

Computational and Biophysical Models of the Brain

Faysal B. Ahmad

Balliol College



Analysis Group

Oxford Centre for Human Brain Activity

Nuffield Department of Clinical Neurosciences

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

October 2015

Abstract

Widely distributed brain networks display highly coherent activity at rest. In this work, we combined bottom-up and top-down approaches to investigate the dynamics and underlying mechanisms of this spontaneous activity. We developed a realistic network model to simulate resting-state data, which incorporates biophysical regional dynamics, empirical brain connectivity, time delays, and background noise. At moderately weak coupling strengths, the model produces spontaneous metastable oscillatory states and a novel form of frequency depression, resulting in transient synchronizations between brain regions at reduced collective frequencies.

We used fixed and sliding window correlation approaches on the power of band-limited MEG data, and show that brain regions exhibit significant functional connectivity (FC) in the alpha and beta frequency bands on slow ($>1\text{sec}$) time scales. We also show that temporal non-stationarity and bistability in FC occur in the same pairs of brain areas, and in the same frequency bands, as stationary measures of FC. We find that the network model reproduces the same frequency-dependency, time-scales, and non-stationary nature of FC as we found in real MEG data. Furthermore, seed-based correlations and independent component analysis also reveal a similar spatial profile of FC in empirical and simulated data, with the existence of widely distributed transient resting state networks in the same frequency bands. Finally, we used the network model simulations to evaluate a range of network estimation methods, and find that often the simplest linear measures perform best and some of the common non-linear measures can often give erroneous estimates.

Overall, our results suggest that structured interactions between brain regions in the presence of delays and noise result in spontaneous synchronizations leading to the organized power fluctuations across brain regions, and some of the simplest statistical measures provide excellent estimates of this connectivity. Our work also highlights the potential of computational models in exploring neural mechanisms.

Acknowledgements

This thesis would have hardly been possible without support and data from many people. I would particularly like to thank:

Mark Woolrich for outstanding supervision, unconditional support and motivation. During my time at Oxford, he always made time for me and guided me through the paths of a research life;

Steve Smith for advice and discussions on the network analysis methods and how this all fits into the bigger picture;

Sven Braeutigam for being an objective and valued sounding board;

Morten Kringelbach for fruitful discussions and providing the DTI derived human brain anatomical network, a key ingredient for this work;

Adam Baker for providing the code for sliding window wavelet correlation and reviewing some of the content of the dynamical functional connectivity chapter;

Henry Luckhoo and Hamid Moseni for providing the outstanding MEG data;

Joana Cabral, whose previous research and thesis served as the foundation of some of my work;

My classmates from OHBA who stuck together through the years and regularly remind me that there's more to academia than just neuroimaging;

The organizers and tutors of the FMRIB Graduate Course for preliminary training

The MSD Skills Training program for preparing me for anything and opening my eyes to how much more there is yet to learn;

The Department of Clinical Neurosciences for funding my studies at Oxford;

My sweetheart baby girl, Mini, for her endless love for me; and

Finally, I wish to acknowledge the Members of the Committee, for their time and interest, and all the future readers of this thesis.

Contents

1	Summary	13
1.1	Resting state brain activity	13
1.2	Computational and biophysical modelling	14
1.3	Biophysical network model	15
1.4	Biophysical network dynamics	17
1.5	Static functional connectivity	18
1.6	Spatial functional connectivity	20
1.7	Dynamic functional connectivity	22
1.8	Network analysis methods	24
2	Introduction	27
2.1	Computational modelling of brain activity	27
2.2	Brain activity during rest	28
2.2.1	FMRI	29
2.2.2	Electrophysiology	31
2.2.3	MEG	34
2.3	Structural networks	37
2.3.1	Structural neuroimaging	38
2.3.2	Estimation of structural networks	39
2.4	Functional networks	41
2.4.1	Functional neuroimaging	42
2.4.2	Estimation of functional networks	44
2.5	Models of brain dynamics	47
2.5.1	Micro and macro-scale models	48
2.5.2	Mean-field and neural mass models	50

2.5.3	Conductance-based model.....	51
2.5.4	FitzHugh-Nagumo model.....	54
2.5.5	Wilson-Cowan model	55
2.5.6	Spiking neuron model	57
2.5.7	Kuramoto model.....	60
3	Biophysical Network Dynamics	63
3.1	Introduction.....	63
3.2	Methods.....	64
3.2.1	Structural brain connectivity	64
3.2.2	Neurodynamical model.....	65
3.2.3	Mean-field dynamics	70
3.2.4	Delay differential equations.....	73
3.2.5	Global network synchronization.....	74
3.3	Results	75
3.3.1	Structural brain connectivity	75
3.3.2	Biophysical network dynamics	76
3.3.3	Modular connectivity and synchronization patterns.....	79
3.3.4	Frequency specificity	80
3.3.5	Transient synchronizations and power fluctuations.....	82
3.3.6	Effect of network parameters	84
3.4	Discussion.....	87
4	Static Functional Connectivity	92
4.1	Introduction.....	92
4.2	Methods.....	93
4.2.1	MEG data collection and pre-processing	93
4.2.2	Signal orthogonalization	96
4.2.3	Biophysical network dynamics	97

4.2.4	Network performance	97
4.2.5	Frequency specific functional connectivity	98
4.2.6	Haemodynamic model	100
4.3	Results	101
4.3.1	Network performance	101
4.3.2	Frequency specific functional connectivity	103
4.3.3	Standard coherence	106
4.3.4	Haemodynamic responses	107
4.4	Discussion	108
5	Spatial Functional Connectivity	114
5.1	Introduction.....	114
5.2	Methods.....	115
5.2.1	Biophysical network dynamics	115
5.2.2	MEG forward model.....	115
5.2.3	Global correlation analysis.....	117
5.2.4	Seed based correlation analysis.....	118
5.2.5	Independent component analysis	118
5.3	Results	120
5.3.1	MEG source reconstruction	120
5.3.2	Global functional connectivity	121
5.3.3	Seed based functional connectivity.....	123
5.3.4	Spatial correlation maps.....	125
5.3.5	Connectivity spectra for functional networks.....	126
5.3.6	Functional resting state networks	127
5.4	Discussion.....	128
6	Dynamic Functional Connectivity.....	134
6.1	Introduction.....	134

6.2	Methods.....	135
6.2.1	Biophysical network dynamics	135
6.2.2	Dynamic functional connectivity	135
6.2.3	Wavelet transform.....	136
6.2.4	Sliding window correlation	137
6.2.5	Signal orthogonalization	138
6.2.6	Summary statistics	138
6.2.7	Significance testing.....	139
6.2.8	Effect of time window.....	141
6.3	Results	143
6.3.1	Transient synchronizations and bimodality.....	143
6.3.2	Summary measures of bimodality	146
6.3.3	Dynamical interactions in functional networks.....	147
6.3.4	Effect of time window.....	148
6.3.5	Static vs. dynamic measures of functional connectivity.....	150
6.4	Discussion.....	151
7	Network Analysis Methods.....	156
7.1	Introduction.....	156
7.2	Methods: network dynamics	158
7.2.1	Network connectivity	158
7.2.2	DCM connection structure.....	159
7.2.3	Neural dynamics	160
7.2.4	Haemodynamic responses	162
7.2.5	Network simulations	164
7.3	Methods: network analysis	167
7.3.1	Correlation.....	167
7.3.2	Coherence	167
7.3.3	Covariance.....	168

7.3.4	Mutual information.....	169
7.3.5	Granger causality	169
7.3.6	Generalized synchronization.....	170
7.3.7	Graphical models	171
7.4	Results	172
7.4.1	FMRI network simulations	172
7.4.2	Variants of the biophysical network model	175
7.4.3	MEG network simulations	178
7.4.4	Summary across all simulations.....	180
7.5	Discussion.....	181
8	General Discussion	189
A	AAL Brain Parcellation Template.....	202
B	DTI Anatomical Connectivity.....	204
C	Network Analysis Methods	208

1 Summary

1.1 Resting state brain activity

Modern theories of brain function emphasize the importance of coordinated activity between multiple brain regions in mediating cognitive behaviour. Recent findings also demonstrate spatiotemporally organized activity while the subject is at rest, in the absence of stimulus presentation or task execution (Beckmann et al., 2005; Brookes et al., 2011b; Hipp et al., 2012; de Pasquale et al., 2010). Moreover, these activity patterns have been associated with subject's task performance (Kahn et al., 2008), and breakdown of this activity has been correlated with neurological deficits (Hawellek et al., 2011).

This spatiotemporally organized activity, termed resting state networks (RSNs), was initially observed using functional magnetic resonance imaging (fMRI), when the blood oxygen level-dependent (BOLD) signal was found to be correlated between remote brain regions (Biswal et al., 1995; Beckmann et al., 2005). Several such networks have since been identified independently using magnetoencephalography (MEG) (Brookes et al., 2011b), which is a direct measure of neural activity offering high temporal resolution as well as the frequency specificity of the underlying oscillations. Spectral analysis of these networks has revealed that resting state functional connectivity (FC) between brain regions utilizes frequency-specific bindings between spatially specific brain regions (Hipp et al., 2012; Brookes et al., 2011b). Furthermore, it has been found that this activity displays temporal non-stationarities in FC with transient periods of heightened synchronizations (de Pasquale et al., 2010). Information about the fundamental neural mechanisms underlying these dynamics can be gained from computational models of the human brain.

1.2 Computational and biophysical modelling

Computational modelling of the brain offers a powerful approach to study the dynamics of the brain in health and disease. These models range in complexity from totally abstract models, such as artificial neural networks (ANNs), to biophysically based ones, which incorporate the known physiology of the human brain (Rosenblatt, 1958; Hodgkin and Huxley, 1952; Wilson and Cowan, 1972; Breakspear et al., 2003; Brunel and Wang, 2003). Several such models have been successfully applied to study the dynamics of different brain functions including memory, learning, and decision-making (Oberauer et al., 2010; Valentin et al., 2014; Lee and Wang, 2008).

In this work we use a biophysical network modelling approach to explore the mechanisms underlying the emergence of coherent spontaneous activity in large-scale brain networks during rest. In this approach (Figure 1.1), an anatomical network is constructed from structural neuroimaging data. This network is coupled with a dynamical model of neural interactions that describes the activity of the network nodes (brain regions). The simulated network responses are used to estimate the functional relations between brain regions, which can be compared with the ones estimated using functional neuroimaging modalities (empirical data). This enables us to investigate the effects of particular assumptions (such as properties of anatomical connectivity or noise levels) on the network dynamics by examining parameter conditions that increase the model predictions. This is particularly appealing in biophysical models where the parameters have an explicit neural interpretation. In principle, if we had sufficient knowledge of the system, including molecular and synaptic details, we could predict empirical dynamics or functional characteristics in a context-dependent way. The understanding we gain about the physiology of network dynamics can be used to elucidate the mechanisms underlying brain function in health and disease. By moving closer to

the actual mechanisms of brain function, this approach should ultimately be more sensitive and more interpretable than descriptive or normative approaches.




Figure 1.1 – Exploring the brain mechanisms using biophysical network models. An anatomical network is combined with a neurodynamical model to simulate functional responses from which the simulated functional networks are derived. These networks can be compared with the functional networks estimated from empirical studies. A similarity metric can be used as a feedback to fine-tune the model parameters (dashed line). Adapted from Cabral (2012).

1.3 Biophysical network model

In the first part of this work, we developed a realistic large-scale biophysical network model of (~90) brain regions in the resting state to simulate spontaneous MEG activity, and we compared the simulated responses with those derived from empirical subjects. The model combines key aspects of brain anatomy and physiology by using a biophysical model for intrinsic node dynamics with an empirically derived brain anatomical network (see Figure 1.2).

The intrinsic dynamics of individual brain regions were modelled using the biophysically based dynamic causal model (DCM, Marreiros et al., 2008, see Figure 1.2). In this model, each source is divided into three different populations: pyramidal cells, spiny cells, and inhibitory inter-neurons. These populations are inter-connected via intrinsic connections within each region, which highlight the anatomy of regional connectivity. Different cortical regions are connected to each other via extrinsic connections among specific populations, with the specific connected populations determining the type of connection. This model captures the dynamics of each node by a conductance-based voltage model expressing the

coupling between inter-connected populations of excitatory, inhibitory and pyramidal neurons (Morris and Lécuyer, 1981).




Figure 1.2 – Biophysical network model. DTI data is acquired from healthy participants and parcellated into 90 brain regions using the AAL parcellation template. The anatomical connectome between these regions is estimated using probabilistic tractography. The estimated structural connectivity is used in combination with a neurodynamical model to simulate the activity of brain regions, where the dynamics of an individual brain region are described by a conductance based voltage model expressing the coupling between different populations of excitatory and inhibitory neurons.

The cortical regions were coupled together using a realistic large-scale brain anatomical network of 90 automated anatomical labelling (AAL) defined brain regions (Tzourio-Mazoyer et al., 2002), derived from diffusion tensor imaging (DTI) data from 21 healthy participants (see Figure 1.2). The structural biophysical attributes of this network are encapsulated by the network parameters coupling strengths, C (Figure 1.3A), which specify how strongly brain regions are connected with each other; and connection lengths, L (Figure 1.3B), from which the delays that signals encounter due to finite transmission velocities are calculated.




Figure 1.3 – Biophysical attributes of the brain anatomical network. (A) Connection strength matrix of 90 AAL brain regions. The coupling strength matrix was based on the total number of fibers connecting different pairs of brain regions. (B) Connection length matrix for the same regions. The connection lengths were based on the 3-D Euclidean distance between pairs of brain regions. A delay matrix is derived from the connection length matrix by the assumption of a finite impulse propagation velocity. The red lines on the matrices show the separation of the left and right hemispheres. Different lobes in the each hemisphere are also marked.

1.4 Biophysical network dynamics

We explore the effect of changing the global scaling, k , of node-node coupling strength (Figure 1.4). The global scaling homogenously scales all the different extrinsic connections specified in the network. We observe that the network shows a phase transition from random dynamics to ordered dynamics, distinguished by the coherence of the activity between the modelled cortical populations (Figure 1.4A). At moderately weak coupling strengths and realistic transmission delays (estimated from the connection lengths), we are able to show that the model leads to spontaneous emergence of metastable oscillatory states (Figure 1.4B), where subsets of brain regions engage in transient synchronizations to form dynamic clusters. By comparing functional connectivity characteristics derived from the simulated responses and resting-state MEG data from 10 healthy participants, we found that the model best reproduces the empirical MEG responses in a region of weak couplings and intermediate delays where high metastability occurs (Figure 1.4C). This regime (optimal working point) is consistent with previous findings

where similar dynamical clusters of nodes have been shown to reproduce functional connectivity patterns observed in fMRI (Deco et al., 2009; Cabral et al., 2011) and MEG (Cabral et al., 2014).




Figure 1.4 – Biophysical network dynamics. Network-level dynamics and model performance as a function of global coupling strengths, k . (A) The global synchrony of the network as the global coupling strength is varied (scaling all connections homogenously). (B) The global metastability (the variability over time of the synchrony) of the network as a function of global coupling. (C) The models performance at different coupling strengths as measured by the Pearson correlation coefficient between the real and simulated data characteristics. The dashed line indicates a coupling of ~ 1.8 where the best fit with the empirical data is achieved.

1.5 Static functional connectivity

Interactions between brain regions of the network lead to a novel form of frequency depression with the production of slow neural activity fluctuations, which fall in the alpha and beta frequency range (8-30 Hz) at the optimal working point (Figure 1.5A,B). Frequency-specific functional connectivity analysis (based on Hilbert envelope correlation) revealed significant power correlations between brain regions in these frequency bands (Figure 1.5C,D), consistent with what has been found in empirical MEG data (Hipp et al., 2012; Brookes et al., 2011a, Brookes et al., 2011b; Luckhoo et al., 2012). We spectrally resolved the envelope correlations to assess its temporal scale, i.e. down-sampling (Figure 1.5E,F). We

find that these correlations are prominent in the low frequency bands and on slow time-scales ($<0.1\text{Hz}$), as observed in previous fMRI studies (Biswal et al., 1995).




Figure 1.5 – Static functional connectivity. (A-B) Frequency specific functional connectivity of band limited neural amplitude signals. Amplitude envelopes were temporally down-sampled by 1s. The histograms reveal the frequency bands in which the pair-wise correlation between the regions is highest for real data (A) and simulated data (B). (C-D) Band limited power (BLP) correlations between functionally connected node-pairs in different frequency bands at four different smoothing (down sampling) windows for real data (C) and simulated data (D). (E, F) Correlation between the top functionally connected node pairs as a function of frequency band and moving average window widths (downsampling) for real data (E) and simulated data (F). All empirical scores are averaged across all subjects.

1.6 Spatial functional connectivity

We analyzed the empirical and simulated MEG data to investigate the spatial profile of functional connectivity. Since we are interested in the spatial patterns here, we explicitly modelled the effects of head volume conductance for simulated data. For this, we first estimated the field distribution across MEG sensors resulting from the source space signals using a single shell head model, and then source space activity was re-estimated using a beamformer spatial filter (Van Veen et al., 1997). We used the source space activity to derive the down-sampled amplitude envelope time courses across frequency bands. We computed the pairwise amplitude envelope correlations between all pairs of brain regions (FC matrices) in different frequency bands, and found different spatial patterns of FC in different frequency bands (Figure 1.6A,B). These patterns were strongest in the alpha and beta frequency bands. We then used a seed-based approach to generate frequency-specific spatial correlation maps (Figure 1.6C-F). These maps reveal a spatial organization of FC similar to those observed in the previous MEG studies (Brookes et al., 2011b, Hipp et al., 2012). We finally used temporal independent component analysis (ICA) to estimate the independent components (ICs) generating the source space activity, where each IC corresponds to a separate RSN. We found that the model can reproduce the spatial signature of some of the RSNs previously found using fMRI (Beckmann et al., 2005) and MEG (Brookes et al., 2011b).




Figure 1.6 – Spatial functional connectivity. (A-B) Frequency-specific functional connectivity matrices in different frequency bands for empirical (A) and simulated data (B). (C-F) Spatial correlation maps for the visual and motor cortex. Seed based correlation maps of the right motor cortex for empirical (C) and simulated (D) data. Seed based correlation maps of the left visual cortex for empirical (E) and simulated (F) data. All maps were obtained in the beta frequency band. Empirical results are averaged over all subjects.

1.7 Dynamic functional connectivity

To investigate time-varying changes in amplitude (power) correlations between neural time series, we computed sliding-window wavelet correlation between different pairs of brain regions in different frequency bands (Figure 1.7A). We find that brain regions exhibit significant temporal non-stationarity in FC primarily in the low frequency bands (8-30 Hz), which is captured by the bi-state nature of correlation distribution at these frequencies (Figure 1.7B). This bimodality is due to the presence of two underlying states (high and low FC), and this distribution is considerably lost at higher frequencies due to a single underlying state (low FC). To explore how this bimodality generalizes across multiple brain regions and subjects, we computed the standard deviation and skewness of the wavelet correlation distribution as summary measures of bimodality for the functionally connected node pairs (Figure 1.7C-F). The significance of the summary statistics was assessed by comparing the results with a null distribution derived from Monte Carlo simulations. The null distribution corresponds to the null hypotheses that the underlying functional connectivity of resting state activity is temporally stationary. We find that both measures reveal high values in the alpha and beta frequency range (due to bimodality) indicating the presence of transient synchronizations in the network, as has been found in empirical MEG data (de Pasquale et al., 2010).




Figure 1.7 – Dynamic functional connectivity. (A) Sliding-window wavelet power correlation between time courses of brain regions. (B) Standard deviation of the wavelet power correlation at different frequencies captures transient period of high correlation in the time-frequency plane. The distribution of correlation values is predominantly bimodal in presence of transient events at lower frequencies and this bimodality is lost if there is only single underlying state. (C-F) Summary bimodality measures of the wavelet correlation distribution averaged over functionally connected brain regions. (C, D) Standard deviation of the wavelet power correlation at different frequencies for real data (C) and simulated data (D). (E, F) Skewness of the wavelet power correlation at different frequencies for real data (E) and simulated data (F). All empirical scores are averaged across all subjects.

1.8 Network analysis methods

We also investigated the performance of many different methods tasked with estimating the network structure between brain regions. In this case, we are not only interested in establishing the statistical dependencies between brain regions as before (functional connectivity) but also the causal properties of these regions (effective connectivity), such as inter-region connectivity and signal directionality. Specifically, we are interested in estimating the structure of the underlying network generating observed (fMRI and MEG) responses. In a recent study, a number of network estimation methods were applied to simulated fMRI resting state data (Smith et al., 2011). However, these simulations were based on an abstract model of neural dynamics having no direct biological relevance. In this work we simulated realistic fMRI BOLD activity for a range of network scenarios using the biophysical network model (for neural activity) in combination with the nonlinear balloon model for vascular dynamics (Figure 1.8A). We compared our results with those reported in Smith et al. (2011), and found that using a more sophisticated underlying generative model does not offer any significant advantage over a linear neural model in fMRI. This may primarily be due to the haemodynamic transformation taking away much of the structural information from the original signal. Similar to results reported in Smith et al. (2011), we found that some of the simplest linear methods (based on correlation and covariance) are the most sensitive to detecting a connection (Figure 1.8B). Bayesian network methods are equally sensitive but are not very robust against different network conditions. Furthermore, these measures only provide chance-level accuracy in estimating directionality. Most of the other methods have average (or below average) performance both in detecting a connection and estimating directionality (Figure 1.8C). Patel's tau performs best for estimating directionality but is relatively insensitive to detecting a connection. All the lag-based methods

(Granger, PDC, DTF) have the worst performance. We also applied the same methods directly on MEG data simulated using the biophysical network model (i.e. without the balloon model), and found that the performance of all methods increases considerably. The biggest increase was seen for the lag-based methods that estimated directionality highly accurately, and with a very high sensitivity.




Figure 1.8 – Network analysis methods. (A) An underlying anatomical network is used in combination with the DCM model for node dynamics to simulate neural activity of the individual regions. This activity is then converted to BOLD space using the haemodynamic balloon model. (B-C) Different methods are applied to the haemodynamic responses to estimate the properties of the underlying network. We compare the sensitivity of a method in detecting a true connection (B) and the true direction (C). The performance of the network estimation method is characterized by comparing the true and estimated networks. See section 7.3 for a description of the methods.

2 Introduction

2.1 Computational modelling of brain activity

The human brain consists of billions of nerve cells, neurons, which are the main computational units of the brain. Neurons connect with each other by wire-like extensions, axons and dendrites, which enables communication between the individual cells. This also enables the neurons to organize into networks that are both anatomical and functional entities. Each neuron can connect to thousands of other cells giving rise to an inconceivable number of network arrangements that are widely distributed throughout the brain. These networks process and sort through the barrage of sensory and internal stimuli we are constantly bombarded with and mediate the appropriate response. Breakdown of these networks have been associated with some of the most devastating mental disorders, such as schizophrenia and dementia.

A central preoccupation of neuroscience is deducing how the dynamics of these networks enable information processing and the emergence of the human mind. One approach to this problem is to develop models of brain functions that are tuned to empirical responses to explore the underlying mechanisms. However, most brain processes are widely distributed across the entire brain with multiple regions taking part in the process, and the architecture of the network itself shapes its pattern of activity. Empirical research has shown this architecture is not random, and brain networks follow consistent organizational principles. Therefore, this approach requires a map of the connections between regions across the entire brain (structural networks) as well as the pattern of activations (functional networks) to capture the brain's full complexity.

In this chapter we describe the details of the computational modelling paradigm to study the network dynamics underlying resting state activity. Of all the brain states, the state during rest is of particular interest as it has been demonstrated to be composed of a very consistent set of spatiotemporal networks that emerge from the intrinsic dynamics of the brain.

In the following sections of this chapter, we first illustrate what constitutes resting state brain activity, and report recent findings of resting state brain imaging studies. We then explain the construction of structural and functional brain networks. Finally, we provide details of some of the computational models recently developed to explore the origins of resting state activity.

2.2 Brain activity during rest

When a person who is awake is not engaged in any mental or physical activity, that person is said to be in a state of rest. Sleeping is not classed under such state because during the resting phase the person is still conscious, although silent or static but still aware of its surroundings, and is ready to react to any peripheral stimulus such as responding to a doorbell or answering a telephone call.

During such a still state of the body, the brain is still engaged. Functional neuroimaging studies have revealed that the resting state activity of the brain can be decomposed into networks of brain regions. Most of the resting state networks significantly overlap with the functional architectures active during specific activities, such as linguistic or visual processing. Spatiotemporally organized brain activity at rest has been consistently found using fMRI (Biswal et al., 1995; Beckmann et al., 2005; Smith et al., 2009), electrophysiology (Laufs et al., 2003; Leopold, 2003), and more recently MEG (de Pasquale et al., 2010; Brookes et al., 2011a; Brookes et al., 2011b). This activity emerges spontaneously during the resting state and disappears as soon as the person phases into the sleeping state or starts performing any other activity (Larson-Prior et al., 2011).

Several studies have tried to link this activity with higher-level cognitive processes such as thinking about the past or future, planning long-term goals, and self-consciousness (Saxe and Kanwisher, 2003; Wagner et al., 2005; Morcom and Fletcher, 2007; Buckner et al., 2008). However, analysis of the patterns observed from the resting brains of rats (Lu et al., 2012) and monkeys (Vincent et al., 2007), strongly support a more basic source of this activity (Figure 2.1).




Figure 2.1 – Comparison of the default mode network (DMN) in rats, monkeys and humans. The different brain regions comprising this network show significant correlated activity at rest. Adapted from Cabral (2012) who has adapted this from Lu et al. (2012).

In the following subsections, we will look at some of the key findings of state of the art neuroimaging studies of resting state brain activity conducted using fMRI, electrophysiology, and MEG.

2.2.1 fMRI

Biswal et al. (1995) accidentally discovered the phenomena of resting state activity using BOLD fMRI (Ogawa et al., 1990), while they were trying to identify the brain regions activated during a bilateral finger-tapping task. They found that these regions were not only active during the task, but also after the stimulus had been removed, even after the removal of the frequencies associated with cardiac

and respiratory fluctuations. They found slow oscillations of less than 0.1Hz in the sensorimotor cortex that were correlated across the hemispheres. Using positron emission tomography (PET) to measure the oxygen demands in the brain during the resting state, Raichle et al. (2001) validated the physiological foundation of BOLD signal fluctuations, proposing that this task-independent activity be considered the default mode of brain activity.

Later, Greicius et al. (2003) identified the default mode network (DMN), by performing connectivity analysis of different brain regions active during the resting state. Brain regions in the default mode network were found to have correlated activation patterns while the brain is in the resting state, and showed significantly lower activity while the subject started performing a task. Due to such observations, default mode network is also known as a task-negative resting state network.

Further resting state studies revealed several other networks similar to those found in task related experimentations (e.g. Figure 2.2) (Beckmann et al., 2005; Damoiseaux et al., 2006; Smith et al., 2009). These networks often consist of brain regions active during specific tasks such as auditory processing, visual processing or motor function (Damoiseaux et al., 2006; De Luca et al., 2006). The distinct networks were found to be highly consistent across trials and subjects (Beckmann et al., 2005; Smith et al., 2009). Using high temporal resolution FMRI data in combination with temporal ICA, it has been further demonstrated that traditional resting state networks are underpinned by transiently synchronizing networks (temporal functional modes) with spatially overlapping subcomponents (Smith et al., 2012).




Figure 2.2 – Resting state networks detected using FMRI. Different RSNs estimated from a group of 10 healthy subjects using independent component analysis (ICA). Each network shows the spatial pattern of activation of the brain regions involved in the network. Adapted from Cabral (2012) who has adapted this from Beckmann et al. (2005).

2.2.2 Electrophysiology

Functional MRI remained popular in resting state studies mainly due to its accuracy in mapping the spatial patterns of resting state networks. This is in spite of its inherent time-delay and temporal blurring (due to measuring a vascular response to brain activity), which gives it a lower temporal resolution on the order of seconds. However, neurophysiological events happening at faster time scales and the dynamic structure of connectivity between the cortical areas cannot be assessed using this technique. By contrast, electrophysiology offers superior temporal resolution as compared to FMRI, and can detect fast dynamic interactions in neural networks. Furthermore, it can serve as a means to confirm the findings of the resting state FMRI studies. If dynamic BOLD oscillations are reflections of the underlying neural activities, then we may expect that oscillations produced from electrophysiological data will exhibit large-scale correlation patterns similar to those identified with FMRI techniques.

To investigate the relationship between the resting state electrical and BOLD recordings at the whole brain level, Laufs et al. (2003) performed simultaneous electroencephalography (EEG) and fMRI recordings on awake subjects in an eyes closed scan. EEG fluctuating signals were filtered into multiple frequency bands¹ and the band-limited power (BLP) of the signal in each frequency band (Figure 2.3) was compared with the BOLD signal time course. The researchers found a high correlation between the EEG signals in the beta frequency band and the fMRI signals originating from selected brain regions. The researchers noted that the regions showing these correlations were very similar to the ones that Greicius et al. (2003) identified as the default mode network. Furthermore, brain regions responsible for attention showed prominent negative correlations between the BOLD signals and the EEG signals in the alpha band. These results suggest that the BOLD signal may be reflecting the BLP fluctuations happening at faster time scales.




Figure 2.3 – Illustration of band-limited power signals in different frequency bands. The black signal shows the filtered brain signals in each frequency band and the red curve shows the amplitude envelope of the associated signal estimated using the Hilbert transform. Adapted from Cabral (2012).

¹ The frequency bands usually considered are delta (1-4 Hz), theta (4-8 Hz), alpha (8-12 Hz), beta (13-30Hz), and gamma (30-70 Hz).

In Mantini et al. (2007) simultaneous fMRI and EEG recordings were taken from 15 subjects followed by ICA analysis on the fMRI data, which revealed six robust RSNs. The temporal signals associated with individual RSNs were compared with band-limited power signals in different frequency bands taken from the EEG data. Each RSN was found to be characterized by a specific electrophysiological signature involving a combination of different rhythms. Similar to results from Laufs et al. (2003), they found the DMN to be correlated with the power of the beta band signal. However, the analysis was limited to the time domain, due to the low spatial resolution of EEG, and correspondence between spatial activation maps was not studied.

Shmuel and Leopold (2008) gathered simultaneous fMRI recordings with intra-cortical electrophysiological recording from just one location in the visual cortex of anesthetized monkeys. The analysis revealed strong correlations between the neural activity in the gamma frequency band and the slow fluctuations in BOLD signal.

Later, Nir et al. (2008) performed intracranial electrophysiological recordings in humans. Similar to Shmuel and Leopold (2008), they found slow spontaneous fluctuations of neural activity in the gamma band, with significant inter-hemispheric correlations. These results point to slow spontaneous modulations in firing rate and gamma local field potential (LFP) as the likely correlates of spontaneous fMRI fluctuations in the human sensory cortex.

Finally, Scheeringa et al. (2011) recently investigated the relationship between the BOLD signal and band-limited neuronal synchronizations while the subjects are performing a cognitive task. They simultaneously collected EEG and BOLD data while subjects engaged in a visual attention task known to induce sustained changes in neuronal synchronization across a wide range of frequencies. The BOLD fluctuations correlated positively with EEG gamma power and

negatively with alpha and beta power. These results suggest that the BOLD-gamma coupling observed in animals can be extrapolated to humans performing a task and that neuronal dynamics underlying high- and low-frequency synchronization contribute independently to the BOLD signal.

2.2.3 MEG

Resting state neuronal fluctuations can also be analyzed using magnetoencephalography (MEG), a technique that measures the perpendicular magnetic fields generated by synchronous oscillations of large populations of neurons. MEG has a similar temporal resolution to EEG (~ ms) but a higher spatial resolution (~ mm), since the skull distorts the magnetic fields less than electric fields. MEG has much potential in advancing our empirical understanding of these networks due to its high temporal fidelity as well as the frequency specificity of the underlying neuronal oscillations (Siegel et al., 2012).

Nikouline et al. (2001) were the first to investigate the underlying neuronal source of resting state BOLD signal fluctuations using MEG. They investigated inter-hemispheric phase synchrony and BLP correlations in the beta oscillations. They found that the beta oscillations exhibited transient synchronized activity, and the synchronization was strongest when the signals had large amplitudes. The power of the beta oscillations was found to correlate across hemispheres over long time intervals (>1s). The authors proposed this amplitude modulation of the oscillations to be the underlying source of low-frequency correlated activity observed in fMRI studies.

Subsequently, de Pasquale et al. (2010) used a seed-based and temporally nonstationary approach to characterize the MEG signatures of two common resting state networks (default mode network and dorsal attention network). They found that low frequency band (alpha, beta) amplitude envelopes of the seeds show periods of high BLP correlation followed by periods of low correlation. They

further demonstrated that inter-hemispheric coupling within the default mode network (DMN) is enhanced when this non-stationarity is taken into account. By identifying time points corresponding to high within-network connectivity, the same group also showed that networks transiently engage with other networks during periods of high internal correlation, with the DMN acting as a hub of cross network interaction (de Pasquale et al., 2012). This observation suggests that the brain routinely explores several networks during rest, which in turns results in a multi-stable regime, where multiple networks (states) that are transiently stable coexist. However, this approach relied upon the a priori specification of multiple seeds taken from FMRI RSNs.

In order to identify resting state networks in MEG without making a priori assumptions using FMRI-derived RSNs, Brookes et al. (2011b) performed an ICA based analysis on resting state MEG data. Temporal ICA was applied to source-space band-limited amplitude envelopes to extract the independent components (i.e. RSNs). Their analysis revealed networks with significant similarity to those observed using FMRI, and most of these networks were found to be in beta frequency band (Figure 2.4). This was the first time RSNs were independently identified in MEG. This strongly supports the neural basis of FMRI resting state networks, and hints at the potential of MEG to further investigate the underlying spatiotemporal dynamics.




Figure 2.4 – Resting state networks detected using MEG and FMRI. (A) Default mode network; (B) left lateral frontoparietal network; (C) right lateral frontoparietal network; (D) sensorimotor network; (E) medial parietal regions; (F) visual network; (G) frontal lobes including anterior cingulate cortex; (H) cerebellum. The networks are estimated using ICA on resting state data. In the case of MEG, most networks were estimated in the beta frequency band except the default mode network (alpha). Adapted from Cabral (2012) who has adapted this from Brookes et al. (2011b).

More recently, Hipp et al. (2012) investigated the frequency-specific correlations in spontaneous brain activity. Similar to Brookes et al. (2011b), the MEG data were filtered into different frequency bands and spatially specific correlations between band-limited signals were calculated. They found that the correlations between brain regions depended on the chosen frequency band, and the strongest correlations were observed between homologous regions in the alpha and beta frequency bands. Their work also revealed the presence of several hubs in different brain regions that also displayed strongest correlations in specific frequency bands. The work further suggests that power correlations between spatially distributed brain regions exhibit frequency specificity.

In summary, a consistent view has emerged that the slow amplitude modulations of oscillatory signals are, at least in some part, representing an underlying source of the observed FMRI signals. These modulations manifest as transient increases in the power in different frequency bands, and brain regions

engage with each other in a frequency-specific way. This suggests a complex interplay between different frequencies, each reflecting different aspects of neural processing (Womelsdorf et al., 2007; Palva and Palva, 2012). The EEG and MEG studies have primarily linked the alpha and beta oscillations with observed fMRI resting state fluctuations (Nikouline et al., 2001; Liu et al., 2010; Brookes et al., 2011b). However, intra-cortical recordings also suggest a strong role of gamma oscillations in spontaneous brain activity (Leopold et al., 2003; Nir et al., 2008; Shmuel and Leopold, 2008). It has been found that a negative correlation exists between the beta and gamma oscillations (Scheeringa et al., 2011), and it might be that oscillations in both frequencies are associated with the BOLD signal.

2.3 Structural networks

Computational models of brain activity can be used to investigate the conditions under which such types of dynamics can emerge. Recently, there has been increasing interest in models that make use of the information about the anatomical connectivity of the brain. The brain's anatomical connectivity, or connectome, is comprehensive map of all neural connections in the brain, and can be thought of as the wiring diagram of the brain that entails its computational architecture (Hagmann, 2005; Sporns et al., 2005). This map may range in scale from a comprehensive mapping of all individual synapses and neurons to macro/large-scale description of structural connectivity between all cortical and subcortical regions. Since we are only interested in large-scale patterns of functional connectivity, which can also be estimated using current functional neuroimaging techniques, it is sufficient to have a macro-scale connectome. The large-scale brain networks are formed by white matter tracts that allow the functional integration of information from widely distributed brain regions (Tononi et al., 1992, 1994; Jirsa and Kelso, 2000). Due the large distances between these regions and a finite propagation velocity of the impulses, signals take

significant time to travel resulting in time-delayed interactions. The anatomical connectivity together with the propagation delays defines a space-time structure for the evolution of brain dynamics (Ghosh et al., 2008; Deco et al., 2011).

In the next subsections we explain how this space-time structure can be estimated from structural neuroimaging data. In particular, we will explain the imaging of white matter tracts, specification of individual brain regions (parcellation), and the estimation of connectivity between these regions.

2.3.1 Structural neuroimaging

Axonal tracing. Historically, the mapping of white matter connections has been performed using anterograde or retrograde axonal tracing techniques. These techniques rely on the transport of molecular or genetic markers along the axon. By visualizing these markers, detailed descriptions of neural projections can be constructed. Kötter and Wanke (2005) derived a realistic Macaque brain network of 38 cortical regions in a single hemisphere (Figure 2.5) using white matter tracing data from the CoCoMac connectivity database (Stephan et al., 2001). This map has been widely used in computational studies to model large-scale brain activity at rest (Honey et al., 2007; Ghosh et al., 2008; Deco et al., 2009). However, these techniques are very time consuming, applied to only certain regions of the brain, and limited in application to humans.




Figure 2.5 – Structural connectivity of the Macaque brain using the CoCoMac database. The coupling matrix captures the anatomical connectivity defining the connections between the 38-parcellated brain regions. The strength of the connections is captured with three different levels (1-low, 2-medium, 3-high). Adapted from Cabral (2012).

