might do as well from being inferred from the envelope of the data. An advantage of this approach is that it permits temporal concatenation of data across multiple subjects and thus enables the HMM to be inferred for a group of subjects as in the case of ICA (Brookes et al., 2011b; Luckhoo et al., 2012). The extension to group concatenated data is the focus of **Chapter 6**.

To confirm that the same states may be inferred from the amplitude envelopes, we inferred an 8 state HMM for a single subject using both the source space data and the amplitude envelope of these data. Partial correlation maps for the default mode, visual, sensorimotor, and temporal networks derived using both approaches are shown in Figure 5.10. It is clear that the networks inferred in either case are equivalent.

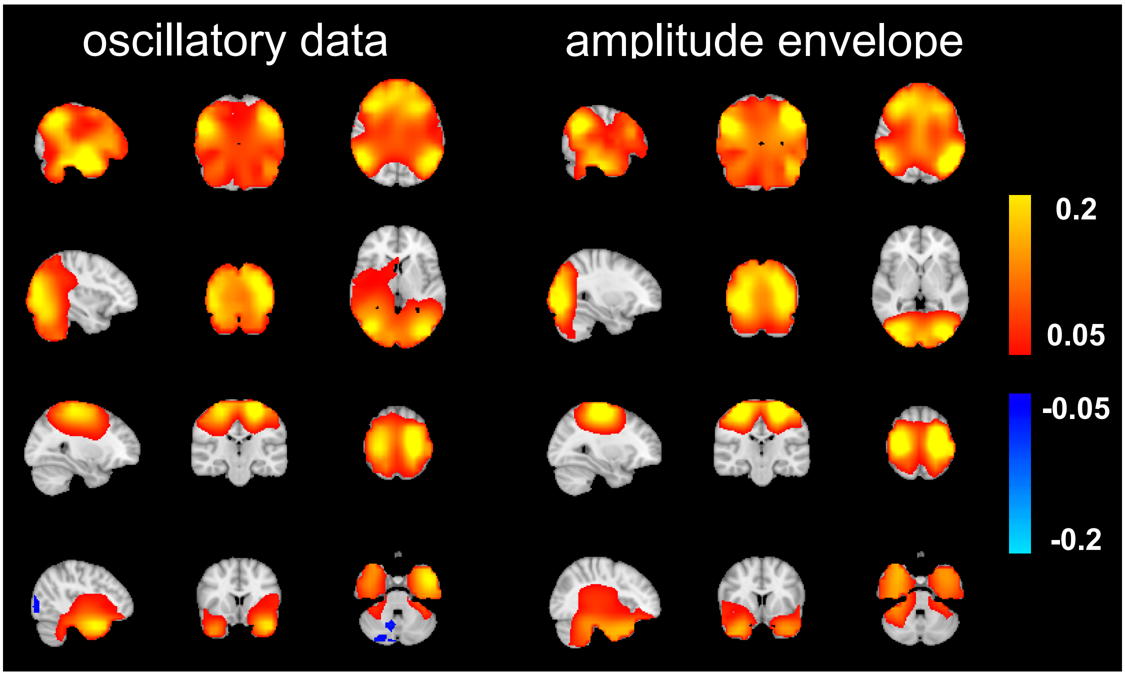

Figure 5.10: Partial correlation maps computed between each state time course and the source space band-limited (4-30 Hz) amplitude envelopes for four states (from top to bottom: DMN, visual, sensorimotor, bilateral temporal lobe) for a single subject. Inference of the HMM states was based on either the oscillatory activity in source space or the amplitude envelope.

5.7.3 Discussion

Application of the HMM to resting state data revealed short-lived (50-100 ms) states corresponding to unique spatiotemporal patterns of spontaneous activity. The spatial topographies of the inferred states represent networks of spatially segregated brain regions that bear a strong similarity to seed based correlation or ICA derived RSNs in both MEG and fMRI (Brookes et al., 2011b; Smith et al., 2009). These results suggest that resting state networks, which are typically associated with time scales below 0.1 Hz, are underpinned by much faster temporal dynamics. By inferring the HMM separately for each subject, it was demonstrated that these

RSN-like states could be consistently identified across different individuals. To our knowledge this is the first time that multiple RSNs have been identified from single session MEG recordings with no prior spatial information, and demonstrates the strength of the HMM for investigating resting state connectivity. Finally, by inferring the HMM from the raw signal we have shown that RSNs may be identified without preprocessing step of computing the amplitude envelope that is typically applied. However, we propose that these two methods are capturing the same dynamics, i.e. changes in the instantaneous power of the MEG signal.

5.8 Summary

In this chapter we introduced the hidden Markov model as a data driven technique for capturing transient functional connectivity. The utility of this model for providing insight into the temporal dynamics of spontaneous brain activity was demonstrated in three stages. Firstly, in Section 5.4 we demonstrated the ability of the HMM to recover non-stationary patterns of variance and covariance in simulated data. Importantly, these simulations were based on simulated alpha oscillations where the covariance structure was known in advance. While the simulation was fairly trivial, it showed that the HMM could identify both changes in signal variance and changes in covariance without any associated variance changes.

Secondly, in Section 5.6 we demonstrated the application of the HMM to real MEG data acquired during a face processing task. The rationale for using task data initially is that it provides a degree of ground truth, in that the states inferred from the HMM could be compared against the timing of the task. In particular, we used a traditional trial-wise GLM analysis to first identify the source and timing of the visual evoked response. By performing a similar analysis based on the HMM state time courses rather than the actual data, the task dependency of the states could be investigated. This showed that the HMM is able to capture the spatial and temporal properties of the P100 evoked response, demonstrating the ability of this model to capture transient changes in activity.

Finally, the HMM was applied to resting state data for individual subjects. This yielded a number of transient states whose spatial topographies strongly resemble RSNs. The life times of these states were short, at between 50 and 100 milliseconds, suggesting that these networks are underpinned by fluctuations in neuronal activity that occur much faster than previously believed. We also demonstrated that equivalent HMM states may be inferred from either the raw oscillatory data, or their amplitude envelopes. This approach is crucial to the application to group-concatenated data in **Chapter 6**.

The results from this chapter suggest that the HMM is a valuable tool for characterising non-stationary functional connectivity in MEG. The discussion of the results from the resting state data have been kept deliberately brief, as a more detailed analysis of resting state data will be the focus of **Chapter 6**. Having shown initial results from single subjects, **Chapter 6** will build on this work by applying the HMM to group-concatenated data. This enables the identification of states that represent a cohort of subjects, and allows further analyses to be performed at the group level. A rigorous discussion of the results, and in particular their importance in terms of our understanding of spontaneous brain activity will also be provided in **Chapter 6**.

Chapter 6

Inferring transient networks from group-concatenated MEG data

---

6.1 Chapter Abstract

In this chapter we apply the HMM to group-concatenated data in order to identify transient states that are consistent within and across subjects [1]. Using this approach we identify short lived (100-200 ms) states with spatial topographies similar to those of well-known resting state networks. By assessing temporal changes in the occurrence of these states, we demonstrate that within-network functional connectivity is underpinned by coordinated neuronal dynamics that fluctuate much more rapidly than has previously been shown. We further evaluate cross-network interactions, and show that anticorrelation between the default mode network and parietal regions of the dorsal attention network is consistent with an inability of the brain to transition directly between two transient states. Finally, we demonstrate that changes in the occurrence of these states reflect temporal variability in

---

[1] The analysis in this chapter has been published in full in (Baker et al., 2014)

within-network functional connectivity.

6.2 Introduction

In **Chapter 5** we applied the hidden Markov model (HMM) to resting state data from single subjects and demonstrated that the inferred states represent short-lived patterns of activity that resemble resting state networks. A limitation of inferring subject-specific HMMs is that it is difficult to compare states identified from different subjects, making group analyses challenging. This arises out of the fact that there will not necessarily be a one-to-one correspondence between states inferred from different individuals. In this chapter, we apply the HMM to *group concatenated* data, to identify states that are consistent both within and across subjects. In this approach, data from multiple subjects are concatenated temporally, resulting in a single aggregated data matrix from which the HMM is inferred. Importantly, PCA dimensionality reduction is performed after this concatenation, such that all subjects are mapped into the same subspace. This approach is similar to that of some group ICA approaches (Smith et al., 2009; Calhoun et al., 2009; Brookes et al., 2011b).

An implication of performing temporal concatenation over multiple subjects is that we must use a rectified form of the data such as the amplitude envelope. This requirement arises out of the fact that an ambiguity exists in the orientation of the estimated dipoles, such that a particular voxel's time course may have arbitrarily different signs across separately beamformed sessions. As a result, source space covariance matrices for different sessions will contain different patterns of positive and negative values thus preventing the HMM from finding consistent multivariate normal distributions across the concatenated data. Taking the envelope of the data removes this ambiguity, since a signal will have the same amplitude envelope regardless of its polarity. In Section 5.7.2.3 it was shown that states inferred from the amplitude and the raw beamformed data are equivalent, demonstrating to some extent that the phase of the signal is not necessary to infer transient states.

In this chapter the previous single subject results are repeated for the group-concatenated data. We then extend this work to assess the temporal properties of the inferred states in more depth. In particular, we address the question of whether the short life times of the states reflect true transient synchronisation, or whether this fast switching is driven by the assumptions underlying the HMM. In addition, we investigate cross-network interactions, by assessing the transitions between states and temporal changes in the rate at which different states are visited. Finally, we provide a detailed discussion of what these states mean in terms of the underlying network dynamics, and how they might relate to other measures of non-stationary functional connectivity such as the sliding window correlation used in **Chapter 4**.

6.3 Dataset and preprocessing

The same dataset as described in Section 5.7 was used in this chapter. As before, the data were band-pass filtered in the 4 to 30 Hz band and projected onto a regular 8 mm grid spanning the entire brain using a scalar beamformer. Following source reconstruction, the oscillatory amplitude envelope at each voxel was derived by computing the magnitude of the Hilbert transform of the source reconstructed data. For computational efficiency, the envelopes were down-sampled to 40 Hz by temporally averaging within sliding windows with a width of 100 ms and 75% overlap between consecutive windows. The amplitude envelopes were concatenated temporally across all subjects after spatially smoothing with a Gaussian kernel (FWHM 9.4 mm). The envelope data for each subject were demeaned and normalised by the global (over all voxels) variance prior to concatenation. The group-concatenated envelopes were demeaned and pre-whitened to reduce the data to 40 principal (temporal) components with unit variance and zero mean. From these reduced data an HMM with 8 states was inferred.

6.4 Results

6.4.1 Large scale transient networks of spontaneous oscillatory activity

State-specific changes in oscillatory amplitude revealed spatial patterns of activity with good similarity to several well-known networks, previously associated with brain wide correlations at slow (<0.1 Hz) time scales (Figure 6.1). Each map represents the partial correlation between the state time courses and the group-concatenated amplitude envelope at each voxel. Accordingly, state-specific increases and decreases in amplitude are represented by red/yellow and blue colours respectively. State 1 shows increased activity in nodes of the default mode network (DMN) including left and right inferior parietal lobule, medial frontal gyrus and medial temporal lobe (but notably, not in the posterior cingulate/precuneus). States 2-6 show increased activity in the visual cortex (states 2 & 6), the sensorimotor network (state 3), and the left and right lateralised temporal lobes (states 4 & 5). States 7-8 show decreases in activity (blue in Figure 6.1) in parietal regions including the intraparietal sulcus (IPS; state 7), and visual cortex (state 8).

The temporal properties of each state were characterised from the state time courses, which indicate the points in time at which each state is active. By inspection it is evident that the states are short lived (Figure 6.2a). As with the single-subject HMM in **Chapter 5**, the temporal characteristics of each state were quantified in terms of their fractional occupancy (fraction of the total time spent in a state; Figure 6.2b), life times (time spent in a state before making a transition; Figure 6.2c) and interval lengths (time between consecutive state visits; Figure 6.2d). Average life times are between 100 and 200 ms. These life times are markedly shorter in duration than the time scales typically associated with resting state networks, which have previously been shown to be dominated by frequencies below 0.1 Hz (Cordes et al., 2001), although slightly longer than those inferred from the single-subject data (perhaps reflecting the increased temporal smoothness of the enveloped data). We also characterised variations in the rate at which these states are visited by computing the fractional occupancy within a ten second sliding window. These fractional occupancy time courses reveal slower temporal changes in the occurrence of the HMM states. These time courses are shown for all subjects in Figure 6.2e. It is clear that each state was represented in all subjects.