Diffusion MRI. Technological advancements in acquisition and analysis of diffusion MRI data have led to a revolution in the estimation of the macro-scale connectomes. This imaging technique is based on the imaging of diffusion of molecules (mainly water) in biological tissues, which depends on the tissue microstructure. Diffusion tensor imaging (DTI) is a diffusion MRI method that is sensitive to the rate and direction of diffusion of water molecules in white matter fibers (Wedeen et al., 1995). The water molecules can rapidly diffuse along the length of the fiber (anisotropic diffusion), whereas their diffusion perpendicular to the fiber is limited due to the myelin sheath. The detection of white matter fibers is based on this anisotropic diffusion of water molecules. An important limitation of DTI is its inability to resolve crossing fibers that is overcome by diffusion spectrum imaging (DSI) that explores the anisotropy in multiple directions (Wedeen et al., 2005). The directional information from diffusion data can be used to infer white matter tracts (i.e. tractography) by computational construction and analysis of curves of maximum diffusion coherence. This allows automatic and fast estimation of structural connections among brain regions. The first large-scale human brain anatomical network of 66 regions estimated using diffusion MRI tractography was published by Hagmann et al. (2008), which has later been used in a several resting state modelling studies (Honey et al., 2009; Deco and Jirsa, 2012; Cabral et al., 2014).

2.3.2 Estimation of structural networks

Brain parcellation. To estimate the structural connectivity at the macro-scale using structural neuroimaging data (e.g. diffusion MRI), it is first necessary to define brain regions between which this connectivity is to be estimated (see Figure 2.6). There are different strategies for dividing the brain into different regions based on genetics, synaptic profiles, short and large-scale anatomical connectivity, and functional specialization. Several schemes to parcellate the brain into distinct

regions have been proposed, ranging from canonical classification into hemispheres and lobes to classification into several thousand regions. In an attempt to standardize the anatomical labelling of brain regions, Tzourio-Mazoyer et al. (2002) proposed a brain parcellation scheme called automated anatomical labelling (AAL, See Appendix A - *AAL Brain Parcellation Template* for ROI names). In this approach, a structural MRI scan from a single healthy male was first segmented into grey matter, glial matter, white matter, cerebrospinal fluid, fat, muscle, skull, and skin. The grey matter was then divided into 90 cortical and subcortical regions (45 in each hemisphere) based on boundaries of different sulci and gyri. Since its creation, it has been used in a number of empirical and computational modelling studies of resting state functional connectivity (Liu et al., 2008; Lynall et al., 2010; Deco et al., 2009; Cabral et al., 2014).

Network connectivity. Once the brain regions have been defined, the connectivity between these regions is captured using two different matrices (Figure 2.6). The coupling strength matrix defines how strongly brain regions are coupled and is based on the count of fibers connecting pairs of brain regions. Each element of the matrix specifies the weight of connection between the associated pair of brain regions. A distance matrix is calculated based on the length of fibers connecting pairs of brain regions. It is also common to calculate this matrix from the Euclidean distances between the centers of gravity of specified brain regions. From the distance matrix, a delay matrix is derived by the assumption of a homogenous signal propagation velocity.




Figure 2.6 – Estimation of structural connectivity using diffusion MRI. Diffusion MRI data is collected for the whole brain and the directional information is used to establish the tracts of the brain. A brain parcellation scheme is used to define regions of interest among which the structural connectivity is to be calculated. The anatomical network is derived using the diffusion MRI data and the AAL parcellation scheme. The strength of connectivity is based on number of fibers linking pair of brain regions, and the time delays are based on estimates of spatial distances between regions and an assumption of a finite signal propagation velocity. Adapted from Cabral (2012).

2.4 Functional networks

Computational models that combine anatomical networks with dynamic models of neural interactions can be used to simulate the spatiotemporal activities of the network nodes. To examine how realistic the simulated network activity is we can compare this with data acquired using functional neuroimaging modalities. We also want to be able to go beyond this simple comparison, to understand the mechanisms underlying empirical responses. A common approach that addresses both issues is to characterize the network activity using measures of statistical dependency, otherwise known as functional connectivity. These methods range

from simple linear measures, such as correlation, covariance, coherence, and ICA, to non-linear measures, including phase synchronization, mutual information, and Granger causality. It has been observed that the functional connectivity broadly reflects the underlying structural connectivity (Damoiseaux and Greicius, 2009; van den Heuvel and Hulshoff Pol, 2010). However, structural connectivity alone is not a sufficient or complete description of connectivity, and robust functional connectivity has been observed in the absence of an anatomical connection (Koch et al., 2002). Analysis of functional connectivity patterns of the simulated data not only enable a comparison with the (patterns of) empirical data, but can also lead to additional insights into the mechanisms and functional organization of the brain.

In the next subsections, we compare the relevant functional neuroimaging techniques for measuring empirical brain activity, and then describe the most common methods for estimating functional connectivity.

2.4.1 Functional neuroimaging

Electrophysiology. The imaging of localized functional activations in brain regions requires techniques that are sensitive to the level of neural activity. For more than a century, electrophysiological measurements were the primary method of detecting brain activity. These techniques measure the changes in voltage levels induced by the ionic currents at synapses which are caused by the simultaneous firing of large number of neurons (Niedermeyer and Lopes da Silva, 1982). Electrophysiological measurements are made by placing electrodes at the recording site, and can be made at the level of individual synapses, cells or at the level of brain regions depending on the type, sensitivity, and recording site of the electrodes. For synaptic and single cell measurements, highly sensitive electrodes are normally implanted next to the recording site. The measurements at the regional level can be made by placing electrodes directly on the cortical surface of exposed brains, a technique called Electrocorticography (ECoG), or by implanting

intra-cortical electrodes to detect activity from deeper brain regions. These techniques allow measurements with high spatial and temporal resolution but are limited in practice due to their high invasiveness, and are generally performed only in certain clinical conditions such as epilepsy. Furthermore, the invasiveness of these techniques makes it very difficult to perform simultaneous recordings from a large population of neurons distributed throughout the brain.

Electroencephalography (EEG). Another approach is to take the measurements by placing electrodes at the scalp level, using EEG. Further to its non-invasiveness, this technique is relatively cheaper than other functional imaging modalities such as MEG and MRI, and is highly portable allowing for recordings in freely behaving conditions. However, there are some important limitations of this technique. The electric fields emanating from the brain are highly distorted by the scalp and the skull. Electrical activity from deeper brain regions, such as the thalamus, is very difficult to detect due to interference and attenuation of electric fields. Furthermore, mapping the observed signals to brain regions generating this activity (source-reconstruction) requires an ill-solved inverse problem with many different solutions, therefore, significantly limiting its spatial resolution (up to a few centimetres). However, the extraordinary temporal resolution of this technique (around 1ms) still makes it a valuable tool for the temporal and spectral analysis of brain activity.

Magnetoencephalography (MEG). Based on the Maxwell's equations, any electric current has an orthogonally oriented magnetic field. The electric currents in the brain induce magnetic fields that can be detected using magnetometers, a technique called magnetoencephalography (MEG). As magnetic fields are less distorted by the scalp and skull, MEG provides a better spatial resolution than EEG (a few millimetres). The temporal resolution is similar to EEG allowing temporal and spectral analysis of high-frequency brain oscillations. However, the magnetic fields generated by the brain are considerably weaker than the

background magnetic noise. Therefore, current MEG sensors require construction of magnetically shielded rooms and hundreds of highly sensitive detectors (requiring cryogenic freezing) to detect brain signals with high precision. Furthermore, as electric currents only produce magnetic fields perpendicular to the direction of current flow, MEG can only detect the tangential components of cortical activity, whereas EEG measures both tangential and radial components.

Functional MRI. Neuronal activity in the brain causes changes in the blood flow that can be imaged using functional MRI (fMRI). The most common form of fMRI uses the blood-oxygen-level dependent (BOLD) contrast, discovered by Seiji Ogawa (Ogawa et al., 1990). This technique images the changes in blood flow (hemodynamic response), which are related to the energy use of neurons. Specifically, it relies on the changes in magnetization between oxygen-poor and oxygen-rich blood as its basic contrast. Brain regions that are especially active are associated with high levels of oxygenated haemoglobin (Hb) (Logothetis et al., 2001). With this increase, the fraction of de-oxygenated haemoglobin (dHb) decreases resulting in a higher MR signal from that region. This method does not require any type of external radiation or injection of chemicals. BOLD-fMRI offers significantly higher spatial resolution than either EEG or MEG (on the order of a few millimetres). However, fMRI is an indirect measure of neural activity and suffers from poor temporal resolution due to the delayed neurovascular coupling. This fundamentally limits the amount of information that can be extracted using fMRI studies (Heeger and Ress, 2002). Despite this, it is the most widely used functional neuroimaging technique (Biswal et al., 1995; Friston et al., 1995; Buckner et al., 1996; Bullmore et al., 1996).

2.4.2 Estimation of functional networks

Pearson correlation. There are several ways to estimate functional interactions between brain regions (i.e. functional connectivity) from their corresponding time

courses. Computing the linear (Pearson) correlation between pair of signals is the most common and widely used method of estimating the strength of linear interactions (Bandettini et al., 1993; Biswal et al., 1995). If the regions are activated simultaneously, the correlation is high implying a functional connection. In EEG and MEG, it is common to compute the correlations between the band-limited power (BLP) of the source signals.

Seed based correlation. To investigate resting-state functional connectivity using correlation, one approach is to compute seed-based correlations by defining a seed region and computing its correlation with the rest of the brain. The seed region is chosen based on the parcellation scheme or the coordinates of the main activation sites. The pairwise correlation between this seed region and all remaining regions is calculated to give a correlation map for the individual seed. This process is repeated for all the seeds belonging to a particular RSN, and the correlation maps of each seed are overlapped to give seed co-activations maps (Fox et al., 2005). Brain regions that have significant correlations with all the seeds from an RSN are recognized as part of that RSN. The seed-based correlation analysis can be used to construct spatial maps of functional brain activation.

Full correlation matrix. Another approach is to compute the pairwise correlations between all the brain regions resulting in a correlation matrix. This matrix typically includes all significant pairwise correlations between all brain regions. Computation of the BOLD correlation matrices at the level of individual voxels can be computationally intensive, and it is common to average the BOLD signals across voxels in the same brain area (as defined by the parcellation template) to analyze the correlation matrix at the regional level (Honey et al., 2009). Computing the correlation matrix can also be thought of as computing the seed-based correlations of all the seed regions with the rest of the brain, for e.g., the first row of the matrix corresponds to the correlations of the first seed with all the brain regions.

Independent component analysis (ICA). Independent Component Analysis is another method to estimate RSN spatial maps from coordinated signal fluctuations (Beckmann et al., 2005; Damoiseaux et al., 2006; De Luca et al., 2006; Mantini et al., 2007; Beckmann, 2012; Brookes et al., 2011b). This method assumes that an observed signal is a mixture of independent spatial or temporal signals. It is a special case of blind source separation and tries to separate the observed multivariate signal into spatially or temporally independent additive subcomponent signals. The observed signal in each voxel or region $x_n(t)$ can be written as a sum of different independent components IC_k with mixing coefficients $\alpha_{n,k}$

$$x_n(t) = \alpha_{n,1}IC_1(t) + \alpha_{n,2}IC_2(t) + \dots + \alpha_{n,k}IC_k(t). \quad (2.1)$$

In theory, the number of independent components (IC) generating the data is equal to the number of sources (voxels or brain regions). However, it is common to reduce the dimensionality of the data prior to performing the ICA. This is usually done using principal component analysis (PCA), a statistical procedure that performs a linear mapping of the original data into a low dimensional space such that the variance of the data is maximized. Each IC is assumed to be a separate RSN and its mixing coefficients indicate the contribution of each brain region to the network. This way a characteristic spatial map can be constructed for each IC, which corresponds to a functional network, showing the relative contribution of the different brain regions. The extraction of different functional networks from resting state data using ICA is particularly appealing since it does not rely on a priori information regarding the spatial characteristic of RSNs.

Nonlinear measures. Both correlation and ICA are linear measures of functional connectivity. Several other non-linear measures for estimating functional interactions have been applied to resting state fMRI and MEG data. The most common of these are mutual information, Granger causality, and phase

synchronization. Mutual information (MI) (Shannon, 1948) aims to quantify the amount of shared information between two signals. Granger causality (Granger, 1969) aims to predict whether patterns in one signal are approximately repeated in another signal after a certain time lag. Phase synchronization occurs when coupled oscillatory signals keep their phase difference bounded while their amplitudes remain uncorrelated.

Most of these methods do not offer any significant advantage over linear measures in fMRI since the non-linear effects are very small (Hlinka et al., 2011; Lynall et al., 2010). Much of the observed linear dependency between brain regions may be due to the low temporal scale of the BOLD signal (a few seconds), which does not capture fast causal interactions among brain regions. This is also the reason why methods based on coherence are not typically applied to the fMRI time courses, since phase-lagged interactions can only be detected at much higher resolutions.

2.5 Models of brain dynamics

The brain can be considered a complex dynamical system when it comes to understanding the fundamental processes allowing the exploratory activities taking place during the resting state (Siegelmann, 2010; Olbrich et al., 2011). The complex space-time structure of the brain's anatomical network, together with the numerous electrical and chemical processes taking place in parallel, forms a framework with a capability to hold an endless number of mind states over which cognition develops (Tononi et al., 1994; Kelso, 2012). By formulating the activity of the brain's network using differential equations, which can then be studied in computer models, we can investigate the conditions under which such types of dynamics can emerge.

2.5.1 Micro and macro-scale models

An important consideration in computational modelling of the brain is the required level of detail (Figure 2.7). This can range from the description of activity at the synaptic and cellular levels to large-scale models of circuits and regions. Although it is becoming feasible to simulate the activity at the level of individual neurons (Markram, 2006; Izhikevich and Edelman, 2008), these models are too computationally expensive to explore the dynamics at the large-scale because this would require simulating the activity of a rather large number of neurons. There is also the concern of the large number of parameters in the model that would need to be tuned to give realistic dynamics. Furthermore, non-invasive functional data is not sufficiently resolved to consider the activity of individual neurons.




Figure 2.7 – Micro or macro-scale computational models of the brain. Adapted from Markram (2012).

On the other hand, we can consider that the individual brain regions are functional entities themselves and then model the large-scale interactions between these regions. It is common to divide the individual regions into multiple populations, corresponding to sub-populations of similar neurons within a brain area. This is motivated by neuroanatomy and neurophysiology observations that neurons in the same population tend to have similar synaptic, structural, and physiological properties, and often show strong dynamical correlations. Typically,

the models consider populations of excitatory and inhibitory neurons separately. This is based on previous findings that most self-sustained oscillations in the brain emerge from the interaction of large numbers of excitatory and inhibitory neurons (Brunel and Wang, 2003). The different populations in this case are coupled with each other with intrinsic connections (in each region), and the cortical regions are coupled to each other via the excitatory populations in each region (i.e. excitatory - excitatory connections), which mimic the glutamatergic synaptic currents (Breakspear et al., 2003) (see Figure 2.8). Such models have proven to be very beneficial as they enable simulations of the temporal activities of a large group of neurons without the need for excessively expensive computation. Recent studies have also demonstrated that the key dynamics of resting state activity can be reproduced using such population models (Deco et al., 2011).

2.5.2 Mean-field and neural mass models

The activity of the populations can be described using a set of differential equations that dictate the time evolution of their state variables. These descriptions can be based on mean-field or neural mass approximations. Mean field models (MFM) capture the behaviour of the population using probability distributions over their state variables (Deco et al., 2008). Most commonly a Gaussian distribution is assumed and the activity is summarized by the mean and variance of the population states. This is achieved by applying the Fokker-Planck equation to the underlying dynamical equations (Deco et al., 2008). These mean fielded models are only useful when the local dynamics are nonlinear (Marerios et al., 2010). Neural mass models (NMM) are a special case of mean field models where the activity of the population is summarized by only the mean of the state variables.

When populations of neurons are put together in a neuroanatomical network, interactions depend on the strength of the connections. Due to the large-

scales distances considered, time delays in the signal conduction become significant. In addition, there is also a presence of background noise in the brain, which introduces further complexity into the system. These factors can be included as additional parameters in the model to increase its biological plausibility. A number of studies have developed different population models of resting state activity to investigate the role of these factors in shaping network dynamics (Ghosh et al., 2008; Deco et al., 2009; Honey et al., 2009; Deco and Jirsa, 2012).

In the following sections of this chapter, we describe some of the recently developed resting state models and comment on their main findings, including the role of key parameters, such as connectivity, couplings, delays and noise.

2.5.3 Conductance-based model

Honey et al. (2007) were the first to investigate the dynamics of neural systems coupled through a realistic anatomical network. In this work, they used the anatomical connectivity of the Macaque brain (Kötter, 2004). Later they investigated the human resting state functional connectivity patterns by incorporating a realistic human brain anatomical network (Honey et al., 2009; Hagmann et al., 2008).

The dynamics of the individual regions were based on a biophysical neural mass model introduced earlier by Breakspear et al. (2003), consisting of populations of inhibitory and excitatory neurons. The dynamics of the populations conform to a conductance-based model (Morris and Lecar, 1981). This sort of model is particularly attractive because this is directly related to specific synaptic process (the source of MEG signals) while having an intermediate complexity. The main differential equations describing the state dynamics of excitatory (V) and inhibitory (Z) populations at time t are given by (where the time index has been dropped for clarity)

$$\dot{V} = -(g_{ca} + a_{ee}Q_V)m_{ca}(V - V_{ca}) - (g_{Na}m_{Na} + a_{ee}Q_V)(V - V_{Na}) - g_KW(V - V_K) - g_L(V - V_L) + a_{ie}ZQ_Z + a_{ne}I_\delta, \quad (2.2)$$

$$\dot{Z} = b(a_{ei}VQ_V + a_{ni}I_\delta), \quad (2.3)$$

where V , membrane potential of excitatory (pyramidal) cells, is the main state variable. The change in the membrane potential depends on the total current across the cell membrane. The current in the cell is a function of voltage-gated ion channels, and depends on the conductance and reverse potential of calcium (g_{ca}, V_{ca}), sodium (g_{Na}, V_{Na}), and potassium (g_K, V_K) ions, as well as the proportion of ion channels that are open at any time (m_{ca}, m_{Na}, m_K). The firing rate of the excitatory population in another region (Q_V) alters the conductance of calcium and sodium channels; weighted by the coupling strengths (a_{ee}). The total current also depends on the leaky conductance (g_L, V_L), which accounts for the leakage in the ion current. There is an additional current in each population that corresponds to the connections between the populations within the same region. This current depends on the firing rate of both populations (Q_V, Q_Z) and the intrinsic coupling weights (a_{ie}, a_{ei}). Finally, there is non-specific current in each population (I_δ) with specific weights a_{ne}, a_{ni} .

The opening of each voltage-gated channel depends on a threshold potential (T_{ion}), which needs to be overcome for the current to flow. To model the fraction of open channels for a large number of ion channels, a Gaussian distribution is assumed, with variance δ_{ion} . This is given by

$$m_{ion} = 0.5 \left[1 + \tanh \left(\frac{V - T_{ion}}{\delta_{ion}} \right) \right]. \quad (2.4)$$

However, the potassium channels relax from one state to another at an exponential rate and this affects the proportion of open channels. Therefore, this is defined differently as

$$\frac{dW}{dt} = \frac{\phi(m_K - W)}{\tau_K}, \quad (2.5)$$

where τ_K is the relaxation time constant of the potassium channels and ϕ is a temperate-dependent scaling factor.

The average firing rates of the excitatory (Q_V) and inhibitory (Q_Z) populations are calculated from the mean membrane potential and the threshold potential. An action potential is fired whenever the membrane potential exceeds a certain threshold. For a population of cells, a Gaussian distribution is assumed and the firing rates are calculated as

$$Q_V = 0.5 \times Q_{V \max} \left(1 + \tanh \left(\frac{V - V_T}{\delta_V} \right) \right), \quad (2.6)$$

$$Q_Z = 0.5 \times Q_{Z \max} \left(1 + \tanh \left(\frac{Z - Z_T}{\delta_Z} \right) \right), \quad (2.7)$$

where V and Z denote the voltages for excitatory and inhibitory channels respectively, and $Q_{\max}$ are the maximum firing rates for these channels. The firing rates serve as the output of the populations and facilitate synaptic interactions between them.

The simulated responses were compared against empirical resting functional connectivity patterns derived from fMRI data from healthy subjects. In studies investigating resting state functional connectivity using the macaque (Honey et al., 2007) or the human brain (Honey et al., 2009), the inter-regional coupling strengths were set to be moderately low, thus allowing the transient switching between the synchronized and unsynchronized states. The dynamics emerging from this model reflected spontaneously forming self-organizing patterns that resemble the resting-state functional connectivity patterns observed in empirical BOLD studies. In this model, the global dynamics emerging at the large scale are purely a function of the structural connectivity and non-linearities in the model, as other key parameters, such as noise and delays, are excluded.

2.5.4 FitzHugh-Nagumo model

Ghosh et al. (2008) explored the interactions of cortical regions in a dynamical framework by incorporating realistic signal transmission delays and background noise. In this work, the cortical regions were coupled according to the anatomical connectivity of the Macaque brain, and time-delays, proportional to the distance between regions, were based on the human brain AAL parcellation template.

The dynamics of the individual regions were modelled using the FitzHugh-Nagumo model (Fitzhugh, 1961; Nagumo et al., 1962). By making extensive simplifications, the FitzHugh-Nagumo model was derived from the Hodgkin-Huxley model (Hodgkin and Huxley, 1952), which describes the activation and deactivation dynamics of individual neurons. In this model, the state of a single cortical region is defined using two state variables, which correspond to the membrane potential ($u_n(t)$) and recovery potential ($v_n(t)$). The dynamics of the state variables at time t are given by the following differential equations

$$\dot{u}_n(t) = g(u_n, v_n, t) - c \sum_{p=1}^N C_{np} u_p(t - \Delta t_{np}) + n_u(t), \quad (2.8)$$

$$\dot{v}_n(t) = h(u_p, v_n, t) - n_v(t), \quad (2.9)$$

where N is the number of regions; C_{np} and Δt_{np} specify, respectively, the coupling and delays between regions n and p ; c is the global scaling of the coupling matrix, which homogenously scales all connection strengths without altering the connection matrix nor affecting the associated time delays; $n_u(t)$ and $n_v(t)$ are additive Gaussian noise in each state.

The functions g and h are defined as

$$g(u_n, v_n, t) = \tau(v_n(t) + \gamma u_n(t) - (u_n(t))^3/3), \quad (2.10)$$

$$h(u_n, v_n, t) = -(1/\tau)[u_n(t) - \alpha + \beta v_n(t)], \quad (2.11)$$

where $\alpha, \beta, \gamma, \tau$ are experimentally determined parameters.

In their investigations, damped oscillations emerged out of isolated nodes in the absence of inter-regional connectivity. When the nodes were coupled together using brain inspired connectivity and finite propagation delays, the system displayed the emergence of fluctuating oscillatory dynamics (~10 Hz) at moderate coupling strengths, a restricted interval of delays (5-10ms), and low noise levels. The global dynamics of the entire network were found to be strongly structured by the underlying anatomical connectivity, giving rise to band-limited power fluctuations between specific regions. The simulated time courses were used to compute the BOLD signal, which displayed slow oscillations (<1 Hz) similar to those previously observed in fMRI studies. The BOLD signal between regions of the DMN was also found to have correlation patterns as observed in empirical studies.

2.5.5 Wilson-Cowan model

Deco et al. (2009) developed another important model to investigate the large-scale resting state dynamics. They used a similar space-time structure to Ghosh et al. (2008), with the macaque anatomical connectivity and propagation delays derived from the AAL human brain template.

The dynamics of the individual regions were based on the Wilson-Cowan units (WCU) (Wilson and Cowan, 1972, 1973). After studying the collective behaviour of large number of coupled inhibitory and excitatory neurons, Wilson and Cowan found that the key dynamics of coupled networks can be reproduced by simple models that combine populations of inhibitory and excitatory neurons (Figure 2.8). The dynamical equations governing the activity of the excitatory ($E(t)$) and inhibitory ($I(t)$) populations of the n -th cortical region at time t are given by

$$\tau \dot{E}_n(t) = -E_n(t) + \phi(I_b + \sum_{p=1}^N \alpha C_{np} E_p(t - D_{np}) - I_n(t)), \quad (2.12)$$

$$\tau' \dot{I}_n(t) = -I_n(t) + \phi(w_I E_n(t)), \quad (2.13)$$

where τ and τ' are the membrane time constants set to 1 and 0.2, respectively. D_{np} specifies the time delays, C_{np} is the inter-regional coupling, and α is global scaling of the anatomical connectivity. $C_{nn}(=w_+/\alpha)$ specify the coupling strength for the recurrent feedback to the same excitatory pool, and w_I specifies the coupling strength between the excitatory-to-inhibitory pools. I_b is a constant background current, and an additive Gaussian noise was also added to each population.

ϕ is the nonlinear sigmoid function that transforms the current into discharge rates

$$\phi(x(t)) = c / (1 - \exp(-a(x(t) - b))) , \quad (2.14)$$

where a, b, c are constants with values of 0.1, 40, and 100, respectively.




Figure 2.8 – Schematic of the Wilson-Cowan model. Each region consists of populations of excitatory and inhibitory neurons that are linked to each other via intrinsic connections within a single region. At the network level, multiple Wilson-Cowan units are linked together by (excitatory) connections between excitatory populations of each region. Adapted from Deco et al. (2012).

In this work, the parameters I_b and w_+ were set such that when the regions are linked together spontaneous gamma oscillations (40Hz) emerge. This was based on the observations that individual cortical regions show gamma frequency oscillations when in isolation. In this regime of self-sustained gamma frequency oscillations, the system was found to organize into two competing functional

networks as a result of noise driven fluctuations between different networks. After computing the BOLD signal from the simulated activity, it was found that synchronization level fluctuated in an anti-correlated manner in the low frequency range (< 0.1 Hz), in agreement with previous reports of anti-correlated resting state networks (Fox et al., 2005).

2.5.6 Spiking neuron model

Deco and Jirsa (2012) recently developed a model for investigating resting state dynamics with an increased degree of realism. In their model, the brain regions were coupled according to the human brain anatomical network (Hagmann et al., 2008) used earlier by Honey et al. (2009). The time delays between the regions were excluded.

The dynamics of individual cortical regions were modelled using interconnected populations of spiking excitatory and inhibitory neurons (Figure 2.9). This type of biophysical network has an increased tendency to settle into stable states, characterized by a stable firing pattern. Perturbations due to the internal or external stimulus or noise can temporarily destabilize this state, causing shifts between multiple stable states.

The dynamics of the individual neurons were based on the leaky integrate-and-fire (IF) spiking neuron model, consisting of excitatory (NMDA and AMPA) and inhibitory (GABA) ion channel receptors. The membrane potential ($V(t)$) of an IF neuron at time t is given by

$$C_m \dot{V}(t) = -g_m(V(t) - V_L) - I_{syn}(t), \quad (2.15)$$

where C_m is the membrane capacitance, g_m is the membrane leak conductance, V_L is the resting voltage, and $I_{syn}(t)$ is the incoming synaptic current. The neuron generates a spike when the voltage across the membrane reaches a certain threshold V_{thr} . This spike is propagated down the network to other connected

neurons, and the membrane potential is reset to the resting potential V_{reset} , for a refractory time period τ_{ref} .

The total synaptic current $I_{syn}(t)$ is the sum of the currents received from other neurons via glutamatergic (AMPA) excitatory connections ($I_{A,ext}(t)$), excitatory (AMPA and NMDA) and inhibitory (GABA) current received via feedback connections ($I_{A,rec}(t), I_{N,rec}(t), I_G(t)$), and external background currents received from external neurons generating uncorrelated Poisson spike trains ($I_P(t)$)

$$I_{syn}(t) = I_{A,ext}(t) + I_{A,rec}(t) + I_{N,rec}(t) + I_G(t) + I_P(t). \quad (2.16)$$

For each ion channel, the current is calculated as

$$I_{ion}(t) = g_{ion}(V(t) - V_{ion}) \sum_{j=1}^N w_j s_j^{ion}(t), \quad (2.17)$$

where g_{ion} is the ionic conductance, V_{ion} is the reversal potential for that channel, w_j is the coupling, and $s_j^{ion}(t)$ is the gating variable describing the fractions of open channels

$$\frac{ds_j^{ion}(t)}{dt} = -\frac{s_j^{ion}(t)}{\tau_{ion}} + \sum_k \delta(t - t_j^k), \quad (2.18)$$

where τ_{ion} are time constants, δ is the delta function, and k denotes the sum over all the spikes emitted by the j -th neuron.




Figure 2.9 – Schematic of the spiking network model. Each cortical region consists of two populations of excitatory and inhibitory neurons (large circles), where excitatory neurons are shown using red triangles and inhibitory interneurons using blue circles. All the neurons of the network are fully connected with each other. Black and yellow arrows show connections between populations within a single region, white read arrows show connections between different cortical regions, which are via excitatory populations only. Adapted from Deco and Jirsa (2012).

The dynamics at the regional level are obtained by coupling the dynamical equations describing each neuron. In their work 100 excitatory neurons and 100 inhibitory neurons were used to model the dynamics of each region. The whole network had 66 brain regions resulting in a total of 13,200 neurons. A mean field reduction of the model was performed to explore the parametric conditions under which the model optimally predicts the observed empirical responses. This reduced the number of state variables for each region that need to be integrated (Brunel and Wang, 2001).

By comparing the simulated responses with empirical data, it was found that the model optimally predicts the empirical BOLD responses in a multistable

regime of weak couplings. This regime is characterized by the coexistence of multiple co-occurring states, such as the multiple networks observed in the resting state brain. It was further demonstrated that realistic resting state functional connectivity patterns can emerge from structured noise fluctuations. In particular, the model reproduced seed-based correlations between brain regions observed in previous fMRI studies.

2.5.7 Kuramoto model

More recently, Cabral et al. (2011) developed a large-scale model of the human brain's neural activity to study fMRI resting state fluctuations. The human brain network developed by Hagmann et al. (2008) was used, with the delays between regions based on length of white matter tracts. This work was later extended to investigate the resting state dynamics in MEG (Cabral et al., 2014). In this work, they used another human brain network of 90 brain regions, with propagation delays based on the AAL parcellation template.

The dynamics of the individual regions were based on the Kuramoto model of phase oscillators (Kuramoto, 1984) assumed to spontaneously oscillate in the gamma frequency range. Since coupled populations of neurons produce dynamics that are oscillatory, one can make a priori assumption that the individual brain regions are in an oscillatory regime to circumvent the complexity of the neural model. The interactions between the regions are based on a sine function of the delayed phase difference between the oscillators. The phase of the n -th region at time t ($\theta_n(t)$) is given by

$$\dot{\theta}_n(t) = \omega + k \sum_{p=1}^N C_{np} \sin[\theta_p(t - \tau_{np}) - \theta_n(t)] + \eta_n(t), \quad (2.19)$$

where C_{np} and τ_{np} specify, respectively, the coupling strength and the delays between regions n and p ; k is the global scaling of the coupling matrix; ω is the intrinsic frequency ($2\pi \times 40\text{Hz}$) that is homogenous for all regions; and $\eta_n(t)$ is uncorrelated Gaussian white noise.

A sine function was used to transform the phase responses into neural activity $r_n(t)$

$$r_n(t) = \sin(\theta_n(t)). \quad (2.20)$$

In the study investigating resting state functional connectivity in FMRI (Cabral et al., 2011), the simulated neural activity was transformed into hemodynamic responses using Balloon-Windkessel model (Friston et al., 2003). At sufficiently weak coupling strengths and realistic propagation delays (10-20ms), the interactions between brain regions led to the emergence of slow frequency fluctuations (< 1 Hz). It was found that the synchrony fluctuations modulated the power of the low frequency bands, matching the observations made in empirical FMRI data. Patterns of positive and negative correlations between different regions were also observed as in previous FMRI studies.

In their later study of MEG activity (Cabral et al., 2014), in a similar regime to their FMRI study, interactions between brain regions resulted in metastable oscillations in the alpha and beta frequency bands (8-30 Hz). This regime is characterized by transient periods of stability, or the recurrent expression of different transient states. The delays were found to be the governing factor for the stability of the rhythms. Analysis of the band-limited amplitudes of the alpha- and beta-frequency oscillations revealed similar correlation patterns as in empirical MEG studies.

3 Biophysical Network Dynamics

3.1 Introduction

The mechanisms underlying the spontaneous emergence of coherent brain activity during rest remain largely unknown. A number of computational models have been developed to investigate this activity, and they have provided much insight into its key determinants. These studies suggest that resting state dynamics arise as emergent properties when brain regions are coupled using a realistic neuroanatomical network, and point to the central role of brain connectivity, couplings, delays, and noise in shaping these dynamics. The majority of the recent models have used a realistic brain-inspired connectivity, giving rise to spontaneous BOLD or MEG FC patterns similar to those reported in empirical studies, highlighting the role of network topology (Ghosh et al., 2008; Honey et al., 2009; Deco et al., 2009; Deco and Jirsa, 2012; Cabral et al., 2014). These studies also suggest that moderate couplings between brain regions enable a metastable regime where transient synchronizations occur that give rise to non-stationaries as observed in empirical FMRI studies. The background noise seems to play an essential role in the spontaneous switching between the different metastable states (Ghosh et al., 2008; Deco et al., 2009; Deco and Jirsa, 2012). Finally, the transmission delays seem to have a significant effect on the frequency of the emergent oscillations (Ghosh et al., 2008; Carbal et al., 2014).

Most of the previous modelling studies have focused their investigations on FMRI. The work by Cabral et al. (2014) is one of the few studies to model the spontaneous activity observed in MEG. However, an important limitation of the Kuramoto model used in this study is its high degree of abstraction. A phase oscillator that has no direct biological relevance describes the dynamics of this

model. Furthermore, the model only relies on phase interactions between brain regions, while disregarding their amplitudes, which highly restricts the dynamics of the individual regions.

In this work, we propose a biophysical network model for simulating spontaneous MEG activity that overcomes the limitations of previous approaches. The network model we developed combines key aspects of brain anatomy and physiology by using a biophysical model for node dynamics with empirically derived brain connectivity. The parameters of the model are biologically realistic having an explicit neuronal interpretation, and encode, among other things: inter-region couplings, time delays, background noise, synaptic constants, and inhibition/excitation properties. We study the synchronization behaviour of the model as a function of global coupling strength and delays. We find that in a regime of moderately weak coupling strengths and biologically plausible transmission delays, the model leads to the spontaneous emergence of metastable oscillatory states (transient synchronous oscillations), where subsets of brain regions transiently synchronize although the network is not globally synchronized. We show that when the model is in this regime, it produces oscillatory activity primarily in the low frequency bands (8-30 Hz) that exhibit temporal non-stationarities. This non-stationarity is underpinned by a bistate (active/inactive) behaviour that is prominent on slow time scales (< 1 Hz). This behaviour is similar to what is observed in empirical MEG data, and we will compare the model responses with empirical data in subsequent chapters.

3.2 Methods

3.2.1 Structural brain connectivity

We used a large-scale human brain anatomical network derived from diffusion tensor imaging (DTI) data acquired from 21 healthy participants (see Figure 3.1, see Appendix B - *DTI Anatomical Connectivity* for more details). The data were

parcellated into 90 automated anatomical labelling (AAL) (Tzourio-Mazoyer et al., 2002) defined brain locations. The anatomical connectome at the voxel level was estimated by applying a probabilistic tractography algorithm. At the region level, the connectivity was calculated based on the total number of voxels in the AAL defined seed regions. A weighted network was derived by associating strength with each connection based on the number of white matter tracts connecting pairs of brain regions, resulting in coupling matrix, $\mathbf{C}$, that was averaged across all subjects. The signal conduction delays between the regions were calculated from the Euclidean distances between these regions and a finite signal propagation velocity, resulting in a delay matrix, $\mathbf{D}$.




Figure 3.1 – Estimation of brain connectivity using diffusion MRI. Diffusion MRI data is collected for the whole brain and the white matter tracts are estimated using probabilistic tractography. The AAL parcellation template is used to define seed regions between which the structural connectivity is estimated.

3.2.2 Neurodynamical model

The intrinsic dynamics of individual brain regions were modelled using the biophysically based dynamic causal model (DCM, Marreiros et al., 2008). The DCM model combines key aspects of regional anatomy and physiology to furnish a hierarchical cortical model governing the dynamics of individual brain regions. Based on empirical observations, a single cortical region is divided into three different layers (subpopulations), reflecting the layered structure of human cortex

(Crick and Koch, 1998; Felleman and Van Essen, 1991) (Figure 3.2). The supra-granular layers in the model represent layers II and III of the cerebral cortex, and consist mainly on inhibitory inter-neurons. The granular layer corresponds to layer IV of the cerebral cortex, which consists of spiny stellate cells and serves as the main input site of the region. The infra-granular layers represent deeper layers V and VI of the cerebral cortex that consist of large pyramidal neurons and serve as the main output site of the region. These populations are inter-connected via intrinsic connections within each region that highlight the anatomy of regional connectivity.




Figure 3.2 – Schematic of the DCM model for a single cortical region. Each source is divided into three different subpopulations. The inhibitory inter-neurons correspond to layers I and II of the cortex, spiny cells correspond to layer IV, and pyramidal cells correspond to deeper layers V and IV. These subpopulations are interconnected via intrinsic connections within each region. Connectivity between different sources is mediated via different types of extrinsic connections. The specific source and target subpopulations define the connection as forward, backward or lateral. The dynamics of each subpopulation are based on a conductance-based model of neurons.

Different cortical regions are connected to each other via different types of extrinsic connections (Figure 3.2, 3.3). These long-range cortico-cortical

connections are excitatory and comprise the axonal processes of pyramidal cells. All extrinsic connections, therefore, emerge from the pyramidal population of one cortical region and feed into different populations of another region, with the specific connected populations determining the type of connection.




Figure 3.3 – Extrinsic connections of the DCM model. Different subpopulations are connected to each other via different types of extrinsic connections. The specific subpopulations involved in the connection determine its type. All connections emerge from the pyramidal population of the source region. The connections are defined as forward (grey) if they project only to the spiny stellate cells of another region. The connections are defined as backward (red) if they project to the inhibitory and pyramidal populations of another region. The connections are lateral (green) if they feed to all the subpopulations of another region.

The directed connections are classified as bottom-up or forward if they originate in the agranular layers and terminate in granular layer (spiny stellate) of another region. Top-down or backward connections connect agranular layers of both regions, and thus they feed to both of the agranular layers (inhibitory interneurons and pyramidal cells) of the target region. Lateral connections originate in

the agranular layers and target all three layers of another region. These extrinsic connections reflect the hierarchal organization observed in human cortex (Felleman and Van Essen, 1991).

Due to the way the extrinsic connections are defined, the inhibitory interneurons and spiny cells connect only locally, with spiny cells also receiving external input, while the pyramidal neurons project both within the region and to other regions. The different types of connections can be used to ensure that specific populations couple to specific conductance types (i.e. each population can only release one sort of neurotransmitter).

The DTI-derived anatomical work we used does not contain information regarding the type of connections, and thus it is difficult to set up the network connections in a principled way. We found that the inclusion of lateral connections causes additional inhibition of network activity, which destroys the emergent metastable regime. Therefore, we used only forward and backward connections in our simulations, such that the forward connection type mediated a feed-forward connection from region i to region j whereas the backward connection type mediated a feedback connection from region j to region i . In the case where we test the effect of globally scaling anatomical connectivity, the global scaling, k , homogenously scales all the different extrinsic connections specified in the network.

Within a single region, the dynamics of each population conform to a conductance based voltage model of neurons (Morris and Lecar, 1981). A single region expresses the coupling between populations of excitatory and inhibitory neurons. The state equations governing the dynamics of a single population at time t are given by

$$C\dot{V}(t) = g_L(V_L - V(t)) + g_E(t)(V_E - V(t)) + g_I(t)(V_I - V(t)) + I(t), \quad (3.1)$$

$$\dot{g}_E(t) = \kappa_E(\zeta_E(t) - g_E(t)), \quad (3.2)$$

$$\dot{g}_I(t) = \kappa_I(\varsigma_I(t) - g_I(t)), \quad (3.3)$$

where the states comprise the transmembrane potential $V(t)$, and the excitatory ($g_E(t)$) and inhibitory ($g_I(t)$) conductances corresponding to different types of ion channels. The rate of change of membrane potential (times capacitance, C) depends on the total current across the cell membrane plus exogenous current $I(t)$. The different membrane currents depend on the conductance of different voltage-gated ion channels ($g_L(t), g_E(t), g_I(t)$), and the reverse potential for that channel. This reverse potential determines whether the channel is leaky (V_L), excitatory (V_E), or inhibitory (V_I). These currents can either hyperpolarize or depolarize the cell depending on whether they are mediated by inhibitory or excitatory receptors, respectively. The transmembrane potential across the pyramidal population is the presumed source of observed MEG signals. We modelled random fluctuations in the network as uncorrelated Gaussian white noise that perturbed the dynamics of each cortical region (fed into the spiny stellate population).

The excitatory and inhibitory conductances change dynamically with characteristic rate constants κ_E and κ_I that can be regarded as the number of open channels. Channels open in proportion to presynaptic input $\varsigma(t)$, which depends on the firing rates of other populations. The propagation of activity is based on mean-field effects (see later) where both the mean and the variance of the neuronal states of each population are calculated using the non-linear Fokker-Plank equation (Frank et al., 2000; Marreiros et al., 2008). Under these assumptions, the presynaptic input at time t is given by

$$\varsigma_z^i(t) = \sum_j \gamma_{ij}^z \phi(\mu_V^j(t) - V_R, \Sigma^{jj}(t)), \quad (3.4)$$

for the i^{th} population and z^{th} conductance $z = [E, I]$. $\mu_V^j(t)$ is the mean voltage in population j , $\Sigma^{jj}(t)$ is the variance of the voltage, V_R is the threshold voltage, and γ_{ij}^z is the coupling between populations i and j . These populations can either be

within the same cortical region or in another one. Therefore, the coupling parameter absorbs both intrinsic and extrinsic connections, which couple populations within and between brain areas, respectively.

The firing in source population is a threshold function of depolarization, where the threshold V_R determines the proportion of cells firing (Figure 3.4). Under the assumptions of mean-field effects, the sigmoid function ϕ is a Gaussian cumulative density function

$$\phi(\mu_v, V_R, \Sigma^{jj}, t) = 1/2[1 + \text{erf}((\mu_v(t) - V_R/\Sigma^{jj}(t))/\sqrt{2})], \quad (3.5)$$

where erf is the standard error function.




Figure 3.4 – Synaptic dynamics of the DCM model. (A) The shape of the nonlinear sigmoid function that transforms the transmembrane potential into firing rates. A spike is fired when the voltage exceeds a certain threshold (-40mV). (B) The excitatory synaptic currents across the membrane of the postsynaptic neuron in response to the discharge of another population. The build-up of synaptic currents increases the transmembrane potential of the postsynaptic neuron and this neuron fires in cascade when the voltage exceeds the threshold.

3.2.3 Mean-field dynamics

The equations of motion provided above describe the activity of a single neuron. The model can be recast as an ensemble model using the non-linear Fokker-Plank

equation (Frank et al., 2001). The Fokker-Planck equation prescribes the evolution of the ensemble dynamics, given the equations of motion that constitute our neural model. In the following, we provide a brief summary for deriving a mean-field model using the Fokker-Planck formalism (see Marreiros et al. (2008) for more details).

The Fokker-Planck equation aims to model the evolution of the ensemble density, $\mathbf{q}$, over phase-space

$$\dot{\mathbf{q}} = -\nabla \cdot \mathbf{f} \mathbf{q} + \nabla \cdot \mathbf{D} \nabla \mathbf{q}, \quad (3.6)$$

where $\chi = [V, g_E, g_I]$. This equation comprises flow (left term) and dispersion (right term), which embed the assumptions about the dynamics and random fluctuations, respectively. The flow ($\mathbf{f}_\chi$) and dispersion ($\mathbf{D}_\chi$) constitute the model at the neural level. Under the Fokker-Planck formalism, even if the dynamics of each neuron are very complicated, or even chaotic, the density dynamics remain simple, linear and deterministic.

In a system with many neural states or different types of population, the phase-space can become very large and computationally intractable. To overcome this problem, a mean-field approximation can be used. In this approach, the evolution of ensemble densities of different populations is considered independently, and these ensembles are coupled through mean-field effects so that the flow of one ensemble depends on the others (the Gaussian cumulative density function presented in the previous section). Importantly, this implies that all the neurons in one ensemble are affected identically by the average behaviour of another ensemble.

The individual ensemble densities of different populations can be further parameterized in terms of statistics of higher order than just the mean, resulting in equations describing the evolution of these statistics. If we assume that the ensemble density of each population conforms to a Gaussian density then the

population dynamics are described by equations governing the evolution of the population's mean and covariance (the sufficient statistics) given by

$$\dot{\boldsymbol{\mu}}_i = \mathbf{f}_i + \frac{1}{2}(\boldsymbol{\Sigma}\mathbf{H}_i)^T, \quad (3.7)$$

$$\dot{\boldsymbol{\Sigma}}_{ij} = \mathbf{J}\boldsymbol{\Sigma} + \boldsymbol{\Sigma}\mathbf{J}^T + \mathbf{D} + \mathbf{D}^T, \quad (3.8)$$

for the i -th state. $\mathbf{f}_i$ is the flow of state i (these refer to the equations of motion presented in the previous section). $\mathbf{D}$ is the dispersion that was based on the standard Wiener process, $\boldsymbol{\Sigma}$ is the covariance between states, $\mathbf{J} = \frac{\partial \mathbf{f}}{\partial \boldsymbol{\chi}}$ is the systems Jacobian, and $\mathbf{H}_i = \frac{\partial^2 \mathbf{f}_i}{\partial \boldsymbol{\chi}_j \partial \boldsymbol{\chi}_k}$ is the systems Hessian for the i -th state.

We derived analytical Jacobian and Hessian matrices for the model for computational efficacy. The Jacobian matrix for a single population at time t is given by

$$\mathbf{J}(t) = \begin{bmatrix} (-g_L - g_E(t) - g_I(t))/C & (V_E - V(t))/C & (V_I - V(t))/C \\ \kappa_E \dot{\phi}(\mu_V(t)) & -\kappa_E & 0 \\ \kappa_I \dot{\phi}(\mu_V(t)) & 0 & -\kappa_I \end{bmatrix}. \quad (3.9)$$

The Hessian matrix is only non-zero for the voltage state

$$\mathbf{H}_v = \begin{bmatrix} 0 & -\frac{1}{C} & -\frac{1}{C} \\ -\frac{1}{C} & 0 & 0 \\ -\frac{1}{C} & 0 & 0 \end{bmatrix}, \quad (3.10)$$

$$H_{gE} = H_{gI} = 0. \quad (3.11)$$

These equations, when integrated, generate the dynamics of the whole ensemble. We simulated the network dynamics for 330s using the SPM8 package (FIL, UCL). We always discarded the first 30s of simulations to avoid the dependence of our results on initial conditions. We also performed various network simulations while keeping the covariance between the states fixed. In this case we used the fixed covariance provided in the SPM8 implementation. We additionally performed several simulations without including mean field effects.

In this case the variance of the voltage state (Σ^{jj}) was kept constant. In both these cases, we found that the characteristics of the simulated activity were never better than those using the full mean field model and, hence, these results are not reported further. All parameters and the values used in our simulations are provided in Table 3.1.

Table 3.1 – Parameter values for the DCM model used in the simulations.

Parameter	Interpretation	Value
C^F	Extrinsic forward connections	Structural matrix
C^B	Extrinsic backward connections	Structural matrix transpose
C^L	Extrinsic lateral connections	0
$\gamma_{31}^E, \gamma_{13}^E, \gamma_{23}^E$	Intrinsic excitatory connectivity	1, 0.5, 1
$\gamma_{12}^I, \gamma_{22}^I, \gamma_{32}^I$	Intrinsic inhibitory connectivity	0.5, 0, 2
τ_{ii}	Intrinsic conduction delays	2ms
τ_{ij}	Extrinsic conduction delays	Delay matrix
g_L	Leaky conductance	1mV
$\tau_{E,I} = 1/\kappa_{E,I}$	Postsynaptic rate constants	4ms, 16ms
V_L, V_E, V_I	Reverse potential	-70mV, 60mV, -90mV
V_R	Threshold potential	-40 mV
Σ^{jj}	Population variance	75
C	Capacitance	1 μ F
Γ	Intrinsic noise	Gaussian white noise

3.2.4 Delay differential equations

An important step to make more realistic models of macroscopic brain activity is the incorporation of time delay effects. Time delays in the brain arise from finite action potential propagation velocities as well as synaptic and dendritic processing, and are usually on the order of 10–100ms. However, delay differential equations can be very expensive to compute in non-linear systems of many variables. To solve these equations efficiently, we employed a caching scheme where we cached the past values of all state variables only up to a maximum specified delay time. The time delay matrix was scaled by a mean delay of 15ms,

except in the case where we explicitly test the effect of delays on synchronization dynamics.

3.2.5 Global network synchronization

To characterize the synchronization behaviour of the network at the global level, we calculated the order parameter, $R(t)$. This captures the collective behaviour of interacting sources over time based on their degree of phase coherence. Since the order parameter is calculated from the cortical phase time courses, we first estimated the instantaneous phase $\theta(t)$ of a filtered neural time course from the phase of its analytical signal derived using the Hilbert transform (see Section 4.2.5 for more details). The phase of the signal can also be identified via the phase of the complex wavelet transform using Gabor functions, but both methods have been shown to produce similar results (Le Van Quyen et al., 2001; Quiñero et al., 2002). The frequency bands used were 1-70Hz, 1-4Hz, 4-8Hz, 8-13, 13-30, and 30-70Hz. We linearized the phase sequences by adding π to the instantaneous phase at each reset (branch point) to get the unwrapped analytic phase, which was used to calculate the order parameter.

The order parameter, $R(t)$, is defined as

$$R(t)e^{i\Phi(t)} = \frac{1}{N} \sum_{n=1}^N e^{i\theta_n(t)}. \quad (3.12)$$

At any time point t , the modulus of order parameter captures the phase coherence among a set of N cortical regions with phases θ_n whereas Φ indicates their average phase. Over time, $R(t)$ gives the phase synchrony time course for the whole network. To characterize the dynamics of the network over the whole simulation time interval, we calculated the following three summary measures from the phase synchrony time course.

Synchrony. The mean synchronisation level, R , is estimated as the mean of the phase synchrony time course. This reflects the average global synchronisation in the network over the whole time interval.

Metastability. The metastability, S , is calculated from the standard deviation of the phase synchrony time course. This measure, therefore, reflects fluctuations in synchrony degree over time, as the network undergoes relatively long-lived unstable and transient states.

Asymmetry. The asymmetry, X , is calculated from the skewness of the phase synchrony time courses. This measure captures the asymmetry of the synchrony time course.

3.3 Results

3.3.1 Structural brain connectivity

We consider the network dynamics supported by the structural connectivity of the human brain (Figure 3.5). The connectivity strength matrix together with the corresponding histograms are shown in Figure 3.5A and C, respectively. The connection length matrix together with the histograms are shown in Figure 3.5B and D, respectively. The distribution of coupling strength values decreases sharply and the connection lengths exhibit a gamma distribution. Previous work has shown that brain networks exhibit modularity, with the presence of small-world networks, characterized by local dense clustering (Honey et al., 2007; Bullmore and Sporns, 2009).




Figure 3.5 – Structural connectivity of the human brain. (A) The connection strength matrix. The number of fibre tracts detected that connect each pair of regions was taken as an index of the coupling strength. The matrix is normalized to $[0, 1]$. (B) The length of existing connections based on the 3-D Euclidean distance between pairs of brain regions. (C) Histogram of connection strengths. (D) Histogram of connection lengths. The matrices are ordered in such a way that corresponding contra-lateral regions are symmetrically arranged with respect to the matrix centre.

3.3.2 Biophysical network dynamics

We first look at the effect of changing the global coupling strength, k , on network dynamics. The global coupling uniformly scales the connectivity matrix (*i.e.* $k \times C$). Therefore, only the scaling of the connectivity matrix changes while its structure remains fixed. Figure 3.6A shows the mean global synchrony, R , of the network as a function of global coupling strength, k . At each coupling strength, k , the metastability, S , and skewness, X , are shown in Figure 3.6B and C, respectively.

We observe that the network synchrony remains close to its minimum value of 0, when the phases of different regions are uniformly distributed (see Figure 3.7 Left). In this case, the network is in a stable incoherent state. As the global coupling crosses a critical value k_c (~ 1.8), the synchrony rapidly approaches its maximum value of 1 as the phases of all regions become aligned (fully synchronized network, see Figure 3.7 Right). This is also reflected in the asymmetry index of the network (Figure 3.6C). At low couplings, the synchrony time course is random, and the asymmetry is high. As the couplings are increased, the network achieves full synchronization (resulting in a flat synchrony time course) and the asymmetry of the network drops to zero. Importantly, there is a range of critical couplings in which the model displays its emergent properties (see later). The precise value of the critical coupling we use here is based on our findings that this is where the model optimally predicts empirical data, as demonstrated in the subsequent chapters (see Section 4.3.1).




Figure 3.6 – Effect of coupling strengths on network dynamics. (A) The mean global synchrony of the network. At each global coupling, the synchrony is estimated from the mean of the phase synchrony time course. (B) The global metastability level of the network. At each coupling, this is estimated from the standard deviation of the synchrony time course, and indicates the variability of synchrony over time. (C) The asymmetry index of the network that, at each coupling, is estimated from the skewness of the synchrony time course. The dashed line indicates the critical coupling ($k_c \sim 1.8$).

The most irregular dynamics are found at intermediate couplings, representing an intermediate regime between incoherence and full synchronization. In this regime, the global coupling is sufficiently weak and perturbations in the network lead to the formation of transient stable states. Therefore, multiple partially synchronized states alternate in time with periods of incoherence. As a result, the network synchronization level fluctuates in time, and these fluctuations are captured by the standard deviation of synchrony time course, which indicates the level of metastability of the network (Shanahan, 2010; Figure 3.6B). This results from a stochastically driven exploration of different state space positions around the network's equilibrium state, indicative of self-organised criticality (Turcotte and Rundle, 2002). This can be considered as a critical state at the border of stability and chaos formed by the spontaneous self-organization in a long transient period. Chaos is a non-linear deterministic process (Ricklefs et al., 2007) that apparently looks random but the distinguishing feature of chaos is the high sensitivity on initial conditions (network parameters).




Figure 3.7 – Phase dynamics as a function of global coupling strength. The phases of all oscillators at one time instant are shown for low (left), intermediate (middle), and strong (right) coupling strengths. As the global coupling k increases, nodes go progressively from incoherence (left) to global synchrony (right). In an intermediate regime (middle), regions exhibit varying degrees of synchrony.

3.3.3 Modular connectivity and synchronization patterns

The networks of the brain have a complex modular connectivity architecture, and that this has been shown to impact the synchronization dynamics (Strogatz, 2001). We find that the network's structure expresses itself in the emergent synchronization patterns (Figure 3.8). Brain regions that have a high structural connectivity between each other tend to form synchronized modules at intermediate coupling strengths, and these modules oscillate at a different frequency and phase from other modules. When these clusters get synchronized, the metastability index for these regions drops relative to other regions. However, as the coupling strength increases further, all regions get synchronized, and the metastability for all the regions drops.




Figure 3.8 – Network dynamics at the regional level as a function of global coupling strength. The synchrony (A), and metastability (B), between all pairs of brain regions is shown for low (left), intermediate (middle), and high (right) coupling strengths. The synchrony and metastability between pairs of regions reflect the patterns of structural connectivity. At low couplings, the synchrony between all regions is low. At intermediate couplings, the increase in synchrony is highest for the most strongly connected regions. As the coupling is increased further, the synchrony increases and the metastability drops for most pairs of regions.

3.3.4 Frequency specificity

The presence of time delays can have a huge impact on the dynamics of the network. In the presence of delays, stable synchronized states occur at reduced collective frequencies, as has been described for systems of coupled oscillators with time delays (Niebur et al., 1991). This reduced collective frequency depends on the natural frequency of the oscillators, global coupling strength, mean delay and the number of regions involved.

Figure 3.9 shows the network dynamics and frequency specificity of the network as a function of mean time delays (for $k \sim 1.8$). Here, the whole delay matrix is scaled by the mean delay without any changes in its structure. For sufficiently small delays, the network synchronizes even for a very weak coupling, and the metastability remains low. As the mean delay increases, synchrony fades away, and the metastability level of the network increases. For large delays, there is higher frequency suppression in the network that gives rise to low collective frequencies in the alpha and beta frequency bands (8-30 Hz).




Figure 3.9 – Effect of time delays on network dynamics. (A) The mean global synchrony of the network as a function of mean time delay. (B) The global metastability level of the network as a function of mean time delay. (C) Collective frequency as a function of mean time delay. The dashed line indicates the mean delay of 15ms.

Figure 3.10 shows frequency-specific metastability of the network for increasing coupling strengths, using a fixed mean delay of 15ms. At weak couplings, the metastability in all frequency bands is low. As the coupling strength is increased, the frequency of the emergent oscillations drops, and the metastability in these bands increases. Due to the presence of delays in the network, alpha and beta oscillations start to appear at intermediate couplings, and the metastability in these frequency bands increases. However, at even higher coupling strengths lower frequency oscillations in the delta and theta (<10 Hz) frequency bands emerge and the metastability in the alpha and beta frequency bands drops. Importantly, the heterogeneity of the structural connectivity leads to the formation of multiple clusters which oscillate at their own collective frequency. This gives rise to a spectrum of frequencies resulting in rich system dynamics.




Figure 3.10 – Frequency specificity of network dynamics. The metastability index of the network is shown in 5 different frequency bands (d – delta, t – theta, a – alpha, b – beta, g – gamma) at low (left), intermediate (middle), and high (right) levels of global couplings. At low couplings, the metastability in all frequency bands is low due to the low overall synchrony of the network. As the coupling strength is increased, the frequency of the emergent oscillations drops, and lower frequency oscillations start to emerge. At intermediate couplings, oscillations emerge in the alpha and beta frequency bands due to the presence of delays, and the metastability in these frequency bands goes up. As the coupling strength is increased further, low frequency oscillations in the delta and theta frequency bands emerge, and the metastability in the alpha and beta frequency bands drops.

3.3.5 Transient synchronizations and power fluctuations

During the period of synchronized cluster formation, the band-limited power (BLP) at the reduced collective frequency increases. However, if the coupling is sufficiently weak, this state is only transient and can be easily disrupted by noise or other competitive processes. When this happens, the signal power drops back again. This leads to structured power fluctuations in the network that are a result of transient synchronizations occurring in the network. Figure 3.11 shows the synchrony time course of the network (A) and corresponding power between 10 to 20Hz (B), for different coupling strengths. At intermediate couplings, we observe the emergence of transient synchronized states that result in power modulations in low frequency bands. We also note that these states emerge spontaneously and last for several seconds before disappearing.




Figure 3.11 – Transient synchronizations and power fluctuations. (A) The synchrony time course of the whole network (blue) and low-pass filtered at cut-off frequency of 1Hz (red). (B) Signal power in the 10-20Hz frequency band over time. As the coupling strength is increased, the network synchrony increases from 0 (left) to 1 (right). At intermediate couplings (center), transient synchronizations in the network result in slow power fluctuations in the low frequency bands. In this regime, there is a high correlation between the smoothed synchrony time course (red line) and power fluctuation time course (black), suggesting that power fluctuations are a result of transient network synchronizations.

We further investigate the origin of these BLP fluctuations, by computing the correlation between the synchrony time course and the BLP in two frequency bands. We find that there is high correlation between synchrony time course and BLP fluctuations in low frequency bands (Figure 3.12 (for $k \sim 1.8$)), suggesting that power fluctuations in these bands are a result of transient network synchronizations. The synchrony time course and power at lower frequencies is characterized by a bimodal distribution, which is indicative of multiple underlying states, whereas the power at higher frequencies is mainly Gaussian.




Figure 3.12 – Network synchronizations and power fluctuations. (A) The synchrony time course of the whole network (blue) and low-pass filtered at cut-off frequency of 1Hz (red). The probability distribution of the synchrony time course is qualitatively shown at the fast (B) and slow (C) timescales. (D) Time course of signal power in the 10-20Hz frequency band with the corresponding probability distribution (E). (F) Time course of signal power in the 30-40Hz frequency band with the corresponding probability distribution (G). (H) There is a high correlation between the smoothed synchrony time course (red line) and the power time course in the 10-20 Hz frequency band. The network synchrony and power at low frequencies is characterized by a bimodal distribution whereas power at high frequencies is characterized by a unimodal distribution.

The global coupling strength also determines the stability of the synchronized states. As shown in Figure 3.13, a slight increase in global coupling (from $k \sim 1.8$ to $k \sim 2$) results in transient states that are more stable and last for longer periods (>10 s).




Figure 3.13 – Stability of synchronized states. (A) The synchrony degree of the network over time (blue) and low-pass filtered at cut-off frequency of 1Hz (red). The probability distribution of synchrony degree at fast (B) and slow (C) timescales. The stability of synchronized states (right peak in the likelihood plot) increases with the coupling strengths.

3.3.6 Effect of network parameters

To further investigate the effect of the brain's space-time structure on the observed dynamics, we consider the cases where the structural connectivity, couplings, delays, and noise are completely altered.

Effect of structural connectivity. Figure 3.14 shows the effect of structural connectivity on the dynamics of the network. We find that the human brain connectome (blue) results in a smoother transition from incoherence to global synchrony, and a broader metastability regime, when compared to the macaque (red) or random (green) network. The random network was setup up such that the degree distribution and connectedness of the network (Rubinov and Sporns, 2010) were matched to that of the human brain network. This highlights the role of the modular structure of the structural connectivity matrix in shaping brain dynamics.




Figure 3.14 – Effect of structural connectivity on network dynamics. The network dynamics are shown for the human brain (blue), macaque brain (red), and a random network (green). (A) The global synchrony of the network. (B) The global metastability level of the network. (C) The global metastability index in the 8-30 Hz frequency band.

Effect of inter-regional couplings. Figure 3.15 shows the effect of changing inter-regional couplings between *all* regions on network dynamics. We find that randomizing the coupling strengths (but keeping the same delays) completely destroys the metastable dynamics (red). In the case of homogenous couplings (green), the transition to synchrony is very abrupt, resulting in a very narrow and weak metastable regime. This indicates that the heterogeneity of the human brain network is a key ingredient to obtain metastable states in the network.




Figure 3.15 – Effect of heterogeneity of coupling on network dynamics. The network dynamics are shown for the original couplings (blue), randomized couplings (red), and homogenous couplings (green). (A) The mean global synchrony of the network. (B) The global metastability level of the network. (C) The global metastability index in the 8-30 Hz frequency band.

Effect of time delays. The effect of altering the time delays between *all* regions on the dynamics of the network is shown in Figure 3.16. The original delays are shown in blue. In the case of homogenous delays (15ms), transition to synchrony occurs at weaker couplings, and the metastability regime (particularly in the 8-30Hz range) is much weaker (red). When delays are completely excluded from the network (green), the transition is very abrupt, and the dynamics in the 8-30 Hz frequency range are almost completely lost (green). These results highlight the role of the delays in shaping the metastable dynamics in the alpha and beta frequency bands (8 – 30 Hz).




Figure 3.16 – Effect of heterogeneity of delays on network dynamics. The network dynamics are shown for the original couplings (blue), homogenous couplings of 15ms (red), and the case where there are no delays between regions (green). (A) The mean global synchrony of the network. (B) The global metastability level of the network. (C) The global metastability index in the 8-30 Hz frequency band.

Effect of noise. We found that the noise level has an effect on the network dynamics but this was more related to coupling strengths. Specifically, we had to set up the noise level such that the synchrony in network remained weak at low levels of global coupling strengths. We also investigated the network dynamics where the brain regions were driven by a binary input as opposed to white noise. This signal can be viewed as inputs from other brain regions. The input was based on a Poisson process, controlling the likelihood of switching between the two states. The average of the active state was 2.5 s and that of the inactive state was 10 s,

reflecting the longer average rest periods than active ones. In this case, we found that no prominent spontaneous oscillations emerged in the network. These results show that the space–time structure of the network, as well as the presence of noise, provides an optimal framework for the emergence of robust metastable dynamics.

3.4 Discussion

In this work we developed a biophysical network model to simulate spontaneous MEG activity and investigated its synchronization dynamics. We furnished the model by combining a realistic brain anatomical network with a biophysical model for node dynamics. We found that the model was highly sensitive to changes in coupling strengths and transmission delays, giving rise to simultaneous regions of low synchrony and high metastability. In the presence of delays, the network is more stable at reduced collective frequencies that fall in the alpha and beta frequency bands (8-30 Hz). Since the synchrony is sufficiently weak, noise and other competitive processes in the network disrupt this delicate state giving rising to the observed temporal fluctuations in synchrony time courses. These results are consistent with previous studies that have found a similar dependency of these parameters in shaping network dynamics (Honey et al., 2007; Ghosh et al., 2008; Deco et al., 2009, 2012; Carbal et al., 2011, 2014).

The structural connectivity provides a substrate as well as major constraints on the dynamics of interacting regions. The connectivity matrix revealed a highly intricate structural organization with small-world topology, which has been implicated in network equilibrium stability and generation of novel oscillatory network states (Deco et al., 2011). The small-world attributes of the brain network are also reflected in the synchronization patterns. Due to the small-world architecture, subsets of brain regions can get highly synchronized while the network synchrony is still low enabling the formations of networks composed of

specific brain regions. We also found that the greater modularity of the human brain connectivity results in a broader metastable regime compared to the Macaque brain, which further highlights the essential role of connectivity in shaping the network dynamics.

Besides the structural connectivity, it is important to setup the system at an appropriate working point to display its behaviour. We found that metastable states in the network emerge at certain global coupling strengths. The coupling strengths need to be sufficiently weak so that synchronization between brain regions is only temporary to avoid full synchronization of the network. However, at much lower couplings, the global network integration decreases considerably. This type of de-synchronization between regions has been observed in schizophrenia patients, suggesting this to be a disease of connectivity. In agreement with Niebur et al. (1991), the collective frequency of the metastable states was found to vary inversely with global coupling strengths (i.e. lower frequencies occur at higher coupling). Therefore, the global coupling strengths can be manipulated to fine tune specific characteristics of the network, including frequency specificity and network integration. Furthermore, the duration of the metastable synchronized states (i.e. dwell time) also depends on the coupling strengths, and can be altered to match the time scales of fluctuations seen in FMRI or MEG.

The presence of transmission time delays between brain regions adds a temporal component to the space-time structure of network, where the anatomical network captures the spatial component. In the presence of time delays, synchronizations among brain regions give rise to a novel form of frequency depression where the system decays to stable states, which oscillate at a delay and interaction dependent collective frequency, as predicted by previous studies (Niebur et al., 1991). We found that biologically plausible delays of 10-15ms result in slow neural activity fluctuations in the alpha and beta frequency bands, and

excluding these delays significantly alters the network dynamics in these frequency bands. This is consistent with the findings of Cabral et al. (2014), who identified emergent MEG resting state networks characteristics in this regime. In previous fMRI modelling studies, several resting state characteristics were also identified in a similar parameter range (Ghosh et al., 2008; Deco et al., 2009).

A number of studies have shown that the brain wanders between different states during rest, where different subsets of brain regions form temporary coalitions, which alternate in time (Chang and Glover, 2010; Kiviniemi et al., 2011; Allen et al., 2012; de Pasquale et al., 2012; Brookes et al., 2014). Several previous modelling studies have suggested an important role of noise in the transition between different stable states (Ghosh et al., 2008; Deco et al., 2009; Deco and Jirsa, 2012). In our work we found that when the network is in the metastable regime, temporary collations between brain regions form that are easily disrupted by noise processes with the emergence of new metastable collations. The sub-threshold level of noise induces sufficient fluctuations in individual regions to explore the metastable trajectory of the system. Although several sources of noise are likely present in the brain (e.g., Brownian motion of molecules, spontaneous synaptic vesicle release, stochastic opening of ion channels, etc.), it is unlikely that noise is a robust signal that encodes information. However, converging evidence from empirical (Winterer et al., 1999) and computational (Frank et al., 1999) studies suggests that noise plays an important role in the brain computation.

Our results support that the resting state brain works in a state of self-organised criticality. Self-organized criticality allows the brain to process information by recruiting a large number of specialized processors distributed across the brain, and to quickly adapt to new situations by re-arranging the spatiotemporal synchronizations among the regions. Large groups of regions engage in temporary synchronizations to form clusters that process information at the same frequency (phase locking). This allows widely distributed brain regions

to communicate among themselves without disruptions from other regions. Bistability allows both symmetric solutions, where each region maintains a similar activity level, and winner-take-all behaviours, where one region suppresses another, to exist simultaneously. This way the brain can optimally compute while in the resting state and quickly jump to new states for task based processing.

4 Static Functional Connectivity

4.1 Introduction

Studies of resting state brain activity based on BOLD fMRI has revealed several brain networks, which show highly robust correlated fluctuations on slow time scales (< 0.1 Hz) (Beckmann et al., 2005; Smith et al., 2009). Similar networks have also been found in MEG data that exhibit significant correlations in alpha and beta frequency bands (Hipp et al., 2012; Brookes et al., 2011b). Furthermore, the power in these frequency bands correlates with the BOLD time course (Laufs et al., 2003; Mantini et al., 2007). These results indicate that resting state FC utilizes frequency-specific bindings between remote brain regions.

Previous work using neural models has shown that patterns of FC emerge as the result of neural interactions taking place on the anatomical network (Deco et al., 2009; Cabral et al., 2014). In particular, these studies propose that the space-time structure of connectivity and delays leads to the formation of emergent RSNs, which display similar frequency specific correlations. However, most previous studies have focused their investigations on fMRI and, the mechanisms underlying the frequency specific correlations in MEG and the link between the fMRI BOLD signal and MEG oscillations remain unknown.

To investigate these mechanisms, we analyzed data from healthy subjects and a large-scale biophysically based network model of the brain. We tune the coupling strengths of the model using empirical data, and find that the best fit is obtained in a region of moderately weak coupling strengths and realistic transmission delays. In this region the model produces spontaneous oscillatory activity that displays the same frequency specificity as the empirical data. We show that the network model can reproduce frequency-specific amplitude

(~power) correlations in the alpha and beta frequency bands. By applying temporal averaging to the amplitude envelopes (down-sampling), we found significant correlations on the time scale of several seconds in a spectrally specific way. We finally use the non-linear balloon model to estimate the vascular dynamics resulting from the underlying neural activity, and show that the hemodynamic responses exhibit fluctuations on slow time scales (< 0.1 Hz) as found in previous fMRI studies.

4.2 Methods

4.2.1 MEG data collection and pre-processing

Data acquisition. MEG resting state data were acquired from 10 healthy participants in a 5-minute scan where the subjects were asked to sit still with their eyes-closed. The data were collected at a temporal resolution of 1000Hz on a 306 channel Elekta Neuromag scanner (Elekta Oy, Helsinki, Finland) comprising 102 magnetometers and 204 planar gradiometers. Localization of the head within the MEG helmet was achieved using four head position indicator coils, which were continuously energized allowing the head to be localized within the scanner helmet. Each subject's head shape was also recorded using a three-dimensional Polhemus Isotrak system (Polhemus Corporation, Colchester, VT, U.S.).

Data pre-processing. Signal space separation (Taulu et al., 2005) compensates for external interference and sensor artefacts by projecting the MEG data onto a basis set of spherical harmonics. Harmonics corresponding to sources originating from within the sensor array are preserved whilst interfering sources from outside the environment surrounding the sensor array are rejected. Using visual inspections, the channels containing sensor artefacts were identified and subsequently removed before applying the separation algorithm. This was also used to down sample the data to a resolution of 250 Hz. This was performed using

MaxFilter software (Taulu et al., 2005). The resulting data were converted to SPM8 format for subsequent analysis.

The sensor space MEG data were de-noised using temporal ICA to remove cardiac, 50 Hz mains and, in some subjects, eye movement artefacts. ICA decomposes the sensor space data into temporal independent components. Specifically, the data were decomposed into 64 components using fastICA (Hyvärinen, 1999) (64 is the rank of the MEG data after signal space separation). Prior to the ICA decomposition, each sensor type was normalised by its smallest eigenvalue to ensure that all sensor types have the same noise variance. Each component has a temporally independent time course (with unit variance and zero mean) and associated sensor topography. Eye-blink, cardiac and mains interference components were manually identified by the combined inspection of the spatial topography, time course, kurtosis of the time course and frequency spectrum for all components. Eye-blink artefacts typically exhibit high kurtosis (>20), a repeated blink structure in the time course and very structured spatial topographies. Cardiac component time courses strongly resemble the typical ECG signals, as well as having high kurtosis (>20). Mains interference has extremely low kurtosis (typically <-1) and a frequency spectrum dominated by 50Hz line noise. The sensor data is de-noised by subtracting the artefact components from the raw data (Mantini et al., 2011). We converted this subtraction into a multiplication via a residual forming matrix. This enabled us to simultaneously de-noise the data and correct the lead fields (via the montage function in SPM8).

Each dataset was co-registered into the standard Montreal Neurological Institute (MNI) space by registering the canonical MNI template to the subjects head shape data. A local spheres forward model (Huang et al., 1999) was then estimated using the subject's head shape. Both co-registration and forward model estimation were performed with the SPM8 package.

Source reconstruction. Source space analyses provide better estimates of underlying neuronal activity, as multiple brain regions contribute to single sensor channels, while also increasing the spatial resolution and signal to noise ratio (Brookes et al., 2008). This also makes the results more interpretable as activations may be better related to the underlying anatomy.

A scalar beamformer (Van Veen et al., 1997) was used to project the sensor space MEG data into source space, that is, to reconstruct the activity at the 90 intracranial locations defined by the AAL brain parcellation template. The electrical source strength, $Q_{\theta}(t)$, at time t and at a predetermined location in the brain is estimated as a weighted sum of MEG sensor measurements

$$Q_{\theta}(t) = \mathbf{W}_{\theta}^T \mathbf{m}(t), \quad (4.1)$$

where $\mathbf{m}(t)$ is a vector of magnetic field measurements made at M sensor locations at time t . $\mathbf{W}_{\theta}$ is a vector of weighting parameters tuned to specific source location and current orientation $\theta = [\mathbf{r}, \delta]$, where $\mathbf{r}$ represents location of source, and δ is the angle of the source with respect to the azimuthal direction. The superscript T denotes matrix transpose.

The weighting parameters $\mathbf{W}_{\theta}$ were computed based on power minimization. The signal power of the output signal $Q_{\theta}(t)$ is minimized with a linear constraint that the signal power originating at θ remains in the output signal

$$\min_{\mathbf{W}_{\theta}} [Q_{\theta}^2] \text{ subject to } \mathbf{W}_{\theta}^T \mathbf{L}_{\theta} = 1, \quad (4.2)$$

where Q_{θ}^2 represents output power given by $Q_{\theta}^2 = \mathbf{W}_{\theta}^T \mathbf{C} \mathbf{W}_{\theta}$. $\mathbf{C}$ represents the $M \times M$ data covariance matrix calculated over a time frequency window of interest. $\mathbf{L}_{\theta}$ represents the lead fields containing the magnetic field distribution across all sensor channels due to a unitary current dipole placed at a given position and orientation θ . A solution to this problem is given by (Van Veen et al., 1997)

$$\mathbf{W}_\theta^T = [\mathbf{L}_\theta^T [\mathbf{C} + \mu\sigma\mathbf{I}]^{-1} \mathbf{L}_\theta]^{-1} \mathbf{L}_\theta^T [\mathbf{C} + \mu\sigma\mathbf{I}]^{-1}, \quad (4.3)$$

where μ is a matrix regularization parameter set to $\mu = 4$ as this provides spatial smoothing and increases signal to noise ratio (Brookes et al. 2011a). σ is an estimate of uncorrelated noise power, and $\mathbf{I}$ is the identity matrix. The noise estimate σ was computed as the minimum singular value of the unregularized covariance matrix. The lead fields $\mathbf{L}_\theta$ were computed using a multi-sphere head model (Huang et al., 1999), and the forward model calculation described by Sarvas (1987). The angle δ was found independently at each $\mathbf{r}$ using a non-linear search to compute the orientation of maximum variance $\mathbf{Z}_{opt}$

$$\mathbf{Z}_{opt} = \max_\delta \left(\frac{\mathbf{w}_\theta^T \mathbf{C} \mathbf{w}_\theta}{\mathbf{w}_\theta^T \Sigma \mathbf{w}_\theta} \right), 0^\circ \leq \delta \leq 180^\circ, \quad (4.4)$$

where $\Sigma = \sigma\mathbf{I}$ represents white noise at each MEG sensor.