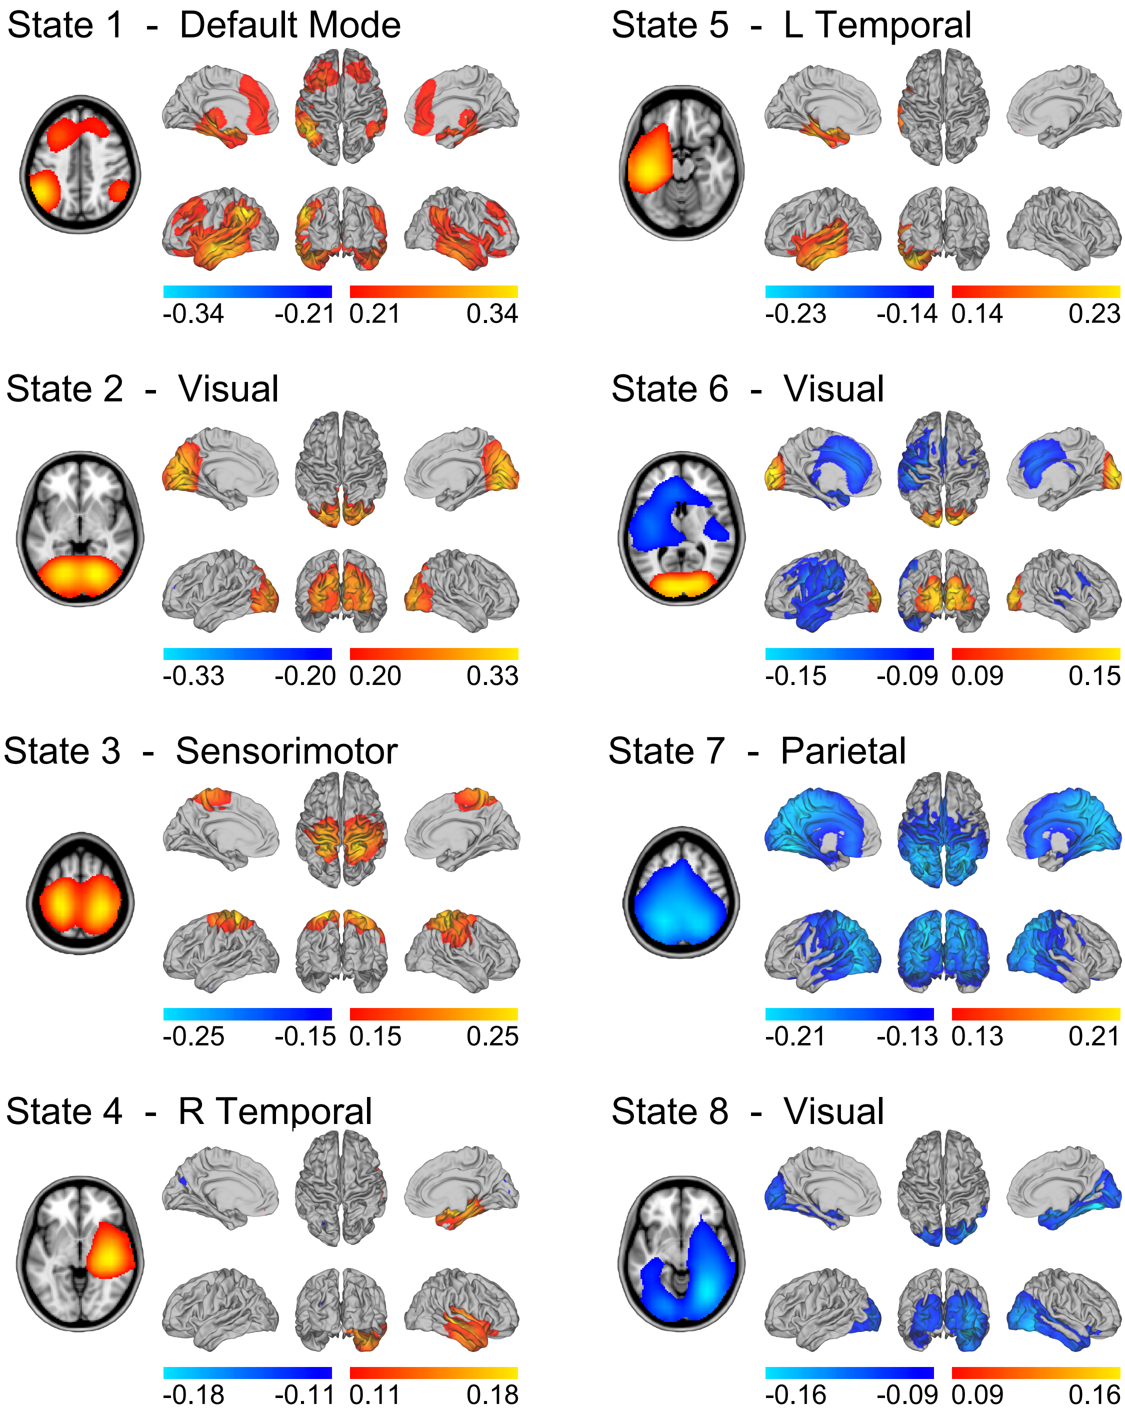

Figure 6.1: State-specific changes in band-limited amplitude. An 8 state HMM was inferred from temporally concatenated band-limited amplitude time courses (concatenated over nine subjects, 10 minutes each). The volumes and surface renderings show the partial correlation of each state time course with the envelope data at each voxel. The correlation values have been thresholded between 60 and 100% of the maximum correlation for each state and the colour maps represent these ranges. Red/yellow and blue colours indicate positive and negative correlations respectively.

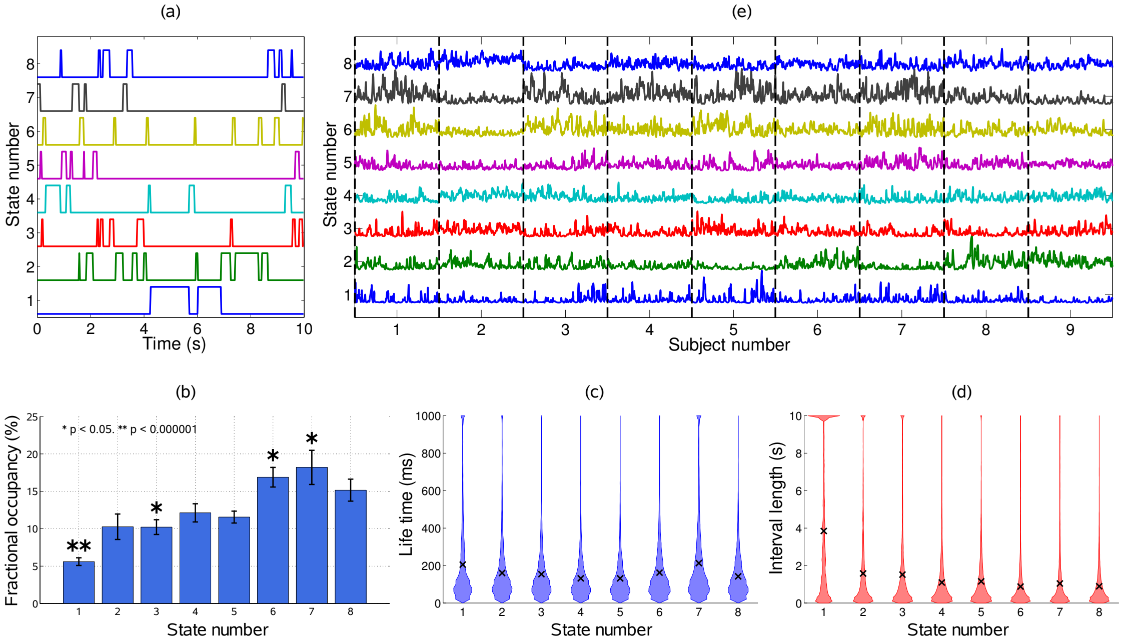

Figure 6.2: Temporal characteristics of the HMM states. (a) State time courses showing the most likely state at each time point for the first 10 seconds of data. (b) Fractional occupancy for each inferred state showing the mean and s.e.m. over subjects. The asterisks denote that the fractional occupancy of a state differs significantly from the other states. (c) life times, and (d) interval lengths for each inferred state. The filled areas in (c) and (d) represent the distribution of values and the black crosses show the mean. (e) Fractional occupancy of each state as a function of time over all subjects, derived by averaging each state time course within a 10 second sliding window (75% overlap between adjacent windows).

6.4.2 Effect of number of states

We assessed the effect of varying the number of states inferred by the HMM from 4 to 15 states. The model evidence (approximated by the negative of the free energy) monotonically increases up to 15 states suggesting that the Bayes-optimal model may require an even higher number of states (Figure 6.3a). However, inferring a larger number of states results in a decrease in minimum fractional occupancy (Figure 6.3b). It has been suggested that fractional occupancy may provide a more meaningful metric, since a sufficient amount of data is required for reliable estimation of the covariance matrix (Woolrich et al., 2013). While this was suggested in the context of deriving spatial filters for beamforming, the argument applies no less here, since the HMM must infer a covariance matrix for each multivariate normal distribution in the observation model. Increasing the number of states also has an effect on the mean life times of the inferred states. Varying the number of states from 4 to 15 results in mean lifetimes decreasing from 230 to 140 ms (Figure 6.3c). This suggests that the splitting of states that arises from increasing the model order does not result in states with fewer occurrences, but rather a splitting into shorter lived events.

We further assessed the effect of the model order on the topographies of the inferred states. Figure 6.4 shows partial correlation maps corresponding to HMMs inferred with 4 to 14 states. Increasing the number of states did not change the topographies of the most prominent RSN-like states nor did it reveal any new RSN-like topographies that were distinct from those inferred with 8 states, but rather resulted in the splitting of states into multiple similar maps. This suggests that there is no advantage to using more than 8 states for the purpose of identifying states corresponding to known RSNs.

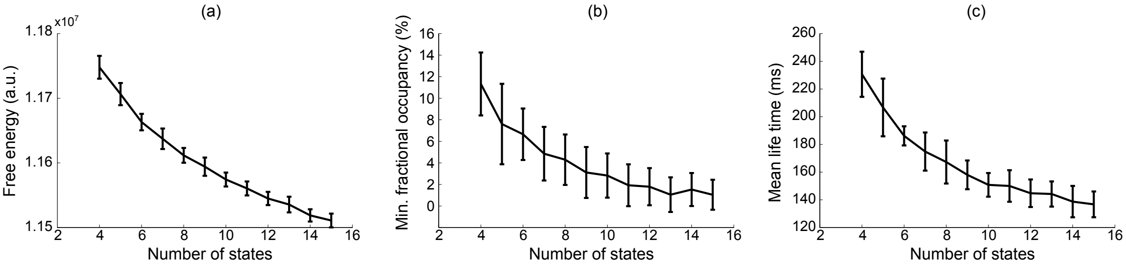

Figure 6.3: Effect of number of states on (a) free energy, (b) minimum fractional occupancy and (c) mean life time, computed over all inferred states and 50 realisations of each HMM inference. Error bars show the mean and s.e.m. over all subjects.

6.4.3 State-specific functional connectivity

Mapping state-specific changes in activity by computing the correlation between the state time courses and the amplitude envelope at each voxel captures changes in the mean and/or variance associated with each state. However, the multivariate normal distribution inferred for each state is defined not only by its mean and variance but also by a full covariance matrix and therefore captures interactions between all source pairs (i.e. functional connectivity). To characterise this, we defined a measure that captures state-specific changes in the global correlation structure. This *connectivity profile* represents differences in the envelope correlation between each voxel and the rest of the brain with respect to the other states.

As described previously, since the observations are principal components, the observation model covariance matrices do not provide an easily interpretable view of the interactions between different anatomical regions. Similarly, the $voxels \times voxels$ matrix that would result from projection back from the eigenspace would be too large to visualise and interpret. Furthermore, in both cases, elements of the covariance matrix would contain information not only about genuine functional connectivity but also covariance due to signal leakage between nearby sources. To this end we defined the *connectivity profile* $\mathbf{P}(k, i)$ for each state k and voxel i as the Euclidean distance between all elements $j \neq i$ from row i of the correlation matrix $\mathbf{C}_k$ (derived from the covariance matrix) for state k and the same row of the correlation matrices $\mathbf{C}_m$ for all other states from all other states $m = \{m_1, ..., m_K, m \neq k\}$:

$$\mathbf{P}(k,i) = \sqrt{\sum_{j \neq i}(|\mathbf{C}_k(i,j) - \frac{1}{K-1}\sum_{m \neq k}\mathbf{C}_m(i,j)|^2)} \tag{6.1}$$

Collapsing over rows of the correlation matrix in this way provides a summary statistic for each state and voxel that can be represented as a spatial map for each state. Moreover, by looking for changes in correlation rather that the correlation itself, the effect of signal leakage between voxels is mitigated since this correlation is stationary over all time points.

Figure 6.5 shows the connectivity profiles for each state alongside the partial correlation maps. These two measures provide complementary descriptions of state-specific neuronal activity by characterising both changes in amplitude and changes in functional connectivity. The topography of the connectivity profile is similar to that of the partial correlation for most states suggesting that states are driven not only by increases in amplitude but also by changes in functional connectivity. However, this measure also provides new information, for instance in state 8, where changes in motor cortex connectivity are apparent despite the main change in amplitude being confined to visual areas.

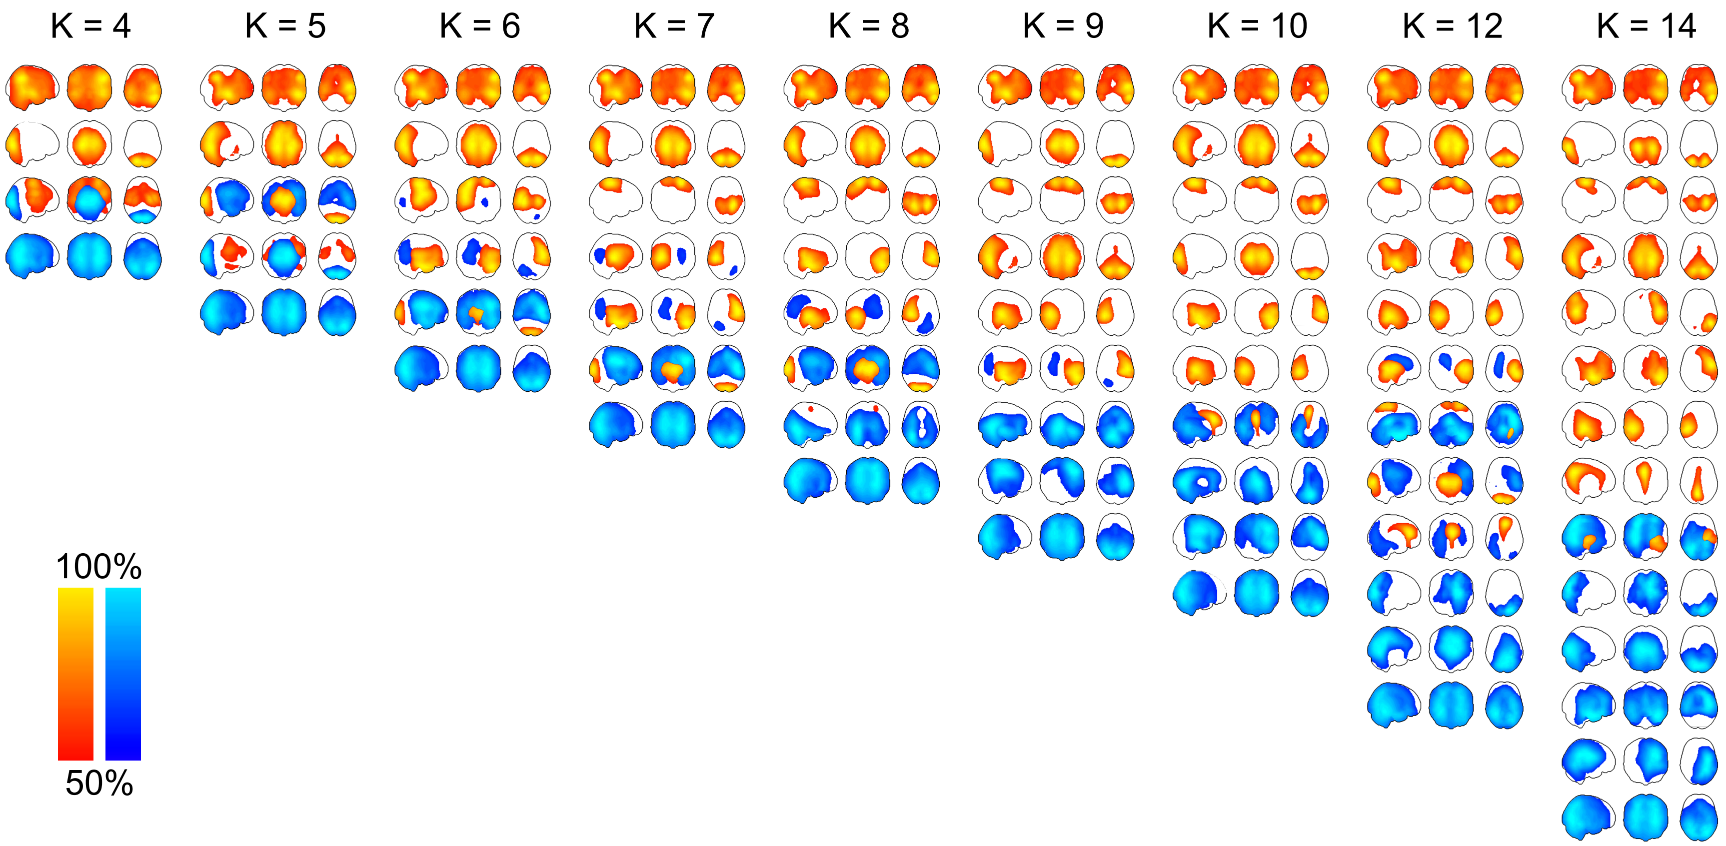

Figure 6.4: Maximum intensity projection maps showing the partial correlation computed between each state time course and the envelope data for a K state HMM for $K = 4$ to $K = 14$.

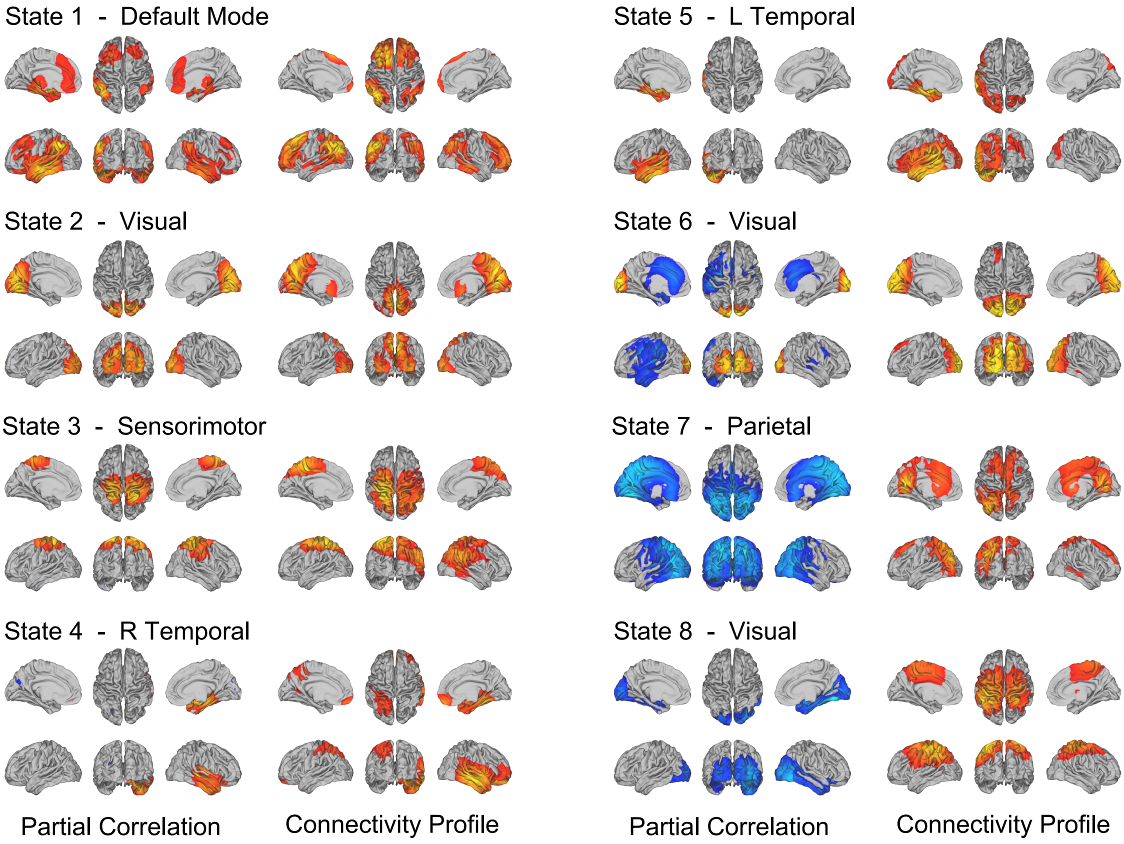

Figure 6.5: Partial correlation and connectivity profiles for each state. The connectivity profile for a particular state is computed for each voxel as the Euclidean distance between that voxel's row of the state-specific correlation matrix and the centroid of the corresponding row of correlation matrices for all other states. The partial correlation maps have been thresholded as before at 60% of the maximum value. The connectivity profiles have been thresholded at their 75th percentile.

6.4.4 Large scale networks of oscillatory activity are underpinned by rapid fluctuations

We have shown that whole brain spontaneous activity may be broken down into a set of distinct connectivity patterns that appear to be stable for periods of 100 to 200 milliseconds. To confirm that these brain states are consistent with coordinated fluctuations at these rapid time scales we performed a follow up analysis using the inferred state time courses. We reasoned that if there were coordinated fluctuations at the fastest time scales in the state time course, then low-pass filtered versions of the state time courses should do worse at explaining fluctuations in the data. Different low-pass filtered versions of the state time courses were obtained by computing fractional occupancy time courses using a range of time windows from 0.1 to 8 s. Note that for the shortest time windows the fractional occupancy time course approximates the state time course. These fractional occupancy time courses were then separately regressed onto the amplitude envelope time course from a representative voxel of the corresponding brain state (the voxel with greatest correlation with the state time course). This analysis reveals a peak in correlation for window widths between 200 and 400 ms demonstrating that, when using the fractional occupancy, the fastest time scales at which we can detect fluctuations in the amplitude envelopes are slower than those suggested by the HMM state life-times ($\sim$100 ms), but still much faster than has been previously shown (Figure 6.6a).

As a control, we repeated the HMM inference after first removing any potential high frequency network interactions from the data. The group-concatenated amplitude envelopes were low-pass filtered below 0.5 Hz to remove any higher frequency dynamics, and an 8 state HMM was inferred from these filtered envelopes. Despite removing the faster envelope fluctuations, a number of states were inferred with similar spatial topographies as the original HMM (Figure 6.6b). The life times of these states were longer, at around 1 second, reflecting the slower time scales of these low-pass filtered signals. Next, to test whether simply introducing high frequency noise could result in shorter life times, band-limited Gaussian noise (low-pass filtered below 10 Hz, reflecting the spectral content of the original envelopes, and therefore with the same non-independence properties between time points), were added to the envelopes prior to inferring the HMM. As before, states with spatial topographies similar to the original HMM states were identified (Figure 6.6c). However, the life times of these states were now much shorter, at around 0.3 s.

As with the real data, we then fitted a GLM with a single regressor that corresponded to the fractional occupancy time course computed using a range of time windows from 0.1 to 8 s. When applied to the low-pass filtered dataset (without added Gaussian noise; Figure 6.6b), this analysis reveals a peak in correlation for a window width of ∼1 s, consistent with the longer state life times. However, for the low-pass filtered dataset with added Gaussian noise (Figure 6.6c), while the state life times were reduced to ∼100 ms, the peak in correlation remains at a window width of ∼1 s. This demonstrates that in this control there are no detectable within-network amplitude fluctuations faster than 1 s, which is consistent with the applied 0.5 Hz low-pass filtering. This is in stark contrast to the real data (Figure 6.6a), where the analysis reveals a peak in correlation for window widths between 200 and 400 ms.

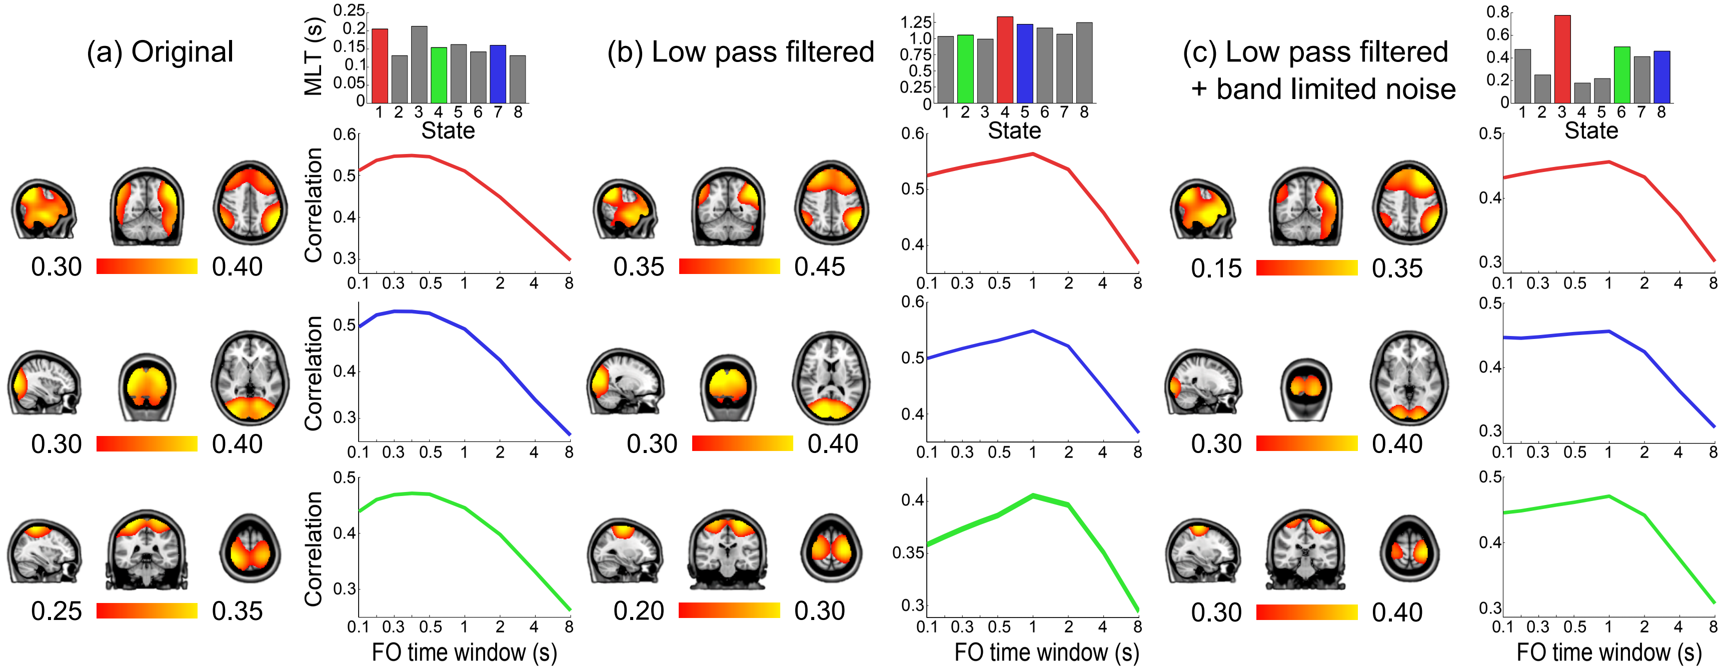

Figure 6.6: Analysis of the time scales that best reflect within-network envelope fluctuations. Correlation maps, mean life times and fractional occupancy time window dependency for 8 state HMMs inferred from (a) the original group-concatenated envelopes, (b) these envelopes low-pass filtered below 0.5 Hz and (c) these low-pass filtered envelopes with uncorrelated Gaussian noise added. Results are shown for the default mode network (top/red), sensorimotor network (middle/green), visual network (bottom/blue). The correlation maps were computed by fitting a GLM to the data with the state time courses for all states as regressors. The fractional occupancy time window dependency was computed by fitting a GLM to the data with the fractional occupancy (computed within different time windows) as a single regressor.

6.4.5 State transitions reveal cross-network interactions

Networks of whole brain spontaneous activity may be characterised not only in terms of their within-network activity, but also in terms of cross-network interactions. To assess the relationship between different functional networks in the context of the HMM we examined the relationship between the points in time at which different states are active. An important result from fMRI studies is the observation that the DMN exhibits anticorrelation with networks associated with attention demanding tasks, such as the dorsal attention network (DAN) (Fox et al., 2005; Smith et al., 2009). Of particular interest therefore is the relationship between state 1, which represents the DMN, and state 7, which we postulate may represent parietal regions of the DAN. While the correspondence between BOLD and electrophysiological activity has yet to be fully understood, it has been shown that alpha and beta power at rest correlate positively with BOLD in the DMN and negatively with BOLD in the DAN (Mantini et al., 2007). Accordingly, increased activity in the DMN (red/yellow in state 1) and decreased activity in parietal regions of the DAN (blue in state 7) would both correspond to an increase in the BOLD signal. Based on this reasoning, we hypothesised that the antagonistic behaviour of these two networks as shown in multiple fMRI studies would manifest in the inferred HMM states as an anticorrelation between the fractional occupancy time courses of these two states. In other words, periods of time in which the DMN state is frequently visited would coincide with periods of time in which the putative DAN state is rarely visited (and vice versa).

To this end we computed the correlation coefficient between the fractional occupancy time courses for each state (computed within 10 second sliding windows as shown in Figure 6.2e). Positive correlations between a pair of states indicate that the two states are visited more frequently during similar periods of time. As predicted we found strong antagonistic behaviour between the DMN (state 1) and the putative DAN (state 7), suggesting an electrophysiological basis to the anticorrelated nature of these networks (Figure 6.7a).

We further assessed the relationship between different functional networks by examining the transitions between the inferred states. These transitions may be represented in the form of a transition matrix, where each row represents the probability of transitioning to any other state given the current state (Figure 6.7b). There is a clear structure to the matrix showing that transitions between certain pairs of states are more likely than others. A number of these transitions are intuitive, for example the strong probability of transitioning between visual states 2 and 6. There is a very low probability of transitioning between the DMN (state 1) and the putative DAN (state 7), raising the intriguing possibility that anticorrelation between these networks may arise from an inability of the system to transition directly between these two transient states.