The magnetometers and gradiometers were scaled by their respective minimum eigenvalue. This effectively gives both classes of sensor equal noise levels. By sequentially applying the beamformer to a single dipole in the center of each of the 90 AAL specified regions of interest, a time course of the neuronal activity within each region was estimated. This source reconstruction using beamforming was only done for the empirical MEG data.

4.2.2 Signal orthogonalization

Investigations of functional connectivity between brain regions based on MEG have important practical problems (Hipp et al., 2012; Schoffelen and Gross, 2009). Due to the limited spatial resolution of MEG, individual brain sources project onto multiple sensors. Furthermore, the ill posed nature of source reconstruction (e.g. beamforming) means that the weights used to project sensor space data to different source locations will be correlated. This results in spurious correlations among brain regions that are an artefact of the source reconstruction (Hipp et al., 2012; Schoffelen and Gross, 2009). However, as the magnetic fields are measured

virtually instantaneously at all MEG sensors, sensor space signals having the same source components are reflected by an identical phase at different sensors (Nolte, 2004) whereas signals from different source regions can have a variable phase relationship. Recently, this information has been utilized to remove spurious correlations among brain regions (Brookes et al., 2011a; Hipp et al., 2012). In this method, the neural signals from *pairs* of brain areas are orthogonalized to remove the signal components that share the same phase before performing further analysis. For a pair of signals $\mathbf{x}$ and $\mathbf{y}$, this is achieved by subtracting the univariate projection of $\mathbf{y}$ on $\mathbf{x}$

$$\mathbf{x}_{\perp} = \mathbf{x} \perp \mathbf{y} = \mathbf{x} - \mathbf{y}(\mathbf{y}^+ \mathbf{x}), \quad (4.5)$$

where $+$ denotes the pseudo-inverse. Subsequent measures of functional connectivity such as band limited power (BLP) correlation computed from these orthogonalized signals, will be unaffected by trivial (zero phase lag) interactions. The signal orthogonalization was performed on both the empirical and simulated MEG data.

4.2.3 Biophysical network dynamics

We used the biophysical network model described in the previous chapter to simulate realistic resting-state MEG activity for a range of coupling strengths. The global scaling of the delay matrix (see previous chapter) was set to 15ms based on our findings that this results in emergent alpha and beta frequency oscillations. Similar values of delay scaling have been shown to reproduce functional connectivity patterns in fMRI and MEG (Deco et al., 2009; Cabral et al., 2014).

4.2.4 Network performance

To assess the performance of the biophysical model in predicting patterns of FC that match those observed in MEG, we computed the correlation between the FC scores estimated from the simulated and empirical data (see Figure 4.1).

As a measure of FC, we computed the correlation between the envelopes of band-limited power time courses, estimated via the Hilbert transform. This approach has previously been shown to provide a robust measure of FC in resting state MEG (Brookes et al., 2011a; Hipp et al., 2012; Luckhoo et al., 2012).

Following pairwise orthogonalization, the empirical and simulated time courses for the 90 brain regions were filtered into a number of different frequency bands, and the amplitude envelopes were computed and temporally averaged using different moving average window widths (see below for details of frequency bands and window widths). The FC scores were computed between all pairs of brain regions and the resulting data vectors were concatenated together. We then computed the Pearson correlation between the simulated and experimental MEG data vectors to assess the performance of model simulations. The correlation values were averaged across subjects giving an estimate of similarity at that coupling strength. For clarity, Figure 4.1 summarizes the analysis pipeline used for computing (and comparing) FC metrics from empirical and simulated data.

In the following, *rr* refers to the correlation between the *empirical* and *simulated* characteristics, and we always report the *least* significant p-value (*P*).

4.2.5 Frequency specific functional connectivity

Functional connectivity scores were computed from the amplitude envelopes for 15 different frequency bands (3-8Hz, 6-11Hz, 9-14Hz, 12-17Hz, 15-20Hz, 18-23Hz, 21-26Hz, 24-29Hz, 27-32Hz, 30-35Hz, 33-38Hz, 36-41Hz, 39-44Hz, 42-47Hz, 45-50Hz). For each given frequency band, the amplitude envelopes were estimated from the filtered signal via the absolute value of its analytical signal derived from the Hilbert transform. The analytical signal $\xi(t)$ of the time course $x(t)$ is a complex function of time defined as (Rosenblum et al., 1996)

$$\xi(t) = x(t) + ix_h(t) = a_\xi(t)e^{i\theta_\xi(t)}, \quad (4.6)$$

where $a_\xi(t)$ and $\theta_\xi(t)$ are the instantaneous amplitude and phase of the analytical signal $\xi(t)$ of $x(t)$. The instantaneous phase $\theta_x(t)$ of $x(t)$ is taken as $\theta_\xi(t)$. $x_h(t)$ is the Hilbert transform of $x(t)$

$$x_h(t) = \text{P.V.} \left[\frac{1}{\pi} \int_{-\infty}^{+\infty} \frac{x(\tau)}{t-\tau} d\tau \right], \quad (4.7)$$

where P.V. denotes the Cauchy principal value of the integral. Analogously the analytical signal is estimated from $y(t)$. The envelopes were temporally averaged (down-sampled) using 4 different moving average window widths (0.5s, 1s, 2s, 4s).




Figure 4.1 – MEG resting state functional connectivity analysis. MEG resting state data from healthy subjects is used to estimate source space activity at brain locations defined in the AAL template. The AAL template is also used to parcellate the DTI data into the same 90 brain locations, and the anatomical connectome between these regions is estimated using probabilistic tractography. The structural connectivity is used in combination with a neuronal model to simulate the activity of brain regions, where the dynamics of an individual brain region are described by a conductance based voltage model expressing the coupling between different populations of excitatory and inhibitory neurons. FC analysis based on amplitude envelope correlations is performed on the empirical and simulated time courses, and the metrics are compared with each other.

4.2.6 Haemodynamic model

To assess the BOLD signals generated by the biophysical model simulations, we used the Balloon–Windkessel haemodynamic model proposed by Friston et al. (2003). The model specifies the coupling of perfusion to the BOLD signal, with a dynamical model of the transduction of neural activity into perfusion changes. In this model, the BOLD signal is taken to be a static nonlinear function of normalized total deoxyhemoglobin voxel content, normalized venous volume, resting net oxygen extraction fraction by the capillary bed, and resting blood volume fraction.

In brief, for the i -th region of interest, neuronal activity z_i causes an increase in a vasodilatory signal s_i that is subject to autoregulatory feedback. Inflow f_i responds in proportion to this signal with concomitant changes in blood volume v_i and deoxyhemoglobin content q_i . The equations relating these biophysical variables are

$$\partial s_i(t)/\partial t = z_i - \kappa_i s_i - \gamma_i(f_i - 1), \quad (4.8)$$

$$\partial f_i(t)/\partial t = s_i, \quad (4.9)$$

$$\tau_i \partial v_i(t)/\partial t = f_i - v_i^{1/\alpha}, \quad (4.10)$$

$$\tau_i \partial q_i(t)/\partial t = f_i(1 - (1 - \rho_i)^{\frac{1}{f_i}})/\rho_i - v_i^{\frac{1}{\alpha}} q_i/v_i, \quad (4.11)$$

ρ being the resting oxygen extraction fraction.

The BOLD signal y_i is taken to be a static nonlinear function of volume and deoxyhemoglobin that comprises a volume-weighted sum of extra- and intravascular signals

$$y_i = V_o(7\rho_i(1 - q_i) + 2(1 - q_i/v_i) + 2(\rho_i - 0.2)(1 - v_i)), \quad (4.12)$$

where $V_o = 0.02$ is resting blood volume fraction. The remaining biophysical parameters were taken as in Friston et al. (2003).

The membrane potential across the pyramidal population was taken as the neural output of each region (z_i) and used to produce the final haemodynamic responses (y_i).

4.3 Results

4.3.1 Network performance

We simulated realistic resting-state MEG activity using the biophysical model for a range of coupling strengths, as in the previous chapter. Figure 4.2A shows the global synchrony between brain regions, and Figure 4.2B indicates the level of metastability at different coupling strengths. To quantitatively compare the simulated responses with experimental data, we identified the region of global coupling strength, k , under which the simulated network best predicts real resting state MEG functional connectivity. We computed the correlation of a vector of FC values concatenated over all pairs of brain areas and frequency bands, between the resting-state FC estimated from the simulated and the resting-state FC estimated from the experimental MEG data (see Figure 4.1).




Figure 4.2 – Network-level dynamics and model performance as a function of global coupling strengths, k . (A) The global synchrony of the network as the global coupling strength is varied. (B) The global metastability of the network as a function of global coupling. (C) The models performance at different coupling strengths as measured by the Pearson correlation coefficient of the frequency (and downsampling) dependent functional connectivity between the real and simulated data. The highest correlation is achieved at the critical coupling ($k_c \sim 1.8$) indicated by the dashed line.

Figure 4.2C shows the correlation coefficient obtained for the different coupling strengths. We find that the model optimally predicts the MEG responses in a range of coupling strengths that coincide with the critical coupling region of the network ($k_c \sim 1.8$, $rr \sim 0.4$, $P < 10^{-20}$). This regime is consistent with previous findings where similar dynamical clusters of nodes have been shown to reproduce functional connectivity patterns observed in FMRI and MEG (Deco et al., 2009; Cabral et al., 2011, 2014). We also find that the simulated responses show a similar power spectrum as the empirical data at the identified coupling strength (Figure 4.3). We used the simulated times series at the identified optimal coupling ($k_c \sim 1.8$) for all further analysis without any change in parameter values.




Figure 4.3 – Power spectra of the empirical and simulated data. Time courses of brain regions of the motor network are shown for the empirical (A) and simulated (C) data. The power spectrum of these networks is shown on the right (B, D). The spectrum of each network is computed by averaging over all pair of brain regions of the network. The values are averaged across all subjects in the case of empirical data.

4.3.2 Frequency specific functional connectivity

To demonstrate the ability of the biophysical network model to predict real MEG data, we examined the frequency dependency of FC scores computed from the real and simulated data. We computed FC (using correlation in amplitude envelopes) for different frequency bands and at different temporal down-sampling windows between all node pairs.

Frequency dependency. We first consider the case where the amplitude envelopes were temporally down-sampled by a window of 1s to identify the frequency bands showing the maximum functional connectivity between brain regions. We focus on the signals alpha and beta bands at this down-sampling window based on findings that it best captures the functional connectivity of RSNs in MEG data (Brookes et al., 2011b; Brookes et al., 2011a; Hipp et al., 2012; Luckhoo et al., 2012).

Figure 4.4A shows the mean functional connectivity of all pairwise connections as a function of different frequency bands. The number of pairwise connections showing maximal functional connectivity in a given frequency band are shown in Figure 4.4B. Both results are averaged over all subjects. Figures 4C-D show the same measures for the simulated responses at the identified optimal working point using the same analysis approach. We find that the model reproduces both the mean ($rr = 0.82$) and the maximal ($rr = 0.9$) frequency dependent functional connectivity of the empirical data. The overall functional connectivity scores are highest in the 10-30 Hz frequency range, with a peak at around 15-17 Hz, and most brain regions engage with each other in these frequency bands.




Figure 4.4 – Frequency specific functional connectivity of band limited neural amplitude signals. (A, C) Mean amplitude correlation in different frequency bands for real (A) data and simulated (C) data. (B, D) The frequency bands in which the pair-wise correlation between the regions is maximum is shown for real (B) and simulated (D) data.

Downsampling Dependency. Anatomically connected brain regions, such as homologous areas in the two hemispheres, exhibit significant correlations in spontaneous oscillatory activity (Cordes et al., 2001; Nir et al., 2008), and brain wide correlation analysis reveals that these regions express strongest correlations in the alpha and beta frequency bands (Hipp et al., 2012; Luckhoo et al., 2012). We, therefore, used the estimated frequency-specific functional connectivity scores at the 1s down-sampling window to pick the top (100) functionally connected brain regions. For these top node pairs, we spectrally resolved the envelope correlation to consider its temporal scale.

Figure 4.5A shows the amplitude envelope correlations for 4 different down-sampling windows. The full correlation profile for different down-sampling

windows is shown in Figure 4.5B. Results are averaged across all subjects. Figures 4.5C,D show the same results obtained from the simulated MEG responses. We find that these regions exhibit strongest correlations in alpha and beta frequency range as observed before ($rr = 0.93$). The full correlation profile reveals that the correlation in all frequency bands is significantly increased for high down-sampling windows (from 1s to 32s). Applying the temporal averaging reveals strong correlations on the time scale of several seconds ($<0.1\text{Hz}$). We, however, observe that simulated responses show some correlation at around 37Hz for high down-sampling windows (32s), and the correlation is also increased in the range of 10-25s with an additional peak at around 10Hz frequency band ($rr = 0.84$).




Figure 4.5 – Band limited power correlations between orthogonalized signals. (A, C) Correlation between the top functionally connected node-pairs in different frequency bands at 4 different moving average window widths for real (A) data and simulated (C) data. The standard error across node pairs is shown with shaded colour. (B, D) Correlation between the top functionally connected node pairs as a function of frequency band and moving average window widths. The values are averaged across all top node pairs (and across subjects for the real data).

4.3.3 Standard coherence

Since we are interested in the frequency-specific interactions between brain regions, we also estimated the statistical dependencies between regional time courses using standard coherence (computed from the cross power spectral density (CPSD), Figure 4.6).




Figure 4.6 – Estimates of standard coherence between functionally connected brain regions. (A-B) Standard coherence (A) and cross power spectral density (B) between the top functionally connected node pairs estimated from the empirical MEG data. The values are averaged over all the subjects. (C-D) Estimates of coherence (A) and CPSD (B) between the same brain regions from simulated data. We find that standard coherence fails to detect functional connectivity, which is captured by the amplitude envelope correlations.

We find that there is low global (over time) coherence (Figure 4.6A) between the regional time courses from the empirical data in all frequencies. However, we found strong functional connectivity between brain regions in the alpha and beta frequency bands, implying that standard coherence based methods

fail to detect this connectivity. Although the cross power spectral density (Figure 4.6B) shows more power in the low frequency bands, it may still be relatively low for robust estimation of correlated patterns. However, for the simulated responses (Figure 4.6C,D), we find that the estimates of coherence are much higher.

4.3.4 Haemodynamic responses

To investigate whether the simulated neural responses can reproduce characteristics of the fMRI resting state networks, we also calculated the BOLD-signal (Figure 4.7) using the Balloon–Windkessel hemodynamic (Friston et al., 2003).




Figure 4.7 – BOLD signal and frequency spectrum. (A) BOLD signal generated from the simulated neural activity of left-right motor cortex using the balloon model, together with the frequency spectrum. (B) BOLD signal for same regions (as in A) with the added red curve showing the synchrony between the BOLD signals in the two regions. The synchrony power spectrum is shown in the right plot. We find that the BOLD synchrony displays fluctuations in < 0.1 Hz as found in previous fMRI studies. These results also complement the ones based on the moving average width shown in the previous section.

Figure 4.7A shows the BOLD signal from a pair of regions at the optimal working point, together with the power spectrum. Figure 4.7B shows the power spectrum of the synchrony time course (red curve) given by the differences between the levels of BOLD signal in the two regions. We find that the synchrony time course reproduces the slow 0.1-Hz oscillations seen in previous fMRI BOLD studies.

4.4 Discussion

In this work, we investigated the spectral profile of static resting state functional connectivity in real and simulated MEG data. We used characteristics of dynamics from real data to tune the network parameters and identify the responses that best match the MEG data. By assessing brain-wide interactions in terms of band-limited amplitude correlations, we demonstrated that a realistic biophysical network model can reproduce frequency-specific functional connectivity patterns reported in previous MEG studies. In line with previous studies, the best correspondence between the simulated and empirical data was achieved when the biophysical model was connected with intermediate coupling strengths and transmission delays (Deco et al., 2009; Cabral et al., 2011, 2014).

An important finding from previous investigations of MEG functional connectivity is that functional connectivity is strongly dependent on frequency, with increased band-limited power correlations within the alpha and beta bands (Brookes et al., 2011a; Brookes et al., 2011b; Hipp et al., 2012). To compare results from real and simulated data with these studies, we computed stationary measures of functional connectivity based on the Hilbert amplitude envelope correlation. However, measures of MEG functional connectivity may be confounded by signal leakage between nearby voxels. We, therefore, orthogonalized the time courses to remove the possibility of FC being driven by signal leakage. This destroys any zero-lag correlations between the signals that

may give rise to artificially high FC. In line with previous studies we found that functionally connected brain regions at rest engage with each other in specific frequencies, with the strongest interactions in the alpha and beta frequency bands.

Applying temporal averaging to the band-limited power time courses revealed significant correlations on the time scale of several seconds in a spectrally specific way, ($<0.1\text{Hz}$), as observed in MEG (Hipp et al., 2012; Luckhoo et al., 2012) and fMRI (Cordes et al., 2001) studies of resting state FC. This behaviour was found not only in the real MEG data but also in data simulated from our biophysical model. Furthermore, using the haemodynamic balloon model, we also found that the network model can reproduce large-scale fluctuations on slow time scales, observed in previous fMRI studies (Beckmann et al., 2005). However, fMRI is based on an indirect link between the hemodynamic signal and the underlying neural activations, whereas our analysis based on the MEG oscillations, which reflect the biophysical properties of underlying networks, revealed similar correlations in the low frequency range providing further evidence in support of fMRI studies. The ability of the model to predict these correlations on slow timescales offers an integrated view about the underlying mechanisms that link the empirical MEG and fMRI responses.

In our simulations, we observed higher correlations in the gamma frequency range that were not seen in the real data. One possibility is that changes in the signal-to-noise ratio across different frequency bands may give rise to an increase in functional connectivity in the low frequency bands. Sensory stimulation tasks have been shown to drive cortical gamma band activity (Logothetis et al., 2001) that can be measured in EEG and MEG (Scheeringa et al., 2011). However, during the resting state the gamma band oscillations may be much weaker. Furthermore, MEG measures the combined activity of large populations of neurons and may be more sensitive to low frequency components with a global spatial coherence. We could, therefore, speculate that the lack of

gamma frequency correlations may be due to the poor signal-to-noise ratio of these weak oscillations.

High frequency neural activity also overlaps with the spectral bandwidth of muscle activity (~20–300 Hz) (Muthukumaraswamy, 2013), and removing muscle artefacts can, therefore, cause attenuation in the observed MEG responses in this range. This would give rise to low correlation values in MEG at higher frequencies than predicted by the underlying neural activity. However, if activity in this frequency band is not completely filtered, it would still remain difficult to assess if the increase in correlation in this band is due to an artefact of poor signal-to-noise ratio or genuine neural activity.

We found that standard measures of coherence fail to detect structured dependencies between functionally connected brain regions. One possible reason for this is that brain regions are more likely to exhibit transient coherence as opposed to remaining coherent for long time periods (Hadjipapas et al., 2004). In this work we estimated the coherence over the whole duration of signal (~300s), which would, therefore, result in low values. Estimates of cross power further revealed that there may be not enough signal power in these frequency bands for robust estimation of functional connectivity. On the other hand, our analysis based on band-limited amplitude envelope correlations revealed strong functional correlations that are consistent with previous findings. However, it has been suggested that these methods may be complementary representing different physiological processes (Brookes et al., 2011a).

Synchronizations among brain regions give rise to a novel form of frequency depression where networks oscillate at reduced collective frequencies. The collective frequency depends on the delays, couplings, and the number of regions in the network (Niebur et al., 1991). The heterogeneity of the couplings and delays enables multiple clusters to form while preventing the full

synchronization of the network. Each of these emergent clusters oscillates at their own collective frequency, giving rise to a range of frequencies. In a regime where the best fit with the empirical data is achieved, these frequencies fall in the alpha and beta frequency bands, explaining the observed correlation patterns. Recent experimental studies implicate a multi-layer feedback loop between the thalamus and the cortex in the generation of these slow frequency rhythms (Schreckenberger et al., 2004). Since this loop is also modelled in our network, it provides a unified network mechanism responsible for the generation of such rhythms.

Our model suggests an important role of delays in determining the collective frequencies and network dynamics. In our work, the delays were estimated from the Euclidean distances between regions. The histograms of spatial distances based on Euclidean distances reveal a gamma distribution with a single peak at approximately 50mm (see Figure 3.5). However, the histograms of actual fiber lengths from the connectome of Hagmann et al. (2008) reveal a bimodal distribution with peaks at approximately 40 and 120mm. The Euclidean distances, therefore, represent a lower bound on the physical connection length since they do not take into account the somewhat convoluted course of the white matter fibers.

The effects of delay distribution ($\sim$ fibre length distribution) on the dynamics of coupled oscillators has been recently addressed by Lee et al. (2009). They found that the critical coupling where metastable states occur depends both on the mean delay (scaling) and the distribution of delays. This suggests that heterogeneity in delays induces increased complexity in the network dynamics, especially in the critical regime. These results further highlight the important role of delays in shaping the dynamics of the network

Other than the limitations of the delays, the DCM model does not fully capture the internal anatomical and physiological detail of cortical regions. The

division of a single region into three cortical populations is somewhat arbitrary and does not capture the 4-6 layered structure of the real human cortex. As a result, connections within and among regions are also ill defined and further constrain the dynamics of the region.

5 Spatial Functional Connectivity

5.1 Introduction

A number of resting state studies have focused on the analysis of FC between segregated brain regions in a spatially specific way. These studies have shown that signals originating in different brain regions follow consistent correlation patterns across the whole brain, which are considerably consistent across sessions and subjects (Smith et al., 2009; Brookes et al., 2011b; Hipp et al., 2012). These networks often resemble the networks seen during the processing of specific tasks (e.g. motor network). Furthermore, these networks show strong frequency dependency in MEG, appearing mainly in the alpha and beta frequency bands.

In the previous chapter we focussed on the extent to which the spectral profile of FC matched globally between simulated and empirical data. In this chapter we now investigate whether the simulations are matching the empirical data in terms of the specific spatial structure. In particular, we look to see if the simulated data contains FC between specific brain areas that we expect to be functionally connected, and if the simulated data can produce canonical resting state networks observed in empirical MEG data.

To this aim, we analyze resting state MEG data from healthy subjects and biophysical model simulations and explore the plausible mechanisms underlying the generation of observed spatial FC patterns and network formation. Since the MEG source estimation can confound the spatial correlations between brain areas, we also apply this procedure to the simulated data by first projecting it onto sensor space and then re-estimating source space activity. We first look at the FC (amplitude envelope correlation) matrices that we used to identify the optimal simulated responses. We find that the empirical data exhibits different spatial

patterns of FC in different frequency bands, and the network model displays significant similarities with the empirical data. We then use a seed-based approach to investigate FC between specific brain areas. Using this approach we produce spatial correlation maps for two functional networks, while also computing the connectivity spectra for these networks. We finally apply ICA to simulated data, and show that the model can reproduce several RSNs whose spatial signature significantly matches the MEG RSNs identified in previous studies.

5.2 Methods

5.2.1 Biophysical network dynamics

We used the same realistic network model described in the previous chapters to simulate the activity of individual regions for a range of coupling strengths (5min). By comparing the amplitude envelope network matrices of the empirical and simulated MEG data, we identified the couplings ($k = 1.8$) where the simulated responses optimally match the empirical data (see Section 4.2.4 and Figure 4.2C). The network activity was then simulated for 60 minutes at the optimal coupling to allow the ICA to robustly estimate functional networks.

5.2.2 MEG forward model

As discussed earlier, the limited resolution of MEG and problems with source estimation can lead to spurious correlations between brain regions that do not reflect genuine connectivity. This will have a strong impact on the observed spatial pattern of correlations. Whereas in the previous chapter we were mainly concerned with the spectral profile of functional connectivity, here we are interested in its spatial profile. We, therefore, applied the same procedure to simulated responses by first computing MEG sensor space signals and then re-estimating source space activity.

We used the MEG forward model to project the simulated time courses onto sensor space. The forward model has three key elements: a source model defining the locations and orientations of the brain areas (dipoles); a sensor model that defines the locations and orientations of sensors; and a volume conductivity model that governs how sources project onto sensor channels.

The scalp fields are computed at a finite set of sensor locations and orientations (sensor configuration) for a given predefined set of source positions and orientations (source configuration). This is based on the lead field matrices which represent the field distribution across all sensor channels due to a unitary current dipole placed at a given position and orientation (i.e. how individual brain sources project onto all sensor channels). The magnetic field $\mathbf{b}$ observed at a sensor location $\mathbf{r}$ due to a dipole at location $\mathbf{r}_q$ is given by the following lead fields

$$\mathbf{b}(\mathbf{r}) = \mathbf{K}(\mathbf{r}, \mathbf{r}_q) \cdot \mathbf{q} \quad (5.1)$$

where $\mathbf{q}$ is the dipole moment of the source and $\mathbf{K}$ is the field kernel.




Figure 5.1 – MEG forward modelling and source estimation. (A) Location of MEG sensors (dots) and head model (grid). (B) Locations of dipoles in one plane, placed at the center of gravity of the AAL defined brain locations.

Here we used a single homogenous shell model as the volume conductivity model (Figure 5.1), which we had developed earlier as part of the network analysis methods framework (see next chapter). The source locations were taken

as the center of gravity of the AAL defined brain locations and the source orientations were computed normal to the surface. We used the 306-channel Elekta Neuromag scanner configuration. An additive subject-matched uncorrelated coloured noise was added to the MEG sensor signals prior to reconstructing the source space activity.

Once the sensor space activity had been estimated using the MEG forward model, it was projected back to the same AAL defined brain locations (source space) using a LCMV beamformer (in the same manner as the empirical data, see Section 4.2.1 for details).

5.2.3 Global correlation analysis

We computed the amplitude envelope correlation between all pairs of brain areas in different frequency bands to estimate the full FC matrices. For each pair, the source space signals were first orthogonalized using linear regression to remove zero-lag correlations induced by spatial leakage (Brookes et al., 2011a; Hipp et al., 2012). The orthogonalized time series were band-pass filtered into 15 different frequency bands (3-8Hz, 6-11Hz, 9-14Hz, 12-17Hz, 15-20Hz, 18-23Hz, 21-26Hz, 24-29Hz, 27-32Hz, 30-35Hz, 33-38Hz, 36-41Hz, 39-44Hz, 42-47Hz, 45-50Hz). The amplitude envelopes were estimated using the Hilbert transform, and temporally down-sampled to a resolution of 1s. We computed the pairwise correlations between the down-sampled envelopes of all regions, and the resulting FC matrices were averaged across subjects in the case of for empirical data.

To take into account artefact correlations generated by intrinsic properties of either the MEG scanner or the source-projection algorithm, the same correlation matrices were computed from an empty-scanner dataset projected onto the same AAL defined brain locations. An additive subject-matched uncorrelated Gaussian noise was added to this dataset prior to reconstructing the source space activity.

5.2.4 Seed based correlation analysis

We performed FC analysis of the source space activity in anatomically defined seed regions of interest to see if the simulations can produce FC between specific brain regions that we expect to be functionally connected. We considered two seed regions: right motor cortex to test for the motor network, and left visual cortex to test for the visual network. Here we compute the correlations between the down-sampled amplitude envelopes of the seed region with that of all other brain regions (in alpha, beta, gamma).

We generated correlation maps for these networks based on the correlations of the seed regions. We computed these maps by thresholding the beta band FC of each seed region while controlling for the false discovery rate (FDR) at $P < 0.05$ (Benjamini and Hochberg, 1995). This procedure consists of sorting the p-values for the voxels being tested in ascending order, and then dividing each observed p-value by its percentile rank to get an estimated FDR. In this way, small p-values that appear far down the sorted list will result in small FDR estimates, and vice versa. This procedure can be used to estimate a threshold value for a specified FDR. For this procedure, we used the implementation available at www-personal.umich.edu/~nichols/FDR.

We further investigated the contribution of different frequencies in the motor and visual networks. This time we computed the pairwise downsampled amplitude envelope correlations between all the nodes of the estimated (thresholded) networks in 15 different frequency bands described earlier. The FC scores were averaged over all the nodes (and subjects for empirical data) to get a final spectral profile of each network.

5.2.5 Independent component analysis

We applied temporal ICA to see if the simulations can produce canonical RSNs similar to those typically seen in MEG data. We used a similar approach as in

Brookes et al. (2011b) for the ICA analysis of the simulated responses, and we compare our estimated networks with those reported in Brookes et al. (2011b).

The simulated sensor space activity was reconstructed on a 7 mm grid spanning the entire brain. This has previously been shown to reproduce RSNs in MEG similar to those observed using fMRI (Brookes et al., 2011b). The source space signals were de-meaned and normalized by an estimation of projected noise. We filtered the data into beta frequency band; estimated the amplitude envelopes using the Hilbert transform; down-sampled the envelopes to a time resolution of 1s, and reduced the dimensionality of the data to 30 principal components (PCs) using Principal Component Analysis. We estimated 25 independent components (ICs) from the reduced dataset using fastICA (research.ics.tkk.fi/ica/fastica). The ICs originate from independent networks, and thus, each IC was taken as a separate RSN. To estimate the contribution of each region to an individual IC, we computed the Pearson correlation between the time course of the IC and the time course of each voxel. To get the final spatial maps at the regional level, we averaged the values of all the voxels belonging to specific brain regions.

To identify similar networks in empirical and simulated data, we computed the spatial correlation between the maps estimated from the simulated data and those reported in Brookes et al. (2011b). For each network pair, the three-dimensional spatial maps were vectorized and we computed the Pearson correlation between these data vectors. We compared all of the 25 simulated and empirical networks, and maps with the highest correlations were grouped together.

We assessed the significance of the spatial correlation using a Monte Carlo approach. The empty scanner data were projected on a 7mm grid using the same spatial filters as for the empirical data to preserve the artefact spatial correlation structure. The source space signals were filtered into beta frequency band, and

down-sampled amplitude envelopes (1s) were estimated yielding a null dataset. We then performed Monte Carlo simulations on the null data. On each iteration of the Monte Carlo algorithm, 25 ICs were estimated (yielding 25 spatial maps), and we computed the maximum spatial correlation between each of these 25 noise spatial maps and the spatial maps estimated from simulated MEG data. A total of 500 iterations were run, across which the data segments were shuffled in time, resulting in 500 maximal correlation coefficients (for each network) representing the best match with the simulation derived spatial maps. We estimated statistical values from the distribution of maximal spatial correlations, as this eliminates the need for multiple comparisons correction.

5.3 Results

5.3.1 MEG source reconstruction

To investigate the effects of MEG source reconstruction (forward modelling and beamforming) on simulated responses, we computed the spatial coherence and cross power spectral density (CPSD) between the top (100) functionally connected brain areas both before and after the MEG reconstruction (Figure 5.2). When computing these measures directly from the simulated responses without source reconstruction, we find high global coherence (Figure 5.2A) and cross power (Figure 5.2B) between the time courses of these regions in the low frequency bands. This is despite the pairwise orthogonalization to correct for spatial signal leakage. However, after the MEG source reconstruction, estimates of coherence between the same regions drop strongly (Figure 5.2C) and the CPSD (Figure 5.2D) also reveals the loss of cross power in the low frequency bands. Importantly, the band limited power correlation between the original and source reconstructed signals was more than 80%. Additionally, as we demonstrate later, these regions still display strong band limited power correlations in specific frequency bands.




Figure 5.2 – Effect of MEG source reconstruction on spatial coherence between functionally connected brain regions. (A-B) Standard coherence (A) and cross power spectral density (B) between the top functionally connected node pairs estimated from the orthogonalized simulated responses before MEG reconstruction. (C-D) Estimates of coherence (C) and CPSD (D) are shown for the same pairs after MEG forward modelling and beamforming. We find that the MEG source reconstruction strongly reduces estimates of coherence and cross power between functionally connected brain regions.

5.3.2 Global functional connectivity

We first look at the amplitude envelope correlation matrices for the empirical and simulated data in a set of different frequency bands (Figure 5.3C,D). The identification of the optimal simulated responses was based on the extent to which these FC matrices matched between empirical and simulated data. The FC matrix of null MEG data obtained with an empty scanner (Figure 5.3A) and the structural connectivity (SC) matrix (Figure 5.3B) are also shown.




Figure 5.3 – Global functional connectivity in resting state MEG data. (A) FC (BPL correlation) matrix of simulated null MEG data. Due to intrinsic limitations of the source projection algorithm, neighbour regions exhibit strong correlated activity. (B) Structural connectivity matrix derived from DTI showing the connection strengths between brain regions. (C-D) Frequency-specific functional connectivity matrices in a set of different frequency bands obtained from empirical data (A) and simulated data (B). The matrices were obtained by computing the pairwise amplitude envelope correlations between all pairs of brain areas. The empirical matrices are averaged over all the subjects. All matrices are ordered in a way such that corresponding contra-lateral regions are symmetrically arranged with respect to the matrix centre, the anti-diagonal revealing the existing connections between these contra-lateral regions. The red lines on the structural connectivity matrix show the separation of the two hemispheres.

At the optimal working point, there is a good agreement between the simulated and empirical FC matrices in all different frequency bands ($rr = 0.33:0.47$). Both the model and empirical MEG data display significant bilateral functional correlations in the alpha and beta frequency bands. The correlations are highest in the 9-17 Hz range, agreeing with previous studies, which have found similar FC peaks in the lower beta frequency range (Brookes et al., 2011a, 2011b; Hipp et al., 2012; Luckhoo et al., 2012). The null MEG data exhibits strong correlations between neighbouring brain regions that were found to be independent of carrier frequency. However, the agreement between the null MEG and simulated data was much weaker ($rr < 0.2$) as the model predicts long-range functional connections that are inexistent in the null MEG data.

5.3.3 Seed based functional connectivity

We computed seed based correlations to investigate patterns of functional connectivity at the regional level. We considered two seed locations: right motor cortex and left visual cortex. Figure 5.4 shows the anatomical and functional connectivity of these seed regions in alpha, beta and gamma frequency bands. We observe that despite the sparseness of the SC for both seeds (<10 non-zero connections), these regions exhibit strong functional correlations with all other brain regions, and the spatial profile of FC depends on frequency. Since we used full correlation to estimate these FC profiles, the observed FC shows the presence of both direct and indirect connections between brain regions. This suggests that most functional correlations between regions are mediated by indirect structural links between them, highlighting the role of the SC in shaping the FC profile.




Figure 5.4 – Seed based functional connectivity in different frequency bands. Horizontal bars indicate the extent of structural (labelled SC) and functional connectivity for empirical (labelled MEG) and simulated data (labelled MODEL). The top plot (A) shows the connectivity of the right motor cortex with all other brain areas and the bottom plot (B) shows this for the left visual cortex. The connectivity of the seed with itself is 1 and is marked in red. The structural connectivity is shown in black. Green, blue and yellow bars indicate the real (left) and modelled (right) envelope correlation for 8–12Hz, 13–30Hz and 30–70Hz frequency bands, respectively. Empirical scores are averaged over all subjects.

Despite the DTI missing many structural connections, the model reproduces some of the key features of FC for both seeds, but performs better for the visual seed ($rr = 0.35:0.57$) as compared to the motor seed ($rr = 0.20:0.45$). One reason for this could be the sparseness of the DTI derived SC for the motor cortex, which results in less structured FC profile of the motor cortex seed region. The visual cortex, however, retains a strong structure in the correlation patterns in the same frequency range. Note that this cannot be due to spatial leakage as the source space signals were orthogonalized to remove any zero-lag correlations.

5.3.4 Spatial correlation maps

In order to better explore the spatial profile of functional connectivity at the regional level, we generated the spatial correlations maps for these seed locations. Figure 5.5 shows the results for the motor and visual networks in the beta frequency band. In the case of empirical data, these maps look very similar to the spatial maps estimated in previous MEG studies (Brookes et al., 2011b). For the simulated data, the model failed to reproduce the bi-laterality of the motor network, which could be due to the beamforming or DTI missing certain tracts. For the visual network, the model displays strong bi-lateral correlations between visual areas similar to those observed in the empirical data.






Figure 5.5 – Whole-brain seed based correlation maps. The top plot shows the maps for the right motor cortex and the bottom plot shows this for the left visual cortex. The maps obtained from empirical MEG data are shown on the left (A,B) and the maps from simulated data are shown on the right (C,D). The seed based BLP correlation (FC) maps are obtained by thresholding the FC profile of the seed regions while controlling the false discovery rate. All maps were obtained in the beta frequency band. The empirical maps are averaged over all subjects.

5.3.5 Connectivity spectra for functional networks

We estimated the contribution of different frequencies in the observed functional connectivity in the motor and visual networks (Figure 5.6). In agreement with previous studies, we find high correlations between the sensorimotor areas in the 10-30Hz range peaking at about 8-15Hz. We further note that the FC in the motor network is mediated by both alpha and beta carrier frequencies, whereas the visual network is more focused on alpha band. The simulations reproduce this frequency dependency of both the motor ($rr = 0.76$) and visual ($rr = 0.73$) networks. These results are in agreement with previous studies that suggest similar connectivity spectra in these networks (Hipp et al., 2012).

A - Motor Network



B - Visual Network





Figure 5.6 – Connectivity spectra of the functional networks. Seed based functional connectivity is shown for the motor network (A) and visual network (B) in different frequency bands. Spatial maps on the left are beta band seed based correlation maps from the empirical (experimental) MEG data. The spectrum of each network is computed by averaging over all pair of brain regions of the network. The standard error across node pairs is shown with shaded colour. The values are averaged across all subjects in the case of empirical data.

5.3.6 Functional resting state networks

We generated spatial maps of the resting state networks estimated from the simulated responses using temporal ICA. We found that 4 of the estimated spatial maps had significant resemblance (correlation range, $rr = 0.36 - 0.57$; least significant p-value, $P < 0.03$) with the RSN maps derived using MEG data in Brookes et al. (2011b). These networks also look similar to the networks estimated from resting state fMRI data. Figure 5.7 shows the spatial maps for these different networks (left lateral frontoparietal network, right lateral frontoparietal network, frontal cingulate network, and the visual network). For these four networks, the model shows good similarity (and similar laterality) with previous empirical MEG and fMRI RSNs. We found all the networks in the beta frequency band.




Figure 5.7 – Resting state networks estimated from simulated neural responses using temporal ICA. Four different networks are shown with significant resemblance to those estimated in Brookes et al. (2011b). They include the frontal cingulate network (Row 1), visual network (Row 2), right lateral frontoparietal network (Row 3), and left lateral frontoparietal network (Row 4). The FMRI and MEG are taken from Brookes et al. (2011b). The FMRI maps (left) are thresholded at a correlation of 0.25. The empirical MEG maps (middle) are thresholded at a correlation of 0.3 apart from the left lateral frontoparietal network, which is thresholded at a correlation of 0.16. The model MEG networks (right) are thresholded at a correlation of 0.3.

5.4 Discussion

In this chapter, we investigated the spatial profile of functional connectivity in resting state MEG using data from healthy subjects and biophysical model simulations. We used the source space activity to derive the amplitude envelope time courses across frequency bands, and these envelopes served as the signals for our analysis. Since the MEG source reconstruction is known to confound estimates of functional connectivity, we projected the simulated MEG data onto sensor space

and then re-estimated the source space activity using the same beamforming technique that we used for empirical data. We also orthogonalized the time courses of brain regions to remove zero-lag correlations induced by spatial leakage.

We first estimated the full functional connectivity matrices in different frequency bands by computing the pairwise amplitude envelope correlations between all the brain regions. We found that the empirical MEG data exhibits different spatial patterns of FC in different frequency bands, with the strongest correlations in the alpha and beta frequency bands. The simulated data reproduced the key characteristics of the empirical data, showing strong bi-lateral functional correlations in the same frequency bands. Using seed based correlation analysis we generated spatial maps of functional connectivity for the motor and visual networks. We found strong correlations between specific brain regions that were highly consistent across subjects, and in agreement with previous studies of FC in these networks (Brookes et al., 2011b; Hipp et al., 2012). The model reproduced the spatial profile of FC for the visual network, but failed to reproduce the bi-laterality of the motor network. Analysis of the connectivity spectra of these networks revealed that interactions in these networks are mediated by different frequencies within the alpha and beta frequency bands, and the model reproduced this frequency dependency of both networks. Application of temporal ICA revealed the existence of several RSNs whose spatial signature significantly matched those identified in previous resting-state MEG studies (Brookes et al., 2011b). Importantly, however, the model failed to reproduce some of the main networks observed in previous fMRI and MEG studies.

The identification of optimal simulated responses based on frequency specific FC matrices is a function of both the spectral and the spatial profile of FC, i.e. the best fit will be achieved when not only the frequency specific behaviour is matched but also when the model has similar spatial FC patterns as the empirical

data. Therefore, at the optimal working point the correspondence between the spatial profiles of FC is the strongest.

An important concern in the assessment of similarity between the spatial FC characteristics of the empirical and simulated data is the artefact correlations generated by the intrinsic properties of the MEG scanner and limitations of the source reconstruction algorithm. We found that the null MEG data exhibits strong functional correlations between spatially neighbouring brain regions. Since we also performed the MEG source reconstruction on the simulated data, the simulated data will also contain some of the same artefact correlation structure present in the empirical data. Due to the strong local connectivity of the human brain connectome, most of these short-range connections also exhibit some degree of real functional connectivity and, therefore, these artefact correlations cannot be simply subtracted from the real and simulated MEG signals. This, however, will lead to biased estimates of spatial similarity, which are not driven by underlying brain activity. To account for these, we compared the simulated FC characteristics with those estimated from the null MEG data. We found that the correlations between the FC characteristics of the simulated and empirical data were much greater than compared to the null data. This was because the model reproduced the long-range functional connections appearing in real MEG data that are not existent in the null MEG data.

The lack of strong correlations between left/right motor cortex, as well as the inability of the model to reproduce canonical RSNs identified in previous studies could be the result of the limitations of beamforming or the DTI missing key connections between brain areas. The beamforming technique used to estimate source space activity has been widely used in previous studies of resting state functional connectivity (Brookes et al., 2011b; Gross et al., 2001; Schoffelen and Gross, 2009; Hadjipapas et al., 2004). An important limitation of beamforming is the suppression of sources whose activity has a high degree of temporal

coherence and failure to localize this activity (Wipf et al., 2010). In the case of empirical data this may not be a big problem as recent studies suggest that coherence between distant brain regions is likely to be transient, and that it must persist for a sufficiently long time period to affect the distribution of sources estimated by beamforming (Hadjipapas et al., 2004). However, for the simulated data, we found that the coherence was much higher, which might result in the failure of beamformer localization (Wipf et al., 2010).

The anatomical network plays an important role in shaping the spatial profile of FC. We found that most of the functional connections between brain regions are mediated by indirect structural links. The structural connectivity matrix, however, does not fully capture the brain connectivity due to the limited resolution of DTI, consistently missing deep brain regions and inter-hemispheric fibers (Hagmann et al., 2008). Resolving crossing fibers also remains difficult, despite recent advances designed to handle this problem (Behrens et al., 2007). Realistic estimates of bi-directional strengths and fibre lengths can also significantly alter the dynamics of the network, which would have considerable effects on the observed functional connectivity patterns. Delays are also crudely modelled in the network, and an appropriate measure of fibre lengths can further increase our understanding of these dynamics.

Despite these limitations, the network model reproduces the key characteristics of the spatial structure of functional connectivity. Our simulations suggest that the time-delayed interactions between brain regions happening in the heterogeneous structural connectivity of the brain may be at the genesis of the spectrally specific spatially structured correlations observed in the resting-state MEG data. On one hand, the heterogeneity of the structural connectivity enables the formation of multiple independent functional networks composed of specific brain regions. On the other hand, the heterogeneity of delays enables these emergent networks to interact at different collective frequencies. Therefore, when

multiple networks form, filtering the signals into different frequency bands can separate the dynamics of these networks. This would explain how the observed frequency dependent functional networks could emerge from the intrinsic dynamics of the brain networks.

6 Dynamic Functional Connectivity

6.1 Introduction

A more complete understanding of resting state FC may be achieved if we relax the assumption that patterns of co-activation are stable over time. Temporal variability in resting state FC has been demonstrated using sliding window techniques in FMRI (Chang and Glover, 2010; Allen et al., 2014). Recent investigations into time varying FC using MEG have found that low frequency amplitude envelopes (alpha, beta) of certain brain regions display periods of heightened synchronizations followed by periods of de-synchronizations, resulting in the formation of transient resting state networks (de Pasquale et al., 2010; de Brookes et al., 2014). Furthermore, it has been shown that these network exhibit temporal non-stationarity at faster time scales than previously observed using FMRI (Liu et al., 2010; Baker et al., 2014).

Recent computational modelling work has provided some insight into the mechanisms of frequency specific temporal synchronizations in resting state data. These models suggest that the presence of noise in weakly coupled networks enables the spontaneous switching between different brain networks, leading to dynamic changes in FC (Deco et al., 2009; Deco and Jirsa, 2012). However, most of these studies have been focused on FMRI or used stationary measures of FC.

In this work we investigate transient FC in resting state MEG by analyzing resting state MEG data from human subjects and biophysical model simulations. We use a sliding window power correlation method that is based on the wavelet transform. In this method, the source space signals are first decomposed into time-frequency representation using the wavelet transform; and then for each frequency band we compute the power correlation within a short time window,

resulting in a representation of the power correlation over frequency and time. We demonstrate that temporal non-stationarity is underpinned by bistate behavior, in which periods of low correlation (inactive state) are interspersed with periods of heightened correlation (active state). By computing summary statistics on the correlation distribution we characterize the extent of temporal non-stationarity in different frequency bands. We combine these measures with a null distribution derived from Monte Carlo simulations to assess the significance of temporal non-stationarity. We show that significant temporal non-stationarity in interregional FC occurs in the same brain areas and frequency bands as traditional measures of stationary FC. We also determine the minimum timescales at which significant transient events can be detected. Finally, we show that the results obtained using this method are in agreement with other static measures of FC.

6.2 Methods

6.2.1 Biophysical network dynamics

We used the biophysical model to simulate the dynamical interaction between brain regions at the optimal coupling strengths identified in the previous chapters ($k \sim 1.8$). Following this, the same analysis was performed on both the simulated and empirical data.

6.2.2 Dynamic functional connectivity

We computed the time varying functional connectivity between pairs of brain regions to assess the temporal dynamics of functional connectivity. Following source space reconstruction using beamforming (for empirical data), the signal power of a brain region is estimated as a function of time and frequency using the wavelet transform. We then compute the power amplitude correlations across multiple frequencies within a sliding window (see Figure 6.2).

6.2.3 Wavelet transform

The wavelet correlation is based on the continuous wavelet transform, which decomposes a single time series into time–frequency space by successively convolving it with scaled and translated versions of a wavelet function. The continuous wavelet transform of a time series x of length N , sampled from an underlying continuous waveform at equal time steps of size Δt , is defined as

$$W^x(s, \tau) = \sum_{t=1}^N x(t) \psi_s[(\tau - t)], \quad (6.1)$$

where t is a time index and s is the wavelet scale. The wavelet function ψ used here is the complex Morlet wavelet

$$\psi_s(t) = \pi^{-1/4} e^{i\omega_0 t} e^{-t^2/2}. \quad (6.2)$$

The Morlet wavelet may be considered as a sinusoid, localised in time by a Gaussian, and can therefore capture local oscillatory components in the time series. Furthermore, it provides the best trade-off between spatial and frequency resolution (Bernardino and Santos–Victor, 2005). The scale parameter s controls the time length of the wavelet and ω_0 governs the number of oscillations of the sinusoid within the wavelet. Therefore, the scale is inversely proportional to frequency for a given ω_0 . The parameter ω_0 controls the relative time and frequency resolution, and we set ω_0 at 6 cycles, which has been shown to provide a good trade-off between time and frequency resolution (Grinsted et al., 2004).

The wavelet transform is a complex quantity whose modulus expresses the amount of power in the signal as function of time at a given frequency. Application of the wavelet transform over different scales provides a time–frequency representation of the signal power envelope. The wavelet transform was chosen over other spectral estimation approaches as it is well suited to nonstationary data and combines spectral decomposition with power envelope

estimation (unlike the Hilbert transform, which first requires filtering of the signal into the desired frequency band).

6.2.4 Sliding window correlation

For a given frequency band, sliding window correlation between node power envelopes was computed to give an estimate of time varying functional connectivity at that frequency. Within a sliding window, $t - \frac{T_f}{2} < t \leq t + \frac{T_f}{2}$, the envelope correlation between two signals is defined as

$$C_{xy}(t, f) = \frac{\sum_{t-\frac{T_f}{2}}^{t+\frac{T_f}{2}} (P^X(t, f) - P^{X*}(t, f))(P^Y(t, f) - P^{Y*}(t, f))}{\sum_{t-\frac{T_f}{2}}^{t+\frac{T_f}{2}} (P^X(t, f) - P^{X*}(t, f))^2 \sum_{t-\frac{T_f}{2}}^{t+\frac{T_f}{2}} (P^Y(t, f) - P^{Y*}(t, f))^2}, \quad (6.3)$$

where $P^X(t, f) = |W^X(t, f)|$ is the mean-magnitude of the wavelet transform of signal $\mathbf{x}$ within the time window.

The size of the time window (T_f) is a function of frequency such that the window size in seconds is given by $T_f = \Delta/f$ where Δ is a constant that corresponds to a fixed number of oscillations at any frequency (Figure 6.1 for $\Delta = 50$).

We used a frequency dependent window size, rather than a fixed duration window, because the effective number of samples in the amplitude 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 controls for this uncertainty such that it is the same across all frequency bands. We used a window size parameter $\Delta = 50$ based on previous findings that MEG FC is nonstationary on a time scale of about 6-10s (de Pasquale et al., 2010).




Figure 6.1 – Window size as a function of frequency for $\Delta = 50$. The time window parameter Δ corresponds to the number of oscillatory cycles over which the power correlation is computed. Since the number of cycles is fixed for any frequency, the window size in seconds is inversely proportional to frequency.

6.2.5 Signal orthogonalization

To remove effects of zero-lag correlations we performed signal orthogonalization separately on the data from each separate sliding window (O'Neill et al., 2015). Since the pairwise orthogonalization can be performed in two directions ($\mathbf{x}$ from $\mathbf{y}$ or $\mathbf{y}$ from $\mathbf{x}$), we repeated the above procedure with the regression performed in reverse and the two correlation scores were averaged to give a single correlation score at each frequency.

6.2.6 Summary statistics

Application of the sliding window method to a pair of regions results in a time-frequency representation of the signal correlation between these regions. We computed summary statistics on the correlation scores to characterize the functional connectivity in regions and different frequency bands, as well as to estimate the extent of temporal nonstationarity.

For each frequency band we computed the standard deviation of the time varying wavelet correlation distribution as a measure of temporal variability in the FC time course. An increase in temporal variability could be driven by increased variance of a single continuous state, which would correspond to a unimodal distribution. Alternatively, increased temporal variability could be underpinned by a bimodal distribution, where one mode represents periods of low functional connectivity (with zero mean) and the other mode represents increased correlation during transient periods of high functional connectivity (see Figure 6.2). This would result in an asymmetric distribution of correlation scores due to the presence of the second mode with positive mean. The extent to which the distribution over time of the FC is uni- or bi- modal can be summarised robustly using skewness (note that this was found to be more robust than comparing Gaussian mixture models fit with either one or two Gaussians). An increase in positive skewness measure would indicate that the underlying data is not unimodal (normally distributed), but instead is bimodal. We calculated both the standard deviation and the skewness separately for each node pair and frequency band by collapsing the computed power correlations over the time dimension.

6.2.7 Significance testing

The limited resolution of MEG can lead to biased estimates of time varying functional connectivity due to the spurious functional connectivity arising from field spread as well as problems in source reconstruction using beamforming. Furthermore, variability in functional connectivity can also be driven by changes in signal to noise ratio or signal power differences in different frequency bands, as well as the high variance of underlying dynamical states. It is also necessary to know whether the variability in the time varying connectivity is significantly greater than which can be explained by the uncertainty in the estimated

correlation value. To account for all this, significance testing of summary statistics was performed against an assumed null hypothesis.

Monte Carlo simulations were used to surrogate realistic null data under the hypothesis that the relationship between the time courses was stationary (time-invariant). As the problem of spurious functional connectivity has been addressed in real data using signal orthogonalization, it was not necessary to account for this in the surrogate data. For each node pair and frequency band, we generated surrogate power envelopes using a phase randomisation technique (Prichard and Theiler, 1994). However, the phase randomisation generates simulated envelopes that follow a Gaussian distribution as opposed to strongly non-Gaussian distribution of the real power envelopes. In order to preserve this characteristic of the real data, we first applied a Box-Cox power transformation to the original envelopes $\mathbf{p}_x$ and $\mathbf{p}_y$

$$\mathbf{p}'_x = \begin{cases} (\mathbf{p}_x^\lambda - 1)/\lambda, & \lambda \neq 0 \\ \log \mathbf{p}_x, & \lambda = 0 \end{cases} \quad (6.4)$$

where the parameter λ determines the distribution of the data.

Phase randomisation of the transformed signal envelopes were performed to generate multiple realisations of amplitude envelopes (simulated envelopes $\widetilde{\mathbf{p}}'_x$ and $\widetilde{\mathbf{p}}'_y$) with no genuine correlation while persevering the spectral characteristics (power spectra and variance) and autocorrelation of the real data. However, this data still follows an approximately Gaussian distribution and, therefore, these envelopes were raised to the inverse of λ to ensure that the envelopes follow the non-Gaussian distribution of the real data

$$\widetilde{\mathbf{p}}_x = \widetilde{\mathbf{p}}'_x^{\frac{1}{\lambda}}. \quad (6.5)$$

To test whether FC between brain regions exhibit significantly greater temporal variability than would be estimated if the underlying association between the time courses were time-invariant, we supplemented the phase

randomisation process with a mixing procedure to ensure that the two surrogate signals retain the same static envelope correlation as the real data. This was achieved by linearly mixing the two surrogate envelopes $\widetilde{\mathbf{p}}_x$ and $\widetilde{\mathbf{p}}_y$

$$\widetilde{\mathbf{p}}_x = (1 - \alpha)\widetilde{\mathbf{p}}_x + \alpha\widetilde{\mathbf{p}}_y, \quad (6.6)$$

where α is a mixing coefficient that was set using an iterative scheme by minimizing the magnitude of the error between $\text{corr}(\mathbf{p}_x, \mathbf{p}_y)$ and $\text{corr}(\widetilde{\mathbf{p}}_x, \widetilde{\mathbf{p}}_y)$. This method ensures that the surrogate envelopes maintain the same static correlation as that of the real envelopes.

This procedure was used to generate 250 realisations of surrogate data, from which the same summary measures of functional connectivity were computed to obtain a null distribution of connectivity measures for each frequency band.

To test the significance at each frequency, the summary measures computed for the empirical (or simulated) data were compared against the null distribution. This was done by computing a Z score for each summary measure, by subtracting the mean and dividing by the standard deviation of the measure's null distribution. The Z scores were estimated for all subjects (and simulation) for the different frequency bands. The subject scores were averaged across all subjects to give a single set of values.

6.2.8 Effect of time window

The estimation of transient synchronizations using sliding window wavelet correlation depends on the chosen time window. For a given window length, random fluctuations in the time courses will result in non-zero correlation values, even if the signals are independent. Since the time varying correlation is computed over time windows of short duration, the uncertainty on the estimate will be relatively high. The choice of window size used will therefore have a strong effect

on the variance of the correlation values and subsequent summary statistics. The selection of an appropriate time window depends on the trade-off between the ability to detect short-lived transient synchronisation events and an increased sensitivity to noise. We use the Monte Carlo simulations described above to test the effect of the window size parameter Δ , by assessing the significance for different window sizes.




Figure 6.2 – Schematic of dynamic functional connectivity analysis. The time courses from a pair of brain regions are extracted and their power envelopes in a given frequency band are estimated using the wavelet transform. Within a sliding window, the correlation between the two power envelopes is estimated. Repeating this over the length of the time courses results in a power correlation time course. For the low frequency bands, the correlation time course is bistate, with periods of active (high FC) and inactive (low FC) states. This can be captured by the bimodal distribution of correlation values. At high frequencies, the correlation time course follows a Gaussian distribution. The non-Gaussianity of the correlation values can be captured by summary measures such as standard deviation or skewness.

It is important to mention that the results presented in this chapter were obtained before the ones presented in the previous (Spatial FC) chapter. It is for this reason that we continue to use the motor network as an exemplar in this

chapter despite previously showing that the model is unable to reproduce its spatial profile. This reordering has been performed since both the previous chapters focus on static aspects/measures of FC whereas this chapter centers on its dynamics aspects. Additionally, as shown in the following sections, despite the limitations of the model, it is still able to reproduce the key dynamic characteristics of functional networks to a high degree.

6.3 Results

6.3.1 Transient synchronizations and bimodality

Application of sliding window correlation reveals temporally and spectrally resolved changes in power correlation (shown for the left-right motor cortex for a single subject in Figure 6.3A). We find that correlation exhibits greater variability at low frequencies (10-30Hz), compared to the high frequencies (30-50Hz).

It is of interest to determine the nature of the underlying distribution of correlation values that gives rise to this variability. One hypothesis is that there are two underlying states where one state corresponds to periods of low FC and the other state corresponds to periods of high FC. Under this hypothesis we would expect a bimodal distribution, where the mean of one distribution (~ 0) corresponds to periods of inactivity, and the mean of other distribution (>0) indicates average FC strength.

Examination of the correlation distribution (Figure 6.3B) via Gaussian mixture modelling (GMM) reveals a bimodal distribution at a frequency of 18Hz. The individual states of this distribution are revealed by the components of GMM fit and are shown with dashed lines. We find that the heightened FC state occurs less frequently than the low FC state. At higher frequencies (35-50Hz), correlation time course have a distribution that is predominantly unimodal, suggesting a single stable underlying state. The standard deviation of the correlation time

course captures the heightened variability in correlation in the 8–30Hz frequency band, showing distinct peaks at both alpha and beta frequencies (Figure 6.3C).




Figure 6.3 – Wavelet power correlation for real MEG data. (A) Sliding window correlation between the left-right motor cortex for a single subject showing transient period of high FC (red patches) in the time-frequency plane. (B) Histogram of correlation values for the whole time period at a frequency of 18Hz and 40Hz. The coloured lines show the histograms while the black lines show the Gaussian mixture model fit. A two-component fit is revealed at 18Hz with the individual components shown in dashed lines. A bimodal distribution is displayed by frequencies exhibiting time varying correlations. (C) Z scores for the standard deviation of the wavelet power correlation at different frequencies, under the null hypothesis that the correlation is temporally stationary.

Data simulated from the network model for a representative node pair (left-right motor cortex) shows similar characteristics as the empirical data (Figure

6.4A), with an underlying bimodality for the same frequency bands (Figure 6.4B). However, we find that the model spends almost equal time in the active and inactive states. The increased variability in the low frequency bands is similarly captured by the standard deviation of the correlation time course (Figure 6.4C).




Figure 6.4 – Wavelet power correlation for the simulated data. (A) Correlation between a pair of brain regions (left-right motor cortex) showing transient periods of high correlation in the time-frequency plane. (B) Histogram of correlation values for the whole time period at a frequency of 18Hz and 40Hz. The coloured show the actual histograms while the black lines show the Gaussian mixture model fit. A bimodal distribution is revealed at 18 Hz with the individual components shown in dashed lines. A bimodal distribution is displayed by frequencies exhibiting time varying correlations. (C) Z-scores for the standard deviation of the wavelet power correlation at different frequencies, under the null hypothesis that the correlation is temporally stationary.

6.3.2 Summary measures of bimodality

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. We therefore also computed the skewness of the correlation time course as a robust summary measure of bimodality (Figure 6.5).




Figure 6.5 – Summary measures of bimodality for the top functionally connected brain regions. The significance of the summary statistics was assessed by comparing the results with the null distribution derived from Monte Carlo simulations. Left: Standard deviation of the wavelet power correlation at different frequencies for real (A) and simulated (C) data. Right: Skewness of the wavelet power correlation at different frequencies for real (B) and simulated (D) data. The standard error across node pairs is shown with shaded colour. The values are averaged across all top node pairs and also subjects in the case of real experimental data.

We find that the bistable nature of FC in the alpha and beta frequency range generalizes across multiple brain regions and subjects, by computing a mixed

effects group average (over 10 subjects) of the standard deviation and the positive skewness of the FC distributions for the top 100 node pairs (Figure 6.5A,B). The summary statistics from the simulated data display similar characteristics, indicating bimodality in the low frequency bands ($rr_{std} = 0.45$; $rr_{skew} = 0.63$). There is also a strong correspondence between the standard deviation and skewness in the all frequency bands, suggesting that increased temporal variability in FC arises out of a bimodal distribution of correlation values rather than due to a broadening of the distribution.

6.3.3 Dynamical interactions in functional networks

Recent research has found that different RSNs show distinct spectral profile of functional connectivity within the alpha and beta frequency bands. In particular, it has been shown that the motor network involves a mix of both alpha and beta frequencies, whereas the visual network relies more on the alpha band (Hipp et al., 2012). To investigate whether nonstationary FC in these networks exhibit similar frequency dependency, we computed the summary measures of bimodality between the nodes these networks (Figure 6.6). We find that these networks display significant bimodality (nonstationarity) in the same frequency bands as those found in previous MEG studies using stationary measures of FC. The simulated data reproduces the frequency dependency of nonstationary FC in both the motor ($rr_{motor} = 0.52:0.75$) and visual networks ($rr_{visual} = 0.31:0.53$).




Figure 6.6 – Dynamical interactions in functional networks. The summary measures of bimodality are shown for the motor (top) and visual (bottom) networks. The Z score of standard deviation is shown on the left (A,C) and Z score of skewness is shown on the right (B,D). The blue plots show the results for the empirical data whereas the grey plots show the results for the simulated data. The standard error across node pairs is shown with shaded colour. The values are averaged across all the nodes of the network and also subjects in the case of empirical data.

6.3.4 Effect of time window

Figure 6.7 shows the effect of the window size parameter Δ on the Z score of standard deviation for the motor and visual networks ($rr_{motor} = 0.54$; $rr_{visual} = 0.42$). The window size that best captures an increase in standard deviation is in the range 50 – 250 cycles. For the alpha and beta frequency bands (8-30 Hz), this corresponds to a time window of 6-12s. This is in agreement with previous work

by de Pasquale et al. (2010) showing that amplitude envelope correlations in the default mode and dorsal attention networks are nonstationary on a time scale of about 10 seconds. Below this range, it becomes increasingly difficult to differentiate the variability resulting from the transient increases in FC from that caused by random fluctuations in the power of both signals.




Figure 6.7 – Effect of wavelet time window parameter Δ on detection of transient functional connectivity. The effect of parameter Δ on the Z score of the standard deviation is shown for the motor (top) and visual (bottom) networks. The results for the empirical data (averaged over subjects) are shown on the left (A, C) and that for the simulated data are shown on the right (B,D). 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 standard 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.

6.3.5 Static vs. dynamic measures of functional connectivity

In the previous chapters, we established that the amplitude (Hilbert) envelope correlation could be used to estimate stationary functional connectivity across frequency bands. Figure 6.8 compares estimates of functional connectivity based on Hilbert envelope correlation (from chapter 4,5) with the time varying measures of connectivity (described in this chapter), as a function of frequency.




Figure 6.8 – Static vs. dynamic measures of functional connectivity. FC between the nodes of the motor network shown for a single subject (A) and simulated (B) data. The green lines show estimates of stationary functional connectivity using the Hilbert envelope correlation. The blue lines show the Z scores of the standard deviation of the sliding window power correlation. There is a strong correspondence between the scores estimated using the different approaches. The sliding window also examines the transient nature of FC not assessed by the static measure. The static method can be thought of a special case of sliding window correlation where the window is over the whole length of the data.

For the temporally stationary envelope correlation, maximum functional connectivity is observed in the alpha and beta bands, consistent with previous findings. The standard deviation of the time varying envelope correlation shows strong agreement with the Hilbert envelope correlations, with prominent peaks in the alpha and beta frequency bands. However, this measure also shows that functional connectivity in these frequencies is significantly nonstationary, due to transient synchronizations. In addition to the standard deviation of the wavelet correlation, the mean of the correlation distribution can be used to detect stationary functional connectivity at shorter time scales than the one-second envelope down sampling used in the Hilbert envelope correlations.

6.4 Discussion

We explored the dynamic behaviour of resting state functional connectivity in real and simulated MEG data. We assessed the time varying nature of resting state FC using sliding window envelope correlation across different frequency bands. We characterised temporal variability in these FC time courses using the standard deviation and skewness as summary measures of this variability. Crucially, by comparing these measures to those computed from realistic null data with imposed stationarity, we found significant variability in resting state FC in alpha and beta frequency bands. We demonstrated that the observed temporal variability arises out of a skewed distribution of correlation scores, suggesting that resting state FC is driven by short-lived periods of high within-network connectivity.

Our results support the hypothesis that brain regions engage with each other in a transient way, with significant temporal variability in the same frequency bands as have been previously shown to exhibit increased FC using stationary measures. Interestingly these non-stationary properties were also evident in data simulated by the biophysical model, suggesting that this behaviour

is an inherent property of weakly coupled dynamical oscillators. 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).

Recent studies have shown qualitatively that resting state MEG functional connectivity is temporally non-stationary (de Pasquale et al., 2010; de Pasquale et al., 2012; Baker, et al., 2014). Here we have attempted to quantify this temporal variability by using the standard deviation of the FC time course as a summary measure of variability. One concern is that changes in FC may be driven by differences in the power and/or temporal autocorrelation structure of the signals of interest. Moreover one would expect a certain degree of variability in pairwise FC to exist as a result of random fluctuations in the two analyzed time series. When using sliding window correlation, this variability will scale inversely with the length of the window used, due to the reduction in available time points with which to compute the correlation. In order to appropriately interpret the observed variability we performed significance testing against data with the same spectral characteristics as the real data. In addition these surrogate data were generated as pairs of data with the same stationary envelope correlation as that of the real data. This enables us to assess whether the observed FC was significantly greater than that which would be expected if the coupling was stationary about a given mean i.e. that the variation cannot be solely explained by the uncertainty in the estimated correlation value. This analysis revealed transient increases in functional connectivity between brain regions in the low frequency range as observed in previous studies.

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 synchronization events. To discriminate these scenarios we assessed the

nature of the distribution of correlation values across different frequency bands. As demonstrated, transient synchronisations at alpha and beta frequencies manifest as two distinct states in the distribution, where one state corresponds to an active state in which transient events take place, while the other state represents an inactive state driven by noise. We found that this bimodal distribution is lost outside of these frequency bands, both in the real and simulated data. To provide a quantitative description of this bimodal behavior, we computed the skewness of the correlation distribution. This revealed that this bimodal behaviour generalizes across functionally connected brain regions, supporting the hypotheses that resting state functional connectivity between brain regions is driven by periods of transient synchronizations.

As with the temporally stationary measures of FC, the results from the network model closely matched those of the real data. In a regime of high metastability, transient events between brain regions result in the formation of dynamical networks, which oscillate at reduced collective frequencies (alpha, beta). Noise and competitive processes in the network cause jumps in the synchronized states resulting in the formation of new networks, and giving rise to observed nonstationary functional connectivity. Recent studies similarly suggest that traditional resting state networks are time-averaged manifestations of more meaningful transiently synchronizing networks where the components of the network can overlap spatially (Smith et al., 2012).

Temporal nonstationarity may play a role in mediating network interactions by allowing particular nodes to participate in multiple networks in different frequency bands or switch between different networks based on their degree of synchrony (de Pasquale et al., 2012). These results lead to the hypothesis that multiple networks consisting of the same regions are active during rest. Stimuli or other internal events may drive the brain into a single stable network for task processing by increasing the synchrony level between certain brain

regions. This may allow for the brain to quickly and optimally process information by recruiting a large number of specialized processors distributed across the brain, and to quickly adapt to new situations by re-arranging the spatio-temporal synchronizations among the regions.

7 Network Analysis Methods

7.1 Introduction

An area gaining increasing interest is the mapping of functional brain networks using functional neuroimaging data. This is typically referred to as effective connectivity, and aims to estimate the causal structure between brain regions (Friston et al., 1993a). This investigation requires three levels of information: identifying the brain regions (nodes) of the network, the connections (edges) between these regions, and the direction of information flow between them (directionality) (Friston et al., 1993a).

Nodes are typically defined as spatial regions observed during some task-based activity, or derived from brain parcellation schemes (Hagmann et al., 2008). Other methods include analysis of activation maps (Huettel et al., 2004) or correlation maps (Biswal et al., 1995), as well as methods based on principal or independent component analysis (Friston et al., 1993b; Arfanakis et al., 2000).

Once the nodes of the network have been defined, the aim is to estimate the connections between these nodes. Simple methods such as correlation can show statistical relationships between the time courses of brain regions but fail to make any distinction between direct or indirect connections. However, regressing the time course of one region from another before computing the correlation (i.e. partial correlation) can indicate the presence of indirect connections (Marallec et al., 2006). This distinction between simple correlation and causal structure is crucial if one aims to estimate the underlying structure of the network.

Another property of interest in network estimation is the direction of the information flow. Methods that attempt to estimate this fall into three broad categories. The first are lag-based methods, such as Granger causality (Granger,

1969), that predict if the patterns in the time course of one region are approximately repeated in the time course of another region after a time lag. The second class of methods is based on conditional independence, and includes Bayesian network methods (Ramsey et al., 2010). The final class of methods is based on statistics of higher order than just the covariance, and include Patel's methods (Patel et al., 2006).

In a recent study (which we refer to as *netsims*) a number of network estimation methods were applied to simulated FMRI resting state datasets (Smith et al., 2011). It was found that certain methods provide better estimates of the network structure under certain network configurations, longer datasets, or in different confounding conditions. However, these simulations were based on an abstract model of neural dynamics having no direct biological relevance.

In this work, we investigate the performance of many different network estimation methods on multiple simulated resting state FMRI and MEG datasets. Here we use a biophysical network model to simulate realistic resting state neural activity for a range of network scenarios. For each network setup, we derive FMRI responses from this activity using the haemodynamic model (Friston et al., 2003). As in Smith et al. (2011) we apply different methods to the simulated FMRI data to estimate the underlying networks. By comparing the estimated networks with the true networks we calculate various performance scores such as the sensitivity and robustness of the method for each network scenario. We find that some of the simplest methods (such as correlation) provide the best results, and more complicated methods can easily give erroneous results. We also compare our results with those reported in Smith et al. (2011), and find that the rich neural dynamics offered by the more biophysically based model had no significant effect on the performance of the different methods. Applying the methods directly on simulated MEG data we find that the biggest gain in performance is achieved by all the lag-based methods.

7.2 Methods: network dynamics

The network simulations were based on a combination of the DCM model and the haemodynamic model described in chapter 3. As in chapter 3, for each network setup, the DCM model is first used in conjunction with an underlying anatomical network to simulate the neural activity of the individual regions. This neural activity is then transformed into FMRI space (BOLD signal) using the Balloon–Windkessel haemodynamic model. We used a default set of parameters in general, except when the effect of a certain parameter was being tested.

7.2.1 Network connectivity

As in *netsims*, we tested several different network topologies in our simulations. We used a 5-node network (skeleton) as default for most simulations (subnet 1 in Figure 7.1), and formed bigger networks by connecting more of them together (such that the network size was always a multiple of 5). The default 5-node network is an acyclic ring (i.e. all the arrows follow a cyclic path except one which is reversed). The 10-node network was formed by connecting two 5-node networks together using only one external connection going from sub-network 1 to sub-network 2. Similarly, connecting three 5-node networks together, as shown in Figure 7.1, formed the 15-node network. This connectivity reflects the small-world topology of real brain networks, and larger networks had similar small-world topology.




Figure 7.1 – Connectivity for the directed 15-node network. The network consists of 3 sub-networks connected to each other. Each sub-network is composed on 5 nodes connected in an acyclic ring. The corresponding connection matrix for the 15-node network is also shown where the 3 sub-networks are also highlighted.

7.2.2 DCM connection structure

The neural signals propagate around the network as defined by the three types of extrinsic connections in the DCM model (see Figure 3.3). The FMRI model tested in *netsims* had a single connection matrix containing all the different types of connections (Figure 7.2A). In our simulations, we used these different types of connections in the following way (Figure 7.2B-D). The forward connections were used with the default anatomical connectivity profile (skeleton) for all the simulations. For the case where we test the effect of symmetric networks we utilized the backward connections. These connections project to both the pyramidal population and the inter-neurons of another population, reflecting the hierarchal connectivity structure observed in human cortex. We used the lateral

connection type when we test the effect of dense networks, which contained any connection other than the default or symmetric network. These connections cause additional inhibition of network activity, and we utilized these connections to test how this would affect the estimation of network structure.




Figure 7.2 – DCM connection structure used in the simulations. (A) The FMRI model tested in *netsims* had a single connection matrix containing all the different types of connections. (B-D) The model used in this work has three different types of extrinsic connections. The forward connection matrix shown (for the 5 node network) is the default profile (B). Backward and lateral connection matrices are only used in certain cases. The backward connection matrix is used if we are testing the presence of backward connections (C). In this case the backward connection is always the transpose of the forward connection matrix. Anything other than forward (default) or backward is added to the lateral connection matrix (D), such as when testing the effect of more connections. In this case only the extra connections are added to the lateral connection matrix.

7.2.3 Neural dynamics

An external input was fed into each node. We tested two different types of input in our simulations. As in *netsims*, we used a binary (active/inactive) input based a Poisson process, which controls the likelihood of switching between the two states (Figure 7.3A). Neural noise was added to this input (std. = 1/20 of the difference in

height between the two states). The average durations of the states were 2.5 s (active) and 10 s (inactive), reflecting the longer average rest periods than active ones. We found that the bistate inputs do not produce any prominent oscillatory activity, whereas, as shown in previous chapters, nodes driven by noise can lead to emergent resting state oscillations in the network with characteristics similar to the empirical MEG data. We, therefore, also tested the effect of using Gaussian white noise as inputs. These inputs can be viewed as uncorrelated neural noise at each node, an external stimulus feeding directly into each node, or inputs from other brain regions.

The neural activity was simulated for a range of coupling strengths and we identified the coupling strengths that give rise to transient RSNs. This was done by calculating the synchronization dynamics for the particular network, which are related to the global time varying coherence between nodes, as described in Section 3.2.5.

Figure 7.3B shows example simulated neural activity using the DCM model for the case of the 5-node network (time courses for only 3 nodes are shown). The neural output of each node was taken as the membrane potential across the pyramidal cells, since these cells project to other brain regions (nodes) whereas the inhibitory interneurons and spiny stellate cells connect locally. For comparison, the neural output from the bilinear model used in *netsims* is also shown. We explicitly specified a neural lag of 50 ms between directly connected nodes. Although this is more towards the upper end of the neural lags between distant brain sources, this value was used to give the lag-based methods a better chance of estimation while still remaining realistic about the network.




Figure 7.3 – Networks dynamics for 3 brain regions. (A) Binary input to each node showing periods of activity and rest. (B) Output neural activity for the three nodes simulated using the DCM model. The simulated output from the biophysical model used here is shown in green whereas yellow shows the output from the bilinear model used in *netsims* (C) BOLD activity for the three nodes estimated using the nonlinear balloon model. The BOLD responses for the biophysical model and bilinear model are shown for comparison.

7.2.4 Haemodynamic responses

We used the nonlinear balloon model (Friston et al., 2003) to simulate the vascular dynamics from the neural activity. The membrane potential of the pyramidal population, taken as the neural output of the node, was used to produce the final haemodynamic responses (Figure 7.3C, see Section 4.2.6). As in *netsims*, we set the amplitudes of the neural time courses such that the amount of nonlinearity matched signals observed in a 3T FMRI data, with BOLD % signal of about 4%

relative to the mean. We find subtle differences in the haemodynamic time courses from the two models within this limited signal that result from the greater complexity of the more biophysical model.