6.4.6 Occurrence of transient states reflect temporal variability in functional connectivity

In **Chapter 4** we demonstrated that resting state functional connectivity exhibits temporal variability on a time scale of several seconds by computing envelope correlation within sliding time windows. Similar results have also been demonstrated in previous work, both from MEG and fMRI (de Pasquale et al., 2010; Allen et al., 2012; Chang & Glover, 2010). To investigate whether any relationship exists between temporal variability in amplitude correlation and the occurrence of states inferred from the HMM, we computed the sliding window envelope correlation between different nodes of the DMN and compared the resulting correlation time courses with changes in the fractional occupancy of the inferred DMN state. Re-

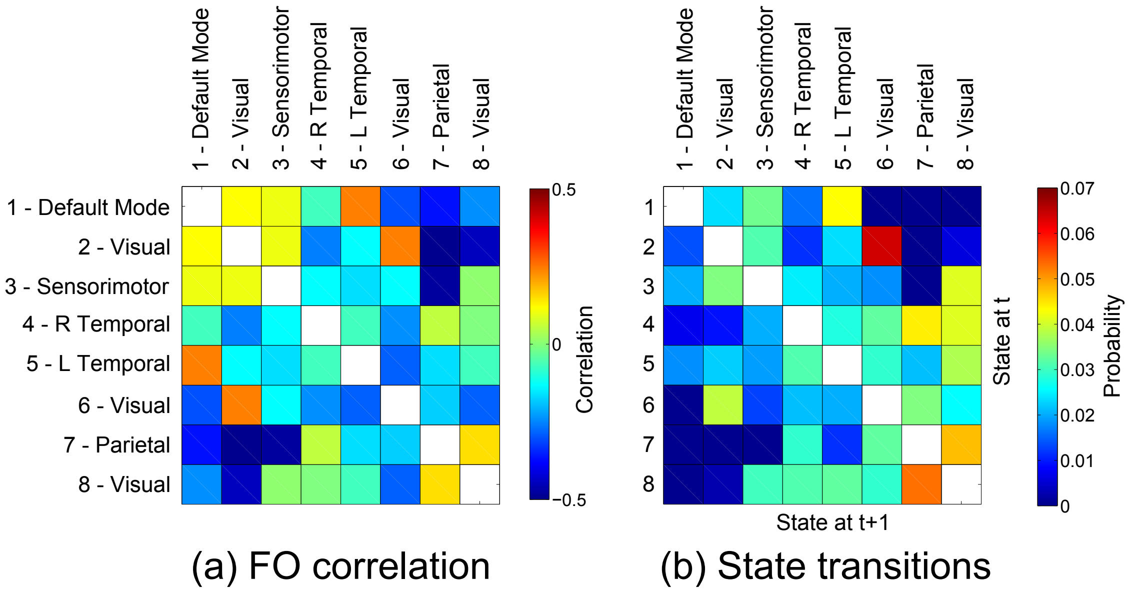

Figure 6.7: Relationship between states. (a) Correlation matrix between the fractional occupancy time courses of each state. Positive correlations between a pair of states indicate that the two states are visited more frequently during similar periods of time (b) State transition matrices for the group HMM. The matrix shows the probabilities of transitioning to any particular state given the current state. The probability of remaining in the same state has been excluded from each matrix (shown in white).

gions of interest (ROIs) were defined in the centres of the nodes corresponding to the inferior parietal lobule (IPL), medial frontal gyrus (MFG) and medial temporal lobe (MTL) in both hemispheres (Figure 6.8a). The envelope of oscillatory activity was computed at these six locations for each subject. With reference to previous findings (de Pasquale et al., 2010), envelope correlation was computed within ten second sliding windows between all six ipsilateral ROI pairs and all three contralateral ROI pairs. The HMM state corresponding to the DMN was identified and the fractional occupancy was computed within the same ten second sliding window. The time courses of these two measures are shown for a single subject in Figure 6.8b, where the DMN pair shown is the right IPL and right MFG. In line with previous findings, it is evident that the envelope correlation between different network nodes alternates between periods of high and low correlation (de Pasquale et al., 2010). Interestingly, periods of high envelope correlation are generally associated with an increase in the fractional occupancy of the DMN state, suggesting that fluctuations in interregional functional connectivity represent periods of time in which a particular transient network is frequently visited. This relationship may be quantified as the correlation between the two time series for each subject and node pair. There is a clear positive correlation between the two measures for all inter- and intra- hemispheric pairs (Figure 6.8c).

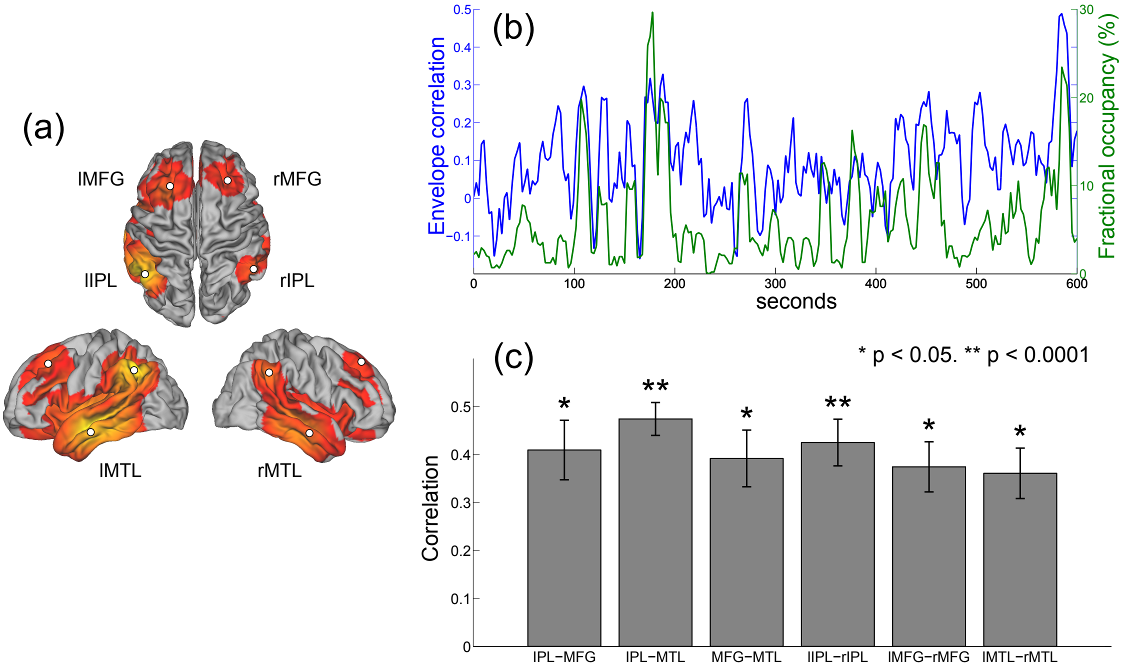

Figure 6.8: Comparison between HMM state occupancy and sliding window envelope correlation. (a) Six nodes identified from the DMN state (rIPL, lIPL, rMFG, lMFG, rMTG, lMTG). (b) time courses of rIPL-rMFG envelope correlation (blue) and DMN state fractional occupancy (green) for a single subject computed using a 10 second sliding window (75% overlap between adjacent windows). (c) correlation between the sliding window envelope correlation and fractional occupancy time courses for each ipsilateral pair (bilaterally homologous pairs from the left and right hemispheres have been averaged together) and each contralateral pair (mean and s.e.m. over subjects).

6.5 Discussion

We have characterised whole-brain spontaneous activity based on time-varying behaviour of source localised MEG signals. Using a distinct methodology that is able to resolve changes in functional connectivity at high temporal resolution, we identified short-lived (100-200 ms) states that represent unique spatiotemporal patterns of neural activity. The inferred states revealed spatial topographies that represent networks of spatially segregated brain regions that bear a strong similarity to seed based correlation, or ICA derived RSNs, in both MEG and fMRI. By assessing temporal changes in the occurrence of these states, we demonstrated that within-network functional connectivity is underpinned by neuronal dynamics that fluctuate much more rapidly than has previously been shown.

6.5.1 Large scale transient networks of spontaneous oscillatory activity

By inferring an HMM from the amplitude envelopes of group-concatenated data we identified spatial maps representing state-specific increases in oscillatory amplitude in anatomical locations corresponding to well-known RSNs. These networks include the default mode, visual and sensorimotor networks and parietal regions of the DAN.

In addition to inferring the spatial covariance structure, the HMM also infers the time points at which each state is active. Each state was well represented over the group with all occupancies between 5 and 20 percent of the total time. The state life times were short, between 100 and 200 ms on average. These findings demonstrate that RSNs derived from conventional techniques such as ICA or seed-based correlation may be identified from discrete transient epochs that represent only a small fraction of the total recording, confirming recent studies from both fMRI and MEG (de Pasquale et al., 2012; Liu & Duyn, 2013).

Due to the assumption of mutual exclusivity of the HMM states we should be cautious in automatically interpreting the fast switching between states in the form of short life times as a property of the underlying spontaneous activity. After all, an attractive hypothesis is that the brain's resting state activity consists of a finite number of unique networks that can overlap spatially and temporally (Smith et al., 2012). However, these relatively weak assumptions render the hypothetical networks unidentifiable using existing decomposition techniques (to the best of our knowledge), and so further constraints are currently needed to proceed. For example, in temporal ICA the assumption of temporal independence will discourage temporal overlap. In the HMM, the mutual exclusion approach will also prohibit temporal overlap. As a result the HMM states can perhaps be best thought of as representing the most dominant unique configurations of these hypothetical networks. Based on this interpretation, the state time courses provide a meaningful window on the underlying network dynamics by indicating the most dominant state at each point in time. Indeed, we have shown that these time courses can provide an insight into the relationship between different functional networks as demonstrated in Figure 6.7.

Nonetheless, the rate of occurrence of the inferred states as described by the fractional occupancy allows us to assess the time scales at which within-network amplitude fluctuations are detectable. By modelling these fluctuations by the fractional occupancy time course computed using a range of different window widths, we show that they are best described by sub-second temporal dynamics with coordinated fluctuations on the order of 200-400 ms. Importantly, these time scales were not found for control surrogate data where the fast network interactions were replaced by uncorrelated noise (Figure 6.6), demonstrating that within-network functional connectivity is underpinned by neuronal dynamics that fluctuate much more rapidly than has previously been shown.

6.5.2 Cross-network interactions

The HMM assumes that the states themselves are mutually exclusive. It might be thought that this assumption negates the possibility of the HMM providing insight into cross-network (i.e. across-state) interactions. While this is true at the fast (below 100 ms) within-state time scales, this does not mean that we cannot assess cross-network interactions at all. In particular, the HMM furnishes us with the probability of transitioning between different networks. This can tell us if there is a preference for the brain to move between two different networks, or an antagonism such that there is rarely a transition between two different networks. Indeed, the relationship between the DMN (state 1) and parietal areas of the DAN (state 7) was shown to be an example of the latter (Figure 6.7a-b). We have also demonstrated how cross-network interactions can be assessed using fractional occupancy time courses, which describe how frequently states are visited within 10 second time windows. By computing correlations between these slower time courses we can assess relationships between networks at the time scales more typically associated with previous investigations into long-range resting state interactions (Fox et al., 2005). Of particular interest is the anticorrelation between the fractional occupancy time courses of the DMN state and the putative DAN state (Figure 6.7a).

The fact that cross-network relationships present in the fractional occupancy time courses are maintained when looking at the faster time scales of the state transitions, suggests a link between these two time scales. However, with the current methodology it is not possible to establish a causal link between these time scales (i.e. whether infrequent transitions arise because these networks are anticorrelated at longer time scales, or that this anticorrelation arises due to these infrequent transitions).

6.5.3 Correspondence with BOLD RSNs

The states inferred by the HMM show some consistency with RSNs measured using fMRI. The spatial correspondence between the spatial maps obtained via the HMM and RSNs derived from application of ICA to BOLD are shown in Figure 6.9. A notable difference between these maps is the absence of the posterior cingulate cortex (PCC) or precuneus in the DMN state. One possible explanation is that the PCC may act as a functional "hub", such that it is not strongly represented in any one state. While the HMM allows spatial overlap, we have visualised the states in terms of state-specific amplitude changes, such that any node that is active during multiple states will be suppressed in the spatial maps. It is possible that, as a hub, the PCC has membership in the majority of states and is therefore poorly identified by mapping state-specific activity. However, it is not trivial to identify nodes that are active in multiple states because this is confounded by the spatial leakage due to the ambiguities in the source reconstruction (Schoffelen & Gross, 2009; Van Veen et al., 1997).

While there is strong evidence for hub-like behavior of the PCC/precuneus in terms of structural (Hagmann et al., 2008; Sporns et al., 2007) and functional (Buckner et al., 2009; Tomasi & Volkow, 2011) connectivity, evidence in MEG is more varied. Graph theoretical analyses of MEG band-limited power have revealed strong hubs in dorsal prefrontal cortex, lateral parietal cortex and temporal cortex (in essence those areas represented by state 1), but notably not the PCC (Hipp et al., 2012). Furthermore, the PCC was not found to be present in RSNs derived from application of temporal ICA to MEG data (Brookes et al., 2011b). Interest-

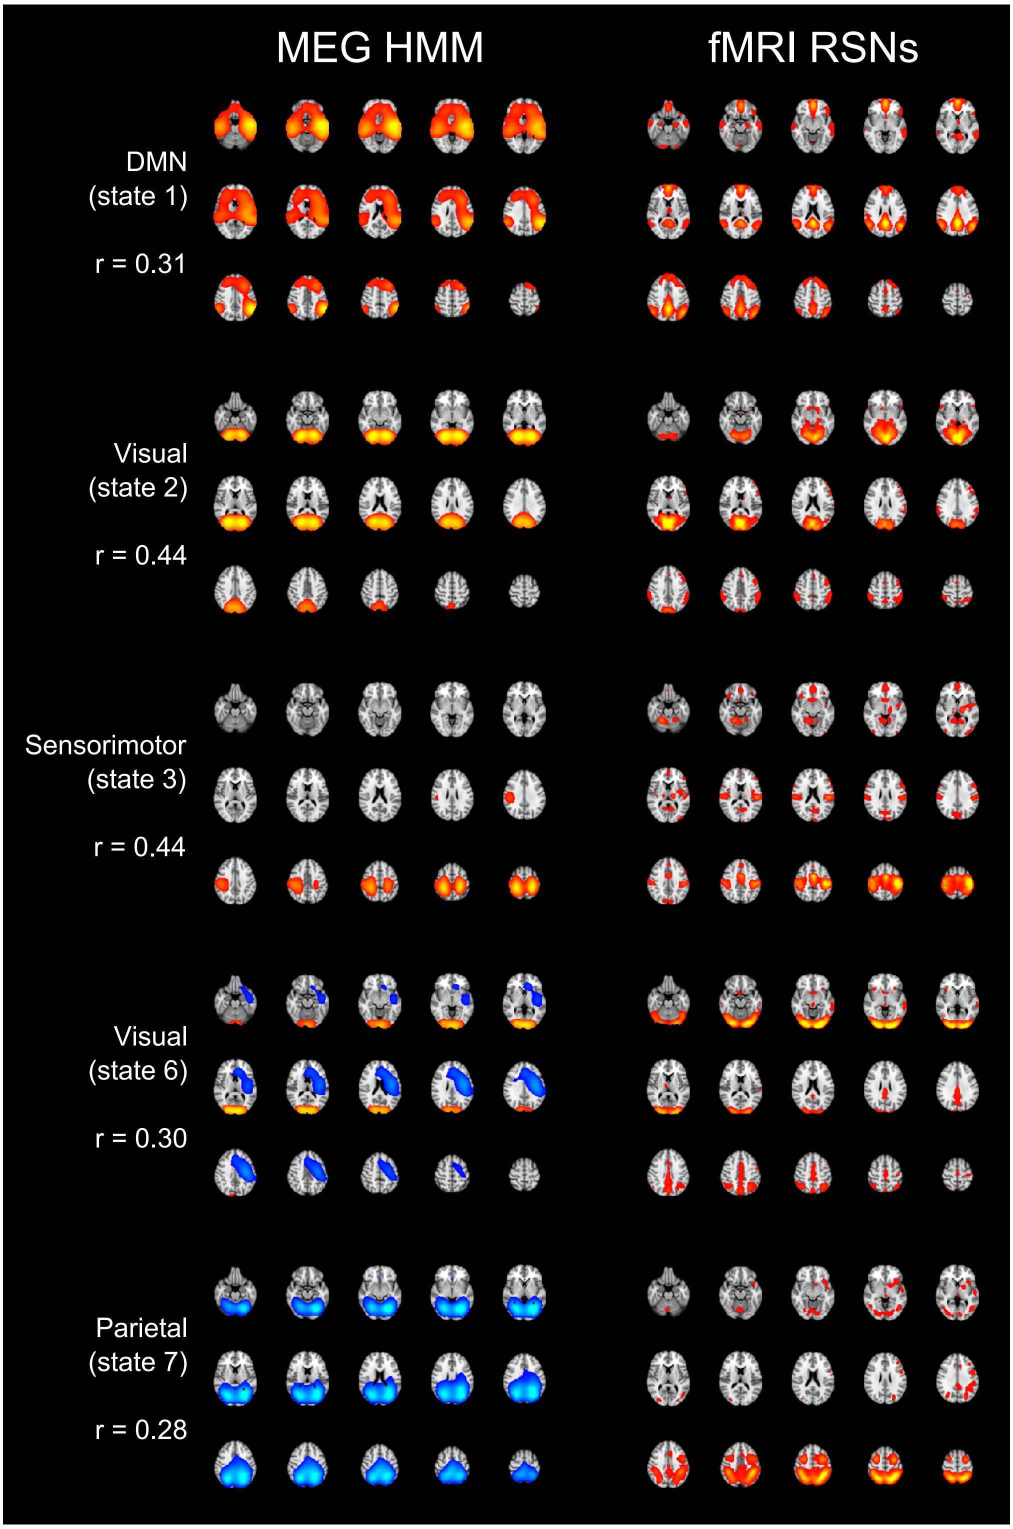

Figure 6.9: Spatial maps of five of the inferred states alongside a matched fMRI RSN (fMRI maps reproduced from (Smith et al., 2009)). For each HMM-RSN pair, the spatial correlation is shown alongside the maps.

ingly, a MEG-derived DMN network comprising the PCC has been shown using seed-based correlation with the PCC as a seed, but only when restricted to time points in which the network was maximally correlated (de Pasquale et al., 2012).