The haemodynamic responses are known to vary between brain regions and subjects, resulting in different lags between neural activity and the BOLD responses (Handwerker et al., 2004; Chang et al., 2008). We therefore used a distribution of values for the haemodynamic model parameters for each node (Figure 7.4).




Figure 7.4 – Haemodynamic response variation. (A) Impulse showing the timing of the input signal to indicate the lag in the haemodynamic responses. (B) Prior distribution on the different parameters of the balloon model. The parameters are signal decay (kappa), auto-regulation (gamma), transit time (tau), alpha, resting oxygen extraction (E0), epsilon, resting venous volume (efficacy). (C) Variation in haemodynamic responses for different values of parameters sampled from the prior distribution. The average response is overlaid in thick black line.

We found that some of the parameters of the haemodynamic model have a stronger effect on the shape of the haemodynamic response than others. For example, we found that the transit time affects both the time-to-peak and undershoot. Variations in the transit-time can change the level of blurring which alters the performance of the different methods. We varied this parameter for each brain node resulting in variations in the haemodynamic response function delay (std. = 0.5s). The distribution over all parameters is shown in Figure 7.4B and the variation in responses to an impulse input are shown in Figure 7.4C.

7.2.5 Network simulations

We performed various simulations to test the estimation methods for a range of different network scenarios. In total, we had 30 different scenarios with the network specifications for these shown in Table 7.1. The first 28 of these scenarios are the same as those tested in *netsims*. More details of the motivation for testing a specific parameter and its relevance in FMRI settings can be found in Smith et al. (2011). In addition to these, we also tested the effect of white noise inputs (Sim29), and the direct application of the network analysis methods on simulated neural time courses (MEG) (Sim30).

For each network scenario, we had 50 simulations (representing trials/subjects) and all had the same parameter values except those that had variation added by default, such as the external input and the haemodynamic parameters for each node. This aspect of our simulation framework is somewhat similar to greedy equivalence search (GES) method in Ramsey et al. (2010). However, in that work the different realizations were used for cross-subject network modelling, whereas we are interested in average (single subject) results and utilize the multiple subjects to characterize the variability of the methods performance. As shown in Table 7.1, we simulated 10-min resting state FMRI data for each subject in most simulations, and the data were sampled with a TR of 3s

except in simulations 19 and 20. We also added white noise to the time course of each node (std. = 0.1–1% of mean signal level).

Table 7.1 – Specifications for all network scenarios tested.

Sim	No. nodes	Session length (min)	Noise	Other factors
1	5	10	1.0	
2	10	10	1.0	
3	15	10	1.0	
4	50	10	1.0	
5	5	60	1.0	
6	10	60	1.0	
7	5	250	1.0	
8	5	10	1.0	Shared inputs
9	5	250	1.0	Shared inputs
10	5	10	1.0	Global mean confound
11	10	10	1.0	Bad ROIs (mixed signals)
12	10	10	1.0	Bad ROIs (added random signals)
13	5	10	1.0	Backward connections
14	5	10	1.0	Cyclic connections
15	5	10	0.1	Stronger connections
16	5	10	1.0	More connections
17	10	10	0.1	
18	5	10	1.0	No HRF variability
19	5	10	0.1	TR reduced to 0.25
20	5	10	0.1	TR reduced to 0.25, No HRF variability
21	5	10	1.0	2-group test
22	5	10	0.1	Nonstationary connection strengths
23	5	10	0.1	Stationary connection strengths
24	5	10	0.1	Only one strong external input
25	5	5	1.0	
26	5	2.5	1.0	
27	5	2.5	0.1	
28	5	5	0.1	
29	5	10	1.0	White noise only inputs
30	5	10	0	Neural time series

As in *netsims*, for each method we estimated the false-positive and false-negative rates from the distribution of networks estimated for the 50 subjects. We averaged the results over all subjects to derive the final estimation scores. To make the results qualitatively comparable across the range of methods, we converted the

estimated connection strengths into Z scores by using an empirically derived null distribution of connection strengths. We estimated the null distribution of connection strengths by feeding null time courses into each method. These time courses were created by testing for connections between time series from different subjects, which have no relations between them. The raw connection strengths were converted into Z scores by using the cumulative density function of an analytical approximation to the estimated null distribution. Figure 7.5 shows the schematic summarizing the characterization of methods performance from network simulations.




Figure 7.5 – Schematic of the network estimation pipeline. For each network setup, an underlying anatomical structural network is used in combination with DCM model for node dynamics. The simulated neural activity of the individual regions is converted to BOLD signals using the haemodynamic balloon model. Different methods are applied to the haemodynamic responses to estimate the properties of the underlying network. The performance is characterized by comparing the true and estimated networks.

7.3 Methods: network analysis

In this section we provide a brief description of each of methods we tested. While we tested several variants of the main methods, including alternative choice of parameters, we only report those variants and parameters that provided the best results. These are the same suite of network analysis methods that were tested with the much simpler network model in *netsims*.

7.3.1 Correlation

The Pearson correlation measures the linear dependence in the time domain.

Full correlation. We computed the full correlations to assess the pairwise similarity between time courses.

Bandpass correlation. We evaluated the full correlation after band pass filtering the fMRI data. In the first case we filtered the data into two equal frequency bands and in the second case we divided the data into eight equal frequency bands. We report the results for the lower half of the two frequency bands (bandpass1/2) and the second lowest of the eight frequency bands (bandpass2/8).

Partial correlation. We computed the correlation between two time courses after regressing out all other time courses (from other nodes) in the data.

7.3.2 Coherence

The coherence measures the linear correlation (at fixed, but potentially non-zero, lag) between two signals as a function of frequency.

Sliding window wavelet coherence (Coherence A). We computed the coherence between two signals as a function of both time and frequency using the wavelet coherence toolbox (Grinsted et al., 2004) (www.pol.ac.uk/home/research/waveletcoherence). This is similar to the sliding window wavelet

correlation method we used to investigate nonstationary effects in MEG data expect that we measured the correlation within a time window. We averaged the results over all estimated time windows.

Fixed window wavelet coherence (Coherence B). We also computed the coherence using estimates of the cross spectral density using the functional connectivity toolbox (Zhou et al., 2009) (groups.google.com/group/fc-toolbox). In this case we set the time window for spectral estimation to be equal to the length of the time course.

We estimated the average coherence (over time) for a range of different frequencies but only report the results for average over all frequencies as they provided the best estimates. These are referred to as Coherence A3 (sliding) and Coherence B3 (fixed) in our results.

7.3.3 Covariance

The covariance is the simplest measure of pairwise similarity between two signals.

Regularised inverse covariance (ICOV). The elements of inverse of the covariance matrix (also called precision matrix) can be interpreted in terms of partial correlations and partial variances (Marrelec et al., 2006), and therefore show the relation of nodes with their neighbours. Regularization can be further applied to this matrix (under the assumption of sparsity) (Banerjee et al., 2006; Friedman et al., 2007) to eliminate elements that are close to zero providing better results. We computed the inverse covariance using L1precision toolbox (www.cs.ubc.ca/~schmidtm/Software/L1precision.html) for different regularisation values and report the results where this is set to 5 (ICOV $\lambda = 5$) and 100 (ICOV $\lambda = 100$).

7.3.4 Mutual information

Mutual information (MI) is an information-theoretic measure that quantifies the amount of shared information between two signals (Shannon, 1948).

Mutual information (MI). We computed the mutual information to measure linear and nonlinear dependencies using the functional connectivity toolbox mentioned above.

Partial mutual information (Partial MI). We also computed the partial mutual information between two time courses after regressing out all other time courses (from other nodes) in the data.

7.3.5 Granger causality

Granger causality (Granger, 1969) is a lag-based method that predicts whether patterns in one signal are approximately repeated in another signal after a certain time lag, and is based on the multivariate vector autoregressive modelling (MVAR) (Asteriou and Hall, 2011). We tested the following different variants of the Granger causality based methods. In all methods, we tested five different model orders (1, 2, 3, 10 and 20) but only report results for the ones that performed the best. Except the first two, all methods were implemented using the BioSig toolbox (biosig.sourceforge.net). Additionally, all methods were performed pairwise except the first.

Conditional Granger causality (Granger An). In this method, one variable is assumed to cause another, if the residual error of the first is reduced by incorporating the second variable (Geweke, 1984; Guo et al., 2008). For this we used the causal connectivity analysis toolbox (www.anilseth.com).

Bayesian Granger causality (Granger Bn). In this implementation, a Bayesian information criterion is used to estimate the model order up to a

specified maximum (www.mathworks.co.kr/matlabcentral/fileexchange/25467-granger-causality-test).

Directed Granger causality (DC). This method computes the directed GC.

Geweke's Granger causality (GGC). This method can be restricted to specific frequencies. We found that the highest frequencies give the best results and only report these results for different model orders.

Partial directed coherence (PDC). This method estimates causality in the frequency domain (Baccalá and Sameshima, 2001). For this we found that all frequencies give very similar results and only report one set for different model orders.

Directed transfer function (DTF). This is another frequency domain method, which also enables to identify the different propagation of different frequencies (Kamiński et al., 1997). For this we also found that all frequencies give very similar results and only report one set for each model order.

7.3.6 Generalized synchronization

Generalized synchronization is based on the concept of a state-space (Dauwels et al., 2010). If two signals are related, then if temporal patterns in a one signal are similar at certain times then this would also be true for the other signal, and this can be quantified using different state-space indexes. Here, we computed the S-index and the H-index (Arnhold, 1999) using the implementation available at: www.vis.caltech.edu/~rodri/programs/synchro.m. We refer to the S-index as Gen Synch S, and H-index as Gen Synch H in our results. We tested several different parameter values but only report the results for the cases where the delay parameters are set to 1 and 2.

7.3.7 Graphical models

In a graphical causal model a graph encodes the structure and relationships between the different signals. The relationships between the nodes are quantified by specifying a conditional probability distribution for each node, and the joint probability distribution. The graph structure is estimated by computing the unconditional and conditional independencies for all nodes.

Directed graphical models (DAG) are called Bayesian networks that take into account the direction associated with a connection and are typically based on assumption of Gaussian distributions. The first five methods reported below fall into this category. We used the Tetrad IV toolbox (www.phil.cmu.edu/projects/tetrad/tetrad4.html) for the Bayesian network methods with default parameter values.

Peter and Clark method (PC). This method searches for causal graphs under the assumptions that there are no cycles in the graph or no hidden variables (Meek, 1995).

Conservative PC (CPC). This method uses an adjacency search similar to PC but in a more conservative way limiting false positives (Ramsey et al., 2006).

Greedy equivalence search (GES). This method used the Bayesian information criterion score to add or eliminate network edges until the score cannot be improved any further (Chickering and Boutilier, 2002; Ramsey et al., 2010).

Fast causal inference (FCI). This method allows existence of latent (unmeasured) variables (Zhang, 2008).

Cyclic causal discovery (CCD). This method allows the existence of cycles in the graph (Richardson and Spirtes, 2001).

Lingam. LiNGAM (Linear, Non-Gaussian, Acyclic causal Models) is a non-Gaussian variant of Bayesian networks. This method used higher-order statistics, via ICA, to estimate the network connections (Shimizu et al., 2006). Under the assumption that each node has an independent input of non-Gaussian distribution, temporal ICA can be used to estimate the inputs and the ICA mixing matrix can be utilized to estimate network connections. We used the LiNGAM package (www.cs.helsinki.fi/group/neuroinf/lingam) for testing the Lingam method with default parameter values.

Patel's conditional dependence. This method looks at the difference between $P(\mathbf{x}|\mathbf{y})$ and $P(\mathbf{y}|\mathbf{x})$ to estimate connectivity and causality (Patel et al., 2006). From the conditional dependencies, a measure of connection strength, tau (τ), and a measure of connection directionality, kappa (κ), is derived. We applied this method to continuous and binary time courses. For the binary case, data under 10th percentile was limited to 0, data over the 90th percentile to 1, data in between was mapped onto a range of 0-1, and different thresholds. We report the results for the continuous data and binary data with a threshold at 0.75.

7.4 Results

7.4.1 FMRI network simulations

We first look at the results from one of the most typical network scenarios where the network has 15 nodes and 10 minutes of FMRI data is simulated for each subject. Figure 7.6 row1 shows the distribution of estimated network Z scores for the network matrix elements that are true connections (True Positives - TP). For methods that produce asymmetric network matrices (directionality) the higher value is used. Figure 7.6 row2 shows the distribution of Z scores for the network matrix elements that should be empty (False Positives - FP).

Figure 7.6 row3 shows the fraction of TP estimated with higher connection strengths than the 95th percentile of the FP distribution. As in *netsims*, this is a measure of the method's sensitivity to estimating the presence of true connections (called c-sensitivity). We find that some of the simplest linear methods (Full correlation, Partial correlation, and ICOV) have a high c-sensitivity of more than 85%. Information-theoretic measures (MI, PMI) and generalized synchronization indexes (Gen Synch) have an average c-sensitivity. All of the lag-based based (Granger, DC, GGC, PDC, DTF) perform very poorly with a c-sensitivity of less than 20%. For coherence and Patel's methods, the results vary depending on the variant or control parameter. However, for these methods even the best results do not show a high c-sensitivity (Patel's κ , Patel's κ bin, Coherence B3). The Bayesian network based methods show the highest c-sensitivity (about 90%) whereas its non-Gaussian variant, Lingam, performs poorly with a c-sensitivity of less than 20%.

Figure 7.6 row4 shows the accuracy of estimating the correct directionality of a connection (called d-accuracy). For the methods that output asymmetric matrices, containing directionality information, the direction accuracy is estimated by subtracting the Z score of the false connection from the Z score of true connection. We again find that that all the lag-based methods (Granger, DC, GGC, PDC, DTF) have only chance level (50%) d-accuracy. Gen sync S gives better results than the H-index, with an accuracy of nearly 60%. Patel's τ provides the best results with a d-accuracy of about 65%. However, we note that this method is not very sensitive to the detection of connections (i.e. low c-accuracy). Bayes net methods perform around average, whereas Lingam provides slightly better results.




Figure 7.6 – Performance of network estimation methods in FMRI for a 15-node network. (row1) The distribution of Z scores for correctly identified connections. (row2) The distribution of Z scores for falsely identified connections. The true connections increase FP estimation resulting in non-zero mean of some distributions. In row1 the TP distribution is shown in black, with the FP distribution shown in orange and reversed in row2. The distributions are over all 50 subjects, and true (TP) or empty (FP) network edges. (row3) This row is estimated from the top two rows, showing the fractional rate (sensitivity) of true connection detection after correcting for false positives. The distributions are over all subjects and the blue line plots the mean. (row4) This distribution of subtracted Z is over all subjects and true connections, and the blue dots show percentage of correct directions.

When comparing our results to those reported in Smith et al. (2011) employing a simpler neural model (which we refer to as *netstim results*), we find that the overall pattern still holds and the performance of most methods is only slightly increased with the more biophysically based model. The c-sensitivity of all methods is increased except for Patel's κ and Bayes nets. We find that this increase is highest for Coherence, Gen Synch and Lingam. We also observe a similar d-accuracy pattern as in *netstim results*. The d-accuracy of Granger, and DTF has slightly decreased. PDC scores have fallen by more than 10% whereas those of Lingham and Patels τ have increased by more than 10%. The scores for Gen Synch and Bayes Nets (CCD, CPC, FCI, PC, GES) are mostly similar to *netstim results*.

7.4.2 Variants of the biophysical network model

We now look at variants of the biophysical network model to see what impact particular assumptions have on the assessment of the different network analysis methods. We refer to the previous simulation as *baseline* against which we compare these new results. The plots for all the simulations are provided in the Appendix C - *Network Analysis Methods*, together with the set of simulations we used to compare each network aspect (Table C.1). For most network scenarios we found that the performance of the methods was very similar to *netstim results*, and we only report the cases where we found particular differences to the *netstim results*.

Network size and noise. We found that increasing the network size (Sim1, Sim2, Sim3, Sim4), reducing the noise level (Sim17), or adding a global additive confound (same random signal added to all nodes with 50% amplitude) (Sim10) did not produce any significant deviations from the *baseline* results. In all cases the overall pattern is very similar to the *baseline* results. The c-sensitivity remains highest for linear methods, Patel's κ , and Bayes nets (85-95%), intermediate for MI and Gen Synch (50-70%), and poor for lag-based methods (<20%). Patel's τ

displays the best d-accuracy (60-70%) while most other methods have average or poor performance. For the case of global mean confound, we observe that the TP and FP distributions have shifted upwards, as expected from adding a global confound. When compared to the *netsim* results, in all cases, both the c-sensitivity and d-accuracy of the lag based methods (DC, GGC, PDC, DTF) has decreased by 3-5%, whereas those of Bayes, Lingam, Patel's κ have increased by about 3-5%.

Session length. As the session length is changed, we find that the overall pattern is retained and generally just shifted up or down (i.e. improved or poorer). At 60 min session length (Sim5, Sim6), the c-sensitivity of most methods is above 80% except for Lingam and lag-based methods, whereas the d-accuracy is slightly increased. Lingam performs the best (77%) while most other methods are under 70%. At 250 min (Sim7), the c-sensitivity of the best methods approaches 97%. The d-accuracy of Lingam (90%), Patel's τ (80%), and CCD (75%) is the highest, with most other methods performing no better than chance. Compared to the *netsim* results, we find that at long session lengths Bayes net methods and Lingam are performing much better; approaching more than 95% c-sensitivity compared to 80% in the *netsim* results. The d-accuracy of Lingam and Patel's τ is also increased by 5% compared to *netsim* results.

One Strong external input. We consider the case where we have only one strong external input feeding into the network (Sim24). In this case, only one input is propagated around the network resulting in high correlations between nodes; making it harder to estimate causality. We find that the c-sensitivity of most linear and Bayes nets methods falls to around 60%. The d-accuracy of Gen Synch S has improved to more than 65% but remains similar for most other methods. Compared to the *netsim* results, the c-sensitivity of lag-based methods and Gen Synch have increased by 5%; Full correlation, ICOV, MI, and Lingam by 20%; and that of Bayes nets by 40%. The d-accuracy of most methods has also increased by 10% compared to *netsim* results.

Inaccurate ROIs. We consider the effect of mixing the FMRI time courses of different nodes. In this case, we mix a small amount of one nodes time course to another randomly chosen one (Sim11). We find extremely poor results, with the c-sensitivity of all methods below 50%. Compared to the *netstim* results, the c-sensitivity of all methods is increased by 20%, except for Full Correlation, ICOV, and Bayes nets for which the gain is more than 30%.

Connection strength changes. We tested the sensitivity of different methods at detecting changes in connection strengths, by halving the network connection strengths for half of the subjects (Sim23). Here we again find a very similar pattern to the *baseline* results. Compared to the *netstim* results, the c-sensitivity of lag-based methods and Coherence has increased by 5%, whereas for the remaining methods it has increased by 10%. The d-accuracy of Patel's τ has increased by 5%.

Backward connections. In the case of backward connections (Sim13), the overall pattern of results is similar to *baseline* but there is a 5% drop in the c-sensitivity of all methods, and a 3-5% drop in the d-accuracy of Gen Synch and Patel's τ . Compared to *netstim* results, the c-sensitivity of the MI, linear and Bayes nets methods has increased by around 15-20%, whereas for lag-based methods and Gen Synch the increase is slightly lower at around 10%. The d-accuracy of Patel's τ , Bayes nets, and Lingam has increased by 5%.

Lateral Connections. We investigated the effect of having dense (lateral) connections in the network (Sim16). We find that the c-sensitivity of all methods has dropped by 10% compared to the *baseline*, whereas the d-accuracy is mostly similar. Compared to the *netstim* results, there is a decrease of more than 5% in the c-sensitivity of all methods. The d-accuracy of all methods has also slightly decreased.

7.4.3 MEG network simulations

We also applied the network analysis methods directly to the simulated neural time courses without the haemodynamic estimation (i.e. simulating MEG not fMRI). With many time points, the Gen Sync methods are computationally intensive and were therefore excluded. When compared to the fMRI results, we find that (Figure 7.7) the performance of all methods is significantly increased except for linear and Bayes nets methods. For Bayes nets, the c-sensitivity has dropped significantly to 10% (80% change) without any change in d-accuracy (chance level). The c-sensitivity of linear methods (correlation, covariance) is only slightly increased to 90% (up 5%). The c-sensitivity of information theoretic measures (MI, Partial MI) is now approaching that of the best linear measures (90%). The biggest gain is observed in all lag-based methods and Lingam, where the d-accuracy has increased to around 95% for most methods, together with a very high c-sensitivity of more than 75% for most methods. The d-accuracy of Patel's τ has increased by 15% to 80%. However, this method still remains somewhat insensitive to detecting a connection in the first place (c-sensitivity $\sim$ 30%). Patel's κ , on the other hand, is now providing high sensitivity to detecting a connection (c-sensitivity $\sim$ 75%), but this measure does not provide information about directionality. We tested another Patel's method, Patel's $\tau \times \kappa$, in which the values from the two measures are multiplied together. We find that this method also offers a high d-accuracy ($\sim$ 80%) but its c-sensitivity remains low ($\sim$ 45%). The c-sensitivity of Coherence based methods has increased to more than 75%. For Coherence measures, we also estimated the 95th percentile (across the entire time and frequency range) instead of the mean to see if this increases sensitivity to nonstationarities in the data; however, we find that this does not increase the c-sensitivity of these methods. We tested the performance of all methods both in the case of bistate inputs and white noise inputs, and observed a very similar pattern.




Figure 7.7 – Performance of network estimation methods in MEG for a 15-node network. (row1) The distribution of Z scores for correctly identified connections. (row2) The distribution of Z scores for falsely identified connections. The true connections increase FP estimation resulting in non-zero mean of some distributions. In row1 the TP distribution is shown in black, with the FP distribution shown in orange and reversed in row2. The distributions are over all 50 subjects, and true (TP) or empty (FP) network edges. (row3) This row is estimated from the top two rows, showing the fractional rate (sensitivity) of true connection detection after correcting for false positives. The distributions are over all subjects and the blue line plots the mean. (row4) This distribution of subtracted Z is over all subjects and true connections, and the blue dots show percentage of correct directions.

7.4.4 Summary across all simulations

The specifications for all the simulations and network scenarios are provided in Table 7.1 and Table C.1. Figures 7.8 and 7.9 summarize the results for all simulations and methods, except simulation 4, 29, and 30 that was not performed for all methods.




Figure 7.8 – Summary of connection sensitivity across all simulations. The sensitivity to correctly detecting a network connection is shown for all simulations and methods with the mean overlaid in a thick black line.

To generate this summary figure, we took the best results from each class of methods (for each simulation). Figure 7.8 shows the sensitivity of all methods in correctly estimating network connections, and Figure 7.9 shows the accuracy of estimating directionality.




Figure 7.9 – Accuracy of estimating the directionality of network connection. The results are shown for all simulations and methods with the mean overlaid in a thick black line. The grey line shows chance level.

7.5 Discussion

In this work we evaluated the performance of many network estimation methods for a range of different network scenarios. We simulated realistic fMRI data using a biophysical model for neural activity in combination with a haemodynamic

model for vascular dynamics. The methods that we tested range from simple linear methods, such as correlation and covariance, to lag-based methods, such as Granger causality, to graphical models, such as Bayesian networks. We evaluated the methods sensitivity to detecting the presence of connections, ability to estimate correct directionality, and robustness against various issues that commonly arise in fMRI data analysis. Although the methods have different performance scores on some of the different network scenarios, an overall pattern still emerges that should be applicable to most setups typically encountered.

We find that some of the simplest linear methods (correlation, covariance) are the most sensitive to detecting the presence of a connection (90%), but do not provide information regarding the directionality of information flow. Bayesian network methods (CCD, CPC, FCI, PC, GES) are also equally sensitive to connection detection (90%), while also providing information about directionality but with only chance-level results. Additionally, we find that these methods are not as robust as the linear methods against different network setups. Information-theoretic measures (MI, Partial MI) and coherence-based methods have an average sensitivity to connection detection (50%) and do not provide any information regarding directionality. Gen Synch methods also have an average sensitivity (50%) as well as chance-level accuracy in estimating the correct directionality. Gen sync S gives better results than the Gen sync H. Patel's τ performs the best for estimating directionality (65%) but is very insensitive to detecting a connection in the first place, and also suffers in complicated network situations. All of the lag-based methods (Granger, DC, GGC, PDC, DTF) perform very poorly with sensitivity of less than 20% and often less than chance-level d-accuracy. The non-Gaussian variant of Bayesian networks, Lingam, also performs poorly both in detecting a connection and estimating directionality.

We compared our results to those reported in *netsims* testing similar network configurations using a simple bilinear model. We found that in most

network scenarios the performance of linear, information-theoretic, and coherence-based methods improved slightly. Performance scores for Patel's τ and Lingam increased by more than 5-10%. The performance of Gen Synch and Bayes nets was found to be mostly similar with only modest changes. All of the lag-based methods perform very poorly, with a drop of 5-10% in their performance compared to *netsim* results.

When the methods are applied directly to simulated neural time courses, the sensitivity of Bayes nets falls considerably to ~10%. Information-theoretic and coherence methods experience more than 20% increases in their sensitivities. The biggest gain is observed in the performance of all the lag-based methods and Lingam, with a d-accuracy approaching 99%. The d-accuracy of Patel's τ has also increased to 80% but this method remains somewhat insensitive to connection detection.

In almost all of the network scenarios the performance of the lag-based methods was just about average and slightly below that reported in the *netsims*. This difference in the performance of these methods compared to *netsim* results is due to the different prior mean/distribution of the haemodynamic parameters used in this work. These parameters affect the haemodynamic responses in different ways such as altering the maximum change, peak sharpness, smoothing level, and delay. Particularly, the distribution on the transit time we used resulted in a larger mean time-to-peak and undershoot amplitude in the haemodynamic responses compared to *netsims*. This increases the level of blurring and delay in the haemodynamic responses making it difficult to estimate causality.

After removing the haemodynamic variability (Sim18), the performance of all the lag-based methods becomes similar to that reported in *netsims*. A similar effect is seen after increasing the neural lag with no haemodynamic variability (Sim20): the c-sensitivity and d-accuracy of some of the lag based

methods increases slightly compared to *netsim* results. However, the increase in d-accuracy is not significant. Even in longer sessions the d-accuracy of most lag-based methods remained at chance level. We therefore conclude that these methods give spurious directionality estimates that are not driven by lag-based causality, but a few other issues highlighted in the Granger literature (Tiao and Wei, 1976; Wei, 1978; Weiss, 1984; Nalatore et al., 2007; Nolte et al., 2008). These include the temporal smoothing (causing correlated signals to be causally linked), and the measurement noise (that can reverse the estimated directionality). On the other hand, these methods perform very well when estimating directionality of the information flow in MEG data. Here these methods utilize the structured information available in the high temporal resolution MEG data, which is lost in BOLD due to haemodynamic induced temporal blurring.

Since the linear methods perform significantly better than non-linear methods, it is indicative that the BOLD signal has a relatively Gaussian distribution. This is expected because the haemodynamic response removes much of the structured information from the signal. This could be one reason why Gen Synch and Bayes nets give mostly similar estimates of directionality compared to *netsim* results. On the other hand, Lingam and Patel's τ utilize higher-order statistics rather than structured information that aids the directionality estimation, and hence observed difference in their scores. This is evident from long session lengths, since there is no significant increase in the d-accuracy of Gen Synch and Bayes nets methods, whereas Lingam and Patel's τ perform much better compared to *netsim* results. This could also explain why the performance of Lingam and Patel's τ increased in MEG data, whereas that of Bayes nets dropped.

Inaccurate ROI specification and connection strength changes result in a drop in performance of most methods compared to *baseline*. This emphasizes the importance of the defining correct functional ROIs to maximize estimates of causality. However, the performance of most methods remains significantly better

than reported in *netsim* results. The sophisticated dynamics emerging from the more biophysical model may give some of these methods a better chance in finding complex interactions.

In the case of backward connections, all methods perform significantly better compared to *netsim* results. The c-sensitivity of most methods is increased by 10%, while the d-accuracy is increased by 5%. The reason for this increase is the organization of connections within a single region. The forward connections in the model influence the output population indirectly via the spiny stellate cells. On the other hand, the backward connections project directly to the pyramidal cells (output) of another region. Furthermore, the backward connections also project to the inhibitory interneurons that have strong inhibitory connections with the population (spiny stellate) receiving the external input. Due to this hierarchical organization, the backward connections directly amplify the output from another node while inhibiting the input received locally. Due to this the activity of other regions get propagated across the network whereas the random input is suppressed.

In the case of dense networks, the extra connections in the bilinear model used in *netsims* were mediated by the same excitatory connection type. However, in our model these connections were mediated by lateral connections that cause inhibitory effects. It is, therefore, not surprising that the performance of all methods compared to *netsim* results has dropped. A more proper comparison would require that the connections are setup such that they cause similar effects in both models. However, there is no explicit connection or population for mediating inhibitory effects in the bilinear model. Although negative connections can result in inhibitory effects in this model, this scenario was not tested in *netsims*.

In the case where the network is driven by one strong external input, the c-sensitivity is on average 20% higher whereas the d-accuracy of the top methods is

10% better compared to *netsim* results. This might again result from the connection structure of the biophysical model. The pyramidal cells (output) of one region project to spiny stellate cells of another region, which then connect internally to the pyramidal cells via strong intrinsic excitatory connections. Due to this, there is a larger amplification of the output received from another node that results in better performance of all methods.

After having compared the results for a range of different network configurations with the *netsim* results, the primary result is that even with much more sophisticated underlying generative model, the main results reported in *netsims* using a simple neural model still hold. Most network configurations that produce major differences are likely due to the connectivity structure of the biophysical model rather than the fMRI setup parameters. This is an interesting result not just for network estimation methods but also for the biophysical network modelling.

One reason for the similar scores is that the haemodynamic responses cause significant blurring reducing a lot of structured information, and most complex dynamics may actually be lost before the methods are even performed. This is evident as the performance of all methods increases significantly when they are applied directly to simulated neural activity without the haemodynamic estimation. This would also explain why driving the network with white Gaussian noise did not produce significant differences compared to *netsim* results.

Another reason is that the model parameters are set up such that they resemble empirical responses to make them biologically plausible. As these models are in same dynamical regime, as opposed to simulating these models with some default parameters, they may result in similar dynamics and hence similar results. Importantly, however, most of the network analysis methods tested here are based on assumptions of stationary FC in the network. Other

approaches that look at network interactions in a dynamic way, such as hidden Markov models (HMMs) (Baum and Petrie, 1966), might perform better at finding complex interactions resulting from the more sophisticated model.

The connectivity structure of the model has a significant effect on the methods performance even for the same fMRI settings. We want to set up the connections such that the results remain comparable, for e.g. by relating some aspects of the bilinear model (in *netsims*) to that of the biophysical model used here. We can try to map the network connection parameters of the biophysical model to the bilinear model to understand which key parameters contribute to the observed dynamics. One approach to this is to fit the biophysical model to the output of the simple bilinear model, and see which parameters in the biophysical model control the gross network parameters of the bilinear model. However, there is no unique set of parameters that can be mapped (i.e. many combinations of parameters will still produce the same BOLD response). Due of this it is generally difficult to set-up the networks such that they remain comparable.

A related problem is the way in which the nodes are defined in both models. Each node in the biophysical model consist of three populations but the generic form of the differential equation that goes into each of them is the same, and can be thought of as an extended version of the bilinear model. This means that the biophysical model effectively has three nodes for each region, whereas in bilinear model each region is actually a single node. The network topology, therefore, is quite different in the two models.

8 General Discussion

In this study we combined experimental and theoretical work to investigate the dynamics and underlying mechanisms of spatiotemporally organized activity in the resting state brain. We analyzed resting state MEG data from 10 healthy participants and found that specific brain regions exhibit significant transient power correlations (functional connectivity) in alpha and beta frequencies on slow time scales, in agreement with previous studies of resting-state functional connectivity in fMRI (Beckmann et al., 2005; Smith et al., 2009) and MEG (Brookes et al., 2011a; Brookes et al., 2011b; de Pasquale et al., 2010; Hipp et al., 2012; Nikouline et al., 2001).

We developed a biophysical network model to explore the mechanisms underlying resting state responses. The model included biophysical node dynamics, empirical brain connectivity, biologically plausible transmission delays, and realistic background noise. We found that the model gives rise to metastable oscillatory states where the best fit with empirical data is also observed. Frequency-specific functional connectivity analysis revealed power correlations in alpha and beta frequency bands prominent on slow time scales (<1 Hz), as observed in empirical data. Using sliding-window correlation we found that the same regions exhibit significant temporal non-stationaries in functional connectivity in the same frequency bands. We also observed the existence of several RSNs that had significant similarity with those reported in previous fMRI and MEG studies.

We used model-based simulations to evaluate the performance of many different network estimation methods. Our evaluation of methods on fMRI data suggests that simple linear measures, such as correlation, can provide robust estimates of connectivity while inferring directionality remains hard for all

methods. When applied directly to MEG data, the performance of most methods increased significantly. The lag-based methods perform particularly well when estimating directionality of the information flow, as they utilize the high temporal resolution of MEG.

The simple analysis of the empirical MEG data has provided much insight into the functional organization of the brain. However, with these methods, it has remained difficult to uncover the physiological mechanisms leading to the observed responses. The employments of brain-inspired computational models can potentially reveal the origin of these functional networks. By using a biophysical network model, we were able to make falsifiable predictions such as whether time delays cause frequency suppression in the network, does noise result in transient synchronizations between regions, and whether spatial patterns in the structural connectivity are also reflected in the simulated dynamics. By investigating such questions we have been able to extract the dynamical scenario where the model best explains the MEG responses. To date, no theoretical mechanisms underlying the generation of these complex dynamics has been suggested.

Our model offers a coherent mechanism for the generation of time-varying frequency specific correlation patterns observed in resting state FMRI and MEG data. The brain anatomical network together with the time delays and noise enables a broad metastable regime (Figure 3.6), between incoherence and full network synchrony, where subsets of brain regions (Figure 3.7, 3.8) transiently synchronize to form dynamic networks. Due to the presence of delays, a novel form of frequency depression occurs resulting in the emergence of low collective frequencies in the alpha and beta bands (Figure 3.9, 4.3). The increase in synchrony between regions, therefore, increases the power in these low collective frequency bands (Figure 3.11), and the regions show structured power correlations in these reduced collective frequencies (Figure 4.4, 5.6). Due to the heterogeneity

of the structural connectivity and delays, multiple dynamic clusters oscillate at slightly different frequencies explaining the observed range of frequencies (within alpha and beta bands) (Figure 4.4). These results are in agreement with previous studies suggesting such reduced collective frequencies in weakly coupled oscillators (Niebur et al., 1991). This theoretical scenario explains how interactions between the cortical regions can lead to the emergence of low frequency oscillations seen in MEG.

Since the regions are only weakly coupled, they are highly metastable and prone to natural disruption (Figure 3.10). Noise and other competitive processes in the network disrupt these delicate synchronized states resulting in the loss of signal power at low frequencies. As these regions form new networks, the power at these low frequencies increases again giving rising to the observed power fluctuations in synchrony/correlations (Figure 3.11, 6.3, 6.4). These synchrony/correlation fluctuations are captured by the bimodal distribution of the correlation values, which is lost at higher frequencies (Figure 3.12, 6.3, 6.4). Therefore, the model explains the band-limited MEG power fluctuations as resulting from transient synchronizations in the underlying networks.

The structure of the anatomical network plays a major role in shaping the spatial profile of functional connectivity. The brain anatomical network is organized in a modular way with the existence of small-world networks (Figure 3.5). The brain regions in these modular networks are strongly connected to each other, than to other regions. These regions, therefore, have a higher tendency to synchronize with each other (Figure 3.8) resulting in functional networks composed of specific brain regions (Figure 5.5, 5.7). The modularity of the anatomical network enables the formation of multiple simultaneous networks that are independent of each other (Figure 5.5, 5.7). Due to heterogeneity of couplings and delays, these networks interact at different collective frequencies (Figure 5.3,

5.6). This scenario would explain how multiple (alpha/beta) networks observed in MEG data can emerge from the intrinsic dynamics of the brain.

An important implication of our findings is the direct relation between degree of synchrony between brain regions and the power in alpha and beta frequency bands. Indeed, it has been observed that beta band synchronization between the left and right hemispheres increases if the oscillations have larger amplitudes (Nikouline et al., 2001). However, the amplitude-based measures can be confounded by several factors, including source reconstruction and volume conduction effects, resulting in lower estimates of in-phase synchronization. As a more robust measure, Hillebrand et al. (2012) estimated the functional connectivity in MEG using the phase lag index (PLI), and observed similar higher synchronization in the alpha and beta frequency bands when there was more power in these frequencies.

By reproducing the synchrony fluctuations (Figure 3.11, 3.12, 4.7) and correlations (Figure 4.5) on slow timescales, the model links the same underlying phenomena with fMRI BOLD fluctuations. A previous study using the Kuramoto model had found that fluctuations in the synchrony degree could explain the resting-state BOLD signal fluctuations (Cabral et al., 2011). Here, we extended this work to show that these transient synchronizations link the MEG power fluctuations with BOLD fluctuations.

Recent intra-cortical recordings suggest that spontaneous modulations in firing rate and gamma power are likely correlated with spontaneous BOLD signal fluctuations (Nir et al., 2008; Shmuel and Leopold, 2008). A previous computational modelling study by Deco et al. (2009) found that Wilson-Cowan units that exhibit self-sustained gamma frequency oscillations give rise to low frequency fluctuations ($< 1\text{Hz}$) observed in the fMRI data. More recently, the study by Carbal et al. (2014), employing Kuramoto oscillators set to oscillate in the

gamma range, reproduced the alpha and beta frequency fluctuations observed in resting state MEG data. However, in these studies, the isolated brain regions were assumed a priori to oscillate in the gamma frequency range, which in turn affects the collective frequencies of the emergent responses. In our model, the isolated regions display gamma oscillations at weak couplings, and this frequency drops when brain regions start to interact with each other. The gamma oscillations in our model emerge naturally from the intrinsic excitatory and inhibitory population dynamics of the individual regions. Therefore, our model further supports the electrophysiological findings reported in previous studies.