Another reason why the PCC may not be present in the HMM is due to the relative insensitivity of MEG to deeper sources. This hypothesis is supported by the fact that in the present study, and in (Brookes et al., 2011b) and (Hipp et al., 2012), the sensor array comprised only axial gradiometers, which measure the spatial gradient of the field and are relatively insensitive to deep sources. Conversely, in (de Pasquale et al., 2012) the array comprised magnetometers which have increased depth sensitivity. It is therefore possible that the PCC is not present simply because it does not generate a measurable signal.

Finally, it is worth considering to what extent the DMN or DAN, and the particular spatial nodes they incorporate, are definitive networks. The particular form of such networks is tied to the approaches used to identify them, e.g. spatial ICA. Indeed when other approaches with different assumptions are made, then different networks can be inferred. One good example of this is that, when using temporal ICA on fMRI data, no single DMN network is found (Smith et al., 2012). In short, the concept of a “network” and the form they take depends on the assumptions made in the data decomposition approach, and none should be considered necessarily better than the other (assuming they show similar objective performance, e.g. Bayesian model evidence and reproducibility), but instead they each offer different perspectives on the nature of the brain’s activity.

6.5.4 Relation to EEG microstates

The short life times inferred from the HMM raises the question of whether there exists a relationship between the states shown here and the *microstates* found from

EEG studies (Koenig et al., 2005; Lehmann et al., 1998). Microstates are quasi-stable topographies in which the distribution of EEG power over the scalp remains stable for periods of around 100 milliseconds. Clustering of these topographies into a limited number of classes has revealed that relatively few (typically four) unique maps are consistently identified across multiple time points and subjects (Lehmann et al., 1998). Segmentation of EEG scalp maps into microstates is based on finding repeating distributions of power across multiple recording sites and therefore captures similar interactions to those that drive the HMM. A clear distinction is that the HMM explicitly models the temporal dynamics and is therefore tuned to finding states that repeat in a predictable way. Another difference is that our approach exploits the superior spatial resolution available using MEG, by basing the inferred states on source space projections of MEG data. This makes the results more directly interpretable in the context of fMRI resting state networks.

Recently it has been suggested that EEG microstates may represent the electrophysiological signatures of resting state networks. A number of recent studies have sought to investigate this by correlating fMRI RSNs with simultaneously acquired EEG recordings (Britz et al., 2010; Musso et al., 2010; Yuan et al., 2012). Directly relating the EEG topologies of microstates to the underlying anatomy is challenging due to the inhomogeneous conductivity profile of the head. In this work, we inferred states from source reconstructed MEG activity and independently found transient states on the same temporal scale as EEG microstates, and with spatial patterns that match well known RSNs. Based on these findings and methodological similarities, we suggest that the HMM states inferred in this work may represent the source space counterparts of EEG microstates. However, without measuring EEG and MEG concurrently, a direct relationship cannot be confirmed.

6.5.5 Relationship with functional connectivity at slower time scales

The states inferred using the HMM represent networks of activity that show some consistency with those found previously using spatial ICA on FMRI data (Beckmann et al., 2005; Smith et al., 2009) and using temporal ICA on MEG data (Brookes et al., 2011b; Luckhoo et al., 2012). In both the ICA and HMM approaches, data from all voxels are used and therefore no prior spatial localisation assumptions (e.g. seed voxels) are required. In the case of temporal ICA, two regions will tend to be strongly represented in the same network if their time courses exhibit a strong time-averaged correlation with the same independent component. By contrast, in the case of the HMM, membership of two regions to a particular state ("network") depends only on there being a repeated pattern of covariance at those points in time at which the state is active, i.e. it is not time averaged over all time points. The results in this work, e.g. the fact that the two approaches result in similar networks, are consistent with the idea that ICA RSNs derive from the transient time-varying behaviour captured by the HMM states. This idea is supported by the findings that seed based functional connectivity is increased if evaluated only within those time points at which the rest of the network is synchronised (de Pasquale et al., 2010). In this chapter we provide evidence of a relationship between the sliding window correlation computed between nodes of the DMN and the occurrence rate of the DMN state (Figure 6.8). Accordingly, changes in within-network envelope correlation may reflect variations in the frequency at which a particular connectivity state is visited. One explanation for this observation is the idea that electrophysiological data are characterised by scale-free dynamics that span from hundreds of milliseconds to tens of seconds (Van de Ville et al., 2010). This fractal property of neural dynamics may provide an explanation for the similar spatial correlation structure that exists between signals at different temporal scales. Hence, slow fluctuations in band-limited amplitude correlations that underlie MEG RSNs may capture similar physiological phenomena as the HMM states, but seen through different temporal filters.

6.5.6 Methodological considerations

In this chapter the data for individual subjects were temporally concatenated, yielding a single combined dataset from which the HMM was inferred. Group concatenation is widely used in unsupervised analyses of resting state activity, particularly in the case of ICA (Calhoun et al., 2009). Nonetheless it is important to bear in mind that this approach assumes there is anatomical correspondence between subjects. In the case of ICA, this assumption means that individuals share common group maps (Calhoun et al., 2009). In the case of the HMM approach, this assumption extends to the requirement that the data from different individuals within a particular state are drawn from the same multivariate normal distribution, and that the transitions between these states occur in a repeatable manner over subjects. This being said, there is no requirement that all states should be present in all subjects, however it is clear from the fractional occupancy time courses in Figure 6.2e that they are. In light of this assumption, states inferred from the group data may best be thought of as representing the spatiotemporal patterns of activity that occur most consistently over the group. An alternative strategy to group concatenation is to infer separate HMMs for each subject individually. This would allow the model to more freely adapt to individual subjects' patterns of activity and functional connectivity. However, a severe limitation of this approach is that there will not necessarily be a correspondence between states inferred from different subjects, making it difficult to perform subsequent analyses at the group level.

A limitation of the proposed technique is that HMM inference requires an *a priori* specification of the number of states, K. Bayesian inference techniques provide a means to test model order selection, by providing an approximation to the Bayesian model evidence via the free energy. In theory, it should be possible to pick the optimal number of states by selecting the model with the greatest (negative) free energy. In practice however, we observe that the free energy increases monotonically up to $K = 15$ states, suggesting that the Bayes-optimal model may require an even higher number of states (Figure 6.3). In the absence of a straightforward data driven approach to model order selection we opted instead to repeat the analysis for values of K from 4 to 15 and arbitrarily chose 8 states as the case to present here, which we believe represents a good trade-off between richness and redundancy. Results for different model orders are shown in Figure 6.4. Varying the number of states between 4 and 14 did not change the topographies of the most prominent RSN-like states. It is worth noting that a similar limitation exists for more established data driven decompositions such as ICA, in which the choice of model order is driven by the application; e.g. lower model orders are used to obtain the classic RSNs, and higher model orders are used to obtain finer grained parcellations for use in subsequent network analysis (Smith et al., 2011).

Finally, the observation model used in this work corresponded to a multivariate normal distribution. The assumption of a Gaussian observation models allows inference of the HMM to be made tractable to variational Bayesian inference and thus permits its application to large amounts of data (40 observations and 1.5 hours of time points). However, it should be recognised that modelling only the first and second order statistics under a Gaussian assumption is likely to be an oversimplification of the underlying network dynamics. The multivariate normal distribution is just one of many potential observation models that may be used in the context of the HMM. For example, binary observation models have previously been used in the context of modelling interictal spikes (Ossadtchi et al., 2005). Future work will focus on implementing a multivariate autoregressive (MVAR) model which can model time lagged dependencies between observations.

6.5.7 Functional role of transient synchronisation

The work presented here rests on the underlying assumption that resting state activity may be broken down into a set of distinct connectivity patterns that repeat over time and where only one functional state may be active at any one time. In other words, the states inferred by the HMM are mutually exclusive. While this assumption may be an oversimplification of the underlying network dynamics, the concept that there exist distinct functional connectivity states that recur at different points in time is compatible with computational models of neuronal connectivity (Deco et al., 2011) and observations from both fMRI (Allen et al., 2012) and EEG (Britz et al., 2010; Musso et al., 2010; Yuan et al., 2012). The idea that RSNs represent states in which distributed cortical areas synchronise transiently is also compatible with the idea of a "dynamic repertoire" of states that are continuously explored in order to more quickly adopt the network configuration optimal for a given impending input (Deco & Corbetta, 2011). This organisation of dynamic activity through transient spatial patterns of coordination may provide the flexibility required to adapt to the rapidly changing computational demands of cognitive processing (Bressler & Tognoli, 2006).

6.6 Summary

In this chapter we have described a novel approach to characterising the time varying nature of resting state connectivity using a hidden Markov model. By applying the method to source projected MEG recordings of resting state activity we identified transient functional connectivity states that reveal spatial topographies similar to those of several well-known RSNs. We assessed temporal changes in the occurrence of these states to demonstrate that within-network functional connectivity is underpinned by neuronal dynamics that fluctuate much more rapidly than has previously been shown. By further assessing cross-network interactions, we showed that the anticorrelation between the default mode and dorsal attention networks previously observed in fMRI RSNs is consistent with an inability of the system to transition directly between these two transient states. Finally we showed that temporal changes in within-network envelope correlation are related to changes in the occurrence rate of the inferred DMN state and suggested a mechanism by which rapidly changing functional states are related to the slower fluctuations in oscillatory power that we propose may extend to BOLD activity.

CHAPTER 7

Conclusions and future work

7.1 Summary

The focus of the work presented in this thesis was to develop methods for characterising functional connectivity during the resting state using MEG. In particular, we have argued that traditional analyses of functional connectivity (such as are typically applied to fMRI recordings) limit our ability to assess the temporal dynamics that underlie resting state networks, and considered how the superior temporal resolution of MEG may be exploited to better understand these dynamics.

Altered functional connectivity has been implicated in a number of neurological and psychological disorders including Alzheimer's diseases, Parkinson's disease and schizophrenia, among others. The resting state has emerged as a unique paradigm in which to study brain function, since its simplicity facilitates its use across both healthy and patient populations, as well as being readily applicable to different imaging modalities. We hope that the development of better tools for characterising brain connectivity will lead to a more complete understanding of the brain in both health and disease.

7.2 Summary of novel contributions

In **Chapter 3** we presented a framework for preprocessing MEG data, with a specific focus on removing sources of interference that may confound brain connectivity measures. In contrast to trial-wise data, where individual trials may be discarded by searching for trials that strongly depart from the norm, resting state data have no natural partitioning into epochs. Our framework is specifically designed for preprocessing continuous streams of resting state data and is based on a combination of manual artefact identification, ICA denoising, and muscle artefact identification and removal. This preprocessing procedure was used in conjunction with a pipeline for measuring functional connectivity to reproduce two key results that have been shown in previous studies. Firstly, we performed a whole-brain seed-based connectivity analysis of three resting state networks to show that we could identify the same spatial patterns of envelope correlation shown in previous studies (Brookes et al., 2011a; Hipp et al., 2012). Secondly, we reproduced findings that demonstrate the frequency dependency of this functional connectivity.

In **Chapter 4** we extended this analysis pipeline to characterise *time varying* functional connectivity between nodes of resting state networks. In particular, we showed that temporal variability in functional connectivity is significantly greater than would be expected if the coupling strength was stationary over time, and that this variability is greatest for those frequencies at which strong within-network coupling has previously been shown to exist. By computing measures of bimodality from the time course of envelope correlation, we further demonstrated that this connectivity is associated with bi-state behaviour, in which there is dynamic switching between periods of low functional connectivity and transient periods of high functional connectivity. A final key contribution of this chapter was the use of phase randomisation to assess the window widths at which this temporal variability may be resolved. Using this approach we showed that time windows of at least 50 cycles were required to identify significant temporal variability in resting state functional connectivity.

In **Chapter 5** we developed a novel methodology for assessing non-stationary functional connectivity based on a hidden Markov model (HMM). This approach provides a data driven decomposition of MEG data into a number of discrete states that represent transient patterns of spontaneous activity. By applying this approach to simulated data in which the spatial and temporal characteristics were known, we showed that the HMM is able to recover transient patterns of covariance, and the time points at which they occur. We further validated the method by applying it to real MEG data for a simple visual task. Using this simple task we demonstrated that the method is able to identify the spatial and temporal signatures of the P100 evoked response. Finally, the HMM approach was applied to resting state data for single subjects, which revealed short-lived states with spatial patterns of coactivation that resembled RSNs.

In **Chapter 6** we demonstrated how the HMM may be applied to group-concatenated data and identified transient states similar to those shown in **Chapter 5**, but which are consistent both within and across individuals. These results provide evidence that resting state functional connectivity is underpinned by neuronal dynamics that fluctuate much more rapidly than has previously been shown. We further characterised these transient states in terms of their temporal properties, and assessed cross-network interactions. In particular we showed that antagonism between the DMN and DAN, a prominent feature of fMRI RSNs, is captured by the HMM state transitions.

7.3 Future work

In this section, we consider how the work presented in this thesis may be extended in the future.

7.3.1 Modelling inter-subject variability

In **Chapter 6** we demonstrated the application of the HMM to identify brain states that were consistent across a group of subjects. This was achieved using temporal concatenation of the individual subjects' time courses, as is commonly used in ICA, followed by a single inference of the HMM from this aggregated data matrix. A disadvantage of this approach is that it does not produce an estimate of individual functional connectivity patterns, and the inter-subject variability is not explicitly modelled. Furthermore, in order to concatenate source reconstructed MEG data, it was necessary to work with the amplitude envelope, rather than the raw time series themselves. This is because of the ambiguity in the polarity of the estimated dipole.

An alternative approach to temporal concatenation would be to infer states/networks for individual subjects, and then combine the subject level network maps, for instance using clustering techniques. However, there will not necessarily be a straightforward mapping between different subjects. Moreover, this sequential approach is likely to be less robust, since the inference of each subject's networks does not benefit from knowledge of the other subjects' data.

A better approach would be to adapt the HMM to use a hierarchical Bayesian model that includes both group and subject levels. This approach would allow a statistical dependency between the subject and group level to be included which would provide a natural regularisation of the inference of a single subject using information from the entire group. The development of such approaches is currently an active research topic in neuroimaging.

7.3.2 Application to task

The focus of this thesis has been the identification and characterisation of patterns of spontaneous activity that occur in the resting brain. However, the techniques developed could equally well be applied to task data. In **Chapter 5** we demonstrated the application of the hidden Markov model to a simple visual task and demonstrated that it was possible to identify a state corresponding to the P100 evoked response. While this demonstrated the ability of the HMM to extract task-related brain states, this result was for a single session of data and more work is required to extend this to capture group-averaged behaviour.

One interesting avenue of research is the identification of functional networks that are associated with, or modulated by, cognitive tasks. Recent studies using MEG have shown how spatial maps derived using ICA may be used as a spatial basis set to identify event related time courses or time-frequency decompositions that represent activity within a particular network (Brookes et al., 2012a; Luckhoo et al., 2012). These techniques may be used to characterise how a particular task affects activity within a network, rather than within a particular voxel. A similar approach could be employed in the context of the HMM, by using the state-specific maps as a spatial basis set.

7.3.3 Application to disease

Another interesting extension of this work would to apply the methods developed herein to resting state activity recorded from different disease populations. Disrupted brain connectivity has been shown to be a feature of many neurological and psychiatric diseases. Accordingly, comparing functional connectivity patterns between healthy and disease populations may help to better understand the causes or symptoms of these diseases.

Recent theses by Luckhoo (2013) and Heise (2013) have demonstrated differences in functional brain networks between populations of carriers of a risk gene for Alzheimer's disease and non-carriers. Their approach used spatial basis sets derived from application of ICA to MEG or fMRI data concatenated across both populations. Differences between the two populations were then identified by reconstructing network maps for each sub-group. Such an approach is readily applicable in the context of the HMM, where it would be possible to interrogate not only spatial differences between different populations, but also temporal differences in terms of the state life times or fractional occupancy.

Finally, the non-stationary measures of functional connectivity developed in this thesis could offer new insights into a variety of neurological disorders, by assessing how the temporal dynamics of spontaneous activity are altered in disease. Emerging work applying dynamic functional connectivity measures to data collected from patients with schizophrenia (Damaraju et al., 2012) and bipolar disorder (Rashid et al., 2013) have shown that connectivity states identified from healthy controls switch more frequently than those identified from patients. Similarly, a study of Alzheimer's disease showed differences in time-varying functional connectivity within the DMN between patients and age-matched controls (Jones et al., 2012). These findings highlight the importance of evaluating changes in connectivity as well as simply characterising connection strength.

7.3.4 Multimodal analysis

This thesis and other studies have shown that MEG is a valuable tool for characterising resting state functional connectivity, particularly in light of its excellent temporal resolution. However, the spatial resolution of MEG is still poor compared to fMRI. Methods for combining datasets acquired using both modalities will likely be of great importance in the future. Such multimodal approaches are already possible, for instance using concurrent EEG and fMRI. In MEG, simultaneous acquisition is not yet feasible but the spatial information provided by fMRI may still be used to inform MEG analyses, for instance through the use of fMRI derived spatial priors for source localisation or functional parcellation. The development of such techniques is ever more pressing with the emerging availability of large neuroimaging datasets such as the NIH Human Connectome Project. This project will provide researchers with highly detailed connectivity data from multiple modalities including fMRI, diffusion tractography and MEG for over a thousand subjects (Smith et al., 2013a; Larson-Prior et al., 2013). The ability to combine these modalities in a principled way will likely provide insights into brain connectivity that cannot be achieved using one modality alone.

7.4 Concluding remarks

How the brain integrates and processes information remains one of the leading mysteries of science. The shift from perceiving the brain as a passive analyser to a complex dynamical system highlights the need to better understand the role of spontaneous brain activity and the interactions between different cortical regions. In recent years MEG has emerged as a valuable tool for characterising brain functional connectivity. This thesis represents significant work in developing methods to exploit the temporal resolution of MEG in order to assess the temporal dynamics of resting state functional connectivity at time scales faster than was previously possible.

APPENDIX A

List of journal and conference publications

Baker, A., Brookes, M., Rezek, I., Smith, S. M., Behrens, T. E. J., Probert Smith, P., & Woolrich, M. W. (2014). Fast transient networks in spontaneous human brain activity. eLife, In Press

Brookes, M. J., O'Neill, G. C., Hall, E. L., Woolrich, M. W., Baker, A., Palazzo-Corner, S., Robson, S. E., Morris, P. G., & Barnes, G. R. (2014). Measuring temporal, spectral and spatial changes in electrophysiological brain network connectivity. NeuroImage, In Press

Mohseni, H. R., Kringelbach, M. L., Woolrich, M. W., Baker, A., Aziz, T. Z., & Probert-Smith, P. (2013). Non-Gaussian probabilistic MEG source localisation based on kernel density estimation. NeuroImage, 87, 444–464

Cabral, J., Luckhoo, H., Woolrich, M., Joensson, M., Mohseni, H., Baker, A., Kringelbach, M. L., & Deco, G. (2013). Exploring mechanisms of spontaneous functional connectivity in MEG: How delayed network interactions lead to structured amplitude envelopes of band-pass filtered oscillations. NeuroImage, 90, 423–435

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

Baker, A., Brookes, M. J., Luckhoo, H., Rezek, I., Brodersen, P., Palazzo-Corner, S., Probert Smith, P., & Woolrich, M. (2013). Inferring Transiently Synchronising Networks using a Hidden Markov Model. In 19th Annual Meeting of the Organization for Human Brain Mapping, Seattle

Baker, A., Luckhoo, H., Brookes, M., Smith, P. P., & Woolrich, M. (2012b). Investigating the temporal dynamics of resting state connectivity using MEG. In 18th International Conference on Biomagnetism, Paris

Baker, A., Luckhoo, H., Brookes, M., Smith, P. P., & Woolrich, M. (2012a). Investigating the temporal dynamics of resting state connectivity using magnetoencephalography. In 18th Annual Meeting of the Organization for Human Brain Mapping, Beijing

APPENDIX B

List of acronyms

ACC	Anterior Cingulate Cortex
AD	Alzheimer's Disease
AR	Autoregressive
BEM	Boundary Element Model
BIC	Bayes Information Criterion
BLP	Band Limited Power
BOLD	Blood Oxygenation Level Dependent
CAPs	Coactivation Patterns
DAN	Dorsal Attention Network
DCM	Dynamic Causal Modelling
DMN	Default Mode Network
DTI	Diffusion Tensor Imaging
ECG	Electrocardiography / Electrocardiogram
EEG	Electroencephalography / Electroencephalogram
EMG	Electromyography / Electromyogram
EOG	Electrooculography / Electrooculogram
EPSP	Excitatory Post-Synaptic Potential
ERF	Event Related Field
FC	Functional Connectivity
FEM	Finite Element Model
fMRI	Functional Magnetic Resonance Imaging
FWHM	Full Width at Half Maximum
GABA	Gamma-Aminobutyric Acid
GLM	General Linear Model

HMM	Hidden Markov Model
HPI	Head Position Indicator
ICA	Independent Component Analysis
IPL	Inferior Parietal Lobule
IPS	Inferior Parietal Sulcus
IPSP	Inhibitory Post-Synaptic Potential
LCMV	Linearly Constrained Minimum Variance
LFP	Local Field Potential
MCE	Minimum Current Estimate
MEG	Magnetoencephalography / Magnetoencephalogram
MFG	Medial Frontal Gyrus
MNE	Minimum Norm Estimate
MNI	Montreal Neurological Institute
MRI	Magnetic Resonance Imaging
MSR	Magnetically Shielded Room
MTL	Medial Temporal Lobe
MVAR	Multivariate Autoregressive
NIH	National Institutes of Health
OEF	Oxygen Extraction Fraction
PCA	Principal Component Analysis
PCC	Posterior Cingulate Cortex
PD	Parkinson's Disease
PET	Positron Emission Tomography
PSP	Post-Synaptic Potential
ROI	Region Of Interest
RSN	Resting State Network
SNR	Signal-to-Noise Ratio
SPL	Superior Parietal Lobule
SPM	Statistical Parametric Mapping
SQUID	Superconducting Quantum Interference Device
SSS	Signal Space Separation
tIC	temporal Independent Component
VB	Variational Bayes

References

Ahlfors, S. P., Han, J., Belliveau, J. W., & Hämäläinen, M. S. (2010). Sensitivity of MEG and EEG to source orientation. Brain topography, 23, 227–32.

Ahonen, A. I., Hämäläinen, M. S., Ilmoniemi, R. J., Kajola, M. J., Knuutila, J. E., Simola, J. T., & Vilkman, V. A. (1993). Sampling theory for neuromagnetic detector arrays. IEEE Transactions on Biomedical Engineering, 40, 859–869.

Allen, E. a., Damaraju, E., Plis, S. M., Erhardt, E. B., Eichele, T., & Calhoun, V. D. (2012). Tracking Whole-Brain Connectivity Dynamics in the Resting State. Cerebral cortex, pp. 1–14.

Arieli, a., Sterkin, a., Grinvald, a., & Aertsen, a. (1996). Dynamics of ongoing activity: explanation of the large variability in evoked cortical responses. Science (New York, N.Y.), 273, 1868–71.

Attal, Y. & Schwartz, D. (2013). Assessment of subcortical source localization using deep brain activity imaging model with minimum norm operators: a MEG study. PloS one, 8, e59856.

Baccalá, L. a. & Sameshima, K. (2001). Partial directed coherence: a new concept in neural structure determination. Biological cybernetics, 84, 463–74.

Baillet, S., Mosher, J. C., & Leahy, R. M. (2001). Electromagnetic Brain Mapping. IEEE Signal Processing Magazine, 18, 14–30.

Baker, A., Brookes, M., Rezek, I., Smith, S. M., Behrens, T. E. J., Probert Smith, P., & Woolrich, M. W. (2014). Fast transient networks in spontaneous human brain activity. eLife, In Press.

Baker, A., Brookes, M. J., Luckhoo, H., Rezek, I., Brodersen, P., Palazzo-Corner, S., Probert Smith, P., & Woolrich, M. (2013). Inferring Transiently Synchronising Networks using a Hidden Markov Model. In 19th Annual Meeting of the Organization for Human Brain Mapping, Seattle.

Baker, A., Luckhoo, H., Brookes, M., Smith, P. P., & Woolrich, M. (2012a). Investigating the temporal dynamics of resting state connectivity using magnetoencephalography. In 18th Annual Meeting of the Organization for Human Brain Mapping, Beijing.

Baker, A., Luckhoo, H., Brookes, M., Smith, P. P., & Woolrich, M. (2012b). Investigating the temporal dynamics of resting state connectivity using MEG. In 18th International Conference on Biomagnetism, Paris.

Barnes, G. R. & Hillebrand, A. (2003). Statistical flattening of MEG beamformer images. Human brain mapping, 18, 1–12.

Bassett, D. S., Meyer-Lindenberg, A., Achard, S., Duke, T., & Bullmore, E. (2006). Adaptive reconfiguration of fractal small-world human brain functional networks. Proceedings of the National Academy of Sciences of the United States of America, 103, 19518–23.

Beckmann, C. F., DeLuca, M., Devlin, J. T., & Smith, S. M. (2005). Investigations into resting-state connectivity using independent component analysis. Philosophical transactions of the Royal Society of London. Series B, Biological sciences, 360, 1001–13.

Berger, H. (1929). Uber das Elektrenkephalogramm des Menschen. Archiv fur Psychiatrie und Nervenkrankheiten, 87, 527–580.

Biswal, B., Yetkin, F. Z., 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.

Biswal, B. B. (2012). Resting state fMRI: a personal history. NeuroImage, 62, 938–44.

Bressler, S. L. & Tognoli, E. (2006). Operational principles of neurocognitive networks. International Journal of Psychophysiology, 60, 139–48.

Britz, J., Landis, T., & Michel, C. M. (2009). Right parietal brain activity precedes perceptual alternation of bistable stimuli. Cerebral cortex, 19, 55–65.

Britz, J., Van De Ville, D., & Michel, C. M. (2010). BOLD correlates of EEG topography reveal rapid resting-state network dynamics. NeuroImage, 52, 1162–70.

Brodmann, K. (1909). Vegleichende Lokalisationslehre der Grosshirnde. (Leipzig: Johann Ambrosius Barth).

Brookes, M. J., Gibson, A. M., Hall, S. D., Furlong, P. L., Barnes, G. R., Hillebrand, A., Singh, K. D., Holliday, I. E., Francis, S. T., & Morris, P. G. (2004). A general linear model for MEG beamformer imaging. NeuroImage, 23, 936–46.

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, 1082–104.

Brookes, M. J., Liddle, E. B., Hale, J. R., Woolrich, M. W., Luckhoo, H., Liddle, P. F., & Morris, P. G. (2012a). Task induced modulation of neural oscillations in electrophysiological brain networks. NeuroImage, 63, 1918–30.

Brookes, M. J., O'Neill, G. C., Hall, E. L., Woolrich, M. W., Baker, A., Palazzo-Corner, S., Robson, S. E., Morris, P. G., & Barnes, G. R. (2014). Measuring temporal, spectral and spatial changes in electrophysiological brain network connectivity. NeuroImage, In Press.

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, 1788–802.

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. Proceedings of the National Academy of Sciences of the United States of America, 108, 16783–8.

Brookes, M. J., Woolrich, M. W., & Barnes, G. R. (2012b). Measuring functional connectivity in MEG: A multivariate approach insensitive to linear source leakage. NeuroImage, 63, 910–20.

Brown, P. (2003). Oscillatory Nature of Human Basal Ganglia Activity : Relationship to the Pathophysiology of Parkinson's Disease. Movement disorders, 18, 357–363.