In our work, we have found a negative correlation between power at low and high frequencies. In the metastable regime, the power at the reduced collective frequency increases resulting in a state with less energy demands (less spikes per second) and a low BOLD signal. However as this synchrony is disrupted, the power in the higher frequencies increases leading to more energy consumption and an increase in the BOLD signal. Therefore, the BOLD signal has a positive correlation with higher frequencies and negative correlation with alpha and beta power, as found in previous studies (Scheeringa et al., 2011). These results suggest the metastable transient synchronizations between regions not only modulate the BOLD signal but also determine the collective frequencies of the regions.

There are other theoretical models that can produce alpha oscillations in brain regions. Most of these studies also emphasize the role of macroscopic connectivity in the generation of these oscillations (Freyer et al., 2011; Lopes da Silva et al., 1974; Nunez, 2000). In particular, Freyer et al. (2011) used to a two state model with cortico-thalamic connectivity to reproduce the bistable alpha band oscillations observed in resting state cortical recordings. In our work, we have included these cortico-thalamic projections and gone further to show that the

whole brain connectivity at the large scale can reproduce many other aspects of resting state activity.

There are other studies that have developed computational models to investigate the mechanisms of resting state brain activity (Honey et al., 2007, 2009; Ghosh et al., 2008; Deco et al., 2009; Deco and Jirsa, 2012; Cabral et al., 2011, 2014). With the exception of the study by Cabral et al. (2014), all these studies have focused their investigations on fMRI. The computational models used in these studies differ quantitatively in their complexity, dynamical regimes, and the interpretation of the emerging activity.

Abstract neural mass models, such as the FitzHugh-Nagumo model (Ghosh et al., 2008), Wilson-Cowan model (Deco et al., 2009), or the Kuramoto model (Cabral et al., 2011, 2014) have been shown to reproduce some characteristics of resting state fMRI (and MEG) data. However, the low degree of biological relevance of these models makes it difficult to interpret and link the underlying mechanisms with neurophysiology. The dynamics emerging from the model may also be generated by complex hidden mechanisms that have no biological plausibility (Ghosh et al., 2008). Moreover, due to the high degree of abstraction at the node level only the network level interactions between regions govern dynamics. Therefore, the contribution of the underlying cortical dynamics to the network dynamics cannot be assessed using these models.

On the other hand, the spiking neuron model (Deco and Jirsa, 2012) explicitly models the spiking behaviour of individual neurons. Having such amount detail can make the simulations more realistic but this comes at a cost of increased complexity (interpretability) and computability. Another concern is the large parameter space, which can be tweaked sufficiently to match almost any empirical observation. While the biophysical basis of the model can restrict this to what remains biologically plausible, this can still remain an issue. More

importantly, fMRI and MEG measure the combined effect of large numbers of spikes and neurons. There is an additional temporal blurring in fMRI due to measuring a vascular response. Since the activity of individual neurons cannot be resolved from non-invasive functional neuroimaging data, it seems implausible that the contribution of individual neurons can be examined using such models.

In contrast, the biophysical model we have developed in this work is based on DCM that offers an intermediate level of detail and complexity. The hierarchical organization and dynamics of the model are based on empirical observations of the brain. The model consists of the main types of neural populations in cortical regions with realistic connectivity. The dynamics of the model capture the activity of main types of ion channels and synaptic currents. Furthermore, the description of activity is based on the population voltage, which makes it particularly attractive to investigate macroscopic network dynamics using MEG. The DCM model can produce oscillatory activity in the whole MEG spectrum by simply changing the physiologically relevant model parameters or population kinetics, and has been used to reproduce a wide variety of empirical neuroimaging data. Therefore, the DCM also offers a more biologically interpretable way of investigating functional brain networks.

The main aim of this work was to investigate whether a biophysical model can reproduce characteristics of empirical MEG resting state data. We investigated interactions between brain regions in the presence of realistic connectivity, couplings, delays, and noise, and find that the model is able to reproduce the empirical characteristics in spectral, temporal, and spatial domains. Our findings suggest that realistic biophysical models can be used to successfully uncover the mechanisms underlying brain activity during rest and otherwise.

Different neuroimaging modalities provide information at a variety of spatial and temporal scales, and it is often useful to combine data from multiple

techniques in identical experimental conditions. However, a systematic approach for combining data across multiple modalities and automatically finding patterns of related changes remains a major challenge. Conventional techniques for data integration often suffer due to missing brain activity, different spatiotemporal scales as well as physiological factors. Combining bottom-up neural generative models with top-down summary metrics offers a unique way to utilize information from different modalities. In this work, we found that the model reproduced MEG BLP fluctuations and fMRI BOLD fluctuations and explained them as due to the same underlying phenomena. This highlights the importance of generative modelling in multimodal data integration, and motivates further analysis of the simulated BOLD data to see if the model can predict other key characteristics. Due to the high spatial resolution of fMRI, it would be particularly attractive to perform spatial FC analysis (e.g. using BLP correlation or ICA) to see if the model can reproduce the same networks that it did in MEG (and which have also been previously reported in fMRI studies (Beckmann et al., 2005; Smith et al., 2009)). On the other hand, our model also reproduced local gamma oscillations that have been reported in previous LFP studies (Nir et al., 2008; Shmuel and Leopold, 2008). However, the role of distributed gamma oscillators in long-range communications and the emergence of low frequencies still remains unclear, and can be explored further by supplementing the model with data from LFP and single neuron studies.

Choosing an appropriate level of model detail is a very important consideration in computational modelling. Having too much detail can make the model computationally intractable whereas having too little will miss result in missing dynamics. A multimodal data analysis approach that combines data from different spatiotemporal scales of activity (e.g. cellular to whole brain) can be used to investigate/select the neuronal elements that produce observable behaviour in a particular neuroimaging modality. For example, synaptic plasticity plays a key

role in neuronal computation. Short-term plasticity (e.g. adaptation) allows synapses to decode spike trains whereas long-term plasticity (e.g. potentiation) underlies learning and memory, but which of these produce an observable effect in fMRI or MEG is mostly unknown. By developing multiple realistic models that incorporate different neuronal elements, verifying them with single cell recordings, and comparing them with non-invasive neuroimaging data the right set of elements can be selected. Further improvements in the model can be made by studying the populations within individual regions and the local circuitry between them to derive more realistic models of population/regional dynamics.

The structural connectivity plays a key role in shaping the spatiotemporal dynamics of brain networks. The connectivity matrix we used in this work suffers from multiple limitations, such as missing brain regions, missing fiber tracts, and non-realistic time delays. It is highly desirable to estimate more realistic whole brain structural connectivities by combining high-resolution structural neuroimaging data from multiple modalities with better parcellation schemes. Estimates of actual fiber lengths would lead to more realistic time delay estimates. Other fiber properties that can affect signal transmission, such as fiber diameter and myelin, could also be taken into account. The structural connectivities can then be used in the model to see if it can better predict the spatiotemporal activity of empirical brain networks. It would also be valuable to characterize how the structural connectivity effects brain dynamics. For e.g. different brain parcellation schemes would result in different cluster sizes, which would affect the dynamic characteristics such as the level of metastability in the network.

In this work, we investigated the interactions between brain regions within single frequency bands. However, recent studies have found that brain regions can interact in different frequency bands (Jensen and Colgin, 2007). One particular example of this is that the phase of a low frequency oscillation can modulate the power of a faster oscillation. These studies highlight the need to examine

interactions between regions in different frequency bands. Previous work has shown that the DCM model can reproduce such phase-amplitude coupling under realistic brain connectivity (Marreiros et al., 2008). Therefore, the model also offers a means to investigate how such type of dynamics can emerge in brain networks.

In our dynamic FC analysis, we used a frequency dependent window size, giving a wide time window for low frequency signals and a narrow time window for high frequency signals. The rationale for using this was to control the uncertainty in the value of the correlation coefficient such that it is the same across all frequency bands (see Section 6.2.4 - *Sliding window correlation* for the discussion). However, this approach is not very well suited for the analysis of brain activity because signals from a brain area may appear in multiple frequency ranges but in a similar time window. Further improvements in our proposed algorithm could solve this problem by adjusting the time window by dynamically changing the number of waves so as to match the neurophysiological patterns.

We investigated the nonstationary behaviour using summary measures without characterizing other temporal properties of the underlying networks, such as the lifetime of network states and transition probabilities between states. A recent study has shown that the activity of resting state networks is shorter lived, and that certain network transitions are more likely than others (Baker et al., 2014). Another study suggests that networks observed in MEG are underpinned by even faster transiently synchronizing networks (Smith et al., 2012). Further investigations of the simulated temporal dynamics would have to be carried out to investigate the extent to which the model is able to reproduce such temporal dynamics.

We found that the model failed to reproduce the spatial profiles of some networks. One reason for this could be the limitations of beamforming. We found that there was high coherence in the simulated data, which might have resulted in

the failure of beamformer localization (Wipf et al., 2010). This would need to be investigated further to see if this problem can be overcome by performing source reconstruction using other methods or altering some properties of the model. Although FC analysis can also be performed in sensor space, it is more desirable to do it in source space as this increases the interpretation. Importantly, our simulation based framework, that combines realistic simulated MEG data and MEG forward modelling, can also be used to systematically assess the conditions under which different source estimation methods might fail, and how they affect FC estimates in source space.

We found that certain network estimation methods perform better than others on MEG data. Our evaluation of network estimation methods on fMRI data also suggests that these methods perform differently in different network setups. However, we did not extensively test the performance of the methods in different MEG network setups, such as the effect of network size (due to a chosen parcellation scheme), effect of session time, or the number of trials. We believe that such evaluation would be very valuable as this would enable someone to pick the right set of methods for given experimental setups, and associate a confidence level with the methods estimates. This can also be used to aid Bayesian model selection (hypothesis testing), where the most likely connectivity model generating some observed data is selected based on error minimization (David et al., 2006). The current approach has key limitations including a lack of a priori connectivity information, need to evaluate a large number of models, and that multiple models can reproduce the observed activity. However, some of these problems can be overcome supplementing this approach with the different connectivity measures to first estimate the most likely connectivity profiles, and associate confidence with them, before further model selection is performed.

There is a rapidly increasing amount of clinical data of the nervous system, with many neurological disorders (e.g. schizophrenia) characterized by changes in

brain connectivity (Bullmore and Sporns, 2009) and neural activity (Lynall et al., 2010). The computational modelling approach presented here can be used to explore how changes in biophysical parameters (e.g. connectivity changes) alter brain activity. This can lead to new biomarkers or insights into the mechanisms underlying disease dynamics, or the generation of new hypothesis about brain function that may be experimentally tested and ultimately translated into the clinic. Moreover, the computational models can also be used to predict the extent to which the brain can adapt in these conditions, and any possible intervention that can aid in brain reorganization.

We believe that the work performed here will serve as groundwork for further investigations into biophysical computational models of the brain and the analysis of brain activity.

Appendix A

AAL Brain Parcellation Template

Table A.1 – List of brain regions in the AAL brain parcellation template.

Brain Region	Left	Right
Precentral gyrus	1	90
Superior frontal gyrus, Dorsolateral part	2	89
Superior frontal gyrus, Orbital part	3	88
Middle frontal gyrus	4	87
Frontal Middle Orbital	5	86
Inferior frontal gyrus, Opercular part	6	85
Inferior frontal gyrus, Triangular part	7	84
Inferior frontal gyrus, Orbital part	8	83
Rolandic Operculum	9	82
Supplementary Motor Area	10	81
Olfactory cortex	11	80
Superior frontal gyrus, Medial part	12	79
Middle frontal gyrus, Orbital part	13	78
Gyrus rectus	14	77
Insula	15	76
Anterior cingulate gyrus	16	75
Median cingulate gyrus	17	74
Posterior cingulate gyrus	18	73
Hippocampus	19	72
ParaHippocampal gyrus	20	71
Amygdala	21	70
Calcarine fissure	22	69
Cuneus	23	68
Lingual gyrus	24	67
Superior occipital gyrus	25	66
Middle occipital gyrus	26	65
Inferior occipital gyrus	27	64
Fusiform gyrus	28	63
Postcentral gyrus	29	62
Superior parietal gyrus	30	61
Inferior parietal but supramarginal and angular gyri	31	60
Supramarginal gyrus	32	59
Angular gyrus	33	58
Precuneus	34	57
Paracentral Lobule	35	56
Caudate nucleus	36	55
Putamen	37	54
Pallidum	38	53
Thalamus	39	52
Heschl gyrus	40	51
Superior temporal gyrus	41	50
Temporal pole: superior temporal gyrus	42	49
Middle temporal gyrus	43	48
Temporal pole: middle temporal gyrus	44	47
Inferior temporal gyrus	45	46

Appendix B

DTI Anatomical Connectivity

Data acquisition. The DTI data² were acquired from 21 healthy, normal participants (11 males and 10 females, age: 22 – 45 years). All scans were performed on a Philips Achieva 1.5 Tesla Magnet. Diffusion MRI data were acquired using a single-shot echo planar imaging-based sequence covering the whole brain (repetition time, $TR = 9390\text{ ms}$; echo time, $TE = 65\text{ ms}$). DTI images utilized 32 optimal nonlinear diffusion weighting directions ($b = 1200\text{ s/mm}^2$) and 2 non-diffusion weighted volumes, giving a reconstructed voxel size of $2\text{mm} \times 2\text{mm} \times 2\text{mm}$. We also acquired $T1$ weighted structural images with a three-dimensional FLASH sequence ($TR = 12\text{ ms}$, $TE = 5.6\text{ ms}$, $\text{flip angle} = 19^\circ$), with elliptical sampling of k-space, giving a voxel size of $1\text{mm} \times 1\text{mm} \times 1\text{mm}$ in 5.05 minutes.

Brain parcellation. We used the automated anatomical labelling (AAL) template (Tzourio-Mazoyer et al., 2002) to parcellate the diffusion MRI data into 90 cortical and subcortical regions, where each region represents a node of the brain network. For each participant, we conducted the parcellation in the diffusion MRI native space.

We used the FLIRT tool (FMRIB, Oxford) (Jenkinson et al., 2002) to linearly co-register the $b0$ image in diffusion MRI space to the $T1$ weighted structural image. The transformed $T1$ -weighted image was then mapped to the $T1$ template of ICBM152 in MNI space (Collins et al., 1994). The resulting transformation was inversed and further applied to warp the AAL mask from MNI space to the

² The data acquisition and analysis of regional connectivity was performed by Prof. Morten Kringelbach of the Department of Psychiatry, University of Oxford, UK.

diffusion MRI native space, where interpolation using nearest-neighbour method ensured that the discrete labelling values were preserved.

Regional connectivity. The diffusion-weighted images from all subjects were co-registered on the same reference volume using affine transforms. The data were pre-processed to correct for head motions and eddy current induced image distortions. We used the FDT toolbox in FSL (<http://www.fmrib.ox.ac.uk/fsl/>, FMRI, Oxford) to carry out the various stages of processing. Following this, we estimated the local probability distribution of fibre direction at each voxel (Behrens et al., 2003). We then used the probtrackx algorithm allowing for automatic estimation of two fibre directions within each voxel, which can significantly improve the tracking sensitivity of non-dominant fiber populations in the human brain (Behrens et al., 2007).

The links between nodes were weighted in proportion to the number of white matter tracts detected between brain areas. The anatomical connectome at the voxel level was estimated by applying a probabilistic tractography algorithm using a sampling of 5000 fibers per voxel. The connectivity probability from voxel i to j was calculated by the total number of fibres passing through voxel j divided by the total number of fibers sampled from voxel i . At the region level, the connectivity was calculated based on the total number of voxels in AAL defined seed regions (i.e. for a region consisting of n voxels, $5000 * n$ fibers were sampled). Using this scheme the connectivity probability between all the 90 AAL defined brain regions was estimated in an all-all way. However, these estimates of regional connectivity depend on the chosen seed locations and are therefore bi-directional. We performed the procedure in both directions (from region j to i , and vice-versa), and the unidirectional connectivity probability between regions i and j was calculated by averaging these two probabilities. Regional connectivity analysis was performed using FSL.


Network Parameters. A weighted network was derived by associating connection strength with each edge, which was based on the average number of white matter tracts connecting the pair of brain regions. This resulted in a symmetric 90-by-90 matrix of coupling strengths, $\mathbf{C}$, for each subject. Coupling matrices were averaged across all subjects resulting in a single group matrix, representing the anatomical organization of the brain. The distance between the 90 brain regions were taken as the Euclidean distance separating the different pairs of brain regions. The distances between brain regions were used to calculate the corresponding time delays between brain regions, arising due to finite action potential propagation velocities. The pair-wise delays between regions i and j was calculated as the physical distance separating these regions (L_{ij}) divided by a signal propagation velocity V homogenous for all connections ($\tau_{ij} = L_{ij}/V$), giving a 90-by-90 delay matrix, $\mathbf{D}$.

Appendix C


Network Analysis Methods

Table C.1 – List of the network aspects tested together with the simulation number used for comparisons with the baseline and the values of parameters used.


Aspect	Simulation No.	Parameter Values tested
No of nodes	1, 2, 3, 4	5, 10, 15, 50
Noise level	2,17	1.0, 0.1
Session length	1, 5, 7, 25, 26, 27, 28	10, 60, 250, 5, 2.5 (min)
Global confound	1, 10	0, 50 (%)
Input	1, 8, 9	Shared inputs
Input	15, 24	One Strong input
Input	2, 29	White noise only inputs
Inaccurate ROIs	1, 11, 12	Mixed signals
Low TR	1, 19, 20	3, 0.25 (s)
HRF variability	1, 18, 20	0.5, 0
Connections	1, 21	Connection strength changes
Connections	1, 22, 23	Nonstationary strengths
Network topology	1, 13	Backwards connections
Network topology	1, 14	Cyclic connections
Network topology	1, 15	Stronger connections
Network topology	1, 16	More connections
Neural activity	1, 30	Neural time series




Simulation 1 (5 nodes, 10 minute sessions, TR=3.00s, noise=1.0%, HRFstd=0.5s, testrun 1: 1 subnet)




Simulation 2 (10 nodes, 10 minute sessions, TR=3.00s, noise=1.0%, HRFstd=0.5s, testrun 2: 2 subnets)




Simulation 3 (15 nodes, 10 minute sessions, TR=3.00s, noise=1.0%, HRFstd=0.5s, testrun 3: 3 subnets)







Simulation 5 (5 nodes, 60 minute sessions, TR=3.00s, noise=1.0%, HRFstd=0.5s, testrun 5: more timepoints)




Simulation 6 (10 nodes, 60 minute sessions, TR=3.00s, noise=1.0%, HRFstd=0.5s, testrun 6: more timepoints with 2 subnets)




Simulation 7 (5 nodes, 250 minute sessions, TR=3.00s, noise=1.0%, HRFstd=0.5s, testrun 7: best directionality info)




Simulation 8 (5 nodes, 10 minute sessions, TR=3.00s, noise=1.0%, HRFstd=0.5s, shared inputs)




Simulation 9 (5 nodes, 250 minute sessions, TR=3.00s, noise=1.0%, HRFstd=0.5s, shared inputs and more timepoints)




Simulation 10 (5 nodes, 10 minute sessions, TR=3.00s, noise=1.0%, HRFstd=0.5s, global mean confound)




Simulation 11 (10 nodes, 10 minute sessions, TR=3.00s, noise=1.0%, HRFstd=0.5s, bad ROIs (timeseries mixed with each other))




Simulation 12 (10 nodes, 10 minute sessions, TR=3.00s, noise=1.0%, HRFstd=0.5s, bad ROIs (new random timeseries mixed in))




Simulation 13 (5 nodes, 10 minute sessions, TR=3.00s, noise=1.0%, HRFstd=0.5s, backwards connections)




Simulation 14 (5 nodes, 10 minute sessions, TR=3.00s, noise=1.0%, HRFstd=0.5s, cyclic connections)




Simulation 15 (5 nodes, 10 minute sessions, TR=3.00s, noise=0.1%, HRFstd=0.5s, stronger connections)




Simulation 16 (5 nodes, 10 minute sessions, TR=3.00s, noise=1.0%, HRFstd=0.5s, more connections via lateral connections)




Simulation 17 (10 nodes, 10 minute sessions, TR=3.00s, noise=0.1%, HRFstd=0.5s, low BOLD noise)




Simulation 18 (5 nodes, 10 minute sessions, TR=3.00s, noise=1.0%, HRFstd=0.0s, zero HRF Covariance)




Simulation 19 (5 nodes, 10 minute sessions, TR=0.25s, noise=0.1%, HRFstd=0.5s, neural lag=100ms)




Simulation 20 (5 nodes, 10 minute sessions, TR=0.25s, noise=0.1%, HRFstd=0.0s, neural lag=100ms)




Simulation 21 (5 nodes, 10 minute sessions, TR=3.00s, noise=1.0%, HRFstd=0.5s, 2-group test)




Simulation 22 (5 nodes, 10 minute sessions, TR=3.00s, noise=0.1%, HRFstd=0.5s, nonstationary connection strengths)







Simulation 24 (5 nodes, 10 minute sessions, TR=3.00s, noise=0.1%, HRFstd=0.5s, only one strong external input)




Simulation 25 (5 nodes, 5 minute sessions, TR=3.00s, noise=1.0%, HRFstd=0.5s, Nsecs:150 with low BOLD noise)




Simulation 26 (5 nodes, 2.5 minute sessions, TR=3.00s, noise=1.0%, HRFstd=0.5s, Nsecs:150)




Simulation 27 (5 nodes, 2.5 minute sessions, TR=3.00s, noise=0.1%, HRFstd=0.5s, Nsecs:300 with low BOLD noise)



Simulation 28 (5 nodes, 5 minute sessions, TR=3.00s, noise=0.1%, HRFstd=0.5s, Nsecs:150 with low BOLD noise)





Bibliography

- Allen, E., Damaraju, E., Plis, S., Erhardt, E., Eichele, T., and Calhoun, V. (2012). Tracking Whole-Brain Connectivity Dynamics in the Resting State. *Cerebral Cortex*, 24(3), 663-676. <http://dx.doi.org/10.1093/cercor/bhs352>
- Arenas, A., Díaz-Guilera, A., and Pérez-Vicente, C. (2006). Synchronization Reveals Topological Scales in Complex Networks. *Physical Review Letters*, 96(11). <http://dx.doi.org/10.1103/physrevlett.96.114102>
- Arfanakis, K., Cordes, D., Haughton, V., Moritz, C., Quigley, M., and Meyerand, M. (2000). Combining independent component analysis and correlation analysis to probe interregional connectivity in fMRI task activation datasets. *Magnetic Resonance Imaging*, 18(8), 921-930. [http://dx.doi.org/10.1016/s0730-725x\(00\)00190-9](http://dx.doi.org/10.1016/s0730-725x(00)00190-9)
- Arnhold, J., Grassberger, P., Lehnertz, K., and Elger, C. (1999). A robust method for detecting interdependences: application to intracranially recorded EEG. *Physica D: Nonlinear Phenomena*, 134(4), 419-430. [http://dx.doi.org/10.1016/s0167-2789\(99\)00140-2](http://dx.doi.org/10.1016/s0167-2789(99)00140-2)
- Asteriou, D. and Hall, S. (2011). *Applied econometrics*. 2nd ed. Basingstoke [England]: Palgrave Macmillan, pp.319–333.
- Baccalá, L., and Sameshima, K. (2001). Partial directed coherence: a new concept in neural structure determination. *Biological Cybernetics*, 84(6), 463-474. <http://dx.doi.org/10.1007/pl00007990>
- Baker, A., Brookes, M., Rezek, I., Smith, S., Behrens, T., Probert Smith, P., and Woolrich, M. (2014). Fast transient networks in spontaneous human brain activity. *Elife*, 3. <http://dx.doi.org/10.7554/elife.01867>
- Bandettini, P., Jesmanowicz, A., Wong, E., and Hyde, J. (1993). Processing strategies for time-course data sets in functional mri of the human brain. *Magnetic Resonance In Medicine*, 30(2), 161-173. <http://dx.doi.org/10.1002/mrm.1910300204>
- Banerjee, O., Ghaoui, L., d'Aspremont, A., and Natsoulis, G. (2006). Convex optimization

- techniques for fitting sparse Gaussian graphical models. *Proceedings Of The 23Rd International Conference On Machine Learning - ICML '06*, 9. <http://dx.doi.org/10.1145/1143844.1143856>
- Baum, L. and Petrie, T. (1966). Statistical Inference for Probabilistic Functions of Finite State Markov Chains. *The Annals of Mathematical Statistics*, 37(6), pp.1554-1563. <http://dx.doi.org/10.1214/aoms/1177699147>
- Beckmann, C. (2012). Modelling with independent components. *Neuroimage*, 62(2), 891-901. <http://dx.doi.org/10.1016/j.neuroimage.2012.02.020>
- Beckmann, C., DeLuca, M., Devlin, J., and Smith, S. (2005). Investigations into resting-state connectivity using independent component analysis. *Philosophical Transactions Of The Royal Society B: Biological Sciences*, 360(1457), 1001-1013. <http://dx.doi.org/10.1098/rstb.2005.1634>
- Behrens, T., Berg, H., Jbabdi, S., Rushworth, M., and Woolrich, M. (2007). Probabilistic diffusion tractography with multiple fibre orientations: What can we gain?. *Neuroimage*, 34(1), 144-155. <http://dx.doi.org/10.1016/j.neuroimage.2006.09.018>
- Behrens, T., Woolrich, M., Jenkinson, M., Johansen-Berg, H., Nunes, R., and Clare, S. et al. (2003). Characterization and propagation of uncertainty in diffusion-weighted MR imaging. *Magnetic Resonance In Medicine*, 50(5), 1077-1088. <http://dx.doi.org/10.1002/mrm.10609>
- Benjamini, Y., and Hochberg, Y. (1995). Controlling the False Discovery Rate: A Practical and Powerful Approach to Multiple Testing. *Journal Of The Royal Statistical Society. Series B*, 57(1), 289-300. Retrieved from <http://www.jstor.org/stable/2346101>
- Bernardino, A., and Santos-Victor, J. (2005). A real-time gabor primal sketch for visual attention. *Proc. of the 2nd IBPRIA*, Estoril, Portugal.
- Biswal, B., Zerrin Yetkin, F., Haughton, V., and Hyde, J. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic Resonance In Medicine*, 34(4), 537-541. <http://dx.doi.org/10.1002/mrm.1910340409>
- Breakspear, M., Terry, J., and Friston, K. (2003). Modulation of excitatory synaptic coupling facilitates synchronization and complex dynamics in a biophysical model of

- neuronal dynamics. *Network: Computation In Neural Systems*, 14(4), 703-732.
<http://dx.doi.org/10.1088/0954-898x/14/4/305>
- Brookes, M., Hale, J., Zumer, J., Stevenson, C., Francis, S., and Barnes, G. et al. (2011a). Measuring functional connectivity using MEG: Methodology and comparison with fcMRI. *Neuroimage*, 56(3), 1082-1104.
<http://dx.doi.org/10.1016/j.neuroimage.2011.02.054>
- Brookes, M., Liddle, E., Hale, J., Woolrich, M., Luckhoo, H., Liddle, P., and Morris, P. (2012). Task induced modulation of neural oscillations in electrophysiological brain networks. *Neuroimage*, 63(4), 1918-1930.
<http://dx.doi.org/10.1016/j.neuroimage.2012.08.012>
- Brookes, M., Vrba, J., Robinson, S., Stevenson, C., Peters, A., and Barnes, G. et al. (2008). Optimising experimental design for MEG beamformer imaging. *Neuroimage*, 39(4), 1788-1802. <http://dx.doi.org/10.1016/j.neuroimage.2007.09.050>
- Brookes, M., Woolrich, M., Luckhoo, H., Price, D., Hale, J., and Stephenson, M. et al. (2011b). Investigating the electrophysiological basis of resting state networks using magnetoencephalography. *Proceedings Of The National Academy Of Sciences*, 108(40), 16783-16788. <http://dx.doi.org/10.1073/pnas.1112685108>
- Brookes, M., O'Neill, G., Hall, E., Woolrich, M., Baker, A., & Palazzo Corner, S. et al. (2014). Measuring temporal, spectral and spatial changes in electrophysiological brain network connectivity. *Neuroimage*, 91, 282-299.
<http://dx.doi.org/10.1016/j.neuroimage.2013.12.066>
- Brunel, N., and Wang, X. (2001). Effects of neuromodulation in a cortical network model of object working memory dominated by recurrent inhibition. *Journal Of Computational Neuroscience*, 37(3), 593-594. <http://dx.doi.org/10.1007/s10827-014-0506-8>
- Brunel, N., and Wang, X. (2003). What Determines the Frequency of Fast Network Oscillations With Irregular Neural Discharges? I. Synaptic Dynamics and Excitation-Inhibition Balance. *Journal Of Neurophysiology*, 90(1), 415-430.
<http://dx.doi.org/10.1152/jn.01095.2002>
- Buckner, R., Andrews-Hanna, J., and Schacter, D. (2008). The Brain's Default Network:

- Anatomy, Function, and Relevance to Disease. *Annals Of The New York Academy Of Sciences*, 1124(1), 1-38. <http://dx.doi.org/10.1196/annals.1440.011>
- Buckner, R., Bandettini, P., O'Craven, K., Savoy, R., Petersen, S., Raichle, M., and Rosen, B. (1996). Detection of cortical activation during averaged single trials of a cognitive task using functional magnetic resonance imaging. *Proceedings Of The National Academy Of Sciences*, 93(25), 14878-14883. <http://dx.doi.org/10.1073/pnas.93.25.14878>
- Bullmore, E., and Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nature Reviews Neuroscience*, 10(3), 186-198. <http://dx.doi.org/10.1038/nrn2575>
- Bullmore, E., Brammer, M., Williams, S., Rabe-Hesketh, S., Janot, N., and David, A. et al. (1996). Statistical methods of estimation and inference for functional MR image analysis. *Magnetic Resonance In Medicine*, 35(2), 261-277. <http://dx.doi.org/10.1002/mrm.1910350219>
- Buxton, R., Wong, E., and Frank, L. (1998). Dynamics of blood flow and oxygenation changes during brain activation: The balloon model. *Magnetic Resonance In Medicine*, 39(6), 855-864. <http://dx.doi.org/10.1002/mrm.1910390602>
- Cabral, J. (2012). *Brain activity during rest* (Ph.D.). Universitat Pompeu Fabra. Barcelona.
- Cabral, J., Hugues, E., Sporns, O., and Deco, G. (2011). Role of local network oscillations in resting-state functional connectivity. *Neuroimage*, 57(1), 130-139. <http://dx.doi.org/10.1016/j.neuroimage.2011.04.010>
- Cabral, J., Luckhoo, H., Woolrich, M., Joensson, M., Mohseni, H., and Baker, A. et al. (2014). 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. <http://dx.doi.org/10.1016/j.neuroimage.2013.11.047>
- Chang, C., and Glover, G. (2010). Time-frequency dynamics of resting-state brain connectivity measured with fMRI. *Neuroimage*, 50(1), 81-98. <http://dx.doi.org/10.1016/j.neuroimage.2009.12.011>

- Chang, C., Thomason, M., and Glover, G. (2008). Mapping and correction of vascular hemodynamic latency in the BOLD signal. *Neuroimage*, 43(1), 90-102. <http://dx.doi.org/10.1016/j.neuroimage.2008.06.030>
- Chickering, D., and Boutilier, C. (2002). Optimal structure identification with greedy search. *Journal Of Machine Learning Research*, 3, 507--554. <http://dx.doi.org/10.1.1.80.9840>
- Collins, D., Neelin, P., Peters, T., and Evans, A. (1994). Automatic 3D Intersubject Registration of MR Volumetric Data in Standardized Talairach Space. *Journal Of Computer Assisted Tomography*, 18(2), 192-205. <http://dx.doi.org/10.1097/00004728-199403000-00005>
- Cordesa, D., Haughtona, V., Arfanakisa, K., Carewa, J., Turskia, P., and Moritza, C. et al. (2001). Frequencies contributing to functional connectivity in the cerebral cortex in "resting-state" data. *American Journal Of Neuroradiology*, 22(7), 1326-1333. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/11498421>
- Crick, F., and Koch, C. (1998). Constraints on cortical and thalamic projections: the no-strong-loops hypothesis. *Nature*, 391(6664), 245-250. <http://dx.doi.org/10.1038/34584>
- Damoiseaux, J., and Greicius, M. (2009). Greater than the sum of its parts: a review of studies combining structural connectivity and resting-state functional connectivity. *Brain Structure And Function*, 213(6), 525-533. <http://dx.doi.org/10.1007/s00429-009-0208-6>
- Damoiseaux, J., Rombouts, S., Barkhof, F., Scheltens, P., Stam, C., Smith, S., and Beckmann, C. (2006). Consistent resting-state networks across healthy subjects. *Proceedings Of The National Academy Of Sciences*, 103(37), 13848-13853. <http://dx.doi.org/10.1073/pnas.0601417103>
- Dauwels, J., Vialatte, F., Musha, T., and Cichocki, A. (2010). A comparative study of synchrony measures for the early diagnosis of Alzheimer's disease based on EEG. *Neuroimage*, 49(1), 668-693. <http://dx.doi.org/10.1016/j.neuroimage.2009.06.056>
- David, O., Kiebel, S., Harrison, L., Mattout, J., Kilner, J., and Friston, K. (2006). Dynamic

- causal modeling of evoked responses in EEG and MEG. *Neuroimage*, 30(4), 1255-1272. <http://dx.doi.org/10.1016/j.neuroimage.2005.10.045>
- De Luca, M., Beckmann, C., De Stefano, N., Matthews, P., and Smith, S. (2006). fMRI resting state networks define distinct modes of long-distance interactions in the human brain. *Neuroimage*, 29(4), 1359-1367. <http://dx.doi.org/10.1016/j.neuroimage.2005.08.035>
- de Pasquale, F., Della Penna, S., Snyder, A., Lewis, C., Mantini, D., and Marzetti, L. et al. (2010). Temporal dynamics of spontaneous MEG activity in brain networks. *Proceedings Of The National Academy Of Sciences*, 107(13), 6040-6045. <http://dx.doi.org/10.1073/pnas.0913863107>
- de Pasquale, F., Della Penna, S., Snyder, A., Marzetti, L., Pizzella, V., Romani, G., and Corbetta, M. (2012). A Cortical Core for Dynamic Integration of Functional Networks in the Resting Human Brain. *Neuron*, 74(4), 753-764. <http://dx.doi.org/10.1016/j.neuron.2012.03.031>
- Deco, G., and Jirsa, V. (2012). Ongoing Cortical Activity at Rest: Criticality, Multistability, and Ghost Attractors. *Journal Of Neuroscience*, 32(10), 3366-3375. <http://dx.doi.org/10.1523/jneurosci.2523-11.2012>
- Deco, G., Jirsa, V., and McIntosh, A. (2011). Emerging concepts for the dynamical organization of resting-state activity in the brain. *Nature Reviews Neuroscience*, 12(1), 43-56. <http://dx.doi.org/10.1038/nrn2961>
- Deco, G., Jirsa, V., McIntosh, A., Sporns, O., and Kötter, R. (2009). Key role of coupling, delay, and noise in resting brain fluctuations. *Proceedings Of The National Academy Of Sciences*, 106(25), 10302-10307. <http://dx.doi.org/10.1073/pnas.0901831106>
- Deco, G., Jirsa, V., Robinson, P., Breakspear, M., and Friston, K. (2008). The Dynamic Brain: From Spiking Neurons to Neural Masses and Cortical Fields. *Plos Comput Biol*, 4(8), e1000092. <http://dx.doi.org/10.1371/journal.pcbi.1000092>
- Felleman, D., and Van Essen, D. (1991). Distributed Hierarchical Processing in the Primate Cerebral Cortex. *Cerebral Cortex*, 1(1), 1-47. <http://dx.doi.org/10.1093/cercor/1.1.1>

- Filippini, N., MacIntosh, B., Hough, M., Goodwin, G., Frisoni, G., and Smith, S. et al. (2009). Distinct patterns of brain activity in young carriers of the APOE- ϵ 4 allele. *Proceedings Of The National Academy Of Sciences*, 106(17), 7209-7214. <http://dx.doi.org/10.1073/pnas.0811879106>
- FitzHugh, R. (1961). Impulses and Physiological States in Theoretical Models of Nerve Membrane. *Biophysical Journal*, 1(6), 445-466. [http://dx.doi.org/10.1016/s0006-3495\(61\)86902-6](http://dx.doi.org/10.1016/s0006-3495(61)86902-6)
- Fox, M., and Raichle, M. (2007). Spontaneous fluctuations in brain activity observed with functional magnetic resonance imaging. *Nature Reviews Neuroscience*, 8(9), 700-711. <http://dx.doi.org/10.1038/nrn2201>
- Fox, M., Snyder, A., Vincent, J., Corbetta, M., Van Essen, D., and Raichle, M. (2005). From The Cover: The human brain is intrinsically organized into dynamic, anticorrelated functional networks. *Proceedings Of The National Academy Of Sciences*, 102(27), 9673-9678. <http://dx.doi.org/10.1073/pnas.0504136102>
- Frank, T., Daffertshofer, A., and Beek, P. (2000). Multivariate Ornstein-Uhlenbeck processes with mean-field dependent coefficients: Application to postural sway. *Physical Review E*, 63(1). <http://dx.doi.org/10.1103/physreve.63.011905>
- Frank, T., Daffertshofer, A., Beek, P., and Haken, H. (1999). Impacts of noise on a field theoretical model of the human brain. *Physica D: Nonlinear Phenomena*, 127(3-4), 233-249. [http://dx.doi.org/10.1016/s0167-2789\(98\)00294-2](http://dx.doi.org/10.1016/s0167-2789(98)00294-2)
- Freyer, F., Roberts, J., Becker, R., Robinson, P., Ritter, P., and Breakspear, M. (2011). Biophysical Mechanisms of Multistability in Resting-State Cortical Rhythms. *Journal Of Neuroscience*, 31(17), 6353-6361. <http://dx.doi.org/10.1523/jneurosci.6693-10.2011>
- Friedman, J., Hastie, T., and Tibshirani, R. (2007). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3), 432-441. <http://dx.doi.org/10.1093/biostatistics/kxm045>
- Friston, K. (1994). Functional and effective connectivity in neuroimaging: A synthesis. *Human Brain Mapping*, 2(1-2), 56-78. <http://dx.doi.org/10.1002/hbm.460020107>