Buckner, R. L., Andrews-Hanna, J. R., & Schacter, D. L. (2008). The brain's default network: anatomy, function, and relevance to disease. Annals of the New York Academy of Sciences, 1124, 1–38.

Buckner, R. L., Sepulcre, J., Talukdar, T., Krienen, F. M., Liu, H., Hedden, T., Andrews-Hanna, J. R., Sperling, R. a., & Johnson, K. a. (2009). Cortical hubs revealed by intrinsic functional connectivity: mapping, assessment of stability, and relation to Alzheimer's disease. The Journal of Neuroscience, 29, 1860–73.

Buzsáki, G. & Draguhn, A. (2004). Neuronal oscillations in cortical networks. Science (New York, N.Y.), 304, 1926–9.

Cabral, J., Hugues, E., Sporns, O., & Deco, G. (2011). Role of local network oscillations in resting-state functional connectivity. NeuroImage, 57, 130–9.

Cabral, J., Luckhoo, H., Woolrich, M., Joensson, M., Mohseni, H., Baker, A., Kringelbach, M. L., & Deco, G. (2013). Exploring mechanisms of spontaneous functional connectivity in MEG: How delayed network interactions lead to structured amplitude envelopes of band-pass filtered oscillations. NeuroImage, 90, 423–435.

Calhoun, V. D., Liu, J., & Adali, T. (2009). A review of group ICA for fMRI data and ICA for joint inference of imaging, genetic, and ERP data. NeuroImage, 45, S163–72.

Chang, C. & Glover, G. H. (2010). Time-frequency dynamics of resting-state brain connectivity measured with fMRI. NeuroImage, 50, 81–98.

Chang, C., Liu, Z., Chen, M. C., Liu, X., & Duyn, J. H. (2013). EEG correlates of time-varying BOLD functional connectivity. NeuroImage, 72, 227–236.

Christoff, K., Gordon, A. M., Smallwood, J., Smith, R., & Schooler, J. W. (2009). Experience sampling during fMRI reveals default network and executive system contributions to mind wandering. Proceedings of the National Academy of Sciences of the United States of America, 106, 8719–24.

Cohen, D. (1968). Magnetoencephalography: evidence of magnetic fields produced by alpha-rhythm currents. Science, 161, 784–786.

Cohen, D. (1972). Magnetoencephalography: detection of the brain's electrical activity with a superconducting magnetometer. Science, 175, 664–6.

Cordes, D., Haughton, V. M., Arfanakis, K., Carew, J. D., Turski, P. a., Moritz, C. H., Quigley, M. a., & Meyerand, M. E. (2001). Frequencies contributing to functional connectivity in the cerebral cortex in "resting-state" data. AJNR. American journal of neuroradiology, 22, 1326–33.

Dalal, S. S., Guggisberg, A. G., Edwards, E., Sekihara, K., Findlay, A. M., Canolty, R. T., Berger, M. S., Knight, R. T., Barbaro, N. M., Kirsch, H. E., & Nagarajan, S. S. (2008). Five-dimensional neuroimaging: localization of the time-frequency dynamics of cortical activity. NeuroImage, 40, 1686–700.

Damaraju, E., Turner, J., Preda, A., Van Erp, T., Mathalon, D., Ford, J., Potkin, S., & Calhoun, V. (2012). Static and Dynamic Functional Network Connectivity

During Resting State in Schizophrenia. In American College of Neuropsychopharmacology, Hollywood, CA.

Damoiseaux, J. S., Rombouts, S. A. R. B., Barkhof, F., Scheltens, P., Stam, C. J., Smith, S. M., & Beckmann, C. F. (2006). Consistent resting-state networks across healthy subjects. Proceedings of the National Academy of Sciences of the United States of America, 103, 13848–13853.

de Pasquale, F., Della Penna, S., Snyder, A. Z., Lewis, C., Mantini, D., Marzetti, L., Belardinelli, P., Ciancetta, L., Pizzella, V., Romani, G. L., & Corbetta, M. (2010). Temporal dynamics of spontaneous MEG activity in brain networks. Proceedings of the National Academy of Sciences of the United States of America, 107, 6040–5.

de Pasquale, F., Della Penna, S., Snyder, A. Z., Marzetti, L., Pizzella, V., Romani, G. L., & Corbetta, M. (2012). A cortical core for dynamic integration of functional networks in the resting human brain. Neuron, 74, 753–64.

Deco, G. & Corbetta, M. (2011). The dynamical balance of the brain at rest. The Neuroscientist, 17, 107–23.

Deco, G., Jirsa, V. K., & McIntosh, A. R. (2011). Emerging concepts for the dynamical organization of resting-state activity in the brain. Nature reviews. Neuroscience, 12, 43–56.

Deco, G., Ponce-Alvarez, A., Mantini, D., Romani, G. L., Hagmann, P., & Corbetta, M. (2013). Resting-state functional connectivity emerges from structurally and dynamically shaped slow linear fluctuations. The Journal of Neuroscience, 33, 11239–52.

Engel, A. K., Fries, P., & Singer, W. (2001). Dynamic predictions: oscillations and synchrony in top-down processing. Nature Reviews Neuroscience, 2, 704–716.

Engel, A. K., Gerloff, C., Hilgetag, C. C., & Nolte, G. (2013). Review Intrinsic Coupling Modes : Multiscale Interactions in Ongoing Brain Activity. Neuron, 80, 867–886.

Fox, M. D., Snyder, A. Z., Vincent, J. L., Corbetta, M., Essen, D. C. V., & Raichle, M. E. (2005). The human brain is intrinsically organized into dynamic, anticorrelated functional networks. Proceedings of the National Academy of Sciences of the United States of America, 102, 9673–9678.

Friston, K. J. (1994). Functional and effective connectivity in neuroimaging: A synthesis. Human Brain Mapping, 2, 56–78.

Friston, K. J. (1997). Transients, metastability, and neuronal dynamics. NeuroImage, 5, 164–71.

Friston, K. J. (2011). Functional and effective connectivity: a review. Brain connectivity, 1, 13–36.

Friston, K. J., Harrison, L., Daunizeau, J., Kiebel, S., Phillips, C., Trujillo-Barreto, N., Henson, R., Flandin, G., & Mattout, J. (2008). Multiple sparse priors for the M/EEG inverse problem. NeuroImage, 39, 1104–1120.

Friston, K. J., Li, B., Daunizeau, J., & Stephan, K. E. (2011). Network discovery with DCM. NeuroImage, 56, 1202–21.

Friston, K. J., Stephan, K. M., Heather, J. D., Frith, C. D., Ioannides, A. A., Liu, L. C., Rugg, M. D., Vieth, J., Keber, H., Hunter, K., & Frackowiak, R. S. (1996). A multivariate analysis of evoked responses in EEG and MEG data. NeuroImage, 3, 167–74.

Fukunaga, M., Horovitz, S. G., van Gelderen, P., de Zwart, J. a., Jansma, J. M., Ikonomidou, V. N., Chu, R., Deckers, R. H. R., Leopold, D. a., & Duyn, J. H. (2006). Large-amplitude, spatially correlated fluctuations in BOLD fMRI signals during extended rest and early sleep stages. Magnetic resonance imaging, 24, 979–92.

Granger, C. W. J. (1969). Investigating Causal Relations by Econometric Models and Cross-spectral Methods. Econometrica, 37, 424–438.

Greicius, M. (2008). Resting-state functional connectivity in neuropsychiatric disorders. Current opinion in neurology, 21, 424–30.

Greicius, M. D., Krasnow, B., Reiss, A. L., & Menon, V. (2003). Functional connectivity in the resting brain: a network analysis of the default mode hypothesis. Proceedings of the National Academy of Sciences of the United States of America, 100, 253–8.

Greicius, M. D., Srivastava, G., Reiss, A. L., & Menon, V. (2004). Default-mode network activity distinguishes Alzheimer's disease from healthy aging: Evidence from functional MRI. Proceedings of the National Academy of Sciences of the United States of America, 101, 4637–4642.

Greicius, M. D., Supekar, K., Menon, V., & Dougherty, R. F. (2009). Resting-state functional connectivity reflects structural connectivity in the default mode network. Cerebral cortex, 19, 72–8.

Grinsted, A., Moore, J. C., & Jevrejeva, S. (2004). Nonlinear Processes in Geophysics Application of the cross wavelet transform and wavelet coherence to geophysical time series. Nonlinear Processes in Geophysics, 11, 561–566.

Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C. J., Wedeen, V. J., & Sporns, O. (2008). Mapping the structural core of human cerebral cortex. PLoS biology, 6, e159.

Hämäläinen, 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. Reviews of Modern Physics, 65, 413–497.

Hämäläinen, M. S. & Ilmoniemi, R. J. (1994). Interpreting magnetic fields of the brain: minimum norm estimates. Medical & Biological Engineering & Computing, 32, 35–42.

Hammond, C., Bergman, H., & Brown, P. (2007). Pathological synchronization in Parkinson's disease: networks, models and treatments. Trends in neurosciences, 30, 357–64.

Hampson, M., Olson, C. A. I. R., Leung, H.-c., Skudlarski, P., & Gore, J. C. (2004). Changes in functional connectivity of human MT / V5 with visual motion input. Neuroreport, 15, 1315–1319.

Hampson, M., Peterson, B. S., Skudlarski, P., Gatenby, J. C., & Gore, J. C. (2002). Detection of Functional Connectivity Using Temporal Correlations in MR Images. Human Brain Mapping, 15, 247–262.

Hansen, P. C., Kringelbach, M., & Salmelin, R. (2010). MEG: An Introduction to Methods. (Oxford University Press).

Heise, V. (2013). How can magnetoencephalography and magnetic resonance imaging improve our understanding of genetic susceptibility to Alzheimer's disease? Ph.D. thesis, University of Oxford.

Hillebrand, A., Barnes, G. R., Bosboom, J. L., Berendse, H. W., & Stam, C. J. (2011). Frequency-dependent functional connectivity within resting-state networks: An atlas-based MEG beamformer solution. NeuroImage.

Hillebrand, A., Singh, K. D., Holliday, I. E., Furlong, P. L., & Barnes, G. R. (2005). A new approach to neuroimaging with magnetoencephalography. Human Brain Mapping, 25, 199–211.

Hipp, J. F., Hawellek, D. J., Corbetta, M., Siegel, M., & Engel, A. K. (2012). Large-scale cortical correlation structure of spontaneous oscillatory activity. Nature neuroscience, 15, 884–90.

Honey, C. J., Kötter, R., Breakspear, M., & Sporns, O. (2007). Network structure of cerebral cortex shapes functional connectivity on multiple time scales. Proceedings of the National Academy of Sciences of the United States of America, 104, 10240–5.

Honey, C. J., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J. P., Meuli, R., & Hagmann, P. (2009). Predicting human resting-state functional connectivity. Proceedings of the National Academy of Sciences of the United States of America, 106, 1–6.

Huang, M. X., Mosher, J. C., & Leahy, R. M. (1999). A sensor-weighted overlapping-sphere head model and exhaustive head model comparison for MEG. Physics in Medicine and Biology, 44, 423–440.

Huang, M. X., Shih, J. J., Lee, R. R., Harrington, D. L., Thoma, R. J., Weisend, M. P., Hanlon, F., Paulson, K. M., Li, T., Martin, K., Millers, G. a., & Canive, J. M. (2004). Commonalities and differences among vectorized beamformers in electromagnetic source imaging. Brain topography, 16, 139–58.

Hutchison, R. M., Womelsdorf, T., Allen, E. a., Bandettini, P. a., Calhoun, V. D., Corbetta, M., Della Penna, S., Duyn, J. H., Glover, G. H., Gonzalez-Castillo, J., Handwerker, D. a., Keilholz, S., Kiviniemi, V., Leopold, D. a., de Pasquale, F., Sporns, O., Walter, M., & Chang, C. (2013). Dynamic functional connectivity: Promise, issues, and interpretations. NeuroImage, 80, 360–78.

Hyvärinen, a. & Oja, E. (2000). Independent component analysis: algorithms and applications. Neural networks, 13, 411–30.

Johansen-Berg, H., Behrens, T. E. J., Robson, M. D., Drobnjak, I., Rushworth, M. F. S., Brady, J. M., Smith, S. M., Higham, D. J., & Matthews, P. M. (2004). Changes in connectivity profiles define functionally distinct regions in human medial frontal cortex. Proceedings of the National Academy of Sciences of the United States of America, 101, 13335–13340.

Jones, D. T., Vemuri, P., Murphy, M. C., Gunter, J. L., Senjem, M. L., Machulda, M. M., Przybelski, S. a., Gregg, B. E., Kantarci, K., Knopman, D. S., Boeve, B. F., Petersen, R. C., & Jack, C. R. (2012). Non-stationarity in the "resting brain's" modular architecture. PloS one, 7, e39731.

Kaminski, M. J. & Blinowska, K. J. (1991). A new method of the description of the information flow in the brain structures. Biological cybernetics, 65, 203–210.

Kiviniemi, V., Vire, T., Remes, J., Elseoud, A. A., Starck, T., Tervonen, O., & Nikkinen, J. (2011). A sliding time-window ICA reveals spatial variability of the default mode network in time. Brain connectivity, 1, 339–47.

Koenig, T., Prichep, L., Lehmann, D., Sosa, P. V., Braeker, E., Kleinlogel, H., Isenhart, R., & John, E. R. (2002). Millisecond by millisecond, year by year: normative EEG microstates and developmental stages. NeuroImage, 16, 41–8.

Koenig, T., Studer, D., Hubl, D., Melie, L., & Strik, W. K. (2005). Brain connectivity at different time-scales measured with EEG. Philosophical transactions of the Royal Society of London. Series B, Biological sciences, 360, 1015–23.

Larson-Prior, L. J., Oostenveld, R., Della Penna, S., Michalareas, G., Prior, F., Babajani-Feremi, a., Schoffelen, J.-M., Marzetti, L., de Pasquale, F., Di Pompeo, F., Stout, J., Woolrich, M., Luo, Q., Bucholz, R., Fries, P., Pizzella, V., Romani, G. L., Corbetta, M., & Snyder, a. Z. (2013). Adding dynamics to the Human Connectome Project with MEG. NeuroImage, 80, 190–201.

Laufs, H., Krakow, K., Sterzer, P., Eger, E., Beyerle, a., Salek-Haddadi, a., & Kleinschmidt, a. (2003). Electroencephalographic signatures of attentional and cognitive default modes in spontaneous brain activity fluctuations at rest. Proceedings of the National Academy of Sciences of the United States of America, 100, 11053–8.

Leahy, R. M., Mosher, J. C., Spencer, M. E., Huang, M. X., & Lewine, J. D. (1998). A study of dipole localization accuracy for MEG and EEG using a human skull phantom. Electroencephalography and Clinical Neurophysiology, 107, 159–173.

Lehmann, D. (1971). Multichannel topography of human alpha EEG fields. Electroencephalography and Clinical Neurophysiology, 31, 439–449.