- Friston, K., Frith, C., and Frackowiak, R. (1993a). Time-dependent changes in effective connectivity measured with PET. *Human Brain Mapping*, 1(1), 69-79. <http://dx.doi.org/10.1002/hbm.460010108>
- Friston, K., Frith, C., Frackowiak, R., and Turner, R. (1995). Characterizing Dynamic Brain Responses with fMRI: A Multivariate Approach. *Neuroimage*, 2(2), 166-172. <http://dx.doi.org/10.1006/nimg.1995.1019>
- Friston, K., Frith, C., Liddle, P., and Frackowiak, R. (1993b). Functional Connectivity: The Principal-Component Analysis of Large (PET) Data Sets. *Journal Of Cerebral Blood Flow and Metabolism*, 13(1), 5-14. <http://dx.doi.org/10.1038/jcbfm.1993.4>
- Friston, K., Harrison, L., and Penny, W. (2003). Dynamic causal modelling. *Neuroimage*, 19(4), 1273-1302. [http://dx.doi.org/10.1016/s1053-8119\(03\)00202-7](http://dx.doi.org/10.1016/s1053-8119(03)00202-7)
- Geweke, J. (1984). Measures of Conditional Linear Dependence and Feedback Between Time Series. *Journal Of The American Statistical Association*, 79(388), 907. <http://dx.doi.org/10.2307/2288723>
- Ghosh, A., Rho, Y., McIntosh, A., Kötter, R., and Jirsa, V. (2008). Noise during Rest Enables the Exploration of the Brain's Dynamic Repertoire. *Plos Comput Biol*, 4(10), e1000196. <http://dx.doi.org/10.1371/journal.pcbi.1000196>
- Granger, C. (1969). Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica*, 37(3), 424. <http://dx.doi.org/10.2307/1912791>
- Greicius, M., Krasnow, B., Reiss, A., and Menon, V. (2003). Functional connectivity in the resting brain: A network analysis of the default mode hypothesis. *Proceedings Of The National Academy Of Sciences*, 100(1), 253-258. <http://dx.doi.org/10.1073/pnas.0135058100>
- Greicius, M., Srivastava, G., Reiss, A., and 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*, 101(13), 4637-4642. <http://dx.doi.org/10.1073/pnas.0308627101>
- Grinsted, A., Moore, J., and Jevrejeva, S. (2004). Application of the cross wavelet transform and wavelet coherence to geophysical time series. *Nonlin. Processes*

- Geophys.*, 11(5/6), 561-566. <http://dx.doi.org/10.5194/npg-11-561-2004>
- Gross, J., Kujala, J., Hamalainen, M., Timmermann, L., Schnitzler, A., and Salmelin, R. (2001). Dynamic imaging of coherent sources: Studying neural interactions in the human brain. *Proceedings Of The National Academy Of Sciences*, 98(2), 694-699. <http://dx.doi.org/10.1073/pnas.98.2.694>
- Guo, S., Seth, A., Kendrick, K., Zhou, C., and Feng, J. (2008). Partial Granger causality, Eliminating exogenous inputs and latent variables. *Journal Of Neuroscience Methods*, 172(1), 79-93. <http://dx.doi.org/10.1016/j.jneumeth.2008.04.011>
- Hadjipapas, A., Hillebrand, A., Holliday, I., Singh, K., and Barnes, G. (2004). Assessing interactions of linear and nonlinear neuronal sources using MEG beamformers: a proof of concept. *Clinical Neurophysiology*, 116(6), 1300-1313. <http://dx.doi.org/10.1016/j.clinph.2005.01.014>
- Hagmann, P. (2005). *From diffusion MRI to brain connectomics* (Ph.D.). Ecole Polytechnique Fédérale de Lausanne (EPFL).
- Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C., Wedeen, V., and Sporns, O. (2008). Mapping the Structural Core of Human Cerebral Cortex. *Plos Biology*, 6(7), e159. <http://dx.doi.org/10.1371/journal.pbio.0060159>
- Hagmann, P., Meuli, R., Thiran, J., Kurant, M., Gigandet, X., Thiran, P., and Wedeen, V. (2007). Mapping Human Whole-Brain Structural Networks with Diffusion MRI. *Plos ONE*, 2(7), e597. <http://dx.doi.org/10.1371/journal.pone.0000597>
- Handwerker, D., Ollinger, J., and D'Esposito, M. (2004). Variation of BOLD hemodynamic responses across subjects and brain regions and their effects on statistical analyses. *Neuroimage*, 21(4), 1639-1651. <http://dx.doi.org/10.1016/j.neuroimage.2003.11.029>
- Hawellek, D., Hipp, J., Lewis, C., Corbetta, M., and Engel, A. (2011). Increased functional connectivity indicates the severity of cognitive impairment in multiple sclerosis. *Proceedings Of The National Academy Of Sciences*, 108(47), 19066-19071. <http://dx.doi.org/10.1073/pnas.1110024108>

- Heeger, D., and Ress, D. (2002). What does fMRI tell us about neuronal activity?. *Nature Reviews Neuroscience*, 3(2), 142-151. <http://dx.doi.org/10.1038/nrn730>
- Hillebrand, A., Barnes, G., Bosboom, J., Berendse, H., and Stam, C. (2012). Frequency-dependent functional connectivity within resting-state networks: An atlas-based MEG beamformer solution. *Neuroimage*, 59(4), 3909-3921. <http://dx.doi.org/10.1016/j.neuroimage.2011.11.005>
- Hipp, J., Hawellek, D., Corbetta, M., Siegel, M., and Engel, A. (2012). Large-scale cortical correlation structure of spontaneous oscillatory activity. *Nature Neuroscience*, 15(6), 884-890. <http://dx.doi.org/10.1038/nn.3101>
- Hlinka, J., Paluš, M., Vejmelka, M., Mantini, D., and Corbetta, M. (2011). Functional connectivity in resting-state fMRI: Is linear correlation sufficient?. *Neuroimage*, 54(3), 2218-2225. <http://dx.doi.org/10.1016/j.neuroimage.2010.08.042>
- Hodgkin, A., and Huxley, A. (1952). A quantitative description of membrane current and its application to conduction and excitation in nerve. *The Journal Of Physiology*, 117(4), 500-544. <http://dx.doi.org/10.1113/jphysiol.1952.sp004764>
- Honey, C., Kötter, R., Breakspear, M., and Sporns, O. (2007). Network structure of cerebral cortex shapes functional connectivity on multiple time scales. *Proceedings Of The National Academy Of Sciences*, 104(24), 10240-10245. <http://dx.doi.org/10.1073/pnas.0701519104>
- Honey, C., Sporns, O., Cammoun, L., Gigandet, X., Thiran, J., Meuli, R., and Hagmann, P. (2009). Predicting human resting-state functional connectivity from structural connectivity. *Proceedings Of The National Academy Of Sciences*, 106(6), 2035-2040. <http://dx.doi.org/10.1073/pnas.0811168106>
- Huang, M., Mosher, J., and Leahy, R. (1999). A sensor-weighted overlapping-sphere head model and exhaustive head model comparison for MEG. *Physics In Medicine And Biology*, 44(2), 423-440. <http://dx.doi.org/10.1088/0031-9155/44/2/010>
- Huettel, S., Song, A., and McCarthy, G. (2004). *Functional magnetic resonance imaging*. Sunderland, Mass.: Sinauer Associates, Publishers.
- Hyvärinen, A. (1999). Fast and robust fixed-point algorithms for independent component

- analysis. *IEEE Trans. Neural Netw.*, 10(3), 626-634.
<http://dx.doi.org/10.1109/72.761722>
- Izhikevich, E., and Edelman, G. (2008). Large-scale model of mammalian thalamocortical systems. *Proceedings Of The National Academy Of Sciences*, 105(9), 3593-3598.
<http://dx.doi.org/10.1073/pnas.0712231105>
- Jenkinson, M., Bannister, P., Brady, M., and Smith, S. (2002). Improved Optimization for the Robust and Accurate Linear Registration and Motion Correction of Brain Images. *Neuroimage*, 17(2), 825-841. <http://dx.doi.org/10.1006/nimg.2002.1132>
- Jensen, O., Colgin, LL. (2007). Cross-frequency coupling between neuronal oscillations. *Trends in Cognitive Science* 11, 267–269.
- Jirsa, V., and Kelso, J. (2000). Spatiotemporal pattern formation in neural systems with heterogeneous connection topologies. *Physical Review E*, 62(6), 8462-8465.
<http://dx.doi.org/10.1103/physreve.62.8462>
- Kahn, I., Andrews-Hanna, J., Vincent, J., Snyder, A., and Buckner, R. (2008). Distinct Cortical Anatomy Linked to Subregions of the Medial Temporal Lobe Revealed by Intrinsic Functional Connectivity. *Journal Of Neurophysiology*, 100(1), 129-139.
<http://dx.doi.org/10.1152/jn.00077.2008>
- Kamiński, M., Blinowska, K., and Szelenberger, W. (1997). Topographic analysis of coherence and propagation of EEG activity during sleep and wakefulness. *Electroencephalography And Clinical Neurophysiology*, 102(3), 216-227.
[http://dx.doi.org/10.1016/s0013-4694\(96\)95721-5](http://dx.doi.org/10.1016/s0013-4694(96)95721-5)
- Kelso, J. (2012). Multistability and metastability: understanding dynamic coordination in the brain. *Philosophical Transactions Of The Royal Society B: Biological Sciences*, 367(1591), 906-918. <http://dx.doi.org/10.1098/rstb.2011.0351>
- Kiviniemi, V., Kantola, J., Jauhiainen, J., Hyvärinen, A., and Tervonen, O. (2003). Independent component analysis of nondeterministic fMRI signal sources. *Neuroimage*, 19(2), 253-260. [http://dx.doi.org/10.1016/s1053-8119\(03\)00097-1](http://dx.doi.org/10.1016/s1053-8119(03)00097-1)
- Kiviniemi, V., Starck, T., Remes, J., Long, X., Nikkinen, J., and Haapea, M. et al. (2009). Functional segmentation of the brain cortex using high model order group PICA.

- Human Brain Mapping*, 30(12), 3865-3886. <http://dx.doi.org/10.1002/hbm.20813>
- Kiviniemi V, Vire T, Remes J, Abou elseoud 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*. 1:339–347.
- Koch, M., Norris, D., and Hund-Georgiadis, M. (2002). An Investigation of Functional and Anatomical Connectivity Using Magnetic Resonance Imaging. *Neuroimage*, 16(1), 241-250. <http://dx.doi.org/10.1006/nimg.2001.1052>
- Kötter, R. (2004). Online Retrieval, Processing, and Visualization of Primate Connectivity Data From the CoCoMac Database. *NI*, 2(2), 127-144. <http://dx.doi.org/10.1385/ni:2:2:127>
- Kötter, R., and Wanke, E. (2005). Mapping brains without coordinates. *Philosophical Transactions Of The Royal Society B: Biological Sciences*, 360(1456), 751-766. <http://dx.doi.org/10.1098/rstb.2005.1625>
- Kuramoto, Y. (1984). *Chemical oscillations, waves, and turbulence*. Berlin: Springer.
- Lachaux, J., Lutz, A., Rudrauf, D., Cosmelli, D., Le Van Quyen, M., Martinerie, J., and Varela, F. (2002). Estimating the time-course of coherence between single-trial brain signals: an introduction to wavelet coherence. *Neurophysiologie Clinique/Clinical Neurophysiology*, 32(3), 157-174. [http://dx.doi.org/10.1016/s0987-7053\(02\)00301-5](http://dx.doi.org/10.1016/s0987-7053(02)00301-5)
- Laufs, H., Krakow, K., Sterzer, P., Eger, E., Beyerle, A., Salek-Haddadi, A., and 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*, 100(19), 11053-11058. <http://dx.doi.org/10.1073/pnas.1831638100>
- Le Van Quyen, M., Foucher, J., Lachaux, J., Rodriguez, E., Lutz, A., Martinerie, J., and Varela, F. (2001). Comparison of Hilbert transform and wavelet methods for the analysis of neuronal synchrony. *Journal Of Neuroscience Methods*, 111(2), 83-98. [http://dx.doi.org/10.1016/s0165-0270\(01\)00372-7](http://dx.doi.org/10.1016/s0165-0270(01)00372-7)
- Lee, D., and Wang, X. (2008). Mechanisms for Stochastic Decision Making in the Primate Frontal Cortex: Single-neuron. Recording and Circuit Modeling. In P. Glimcher, C.

- Camerer, E. Fehr and R. Poldrack, *Neuroeconomics: Decision Making and the Brain* (1st ed.). Academic Press. Retrieved from http://www.cns.nyu.edu/wanglab/publications/pdf/neuroeconomics_chap.pdf
- Lee, W., Ott, E., and Antonsen, T. (2009). Large Coupled Oscillator Systems with Heterogeneous Interaction Delays. *Physical Review Letters*, 103(4). <http://dx.doi.org/10.1103/physrevlett.103.044101>
- Leopold, D. (2003). Very Slow Activity Fluctuations in Monkey Visual Cortex: Implications for Functional Brain Imaging. *Cerebral Cortex*, 13(4), 422-433. <http://dx.doi.org/10.1093/cercor/13.4.422>
- Liu, Y., Liang, M., Zhou, Y., He, Y., Hao, Y., and Song, M. et al. (2008). Disrupted small-world networks in schizophrenia. *Brain*, 131(4), 945-961. <http://dx.doi.org/10.1093/brain/awn018>
- Liu, Z., Fukunaga, M., de Zwart, J., and Duyn, J. (2010). Large-scale spontaneous fluctuations and correlations in brain electrical activity observed with magnetoencephalography. *Neuroimage*, 51(1), 102-111. <http://dx.doi.org/10.1016/j.neuroimage.2010.01.092>
- Logothetis, N., Pauls, J., Augath, M., Trinath, T., and Oeltermann, A. (2001). Neurophysiological investigation of the basis of the fMRI signal. *Nature*, 412(6843), 150-157. <http://dx.doi.org/10.1038/35084005>
- Lopes da Silva, F., Hoeks, A., Smits, H., and Zetterberg, L. (1974). Model of brain rhythmic activity. *Kybernetik*, 15(1), 27-37. <http://dx.doi.org/10.1007/bf00270757>
- Lu, H., Zou, Q., Gu, H., Raichle, M., Stein, E., and Yang, Y. (2012). Rat brains also have a default mode network. *Proceedings Of The National Academy Of Sciences*, 109(10), 3979-3984. <http://dx.doi.org/10.1073/pnas.1200506109>
- Luckhoo, H., Hale, J., Stokes, M., Nobre, A., Morris, P., Brookes, M., and Woolrich, M. (2012). Inferring task-related networks using independent component analysis in magnetoencephalography. *Neuroimage*, 62(1), 530-541. <http://dx.doi.org/10.1016/j.neuroimage.2012.04.046>
- Lynall, M., Bassett, D., Kerwin, R., McKenna, P., Kitzbichler, M., Muller, U., and

- Bullmore, E. (2010). Functional Connectivity and Brain Networks in Schizophrenia. *Journal Of Neuroscience*, 30(28), 9477-9487. <http://dx.doi.org/10.1523/jneurosci.0333-10.2010>
- Mantini, D., Penna, S., Marzetti, L., de Pasquale, F., Pizzella, V., Corbetta, M., and Romani, G. (2011). A Signal-Processing Pipeline for Magnetoencephalography Resting-State Networks. *Brain Connectivity*, 1(1), 49-59. <http://dx.doi.org/10.1089/brain.2011.0001>
- Mantini, D., Perrucci, M., Del Gratta, C., Romani, G., and Corbetta, M. (2007). Electrophysiological signatures of resting state networks in the human brain. *Proceedings Of The National Academy Of Sciences*, 104(32), 13170-13175. <http://dx.doi.org/10.1073/pnas.0700668104>
- Markram, H. (2006). The Blue Brain Project. *Nature Reviews Neuroscience*, 7(2), 153-160. <http://dx.doi.org/10.1038/nrn1848>
- Markram, H. (2012). The Human Brain Project. *Scientific American*, 306, 50-55. <http://dx.doi.org/10.1038/scientificamerican0612-50>
- Marreiros, A., Daunizeau, J., Kiebel, S., and Friston, K. (2008). Population dynamics: Variance and the sigmoid activation function. *Neuroimage*, 42(1), 147-157. <http://dx.doi.org/10.1016/j.neuroimage.2008.04.239>
- Marreiros, A., Kiebel, S., and Friston, K. (2010). A dynamic causal model study of neuronal population dynamics. *Neuroimage*, 51(1), 91-101. <http://dx.doi.org/10.1016/j.neuroimage.2010.01.098>
- Marrelec, G., Krainik, A., Duffau, H., Pélérini-Issac, M., Lehericy, S., Doyon, J., and Benali, H. (2006). Partial correlation for functional brain interactivity investigation in functional MRI. *Neuroimage*, 32(1), 228-237. <http://dx.doi.org/10.1016/j.neuroimage.2005.12.057>
- Meek, C. (1995). Causal inference and causal explanation with background knowledge. *Proceedings Of The 11Th Annual Conference On Uncertainty In Artificial Intelligence*, 403-410. Retrieved from <http://dl.acm.org/citation.cfm?id=2074204>
- Morcom, A., and Fletcher, P. (2007). Does the brain have a baseline? Why we should be

- resisting a rest. *Neuroimage*, 37(4), 1073-1082.
<http://dx.doi.org/10.1016/j.neuroimage.2006.09.013>
- Morris, C., and Lecar, H. (1981). Voltage oscillations in the barnacle giant muscle fiber. *Biophysical Journal*, 35(1), 193-213. [http://dx.doi.org/10.1016/s0006-3495\(81\)84782-0](http://dx.doi.org/10.1016/s0006-3495(81)84782-0)
- Muthukumaraswamy, S. (2013). High-frequency brain activity and muscle artifacts in MEG/EEG: a review and recommendations. *Frontiers In Human Neuroscience*, 7. <http://dx.doi.org/10.3389/fnhum.2013.00138>
- Nagumo, J., Arimoto, S., and Yoshizawa, S. (1962). An Active Pulse Transmission Line Simulating Nerve Axon. *Proceedings Of The IRE*, 50(10), 2061-2070. <http://dx.doi.org/10.1109/jrproc.1962.288235>
- Nalatore, H., Ding, M., and Rangarajan, G. (2007). Mitigating the effects of measurement noise on Granger causality. *Physical Review E*, 75(3). <http://dx.doi.org/10.1103/physreve.75.031123>
- Niebur, E., Schuster, H., and Kammen, D. (1991). Collective frequencies and metastability in networks of limit-cycle oscillators with time delay. *Physical Review Letters*, 67(20), 2753-2756. <http://dx.doi.org/10.1103/physrevlett.67.2753>
- Niedermeyer, E., and Lopes da Silva, F. (1982) *Electroencephalography: basic principles, clinical applications, and related fields*. Urban and Schwarzenberg.
- Nikouline, V., Linkenkaer-Hansen, K., Huttunen, J., and Ilmoniemi, R. (2001). Interhemispheric phase synchrony and amplitude correlation of spontaneous beta oscillations in human subjects: a magnetoencephalographic study. *Neuroreport*, 12(11), 2487-2491. <http://dx.doi.org/10.1097/00001756-200108080-00040>
- Nir, Y., Mukamel, R., Dinstein, I., Privman, E., Harel, M., and Fisch, L. et al. (2008). Interhemispheric correlations of slow spontaneous neuronal fluctuations revealed in human sensory cortex. *Nature Neuroscience*, 11(9), 1100-1108. <http://dx.doi.org/10.1038/nn.2177>
- Nolte, G., Bai, O., Wheaton, L., Mari, Z., Vorbach, S., and Hallett, M. (2004). Identifying true brain interaction from EEG data using the imaginary part of coherency. *Clinical*

- Neurophysiology*, 115(10), 2292-2307. <http://dx.doi.org/10.1016/j.clinph.2004.04.029>
- Nolte, G., Ziehe, A., Nikulin, V., Schlögl, A., Krämer, N., Brismar, T., and Müller, K. (2008). Robustly Estimating the Flow Direction of Information in Complex Physical Systems. *Physical Review Letters*, 100(23). <http://dx.doi.org/10.1103/physrevlett.100.234101>
- Nunez, P. (2000). Toward a quantitative description of large-scale neocortical dynamic function and EEG. *Behavioral And Brain Sciences*, 23(3), 371-398. <http://dx.doi.org/10.1017/s0140525x00003253>
- Oberauer, K., and Lewandowsky, S. (2010). Modeling working memory: a computational implementation of the Time-Based Resource-Sharing theory. *Psychonomic Bulletin and Review*, 18(1), 10-45. <http://dx.doi.org/10.3758/s13423-010-0020-6>
- Ogawa, S., Lee, T., Kay, A., and Tank, D. (1990). Brain magnetic resonance imaging with contrast dependent on blood oxygenation. *Proceedings Of The National Academy Of Sciences*, 87(24), 9868-9872. <http://dx.doi.org/10.1073/pnas.87.24.9868>
- Olbrich, E., Achermann, P., and Wennekers, T. (2011). The sleeping brain as a complex system. *Philosophical Transactions Of The Royal Society A: Mathematical, Physical And Engineering Sciences*, 369(1952), 3697-3707. <http://dx.doi.org/10.1098/rsta.2011.0199>
- O'Neill, G.C., Bauer, M., Woolrich, M.W., Morris, P.G., Barnes, G.R., Brookes, M.J., (2015). Dynamic recruitment of resting state sub-networks. *NeuroImage* 115, 85–95
- Palva, S., and Palva, J. (2012). Discovering oscillatory interaction networks with M/EEG: challenges and breakthroughs. *Trends In Cognitive Sciences*, 16(4), 219-230. <http://dx.doi.org/10.1016/j.tics.2012.02.004>
- Patel, R., Bowman, F., and Rilling, J. (2006). A Bayesian approach to determining connectivity of the human brain. *Human Brain Mapping*, 27(3), 267-276. <http://dx.doi.org/10.1002/hbm.20182>
- Prichard, D., and Theiler, J. (1994). Generating surrogate data for time series with several simultaneously measured variables. *Physical Review Letters*, 73(7), 951-954. <http://dx.doi.org/10.1103/physrevlett.73.951>

- Quian Quiroga, R., Kraskov, A., Kreuz, T., and Grassberger, P. (2002). Performance of different synchronization measures in real data: A case study on electroencephalographic signals. *Physical Review E*, 65(4). <http://dx.doi.org/10.1103/physreve.65.041903>
- Raichle, M., MacLeod, A., Snyder, A., Powers, W., Gusnard, D., and Shulman, G. (2001). A default mode of brain function. *Proceedings Of The National Academy Of Sciences*, 98(2), 676-682. <http://dx.doi.org/10.1073/pnas.98.2.676>
- Ramsey, J., Hanson, S., Hanson, C., Halchenko, Y., Poldrack, R., and Glymour, C. (2010). Six problems for causal inference from fMRI. *Neuroimage*, 49(2), 1545-1558. <http://dx.doi.org/10.1016/j.neuroimage.2009.08.065>
- Ramsey, J., Zhang, J., and Spirtes, P. (2006). Adjacency-faithfulness and conservative causal inference. *Proceedings Of The 22Nd Annual Conference On Uncertainty In Artificial Intelligence*, 401-408. Retrieved from <http://arxiv.org/abs/1206.6843>
- Richardson, T., and Spirtes, P. (2001). Automated discovery of linear feedback models. In C. Glymour and G. Cooper, *Computation, Causation, and Causality* (1st ed.). MIT Press. Retrieved from <http://www.stat.washington.edu/tsr/papers/autodisc.html>
- Rickles, D., Hawe, P., and Shiell, A. (2007). A simple guide to chaos and complexity. *Journal Of Epidemiology and Community Health*, 61(11), 933-937. <http://dx.doi.org/10.1136/jech.2006.054254>
- Rosenblatt, F. (1958). The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6), 386-408. <http://dx.doi.org/10.1037/h0042519>
- Rosenblum, M., Pikovsky, A., and Kurths, J. (1996). Phase Synchronization of Chaotic Oscillators. *Physical Review Letters*, 76(11), 1804-1807. <http://dx.doi.org/10.1103/physrevlett.76.1804>
- Rubinov, M., and Sporns, O. (2010). Complex network measures of brain connectivity: uses and interpretations. *Neuroimage*, 52, 1059–1069.
- Sarvas, J. (1987) Basic mathematical and electromagnetic concepts of the biomagnetic inverse problem. *Phys Med Biol* 32, 11–22.

- Saxe, R., and Kanwisher, N. (2003). People thinking about thinking people: The role of the temporo-parietal junction in "theory of mind". *Neuroimage*, 19(4), 1835-1842. [http://dx.doi.org/10.1016/s1053-8119\(03\)00230-1](http://dx.doi.org/10.1016/s1053-8119(03)00230-1)
- Scheeringa, R., Fries, P., Petersson, K., Oostenveld, R., Grothe, I., and Norris, D. et al. (2011). Neuronal Dynamics Underlying High- and Low-Frequency EEG Oscillations Contribute Independently to the Human BOLD Signal. *Neuron*, 69(3), 572-583. <http://dx.doi.org/10.1016/j.neuron.2010.11.044>
- Schnitzler, A., and Gross, J. (2005). Normal and pathological oscillatory communication in the brain. *Nature Reviews Neuroscience*, 6(4), 285-296. <http://dx.doi.org/10.1038/nrn1650>
- Schoffelen, J., and Gross, J. (2009). Source connectivity analysis with MEG and EEG. *Human Brain Mapping*, 30(6), 1857-1865. <http://dx.doi.org/10.1002/hbm.20745>
- Schreckenberger, M., Lange-Asschenfeld, C., Lochmann, M., Mann, K., Siessmeier, T., and Buchholz, H. et al. (2004). The thalamus as the generator and modulator of EEG alpha rhythm: a combined PET/EEG study with lorazepam challenge in humans. *Neuroimage*, 22(2), 637-644. <http://dx.doi.org/10.1016/j.neuroimage.2004.01.047>
- Shanahan, M. (2010). Metastable chimera states in community-structured oscillator networks. *Chaos*, 20(1), 013108. <http://dx.doi.org/10.1063/1.3305451>
- Shannon, C. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, 27(4), 623-656. <http://dx.doi.org/10.1002/j.1538-7305.1948.tb00917.x>
- Shimizu, S., Hoyer, P., Hyvärinen, A., and Kerminen, A. (2006). A Linear Non-Gaussian Acyclic Model for Causal Discovery. *Journal Of Machine Learning Res*, 7, 2003-2030. Retrieved from <http://www.cs.helsinki.fi/u/ahyvarin/papers/JMLR06.pdf>
- Shmuel, A., and Leopold, D. (2008). Neuronal correlates of spontaneous fluctuations in fMRI signals in monkey visual cortex: Implications for functional connectivity at rest. *Human Brain Mapping*, 29(7), 751-761. <http://dx.doi.org/10.1002/hbm.20580>
- Siegel, M., Donner, T., and Engel, A. (2012). Spectral fingerprints of large-scale neuronal interactions. *Nature Reviews Neuroscience*. <http://dx.doi.org/10.1038/nrn3137>

- Siegelmann, H. (2010). Complex Systems Science and Brain Dynamics: A Frontiers in Computational Neuroscience Special Topic. *Frontiers in Computational Neuroscience*, 4. <http://dx.doi.org/10.3389/fncom.2010.00007>
- Smith, S., Fox, P., Miller, K., Glahn, D., Fox, P., and Mackay, C. et al. (2009). Correspondence of the brain's functional architecture during activation and rest. *Proceedings Of The National Academy Of Sciences*, 106(31), 13040-13045. <http://dx.doi.org/10.1073/pnas.0905267106>
- Smith, S., Miller, K., Moeller, S., Xu, J., Auerbach, E., and Woolrich, M. et al. (2012). Temporally-independent functional modes of spontaneous brain activity. *Proceedings Of The National Academy Of Sciences*, 109(8), 3131-3136. <http://dx.doi.org/10.1073/pnas.1121329109>
- Smith, S., Miller, K., Salimi-Khorshidi, G., Webster, M., Beckmann, C., and Nichols, T. et al. (2011). Network modelling methods for FMRI. *Neuroimage*, 54(2), 875-891. <http://dx.doi.org/10.1016/j.neuroimage.2010.08.063>
- Sporns, O., Tononi, G., and Kötter, R. (2005). The Human Connectome: A Structural Description of the Human Brain. *PLoS computational biology*, 1(4), e42. <http://dx.doi.org/10.1371/journal.pcbi.0010042>
- Stephan, K., Friston, K., and Frith, C. (2009). Dysconnection in Schizophrenia: From Abnormal Synaptic Plasticity to Failures of Self-monitoring. *Schizophrenia Bulletin*, 35(3), 509-527. <http://dx.doi.org/10.1093/schbul/sbn176>
- Stephan, K., Kamper, L., Bozkurt, A., Burns, G., Young, M., and Kötter, R. (2001). Advanced database methodology for the Collation of Connectivity data on the Macaque brain (CoCoMac). *Philosophical Transactions Of The Royal Society B: Biological Sciences*, 356(1412), 1159-1186. <http://dx.doi.org/10.1098/rstb.2001.0908>
- Strogatz, S. (2001). Exploring complex networks. *Nature*, 410(6825), 268-276. <http://dx.doi.org/10.1038/35065725>
- Taulu, S., Simola, J., and Kajola, M. (2005). Applications of the signal space separation method. *IEEE Transactions On Signal Processing*, 53(9), 3359-3372. <http://dx.doi.org/10.1109/tsp.2005.853302>

- Tiao, G., and Wei, W. (1976). Effect of Temporal Aggregation on the Dynamic Relationship of Two Time Series Variables. *Biometrika*, 63(3), 513. <http://dx.doi.org/10.2307/2335729>
- Tononi, G., Sporns, O., and Edelman, G. (1992). Reentry and the Problem of Integrating Multiple Cortical Areas: Simulation of Dynamic Integration in the Visual System. *Cerebral Cortex*, 2(4), 310-335. <http://dx.doi.org/10.1093/cercor/2.4.310>
- Tononi, G., Sporns, O., and Edelman, G. (1994). A measure for brain complexity: relating functional segregation and integration in the nervous system. *Proceedings Of The National Academy Of Sciences*, 91(11), 5033-5037. <http://dx.doi.org/10.1073/pnas.91.11.5033>
- Turcotte, D., and Rundle, J. (2002). Self-organized complexity in the physical, biological, and social sciences. *Proceedings Of The National Academy Of Sciences*, 99(Supplement 1), 2463-2465. <http://dx.doi.org/10.1073/pnas.012579399>
- Tzourio-Mazoyer, N., Landeau, B., Papathanassiou, D., Crivello, F., Etard, O., and Delcroix, N. et al. (2002). Automated Anatomical Labeling of Activations in SPM Using a Macroscopic Anatomical Parcellation of the MNI MRI Single-Subject Brain. *Neuroimage*, 15(1), 273-289. <http://dx.doi.org/10.1006/nimg.2001.0978>
- Uhlhaas, P., and Singer, W. (2010). Abnormal neural oscillations and synchrony in schizophrenia. *Nature Reviews Neuroscience*, 11(2), 100-113. <http://dx.doi.org/10.1038/nrn2774>
- Valentin, V., Maddox, W., and Ashby, F. (2014). A computational model of the temporal dynamics of plasticity in procedural learning: sensitivity to feedback timing. *Frontiers In Psychology*, 5. <http://dx.doi.org/10.3389/fpsyg.2014.00643>
- van den Heuvel, M., and Hulshoff Pol, H. (2010). Exploring the brain network: A review on resting-state fMRI functional connectivity. *European Neuropsychopharmacology*, 20(8), 519-534. <http://dx.doi.org/10.1016/j.euroneuro.2010.03.008>
- Van Veen, B., Van Drongelen, W., Yuchtman, M., and Suzuki, A. (1997). Localization of brain electrical activity via linearly constrained minimum variance spatial filtering. *IEEE Transactions On Biomedical Engineering*, 44(9), 867-880.

<http://dx.doi.org/10.1109/10.623056>

- Vincent, J., Patel, G., Fox, M., Snyder, A., Baker, J., and Van Essen, D. et al. (2007). Intrinsic functional architecture in the anaesthetized monkey brain. *Nature*, 447(7140), 83-86. <http://dx.doi.org/10.1038/nature05758>
- Wagner, A., Shannon, B., Kahn, I., and Buckner, R. (2005). Parietal lobe contributions to episodic memory retrieval. *Trends In Cognitive Sciences*, 9(9), 445-453. <http://dx.doi.org/10.1016/j.tics.2005.07.001>
- Wedeen, V., Davis, T., Weisskoff, R., Toottell, R., Rosen, B., and Belliveau, J. (1995). White matter connectivity explored by MRI. *Proceedings Of The First International Conference For Functional Mapping Of The Human Brain*, 36. Retrieved from <http://onlinelibrary.wiley.com/journal/10.1002/%28ISSN%291097-0193>
- Wedeen, V., Hagmann, P., Tseng, W., Reese, T., and Weisskoff, R. (2005). Mapping complex tissue architecture with diffusion spectrum magnetic resonance imaging. *Magnetic Resonance In Medicine*, 54(6), 1377-1386. <http://dx.doi.org/10.1002/mrm.20642>
- Wei, W. (1978). The effect of temporal aggregation on parameter estimation in distributed lag model. *Journal Of Econometrics*, 8(2), 237-246. [http://dx.doi.org/10.1016/0304-4076\(78\)90032-5](http://dx.doi.org/10.1016/0304-4076(78)90032-5)
- Weiss, A. (1984). Systematic sampling and temporal aggregation in time series models. *Journal Of Econometrics*, 26(3), 271-281. [http://dx.doi.org/10.1016/0304-4076\(84\)90022-8](http://dx.doi.org/10.1016/0304-4076(84)90022-8)
- Wilson, H., and Cowan, J. (1972). Excitatory and Inhibitory Interactions in Localized Populations of Model Neurons. *Biophysical Journal*, 12(1), 1-24. [http://dx.doi.org/10.1016/s0006-3495\(72\)86068-5](http://dx.doi.org/10.1016/s0006-3495(72)86068-5)
- Wilson, H., and Cowan, J. (1973). A mathematical theory of the functional dynamics of cortical and thalamic nervous tissue. *Kybernetik*, 13(2), 55-80. <http://dx.doi.org/10.1007/bf00288786>
- Winterer, G., Ziller, M., Dorn, H., Frick, K., Mulert, C., and Dahhan, N. et al. (1999). Cortical activation, signal-to-noise ratio and stochastic resonance during information

- processing in man. *Clinical Neurophysiology*, 110(7), 1193-1203.
[http://dx.doi.org/10.1016/s1388-2457\(99\)00059-0](http://dx.doi.org/10.1016/s1388-2457(99)00059-0)
- Wipf, D., Owen, J., Attias, H., Sekihara, K., and Nagarajan, S. (2010). Robust Bayesian estimation of the location, orientation, and time course of multiple correlated neural sources using MEG. *Neuroimage*, 49(1), 641-655.
<http://dx.doi.org/10.1016/j.neuroimage.2009.06.083>
- Womelsdorf, T., Schoffelen, J., Oostenveld, R., Singer, W., Desimone, R., Engel, A., and Fries, P. (2007). Modulation of Neuronal Interactions Through Neuronal Synchronization. *Science*, 316(5831), 1609-1612.
<http://dx.doi.org/10.1126/science.1139597>
- Zhang, J. (2008). On the completeness of orientation rules for causal discovery in the presence of latent confounders and selection bias. *Artificial Intelligence*, 172(16-17), 1873-1896. <http://dx.doi.org/10.1016/j.artint.2008.08.001>
- Zhou, D., Thompson, W., and Siegle, G. (2009). MATLAB toolbox for functional connectivity. *Neuroimage*, 47(4), 1590-1607.
<http://dx.doi.org/10.1016/j.neuroimage.2009.05.089>

Abbreviations

AAL	Automated Anatomical Labelling
AMPA	α -amino-3-hydroxy-5-methyl-4-isoazolepropionic acid
ANN	Artificial Neural Network
BLP	Band-Limited Power
BOLD	Blood-Oxygen-Level Dependent
CCD	Cyclic Causal Discovery
CoCoMac	Collation of Connectivity on the Macaque brain
CPC	Conservative PC
CPSD	Cross Power Spectral Density
DAN	Dorsal Attention Network
DC	Directed Granger
DCM	Dynamic Causal Model
DMN	Default Mode Network
DSI	Diffusion Spectrum imaging
DTF	Directed Transfer Function
DTI	Diffusion Tensor Imaging
ECoG	Electrocorticography
EEG	Electroencephalography
FC	Functional Connectivity
FCI	Fast Causal Inference
FDR	False Discovery Rate
FMRI	Functional Magnetic Resonance Imaging
FP	False Positives
GABA	gamma-Aminobutyric acid
GES	Greedy Equivalence Search
GGC	Geweke's Granger Causality
GMM	Gaussian Mixture Model
Hb	Haemoglobin
HMM	Hidden Markov Model
HRF	Haemodynamic Response Function

IC	Independent Component
ICA	Independent Component Analysis
ICOV	Regularised inverse covariance
IF	Integrate-and-Fire
LCMV	Linearly Constrained Minimum Variance
LFP	Local Field Potential
LiNGAM	Linear, Non-Gaussian, Acyclic causal Model
MEG	Magnetoencephalography
MFM	Mean Field Model
MI	Mutual Information
MNI	Montreal Neurological Institute
MRI	Magnetic Resonance Imaging
MVAR	Multivariate Vector Autoregressive Model
NMDA	N-methyl-D-aspartate acid
NMM	Neural Mass Model
PC	Principal Component
PC	Peter and Clark method
PCA	Principal Component Analysis
PDC	Partial Directed Coherence
PET	Positron Emission Tomography
PMI	Partial Mutual Information
ROI	Regions of Interest
RSN	Resting State Network
SC	Structural Connectivity
SPM	Statistical Parametric Mapping
std	Standard Deviation
TP	True Positives
WCU	Wilson-Cowan Unit