Lehmann, D., Strik, W. K., Henggeler, B., Koenig, T., & Koukkou, M. (1998). Brain electric microstates and momentary conscious mind states as building blocks of spontaneous thinking: I. Visual imagery and abstract thoughts. International Journal of Psychophysiology, 29, 1–11.

Liang, M., Zhou, Y., Jiang, T., Liu, Z., Tian, L., Liu, H., & Hao, Y. (2006). Widespread functional disconnectivity in schizophrenia with resting-state functional magnetic resonance imaging. Neuroreport, 17, 209–13.

Liu, X. & Duyn, J. H. (2013). Time-varying functional network information extracted from brief instances of spontaneous brain activity. Proceedings of the National Academy of Sciences of the United States of America, 110, 4392–7.

Liu, Z., Fukunaga, M., De Zwart, J. A., & Duyn, J. H. (2010). Large-scale spontaneous fluctuations and correlations in brain electrical activity observed with magnetoencephalography. NeuroImage, 51, 102–111.

Logothetis, N. K., Pauls, J., Augath, M., Trinath, T., & Oeltermann, a. (2001). Neurophysiological investigation of the basis of the fMRI signal. Nature, 412, 150–7.

Lohmann, G., Erfurth, K., Müller, K., & Turner, R. (2012). Critical comments on dynamic causal modelling. NeuroImage, 59, 2322–9.

Lorente de No (1947). Action potential of the motoneurons of the hypoglossus nucleus. J Cell Comp Physiol, 29, 207–287.

Lu, H., Zou, Q., Gu, H., Raichle, M. E., Stein, E. a., & Yang, Y. (2012). Rat brains also have a default mode network. Proceedings of the National Academy of Sciences of the United States of America, 109, 3979–84.

Luckhoo, H. (2013). Investigating the role of APOE-4, a risk gene for Alzheimer's disease, on functional brain networks using magnetoencephalography. Ph.D. thesis, University of Oxford.

Luckhoo, H., Hale, J. R., Stokes, M. G., Nobre, A. C., Morris, P. G., Brookes, M. J., & Woolrich, M. W. (2012). Inferring task-related networks using independent component analysis in magnetoencephalography. NeuroImage, 62, 530–41.

Luckhoo, H. T., Brookes, M. J., & Woolrich, M. W. (2014). Multi-session statistics on beamformed MEG data. NeuroImage, In Press.

Makeig, S., Debener, S., Onton, J., & Delorme, A. (2004). Mining event-related brain dynamics. Trends in cognitive sciences, 8, 204–10.

Mantini, D., Penna, S. D., Marzetti, L., de Pasquale, F., Pizzella, V., Corbetta, M., & Romani, G. L. (2011). A Signal-Processing Pipeline for Magnetoencephalography Resting-State Networks. Brain Connectivity, 1, 49–59.

Mantini, D., Perrucci, M. G., Del Gratta, C., Romani, G. L., & Corbetta, M. (2007). Electrophysiological signatures of resting state networks in the human brain. Proceedings of the National Academy of Sciences of the United States of America, 104, 13170–5.

Marzetti, L., Della Penna, S., Snyder, a. Z., Pizzella, V., Nolte, G., de Pasquale, F., Romani, G. L., & Corbetta, M. (2013). Frequency specific interactions of MEG resting state activity within and across brain networks as revealed by the multivariate interaction measure. NeuroImage, 79, 172–183.

Matsuura, K. & Okabe, Y. (1995). Selective minimum-norm solution of the biomagnetic inverse problem. IEEE Transactions on Biomedical Engineering, 42, 608–615.

Michalareas, G., Schoffelen, J.-M., Paterson, G., & Gross, J. (2013). Investigating causality between interacting brain areas with multivariate autoregressive models of MEG sensor data. Human brain mapping, 34, 890–913.

Mohr, C., Michel, C. M., Lantz, G., Ortigue, S., Viaud-Delmon, I., & Landis, T. (2005). Brain state-dependent functional hemispheric specialization in men but not in women. Cerebral cortex, 15, 1451–8.

Mohseni, H. R., Kringelbach, M. L., Woolrich, M. W., Baker, A., Aziz, T. Z., & Probert-Smith, P. (2013). Non-Gaussian probabilistic MEG source localisation based on kernel density estimation. NeuroImage, 87, 444–464.

Musso, F., Brinkmeyer, J., Mobascher, a., Warbrick, T., & Winterer, G. (2010). Spontaneous brain activity and EEG microstates. A novel EEG/fMRI analysis approach to explore resting-state networks. NeuroImage, 52, 1149–61.

Muthukumaraswamy, S. D. (2013). High-frequency brain activity and muscle artifacts in MEG/EEG: a review and recommendations. Frontiers in human neuroscience, 7, 138.

Nikouline, V. V., Linkenkaer-Hansen, K., Huttunen, J., & Ilmoniemi, R. J. (2001). Interhemispheric phase synchrony and amplitude correlation of spontaneous beta oscillations in human subjects: a magnetoencephalographic study. Neuroreport, 12, 2487–91.

Nolte, G. (2003). The magnetic lead field theorem in the quasi-static approximation and its use for magnetoencephalography forward calculation in realistic volume conductors. Physics in medicine and biology, 48, 3637–52.

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. Clinical Neurophysiology, 115, 2292–307.

Nyberg, L., McIntosh, a. R., Cabeza, R., Nilsson, L. G., Houle, S., Habib, R., & Tulving, E. (1996). Network analysis of positron emission tomography regional cerebral blood flow data: ensemble inhibition during episodic memory retrieval. The Journal of Neuroscience, 16, 3753–9.

Okada, Y. (1981). Neurogenesis of evoked magnetic fields. In Biomagnetism, S. N. Erné, H. D. Hahlbohm, & H. Lübbig, eds., pp. 399–408. (de Gruyter, Berlin).

Ossadtchi, a., Mosher, J. C., Sutherling, W. W., Greenblatt, R. E., & Leahy, R. M. (2005). Hidden Markov modelling of spike propagation from interictal MEG data. Physics in medicine and biology, 50, 3447–69.

Palva, S. & Palva, J. M. (2012). Discovering oscillatory interaction networks with M/EEG: challenges and breakthroughs. Trends in cognitive sciences, 16, 219–30.

Plonsey, J. M. & Robert (1995). Bioelectromagnetism: Principles and Applications of Bioelectric and Biomagnetic Fields. (New York: Oxford University Press).

Prichard, D. (1994). Generating surrogate data for time series with several simultaneously measured variables. Physical Review Letters, 73.

Raichle, M. E. (2010). Two views of brain function. Trends in cognitive sciences, 14, 180–90.

Raichle, M. E., MacLeod, a. M., Snyder, a. Z., Powers, W. J., Gusnard, D. a., & Shulman, G. L. (2001). A default mode of brain function. Proceedings of the National Academy of Sciences of the United States of America, 98, 676–82.

Rashid, B., Damaraju, E., & Calhoun, V. (2013). Comparison of resting state dynamics in healthy, schizophrenia and bipolar disease. Proc. HBM, Seattle, WA. In 19th Annual Meeting of the Organization for Human Brain Mapping, Seattle.

Rezek, I. & Roberts, S. (2005). Ensemble Hidden Markov Models with Extended Observation Densities for Biosignal Analysis. Probabilistic Modeling in Bioinformatics and Medical Informatics.

Robinson, S. E. & Vrba, J. (1999). Functional neuro-imaging by synthetic aperture magnetometry (SAM). In Recent Advances in Biomagnetism, T. Yoshimoto, M. Kotani, H. Karibe, & N. Nakasato, eds. (Sendai: Tohoku Univ Press), pp. 302–305.

Sarvas, J. (1987). Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem. Physics in medicine and biology, 32, 11–22.

Schlögl, A. & Supp, G. (2006). Analyzing event-related EEG data with multivariate autoregressive parameters. Progress in brain research, 159, 135–47.

Schnitzler, A. & Gross, J. (2005). Normal and pathological oscillatory communication in the brain. Nature reviews. Neuroscience, 6, 285–96.

Schoffelen, J.-M. & Gross, J. (2009). Source connectivity analysis with MEG and EEG. Human brain mapping, 30, 1857–65.

Sekihara, K., Nagarajan, S. S., Poeppel, D., Marantz, a., & Miyashita, Y. (2001). Reconstructing spatio-temporal activities of neural sources using an MEG vector beamformer technique. IEEE transactions on bio-medical engineering, 48, 760–71.

Shirer, W. R., Ryali, S., Rykhlevskaia, E., Menon, V., & Greicius, M. D. (2012). Decoding subject-driven cognitive states with whole-brain connectivity patterns. Cerebral cortex, 22, 158–65.

Shulman, G. L., Fiez, J. a., Corbetta, M., Buckner, R. L., Miezin, F. M., Raichle, M. E., & Petersen, S. E. (1997). Common Blood Flow Changes across Visual Tasks: II. Decreases in Cerebral Cortex. Journal of cognitive neuroscience, 9, 648–63.

Siegel, M., Donner, T. H., & Engel, A. K. (2012). Spectral fingerprints of large-scale neuronal interactions. Nature reviews. Neuroscience, 13, 121–34.

Smith, S. M. (2012). The future of FMRI connectivity. NeuroImage, 62, 1257–66.

Smith, S. M., Beckmann, C. F., Andersson, J., Auerbach, E. J., Bijsterbosch, J., Douaud, G., Duff, E., Feinberg, D. a., Griffanti, L., Harms, M. P., Kelly, M., Laumann, T., Miller, K. L., Moeller, S., Petersen, S., Power, J., Salimi-Khorshidi, G., Snyder, A. Z., Vu, A. T., Woolrich, M. W., Xu, J., Yacoub, E., Uurbil, K., Van Essen, D. C., & Glasser, M. F. (2013a). Resting-state fMRI in the Human Connectome Project. NeuroImage, 80, 144–68.

Smith, S. M., Fox, P. T., Miller, K. L., Glahn, D. C., Fox, P. M., Mackay, C. E., Filippini, N., Watkins, K. E., Toro, R., Laird, A. R., & Beckmann, C. F. (2009). Correspondence of the brain's functional architecture during activation and rest. Proceedings of the National Academy of Sciences of the United States of America, 106, 13040–5.

Smith, S. M., Miller, K. L., Moeller, S., Xu, J., Auerbach, E. J., Woolrich, M. W., Beckmann, C. F., Jenkinson, M., Andersson, J., Glasser, M. F., Van Essen, D. C., Feinberg, D. a., Yacoub, E. S., & Ugurbil, K. (2012). Temporally-independent functional modes of spontaneous brain activity. Proceedings of the National Academy of Sciences of the United States of America, 109, 3131–6.

Smith, S. M., Miller, K. L., Salimi-Khorshidi, G., Webster, M., Beckmann, C. F., Nichols, T. E., Ramsey, J. D., & Woolrich, M. W. (2011). Network modelling methods for FMRI. NeuroImage, 54, 875–91.

Smith, S. M., Vidaurre, D., Beckmann, C. F., Glasser, M. F., Jenkinson, M., Miller, K. L., Nichols, T. E., Robinson, E. C., Salimi-Khorshidi, G., Woolrich, M. W., Barch, D. M., Uurbil, K., & Van Essen, D. C. (2013b). Functional connectomics from resting-state fMRI. Trends in cognitive sciences, 17, 666–82.

Sporns, O., Honey, C. J., & Kötter, R. (2007). Identification and Classification of Hubs in Brain Networks. PloS one.

Sporns, O., Tononi, G., & Edelman, G. M. (2000). Theoretical neuroanatomy: relating anatomical and functional connectivity in graphs and cortical connection matrices. Cerebral cortex, 10, 127–41.

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. Human Brain Mapping, 28, 1178–1193.

Supp, G. G., Schlögl, A., Trujillo-Barreto, N., Müller, M. M., & Gruber, T. (2007). Directed cortical information flow during human object recognition: analyzing induced EEG gamma-band responses in brain's source space. PloS one, 2, e684.

Tagliazucchi, E., Balenzuela, P., Fraiman, D., & Chialvo, D. R. (2012). Criticality in large-scale brain FMRI dynamics unveiled by a novel point process analysis. Frontiers in physiology, 3, 15.

Taulu, S., Simola, J., & Kajola, M. (2005). Applications of the Signal Space Separation Method. IEEE Transactions on Signal Processing, 53, 3359–3372.

Tomasi, D. & Volkow, N. D. (2011). Functional connectivity hubs in the human brain. NeuroImage, 57, 908–17.

Tononi, G., Sporns, O., & Edelman, G. M. (1994). A measure for brain complexity: relating functional segregation and integration in the nervous system. Proceedings of the National Academy of Sciences of the United States of America, 91, 5033–7.

Van de Ville, D., Britz, J., & Michel, C. M. (2010). EEG microstate sequences in healthy humans at rest reveal scale-free dynamics. Proceedings of the National Academy of Sciences of the United States of America, 107, 18179–84.

Van Essen, D. C. (2005). A Population-Average, Landmark- and Surface-based (PALS) atlas of human cerebral cortex. NeuroImage, 28, 635–662.

Van Veen, B. D. & Buckley, K. M. (1988). Beamforming: a versatile approach to spatial filtering. IEEE ASSP Magazine, 5, 4–24.

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 transactions on bio-medical engineering, 44, 867–80.

Vincent, J. L., Patel, G. H., Fox, M. D., Snyder, a. Z., Baker, J. T., Van Essen, D. C., Zempel, J. M., Snyder, L. H., Corbetta, M., & Raichle, M. E. (2007). Intrinsic functional architecture in the anaesthetized monkey brain. Nature, 447, 83–6.

Vrba, J. & Robinson, S. E. (2001). Signal processing in magnetoencephalography. Methods (San Diego, Calif.), 25, 249–71.

Wang, K., Liang, M., Wang, L., Tian, L., Zhang, X., Li, K., & Jiang, T. (2007). Altered functional connectivity in early Alzheimer's disease: a resting-state fMRI study. Human brain mapping, 28, 967–78.

Wipf, D. & Nagarajan, S. (2009). A unified Bayesian framework for MEG/EEG source imaging. NeuroImage, 44, 947–66.

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.

Woolrich, M. W., Beckmann, C. F., Nichols, T. E., & Smith, S. M. (2009). Statistical analysis of FMRI data. In fMRI Techniques & Protocols, M. Filippi, ed. (Humana Press), pp. 179–236.

Wu, L., Eichele, T., & Calhoun, V. D. (2010). Reactivity of hemodynamic responses and functional connectivity to different states of alpha synchrony: a concurrent EEG-fMRI study. NeuroImage, 52, 1252–60.

Yuan, H., Ding, L., Zhu, M., & Bodurka, J. (2013). Microstates, Spatial and Temporal Similarity of EEG Networks, and BOLD Resting State. In 18th Annual Meeting of the Organization for Human Brain Mapping, Beijing.

Yuan, H., Zotev, V., Phillips, R., Drevets, W. C., & Bodurka, J. (2012). Spatiotemporal dynamics of the brain at rest–exploring EEG microstates as electrophysiological signatures of BOLD resting state networks. NeuroImage, 60, 2062–72.